Order Number 8902340

Personal and doxastic variants of epistemic justification and their roles in the theory of knowledge

Engel, Mylan, Jr., Ph.D.

The University of Arizona, 1988

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

PERSONAL AND DOXASTIC VARIANTS OF EPISTEMIC JUSTIFICATION AND THEIR ROLES IN THE THEORY OF KNOWLEDGE

by

Mylan Engel, Jr.

A Dissertation Submitted to the Faculty of the

DEPARTMENT OF PHILOSOPHY

In Partial Fulfillment of the Requirements
For the Degree of

DOCTOR OF PHILOSOPHY

In the Graduate College

THE UNIVERSITY OF ARIZONA

1988

THE UNIVERSITY OF ARIZONA
GRADUATE COLLEGE

As members of the Final Examination Committee, we certify that we have read the dissertation prepared by Mylan Engel, Jr.

entitled PERSONAL AND DOXASTIC VARIANTS OF EPISTEMIC JUSTIFICATION AND THEIR ROLES IN THE THEORY OF KNOWLEDGE

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

Keith Lehrer	August 1, 1988 Date
Henry C. Byerly	August 1,1988 Date
J. Christopher Maloney	August 1, 1988 Date
	Date
	Date

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

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

Dissertation Director Keith Lehrer — Date August 1, 1988

STATEMENT BY AUTHOR

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

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

SIGNED: 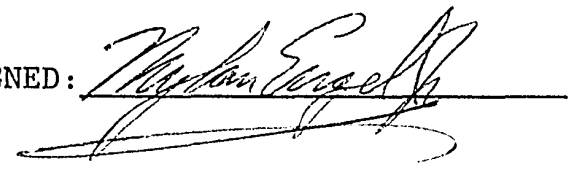

For Mom, Dad, and Jan

ACKNOWLEDGMENTS

First and foremost, I want to thank Keith Lehrer for directing this dissertation. Six years ago, when I began my study at The University of Arizona as a rather green graduate student, Keith took me under his wing and taught me how to do epistemology. Six years and one dissertation later, Keith's influence is everywhere present, especially in Chapter 3. I also want to thank my other committee members, Chris Maloney and Henry Byerly, for their insightful criticisms and helpful stylistic suggestions. The present version is much improved as a result. Another person who deserves special mention is my friend and former professor, Clem Dore, who after numerous discussions finally convinced me to take internalism seriously. I am also indebted to my dear friend Tim Rode who, over the course of many beer-enhanced conversations, helped shape most of my views on epistemology. Tim kindly suffered through each portion of the dissertation as it was produced, always offering sound advice and encouragement.

I also want to thank Alvin Goldman and John Pollock, both of whom contributed indirectly by conducting enlightening courses in epistemology. Their influence is especially obvious in Chapters 2 and 3. Two other professors, Bob Fancher and Michael Hodges, deserve special thanks, for they sparked my interest in philosophy.

My deep gratitude goes out to my typist Barbara Mahan, whose painstaking effort has resulted in a letter-perfect manuscript. I couldn't have done it without you, Barbara. Thanks also to Anne Brunson, who typed portions of the manuscript. I also want to thank Ann Hickman and Lois Day for all their help and secretarial assistance over the years.

Financial support was provided by the National Graduate Fellows Program in 1986-87 and by the Jacob K. Javitts Fellowship Program in 1987-88. Since much of the dissertation was conceived while attending the 1986 NEH Summer Institute on the Theory of Knowledge, I wish to thank the directors, Keith Lehrer and Alvin Goldman, for selecting me to be their graduate assistant participant and also to thank the National Endowment for the Humanities for providing the financial means needed to attend the institute.

Finally, my deepest thanks goes out to my mother and father for their constant encouragement and support, and especially for their confidence in me. They never once doubted that I would complete the dissertation, even when I was not so sure. Equally deep thanks goes to my dear friend and companion Jan Clarke. Jan, thank you for all the encouragement, kindness, and love, for all of the sacrifices on your part, and most of all, for putting up with me while I lived this dissertation. I dedicate my dissertation to these last three people.

TABLE OF CONTENTS

TABLE OF CONTENTS--continued

TABLE OF CONTENTS--continued

ABSTRACT

Most epistemologists agree that epistemic justification is a requirement for knowledge. This requirement is usually formulated in one of two ways:

(JR1) S knows that p only if S is epistemically justified in believing that p.

(JR2) S knows that p only if S's belief that p is epistemically justified.

Surprisingly enough, (JR1) and (JR2) are generally regarded as synonymous, stylistic variants of the justification condition. In Chapter 1, I argue that such a synonymy thesis is simply mistaken and that, in fact, (JR1) and (JR2) specify substantively different requirements. After all, (JR1) requires that the person (or would-be knower) be epistemically justified, whereas (JR2) requires that the belief in question be epistemically justified, and intuitively these constitute different requirements. Thus, it is concluded that (JR1) and (JR2) employ inherently different kinds of epistemic justification in their respective analysantia. I dub them "personal justification" and "doxastic justification", respectively. The remainder of the dissertation is devoted to demonstrating the legitimacy of the personal/doxastic justification distinction and to tracing out its ramifications for the

theory of knowledge. For example, in Chapter 2, we see that the personal/doxastic justification distinction accounts for the divergent intuitions that regularly arise regarding justificatory evaluations in demon world contexts.

In Chapters 2 and 3 I provide analyses for doxastic justification and personal justification, respectively. Chapter 2 spells out an externalist reliabilist account of doxastic justification which safely avoids demon world counterexamples. In Chapter 3, an internalist coherence account of personal justification is advanced. In defending this coherence theory, I argue that all foundation theories are false and that the regress argument on which they are predicated is unsound. With accounts of doxastic and personal justification in hand, I turn to the task of analyzing knowledge.

In Chapter 4, I propose an analysis of ordinary knowledge which only requires doxastic justification. Even so, personal justification has a negative, undermining role to play in the analysis. I then demonstrate that this analysis of knowledge is immune to typical Gettier examples. It also remains unscathed by Harman's beefed-up Gettier cases. Finally, I consider a stronger analysis of knowledge which requires both doxastic and personal justification. Though the latter analysis proves too strong for ordinary knowledge, it remains interesting as an analysis of a more

intellectualistic kind of knowledge.

The final chapter examines the internalist/externalist controversy and demonstrates that this controversy is a direct result of the failure to distinguish personal justification from doxastic justification.

CHAPTER 1

CLARIFYING "EPISTEMIC JUSTIFICATION"

1. The Need for Such Clarification

Many philosophers regard epistemic justification as the most important notion in epistemology.[1] As a result, the current epistemological literature has been inundated with theories which purport to analyze the concept of epistemic justification. However, the interest in epistemic justification is hardly a contemporary phenomenon. It certainly dates back to Descartes, who sought to ground all his knowledge on the firmest justificatory foundation possible, and it may even date back to the query in Plato's Theaetetus of what must be added to true judgment in order to obtain knowledge. This widespread interest in epistemic justification can be accounted for primarily in two ways. First, epistemic justification is generally regarded as the necessary condition for knowledge which rules out lucky

---

[1] E.g., Roderick Chisholm asserts, "It is certainly true that the concept of justification may be thought of as the central concept of the traditional theory of knowledge." [See his "The Place of Epistemic Justification", Philosophical Topics, Vol. 14, No. 1 (Spring, 1986), p. 85.]; John Pollock maintains that epistemic justification is the principal focus of epistemology. [See Chapter 1 of his Contemporary Theories of Knowledge (in manuscript).] Stewart Cohen suggests that the difference between epistemically justified and epistemically unjustified belief marks the central distinction in epistemology. [See his dissertation Justification and Truth, p. 9.].

guesses. So, those interested in providing an account of knowledge have ipso facto been interested in providing an account of epistemic justification suitable for such knowledge. Second, epistemic justification is intrinsically interesting in its own right, since many people want to know when believing a proposition is justified.

Despite the widespread interest in epistemic justification and its analysis, surprisingly little work has been done to clarify just what concept it is that epistemologists have been trying to analyze. This singular hiatus in contemporary epistemology is probably best explained by the fact that most epistemologists have simply failed to see the need for such conceptual clarification. Instead, they have just taken it for granted that there is a common ordinary notion of epistemic justification and have offered various analyses intended to capture this ordinary notion. However, this "capture the ordinary notion" approach is best viewed as a reductio of the claim that there is such a unique notion of epistemic justification which all epistemologists share. The analyses proposed by those using this approach vary so wildly in the beliefs that they count as justified that they cannot plausibly be construed as analyses of the same concept. As Alvin Plantinga puts it,

> The differences among these views are enormous; this is by no means a case of

> variations on the same theme. Indeed, disagreement is so deep and radical it is sometimes hard to be sure the various disputants are discussing approximately the same issue.[2]

Moreover, the intuitions used to bolster these divergent views are so disparate that they must inevitably be driven by competing conceptions of epistemic justification.

Having recognized that most epistemological theorizing has had, as its starting point, these unspecified, competing conceptions of epistemic justification, a few philosophers, most notably Alvin Plantinga and William Alston, have sought to clarify the concept of epistemic justification in a theory neutral way. Their procedure for clarifying the concept has been, roughly, to point out those features of epistemic justification which seem to be shared by the various competing conceptions and then to regard these shared features as constitutive of the ordinary notion of epistemic justification. But, as we shall see, this procedure fares no better in providing a unitary concept of epistemic justification.

Plantinga identifies three elements fundamental to the concept of epistemic justification.[3] First, 'justification'

---

[2]Alvin Plantinga, "Justification and Theism" (in manuscript), p. 1.

[3]Ibid., pp. 2-3.

is a term of positive epistemic appraisal such that "to say that a proposition is justified for a person is to say that his believing or accepting it has positive epistemic status for him."[4] Second, epistemic justification admits of degrees. And third, epistemic justification (or something close to it) is what must be added to true belief to get knowledge. Plantinga sums up what he means by the term 'epistemic justification' as follows:

> Initially, then, and to a first approximation, we can identify justification or positive epistemic status as a normative property that comes in degrees, and that lies in the near neighborhood of what distinguishes true belief from knowledge.[5]

Alston offers an initial conception of epistemic justification slightly different from, but apparently compatible with, Plantinga's conception. Alston begins by distinguishing between "one's being justified in believing that p, and one's justifying one's belief that p, where the latter involves one's doing something to show that p, or to show that one's belief was justified,"[6] and he then asserts that he "will be concentrating on the 'be justified' side of

---

[4] Ibid., p. 2.

[5] Ibid., p. 3.

[6] William Alston, "Concepts of Epistemic Justification", The Monist, Vol. 68, No. 1 (January, 1985), p. 58.

this distinction, since that is of more fundamental epistemological interest."[7] This is already a substantive claim on his part, which is not common to all conceptions of justification, e.g. in Knowledge, Keith Lehrer seems more interested in the "justifying" side of Alston's distinction, but for the sake of exegesis let us suppose, for the moment, that Alston has located the fundamental sense of epistemic justification. Alston identifies the following four features as the common ground of this "be justified" sense of epistemic justification:

> (1) It applies to beliefs, or alternatively to a cognitive subject's having a belief.[8]
>
> (2) It is an evaluative concept, in a broad sense in which this is contrasted with "factual." . . . It is to accord S's believing a positive epistemic status.[9]
>
> (3) It has to do with a specifically epistemic dimension of evaluation. . . . Epistemic evaluation is undertaken from what we might call the "epistemic point of view." That point of view is defined by the aim at maximizing truth and minimizing falsity in a large body of beliefs.[10] And,

---

[7] Ibid.

[8] Ibid.

[9] Ibid.

[10] Ibid., p. 59.

(4) It is a matter of degree.[11]

He also suggests that epistemic justification so construed is a necessary condition of knowledge.[12]

In a recent article, Roderick Chisholm also attempts to clarify the concept of epistemic justification, which he takes to be the central concept of the traditional theory of knowledge. He maintains that "the sense of 'justify' that is central to the traditional theory of knowledge pertains to the question whether the belief may be said to be reasonable."[13] For Chisholm, 'justified belief' and 'reasonable belief' are synonymous expressions, and he specifies three ways of demarcating the sense of "reasonable" which is crucial to the epistemological enterprise. First, "It provides us with the materials by means of which we can answer the question of the Theaetetus: 'What does one add to the concept of true belief to get the concept knowledge?'"[14] Second, "The sense of 'reasonable' with which we are concerned is that which provides us with the means of defining the other fundamental concepts of the

---

[11] Ibid.

[12] Ibid., p. 58.

[13] Roderick Chisholm, op. cit., p. 86.

[14] Ibid.

theory of knowledge."[15] Thus, Chisholm takes 'reasonable' to be a primitive, unanalyzable epistemic term in terms of which all other epistemic terms are to be defined. And third,

> The relevant sense of reasonable belief is one which is such that a believer can ascertain by himself at any time which of his beliefs are reasonable for him at that time. . . . Hence reasonability is properly called an "internal" concept.[16]

Finally, Chisholm contends that there is no direct relation between a belief's being true and its being justified in this "reasonable belief" sense.[17]

Chisholm's last contention is directly at odds with one of the ingredients that Stewart Cohen takes to be essential to the concept of epistemic justification. Cohen locates two features constitutive of such justification: (1) epistemic justification is essentially a normative concept for guiding and evaluating reasoners,[18] and (2) it is internally connected to truth, since this is what distinguishes it from moral and pragmatic justification.[19]

---

[15] Ibid.

[16] Ibid.

[17] Ibid, p. 90.

[18] Stewart Cohen, op. cit., pp. 8 and 13.

[19] Ibid., pp. iv and 2.

Thus, whereas Chisholm thinks that there is no internal connection between epistemic justification and truth, Cohen contends that an indispensable component of epistemic justification which distinguishes it from other senses of "justification" is its conceptual connection with truth. Interestingly enough, Cohen goes on to argue that those theories of epistemic justification that do provide such a connection to truth fail because they cannot accommodate the normativity of justification,[20] which makes one wonder if Cohen has isolated a coherent notion of epistemic justification.

John Pollock, like Chisholm, eschews an internal connection between epistemic justification and truth. For Pollock, epistemic justification is essentially a permissibility notion. As he puts it, "A justified belief is one that it is 'epistemically permissible' to hold. Epistemic justification is a normative notion. It pertains to what you [epistemically] _should_ or _should not_ believe."[21] But unlike Chisholm, Pollock maintains that epistemic justification in this belief-guiding normative sense is not

---

[20] _Ibid._, Chapters 1, 2 and 5.

[21] John Pollock, _op. cit._, p. 9.

all that must be added to true belief to get knowledge, not even in non-Gettier situations.[22]

Alvin Goldman, on the other hand, does think that epistemic justification is intimately connected with truth.[23] In "What Is Justified Belief?" he attempts to provide an account of epistemic justification which results in epistemically justified beliefs being probably true.[24] As Goldman conceives of it, epistemic justification is an evaluative concept, but, in contrast to Chisholm, it is primarily an "external" concept in Goldman's estimation, since he maintains that a belief can be justified for a person without that person being aware that it is justified and without that person "possessing" anything which could be called a "justification".[25] He also maintains that epistemic justification can be successfully reduced to the non-epistemic and that an appropriately deep or revelatory

---

[22] Ibid., see Appendix, pp. 203-218.

[23] In Epistemology and Cognition Goldman contends that the criterion of Justification-rule rightness is a truth-linked criterion, to wit, a system of J-rules is right iff it results in a sufficiently high ratio of true beliefs to total beliefs. Alvin Goldman, Epistemology and Cognition (Cambridge, Mass. and London: Harvard University Press, 1986). See Chapter 5, sections 5.5 - 5.9.

[24] Alvin Goldman, "What Is Justified Belief?", Justification and Knowledge, ed. George Pappas (Dordrecht: D. Reidel Publishing Company, 1979).

[25] Ibid., pp. 1-2.

account of epistemic justification will evince such a reduction.[26]

To my knowledge, this exhausts those philosophers who have actually attempted to clarify the concept of epistemic justification with which they are working. So, I will summarize the "findings" of this section by listing the things that these philosphers have said regarding the concept of epistemic justification.

(1) Epistemic justification is a normative or evaluative concept.

(2) It admits of degrees.

(3) It is to be equated with positive epistemic status.

(4) It is to be equated with epistemic reasonableness.

(5) It is what must be added to true belief to get knowledge.

(6) It lies in the neighborhood of what it is that must be added to true belief to get knowledge.

(7) It is not all that must be added to true belief to get knowledge.

(8) It is an "internal" concept in the sense of being directly accessible to those who possess it.

(9) It is an "external" concept in the sense that cognizers can, and often do, lack access to the justificational status of their beliefs, even when those beliefs are justified.

---

[26] Ibid.

(10) It has an internal, conceptual connection to truth.

(11) It is not internally connected to truth.

(12) It applies to beliefs, or alternatively to a subject's having a belief.

(13) It is an irreducible, unanalyzable epistemic primitive.

(14) It can be successfully reduced to the non-epistemic.

(15) It is a permissibility notion.

Suffice it to say that the waters are quite muddied where the concept of epistemic justification is concerned. In light of the "findings" listed above, there appears to be no unitary concept of epistemic justification. About the only thing which does seem to be universally accepted is that epistemic justification is in some sense a normative notion which admits of degrees. I trust that the need for a conceptual clarification of "epistemic justification" is now patently evident, and in the remainder of this chapter, I will attempt to provide such clarification. But a comment is in order concerning the direction that this conceptual clarification should take. I regard the fundamental lesson of this section to be the realization that there is no unitary notion of epistemic justification. There simply is no single concept of epistemic justification which can do all of the things epistemologists have expected epistemic justification to do. So, I shall not attempt to isolate a

single sense of "epistemic justification", for any such attempt seems doomed at the outset. Instead, my tack will be to isolate a small family of epistemically evaluative concepts, concepts which have heretofore been batted around under the single heading "epistemic justification".

2. A Traditional View

Very frequently, epistemic justification is identified with that which must be added to true belief (at least in non-Gettier situations) in order to obtain knowledge. The motivation for this view stems from the recognition that true belief is not sufficient for knowledge. Mere true belief falls short of knowledge because a person can come to hold a belief for all sorts of ludicrous reasons or can be caused to have a belief by some epistemically illegitimate belief-forming cognitive process such as wishful thinking, and in such cases, when the belief happens to be true, it seems obvious that the person does not know that it is true. Moreover, since knowing a proposition just is knowing that that proposition is true, it seems clear that the person just described does not know that which she believes. An example will illustrate the point. Consider Sally the sports fan. Sally is a die-hard Chicago Cubs fan. Every year prior to the season Sally forms the belief that the Cubs will win the pennant that year, simply out of wishful thinking. Of course, as anyone who follows the Cubs knows,

her beliefs in this regard have always been false. But Sally's faith in the Cubs is unshakable, and so, purely out of wishful thinking, she forms the belief that the Cubs will win the pennant in 1988. Suppose that, as luck would have it, the Cubs do win the pennant in 1988. Stranger things have happened (though not many). Surely, in this situation despite her true belief, we would not want to say that Sally knew that the Cubs would win.[27] Our reluctance in ascribing knowledge to Sally in the case described derives from the fact that it is simply a matter of luck (and, given the Cubs' history, a great deal of luck) that her belief happens to be true this time. There is an overriding intuition that beliefs which only luckily turn out to be true fall short of knowledge.

Since true belief is not enough for knowledge, something else is needed for a person to know a proposition. As we have seen, that something else is generally thought to

---

[27]Sally's lack of knowledge becomes more obvious when we contrast Sally with Ina the informed sports fan. Ina, who has never had much hope for the Cubs before, has followed the Cubs extremely closely. On the basis of their latest recruits, as well as their gradual improvement over the years, Ina feels confident that the Cubs will win the pennant in 1988 and believes accordingly. Again, assuming that the Cubs do win, Ina has a true belief. Even so, given all the contingencies of professional baseball, we are reluctant to count Ina's well-founded belief as knowledge. Surely, if Ina's well-founded true belief does not constitute knowledge, Sally's wishful true belief falls far short of knowledge.

be epistemic justification, and a sufficiently high degree of epistemic justification at that. In the context of our Sally example, epistemic justification's role in an account of knowledge is quite clear. The epistemic justification requirement is intended to rule out lucky guesses as instances of knowledge. The intuition behind the justification requirement is basically that when a person's belief is (very well) justified, it is no mere matter of luck when that belief happens to be true. Put another way, a belief's epistemic justification is thought to be an indication of its truth, thus limiting luck's role in the belief's being true. On this view, the more justified a belief is epistemically, the more likely it is that that belief is true.

It should be obvious, however, that only a certain conception of epistemic justification is properly suited to play the role of the justification requirement for knowledge. Recall that the role of the justification requirement is to rule out (or at least greatly limit) the role of luck in knowledge. In order to limit the element of luck in knowledge, a belief's epistemic justification must be an indication of that belief's truth. But for a belief's epistemic justification to be an indication of its truth, epistemic justification must be in some way conceptually connected with truth. For epistemic justification to be conceptually connected to truth, it must be the case that

for every possible world W, if conditions C make person S's belief B epistemically justified in W, then conditions C make it probable that B is true in W.[28] Lehrer and Cohen call this sort of conceptual connection a "truth connection".[29] The need for a truth connection is straightforward. If there were no connection between epistemic justification and truth, then it would be just as much a matter of luck when a justified belief turned out to be true as when an unjustified belief turned out to be true. Furthermore, a better justified belief would be no more likely to be true than a much less well justified belief, for without a truth connection no amount of epistemic justification is an indication of truth. Thus, if epistemic justification is to be indicative of truth and thereby limit the element of luck in knowledge, it must be internally connected with truth. Laurence Bonjour explains the need for such a truth connection rather eloquently as follows:

---

[28]This spells out the kind of connection required by fallibilist theories of epistemic justification. On an infallibilist theory, epistemic justification is conceptually connected to truth iff for every possible world W, if conditions C make S's belief B justified in W, then conditions C logically entail that B is true in W. Since infallibilism leads directly to skepticism, the kind of justification needed for knowledge must have a fallibilist connection to truth, as specified in the text above.

[29]Keith Lehrer and Stewart Cohen, "Justification, Truth, and Coherence" Synthese, Vol. 55 (1983), p. 191.

> a satisfactory defense of a particular standard of epistemic justification must consist in showing it to be truth-conducive. . . . Without such a meta-justification, a proposed standard of epistemic justification lacks any underlying rationale. Why after all should an epistemically responsible inquirer prefer justified beliefs to unjustified ones, if not that the former are more likely to be true? To insist that a certain belief is epistemically justified, while confessing in the same breath that this fact about it provides no good reason to think that it is true, would be to render nugatory the whole concept of epistemic justification.[30]

Accordingly, anyone wishing to analyze the "what must be added to true belief to get knowledge" conception (hereafter the "knowledge conception") of epistemic justification inevitably faces the onus of specifying the nature of such justification's connection with truth.

What other features besides truth-connectedness are constitutive of the knowledge conception of epistemic justification? As we have already seen, it admits of degrees. Some beliefs are more justified than others. It is a notion of epistemic appraisal in that justified beliefs have positive epistemic status to the degree in which they are justified. Thus, it is broadly speaking a normative or evaluative notion since it characterizes some beliefs as

---

[30]Laurence Bonjour, "Can Empirical Knowledge Have a Foundation?", American Philosophical Quarterly, Vol. 15, No. 1 (January, 1978), p. 5.

being better than others for the purposes of gaining truth and avoiding error. We might summarize the knowledge conception of epistemic justification as follows:

> (KC) Epistemic justification is a graded normative notion of positive epistemic appraisal that bears an essential internal connection with truth, a certain degree of which is necessary for knowledge.

It remains to be seen whether the knowledge conception of epistemic justification, to wit (KC), can be successfully analyzed, for if Cohen is right, then any analysis of epistemic justification which affixes the needed truth connection will lack the required normativity, and hence, the kind of epistemic justification characterized in (KC) will not be satisfactorily analyzable. Analyzable or not, (KC) certainly isolates one conception of epistemic justification which is prevalent in the epistemological literature. I will eventually argue that the (KC) conception of epistemic justification is, in fact, analyzable and that in order to provide the required truth connection such an analysis must be externalist in nature. However, since the task at hand is that of clarifying the concept of epistemic justification, I will use the next section to distinguish two other viable conceptions of such justification.

3. Two Additional Conceptions of Epistemic Justification: Personal and Doxastic

In this Section, I will identify two distinct conceptions of epistemic justification, two conceptions which to my knowledge have been conflated by every epistemologist that has considered the subject. The main contention of this dissertation is that most of the confusion which surrounds epistemic justification, especially that which surrounds the internalist/externalist controversy, is directly traceable to this conflation. In the chapters that follow, I will argue that both kinds of epistemic justification have crucial roles to play in epistemology.

As has been mentioned previously, most epistemologists take epistemic justification to be an essential ingredient of knowledge, even those epistemologists who deny that it is all that must be added to true belief to get knowledge. The justification requirement for knowledge has been formulated in various ways. Two of the most common formulations are:

(JR1) S knows that p only if S is epistemically justified in believing that p.

And,

(JR2) S knows that p only if S's belief that p is epistemically justified.[31]

---

[31]For the sake of simplicity, the required temporal indices have been suppressed in (JR1) and (JR2).

(JR1) and (JR2) have generally been regarded to be synonymous formulations of the justification requirement. Accordingly, which formulation is used is thought to be indicative of nothing more than stylistic preference. The synonymy of (JR1) and (JR2) entails what I call the "equivalency thesis". The equivalency thesis asserts:

> (ET) S is epistemically justified in believing that p iff S's belief that p is epistemically justified.

According to the equivalency thesis, there is no difference between S being justified in believing that p and S's belief that p being justified. Most epistemologists tacitly embrace the equivalency thesis, since they jump back and forth between talking about S being justified and S's belief being justified. However, some are more explicit in their commitment to the equivalency thesis. For example, William Alston says of epistemic justification:

> It applies to beliefs, or alternatively to a cognitive subject's having a belief. I shall speak indifferently of S's belief that p being justified and of S's being justified in believing that p[32] (emphasis added).

I contend that the equivalency thesis is false and that embracing it has led philosophers astray in their

---

[32]William Alston, op. cit., p. 58.

epistemological theorizing and criticism. However, before explaining why I think (ET) is false, I want to discuss its significance.

I have suggested that most (all?) epistemologists endorse the equivalency thesis and its corollary the synonymy of (JR1) and (JR2). Nevertheless, which version of the justification condition they adopt varies, and it varies in a somewhat systematic way. Generally, internalists adopt the (JR1) formulation of the justification requirement. For example, in Knowledge, which spells out an internalist coherence theory of justification, Lehrer formulates the justification condition in the following way:

> If S knows that p, then S is completely justified in believing that p.[33]

In Knowledge and Justification Pollock, also an internalist, speaks indifferently of beliefs being justified and of person's being justified in beliefs, but when he formally presents conditions, they are generally in the style of (JR1), e.g.

> If P is a prima facie reason for S to believe that Q, and S justifiably believes-that-P and believes-that-Q on the basis of his belief-that-P, then S is justified in

---

[33]Keith Lehrer, Knowledge (Oxford: Oxford University Press, 1974), p. 13.

> believing-that-Q iff he does not believe any defeaters for this prima facie reason[34] (emphasis added).

And in explaining what it is to have a reason, he states, "In order for a person to have a reason for believing something, it must be a good reason, and he must be justified in believing that it is true"[35] (emphasis added). In contrast, externalists seem to embrace the (JR2) formulation. The clearest example of this derives from Goldman's "What Is Justified Belief?", where he asserts, "A theory of justified belief will be a set of principles that specify truth-conditions for the schema |S's belief in p at time t is justified| ."[36] I am not claiming that every internalist uses (JR1) and every externalist uses (JR2), since most epistemologists use (JR1) and (JR2) interchangeably. What I am suggesting is that (JR1) captures something central to internalism and (JR2) captures something central to externalism. Now, if, as I maintain, the equivalency thesis is false, then it follows that (JR1) and (JR2) are not synonymous. If (JR1) and (JR2) are not synonymous, it may be that, rather than disagreeing with

---

[34]John Pollock, Knowledge and Justification (Princeton, NJ: Princeton University Press, 1974), p. 44.

[35]Ibid., p. 35.

[36]Alvin Goldman, "What Is Justified Belief?", op. cit., p. 3.

each other, internalists and externalists have for the most part been arguing past each other, which is what I take to be the case. Of course, this latter contention rests on the falsity of the equivalency thesis. So, I shall now explain why I think (ET) is false.

Recall that the equivalency thesis maintains:

> (ET) S is epistemically justified in believing that p iff S's belief that p is epistemically justified.

When I look at both sides of this bi-conditional, they seem to be so remarkably different that it is surprising that anyone has taken them to be extensionally equivalent, much less synonymous. Since they have different domains of evaluation,[37] they do not even purport to be about the same thing. The left-hand side of the bi-conditional is evaluating S, the would-be knower, as being epistemically

---

[37] In Epistemology and Cognition Goldman, an externalist, suggests a variety of domains which are subject to epistemic evaluation. The epistemically evaluatable domains that he recognizes are: beliefs, methods, psychological processes, hypothesis-forming processes, concept-forming processes, search processes, second-order processes, speech acts, institutional arrangements, and social structures and processes. The one domain he overlooks is: persons. This supports my contention that externalists have been concerned with evaluating beliefs, not persons, while internalists have focused on evaluating persons, not beliefs, and that, thus, unbeknownst to them internalists and externalists have not been discussing the same subject. [See Epistemology and Cognition, op. cit., p. 21.]

justified, and hence, it has persons or, more broadly, cognizers as its domain of evaluation. In contrast, the right-hand side of the bi-conditional is evaluating the justificatory status of S's belief, and consequently, it has beliefs for its domain of evaluation.

Given that the two sides of the bi-conditional (ET) have different domains of evaluation, a natural question arises concerning why they have generally been taken to be extensionally equivalent. The only answer forthcoming seems to be that epistemologists must have contended that a believer, qua believer, cannot be evaluated apart from that which she believes and that a belief cannot be evaluated differently than the cognizer who holds the belief. It seems obvious to me that both contentions are mistaken. For one, if we consider standard epistemological practice, we find that people are frequently evaluated in terms of the reasoning which leads to their beliefs, rather than the beliefs on which they actually settle. For example, a person is often thought to be justified if she has reasoned well (or if she has done her best to reason well) regardless of what belief she adopts. Regarding the flipside of this epistemological coin, we hear such things as: "Though her belief is a reasonable one to hold, she came to hold it in an epistemically irresponsible way, and consequently, she is unjustified in believing what she does." An example may be

helpful here: Consider Nancy, another sports fan whose favorite basketball team is the Arizona Wildcats. Midway through the 1987-88 season, Nancy formed the belief that Arizona's basketball team would win the PAC 10. This belief was eminently reasonable at that time, since by then Arizona had beaten every team in the conference, winning by an average of 29 points per game, and was itself undefeated in conference play. Moreover, Nancy bases this belief on Arizona's impressive record. However, Nancy also believes that astrology is a completely reliable science, and her horoscope on the day she formed her belief about Arizona, which she did read, said, "Your favorite basketball team will not win the PAC 10; so, do not bet on them." Here we have a case where Nancy's belief is based on reliable evidence and, hence, is justified, but Nancy is unjustified in believing it, because she has what she takes to be conclusive counterevidence for her belief, counterevidence which she simply chooses to ignore. In this case her belief is evaluated positively while she is evaluated negatively in direct contradiction to the claim that beliefs and believers cannot be evaluated independently.

I can now summarize my argument for (ET)'s falsity as follows: First, the two sides of (ET) clearly have different domains of evaluation, and consequently, they don't even purport to be about the same thing. Therefore, intuitively, they seem to spell out different requirements.

Surely, we can agree that the burden of proof lies with the person who wants to butt heads with the intuitive by maintaining that (ET) is true. But the only proof (if you can call it that) that has been offered for (ET)'s truth is the contention that beliefs and believers are not subject to independent epistemic evaluations, and we have just seen with the Nancy example that this contention is false. In light of the intuitive evidence for (ET)'s falsity and the lack of any compelling evidence for its truth, I submit that (ET) is false and that therefore (JR1) and (JR2) embody inherently different conceptions of epistemic justification. Since (JR1) is concerned with evaluating persons, I will call the kind of epistemic justification associated with it "personal justification". And since (JR2) is evaluative of beliefs, I will call the kind of epistemic justification underlying it "doxastic justification". More precisely, the personal conception of epistemic justification is:

> (PJ) Personal justification is a normative notion in terms of which persons are evaluated from the epistemic point of view.

A succinct statement of the doxastic conception of epistemic justification is:

> (DJ) Doxastic justification is a normative notion in terms of which beliefs are evaluated from the epistemic point of view.

I contend that both personal justification and doxastic justification are viable conceptions of epistemic justification with important roles in epistemology, though I will argue that doxastic justification is more fundamental to the traditional epistemological desideratum of analyzing knowledge. In the chapters that follow, I will attempt to analyze both kinds of epistemic justification and delineate their respective roles in the theory of knowledge. In addition, I will argue that internalists have been concerned primarily with personal justification, whereas externalists have focused on doxastic justification, and that this is what accounts for the radical divergence in their views. However, at this point, I fear that there are still some philosophers who will find my distinction between personal and doxastic justification a spurious one. So, in the next section I will offer further support for the legitimacy of the distinction.

4. Bolstering the Personal/Doxastic Justification Distinction via an Ethical Analogy

As mentioned above, I fear that some epistemologists entrenched in the equivalency thesis tradition will regard the distinction between personal and doxastic justification as nothing more than a false dichotomy. Such theoretical entrenchments aside, the crucial point at issue here seems to be whether or not cognizers and beliefs are subject to

independent and sometimes different epistemic evaluations. I take it to be noncontroversial that cognizers are agents, to wit epistemic agents, and that beliefs can be viewed as epistemic actions. Thus, the question at hand is whether or not epistemic agents and epistemic actions can be evaluated independently. Since one may lack clear intuitions on how this question should be answered, I suggest that we look for guidance to an analogous question in normative ethics, namely, "Are agents and actions subject to independent moral evaluations?"

The answer to this latter question is an uncontrovertable "Yes". Ethicists have readily recognized that agents and actions are open to independent and sometimes discrepant moral evaluations. For example, it is generally acknowledged that an agent may be morally virtuous in performing some action A even though A is morally wrong and, conversely, that an agent may be morally wicked in performing a morally right action. Moreover, in the moral domain we frequently find the criteria for evaluating agents and actions to be entirely distinct. Kant, for example, is plausibly interpreted as offering such distinct criteria of ethical evaluation. For Kant, an agent has moral worth iff she performs her action out of respect for the moral law (i.e. she acts out of a sense of duty), whereas an action is right iff it satisfies the universalizability criterion (or

some other supposedly equivalent formulation of the categorical imperative). Thus, on the Kantian view just described, an agent S who performs action A is to be deemed morally virtuous provided that S does A out of a sense of duty, even if, as misfortune would have it, A is morally wrong; and, conversely, an agent S who performs action A is to be deemed morally vicious if S does A for wicked motives, even if A accords with duty.[38]

Unfortunately, recognizing that agents and actions can in this way be evaluated independently does not resolve our earlier query. All we have shown is that the following two moral theses are false:

(MT1) Agent S is morally virtuous in doing action A iff A is morally right.

(MT2) Agent S is morally wicked in doing action A iff A is morally wrong.

---

[38]The arguments in this section do not depend on Kant's having actually held the view that I am attributing to him. Though I suspect that Kant did indeed hold a view very similar to the one that I have labelled "the Kantian view", I have admittedly oversimplified the view somewhat. For example, regarding attributions of moral virtuosity, the following probably comes closer to Kant's view:

> S is morally virtuous in doing A iff (1) S does A out of a sense of duty and (2) S has done her best to assess that A is her duty.

Such subtleties aside, the point is not whether this is Kant's view, but whether it is legitimate to evaluate agents and actions independently, and I think that "the Kantian view" as I have presented it shows that it is legitimate.

Of course, neither (MT1) nor (MT2) is analogous to the equivalency thesis. (MT1) and (MT2) are analogous to the following two epistemological theses:

(ET1) Cognizer S is epistemically virtuous in believing that p iff the belief that p is epistemically right.

(ET2) Cognizer S is epistemically reprehensible for believing that p iff the belief that p is epistemically wrong.

The only plausible way to cash out epistemic rightness and wrongness in the above context is as truth and falsity, respectively, and so we can rephrase (ET1) and (ET2) as:

(ET1') Cognizer S is epistemically virtuous in believing that p iff the belief that p is true.

(ET2') Cognizer S is epistemically reprehensible for believing that p iff the belief that p is false.

It takes little reflection to see that neither (ET1') nor (ET2') has anything going for it. Moreover, one can consistently maintain that the equivalency thesis (ET) is true while denying the truth of (ET1') and (ET2').

One thing the foregoing considerations show is that the legitimacy of the personal/doxastic justification cannot be established simply by showing that agents and actions are subject to independent moral evaluations. Nevertheless, the above considerations have been fruitful in another way.

They have shown that the analogy between normative ethics and normative epistemology is quite exact. Just as we can evaluate agents as morally virtuous and morally reprehensible, so too can we evaluate cognizers as epistemically virtuous and epistemically reprehensible. Just as some actions are morally right and others morally wrong, some beliefs are epistemically right and others epistemically wrong. Given the exactness of the analogy between ethics and epistemology, we may be able to shed some light on the equivalency thesis by looking at its moral analogue.

The exact moral analogue of the equivalency thesis asserts:

> (MA) Agent S is morally justified in doing action A iff action A is morally justified.

This baldly stated, most ethicists would no doubt maintain that (MA) is false, probably on the grounds that (MA) conflates subjective and objective notions of moral justification. For example, they might well claim that S is morally justified in doing A iff S is subjectively justified in thinking that A is right (or, at least, permissible) and that A is morally justified iff A is justified in some objective sense. Actually, I think that the subjective/objective justification distinction, at least as it is characteristically drawn, is not what is needed to

demonstrate the falsity of (MA), as I will argue momentarily; but what is interesting for our purposes is that in the moral domain (MA) [i.e. the exact moral analogue of the equivalency thesis] is extremely suspect, and accordingly, analogical reasoning suggests that we should be suspicious of (ET), as well. But if the problem with (MA) is not that it conflates subjective and objective justification, then what is wrong with (MA)? To answer this question, I will begin by explaining why I think that the appeal to the subjective/objective justification distinction is inappropriate.

The subjective/objective justification distinction in ethics is normally drawn in such a way that "subjective justification" and "objective justification" have the same domain of evaluation, to wit, actions. For example, Richard Feldman asserts,

> It is widely held that there is a distinction in ethics between those actions that are objectively justified and those that are subjectively justified. Roughly, an action is objectively justified when it is in fact the morally best action open to the agent, while an action is subjectively justified when, in some sense, it seems best from the agent's perspective[39] (emphasis added).

---

[39]Richard Feldman, "Subjective and Objective Justification in Ethics and Epistemology" (in manuscript), p. 1.

But if subjective justification has actions for its domain of evaluation, then it cannot be the kind of justification which is being employed in the left-hand side of (MA), since what is being evaluated as morally justified in the left-hand side of (MA) is the agent S and not some action. In my opinion, what has gone amiss with (MA) is exactly analogous to the earlier objection which I raised to (ET), namely, the left-hand and right-hand sides of (MA) are evaluating different domains. Whereas the left-hand side of (MA) is evaluating agents, the right-hand side is evaluating actions. It is this fact about (MA) which makes it suspect. Since S may be extremely well justified, in terms of the things she believes, in doing A, even though A itself is unjustified, (MA) is false. Of course, one might want to define 'subjective justification' in such a way that it attaches to agents and not actions, and then claim that (MA) is conflating objective justification with subjective justification in this newly defined sense. I would welcome such a move on the part of an ethicist, for then subjective moral justification would be exactly analogous to personal justification and objective moral justification would be the moral analogue of doxastic justification.

I realize, of course, that those who found my appeal to different domains of evaluation to show the falsity of (ET) unconvincing may also be unconvinced by my appealing to different domains of evaluation to explain the falsity of

(MA). However, even prior to such an appeal, (MA) seems intuitively suspect. If my "different domains of evaluation" diagnosis of (MA) is mistaken, then we are still in need of an explanation of what is wrong with (MA). We have already seen that the appeal to the subjective/objective justification distinction fails to explain what is wrong with (MA), and I can think of no viable explanation other than the one I have given. Whatever explanation one settles on, (MA) intuitively seems false, and so, given its analogousness to (ET), we have reason to think that (ET) is likewise false, which is one of the things that I have been attempting to establish.

Nevertheless, I am not only committed to the falsity of (ET), but to the correctness of the "different domains of evaluation" diagnosis of that falsity, as well. I contend that the findings of this section make such a diagnosis eminently plausible. We have seen that agents and actions, as well as cognizers and beliefs, are subject to independent evaluations, for this is entailed by the falsity of (MT1), (MT2), (ET1') and (ET2'). Now since justification is itself an evaluative concept, by parity of reasoning it seems that we should be able to evaluate the justificatory status of agents and actions and of cognizers and beliefs independently, and in the epistemic realm, the personal/doxastic justification distinction provides us with the tools to do just that.

This exhausts the intuitive considerations that I have to offer in support of the distinction between personal and doxastic justification. I urge the reader who does not find the distinction an intuitive one simply to regard it as a stipulative one for the time being. After all, the surest test of a distinction's genuineness is not its intuitiveness, but rather the work that it does. The remaining chapters of this dissertation will demonstrate that the personal/doxastic justification distinction scores especially well on this latter test.

CHAPTER 2

RELIABILISM AS DOXASTIC JUSTIFICATION

1. Introduction

In the previous chapter, I introduced three distinct conceptions of epistemic justification, to wit, the knowledge conception, the doxastic conception, and the personal conception. In this chapter, I will argue that a version of process reliabilism provides the correct account of doxastic justification. Certain needed refinements aside, I will defend the view that a belief is doxastically justified iff it results from one or more reliable belief-forming cognitive processes [BCP's]. As a point of departure for this defense, I devote section 2 to explicating Goldman's historical reliabilism as formulated in his "What Is Justified Belief?", since I think that his formulation therein comes quite close to being the correct analysis of doxastic justification. In section 3, I will present several objections to Goldman's theory which purportedly show that a belief's being produced by a reliable BCP is neither necessary nor sufficient for that belief's being justified. In section 4, I will argue that these objections, which arise out of the failure to distinguish personal justification from doxastic justification, simply do not apply if Goldman's reliability theory is construed as an analysis of doxastic

justification. Finally, in section 5, I will raise my own objection to Goldman's theory which will show where it falls short as an analysis of doxastic justification. This shortcoming will provide the backdrop for the correct account of doxastic justification, which will also be given in section 5.

One caveat is in order regarding sections 2 and 3. Taking themselves to be discussing the same concept, both Goldman and his antagonists present their views using the unclarified, ambiguous term 'epistemic justification'. So, when delineating their positions in sections 2 and 3, I too will employ the ambiguous locution 'epistemic justification'. Also, unless otherwise noted, when the term 'justification' is used in sections 2 and 3, it is just an abbreviation of the more cumbersome 'epistemic justification'. The same applies when other forms of the word, e.g. 'justified', are used.

2. Reliabilism as Epistemic Justification

In "What Is Justified Belief?" Goldman attempts to provide a reductive analysis of the epistemic justification of beliefs in terms of the reliability of the cognitive processes and mechanisms which give rise to those beliefs. In his more recent works, most notably Epistemology and Cognition, Goldman's views concerning justification appear to have changed considerably, largely in response to

objections to his earlier reliability theory. His current approach accounts for justification in terms of a rule framework, where a belief is justified only if it is permitted by a right system of justification rules [J-rules]. Reliability still looms large in Goldman's present theory, but it enters in at a new level, viz. the level of the criterion of rightness for a system of J-rules. Accordingly, a system of J-rules is right iff the ratio of true beliefs to total beliefs, sanctioned by the system as a whole, is sufficiently high. Things are complicated further by the fact that J-rules per se do not sanction beliefs directly at all. Instead, J-rules sanction cognitive processes (and methods). So, a system S of J-rules is right iff the cognitive processes (and methods) which S sanctions have a sufficiently high truth-ratio in terms of the beliefs which they produce, and now the new theory begins to look suspiciously similar to the earlier theory it is intended to replace. By my lights, the major difference between the two theories is that the more recent and supposedly superior theory is rife with excess baggage, making it much more difficult to evaluate, especially since, lacking any J-rules with content, we are not in a position to assess their adequacy as such. My concern here, however, is not to raise detailed criticisms and/or objections to Goldman's "rule-framework reliabilism". Instead, as mentioned at the outset, my overriding concern in this chapter is to provide

an account of doxastic justification, and since I believe that Goldman's earlier theory comes close to doing just that, I shall focus the remainder of this section on his earlier view.

Goldman's stated desideratum in "What Is Justified Belief?" is to determine "a set of substantive conditions that specify when a belief is justified."[1] The conditions sought are to be reductive conditions, i.e. conditions which

---

[1]Alvin Goldman, "What Is Justified Belief?", Justification and Knowledge, ed. George Pappas (Dordrecht: D. Reidel Publishing Company, 1979), p. 1. Actually, a close examination of "What Is Justified Belief?" reveals two distinct desiderata. The first desideratum, as mentioned above, is to provide conditions which specify when a belief is actually justified. The second desideratum is to explain why we count certain beliefs as justified, or as Goldman puts it:

> What we really want is an explanation of why we count, or would count, certain beliefs as justified and others as unjustified. Such an explanation must refer to our beliefs about reliability, not to the actual facts. The reason we count beliefs as justified is that they are formed by what we believe to be reliable belief-forming processes. Our beliefs about which belief-forming processes are reliable may be erroneous, but that does not affect the adequacy of the explanation. Since we believe that wishful thinking is an unreliable belief-forming process, we regard beliefs formed by wishful thinking as unjustified. What matters, then, is what we believe about wishful thinking, not what is true (in the long run) about wishful thinking. [p. 18]

ultimately reduce the epistemic to the non-epistemic, and they are to be couched in a recursive format with one or more base clauses, a set of recursive clauses (if necessary), and a closure clause. The goal of providing reductive conditions for justified belief necessitates the following admissibility constraint on base clauses:

> (AC) A base clause B is admissible only if it is the case that no epistemic predicates appear in the antecedent of B.

It is also important to note that the criteria of justifiedness with which Goldman is concerned are semantic rather than epistemic in nature, i.e. they consist of truth-conditions for the schema |S's belief in p at time t

---

In my opinion, Goldman wavers between these two desiderata, because, not recognizing the distinction between personal and doxastic justification, he feels some intuitive pull in the direction of each of these kinds of justification. After all, we could hardly fault a cognizer for forming a belief using a process which, despite actually being unreliable, everyone herself included believes to be perfectly reliable. In fact, such a cognizer has formed her belief in an epistemically impeccable fashion. The intuition that such a cognizer is epistemically justified in so-forming a belief makes one feel that "counted justification" rather than "actual justification" is what matters in the epistemic realm.

Nevertheless, Goldman's overriding concern throughout the article is to provide an account of actually justified belief, the one passage above to the contrary notwithstanding, for he tells us at the outset that "On the account of justified belief suggested here, it is necessary for knowing, and closely related to it" [p. 1.], and surely

is justified | , not of conditions which let us know when these truth-conditions are met.[2]

Before putting forth his own base clause, Goldman examines various candidate base clauses, all of which prove unsuccessful. Goldman attributes this lack of success to their failure to provide requirements concerning how the belief is caused. The lesson to be learned is that an adequate base clause must stipulate causal requirements, "where 'cause' is construed broadly to include sustainers as

---

what is needed for knowledge is actual, and not merely counted, justification. For this reason, Goldman's "second desideratum" is best regarded as an unintentional and uncharacteristic slip into the personalist/internalist camp, which runs counter to everything else that he is trying to do in the article. Accordingly, I will interpret Goldman's reliabilist account as an attempt to satisfy his first and (pardon the pun) actual desideratum and will make no mention of the somewhat confused "second desideratum" in the body of section 2.

[2] Ibid., p. 3. What motivates this interest in semantic criteria is Goldman's commitment to the view that cognizers may lack "privileged access" to the justificational status of their beliefs, e.g. he maintains that young children [I would go further and include most adults] have justified beliefs without realizing that those beliefs are justified. [pp. 15 and 19] An adequate theory of justifiedness, contends Goldman, must account for such "nonpossessed" justification. If a cognizer can have justified beliefs without "possessing" a statable (or thinkable) justification for those beliefs, then their justification must be accounted for in some non-intellectualist fashion. Goldman accounts for the justifiedness of such beliefs in terms of justification-conferring processes, i.e. processes which, in giving rise to such beliefs, confer justification on them independent of the cognizer's efforts.

well as initiators of belief."[3] Of course, since all beliefs are produced by some sort of causal process or other, but not all beliefs are justified, it follows that only certain belief-forming causal processes confer justification on the beliefs which they produce. The question which arises immediately is, "What kinds of belief-forming causal processes do confer justifiedness on the beliefs to which they give rise?" To answer this question, it is useful to examine the sorts of beliefs which we generally regard as justified, some of which include ordinary perceptual beliefs, memory beliefs, sound inferential beliefs, and introspective beliefs. What feature do all these kinds of beliefs share in virtue of which we regard them as justified? They all share the feature of having been produced by belief-forming cognitive processes [BCP's] that are reliable, i.e. by BCP's which generally produce true beliefs. The answer to our question concerning what kinds of BCP's are justification-conferring is now straightforward. Reliable and only reliable BCP's are justification-conferring, for it is in virtue of their reliability that they possess their justification-conferring status. Goldman makes the point as follows:

---

[3] Ibid, p. 9.

> The justificational status of a belief is a function of the reliability of the process or processes that cause it, where (as a first approximation) reliability consists in the tendency of a process to produce beliefs that are true rather than false.[4]

This finding meshes well with the graded notion of justification, since the more reliable a BCP, the more justification that BCP confers on the beliefs which it produces. At the other end of the justificational spectrum, since unreliable BCP's (i.e. BCP's which tend to produce false beliefs) fail to confer justification on the beliefs that they produce, when beliefs are produced by unreliable BCP's like wishful thinking and confused reasoning, those beliefs are not justified.

Taking himself to have established that reliability is the defining feature of justification-conferring BCP's, Goldman returns to the task of specifying an adequate base clause. To set the stage for his first attempt at offering such a base clause, he draws two related distinctions. First, he distinguishes between two different kinds of reliability, viz. conditional reliability and unconditional reliability. The notion of unconditional reliability is that notion captured by his earlier provisional definition of reliability _per se_, to wit, a BCP is (unconditionally)

---

[4] _Ibid._, p. 10.

reliable if it tends to produce beliefs which are true rather than false. Concerning conditional reliability, he states, "A process is conditionally reliable when a sufficient proportion of its output-beliefs are true given that its input-beliefs are true."[5] Reasoning, say in accordance with modus ponens, is an example of a conditionally reliable BCP, for reasoning in accordance with modus ponens is a reliable guide to truth only if the premises (i.e. input-beliefs) from which one reasons are true.

This distinction between conditional and unconditional reliability leads Goldman to make a second distinction —— the distinction between belief-dependent and belief-independent BCP's. Belief-dependent BCP's are "processes some of whose inputs are belief-states."[6] Belief-independent BCP's are "processes none of whose inputs are belief-states."[7] These two kinds of BCP's correspond to and are interrelated with the two kinds of reliability that Goldman distinguishes, since conditional reliability is the kind of reliability applicable to belief-dependent BCP's and

---

[5] Ibid., p. 13.

[6] Ibid.

[7] Ibid.

unconditional reliability is the kind of reliability that attaches to belief-independent BCP's.

With these distinctions in hand Goldman makes his first attempt at offering a unified reductive theory of justifiedness. He begins by proffering the following base clause:

> (BC1) If S's belief in p at t results ('immediately') from a belief-independent process that is (unconditionally) reliable, then S's belief in p at t is justified.[8]

He couples (BC1) with the following recursive clause:

> (RC1) If S's belief in p at t results ('immediately') from a belief-dependent process that is (at least) conditionally reliable, and if the beliefs (if any) on which this process operates in producing S's belief in p at t are themselves justified, then S's belief in p at t is justified.[9]

Finally, by adding a standard closure clause to (BC1) and (RC1), Goldman completes his first approximation of a reductive theory of justified belief. Goldman summarizes

---

[8] Ibid.

[9] Ibid, p. 14. NOTE: (RC1) is admissible as a recursive clause, for it is permissible for epistemic predicates to appear in the antecedent of recursive clauses. The only place that epistemic predicates are not allowed to appear is in the antecedent of base clauses.

the gist of this theory as follows: "The theory says, in effect, that a belief is justified if and only if it is 'well-formed', i.e., it has an ancestry of reliable and/or conditionally reliable cognitive operations."[10]

The theory just sketched makes the justificational status of a belief exclusively a function of the reliability of the BCP's that produce it. Goldman notes, however, that such a theory is flawed since it will occasionally count as justified some beliefs which intuitively are not justified. To illustrate the point, he has us consider the following counterexample:

> Suppose that Jones is told on fully reliable authority that a certain class of his memory beliefs are almost all mistaken. His parents fabricate a wholly false story that Jones suffered from amnesia when he was seven but later developed pseudo-memories of that period. Though Jones listens to what his parents say and has excellent reason to trust them, he persists in believing the ostensible memories from his seven-year-old past. Are these memory beliefs justified? Intuitively, they are not justified. But since these beliefs result from genuine memory and original perceptions, which are adequately reliable processes, our theory says that these beliefs are justified.[11]

In re. this counterexample Goldman offers the following diagnosis:

---

[10] Ibid.

[11] Ibid., p. 18.

> Jones has strong evidence against certain propositions concerning his past. He doesn't use this evidence, but if he were to use it properly, he would stop believing these propositions. Now the proper use of evidence would be an instance of a (conditionally) reliable process. So what we can say about Jones is that he fails to use a certain (conditionally) reliable process that he could and should have used.[12]

And he concludes:

> This diagnosis suggests a fundamental change in our theory. The justificational status of a belief is not only a function of the cognitive processes actually employed in producing it; it is also a function of processes that could and should be employed.[13]

To accommodate this change in his theory, Goldman proposes the following base clause:

> (BC1') If S's belief in p at t results from a reliable cognitive process, and there is no reliable or conditionally reliable process available to S which, had it been used by S in addition to the process actually used, would have resulted in S's not believing p at t, then S's belief in p at t is justified.[14]

---

[12] Ibid., pp. 19-20.

[13] Ibid., p. 20.

[14] Ibid.

Though Goldman does not explicitly revise recursive clause (RC1), his new theory, according to which the justificational status of a belief is partly a function of cognitive processes which could and (epistemically) should be used, seems to require the following revised recursive clause:

> (RC1') If S's belief in p at t results from a conditionally reliable BCP, if the beliefs on which this BCP operates in producing S's belief in p at t are themselves justified, and if there is no reliable or conditionally reliable BCP available to S which, had it been used by S in addition to the BCP actually used, would have resulted in S's not believing p at t, then S's belief in p at t is justified.

Finally, by adding a standard closure clause to (BC1') and (RC1'), we get Goldman's complete theory of justified belief.[15]

Before considering some of the many objections to Goldman's theory, I want to conclude this section by examining three (potential) virtues of the theory. First, if Goldman's reliabilism (or a slight modification thereof)

---

[15]Whether or not Goldman actually held the theory I am attributing to him [to wit, (BC1'), (RC1'), and a closure clause] is somewhat beside the point, since he no longer espouses this theory. Nevertheless, for lack of a better term, I will refer to this theory as "Goldman's theory" and as "Goldman's reliabilism".

is correct, it has the virtue of evincing a successful reduction of the epistemic to the non-epistemic. This feature alone makes his theory worthy of careful consideration and scrutiny. A second virtue of Goldman's reliabilism is the apparent ease with which it avoids skeptical objections. Skeptical hypotheses, e.g. Descartes' evil demon hypothesis or the more contemporary malevolent neurophysiologist hypothesis, which seem devastating to purely internalist epistemologies, are readily handled by Goldman's reliabilism. After all, such hypotheses only serve to show that it is logically possible that our cognitive processes are unreliable, but if our cognitive processes are in fact reliable, then, skeptical hypotheses or no, on a reliabilist account the beliefs resulting from such processes are justified.[16] The third and perhaps most seductive feature of Goldman's reliabilism is the promise it holds for providing the sought after connection between a belief's truth and its justification in the following way: By definition, a BCP is reliable iff it generally produces true beliefs. But this is just to say that a BCP is reliable iff the indefinite probability of beliefs produced by it being true is high (at the very minimum, greater than

---

[16] A similar observation is made by Lehrer and Cohen. See their "Justification, Truth, and Coherence", Synthese, Vol. 55 (1983), p. 192.

.5). Since Goldman's theory basically asserts that a belief is justified iff it is produced by a reliable BCP, and since, by definition, beliefs produced by reliable BCP's have a high indefinite probability of being true, it follows that justified beliefs have a high indefinite probability of being true. It should be clear from the context that indefinite probabilities are dyadic relations relating classes or properties by specifying the probability of a member of one class being a member of another class. Consequently, by proving that Goldman's reliabilism entails that justified beliefs have a high indefinite probability of being true, we have ipso facto proved that Goldman's reliabilism entails that beliefs belonging to the class of justified belief have a high probability of belonging to the class of true belief. Thus, Goldman's reliabilism affixes a probabilistic connection between justification and truth, and in light of this probabilistic truth connection, it potentially constitutes a correct analysis of the knowledge conception of epistemic justification.

Since these virtues are, of course, contingent on the correctness of Goldman's reliabilism, we need to determine whether it really does provide an accurate analysis of epistemic justification. In order to make such a determination, we need to assess whether and to what extent reliabilism can stand up to the objections and purported counterexamples vied against it. Naturally, the first step

in making this assessment is to look at the objections themselves. Accordingly, the task of the next section is to present several of these objections.

3. Counterexamples to Goldman's Theory

Goldman's brand of reliabilism has incurred objections from numerous epistemologists, to wit, Bonjour, Chisholm, Cohen, Feldman, Lehrer, and Pollock, to name a few.[17] In this section, however, I shall focus exclusively on the objections raised by Lehrer, Cohen, and Bonjour. My doing so should not be taken to suggest that the other objections are of little or no importance. To the contrary, some of them pose extremely difficult problems for reliabilism, but since they have no direct bearing on the goal of this chapter, which is to provide an account of doxastic justification, I have elected to save their discussion for future papers, where I will be able to deal with them in the detail they deserve.

---

[17] Laurence Bonjour, "Externalist Theories of Empirical Knowledge", Midwest Studies in Philosophy, ed. French, Vehling, and Wettstein (University of Minnesota: 1980); Roderick Chisholm, "The Place of Epistemic Justification", Philosophical Topics, Vol. 14, No. 1, (Spring, 1986); Stewart Cohen, in his dissertation Justification and Truth, Chapter 1, and in his "Justification and Truth", Philosophical Studies, Vol. 46 (1984); Richard Feldman, "Reliability and Justification", The Monist, Vol. 68, No. 2 (April, 1985); and John Pollock, in his Contemporary Theories of Knowledge (in manuscript), and in his "Reliability and Justified Belief", Canadian Journal of Philosophy, Vol. XIV, No. 1 (March, 1984).

Lehrer, Cohen, and Bonjour raise a host of astute objections to Goldman's theory. Some of these objections are rooted in purely intuitive considerations, while others are based on logical grounds. My concern here is to present the intuitive objections that they have raised against reliabilism to the effect that a belief's being produced by a reliable BCP is neither necessary nor sufficient for that belief's being epistemically justified. I will begin with their attacks on reliabilism's necessity.

In order to provide a counterexample to the necessity of Goldman's reliabilism, one needs to present a case where intuitively a belief is justified even though that belief was produced by an unreliable BCP. The standard counterexample to reliabilism's necessity runs as follows: Consider a possible world W where unbeknownst to us the evil demon hypothesis is true. In such a world virtually all of our beliefs turn out to be false owing, of course, to the malevolent machinations of the demon. Moreover, the BCP's (e.g. perception, memory, and inference) which produced these beliefs are unreliable in W, since they tend to produce false beliefs in W. Lehrer and Cohen rightly note, "It would follow on reliabilist views that under such conditions the beliefs generated by those processes would

not be justified."[18] However, they maintain that this result is unacceptable, since

> The truth of the demon hypothesis also entails that our experiences and our reasonings are just what they would be if our cognitive processes were reliable, and, therefore, that we would be just as well justified in believing what we do if the demon hypothesis were true as if it were false.[19]

Their point is worth belaboring. Our experiences and reasonings in W are, by hypothesis, phenomenologically indistinguishable from the experiences and reasonings we would have in a verific world W* where we would indeed be justified in holding the beliefs we do. But since our justification for our beliefs, viz. our experiences and reasonings, is exactly the same in both worlds, intuitively we are just as justified in holding the beliefs we do in W as we are in W*. Of course, since intuitively we are justified in holding our beliefs in W despite the fact that they have all been produced by unreliable BCP's (the demon has seen to that), it follows that being produced by a reliable BCP is not necessary for a belief to be epistemically justified.

---

[18] Keith Lehrer and Stewart Cohen, "Justification, Truth, and Coherence", Synthese, Vol. 55 (1983), p. 192.

[19] Ibid.

In his dissertation Justification and Truth, Cohen offers an even more compelling argument against reliabilism's necessity. Once again we are to consider a demon-manipulated world, say W'. We are then asked to imagine two inhabitants of W':

> A who is a good reasoner, i.e., reasons in accordance with the canons of inductive inference, and B who engages in confused reasoning, wishful thinking, reliance on emotional attachments, guesswork, etc.[20]

As is the case in worlds such as W', unbeknownst to our two hapless inhabitants, the demon sees to it that BCP's like reasoning in accordance with the canons of inductive inference are just as unreliable as BCP's like wishful thinking and confused reasoning. Now as we know from section 2, Goldman maintains that reliability is the defining feature of justification-conferring BCP's, from which it follows that in W' the unreliable BCP of reasoning according to the laws of inductive logic is just as non-justification-conferring as the equally unreliable BCP's of wishful thinking and confused reasoning. Accordingly, Cohen rightfully notes:

> Since the beliefs of A & B are both produced by unreliable processes (the evil demon sees to this), a reliabilist theory of

---

[20]Stewart Cohen, Justification and Truth, op. cit., p. 10.

> justification must render identical epistemic appraisals of both sets of beliefs.[21]

To wit, a reliabilist theory must maintain that neither A's beliefs nor B's beliefs are justified in W'. But, Cohen asserts,

> Plainly, This cannot be correct. A's beliefs are conditioned by the evidence whereas B's beliefs are not. A is a good reasoner whereas B is not. A's beliefs are reasonable whereas B's beliefs are not. There is a fundamental epistemic difference between the beliefs of A and the beliefs of B. But the Reliabilist does not have the theoretical means to display this difference.[22]

Cohen maintains that the fundamental epistemic difference between A's beliefs and B's beliefs is that A's beliefs are justified whereas B's beliefs are not, for in his estimation, "Beliefs produced by good reasoning are paradigm cases of justified belief and beliefs arrived at through fallacious or arbitrary reasoning are paradigm cases of unjustified belief."[23] Again, since A's unreliably produced beliefs are intuitively justified, reliability is not necessary for epistemic justification.

---

[21] Ibid., pp. 10-11.

[22] Ibid., p. 11.

[23] Ibid.

Now bordering on overkill, Cohen offers yet another counterexample to reliability as a necessary condition for justification. The counterexample proceeds by again contrasting A and B in W'. This time the belief in question is the non-discursive perceptual belief, which both A and B have, that there is something ϕ before them. Both A and B hold this belief on the basis of being appeared to ϕ-ly, but the epistemically relevant difference between them is that "While A has no evidence to the contrary, B is presented with strong evidence that owing to a clever deception there is nothing ϕ before him."[24] Again, Cohen contends that there is a clear epistemic difference between A's perceptual belief and B's perceptual belief, because, as he puts it,

> from an epistemic point of view, B ought not to have proceeded in the way he did. We might say that contrary to A, B has been epistemically irresponsible in accepting that there is something before him. As a result, while A is justified in his ϕ perceptual belief, B is not.[25]

Of course, since in W', A's perceptual belief is intuitively justified even though his perceptual faculties are just as unreliable as B's, it follows that being reliably produced

---

[24] Ibid., p. 13.

[25] Ibid., p. 14.

is not a necessary condition for a belief's being epistemically justified.

Fairly confident that necessity has fallen by the wayside, Cohen sets his sights on demonstrating that reliabilism fails as a sufficient condition for epistemic justification, as well. He begins by reminding us that

> The recipe for finding a counter-example to reliability as a sufficient condition for justification is to take an intuitively unjustified process (that is a process that intuitively does not produce justified beliefs) and suppose that it were reliable.[26]

The case he discusses is one which Goldman himself raises in "What Is Justified Belief?". We are to imagine a possible world W^+ where a benevolent demon arranges things such that the vast majority of beliefs arrived at by wishful thinking are true in W^+. As a result, wishful thinking is a reliable BCP in W^+. "Thus," Cohen notes, "on Goldman's view, it turns out that such beliefs are justified."[27] To fill out the example, we are to assume that the inhabitants of W^+ are unaware that wishful thinking is reliable in their world. Since they are unaware that wishful thinking is a reliable BCP, Cohen maintains that their wishfully formed, completely

---

[26] Ibid.

[27] Ibid.

reliably formed, beliefs are not justified, reliabilism to the contrary notwithstanding. As he puts it,

> The crucial factor, what by my lights makes the beliefs unjustified in these cases, is the fact that the reliability of the belief forming process is due to facts that are completely outside the ken of the subject. If as far as the subject knows, the state of affairs expressed by P is merely something he wishes for, then he is being epistemically irresponsible in accepting that P.[28]

Thus, these wishfully formed beliefs are intuitively unjustified despite being reliably produced, and consequently, reliabilism is not sufficient for epistemic justification.

To bolster this conclusion, Cohen has us consider our earlier subjects A and B who this time fortunately find themselves in benevolent world W^{+}. Again, A is a good reasoner who reasons in accordance with the laws of induction, whereas B acquires his beliefs via wishful thinking. Thanks to the benevolent demon, both A's BCP and B's BCP are extremely reliable. But, Cohen attests, "If one adheres to the position that reliability is a sufficient condition of justification, then one must give the same epistemic appraisal to the beliefs of A and B."[29] Namely,

---

[28] Ibid., p. 16.

[29] Ibid., p. 20.

on the reliabilist account, both their sets of beliefs turn out to be justified, but since intuitively B's beliefs, unlike A's beliefs, are not justified, reliable production is not a sufficient condition for epistemic justification. And Cohen is not alone in his doubts about reliabilism's sufficiency.[30]

In "Externalist Theories of Empirical Knowledge", Bonjour presents several cases which purportedly demonstrate reliabilism's insufficiency for epistemic justification. In one particularly compelling counterexample, we are asked to consider the following case:

> Suppose that Norman, under certain conditions that usually obtain, is a completely reliable clairvoyant with respect to certain kinds of subject matter. He possesses no evidence or reasons of any kind for or against the general possibility of such a cognitive power, or for or against the thesis that he possesses it. One day Norman comes to believe that the President is in New York City, though he has no evidence either for or against this belief. In fact the belief is true and results from his clairvoyant power, under circumstances in which it is completely reliable.[31]

---

[30]In addition to the Bonjour article which I am about to discuss, see Alvin Plantinga's "Justification and Theism" (in manuscript), pp. 52-53.

[31]Laurence Bonjour, "Externalist Theories of Empirical Knowledge", op. cit., p. 62.

Bonjour maintains that reliabilism entails that Norman's completely reliably produced belief about the President's whereabouts is justified.[32] Hoping to convince us that such a result is intuitively unacceptable, he asks us:

> Is Norman epistemically justified in believing that the President is in New York City, so that his belief is an instance of knowledge? According to the modified externalist position, we must apparently say that he is. But is this the right result? Are there not still sufficient grounds for a charge of subjective irrationality to prevent Norman's being epistemically justified?[33]

Surely, the intuitive answer to this last question is "Yes", and so, once again, we are presented with a case where a completely reliably produced belief is intuitively unjustified, thereby demonstrating that reliable production is not sufficient for epistemic justification. Fortunately for reliabilism, the next section will show that these

---

[32] In point of fact, there appear to be other reliable BCP's available to Norman which, had he used them, would have prevented him from forming the belief that the President is in New York City. Hence, the counterfactual clause in Goldman's (BC1') is unsatisfied, and so, on Goldman's view, Norman's belief is not justified. Since I shall argue in section 5 that it is a mistake on Goldman's part to include the counterfactual clause in (BC1'), I am going to ignore the counterfactual wrinkle of Goldman's theory and simply assume with Bonjour that reliabilism entails that Norman's clairvoyant belief is justified.

[33] Laurence Bonjour, "Externalist Theories of Empirical Knowledge", op. cit., p. 62.

seemingly decisive counterexamples are not as devastating as they prima facie appear.

4. Sorting Things Out

a. Goldman Replies

Goldman has offered two different responses to the objections against necessity.[34] The first, found in Epistemology and Cognition, is an attempt to accommodate within a reliabilist framework the overriding intuition that our demon-world-inhabitant's beliefs are in fact justified, despite being produced by BCP's which are unreliable in that world. The second and more recent reply consists of distinguishing two types of justification, strong and weak, and arguing that our demon-world-inhabitant's beliefs are justified only in the latter sense. In the remainder of this subsection, I will present these two responses in detail. In section 5, I will argue that neither response is satisfactory.

In Epistemology and Cognition Goldman responds to the objections against necessity by first embracing the intuition that the well-reasoned beliefs of demon-manipulated cognizers are in fact justified and then

---

[34]To my knowledge, Goldman has not responded to the objections against reliabilism's sufficiency.

arguing that the brand of reliabilism which he espouses —— what I will call "normal worlds reliabilism" —— accommodates this intuition. To fully understand the nature of this response, some background into normal worlds reliabilism is needed. Recall from section 2 that reliabilism maintains that a belief is justified only if it results from a justification-conferring BCP, where a BCP has the property of being justification-conferring just in case it is reliable, i.e. it tends to produce true beliefs. Expanding on this view, normal worlds reliabilism maintains that the property of being justification-conferring is a necessary property of those BCP's which possess it; which is to say, if a BCP is justification-conferring, then it is necessarily justification-conferring.[35] Accordingly, if the BCP good reasoning is justification-conferring, then good reasoning is justification-conferring in every possible world in which good reasoning occurs. But good reasoning is not reliable in every possible world in which it occurs —— the evil demon world is a case in point. So, how are we to understand the reliabilist claim that the justification-conferring status of a BCP is a function of that BCP's

---

[35] x is necessarily ϕ iff x is ϕ in every possible world in which x exists. Accordingly, a BCP is necessarily justification-conferring just in case it is justification-conferring in every possible world in which it exists.

reliability? According to normal worlds reliabilism, a BCP's justification-conferring status in world W is not determined by its reliability in W. Instead, a BCP's justification-conferring status in world W is a function of its reliability in normal worlds. Goldman uses 'normal worlds' in a technical sense which he explains as follows:

> We have a large set of common beliefs about the actual world: general beliefs about the sorts of objects, events, and changes that occur in it. We have beliefs about the kinds of things that, realistically, do and can happen. Our beliefs on this score generate what I shall call the set of normal worlds.[36]

Simply put, a normal world is one where things are as we think them to be in the actual world.

Since we think that perception, memory, and good reasoning are reliable in the actual world, they are, by definition, reliable in normal worlds (for, as we have just seen, normal worlds are defined by what we think is true in the actual world). Since perception, memory, and good reasoning are reliable in normal worlds, it follows, on the normal worlds reliabilism view, that perception, memory, and good reasoning are justification-conferring in all

---

[36]Alvin Goldman, Epistemology and Cognition (Cambridge, Mass. and London: Harvard University Press, 1986), p. 107.

possible worlds where they occur, and that, of course, includes evil demon worlds.

Now Goldman's reply to the necessity counterexamples is straightforward:

> The justificational status of a W-world belief does not depend on the reliability of the causing processes in W. Rather, it depends on the reliability of the processes in normal worlds. Now an evil demon world is a paradigm case of a non-normal world. So it does not matter that the processes in question are highly unreliable in that world. It only matters whether they are reliable in normal worlds, and that apparently is the case.[37]

Thus, according to normal worlds reliabilism, since in the first counterexample to necessity we use perception, memory, and inference in forming our beliefs in evil demon world W, our beliefs are justified —— just as Lehrer and Cohen contend —— because perception, memory, and inference are reliable in normal worlds. Normal worlds reliabilism handles the other objections against necessity, as well. For according to normal worlds reliabilism, A's beliefs are justified in W' while B's beliefs are not —— just as Cohen maintains —— because the BCP's which A uses to form his beliefs in W' are reliable in normal worlds, where as the BCP's which B uses to form his beliefs in W' are unreliable

---

[37] Ibid., p. 113.

in normal worlds. Thus, reliabilism in the form of normal worlds reliabilism does possess the theoretical means necessary for displaying the different justificatory statuses had by A's and B's beliefs, Cohen's claims to the contrary notwithstanding.

Despite its initial attractiveness in handling evil demon cases, normal worlds reliabilism has been the brunt of so many telling criticisms that Goldman has since decided to abandon the view. In its stead, he introduces a duplex theory of justification which supposedly accords with and accounts for the divergent intuitions that arise when evaluating the justificatory statuses of the beliefs of evil demon world inhabitants. As the name suggests, the duplex theory introduces two conceptions of epistemic justification, one strong, the other weak. To motivate this distinction between strong and weak justification, Goldman presents us with the following case:

> Consider a scientifically benighted culture, of ancient or medieval vintage. This culture employs certain highly unreliable methods for forming beliefs about the future and the unobserved. Their methods appeal to the doctrine of signatures, to astrology, and to oracles. Members of the culture have never thought of probability theory or statistics, never dreamt of anything that could be classed as 'experimental method'. Now suppose that on a particular occasion a member of this culture forms a belief about the outcome of an impending battle by using one of the aforementioned methods, say, by consulting zodiacal signs in a culturally approved fashion. Call this method M. Is

> this person's belief justified, or warranted?[38]

Goldman maintains that in attempting to answer this question we are naturally drawn in two different directions. On the one hand, consulting zodiacal signs is highly unreliable and is, consequently, a very poor way to form beliefs. Moreover, it is natural to regard beliefs formed by poor or inadequate methods as unjustified, which suggests that the belief of our scientifically benighted cognizer, S*, is unjustified. On the other hand, given the plight of living in a wholly unscientific culture, S* has done the best he could in forming his belief. He has used a method which is highly regarded by the members of his community, a method for which he can find no reason to doubt. Given his epistemic situation, we cannot fault him for using method M nor for forming the belief he does. Since his belief is epistemically <u>blameless</u>, we are inclined to say that it is justified, after all.

Which of these two views is correct? The duplex theory acknowledges the legitimacy of each of these epistemic evaluations, claiming that they simply embody different conceptions of justification. On the first conception, Goldman tells us, "a justified belief is (roughly) a

---

[38]Alvin Goldman, "Strong and Weak Justification" (in manuscript), pp. 1-2.

well-formed belief, a belief formed (or sustained) by proper, suitable, or adequate methods, procedures, or processes."[39] On the second, he observes, "a justified belief is a faultless, blameless, or non-culpable belief."[40] Goldman refers to these two conceptions of justification as "strong" and "weak", respectively.

Having thus provided an intuitive case for the strong/weak justification distinction, Goldman attempts to delineate the conditions for strong and weak justification, respectively.[41] Certain subtleties aside (see footnote 41), Goldman maintains that:

(SJ) A belief of person S is strongly justified iff

---

[39] Ibid., p. 3.

[40] Ibid.

[41] Actually, Goldman distinguishes between two levels of justifiedness: primary justifiedness which is justifiedness at the level of cognitive processes and secondary justifiedness which is justifiedness at the level of methods; and he suggests that the strong/weak distinction enters in at each level. [see his "Strong and Weak Justification", op. cit., pp. 3-4.] For a belief to be fully justified, it must be strongly justified at both the primary and secondary levels. So, to simplify his account somewhat, I will combine both levels and provide composite conditions for strong justification (at both levels simultaneously) and for weak justification (at both levels simultaneously). This simplified account remains true to the spirit of Goldman's more complicated theory, and it does not alter, in any essential way, his most recent reply to the counterexamples against necessity.

> (1) it is produced (or sustained) by proper methods and/or processes, where a method or process is proper *just in case* it is reliable,
> (2) the methods (if any) used in producing S's belief have been *acquired* in a suitable way, where suitable method acquisition requires being acquired by other methods or processes that are reliable or metareliable, and
> (3) S's cognitive state at the time the belief is formed does not *undermine* the properness of the methods and/or processes employed, where the properness of a method or process is undermined *just in case* either
> (a) S (mistakenly) *believes* the method or process to be unreliable, or
> (b) S is *justified* in regarding the method or process as unreliable.[42]

Accordingly, the core idea of strong justification can be captured as follows: S's belief that p is strongly justified *just in case* (1) it is produced by processes and/or suitably acquired methods that are reliable, and (2) S's cognitive state when the belief that p is formed does not undermine these processes' and/or methods' reliability.

Goldman notes that beliefs which satisfy these conditions for strong justification will (presumably) also be blameless and, hence, would appear to be weakly

---

[42]Alvin Goldman, "Strong and Weak Justification", *op. cit.*, pp. 4-7 and 11.

justified, as well. But he wants strong and weak justification to be mutually exclusive, opposing notions. So, he modifies the notion of weak justification in such a way that it only attaches to merely blameless beliefs, i.e. ill-formed (strongly unjustified) but blameless beliefs, rather than to beliefs that are blameless per se. To capture this ill-formed-but-blameless sense of weak justification, Goldman provides us with the following set of jointly sufficient (though non-necessary) conditions for such justification:

> (WJ) S's belief in p is weakly justified if
> (1) the method M (or cognitive process C) by which the belief is produced is unreliable, but
> (2) S does not believe that M (or C) is unreliable, and
> (3) S neither possesses, nor has available to him/her, a reliable way of telling that M (or C) is unreliable.[43]

Goldman goes on to suggest that we might need to supplement these three conditions with a fourth condition, to wit, (4) there is no process or method S believes to be reliable which, if used, would lead S to believe that M (or C) is unreliable.[44]

---

[43] Ibid., pp. 8 and 11. Again, this is an oversimplification of Goldman's view in that it does not distinguish between weak justification at the level of primary justifiedness and weak justification at the level of secondary justifiedness.

[44] Ibid.

With the conditions for strong and weak justification before us, let us return to the case of the scientifically benighted cognizer S*. S*'s belief about the outcome of the impending battle is strongly unjustified because it does not satisfy condition (1) of (SJ), since the method S* uses —— that of consulting zodiacal signs —— being quite unreliable, is not a proper method. Nevertheless, S*'s belief about the battle outcome is weakly justified because, although consulting zodiacal signs is unreliable, S* does not believe that is it unreliable, nor does he possess a reliable method or process which would lead him to think that it is unreliable. Thus, S*'s belief satisfies the three conditions of (WJ). Moreover, there is no method or process, which S* believes to be reliable, that would lead him to think that consulting zodiacal signs is unreliable; so, supplementary condition (4) is satisfied, as well.

We are now in a position to see how the duplex theory of justification handles evil demon world counterexamples. It should be fairly obvious that cognizers in a demon-manipulated world are in a situation not unlike S*'s. We demon-manipulated cognizers of the first counterexample form our beliefs using perception, memory, and inference, all of which the demon has rendered unreliable in W. Consequently, our beliefs are strongly unjustified since they fail to satisfy condition (1) of (SJ). Even so, there is a sense in which our beliefs are justified —— as Lehrer

and Cohen maintain. Our demon-manipulated beliefs are weakly justified because, even though they are produced by unreliable processes, we do not believe that perception, memory and inference are unreliable, we do not possess a reliable method or process which would lead us to think that they are unreliable, and there is no method or process that we believe to be reliable which, if used, would lead us to believe that they are unreliable. Thus, on the duplex view, Lehrer and Cohen are right in maintaining that our demon-manipulated beliefs are justified, though the only kind of justification they possess is weak justification.

Finally, Goldman's duplex theory of justification is also capable of handling the other objections against necessity. Regarding strong justification, the duplex theory renders identical assessments of A's beliefs and B's beliefs. Both A's beliefs and B's beliefs are strongly unjustified since they are produced by demon-rendered unreliable processes. Nevertheless, on the duplex view, there remains a marked epistemic difference between the beliefs of A and the beliefs of B. The beliefs of A, although ill-formed, are weakly justified, whereas the beliefs of B, also ill-formed, are not even weakly justified. Accordingly, while A's beliefs are blameless, we can fault B for his beliefs. Thus, Goldman's duplex theory provides the means for rendering the different epistemic

assessments of A's beliefs and B's beliefs, which Cohen's counterexamples require.

b. Personal and Doxastic Justification Revisited

As mentioned at the outset of the previous subsection, I find both of Goldman's responses to the necessity counterexamples unsatisfactory. In section 5, I will explain why I take them to be unsatisfactory. But first, in subsection c, I will offer what is by my lights the proper diagnosis of why the objections presented in section 3 are unsuccessful in refuting reliabilism. Since both the explanation and the diagnosis depend on and are rooted in the personal/doxastic justification distinction, some additional clarificatory remarks concerning these two kinds of justification are in order. The burden of the present subsection is to provide these clarificatory remarks.

Recall from Chapter 1 that according to (DJ), "Doxastic justification is a normative notion in terms of which beliefs are evaluated from the epistemic point of view." So stated, (DJ) does no more than identify doxastic justification as the kind of justification which attaches to beliefs. Not wanting to beg any questions, I formulated (DJ) in such a way that it remains an open question as to when a belief is doxastically justified. Of course, for doxastic justification to be a usable notion, we need at least some idea of when a belief possesses it, that is to

say, we need an answer to the question, "When is a belief doxastically justified?" A trivial, though not wholly uninformative, answer is: A belief is doxastically justified *just in case* it has positive epistemic status. This suggests that to get a less trivial answer to our question, we need to recast it in a new light, viz. "When, from the epistemic point of view, should a belief be evaluated positively?" Since the epistemic point of view is defined by the goal of maximizing truth and minimizing falsity in a large body of beliefs, one might think that the answer to this latter question simply is:

(A1) A belief has positive epistemic status *iff* it is true.

After all, from the epistemic viewpoint, true beliefs are better than false ones. But an answer like (A1) is essentially nothing more than a restatement of the epistemic goal itself. What we want is a way of evaluating beliefs apart from their actual truth-value that will help us to attain our dual-pronged goal of gaining truth and avoiding error.

Probability provides us with such a means of evaluating beliefs. From the epistemic viewpoint, beliefs that are probably true are better than beliefs that are probably false. Accordingly, a natural answer to our question concerning when a belief has positive epistemic status is:

(A2) A belief has positive epistemic status iff it has a sufficiently high probability of being true.

(A2), together with our earlier trivial observation that a belief is doxastically justified just in case it has positive epistemic status, entails the following nontrivial result:

(DJ') A belief is doxastically justified iff it has a sufficiently high probability of being true.

Intuitively, (DJ') seems right to me. It captures what, in my opinion, is the central idea behind doxastic justification by correctly describing which beliefs we would want to count as justified from the epistemic standpoint. That it does so can be seen as follows: When we say a belief is justified, we are appraising that belief positively as being one that is good to hold for the purpose of gaining truth and avoiding error. Now according to (DJ'), highly probable beliefs are doxastically justified, which is to say highly probable beliefs are good ones to hold for the purposes of gaining truth and avoiding error. This seems correct, since highly probable beliefs are, intuitively, the sorts of beliefs best suited for maximizing truth and minimizing falsity within our doxastic corpus. It is important to note that (DJ') only provides us with a semantic criterion of doxastic justifiedness, i.e. it does

no more than describe which beliefs are doxastically justified, i.e. are good ones to hold. It offers no epistemic criterion by which we can tell which beliefs satisfy this semantic criterion. That is to say, (DJ') does not provide us with a means for telling which beliefs are probably true and, hence, good ones to hold. The reason being, a belief can be a good one to hold without our knowing that it is a good one to hold. More to the point, a belief can be doxastically justified without our knowing (or even being able to tell) that it is doxastically justified.

Admittedly, these last remarks have an externalist flavor which those of internalist tastes may find objectionable. They may, for example, object that I have formulated (DJ') in an ad hoc fashion. After all, I have suggested that a form of reliabilism not unlike Goldman's reliabilism provides the correct account of doxastic justification. And I have now just given a wholly externalistic definition of doxastic justification. So, it is hardly surprising that reliabilism, an externalist theory, can account for doxastic justification so-defined. But, surely, defining a notion in such a way that one's chosen theory can account for it is as ad hoc as one can get.

Despite such circumstantial evidence, I plead innocent to this ad hoc-ness charge. In my defense, recall that in Chapter 1 I argued that there are several different senses

of justification in the epistemic realm, all of which currently are batted around under the single heading "epistemic justification". Rather than being seen as offering some esoteric, ad hoc conception of justification, (DJ') should be viewed as distilling out one of the various senses of epistemic justification currently in use. After all, the idea that probability and justification are intimately connected is hardly novel. Moreover, an entire school of epistemology, viz. probabilism, can be seen as an attempt to capture this idea by providing a probabilistic theory of justification. In its simplest form, probabilism maintains that a belief is justified iff it is highly probable. Pollock points out that this contention, which he calls "the simple rule", has been endorsed by the likes of Chisholm, Hempel, Kyburg, Jeffery, Carnap, et. al.[45] (DJ') virtually restates the simple rule. However, (DJ') and the simple rule do differ in that the simple rule employs the unclarified, ambiguous term 'justified', whereas (DJ') is concerned with a specific sense of justification, to wit, doxastic justification. What (DJ') does, in effect, is capture the sense of justification indigenous to probabilism while acknowledging that it is only one of several senses of epistemic justification. Thus, far from

---

[45]John Pollock, Contemporary Theories of Knowledge, op. cit., p. 116 in manuscript.

being an ad hoc contrivance, (DJ') isolates one conception of epistemic justification already extant in the epistemological literature.

Before turning to personal justification, a digression into the motivation behind doxastic justification is in order. To set the stage for this digression, recall from Chapter 1 that the knowledge conception of epistemic justification requires that justification be conceptually connected with truth. To be adequate, an account of the knowledge conception of justification must specify the nature of this conceptual connection. Infallibilist theories of justification embody the strongest connection possible between justification and truth. They maintain that there is a necessary connection between justification and truth such that justification logically entails truth. The problem with such theories is that they not only lead to skepticism, they entail it, since they make human knowledge logically impossible.[46] Given this untenable result, most

---

[46]Infallibilist theories make human knowledge logically impossible because they make justified belief, a necessary condition for knowledge, logically impossible. That they make justified belief logically impossible can be seen as follows: Infallibilist theories require that the conditions which make a belief justified entail that belief's truth. But the evil demon hypothesis demonstrates that no matter what conditions for justified belief one settles on, it is always logically possible for the conditions to obtain and that the belief be false. What this shows is that there are no conditions for justified belief which entail truth.

epistemologists have found infallibilist theories, as well as their assumption that justification and truth are necessarily connected, to be unacceptable.

The demise of infallibilism has led most epistemologists to adopt fallibilist theories of justification. Fallibilist theories admit and maintain that no matter how well a belief is justified, it is still logically possible for that belief to be false (assuming the proposition believed is not a necessary truth). The task facing fallibilists is to make sense of the connection between justification and truth, given this possibility of justified-but-false belief. It seems that the most promising way to do so is probabilistically, for if justification and truth are probabilistically connected, then justified beliefs, though possibly false, have the virtue of being more probable than unjustified beliefs. The probabilists were driven to the simple rule in an attempt

---

Consequently, it is logically impossible for beliefs to possess the kind of truth-entailing justification that infallibilist theories require.

My discussion so far has ignored necessary truths and the cogito, since whether or not these beliefs are infallibly justified remains to be decided and since, even if they turn out to be infallibly justified and hence capable of being known, all other human knowledge of contingent propositions remains logically impossible on an infallibilist theory. Consequently, the possibility of cogito knowledge and/or knowledge of necessary truths does little, if anything, to make infallibilist theories more tenable.

to provide such a probabilistic connection between justification and truth. Moreover, they regarded the simple rule as an analysis of the concept "justified belief". It is on this last point that I part company with the probabilists.

I am not offering (DJ') as an analysis of the concept "doxastic justification", since one cannot analyze a concept before one knows what that concept is. Rather than analyzing doxastic justification, (DJ') is specifying what I take the concept "doxastic justification" to be, for this is the concept which I feel is in need of analysis. With the doxastic conception of justification provided via (DJ') before us, we are in a position to see the motivation behind such a conception. Simply put, it is to have a working conception of justification that is internally connected with truth. Digression ended.

Having provided a working conception of doxastic justification, I now turn to the topic of personal justification. Since my discussion of personal justification may seem overly brief and incomplete, let me preface it with an explanation for its brevity and incompleteness. Chapter 3 is devoted entirely to a discussion of the personal conception of justification. There, expanding upon the comments made below, I will clarify the concept of personal justification fully and will also offer a detailed analysis of this kind of

justification. Not wanting Chapter 3 to be an exercise in repetition, I am purposefully limiting the discussion offered here to that which is most essential for understanding the personal conception of justification. Accordingly, my present goal is simply to give the reader some idea of in what the notion of personal justification consists — or, if you will, to give the reader a feel for the notion of personal justification. Now, on to the discussion, lest my explanation for its brevity exceed it in length.

The expression 'Person S is personally justified in believing that p' is ambiguous. On one reading, it implies that S does, in fact, believe that p and is personally justified in doing so. On a second reading, it asserts that S does not believe that p, but that (given her present cognitive state) she would (or at least could) be personally justified in believing that p, were she to do so.[47] For

---

[47]Both Goldman and Pollock have noted roughly this same sort of ambiguity, though not in the context of personal justification per se. Goldman distinguishes between two uses of the term 'justified', an ex post use and an ex ante use, which correspond respectively to the two readings I distinguish above. [See his "What is Justified Belief?", op. cit., p. 21.]. Pollock distinguishes between justified belief and justifiable belief. "A justifiable belief," he tells us, "is one the believer could become justified in believing if he just put together what he already believes in the right way." [from his Contemporary Theories of Knowledge, op. cit., p. 90.]. Accordingly, 'justifiable belief' applies only when either

example, we might want to say that S is personally justified in believing that p — in the latter sense — when S has adequate evidence for the belief that p, but has not yet come to believe that p. In order the keep these two senses of S's being personally justified in believing that p separate, one might want to adopt the following stipulative terminology: Let the locution 'S is personally justified in believing that p' be used exclusively for the former sense that gives existential import to the belief in question, and let the locution 'S is ex ante personally justified in believing that p' be used to capture the second sense where S does not yet hold the belief that p. Throughout the course of this dissertation when discussing the justificatory status of persons, I will confine myself to the topic of personal justification, since, unlike ex ante personal justification, it has a role to play in the theory of knowledge.[48]

---

(1) S does not yet hold the belief or (2) the belief, which S does hold, is unjustified, because of the way she came to hold it. 'Justified belief' applies only when S does hold the belief. Thus, "justified belief" is like my first reading in that it entails the belief's existence. However, "justifiable belief" differs slightly from my second reading in that "justifiable belief" does not entail the nonexistence of the belief in question.

[48] Ex ante personal justification cannot play a role in an account of knowledge because it entails nonbelief. In entailing nonbelief, it ipso facto entails that a necessary condition for knowledge, viz. belief, is not satisfied.

According to our initial characterization (PJ), personal justification is a normative notion in terms of which persons are evaluated from the epistemic viewpoint. In order to flesh out (PJ) and thereby get a hold on the concept of personal justification, we need to reflect on just what it is we are doing when we make personal justification evaluations. Let us, therefore, start with the obvious. When we evaluate a person S as being personally justified in believing that p, we are evaluating S positively from the epistemic viewpoint, and when we evaluate S as being personally unjustified in believing that p, we are evaluating S negatively from that same viewpoint. Of course, in evaluating S positively from the epistemic viewpoint, we are, in effect, praising S epistemically. Similarly, in making the negative evaluation that S is personally unjustified in believing that p, we are blaming S epistemically for believing that p. Consequently, at least as a first approximation, we can define personal justification in terms of epistemic praiseworthiness and blameworthiness as follows:

(PJ_j) S is personally justified in believing that p iff S is worthy of epistemic praise for believing that p.

(PJ_u) S is personally unjustified in believing that p iff S is deserving of epistemic blame for believing that p.

How, then, do we decide whether a person merits epistemic praise (or blame) for believing a given proposition? We do so on the basis of whether or not she has been epistemically responsible in coming to believe that proposition. If a person comes to believe that p in an epistemically responsible manner (e.g. checking her work, considering defeaters, weighing the evidence), she is worthy of epistemic praise and is, therefore, personally justified in believing that p. If, on the other hand, a person comes to believe that p in an epistemically irresponsible manner (e.g. wrecklessly adopting beliefs, ignoring counter-evidence, trusting Evan Mecham), she is epistemically blameworthy and is, therefore, personally unjustified in believing that p. Thus, a person's personal justificatory status is a function of whether or not she has proceeded in an epistemically responsible way in coming to hold a given belief. Accordingly, we can revise our first approximations (PJ_j) and (PJ_u) as follows:

(PJ_j') S is personally justified in believing that p iff S has come to believe that p in an epistemically responsible fashion.

(PJ_u') S is personally unjustified in believing that p iff S has been

epistemically irresponsible in coming to believe that p.[49]

While there may be other conceptions of personal justification, they are of no interest to me here. The type of personal justification with which I am concerned in the present dissertation is that cashed out by (PJ_j') and (PJ_u'), since personal justification so-conceived —— that brand of justification intimately connected with our notions of epistemic praise, blame, responsibility, and irresponsibility —— is, as we shall see, the kind of justification which has served as the impetus for internalism.

Now that we have working conceptions of both personal justification [(PJ_j') and (PJ_u')] and doxastic justification [(DJ')], progress in epistemology readily awaits us. For example, in the next subsection I will use these two kinds of justification to demonstrate that all of the counterexamples to reliabilism presented in section 3 are ultimately unsuccessful. As a result, reliabilism will reemerge as a viable theory, albeit a theory of doxastic justification, and the personal/doxastic justification distinction will be further legitimated.

---

[49] (PJ_j') and (PJ_u') are not being offered as an analysis of the concept "personal justification". Rather, they are intended to point out what the concept of personal justification is. It is this conception of personal justification which I will attempt to analyze in Chapter 3.

c. The Right Reliabilist Reply

The onus of this subsection is to prove that all of the section 3 objections to reliabilism fail. My argument to this effect proceeds in two steps. First, I argue that the most plausible way to interpret reliabilism is as an account of doxastic justification. Then, I reexamine the objections to reliabilism in order to show that they all conflate personal and doxastic justification, sometimes blatantly. Once they are deconflated, the objections at best only serve to show that reliabilism fails as a theory of personal justification. However, since reliabilism is intended as a theory of doxastic, not personal, justification, it becomes obvious that the objections are fundamentally misguided and simply do not apply to reliabilism properly construed. I now turn to step one, to wit, eliciting the proper construal of reliabilism.

Although Goldman does not definitively distinguish doxastic justification from personal justification, numerous passages in his "What Is Justified Belief?" strongly suggest that he offered his theory with something very much like doxastic justification in mind.[50] The opening sentence of this article states, "The aim of this paper is to sketch a

---

[50]NOTE: Even the title of Goldman's article, to wit, "What Is Justified Belief?" (my emphasis), suggests that he is concerned with doxastic justification, the kind of epistemic justification which attaches to beliefs.

theory of justified belief"[51] (my emphasis). He tells us explicitly what a theory of justified belief will consist of. It will consist of "a set of principles that specify truth-conditions for the schema |S's belief in p at time t is justified| "[52] (my emphasis). Moreover, if we recall the base and recursive clauses that are constitutive of his theory, which were presented in section 2, viz. (BC1), (RC1), and (BC1'), we see that in each case his analysandum is indeed "S's belief in p at t is justified"[53] (my emphasis). These passages clearly suggest that Goldman intends his theory to be a theory for epistemically evaluating beliefs, not persons.

That he offers his theory as a theory of justified beliefs, not justified persons, becomes even more obvious when we consider the following passage:

> Suppose S has a set B of beliefs at time t_0, and some of these beliefs are unjustified. Between t_0 and t_1, he reasons from the entire set B to the conclusion p, which he then accepts at t_1. The reasoning procedure he uses is a very sound one, i.e., one that is conditionally reliable. There is a sense or respect in which we are tempted to say that S's belief in p at t_1 is 'justified'. At any rate, it is tempting to say that the person [his emphasis] is justified in

---

[51] Alvin Goldman, "What is Justified Belief?", op. cit., p. 1.

[52] Ibid., p. 3.

[53] Ibid., pp. 13, 14 and 20.

> believing p at t_{11}. Relative to his antecedent cognitive state, he did as well as could be expected: the transition from his cognitive state at t_0 to his cognitive state at t_1 was entirely sound. Although we may acknowledge this brand of justifiedness —— it might be called 'Terminal-Phase Reliabilism' —— it is not a kind of justifiedness so closely related to knowing.[54]

Since Goldman thinks that the justifiedness of persons has little to do with knowledge and since he contends that "On the account of justified belief suggested here, it is necessary for knowing, and closely related to it",[55] it is clear that he regards his theory as a theory of the justifiedness of beliefs, i.e. as a theory of what I have been calling "doxastic justification". Thus, one extremely compelling reason for interpreting reliabilism as a theory of doxastic justification is that Goldman himself clearly seems to have intended it as such. There is, however, an additional reason, independent of Goldman's intentions, for viewing reliabilism as a theory of doxastic justification.

Recall from subsection b that according to (DJ') a belief is doxastically justified iff it has a high probability of being true. It is quite reasonable to regard

---

[54] Ibid., pp. 15-16. Here Goldman appears to be on the verge of discovering the personal/doxastic justification distinction, but he fails to pursue the distinction any further.

[55] Ibid, p.1.

reliabilism as an account of doxastic justification so-conceived, since the beliefs which reliabilism deems justified do have a high probability of being true beliefs. That reliabilistically justified beliefs do have a high probability of being true beliefs can be demonstrated as follows. In its simplest form, process reliabilism asserts that a belief is justified iff it results from a reliable BCP, where a BCP is reliable just in case it tends to produce true beliefs, i.e. just in case the indefinite probability of beliefs produced by it being true beliefs is high (at least greater than .5).[56] Since, by definition, beliefs produced by reliable BCP's have a high indefinite probability of being true beliefs, it follows on a reliabilist account that justified beliefs have a high indefinite probability of being true beliefs, since reliabilism identifies justified beliefs with reliably produced beliefs. As you may recall from section 2, indefinite probabilities are dyadic relations which relate classes (or properties) by specifying the probability of a

---

[56] If the indefinite probability of beliefs produced by a given BCP being true beliefs is high, then that BCP will tend to produce true beliefs. If, on the other hand, the indefinite probability of beliefs produced by a given BCP being true beliefs is low, then that BCP will tend to produce false beliefs. Thus, a BCP will tend to produce true beliefs iff it is the case that the indefinite probability of beliefs produced by that BCP being true beliefs is high.

member of one class being a member of a second class. Consequently, in demonstrating that reliabilism entails that justified beliefs have a high indefinite probability of being true beliefs, we have ipso facto demonstrated that reliabilism entails that beliefs belonging to the class of justified belief have a high probability of belonging to the class of true beliefs. The former demonstration entails the latter since, by the very definition of indefinite probability, for any two classes A and B, the indefinite probability of A's being B's is high just in case x's belonging to class A have a high probability of belonging to class B. Since reliabilism entails that justified beliefs (members of the class justified belief) have a high probability of being true beliefs (members of the class true belief), reliabilism entails that justified beliefs have a high probability of being true (since true beliefs are true). But notice, this is precisely the sort of probabilistic truth connection which a theory of doxastic justification must affix, since a belief is doxastically justified iff it has a high probability of being true. It should by now be obvious that reliabilism provides exactly the kind of analysis which a theory of doxastic justification must provide, and for this reason, it ought to be interpreted as just such a theory.

Let me conclude step one with a brief summary. We have seen that reliabilism takes beliefs as its domain of

evaluation. We have also seen that Goldman clearly seems to have intended his theory as a theory of doxastic (not personal) justification. And finally, we have just seen that reliabilism provides the sort of probabilistic connection between justified belief and truth that is definitive of doxastic justification. For these reasons, I contend that the only plausible way to construe reliabilism is as a theory of doxastic justification.

We are now in a position to see why none of the section 3 objections apply to reliabilism so-construed. Recall that in the first counterexample to necessity we are to suppose that unbeknownst to us the evil demon hypothesis is true and that consequently all our beliefs have been produced by BCP's which the demon has rendered unreliable. According to reliabilism, none of our beliefs would be justified in such a world. Lehrer and Cohen contend that this result is untenable. I, on the other hand, maintain that, in evaluating our beliefs in such a world as unjustified, reliabilism provides precisely the right result. After all, in an evil demon world where all of our beliefs are produced by highly unreliable BCP's, all of our beliefs are extremely improbable, i.e. to say, all of our beliefs are probably false.[57] It strikes me as antithetical to the entire

---

[57]Of course, it follows by definition (see (DJ')) that such probably false beliefs are not doxastically justified.

epistemological enterprise to regard beliefs which are probably false as having positive epistemic status, i.e. as being epistemically justified. After all, probably false beliefs obviously run counter to the epistemic goal of maximizing truth and minimizing error, since they virtually ensure error, and surely, there is nothing epistemically positive about beliefs which virtually ensure error. Consequently, reliabilism, in evaluating such demon-manipulated, probably false beliefs negatively as being epistemically unjustified, yields exactly the right result, Lehrer and Cohen's intuitions to the contrary notwithstanding.

Nevertheless, it must be admitted that their objection is prima facie intuitively quite appealing, so much so that Goldman felt compelled to modify his theory accordingly. This initial appeal derives from their correct observation that were we to reason in the evil demon world exactly like we reason in the actual world, we would be reasoning as well as could be expected given our unfortunate circumstances. In light of this observation, it is both natural and correct to claim that we would be just as well justified in

---

But such a stipulative truth, if that is all it were, would be neither interesting nor illuminating. What I now hope to show in the body of the text is that the view that such beliefs are not (doxastically) justified is intuitively correct.

believing what we do in the demon world as we are in believing what we do in the actual world. On the basis of this correct claim, Lehrer and Cohen conclude that our beliefs are just as well justified in the evil demon world as they are in the actual world, but in drawing this inference they are clearly conflating doxastic justification with personal justification. This inference amounts to thinking that our beliefs must be epistemically good (i.e. doxastically justified), because we have reasoned well. However, thinking that our beliefs must be epistemically good because we have reasoned well is just as fallacious as thinking that our beliefs must be true because we have reasoned well. That we have reasoned well does, indeed, confer positive epistemic status on us, i.e. makes us justified, but it does not of itself confer positive epistemic status on our beliefs, because our beliefs can still be extremely improbable, despite our having reasoned well. It is precisely this failure on the part of their readers to keep personal and doxastic justification separate that gives their counterexample its initial plausibility. As we shall now see, the other counterexamples manifest this conflation even more clearly.

In the second counterexample Cohen contrasts two inhabitants of evil demon world W', A who is a good reasoner and B who is a confused reasoner, a wishful thinker, etc. Of course, the demon has seen to it that good reasoning is

as unreliable in W' as confused reasoning and wishful thinking are in W'. Cohen rightfully observes that:

> Since the beliefs of A & B are both produced by unreliable processes (the evil demon sees to this), a reliabilist theory of justification must render identical epistemic appraisals of both sets of beliefs.[58]

Cohen finds this result to be unacceptable. Since, in his opinion, "There is a fundamental epistemic difference between the beliefs of A and the beliefs of B",[59] reliabilism, which fails to take this difference into account, must be mistaken. His argument for there being such an epistemic difference between A's beliefs and B's beliefs is roughly that since A is a good reasoner and B is not, the beliefs of A are justified while the beliefs of B are not. In so arguing, Cohen clearly makes the mistake of conflating personal and doxastic justification. His example is instructive, not because it is an objection to reliabilism, but because it demonstrates the importance of keeping these two kinds of justification separate.

Cohen's example does uncover an epistemic difference, but not the one he thinks it does. Since A is a good

---

[58]Stewart Cohen, Justification and Truth, op. cit., pp. 10-11.

[59]Ibid., p. 11.

reasoner and B is a wishful thinker, there is a definite epistemic difference between person A and person B. By reasoning in accordance with the canons of inductive logic, A is presumably adopting his beliefs in an epistemically responsible fashion. B, on the other hand, in forming wishful beliefs, is presumably being epistemically irresponsible.[60] Hence, A is personally justified in his beliefs in W', whereas B is personally unjustified in his beliefs in W'. It is crucial to realize, however, that A's being being personally justified in his beliefs does not entail that his beliefs are themselves doxastically justified. This will become obvious once we see that in the case under discussion A's beliefs are doxastically unjustified.

Due to demon influence, both A's beliefs and B's beliefs have been produced by equally unreliable BCP's, and because of this, A's beliefs are just as improbable as B's beliefs. Accordingly, A's beliefs are no better from the epistemic standpoint than B's beliefs, since they are just as likely to result in error as are B's beliefs. According to (DJ'), since both A's beliefs and B's beliefs are probably false, both A's beliefs and B's beliefs are

---

[60] I am assuming that B believes that wishful thinking is an unreliable way to form beliefs. The situation would be more complicated, were B to believe wishful thinking to be reliable.

doxastically unjustified. Since reliabilism entails that A's and B's unreliably produced beliefs are unjustified, as a theory of doxastic justification it yields exactly the right result.

Cohen's next purported counterexample to necessity can be handled in a similar fashion. Recall that this example has us consider a case where, owing to the demon, both A and B have the unreliably-produced, non-discursive perceptual belief that ϕ , but where B, unlike A, has strong counter evidence which he simply ignores. Here is what Cohen has to say about this case:

> I think it's clear that there is a fundamental epistemic difference between A's perceptual belief and B's perceptual belief -- a difference which again underscores the normative character of epistemic justification. Notice that we need not assume that B disregards the evidence as a result of any discursive process. He may just arbitrarily ignore it. But from an epistemic point of view, B ought not to have proceeded in the way he did. We might say that contrary to A, B has been epistemically irresponsible in accepting that there is something before him. As a result, while A is justified in his ϕ perceptual belief, B is not[61] (emphasis added).

Let me first say that I essentially agree with Cohen's conclusion that "while A is justified in his ϕ perceptual

---

[61]Stewart Cohen, Justification and Truth, op. cit., pp. 13-14.

belief, B is not."[62] But it should be noted that Cohen is making a personal evaluation here, not a doxastic one, since the subjects being evaluated are cognizers A and B. While this personal evaluation is correct, neither it nor anything else in the passage supports his initial contention that "there is a fundamental epistemic difference between A's perceptual belief and B's perceptual belief". That he apparently takes the personalist conclusion above to support his initial contention is yet another manifestation of the personal/doxastic justification conflation, which pervades and also undermines all of his purported counterexamples to reliabilism. While there is a fundamental epistemic difference between person A and person B, this does not demonstrate that there is a similar epistemic difference between A's ϕ belief and B's ϕ belief, Cohen's thoughts to the contrary notwithstanding. Since A's perceptual belief is just as improbable as B's perceptual belief, A's

---

[62] My assertion that I essentially agree with Cohen's conclusion needs to be qualified. For reasons which will become clear in Chapter 3, I think that Cohen should have concluded that while B is personally unjustified in his belief, A is not personally unjustified in his. This conclusion is not equivalent to Cohen's, because it allows for the possibility that A is personally ajustified in his belief, a possibility which I take to be actual in the case at hand. Thus, while strictly speaking I do not agree entirely with Cohen's conclusion as it is stated, I come so close to agreeing with it that feigning such agreement to facilitate the present discussion is warranted.

perceptual belief is just as doxastically unjustified as B's perceptual belief. So, once again, reliabilism yields the right doxastic evaluation.

Cohen's counterexamples to sufficiency rest on the same conflation that his counterexamples to necessity do. To see that this is so, let us consider his example in which, unbeknownst to the inhabitants of W^+, wishful thinking happens to be completely reliable there. It is a consequence of reliabilism that in such a world wishfully-formed beliefs are justified. Cohen, on the other hand, contends that such wishfully-formed beliefs are unjustified, no matter how reliable wishful thinking turns out to be, for, as he explains:

> The crucial factor, what by my lights makes the beliefs unjustified in these cases, is the fact that the reliability of the belief forming process is due to facts that are completely outside the ken of the subject. If as far as the subject knows, the state of affairs expressed by P is merely something he wishes for, then he is being epistemically irresponsible in accepting that P.[63]

If the subject in question believes that wishful thinking is unreliable, then I agree with Cohen that the subject is being epistemically irresponsible in accepting that p on the

---

[63]Stewart Cohen, Justification and Truth, op. cit., p. 16.

basis of wishful thinking, and as a result, he is personally unjustified in his belief that p [recall (PJ_u')]. But it does not follow from this that his belief that p is doxastically unjustified. To the contrary, since wishful thinking is completely reliable in W^+, our subject's belief that p is extremely probable, and so, according to (DJ'), his belief that p is, in fact, doxastically justified, which is just what reliabilism maintains.[64]

Although Cohen is wrong to regard our subject's belief that p as unjustified, his explanation for it supposed unjustifiedness merits further consideration. He contends that it is unjustified because "the reliability of the belief forming process [in this case wishful thinking] is due to facts that are completely outside the ken of the subject." This could mean one of two things. It might mean that the belief is unjustified because the subject is unaware of the actual reliability of the BCP which produced

---

[64]Cohen's second counterexample to sufficiency (see pp. 51-52) collapses for similar reasons. Even though there is a clear epistemic difference between person A the good reasoner and person B the wishful thinker, A's beliefs and B's beliefs remain on equal epistemic footing. After all, since good reasoning and wishful thinking are equally reliable in W+ (the good demon has seen to this), both the beliefs of A and the beliefs of B, being highly probable, are doxastically justified. Accordingly, it is B, and not his beliefs, that is unjustified. Since reliabilism is only concerned with doxastic evaluations, it remains unscathed by this purported counterexample, as well.

it. On this construal, it would turn out that virtually all of our beliefs are unjustified because we are, in a very real sense, unaware of the actual reliability of all our BCP's.[65] Although I think that this is the construal that Cohen most likely intended, given its untoward consequences, the principle of charity dictates that we consider the second thing Cohen's claim might mean. It might mean that the belief is unjustified, not because the subject is unaware of the actual reliability of the BCP which produced it, but because he is unaware of the facts in virtue of which the BCP is reliable; e.g. since he is unaware of the fact that a benevolent demon is making his wishful beliefs turn out to be true, his wishful belief that p is unjustified. Such a view is reminiscent of causal theories which require that the cognizer properly reconstruct the causal chain that led from the fact that p to his belief that p, only the reconstruction required here is much more complex than the reconstruction required by causal theories, since the cognizer is required to reconstruct how the BCP

---

[65] One might contend that sophisticated cognizers are aware of the actual reliability of their BCP's. Such a contention is open to doubt, but even if it is correct, it would turn out that all of the beliefs of young children (and, perhaps, of children who are not all that young) are unjustified, because young children are unaware that they have BCP's and a fortiori they are unaware of their BCP's actual reliability.

works in virtue of which it is reliable. This construal of Cohen's claim is no more plausible than the first, since it too entails that virtually all of our beliefs are unjustified. After all, the most sophisticated cognitive scientists and neurophysicists still are unaware of how most of our BCP's work and in virtue of which they are reliable. So, obviously, the man in the street is unaware of these things. Consequently, on either construal, Cohen's purported explanation of why our subject's wishful belief that p is unjustified entails that virtually all of our beliefs are unjustified. But such an explanation is no explanation at all, since what was to be explained is how this wishful belief differs from ordinary perceptual and inductive beliefs and in virtue of which the wishful belief is unjustified. Since no such difference is forthcoming, my contention that the reliably-produced wishful belief is (doxastically) justified is further vindicated.

Finally, let us turn to Bonjour's counterexample to reliabilism's sufficiency and give it the scrutiny it deserves. The case he has us consider centers around Norman, a perfectly reliable clairvoyant who is entirely unaware of his own clairvoyant power. As you may recall, Norman has absolutely no evidence as to the President's whereabouts, but his completely reliable faculty of clairvoyance, nevertheless, causes him to believe that the President is in New York City. Since his belief results

from a completely reliable BCP, reliabilism yields the result that Norman's belief about the President is justified. Bonjour questions this result:

> Is Norman epistemically justified in believing that the President is in New York City, so that his belief is an instance of knowledge? According to the modified externalist position, we must apparently say that he is. But is this the right result? Are there not still sufficient grounds for a charge of subjective irrationality to prevent Norman's being epistemically justified?[66]

By now you can probably anticipate my response. I most certainly agree with Bonjour that there are sufficient grounds for a charge of subjective irrationality to prevent Norman from being epistemically justified. After all, Norman has no evidence as to the President's whereabouts. Moreover, he is completely unaware that he has reliable clairvoyant power. So, from his own internal standpoint, it must surely seem as if his belief about the President's present location simply popped into his head out of thin air. And, obviously, it is epistemically irresponsible of Norman to continue to hold such a spontaneously occurring ungrounded belief. All that this shows, however, is that Norman is personally unjustified in believing that the

---

[66]Laurence Bonjour, "Externalist Theories of Empirical Knowledge", op. cit., p. 62.

President is in New York City, and by now we know that Norman's personal unjustifiedness is irrelevant to the doxastic justificational status of his belief. His belief is the result of a highly reliable BCP, to wit, completely reliable clairvoyance. As such, his belief is highly probable and is, therefore, a good one to hold from the epistemic viewpoint. Notice: Norman's situation with respect to his faculty of clairvoyance is not unlike a young child's situation with respect to her perceptual faculties. What's more, it is usually agreed that the young child's perceptual beliefs are justified even though she lacks any rationale for them. So, by parity of reason, we should agree that Norman's belief about the President's whereabouts is justified, while nevertheless maintaining that, relative to the other things he believes, he is being epistemically irresponsible in holding the belief and is, therefore, personally unjustified in doing so. Simply put, the fault lies with Norman, not with his belief, and our respective justificatory evaluations should reflect this fact. Therefore, I submit that if reliabilism is viewed as an account of doxastic justification, i.e. as a theory of justified belief, it remains unscathed by Bonjour's purported counterexample, as well.

In short, we have seen that none of the section 3 counterexamples to reliabilism hold up once the personal/doxastic justification distinction is brought to

bear on them. These counterexamples were designed to show that a belief's being reliably produced is neither necessary nor sufficient for that belief's being justified. However, instead of showing this, all that they succeed in showing is that a belief's being reliably produced is neither necessary nor sufficient for a person to be justified in holding that belief.[67] Of course, this success does not serve to refute reliabilism, since reliabilism properly construed is a theory of doxastic evaluation, not a theory of personal evaluation. In fact, rather than refuting reliabilism, these counterexamples actually serve to confirm reliabilism as an account of doxastic justification, since for each case reliabilism provides the correct doxastic evaluation. Nevertheless, in the next section we shall see that reliabilism as delineated in section 2 does not always result in the right doxastic evaluations, and therefore, it must be revised.

5. A Theory of Doxastic Justification: Goldman's Reliabilism Revised

The goal of the present section is to arrive at the correct theory of doxastic justification. Consequently, the

---

[67] That reliable production is neither necessary nor sufficient for personal justification should come as no surprise, since personal evaluations proceed in terms of whether or not a person has been epistemically responsible in coming to hold her beliefs.

present section is truly a pivotal one, since much of the dissertation hangs on the satisfactory attainment of this goal. In particular, the theory of knowledge which I sketch in Chapter 4 depends on the correctness of the theory of doxastic justification that is presented here. I shall ultimately argue that a revised version of Goldman's reliabilism provides the sought after theory. To demonstrate both the need for this revision and the sort of revision needed, I shall present cases where Goldman's theory results in the wrong doxastic evaluations. However, before doing so, I will first explain why both of Goldman's own revisions, viz. normal worlds reliabilism and the duplex theory, are unsatisfactory, since doing so reinforces the importance of the personal/doxastic justification distinction.

Since its inception, normal worlds reliabilism has been under the steady fire of counterexamplers. I will not reiterate their counterexamples here. In fact, my objections to normal worlds reliabilism are not, properly speaking, counterexamples at all. They are, instead, what might be called "theoretical objections", since they point out theoretical shortcomings of the normal worlds approach. My first objection centers around the much sought after truth connection. As we saw in section 2, one of the most seductive and theoretically attractive features of Goldman's reliabilism is that it affixes probabilistic truth

connection such that beliefs that are justified in W are probably true beliefs in W. However, unlike Goldman's reliabilism, normal worlds reliabilism provides no such connection. To see why, recall that according to normal worlds reliabilism a belief is justified in W *just in case* the BCP which produced it in W is reliable *in normal worlds*. On this view, justification does not entail probable truth, since a belief can be justified in W even though that belief is probably false in W. For example, suppose that belief B is produced in W by a BCP which is highly reliable in normal worlds, but terribly unreliable in W. On the normal worlds view, B is justified in W since it has been produced by a BCP that is reliable *in normal worlds*, but because this BCP is unreliable *in W*, B is probably a false belief in W. Hence, normal worlds reliabilism fails to provide the sort of probabilistic truth connection which its predecessor, Goldman's reliabilism, succeeds in providing. Consequently, normal worlds reliabilism lacks the theoretical attractiveness of its predecessor. This is an especially unhappy consequence for an externalist theory like normal worlds reliabilism, since the primary motivation behind externalist theories lies in their unique ability to provide the required truth connection.

There remains another more devastating objection to normal worlds reliabilism, namely, it is an unnecessary modification of Goldman's reliabilism that results in the

wrong justificatory evaluation of demon-manipulated beliefs. According to normal worlds reliabilism, the beliefs of unsuspecting demon-world-inhabitants are justified, since the demon-rendered unreliable BCP's which produced these beliefs are reliable BCP's in normal worlds. But, as we have already seen, to regard such beliefs as justified is a theoretical mistake.[68] Since such demon-manipulated, unreliably-produced beliefs are probably false, they lack positive epistemic status and are, therefore, unjustified. This, of course, is exactly the justificatory evaluation rendered by Goldman's reliabilism without the normal worlds codicil. Consequently, normal worlds reliabilism is both unnecessary and theoretically unfounded.[69]

Goldman has since come to be of the opinion that the beliefs of a demon-manipulated cognizer are, in a very important sense, unjustified. I think that this realization on his part is primarily what led him to abandon the normal worlds approach in favor of the duplex account. Although I do not find the duplex theory to be entirely satisfactory, I am nevertheless quite sympathetic with its underlying motivation. The duplex theory is motivated by a desire to explain a certain epistemological datum in a non-ad hoc way.

---

[68]The conflation of personal justification with doxastic justification is, no doubt, the source of this mistake.

[69]For further objections to Goldman's normal worlds

That datum is the fact that unbiased epistemologists genuinely disagree about the justificatory status of demon-manipulated beliefs. Some epistemologists genuinely aver that such beliefs are clearly unjustified, while others maintain with equal sincerity that such beliefs are entirely justified. Rather than throwing his hands up in the air and saying, "Well, this is just a case of competing intuitions, and if you don't share my intuitions, then we simply cannot discuss the matter any further.", Goldman attempts to isolate the different perspectives or conceptions of justification that lead to these divergent intuitions, which is surely a more appropriate and less painful way of dealing with competing intuitions than the standard philosophical practice of butting heads. Here is roughly what Goldman observes: If we look at demon-manipulated beliefs as the unreliably-produced and, hence, probably false beliefs that they are, it is quite natural to regard them as unjustified, since probably false beliefs do not promote the epistemic goal of gaining truth and avoiding error. However, if we look at demon-manipulated beliefs from the perspective of the hapless demon-world-inhabitant who has exactly the same evidence for his beliefs that he would have were he in a verific non-manipulated world, it is quite natural to regard

---

reliabilism, see my "Coherentism Reliabilized", Acta Analytica, No. 2 (1986).

him as justified in holding those beliefs, since, given what he has to go on, he has done the best he could in forming his beliefs and, therefore, cannot be blamed epistemically for holding the beliefs he does.

Goldman contends that these two perspectives employ inherently different conceptions of justification, conceptions which are captured by his notions of strong and weak justification, respectively. With this contention I disagree. In particular, I think that Goldman's account of weak justification fails to capture the "epistemically blameless" conception of justification. In what is to follow, I shall demonstrate via three objections the inadequacy of Goldman's account of weak justification, which a fortiori will demonstrate the falsity of the duplex theory.[70]

First, when we say a belief is justified, we attribute positive epistemic status to that belief. By parity of reason, when we say a belief is weakly justified, we

---

[70]It is, of course, conceivable for the duplex theory as a whole to be false, and it nevertheless be the case that its account of strong justification is correct. However, I contend that its account of strong justification is also mistaken, albeit only in a relatively minor way. It seems straightforward that Goldman's account of strong justification is a theory of doxastic justification, i.e. is a theory of justified belief. The undermining clause has no business being in such a theory, since it in no way affects the indefinite probabilities of the beliefs produced. I do not deny that the undermining clause is important, but its importance lies in the realm of personal justification. If S (mistakenly) believes a method or process to be

attribute a weak degree of positive epistemic status to that belief, but positive epistemic status nonetheless. Since weakly justified beliefs are ill-formed-but-blameless beliefs, the question which arises is, "Do ill-formed-but-blameless beliefs have any degree of positive epistemic status? I contend that they do not. Ill-formed-but-blameless beliefs are ipso facto ill-formed beliefs. A belief is ill-formed in Goldman's terminology just in case it is produced by an unreliable cognitive process or method. Consequently, all ill-formed beliefs have a high indefinite probability of being false beliefs, and that, of course, includes ill-formed-but-blameless beliefs. I contend that probably false beliefs have no positive epistemic status whatsoever, not even weak positive epistemic status. After all, as we saw in section 4 subsection c, probably false beliefs run counter to the epistemic goal of maximizing truth and minimizing error, and so, from the epistemic point of view there is nothing positive about them. Since ill-formed-but-blameless beliefs, being probably false, have no positive epistemic status, they are not even weakly justified.[71]

---

unreliable, but uses it anyway, then S is being epistemically irresponsible, and it is S that is unjustified, not S's belief.

[71]Goldman may insist that as he is using the locution

Second, taken literally, the notion "blameless belief" does not make sense, since beliefs are not the kinds of things to which blame can be properly ascribed. Beliefs do not do anything in the agency sense of 'do'.[72] Only agents, e.g. persons, do things in this sense. Moreover, when speaking literally, agents are the only kinds of things to which any normative sense of blame can be ascribed.[73] Admittedly, "blameless beliefs" may be some sort of metaphor or abbreviation for saying that the person cannot be blamed for holding the belief. Goldman may even intend it as such an abbreviation, but if so, then notice that it is the person, not the belief, which is free from blame, and this is surely a personal evaluation.

Finally, wanting strong and weak justification to be mutually exclusive, opposing notions, Goldman narrows the notion of weak justification to that of mere epistemic blamelessness. If my second objection is right, then the

---

'weakly justified belief', it does not imply that the belief has weak positive epistemic status. If so, then he owes us an account of just what is being attributed to a belief when it is said to be weakly justified.

[72]Of course, beliefs have causal effects and so can be said to cause certain things, but this kind of causation is not agency.

[73]We can blame a hurricane for the massive destruction left in its wake, but this is a causal use of 'blame', not a normative one.

only way we can make sense out of the "epistemically blameless" conception of justification is to regard it as a kind of personal evaluation. So, on the narrower conception of weak justification, a person is weakly justified in holding a belief just in case he is merely free from blame for holding that belief. But if weak justification is modified in this way, it is far from clear that weak justification captures the kind of justification possessed by the scientifically benighted cognizer and the demon-world-inhabitant. Consider, for example, the case of the scientifically benighted cognizer S^*. If you recall, S^* uses method M, the method of consulting zodiacal signs in a culturally approved fashion, to form a belief about the outcome of an impending battle. In order to motivate the intuition that S^* is indeed weakly justified in his belief, Goldman tells us:

> He [S^*] is situated in a certain spatio-historical environment. Everyone else in this environment uses and trusts method M. Moreover, our believer [S^*] has good reasons to trust his cultural peers on many matters, and lacks decisive reasons for distrusting their confidence in astrology. While it is true that a scientifically trained person, set down in this same culture, could easily find ways to cast doubt on method M, our believer [S^*] is not so trained, and has no opportunity to acquire such training. It is beyond his intellectual scope to find flaws in M. Thus, we can hardly fault him for using M, nor fault him therefore for believing what he does. The belief in question is blameless, and that seems to explain why we are tempted to call it

justified[74] (brackets and their contents added).

Now I ask you, is it the case that S^* is merely epistemically blameless for holding the belief he does on the basis of method M? I contend that it is not. Given S^*'s cultural plight, as Goldman has so aptly described it, S^* has formed his belief exactly as he epistemically should have. Consequently, S^* is not merely epistemically blameless. He is worthy of epistemic praise for having proceeded in such a culturally approved epistemic fashion. In fact, it would have been epistemically irresponsible of S^* to reject M and thereby not form the belief, given that he has no reason to doubt M and every reason culturally available to accept and employ M. It is simply false that S^* is only weakly justified, i.e. only merely blameless, in holding his belief, for he is as fully justified in holding his belief as anyone in his unfortunate situation could be. Therefore, weak justification in the sense of mere epistemic blamelessness fails to capture the kind of justification had by S^*.

Once we acknowledge that weak justification is not the sort of justification that S^* possesses, it becomes clear

---

[74]Alvin Goldman, "Strong and Weak Justification", op. cit., pp. 2-3. It should be noted that in the last two sentences of this passage Goldman shifts from a personal evaluation to a doxastic one.

that Goldman's strong/weak justification distinction does not account for the two opposing justificatory evaluations which the S^* case entices us to make. What does account for the pull we feel toward each of these opposing justificatory evaluations is, I submit, the personal/doxastic justification distinction. On the one hand, we are inclined to evaluate S^* as being epistemically justified in holding his battle belief, because he has done his epistemic best in forming the belief. On the other hand, there is an inclination to evaluate S^*'s belief as being epistemically unjustified, since, having been formed by what is in fact an unreliable method, the belief is probably false, and probably false beliefs have negative epistemic status. Quite clearly, the former constitutes a personal evaluation, while the latter constitutes a doxastic one. Consequently, rather than supporting the strong/weak justification distinction, the S^* case reinforces the need for and the legitimacy of the personal/doxastic justification distinction. It is with this need in mind that I now attempt to provide an account of doxastic justification.

I began this chapter by claiming that a modified version of Goldman's reliabilism provides the correct account of doxastic justification. The time has come for me to defend this claim. To do so, I shall begin by demonstrating that such modification is in fact needed. I shall then modify the theory accordingly.

Most of the philosophers who have objected to Goldman's theory have done so by attacking the final base clause that he puts forward. As a result, relatively little attention and/or criticism has been directed toward the recursive clause that he adopts, though, as I shall now argue, such criticism is certainly warranted. Recall how Goldman has formulated the recursive clause:

> (RC1) If S's belief in p at t results ('immediately') from a belief-dependent process that is (at least) conditionally reliable, and if the beliefs (if any) on which this process operates in producing S's belief in p at t are themselves justified, then S's belief in p at t is justified[75] (emphasis added).

Also recall how Goldman characterizes a conditionally reliable BCP: "A process is conditionally reliable when a sufficient proportion of its output-beliefs are true given that its input-beliefs are true." This rendering of conditional reliability strongly suggests that the following corollary regarding conditional reliability is also true:

> (CCR) A conditionally reliable BCP is unreliable when its input-beliefs are false.

---

[75]For the sake of simplicity, I have elected to use (RC1) rather than the more complicated (RC1').

Admittedly, (CCR) does not follow from Goldman's characterization of conditional reliability. After all, a cognitive process P might satisfy Goldman's criterion for conditional reliability, i.e. might be reliable when its input-beliefs are true, while also being reliable when its input-beliefs are false. Nevertheless, I submit that the context in which Goldman introduces the notion of conditional reliability indicates that (CCR) is true of conditionally reliable BCP's. Recall that Goldman contrasts conditionally reliable BCP's with unconditionally reliable BCP's. But if we take a belief-dependent cognitive process P and assume that it is reliable when its input-beliefs are true and also assume that it is reliable when its input-beliefs are false, then P will turn out to be unconditionally reliable rather than merely conditionally reliable, since it will be reliable no matter what input-beliefs are used. Thus, if P is to be merely conditionally reliable, then it must be reliable when its input-beliefs are true and unreliable when its input-beliefs are false, and this shows that (CCR) is a necessary condition for a BCP to be conditionally reliable.

Having established (CCR), I will use it to show that (RC1) is false, but first a few preliminaries are needed. As we have already observed, when a belief is produced by a reliable BCP, that belief, by definition, has a high indefinite probability of being a true belief, and so it is

doxastically justified. On the other hand, when a belief is produced by an unreliable BCP, that belief, by definition, has a high indefinite probability of being a false belief, and so it is doxastically unjustified. In my opinion, this shows that reliability is the underlying ingredient of doxastic justification. With this in mind let us examine (RC1). According to (RC1), a belief is (doxastically) justified if it results from a conditionally reliable BCP which only has justified beliefs as inputs. It is easy to see that (RC1) does not provide a sufficient condition for doxastic justification. After all, on any plausible theory of (doxastic) justification, it will always be possible to have justified-but-false beliefs. It is, of course, this possibility which leads to the demise of (RC1). For suppose that a belief results from a conditionally reliable BCP which has as its only input a justified-but-false belief. In such a case the analysans of (RC1) is satisfied, but the belief in question is not doxastically justified. The belief is not doxastically justified because it is probably false. The reason that the belief is probably false is because it was unreliably produced. After all, it was produced by a conditionally reliable BCP which has a false belief as its only input, and according to (CCR) such a BCP is unreliable when its input-beliefs are false.

In light of the previous objection, it is obvious where (RC1) goes wrong. (RC1) requires that the input-beliefs to

the conditionally reliable BCP be justified, when it should have required that these input-beliefs be true. I propose that we modify (RC1) in just this way to get:

> (BC2) If S's belief that p at t is produced by a conditionally reliable BCP and if the beliefs on which this BCP operates in producing said belief are true, then S's belief that p at t is doxastically justified.

It should be noted that the resulting sufficient condition for doxastic justifiedness is no longer a recursive clause, but rather an additional base clause.

Before examining Goldman's base clause (BC1'), I want to consider one objection which might be raised to (BC2). The objection goes as follows: In requiring that the input-beliefs be true rather than merely justified, you are making the theory of doxastic justification even more externalistic than it already was, and such rampant externalism cannot be tolerated. Two responses are in order. First, while I admit that the correct theory of personal justification must be internalistic in nature, I aver that only an externalist theory can provide the correct account of doxastic justification, since only an externalistic theory can provide the kind of probabilistic truth connection which doxastic justification conceptually requires. Hence, that a theory of doxastic justification is rampantly externalistic does not constitute an objection to

such a theory. Second, it is simply a mistake to maintain that (BC2) is more externalistic than (RC1). After all, the kind of justification appealed to in (RC1) is Goldman's brand of externalistic justification. Consequently, S no more knows whether his input-beliefs are justified than whether they are true. That is to say, the justificatory status of S's input-beliefs are as external to S as their truth values are. Therefore, I am not guilty of offering a more externalistic condition of doxastic justification, but even if I were, this would not constitute an objection to the base clause presented.

Speaking of base clauses, let us now turn to Goldman's base clause, which in my opinion is also in need of amendment. As we saw in section 2, the base clause which Goldman settles on is:

> (BC1') If S's belief in p at t results from a reliable cognitive process, and there is no reliable or conditionally reliable process available to S which, had it been used by S in addition to the process actually used, would have resulted in S's not believing p at t, then S's belief in p at t is justified.

Goldman adds the counterfactual element to handle the case of Jones, who persists in believing his reliably-produced memory beliefs, despite having strong evidence that his memory is unreliable. In so doing, Goldman makes the justificational status of a belief a function not only of

the BCP's actually used in producing it, but also of the BCP's which could and should have been used. I contend that making the justificational status of a belief partly a function of the BCP's which could and should have been used is a mistake. After all, the existence of BCP's which could and should have been used in no way affects the indefinite probability of being a true belief which the belief has in virtue of having been produced by the BCP which in fact produced it. For example, in the case of Jones, although there are BCP's which Jones could and should have used, since his memory beliefs are the result of highly reliable memory, they have a high indefinite probability of being true beliefs (in virtue of having been so-produced). Since, according to (DJ'), a belief is doxastically justified <u>just in case</u> it is probably true, I submit that Jones' memory beliefs are in fact doxastically justified, Goldman's intuitions to the contrary notwithstanding. Nevertheless, I agree with Goldman that in light of the evidence <u>Jones</u> ought not continue to believe his memory beliefs. But the point which must be stressed is that the fault lies with <u>Jones</u>, not his <u>beliefs</u>. Since Jones has excellent evidence that his memory is unreliable, but persists in holding his memory beliefs despite this evidence, Jones is being epistemically irresponsible. As a result, according to (PJ_u') <u>Jones</u> is personally unjustified in holding his memory beliefs. Of course, as I have been emphasizing all along, Jones'

personal unjustifiedness is perfectly compatible with his memory beliefs being highly probable and therefore doxastically justified. In fact, the Jones example provides a wonderful illustration of such compatibility. For these reasons, I contend that, rather than demonstrating the need for a counterfactual element in doxastic justification, the Jones example demonstrates the need for the personal/doxastic justification distinction.

If (BC1') with its counterfactual element is mistaken, then what is the correct base clause? I contend that Goldman had it right to begin with when he formulated (BC1). However, in order to make it clear that the analysandum is doxastic justification I will rewrite it as:

> (BC1*) If S's belief that p at t is produced by an unconditionally reliable BCP, then S's belief that p at t is doxastically justified.

Since no recursive clause is needed, we can get a complete theory of doxastic justification by adding a standard closure clause to the two base clauses (BC1*) and (BC2). It would be nice for the purposes of theoretical neatness and conciseness, however, if we could formulate the theory in terms of a single base clause with a standard closure clause. This is what I propose to do, but I must first introduce as a technical notion the notion of "actual reliability".

(AR) A BCP is actually reliable iff either (1) it is an unconditionally reliable BCP or (2) it is a conditionally reliable BCP whose input-beliefs are all true.

Given the notion of actual reliability, we can easily formulate the ultimate base clause of doxastic justification as follows:

(UBC) If S's belief that p at t is produced by an actually reliable BCP, then S's belief that p at t is doxastically justified.

Combining (UBC) with a standard closure clause gives us, I submit, the correct theory of doxastic justification, a theory which we might call "doxastic reliabilism". That doxastic reliabilism is correct is evidenced by the fact that any belief which is produced by an actually reliable BCP will probably be a true belief in virtue of such production, and probable truth is what is required by (DJ') for doxastic justification.[76] In addition to being correct, the theory also has the virtue of straightforwardly revealing its commitment to the view that actual reliability is the underlying ingredient of doxastic justification.

---

[76]Similarly, if a belief is produced by an actually unreliable BCP, then it probably is a false belief and, hence, is doxastically unjustified.

I will have more to say about doxastic reliabilism in Chapter 4, where I will argue that doxastic justification, as analyzed by doxastic reliabilism, is a necessary condition for knowledge. But first, Chapter 3 is devoted to a discussion of personal justification, since whether or not a person is personally justified in holding a belief is also relevant to whether or not she knows that which she believes.

CHAPTER 3

A COHERENCE THEORY OF PERSONAL JUSTIFICATION

1. Introduction

The goal of this chapter is to provide an account of personal justification, where "personal justification" is understood in terms of (PJ_j') and (PJ_u'). As the title suggests, the account to be advanced is a coherence theory, but it is a nonstandard coherence theory, since I incorporate into the theory elements from foundation theories, as well as coherence theories. For example, like some foundationalists, I maintain that personal justification "proceeds in terms of reasons"[1] where, by definition, "A belief P is a reason for a person S to believe Q iff it is logically possible for S to become justified in believing Q by believing it on the basis of P."[2] Like some coherence theorists, I contend that "There is no exit from the circle of one's beliefs",[3] and consequently, I hold that beliefs are the only mental states which can serve as reasons for holding other beliefs.

---

[1]John Pollock, Knowledge and Justification (Princeton, NJ: Princeton University Press, 1974), p. 33.

[2]John Pollock, Contemporary Theories of Knowledge (in manuscript), p. 41.

[3]Keith Lehrer, Knowledge (Oxford: Oxford University Press, 1974), p. 188.

Even with such similarities, my account of personal justification will prove to be substantially different from every version of foundationalism and coherentism currently offered. For example, my account rests on a linear picture of reasoning rather than on a holistic one and, thus, parts company with current coherence theories. While most philosophers who have embraced a linear view of reasoning have felt compelled to adopt foundation theories (largely as a last ditch attempt to stay the impending regress), I am not among those philosophers, since I take issue with the central feature of foundationalism. This feature —— the heart of foundationalism —— can be specified as follows: For any theory T, T is a foundation theory only if T requires that a person's justification for any given belief ultimately rests on an epistemically privileged subclass of foundational beliefs. Put another way, for any theory T, T is a foundation theory only if T entails that a person S is justified in believing that p only if the chain of reasoning which leads S to believe that p ultimately derives from a belief (or set of beliefs) which S is immediately justified in believing. Such foundational beliefs are usually called "basic beliefs". I contend that basic beliefs are nothing more than a philosopher's fiction, conjured up to end what is thought to be a potentially vicious regress of reasons. As we shall see momentarily, my account entails that there are no basic beliefs, and thus,

it has the virtue of not being saddled with these unmoved epistemic movers. On my view, personal justification always proceeds in terms of reasons, which is just to say that a person S is personally justified in believing that p only if S has a reason R (some other belief) for believing that p. Consequently, basic beliefs are logically impossible on my view. It may seem that a regress of reasons is inevitable given such a view, but in section 6 I will prove otherwise.

The foregoing remarks make it obvious that I am not a foundationalist, since I reject the existence of basic beliefs. But I also reject the holistic view of reasoning indigenous to coherence theories. So, how can I claim to be offering a coherence theory of personal justification? The answer is quite simple. While it is true that holistic reasoning is a trademark of coherence theories, it need not be. In the next section, where will I classify the different kinds of justification theories available, we shall see that the defining feature of a coherence theory is its being a doxastic theory which denies the existence of basic beliefs, not its being a theory which employs holistic reasoning.

2. Partitioning Theories of Justification into Logical Space

If there is one thing which contemporary epistemology is not lacking, it is categories for classifying different kinds of epistemic theories. We have already been exposed

to some of these categories in the present dissertation, e.g. "Internalism", "Externalism", "Foundationalism", "Coherentism", etc. The purpose of such a taxonomy is not simply to increase one's vocabulary, but rather to facilitate epistemological discussion by grouping similar theories together. Unfortunately, these categories are used so loosely and with so little precision that rather than aiding discussion, they often impede it. As a result, it is not uncommon for philosophers to discover, after arguing at length, that they actually have no substantive disagreement, only a terminological one. The best way of avoiding such pseudodisagreements is to give these epistemic categories stipulative definitions and use them accordingly. Pollock has provided these desperately needed stipulative definitions in his Contemporary Theories of Knowledge.[4] Since his epistemological taxonomy is by far the clearest offered to date, I have elected to reiterate it in the present section. It should be noted that he cashes out all of these categories in terms of the unclarified, ambiguous notion "epistemic justification". However, such ambiguity is beneficial in the present context, for it allows us to map out the logical geography of theories of epistemic justification in general. Once we know the kinds of

---

[4]See Chapter One, Section 4, of his Contemporary Theories of Knowledge, op. cit.

theories of epistemic justification that are available to us, we will be in a position to decide which kind of theory is best suited to serve as a theory of personal justification. My strategy for clarifying these epistemic categories will be to identify and define the broadest categories first. Then, when I define each of the narrower subcategories, I will be able to indicate under which broad category it falls.[5]

Every theory of epistemic justification falls into one of two camps —— internalism or externalism. An internalist theory of epistemic justification is any theory which maintains that epistemic justifiedness is exclusively a function of the cognizer's internal states (e.g. belief states, memory states, perceptual states, etc.). An externalist theory of epistemic justification is any non-internalist theory of epistemic justification. Hence, an externalist theory maintains that epistemic justifiedness is (at least) partly a function of external features (e.g. the actual reliability of the producing process), features to which the cognizer lacks cognitive access.

---

[5]Since all of the stipulative terminology to follow is borrowed from Chapter One, Section 4, of Pollock's Contemporary Theories of Knowledge, op. cit., I will not distract the reader with footnotes for each definition borrowed. Instead, I simply refer the reader to the aforementioned reference for the original presentation of the stipulative definitions presented in this section.

Thus, by stipulation, internalism and externalism are mutually exclusive and jointly exhaustive categories of epistemic justification theories.

Similarly, every theory of epistemic justification is either a doxastic theory or a nondoxastic theory. Doxastic theories of epistemic justification embrace the doxastic assumption, viz. the assumption that epistemic justifiedness is exclusively a function of one's set of beliefs, i.e. one's doxastic corpus. Any theory of epistemic justification which denies the doxastic assumption is a nondoxastic theory. Accordingly, a nondoxastic theory of justification makes epistemic justifiedness (at least) partly a function of nondoxastic states, where nondoxastic states include nondoxastic internal states (e.g. perceptual states) and all external states (e.g. the producing BCP's actual reliability). Thus, the categories doxastic theory and nondoxastic theory provide us with another pair of mutually exclusive and jointly exhaustive categories of epistemic justification theories.

From the preceding definitions, it follows that every doxastic theory is an internalist theory. It also follows that every externalist theory is a nondoxastic theory. This leaves us with only one other possibility, viz. that of a nondoxastic internalist theory. Consequently, every theory of epistemic justification must either be a doxastic theory, a nondoxastic internalist theory, or an externalist theory.

The class of doxastic theories is exhausted by two mutually exclusive theory-types, to wit, foundation theories and coherence theories. Foundation theories are doxastic theories which maintain that there is an epistemically privileged subclass of basic beliefs which serves as the foundation for all other epistemic justifiedness. More precisely, for any theory T, T is a foundation theory iff (1) T is a doxastic theory, i.e. a theory which makes justifiedness exclusively a function of the beliefs one holds, and (2) T asserts that there exists an epistemically privileged subclass of basic beliefs on the basis of which all other epistemic justification proceeds. Coherence theories, on the other hand, deny the existence of a privileged class of basic beliefs. They maintain that all beliefs are capable of conferring epistemic justifiedness, not just the privileged few. Metaphorically, we can regard coherence theories as favoring doxastic democracy over foundational plutocracy. The metaphor is appropriate since it belies coherence theories' commitment to the view that all beliefs are on equal epistemic footing. We can capture this egalitarian commitment as follows: For any theory T, T is a coherence theory iff (1) T is a doxastic theory, and (2) T denies the existence of an epistemically privileged subclass of basic beliefs. Having exhausted the class of doxastic theories, let us turn our attention to nondoxastic theories.

The class of nondoxastic theories is comprised of externalist theories and nondoxastic internalist theories. A nondoxastic internalist theory is any internalist theory which denies the doxastic assumption, the assumption that epistemic justifiedness is exclusively a function of one's doxastic corpus. Direct realism is one such theory. It maintains that epistemic justifiedness is partly a function of the beliefs one holds, but also partly a function of certain other internal states. Thus, on the direct realist view, while beliefs are capable of conferring justification, they are not the only internal states which can confer epistemic justifiedness, since perceptual states (which are themselves neither justified nor unjustified) are also capable of conferring such justifiedness. To date, direct realism is the only nondoxastic internalist theory to have been worked out in any detail. Let us, therefore, consider the remaining class of epistemic theories, namely, externalist theories.

The class of externalist theories is comprised of all those theories which make epistemic justifiedness (at least) partly a function of features outside the scope of the cognizer's awareness. Pollock identifies two such theories —— process reliabilism and probabilism. Goldman's process reliabilism is an externalist theory, since, as we know from Chapter 2, it makes epistemic justifiedness partly a function of the reliability of the producing BCP and partly

a function of alternative BCP's which could and should be used. Probabilism represents another kind of externalist theory. According to probabilism, epistemic justifiedness is a function of the definite probabilities of the cognizer's beliefs, regardless whether the cognizer can assess these probabilities. Although Pollock only considers these two externalist theories, various other externalist theories are possible. For example, Lehrer's most recent theory, which he calls the "Monster Theory", is an externalist theory. The monster theory maintains that epistemic justifiedness is a function of coherence with the given cognizer's ultrasystem, where a proposition p coheres with person S's ultrasystem iff p beats or neutralizes every proposition q, with which p competes, on the basis of every member system of S's indefinitely large ultrasystem.[6] Obviously, human cognizers have no way of determining whether or not this complex coherence relation is satisfied. Similarly, doxastic reliabilism, as I formulated it in Chapter 2, is an externalist theory, because it makes epistemic justifiedness exclusively a function of the actual reliability of the producing BCP.

We now have a complete picture of the logical space of epistemic justification theories, which we can diagram as

---

[6]Keith Lehrer, "Metaknowledge: Undefeated Justification" (in manuscript).

follows:[7]

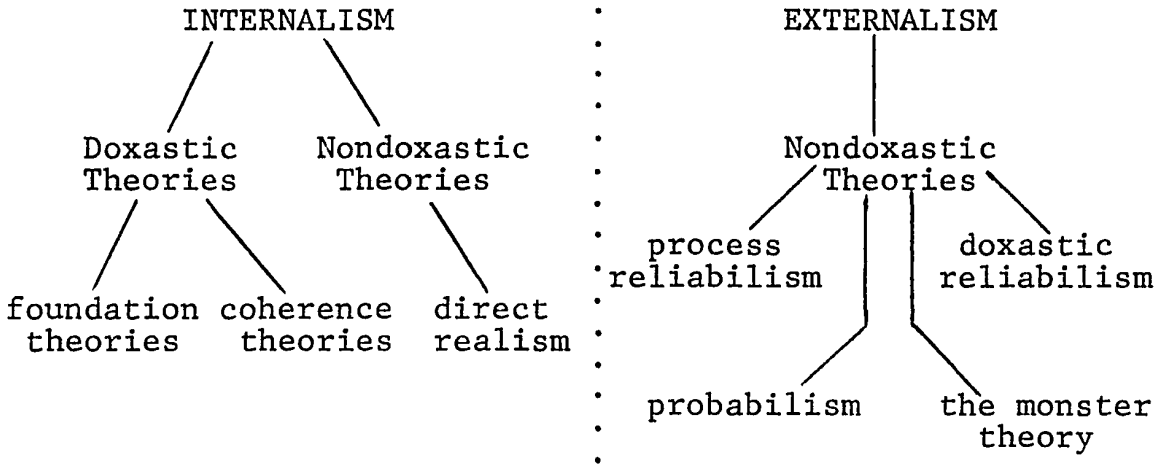

In Chapter 2, I argued that doxastic reliabilism, an externalist theory, provides the correct account of doxastic justification. From the outset of that chapter, I took it as a premise that the correct theory of doxastic justification had to be an externalist theory. Now that we have before us characterizations of the different kinds of epistemic justification theories, we can easily see that

---

[7]Notice that the headings "personal justification" and "doxastic justification" do not appear in the diagram. This is because personal justification and doxastic justification are species of epistemic justification. They are not categories of justification theories. Since the diagram exhausts the justification theories currently available, we must decide, from among these theories, which theory provides the correct account of personal justification. We have already observed that doxastic reliabilism provides the correct account of doxastic justification.

this premise is true. After all, (DJ') requires that there be a probabilistic connection between doxastic justification and truth. It should be obvious that no internalist theory, doxastic or nondoxastic, can provide such a probabilistic truth connection (or any other kind of conceptual conception between justification and truth), since, for whatever internal states a given internalist theory appeals to, there will always be possible worlds where an evil demon or a malevolent neuroscientist has seen to it that we possess those internal states, even though all of our contingent beliefs are false. Hence, no internalist theory can provide an account of doxastic justification. It follows, therefore, that the correct theory of doxastic justification (assuming there is one) must be an externalist theory. As we have seen, that externalist theory is doxastic reliabilism.

It should be just as obvious that only an internalist theory can provide the correct account of personal justification. We know from (PJ_j') and (PJ_u') that personal justification evaluations are based on and reflect whether a person has been epistemically responsible or epistemically irresponsible in deciding to hold a given belief. In deciding whether or not to hold a particular belief, a person can appeal <u>only</u> to those states of affairs to which she has cognitive access. Consequently, it is impossible for a person to appeal to external states of affairs, since

they are, by definition, states of affairs to which one lacks all cognitive access. Assuming that ought implies can in the present context, a person cannot be held responsible (and/or blamed) for failing to take external states of affairs into account, because it is impossible for her to do so. But notice, if we were to adopt an externalist theory, i.e. any theory which makes epistemic justifiedness partly a function of external states of affairs, as a theory of personal justification, we would be, in effect, holding the person responsible for failing to take into account states of affairs to which she lacks all cognitive access, i.e. we would be holding her responsible for failing to do the impossible, which is something she cannot be held responsible for. Thus, externalist theories are self-defeating if taken as theories of personal justification, which ipso facto entails that they cannot provide the correct theory of such justification. As a result, we need only consider internalist theories to determine the correct account of personal justification.

Since the class of internalist theories is exhausted by three mutually exclusive epistemic theories, viz. foundationalism, coherentism, and direct realism, we know that the correct account of personal justification can only be given by one of these theories. In the next section, I shall present an argument from elimination to show that the correct account of personal justification can only be given

by a coherence theory.

3. The Argument from Elimination

Although the three internalist theories are mutually exclusive theories, they do have one feature in common. They all share the view that justification proceeds in terms of reasons. Nevertheless, they disagree over what kinds of internal states count as reasons. Foundation theories and coherence theories, being doxastic theories, maintain that beliefs are the only internal states that can be reasons, whereas direct realism, which rejects the doxastic assumption, maintains that nondoxastic perceptual states, as well as beliefs, can serve as reasons. What is ultimately at issue here is whether or not the doxastic assumption is true. I contend that with regard to personal justification the doxastic assumption is true and that, therefore, direct realism is false. To see why the doxastic assumption is true, we must consider the kind of reasoning which gives rise to personal justification.

I contend that occurrent reasoning is the only kind of reasoning that is capable of conferring personal justification. Some philosophers, e.g. Thomas Reid, have held that occurrent (or conscious) reasoning is the only kind of reasoning that there is.[8] I am sympathetic with

---

[8]Thomas Reid, Essays on the Intellectual Powers of Man,

such a view.[9] However, even if there is such a thing as subconscious or unconscious reasoning, it certainly has no role to play in a theory of personal justification. After all, we have already seen that personal justification evaluations are basically evaluations of epistemic praise or blame, and surely, we cannot praise or blame a person epistemically for her subconscious or unconscious reasonings. Intuitively, it seems that occurrent reasoning is the only kind of reasoning for which we can hold a person directly responsible and hence for which we can praise or blame her, and this is why I maintain that occurrent reasoning is the only kind of reasoning that can confer personal justification. Indeed, the view that occurrent reasoning is the only kind of reasoning that can confer personal justification seems so obvious that it would hardly be worth mentioning, were it not for the fact that most epistemologists of late have rejected it.

Views like the one I am espousing, according to which justification proceeds only from occurrent reasoning, have

---

ed. Baruch A. Brody (Cambridge, Mass.: The M.I.T. Press, 1969).

[9] I am not denying the existence of unconscious computational processing. To the contrary, I think that a vast array of unconscious computational processing is going on inside us much, if not all, of the time. Nevertheless, I think it is a mistake to view these instances of unconscious computational processing as instances of reasoning.

been lumped together under the heading "Intellectualist Model of Justification".[10] Currently, the intellectualist model of justification has fallen into disrepute. By far the most common objection to the intellectualist model has been that it rests on a psychologically unrealistic picture of reasoning. At first glance, it is not at all clear what it is that is supposed to be psychologically unrealistic about its depiction of reasoning. Surely, it is not psychologically unrealistic to claim that we have a capacity for occurrent reasoning which we sometimes employ. To the contrary, that we can and do reason occurrently is a psychological fact. If the occurrence of occurrent reasoning is not itself psychologically unrealistic, what is it about the intellectualist model that is psychologically unrealistic? It is that the intellectualist model supposedly entails that we come to hold most of our beliefs (including our perceptual beliefs) as a result of occurrent reasoning. That it is psychologically unrealistic to

---

[10]Ernest Sosa coined the term 'Intellectualist Model of Justification' in his article "The Raft and the Pyramid: Coherence versus Foundations in the Theory of Knowledge", Midwest Studies in Philosophy, ed. French, Vehling, and Wettstein (University of Minnesota: 1980), p. 8. He built into this model the view that reasons have to entail that for which they are reasons, a view which I reject. Since its introduction, the term's usage has loosened significantly and now it generally refers to any theory which requires a person to proceed occurrently through the steps of reasoning in order to be justified.

maintain that we come to hold most of our beliefs as a result of occurrent reasoning cannot be denied, since most of our beliefs result from automatic belief-forming, cognitive processes over which we have little control. Therefore, if the intellectualist model does entail that most of our beliefs result from occurrent reasoning, then it is a psychologically unrealistic model of justification. So, let us examine the reductio ad absurdum argument which supposedly demonstrates this entailment.

A.	The intellectualist model of justification is correct.	(assump.)
1.	If the intellectualist model of justification is correct, then we are justified in believing that p only if we have occurrently reasoned to the belief that p.	(by def.)
2.	Therefore, we are justified in believing that p only if we have occurrently reasoned to the belief that p.	(A, 1)
3.	We know that p only if we are justified in believing that p.	(assump.)
4.	Therefore, we know that p only if we have occurrently reasoned to the belief that p.	(2,3)
5.	Skepticism is false.	(assump.)
6.	If skepticism is false, then we know most of our beliefs (including most of our perceptual beliefs.	(anal.)
7.	Therefore, we know most of our beliefs (including most of our perceptual beliefs).	(5,6)

8. Therefore, we have occurrently reasoned to most of our beliefs (including most of our perceptual beliefs). (4,7)

It must be admitted that it is psychologically unrealistic to maintain that most of our beliefs (particularly our perceptual beliefs) result from occurrent reasoning, and therefore, conclusion 8 is false. Since the argument is valid, but conclusion 8 is false, either our initial assumption that the intellectualist model is correct is false, or the argument is unsound. Fortunately for the intellectualist model, the argument is unsound, for it rests on a false premise, namely, premise 3.

It would be easy to see that premise 3 is false, were it not for the tendency to conflate different kinds of epistemic justification, one kind of which <u>is</u> necessary for knowledge. But once we realize that premise 3 asserts that personal justification is necessary for knowledge, we can demonstrate that it is false as follows. First, we know from (PJ_j) that a person is personally justified in believing that p <u>iff</u> she is worthy of epistemic praise for holding that belief. Second, we know all sorts of things which we do not deserve epistemic praise for believing. For example, we know many of our perceptual beliefs, but we very rarely reflect on, question, recheck, or amend our perceptual beliefs. Instead, we typically stand by passively and let our perceptual beliefs happen to us,

embracing them without the least bit of reflection into the matter, and surely, such unquestioning acceptance does not constitute epistemically praiseworthy behavior. Since we hold many beliefs which we do, in fact, know, but which we do not deserve epistemic praise for holding, it follows that we have many beliefs which we know, but which we are not personally justified in believing. Hence, personal justification is not necessary for knowledge, and therefore, premise 3 is false.

Without premise 3, the purported psychological unreality of the intellectualist model vanishes. It is not psychologically unrealistic to maintain that a person deserves epistemic praise for holding a belief only if she has occurrently reasoned to that belief, and this is all that an intellectualist model of personal justification entails. Such a view not only accords with psychological reality, but with epistemological reality, as well. For this reason, I submit that the intellectualist model provides the correct model for personal justification, at least if personal justification is understood in terms of (PJ_j) and (PJ_j').

The truth of the intellectualist model entails the truth of the doxastic assumption which, in turn, entails the falsity of direct realism. That the intellectualist model entails the doxastic assumption can be seen as follows. According to the intellectualist model, a person is

personally justified in believing that p only if she has occurrently reasoned to the belief that p. When reasoning occurrently to the belief that p, one comes to believe that p by appealing to information which either entails or inductively supports the truth of p. That just is what occurrent reasoning is. Moreover, the only information to which we can occurrently appeal is information that has been encapsulated in the form of beliefs. This is not to deny that we possess information which has not been encapsulated in beliefs. To the contrary, it seems obvious that we do possess such unencapsulated information, e.g. crude, uninterpreted sense data. What is at issue is whether or not such unencapsulated data can play a role in occurrent reasoning. I submit that it cannot, for it is difficult, if not altogether impossible, to imagine how we could occurrently appeal to such nondoxastic data. For example, suppose I am in such and such a perceptual state, and suppose that in virtue of being in this perceptual state I occurrently reason that "Since I am in such and such a perceptual state, there must be such and such an object before me." For me to reason this way, I must first recognize that I am in such and such a perceptual state, but in order for me to recognize that I am in such and such a perceptual state, I must believe that I am in that perceptual state (i.e., recognition entails belief). As a result, it appears that the only information capable of

playing a role in occurrent reasoning is information in the form of beliefs. This observation, together with our earlier observation that occurrent reasoning is the only reasoning which can give rise to personal justification, entails that beliefs are the only internal states with a role to play in personal justification, which just is the doxastic assumption. Since direct realism denies the doxastic assumption and maintains that nondoxastic internal states have a role to play in justification, as a theory of personal justification it is false. Consequently, the correct theory of personal justification (if there is one) must be a doxastic theory.

Foundation theories are doxastic theories which posit the existence of an epistemically privileged subclass of basic beliefs, beliefs which a person is immediately justified in believing. Being doxastic theories, they hold promise for providing the correct account of personal justification. Nevertheless, in order for any foundation theory to be correct, there must, in fact, be basic beliefs. Are there any basic beliefs? One very natural line of argumentation suggests that there must be. This aptly named "Regress Argument"[11] begins with two seemingly innocuous assumptions:

---

[11]The earliest formulation of the regress argument is

A1. S is justified in believing that p iff either (1) S is immediately justified in believing that p or (2) S is mediately justified in believing that p.

A2. S is mediately justified in believing that p iff S comes to believe that p on the basis of some other belief q which S is justified (either mediately or immediately) in believing.

According to A1, for any belief B_1, if S is justified in believing B_1, then either (1) B_1 is basic or (2) S is mediately justified in believing B_1. If (1), then obviously basic beliefs exist. If (2), then per A2, S must believe B_1 on the basis of some other belief B_2 which S is justified in believing. Of course, if S is justified in believing B_2, then either B_2 is basic or S is mediately justified in believing B_2 by believing it on the basis of B_3 which S is justified in believing, and so on. So, given A2, S is mediately justified in believing B_1 only if she believes B_1 on the basis of a chain of reasoning like the one just described. As we trace it backwards, S's chain of reasoning

---

due to Sextus Empiricus. For more contemporary formulations see: William Alston, "Two Types of Foundationalism", The Journal of Philosophy, Vol. LXXIII, No. 7 (April 8, 1976); Laurence Bonjour, "Can Empirical Knowledge Have a Foundation?", American Philosophical Quarterly, Vol. 15, No. 1 (January, 1978) and "Externalist Theories of Empirical Knowledge", Midwest Studies in Philosophy, ed. French, Vehling, and Wettstein (University of Minnesota: 1980); and Ernest Sosa, "The Raft and the Pyramid: Coherence versus Foundationalism in the Theory of Knowledge", op. cit.

must ultimately do one of four things:

(i) The chain stops at some arbitrary belief B_n which S is not justified in believing.

(ii) The chain continues backward infinitely, each belief-link of which S is mediately justified in believing.

(iii) The chain eventually loops back on itself forming a circle.

(iv) The chain stops at a basic belief which S is immediately justified in believing.

Of these four types of chains, the regress argument concludes that only type (iv) chains are genuinely justification-conferring. For consider type (i) chains. Since S is not justified in believing B_n, A2 entails that S is not (mediately) justified in believing B_{n-1}. Since S is not justified in believing B_{n-1}, S is not justified in believing B_{n-2}, and so on for each succeeding belief-link in the chain, including the terminal belief B_1. We might make the point as follows. A2 entails that chains of reasoning are only justification-transmitting, not justification-generating.[12] Since S is not justified in the initial belief B_n, there is no justification for the chain

---

[12]Chains of reasoning supposedly transmit justification in much the same way that valid arguments preserve truth. For an in depth discussion of justification transmission vs. justification generation, see James Van Cleve's "Foundationalism, Epistemic Principles, and the Cartesian Circle", The Philosophical Review, Vol. 88, No. 1 (1979).

to transmit, and consequently, S is not justified in believing the terminal belief B_1. Therefore, type (i) chains cannot justify S in holding the target belief B_1.

The problem with type (ii) chains is straightforward, once we recall that occurrent reasoning is the only kind of reasoning capable of conferring personal justification. A type (ii) chain cannot justify S in believing B_1 for the simple reason that it is humanly impossible for S to proceed occurrently through infinitely many steps of reasoning. Thus, type (ii) chains fail to be justification-conferring, as well.

Type (iii) chains, regardless of their length, fare no better as sources of mediate justification, because either they fail to be justification-conferring or they are unnecessary. We can see this by considering a very short type (iii) chain, e.g. B_1, B_2, B_3, B_4, B_1. Such a chain justifies S in believing B_1, only if S is already justified in believing B_1. For if S is not justified in believing initial belief B_1, then what we have is essentially a type (i) chain with no justification to transmit, and as a result, S is not justified in believing B_4 nor any other belief-link in the chain, including the terminal belief B_1. On the other hand, if S is already justified (either immediately or mediately via some other chain) in believing B_1, then the circular chain of reasoning is superfluous and does no epistemic work. Consequently, type (iii) chains are

either unnecessary or unsuccessful, as sources of justification.

Unlike the first three types of reasoning chains, type (iv) chains appear to be genuinely justification-conferring. After all, in a type (iv) chain, the initial belief of the chain is a basic belief which S is immediately justified in believing. Since chains of reasoning are justification-transmitting and since S is justified in believing the initial belief of the type (iv) chain, S's justifiedness is transmitted through each belief-link in the chain. As a result, S is justified in believing each belief-link, including the terminal link B_1. Therefore, of the four possible chain-types, only type (iv) chains succeed in conferring personal justification.

The last step of the regress argument is to show that, in light of the above considerations, basic beliefs must exist. If people are ever justified in believing what they do, then, since type (iv) chains are the only reasoning chains capable of conferring justification, it follows that basic beliefs must exist. To see why, recall our earlier observation based on A1 that for any belief B_1, if S is justified in believing B_1, then either (1) B_1 is basic or (2) S is mediately justified in believing B_1. Obviously, (1) entails that basic beliefs exist. But, since type (iv) chains are the only chains which confer mediate justification, (2) entails the existence of basic beliefs,

as well, because if (2), then it must be the case that S believes B_1, on the basis of a type (iv) chain, a chain whose initial belief is basic. So if either (1) or (2), then basic beliefs exist. Therefore, if the regress argument is sound, then basic beliefs must exist if personal justification is ever to obtain.

In light of the regress argument, should we conclude, as many philosophers have, that foundationalism is correct? I do not think we should. For while it is true that the regress argument provides the principal motivation for foundationalism, it does not, by itself, entail that foundationalism is correct. It only entails that basic beliefs exist, and although the existence of basic beliefs is a necessary condition for the correctness of foundationalism, it is not sufficient for such correctness. In order to see that the existence of basic beliefs is not alone sufficient for foundationalism, we need to examine the role which foundation theories assign to basic beliefs.

According to foundation theories, basic beliefs serve as the epistemic foundation or ground needed to justify us in holding nonbasic beliefs, a foundation from which <u>all</u> of our justification for nonbasic beliefs must ultimately derive. In light of the foundational role assigned to basic beliefs, two conditions must be satisfied if a foundation theory is to work: (A) Basic beliefs must exist, and (B) There must be enough basic beliefs to provide us with an

adequate foundation for our nonbasic beliefs.[13] Since the regress argument only entails (A), it does not entail that a foundation theory is correct. So, if we want to determine whether a foundation theory is correct, we must look beyond the regress argument to the theory itself to see whether it provides an adequate number of basic beliefs.

Whether or not a given foundation theory provides enough basic beliefs depends on how it answers the following two questions:

Q1. Which beliefs are basic?

Q2. What is it about those beliefs that makes them basic?

Two different answers have been given to Q1. Some foundation theorists contend that the set of basic beliefs is made up of very simple perceptual beliefs about physical objects. One such physical object belief is the belief that there is a green object in front of me. Supposedly, these beliefs are so safe that a person is immediately justified in holding them. Other foundation theorists, fearing that simple physical object beliefs are not properly suited for the role of foundational beliefs, maintain that the only beliefs which are actually basic are appearance beliefs,

---

[13]William Alston makes a similar observation in his "Two Types of Foundationalism", op. cit., p. 166.

i.e. beliefs about the way we are being appeared to. The belief that I am currently being appeared to as if there were a green object in front of me and the belief that I am now being appeared to greenly are examples of appearance beliefs. In order to decide which of these two proposals (if either) is correct, we first need to know the answer to Q2.

An answer to Q2 must tell us what property basic beliefs have and nonbasic beliefs lack in virtue of which the former are basic. The only plausible answer that has been offered is that basic beliefs have the property of being incorrigible and it is their incorrigibility which makes them basic. What is it for a belief to be incorrigible? As Pollock tells us, "A belief is incorrigible for a person S iff it is [logically] impossible for S to hold the belief and be wrong."[14] The rationale for selecting incorrigibility as the property which makes beliefs basic is straightforward. Foundation theories require that all our justification for our nonbasic beliefs ultimately rests on the epistemic foundation provided by our basic beliefs. If this foundation gives way, then our all mediate justification collapses with it. If we are,

---

[14]John Pollock, Contemporary Theories of Knowledge, op. cit., p. 34.

therefore, to avoid such justificatory vulnerability, the epistemic foundation from which we reason must be absolutely unshakeable. A foundation that consists exclusively of incorrigible beliefs is entirely unshakeable, because, since it is impossible for a person to hold incorrigible beliefs and be mistaken, it is ipso facto impossible for any belief in a foundation of incorrigible beliefs to be false. Let us assume that the incorrigibility account of basicness is correct.[15] We are now in a position to see that neither of the proposed basic belief candidates is satisfactory.

Two problems beset the simply physical object belief proposal. First, we are lacking an account of simplicity which tells us when a physical object belief is simple enough to be basic. Presumably, the belief that there is a green object in front of me is simple enough to qualify as basic. But is the belief that there is a green plant a few feet in front of me simple enough to be basic? What about

---

[15]Actually, the incorrigibility account of basicness is open to two devastating objections. The standard objection concerns necessary truths. Since it is impossible to believe a necessary truth and be mistaken, it follows from the incorrigibility account that every necessary truth which a person happens to believe is basic. This, of course, includes complex mathematical truths. But, the objection continues, surely a person is not immediately justified in believing every complex mathematical truth that he happens to stumble upon. To the contrary, a person is justified in believing such a complex mathematical truth only if he derives that truth from the axioms of the system in question. Therefore, not every incorrigible belief is

the belief that there is a four foot tall Dracaena masangeana next to the grand piano? Surely, no foundation theorist would want to count this last belief as basic. But without an account of simplicity, it is difficult to see how such a belief can be ruled out. In short, the foundation theorist must provide us with a simplicity account if the physical object belief proposal is to be at all tenable.

Second, even with a simplicity account, the physical object belief proposal remains unsatisfactory, since physical object beliefs are not incorrigible.[16] For

---

basic.

The second objection is even more devastating, because it undercuts the incorrigibility account even if the account is limited to contingent truths. Obviously, the average person has no idea what an incorrigible belief is, and as a result, when Mr. Average holds an incorrigible belief, he has absolutely no means of recognizing it as such. What this shows is that a person can hold an incorrigible belief without realizing that the belief is incorrigible. Suppose S holds the incorrigible belief B, but does not realize that B is incorrigible. Also suppose that S thinks that B is unreasonable, but persists in believing B anyway. According to the incorrigibility account, since B is incorrigible, S is immediately justified in believing B. However, since S thinks B is unreasonable, intuitively S is unjustified in believing B. Since the incorrigibility account entails that S is justified in believing B, when in fact he is not, as an account of basicness it is false.

I have chosen to ignore these problems in the body of the text, because the incorrigibility account, which is the only even remotely plausible account of basicness, actually serves to undermine foundation theories.

[16] I can only think of one physical object belief which might be incorrigible. If some version of the identity theory is correct, then the cogito belief that I exist is a physical object belief which looks to be incorrigible.

example, consider the belief that there is a red object in front of me. Since I may believe that there is a red object in front of me, when, in fact, there is a white object bathed in red light in front of me or when I am hallucinating and there is really no object in front of me, it is possible for me to hold the belief and be mistaken. More generally, the demon hypothesis demonstrates that the possibility of error extends to (virtually) all of our physical object beliefs. Since it is possible to hold physical object beliefs but be mistaken, such beliefs are not incorrigible. Since physical object beliefs are not incorrigible, they lack the property requisite for basicness, and thus, they are not basic.

At first glance, appearance beliefs look to be much more promising basic belief candidates, because unlike physical object beliefs, appearance beliefs seem to be incorrigible. While it is easy to be mistaken about the way things <u>are</u>, it is hard to see how we could be mistaken about the way it <u>seems</u> to us that things are. If it really is impossible to hold appearance beliefs and be wrong, then appearance beliefs are incorrigible, making them natural candidates for basic beliefs.

Despite its prima facie plausibility, the appearance

However, if this is our only basic belief, then we do not have an adequate foundation for our nonbasic beliefs, Descartes' meditations to the contrary notwithstanding.

belief proposal is open to three devastating objections. Perhaps the most obvious objection is that we rarely have appearance beliefs. When we walk outside, for example, we form beliefs about cacti, flowers, hummingbirds, clear skies, snow capped mountains, and the like. We do not find ourselves forming beliefs about our phenomenal states, such as the belief that I am being appeared as if there is a saguaro in front of me or the belief that I am experiencing a sensation of yellow. It is not that we are unable to form beliefs of the latter sort. It is just that we usually do not. Simply put, appearance beliefs are the exception, not the rule. As a result, we have far too few appearance beliefs to constitute an adequate foundation.[17]

Another problem with the appearance belief proposal is that appearance beliefs are incapable of performing the very function which foundational beliefs are supposed to perform, namely, the function of justifying us in holding nonbasic beliefs, including nonbasic beliefs about the world. It should be obvious that appearance beliefs alone are incapable of justifying us in believing anything about the external world. The appearance belief that I am currently experiencing a sensation of yellow cannot, by itself,

---

[17]John Pollock raises this same objection in Chapter Two of his Contemporary Theories of Knowledge, op. cit.

justify me in believing that there is something yellow before me. It can only do so if it is conjoined with the belief that things are the way they appear to be. Of course, according to foundation theories, I must be justified in holding this latter belief, if it is to help justify me in believing something else, and since it is not basic, I must be mediately justified in believing it. However, no amount of appearance beliefs can justify me in believing that things are the way they appear to be, and so, I am not mediately justified in believing it (for remember, appearance beliefs are all I have to go on at this stage). But without the belief that things are the way they appear to be, appearance beliefs cannot justify me in believing anything about the way things are. Thus, a foundation of appearance beliefs is totally useless, because no nonbasic beliefs can be built up from it.

Although I regard the previous objection as decisive, the appearance belief proposal is open to yet another equally devastating objection, to wit, appearance beliefs are not incorrigible. Recent psychological experiments have demonstrated that people do make mistakes concerning how they are being appeared to. Paul Churchland describes the results of one such experiment as follows:

> An orange-expectant subject fed lime sherbet may confidently identify her taste-sensation as being the kind normally produced by orange sherbet, only to retract the

> identification immediately upon being given a (blind) taste of the genuinely orange article. Here one corrects one's qualitative identification, in flat contradiction to the idea that mistakes are impossible. Mistakes of this kind are called expectation effects, and they are a standard phenomenon with perception generally.[18]

Expectation effects can even adversely affect our ability to detect color appearances, for, as Pollock tells us,

> A discovery made fairly early by every landscape painter is that snow looks blue (particularly the shadows). Most people think snow looks white to them, but they are wrong.[19]

Presumably, if you ask these people to pay closer attention to the way in which they are being appeared to, they will realize that the snow appears to have a bluish hue and will correct their appearance judgments accordingly. Examples of similar mistakes in appearance judgments abound. Just think of the last time you accidently touched a cold stove, which you expected to be hot. You may have believed that you experienced a burning sensation, when in fact you did not. These examples demonstrate that a person can hold an

---

[18]Paul Churchland, Matter and Consciousness (Cambridge, Mass. and London, England: The MIT Press, 1984), p. 77.

[19]John Pollock, Contemporary Theories of Knowledge, op. cit., p. 70.

appearance belief and be mistaken. This, in turn, demonstrates that appearance beliefs are not incorrigible and, therefore, do not qualify as basic.[20]

In short, neither simple physical object beliefs nor appearance beliefs are properly suited for the role of basic beliefs. Simple phyical object beliefs are not incorrigible, because physical object beliefs in general are not incorrigible, thus making them unsuitable as basic beliefs. Appearance beliefs are ill-suited for the role of basic beliefs in three respects: (1) There are too few of them to provide us with an adequate epistemic foundation, (2) They are incapable of justifying us in holding nonbasic beliefs, and (3) They are not incorrigible. Thus, to date the foundation theorist has not been able to furnish us with a single basic belief. Moreover, since simple physical object beliefs and appearance beliefs are the only even remotely plausible candidates for basic beliefs, it is very unlikely that the foundation theorist will be able to provide us with any basic beliefs in the future.

If there are no basic beliefs, then foundation theories are false. The fact that no foundation theorist has ever been able to generate any basic beliefs strongly suggests

---

[20]John Pollock raises the same objection to appearance beliefs in Chapter Two of his Contemporary Theories of Knowledge, op. cit.

that there are no basic beliefs. But we can go further and demonstrate that there are no basic beliefs by showing that the very idea of basic beliefs is absurd. Recall that a basic belief is a belief which a person is immediately justified in believing. Also recall that a person is justified in holding a belief iff she deserves epistemic praise for holding that belief. Therefore, for a belief B to be basic, it must be the case that a person immediately deserves epistemic praise for holding B —— though she has no reasons or evidence of any kind for B and has not consciously reflected on or questioned B at all —— simply in virtue of the fact that she holds B. Once we realize that this is what is required for basic beliefs to exist, it seems obvious that there are no basic beliefs. How could there be? After all, whether or not a person deserves epistemic praise for holding a given belief has nothing to do with the particular belief she happens to be holding, but rather it has to do with her —— it has to do with whether or not she has been epistemically responsible in coming to hold the belief. Since a person can be epistemically irresponsible in holding any belief whatsoever, it follows that there are no beliefs which a person automatically deserves epistemic praise for believing, just because she happens to believe them. Of course, this entails that there are no beliefs which a person is automatically justified in believing, just because she believes them. Hence, there are

no basic beliefs.

The fact that there are no basic beliefs entails that foundation theories are false. It also entails that the regress argument, which has been so influential in contemporary epistemology, is unsound. Although the regress argument is unsound, we do not as yet know why. In section 6, the regress argument will be reexamined, to see if we can pinpoint where it goes wrong. Our present task, however, is that of providing an account of personal justification.

We began with the observation that the correct account of personal justification could only be given by an internalist theory. Nondoxastic internalist theories, like direct realism, fell by the wayside because the intellectualist model is correct for personal justification. That only left us with doxastic theories to consider. We have just seen that all foundation theories of personal justification are false, because there are no basic beliefs. Since the class of doxastic theories is exhausted by foundation theories and coherence theories, it follows that <u>if</u> there is a correct account of personal justification, then it must be given by a coherence theory.

It is important to note the hypothetical nature of the conclusion just reached, so as to avoid the all too common mistake of taking an argument to prove more than it does. Strictly speaking, the argument from elimination just offered <u>only</u> proves that all noncoherence theories are

false. It does not, in and of itself, prove that a coherence theory of personal justification is correct. It is, after all, possible that every coherence theory is false, in which case we would be left with a "no account" account of personal justification. Since there is no a priori way to rule out the "no account" account possibility, a positive argument needs to be offered to show that a coherence theory of personal justification is correct. The next two sections are devoted to presenting such an argument. In Section 4, I will discuss both the nature of reasons and the linear picture of reasoning. In Section 5, I will develop a coherence theory which incorporates such linear reasons and will show that the personal justification evaluations yielded by this theory are intuitively correct. This, in turn, will show that a coherence theory does provide the correct account of personal justification.

4. Reasons and Linear Reasoning

Coherence theories are doxastic theories which deny the existence of an epistemically privileged subclass of basic beliefs. Being doxastic theories, they maintain that beliefs are the only internal states that a person can appeal to in order to justify her in holding other beliefs. In short, coherence theories hold that justification proceeds from reasoning, where beliefs are the only internal

states capable of being reasons. In order for such a theory to be correct, it must provide us with an account of reasoning which explains how a person's beliefs function to justify her in holding other beliefs, and the account must be correct. Traditionally, coherence theories have provided us with a holistic account of reasoning, according to which a person reasons from her set of beliefs in toto. In what is to follow, I shall argue that such an account of reasoning is psychologically unrealistic and, hence, empirically false and that, therefore, a coherence theory must adopt an alternative account of reasoning, if it is to provide the correct analysis of personal justification. Then, I shall present a linear account of reasoning as the required alternative.

Every coherence theory offered to date has assumed a holistic account of reasoning. I shall refer to such theories as "holistic coherence theories".[21] One of the major difficulties with evaluating holistic coherence theories is that most holistic coherence theories have not been worked out in much detail.[22] Typically, we are only

---

[21] I borrow this terminology from John Pollock, who coins it in Chapter Three of his Contemporary Theories of Knowledge, op. cit.

[22] A notable exception is Keith Lehrer's theory, which I shall consider shortly. See Keith Lehrer, Knowledge, op. cit.

given a sketch of a holistic coherence theory, as, say, a theory which rests on the Neurath metaphor. Since such sketches resist proper evaluation, let us, instead, start by examining a generic holistic coherence theory of personal justification:

> (GT) S is personally justified in believing that p iff S reasons from the set of her beliefs in toto to the belief that p.

A special problem arises for (GT) when S already believes that p, because then the belief that p belongs to the set of S's beliefs, and so, ipso facto the set of S's beliefs entails that p. In such a situation, (GT) seems to entail that S could come to be justified in believing that p simply by realizing that she already believes that p and that, therefore, the set of her beliefs entails that p. I shall not pursue the prior belief problem, however, because (GT) could probably be reformulated in such a way as to avoid the problem and because the resulting reformulation would still be open to the next objection.

Let us consider a case where S, who as yet does not believe that p and does not believe any conjunction containing p as a conjunct, is considering whether or not to believe that p. According to (GT), in order for S to be justified in believing that p, she must come to believe that p by reasoning from her entire set of beliefs to the belief that p. To see the problem with this requirement, recall

from section 3 that the intellectualist model, according to which occurrent reasoning is the only kind of reasoning capable of conferring justification, is correct for personal justification. Given the correctness of the intellectualist model, we must interpret (GT) as asserting that S is personally justified in believing that p just in case S comes to believe that p by occurrently reasoning from the entire set of her beliefs to the belief that p. But on this interpretation, (GT) and its account of reasoning are psychologically unrealistic. (GT) is psychologically unrealistic as a requirement for personal justification, because it is impossible for us to reason occurrently from our entire set of beliefs. (GT)'s account of reasoning is psychologically unrealistic as a description of how people reason, because people never reason from the set of all their beliefs. I will consider each of these problems in turn.

We can see that (GT) is an unrealistic requirement for personal justification in two different ways. First, in order for a person to reason occurrently from any set of beliefs to some other belief, she must first know what the members of that set are. So, in order for a person to occurrently reason from the set of all her beliefs, she would have to know what all her beliefs were. Of course, we all hold some beliefs which we are unaware of holding, e.g. prejudicial beliefs, and thus, we do not know what all our

beliefs are. Therefore, we cannot occurrently reason from the set of all our beliefs. Second, even if we were to know what all our beliefs were (in some dispositional sense), we still would not be able to reason occurrently from the set of those beliefs, because, as Pollock points out, current psychological evidence indicates that people can hold no more than seven occurrent beliefs in the mind at one time and the number tends to diminish as the beliefs become more complicated.[23] Since we can occurrently reason from beliefs only if we occurrently hold those beliefs and since we can only occurrently hold seven (or less) beliefs, we can only occurrently reason from seven (or less) beliefs. Since the set of our beliefs is extremely large, but we can only occurrently reason from seven (or less) beliefs, we obviously cannot occurrently reason from the set of all our beliefs. Therefore, (GT) presents us with an unfulfillable and, hence, unrealistic requirement for personal justification.

(GT)'s account of reasoning is equally psychologically unrealistic, if it is intended as a description of how we do, in fact, reason. This probably seems obvious, since we have just seen that it is impossible for us to reason

---

[23]John Pollock, Contemporary Theories of Knowledge, op. cit., p. 53.

occurrently from the set of all our beliefs. Nevertheless, the present objection does not depend on the impossibility of such reasoning, because even if we could occurrently reason from all our beliefs, in point of fact we never do reason that way. Instead, we reason from individual beliefs (or from very small sets of beliefs) to other beliefs. Thus, (GT)'s account of reasoning is empirically false.

Since we never do reason from the set of our beliefs _in toto_, we never satisfy the _analysans_ of (GT). Consequently, either we are never personally justified in holding any of our beliefs or (GT) is false. Since occasionally some of us do deserve epistemic praise for holding some of our beliefs, occasionally some of us are personally justified in holding some of our beliefs. Thus, it is false that we are never personally justified in holding any of our beliefs, and so, it follows by disjunctive syllogism that (GT) is false. Of course, we probably should have expected as much, since (GT) violates the "ought implies can" dictum by requiring us to carry out the impossible task of occurrently reasoning from all our beliefs.

The failure of (GT) is instructive. It shows that those proposing sketchy holistic coherence theories were wrong to assume that something like (GT) would work. Since (GT) does not work, if we are to take the possibility of a holistic coherence theory seriously, we must be given more than a rough sketch of the theory. We must be given a

concrete proposal that has been worked out in careful detail. Lehrer is the only philosopher to have developed such a detailed holistic coherence theory. Moreover, his theory looks much more promising than the generic theory, since, at least on one interpretation, it does not require us to reason occurrently from the set of all our beliefs and also appears to avoid the prior belief problem. Let us, therefore, turn to his theory as he presents it in Knowledge to see whether it can provide the needed account of personal justification.[24]

In Knowledge, Lehrer maintains that S is completely justified in believing that p iff p coheres with S's corrected doxastic system, where S's corrected doxastic system is a set of statements of the form, S believes that q, S believes that r, and so on, which describes what S believes for the purposes of gaining truth and avoiding error.[25] What is it for a statement p to cohere with such a

---

[24]Lehrer has since abandoned the theory he formulated in Knowledge and has replaced it with the Monster Theory. Nevertheless, I am presenting the earlier theory, because it is a coherence theory, whereas the Monster Theory is an externalist theory. We have already seen that no externalist theory (nor any other noncoherence theory) can provide the correct account of personal justification. Unlike the Monster theory, his earlier theory, being a coherence theory, remains a viable candidate for a theory of personal justification.

[25]For the original presentation of Lehrer's coherence theory, see Chapter 8 of Knowledge, op. cit.

system? In order to answer this question, Lehrer introduces the notion of a "competitor". Simplifying a bit, p competes with q for S iff S believes that q is negatively relevant to p, where q is negatively relevant to p iff p is less probable on the assumption that q than it is without that assumption, i.e. q is negatively relevant to p iff prob (p/q) < prob (p). In short, when S believes that q is negatively relevant to p, then p competes with q, and hence, q is a competitor of p. Armed with this account of competitors, Lehrer provides an explication of the coherence relation, to wit, p coheres with S's corrected doxastic system iff, within S's corrected doxastic system, p is believed to be more probable than any of its competitors. Having explicated the relation of coherence, Lehrer offers the following analysis of complete justification:

> (LJ) S is completely justified in believing that p if and only if, within the corrected doxastic system of S, p is believed to have a better chance of being true than the denial of p or any other statement that competes with p.[26]

Unfortunately, the analysans of (LJ) is ambiguous. Its ambiguity can be captured using the two reformulations of (LJ) that follow:

---

[26]Keith Lehrer, Knowledge, op. cit., p. 198.

(LJ1) S is completely justified in believing that p iff S believes, on the basis of her corrected doxastic system, that p is more probable than any of its competitors.

(LJ2) S is completely justified in believing that p iff for the purposes of gaining truth and avoiding error S believes that p is more probable than any of its competitors, i.e. iff the statement "S believes that p is more probable than any of its competitors" is a member of S's corrected doxastic system.

(LJ1) fares somewhat better than (GT), since it is immune to the prior belief problem.[27] Nevertheless, (LJ1) remains open to the other objections facing (GT). That (LJ1) is a psychologically unrealistic requirement for personal justification can be seen in several ways. First, in order for S to base her belief —— that p is more probable than any of its competitors —— on her corrected

---

[27]To see that (LJ1) is immune to that prior belief problem consider the following: Let p = the chair is red, and let q = a red light is shining on the chair. q is negatively relevant to p, because the probability that the chair is red given that there is a red light shining on the chair is less than the probability that the chair is red without such an assumption, i.e., prob (p/q) < prob (p). Suppose that, for the purposes of gaining truth and avoiding error, S believes that p, believes that q, believes that q is a competitor of p, and believes that q is more probable than p. In the situation just described, S already believes that p, but since, relative to S's corrected doxastic system, S does not believe that p is more probable than its competitor q, S is not completely justified in believing that p. Thus, prior belief does not entail justification.

doxastic system, S would have to know what the member-statements of her corrected doxastic system were. If S is like most people, then she does not even know what a corrected doxastic system is, and so, she obviously does not know what the member-statements of her corrected doxastic system are. However, even if S were an epistemologist and did know what a corrected doxastic system is, she still would not know the membership of her corrected doxastic system, because the set of statements describing what she believes (for the purposes of gaining truth and avoiding error) is much too large a set for her (or anyone else) to comprehend. Since S does not know what the member-statements of her corrected doxastic system are, she cannot base any beliefs on her corrected doxastic system, much less the belief that p is more probable than any of its competitors. Second, even if S were somehow to know the membership of her corrected doxastic system, she still could not base any beliefs on the entire system, because, given human psychological constraints, S can only occurrently attend to seven (or less) of the member-statements of her corrected doxastic system. At best, S could only base beliefs on a small subset of her corrected doxastic system. Since (LJ1) requires S to base beliefs on her corrected doxastic system taken as a whole, it represents an unrealistic requirement.

(LJ1) is also unrealistic as a description of how

people do, in fact, reason. Even if we could reason from our entire corrected doxastic system, as a matter of fact we never do. We never even reason from small subsets of our corrected doxastic system. We need only introspect to realize that we base beliefs on other beliefs, not on statements describing what we believe for the purposes of gaining truth and avoiding error. I contend that the only way such statements can play a role in belief-formation is if we believe those statements. If we do not believe those statements, then it is not at all clear how they could affect our cognitive reflections. If we do believe those statements, then it seems that it is our beliefs of those statements, rather than those statements per se, which are justifying us in holding other beliefs. But even if we take (LJ1) to be suggesting that we base beliefs on beliefs of statements describing what we believe for the purposes of gaining truth and avoiding error, (LJ1) remains unrealistic, because people rarely have such meta-beliefs. In short, no matter how we take it, (LJ1) rests on an unrealistic picture of reasoning. Consequently, the principle of charity dictates that we should not saddle Lehrer with the (LJ1) view, unless it is absolutely necessary. To see whether it is necessary, let us consider (LJ2).

(LJ2) is a more plausible proposal than (LJ1), since it avoids all of the objections raised against (GT). We can see that (LJ2) is not subject to the prior belief problem by

considering Tommy the toddler. Tommy has not yet acquired the concept of probability, but he does believe that p. Tommy also believes that q, and q is negatively relevant to p (though Tommy does not know that it is) and, hence, q is a competitor of p. Although Tommy already believes that p, he does not believe that p is more probable than q. In this situation, (LJ2) entails that Tommy is not completely justified in believing that p. Therefore, on the (LJ2) view, prior belief does not entail complete justification.

That (LJ2) is immune to the other two objections facing (GT) can be seen as follows: According to (LJ2), for S to be completely justified in believing that p, S need only believe (for the purposes of gaining truth and avoiding error) that p is more probable than any of its competitors. S does not need to know what the member-statements of her corrected doxastic system are. She does not even need to know that she has a corrected doxastic system. And most important of all, she does not need to base her belief —— that p is more probable than any of its competitors —— on her corrected doxastic system. The only thing that (LJ2) requires of S is that S believe that p is more probable than any of its competitors, and this is a requirement she can fulfill. Moreover, people do sometimes reason this way. For example, it is not uncommon for a scientist to compare competing alternatives and decide that one alternative is more probable than the others and then come to believe that

alternative. Thus, unlike (GT) and (LJ1), (LJ2) requires us to reason in a way in which people do in fact sometimes reason.

In light of the foregoing, I submit that (LJ2) is the most plausible holistic coherency theory going. Nevertheless, it too is ultimately unsatisfactory, for while people do occasionally compare alternatives and believe that which they take to be most probable, they do not do so very often. To the contrary, we rarely have beliefs of the form that p is more probable than each of its competitors.[28] When we do have such beliefs, they may, in fact, justify us in believing that p, but just because we lack such a belief does not entail that we are not justified in believing that p. In short, while (LJ2) may provide us with a sufficient condition for personal justification, it does not provide us with a necessary one, because there are more ways for a person to be epistemically responsible than just ruling out competitors. Consequently, we must look elsewhere for a complete account of personal justification. That elsewhere starts with the linear picture of reasoning which I shall now detail.

My account of linear reasoning draws heavily from

---

[28]John Pollock raises the same objection to Lehrer's theory in Chapter Three of his Contemporary Theories of Knowledge, op. cit., p. 89.

Pollock, who has formulated a detailed account of reasoning in Chapter Two of his Contemporary Theories of Knowledge. While I do not agree with his account entirely, I find much of it to be correct. For this reason, I will begin by presenting his account of linear reasoning. I will then modify his account somewhat.

Since occurrent reasoning proceeds in terms of reasons, we need to know precisely what a reason is. Pollock tells us, "A belief P is a reason for a person S to believe Q iff it is logically possible for S to become justified in believing Q by believing it on the basis of P".[29] There are two kinds of reasons, conclusive reasons and nonconclusive reasons. A conclusive reason is a reason which entails that for which it is a reason. Nonconclusive reasons rationally support, but do not entail, that for which they are reasons. Pollock contends that "The most important characteristic of nonconclusive reasons is that they are defeasible."[30] A reason is defeasible if the evidence it provides for its conclusion can be defeated by certain other evidential considerations. For example, if every mammal that we have examined has had hair, this provides us with a reason for thinking that all mammals have hair. However, if we later

---

[29] John Pollock, Contemporary Theories of Knowledge, op. cit., p. 41.

[30] Ibid., p. 42.

discover a mammal with no hair, this new information defeats our earlier reason for thinking that all mammals have hair and, thus, constitutes a defeater for that earlier reason. Pollock defines a defeater as follows: "If P is a reason for S to believe Q, R is a defeater for this reason iff R is logically consistent with P and (P&R) is not a reason for S to believe Q."[31] Pollock calls a reason for which there can be defeaters a "prima facie reason".

Pollock indentifies two kinds of defeaters for prima facie reasons, to wit, rebutting defeaters and undercutting defeaters. He defines the former as "If P is a prima facie reason for S to believe Q, R is a rebutting defeater for this reason iff R is logically consistent with P and R is a reason for S to believe $\sim$Q."[32] Thus, a rebutting defeater is a reason for denying the conclusion supported by the prima facie reason. In the mammal example, discovering a mammal with no hair is a rebutting defeater, since it is a reason for thinking that it is false that all mammals have hair.

Whereas rebutting defeaters attack the conclusion of the prima facie reason, undercutting defeaters attack the connection between the prima facie reason and its

---

[31] Ibid.

[32] Ibid., p. 43.

conclusion. For an example of an undercutting defeater, consider the following: Suppose the boxer Killer Queen has won all of his previous 57 fights. This gives us a reason for thinking Killer Queen will win his next fight. If, however, we discover that all of his previous fights were ten rounders and his next fight is a scheduled fifteen rounder, this information undercuts our earlier prima facie reason, because what a fighter can do in ten rounds is not an accurate indicator of what he can do in fifteen rounds. Notice this new information does not give us a reason for thinking Killer Queen will lose. It just gives us a reason for denying that he would not have won all of his previous fights unless he were a good enough fighter to win his next fight. Pollock defines such a defeater as follows: "If P is a prima facie reason for S to believe Q, R is an undercutting defeater for this reason iff R is logically consistent with P and R is a reason for S to deny that P wouldn't be true unless Q were true."[33]

Since defeaters are themselves reasons, they may be subject to defeaters, as well. If a defeater R is itself defeated, this blocks its action as a defeater and, thus, reinstates the prima facie reason P for which R is a defeater (unless, of course, R's defeater is defeated, in

---

[33] Ibid.

which case R is reinstated, thereby redefeating P). In short, there can be defeater defeaters and defeater defeater defeaters, and so on.

Given this account of reasons and defeaters, a linear picture of reasoning naturally emerges. In order for S to be justified in believing that q, S must base his belief that q on a chain of reasons. He may, for example, base his belief that q on reason p_3 which he based on p_2 which was based on p_1. If S reasons in this way, he instantiates the following chain of reasoning:

$$p_1 \rightarrow p_2 \rightarrow p_3 \rightarrow q$$

In virtue of having proceeded through such a chain of reasoning, S is presumably justified in believing that q, unless, of course, S believes an undefeated defeater for one of the reasons leading to q.

I think that this linear picture of reasoning accurately depicts how we do in fact reason.[34] Even so, Pollock's account needs to be modified somewhat. Pollock maintains that nonconclusive reasons are the only reasons capable of being defeated. I, on the other hand, contend that all reasons are subject to being defeated by evidential

---

[34]In fact, Lehrer's coherence theory, on the (LJ2) reading, looks to be an isolated instance of this more general theory, because deciding that q is more probable than all of its competitors is just one way of ruling out potential defeaters.

considerations. To defend this contention, I must introduce a third kind of defeater, which does not satisfy Pollock's definition of defeaters per se. Suppose that p_1 in the chain above is a conclusive reason for S to believe p_2. In such a situation, it still might be the case that S believes that r_1, where r_1 is a reason for thinking $\sim p_1$, and although r_1 is not logically consistent with p_1, it certainly seems to defeat p_1 as a reason for p_2. We might call this kind of defeater a "negating defeater" and define negating defeaters as follows: If p is a reason (either conclusive or nonconclusive) for S to believe that q, r is a negating defeater for p iff r is a reason for thinking that p is false. While only nonconclusive reasons are subject to undercutting and rebutting defeaters, all reasons are subject to negating defeaters, and so, Pollock is wrong to claim that only nonconclusive reasons are capable of being defeated. Like other defeaters, negating defeaters are themselves subject to defeaters which can nullify their effect as defeaters. Hence, the linear picture of reasoning extends to negating defeaters, as well. Now that we have a linear account of reasoning before us, we can use it to formulate a linear coherence theory which correctly analyzes personal justification.

5. An Account of Personal Justification

a. A Tripartite Distinction

Before presenting my account of personal justification, I need to digress briefly to discuss the sorts of personal justification evaluations which can be made. Traditionally, personal justification evaluations have been made in an all or nothing fashion. As a result, on the traditional view, if S believes that p, then either S is justified in believing that p or S is unjustified in believing that p.

I think that the traditional view presents us with a false dichotomy. To see why, recall from Chapter 2 that to say that S is justified in believing that p is to say that S deserves epistemic praise for believing that p and to say that S is unjustified in believing that p is to say that S deserves epistemic blame for believing that p. Incorporating this observation into the traditional view yields the following result: If S believes that p, then either S deserves epistemic praise for believing that p or S deserves epistemic blame for believing that p. Such a result is clearly mistaken. Very often we hold beliefs for which we deserve neither praise nor blame. Perceptual beliefs (at least most of the time) are a case in point. Typically, our perceptual beliefs come to us as a result of automatic belief-forming cognitive processes over which we have little control. We certainly are not doing anything epistemically reprehensible in holding such beliefs, and so,

we are free from epistemic blame for holding them. On the other hand, since we generally do not reflect on our perceptual beliefs or carry out any kind of reasoning procedure to ensure that they are true, we do not deserve epistemic praise for holding them either. Therefore, properly speaking, we are neither justified nor unjustified in holding most of our perceptual beliefs, but rather are personally ajustified in holding them.

The same point can be noted by reconsidering (PJ_j') and (PJ_u'). According to (PJ_j') and (PJ_u'), personal justification is concerned with whether a person has been epistemically responsible or epistemically irresponsible in coming to hold a given belief. Consequently, when a person is neither epistemically responsible nor epistemically irresponsible in holding a particular belief, she falls outside the scope of personal justification. But for a person to fall outside the scope of personal justification just is for that person to be personally ajustified.

What this shows is that, when assessing a person's justificatory status, three evaluations are possible. The person may be personally justified in holding a certain belief. She may be personally unjustified in holding that belief. Or she may be personally ajustified in holding the belief. The failure to recognize this third possibility has led philosophers to create theories of justification which count persons as being justified, when in fact they are not.

In short, if we want to provide the correct account of personal justification, we must recognize this tripartite distinction of personal justification evaluations, especially since, if I am right, we are ajustified in holding most of our beliefs. In the next subsection, I shall offer an analysis of personal justification which takes these three evaluative possibilities into account.

b. A Linear Coherence Theory of Personal Justification

In section 4 of the present chapter, I presented a slightly modified version of Pollock's account of linear reasoning. According to Pollock's account, a person S is justified in believing that q only if S bases her belief that q on an ultimately undefeated chain of reasoning of the form:

$$P_1 \rightarrow P_2 \rightarrow \ldots \rightarrow P_n \rightarrow q,$$

where a chain of reasoning is ultimately undefeated for S iff either (1) there is no belief-link in the chain for which S believes a defeater or (2) there are some belief-links for which S believes defeaters, but these defeaters are themselves defeated by defeater defeaters which S believes. Of course, this only provides us with a necessary condition for personal justification. To get an analysis which is both necessary and sufficient for personal

justification, we must look further. Pollock thinks that in order for an ultimately undefeated chain of reasoning to justify S in believing its conclusion, S must be justified in believing each belief-link of the chain. Thus, Pollock concludes that S is jusitifed in believing that q iff S bases her belief that q on an ultimately undefeated chain of reasoning, each belief-link of which S is justified in believing.[35]

At this point, I part company with Pollock, since on a view such as his we must embrace basic beliefs in order to prevent an infinite regress of reasons. We have already seen in section 3 that there are no basic beliefs. So, we must look to a different theory, if we want to provide the correct account of personal justification. I will now set the stage for that different theory.

When we reason, we reason from our beliefs in general. We do not first partition our beliefs into those which we are justified in believing and those which we are not, and

---

[35] In Chapter Two of his Contemporary Theories of Knowledge, op. cit., Pollock formulates this conditon as "S is justified in believing P iff S instantiates an ultimately undefeated argument supporting P" (p.53), where an argument is "a finite sequence of propositions ordered in such a way that for each proposition P in the sequence, either (1) P is epistemologically basic or (2) there is a proposition (or set of propositions) earlier in the sequence that is a reason for P" (p. 51) and where "A person instantiates an argument iff he believes the propositions comprising the argument and he believes each nonbasic proposition in it on the basis of reasons for it that occur earlier in the argument" (p. 51).

then reason only from the former. To think that we do partition our beliefs in this way is unrealistic, since for the most part we do not even know which of our beliefs we are justified in believing and which of them we are not justified in believing. Instead, when we reason, we take our beliefs for granted and proceed from there. If you have doubts about this, consider the last time you occurrently reasoned to some conclusion q on the basis of some chain $p_1 \ldots p_n$. Did you stop to ask yourself if you were justified in believing each belief-link of the reasoning chain? I suspect that you did not. This just is not the way we usually reason. Rather, once a belief gets into our doxastic corpus, we simply regard it as belief-worthy and reason from it. We never even consider whether we are justified in believing it, unless some counterevidence to it comes to the fore.

In this regard, private reasoning is very much like public reasoning. When publically reasoning to (or arguing for) a conclusion, we simply state our premises. We do not attempt to prove our premises, unless someone calls them into question. If someone does call our premises into question, we offer other premises in support of our initial premises, which we do not attempt to prove unless they too are called into question. Similarly, when privately reasoning, we reason from our beliefs (our internalized premises). We do not question whether we are justified in

holding these beliefs, unless some counterevidence to them arises. If such counterevidence does arise, we look to other beliefs to justify us in continuing to hold those earlier beliefs (looking for defeater defeaters is one way of doing this), and we do not question whether we are justified in holding these other beliefs, unless counterevidence surfaces which makes them suspect, as well.

Once we realize that personal justification is a function of our beliefs and not a function exclusively of the beliefs we are justified in holding, a linear coherence theory becomes extremely plausible. Our chains of reasoning do not need to start with basic beliefs which we are in some mysterious way immediately justified in believing, as foundation theorists maintain, because our chains of reasoning do not need to start with beliefs which we are justified in believing. Instead, as coherence theorists rightly maintain, all of our beliefs are on equal epistemic footing, at least in the sense of being able to provide us with reasons for other beliefs, and so, constitute legitimate starting points from which to reason.

We are now in a position to formulate a linear coherence theory of personal justification. We have just seen that all of our beliefs are capable of providing us with reasons for other beliefs and, hence, can initiate justification-conferring chains of reasoning. This observation, together with our earlier observation that the

intellectualist model is correct for personal justification, suggests that a person is justified in holding a given belief just in case she comes to hold that belief by occurrently reasoning through an ultimately undefeated chain of reasoning which supports it, each belief-link of which she believes. Unfortunately, this suggestion is not quite right. Very often we find ourselves holding beliefs which were not the result of occurrent reasoning. On any adequate theory of personal justification, it must be possible for us to come to be justified in holding these beliefs, as well. Consequently, in order to provide the correct account of personal justification, we must modify the above suggestion in such a way that it satisfies this adequacy constraint. I submit that the following modification provides us with the correct analysis of personal justifiedness.

(PJ_j*) S is personally justified in believing that p iff either
(1)(a) S comes to believe that p by occurrently reasoning through an ultimately undefeated chain of reasoning which supports p, and
(b) S believes each belief-link of C; or
(2)(a) S comes to believe that p via a BCP other than occurrent reasoning,
(b) S occurrently notices an ultimately undefeated chain of reasoning which supports p, and

(c) S believes each belief-link of C.[36]

In short, (PJ_j*) asserts that a person is justified in holding a belief iff she has occurrently noticed that she has evidence (in the form of other beliefs) for that belief.[37] On this view, a person's evidence (or reasons) for p need not play a causal role in the genesis (or sustenance) of her belief that p in order for her to be justified in believing that p.

Lehrer has argued that the reasons which justify a person in holding a belief need not be causally responsible

---

[36]Where, modifying Pollock's definition, a chain of reasoning is ultimately undefeated for S iff either (1) there is no belief-link in the chain for which S is aware of having a defeater or (2) for each belief-link in the chain for which S is aware of having a defeater D, S is also aware of having a defeater defeater for D. This modification is necessary, because if S believes a defeater D for reason R, but is not aware that D is a defeater for R, then D does not defeat R for S.

[37]It might be objected that, since noticing that one has evidence E for p entails knowing that one has evidence E for p, (PJ_j*) threatens to start us on a vicious regress, because for S to know that she has evidence E, S must be justified in believing that she has evidence E, but for S to be justified in believing that she has evidence E, S must know that she has evidence E' for the belief that she has evidence E, which entails that S is justified in believing that she has evidence E' for the belief that she has evidence E, and so on. Such an objection is misguided, however, because personal justification is not necessary for knowledge. Consequently, there is nothing objectionable about requiring S to know that she has evidence (in the form of an ultimately undefeated chain of reasoning) for p in order for her to be personally justified in believing that p.

for her holding that belief.[38] He has us consider a lawyer whose client is accused of committing eight murders. There is conclusive evidence that his client committed the first seven murders. The lawyer, like everyone else, believes that his client is guilty of the eighth murder, as well, since it has exactly the same modus operandi. However, the lawyer is a gypsy who believes that the cards never lie. After consulting the cards, which repeatedly say that his client is innocent, he comes to believe that his client did not commit the eighth murder. This causes him to dig deeper into the case, whereupon he discovers evidence of his client's innocence, to wit, his client could not have obtained the eighth murder weapon. This newfound evidence justifies the lawyer in believing that his client did not commit the eighth murder, but it does not in any way cause him to hold this belief. Instead, his belief is caused by his trust in the cards, for if he did not believe the cards, the evidence he uncovered would no longer convince him of his client's innocence. Thus, we are told, the reasons which justify the lawyer in holding his belief are not causally responsible for his holding it.

---

[38]Keith Lehrer, "How Reasons Give Us Knowledge, or The Case of the Gypsy Lawyer", The Journal of Philosophy, Vol. LXVIII, No. 10 (May 20, 1971), pp. 311-313.

Unfortunately, Lehrer's example breaks down upon closer inspection, at least if (PJ_j*) is correct. His example is supposed to show that the reasons which justify a person in holding a belief need not cause that belief. All that his example actually shows (if (PJ_j*) is correct) is that a person can have additional reasons for some belief that p which play no causal role in her believing that p. Given (PJ_j*), the lawyer is justified in believing that his client is innocent, before he discovers that his client could not have obtained the eighth weapon. After all, he has based his belief about his client's innocence on his beliefs that the cards never lie and that the cards say his client is innocent. From what we have been told about the lawyer, these latter two beliefs provide him with an ultimately undefeated chain of reasoning for his client's innocence. It is this chain of reasoning which justifies him in his belief about his client's innocence, and, of course, this chain of reasoning is also what causes him to hold the belief. Thus, we have not been given a case where the reasons which justify a person in holding a belief do not cause that belief. In a moment, I will present a case which does illustrate that a person can be justified in holding a belief on the basis of reasons which do not cause him to have that belief, but first I want to consider a potential objection to (PJ_j*).

In objecting to Lehrer's example, I claimed that the

lawyer is justified in believing that his client is innocent, given (PJ_j*), because he has based this belief on his beliefs that the cards never lie and that the cards say that his client is innocent. I claimed this, even though I believe, as I expect you do, that basing beliefs on tarot predictions is a terribly unreliable way to form beliefs. But if we assume that the tarots are unreliable indicators, then how can they justify our gypsy lawyer in his belief about his client's innocence? The answer, which reveals the coherence theoretic implications of (PJ_j*), is straightforward. The lawyer does not believe that the tarots are unreliable. Unlike us, he believes that they are perfectly reliable guides to truth. Given what he believes, namely, that the tarots are reliable and that they say his client is innocent, it is entirely reasonable of him to believe that his client is innocent.[39] Although our gypsy lawyer is fundamentally misguided in some of his beliefs, given these beliefs, it is not profoundly irrational of him to conclude that his client is innocent. It is a virtue of the coherence theory that it allows us to recognize this

---

[39] I am assuming that he does not believe any undefeated counterevidence to either of his reasons. If he were to believe some undefeated counterevidence to either of these reasons, then his chain of reasoning would not be ultimately undefeated, and so, according to (PJ_i*), he would not be personally justified in believing that his client is innocent.

difference.

Let us now return to the task of illustrating that the reasons which justify a person in holding a given belief need not be the cause of his having that belief. Consider the case of the paranoid husband: Max the millionaire is an extreme paranoiac who, throughout his twenty years of marriage, has lived with the unfounded, paranoiac belief that his wife is planning to kill him off for his millions. Fearing that his wife and the butler have been having an affair, he starts to eavesdrop on their conversations. One day, to his dismay, he overhears his wife tell the butler of her plans to get rid of her husband once and for all by injecting a lethal dose of air into his veins while he is sleeping. Hearing this, Max heads to the medicine cabinet, where he finds a brand new syringe, complete with hypodermic needle, which was not there two days ago. His beliefs about the presence of the syringe and about what his wife has said provide him with reasons which justify him in believing that his wife is planning to kill him, but they are not what causes him to hold this belief. His belief that his wife is planning to kill him is entirely the result of his paranoia. After all, had he not acquired the new evidence, he would have still believed out of paranoia that his wife is planning to kill him. Moreover, had he not been paranoid, he would have dismissed his wife's remarks and the presence of the syringe as nothing more than a somewhat morbid

practical joke. Given the above scenario, Max is justified in his belief about his wife's homocidal intentions, but the reasons in virtue of which he is justified are in no way causally responsible for his belief.

It is not uncommon for theories of justification to maintain that a person is justified in believing that p only if her reasons for p are what cause her to believe that p.[40] Such theories are unable to handle the case of the paranoid husband, since they entail that Max is not justified in believing that his wife is planning to kill him, when intuitively he is justified in this belief. Unlike these theories, the linear coherence theory suggested by (PJ_j*) correctly evaluates Max as being justified in his belief about his wife's deadly plans, and so, once again (PJ_j*) yields the right result.

Presently, our account of personal justification is incomplete. (PJ_j*) only tells us when a person is personally justified in holding a belief. It does not tell us when a person is personally unjustified in holding a belief. Therefore, in order to complete our account of personal justificaton, we need to provide an analysis of personal unjustifiedness. We already know from (PJ_u) that S

---

[40] John Pollock's theory is one such theory. See his Contemporary Theories of Knowledge, op. cit., p. 41, where he asserts that the basing relation is at least partly a causal relation.

is personally unjustified in believing that p iff S is being epistemically irresponsible in believing that p. Epistemic irresponsibility is a function of knowingly making illegitimate epistemic moves, e.g. consciously ignoring counterevidence. Consequently, S is personally unjustified in believing that p just in case S knowingly makes epistemically illegitimate moves in coming to (or continuing to) believe that p. It is with this conception of personal unjustifiedness in mind that I offer the following analysis:

(PJ_u*) S is personally unjustified in believing that p iff either
(1) S comes to believe that p by occurrently reasoning to p from an ultimately defeated chain of reasoning C, where a chain of reasoning C is ultimately defeated for S iff C is not ultimately undefeated for S (i.e. to say, iff S is aware of an undefeated defeater for at least one of the belief-links of C); or
(2)(a) S comes to believe that p via a BCP other than occurrent reasoning,
(b) S is occurrently aware of an ultimately undefeated chain of reasoning C which supports ∿p, and
(c) S believes each belief-link of C.

It might be objected that (PJ_u*)'s analysans is so strong that no person will ever satisfy it and that, therefore, given the (PJ_u*) analysis, no person will ever turn out to be personally unjustified in any of her beliefs. I think such an objection is simply mistaken. People rise

to the level of irrationality required by (PJ_u*) all to often. For example, cancer patients often persist in believing that they do not have cancer, despite their having been told by various medical authorities that they do have cancer. Parents of paralyzed children often continue to believe that their child will walk again, despite being inundated with conclusive evidence that their child's condition is permanent. Violent alcoholics frequently persist in believing that they can control their drinking, while knowing that they have repeatedly beaten their wives and children when intoxicated. Many a cigarette smoker believes that smoking is not harmful to one's health, in the face of warning labels, recommendations by the AMA, and reports by the surgeon general. And the list goes on. Unfortunately, there is no doubt that people can achieve the level of irrationality needed to satisfy the analysans of (PJ_u*), and when they do, they are personally unjustified in their beliefs.

Given (PJ_j*) and (PJ_u*), the only thing needed to round out our account of personal justification is an analysis of personal ajustifiedness. We can formulate such an analysis quite simply as follows:

(PJ_a*) S is personally ajustified in believing that p iff
(1) S believes that p, and
(2) S is neither personally justified nor personally unjustified in believing that p.

(PJ_a*) yields the result that we are personally ajustified in holding most of our perceptual beliefs, since (1) we rarely occurrently notice chains of reasoning which support our perceptual beliefs and, so, usually are not personally justified in believing them, and (2) we rarely consciously ignore counterevidence to our perceptual beliefs and, hence, generally are not personally unjustified in believing them, either. I submit that this is precisely the right result, since typically we deserve neither epistemic praise nor epistemic blame for our perceptual beliefs.

Taken together, (PJ_j*), (PJ_u*), and (PJ_a*) provide us with a sophisticated linear coherence theory.[41] This theory, unlike the other theories of justification currently offered, allows us to make all three of the needed types of personal justification evaluations. For this reason alone, it is theoretically superior to these other theories. In light of its theoretical superiority and its ability to provide the intuitively correct personal justification evaluations, I submit that the linear coherence theory developed herein provides us with the correct account of personal justification.

6. Laying the Regress Argument to Rest

With the linear coherence theory of personal

---

[41]To my knowledge, this is the first linear coherence theory to be offered in print.

justification now before us, we can easily see where the regress argument goes awry. The regress argument is unsound, because one of its starting assumptions, viz. A2, is false. A2 asserts:

> S is mediately justified in believing that p iff S comes to believe that p on the basis of some other belief q which S is justified (either mediately or immediately) in believing.

We considered such an assumption in section 5, subsection b, and rejected it on the grounds that it fails to take into account the way in which we actually reason. When we reason, we reason from our beliefs in general, not just from the beliefs we are justified in believing. As a result, we concluded that personal justification is a function of our beliefs, not just a function of the beliefs we are justified in believing. Put another way, personal justification results via reasoning from our beliefs, not via reasoning exclusively from beliefs we are justified in holding, and therefore, A2 is false. Obviously, the falsity of A2 undermines the regress argument.

There is one other point worth mentioning, while we are reconsidering the regress argument. On the basis of the regress argument, it has been thought that coherence theories (especially linear coherence theories) must embrace circular chains of reasoning, i.e. type (iii) chains. This view has led many philosophers to dismiss coherence theories

as being unsatisfactory outright. Such an objection is misguided, however, since a linear coherence theory need not be committed to type (iii) chains. Instead, as we saw in section 5, a linear coherence theory can embrace type (i) chains, and type (i) chains are genuinely justification-conferring, the regress argument to the contrary notwithstanding. I submit that a properly understood linear coherence theory, complete with type (i) chains, is much more plausible than its foundational rivals, since, unlike these rivals, it does not posit the existence of those mysterious unmoved epistemic movers known as "basic beliefs".

7. Looking Ahead

In Chapters 1 and 2, I claimed that personal justification has a crucial role to play in the theory of knowledge. However, in the present chapter, I have claimed repeatedly that personal justification is not necessary for knowledge. I still stand by both of these seemingly incongruous claims. Personal justification is not necessary for knowledge, because a person can know that p without deserving epistemic praise for believing that p. To think otherwise would be to over-intellectualize our ordinary notion of knowledge. Thus, it is true that personal justification has no positive role to play in the theory of knowledge. This, of course, is consistent with its having

some other kind of role to play in the theory. In the next chapter, we shall see that personal justification does have a crucial, albeit negative, role to play in an account of knowledge.

CHAPTER 4

A PARTIAL ANALYSIS OF KNOWLEDGE

1. Introduction

Chapters 2 and 3 were devoted to providing accounts of doxastic justification and personal justification. The present chapter is devoted to delineating the roles which these two types of justification play in the theory of knowledge. In order to specify their respective roles, I shall present a partial analysis of knowledge, an analysis consisting of necessary conditions for S to know that p.[1] Ultimately, I argue that, of these two types of justification, only doxastic justification is necessary for knowledge. Nevertheless, personal justification is not irrelevant to knowledge. To the contrary, it plays a negative, undermining role in the theory of knowledge.

The chapter proceeds in four stages. In section 2, I reconsider the knowledge conception of epistemic justification discussed in Chapter 1 and argue that the kind of justification needed for knowledge turns out to be doxastic justification. In section 3, I propose an analysis of "S knows that p" which takes into account doxastic

---

[1] I leave it an open question as to whether or not these necessary conditions are jointly sufficient for S to know that p.

justification's necessity. In section 4, I demonstrate that the analysis of knowledge offered in section 3 satisfactorily handles some of the miriad of Gettier examples currently afloat. The chapter concludes with section 5, wherein I explore an alternative analysis of knowledge which is considerably stronger than the analysis proffered in section 3. The alternative analysis, unlike the weaker analysis of section 3, is found to be unsatisfactory as an analysis of our ordinary notion of knowledge.

2. The Knowledge Conception of Epistemic Justification Revisited

In Chapter 1, I identified three distinct conceptions of epistemic justification, viz. the doxastic conception, the personal conception, and the knowledge conception. I subsequently provided accounts for the kinds of epistemic justification associated with the former two conceptions. Since the goal of the present chapter is to provide an acceptable (if not complete) account of knowledge, it is only fitting that we return to the knowledge conception of epistemic justification and examine it more closely. Let us start with a brief review.

According to the knowledge conception, epistemic justification is that which must be added to true belief (in non-Gettier situations) to get knowledge. Of course, we

get knowledge only when a particular kind of epistemic justification is added to true belief. That kind of epistemic justification can be characterized as follows:

> (KC) Epistemic justification is a graded normative notion of positive epistemic appraisal that bears an essential internal connection with truth, a certain degree of which is necessary for knowledge.

Quite clearly, truth-connectedness is the essential feature of epistemic justification so-conceived. Thus, in a nutshell, (KC) tells us that the kind of epistemic justification needed for knowledge must be in some way conceptually connected with truth, i.e. it must possess a truth connection.

In light of (KC), we can demonstrate that personal justification is not the kind of justification necessary for knowledge in the following way:[2] Given (KC), personal justification can be the kind of justification necessary for knowledge only if it has a truth connection. Personal justification has a truth connection iff for every possible world W, if conditions C make S personally justified in

---

[2]The argument to be advanced is predicated on the assumption that only one kind of justification is necessary for knowledge. Were this assumption incorrect, the argument would only show that personal justification is not the kind of justification characterized by (KC).

believing that p in W, then conditions C make it probable that p is true in W. For personal justifiedness, the relevant conditions C are specified by the analysans of (PJ_j*). Simplifying (PJ_j*) somewhat, S is personally justified in believing that p just in case the following condition is satisfied:

(C1) S has occurrently noticed an ultimately undefeated chain of reasoning which supports p.

Obviously, there are possible worlds W such that satisfying (C1) in W does not make it probable that p is true in W. The demon world is a case in point. There, the demon sees to it that for every contingent proposition p, if S occurrently notices an ultimately undefeated chain of reasoning which supports p, then p is false. In addition, the demon sees to it that for every necessary proposition p, if S believes that p on the basis of some occurrently noticed chain of reasoning R, then S also occurrently notices an undefeated defeater for R. In such a world, whenever (C1) is satisfied, p is false. Consequently, the conditions (C1), which make S personally justified in believing that p in the demon world, make it probable that p is false in that world. Since there is a possible world W where the conditions which make S personally justified in believing that p in W do not make it probable that p is true in W, personal justification lacks a truth connection.

Therefore, it follows that personal justification is not the kind of justification necessary for knowledge.[3]

We are now in a position to see that doxastic justification as analyzed by doxastic reliabilism is the kind of justification required for knowledge. We know from (KC) that the kind of justification needed for knowledge must have a truth connection. That doxastic justification possesses the requisite truth connection can be demonstrated as follows:[4] Doxastic justification has a truth connection iff for every possible world W, if conditions C make S's belief that p doxastically justified in W, then conditions C make it probable that p is true in W. S's belief that p is doxastically justified in W iff the following condition is satisfied:

(C2) S's belief that p is produced by a BCP that is actually reliable in W.

A BCP is actually reliable in W just in case the indefinite probability of beliefs produced by it in W being true

---

[3]An analogous argument can be advanced for any type of justification for which the conditions of justifiedness are formulated exclusively in terms of states internal to the cognizer. Hence, no brand of internalist justification can be the kind of justification necessary for knowledge.

[4]A similar argument was advanced in Chapter 2 to illustrate that process reliabilism successfully affixes a probabilistic connection between justification and truth.

beliefs in W is high. Thus, for any possible world W, whenever (C2) is satisfied in W, it is highly probable that S's belief that p belongs to the class of true beliefs in W, and so, it is highly probable that S's belief that p is true in W (because members of the class of true beliefs in W are true in W). Consequently, for every possible world W, the condition which makes S's belief that p doxastically justified in W, viz. (C2), makes it probable that p is true in W. Therefore, it follows that doxastic justification has a truth connection.[5]

We have just seen that doxastic justification is conceptually connected with truth. Since doxastic justification is conceptually truth-connected, it possesses the essential feature of the kind of justification necessary for knowledge. This observation, when combined with the general truth that if f is the essential feature of y and if x possesses f, then x is y, entails that doxastic justification is the kind of epistemic justification necessary for knowledge.[6]

---

[5]That doxastic justification has a probabilistic truth connection can also be demonstrated as follows: Doxastic reliabilism identifies doxastically justified beliefs in W with reliably produced beliefs in W. Since, by definition, reliably produced beliefs in W have a high probability of belonging to the class of true beliefs in W, it follows that doxastically justified beliefs in W have a high probability of belonging to the class of true beliefs in W and, hence, are probably true in W.

[6]In Chapter 1, I claimed that the kind of epistemic

We have just found doxastic justification to be the kind of epistemic justification necessary for knowledge. Given this finding, it follows that any acceptable account of knowledge will have a doxastic justification component. The burden of the next section is to provide such an account of knowledge.

3. Analyzing "S knows that p."

My aim in the present section is to provide an account of knowledge which specifies several conditions necessary for S to know that p. It may turn out that these conditions are jointly sufficient for S to know that p, as well. However, for the purposes of the present discussion, I am only committed to their necessity.[7] So much by way of disclamatory preamble.

---

justification needed for knowledge is analyzable and that in order to provide the required truth connection its analysis must be given by an externalist theory. We are now in a position to see that both claims are correct. Since doxastic reliabilism provides the correct analysis of doxastic justification and since doxastic justification is the kind of epistemic justification needed for knowledge, it follows that doxastic reliabilism is the correct analysis of the kind of epistemic justification needed for knowledge. We know from Chapter 3 that doxastic reliabilism is an externalist theory. Moreover, we saw in Chapter 3 that no internalist theory is capable of providing any kind of conceptual connection between justification and truth.

[7] I do contend that the analysis to be offered is

Most epistemologists agree that true belief is necessary for knowledge.[8] In fact, the traditional, knowledge conception of epistemic justification, according to which epistemic justification is that which must be added to true belief (in non-Gettier situations) to get knowledge, is predicated on the assumption that true belief is necessary for knowledge. In light of its nearly universal acceptance, the assumption that true belief is necessary for knowledge shall be taken as a starting assumption without argument. We can capture this starting assumption with the following partial analysis of knowledge:

---

sufficient in non-Gettier situations for S to know that p. In fact, in the next section, I go even further and demonstrate that this analysis does handle some Gettier cases. Even so, in the face of increasingly complex Gettier examples, I am unwilling to commit to its sufficiency per se.

[8]See, e.g. Roderick Chisholm, Theory of Knowledge, second ed. (Englewood Cliffs, NJ: Prentice-Hall, Inc., 1977); Keith Lehrer, Knowledge (Oxford: Oxford University Press, 1974); Marshall Swain, "Knowledge, Causality, and Justification", Essays on Knowledge and Justification, ed. Pappas and Swain (Ithaca, NY: Cornell University Press, 1978); Alvin Goldman, "A Causal Theory of Knowing", Essays on Knowledge and Justification, op. cit., and Epistemology and Cognition (Cambridge, Mass, and London: Harvard University Press, 1986); Robert Nozick, Philosophical Explanations (Cambridge, Mass: The Belknap Press of Harvard University Press, 1981); and John Pollock, Contemporary Theories of Knowledge (in manuscript).

(AK) S knows that p only if
(1) p is true, and
(2) S believes that p.

Thus, our task is to specify additional conditions which, when conjoined with (AK)'s analysans, yield an illuminating account of knowledge capable of handling all non-Gettier situations.

According to what is sometimes called the "traditional analysis" of knowledge,[9] in addition to truth and belief, some kind of epistemic justification is necessary for knowledge.[10] In the previous section, we concluded that the kind of epistemic justification needed for knowledge is doxastic justification, from which it follows that the correct account of knowledge must have a doxastic justification requirement. In order to provide this requirement, we need to modify (AK) in the following way:

(AK') S knows that p only if
(1) p is true,
(2) S believes that p, and
(3) S's belief that p is doxastically justified.

---

[9] George Pappas and Marshall Swain, "Introduction", Essays on Knowledge and Justification, op. cit., p. 1.

[10] In Chapter 1, we saw that some kind of epistemic justification is needed to rule out lucky guesses as instances of knowledge and that in order to do so, that kind of epistemic justification must have a truth connection.

While conditions (1) - (3), as specified by (AK')'s analysans, are individually necessary for S to know that p,[11] they are not alone sufficient for S to know that p, not even in non-Gettier situations. (AK')'s analysans is insufficient for knowledge, because there can be situations where it is subjectively irrational for a person to believe that p, even though, unbeknownst to him, conditions (1) - (3) are satisfied, and in such situations, if the person continues to believe that p, his subjective irrationality prevents him from knowing that p. An example will illustrate the point. Consider once again the case of Jones, whom we met in Chapter 2.[12] This case, as you may recall, goes roughly as follows: Jones has been told that certain of his memory beliefs are completely erroneous. His parents, whose testimony is usually quite reliable, have fabricated an entirely false story to the effect that Jones suffered amnesia when he was seven, but later developed pseudo-memories of that period. Having no reason to doubt his parents, Jones believes what they have said. Nevertheless, he persists in believing the ostensible memories from his seven-year-old past. Since his memory

---

[11]Conditions (1) and (2) are taken to be necessary by assumption. Condition (3) is necessary to rule out merely lucky guesses as instances of knowledge.

[12]See Chapter 2, p. 45.

beliefs result from genuine memory and original perceptions, which in the case at hand are actually reliable BCP's, doxastic reliabilism says that his memory beliefs are doxastically justified.[13]

In the case just described, (AK')'s conditions (1) - (3) are all satisfied. Jones' memory beliefs are doxastically justified, since they result from actually reliable memory, and so, condition (3) is satisfied. Since Jones continues to hold these memory beliefs, condition (2) is satisfied. Moreover, his memory beliefs are true, his parents testimony to the contrary notwithstanding. Hence, condition (1) is also satisfied. Nevertheless, Jones does not know that which he believes about his seven-year-old past. After all, Jones has compelling counterevidence for these memory beliefs, in light of which it is subjectively irrational of him to continue to hold them. Ironically, if Jones were to stop ignoring this counterevidence, he would

---

[13]The case is due to Alvin Goldman and appears in his "What Is Justified Belief?", Justification and Knowledge, ed. George Pappas (Dordrecht: D. Reidel Publishing Company, 1979), p. 18. Goldman presents the case as a counterexample to reliabilist theories like doxastic reliabilism. In Chapter 2, I argued that Goldman is mistaken to regard it as a counterexample, since the fault lies with Jones, not with his memory beliefs. There, it was concluded that Jones' memory beliefs are doxastically justified, just as doxastic reliabilism asserts, but that Jones, given his counterevidence, is personally unjustified in believing them.

worsen his epistemic situation by giving up many true beliefs. But this ironic wrinkle hardly exculpates Jones, since it could only do so, if Jones himself were aware of the wrinkle, which he is not. From his own internal standpoint, Jones is continuing to hold beliefs that he has good reason to think are false and, thus, is being epistemically irresponsible in continuing to hold them. There is, I submit, an overriding intuition that epistemically irresponsible behavior cannot yield knowledge, not even in cases where it does accidentally yield doxastically justified true belief. Thus, we must conclude that Jones lacks knowledge about his seven-year-old past, even though his beliefs in this regard are doxastically justified true beliefs. Consequently, (AK')'s analysans is not sufficient for knowledge.

We have just observed that doxastically justified true belief is not sufficient for knowledge, because it can be undermined by epistemically irresponsible behavior on the part of the would-be knower. Thus, in order to provide an adequate account of knowledge, a further condition must be added to (AK')'s analysans. Appealing to the Jones case can help us determine what this further condition should be.

We concluded that Jones lacks knowledge of certain of his memory beliefs, because it is epistemically irresponsible of him to continue to hold those beliefs. However, this diagnosis of Jones' lack of knowledge, despite

its intuitiveness, is somewhat superficial. What ultimately prevents Jones from knowing certain of his memory beliefs is that which makes it epistemically irresponsible of him to continue to hold those beliefs. What makes it epistemically irresponsible of him to continue to hold his memory beliefs is the fact that he is aware of undefeated counterevidence in the form of rebutting defeaters for these beliefs. This suggests that in order for S to know that p, S must <u>not</u> be aware of any defeaters for p or for the chain of reasoning which supports p (assuming there is such a chain). But notice, this requirement is equivalent to the requirement that S <u>not</u> be personally <u>un</u>justified in believing that p. It is, I submit, this latter requirement which should be built into our account of knowledge. So supplementing (AK'), we get:

(AK*) S knows that p <u>only if</u>
(1) p is true,
(2) S believes that p,
(3) S's belief that p is doxastically justified, and
(4) S is not personally unjustified in believing that p.

I contend that (AK*) provides us with an analysis of knowledge which adequately accounts for most (if not all) of our knowledge. In the next section, I shall show that (AK*) is even capable of handling some Gettierized cases. However, before turning to the Gettier problem, I want to point out two virtues of the (AK*) account.

Perhaps the most impressive virtue of the (AK*) theory is the unique way in which it accounts for our perceptual knowledge. (AK*) accounts for our perceptual knowledge in the following way: Usually, our perceptual beliefs result from actually reliable perceptual processes[14] and, hence, are doxastically justified, just as, usually, we are personally ajustified in believing our perceptual beliefs and, thus, are not personally unjustified in believing them. When, as is often the case, we are personally ajustified in believing our doxastically justified perceptual beliefs and they happen to be true, then we have perceptual knowledge. Thus, our theory asserts that we can and typically do have perceptual knowledge without being personally justified in our perceptual beliefs. That this is a virtue of the theory can be seen by contrasting it with those theories that do require personal justification for knowledge. Theories that do require personal justification for knowledge are forced to embrace ad hoc theories of personal justification, which say that we are personally justified in believing most of our perceptual beliefs, so that they can yield the result

---

[14] I am assuming that the actual world is not a demon world. If, however, the actual world is a demon world (or some other kind of manipulated world) such that our perceptual processes are not actually reliable, then (AK*) rightly entails that we have no perceptual knowledge.

that we do have perceptual knowledge.[15] The problem with such theories is that once they embrace their ad hoc theories of personal justification, they no longer possess the theoretical means to distinguish the person who actively uncovers reasons for her perceptual beliefs from the person who simply holds her perceptual beliefs willy-nilly. After all, since both persons have perceptual knowledge, these theories are forced to assume that both persons are personally justified in their perceptual beliefs. Our theory avoids this unhappy result, because it allows us to maintain that while both persons have perceptual knowledge, only the former is personally justified in her perceptual beliefs.

A second and related virtue of (AK*) is that it does not require any grandiose intellection on the part of the would-be knower. As a result, it allows us to account for the knowledge had by young children and even for the knowledge had by non-humanoid animals. I take it to be an adequacy constraint on a theory of (ordinary) knowledge that it be able to account for children and non-humanoid animal knowledge, especially since a theory which is unable to

---

[15] If they do not embrace a theory of personal justification, according to which we are personally justified in most of our perceptual beliefs, then they yield skepticism with regard to our perceptual beliefs.

account for such knowledge probably cannot account for the man-in-the-street's knowledge, either. That our theory (AK*) satisfies this constraint lends further credence to its being a correct (if not complete) analysis of knowledge.

In light of these virtues, I submit that (AK*) provides us with a very plausible account of knowledge. Let us now turn to see how (AK*) fares in the wake of Gettier examples.

4. The Gettier Problem

a. (AK*) and Two Easy Cases

In his article "Is Justified True Belief Knowledge?",[16] Edmund Gettier single-handedly altered the course of contemporary epistemology by conclusively demonstrating that justified true belief is not sufficient for knowledge. To set the stage for his counterexamples to the justified true belief analysis, Gettier avers that

> in the sense of 'justified' in which S's being justified in believing P is a necessary condition of S's knowing that P, it is possible for a person to be justified in believing a proposition that is in fact false.[17]

---

[16] Edmund Gettier, "Is Justified True Belief Knowledge?", Analysis, Vol. xxiii (1963), pp. 121-123.

[17] Ibid., p. 121. Interestingly enough, here Gettier is employing personal justification without realizing it. As we shall soon see, distinguishing personal and doxastic justification provides us with the means needed to solve at least some of the Gettier problems that have been raised.

Having so avowed his fallibilism, Gettier presents us with a case where Smith has strong evidence for the proposition:

(p) Jones owns a Ford.

As far back as Smith can remember, Jones has always owned a Ford. Moreover, Jones has just offered Smith a ride while driving a Ford.[18] Smith realizes that (p) entails the following proposition:

(q) Either Jones owns a Ford or Brown is in Barcelona.

Although Smith has no evidence as to his friend Brown's whereabouts, having seen the connection between (p) and (q), he believes (q) on the basis of (p) and this recognized entailment, and so, he is justified in believing that (q). Now here is the rub. Unbeknownst to Smith, Jones has sold his Ford and is now driving a rental car. But, as luck would have it, Brown happens to be in Barcelona. Thus, (q) is true, Smith believes that (q), and Smith is justified in believing that (q). Nevertheless, Smith does not know that (q), since from his vantage point, it is simply a matter of

---

[18] If need be, the case can be strengthened by assuming that Smith was with Jones a few months earlier when he purchased a Ford exactly like the one he was driving when he offered Smith a ride.

luck that (q) is true.[19]

This original Gettier case is easily accommodated by our theory (AK*). Since (q) is true, Smith believes that (q), and Smith is personally justified in believing that (q),[20] conditions (1), (2), and (4) of (AK*)'s analysans are all satisfied. Nevertheless, (AK*)'s condition (3) is not satisfied. The process that gives rise to Smith's belief that (q) is reasoning from [(p) and (p) entails (q)] to (q). But notice, the process of reasoning from [(p) and (p) entails (q)] to (q) is a conditionally reliable process. Moreover, we know from Chapter 2, that a conditionally reliable process is actually reliable just in case its input beliefs are true. Since (p) is false, the process of reasoning from [(p) and (p) entail (q)] to (q) is not actually reliable. Furthermore, since a belief is

---

[19]To stress the lucky nature of Smith's justified true belief that (q), Gettier tells us that Smith also believes the following propositions:

(n) Either Jones owns a Ford or Brown is in Boston.
(s) Either Jones owns a Ford or Brown is in Brest-Litovak.

Smith has exactly the same evidence for (q), (n), and (s), namely, (p) which he is justified in believing and (p) entails the propositions (q), (n), and (s), which he is also justified in believing. Surely, that (q) turns out to be true is a matter of luck from Smith's standpoint.

[20](PJ_i^*) entails that Smith is personally justified in believing that (q).

doxastically justified iff it results from an actually reliable BCP, Smith's belief that (q) is not doxastically justified, for it results from a conditionally reliable BCP with a false input belief. Since Smith's belief that (q) is not doxastically justified, condition (3) of (AK*)'s analysans is not satisfied. Therefore, (AK*) yields the right result, namely, that Smith does not know that (q).[21] Let us now consider another counterexample to the justified true belief analysis of knowledge, which, for reasons that will soon become apparent, it seems inappropriate to call a "Gettier example".

It cannot be denied that Gettier deserves much credit for convincing the epistemological community that justified true belief is insufficient for knowledge. Even so, nearly a century earlier, Alexius Meinong (1853-1920) presented his own decisive counterexample to the justified true belief analysis of knowledge, a counterexample which went largely unnoticed by the philosophical community. Meinong's counterexample goes roughly as follows: A man walks into a lecture hall and, wanting to know the time, checks the clock on the wall, which says that it is twelve o'clock.

---

[21]Gettier's second example can be dealt with in exactly the same fashion and, thus, poses no problem for (AK*).

Naturally, the man believes that it is twelve o'clock, and he is justified in holding this belief on the basis of what the clock says. After all, clocks justify us in temporal beliefs all the time. Moreover, it really is twelve o'clock, and so, his belief is true. Thus, the man is justified in believing his true belief that it is twelve o'clock. But, the story continues, the man does not know that it is twelve o'clock, because, unbeknownst to him, the clock stopped at twelve midnight the night before. Had he looked at the clock, say, an hour earlier, he still would have believed it to be twelve o'clock. Since it is just a lucky coincidence that he looked at the clock when he did, he does not know that it is twelve o'clock.

Our theory handles Meinong's example, as well. In his example, it is true that it is twelve o'clock, the man believes that it is twelve o'clock, and he is personally justified in this belief (since the clock gives him good reason to think that it is twelve o'clock). Thus, conditions (1), (2), and (4) of (AK*)'s <u>analysans</u> are satisfied. However, once again, condition (3) is unsatisfied. The process that gives rise to the man's temporal belief, to wit, inferring the time from a stopped clock, is an actually unreliable BCP, for it tends to produce false beliefs. Since the man's belief results from an actually unreliable BCP, it is not doxastically justified. Since his belief is not doxastically justified,

(AK*)'s condition (3) does not obtain. Consequently, (AK*) entails that the man does not know that it is twelve o'clock, and thus, once again, (AK*) yields the correct result.

In this subsection, we have seen that (AK*) possesses the theoretical means needed to handle not only Gettier's original examples, but also the original "Gettier example" raised by Meinong, and this suggests that (AK*) may be sufficient for knowledge, after all. However, the examples considered so far are, admittedly, some of the easier Gettier examples to deal with, and so, it would be premature to conclude that (AK*) is sufficient for knowledge. In the next subsection, I shall examine two much more difficult Gettier cases, which, if correct, do seem to undermine (AK*)'s sufficiency.

b. The Harman Cases and Reflective Equilibrium

In his book Thought,[22] Gilbert Harman presents three Gettierized cases which threaten to undermine (AK*) as a sufficient condition for knowledge. Each of these cases presents us with a person who is personally justified in believing some doxastically justified true belief that p,

---

[22]Gilbert Harman, Thought (Princeton, NJ: Princeton University Press, 1973).

but who supposedly does not know that p because there exists readily available misleading evidence, which the person does not yet possess. It should be noted that if Harman is right in his assumption that a person so situated lacks knowledge, then his cases do provide counterexamples to (AK*)'s sufficiency. I shall present two of these purported counterexamples, but before doing so, a few methodological remarks are in order.

At the outset of this chapter, I disavowed any commitment to (AK*)'s sufficiency for knowledge. That disavowal still stands and should be taken as the context for these methodological remarks. I have already admitted that if Harman's examples are intuitively correct, then they do constitute legitimate counterexamples to (AK*)'s sufficiency. With that said, I must also admit that I, for one, do not share Harman's intuitions regarding these cases.[23] More importantly, I have discussed Harman's cases with numerous persons, epistemologists and non-epistemologists alike, and their intuitions are split

---

[23]Two comments are in order. First, I am not just claiming to have contrary intuitions in order to save (AK*)'s sufficiency, since I genuinely am not committed to its sufficiency at this time. I am claiming contrary intuitions, because my intuitions really do conflict with Harman's. Second, my intuitions regarding Harman's cases conflicted with his, long before I ever conceived of (AK*).

along three lines. Some agree with Harman that the persons in his examples do not possess knowledge. Others think Harman is mistaken and maintain that the persons in question do have knowledge. Still others admit that they have no clear intuitions one way or the other. In such a situation, one seems ill-advised to abandon or modify one's theory in immediate response to the counterexamples. Instead, one should proceed with caution, keeping the following legitimate, methodological principle in mind:

> (MP) One need not abandon or modify a philosophical theory in response to a purported counterexample unless
> (1) the counterexample rests on clear intuitions shared by the bulk of the philosophical community,
> (2) the counterexample is taken to a be a genuine counterexample by the bulk of said community, and
> (3) one cannot show said community to be mistaken regarding (1) and/or (2).[24]

I realize, however, that if a person strongly shares Harman's intuitions, then she will, more than likely, think (AK*) in need of modification. Her thinking this is not directly at odds with (MP). (MP) simply asserts that it is not incumbent upon one to make such modifications, unless

---

[24] I want to thank Ann Levey for suggesting such a principle to me. The current version of (MP), however, is entirely my own, as are any mistakes that may have been made in its formulation.

conditions (1)-(3) are satisfied. When conditions (1)-(3) remain unsatisfied, as they do for the Harman examples, one is free to exercise discretion concerning whether or not to modify the theory. We may, for instance, find it more reasonable to modify (or abandon) the intuitions than to modify the theory. Fortunately, there is a rational procedure for deciding whether we should modify the theory or modify the intuitions, instead. That procedure is to seek a Rawlsean/Goodmanesque state of "reflective equilibrium"[25] between theory and intuition.

The method of reflective equilibrium was first propounded by Nelson Goodman in his "The New Riddle of Induction".[26] There, he is concerned with explaining how we determine the correctness of general rules of inference and of the particular inferences they yield. He tells us that the correctness of rules and of particular inferences is established by bringing the rules and particular inferences into agreement with each other in the following way:

---

[25]John Rawls coined the term 'reflective equilibrium' in his A Theory of Justice (Cambridge, Mass: The Belknap Press of Harvard University Press, 1971). For a shorter presentation, see "A Theory of Justice", excerpted and reprinted in Justice and Economic Distribution, ed. Arthur and Shaw (Englewood Cliffs, NJ: Prentice-Hall, Inc., 1978), pp. 23-26.

[26]Nelson Goodman, "The New Riddle of Induction", Fact, Fiction, and Forecast, fourth ed. (Cambridge, Mass. and London: Harvard University Press, 1983), pp. 59-83.

> A rule is amended if it yields an inference we are unwilling to accept; an inference is rejected if it violates a rule we are unwilling to amend.[27]

By so proceeding, we eventually reach (or at least approach) a state of reflective equilibrium where the rules we accept only yield inferences we accept and the inferences we make accord with rules we accept. At this point, we are justified in taking both the rules and the inferences they yield to be correct.

I submit that we can (and should) use a similar process of mutual adjustment to help us decide when to modify our theory of knowledge to fit our intuitions and when to modify our intuitions to fit our theory of knowledge. In order to describe this process, I shall define a clear case as a case where we share very clear intuitions concerning whether or not the person in question knows, an unclear case as a case where we have unclear or divided intuitions as to whether or not the person in question knows, and a clear theory as a theory capable of accommodating all the clear cases. In order to bring our theory of knowledge and our intuitions into agreement, we should proceed as follows: Our theory of knowledge is to be modified (or abandoned) if it cannot accommodate the clear cases. Once we arrive at a clear

---

[27] Ibid., p. 64.

theory of knowledge, i.e. a theory which does account for the clear cases, our intuitions in unclear cases are to be adjusted so as to accord with the dictates of our clear theory. Simply put, clear cases take precedence over any theory of knowledge, and a clear theory of knowledge takes precedence over unclear cases. By proceeding in this way, we should gradually reach (or approach) a point where our theory of knowledge accommodates our intuitions and where our intuitions accord with our theory of knowledge, a point of epistemic equilibrium as it were. This concludes my methodological remarks. I shall now present two of Harman's purported counterexamples to (AK*)'s sufficiency.

In one example your friend Donald has gone off to Italy for the summer. He told you that he was going to Italy, and you saw him off at the airport. As a practical joke, Donald decides to fool you into believing that he is in California, rather than Italy. So, he writes several letters saying that he has gone to San Francisco for the summer, and he sends them to someone he knows there, who, in turn, sends them to you with San Francisco postmarks, one at a time. Having been out of town a week yourself, you have not as yet read any of the letters. You are now about to open your mail from the past week, which includes two of the phony letters. Just then a mutual friend calls and asks you if you know where Donald is. You reply, "Yes, I know that he is in Italy." As Harman sees it,

> You are right about where Donald is and it would seem that your justification for believing that Donald is in Italy makes no reference to letters from San Francisco. But you do not know that Donald is in Italy. Your knowledge is undermined by evidence you do not as yet possess.[28]

In another example, we are to imagine the following scenario:

> A political leader is assassinated. His associates, fearing a coup, decide to pretend that the bullet hit someone else. On Nationwide television they announce that an assassination attempt has failed to kill the leader but has killed a secret service man by mistake. However, before the announcement is made, an enterprising reporter on the scene telephones the real story to his newspaper, which has included the story in its final edition. Jill buys a copy of that paper and reads the story of the assassination. What she reads is true and so are her assumptions about how the story came to be in the paper.[29]

Harman concludes that even though Jill has justified true belief, she does not know that the political leader has been assassinated, because of the undermining television evidence she does not possess.[30] Let us now examine each of these examples in turn, keeping in mind our goal of epistemic

---

[28]Gilbert Harman, Thought, op. cit., p. 143.

[29]Ibid., pp. 143-144.

[30]Ibid., p. 144.

equilibrium.

In the first example, it is true that Donald is in Italy, and you believe that Donald is in Italy. Moreover, since your belief that Donald is in Italy is based on Donald's testimony and on your belief that you saw him board a plane bound Italy, which confirms his testimony, and since you have no reason to think that Donald was lying, you are personally justified in believing that he is in Italy. Hence, we have another case where (AK*)'s conditions (1), (2), and (4) are satisfied. However, in this case it also looks as if condition (3) is satisfied, because your belief that Donald is in Italy is the result of inference based on testimony, which presumably is an actually reliable process.[31] Since (AK*)'s analysans is satisfied, it follows that either you do know that Donald is in Italy or else (AK*)'s analysans is not sufficient for knowledge. Since, at least among the people I have consulted, intuitions are genuinely divided concerning whether or not you know that Donald is in Italy, the Donald case is an unclear case. Therefore, since (AK*) has been able to handle the clear cases so far examined, our method of seeking epistemic equilibrium dictates that we bring our intuitions in line

---

[31]Thomas Reid observed, and I think rightly so, that people have a natural propensity towards honesty, and as a result, their testimony tends to be true. See his Inquiry and Essays, ed. Lehrer and Beanblossom (Indianapolis, IN: The Bobbs-Merrill Company, Inc., 1975), pp. 93-95.

with (AK*) and assume that you do know Donald's whereabouts. While I do believe that this is the most rational way to deal with unclear cases, it may strike some as question-begging.[32] It would, at any rate, be more satisfying intellectually, if we could explain away Harman's contrary intuitions to the effect that you do not know.

Harman does not actually explain why he thinks that you lack knowledge of Donald's whereabouts, except that there is undermining evidence which you do not possess. Nevertheless, he presumably has something like the following in mind. You hang up the phone and start opening your mail from the past week. A few minutes pass, and sure enough, you stumble upon and read one of Donald's phony letters. Then, you find the second letter, as well. On the basis of these letters, you come to believe that Donald is in San Francisco. Thus, one thing is clear: You do not now know that Donald is in Italy. Moreover, you now think you know that you did not know that Donald is in Italy, after all. But notice, your thinking this is consistent with your having actually known that Donald is in Italy. After all,

---

[32]Actually, such question-beggingness seems to be built into the method of seeking epistemic equilibrium (or any other kind of reflective equilibrium for that matter). In short, just as we beg the question in favor of clear cases, so too we beg the question in favor of clear theories over unclear cases.

knowing that you didn't know is tricky business, just as tricky, I dare say, as knowing that you know.[33] Since you could easily be mistaken in thinking that you now know that you didn't know that Donald is in Italy, your thinking this hardly constitutes good evidence that you did not know. So, Harman must have had something else in mind. Perhaps, he thinks that since you do not now know that Donald is in Italy, you could not have known that Donald is in Italy ten minutes earlier before you read the letters. But this would only be true, if the following principle were true:

(P ~K) If S does not know that p at t, then S did not know that p at t-1.

But this principle is false. S might know that p at t-1 and then forget that p at t and, so, not know that p at t. Once again, we are left without a good reason for thinking that you did not know that Donald is in Italy.

On the other hand, (AK*), offers a natural suggestion for why you do know that Donald is in Italy, before you read

---

[33] In order for you to know that you didn't know, you would have to believe that you didn't know, this belief would have to be doxastically justified, you would have to not be personally unjustified in believing it, and most important of all, it would have to be true that you didn't know, which is the very thing in question. Thus, assuming that you know you didn't know equally begs the question.

the letters. Your belief that Donald is in Italy is reliably produced and ipso facto doxastic justified, thereby giving it a high probability of being true. Moreover, you are personally justified (and hence, not personally unjustified) in believing that Donald is in Italy. Given your situation, it would have been epistemically irrational of you not to believe that Donald is in Italy. That you are being epistemically rational in holding the belief, a belief which has been reliably produced, makes it at least plausible to think that you do know that Donald is in Italy. In short, in the absence of any good reason for thinking you don't know Donald's whereabouts, it is at least as reasonable (if not more so) to think that you do possess such knowledge. This being the case, sticking with our theory (AK*), as epistemic equilibrium requires, is certainly a reasonable thing to do. Now, let us turn to the Jill example.

Regarding this example, Harman does offer an explanation for why Jill lacks knowledge, despite having justified true belief. His explanation is as follows:

> [Jill] does not know that the political leader has been assassinated. For everyone else has heard about the televised announcement. They may also have seen the story in the paper and, perhaps, do not know what to believe; and it is highly implausible that Jill should know simply because she lacks evidence everyone else has. Jill does not know. Her knowledge is undermined by evidence she does not

possess[34] (emphasis added).

Is it really so implausible that Jill should know just because she lacks evidence everyone else has? I do not think so. After all, it is not at all implausible that a person should know just because he is in a different evidential situation than everyone else. In fact, this happens all the time. Admittedly, the normal case is when the person has more evidence than everyone else and, thus, knows, where everyone else lacks knowledge. But not always. Sometimes a person knows, where everyone else lacks knowledge, not because he has more evidence, but because he has different evidence. This suggests, and I think rightly so, that whether or not a person knows is more a function of the quality of his evidence than the quantity of his evidence. Jill has less evidence than everyone else, but she has good evidence. After all, reputable newspapers are reliable sources of information. Is it implausible to think that Jill knows just because she has good evidence in the form of a reputable newspaper's column? Of course not. And this description of Jill's situation is just as fair and accurate a description as Harman's description that she knows "just because she lacks evidence that everyone else has". It is, I submit, a mistake on Harman's part to focus

---

[34]Gilbert Harman, Thought, op. cit., p. 144.

on the quantity, rather than the quality, of Jill's evidence.

Jill bases her belief about the political leader's assassination on good evidence, for which she has no defeaters, and so, she is personally justified in believing that the political leader has been assassinated. Moreover, since inferring beliefs from reputable newspapers is a reliable way to form beliefs, her belief is doxastically justified, as well. And since her belief is also true, it seems quite reasonable to maintain that she knows whereof she believes. True, if she learns of the televised announcement, then she will not know; but, and here lies the crux of the biscuit, this is not the evidential situation that she currently is in. In her current situation, where she has no defeaters for her evidence, it is plausible to think that she knows. Thus, unless we are given a better explanation of Jill's purported lack of knowledge, an explanation which converts this unclear case into a clear case, we should seek epistemic equilibrium by agreeing with (AK*)'s assessment that Jill does, in fact, know that the political leader has been assassinated.

It might be objected that my discussion of Jill's situation has overlooked the most important feature of the Jill example, namely, that Jill is extremely lucky to be in the evidential situation she is in, and since knowledge is not a matter of luck, Jill does not know that the political

leader has been assassinated. Such an objection is misguided, however, because there is more than one way for a person to be lucky and not all of these are incompatible with knowledge. There is, I submit, an epistemologically relevant difference between the person who is lucky because, given her evidential situation, it is a matter of luck that her belief turns out to be true, and the person who is lucky to be in the evidential situation she is in, but that, once in this situation, it is not a matter of luck that her belief is true. Let us say the former person is lucky_1 and the latter person is lucky_2. Examples of persons who are lucky_1 include the person who holds a true belief by merely guessing and the person in typical Gettier cases whose justification (i.e. what justifies her) in holding a belief has nothing to do with that belief's being true.[35] Jill, on the other hand, is lucky_2, since she is lucky to be in the evidential situation she is in vis-a-vis the political leader's assassination, but, given that she is in this evidential situation and has a reliably produced belief, it is not a matter of luck that her belief is true. Clearly,

---

[35]For example, in Gettier's example, Smith's evidence that Jones owns a Ford, which justifies him in believing that either Jones owns a Ford or Brown is in Barcelona, has nothing to do with the latter belief's being true, since it is true because Brown is in Barcelona and not because Jones owns a Ford.

the person who is $lucky_1$ lacks knowledge. Should we likewise conclude that people like Jill who are $lucky_2$ lack knowledge? I do not think we should. To see why, consider the following case: Joe now knows that it is storming outside, because he sees the rain and wind whipping around outside his study's window. Surely, this is a typical case of perceptual knowledge. Even though Joe knows that it is storming, he could have just as easily not known this, since he could have been working in his windowless carrel in the bowels of the library, instead of in his study. Moreover, since Joe works in his carrel at least as often as he works in his study, it is largely a matter of luck that he is in his present evidential situation vis-a-vis the storm, rather than the evidential situation he would have been in, had he been in his carrel. While it is true that Joe is lucky to be in his present evidential situation (for he could have easily been in a much less fortuitous evidential situation vis-a-vis the storm), surely, we do not want to deny that he knows that it is storming, when he sees the wind and rain whipping around outside his window, for what could be a more standard case of perceptual knowledge?

Of course, once we admit that Joe knows that it is storming, despite the luck involved in his being in a evidential situation appropriate for such knowledge, i.e. despite his being $lucky_2$, we should, by parity of reason, also admit that Jill knows that the political leader has

been assassinated, despite her being lucky$_2$. Since Jill's being lucky$_2$ does not prevent her from knowing that the political leader has been assassinated, we are once again left with no good reason for thinking that she lacks such knowledge. Without a clearly intuitive reason for thinking that Jill does not know, our method of seeking epistemic equilibrium dictates that we should side with (AK*) and maintain that she does, in fact, have knowledge of the political leader's assassination.

We began, in subsection a, by considering two fairly typical Gettier-type cases where the persons in question were lucky$_1$ and, hence, lacked knowledge, and we saw that (AK*) concurred and, thus, correctly handled both of these clear cases. In the present subsection, we have just examined two extremely controversial cases, due to Harman, both of which threatened to undermine (AK*)'s sufficiency. However, since both of these examples were found to rest on widely disputed intuitions, both failed to offer convincing proof of (AK*)'s insufficiency. Thus, I submit that, in light of our method of epistemic equilibrium, since (AK*) does accurately account for the clear cases so far examined, we are justified in regarding (AK*)'s <u>analysans</u> as both necessary and sufficient, at least until a clear counterexample to its sufficiency is presented.[36] Let us,

---

[36]I remain open to the possibility that such a

therefore, agree that (AK*) provides us with a prima facie correct theory of knowledge.

5. Ordinary versus Intellectual Knowledge

In section 3 of the present chapter, I claimed that a theory of ordinary knowledge is adequate only if it is able to account for the knowledge had by young children and non-humanoid animals. I then suggested that it is a virtue of (AK*) that it satisfies this adequacy constraint. Some philosophers would disagree. For example, Lehrer and Cohen would deny that the constraint I proposed is really an adequacy constraint on a theory of knowledge, since they do not think that young children and non-humanoid animals have knowledge, and would, therefore, deny that (AK*)'s satisfying such a constraint constitutes a virtue. Regarding the question of young children and non-humanoid animal knowledge, they tell us,

> Here we must, of course, be wary of the sympathetic fallacy. The charm of very small children and animals naturally disposes us to attribute cognitive accomplishments to them of which they are entirely incapable. We prefer to say that such beings have information but lack knowledge.[37]

---

counterexample does exist, and for this reason, I am not fully committed to (AK*)'s sufficiency.

[37]Keith Lehrer and Stewart Cohen, "Justification, Truth, and Coherence", Synthese, Vol. 55 (1983), p. 200.

They then suggest,

> To avoid a verbal impasse, however, one might choose to speak of such beings as having a primitive form of knowledge which lacks the usual justification that is a constituent of a more advanced form of knowledge.[38]

Their theory is offered as an analysis of advanced knowledge, rather than its more "primitive" counterpart. That they have identified different conceptions of knowledge and have chosen to analyze their favored conception is unobjectionable. If, however, they mean to imply that advanced knowledge is the kind of knowledge that adult human beings ordinarily have,[39] then their suggestion is far less benign, for whether or not non-humanoid animals have the same kind of knowledge as that ordinarily had by adult human beings is precisely what is in question. Just as Lehrer and Cohen have cautioned us not to commit the sympathetic fallacy in answering this question, so too should we caution them, when answering this question, not to commit the superiority fallacy, the fallacy human beings are wont to make of taking themselves to have cognitive capabilities

---

[38] Ibid.

[39] It looks as if they do mean to imply this, given the pejorative nature of "primitive knowledge".

radically superior (even to the extent of being radically different in kind) to those had by other animals, despite the fact that we recently evolved from such animals and still retain significant cerebral similarities.[40] In attempting to avoid both of these fallacies, it seems that, on the one hand, we must admit that human beings are clearly capable of some cognitive achievements that other animals are not capable of, while, on the other hand, we must also admit that perceptual knowledge is not one of these exclusively human cognitive achievements. Thus, we can agree that complex mathematics is limited to humans (and only some humans at that), while also agreeing that the perceptual belief that there is a threatening dog chasing me can be known by humans and non-humanoid animals alike. Moreover, I think we should agree that when such a perceptual belief is known, it tends to be known in the same way, regardless whether it is known by a human or some other animal. Hence, a less disparaging, and I think more accurate, way of making the distinction that Lehrer and

---

[40]There is even reason to think that many humans employ the superiority fallacy to justify their mistreatment and exploitation of other animals, and that it is only fairly recently that the fallacy has started being limited to non-humanoid animals. In the antebellum era, for instance, some Southern whites used their belief that blacks were cognitively inferior animals to justify the mistreatment and exploitation of their slaves. Perhaps, this is still going on in apartheid South Africa.

Cohen were driving at is to distinguish between intellectual knowledge, such as complex mathematical knowledge, and ordinary knowledge, such as standard perceptual knowledge.

Once we make the distinction in this way, it becomes clear that the adequacy constraint I suggested in section 3 really is an adequacy constraint on a theory of ordinary knowledge. Moreover, since (AK*) satisfies this constraint and has so far been immune to counterexamples, it is reasonable to think that (AK*) provides the correct analysis of ordinary knowledge. What, then, is the correct analysis of intellectual knowledge? I contend that the correct account of intellectual knowledge can also be formulated in terms of doxastic and personal justification as follows:

(AK_I) S intellectually knows that p only if
(1) p is true,
(2) S believes that p,
(3) S's belief that p is doxastically justified, and
(4) S is personally[41] justified in believing that p.

It might be objected that since (AK_I), unlike (AK*), requires personal justification, it threatens to start us on a vicious regress in a way that (AK*) safely avoids. After all, according to (AK_I), in order for a person to know that p, she must be personally justified in believing that p.

---

[41] Again, I leave it open as to whether (AK_I)'s analysans is sufficient for intellectual knowledge.

But, according to (PJ_j*), in order for a person to be personally justified in believing that p, she must occurrently notice and, hence, know that she has evidence E for p. Of course, according to (AK_I), in order for her to know that she has evidence E for p, she must be personally justified in believing that she has evidence E for p; and in order for her to be personally justified in believing that she has evidence E for p, she must occurrently notice and, hence, know that she has evidence E' for believing that she has evidence E for p; and so on ad infinitum.

Fortunately, such an objection is easy to forestall. The kind of knowledge which personal justification requires is ordinary knowledge. Thus, in order for a person to be personally justified in believing that p, she must occurrently notice and, hence, know in the ordinary sense that she has evidence E for p. Of course, she can know in the ordinary sense that she has evidence E for p, without being personally justified in believing that she has evidence E' for believing that she has evidence E for p, because personal justification is not necessary for ordinary knowledge. As a result, (AK_I) does not start us on a vicious regress. (AK_I) does require a person to be personally justified in believing that p, in order to know in the intellectual sense that p, and so, indirectly requires the person to know in the ordinary sense that she has evidence E for p. But, (PJ_j*) does not require the

person to know (in either sense) anything else, in order to be personally justified in believing that p, and therefore, the regress stops here.

Our response to the regress objection is instructive because it shows that (AK_I) is parasitic on (AK*). Such parasitism should probably come as no surprise, since all throughout history those epistemologists who have embraced overly intellectual accounts of knowledge have, in turn, been forced to acknowledge some other kind of non-inferential immediate knowledge to stay the impending regress of reasons for one's reasons.[42] Their mistake is that they have traditionally tried to account for this non-inferential immediate knowledge in terms of their intellectual account of knowledge. To this end, they have been compelled to embrace basic beliefs, which the knower is somehow immediately justified in believing, but have never been able to give an adequate account of how basic beliefs are possible nor of which beliefs are basic. Thus, their accounts of intellectual knowledge have remained unsatisfactory. However, once we realize that there exists a considerably less intellectual, more fundamental kind of

---

[42]Descartes, for example, is forced to distinguish between that which is known directly and immediately, our sensations and ideas, from that which is only known indirectly and problematically on the basis of these sensations and ideas, e.g. external objects.

knowledge —— the kind of knowledge we ordinarily have —— we can, then, easily construct an account of intellectual knowledge upon it. This is precisely what (AK_I) does and is the reason for its success in avoiding spurious circles and repugnant regresses.

While (AK_I) is interesting in its own right and does provide us with a philosophically satisfying account of intellectual knowledge, it yields the result that much of which we ordinarily know we do not intellectually know. For example, given (AK_I), we lack intellectual knowledge of most of our perceptual beliefs, of the external world in general, and of the existence of other minds, all of which are things we ordinarily know. On the other hand, we are not entirely without intellectual knowledge, either, and some of us possess a great deal of intellectual knowledge. For example, scientists frequently have intellectual knowledge in their respective domains, as do other specialists, such as economists, mathematicians, logicians, physicians, philosophers, engineers, etc. Intellectual knowledge is not limited to the intellectual elite, however, for we even occasionally have intellectual knowledge of beliefs concerning day to day things, when we base those beliefs on ultimately undefeated chains of reasoning. Even so, it must be stressed that most of what we know we do not intellectually know, since we do not usually take occurrent note of our reasons. Thus, while (AK_I) analyzes a kind of

knowledge worthy of striving for, it does not analyze the kind of knowledge we ordinarily have. This latter kind of knowledge is analyzed by (AK*), instead. One more observation and this chapter is finished.

The goal of this chapter was to provide an analysis of knowledge. In the end, we provided two such analyses: (AK*), as an analysis of ordinary knowledge, and (AK_I), as an analysis of intellectual knowledge. In closing, we should note that neither of these analyses could have been formulated without the aid of the personal/doxastic justification distinction. Hence, once again, the personal/doxastic justification distinction proves indispensable for making progress in epistemology.

CHAPTER 5

CONCLUDING REMARKS

1. A Look at the Internalist/Externalist Controversy

The controversy over whether internalism or externalism provides the correct account of epistemic justification has been and continues to be one of the most heavily discussed topics in recent contemporary epistemology. Unfortunately, this discussion, which largely consists of advocates for one of the two positions extolling the virtues of their favored position, while citing the shortcomings of their opponents' position,[1] has done little, if anything to resolve the controversy. As mentioned in Chapter 1, I contend that much of the confusion surrounding the internalist/externalist controversy is directly traceable to the failure to distinguish personal and doxastic justification. It is now time to defend this contention.

In Chapter 1, I noted that the justification requirement for knowledge is usually formulated in one of the following two ways:

(JR1) S knows that p only if S is epistemically justified in believing that p.

---

[1] A notable exception is Alvin Goldman's attempt to resolve the controversy in his "Strong and Weak Justification" (in manuscript).

(JR2) S knows that p only if S's belief that p is epistemically justified.

I then observed that most epistemologists regard (JR1) and (JR2) to be synonymous, stylistic variants of the justification requirement, and in so doing, tacitly embrace the equivalency thesis, to wit,

(ET) S is epistemically justified in believing that p iff S's belief that p is epistemically justified.

But notice, one can embrace (ET) and yet consistently maintain that epistemic justification's primary domain of evaluation is persons. Likewise, one can embrace (ET) and consistently maintain that epistemic justification's proper domain of evaluation is beliefs. I submit that, without realizing it, internalists have done the former, and externalists have done the latter. After all, we know from Chapter 1 that internalists, like Lehrer and Pollock, have gravitated towards the (JR1) formulation of the justification requirement, whereas externalists, like Goldman, have opted for the (JR2) formulation. And we know from Chapter 2 that, when faced with demon world examples, internalists focus on the justificatory status of the demon-world inhabitant himself, while externalists focus on the justificatory status of his beliefs. Thus, it seems clear that internalists have been primarily concerned with personal justification and that externalists have, instead,

been interested in doxastic justification. However, since both internalists and externalists have embraced (ET), which conflates personal and doxastic justification, both have failed to realize that they have been talking about different kinds of epistemic justification. Interestingly enough, once we do realize that internalists have been interested in personal justification and that externalists have been interested in doxastic justification, we can see that both have been right all along, for, as we saw in Chapter 3, the correct account of personal justification can only be given by an internalist theory and the correct account of doxastic justification can only be given by an externalist theory. Thus, the internalist/externalist controversy resolves in both of their favors.

Despite their mutual correctness, however, the externalist ultimately wins out in the following way. Both the internalist and the externalist take themselves to be talking about the kind of epistemic justification necessary for knowledge. In this regard, only the externalist is correct, because, as we know from (AK*), only doxastic justification is necessary for knowledge.

2. Summary of the Dissertation: A Vindication of the Personal/Doxastic Justification Distinction

I began this dissertation by demonstrating that there is no single, unitary notion of epistemic justification, but

rather a family of notions that are currently batted around under the single heading "epistemic justification". I then suggested that we could isolate out two very important kinds of epistemic justification in terms of their respective domains of evaluation, to wit, doxastic justification which takes beliefs as its domain of evaluation and personal justification which has persons as its domain of evaluation. The remainder of the dissertation was devoted to analyzing these two types of epistemic justification and to tracing out their ramifications for epistemology.

In Chapter 2, I formulated an externalist account of doxastic justification, namely, doxastic reliabilism. In the course of defending this account, it was discovered that the personal/doxastic justification distinction provides us with the theoretical means needed to account for the divergent intuitions that regularly arise regarding justificatory evaluations in demon world contexts. This, in turn was seen to secure the distinction.

In Chapter 3, the logical geography of justification theories was mapped out. Then, using this geography as the basis for an argument from elimination, I demonstrated that the correct account of personal justification can only be given by a coherence theory. And then I formulated and defended a linear coherence theory of personal justification.

Having thus provided accounts of both personal and

doxastic justification, I proceeded to offer an analysis of knowledge in terms of these two kinds of justification. It was argued that only doxastic justification is necessary for ordinary knowledge, but that personal justification, nevertheless, has a negative, undermining role to play in such knowledge. I then illustrated that the personal/doxastic justification distinction helps us to account for typical Gettier cases. Finally, I presented an even stronger analysis of knowledge which requires both personal and doxastic justification. This latter analysis, while interesting in its own right, proved too strong for an analysis of ordinary knowledge.

Finally, in the concluding chapter, I argued that the failure, on the part of most epistemologists, to distinguish personal justification from doxastic justification is directly responsible for the current internalist/externalist controversy in contemporary epistemology, and that once this distinction is brought to bear on the controversy, the controversy itself simply dissolves.

Thus, in the course of the present dissertation alone, the personal/doxastic justification distinction has proved indispensable in at least four ways. It has proved indispensable: (1) for explaining away the conflicting intuitions in demon world examples, (2) for providing the correct account of ordinary knowledge, (3) for allowing us

to account for at least some of the Gettier examples currently afloat, and (4) for leading to the dissolution of the internalist/externalist controversy. Given its indispensability in these respects, I submit that if we want to make progress in epistemology, we must keep an eye toward the personal/doxastic justification distinction and tailor our theorizing and criticism accordingly.

LIST OF REFERENCES

Alston, William, "Concepts of Epistemic Justification", The Monist, Vol. 68, No. 1, January, 1985.

———————, "Two Types of Foundationalism", The Journal of Philosophy, Vol. LXXIII, No. 7, April 8, 1976.

Bonjour, Laurence, "Can Empirical Knowledge Have a Foundation?", American Philosophical Quarterly, Vol. 15, No. 1, January, 1978.

———————, "Externalist Theories of Empirical Knowledge", Midwest Studies in Philosophy, ed. French, Vehling, and Wettstein, University of Minnesota, 1980.

Chisholm, Roderick, "The Place of Epistemic Justification", Philosophical Topics, Vol. 14, No. 1, Spring, 1986.

———————, Theory of Knowledge, second ed., Prentice-Hall, Inc., Englewood Cliffs, NJ, 1977.

Churchland, Paul, Matter and Consciousness, The MIT Press, Cambridge, Mass. and London, England, 1984.

Cohen, Stewart, Justification and Truth, Doctoral Dissertation on File with The University of Arizona Library, 1983.

———————, "Justification and Truth", Philosophical Studies, Vol. 46, 1984.

Dretske, Fred, "Conclusive Reasons", Essays on Knowledge and Justification, ed. Pappas and Swain, Cornell University Press, Ithaca, NY, 1978.

Engel, Jr., Mylan, "Coherentism Reliabilized", Acta Analytica, No. 2. 1986.

Feldman, Richard, "Reliability and Justification", The Monist, Vol. 68, No. 2, April, 1985.

———————, "Subjective and Objective Justification in Ethics and Epistemology", in manuscript, 1986.

Gettier, Edmund, "Is Justified True Belief Knowledge?", Analysis, Vol. xxiii, 1963.

Goldman, Alvin, "A Causal Theory of Knowing", Essays on Knowledge and Justification, ed. Pappas and Swain, Cornell University Press, Ithaca, NY, 1978.

———————, Epistemology and Cognition, Harvard University Press, Cambridge, Mass. and London, England, 1986.

———————, "Strong and Weak Justification", in manuscript, 1986.

———————, "What Is Justified Belief?", Justification and Knowledge, ed. G. Pappas, D. Reidel Publishing Company, Dordrecht, 1979.

Goodman, Nelson, "The New Riddle of Induction", Fact, Fiction, and Forecast, fourth ed., Harvard University Press, Cambridge, Mass. and London, England, 1983.

Harman, Gilbert, Thought, Princeton University Press, Princeton, NJ, 1973.

Lehrer, Keith, "How Reasons Give Us Knowledge, or The Case of the Gypsy Lawyer", The Journal of Philosophy, Vol. LXVIII, No. 10, May 20, 1971.

———————, Knowledge, Oxford University Press, Oxford, England, 1974.

———————, "Knowledge, Truth, and Ontology", Proceedings of the 6th International Wittgenstein Symposium, August, 1981.

———————, "Metaknowledge: Undefeated Justification", in manuscript, 1987.

——————— and Stewart Cohen, "Justification, Truth, and Coherence", Synthese, Vol. 55, No. 2, May, 1983.

Nozick, Robert, Philosophical Explanations, The Belknap Press of Harvard University Press, Cambridge, Mass., 1981.

Pappas, George and Marshall Swain, "Introduction", Essays on Knowledge and Justification, ed. Pappas and Swain, Cornell University Press, Ithaca, NY, 1978.

Plantinga, Alvin, "Justification and Theism", in manuscript, presented at the APA Conference, December, 1984.

Pollock, John, Contemporary Theories of Knowledge, in manuscript, 1985; Forthcoming in Rowman and Allenheld.

————————, Knowledge and Justification, Princeton University Press, Princeton, NJ, 1974.

————————, "Reliability and Justified Belief", Canadian Journal of Philosophy, Vol. XIV, No. 1, March, 1984.

Rawls, John, A Theory of Justice, The Belknap Press of Harvard University Press, Cambridge, Mass., 1971.

————————, "A Theory of Justice", Justice and Economic Distribution, ed. Arthur and Shaw, Prentice-Hall, Inc., Englewood Cliffs, NJ, 1978.

Reid, Thomas, Essays on the Intellectual Powers of Man, ed. Baruch A. Brody, The MIT Press, Cambridge, Mass., 1969.

————————, Inquiry and Essays, ed. Lehrer and Beanblossom, The Bobbs-Merrill Company, Inc., Indianapolis, IN, 1975.

Sosa, Ernest, "The Raft and the Pyramid: Coherence versus Foundations in the Theory of Knowledge", Midwest Studies in Philosophy, ed. French, Vehling, and Wettstein, University of Minnesota, 1980.

Swain, Marshall, "Knowledge, Causality, and Justification", Essays on Knowledge and Justification, ed. Pappas and Swain, Cornell University Press, Ithaca, NY, 1978.

Van Cleve, James, "Foundationalism, Epistemic Principles, and the Cartesian Circle", The Philosophical Review, Vol. 88, No. 1, 1979.