THE PROBLEM OF INDUCTION:

AN EPISTEMOLOGICAL AND METHODOLOGICAL RESPONSE

BY

Alan Robert Rhoda

BA, University of Nevada, Las Vegas, 1993

MA, Fordham University, 1996

DISSERTATION

SUBMITTED IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

FOR THE DEGREE OF DOCTOR OF PHILOSOPHY

IN THE DEPARTMENT OF PHILOSOPHY

AT FORDHAM UNIVERSITY

NEW YORK

DECEMBER 2003

ACKNOWLEDGEMENTS

The number of people who have significantly contributed to my graduate studies, which culminate in this dissertation, is numerous. I thank Fordham University and the Department of Philosophy in general for the opportunity to study at a great institution under knowledgeable professors who really care about their students. More specifically, I thank my dissertation advisor, John Greco, for his encouragement and for his ever insightful and challenging feedback on my work. I also thank my dissertation readers, Dominic Balestra and Judith Green, for their involvement and support. Thanks also to Paul Schollmeier and Michael Theemling for helpful comments on drafts of this dissertation; to Sara Penella for introducing me to Peirce; to Danny Green, my high school chemistry teacher, for introducing me to philosophy; and to Maria Terzulli, that department-secretary-than-which-a-greater-cannot-be-conceived. Special thanks to my mom and step-father, Ursula and Ronald Limuti, and to my dad and step-mother, R. Carl and Carmen Rhoda, for their constant love, support, and encouragement. It is to my parents that I dedicate this work. Finally, I am grateful to my wife and best friend, Heather, for loving me and believing in me even when I have doubted myself.

TABLE OF CONTENTS

INTRODUCTION

This dissertation concerns the problem of induction, or rather, two related *problems* of induction. The first is what I call the *negative problem of induction*. This is the problem of rebutting the skeptical argument, first articulated by Hume, which purports to show that inductive inferences cannot in principle result in epistemically justified beliefs. A solution to this problem, i.e., a *negative justification* of induction, merely shows that the skeptical conclusion can be plausibly denied, specifically, that it is *possible* that induction results in epistemically justified beliefs. It does not show that this is *probable*, or that it is *reasonable* to rely on induction as a source of epistemically justified beliefs. This is the second or *positive problem of induction*. In this dissertation I develop solutions to both of these problems. By way of introduction, I will briefly outline the structure of my argument and then say a few words about the philosophical significance of my results.

I. An Outline of the Argument

A successful response to the negative problem of induction must show either (a) that the skeptical argument is invalid, or (b) that it is unsound, i.e., that it relies on some premise that we can plausibly reject. In chapter one I outline the skeptical argument and argue that it is both valid and *apparently*, at least, sound. I do not think it *is* sound, but it is important to pose a problem as forcefully as possible before attempting to solve it, lest one succeed only in knocking down a straw man. Indeed, if the skeptical argument were not at least apparently sound to a great many people, including many of the learned, then there wouldn't be a philosophical *problem* of induction, for if any of the premises were obviously implausible then we could avoid the skeptical result simply by rejecting that

premise. The very fact that so many have grappled with the problem over the past two-and-a-half centuries, and the fact that there is still nothing even close to a consensus on how to solve the problem,[1] should caution us against cavalier dismissals. If anyone thinks the skeptical problem easy to solve, odds are they haven't fully appreciated its force. In chapter two I underscore the difficulty by examining several attempts to solve the negative problem, arguing that all of them fail. There is, I contend, no way to show that inductive inferences can yield epistemically justified beliefs without making substantive assumptions—assumptions that, it seems, would have to be inductively justified.

This negative result is not conclusive, however, and in chapter three I develop a response to the skeptical argument. Since the problem is to provide a negative justification of inductive inference, that is, to show that the skeptical conclusion that inductive inferences cannot yield epistemically justified beliefs can be plausibly denied, I take a close look at both the nature of inference and the nature of epistemic justification. What emerges are accounts of inferential and noninferential justification in the light of which it becomes plausible to deny a key premise of the skeptical argument—namely, that the substantive assumptions upon which inductive inference depends cannot be *noninferentially* justified. More specifically, I articulate and defend an *internalist* definition of inference, according to which to infer requires having a conscious perspective on the relation between premises and conclusion. In defending this way of understanding inference, I develop a general account of epistemic justification in terms of *adequate grounding* and *epistemic responsibility* that plausibly captures, I believe, the core intuitions of both internalist and externalist theories of epistemic justification. One's

---

[1] This is apparent from an examination of the literature, e.g., Salmon 1966, Skyrms 1975, Rescher 1980, and Howson 2000.

grounding for a belief is the objective probability that the belief is formed and maintained in such a way as to be true. This reflects the externalist intuition that epistemic justification should have a *de facto* connection with truth. And one is epistemically responsible in believing to the extent that one exercises control over his epistemic resources in a manner that, as far as he can tell, is conducive to adequate grounding. This reflects the internalist intuition that epistemic justification is significantly dependent on one's conscious, first-person access to information. Combining this account of epistemic justification with an internalist definition of inference leads to accounts of inferential and noninferential justification in view of which it can be plausibly maintained that the substantive assumptions upon which inductive inference depends *can* be noninferentially justified. The plausible rejection of a key premise of the skeptical argument reveals the argument to be unsound and constitutes, therefore, a solution to the negative problem of induction.

What this solution to the negative problem shows is that it is *possible*, for all we know, that induction result in epistemically justified beliefs. This is definitely an encouraging result in the face of the skeptical challenge, but it does not suffice to underwrite our actual inductive practices. To do that we need to solve the *positive problem of induction*, which is to show that it is reasonable to rely on induction as a source of epistemically justified beliefs. A successful response to this constitutes a *positive justification* of induction. This is what I pursue in chapters four through six.

In chapter four, I argue that inferences essentially occur within a context of problem-solving inquiry—as inquiry is the search for answers to our questions, inference is a cognitive event whereby we arrive at an answer to a question on the basis of the information at our disposal. An examination of the structure of the problem-solving

process reveals a correspondence between three stages of inquiry, three types of question, and three types of inference. Peirce called these types of inference abduction, induction, and deduction, but I call them explanatory, evaluative, and explicative inference, respectively, in order to emphasize the role they play in the process of inquiry. Roughly speaking, explanatory inference reasons from effect to a hypothetical cause, explicative inference reasons from a hypothesis to its consequences, and evaluative inference assesses the truth or falsity of a hypothesis by examining the truth or falsity of its consequences. I argue that explicative inference is essentially deductive in character, while explanatory and evaluative inference are *distinct* types of inductive inference. A response to the positive problem of induction, therefore, needs to deal with both. Since inferences are essentially embedded within a process of inquiry and are defined by the role they play in that process, the positive problem of induction takes on a decidedly *methodological* character. That is, our reliance on explanatory and evaluative inference is positively justified if we can responsibly trust them to fulfill their role in the process of inquiry.

In chapters five and six I examine explanatory and evaluative inference, respectively, and argue that we are positively justified in relying on them as sources of epistemically justified answers to our questions. My basic strategy is as follows: First, I argue that in order for explanatory inference to get off the ground we have to accept what I call the *correspondence thesis*, that there exists a significant and fundamental correspondence between knower and known, between our cognitional makeup and the intelligible structure of reality. I argue that acceptance of this thesis is epistemically responsible because it is methodologically necessary—any attempt to responsibly pursue adequate grounding for our beliefs has to assume that reality is *knowable*, and knowable

by us. Second, from the correspondence thesis I derive several important epistemological corollaries, among which are that our cognitive faculties are generally reliable when responsibly deployed and that responsibly obtained appearances are reliable guides to reality. Third, I employ Bayes' Theorem to express the logic of explanatory and evaluative inference in terms of probabilities. Fourth, I use the correspondence thesis and its corollaries to argue that we can make epistemically responsible estimations of those probabilities. It follows that we can responsibly rely on explanatory and evaluative inference as sources of epistemically justified beliefs. Fifth and finally, I try to specify some general conditions under which both explanatory and evaluative inferences lead to epistemically justified beliefs. Thus concludes my response to the positive problem of induction.

II. The Philosophical Significance of These Results

As my mentor, John Greco, points out, "a close analysis of skeptical problems drives positive epistemology".[2] This is evidenced by the history of philosophy, where we see that the epistemological innovations of figures like Descartes, Kant, and Reid were driven in large measure by a concern to meet the challenge of skepticism. And it is evident in contemporary epistemology, where skeptical arguments play an important methodological role by providing test cases for epistemological positions. Given that we do have knowledge or justified belief in most of the cases in which we think we do, any adequate epistemology should be immune to wide-ranging skeptical arguments and should have the resources to show us where such skeptical arguments go wrong. If an

[2] Greco 2000a, p. 4. This claim is a corollary of the more basic truism that *problems drive positive inquiry*, in all areas, not just epistemology.

epistemological position cannot do that then it is, at the very least, incomplete. And if it should lead to skepticism, then that is a decisive reason for rejecting it.

What makes the skeptical problem of induction an interesting problem is that it starts from assumptions that most everyone either accepts or finds plausible and constructs a valid argument for the utterly incredible conclusion that none of our inductively inferred beliefs amount to knowledge or are even so much as epistemically justified. If we are to reject the conclusion, then we must reject something in the argument. But as the premises seem plausible, it is not clear what we can plausibly reject—hence the problem. The problem, moreover, has proven to be highly robust. Many inventive solutions have been many proposed, but it seems that most if not all have failed under close scrutiny. Still, failed solutions can succeed indirectly by teaching us more about the problem and how to solve it. In this dissertation I have drawn together several ideas from others who have struggled with this problem. From Peirce and Rescher[3] I learned that methodological factors are relevant for epistemic justification. From Salmon[4] and Howson[5] I learned that Bayes' Theorem is a useful tool for modeling inductive inference. From Greco[6] I learned that any adequate response to skepticism has to adopt at least a partially externalist position on epistemic justification. And from Williams[7] I learned the importance of distinguishing clearly between the grounding and responsibility aspects of epistemic justification. I have combined these lessons, and others, in a way that, it seems to me, successfully rebuts the skeptical argument and exhibits the rationality of inductive reasoning more clearly, in some respects, than has

---

[3] Rescher 1980.
[4] Salmon 1966.
[5] Howson 2001.
[6] Greco 2000a.
[7] Williams 2001.

been done before. Whether my proposal ultimately succeeds or not, time will tell. But even if it does not, I think a couple lessons can be learned that make a substantive contribution to the ongoing epistemological discussion.

First, discussions of the problem of induction rarely pause to consider the inferential / noninferential distinction. More often then not, it is simply *assumed* that the relevant conception of inference should be what I call an *externalist* one, according to which an inference is, roughly, any cognitive event that takes beliefs as input and generates beliefs as output. As my discussion in chapter three shows, it is far from obvious that this is the conception of inference that we ought to be working with in dealing with this problem. Instead, I argue that an *internalist* definition, according to which inference essentially involves having a conscious perspective on the relation of support between premises and conclusion, is the definition that we ought to work with. Even if I am wrong on that point, it is an issue that to my knowledge has not been seriously raised before and one that, as my solution to Fumerton's dilemma and my rebuttal to Hume show, has significant epistemological and antiskeptical ramifications. Because an internalist definition of inference is narrower than the usual externalist definition, it restricts the scope of inferential justification and expands the scope of noninferential justification. This opens additional space for a noninferential justification of the substantive assumptions on which inductive inferences depends.

Second, nearly all epistemological discussions of inference in general and inductive inference in particular overlook the pragmatic dimension of inference that Peirce noticed and emphasized,[8] and which I explore in chapter four. I argue there that

[8] See Levi 1997 for details.

inferences do not occur in a disinterested logical space, but are essentially embedded within a process of inquiry the aim of which is to solve epistemic problems, and more specifically, to answer questions. When we look at inferences from this angle, I contend, we arrive at a *three-fold* classification of inferences distinguished by the role they play in the process of inquiry in virtue of the type of question that they answer. This is significant for at least a couple reasons. The first is that if inferences have an essentially pragmatic dimension oriented toward the solution of *epistemic* problems, then methodological considerations are relevant to the inferential justification of belief. For in that case the practice of a sound methodology, of good problem-solving technique in the conduct of inquiry, just *is* the epistemically responsible thing to do. The second is that we discover two pragmatically distinct types of induction: explanatory inference and evaluative inference. Because they play different roles in a process of inquiry, different methodological issues arise in each case, which means that a complete solution to the positive problem of induction needs to address both. Most treatments of induction miss this distinction because they approach inference from a primarily or even exclusively formal perspective, in which case, as I point out, the distinction collapses. But the distinction is nonetheless epistemically relevant and thus highlights the need to look at not only the form of an inference but its function as well. If this is correct, then philosophers need to pay more attention to the pragmatics of inference.

More could be said by way of introduction, but enough has already been said to give the reader a sense of the broad outline of the argument. So without further ado, let us get underway by examining the skeptical problem of induction.

CHAPTER ONE

THE PROBLEM OF INDUCTION

Since the publication of Hume's *Treatise* in 1739, what has come to be called the 'problem of induction' has challenged and frustrated generations of philosophers. In that work and in the *Enquiry* that followed, Hume seemingly demolished any pretensions we might have to knowledge or even justified belief concerning unobserved matters of fact. Justified beliefs in unobserved matters must presumably be arrived at by some kind of rational inference, but Hume's argument, if sound, shows that there is no rational reasoning by which we could arrive at such beliefs. Since we think we are justified in believing some unobserved matters of fact (e.g., that the sun will rise tomorrow), either we are wrong in so thinking or Hume's argument must be unsound. The first option is quite hard to accept—surely *some* of our beliefs about unobserved matters of fact are justified? Accordingly, for the past couple centuries philosophers have typically focused on trying to refute Hume's argument. There have been many responses, some highly ingenious, but unfortunately each has met with strong criticism, and there is still no consensus about which response, if any, successfully shows that our inferential practices rest on a firm, rational basis. In this chapter I present and defend a version of Hume's skeptical argument. I do think that the skeptical conclusion is false, and that the Humean argument must therefore be flawed, but I will argue that there are no *obvious* mistakes in the reasoning or in the truth of the premises—*prima facie* the argument looks sound. Both in this chapter and the next I will consider several types of responses and argue that, for one reason or another, each is inadequate, thereby setting the stage for my own proposal.

Hume's Critique of Inductive Inference

We all take ourselves to have knowledge (or, at the very least, justified beliefs) concerning a wide variety of facts the truth of which is not open to our direct inspection. Most of our beliefs about the world are of this variety. Limited as we are by space and time and by the acuity and reliability of our senses, if our knowledge were confined to our present experiences (along with memories of past experiences) we would not know very much. We would know nothing of the future, of the past apart from what we remember, or of present events beyond the current reach of our senses. Yet we do take ourselves to possess knowledge of many unobserved matters of fact, for example, that Napoleon once ruled in France, that the sun will rise tomorrow, that there are craters on the far side of the moon. In each case, the evidence in hand does not *entail* the truth of the belief—the historians' documents might be elaborate forgeries, the sun may go nova tonight, and the moon's surface geology might have undergone some radical transformation since last observed. The truth of unobserved matters of fact is not simply *given* by the evidence in hand—otherwise they wouldn't be *unobserved* matters of fact. Often we arrive at beliefs about unobserved matters of fact through reliance on the testimony of others, but ultimately, it would seem, such beliefs must be grounded in some kind of inference.

Philosophers typically classify inferences as either deductive or inductive. 'Deduction' refers to an inference in which the conclusion is thought to follow *necessarily* from the premises. In a *valid* deductive inference, the conclusion is not only *thought* to follow with necessity from the premises, but actually does so, such that if the premises are true then the conclusion must also be true. The necessity arises from the fact that the conclusion is implicitly affirmed *in* the premises, so that accepting the premises and denying

the conclusion amounts to a contradiction. 'Induction' refers an inference in which the conclusion purports to go beyond what is given in the premises and, therefore, is not regarded as following necessarily from them. So defined, deduction and induction are mutually exclusive and jointly exhaustive; every inference is either deductive or inductive. Deductive reasoning purports to make explicit what is *already* given in the premises. Inductive reasoning purports to carry us *beyond* what is given in the premises. To call a matter of fact unobserved is just to say that its truth is not given to us. Hence if there is a type of inference that can lead us to knowledge or justified belief regarding unobserved matters of fact, it must be inductive.

Hume argues, however, that beliefs arrived at via inductive inference cannot be justified. He is unperturbed by this because he thinks we don't really use or need inductive inference anyway. On his view, beliefs in unobserved matters of fact result from an essentially non-rational process of associating ideas based on custom or habit.[1] This answer looks plausible when dealing with beliefs in unobserved matters of fact that we just seem to arrive at apart from any conscious, deliberate reasoning process, like when a person hears a knock at the door and, immediately, it seems, forms the belief that someone is there. When it comes to other beliefs in unobserved matters of fact, however, such as the belief in a Big Bang, the belief in a heliocentric solar system, or the belief that a meteorite impact led to the extinction of the dinosaurs, Hume's purely psychological account seems obviously inadequate. Such beliefs may turn out to be false, but the processes by which scientists arrived at them clearly involved conscious and deliberate reasoning, and the beliefs do in

[1] "If the mind be not engaged by argument to make this step, it must be induced by some other principle of equal weight and authority. . . . This principle is Custom or Habit. . . . [A]fter the constant conjunction of two objects . . . we are determined by custom alone to expect the one from the appearance of the other. . . . All inferences from experience . . . are the effects of custom, not of reasoning." (Hume 1975, pp. 34-36).

fact appear to be rationally justified.[2] So even if, à la Hume, non-rational processes account for *some* of our beliefs in unobserved matters of fact, that cannot be the whole story. The purely descriptive, psychological account of belief formation needs to be supplemented with a normative and logical account. For not only do we in fact make inductive inferences resulting beliefs about unobserved matters of fact, but we often place high confidence in those results and take ourselves to be rational, reasonable, warranted, or justified in doing so, such that anyone confronted with similar evidence would, on pain of irrationality, be constrained to place a similar amount of confidence in that belief.

The *locus classicus* for Hume's celebrated argument is Book IV of his *Enquiry Concerning Human Understanding*. There he inquires into the nature of the evidential connection between experiential evidence and beliefs about unobserved matters of fact. Hume asks by what reasoning, if any, we could arrive at such beliefs. Starting from plausible intuitions about what that kind of reasoning would have to look like, he is led to conclude that such inferences inevitably self-destruct by falling into vicious circularity:

> These two propositions are far from being the same, I have found that such an object has always been attended with such an effect, and I foresee, that other objects, which are, in appearance, similar, will be attended with similar effects. I shall allow, if you please, that the one proposition may justly be inferred from the other: I know, in fact, that it always is inferred. But if you insist that the inference is made by a chain of reasoning, I desire you to produce that reasoning. The connexion between these propositions is not intuitive. . . .
>
> All reasonings may be divided into two kinds, namely, demonstrative reasoning, or that concerning relations of ideas, and moral reasoning, or that concerning matter of fact and existence. That there are no demonstrative arguments in the case seems evident; since it implies no contradiction that the

[2] There are some, such as Paul Feyerabend and advocates of the so-called 'strong program' in the sociology of science, who contend that science is not fundamentally a rational enterprise. On this view, scientific discoveries owe more to nonrational factors and serendipity than to good logic. For a rebuttal to such views, see Newton-Smith 1981, especially chapters 6 and 10.

> course of nature may change, and that an object, seemingly like those which we have experienced, may be attended with different or contrary effects. . . .
>
> If we be, therefore, engaged by arguments to put trust in past experience, and make it the standard of our future judgement, these arguments must be probable only, or such as regard matter of fact and existence. . . . But that there is no argument of this kind, must appear, if our explication of that species of reasoning be admitted as solid and satisfactory. We have said that all arguments concerning existence are founded on the relation of cause and effect; that our knowledge of that relation is derived entirely from experience; and that all our experimental conclusions proceed upon the supposition that the future will be conformable to the past. To endeavour, therefore, the proof of this last supposition by probable arguments, or arguments regarding existence, must be evidently going in a circle, and taking that for granted, which is the very point in question.[3]

Hume supposes that beliefs about unobserved matters of fact must be based on what we have observed, namely, past and present observations. What we observe, however, can tell us nothing about what we haven't observed, unless we introduce an *auxiliary assumption* (AA) such as "what is observed is a reliable guide to what is unobserved." For example, if we are to infer any beliefs about the future based on past and present observations, we must suppose the auxiliary assumption that the future will resemble the past. But how could we know that such an AA is true? That the future will resemble the past and that what is observed is a reliable guide to what is unobserved are themselves unobserved matters of fact. As such, our alleged knowledge of those facts must be grounded in an inference from past and present observations, together with a suitable auxiliary assumption. We can't, however, invoke the very same AA that we employed in the original inference to justify the AA itself, for that would be viciously circular and render the inference epistemically worthless. What if we employ a different auxiliary assumption (call it AA') in the process of trying to justify the original auxiliary assumption, AA? In this case the same problem arises with respect to AA'—how do we know it to be true? Since it has to say something about what we *haven't*

[3] Hume 1975, pp. 78-80.

observed in order to connect our observations with the unobserved matter of fact (namely, AA), it cannot be known simply by observation. Thus we need another inference with another AA (call it AA") to justify AA'. And so on. We wind up in either vicious circularity or an infinite regress. The upshot is that the AA cannot be justified and that, therefore, the inferential process leading from observed to unobserved matters of fact cannot yield justified beliefs. Our beliefs about unobserved matters of fact derive from custom or habit, not sound reasoning. In terms of their relation to the available evidence, such beliefs are no more rational than their denials.

I offer here a reconstruction of Hume's argument. Actually, it is more precisely described as "broadly Humean," for my concern in formulating the argument is not so much to develop accurately what Hume actually said, but more so to formulate the strongest case against the inductive justification of belief that I can muster, with Hume serving as the guiding inspiration.[4]

(SA1) 1. All inductive inferences require as a premise a contingent auxiliary assumption (AA).
2. A belief is inferentially justified only if belief in the premises is justified.
3. Therefore, a belief is inductively justified only if belief in the AA is justified. (1,2)
4. There is no way of justifying belief in the AA.
   a. All justification of belief is either inferential or noninferential.
   b. Belief in the AA cannot be noninferentially justified (it is not self-evident, evident to the senses, incorrigible, or otherwise properly basic).
   c. Inferential justification is either deductive or inductive.
   d. Belief in the AA cannot be justified by deduction.
   e. Belief in the AA cannot be justified by induction.
   f. Therefore, the AA cannot be inferentially justified. (c,d,e)
   g. Therefore, the AA cannot be justified. (a,b,f)
5. Therefore, beliefs arrived at via inductive inferences are unjustified. (3,4)

---

[4] There are other ways of developing Hume's argument. Greco, for example, frames it as an argument against *knowledge* of unobserved matters of fact (Greco 2000a, pp. 140ff.). I frame it as an argument against *epistemic justification* because that seems to me the strongest way to state the argument, for if we do not even have justified beliefs concerning unobserved matters of fact, then *a fortiori* neither do we have knowledge.

Argument (SA1) is clearly valid. We now have to examine the truth of the premises to see whether it is also sound. The main premises that we need to examine are (1), (2), and (4). In the subargument for (4), premises (a) and (c) are true by definition of terms, leaving us (b), (d), and (e) as premises to examine. Therefore, if there is a flaw in the argument, as I suppose there must be since I think the conclusion is false, it must lie in (1), (2), (4.b), (4.d), or (4.e). I will argue, however, that each of these premises is at least *prima facie* plausible and that, therefore, it is not obvious where the fault in the argument lies.

Premise (1): Auxiliary Assumptions Are Necessary

Let us consider the matter abstractly. Let us say that evidence (E) is relevant to conclusion (C) if and only if the truth of E makes a difference to the probability of C, i.e., Prob(C|E) ≠ Prob(C). Let us say further that E is *positively relevant* to C if and only if Prob(C|E) > Prob(C). Now, in every inductive inference the conclusion goes beyond the information in hand, that is to say, there is some set of propositions, X, entailed by the conclusion that is not entailed by the premises. Let us call X the *inductive extension* because it denotes that respect in which the inductive conclusion goes beyond the evidence. With regard to any particular induction, it seems reasonable to ask whether the evidence is in fact positively relevant to the affirmation of X in the conclusion. What warrants our inductively extending the conclusion beyond the evidence *in the way that we do*? In making the inference from evidence to X, something not given in the evidence, we clearly suppose that they are connected somehow. Since both evidence and the conclusion represent putative facts about the world that *ex hypothesi* are only contingently related, and since a merely

accidental connection would be too weak,[5] a natural way to conceive of the connection is in terms of some nomic regularity. For example, if we conclude from a sample of black ravens that all ravens are black, we assume that blackness is not simply an accidental characteristic of the ravens we have observed. Instead, we suppose that there is some kind of real connection between being a raven and being black, i.e., that ravens, as such, tend to be black. Furthermore, when we infer from evidence to conclusion, we clearly suppose that our evidence is reliable (i.e., probably and approximately true) and, by extension, that the means, methods, and processes by which we acquired the evidence are reliable.[6] If we thought that our data-acquisition methods were systematically biased or unreliable we would not regard the results as good evidence for inferring anything. So it seems that in making inductive inferences we assume as an AA both that our data (and therefore our data-acquisition methods) are reliable and that they are positively relevant to the inductive extension by virtue of some nomic regularity, or something functionally equivalent.[7] Evidence that is either unreliable or irrelevant affords us no reason *at all* for trusting our inductive extensions beyond the evidence.

The notions of reliability and regularity require further analysis in two respects. First, they admit of variations by *degree*. For something to be reliable means, minimally, that *most*

---

[5] If it were mere coincidence that the sun has arisen every day in the past, then that would afford us no more reason to think it will come up tomorrow than a fair die landing six ten times in a row affords us a reason to think it will land six the next time.

[6] According to the standard usage of the term among analytic epistemologists, a process, cognitive faculty, method, or inference-type is *reliable* if it tends to produce true beliefs or accurate results. When I describe beliefs, premises, or evidence as *reliable* I mean that they are true or approximately true and therefore provide a trustworthy evidential basis for making inferences. These two uses of the term are closely related: a belief-forming process is reliable (i.e., truth-conducive) only if it generally produces beliefs that are reliable (i.e., probably and approximately true). Conversely, while the *de facto* reliability of a belief does not require that it be produced by a reliable process or faculty, our *confidence* in the reliability of our beliefs implies confidence in the reliability of the processes or faculties through which the beliefs are formed.

[7] The assumptions of reliability and relevance correspond to the two necessary conditions for inductive cogency: true premises and strong support for the conclusion.

of the time when we *rely* on it the results are positive. Thus, in this sense, a belief is reliable if it is probably true, and a belief-forming process is reliable if it yields true beliefs most of the time. We could, however, understand reliability in a stronger sense to require positive results all or almost all of the time. And regularities can be understood as exceptionless laws or as statistical generalities that admit of exceptions. Second, reliability and regularity can be understood differently with respect to their generality or *scope*. We could, for example, suppose when inferring that our data-acquisition methods are reliable *in general*, or we could suppose further that the specific methods employed are reliable are this *particular* occasion. Similarly, regularities can be highly general, like "the future will resemble the past," or specific, like "ravens, as such, tend to be black." Given that inductive inferences require an AA to the effect that the evidence is reliable and relevant to the conclusion by virtue of a regularity, it makes sense to ask how strong that assumption must be: What minimal degree of reliability and regularity must we suppose? And how specific must the AA be? With respect to these questions, I will argue that the AA need not commit us to a maximal or near-maximal degree of reliability and regularity—it is only necessary that our evidence be trustworthy for the most part and that the regularity that grounds their relevance hold for the most part. I will also argue the AA must be specific—vague generalities like "the future will resemble the past" are inadequate for grounding the particular inferences that we make.

With respect to the question of degree, it seems that we need not require both a high degree of reliability for the evidence and a high degree of regularity to ground their relevance. Intuitively, our confidence in the conclusion of an inference should be directly proportionate to our confidence in the evidence (how reliable we think it is, i.e., the probability that the evidence is true) and to how strongly the evidence supports the

conclusion (its degree of relevance) if it is true. To see this, consider that inference is, among other things, a transmitter of confidence from evidence to conclusion. The assumption of an exceptionless regularity will secure maximal relevance and thus transmit all of our confidence in the evidence to the conclusion. A weaker regularity means less relevance, and thus the transmittal of only part of that confidence. But one cannot transmit what is not there to begin with—thus, we require some degree of confidence in the evidence. The greater that confidence, the more is available to transmit. According to this simplified model, our degree of rational confidence in conclusion (C) based on evidence (E) should approximately follow the equation: $\text{Prob}(C) = \text{Prob}(E) \times \text{Prob}(C \mid E)$.[8] In other words, degree of reliability of the evidence, Prob(E), times its degree of relevance, Prob(C | E) (which depends on the strength of the regularity tying evidence and conclusion), equals the degree of rational confidence in the conclusion, Prob(C). On this formula, Prob(E) and Prob(C | E) must both be greater than 0.5 if the conclusion is to be probable, but they may also both be less than 1. Inductive inferences, therefore, do not presuppose maximally reliable evidence or exceptionless regularities.

As to how specific the AA must be, Hume and many following him (e.g., Mill, Russell, etc.) have apparently thought that a vague and general principle like "the future will resemble the past" is all we need to make our inductive inferences. If only we could justify

---

[8] According to the probability calculus, the exact equation should be $\text{Prob}(C \,\&\, E) = \text{Prob}(E) \times \text{Prob}(C \mid E)$, but this difference does not negatively affect my conclusion that inductive inference does not require either maximal reliability or maximal regularity. On the one hand, when C entails E, Prob(C & E) = Prob(C), thus reducing to my equation. In many cases of induction this simplification will be appropriate. For example, in inference to the best explanation we would like, all other things being equal, our concluding hypothesis C to entail the evidence E that we are looking to explain. On the other hand, when C does not entail E, we can represent Prob(C) by an inversion of Bayes' Theorem thus: $\text{Prob}(C) = [\text{Prob}(E) \times \text{Prob}(C \mid E)] / \text{Prob}(E \mid C)$. Because in this case C does not entail E, the denominator must be less than one, meaning that Prob(C) can be greater than 0.5 even when Prob(E) and Prob(C | E) are *less than* 0.5. So in neither case do Prob(E) and Prob(E | C) have to equal one in order for Prob(C) to be greater than 0.5.

some grand Principle of the Uniformity of Nature (Mill) or a Principle of Limited Independent Variety (Keynes), then the problem Hume posed would be solved. Others have apparently thought that Hume required a specific AA, and that this would make the resulting inference deductive rather than inductive. For example, David Stove writes:

> Sometimes when we say of an argument from *p* to *q*, that it presupposes *r*, our meaning is as follows: that, as it stands, the argument from *p* to *q* is not valid, and that, in order to turn it into a valid argument, it would be necessary to add to its premisses the proposition *r*. I believe that this is the sense in which "presuppose" occurs in . . . Hume's argument.
>
> Hume's argument . . . may therefore be summed up in the following way: from premises which prove at most the invalidity of predictive-inductive inferences, along with the unstated premiss that an inference is unreasonable if it is invalid, Hume concluded that predictive-inductive inferences are unreasonable.[9]

To require that all inductions be turned into deductions is tantamount to assuming that beliefs arrived at via inductive inferences, as such, are not justified. In a context where one is supposed to be trying to prove this, such an assumption is question begging. Accordingly, many have thought it a sufficient rebuttal to Hume to simply point this out. Robert Audi, for example, writes:

> When we come to the problem of induction, it seems clear that one assumption the skeptic is making is that if we believe something on the basis of one or more of the premises, then we can know it on the basis of those premises only if it follows from them, in the sense that they entail it. Call this the *entailment principle*. It says in effect that *knowledge can be transmitted only deductively*. . . . I do not see that skeptics give us any good reason to believe . . . the entailment principle.[10]

I think these views are wrong. Induction requires a specific AA, but addition of a specific AA does not suffice to make the inference deductive—for that we would need not merely specific regularities but *exceptionless* ones. Attempts to rebut Hume via some grand

---

[9] Stove 1973, pp. 43, 51.
[10] Audi 1988, pp. 300, 302.

metaphysical principle like the Uniformity of Nature as well as claims that premise (1) of (SA1) immediately implies deductivism are, therefore, mistaken.[11]

I maintain that induction requires a specific AA for two reasons. First, it is very difficult, if not impossible, to formulate a general Uniformity Principle in such a way that it remains both plausible and useful, as Peirce points out:

> Let us now inquire what the "uniformity of nature," with its synonymous expressions that "the future resembles the past," and so forth, can mean. Mill says that it means that if all the circumstances attending two phenomena are the same, they will be alike. But taken strictly this means absolutely nothing, since no two phenomena ever can happen in circumstances precisely alike, nor are two phenomena precisely alike. It is, therefore, necessary to modify the statement in order to give it any meaning at all; and it will be found that, however it may be so modified, the moment it begins to carry a definite meaning, one of three things results: it becomes either, first, grossly false, or, second, an assertion which there is really no good reason to believe even approximately true, or, thirdly, it becomes a quasi-subjective truth, not lending any colour of validity to induction proper. . . . In fact, the great characteristic of nature is its diversity. For every uniformity known, there would be no difficulty in pointing out thousands of non-uniformities.[12]

Take "the future will resemble the past," for instance. What do we mean by "resemble"—complete sameness, partial sameness (how partial?), sameness in those aspects that are most fundamental (such as the laws that govern reality), analogical similarity (i.e., past and future vary in certain respects along some continuum), or something else? Do we mean the near/distant future will resemble the recent/distant/observed past? As it stands, this principle is so utterly vacuous that nothing determinate follows from it whatsoever. Thus, it cannot

---

[11] In the following chapter I will show that premise (1) as stated in (SA1) does *implicitly* lead to deductivism, but the skeptical argument can be easily modified to avoid this problem. My point at present is simply that the route from premise (1) to deductivism is not immediate, i.e., inductive inference does not require as a premise an AA that, simply by its inclusion, would suffice to turn the inductive inference into a valid deduction.

[12] CP 6.100. References to Peirce's writings will be given in what has become the standard scholarly form. "CP" refers to the eight-volume set of his *Collected Papers*, "W" refers to the six-volume (so far) chronological edition of his *Writings*, and "EP" refers to the recent two-volume set entitled *The Essential Peirce*. In each citation the appropriate abbreviation of the work will be given, followed by the volume number, a period, and then the page number (or paragraph number in the case of the *Collected Papers*).

perform its job of grounding the relevance of evidence to conclusion. For that we have to clarify what we mean by "resemble" and define the proper scope of the past and the future in such a way that the resulting Uniformity Principle remains both plausible and useful. Peirce's point is that it is far from obvious how this is to be done. Moreover, it is not clear that there is any meaningful characterization of the "uniformity of nature" other than as a summary term denoting the set of specific regularities that pervade it—laws of physics, chemistry, biology, psychology, etc.

Second, it seems to be a more accurate description of how we actually reason to say that inductive inferences invoke *specific* AA's. When we reason inductively, we reason from some particular set of evidence to some particular matter of fact as our conclusion. If I conclude from my observations of ravens that all ravens are black, it is because I think that *my raven evidence* is fairly reliable and that *the color of ravens* is somehow closely tied to their ravenhood. If I reason that the sun will rise tomorrow based on past experience, it is because I think that *my past experience of the sun's rising* is reliable and that *the relative motion of the sun with respect to the earth* is grounded in some kind of regularity or natural law. Furthermore, when we criticize particular inductive inferences (e.g., an extrapolation from polling evidence), we do so by questioning either the reliability of the specific evidence in question (e.g., the questions were misleading, the respondents had some incentive to deceive, etc.) or their relevance (e.g., the number polled is too small, the types of respondents were not sufficiently varied, etc.). This is not to say that we do not in some sense assume that the world is *generally* intelligible and therefore pervaded by regularities, or that certain means of acquiring evidence are *generally* reliable (e.g., perception, certain authorities, a specific sampling methodology, etc.). My argument is that, whatever more

general regularity or reliability assumptions may be tacitly involved in our reasoning, such reasoning must *also* assume the reliability of the *particular* evidence on which the inference is based and its relevance to the conclusion in question.

In sum, premise (1) seems correct. Inductive inferences presuppose an AA to the effect that the evidence is reliable and positively relevant to the inductive extension by virtue of some regularity. This AA must be case-specific, reflecting on the reliability and relevance of the specific evidence involved in the inference. Both reliability and relevance admit of degrees, but inductive inferences do not require assumptions of maximal reliability or relevance.

Premise 2: The Premises of an Inference Must be Justifiably Believed

Premise (2) states "A belief is inferentially justified only if belief in the premises is justified." This is plausible as it stands and I am not aware of any serious challenges to it. In fact, it is generally taken for a platitude. Laurence BonJour, for example, writes, "it is obviously a necessary condition for inferential justification that the beliefs appealed to as premises be themselves *already* justified in some fashion."[13] This makes sense. Suppose that premise (2) is false, that is, suppose that a conclusion could be justified by inference from premises that we had no good reason for thinking were true. In that case it seems that inferential justification would be far too easy, for one could then just help oneself to whatever premises were necessary to support the conclusion without needing to defend those premises. As far as I can see, the only context in which the premises of an inference may be exempted from the need for justification is in the case of suppositional reasoning, where we *assume* that certain premises are true and reason out the consequences. The strength of

[13] BonJour 1986, p. 117.

suppositional reasoning is that we don't have to justify the premises, the weakness is that by it we cannot justify belief in the *truth* of the conclusion—all we can say is that *if* our assumptions are correct, then the conclusion would be justified. In all other cases, where we reason to the truth of some conclusion (some matter of fact), suppositional reasoning is not relevant for, in Humean terms, it only deals with relations of ideas. Insofar as inductive reasoning concerns matters of fact (and it always does), premise (2) is correct.

Premise 4: There Is No Way to Justify Belief in the Auxiliary Assumption

Given the truth of premises (1) and (2), it follows that a belief is inductively justified only if belief in the AA is already justified. Premise (4) claims that there is no way of justifying belief in the AA. Premise (4.a) states the platitude that justification is either inferential or noninferential. In standard foundationalist terms, let us say that a belief that is noninferentially formed is 'basic' and that a belief that is basic and justified is 'properly basic.' Premise (4.b) claims that belief in the AA is not properly basic. Philosophers have proposed different types of properly basic beliefs, which fall mainly into one of three categories: (i) self-evident (e.g., 2+2=4), (ii) evident to the senses (e.g., "Look, there is a cat on the mat"), and (iii) incorrigible (e.g., "I feel pain" or "I seem to see something blue").[14] Now, whatever AA is involved in a particular inductive inference, it does not seem that it could be any of these three types. For the sake of argument, let's take the AA to be the proposition "Observed cases of X are a reliable guide to unobserved cases of X." With regard to (i), it is at least conceivable, as Hume points out, that the course of nature might undergo some sudden revolution, such that what has been observed to hold may cease to provide any reliable evidence regarding what has not been observed. Accordingly, the AA is

[14] Cf. Plantinga 1983, pp. 55-59.

a contingent truth, if true at all. Since contingent truths, as such, cannot be known to be true simply by understanding them, they are not self-evident. Hence, the AA is not self-evident. With regard to (ii), the AA is needed to link the evidence (observed cases) with a conclusion that goes beyond the evidence (unobserved cases). It must, therefore, be in part about something unobserved. As such it cannot be a perceptual belief, or something evident to the senses. With regard to (iii), the AA is clearly not incorrigible. The standard examples of incorrigible beliefs are generally introspective reports (e.g., "I am in pain" or "I am being appeared to redly"), which the AA is not. It is easy to imagine someone believing the AA and being mistaken; hence, it is not incorrigible. In principle there may be other types of properly basic beliefs. If the AA's required by inductive inferences can be shown to fall into one of those other types, then premise (4.b) would have to be rejected. But it is not obvious what other types of properly basic beliefs there might be, and for any proposal along those lines we would need some principled account of why beliefs of that sort are properly basic. *Prima facie*, then, premise (4.b) looks to be true.

What about inferential justifications of the auxiliary assumptions? We have two possibilities here: justify the AA by deduction, or justify the AA by induction. The problems with the deductive approach are twofold. First, deductive justification by itself is never *final*—it merely passes the justificatory buck, so to speak. The argument runs as follows:

(4.d.i) 1. Deduction can only draw out what is already in the premises.
2. Therefore, any deductive proof of the AA must assume premises that have at least as much content as the AA.
3. Therefore, the problem of justifying those premises must be at least as great as that of justifying the AA in the first place.
4. Therefore, deductive proofs of the AA merely postpone the problem and may even make it worse.

Since valid deduction can only lead us to a conclusion that is already contained in the premises, any attempt to argue deductively for an AA like "observed cases of X are a reliable guide to unobserved cases of X," would have to assume premises that have at least as much content as the AA in question. Those premises would, therefore, be at least as much in need of justification as the AA itself. It the premises happen to have more content than the AA, then the problem of justification becomes worse than it was at the outset. Mathematicians get around this problem by simply stipulating the truth of their ultimate premises, the axioms. But whatever form the AA takes, it will have to be a *substantive* thesis about the world, bearing on the reliability of the evidence and its factual relevance to the conclusion by virtue of some regularity. Accordingly, the ultimate premises of any deductive argument for the AA must also be substantive theses about the world—and we can't establish substantive truths (or any truths for that matter) by stipulation. Could we establish these ultimate premises by some noninferential means like observation? Unless we can find a way around premise (4.b), it appears not, which brings us to a second problem:

(4.d.ii) 1. The AA has to take us beyond the evidence because it serves to connect that evidence with an inductive conclusion that transcends the evidence.
2. Therefore, any deductive proof of the AA will require us to assume premises that go beyond the evidence.
3. Since deduction can never take us beyond the evidence, the truth of those premises cannot be established by deduction.
4. Therefore, deductive proofs of the AA are in principle *incomplete*.

We can't establish the ultimate premises of a deductive justification of the AA by observation because to perform its function the AA must be to some extent an *unobserved* matter of fact—it must go beyond what is entailed by the available evidence. But since for deduction to be valid all the content of the conclusion must be given in the premises, at least some of those premises must also be unobserved matters of fact. Because deduction by its

very nature is nonampliative—it cannot take us beyond the available evidence—the truth of those premises, and therefore the truth of the AA, cannot be established by deduction. In other words, any purely deductive justification of the AA must be *incomplete*, requiring us to fall back on some non-deductive means of establishing its premises. Since observation and stipulation won't suffice, we seem to be left with induction, which is just what we were trying to avoid. In sum, deductive justifications at best postpone the problem, may even make it worse, are in principle incomplete, and seem to require induction to establish at least some of the premises. Premise (4.d) looks good.

Suppose we try to justify our AA inductively. Here, too, seemingly insurmountable problems arise. Let's lay out the basic problem abstractly and then illustrate with a specific example.

(4.e.i)
1. All inductive inferences require as a premise a contingent auxiliary assumption (AA).
2. Therefore, any inductive justification of the AA must also presuppose an AA (AA').
3. AA' must either be the same as AA or different.
4. If AA' is the same as AA, then the argument is viciously circular; AA is being assumed in the course of proving AA.
5. If AA' is different from AA, then AA' itself stands in need of justification by a prior inductive inference that must also presuppose an AA (AA"). (1)
6. And so on. There results an infinite regress, unless the same AA recurs, at which point the argument becomes viciously circular.
7. Therefore, the AA cannot be justified by induction.

Since it is only contingently true, whatever justification the AA possesses must come from factual, presumably empirical, evidence—i.e., past and present observations. Our inductive argument for the AA would then look like this:

> We have often observed that observed cases of X have been a reliable guide to unobserved cases of X.
> Therefore, (AA) observed cases of X are a reliable guide to unobserved cases of X.

But perhaps our past successes with the AA have just been good luck. If we think that our observed successes give us good reason to expect successful application of the AA in as yet unobserved cases, it must be because we think (AA') that observed cases of the application of AA are a reliable guide to the application of AA in unobserved cases. Now, we can take AA' in either of two ways. First, as just an instance of the original assumption (AA). Thus:

> We have often observed that observed cases of X have been a reliable guide to unobserved cases of X.
> (AA) Observed cases of X are a reliable guide to unobserved cases of X.
> Therefore, (AA) observed cases of X are a reliable guide to unobserved cases of X.

In this case the inference is viciously circular. Or, we can take AA' as a second-order auxiliary assumption, distinct from the first-order auxiliary assumption (AA) that we are trying to justify. Thus:

> We have often observed that observed cases of X have been a reliable guide to unobserved cases of X.
> (AA') Observed cases of the application of AA are a reliable guide to the application of AA in unobserved cases.
> Therefore, (AA) observed cases of X are a reliable guide to unobserved cases of X.

This version is not circular, but it threatens an infinite regress, for the problem of justifying AA' is exactly the same that we faced when trying to justify AA. It seems that we will need an inductive argument to support AA' with a third-order auxiliary assumption (AA"), and so on. If the auxiliary assumption employed at any level is logically identical to that employed at any subsequent level, viciously circularity results. If not, we are stuck in an infinite regress. The problem here results from acceptance of the first premise, which is the same as premise (1) of the main skeptical argument. But we have already argued that premise (1) is correct. Therefore, we cannot, it would seem, inductively justify our AA's.

Summing up our discussion of premise (4), if there a way to justify belief in an AA, it must be either inferential or noninferential. If inferential, it must be either deductive or inductive. There are no other possibilities. But none of the possibilities seems adequate to the task. Premise (4), therefore, looks plausible.

Conclusion

Hume's argument is deductively valid—if we concede the premises, then we have no choice but to accept the conclusion as well. Furthermore, each of the premises looks plausible. Yet the conclusion that beliefs arrived at via inductive inferences are *ipso facto* unjustified strikes most of us as clearly unacceptable. This is why philosophers speak of the "problem" of induction. Any solution of it will necessarily require us to reject a seemingly plausible premise, which means that we will have to defend a more or less counterintuitive thesis. Whatever approach one takes will therefore require us to rethink some of our common assumptions about inductive inference and about inferential justification. In this way, as Greco rightly emphasizes, skeptical arguments drive positive epistemology.[15] In the following chapter I will look at several common strategies for responding to Hume's argument and will argue that, unfortunately, none of them work. Their failure, however, will motivate taking a step back to rethink the problem and thereby point the way toward a different sort of approach.

We can, perhaps, make one forward-looking observation at this point. At the outset I identified five premises of (SA1) that required defense: (1), (2), (4.b), (4.d), and (4.e). All of these, I have argued, look plausible, but if we are trying to find one to reject so that we can avoid Hume's skeptical conclusion, some of these seem to be more promising candidates for

[15] Greco 2000a, pp. 2-4.

reexamination, particularly premises (1) and (4.b). Premise (1) deserves further scrutiny simply because it is the primary lynchpin of the whole argument—if we can plausibly deny that inductive arguments require AA's as premises then it is no longer clear why we should need to justify them, and it also opens up the possibility for a noncircular inductive justification of inductive presuppositions. As we will see in the next chapter, this is the approach that reliabilism takes. Premise (4.b) deserves further scrutiny because the case for it is inconclusive—I argued that AA's do not fit into any of the standard categories of properly basic beliefs but left it open whether there might be some other, hitherto overlooked, category of properly basic beliefs that could either include our AA's or that could supply premises for an inferential justification of our AA's.

CHAPTER TWO

AN EXAMINATION OF SOME RESPONSES TO THE PROBLEM

There have been many different responses to the problem of induction and many critical discussions of the various responses.[1] In this chapter I examine the major responses and argue that each is inadequate. Since my remarks largely follow what others have already said, I generally keep my comments fairly brief and to the point. The purpose of this chapter is not to give a detailed treatment of each of these proposals—that would take a book in itself—but to clarify the nature of the problem and to clear the ground, as it were, before laying the foundation for my own proposal in the following chapters. One response, reliabilism, will occasion a somewhat lengthy discussion of inference rules, leading to a significant refinement in the skeptical argument. In classifying the various types of responses I follow the handy catalog prepared by Nicholas Rescher.[2] He groups responses under four main categories, which can be defined by the following 2×2 matrix:

Induction . . .	Doesn't need to be justified	Needs to be justified
Can't be justified	I	
Can be justified	II	III

The grayed out section is of no interest to us here because it represents the skeptical position. The other three sections classify different ways in which philosophers have responded to the skeptical argument. I will now examine these responses in turn and argue that neither the

[1] See, for example, Salmon 1966, Skyrms 1975, Rescher 1980, and Howson 2000.
[2] Rescher 1980, p. 187.

view that induction doesn't need to be justified (I-II) nor the major attempts so far to justify it (II-III) are adequate.[3]

I. Inductive Inference, Who Needs It?

First, there is the view championed by Karl Popper that inductive inference cannot be justified and does not need to be justified because it is *dispensable*. Popper embraces Hume's skeptical conclusion and argues that we should replace our reliance on inductive forms of reasoning with a falsificationist methodology that employs only deductive reasoning. In response to a problem situation, says Popper, we conjecture tentative theories (which are not in any way supported by the data of the problem situation) and then attempt to falsify them by obtaining experimental results that conflict with what the theories predict. Theories that withstand our best attempts at falsification are 'corroborated' and worthy of tentative endorsement.[4] Corroboration, however, is "just a report about the state of discussion at the time *t*, concerning the logical and empirical preferability of competing theories."[5] It tells us nothing about the truth of the theory or its reliability in the future: "[W]e can determine at most the falsity of theories."[6] "By this method of elimination, we may hit upon a true theory. But in no case can the method *establish* its truth, even if it is true; for the number of *possibly* true theories remains infinite, at any time and after any number of crucial tests."[7]

There are several problems with this proposal. To begin, it does not square with actual inferential practice. Scientists do not in fact strive to refute their theories in the manner

[3] Here and throughout this chapter, when I speak of a 'justification of induction' I mean a *negative* justification, i.e., a rebuttal to the skeptical argument showing that it is *possible*, for all we know, for induction to yield epistemically justified beliefs. So to say that induction can be justified is to say that it is possible to rebut the skeptic, and to say that induction needs to be justified is to say that it is epistemically important to rebut the skeptic because otherwise it undermines the justification of our inductively formed beliefs.
[4] Popper 1979, pp. 18-19, 22.
[5] Popper 1979, p. 19.
[6] Popper 1979, p. 12.
[7] Popper 1979, p. 15.

Popper outlines, nor do they reject a theory just because it makes some predictions that do not square with experimental results, nor do they regard their theories as lacking positive evidential support, nor is there any compelling reason for thinking that they *ought* to do so.[8] Additionally, according to the well-known Duhem–Quine thesis, any hypothesis can be saved from falsification provided we are prepared to make suitable modifications to other parts of our noetic structure.[9] When we test a theory we derive predictions from the theory *in conjunction with* various background beliefs. If the experimental test fails, should we reject the theory or one or more of the background beliefs? Deductive logic alone cannot answer this question. Whatever rational grounds we have for our decision have to be based on our positive evidence for the theory in conjunction with background beliefs. But insofar as theory and background beliefs are more than just summary statements of observations, their relation to the evidence is logically ampliative and therefore inductive. Popper's rejection of inductive reasoning leaves the decision to reject a theory not only ungrounded but ungroundable.[10] Since Popper insists that there always remains an infinite number of competing, unfalsified theories, and since corroboration tells us nothing about the future reliability of a theory, it gives us no rational grounds for preferring a theory over any of its unfalsified competitors. If it is to provide a rational basis for theory preference, corroboration must necessarily involve inductive reasoning. Wesley Salmon aptly sums up the situation:

> Popper has not succeeded in purging the logic of science of all inductive elements. . . . Corroboration is a nondemonstrative form of inference. It is a way of providing for the acceptance of hypotheses even though the content of these hypotheses goes beyond the content of the basic statements [i.e., observations and test results]. *Modus tollens* without corroboration is empty; *modus tollens* with corroboration is induction.[11]

---

[8] Newton-Smith 1981, pp. 72-75.
[9] Quine 1980, pp. 64-65. See also Duhem 1991, pp. 183-188.
[10] Newton-Smith 1981, pp. 61-63.
[11] Salmon 1966, p. 26.

In conclusion, then, induction is not dispensable.

II. Inductive Inference Need Not Be Justified

The second type of response rejects the skeptical demand for a justification of induction while maintaining, in contrast to Popper, that induction can be (or just is) justified. There are two positions in this category that I will consider. According to the *analytic rationalism* of Strawson and Ayer, all justification must terminate somewhere and inductive inference is simply part of our epistemic foundation. We use it to justify other beliefs, but induction itself does not need to be justified. It *is* justified because to reason inductively is part of what it means to be rational. According to *reliabilism*, a relatively recent view that has been defended by Van Cleve, Papineau, and Greco, among others, our inductive procedures are *in fact* reliable, and this is sufficient for justification. Since reliabilist justification is externalist,[12] we do not need to know or be able to show that it is reliable to be justified in using it. Given this, it is argued, we can show that induction is reliable *by induction*, without falling into vicious circularity. Because reliabilism offers a direct challenged to premise (1) of (SA1), I will discuss it at greater length.

II.a Analytic Rationalism

Analytic rationalism says that attempts to justify inductive inference are misguided because the epistemic norms governing inductive inference *define* rationality in those

[12] Externalism and internalism are two broad perspectives on the nature of epistemic justification. Roughly speaking, internalists say that whether S is justified in believing *p* is a matter of whether the belief is well-formed from a *first-person* perspective. If S is suitably acquainted with the evidence on which his belief that *p* rests, then S is justified in believing *p*. In contrast, externalists say that whether S is justified in believing *p* is a matter of whether the belief is well-formed from a *third-person* perspective. Thus, if S's belief that *p* was, e.g., formed by a reliable process, then it is justified. For a good introduction to the internalist/externalist controversy in epistemology see Kornblith 2001.

contexts. To try to justify inductive inference is thus rather like asking "Is the law legal?" In Strawson's words:

> [T]o ask whether it is reasonable to place reliance on inductive procedures is like asking whether it is reasonable to proportion the degree of one's convictions to the strength of the evidence. Doing this is what 'being reasonable' *means* in such a context.[13]

But even if 'being reasonable' does entail 'relying on induction', all that tells us is something about *how we use the words* 'induction' and 'reasonable'. It tells us nothing about whether our inductive procedures reliably put us in touch with the non-linguistic world.[14] There are certainly imaginable alternatives to inductive inference as a means for answering our questions about the world—methods like crystal ball gazing, random guessing, consulting a ouija board, asking one's closest neighbor, etc. Given such alternatives it seems reasonable to ask for a justification of inductive inference that, beyond simply *stating* that our reliance on inductive inference is rational, *explains why* it is so:

> No doubt induction is reasonable by the *established* canons of rationality, but we may still wonder whether in this instance these canons are themselves *appropriate*. To revert to Strawson's analogy: it is indeed senseless to ask why the legal system is *legal*, but it is perfectly natural and proper to wonder whether the legal system is *just*—that is, whether it is able to attain the ends for which legality is instituted in the first place.[15]

II.b Reliabilism

An increasingly popular line of response to inductive skepticism is to invoke a reliabilist conception of justification.[16] In its simplest form, reliabilism is the view that epistemic justification is purely a function of the *de facto* reliability of the cognitive faculties

[13] Strawson 1952, ch. 9. Italics in original.
[14] Howson 2000, p. 16.
[15] Rescher 1980, p. 215.
[16] See, for example, Goldman 1986, Van Cleve 1984, Papineau 1993, and Greco 2000a.

and processes that give rise to our beliefs.[17] Accordingly, reliabilists argue that a person is inductively justified in a belief if the belief is formed via induction, the person is justified in believing the premises, and *in fact* induction reliably produces true beliefs given true premises.[18] Since reliabilist justification is externalist, one doesn't have to know or show or (on some accounts) even believe that induction is reliable to be justified in using it. Accordingly, if induction is in fact reliable, and reliability is sufficient for justification, we can justifiably infer inductively from "Most A's have been found to be B's" to "Most A's are B's" without invoking *as premises* any contingent AA's.[19] Reliabilism thus rejects premise (1) of (SA1). Moreover, by means of what Van Cleve labels 'Argument A', the same principles allow us to establish the reliability of induction inductively *without falling into vicious circularity*:

> Most of the inductive inferences I have drawn in the past from true premises have had true conclusions.
> Hence, the majority of *all* inductive inferences with true premises have true conclusions.[20]

Thus, we don't need to be justified in believing that induction is reliable, but we can be.[21] While it is true that the inductive argument for the reliability of induction makes use of an inductive inference rule (in this case, the so-called 'straight rule' of enumerative induction) to argue for the reliability of that very inference rule, such 'rule-circularity,' it is argued, is not vicious.[22] If it were, argue Van Cleve and Papineau, then we couldn't justify induction *or*

---

[17] More precisely, we want our sources of knowledge to be reliable not only in the actual world but also in nearby possible worlds, i.e. worlds that satisfy some minimum standard of probability *relative to* our current information. See Papineau 1993, pp. 150-152.
[18] Van Cleve, 1984, pp. 556, 559.
[19] Greco 2000a, pp. 172-173.
[20] Van Cleve 1984, p. 557.
[21] Van Cleve 1984, p. 559.
[22] Van Cleve 1984, 558-560; Papineau 1993, 157-158. Van Cleve distinguishes 'rule-circularity' from 'premise-circularity'. Only the latter, he argues, is viciously circular.

even deduction. But we tolerate deductive proofs of deduction, so why not extend the same courtesy to induction?

This approach to inductive skepticism is attractive, but inadequate. To begin, the claim that we need not believe in the reliability of our inferences to be justified in making them depends on too narrow a conception of belief. The claim is undoubtedly true if by 'believe that *p*' we mean 'consciously attend to and assent to *p*'. But Freud and his successors have taught us that beliefs need not be conscious to be real, i.e., to exert a determinative influence on our behavior. Belief informs our behavior. To believe something is, among other things, to be disposed to behave *as if* that something were true. If, when placed in the appropriate circumstances, a person fails to act in accordance with their professed beliefs, we take that as *prima facie* evidence against the genuineness of their profession. Moreover, we explain or render the behavior of persons coherent and intelligible by imputing beliefs to them. Where we cannot find a relatively plausible[23] set of beliefs and desires to account for a person's behavior, we judge their actions to be unintelligible or irrational. Now, when we try to account for our inductive behavior, it seems that we do, indeed must, believe that induction is a reliable mode of inference. Our predisposition to make such inferences is unintelligible except under the assumption that we believe them to be reliable. For if we did not regard the inference as reliable, then we wouldn't make the inference or act as if the conclusion we arrived as was true.

Given that we do need to believe in the reliability of our inductive inferences, it makes sense to ask whether we are justified in so believing. The reliabilist answers affirmatively. We are justified because the inductive inference rules we apply are in fact reliable, and we can show this by giving an inductive argument in favor of induction. This

[23] That is, plausible relative to information that we suppose the person to possess.

looks rather obviously question begging, however. The question at issue is the reliability of induction, so surely we can't appeal to induction to justify induction, at least not by using the very inference rule in question as part of its own justification. The idea that "self-authentication is no authentication"[24] is highly intuitive and cannot be lightly dismissed. But the reliabilist has a seemingly decisive rejoinder: In deductive logic, if an inference from A to B is valid, then the truth of B can be justifiably inferred from the truth of A without requiring *as a premise* a statement that the inference is valid. If this were not so, then, as Lewis Carroll famously argued,[25] this would land us in an infinite regress, for the amended inference would require yet another premise asserting its validity, and so on. To escape the regress we must conclude that deductive inference rules are not premises. And if this is the case with respect to deduction, why can't it also be true with respect to induction? If inductive inference rules are not premises, then even if we do need to be justified in believing in the reliability of the inference rule, we do not wind up in a vicious premise-circularity. There is an answer to this rejoinder, however. In the deductive case we can prove the validity of any given inference rule *without invoking the rule in question*.[26] Indeed, any rigorous mathematics text will *insist* that this be done. One is never entitled to use the theorem in question (either as a premise or as an inference rule) or premises that presuppose the truth of the theorem in question as part of the theorem's own proof. So we can turn the reliabilist rejoinder around: If you wish us to take the parallel with deduction seriously, then

[24] Howson 2000, p. 29.
[25] Carroll 1885, pp. 278-280.
[26] Howson 2000, pp. 27-29. In comments on a draft of this chapter, John Greco suggested that this just widens the circularity and that, therefore, allegedly non-rule-circular deductive proofs of deductive inference rules are in fact still rule-circular at a higher level. The suggestion is mistaken, however. Greco's result would follow only if *all* the presuppositions of deductive inferences (i.e., premises and inference-rules) *had to be justified by deduction*. If that were true, then either rule-circularity, premise-circularity, or an infinite regress would be unavoidable. But it isn't true. The ultimate premises and inference-rules of any deductive inference are always established nondeductively. Mathematical proofs, the paradigmatic examples of deductive inference, are grounded in axioms, which are established either by stipulation or by direct rational intuition.

give us an inductive argument for an inductive inference rule that does not use the rule in question. Could this be done by distinguishing different types of induction and using one type to justify another? Well, first, reliabilists need to show how this can be done. Second, we would still need to avoid higher-level rule-circularity (e.g., inference rule A is established by rules, B, C, and D; but rule B is in turn established by a set of rules including A), and that, I submit, can't be done without generating an infinite regress. To lay out this problem more fully, and to show why it isn't also a problem for deduction, we need to look more closely at the differences between deductive and inductive inference rules.

Excursus on Deductive and Inductive Rules

The job of an inference rule is to license the move from premises to conclusion. Thus it will have the form of an if–then: "if this premise is true, then this conclusion must be or is probably true." So given an inference with premise P and conclusion C, the inference rule that links them will be of the form 'If P then (probably) C'. That point of commonality aside, however, deductive and inductive inference rules necessarily have very different properties.

By its very nature, deduction can only validly arrive at conclusions that are already contained in the premises. Deductive inference rules are valid, therefore, only if they do not add anything to the conclusion that is not already in the premises. Since they add nothing to the conclusion, *they have no content*; i.e., they make no substantive claims about the world. Instead, they assert a relationship between *ideas*. Since they require that the conclusion follow with necessity from the premises, *they must be necessary truths*. Because they are necessary truths concerning only relations of ideas, their justification is straightforward—we can prove the validity of the deductive inference rule employed by any particular deductive inference by simply analyzing the premises until the conclusion is evident. The process of

analysis may involve the use of other deductive inference rules, but it will not continue *ad infinitum* or inevitably mire us in rule-circularity because it has a determinate stopping point. Once the connection between premises and conclusion has been made perspicuous to rational intuition we are done. In other words, the most basic deductive inference rules (e.g., *modus ponens*) don't need to be justified at all by appeal to other deductive inference rules; they are axiomatic. We can just *see* that they are true.

Unlike deduction, induction projects a conclusion that goes beyond what is given in the premises. Inductive inference rules, therefore, license the importation of *new content* into the conclusion. This means that they are *substantive* in a manner that deductive inference rules are not. Because they are substantive, inductive inference rules are also *contingent*—for any set of premises and any conclusion that goes beyond the premises, we can always specify some contingent situation in which the conditional 'if P then probably C' would be false.[27] Because they are substantive and contingent, inductive inference rules can't be justified *simply* by rational intuition as with deduction. We can analyze the premises forever and we will never be able to *see* that the truth or probable truth of the conclusion follows in any necessary way from the premises precisely because it doesn't follow in any necessary way from the premises. An important corollary of the preceding points is that inductive inference rules must not be construed as premises in the inference. On this point the reliabilists are correct. Because the inference rule is substantive, inclusion into the inductive argument *as a premise* would fundamentally change the inference by augmenting the evidence that we have to work with. If for some reason we *had* to include such rules as premises then if the resulting inference was still inductive, we would have to incorporate yet another inference rule as a premise, and another, and so on . . . until the inference was no

[27] For an in-depth argument on just this point, see Howson 2000, ch. 4.

longer inductive but deductive. Therefore, to avoid changing a given inductive inference into a fundamentally different inference, and to preserve the autonomy of induction as a mode of inference distinct from deduction, we must not construe inductive inference rules as premises. With these observations in mind, we need to consider the relation between inductive inference rules and the auxiliary assumptions that we talked so much about in the previous chapter.

First, we have seen that both inductive inference rules and auxiliary assumption are substantive and contingent. We can further show that, just like auxiliary assumptions, inductive inference rules need to be case-specific. Why? Consider the generic inductive inference rule 'If P then probably C'. The letters P and C are simply placeholders for a specific set of premises and a specific conclusion. The inference rule cannot license the inference from *these* premises to *this* conclusion until those placeholders have been specified.[28] This function of licensing the inference from premises to conclusion is yet another property that inductive inference rules and auxiliary assumptions have in common. The auxiliary assumption in an inductive inference indicates precisely how our evidence is relevant to the conclusion and thereby licenses the inferential move from one to the other with a degree of confidence proportional to our confidence in the evidence. The problem of

[28] We cannot avoid this conclusion by construing inductive inference rules as expressing a relation between premise and conclusion *types*. Suppose we represent the inference rule in this manner: 'given evidence *of a certain type*, infer the probable truth of a conclusion *of a certain type*' (e.g., given a large, random sample from a population in which X% has property P, infer that, probably, X% of the population has property P). The problems with this type of generic inference rule are twofold. First, we can't *apply it* in practice until we know that the particular evidence in question is of the sort required by the inference rule, and that determination requires knowing a number of *specific* things about the evidence in question. For example, how do we know that *this* sample is random? (Whether any given sampling-methodology will in fact yield a random sample is a contingent matter.) And how do we know that *this* sample is large enough? (That requires knowledge about the probable size of the population in question.) Second, suppose we've satisfactorily determined that our evidence satisfies the conditions of the inference rule. We are now licensed to infer a certain type of conclusion. But which particular conclusion of that type should we infer? This problem raised by this question is exacerbated by the fact that in many (and maybe all) such cases there will be an infinite number of mutually incompatible conclusions of that type, as Goodman's grue paradox (described below) illustrates.

justifying inductive inference rules and the problem of justifying inductive auxiliary assumptions amount, therefore, to the *same* thing, namely, the problem of justifying inductive inference. A solution to either would tell us that we are licensed to inductively infer a given conclusion from a given set of evidence. This implies that inductive inference rules and auxiliary assumptions are intimately related in some way. It seems to me that the right way to express this relationship is to say that the auxiliary assumption grounds the truth of the inference rule by indicating precisely how and to what extent the truth of the premises makes probable the truth of the conclusion. In other words, the inference rule (a specific instance of 'if P then probably C'), as a substantive and contingent principle, presupposes that P and C are really connected in some way and that this connection is not merely coincidental or fortuitous but regular and lawlike. The close relationship between inductive inference rules and auxiliary assumptions and the fact that the inference rules are not premises, suggests strongly that we should not construe auxiliary assumptions as premises either. In fact, this must be so if we want to avoid inductive skepticism. Suppose we say as in premise (1) of (SA1) that *all* inductive inferences presuppose an AA *as a premise*. If so, then if the inference that results after adding the AA as a premise is still inductive, we will have to presuppose yet another AA (AA') as a premise, and so on . . . until the inference is no longer inductive but deductive. Premise (1) of (SA1) thus implicitly assumes the deductivism (the view that only deductive inferences are good inferences) that so many of Hume's critics charge him with.[29] We have identified, therefore, a principled reason for rejecting premise (1) of (SA1) as false—inductive inferences do not require a contingent AA *as a premise*.

[29] It must be emphasized, however, that while cashing out Hume's skeptical argument as we did in (SA1) does lead to deductivism, that is *not* how Hume himself argues. The problem he raises is not that AA's are needed to turn inductions into deductions but that we have no way of licensing the *inductive* move from evidence to conclusion. The revised statement of the skeptical argument (SA2) below is, therefore, a more accurate reconstruction of Hume's reasoning and one which is not open to the charge of assuming deductivism.

Does this mean that we have found a fatal flaw in the skeptical argument? Unfortunately, no, for (SA1) can easily be rewritten to avoid the problem:

(SA2) 1. All inductive inferences require a contingent inference rule (IR).
2. A belief is inferentially justified only if belief in the inference rule is justified.
3. Therefore, a belief is inductively justified only if belief in the IR is justified. (1,2)
4. There is no way of justifying belief in the IR.
   a. All justification of belief is either inferential or noninferential.
   b. Belief in the IR cannot be noninferentially justified (it is not self-evident, evident to the senses, incorrigible, or otherwise properly basic).
   c. Inferential justification is either deductive or inductive.
   d. Belief in the IR cannot be justified by deduction.
   e. Belief in the IR cannot be justified by induction.
   f. Therefore, belief in the IR cannot be inferentially justified. (c,d,e)
   g. Therefore, belief in the IR cannot be justified. (a,b,f)
5. Therefore, beliefs arrived at via inductive inference are unjustified. (3,4)

Because (SA2) works in substantially the same way as (SA1), a brief examination of the premises will suffice. Premise (1) states that inductive inferences require a contingent IR. That they require an inference rule seems hard to deny. That the rule must be *contingent* follows from the nature of induction. An IR has the form "If P then probably C." For any P and any C not entailed by P, there exists a logical gap between P and C which we denote by X, the inductive extension of C beyond P. X is the set of all propositions entailed by C that are not entailed by P. Whether C is probable given P necessarily depends on the relation between P and X. If there were *no relation* between P and X then P would simply be *irrelevant* to the affirmation of X in C and would not, therefore, be adequate support for C. But since *ex hypothesi* P says nothing about X, it can't tell us whether or not X is probable given P. As far as P goes, X may or may not obtain. That is, the relation between P and X is not necessary, but contingent. But then the relation between P and C must be contingent as well. If we reason inductively from P to C, therefore, it must be via a contingent principle

that establishes the relevance of P to X. As we saw in the previous chapter, that's what the AA does. In so doing, it grounds the truth of the IR.

Premise (2) of (SA2) is a mere platitude. Since an inference rule is what licenses the move from premises to conclusion, to infer without a justified inference rule is just to infer without a proper license to infer, which is to say that the inference is not justified. As for premise (4), because IR's are substantive and contingent just like AA's, the arguments for premises (4.b) and (4.d) run exactly like they did in the previous chapter. The argument for premise (4.e) needs some revision, however, since it depends on premise (1). The revised argument looks goes like this:

(4.e.ii) 1. All inductive inferences require a contingent inductive inference rule (IR).
2. Therefore, any inductive justification of the IR must also presuppose an IR (IR').
3. IR' must either be the same as IR or different.
4. If IR' is the same as IR, then the argument is viciously circular; IR is being assumed in the course of proving IR.
5. If IR' is different from IR, then IR' itself stands in need of justification by a prior inductive inference that must also presuppose an IR (IR"). (1)
6. And so on. There results an infinite regress, unless the same IR recurs, at which point the argument becomes viciously circular.
7. Therefore, the IR cannot be justified by induction.

The main difference between (4.e.i) and (4.e.ii) is the kind of circularity in view. The earlier argument was concerned with premise-circularity, which everyone agrees is viciously circular. This one is concerned with rule-circularity, which reliabilists like Van Cleve and Papineau want to say is not viciously circular. Hence, they would reject premises (4) and (6) of (4.e.ii). With that, let us return to our discussion of the reliabilist response to the problem of induction.

The argument for saying that rule-circular inferences are not viciously circular is a *reductio*—if they were, we wouldn't be able to justify either induction *or deduction*. The

conclusion doesn't follow, however, because we can give a non-rule-circular deductive proof for any given deductive rule. Moreover, we can justify any deductive inference (and its associated inference rule) directly by analyzing the premises until their necessary relation to the conclusion is evident to rational intuition. This is possible because deductive inferences rules are not substantive. Inductive inference rules, in contrast, are substantive and contingent. We cannot, therefore, establish their truth by rational intuition. No matter how long we analyze the premises, the inductive conclusion will never be evident because it isn't *in* the premises. Because inductive inference rules are substantive, it is impossible to show that our inductive projections are reliable without invoking substantive assumptions.

> Let P be the conjunction of all factual statements known to be true. Suppose that the inference from P to a statement Q describing some event not known to be true is not deductive. . . . It follows immediately from the definition of deductive validity[30] that in some subset W of all the possible worlds . . . P is true and Q is false. . . . [W]hat further information could be appealed to which would make it more likely that our world is not in W? . . . [T]he only information we have about this world that is known to be true is already in P. . . . All we know is that in our world Q may be true or it may be false. . . . Hence any principle claiming to justify the inference from P to the truth *or even the probable truth* of Q must beg the question.[31]

The problem of justifying inductive inference rules is thus a real one, and one that is both fundamentally different from and more difficult than the problem of justifying deductive inference rules. The attempt by Van Cleve and Papineau to legitimize rule-circular inferences by appealing to the parallel between deductive and inductive inference rules is misguided. The necessary parallel doesn't exist, and without that we are left with no good reason to suspend the highly intuitive principle that self-authentication is no authentication in favor of permitting rule-circularity.

---

[30] A inference is deductively valid if there is no possible world in which the premises are true and the conclusion false.
[31] Howson 2000, pp. 11-12.

Perhaps we can, with Van Cleve[32] and Papineau,[33] take comfort in the thought that induction has not been shown to be objectively 'problematic' and that, therefore, given the *de facto* reliability of induction, we ought to be entitled to use it in the absence of any positive justification for it. Unfortunately, this is not so. Enumerative induction (which is what both Van Cleve and Papineau have principally in mind) is not only problematic, but *inconsistent*.[34] The argument for this is due to Goodman[35] and is justly famous: Define the predicate 'grue' to mean 'has been observed up to now and is green or has not been observed up to now and is blue.' Accordingly, any emerald that has been observed up to now and is green may with equal warrant be described as grue. Enumerative induction tells us that from a large enough sample we can justifiably extrapolate from the characteristics of the sample to the whole reference class (or to the next instance). But if we have observed a large number of green emeralds, enumerative induction licenses us to infer both that the next emerald is green and that the next emerald is grue. In other words, the next emerald to be observed is now both green and blue, which is to say that extrapolating from the past is not just unreliable, but maximally so—it yields incompatible projections.[36] In like manner, we can use the inference rule to *prove its own unreliability*:

> Call an inductive inference made by the rule 'right' if it concludes to a true assertion, and 'wrong' if it concludes to a false one. Call an inductive inference 'ring' if it has been checked and found to be right, or not been checked and is wrong. Suppose the majority of checked inductive inferences have been found to be right. It follows that they are also ring. Using as the justified premise 'The majority of checked inductive inferences are ring' in van Cleve's argument A, we conclude from A that the majority of all inductive inferences are ring. But only a finite number have been checked,

[32] Van Cleve 1984, p. 557.
[33] Papineau 1993, p. 160.
[34] Howson 2000, p. 29.
[35] Goodman 1983, pp. 72-81.
[36] Howson 2000, p. 30.

> leaving a potential infinity unchecked. It follows that the majority of all inductive inferences are wrong, and hence that the rule is unreliable![37]

The only way to avoid these consequences is to place external constraints on enumerative induction so that only certain types of projections are allowed. So even if we could somehow justify the inference rule, we couldn't consistently *apply* it without invoking additional substantive assumptions. And, of course, this raises the problem of how we are to justify those assumptions.

Summing up, we have seen that the reliabilist attempt to justify induction by applying an externalist standard of justification is fraught with difficulties. The claim was that (a) we *are* justified in reasoning inductively if induction is *de facto* reliable; (b) since induction just has to *be* reliable, we don't need to *believe* that it is, and therefore we don't need to *justify* the belief that it is; and (c) we can justify the belief that induction is reliable via a rule-circular inductive procedure. I have argued against (b) that we do need to believe that induction in reliable because if we didn't our acts of inferring would be unintelligible. Consequently, we do need to justify that belief. I have argued against (c) that the proposed method for justifying belief in the reliability of induction fails—one cannot justify a rule by invoking the rule in question. Against (a) I have argued that induction is in fact *unreliable*, unless we help ourselves to additional substantive assumptions, in which case we need to justify those. In chapter three I will add another objection to this list, against the adequacy of *purely* externalist accounts of inference, and by extension, inferential justification. On a positive note, however, the reliabilist defense of rule-circularity afforded us a useful opportunity to explore the nature of inference rules and to make a significant refinement to the skeptical argument through the observation that inference rules are not premises.

---

[37] Howson 2000, pp. 30-31.

Furthermore, as I argue in the following chapter, epistemic justification is *partly* externalist, and this is significant for responding to Hume.

III. Inductive Inference Can and Should Be Justified

The third and final approach is to accept the skeptical demand for a justification of induction and try to satisfy it by supplying a justification—an implicit rejection of premise (4) of the skeptical argument (either version). Several methods have been proposed for accomplishing this: (a) *categorical apriorism* (early Peirce, Stove, D. C. Williams) is the view that inductive inference *must work* based on strictly theoretical grounds; (b) *conditional apriorism* (late Peirce, Reichenbach) maintains that *if* anything can work, then induction can; (c) *probabilism* (middle Peirce, Carnap, Bayesianism) holds that the mathematical theory of probability can provide us with unambiguous assessments of the likelihoods of hypotheses in relation to a given body of evidence; (d) *inductivism* (Braithwaite, Black) says that induction can be justified inductively; and (e) *uniformitarianism* (Mill, Russell, Keynes) is the view that induction can be justified by postulating some broad metaphysical principle like the Uniformity of Nature.

These last two approaches were discussed in the previous chapter, so a brief recap here will suffice. (d) fails because inductive arguments for induction either generate an infinite regress or are viciously circular. As I argued above, whether that circularity concerns premises (4.e.i) or rules (4.e.ii), it is still vicious. (e) is problematic because (i) it is very difficult if not impossible to articulate a Uniformity Principle that is both plausible and useful, (ii) there remains the problem of justifying the Principle itself (and how could one hope to do that without using induction?), and (iii) inductive inferences require *specific* AA's

(and IR's), so even if one manages to articulate and justify a useful general Principle, the problem of induction remains.

III.a Categorical Apriorism

Like reliabilism, this approach amounts to a rejection of premise (1) of (SA1), but whereas reliabilism denies that we require contingent AA's *as premises*, categorical apriorism denies that induction requires *contingent* AA's (or IR's)—induction, it is argued, *must* work in any possible world. There are various ways in which this position has been developed. Peirce argued that every possible world must have some internal order that application of induction methods would eventually enable us to discover:

> If men were not to be able to learn from induction it might be because as a general rule, when they had made an induction, the order of things . . . would then undergo a revolution. . . . But this general rule would be capable of being itself discovered by induction; and so it must be a law of such a universe, that when this was discovered it would cease to operate. But this second law would itself be capable of discovery. And so in such a universe there would be nothing which would not sooner or later be known; and it would have an order capable of discovery by a sufficiently long course of reasoning. But this is contrary to the hypothesis, and therefore that hypothesis is absurd.[38]

But Peirce underestimates how disordered things can be. Rescher asks us to imagine a *nomic* chaos, in which the laws of nature change randomly with randomly varying half-lives. Couple that idea with the notion of a Cartesian demon working to frustrate our inductive efforts and the mathematical fact that any given specification of order in a finite series can be defeated as the series is extended, and it certainly seems *conceivable* that there could be a world in which our best inductive efforts would fail.[39]

---

[38] Peirce, CP 5.352.
[39] Rescher 1980, p. 191.

According to Stove, that all of the many observed ravens have been black is, by itself, *necessarily supportive* of the claim that all ravens are black:

> The proposition,
>
> [A] That [B] all the many observed ravens have been black is a reason to believe [C] that all ravens are black,
>
> . . . mentions the contingent propositions [B] and [C]. But . . . it does not *assert* either of those propositions. Its truth, therefore, does not depend on what their truth-values happen to be. Suppose that those truth values are the 'worst' possible: that is, suppose the conjunction of [B] with the negation of [C]. Even this contingent proposition, it should be evident, is not inconsistent with [A]. . . . Every other contingent proposition, *a fortiori*, is consistent with [A]. A proposition is contingent, however, only if there is some contingent proposition which is inconsistent with it. So [A] is not contingent.[40]

He concludes that "*p* is a reason to believe *q*" has no contingent instances, where *p* is a proposition about the observed, and *q* is a proposition about the unobserved.[41]

Stove's argument for the necessary truth of [A] moves too quickly, however. It shows at most that the logical relation between [B] and [C], *whatever it is*, is necessary. It does not show that [B] *supports* [C], for one can easily construct a parallel argument in which the support relation between the component propositions does not obtain. For example,

> [D] That [B] all the many observed ravens are black is a reason to believe [E] that all dogs hate cats.

If Stove's argument for the necessary truth of [A] is right, it works equally well *mutatis mutandis* to establish the necessary truth of [D]. But that is absurd. [B] clearly is not a reason to believe [E].

Williams appeals to the strong law of large numbers, which states that "the average of a sequence of independent random variables having a common distribution will, with

---

[40] Stove 1986, pp. 4-5.
[41] Stove 1986, pp. 5-6.

probability 1, converge to the mean of that distribution."[42] In other words, continued random sampling from a population with a fixed distribution will eventually yield a distribution of results that is arbitrarily close to the actual distribution of the population. Of course, results might fluctuate wildly for a long time before settling down to a reasonably accurate distribution, so how can we ever justifiably make projections from the current sample distribution if we don't know how close it is to the actual distribution? Williams has an answer: Consider the 'hyperpopulation' of all *n*-membered subsets of the population. For any specified population size *n* we can demonstrate statistically, with a calculable margin of error, that a certain percentage of those *n*-membered subsets are (approximately) representative of the population. When that percentage is sufficiently high for plausible values of *n*, then we are justified in our projections.[43] This result, he insists, holds independently of any assumption about whether the *actual* sample is a representative one.

Unfortunately, Williams is mistaken. In order to apply the law of large numbers to concrete cases we have to make some *substantive* assumptions. The law only applies *if* sampling is random and the population has a fixed distribution. That sampling must be random is tantamount to a reliability assumption. It means that our sampling methods are not systematically biased, tending to select some members of the population more frequently than others. That the population has a fixed distribution is tantamount to a regularity assumption. If the actual distribution of the population occasionally underwent abrupt changes (à la Rescher's nomic chaos), then the sample data that we had accumulated would suddenly cease to provide any solid basis for projections, and would become irrelevant and even positively misleading. Furthermore, even if we help ourselves to these assumptions, we

[42] Ross 1998, p. 407.
[43] Williams 1963, ch. 4.

can't assess the propriety of any projections based on the data without making some assumption about the size of the population. In sum, therefore, the law of large numbers is insufficient to guarantee that induction must work.

The attempts of Peirce, Stove, and Williams to justify induction without relying on contingent assumptions fail. Indeed this failure could have been predicted from the outset. Inductive inference aims to extend our knowledge base by projecting from the data in hand to a conclusion that goes beyond the data. Since *ex hypothesi* our current knowledge base is *all we know*, any attempt to prove that induction reliably leads us from true premises to true or probably true conclusions is an attempt to establish a relationship between the known and the unknown. Inasmuch as the conclusion goes beyond the data, any proposed relationship is necessarily synthetic, and is therefore *contingently* true, if true at all.[44]

III.b Conditional Apriorism

Another approach to justifying inductive inference is *conditional apriorism*, which attempts to show that *if* any method of making projections from data can work, then (enumerative) induction must work. According to Reichenbach, if there is a definite probability that an *X* will have a given trait (a regularity assumption), then as sampling of *X*'s continues indefinitely, the relative frequency of *X*'s with that trait must eventually converge on the true probability (a reliability assumption), for probability is simply a long-range

[44] Kant would object that there are synthetic statements that are necessarily true. Unfortunately, Howson points out, we can conceive of alternatives to principles that Kant held to be necessary conditions of human thought (e.g., deterministic causality and Euclidean geometry as the only possible geometry of space), some of which are in fact now regarded as false. Kant's 'transcendental deduction' thus appears to be unsound. Furthermore, any sound deduction would have to employ some non-tautological premises, raising the issue of the justification for those premises, at which point Hume enters the scene again (Howson 2000, p. 18). For further criticism of the notion of the synthetic *a priori* as it pertains to the problem of induction, see Salmon 1966, pp. 27-40.

relative frequency.[45] In other words, the claim is that insofar as the world is characterized by stable regularities, then to that extent inductive projections based on continued sampling are reliable.

While this approach is an improvement over categorical apriorism in that it appreciates the fundamental contingency of inductive reliability, it fails to appreciate the full depth of that contingency. Following Rescher, let Q* be the relative frequency of some trait in the whole *observable* population and let Qt be the relative frequency of that trait in the actually *observed* subpopulation at time *t*. The claim is that $Q^* = \lim_{t\to\infty} Qt$, assuming that the limit exists. This claim is problematic for several problems. First,

> the consideration that statistical induction is "ultimately bound to succeed" in the long run does not underwrite any conclusion whatsoever about the appropriateness of its use here and now. . . . After all, . . . as Keynes remarked, in the long run we are all dead. And, unhappily, the long-run behavior of such a limit function is compatible with *any* short-term behavior whatsoever.[46]

Thus, the rationality of our *present* reliance on the results of inductive reasoning remains in question. Second, at no time *t* can we guarantee that there is any fixed relationship between Q* and Qt. Thus, even if we are indeed converging on the actual limit, we can never be certain that we are. Third, observation biases may prevent the limiting value of the relative frequency in the *observed* population, Qt, from ever matching the limiting value in the *observable* population, Q*. In other words, for induction to work not only must the world in fact be characterized by regularities, but also our data acquisition methods must be in fact reliable (i.e., unbiased). Fourth, even if we do obtain Q*, that may not be the same as the *actual frequency in nature*, Q, which may not be completely observable.[47] Fifth, as Salmon

[45] Reichenbach 1938, p. 350, 355.
[46] Rescher 1980, p. 101.
[47] Rescher 1980, pp. 101-104.

points out, there is an infinite number of limit functions that will converge to the same limit in the long-run but may differ to any extent in their short-term predictions. If Reichenbach's argument works at all, it works equally well for all of these limit functions.[48] What, then, justifies Reichenbach in choosing the limit function that he does? Appeal to 'descriptive simplicity'[49] is appropriate to decide between theories, statements, or rules that are empirically equivalent, but these competing limit functions are not.[50]

III.c Probabilism

A third justificatory approach is *probabilism* (Laplace, middle Peirce, Carnap, Bayesianism), in which appeal is made to the mathematical theory of probability to underwrite the validity of inductive inference. Peirce argued that certain forms of inductive inference necessarily lead from true premises to true conclusions *most* of the time. Arguments of the form "Most A's are B's, x is an A; therefore, x is a B" seem to fit the bill, but appearances are deceiving, as Peirce later realized. The conclusion (x is a B) only follows probabilistically if x is an *appropriately random* instance of A.[51] If x is an A identified by some biased selection procedure, it remains a distinct possibility that most A's may be B's and yet x fail to be a B most of the time. The only way to secure against the possibility of observational bias and ensure the randomness of the sample is to invoke some metaphysical principle like the Uniformity of Nature. And since such a principle cannot be justified *a priori*, neither can one justify *a priori* the appropriateness of probabilistic or statistical types of inference along these lines. Moreover, how is the premise "most A's are B's" to be established without recourse to the very inductive procedures in question? Rescher writes:

---

[48] Salmon 1966, p. 88.
[49] Reichenbach 1938, sec. 42.
[50] Salmon 1966, p. 89.
[51] Peirce, CP 2.726.

> Authentic probability values are never *given*, but always inferred, and inferred by processes in whose operation induction must play a part. . . . We cannot use probabilities to justify induction, because we need induction (or else some randomness-assuring metaphysical assumptions) to obtain probabilities.[52]

Another probabilist strategy is to invoke Bayes' Theorem to compute the probabilities of hypotheses (H) relative to the evidence (E). In simplified form, the theorem says that

$$\text{Prob}(\text{H} \mid \text{E}) = \frac{\text{Prob}(\text{H}) \times \text{Prob}(\text{E} \mid \text{H})}{\text{Prob}(\text{E})}.$$

Accordingly, we can compute the probability of H given E if we know the probabilities of H, E, and E given H. The formula clearly assumes that Prob(E) and Prob(H) ≠ 0. Indeed, if Prob(H) equaled zero then the right-hand side of the equation would equal zero *no matter what the values of Prob(E) or Prob(E | H) are*, which is to say that unless the prior probability of H is nonzero then the posterior probability of H is *completely insensitive to the evidence E*. Where do these nonzero prior probabilities come from? One might plausibly claim to know that Prob(E) is nonzero on the basis of observation, but the same cannot be said for Prob(H). As Goodman's new riddle easily shows, for any set of observations there is an *infinite* set of hypotheses that is consistent with the evidence. But, says the Bayesian, we can assign to each of these hypotheses a nonzero prior probability (e.g., $\text{Prob}(\text{H}_1) = 1/2$, $\text{Prob}(\text{H}_2) = 1/4$, $\text{Prob}(\text{H}_3) = 1/8$, $\text{Prob}(\text{H}_4) = 1/16$, and so on), the specific distribution of probability assignments doesn't matter, and we can then use Bayes' Theorem to show that, necessarily, if Prob(E | H) / Prob(E) > 1, then Prob(H | E) > Prob(H). In other words, if E is more likely to be true given H than given the complement of H, then E must increase the probability of H by the factor Prob(E | H) / Prob(E). So far it appears that we may have the

[52] Rescher 1980, p. 202.

makings of a genuine solution to the problem of induction—we can show, it would seem, that suitable evidence must raise the probability of the 'correct' inductive conclusion and diminish that of its competitors. But, alas, matters are not so fortunate, as Howson explains:

> It is not very difficult to see that the additivity principle implies that there [are] at most n elements of any possibility-space with probability at least 1/n, from which it follows that *at most* denumerably many members can have positive prior probability. Since the difference between an uncountable and a countable set can itself be shown to be uncountable, this implies that *uncountably many possibilities in an uncountable possibility space must have probability 0*.[53]

The problem is that the set of competing hypotheses is in general nondenumerable,[54] so we cannot assign a nonzero prior probability to each hypothesis on the grounds that to do so would be to dogmatically render them insensitive to the evidence. Indeed, almost all of them (nondenumerably many) *must* be assigned probability zero. That the 'correct' hypothesis H can be confirmed at all is therefore based on an apparently arbitrary decision to assign it a nonzero prior probability in the first place! Our initial probability assignments are thus tantamount to substantive assumptions, and the problem of how to justify those assumptions remains.

IV. The Way Forward

C. D. Broad once referred to induction as "the glory of Science, [and] . . . the scandal of Philosophy."[55] On the one hand, we make very successful use of induction both in science and in daily life; on the other hand, Hume's problem has proven itself robust enough to resist more than a couple centuries of attempts at resolution. Neither the claim that we don't *need* induction (Popper), nor the claim that we don't need to *justify* induction (analytic rationalism

[53] Howson 2000, p. 75.
[54] Howson 2000, p. 76.
[55] Broad 1952, p. 143.

and reliabilism), nor the major attempts to justify induction have proven adequate as responses to Hume.[56] So what are we to do? In his classic work on problem solving, mathematician G. Polya remarks that when the direct approach to solving a problem doesn't work, one should try to come at it from a different angle. Sometimes taking a step back to first answer a more general problem can point the way forward to a solution of the problem at hand.[57] And sometimes it helps to reconsider our formulation of the problem. Let us consider this tactic.

The problem facing us is the justification of inductive inference. Since we find ourselves in a quandary regarding how to answer the problem, we should consider whether our conception of the problem might tacitly conceal some erroneous assumptions. Accordingly, we should examine the terms of the problem. First, what exactly do we mean by a *justification* of induction? We have seen above different standards of justification employed in an attempt to answer the problem (e.g., that induction must work, that induction will work if anything can work, that induction is reliable, etc.). I have argued above that attempts to show that induction satisfies these standards ultimately require us to assume substantive claims that it seems we could justify only by induction. So either we have to find another defensible way to justify these substantive assumptions, or we need to find some other appropriate standard of justification to apply to induction. In chapter three I explore the topic of epistemic justification and argue that Hume's argument implicitly invokes a standard of *noninferential* justification that we need not accept. This will constitute a *negative* response to the skeptical argument, by which I mean a response that rebuts the skeptical

[56] This is not to say that none of these proposals have useful insights to contribute—indeed, I will eventually argue that both reliabilism and Bayesianism supply important components of a *positive* solution to the problem of induction even though they don't suffice for a negative solution.
[57] Polya 1957, pp. xvi-xvii, 108-109.

claim to have shown that induction is not justified. In subsequent chapters I go further and develop a *positive* response, i.e., one that shows that our reliance of induction *is* justified.

Second, what exactly is this thing called 'inductive inference' that we are trying to justify? Answering this question is essential for developing a positive response to inductive skepticism. We have defined induction over against deduction, but we have not yet considered what it means for something to be an *inference*. Nor have we considered whether or not there are distinctive types of inductive inference, each with its own unique justificatory issues. I take up the first of these issues in chapter three and the second in chapter four. One thing seems clear: inferring is something we *do*, which prompts a third question: *Why* do we infer? What is the purpose or function of inference? If we can answer that, then we can construe the *positive* task of justifying inference in general and induction in particular as one of assessing whether it adequately serves its intended function. Is induction an effective means to achieve the ends for which we employ it? As the study of the adequacy of means to ends is called methodology, we arrive at the idea of a *methodological justification of induction*. Whether this strategy can actually be carried out and avoid the pitfalls of other approaches remains to be seen, but it is the strategy that I pursue in the following chapters.

CHAPTER THREE

INFERENCE IN PERSPECTIVE

Questions were raised at the end of chapter two concerning the nature and purpose of inference and the nature of inferential justification, i.e., the epistemic justification of beliefs arrived at via inference. These questions are closely related. A definition of what inference is tells us what we need to justify and constrains the types of explanations we can give as to why we infer. An account of the goal or purpose of inference is partially constitutive of our understanding of what inference is and may provide a basis for a partly methodological understanding of inferential justification. And, finally, the issue of inferential justification constrains what we say about the nature of inference. For unless we want to fall prey to radical skepticism, we should make sure that our definition of inference is one that can satisfy a plausible standard of justification. My goal in this chapter is to answer the definitional question in abstraction from the question of purpose, which I take up in chapter four. In the first section I unpack a standard, though not entirely uncontroversial, definition of inference. In section two I contrast this 'internalist' definition with a rival 'externalist' definition and argue that the internalist definition is superior for dealing with issues of inferential justification. Along the way, I develop a general account of epistemic justification. In the third section I employ the distinction between internalist and externalist *definitions* of inference to defuse an objection that inferential internalism (i.e., internalism with respect to inferential *justification*) leads to unavoidable skeptical consequences. Finally, in section four, I develop complementary accounts of inferential and noninferential justification and show how they can be used to rebut Hume's inductive skepticism.

I. Inference: A Preliminary Analysis

Let's start with the following standard definition of 'inference'[1]:

(D1) An inference is an act of thought in which one passes from a set of propositions (the premises) to a proposition (the conclusion) because it appears to one that the latter must be or is probably true if the former is.

This definition is not uncontroversial, as I point out in the following section, but for present expository purposes it makes a good place to start. My goal in this section is to unpack (D1) and to clarify some ambiguities. Several aspects of (D1) bear emphasizing. First, inference is an *act*; it is something we do. This should not be construed as implying that we have direct voluntary control over our inferences—we can't just infer in the same sense that we can, say, cast a vote. Rather, inference is an act in the same way that sneezing is an act—we do it, but not (as least not ordinarily) as the result of an intention to do so. To avoid confusion on this point, I will speak of inference not as an act but as an *event*. Unlike sneezing, however—and this is my second point—inference is a *cognitive* event. Third, inference is *discursive*; we *pass* from one mental state to another mental state. Fourth, the mental states in question have *propositional content*; that is, the content of the mental states, what we are thinking about, is something that can be either true or false. For example, I can think about a concept, say 'blue', and pass from that to thinking that 'the sky is blue'. This is indeed a discursive act of thought, but it is not inferential because the concept 'blue' is not propositional; it makes no sense to ascribe to it a truth-value, i.e., to say 'it is true that blue'. Fifth, in inference there is a *perceived connection between the truth-values* of the premises and conclusion. In reverie, by way of contrast,

[1] Cf. Robert S. Tragresser, "Inference," in Dancy and Sosa 1992, p. 206.

one may pass from the proposition ‘the sky is blue’ to ‘violets are blue’ and from there to ‘roses are red’ as one thought triggers another, without ever thinking that the truth of one proposition has anything to do with the truth of another. Sixth, inferential transitions from one mental state to another take place *because of* the perceived connection between their truth-values. I take this to mean two things: (i) a passage of thought is inferential *only to the extent* that it is driven by this perceived connection, and (ii) the resultant confidence in the conclusion is directly proportional to the strength of the perceived connection.[2] Seventh, this perceived connection between truth-values must meet a minimum qualitative standard: The conclusion (C) should be seen to be *probable* in relation to the premises (P), i.e., Prob(C | P) > 0.5.[3]

Having unpacked definition (D1) somewhat, we now need to refine it. To begin, whenever we “pass” from a set of propositions to a proposition, we always regard these propositions in some manner, i.e., we consider them with a certain ‘propositional attitude’. Take the proposition *p*. We can believe that *p*, doubt that *p*, desire that *p*, etc.

---

[2] (i) implies that a mental transition may be *partially inferential*. For example, suppose someone strongly desires that *q* be true and thus is disposed to form the belief that *q* (this is called “wishful thinking”, a common phenomenon). Suppose that this person also has strong evidence *p* for believing in *q* and is aware of that fact. In this case two different causal factors are working together to form the belief that *q*. According to (i), the passage of thought from *p* to *q* is inferential only insofar as the evidence is driving it and not the desire. (ii) captures the intuition that weaker perceived evidential support from the premises should lead to lower confidence in our conclusions and stronger perceived evidential support to greater confidence. If a person’s confidence in their conclusions was unresponsive to (or worse, in opposition to) the perceived strength of the evidence, that would strongly suggest that their confidence was rooted in something *other* than their appreciation of the evidential support from premises to conclusion. To that extent, by (i), their mental passage from premises to conclusion would not be inferential.

[3] Suppose that Prob(C | P) < 0.5. In that case, C is perceived to be *improbable* in relation to P, which is the same as ~C being perceived to be probable in relation to P. But then any passage of thought to C rather than ~C would not be *because of* the perceived connection but *in spite of it*, thereby violating the sixth criterion for inference. Suppose that Prob(C | P) = 0.5. In that case, C and ~C are perceived to be equally probable given P, which is to say that the perceived connection between P and C doesn’t license either conclusion. But then, again, any passage of thought to one or the other would not be *because of* that perceived connection and thus would not be inferential.

We have to take some kind of mental posture toward the proposition. Now, it is sometimes proposed that inference must pass from belief to belief.[4] After all, if we don't believe our premises, then why would we trust them as a basis for inferring anything? And what would it mean to infer to a conclusion but not believe it? There are a couple important complications, however. First, we can, and often do, reason in a manner that is strictly hypothetical. That is, we can *suppose* that *p* is true for the sake of argument (even if we don't believe *p*), note that *p* entails *q*, and infer the conditional proposition *if p then q*. Thus, inference may start from supposition rather than belief. Second, it is questionable whether the mode of acceptance of the conclusion must always amount to full belief. Suppose my car won't start and I am trying to figure out what is wrong. Even with my limited mechanical knowledge I can come up with a short list of plausible explanations: dead battery, bad starter, faulty ignition switch, etc. Without further exploration I would not be justified in concluding that, say, the ignition switch is at fault. That is, I am not entitled to *believe* it in the sense of accepting it *as true*. But I am entitled to *pursue* it as a plausible working hypothesis.[5] One way of regarding explanatory inference (to be discussed in chapters four and five) is to see it as the adoption of a hypothesis as worthy of pursuit based on its being a plausible explanation for some phenomenon. So it seems that inference can result in modes of acceptance weaker than belief. Matters are not clear cut, however, because we have two ways of interpreting this result. On the one hand, we can view it as adopting a *weaker propositional attitude*

---

[4] E.g., Greco 2000a, p. 95. Cf. also Peirce: "When it happens that a new belief comes to one as consciously generated from a previous belief, . . . I call the event an *inference* or *reasoning*" (EP 2.463).

[5] Larry Laudan distinguishes between acceptance and pursuit as two "modalities of appraisal" (Laudan 1977, pp. 108-114).

toward the conclusion. Thus, instead of regarding 'the ignition switch is faulty' as *true* I might regard it as *possibly true*. On the other hand, we can view it as adopting a *weaker proposition* as the conclusion. Thus, instead of concluding that 'the ignition switch *is* faulty' I might conclude that 'the ignition switch *might be* faulty'. The first approach allows inference to result in modes of acceptance weaker than belief, and also allows us to speak of stronger and weaker beliefs. The second approach allows us to say that inference always results in full belief and fixes the propositional attitude so that we can no longer talk of stronger and weaker beliefs but only of stronger and weaker propositions believed. As far as I can see, which approach we take is a matter of convenience, as the two modes of discourse are intertranslatable.[6] I will adopt the second approach because it simplifies matters, allowing us to view inference as resulting only in full belief and not in gradations thereof. Specifying (D1) in terms of propositional attitudes and replacing "act of thought" with "cognitive event" we get:

> (D2) An inference is a cognitive event in which one passes from a set of believed or supposed propositions (the premises) to belief in a proposition (the conclusion) because it appears to one that the latter must be or is likely to be true if the former is.

(D2) introduces the concept of belief into our definition. A detailed analysis of the concept of belief is not necessary for our purpose, but it may prove helpful to make a few miscellaneous observations. First, belief is a cognitive *state*, not a cognitive event like inference. Thus inference effects a transition between cognitive states. Second, as

---

[6] Let $C(p, \phi)$ be a confidence function that ascribes to proposition p a probability ϕ, such that $\phi = \text{Prob}(p)$. Further, let q be the proposition that p has probability φ, i.e., $q \equiv [\text{Prob}(p) = \varphi]$. Thus, $\varphi = \phi$. It follows that $q \equiv [\text{Prob}(p) = \varphi] \equiv (\phi = \varphi) \equiv (\phi = \phi)$. q is, therefore, a tautology, hence its probability is fixed at 1. So for any proposition p held with a degree of confidence ϕ we can define a proposition q that has maximal confidence, thus: $C(p, \phi) \equiv C([\text{Prob}(p) = \phi], 1) \equiv C([\text{Prob}(p) = \varphi], 1) \equiv C(q, 1)$.

commonly conceived beliefs can vary in *strength*, i.e., a person can be more or less confident that *p*. Nevertheless, the convention adopted above requires us to avoid this manner of speaking and to construe differences in confidence as differences in some modal or probabilistic qualifier of the proposition believed. This departure from common usage can lead to some awkward locutions,[7] but the reader should bear in mind the intertranslatability of these modes of discourse and the fact that this manner of speaking is merely an instrumental device so that we only have to deal with full belief as the propositional attitude resulting from inference. Third, while beliefs (and propositions in general) are *expressible* in sentential form, that does not of itself imply that beliefs *are* sentential in form, i.e., that they are linguistically encoded in our minds. Instead, the propositional content may be encoded in something more like a picture or model of a possible state-of-affairs.[8] Fourth, for something to count as a belief at all, it must have some bearing on how we are disposed to act. Whether the belief is ever manifested in overt behavior depends, of course, on many factors: other beliefs, desires, opportunities for action, etc. But a belief must make some *difference* in one's behavioral dispositions. We implicitly recognize this whenever we dismiss a person's profession of belief as disingenuous because they don't act in accordance with what they profess. Fifth and finally, Freud taught us that beliefs can exist at any level of consciousness and continue to influence our behavior. Some we are explicitly aware of. Some we are aware of from

---

[7] For example, if we construe confidence levels in probabilistic terms, high confidence that *p* is the belief that p *is highly probable* and uncertainty that *p* is the belief that p *has mid-range probability* (i.e., the probability is approximately 0.5).

[8] For an extended defense of this view, see Stalnaker 1987, ch. 4. One advantage of this position is that it allows us to meaningfully ascribe beliefs to pre-linguistic children and higher animals.

time to time, but are not currently attending to. And some never emerge from the depths of subconsciousness.[9]

Another aspect of (D2) that deserves comment is the claim that it must *appear* to the one inferring that the conclusion is likely to be true if the premises are. This it implies that inference is, on some level, a conscious event. For, in order for someone to be appeared to, that which appears must enter into his consciousness. Let us, therefore, call this aspect of (D2) the *consciousness thesis*. This raises a question: Since consciousness comes in degrees[10]—we can be more or less conscious, more or less aware of ourselves and our surroundings—if inference must be a (somewhat) conscious event, how conscious must it be? One thing is clear from (D2): One must be sufficiently conscious for it to appear *that* the conclusion is likely to be true if the premises are, meaning that one's conscious state has representational, indeed propositional, content. In other words, inference requires having a cognitive perspective on the logical relationship between premises and conclusion. Call this the *perspective thesis*.

Now a perspective is a vantage point from which one sees. In this case it is the vantage point from which one sees *that* a certain proposition is true. A perspective can have more or less content, depending on how much it allows one to see; and it can be

---

[9] Perceptual beliefs are often like this. A illustrative example is the driver who travels down a familiar road for several miles and then suddenly "wakes up" and realizes that he doesn't remember any of the details, and yet still made all the proper turns, stopped at the lights, avoided other drivers, and so on. His ability to navigate successfully shows that he had beliefs about, say, whether a given light was red or green, but the beliefs never emerged into consciousness.

[10] On one end of the continuum lies consciousness proper (i.e., waking consciousness). Below that lies the preconscious, that which lies on the fringe of waking consciousness and which one can easily become aware of. Below that lies the unconscious, what Freud called the subconscious. Here lie thoughts, feelings, desires, etc. that can *in principle* be brought into waking consciousness, but only with difficulty and not by any direct effort. Finally, the continuum terminates in the nonconscious, that which cannot, even in principle, become an object of direct conscious awareness.

more or less clear, depending on the extent to which it enables one to see without obstruction or distortion. The consciousness thesis ties in at this point, for one's capacity to have a perspective, especially a clear one, directly correlates with one's degree of consciousness. Thus, waking consciousness is characterized by having a perspective on oneself and on one's surroundings, whereas unconsciousness is characterized by not having one. Additionally, a wide-awake person can generally perceive more things, and perceive them more clearly and distinctly than someone who is drowsy.

Again, a question arises. Perspectives come in degrees both with respect to content and with respect to clarity, so if inference requires having a conscious perspective, how much content and how clear of a perspective does it require? As for content, (D2) sets a clear minimum threshold—one must see *that* the conclusion must be or is likely to be true if the premises are. As for clarity, we should note that to have a perspective, or at any rate one with propositional content, is both to distinguish—self from non-self, subject from object, perceiver from perceived—and to relate—self to non-self, subject to object, perceiver to perceived. For example, to see *that* the cat is on the mat, one has to distinguish the cat from the mat and see that the former is related to the latter in the manner of "being on". At a minimum, therefore, inference according to (D2) requires that one's perspective be clear enough to distinguish premises from conclusion and to grasp how they are related.

Summing up the last three paragraphs, we arrive at the *conscious perspective (CP) thesis*—to infer is to have a cognitive perspective on the relation between premises and conclusion that is conscious and clear enough to see that the conclusion must be or is likely to be true if the premises are. The CP thesis sets forth a minimum standard for

cognitive events to qualify as inferential, but doesn't give us a sharply precise demarcation where everything on one side is equally inferential and everything on the other equally noninferential. Because consciousness, clarity, and content are matters of degree, and because beliefs can exist at all levels of consciousness, what we get is a continuum of belief-forming cognitive events taking propositional input in which perspectives on the relation between premises and conclusion range from nonexistent to vague to luminously clear, from seeing nothing to seeing merely that this premise supports this conclusion to seeing that it does so as an instance of a general logical principle (e.g., *modus ponens*). In other words, some belief-forming cognitive events are more inferential than others, depending on quality of the epistemic agent's perspective on the relation between premises and conclusion. Thus, the inferential shades into the noninferential such that there exist cognitive events that are analogous to inferences in that they are belief-forming cognitive events taking propositional input, but are nevertheless noninferential because they lack sufficient appreciation of the relation between premises and conclusion and therefore do not take place *because of* that appreciation. Peirce calls these 'associational suggestions of belief' and contrasts them both with 'reasonings'—in which we are "conscious, not only of the conclusion, and of our deliberate approval of it, but also of its being the result of the premiss [*sic*] from which it does result, and furthermore that the inference is one of a possible class of inferences which conform to one guiding principle."—and with 'acritical inferences'—in

which "we are conscious that a belief has been determined by another given belief, but are not conscious that it proceeds on any general principle."[11]

Modifying (D2) in accordance with the CP thesis, to the defense of which I now turn, we get:

> (D3) An inference is a cognitive event in which one passes from a set of believed or supposed propositions (the premises) to belief in a proposition (the conclusion) because from one's cognitive perspective it consciously and clearly appears that the latter must be or is likely to be true if the former is.

II. Internalist vs. Externalist Definitions of Inference

Let us call any definition of inference that affirms the CP thesis an *internalist* definition because it requires having an internal, first-person appreciation of the relation between premises and conclusion. Such definitions are controversial.[12] Let us call definitions of inference that deny the CP thesis *externalist*. An externalist approach will typically say that for a cognitive event to be inferential it is sufficient that it result in the formation of a belief *on the basis of* propositional input, for some suitable basing relation. Externalist definitions may construe the basing relation in different ways, but in no case will it be taken to imply having a conscious perspective on the relation between premise and conclusion. The simplest form of externalist definition says that inference is any cognitive process that takes propositional input and generates beliefs as outputs where what goes on in between is irrelevant to its inferential status.

---

[11] Peirce, EP 2.348.

[12] Fumerton, for example, explicitly rejects an internalist definition of inference: "[E]ngaging in conscious consideration of some set of premises on the way to reaching a conclusion that is consciously thought of as following from those premises is not a necessary condition for the justification supporting that belief to be inferential" (Fumerton 1995, p. 39).

A useful way of contrasting internalist and externalist definitions of inference is to say that the former takes a top-down approach to understanding inference, while the latter takes a bottom-up approach.[13] Internalist definitions are top-down in that they focus on paradigm cases of inference, such as those that result from the conscious, self-critical deliberations of a professional scientist or philosopher. Other cases are then understood *analogically* in relation to the paradigm, with the results that non-paradigmatic cases are seen as less inferential than paradigm cases and the inferential shades into the noninferential without any sharp line of demarcation between them. In contrast, externalist definitions typically work from the bottom-up, looking for the greatest common factor among all cases of reasoning. Since small children and higher animals are capable of reasoning, at least on a rudimentary level, externalist definitions will focus on cases where reasoning is automatic and instinctual rather than conscious and deliberate. The result will be a concept of inference that is *univocal*, one where Fido-type cases and Einstein-type cases are equally inferential, even though Einstein has a much better perspective on his reasoning.

Having distinguished between internalist and externalist definitions of inference, is there any principled reason for preferring one over the other? Clearly both are legitimate approaches to trying to understand the concept of inference. Which one is better depends on what purpose the definition is to serve. An artificial intelligence researcher who is trying to model human reasoning on a computer will likely find an

---

[13] Here I am adapting Dretske's distinction between top-down and bottom-up approaches to understanding the concept of knowledge. See Dretske, "Two Conceptions of Knowledge: Rational vs. Reliable Belief" in Dretske 2000, pp. 80-93.

externalist definition more useful because it allows inference to be understood as a black box input–output process, the sort of thing that can be modeled on a nonsentient computer. As long as the computer is able to efficiently and reliably generate the right sort of outputs for a set of inputs, everyone's happy. For this paper, however, we need to answer questions concerning inferential justification—and *for that purpose*, I now want to argue, an internalist definition is preferable. My argumentative strategy is as follows:

ID 1. For the purpose of answering questions of inferential justification (i.e., the epistemic justification of beliefs arrived at via inference), it is preferable to work with a definition of inference that best accommodates the factors upon which inferential justification depends.
2. Internalist definitions accommodate the factors upon which inferential justification depends better than externalist definitions.
   a. Epistemic justification (and therefore inferential justification) is a function of two distinct factors: epistemic responsibility and adequate grounding.
   b. Internalist definitions accommodate grounding as well and responsibility better than externalist definitions.
3. Therefore, internalist definitions of inference are preferable to externalist definitions with respect to inferential justification. (2,3)

When I speak in (ID.1) of a definition of inference "accommodating" the factors upon which inferential justification depends, I mean this in both a negative and a positive sense. Negatively, a definition of inference accommodates the factors upon which inferential justification depends if it *does not exclude* any of those factors—indeed, were that so, then we would be in the awkward position of having defined inference in such a way that few if any inferences resulted in justified beliefs. Positively, a definition of inference accommodates the factors upon which inferential justification depends if it *includes* those factors. Since we are looking for a general definition of inference, we are only interested in including those justificatory factors that make a difference to the intrinsic character of the cognitive event, i.e., those factors that affect whether or to what

extent a cognitive event is inferential in the first place. If any such justificatory factors exist, then definitions of inference that include those factors will be more useful than, and thus preferable to, ones that don't, at least for present purposes and all other things being equal. The burden of the rest of the argument is to show that such factors do exist and that having a conscious perspective on the relationship between premises and conclusion is one such factor. Since internalist definitions include that factor in the form of the CP thesis and externalist definitions don't, the former are preferable.

In support of (ID.2) I need, first, to identify some of the factors upon which inferential justification depends and, second, to show that internalist definitions of inference fare better than externalist definitions in accommodating those factors. I begin with Williams' observation that the phrase 'justified belief' is ambiguous: "What exactly is supposed to be 'justified': a *person's believing* some particular proposition, or the *proposition* that he believes?"[14] As he points out, these possibilities define two different standpoints of epistemic assessment.[15] The first standpoint focuses on whether a belief has been responsibly formed or is responsibly held, that is, on whether a person's believing is justified. Thus, we can ask whether a person negligently overlooked relevant counter-evidence in forming the belief, whether they employed reasonable methods of data acquisition and analysis, whether they unduly allowed their emotions to sway their thinking, whether their holding of a belief is dogmatic or remains sensitive to the evidence, and so on. Williams calls this type of justification 'epistemic responsibility' ('responsibility' for short)—the formation and maintenance of belief in conformance with

---

[14] Williams 2001, p. 21.
[15] Williams 2001, pp. 21-22. The same point is made in Fogelin 1994, ch. 1.

one's epistemic duties as determined by the epistemic goals of obtaining significant true beliefs and avoiding significant error.[16]

The second standpoint is concerned with whether the proposition believed has what Williams calls 'adequate grounding' ('grounding' for short). From this angle, a belief is justified if it is formed and maintained in such a way that it is objectively probable that the belief is true.[17] For example, an inferentially formed belief has adequate grounding if, given the objective probability that the premises are true and the degree to which they support the conclusion, the conclusion is at least probably true. And a perceptually formed belief has adequate grounding if one's perceptual faculties are reliable in the circumstances in which they were employed. Grounding is an external type of justification. As such, whether our beliefs are well-grounded or not may be ultimately outside of our purview. In other words, while we may have a first-person perspective on the grounding of our beliefs, and while that perspective may even play a role in the formation and maintenance of the belief in question, the belief's *being* grounded need not have anything to do with our having such a perspective. Regardless of whether a belief *seems* grounded or not, it *is* grounded if and only if its truth is objectively probable given the conditions under which it is formed and held.

---

[16] To be "in conformance with" one's epistemic duties means (a) that it does not violate any negative epistemic duties (i.e., duties regarding what ought *not* to be done), and (b) that it fulfills any epistemic positive duties (i.e., duties regarding what *ought* to be done).

[17] Williams 2001, p. 22. Two points of clarification: First, as I argued in chapter 2, 'objectively probable' cannot be interpreted in a purely logical sense, that is, apart from any substantive commitments. As there are nondenumerably many logical possibilities, probabilities are ill-defined on such a sample space. So 'objectively probable' has to mean probable relative to a sample space that is restricted to a denumerable set of relevant possibilities. How relevant possibilities should be characterized is, however, a difficult question that I will not here pursue. For a recent proposal, see Greco 2000a, ch. 8.

The distinction between these two standpoints emerges from a third-person perspective.[18] With respect to my own beliefs all I can do is strive to be as responsible as possible by doing what I can to ensure that my beliefs have adequate grounding. In other words, I can use the epistemic resources at my disposal in a manner that, as far as I can tell, is likely to generate the best epistemic payoff in terms of generating significant true beliefs and avoiding significant error. I cannot simultaneously regard myself as responsible in believing *p* and as lacking adequate grounds for *p*. With respect to someone else's beliefs, however, I may recognize both that he is responsible in believing *p* and that he lacks adequate grounds for *p* because I may possess information he doesn't. For example, prior to the Copernican revolution, it was possible to be responsible in believing that the Earth was at the center of the solar system.[19] Today, after the work of Galileo, Kepler, Newton, and the discovery of the concept of inertia, stellar parallax, the phases of Venus, and so forth, we see the geocentric hypothesis as poorly grounded and as something that no educated person could be responsible in believing.

While distinct, justification as responsibility and justification as grounding are closely related. First, the standards we employ in evaluating responsibility are those we deem conducive to grounding:

> The point of setting standards for epistemic responsibility is to reduce the risk of error. Accordingly, epistemically responsible behaviour is itself a kind of grounding: by behaving in an epistemically responsible way, I increase the likelihood that the beliefs I form are true.[20]

---

[18] Williams 2001, pp. 22-23.
[19] The anti-Copernicans had several powerful arguments in their arsenal. For a summary account see Finocchiaro 1989, pp. 15-25.
[20] Williams 2001, p. 22.

In other words, to be epistemically responsible is to seek adequate grounding for one's beliefs. We criticize someone as epistemically irresponsible when they employ methods of fixing belief that they ought to have recognized as unreliable, such as astrology, crystal ball gazing, or reading tabloid newspapers.

Second, the demands of responsibility determine how much grounding is *adequate*. While it seems reasonable to stipulate that adequate grounds must establish *at least* an objective probability greater than 0.5, the dialectical and/or moral context may raise the standards of responsibility much higher. For example, in a criminal trial the prosecution accepts a burden of proof to show beyond a reasonable doubt that the accused is guilty. In such a dialectical context, we might expect the grounds cited to establish an objective probability of greater than 0.95. Additionally, where grave consequences threaten if we are wrong, it may be morally necessary to seek more grounding. In other words, moral duties can augment epistemic duties.[21] For example, we wouldn't want a new medication to be approved merely because it was "probably safe".

Third, epistemically responsible or irresponsible behavior can drastically augment or decrease our chances of forming true beliefs and avoiding false beliefs. Thus, grounding is often contingent upon responsibility. For example, even if perceptual conditions are highly favorable (e.g., there is good lighting, no obstructions, etc.), a perceiver who barely opens her eyes or takes only a cursory glance will not have as much grounding as if she had taken a good look. And deliberative methods and instruments (e.g., a test for blood type) are reliable only if properly (i.e., responsibly) used.

---

[21] On the relation between epistemic and moral duty see Susan Haack, "'The Ethics of Belief' Reconsidered" and Bruce Russell, "Epistemic and Moral Duty", both in Steup 2001.

Fourth, knowledge requires both kinds of justification. That is, both are *epistemically* relevant. First, one can have a true, epistemically responsible belief that fails to be knowledge for want of adequate grounding. For example,

> I believe that the match will kick off at 3:30 because I looked up the time in the newspaper. But you have just been on the phone to the editor, who has told you that there has been a misprint and the match is actually slated to start at 2:30. So you recognize that, although I am personally justified in my belief—I have not been irresponsible—the evidence I have for it is misleading. So even if there is an unexpected delay and the match does kick off at 3:30, we would not say that I knew it would be so.[22]

Similarly, one can have adequate grounding by employing a belief-fixing procedure that is in fact reliable but fail to have knowledge because one ought to have distrusted the procedure. BonJour offers the example of a perfectly reliable clairvoyant confronted with cogent undefeated evidence against the possibility of clairvoyance. In such a case it seems that the clairvoyant would be irresponsible in trusting that faculty and that any beliefs he formed based on that faculty would thus not amount to knowledge.[23] Epistemic justification is therefore a function of both responsibility and grounding. Since inferential justification is a subtype of epistemic justification, it too is a function of those factors.

For inferential justification to be a function of responsibility and grounding does not imply, however, that either of those factors is related to inference in such a way that we ought to positively accommodate them in a definition of inference. To establish (ID.2) I also need to show that internalist definitions of inference, i.e., definitions that include the CCP thesis, are *for that reason* better able to accommodate the factors upon which inferential justification depends. More specifically, I need to show that having a

---

[22] Williams 2001, p. 23.
[23] Laurence BonJour, "Externalist Theories of Empirical Knowledge," in Kornblith 2001, pp. 17-20.

conscious perspective on the relation between premises and conclusion is relevant for either responsibility, grounding, or both. I will focus on responsibility. Here's my argument:

ER
1. Epistemic responsibility implies the existence of an epistemic duty to utilize one's resources in such a manner that, as far as one can tell, is conducive to securing adequate grounding for one's beliefs..
2. Duties only pertain to an agent insofar as it is really possible for him to fulfill them.
3. It is really possible for an agent to fulfill a duty only if he has a measure of control over its fulfillment.
4. Therefore, epistemic responsibility applies to an agent only to the extent that he can exercise control over the grounding of his beliefs. (1,2,3)
5. An agent can exercise voluntary control over something only if he has a conscious perspective on it.
6. Therefore, an agent can exercise voluntary control over the grounding of a belief only if he has a conscious perspective on its grounding. (5)
7. The ability to exercise voluntary control over the grounding of a belief is relevant for epistemic responsibility. (4)
8. Therefore, having a conscious perspective on the grounding of a belief is relevant for epistemic responsibility. (6,7)
9. Inferential grounding is a function of the truth of the premises and of the relation between premises and conclusion.
10. Therefore, having a conscious perspective on the relation between premises and conclusion is relevant for epistemic responsibility. (8,9)

This is a valid argument, with the logic as I have indicated. It remains to examine the independent premises (viz., 1, 2, 3, 5, and 9).

First, we have (ER.1). Epistemic responsibility, like moral responsibility, is a deontological notion. It implies that we have epistemic duties. As mentioned above, the norms for epistemic responsibility relate to our need to secure adequate grounding for our beliefs, and we all recognize that there are things we should do and things we should not do if we want to put ourselves in a good position to have well-grounded beliefs. For example, we should try to find reliable sources of information, make careful observations as opposed to cursory ones, check ourselves where possible, look at the arguments and

evidence on both sides of an issue, not reason too hastily, cultivate good cognitive habits, employ sound research methods, and so on. Thus, we have an epistemic duty to utilize our resources (e.g., cognitive and perceptual faculties, available evidence, research equipment, testing procedures, etc.) in such a way that what we believe is, as far as we can tell, adequately grounded.

(ER.2) is implied by the principle that ought implies can, or more precisely, can or could have.[24] Duties only apply to an agent if it is, or at some point was, a real possibility for the agent to fulfill the duty. Conversely, cannot implies need not. If something is not and never was a real possibility for an agent, then the agent cannot have an obligation with respect to that something. For example, I can't have epistemic duties to secure apodictic proofs for contingent propositions or to complete an infinite regress in supplying grounds for my beliefs or to exercise logical omniscience (i.e., be able to see *all* the deductive entailments of a given proposition or set of propositions), for those aren't real possibilities, at least not for epistemic agents like us.

For it to be a real possibility *for an agent* to fulfill a duty, the agent must have a measure of *control* over fulfillment of the duty. Thus (ER.3). In other words, we can only have responsibility for things (i.e., be properly subject to praise or blame regarding them) to the extent that we can be responsible for them (i.e., be causally efficacious in bringing them about). More precisely, since it is agents that are the bearers of duties, the relevant sense of control is that of *agency*. To be an agent is to be a source of power that is

[24] An agent can be culpable for an action even though they cannot now refrain because they could have made different choices earlier and, had they done so, then they would have been able to refrain now. For example, Bob, a longtime smoker, may now find it impossible to quit. But his smoking now is still culpable (though perhaps in a mitigated sense) because he could have not started smoking in the first place, or he could have quit earlier while it was still possible for him to do so.

personal, characterized by a mind and a will. Thus we are not interested in the kind of control that a thermostat has over the temperature in a house, for that is not an exercise of *personal* power. Nor are we interested in cases where a person is merely an instrumental cause, such as we would be if our mind and will were externally controlled. Such a person is not a genuine agent, for they are not functioning as a *source* of personal power. Now, if we have no control, in the sense of agency, over whether something occurs, then we cannot be subject to duties concerning it. Conversely, the more control we have, the more responsible we are for things turning out one way rather than another, and the more that may be required of us in terms of duty fulfillment.

As far as epistemic duties are concerned, the relevant scope of control concerns whatever we can do or refrain from doing that, as far as we can tell, positively or negatively impacts the grounding of our beliefs. The reason for the "as far as we can tell" qualification lies in the fact that what does *in fact* impact the grounding of our beliefs, whether for better or for worse, may differ markedly from what we *think* impacts it. If we happen to live in a universe ruled by a Cartesian demon, or if we are brains in a vat, then our most diligent cognitive efforts may lead to far more error than if we were simply to, say, form beliefs through wishful thinking and self-hypnosis. But, as far as we can tell, things are not that way. If we are wrong, we are wrong, but even so our error is nonculpable because we cannot reasonably be expected to know better if indeed we live in such a universe. Brains in a vat may lack knowledge for want of true belief and adequate grounding, but they can still be epistemically responsible by doing the best they can, given what they have to work with.

This brings us to (ER.5). To *exercise* control in the sense of agency is to bring something about through a deliberate act of will. What is brought about may either be the intended effect—call this *voluntary* control. Or it may be an unintended effect—call this *nonvoluntary* control. Or it may include both. For example, thinking the gun is unloaded, I may exercise voluntary control in pulling the trigger and nonvoluntary control in shooting myself. Furthermore, one cannot exercise nonvoluntary control without at least *attempting* to exercise voluntary control—as someone might (voluntarily) try to lift a heavy boulder and only succeed (nonvoluntarily) in giving himself a hernia. Thus to exercise any control, whether voluntary or nonvoluntary, one must at least attempt to exercise voluntary control over something. But to exercise, or even attempt to exercise, voluntary control over something requires that one have a conscious perspective on that something. To see this, consider on the one hand an unconscious person, perhaps fast asleep. As unconscious he has no perspective on himself or his surroundings. On the other hand, consider a fully conscious person who is aware of herself and her environment and alert to changes in either. Now the unconscious person as such has no control over events. He may have had some say over when and where he went to sleep, but that say has to be exercised while he is conscious, i.e., *before* going to sleep. The wide-awake person, on the other hand, has a conscious perspective on herself and her environment, and this makes the exercise (or even the attempted exercise) of voluntary control possible. For, to have a perspective, as noted above, is both to distinguish and to relate; and both are necessary for the exercise of voluntary control. For example, to exercise voluntary control over getting from point A to point B, I have to distinguish point B from where I am, and I have to relate it to where I am so that I can see how to get

from here to there. In other words, I have to have a cognitive perspective that encompasses my current state, my target state, and a path between them. Thus the exercise of voluntary control requires having a conscious perspective that encompasses one's current state, one's target state, and a path between them. It follows that to exercise voluntary control over the grounding of a belief I must have a conscious perspective encompassing its current degree of grounding, the target degree of grounding, and a path between them, that is, I have to have some idea of the impact different actions are likely to have on the belief's grounding.

Furthermore, the possibilities for exercising voluntary control over the grounding of one's beliefs vary with the quality of one's perspective—the more conscious, content-rich, and clear our perspective, the greater the amount of voluntary control we are able to exercise. In some cases, e.g., perception, voluntary control is often minimal or nonexistent because we typically are not even aware of having formed a belief. You open you eyes and, voilá, form the belief that the cat is on the mat. Only after you acquire a conscious perspective on that belief can you begin to critically assess it—is it *really* a cat, or just a stuffed animal? In other cases, beliefs result from conscious deliberation and active investigation informed by a rich and clear perspective. Here the possibilities for control are numerous. An experimental scientist, for example, has to take considerable care in designing, conducting, and evaluating an experiment. Decisions have to be made with respect to which variables to consider relevant and which to ignore, which instruments to use and how to use them, how to record data and how to analyze it, how much data to obtain, etc. Bad decisions in any of these areas could greatly undermine the grounding of his conclusions, rendering them unreliable. In general, then, the more fully

conscious the processes involved in belief formation and maintenance, and the better one's perspective on those processes, the more controllable they are, and the greater the importance of epistemic responsibility. Because these factors are matters of degree, the requirements for the responsibility component of epistemic justification should scale along with one's degree of consciousness and the quality of one's perspective.

It follows from what has been said so far that having a conscious perspective on the grounding of one's beliefs is relevant for epistemic responsibility. For epistemic responsibility pertains only to what we have control over in the way of grounding, and having a conscious perspective on the grounding of one's beliefs is a necessary condition for the exercise of voluntary control over their grounding. Grounding, however, refers to the objective probability that a belief is true. Since inferentially formed beliefs are conclusions, inferential grounding refers to the probability that the conclusion is true, which, as (ER.9) states, is a function of both the truth of the premises and the relation between premises and conclusion. Having a conscious perspective on the relation between premises and conclusion is, therefore, a prerequisite for exercising voluntary control over inferential grounding and thus is relevant for epistemic responsibility.

A couple additional considerations buttress this result. First, it is clear that in at least *some* cases of inference we do have a conscious perspective on the relation between premises and conclusion—anyone who has ever grasped the validity of *modus ponens* can testify to that. And it is clear that at least *some* cases of inference result from a conscious and controlled deliberative process. Further, paradigmatic cases of inference satisfy those conditions. This should not be surprising. Paradigmatic cases of anything have to be *clear* because we use them as reference points to help us understand the not-

so-clear cases. Inference is a cognitive event. And since each of us has a kind of privileged introspective access to his own mental life, the clearest examples of any type of cognitive event that we could possibly find would be those in which we are consciously engaged, if any such examples exist. But, as just noted, such examples do exist in the case of inference; therefore, paradigmatic cases of inference will be those cases in which we are the most consciously engaged, which means those cases resulting from a controlled process of deliberation in which we have a conscious perspective on the relation between premises and conclusion.

Second, the quality of one's perspective on the relation between premises and conclusion can affect one's degree of inferential justification with respect to *grounding*, not just responsibility. Suppose, for example, that *p* entails *q*. Bob believes *p* and infers *q* because of a vague notion that it is entailed by *p*. Mary believes *p* and infers *q* because she perceives clearly and distinctly that *p* entails *q*. Clearly, Mary has greater justification, in the sense of grounding, for believing *q* than Bob precisely because she has a better cognitive perspective on the relation between premise and conclusion. She has more grounding because one who has a clear and distinct perspective on a logical relation is less likely, all other things being equal, to make a mistake in reasoning. Thus the objective probability that Mary has reached a true conclusion is higher than with Bob.

Summing up, having a conscious perspective on the relation between premises and conclusion is relevant to inferential justification. Therefore, definitions of inference that positively accommodate this by including the CP thesis are preferable (for present purposes) to ones that don't. Accordingly, internalist definitions are preferable to externalist ones.

III. Does Inferential Internalism Lead to Skepticism?

There is an objection that may be raised at this point. It turns on the fact that my argument in favor of internalist *definitions* of inference invokes a partially internalist position on inferential *justification*. For I have argued that internal considerations (e.g., having a conscious perspective on the relation between premises and conclusion) are relevant to inferential justification. And since I have *defined* inference as a cognitive event in which we have such a perspective—more specifically, one in which the truth of premises appears to render likely the truth of the conclusion—I am committed to the thesis that having a conscious perspective of the right sort is a necessary condition for inferential justification. Let us call one who accepts this condition on inferential justification an *inferential internalist* and one who denies it an *inferential externalist*. The objection, which has been forcefully presented by both Richard Fumerton and John Greco, is that inferential internalism leads to skepticism.[25] In this section I argue that the objection fails and does so in a way that, ironically, gives us an additional reason for preferring internalist definitions of inference.

> We start with Fumerton's proposed Principle of Inferential Justification (PIJ):
>
> > To be justified in believing one proposition *P* on the basis of another proposition *E*, one must be (1) justified in believing *E* and (2) justified in believing that *E* makes probable *P*.[26]

As Fumerton points out, clause (1) of PIJ is a platitude that virtually everyone accepts. Clause (2), however, is disputed. Inferential internalists accept it, but inferential

---

[25] Tracing out the skeptical ramifications of PIJ is a central theme of Fumerton 1995. Greco's argument against PIJ is contained in Greco 1999 and Greco 2000b.
[26] Fumerton 1995, p. 36.

externalists do not.[27] Fumerton, who accepts PIJ, argues that it has serious skeptical ramifications, but continues to hold on to a slim hope that skepticism may be avoided. At any rate, he thinks that we cannot give up PIJ without giving up on the traditional epistemological enterprise.[28] Greco, on the other hand, argues that the skeptical consequences of PIJ cannot be avoided and takes that to be a decisive reason for rejecting PIJ, and with it, inferential internalism.[29]

How is it that PIJ leads to skepticism? Very easily. If the belief referred to in clause (2) of PIJ, i.e., the belief that *E makes probable P*, must be *inferentially* justified, then it must be the conclusion of an inference from some body of evidence E'. Therefore, by PIJ, we must be inferentially justified in believing that *E' makes probable (E makes probable P)*. But that, then, must be the conclusion of an inference from E", and so on. We fall into either an infinite regress or vicious circularity, which means that inferential justification is impossible and radical skepticism results. There are two ways to avoid this skeptical conclusion: (i) reject clause (2) of PIJ and adopt an externalist position on inferential justification (Greco's approach), or (ii) maintain that the belief that 'E makes probable P' can be *noninferentially* justified (Fumerton's approach). The problem with (ii) is that it seems very difficult to defend for inductive inferences.

In the case of deductive inference, the relationship between E and P is presumed to be *necessary*. Accordingly, it is plausible to suppose that we can just 'see' whether or not the conclusion is in fact contained in the premises via rational intuition. As noted

---

[27] Fumerton 1995, p. 37.
[28] Fumerton 1995, p. 220.
[29] Greco 2000b.

above, anyone who has grasped the validity of *modus ponens* has phenomenological confirmation of our ability to 'see' such relations. In more complex cases, such as long mathematical proofs, we may not be able to directly appreciate the relation between our starting premises and the final conclusion, but we can reach the conclusion through a series of deductive inferences in each of which we can directly 'see' the relation. In the case of inductive inference, however, the conclusion is presumed to reach beyond the premises. It is not so clear, therefore, how we could just 'see' by reflection that 'E makes probable P'. Fumerton realizes this and concludes that the only way that we could be noninferentially justified in believing that 'E makes probable P' and thereby avoid inductive skepticism is if the sort of *probability* involved in that claim represents an internal relation holding between propositions. In other words, we must

> understand the concept of nondeductive epistemic probability as being . . . like the concept of entailment, and . . . subsequently convince ourselves that epistemic principles are necessary truths knowable *a priori*.[30]

Fumerton examines several different conceptions of probability and concludes that each is inadequate to the task. Relative frequency and propensity theories render 'E makes probable P' a complex *contingent* truth, and it is very implausible to suppose that we could 'see' via rational intuition that such propositions are true.[31] Analytic conceptions of probability seem wrong because it is not at all clear how skeptics who deny that our evidence for belief in the external world makes probable that there is an eternal world are contradicting themselves.[32] Subjective conceptions of probability lack a

---

[30] Fumerton 1995, p. 198.
[31] Fumerton 1995, p. 193.
[32] Fumerton 1995, pp. 194-195.

definite connection with truth and for that reason Fumerton thinks they are too weak to handle skeptical challenges.[33] The most promising candidate, thinks Fumerton, takes 'E makes probable P' to be a synthetic *a priori* truth. The problem here is that there seems to be no way to ground the claim that 'E makes probable P' is synthetic *a priori* without appealing to contingently reliable evidence of the very sort that skepticism calls into question.[34] He concludes on a skeptical tone:

> I cannot quite bring myself to believe that I am phenomenologically acquainted with this internal relation of making probable bridging the problematic gaps. . . . [I]n the end, I strongly suspect that the probability relation that philosophers *do seek* in order to avoid skepticism concerning inferentially justified beliefs is an illusion.[35]

By way of response to Fumerton, I argued in chapter 2 that inductive inference rules are, by their very nature, contingent. Fumerton's failure to find an objective and necessary link between E and P is, therefore, unsurprising. Greco argues convincingly that no kind of probability can do all the work that Fumerton needs it to do:

> Let us say that a kind of probability is "subjective" if it does not imply reliability. . . . Let us call a kind of probability "objective" if it does imply reliability. . . . And now the problem is this: If we think of probability as subjective, then it is plausible that propositions of the form "E makes probable P" can be necessary truths, and so knowable non-inferentially. However, knowledge of subjective probability relations will not involve an awareness of the reliability of one's evidence. If we think of probability as objective, then knowledge that "E makes probable P" does give one knowledge that E is a reliable indication of P. However, it will now be entirely implausible that such propositions are necessary, and so knowable non-inferentially.[36]

---

[33] Fumerton 1995, p. 196.
[34] Fumerton 1995, pp. 201, 204-218.
[35] Fumerton 1995, p. 218.
[36] Greco 1999, p. 283.

Greco concludes that we should abandon PIJ in favor of reliabilism, which does not require that evidential relations be necessary or that one be able to know or even justifiably believe that one's evidence is reliable.[37] With Greco and against Fumerton, I don't think that evidential relations need be necessary. Against both, I deny that "E makes probable P" must be necessary in order to be noninferentially justifiable.

Why think that evidential relations must be *necessary* in order for the belief that they obtain to be noninferentially justifiable? Fumerton only gives one argument: If the claim that "E makes probable P" was contingent, it would have to be highly complex, and it is highly implausible that we could 'see' that such a complex claim was true.[38] But is this really so implausible? I don't think so. A well-trained, experienced mind can grasp contingent evidential relations very quickly. For example, an experienced auto mechanic can tell by the sound of an engine that a piston rod needs replacing. The relation between the evidence (the engine sound) and the conclusion (there's a bad piston rod) is clearly contingent, for it is certainly possible that the sound could be caused by something else.[39] What makes the grasp of this contingent relation possible is the presence of a large and complex body of background beliefs: beliefs about the internal structure of a car engine, beliefs about what a healthy engine sounds like, beliefs about what different types of unhealthy engines sound like, etc. These background beliefs define a finite probability

---

[37] Greco 1999, p. 284.
[38] Fumerton 1995, p. 193. Greco merely cites Fumerton with approval on this point (Greco 1999, p. 283).
[39] But couldn't the evidential relation be *necessarily probable* rather than contingent? No, for two reasons. First, in terms of the standard possible worlds semantics, probabilities aren't well-defined. For P to be necessarily probable given E, would mean that in *most* possible worlds in which E is true, P is true. But there are nondenumerably many possible worlds, and it is not at all clear what "most" means when comparing denumerable infinite sets, much less nondenumerable ones. Second, even if a clear meaning can be given to "necessarily probable," how could we ever be in a position to *know* whether E is true in most possible worlds or not? We have no way of sampling possible worlds to ascertain relative frequencies.

space in which the conditional probability of the piston rod hypothesis given the evidence is higher than any other available hypothesis. It is precisely because this complex body of background beliefs is *in the background* that the mechanic is able to grasp the contingent connection between the engine's sound and the piston rods. His background beliefs enable him to noninferentially 'see' contingent evidential relations without themselves being seen.

The reason grasping contingent evidential relations seems psychologically implausible to Fumerton is that he holds to an externalist definition of inference. We see this both in his formulation of PIJ, in which he characterizes inference simply as "believing one proposition *P* on the basis of another proposition *E*," and in his explicit denial that inference requires having a conscious perspective on the relation between premises and conclusion.[40] But if inference is forming a belief "based on" other beliefs where the basing relation has no essential tie to one's conscious perspective, then it becomes nearly impossible to avoid conflating the distinction between evidence and background beliefs. The reason is simply that the distinction between evidence and background beliefs just is one of consciousness. Background beliefs are "in the background" because we are not conscious of them. If we were then they would be in the foreground along with our evidence. With that conflation made, the mechanic's inference from evidence E (the engine's sound) to conclusion P (piston rod is broken) by means of background beliefs B becomes indistinguishable from an inference from (E and B) to P. But B is a highly complex network of propositions, which raises serious questions about

---

[40] Fumerton 1995, p. 39. See footnote 12 above for quote.

the psychological plausibility of grasping that (E and B) makes probable P. This complexity and the resulting psychological implausibility are avoidable, thinks Fumerton, only if "E makes probable P" is grounded on a necessary and thus purely *internal* relation between E and P—in which case we could grasp the relation without the help of background beliefs. But, as both Fumerton and Greco indicate, this won't work, for we lack a suitable account of probability on which "E makes probable P" is both a priori necessary and expressive of an objective truth about the world.

All this changes, however, if we adopt an internalist definition of inference. If inference requires having a conscious perspective on the relation between premises and conclusion, then a clear distinction emerges between a person's evidence and their background beliefs. Because background beliefs are in the background, and therefore outside of consciousness, we lack a conscious perspective on their relation to the conclusion. Therefore, the relation of background beliefs to the conclusion is *not inferential*. In other words, background beliefs are not premises in an inference. A person's evidence, however, is that which is evident to them, that which stands out in sharp relief against the background.[41] It is therefore part of their conscious perspective and thus can serve as a premise in an inference. Accordingly, a person's background beliefs are not evidence for them. Rather, they are what allow contingent evidential relations to be seen. Thus, adopting an internalist definition of inference enables us to avoid the skeptical problem that plagues Fumerton's PIJ by allowing us to preserve the

[41] I use the term "evidence" here in the subjective sense of "evidence *for a person S* that *p*" to denote information that one has a conscious perspective on. Frequently, however, the term is used in objective sense of "evidence for *p*" to denote information of any sort that bears on the truth or falsity of *p*.

distinction between evidence and background beliefs. Because internalist definitions include the CP thesis, whereas externalist definitions do not, they give us a narrower scope for inference—simultaneously restricting the scope of PIJ and expanding the possibilities for noninferential justification. This creates a space through which we can escape Fumerton's problem. Given the utility of the evidence/background belief distinction for defusing this skeptical problem, the fact that internalist definitions of inference enable us to preserve that distinction while externalist definitions undermine it is another reason for preferring internalist definitions of inference.

IV. Inferential and Noninferential Justification: A Rejoinder to Hume

Having defused one skeptical problem (Fumerton's) by means of an internalist definition of inference, can we defuse another? What about Hume's? The answer, I submit, is yes. As characterized in chapter 2, Hume's skeptical argument has the following structure (though we could substitute AA's for IR's with equal effect):

(SA2)
1. All inductive inferences require a contingent inductive inference rule (IR).
2. A belief is inferentially justified only if belief in the inference rule is justified.
3. Therefore, a belief is inductively justified only if belief in the IR is justified. (1,2)
4. There is no way of justifying belief in the IR.
    a. All justification of belief is either inferential or noninferential.
    b. Belief in the IR cannot be noninferentially justified (it is not self-evident, evident to the senses, incorrigible, or otherwise properly basic).
    c. Inferential justification is either deductive or inductive.
    d. Belief in the IR cannot be justified by deduction.
    e. Belief in the IR cannot be justified by induction.
    f. Therefore, belief in the IR cannot be inferentially justified. (c,d,e)
    g. Therefore, belief in the IR cannot be justified. (a,b,f)
5. Therefore, beliefs arrived at via inductive inference are unjustified. (3,4)

As we have just seen, however, adopting an internalist definition of inference limits the scope of inferential justification to cases where we have a conscious perspective on the relation between premises and conclusion and expands the possibilities for noninferential

justification by allowing background beliefs to play a role. This suggests that we ought to reexamine premise (4.b) of SA2. I claim that there is no good reason to accept this premise. More specifically, there is no good reason to limit noninferential justification to what is self-evident, evident to the senses, or incorrigible, and there is good reason to think that inductive inference rules (and auxiliary assumptions) can be properly basic.

To show this we need to articulate principles of both inferential and noninferential justification that reflect the importance of having a conscious perspective on the grounding of one's beliefs. If we revise PIJ to accommodate the responsibility and grounding aspects of epistemic justification as well as the CP thesis, we get the following:

> RPIJ: The belief that *p* is *inferentially* justified for a person S iff (1) S's grounds[42] for believing *p* are in fact (a) adequate and (b) responsibly obtained, and (2) S has or has had[43] a conscious perspective on the adequacy of his grounds for believing *p*, in which perspective his grounds appear adequate.

Clause (1) of RPIJ (for *revised* PIJ) reflects the fact that inferential justification, as a species of epistemic justification, is a function of both grounding and responsibility. Clause (2) incorporates the CP thesis, reflecting the fact that to have a conscious perspective on the adequacy of one's grounds is to see them as evidence, as premises supporting a conclusion. And for one's grounds to appear adequate is, at the very least,

---

[42] Grounds are anything that, in contributing to the formation and maintenance of a belief, contributes to its grounding, i.e., anything that significantly affects the objective probability that the belief is formed and maintained in such a way as to be true. This includes S's evidence for *p*, the reliability of S's cognitive faculties, and the methods or procedures by which S acquired his evidence.

[43] I say "has or has had" a conscious perspective to accommodate beliefs that were formed via an inference, and therefore have inferential justification, but for which we do not currently have a conscious perspective on the adequacy of their grounds. For example, I want to say that my belief in the Pythagorean theorem is inferentially justified because I have worked through the mathematical proof, even though I do not currently have the proof in mind.

for it to appear to one that the conclusion is likely to be true if the premises are. Now, on an internalist definition of inference, noninferential cases are those in which one forms a belief *without* having a conscious perspective on the adequacy of one's grounds. But like all forms of epistemic justification, it too is a function of both grounding and responsibility. Thus we arrive at a principle of noninferential justification (PNJ):

> PNJ: The belief that *p* is *noninferentially* justified for a person S iff (1) S's grounds for believing *p* are in fact (a) adequate and (b) responsibly obtained, and (2) S has never had a conscious perspective on the adequacy of his grounds for believing *p*.

Note that RPIJ and PNJ are mutually exclusive with respect to a particular set of grounds, but not with respect to a given belief. Thus if S has or had had a conscious perspective on part of his grounds for believing *p* and has never had a conscious perspective on the other part of his grounds, and if clause (1) is satisfied, then the justification for S's belief that *p* may be partially inferential and partially noninferential.

Now let's look at premise (4.b). The skeptical claim is that belief in inductive auxiliary assumptions (AA's) or inductive inference rules (IR's) cannot be noninferentially justified. More broadly and to the point, the skeptical claim is that belief in a contingent, possibly complex proposition that goes beyond the information in hand—this is what both AA's and IR's are—cannot be noninferentially justified. My response is the same as it was to Fumerton's claim that contingent evidential relations cannot be noninferentially justified—why not? To be noninferentially justified according to PNJ a belief needs to have adequate grounds responsibly obtained under conditions where the epistemic agent has never had a conscious perspective on the adequacy of his grounds. Why can't these conditions be satisfied in the case of beliefs in complex contingent propositions going beyond the information in hand?

First, since we have already established that neither AA's nor IR's are properly construed as premises in an inductive inference, we can suppose that they lie among our background beliefs, outside of our conscious perspective. There is nothing problematic in this. Lots of beliefs in complex contingent propositions lie among our background beliefs, and many (indeed most) of these refer to things that go beyond our information in hand and pertain to unobserved matters of fact. Thus it is plausible that we can lack a conscious perspective on the grounding of such beliefs. Furthermore, it is plausible that some of these complex contingent background beliefs have always been in the background, such that we have never had a conscious perspective on their grounding. For example, it appears that some animals have complex contingent background beliefs:

> Without something like geometrical, kinetic and mechanical conceptions, no animal could seize his food or do anything which might be necessary for the preservation of the species.[44]

Many animals display remarkably complex instincts related to mating, migration, the rearing of young, what plants to eat and which to avoid, and so on. And animals presumably have never had, nor are capable of having, a conscious perspective on the grounding of their instincts. So there doesn't seem to be any absurdity in supposing that we humans have complex contingent background beliefs that have always been in the background and that occasionally function for us as AA's and IR's. In fact, there is reason to think that we have many such beliefs. Some may be innate, even instinctual, but many may be formed via the subconscious or nonconscious functioning of our cognitive faculties. Think again of the driver mentioned in note 9 above, who travels down a

---

[44] Peirce, CP 6.418.

familiar road for several miles and then suddenly “wakes up” and realizes that he doesn’t remember any of the details, and yet still made all the proper turns, stopped at the lights, avoided other drivers, and so on. The fact that he was able to navigate successfully strongly suggests that he had beliefs about, say, whether a given light was red or green, even though he was not conscious of those beliefs. Thus, it is plausible that we possess a great many beliefs that are similarly formed in the background through the automatic functioning of our cognitive faculties and then remain there, such that we have never had a conscious perspective on their grounding.

Second, having adequate grounds just means that the belief is formed and maintained in such a way that it is objectively probable that the belief is true. Why can’t this be true of AA’s and IR’s? Again, we have lots of adequately grounded background beliefs in complex contingent propositions that go beyond our information in hand. For example, I believe that if I were to combine hydrochloric acid and sodium hydroxide in equal proportions that I would wind up with an aqueous solution of sodium chloride (table salt). This belief is well grounded—I have the authority of my chemistry textbook to back it up. Since I rarely think about chemistry anymore it is almost always in the background. It is certainly complex and contingent—it is conceivable that the laws of chemistry could change. And it goes beyond the information in hand in that it pertains to a possible experiment, one that I haven’t performed. Now why can’t the same thing be true for a complex contingent belief that has always been in the background? To return to our semi-conscious driver, if his cognitive faculties are in fact functioning reliably in forming beliefs about the color of stop lights and so forth then, all other things being equal, the beliefs are adequately grounded. And if we have any instinctual beliefs, it is

similarly possible for them to be adequately grounded if our instincts are in fact sufficiently attuned to reality.

Third and finally, there is no good reason to think that one cannot have *responsibly obtained* adequate grounds for a background belief in a complex and contingent proposition about unobserved matters of fact. My belief in the combinatory properties of HCl and NaOH is a case in point. Furthermore, if there are complex contingent beliefs that have always been in the background, that are either instinctive or results of the subconscious or nonconscious operation of our cognitive faculties, then we have never had the conscious perspective on them necessary to exercise or attempt to exercise voluntary control over their grounding. So the only sort of control that we could over their grounding would be the nonvoluntary sort that we exercise when we are trying to do something else. Thus, I could train my perceptual faculties to be more discriminating or I could impair my cognitive faculties by taking drugs, and so on. Now suppose that our cognitive faculties are generally reliable *at the outset*, whether through evolutionary attunement or the design of our creator, before we have had a chance to voluntarily do anything that might affect the grounding of the beliefs they produce. To suppose that our cognitive faculties are *generally* reliable is just to say that, when they do yield beliefs, they yield true ones most of the time in the vast majority of the types of circumstances in which we typically find ourselves. Thus the supposition of general reliability is perfectly compatible with the observation that we can and frequently do err. In chapter five I will argue that we have compelling reasons to think this supposition is true, but for now it suffices that there be no compelling reasons to think it is false, since all we are trying to show here is that premise (4.b) of the skeptical argument can be

plausibly denied. Now if we suppose that we start out as generally reliable cognitive agents, then, absent standards raising contextual factors, it follows that the requirements of epistemic responsibility are initially fulfilled. For if our cognitive faculties are generally reliable, then the beliefs they spontaneously yield will normally be adequately grounded. Since epistemic responsibility requires using whatever control we have to secure adequate grounding for our beliefs, if those beliefs are already adequately grounded, then epistemic responsibility requires nothing further of us except not doing anything that would seriously undermine their grounding. Therefore, absent compelling reasons for rejecting the general reliability of our cognitive faculties, it is plausible to think that we can have responsibly obtained adequate grounds for beliefs in complex and contingent propositions about unobserved matters of fact even when we have never had a conscious perspective on their grounding.

In sum, there seems to be nothing inherently implausible in supposing that the beliefs that serve as our initial AA's and IR's cannot be noninferentially justified. Thus premise (4.b) can be plausibly denied and with it the conclusion of Hume's skeptical argument.

V. Conclusion

I have covered a lot of ground in this chapter. I have (a) articulated and defended an internalist definition of inference according to which to infer requires having a conscious perspective on the relation between premises and conclusion, (b) outlined a general account of epistemic justification as a function of two aspects, epistemic responsibility and adequate grounding, (c) developed an account of both inferential and noninferential justification corresponding to the results of (a) and (b), and (d) used these

results to defuse Hume's skeptical problem of induction and thereby provide a negative justification of inductive.

Our task is not yet over, however. A negative justification of induction only shows that induction has *not* been shown to be unjustified. This only tells us that it is *possible*, for all we know, that some inductive inference results in an epistemically justified belief. What we want is something more, a positive justification, an account of inductive inference and of inferential justification that makes it reasonable to believe that inductive inference *reliably* leads to justified beliefs. In other words, we want to show that it is in accord with epistemic responsibility, i.e., not a violation of any epistemic duties, to rely on induction as a source for epistemically justified beliefs. The account of inferential justification that I have given in this chapter goes part of the way toward that goal, but it is not sufficient in two respects. First, as a principle of inferential justification, RPIJ is too generic. What we want is an account of inferential justification that is specific to *induction*, or even more specifically to the various *types* of induction. Second, what does it mean for adequate grounds to be, as RPIJ states, *responsibly obtained*? There are two related issues here that have not yet been addressed. One is control. Since responsibility is limited by control, we need to clarify the type and extent of control that pertains to inference. The other is purpose. The exercise of control is responsible when it is properly directed to a proper end. So what is the end toward which inferentially relevant control is to be exercised? What is the purpose of inference? In the following chapter I tackle both of these questions. As the purposeful exercise of control is a *method*, the result that we are working toward is a positive account of inductive justification that has a strong methodological component.

CHAPTER FOUR

INFERENCE AND INQUIRY

Two issues were left hanging at the end of chapter three: whether and in what sense inference is subject to control, and whether there is a general purpose to inference, and if so, what that might be. The goal of this chapter is to deal with these issues and to ascertain what bearing, if any, they have on the nature of inference. In section one I argue that we have indirect voluntary control over our inferences and the resulting beliefs. But having voluntary control raises the question of the ends in the interests of which that control is exercised, which brings us to the question of purpose. In section two, after giving a general description of the human epistemic condition, I argue that the purpose of inference is to overcome in some particular way the current limitations of that condition. More specifically, the purpose of inference is to solve problems—epistemic problems—so as to enlarge the scope and enhance the security of our knowledge. Epistemic problems can be expressed as questions. Since the asking of questions with a view to obtaining an answer is inquiry, and since the purpose of inference is to answer questions, inference is essentially related to inquiry. In section three I argue that, while there are many different types of problems, there is nevertheless a pervasive general pattern to the problem-solving process called inquiry. To show this and to ascertain what that pattern is, I present independent analyses of the process from several different authors. Comparing and combining these results into a general model of inquiry reveals a correspondence between three different stages of inquiry, three different types of question, and three different types of inference. I argue that each type of inference plays a distinctive role in the process of inquiry and, in the course of performing that role, answers a particular type

of question. Finally, in the fourth and last section, I clarify the nature of each type of inference by analyzing the types of question that each is geared to answer.

I. Inference Under Control

As noted in chapter three, Peirce thinks the main difference between inferences and those pseudo-inferential cognitive events that he calls 'associational suggestions of belief' lies in whether we have a conscious perspective on the relation between premises and conclusion—below a certain threshold of consciousness all we have are quasi- and pseudo-inferences, not inferences proper. As to where that threshold lies, Peirce thinks that it has a lot to do with controllability.

> [A]bductive inference shades into perceptual judgment without any sharp line of demarcation between them; or, in other words, our first premisses [*sic*], the perceptual judgments, are to be regarded as an extreme case of abductive inferences, from which they differ in being absolutely beyond criticism. . . . [T]he perceptive judgment is the result of a process, although of a process not sufficiently conscious to be controlled, or, to state it more truly, not controllable and therefore not fully conscious. If we were to subject this subconscious process to logical analysis, we should find that it terminated in what that analysis would represent as an abductive inference.[1]

Unlike perceptual judgments, says Peirce, inference is the result of a process that is "sufficiently conscious to be controlled." Elsewhere, he defines inference as "the conscious and controlled adoption of a belief as a consequence of other knowledge" and goes on to describe it as an active deliberative process involving the three essential steps of colligation, observation, and judgment.[2]

A couple problems arise for Peirce at this point. The first has to do with internal consistency: He describes inference both as a deliberative *process* having distinct steps

---

[1] Peirce, CP 5.181.
[2] Peirce, EP 2.22-24.

and he identifies it with the *result* of that process (viz., "the conscious and controlled adoption of a belief . . ."). The second has to do with psychological plausibility. Controlled *adoption* of a belief seems to imply doxastic voluntarism, the thesis that belief is sometimes under direct voluntary control. But can we form a belief simply be deciding to do so? That seems dubious. These problems are connected in that failure to distinguish clearly between the deliberative process and the result of that process makes the idea that we can sometimes directly adopt a belief through inference seem much more plausible than it is. For while it is noncontroversial that we have direct voluntary control over the deliberative process—we can choose where to look for evidence, whether to check our reasoning, and so on—if we conflate process and result, then it will seem that we have direct voluntary control over the result. That, I submit, would be a mistake.

Direct voluntary control over the deliberative process only amounts to *indirect* voluntary control over the resulting beliefs. This control is indirect because we exercise it by doing *something else*, e.g., consulting a particular authority, looking in this or that direction, focusing attention on skeptical possibilities, etc. I concur with Alston and with Audi that, once the proper distinctions are in place, there is no good reason for thinking that our beliefs themselves are ever subject to our direct voluntary control.[3] Ordinarily, when we notice that premises we are committed to imply or strongly support some conclusion, we simply come to believe the conclusion as a result, without in any sense choosing to believe it. In some cases, such as where a conclusion is highly counterintuitive or where we have a strong desire to disbelieve it, we may be able to

[3] For criticism of doxastic voluntarism see Alston 1989, pp. 115-152; and Audi 2001, pp. 93-111.

withhold assent, but, as Audi points out, that is only a negative form of control and does not imply that we have the positive ability to form or adopt beliefs at will:

> There may be cases in which *withholding* assent suffices to prevent belief formation, and these are significant for such Cartesian purposes as preventing the formation of false beliefs. But . . . [that] point lends no support to . . . voluntarism. There may be times, as where one is deliberating about evidence, when one is in a withholding mode and will not believe the proposition in question unless one does assent to it; here, assenting is necessary for belief formation. But it does not follow (and seems false) that its occurrence entails belief formation. Opening a gate need not bring anything in; and if it does, the entry is not accomplished at will but as a result of the forces that move the entrant. . . . We perhaps picture ourselves as agents of belief formation when what we have really done is to create (or enter) circumstances in which it occurs as a nonvoluntary response to a pattern of evidence. We have accepted *p*, but not because we assented to it or performed an act of acceptance. Rather, the pattern of evidence produced the belief; the belief is more like a response to external grounds than a result of an internal volitive thrust. The belief formation is like the conviction that underlies a competent judge's verdict; it is not like the giving of the verdict: The latter is volitional, the conviction is prevolitional.[4]

Thus, there seems to be no good reason for thinking that beliefs are ever subject to direct voluntary control. Moreover, as Michael Williams points out, we wouldn't want believing to be directly subject to the will: If we are concerned for truth, then we will want our beliefs to be constrained by evidence, to be determined by the state of the world, not by our whim.[5] In conclusion, if we do not have direct voluntary control over our beliefs then inference, which leads to the formation of a belief, does not involve that kind of control. We do have direct voluntary control over the deliberative process, however, and since, as I shall argue in the next two sections, inferences are *essentially* embedded within such a process, we have indirect voluntary control over them and over the beliefs

---

[4] Audi 2001, pp. 97-98.
[5] Williams 2001, p. 26.

they give rise to. But given that we do have indirect voluntary control over at least some inferences, toward what end (or ends) do we typically exercise that control?

II. The Purpose of Inference

Why do we infer? More precisely, why would we care to exercise whatever indirect voluntary control we have over the belief-forming cognitive events called inference so as to ensure that our conclusions are well supported by our premises? The answer to this turns on two general facts about our epistemic situation.

First, our epistemic situation is far from perfect. In his *Enquiry Concerning Human Understanding* Hume writes, "so narrow are the bounds of human understanding, that little satisfaction can be hoped for in this particular, either from the extent or security of [our] acquisitions."[6] He calls attention here to two ways in which our epistemic situation falls short of perfection. First, the *extent* of our knowledge, while in some respects considerable, is still very limited. Compared to what we do know, there is vastly more that we do not know, and perhaps much that, given our limitations, we cannot know. Furthermore, much of what there is to know is not directly accessible to us—we can't just take a good look and thereby come to know how things stand. This is the case with natural laws, the distant past, things too small to be seen (e.g., electrons), and so on. Thus, unless there is some indirect (i.e., inferential) means of coming to know them, the extent of our knowledge will in such matters remain forever impoverished. Second, the *security* of our knowledge rarely, if ever, escapes altogether the possibility of error. Even our most confidently held beliefs, with the possible exception of certain types of introspective reports (e.g., "I am in pain") and clear, self-evident *a priori* truths (e.g.,

---

[6] Hume 1975, p. 8.

2+2=4) may yet prove false. That human fallibility is a pervasive fact the history of science and individual experience readily attest. Furthermore, we often have no direct way to check or verify how things stand given our perceptual limitations. Thus, unless there is some indirect (i.e., inferential) means of checking the truth of knowledge claims, much of our knowledge will remain forever unsecurable. A third type of epistemic imperfection that Hume does not mention concerns the lack of *permanence* of our knowledge. Genuine knowledge once obtained can be lost as memories fade and records are destroyed. We often have to relearn what we once knew. This third type of epistemic imperfection is perhaps reducible to the other two. It is related to the first in that the loss of knowledge leads to a decrease in the extent of our knowledge. It is related to the second in that the loss of knowledge may undermine the evidential security of beliefs that rest upon the knowledge that was lost. In these ways our epistemic situation falls short of perfection and leaves us considerable room for improvement.

The second general truth about our epistemic situation is that we want to improve it. We desire to know, to understand, and, as Aristotle observed, we do so by nature.[7] One reason for this is sheer curiosity or wonder. It gives us great delight to figure things out, make discoveries, and enlarge our understanding of the world. Another, more practically urgent reason is that we live in a changing and often dangerous world. If we do not make an effort to understand it so as to be able to predict events and thereby in some measure control or navigate around them, both the duration and quality of our lives will be seriously diminished. Here's an analogy. A set of beliefs limited in extent, security, and permanence is like a road map. The map is limited in extent both because of the region it

---

[7] Aristotle, *Metaphysics* A.1.

restricts itself to and because of its level of detail. It is limited in security because of the inevitable inaccuracies—inevitable because even the best map offers a static representation of a continually changing world, so that eventually the map and the world it represents will be out of sync. It is limited in permanence because occasionally maps get torn, lose pieces, acquire coffee stains, etc. The worse the map, the less useful and the more misleading it is likely to be. A bad map can make it very difficult to navigate the world successfully and safely. Conversely, the better a map is, the more reliably it will enable you to get to where you want to go and to do so without unnecessary backtracking, detours, risks, etc. Hence, there is good reason to prefer maps that are more informative and accurate. Similarly, we have good reason to desire improvement in both the extent and security of our knowledge.

These two facts—recognition of some specific deficiency in either the extent or security of our knowledge and a desire to eliminate or at least diminish that deficiency—give rise to inquiry. When we desire to know something, we go on a quest for knowledge, we ask questions, that is to say, we inquire or engage in inquiry. Inquiry is the search for answers to our questions. If our epistemic situation was not lacking in extent or security, inquiry would be unnecessary. If we didn't want to improve our epistemic situation, then we wouldn't try to find answers to our questions and inquiry wouldn't happen. The same is true of inference. If our epistemic situation were already perfect, then we would have no need of inference for we would be like God, already knowing everything that we could possibly hope to learn from inference. If we didn't want to improve our epistemic situation, then we wouldn't try to extend our knowledge by deriving new beliefs and we wouldn't care to secure our ideas against error by examining evidence.

Both inquiry and inference arise in response to the limitations of our epistemic situation and are part of our endeavor to overcome, in some measure, those limitations. What, then is the connection between them? We have already said that inquiry is the search for answers to questions. Questions, however, have presuppositions.[8] That is, every question starts from a set of propositions that is either believed or supposed. Furthermore, inquiry ends when the question is satisfactorily answered. A satisfactory answer—satisfactory, that is, from the point of view of the inquirer—is one that the inquirer is prepared to accept as true, i.e., believe. Inquiry, then, begins from something believed or supposed and ends in a state of belief, just like inference. Inquiry and inference are, therefore, correlative. As inquiry is the search for answers to our questions, *inference is arriving at answers to our questions on the basis of the information in hand.* In other words, to inquire is to ask a question with a view to obtaining a satisfactory answer; to infer is to arrive at an answer to a question by reasoning from the information in hand. The conclusion of every inference is, therefore, the answer to a question.

With respect to the claim that all inferences answer questions, a couple points of clarification require mention. First, the converse does not hold. That is, questions may be answered by means other than inference. One common means is perception, another is consulting a trusted authority. For example, wanting to know what time it is I glance at my wristwatch, or wanting to know what 'obsequious' means I consult a dictionary. Second, questions answered by inference need not be explicitly formulated to ourselves, nor need we be acutely conscious of the fact that in inferring we are answering a

---

[8] Walton 1989, pp. 54-55. For example, yes-no questions ask which of a list of alternatives is true and presuppose, therefore, that one of the alternatives is true. Why-questions ask for an explanation of some state of affairs that is that is presupposed to obtain. And every genuine question assumes *in the very asking of it* that an adequate answer can in principle be had.

question. The extent to which a question is explicitly or consciously entertained is a matter of degree because consciousness is a matter of degree. The more aware we are of what question we are seeking to answer, the greater the amount of control we can exercise in the answering, the better the perspective we will have on the adequacy of any proposed answer, and thus the more inferential the answering process will be. But if we are not typically consciously aware of answering a question when inferring the reason is because it is of the nature of a question to direct attention away from itself and toward the answer sought. The only time the question itself comes into focus is when we pause to question the question, i.e., when we reflect on the nature of the problem we are trying to solve.

What has emerged so far is that our epistemic situation leaves us confronted will all sorts of epistemic problems. Desire to overcome these results in a problem-solving process of inquiry, the asking of questions with a view toward finding answers. Inference is related to inquiry as answering is to questioning; it is the cognitive event by which we arrive at an answer to a question on the basis of the information in hand. To go further in our analysis we need to look more closely at this problem-solving process called inquiry. In the following section I develop a general account of the structure of inquiry. Antecedently one might expect that such a general account is not to be had. After all, there seem to be so many different types of problems and different kinds of questions that one might wonder whether any theoretical unity can be found. Despite initial appearances to the contrary, I argue that inquiry as a problem-solving process follows a pervasive general pattern; it has a characteristic structure. Before embarking on this enterprise to find a general pattern to inquiry let me emphasize that I am concerned with *epistemic*

problems, problems concerning the extent and security of knowledge. Consequently, my account must be somewhat idealized. I do not claim that all actual inquiries are or should be conducted as I propose. We have other aims—*nonepistemic* ones—that compete with the aims of inquiry for our limited resources. For practical purposes it is often necessary to cut short a process of inquiry before an unambiguous answer is indicated by the evidence. The examples of inquiry that matter most for my purposes, therefore, are those in which epistemic aims dominate over other considerations and are most deliberately and consciously pursued, as is perhaps most frequently the case in theoretical disciplines like the sciences, math, and philosophy.

III. The Structure of Inquiry

John Dewey comments that "inquiry, in spite of the diverse subjects to which it applies, and the consequent diversity of its special techniques has a common structure or pattern."[9] A cursory glance at the literature on problem-solving reveals such a pattern. For example, in his book *How to Solve It*, mathematician G. Polya describes problem-solving as a four-stage process:[10]

1. Understanding the Problem
2. Devising a Plan
3. Carrying Out the Plan
4. Looking Back [i.e., checking the solution for adequacy]

Compare this to cognitive psychologist John R. Hayes' analysis in *The Complete Problem Solver*:[11]

---

[9] Dewey 1991, p. 105.
[10] Polya 1957, pp. xvi-xvii.
[11] Hayes 1981, pp. 2-3.

1. Finding the Problem: recognizing that there is a problem to be solved.
2. Representing the Problem: understanding the nature of the gap to be crossed.
3. Planning the Solution: choosing a method for crossing the gap.
4. Carrying Out the Plan
5. Evaluating the Solution: asking "How good is the result?" once the plan is carried out.
6. Consolidating Gains: learning from the experience of solving.

Hayes' stages 2 through 5 correspond directly to Polya's four stages. His stages 1 and 6 define the start and end of the process more clearly. It is recognized by both Polya and Hayes, of course, that the process very rarely occurs in simple 1-2-3 fashion. One reason is that later stages often yield new insights that lead us to reexamine the results from earlier stages. For example, in the course of carrying out a plan I may encounter unexpected complications that cause me to reevaluate my solution strategy. Another reason is that complex problems often consist of numerous subproblems, requiring a recursive application of the process.

Parallel analyses of problem-solving can easily be found,[12] but I want to focus on two in particular: those given by Bernard Lonergan and Charles Peirce, respectively. Their analyses corroborate those of Polya and Hayes but go further in important ways. Lonergan's analysis is helpful because it points out how different types of questions dominate at different stages of inquiry, and Peirce's analysis is helpful because it shows how different types of inference relate to the process of inquiry.

Lonergan's Theory of Inquiry

According to Lonergan, inquiry proceeds through three main stages—experience, understanding, and judgment—the transition between which is effected by a cognitive

[12] For example, Popper 1979, p. 243; Dewey 1991, chs. 6 and 7; and Polanyi 1962, pp. 120-121.

event that he calls an 'insight.' Experience includes all that is present to our awareness apart from any interpretation or understanding of it. When confronted with the data of experience, we spontaneously desire to know what it is an experience *of*, we want to *understand* it. Hence we are prompted to inquiry, to ask what Lonergan calls 'questions for intelligence', which he typifies with the Latin *quid est?*—What is it?

In response to questions for intelligence, a direct insight[13] may occur. That is, one may get a grasp on the intelligibility immanent in one's experience. Insight occurs at the "ah-ha!" moment when "the light goes on" or one "puts two and two together" and, to borrow a phrase from William James, the "blooming, buzzing confusion" of experience resolves into recognizable objects and intelligible patterns. It involves an act of abstraction that takes us from concrete experience to an understanding or conceptualization of experience.[14] As insights accumulate they are combined and systematized with other insights, enlarging our understanding. Thus in reading there are acts of insight involved in recognizing the individual words, grasping the significance of the words, of their arrangement, of the larger context of the sentence, and so on, such that our understanding of the text is continually enlarged and refined as new insights are added.

Understanding, however, is not enough. We want, says Lonergan, not merely a possible or plausible understanding of experience; we want a *true* understanding. Thus, there now arises a 'question for reflection' (*an sit?*, Is it so?). The answer takes the form of a judgment: "Yes, it is" or "No, it is not." Like the transition from experience to

---

[13] Lonergan 1958, pp. 3-6. Lonergan also discusses "inverse insight" (pp. 19-25), i.e., a grasp of the *absence* of some expected intelligibility.
[14] Lonergan 1958, p. 88.

understanding, the transition from understanding to judgment occurs through an insight—in this case a 'reflective insight,' a grasp that the evidence is sufficient to warrant affirmation of the prospective judgment.[15]

Peirce's Theory of Inquiry

Peirce approached inquiry from two complementary perspectives: one psychological, the other, logical. According to the first, inquiry is the mental struggle to eliminate the irritation of doubt and attain the 'fixation of belief.'[16] According to the second, inquiry progresses toward a knowledge of reality through an inferential triad of abduction, deduction, and induction. It is this second account that we want to look at.

Peirce identifies three distinct modes of inference—abduction, deduction, and induction.[17] Their formal structures may be roughly illustrated as follows:[18]

Abduction:	Deduction:	Induction:
E	H	$H \rightarrow C$
$\underline{H \rightarrow E}$	$\underline{H \rightarrow C}$	$\underline{C}$
H	C	H

But, for Peirce, the form of an inference is less important than its function in the process of inquiry. Abduction, deduction, and induction each play a specific role in the larger

---

[15] Lonergan 1958, p. 287.

[16] For Peirce's psychological account, see his famous 1877 essay "The Fixation of Belief" in W 3.242-257.

[17] A brief terminological note: I have been using the term 'induction' to refer to non-deductive inference generally. For Peirce, 'induction' refers to one kind of non-deductive inference, and 'abduction' another. For his mature descriptions of these types of inference, see his 1908 essay "A Neglected Argument for the Existence of God" in EP 2, esp. pp. 440-445.

[18] Two general comments on these diagrams: First, the deductive pattern is logically valid as it should be (here it is an instance of *modus ponens*). The abductive and inductive patterns are deductively invalid (here they are instances of affirming the consequent). Again, this is as it should be—because they are non-deductive modes of inference, the conclusion goes beyond what is given in the premises, and that guarantees their invalidity. Second, as represented here, abduction and induction appear only superficially distinct, the only difference being the order of the premises. Bear in mind that, for Peirce, the differences between types of inference are primarily functional, and only secondarily formal. While the difference in the order of abductive and inductive premises, as depicted here, is not formally significant, it is functionally significant, reflecting the difference between *accommodating* old data and *predicting* new data.

process of inquiry and apart from that role have no legitimate application.[19] Abduction begins with a problem situation or explanandum (E), reasons backward to a hypothesis (H) that, if true, would explain or make sense out of the situation (H→E). Deduction takes up the proposed explanation (H) and traces out additional consequences (C) that would follow if the hypothesis were true (H→C). Induction takes up the deductively explicated hypothesis (H→C), and evaluates whether it is true or false by checking whether the expected results (C) obtain. In short, abduction explains, deduction explicates, and induction evaluates.[20]

Toward a Unified Account of Inquiry

Comparison of Lonergan's and Peirce's theories of inquiry reveals significant parallels. Peirce's abduction, which yields plausible explanations for puzzling phenomena, corresponds to Lonergan's direct insight, which takes us from experience to an understanding of experience. Peirce's induction, which yields an assessment of the truth or falsity of a hypothesis, corresponds to Lonergan's reflective insight, which results in a judgment on the truth or falsity of a possible understanding. Now this correspondence is not exact, for an insight is not necessarily inferential,[21] but it is suggestive of a hypothesis. If Lonergan's direct and reflective insights answer questions for intelligence and questions for reflection, respectively, and if Peircean abduction and induction correspond to Lonergan's direct and reflective insights, respectively, then abduction and induction answer different types of questions. And since inquiry is the

---

[19] Peirce, CP 1.65. Cf. Levi 1997, pp. 34–56.
[20] Peirce, CP 5.171.
[21] For Lonergan, "insight" refers to a type of cognitive event that occurs in response to the basic interrogative intentionality of consciousness. While all inferences involve an insight in the grasp of the relation between premises and conclusion, not all insights are inferences because insights don't require a *conscious* appreciation of one's grounds.

search for answers to our questions and to infer is to arrive at an answer to a question by reasoning from the information in hand, this suggests that answering different kinds of inference answer different kinds of question and that there are as many basic kinds of inference as there are basic kinds of questions asked in the context of inquiry. Whether this is so remains to be seen. To verify it, we will have to (1) examine the process of inquiry to identify the fundamental kinds of questions that occur in that context, and (2) establish a one-to-one correspondence between types of question and types of inference, showing how each type of inference answers its related question-type.[22] Let us see if we can carry out this program.

If a given episode of inquiry begins with a problem situation resulting from a particular deficiency in our knowledge, the end of that episode will come when the problem is solved and the deficiency rectified. What begins with uncertainty and ignorance ends with the confident judgment "S *is* P." But what happens in between? We can answer that question by reflecting on what is presupposed in making a concrete judgment of fact.[23] As to inquire is to ask a question, the judgment that terminates an episode of inquiry will be the answer to a question. If that answer is reasonable it will be based on sufficient evidence grasped as sufficient. In other words, the judgment "S is P" will be an answer to the critical question "*Is* S P?". But if one had possessed sufficient evidence at the outset and grasped it as sufficient, inquiry would have been unnecessary. Since we are supposing that the judgment was reached via a process of inquiry, it must be the case that at some *earlier* point in the process there was a lack of sufficient evidence

---

[22] One might immediately wonder whether this is plausible, for Peirce identifies three types of inference, and Lonergan only two types of question. What type of question does deduction answer? On the other hand, deduction, which draws out the consequences of ideas, corresponds rather nicely to Lonergan's idea that at the level of understanding a progressive elaboration and development of insights occurs.

[23] My argument here is inspired by chapter 12 of Lonergan 1958.

recognized as such. Hence, prior to the judgment, further evidence must have been acquired or recognized. But in order to know where to look for evidence and ascertain its relevance, one must have *criteria* for determining whether S is P. How can we identify appropriate criteria? Peirce's pragmatic maxim tells us that if we want to better understand what a concept means, we should ask ourselves "What would we expect to be true *if* this concept applies?" So if we want to know whether S is P we should ask the conditional question "*If* S is P, what follows?" What would we expect to be true, if S is P? But one cannot meaningfully ask and answer that question without some prior understanding of both S and P. That is to say, one must already possess informative answers to the questions "what is S?" and "what is it to be P?" If one's answers to those questions are to be more than stipulations, they require an initial body of data to draw from, whether that be the given of experience or the given of already possessed background beliefs. Thus we come to the starting point of inquiry in a given or presupposed body of data. All inquiry into concrete matters of fact must take something as given, at least provisionally.

Reversing the order of this analysis, we see that the process of inquiry involves six distinct stages. First is the stage of *epistemic dissatisfaction*, where, against the backdrop of experience and background beliefs, some item stands out as problematic. Second is the *explanatory* stage. One struggles to make sense of the problem situation by asking questions like "what is S?" or "why *p* rather than *q*?" I call these *interpretive* questions because their aim is to "make sense" out of the problem situation. When such questions have been answered, one has reached the third or *explicative* stage. Having grasped a possible solution to the problem, one faces the critical task of ascertaining

whether the solution really works. To be tested, a vague conception needs to be explicated and made more precise, its consequences delineated so they can be checked. The struggle to clarify one's ideas involves asking *conditional* questions like "if S were P, what would follow?" Having answered that question one can proceed to the fourth or *experimental* stage of gathering further data and checking whether the predicted consequences of S's being P hold. Fifth is the *evaluative* stage of judging whether S really is P in light of the experimental results. This involves answering the *categorical*, yes/no question "Is S really P?" Sixth and finally, the result of this process is a stage of (relative) *epistemic satisfaction*, provided that a positive evaluation is forthcoming at stage 5. If not, then we are driven back to the drawing board, so to speak, and have to come up with another explanation (stage 2) or, perhaps, conclude that there just is no systematic explanation. In summary, the basic pattern of inquiry appears to be this:

1. Epistemic dissatisfaction (problem situation)
2. Explanation—come up with possible solution. (answers interpretive question)
3. Explication—derive testable predictions. (answers conditional question)
4. Experimentation—gather further information, run tests.
5. Evaluation—assess truth/falsity in light of new data. (answers categorical question)
6. Epistemic satisfaction (if a positive evaluation at stage 5; otherwise, back to stage 2)

I only claim that this is only the *basic* pattern of inquiry. Actual inquiries may be very extensive, and include many, many sub-inquiries and sub-sub-inquiries. The process at any level can be cut short or interrupted as new and more pressing questions rise to the forefront of our attention. Sometimes we skip steps. For example, if only one plausible explanation comes to mind in stage 2 we may have enough confidence to believe the explanation without needing to engage in any experimentation. Sometimes the process

may be run through subconsciously,[24] but as inquiry becomes more self-conscious and deliberate, the different stages become more distinguishable until, in the conduct of the sciences for example, we can see inquiry "writ large," as it were.[25] In such cases, I submit, we will be able to discern this pattern of stages.

In summary, inquiry is a problem-solving process that has a recognizable structure consisting of several distinct stages. Corresponding to three of these stages—the explanatory, explicative, and evaluative stages—are three types of inference (abduction, deduction, and induction), which answer three different types of question (interpretive, conditional, and categorical). The relation between the last two triads (types of inference and types of question) will be developed more in the next section, but the following table summarizes these results:

Stage of Inquiry	Explanation	Explication	Evaluation
Type of Inference	Abduction	Deduction	Induction
Type of Question Answered	Interpretive	Conditional	Categorical

To emphasize the role that the different types of inference play in a process of inquiry, in succeeding chapters I will generally refer to them in terms of that role. Thus, I will refer to abduction, deduction, and induction as explanatory, explicative, and evaluative inference, respectively. At times, I will revert to Peirce's terminology or to the traditional deduction/induction distinction, such as when quoting from or discussing sources that use those classifications. Context should make my usage clear, but for the reader's benefit, here is a translation table:

---

[24] Polanyi (1962, pp. 120-124) cites W. Köhler's research on the problem-solving abilities of chipanzees and shows how it fits the broad pattern I have outlined here.

[25] For example, The pattern is clearly evident in the work of Semmelweis, who from 1844-1848 conducted research that led to the discovery that "cadaveric matter" was responsible for the childbed fever that regularly took the lives of woman in a Viennese hospital maternity ward. Semmelweis's research is frequently referred to by philosophers of science (e.g., Hempel 1966, pp. 3-6; Lipton 1991, ch. 5) as a paradigmatic instance of scientific inquiry.

Traditional Classification	Peircean Classification	My Classification
Deduction	Deduction	Explication
Induction	Abduction	Explanation
	Induction	Evaluation

IV. On the Different Types of Questions

> The path of all knowledge leads through the question. . . . All questioning and desire to know presuppose a knowledge that one does not know; so much so, indeed, that a particular lack of knowledge leads to a particular question.
>
> —Gadamer, *Truth and Method*[26]

> Every step in [an] argument depends on asking a question. The question is the . . . motive force.
>
> —Collingwood, *The Idea of History*[27]

In the account of inquiry presented so far, I have identified three types of questions (interpretive, conditional, and categorical), which correlate with three types of inference (abduction, deduction, induction), and with three stages in the process of inquiry (explanation, explication, evaluation). In this section I clarify the three types of questions and show that the kind of intentionality represented by each is distinct. In other words, they expect distinct kinds of answer. Since it is the purpose of inferences to answer questions, this result supports the Peircean thesis that there are three distinct types of inference and sheds some light on the nature of each. I begin with a preliminary description of each type of question, and then refine these descriptions by way of a couple important distinctions. Lastly, I present a logical diagram of each type of inference, showing how they serve to answer the three types of questions.

[26] Gadamer 1988, pp. 363, 365-366.
[27] Collingwood 1993, p. 273.

A Preliminary Description

Interpretive questions typically begin with one of the standard question-words: why? what? how? etc. These are the sorts of questions we ask when we wish to better understand something. An unexpected event occurs. Why? I am served an unfamiliar dish. What is it? I see a suspicious person across the way. Who might that be? I get lost. Where am I? A magician makes a rabbit appear. How did he do that? I call them 'interpretive' because they emerge from our trying to *make sense* of things, as in interpreting a text one tries to make sense of the words.

Conditional questions, as the name implies, have an if–then structure. We ask these when we want to explore the consequences of some state-of-affairs' obtaining or of something's being true: If I were to mix these two chemicals, what would happen? What might world history have looked like if Hitler had won? In answering such questions, deductive inference[28] explicates, spells out the significance of, the premises taken in conjunction with our existing background beliefs.

Categorical questions expect a yes/no, true/false kind of answer and usually begin with some form of the verb *to be*. Whereas interpretive questions leave the range of possible answers open-ended, categorical questions begin with a more-or-less well-defined set of possible answers and seek to restrict it further. Is π greater than, less than, or equal to 3.1416? Is abortion murder (or not)?

---

[28] Why *deductive* inference? A conditional question asks us to start from given assumptions and then trace out consequences that would obtain if those assumptions were true. Insofar as we are tracing out the consequences of these assumptions we cannot go beyond them. Thus, the only kind of inference that can properly answer a conditional question has to be non-ampliative. That excludes all forms of induction.

The Descriptions Refined

Wittgenstein distinguishes between ‘surface grammar’ and ‘depth grammar.’ The accurate analysis of language, he argues, requires that we look beyond or behind the surface grammatical form of an utterance and appreciate the use or function of the utterance within the larger conversational context.[29] This parallels my claim that the function of an inference in the context of inquiry is more important than its logical form. In the present context this distinction is important because not all utterances that look like questions really are questions. So-called rhetorical questions are a case in point. Rhetorical questions are linguistic expressions that look like questions but aren’t because they function like statements, commands, or exclamations. For example, suppose a boorish acquaintance has been trying my patience so much that I finally say, “Why don’t you go home?” This looks like an interpretive question, but I am not really interested in receiving an explanation as to why my acquaintance doesn’t leave; rather, I am issuing a request: Please leave me alone. Furthermore, sometimes utterances that look non-interrogative may in fact be questions. For example, depending on intonation, “You are going” may either make a statement of fact or ask the question “Are you going?” While surface grammar may help us distinguish questions from non-questions, whether an utterance is a question or not ultimately depends on its function in a context and not on the form.

Similar comments apply to the problem of distinguishing between types of questions. My preliminary description of interpretive, conditional, and categorical questions was in terms of surface grammar. One might easily get the impression from

---

[29] Cf. Wittgenstein 1958, §664.

that account that any question that is conditional in form is therefore a conditional question, or that any question that is not conditional in form and begins with a standard question word like ‘what’ or ‘why’ is therefore an interpretive question. Such is not the case, and counterexamples are easy to come by. Here’s a couple: “What flavor is it—chocolate or vanilla?” and “Why are you here—to work or to play?” These start out looking like interpretive questions but are in fact equivalent to the categorical questions “Is the flavor chocolate or vanilla” and “Are you here to work or to play?” And some interpretive and categorical questions are conditional in form: If something is cold, white, and made of 6-pointed crystals, what is it? (an interpretive question) If the litmus paper turns red, is it an acid? (a categorical question).

What matters, therefore, is not the verbal form of the utterance, but its *intentional quality*, or the *type of answer* that is being sought. In this regard, John Bruin introduces a useful distinction between ‘hermeneutical’ and ‘predicative’ questions:

> [T]he predicative question starts with the “name” and then proceeds to determine the “predicate.” Typically, the possibilities open to a predicative question are clear: “What kind of flower is it—a rose or a poppy?” “What color is it—red or blue?” . . . The hermeneutical question, by contrast, proceeds the other way about. Starting with the predicate-clues (*—is x, —is y, —is z*), it then looks for the “name,” or the answer, to the “riddle” “What is it?” (*? is x, y, and z*).[30]

What I call interpretive and categorical questions correspond to Bruin’s ‘hermeneutical’ and ‘predicative’ questions, respectively. When we ask an interpretive question, we seek to make sense of a puzzling situation or solve a riddle, so to speak, by studying the available clues. We look for a suitable ‘subject’ to which the available predicate-clues can be applied. The direction of thought is retrospective, from effect to cause, from what

[30] Bruin 2001, pp. 62-63.

is conditioned to that which conditions. Interpretive or hermeneutical questions are asked with the aim of *generating* a range of *possible* solutions: What *might* this mean? How *might* we make sense of this? In contrast, when we ask a categorical question, we seek to render an *already identified* subject more determinate by assigning a predicate—yes or no? true or false? black or white or gray? And so on. The direction of thought is, for lack of a better word, *in*spective, from partially conditioned to further conditioned. Categorical or predicative questions are asked with the aim of *restricting* possibilities, of narrowing a set of pre-delimited possibilities to what is *actually* the case.

Conditional questions, in the sense in which I intend them to be understood, do not fit into Bruin's classification. As noted above, there are questions that *appear* conditional in form that are really interpretive/hermeneutical or categorical/predicative. The relevant criterion, as always, is how the question *functions* in the context of inquiry. Conditional questions posit a set of conditions (if …) and ask for the consequences (then …). Like interpretive questions and unlike categorical questions they are open-ended; "What follows?" has no pre-envisioned range or set of possible answers. Unlike both interpretive and categorical questions, they are not retrospective or inspective but forward-looking or *pro*spective, from condition to conditioned. Furthermore, whereas interpretive questions ask for explanatory *possibilities* and categorical questions ask for which of a set of possibilities is *actual*, conditional question ask for what is *necessary* given the conditions.[31] Thus there is a distinct modal shift between these three types of

[31] Why necessary? Because conditional questions posit a set of conditions (some of these may be unstated background assumptions taken for granted in the context in which the question is uttered) and ask for a set of results consequent upon those conditions, the consequences cannot go beyond the conditions. If they did, then they would not truly be consequences of those conditions. But if the consequences do not go beyond the conditions then they must be entailed by them.

questions. The following table summarizes the distinctive intentionality of each type of question:

Question Type	Standard Examples	Intentional Diagram	Temporality	Modality
Interpretive	Why? What is it?	If ? then X	past - retrospective	possibility
Conditional	If so, then what?	If X then ?	future - prospective	necessity
Categorical	Is it?	X? or Y?	present - inspective	actuality

Questioning and Inferring

We ask questions in hopes of finding answers. An inference is a cognitive event whereby we arrive at an answer to a question. As different types of questions look for different types of answers, so also there are different inferential processes by which we arrive at these answers. The following logical diagrams show how each type of inference provides an answer to its corresponding question type:

Explanatory:	E	Explicative:	H	Evaluative:	Is H true?
	Why E?		If H, what follows?		H→C
	H→E		H→C		C
	H		C		H

In an explanatory inference, some explanandum E is the occasion for an interpretive question, which expresses the desire to understand the situation. We reason backward from explanandum to a putative explanans, a hypothesis H. In an explicative inference, we start with some posit H and ask the conditional question: If H were true what consequences C would follow? The result is a prediction of those consequences. In an evaluative inference, we examine criteria C that are conditional predictions of a hypothesis to ascertain whether that hypothesis is true. In each case, the conclusion gives an answer to the question: Why E? H. If H, what follows? C. Is H true? Yes. H.

V. Retrospect and Prospect

What is the purpose of inference? To answer questions. To improve our epistemic situation with respect to both extent and security. Before questions can be answered, however, they must be asked. Thus it is inquiry, the asking of questions, that establishes the context within which inferences take place. Analysis of inquiry revealed that it is a multi-stage process in which there is a correlation between three stages of inquiry, three kinds of inference, and three kinds of question. The different types of inference are distinguished primarily by their role in a process of inquiry.

The bearing of these results on a positive justification of induction is twofold. First, because each inference plays a definite role a process of inquiry, the justification of belief by any particular inference is in part of function of whether and how well the inference fulfills its role, of whether and how well it answers the question that motivated it. Second, because there are two fundamentally different types of nondeductive inference, one explanatory in function and the other evaluative, the positive problem of induction divides into the problems of justifying these two types of inference, of showing that our confidence in the reliability of each as a source of epistemically justified beliefs is well placed. That, of course, necessitates taking a closer look at explanatory and evaluative inference. This I do in the next two chapters.

CHAPTER FIVE

EXPLANATORY INFERENCE

My goal in this chapter and the next is to give a positive justification of explanatory and evaluative inference. In other words, we want to show that it is reasonable, in accord with epistemic responsibility, to trust each as a source of epistemically justified beliefs. In chapters three and four I developed an account of inference and of inferential justification that goes part of the way toward this goal. Thus, an inference results in a justified belief just in case it results in an *inferentially* justified belief, which, according to RPIJ, requires that one's grounds for the belief be adequate and responsibly obtained. Grounds are adequate when the belief is formed and maintained in such a way that it is objectively probable that the belief is true. Responsibility requires that we use whatever control we have over our epistemic resources to secure grounding for our beliefs that is, as far as we can tell, adequate. Since inferences are essentially embedded within a process of inquiry, obtaining grounding in an epistemic responsible way depends upon the proper conduct of inquiry, i.e., it depends upon a sound methodology. So far so good, but the foregoing applies to inference generally and we want a positive justification of explanatory and evaluative inference in particular.

This chapter deals with explanatory inference. In section one I identify three aspects of explanatory inference: hypothesis generation, hypothesis screening, and hypothesis preferencing. The reliability of explanatory inference depends on our ability to carry out these three tasks in a manner that yields adequately grounded beliefs. For a positive justification of explanatory inference, therefore, we need to show that responsible use of our epistemic resources in these areas entails that the resulting beliefs are, as far as we can tell, adequately grounded. In section two I look at hypothesis generation more closely and argue

that explanatory inference can generate adequately grounded beliefs only if what I call the *correspondence thesis* is true. This thesis says that there is a fundamental correspondence between knower and known, between the way we think and the intelligible structure of the world. From this thesis I derive several important epistemological corollaries, and I argue that acceptance of the thesis and its corollaries is epistemically responsible. In section three I look at hypothesis screening and preferencing and clarify the notion of *explanatory plausibility* on which they depend, arguing that there exist criteria that enable us to make epistemically responsible plausibility assessments. Finally, in section four, I draw these results together into a positive justification of explanatory inference.

I. Three Aspects of Explanatory Inference

What makes an inference explanatory is the role that it plays in the process of inquiry in virtue of answering an interpretive question. To understand explanatory inference, therefore, we need to look at what is involved in answering an interpretive question. As pointed out in chapter four, interpretive questions are open-ended and retrospective—in asking them we seek to generate a range of conditions for that which is conditioned. In other words, we seek to identify causes from effects; to identify a working hypothesis or explanation for a given explanandum. The first task of explanatory inference is, therefore, *hypothesis generation*—we need to come up with possible answers to the interpretive question. But not just any answers will do. We want plausible answers, ones that, as far as we can tell, stand a good chance of being correct. Of those hypotheses that are generated, some may be sufficiently plausible to merit serious consideration, and some may not be worth pursuing. Thus, a second task of explanatory inference is *hypothesis screening*—we need to

assess the plausibilities of hypotheses to know which to most seriously consider. Finally, if screening leaves us with several plausible candidate hypotheses, then we need to decide which of the candidate hypotheses to accept as true and/or to pursue as our working theory.[1] Thus, a third task of explanatory inference is *hypothesis preferencing*[2]—we need to compare the plausibilities of rival hypotheses in order to isolate the best (working) explanation. When these three aspects of explanatory inference are taken into account, its formal structure looks like this:

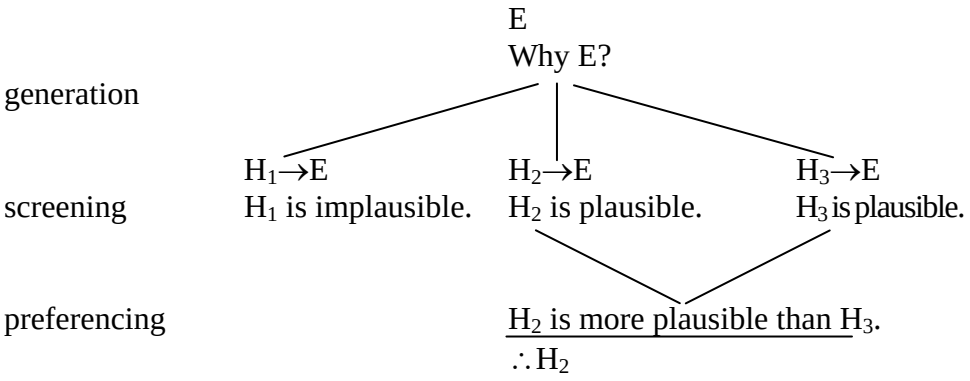

We now have to look at these three aspects of explanatory inference to see whether it is reasonable to believe that they can be carried out in a manner that yields adequately grounded beliefs.

II. Hypothesis Generation and the Correspondence Thesis

Following Hume, Popper and many philosophers of science have thought that hypothesis generation is fundamentally nonrational.

> The initial stage, the act of conceiving or inventing a theory, seems to me neither to call for logical analysis nor to be susceptible of it. The question how it happens that a new idea occurs to a man—whether it is a musical

[1] The acceptance/pursuit distinction comes from Laudan 1977, pp. 108-114. His point is that some theories are worth pursuing, i.e., developing and testing to determine their adequacy, but not (yet) worth accepting as true.
[2] The terms 'hypothesis generation' and 'hypothesis preferencing' are from Delaney 1993, pp. 15-16.

> theme, a dramatic conflict, or a scientific theory—may be of great interest to empirical psychology; but it is irrelevant to the logical analysis of scientific knowledge. . . . [T]here is no such thing as a logical method of having new ideas, or a logical representation of this process.[3]

There is some initial plausibility to this view, for it does often seem that hypothesis generation is just a matter of imagination, creativity, and serendipity. Ideas often seem to come to us "in a flash," as it were. What could there possibly be to logically analyze in *that*?

The central difficulty with this view is known as the problem of the underdetermination of theory by data. Peirce lays the problem out nicely:

> Consider the multitude of theories that might have been suggested. A physicist comes across some new phenomenon in his laboratory. How does he know but the conjunctions of the planets have something to do with it or that it is not perhaps because the dowager empress of China has at that same time a year ago chanced to pronounce some word of mystical power or some invisible jinnee may be present. Think of what trillions of trillions of hypotheses might be made of which one only is true; and yet after two or three or at the very most a dozen guesses, the physicist hits pretty nearly on the correct hypothesis. By chance he would not have been likely to do so in the whole time that has elapsed since the earth was solidified.[4]

Because the number of logically possible answers that can be given to any interpretive question is nondenumerably infinite,[5] without constraints to sharply narrow the search space, the chance of finding the true hypothesis, or even a close approximation to it, would be slim to none. That the human quest for knowledge has apparently gotten off the ground, progressed quite far, and continues to do so at an impressive rate is, therefore, strong evidence that there exist search constraints that guide inquiry along paths that, by and large,

---

[3] Popper 1959, pp. 31-32.
[4] Peirce, CP 5.172. Cf. also CP 5.431: "Truths, on the average, have a greater tendency to get believed than falsities have. Were it otherwise, considering that there are myriads of false hypotheses to account for any given phenomenon, against one sole true one (or if you will have it so, against every true one), the first step towards genuine knowledge must have been next door to a miracle." See also Polanyi 1962, p. 30.
[5] Howson 2000, p. 76.

keep us pointed in the direction of truth rather than error.[6] Given that this is so, it follows that the generation of hypotheses is not random or haphazard but is the result of processes that are significantly truth oriented. If these processes are to any extent subject to control, then to that extent there can be a rational method of generating plausible hypotheses, contrary to Popper.

But even apart from the evidence of past and present epistemic success we must assume, argues Peirce, that truth-oriented search constraints exist, for

> there is . . . a hypothesis which we must embrace at the outset, however destitute of evidentiary support it may be. That hypothesis is that the facts in hand admit of rationalization, and of rationalization by us. That we must hope they do, for the same reason that a general who has to capture a position or see his country ruined, must go on the hypotheses that there is some way in which he can and shall capture it.[7]

Call the hypothesis that "the facts in hand admit of rationalization and of rationalization by us" the *correspondence thesis*, because it posits a fundamental correspondence between knower and known, between the way we think and the intelligible structure of the world.[8] If we do not accept this thesis, then it is impossible to make sense of the practice of inquiry, for to ask a question with a view to obtaining an answer is to suppose that there is a true answer to be had, and had by us. If a true answer *about the world* is to be had then the world must be, at least to that extent, intelligible. And if it is to be had *by us*, then the way we think and

---

[6] In the Appendix I argue that there are in fact at least four such search constraints, four clues that guide hypothesis generation: understanding of the problem, background beliefs and concepts, experience, and correction of prior hypotheses.

[7] Peirce, CP 7.219.

[8] There are two ways to account for this correspondence. One is to suppose that we are the products of intelligent design, that the Author of Nature and of human nature has created an orderly and intelligible world and designed our cognitive faculties so that we could understand it. The other is evolutionary—given that the ability to guess right about the world has real adaptive value, we can suppose that Nature has shaped us so that we can understand her.

the functioning of our cognitive faculties, must be suitably in sync with the intelligible structure of the world. Several important corollaries follow from this thesis.

The first corollary is that our cognitive faculties are generally reliable (hereafter just 'reliable') when deployed in an epistemically responsible manner. The reason lies in the fact that inquiry aims at *knowledge*, not just true belief. Accordingly, inquiry presupposes not merely that it is *possible* for us to arrive at true answers but that it is *likely* given epistemically responsible behavior. This is so because knowledge requires adequately grounded belief, i.e., a belief that is formed and maintained in such a way that it is objectively probably true. If arriving at true answers was unlikely, then it would not be objectively probable that the belief is true, thus it could not be adequately grounded, or amount to knowledge. So a necessary presupposition of knowledge-oriented inquiry is that arriving at true answers is likely given epistemically responsible behavior. But that can only be so if our cognitive faculties are reliable when responsibly deployed.

The second corollary is that beliefs formed and maintained in a responsible way have what Williams calls 'default justification'.[9] For if our cognitive faculties are reliable when responsibly deployed then there is an initial presumption that what we believe in a responsible way is probably true. Thus, the presumption is that such beliefs are already adequately grounded. This presumption can be overcome by relevant defeaters, or if contextual factors raise the standards for grounding beyond what is warranted by the presumption. But such beliefs are "innocent until proven guilty", i.e., unless and until a relevant defeater is on the table or standards-raising contextual factors are in place, there is no epistemic obligation to show that such beliefs are true by citing further evidence. In the

[9] Williams 2001, pp. 148-150.

case of noninferentially formed and maintained beliefs—most of our ordinary perceptual judgments would fall into this category—default justificatory status is common. For in such cases we lack a conscious perspective on the grounding of the beliefs and therefore have much less control over their grounding than we would have had otherwise. Having less control means that less is required of us in the way of epistemic responsibility, so default justificatory status is more easily attained.

The third corollary is that responsibly obtained appearances are reliable guides to reality. For, according to the first corollary, our cognitive faculties are reliable when responsibly deployed. But it is through the operation of those faculties that things appear to us. Indeed, for many of our cognitive faculties (e.g., our perceptual faculties) to be reliable just is to give rise to appearances that are reliable guides to reality. Thus, if responsibly obtained appearances are not reliable then neither are our cognitive faculties. As that is contrary to the first corollary, the antecedent must be false. Furthermore, as argued in chapter three, epistemic responsibility requires only that we seek grounding that is adequate *as far as we can tell*. But responsibly obtained appearances are simply those that, as far as we can tell, *are* reliable guides to reality. For if we conclude anything about reality it is because it *seems* to us that it is so. Thus we can never conclude that some appearances are unreliable without crediting the reliability of other appearances. Hence, it is not irresponsible, i.e., not a violation of epistemic duties, to trust responsibly obtained appearances as reliable guides to reality.

The fourth corollary is that the belief that we are epistemically responsible in deploying our cognitive faculties and in obtaining appearances on any particular occasion has default justification. In other words, it is epistemically responsible to think you are being

epistemically responsible unless you have a particular reason for thinking otherwise or contextual factors require exercising additional scrutiny. To deny this is to suppose that we have an unrestricted epistemic duty to make sure that we have fulfilled our epistemic duties, which immediately leads to an infinite regress of epistemic duties. If making sure that we have fulfilled our duties were an unrestricted duty then we would have a duty to make sure that we had fulfilled that duty, and a duty to make sure that we had fulfilled that duty, and so on … *ad infinitum*. But this is an impossible task, and since ought implies can (or could have), we cannot have an unrestricted epistemic duty to make sure that we have fulfilled our epistemic duties.

Now the skeptic will quickly remind us that neither the correspondence thesis nor its corollaries have been proven. True, but I haven't tried to prove them. Indeed, any attempt to do so would invariably take for granted one or more of the propositions at issue. For example, any argument we might mount that appearances are reliable will have to start from appearances that are presumably reliable. I have argued instead that they are necessary presuppositions of knowledge-oriented inquiry. Our justification for these presuppositions is, therefore, not evidential but methodological: if we are to have any chance of reaching the objectives of inquiry, then we have to hope that these things are true and conduct inquiry accordingly. The epistemic relevance of this methodological justification is that it implies that accepting the correspondence thesis and its corollaries at the outset of inquiry is epistemically responsible. To see this, recall that epistemic responsibility is doing what you can to secure grounding for your beliefs that is, as far as you can tell, adequate. But knowledge-oriented inquiry is, among other things, the pursuit of adequately grounded answers to questions. Since the correspondence thesis and its corollaries are *necessary*

presuppositions of knowledge-oriented inquiry, it follows that to conduct inquiry as if they are true is the *only* way to do what you can to secure adequately grounded answers to questions. Thus, accepting them at the outset of inquiry is epistemically responsible, i.e., they have default justificatory status. Hence, we don't owe the skeptic anything further.

III. Screening, Preferencing, and Plausibility

We have seen so far that for explanatory inference to have any chance of yielding true answers to interpretive questions there must exist search constraints that make the generation of true hypotheses significantly more likely than false ones. But we are not interested in answers that merely happen to be true. We want epistemically justified true answers; we want knowledge. And so in addition to search constraints we need epistemically principled criteria for sifting through the hypotheses generated, for screening out those that are not plausible and for deciding among those that are. Plausibility is, perhaps, most naturally construed as a degree of probability. At any rate, that is how I construe it. We should note, furthermore, that in the present context 'plausibility' means *explanatory* plausibility. That is, we want criteria to ascertain whether and to what degree hypothesis H is plausible *as an explanation for E*, the explanandum. Fortunately, there are several such criteria, of which two are paramount: *predictiveness* and *relative likelihood*.[10] Briefly, by predictiveness I mean the objective conditional probability of explanandum E given H in the light of background beliefs B, i.e., $\text{Prob}_B(E \mid H)$.[11] By relative likelihood I mean the objective probability of H in the light of B in relation to the objective probability of E in the light of B,

[10] Cf. Peirce, CP 7.202. Other criteria (e.g., simplicity, completeness, informativeness, etc.) are subsumable under these, as I explain below.

[11] The subscript 'B' signifies that the probability is to be evaluated on a sample space defined by background beliefs B, or more precisely, by that subset of B that is relevant to the probability in question.

i.e., $Prob_B(H) / Prob_B(E)$. A brief discussion of both predictiveness and relative likelihood is in order, after which I combine them into a single plausibility measure via Bayes' Theorem. After addressing an issue with consistency, I then show how plausibilities may be estimated and compared in an epistemically responsible way.

Predictiveness

A hypothesis is predictive to the extent that antecedent knowledge of the hypothesis, in the light of relevant background beliefs, would have enabled us to predict the explanandum before we knew about it. For example, someone who understands the law of gravity can predict with a very high degree of certainty a rock dropped from a height will fall. Thus we can say that the gravity hypothesis is highly predictive of the falling of dropped rocks.

Predictiveness depends in part on the *completeness* of the hypothesis as an explanation. A hypothesis is complete to the extent that it accounts for (i.e., predicts) all significant details of the explanandum. If explanandum E contains two significant details (E_1 and E_2) and hypothesis H predicts E_1 but not E_2 then H is incomplete and in that respect nonpredictive of E. For example, a complete explanation of the custom of wearing wedding rings on the fourth finger of the left hand would need to account for the wearing of wedding *rings* (as opposed to, say, necklaces), the wearing of the ring on the *fourth* finger (as opposed to, say, the third), and the wearing of the ring on the *left* hand (as opposed to the right). All other things being equal, the more details of the explanandum that are accounted for, the more predictive the hypothesis.

Predictiveness is also related to the *type* of hypothesis employed to do the explaining, for some hypotheses posit a stronger correlation between explanandum and explanans. For example, explanations in terms of universal laws are inherently more predictive than explanations in terms of stochastic (probabilistic) laws. Thus the law of gravity predicts that rocks dropped from a height will *always* fall (barring interference by some external force), whereas stochastic laws, like those governing radioactive decay, can only predict that a radioactive atom of a given type will probably decay in a certain time frame. This is not to say that universal laws are always better as explanations than stochastic laws or non-lawlike explanations such as appeals to brute chance or intelligent design, for lesser predictiveness may be counterbalanced by greater relative likelihood.[12]

Relative Likelihood

As I have represented it, the relative likelihood of hypothesis H as an explanation for E is a function of two quantities: $Prob_B(H)$ and $Prob_B(E)$. More specifically, it is the quotient $Prob_B(H) / Prob_B(E)$. To see the rationale behind this definition of relative likelihood, consider that when we ask an interpretive question about a fact E it is because there is some respect in which E is not fully intelligible in terms of our existing background beliefs B. It is the job of the hypothesis to answer the interpretive question and thus render E more intelligible by bringing E and B together, by linking what we don't understand to what we already do understand. In other words, there is a gap of understanding between E and B that we want to bridge by means of H. Now, if E does not go beyond B, i.e., if there is no gap of understanding, then $Prob_B(E) = 1$ and we have within B a fully adequate explanation H such

[12] On the relations between law, chance, and design as explanations, see Dembski 1998.

that $Prob_B(H) = 1$. But if there is a gap of understanding then $Prob_B(E) < 1$ and, as our hypothesis will have to stretch beyond B to bring E and B together, $Prob_B(H) < 1$ as well. The wider the gap, the lower $Prob_B(E)$, the further H will need to stretch in order to be an adequate explanation, and consequently the lower $Prob_B(H)$. In general $Prob_B(H)$ will have to be lower than $Prob_B(E)$ because H needs to have more content—in order to bring E and B together it needs to include E as well as state how E is related to B. Thus, if $Prob_B(H)$ is not lower than $Prob_B(E)$, then we have an explanation that is too weak and/or not informative.[13] But we don't want H to extend further than necessary to adequately explain E. All other things being equal, we want the quotient $Prob_B(H) / Prob_B(E)$ to be less than 1 but still as close to 1 as possible consistent with H being an informative explanation of E. This brings us to the explanatory virtue of simplicity (in at least one sense of the term)—we don't want H to claim too much, to posit more than is necessary to adequately account for E.

The Consistency Problem

Predictiveness and relative likelihood combine, by means of Bayes' Theorem, into the following measure of the objective probability of H in relation to E, viz., $Prob_B(H \mid E)$:

$$(1) \qquad \mathrm{Prob}_B(H \mid E) = \frac{\mathrm{Prob}_B(E \mid H)\,\mathrm{Prob}_B(H)}{\mathrm{Prob}_B(E)}.$$

According to (1), we want the product of predictiveness and relative likelihood to be as high as possible. Now as a formal principle Bayes' Theorem is unimpeachable. In practical application, however, two issues arise. The first is to show how we can derive epistemically

[13] The classic example of a non-informative explanation comes from a satire by Molière in which a doctoral student "explains" the sleep-inducing properties of opium by saying that it has "dormitive power". Clearly, the hypothesis just restates the explanandum. In other words, $Prob_B(H) = Prob_B(E)$.

responsible probability estimates (i.e., estimates that are, as far as we can tell, reliable) for the terms on the right-hand side of the equation. Call this the *estimation problem*. The second is that, while objectively and logically speaking the probabilities on the right-hand side of the equation must be mutually consistent and so must yield a result in the range [0,1], our estimates of these probabilities may not be mutually consistent and so may result in a value for $Prob_B(H \mid E)$ that is greater than one. For example, if I estimate $Prob_B(E \mid H)$ at 0.9, $Prob_B(H)$ at 0.5, and $Prob_B(E)$ at 0.2, then I will get the result that $Prob_B(H \mid E) = 2.25$, an impossible result. Call this the *consistency problem*. Let's deal with the second problem first.

If I plug my probability estimates into the right-hand side of (1) and arrive at a result greater than one, this shows me that my probability estimations are inaccurate, but it doesn't tell me where I'm off or by how much. And even if my result is less than one, that doesn't rule out the possibility that my estimates are based on beliefs that are not mutually consistent, albeit unrecognized by me. So how should we deal with this problem? One possible answer is to insist that we have an unrestricted epistemic duty to ensure that our beliefs are mutually consistent. The problem with this response is that it requires something of us that seems to be psychologically impossible. For the only way we could be certain that all of our beliefs were mutually consistent would be to have logical omniscience, the ability to see *all* of the entailments of our beliefs. But surely this is asking too much, especially since we are unconscious of the vast majority of our beliefs at any given time. And since ought implies can or could have, if logical omniscience is not psychologically possible for us then we cannot have an epistemic duty to make certain that all of our beliefs are mutually consistent. It seems the most that we can realistically demand is that we strive for consistency in our thinking and root out inconsistency whenever we find it. With respect to

the application of Bayes' Theorem we can facilitate this by stating H, E, and relevant elements of B as clearly and precisely as possible and by rewriting (1) in a way that makes it less likely that we'll inadvertently assign inconsistent probability values. Using the identity that $Prob_B(H \& E) = Prob_B(E \& H) = Prob_B(E \mid H) \cdot Prob_B(H)$ we get

$$(2) \qquad Prob_B(H \mid E) = \frac{Prob_B(H \& E)}{Prob_B(E)}.$$

In (2) it is clear that H & E claims at least as much as E. So the probability in the numerator must be less than or equal to the probability in the denominator, and $Prob_B(H \mid E)$ must be less than or equal to 1. We can also rewrite (2) using the identity $Prob_B(E) = Prob_B(H \& E) + Prob_B(\sim H \& E)$ to get

$$(3) \qquad Prob_B(H \mid E) = \frac{Prob_B(H \& E)}{Prob_B(H \& E) + Prob_B(\sim H \& E)}.$$

Again, (3) ensures that $Prob_B(H \mid E)$ is less than or equal to 1 because the numerator clearly cannot be greater than the denominator. Besides making it harder for us to assign incorrect probability values, another advantage of (2) and (3) is that we only have to estimate two values, not three. At any rate, whether we use (1), (2), or (3), we should be alert to the possibility of internal contradictions in our beliefs that could vitiate the accuracy of our probability estimates. The consistency problem cannot be altogether avoided, but we can still make epistemically responsible probability estimates by doing what we can to ensure that our estimations are, as far as we can tell, reliable.

The Estimation Problem

So how are we to make epistemically responsible probability assessments? Fortunately, we have several estimation methods at our disposal, from haphazard guessing to

highly sophisticated statistical techniques and everything in between, e.g., we can appeal to an authority, make an educated guess based on background beliefs and gut-level intuitions, extrapolate from personal experience, carefully measure relative frequencies in a sample. A detailed discussion of methods of probability estimation would be beyond the scope of this work, and it would be unnecessary. For there is one method that is fundamental to every other rational method, and this method is sufficient to yield estimations of probabilities that are default justified. When epistemic standards rise beyond the default level we can employ other methods (e.g., statistics) to check our estimations and render them more precise.

So what is this most fundamental probability estimation method? It is simply what we call educated guessing. It is educated because it is informed by background beliefs, experience, and gut-level intuitions. Is it guessing because it involves a subjective judgment. It is fundamental because all other rational probability estimation methods require educated guesses at some point. For example, I can calculate that the probability of rolling double-sixes in dice is 1/36, but only if I have made a *prior* judgment that the dice are fair. I may consult an authority to obtain information from which I can calculate probabilities, but in doing so I express a *prior* judgment that the authority is reliable. I may make sophisticated statistical extrapolations from a large and varied sample using confidence intervals, binomial distributions, control groups, and so on. Still, my results depend on a number of *prior* judgments: that all relevant factors have been identified and controlled for, that my sampling methods are reliable and not systematically biased, that the distribution is properly binomial, and so on. Unless we are dealing with tautologies and contradictions, there is no such thing as a purely objective way of estimating probabilities because every estimation procedure

rests upon subjective judgments that certain contingent assumptions are true. This point has been amply demonstrated in the literature.[14]

So the question is whether making educated guesses is an epistemically responsible thing to do. That it is seems obvious. Indeed, if making educated guesses is irresponsible, then guessing of any sort is irresponsible. But as our above discussion of the problem of underdetermination showed, if guessing is ruled out then the whole epistemological enterprise of rational inquiry cannot get off the ground. Thus, it is methodologically necessary and therefore epistemically responsible to put our default trust in the reliability of our educated guesses. In other words, particular educated guesses are always "innocent until proven guilty"; we are justified in trusting them until a relevant undefeated defeater is on the table or contextual factors raise the level of scrutiny beyond what is warranted by default justification.

Comparing Plausibilities

Once we have some hypotheses on the table and have evaluated their plausibilities, screening out those that are implausible, we can compare their plausibilities by the formula:

$$\frac{\text{Prob}_B(H_1 \mid E)}{\text{Prob}_B(H_2 \mid E)} = \frac{\text{Prob}_B(E \mid H_1)\,\text{Prob}_B(H_1)}{\text{Prob}_B(E \mid H_2)\,\text{Prob}_B(H_2)},$$

where H_1 and H_2 are two distinct hypotheses. If the fraction is greater than one, then H_1 is more plausible than H_2; if less, then the opposite is true.

Plausibility Thresholds

---

[14] See Polanyi 1958, chs. 1-3; Duhem 1991; and Howson 2000, ch. 4. Duhem argues that acceptance or rejection of a theory ultimately depends on "good sense" (pp. 216-218). Howson's work is a recent *tour de force* on this issue.

I have described how to assess and compare plausibilities, but how high does a hypothesis have to score to count as "plausible"? What is the plausibility threshold below which we should consider a hypothesis to be "implausible"? The short answer is: there is no fixed threshold. It rises or falls depending on the dialectical context. For example, we want to find a true answer to an interpretive question. So we brainstorm. If we come up with a highly plausible hypothesis with no close competitors on the horizon, we may conclude straightaway that that's the right answer. But if we have trouble coming up with a hypothesis that is plausible according to our current threshold, the problem that prompted the interpretive questions still remains. So we either have to give up trying to answer the question, keep looking for a hypothesis that meets our current plausibility threshold, or lower our threshold and start looking at less plausible hypotheses.

IV. A Positive Justification of Explanatory Inference

In light of what has been said, is it reasonable, in accord with epistemic responsibility to trust explanatory inference as a reliable means of answering interpretive questions? At least sometimes the answer is yes. As I have argued, there are several methodological presuppositions of explanatory inference that we are justified (i.e., epistemically responsible) in accepting because they are methodologically *necessary*. If the functioning of our cognitive faculties was not in sync with the intelligible structure of the world and sufficiently so that we could obtain adequately grounded beliefs, then the goals of inquiry could not be realized. Thus we must hope that this correspondence holds. We are epistemically responsible in doing so because it is only by acting in accordance with that hope that we are doing what we can to secure beliefs that are, as far as we can tell, adequately grounded. Since we are

epistemically entitled to trust the correspondence thesis and its corollaries, we can make epistemically responsible plausibility assessments of the hypotheses that we generate. That being the case we can make epistemically responsible explanatory inferences. Whether explanatory inference results in epistemically justified beliefs, however, depends on whether the correspondence thesis and the general reliability of our cognitive faculties are *in fact* true, for these make adequate grounding possible. We cannot show that they are true without begging the question, but we are nonetheless justified (i.e., epistemically responsible) in believing that they are.

There are, however, some important cases where explanatory inference will not result in epistemically justified beliefs even if the correspondence thesis and the general reliability of our cognitive faculties are true. First, sometimes we have more than one plausible hypothesis, or comparison of the plausibilities of two or more competing hypotheses yields no clear winner. Under such conditions it seems premature to conclude that one of those hypotheses is true. Instead, we should try to discriminate between the competitors through further evidence gathering and experimental testing. Second, to yield an epistemically justified belief the assessed plausibility of the hypothesis must be high, i.e., greater than 0.5. That is, we must judge that the hypothesis is more likely to be true than false. This follows from the account of inferential justification given in chapter three, which required grasping that the conclusion must be or is likely to be true if the premises are. But sometimes we are not able to come up with any highly plausible hypotheses. Sometimes even our best hypothesis is too improbable to rely on with confidence. In such cases we can lower our plausibility threshold and start looking at less plausible hypotheses to see if, perchance, one

of them might just happen to be true, but we will not yet be justified in accepting it as the truth.

In conclusion, explanatory inference can yield epistemically justified beliefs when we have a highly plausible hypothesis that is clearly superior to any extant rivals.

CHAPTER SIX

EVALUATIVE INFERENCE

In the last chapter I developed a methodologically oriented positive justification of explanatory inference; in this chapter I do the same for evaluative inference. As we will see, the issues are somewhat different, and somewhat similar. With explanatory inference we had to generate hypotheses and then assess and compare their plausibilities. With evaluative inference we have to decide which of a set of pre-specified values to assign. For example, is the hypothesis true (value one) or false (value two)? But if evaluative inference is to arrive at epistemically justified beliefs, then the assignment of values cannot be arbitrary but must be based on epistemically principled criteria, which inevitably involves experimental testing. All this I explain in section one. In section two I show how evaluative inference can be modeled in terms of Bayes' Theorem just like explanatory inference. This means that making epistemically responsible evaluative judgments also turns on whether we can make epistemically responsible probability assessments, and so we can give a positive justification of evaluative inference along similar methodological lines as in chapter five. Finally, in justifying this application of Bayes' Theorem to issues of confirmation and falsification, I respond to a couple common objections to its use in evaluative contexts.

I. The Nature of Evaluative Inference

As explained in chapter four, evaluative inferences answer categorical questions and in doing so play a judgmental or critical role in the process of inquiry. Categorical questions ask which of a set of available predicates to apply to a subject. For example: Is

the light *red*, *yellow*, or *green*? Is the hypothesis *true* or *false*? Did Oswald *shoot J.F.K.* or *not*? And so on. Whereas explanatory inference involves generating plausible hypotheses and thereby expanding the set of cognitive possibilities, evaluative inference involves *restricting* a more-or-less well-defined set of possibilities so as to determine which is actually the case. In so doing, evaluative inference evaluates. That is, it assigns a value from a value set to the subject being evaluated (the *evaluandum*). The result is a *judgment* in which a value is predicated of the evaluandum. For example, in response to the question "Is the light red, yellow, or green?", the judgment "The light is red" assigns the value 'red' from the value set {red, yellow, green} to the evaluandum (the light).

So evaluative inference assigns a value from a value-set to the evaluandum, resulting in a judgment. The content of the value-set is determined *by the question*. Anything that would count as an intelligible answer to the question is a possible value, and therefore a member of the value-set. In the question "Did Oswald shoot J.F.K." one value is explicitly stated (shot J.F.K.), and its negation (did not shoot J.F.K.) is implied. In general, evaluation requires at least two possible value choices. If one value is V then another has to be not-V. Things can get more complicated. In the stop light example we have three explicitly stated values (red, yellow, green), their implied negations (not red, not green, and not yellow), and—because the alternatives are not jointly exhaustive—the joint denial of all of them (none of the above). If the stated values are not mutually exclusive, we can also have members of the value set that are combinations of other values. For example, "Is the cake tasty or sweet?" could be answered by the judgment, "The cake is both tasty *and* sweet." To simplify matters for discussion, I will stick with the most basic case, in which the value-set contains only two values (V and not-V) that

are mutually exclusive and jointly exhaustive. Assuming bivalence and excluded middle, “Is the hypothesis true or false?” is just such a case.

Given a categorical question and a value set determined by that question, the problem we face is to decide which member of the value set to predicate of the evaluandum. There are many methods we could use for making this decision. We could, for example, flip a coin or pose the question to the first person we meet, but in many if not most cases such methods are not very reliable. What we need are decision criteria related to the truth of the hypothesis. Consider the judgment “this apple is red.” The judgment makes a factual claim, that is, it purports to assert a truth. Since the claim assigns the value ‘red’ to the apple, the claim is true if and only if the apple is in fact red. So any consequences that we can derive from the claim, assuming it to be true, give us truth-relevant criteria for assessing whether the apple is in fact red. For example, *if* the apple is in fact red, then it should appear red under good observational conditions. And *if* the apple is red, then it will appear closely similar in hue to other things already known to be red, e.g., stop signs. And so on. Our criteria come from supposing that the proposed evaluation is true and then asking ourselves what would follow. What else would we expect to be true if the proposed evaluation is?

Let’s generalize and formalize this result. Suppose we want to answer the categorical question “Is S P?” We have to decide whether predicate P properly applies to subject S or not. That is, we have to decide between the judgments “S is P” and “S is not P”. Let’s call the first our hypothesis H and the second ~H. To decide between H and ~H we need criteria, we need to identify possible evidence that could count in favor of H and against ~H, or vice-versa. Through explicative inference and in the light of background

beliefs B, we can derive additional consequences if H is true or if it is not. Thus, we predict that if (H and B) then evidence E would (probably) obtain if conditions C were fulfilled.

For example, suppose we want to evaluate whether the combination of Boyle's Law and Charles's Law (a simplified form of the Ideal Gas Law) is true. Our hypothesis H is that

$$V = k\frac{T}{P},$$

where V is volume, T is temperature, P is pressure, and k is a constant. The primary background assumption B is that the same quantity of gas is maintained. In order to evaluate the law, we have to identify criteria that would point us in one direction or the other as far as the truth of H is concerned. So we ask, "Given H and B, what evidence E would we expect if conditions C were fulfilled?" Here's one answer to the question: If H and B, then C, putting a fully inflated balloon at normal room temperature in a freezer, should E, cause it to shrink. The freezer lowers the air temperature, and the pressure is kept fairly constant by the balloon, so the volume should decrease. With this prediction in hand we can test the hypothesis by setting up the appropriate conditions and observing the results. If the results are as predicted (volume decreases), then H is confirmed. If the results are not as predicted (volume increases or stays the same) then ~H is confirmed, which is to say, H is falsified. Thus we have two possible cases:[1,2]

---

[1] This way of characterizing the logic of confirmation and falsification is adapted from Balestra 1994.

[2] Notice that neither confirmation nor falsification is deductively valid. Confirmation is invalid because it affirms the consequent. More specifically, it is invalid because even though E occurs with C it is possible that E did not in fact result from (H & B & C) but from some other cause, perhaps (~H & B & D). Falsification is invalid because while (C & ~E) by *modus tollens* on the first premise implies ~(H & B), it doesn't suffice to tell us whether H is false or whether one or more of our background assumptions B is.

Confirmation:	(H & B)→(C→E)	Falsification:	(H & B)→(C→E)
	C & E		C & ~E
	H		~H

II. A Positive Justification of Evaluative Inference

What we need to show for a positive justification of evaluative inference is that such inferences can give rise to beliefs that are, as far as we can tell, epistemically justified. In other words, we need to show that it is epistemically responsible to trust evaluative inference as a source of epistemically justified beliefs. Specifically, we need to show that we can make epistemically responsible judgments to the effect that the evaluative conclusion is probable given the premises. My strategy is to develop a probabilistic interpretation of confirmation and falsification in terms of Bayes' Theorem. Thus, making epistemically responsible evaluative judgments reduces to the problem of making epistemically responsible probability estimations, which can be handled in the same way as in chapter five.

A Probabilistic Interpretation of Evaluative Inference

When we say that evidence E confirms or falsifies hypothesis H, we mean that H is more or less likely given E than it would be otherwise.

$$\text{Confirmation:} \quad \text{Prob}_{BC}(H \mid E) > \text{Prob}_{BC}(H)$$
$$\text{Falsification:} \quad \text{Prob}_{BC}(H \mid E) < \text{Prob}_{BC}(H)$$

Thus, evidence E confirms hypothesis H if and only if the conditional probability of H given E is greater than the unconditional probability of H, and E falsifies H if and only if the conditional probability of H given E is less than the unconditional probability of H. If the probabilities are equal, then that tells us that E is irrelevant to H. The subscript 'BC'

means that the probabilities are to be evaluated on a sample space defined by background beliefs B and by the fulfillment of experimental conditions C. This is not to say that B and C cannot be questioned, but to do so would be to change the subject of evaluation from H to some other claim. As long as we are evaluating H in relation to E, other relevant factors like B and C have to stay in the background.

We can evaluate $\text{Prob}_{BC}(H \mid E)$ with Bayes' Theorem:

$$\text{(EQ1)} \qquad \text{Prob}_{BC}(H \mid E) = \frac{\text{Prob}_{BC}(E \mid H)\,\text{Prob}_{BC}(H)}{\text{Prob}_{BC}(E)}$$

But for present purposes it will be more helpful to rewrite EQ1 as follows:[3]

$$\text{(EQ2)} \qquad \text{Prob}_{BC}(H \mid E) = \frac{\text{Prob}_{BC}(H)}{\text{Prob}_{BC}(H) + \frac{\text{Prob}_{BC}(E \mid \sim H)}{\text{Prob}_{BC}(E \mid H)}\text{Prob}_{BC}(\sim H)}$$

The quotient $\text{Prob}_{BC}(E \mid \sim H) / \text{Prob}_{BC}(E \mid H)$ is known as the *Bayes Factor*.[4] It represents a comparison of the probability of E given that the hypothesis is false with its probability given that the hypothesis is true. Since $\text{Prob}_{BC}(\sim H) = 1 - \text{Prob}_{BC}(H)$, the value on the right side of EQ2 is a function of two quantities: the prior probability of H and the Bayes Factor. We can see from EQ2, $\text{Prob}_{BC}(H \mid E)$ is inversely proportional to the Bayes Factor. As the Bayes Factor approaches 0, $\text{Prob}_{BC}(H \mid E)$ will approach 1; and as the Bayes Factor increases without bound, $\text{Prob}_{BC}(H \mid E)$ will approach 0. This accords well with our evaluative intuitions. If E is more probable given ~H than given H, then the occurrence of E confirms ~H more than H, which is to say that it falsifies H. If E is more probable given H than given ~H, then of course H is confirmed. Also, if the prior

---

[3] Howson 2000, p. 179.
[4] More specifically, it is called the *Bayes Factor in favor of ~H against H*.

probability of H is low, then it is clear from EQ2 that we will need a higher Bayes Factor to raise the posterior probability of H to a point where we would say that H is more likely true than false and thus a candidate for epistemic justification. In other words, the less probable our hypothesis is to start with, the better our evidence needs to be.

It is important to note that the probabilities in EQ1 and EQ2 are *objective*; they reflect logically necessary relations among propositions. If $Prob_{BC}(H \mid E)$ is in fact greater than 0.5, then H is more likely true than false, and a *belief* in H formed via inference from E will be adequately grounded (unless contextual factors raise the adequacy standard). Subjective probabilities only come into play when we try to *apply* Bayes' Theorem, for then we have to *estimate* the probabilities on the right-hand side of the equation.

As pointed out in chapter five, casual estimations of probabilities can lead to results that violate the probability calculus. Fortunately, there are ways to rewrite Bayes Theorem in such a way that contradictory estimations are easier to spot. Unfortunately, there is no way of ensuring *a priori* that our estimations won't be mutually contradictory short of imposing a completely unrealistic demand for logical omniscience. Epistemic responsibility only requires that we be alert to possible inconsistencies in our beliefs and root them out as we become aware of them.

The question is whether we can estimate probabilities in an epistemically responsible manner, i.e., in such a way that we are entitled to rely on them in judging the extent to which E confirms or falsifies H. In the light of the results of chapter five, the answer to that is definitely affirmative, for the necessary methodological presuppositions of inquiry, of which evaluative inference is a part, give us a default entitlement to trust our educated guesses in estimating these probabilities. Thus, it is in accord with epistemic

responsibility to trust evaluative inference as a reliable source of answers to categorical questions. In other words, we are positively justified in trusting evaluative inference as a source of epistemically justified beliefs.

Of course evaluative inferences won't always result in epistemically justified beliefs. For one thing such beliefs must be adequately grounded, and that is a matter of the *de facto* reliability of the belief forming process. But even if the belief is not adequately grounded we can still be epistemically responsible in believing it to be adequately grounded. For another thing, our default entitlement to confidence in our probability estimations can be overridden by relevant defeaters, such as the recognition that some of our estimations are mutually inconsistent, or by standards-raising contextual factors. In such cases we may have to test our estimations against further evidence involving further estimations, and so on. (That's why our estimations have to have default justification—if they didn't they we'd be stuck once our initial estimations were called into question or were contextually inadequate.) Finally, to have an inferentially justified belief requires that the posterior probability of the hypothesis be greater than 0.5. But its posterior probability could be greater than 0.5 even if the experiment counted against the hypothesis, provided that its prior probability was sufficiently high. In such a case we shouldn't say that one is justified in *inferring* H from E because E undermines rather than supports H. What we should say is that the belief that H is inferentially justified given the *totality* of the available evidence bearing on H, which evidence includes the anomalous result E. In other words, evaluative inference can only lead to the *formation* of an epistemically justified belief if $\text{Prob}_{BC}(H \mid E) > 0.5$ *and* E confirms H, but one can still

maintain an epistemically justified belief in the face of recalcitrant evidence as long as $Prob_{BC}(H \mid E) > 0.5$.

III. Some Objections Answered

There are two common objections to the use of Bayes' Theorem in the context of hypothesis confirmation and falsification that should be addressed. The first is the so-called *old evidence problem*. The main proponent of this objection, Clark Glymour, has argued that, on a Bayesian account, hypotheses cannot be confirmed by evidence that we already know to be true. And since many scientific theories are regarded as having been confirmed by old evidence, there must be something wrong with the Bayesian model of hypothesis confirmation and falsification. Glymour writes:

> Newton argued for universal gravitation using Kepler's second and third laws, established before the *Principia* was published. The argument that Einstein gave in 1915 for his gravitational field equations was that they explained the anomalous advance of he perihelion of Mercury, established more than half a century earlier. . . . Old evidence can in fact confirm new theory, but according to [Bayes' Theorem] it cannot.[5]

The reason that old evidence cannot confirm a hypothesis on the Bayesian account, says Glymour, is that $Prob_{BC}(E)$, where E is already known, must equal 1. If so, then $Prob_{BC}(E \mid H) = 1$, and it follows that $Prob_{BC}(H \mid E) = Prob_{BC}(H)$. Thus, old evidence E cannot raise the posterior probability of H.

If the old evidence objection were cogent, it would vitiate my use of Bayes' Theorem to transform the justification of evaluative inference into one of probability estimation subject to the regulative constraint of logical consistency. The

[5] Glymour 1980, p. 88.

objection fails, however, because it does not follow from the mere fact that the evidence is old that $Prob_{BC}(E) = 1$. Howson and Urbach explain:

> [T]he mistake lies in relativising all the probabilities to the *totality* of current knowledge: they should have been relativised to current knowledge minus *e*. The reason for the restriction is, of course, that *your current assessment of the support of* h *by* e *measures the extent to which the addition of* e, *to the remainder of what you currently take for granted, would cause a change in your degree of belief in* h.[6]

In other words, because it is the relation between E and H that we are trying to assess, the probabilities of both E and H have to be assessed *in light of* our background beliefs, which means that we cannot, for the purposes of this assessment, treat either of them as part of our background beliefs. If we already know E, to assess how much support it gives H we have to estimate its probability *as if* we didn't already know it.

A second objection may be called the *new evidence problem*. Richard W. Miller has argued that when confronted with new evidence E for some hypothesis H we can either apply the Bayesian formula and calculate a posterior probability for H, or we can fix $Prob_{BC}(H \mid E) = Prob_{BC}(H)$ and make *ad hoc* revisions to the other priors to maintain consistency. Miller's complaint is that Bayes' Theorem gives us no insight as to which response we should make to the new evidence.[7] His point is akin to the observation by Harman that deductive rules of entailment like *modus ponens* do not tell us what to believe but only what our logical options are.[8] It is also akin to the observation by Quine

---

[6] Howson and Urbach 1989, p. 271, italics in original.
[7] Miller 1987, pp. 297-319.
[8] Harman 1986, ch. 1-2, app. A. Thus one person may accept *if* p *then* q and *p* and infer *q*, while another person may take that as *reductio* and infer from *if* p *then* q and *not*-q to *not*-p—one person's *modus ponens* may become another's *modus tollens*.

that any evidence, positive or negative, can be neutralized by making suitable revisions to our background beliefs.[9]

In response to Miller, we need to realize that if we decide the fix $Prob_{BC}(H \mid E) = Prob_{BC}(H)$, then we are no longer evaluating H with respect to E. Instead we have *changed the subject* to an evaluation of B and/or C with respect to E. We are answering a different categorical question. In other words, Miller's observation about the ambiguity of Bayes' Theorem only applies *if we abstract its application from the context of inquiry*. Within the context of inquiry there is no ambiguity. We first evaluate H with respect to E *and then*, if that result is unacceptable, we can turn around and reevaluate our background beliefs. There are two separate evaluations going on here, not one evaluation with an ambiguous result.

IV. Conclusion

We have seen that evaluative inference, like its explanatory cousin, can be modeled in probabilistic terms using Bayes' Theorem and, therefore, is positively justified provided that we can make epistemically responsible probability estimations. As argued in chapter five, we can do this by making educated guesses and striving for consistency among our estimations. Our educated guesses have default justification, and thus it is in accord with epistemic responsibility to trust them as providing adequately grounded estimations in the absence of relevant undefeated defeaters or standards-raising contextual factors. An evaluative inference *results in* an epistemically justified belief in H if and only if, consequent upon epistemically responsible probability estimations, it

[9] Quine 1980, p. 65.

follows that E confirms H and that our estimate of the probability of H conditional on E is greater than 0.5 *and* our belief that H is in fact adequately grounded, i.e., formed and maintained in a such a manner that it is objectively probable that H is true.

CONCLUSION

This study has covered a lot of ground, and there is undoubtedly much more to be said with respect to the epistemological issues surrounding induction. But when the particular questions driving an episode of inquiry have been satisfactorily answered, it is, perhaps, time for that episode to come to an end, and for a new episode to begin as energies are redirected to as-yet unsolved problems. In conclusion, then, I briefly recap the argument of the preceding chapters and then indicate what I think to be the most important lessons of this study.

In the preceding chapters I proposed responses to both the positive and negative problems of induction. Briefly, my response to the negative problem of induction, i.e., the problem of rebutting the skeptical argument, was based on an internalist definition of inference coupled with an analysis of epistemic justification in terms of adequate grounding and epistemic responsibility. This led to principles of inferential and noninferential justification in the light of which, I argued, it can be plausibly maintained that the substantive assumptions upon which inductive inference depends can be *noninferentially* justified. As for my response to the positive problem of induction, i.e., the problem of showing that our reliance on induction as a source of epistemically justified beliefs is reasonable, my approach was primarily methodological. I argued that inferences are essentially embedded within a process of inquiry, the aim of which is to solve epistemic problems. Because inferences are embedded within such a practical context that is epistemically oriented, methodological considerations bearing on the conduct of inquiry are relevant to the epistemic justification of inferentially formed beliefs because the employment of a sound methodology is the *epistemically responsible* thing to do. But there are

substantive assumptions that are necessary presuppositions of any problem solving methodology. The most fundamental, which I have called the *correspondence thesis*, expresses the hope that a solution can realistically be found and found by us. From this, I argued, follow several important epistemological corollaries that, together with the correspondence thesis allow us to show that reliance on induction is an epistemically responsible thing to do.

There are several lessons to be learned from this study, but three are, I think, of paramount significance. First, as I show in the first two chapters, the challenge of inductive skepticism cannot be easily dismissed or overcome. The problem arises out of epistemological commitments that *seem* quite plausible, so that it is not at all obvious where the flaw in the skeptical argument lies. Accordingly, giving serious attention to the argument drives positive epistemology. For showing that a seemingly plausible premise of the skeptical argument can be plausibly rejected requires development of a substantive epistemological position in the light of which the apparent plausibility of the premise can be more accurately assessed. Second, as my solution to the negative problem of induction shows, how we understand *inference* has significant epistemological ramifications, for it determines the scope of both inferential and noninferential justification. This also has significant bearing on the internalism / externalism debate, for as my discussion of Fumerton's dilemma reveals, internalists with respect to inferential justification will find it very hard to avoid skepticism unless they adopt an internalist definition of inference. Unfortunately, how to properly construe inference is an issue that epistemologists have, for the most part, paid scant attention to. Third, as my discussion in chapter four shows, inferences have a fundamental pragmatic dimension because they arise in the context of a

process of inquiry, the aim of which is to solve epistemic problems. Appreciation of this point leads to the recognition that there are two types of inductive inference (explanatory and evaluative), which are distinguished by the roles they play in the process of inquiry. Even more significantly, it means that methodological considerations are epistemically relevant to inferential justification. For, to conduct the process of inquiry in a methodologically sound manner is to be epistemically responsible with respect to the formation of beliefs by inferences embedded within that process. For these reasons, what we might call the *pragmatics* of inference is an issue that epistemologists need to pay attention to.

BIBLIOGRAPHY

Alston, William. 1989. *Epistemic Justification*. New York: Cornell Univ. Press.

Aristotle. 1984. *The Complete Works of Aristotle*. 2 vols. Edited by Jonathan Barnes. Princeton, NJ: Princeton Univ. Press.

Audi, Robert. 1998. *Epistemology: A Contemporary Introduction to the Theory of Knowledge*. London: Routledge.

Audi, Robert. 2001. "Doxastic Voluntarism and the Ethics of Belief." In Steup 2001, pp. 93-111.

Balestra, Dominic. 1994. "Galileo's Unfinished Case and Its Cartesian Product: Method, History, and Rationality." *International Philosophical Quarterly* 34, 307-322.

Blachowicz, James. 1998. *Of Two Minds: The Nature of Inquiry*. Albany: SUNY Press.

BonJour, Laurence. 1986. "The Coherence Theory of Empirical Knowledge." In Moser 1986. Originally published in *Philosophical Studies* 30 (1976): 281-312.

Broad, C. D. 1952. *Ethics and the History of Philosophy*. London: Routledge and Kegan Paul.

Bruin, John. 2001. *Homo Interrogans: Questioning and the Intentional Structure of Cognition*. Ottawa: Univ. of Ottawa Press.

Brunning, J., and P. Forster, eds. 1997. *The Rule of Reason*. Toronto: Univ. of Toronto Press.

Carroll, Lewis. 1885. "What Achilles Said to the Tortoise." *Mind* 4, 278-280.

Collingwood, R. G. 1993. *The Idea of History*. Revised edition. Oxford: Oxford Univ. Press.

Dancy, Jonathan and Ernest Sosa, eds. 1992. *A Companion to Epistemology*. Oxford: Blackwell.

Delaney, C. F. 1993. *Science, Knowledge, and Mind*. Notre Dame: Univ. of Notre Dame Press.

Dembski, William. 1998. *The Design Inference: Eliminating Chance Through Small Probabilities*. Cambridge: Cambridge Univ. Press.

Dewey, John. 1991. *Logic: The Theory of Inquiry*. Carbondale, IL: Southern Illinois Univ. Press.

Dretske, Fred. 2000. *Perception, Knowledge and Belief: Selected Essays*. Cambridge: Cambridge Univ. Press.

Duhem, Pierre. 1991. *The Aim and Structure of Physical Theory*. Translated by Philip P. Wiener. Princeton, NJ: Princeton Univ. Press.

Finocchiaro, Maurice A. 1989. *The Galileo Affair: A Documentary History*. Berkeley: Univ. of California Press.

Fogelin, Robert. 1994. *Pyrrhonian Reflections on Knowledge and Justification*. New York: Oxford Univ. Press.

French, P., T. Uehling, and H. Wettstein, eds. 1984. *Midwest Studies in Philosophy 9*. Minneapolis: Univ. of Minnesota Press.

Fumerton, Richard. 1995. *Metaepistemology and Skepticism*. Lanham, MD: Rowman & Littlefield.

Gadamer, Hans-Georg. 1988. *Truth and Method*. 2nd revised edition. New York: Continuum.

Glymour, Clark. 1980. *Theory and Evidence*. Princeton, NJ: Princeton Univ. Press.

Goldman, Alvin. 1986. *Epistemology and Cognition*. Cambridge, MA: Harvard Univ. Press.

Goodman, Nelson. 1983. *Fact, Fiction and Forecast*. 4th edition. Cambridge, MA: Harvard Univ. Press.

Greco, John. 1999. "Agent Reliabilism." In Tomberlin 1999, pp. 273-296.

Greco, John. 2000a. *Putting Skeptics in Their Place: The Nature of Skeptical Arguments and Their Role in Philosophical Inquiry*. Cambridge: Cambridge Univ. Press.

Greco, John. 2000b. "Scepticism and Epistemic Kinds." In Sosa and Villanueva 2000, pp. 366-376.

Haack, Susan. 2001. "The Ethics of Belief Reconsidered." In Steup 2001, pp. 21-33.

Harman, Gilbert. 1986. *Change in View*. Cambridge, MA: MIT Press.

Hayes, John R. 1981. *The Complete Problem Solver*. Philadelphia: The Franklin Institute Press.

Hempel, Carl. 1966. *Philosophy of Natural Science*. Upper Saddle River, NJ: Prentice Hall.

Howson, Colin. 2000. *Hume's Problem: Induction and the Justification of Belief*. Oxford: Clarendon Press.

Howson, Colin, and Peter Urbach. 1989. *Scientific Reasoning: The Bayesian Approach*. La Salle, IL: Open Court.

Hume, David. 1975. *Enquiries Concerning Human Understanding and Concerning the Principles of Morals*. 3rd edition. Oxford: Clarendon Press.

Kelley, David. 1986. *The Evidence of the Senses: A Realist Theory of Perception*. Baton Rouge: Louisiana State Univ. Press.

Kornblith, Hilary, ed. 1994. *Naturalizing Epistemology*. 2nd edition. Cambridge, MA: MIT Press.

Kornblith, Hilary, ed. 2001. *Epistemology: Internalism and Externalism*. Malden, MA: Blackwell.

Kuhn, Thomas. 1977. *The Essential Tension*. Chicago: Univ. of Chicago Press

Laudan, Larry. 1977. *Progress and Its Problems: Towards a Theory of Scientific Growth*. Berkeley: Univ. of California Press.

Levi, Isaac. 1997. "Inference and Logic According to Peirce." In Brunning and Forster 1997, pp. 34–56.

Lipton, Peter. 1991. *Inference to the Best Explanation*. London: Routledge.

Lonergan, Bernard. 1958. *Insight: A Study of Human Understanding*. San Francisco: Harper & Row.

Miller, Richard W. 1987. *Fact and Method*. Princeton, NJ: Princeton Univ. Press.

Morick, Harold, ed. 1980. *Challenges to Empiricism*. Indianapolis: Hackett.

Moser, Paul, ed. 1986. *Empirical Knowledge*. Totowa, NJ: Rowman and Littlefield.

Newton-Smith, W. H. 1981. *The Rationality of Science*. London: Routledge.

Papineau, David. 1993. *Philosophical Naturalism*. Oxford: Blackwell.

Peirce, Charles S. 1931-1958. *Collected Papers of Charles Sanders Peirce*. 8 vols. Edited by C. Hartshorne, P. Weiss, and A. Burks. Cambridge, MA: Harvard Univ. Press.

Peirce, Charles S. 1982-2000. *Writings of Charles S. Peirce*. 6 vols. to date. Edited by M. H. Fisch, E. Moore, and C. J. W. Kloesel, et al. Bloomington, IN: Indiana Univ. Press.

Peirce, Charles S. 1992-1998. *The Essential Peirce: Selected Philosophical Writings*. 2 vols. Edited by N. Houser, C. Kloesel, and the Peirce Edition Project. Bloomington, IN: Indiana Univ. Press.

Plantinga, Alvin. 1983. "Reason and Belief in God." In Plantinga and Wolterstorff 1983, pp. 16-93.

Plantinga, Alvin, and Nicholas Wolterstorff, eds. 1983. *Faith and Rationality: Reason and Belief in God*. Notre Dame, IN: Univ. of Notre Dame Press.

Polanyi, Michael. 1962. *Personal Knowledge*. Chicago: Univ. of Chicago Press.

Polya, G. 1957. *How to Solve It: A New Aspect of Mathematical Method*. 2nd edition. Princeton, NJ: Princeton Univ. Press.

Popper, Karl. 1959. *The Logic of Scientific Discovery*. London: Routledge.

Popper, Karl. 1979. *Objective Knowledge*. Revised edition. Oxford: Oxford Univ. Press.

Quine, W. V. O. 1980. "Two Dogmas of Empiricism." In Morick 1980, pp. 46-70.

Quine, W. V. O. 1994. "Epistemology Naturalized." In Kornblith 1994, pp. 15-31.

Reichenbach, Hans. 1938. *Experience and Prediction*. Chicago: Univ. of Chicago Press.

Reid, Thomas. 1983. *Inquiry and Essays*. Edited by Ronald E. Beanblossom and Keith Lehrer. Indianapolis: Hackett.

Rescher, Nicholas. 1980. *Induction*. Pittsburgh: Univ. of Pittsburgh Press.

Ross, Sheldon. 1998. *A First Course in Probability*. 5th edition. Upper Saddle River, NJ: Prentice Hall.

Russell, Bruce. 2001. "Epistemic and Moral Duty." In Steup 2001.

Salmon, Wesley. 1966. *The Foundations of Scientific Inference*. Pittsburgh: Univ. of Pittsburgh Press.

Skyrms, Brian. 1975. *Choice and Chance*. 2nd edition. Encino, CA: Dickenson Publishing Co.

Sosa, Ernest, and Enrique Villanueva, eds. 2000. *Philosophical Issues, 10, Skepticism*. Cambridge, MA: Blackwell.

Stalnaker, Robert C. 1987. *Inquiry*. Cambridge, MA: The MIT Press.

Steup, Matthias, ed. 2001. *Knowledge, Truth, and Duty*. Oxford: Oxford Univ. Press.

Stove, David C. 1973. *Probability and Hume's Deductive Scepticism*. Oxford: Clarendon Press.

Stove, David C. 1986. *The Rationality of Induction*. Oxford: Oxford Univ. Press.

Strawson, Peter. 1952. *Introduction to Logical Theory*. Oxford: Oxford Univ. Press.

Tomberlin, James E., ed. 1999. *Philosophical Perspectives, 13, Epistemology*. Cambridge, MA: Blackwell.

Tragresser, Robert S. 1992. "Inference." In Dancy and Sosa 1992.

Van Cleve, James. 1984. "Reliability, Justification, and the Problem of Induction." In French, Uehling, and Wettstein 1984.

Walton, Douglas N. 1989. *Informal Logic: A Handbook for Critical Argumentation*. Cambridge: Cambridge Univ. Press.

Williams, Donald C. 1963. *The Ground of Induction*. New York: Russell & Russell.

Williams, Michael. 2001. *Problems of Knowledge: A Critical Introduction to Epistemology*. Oxford: Oxford Univ. Press.

Wittgenstein, Ludwig. 1958. *Philosophical Investigations*. 3rd edition. Translated by G. E. M. Anscombe. Englewood Cliffs, NJ: Prentice Hall.

APPENDIX

FOUR CLUES TO GUIDE HYPOTHESIS GENERATION

There are at least four types of clues or constraints that guide and inform hypothesis generation: (a) the problem situation, (b) background beliefs and concepts, (c) experience, and (d) correction of prior hypotheses. Let us look briefly at each.

I. The Problem Situation

Inquiry is rooted in a problem situation. Initially, our grasp of the problem situation may be no more than a vague feeling that "something" is amiss. Vague somethings, however, aren't much to go on. Without a more precise grasp of the problem situation, inquiry lacks direction. But if the problem situation can be brought into focus and transformed into a determinate problem, then real progress toward a solution may be possible. Dewey explains:

> A problem represents the partial transformation by inquiry of a problematic situation into a determinate situation. It is a familiar and significant saying that a problem well put is half-solved. To find out what the problem and problems are which a problematic situation presents to be inquired into, is to be well along in inquiry. To mis-take the problem involved is to cause subsequent inquiry to be irrelevant or to go astray. Without a problem, there is blind groping in the dark. The way in which the problem is conceived decides what specific suggestions are entertained and which are dismissed; what data are selected and which rejected; it is the criterion for relevancy and irrelevancy of hypotheses and conceptual structures.[1]

As Polya points out, the first step in trying to solve any problem is to make sure that you *understand* it.[2] The more clear we are on what the problem is, the more clear we will be on where to look for solutions and on what would count as an adequate solution.

---

[1] Dewey 1999, pp. 111-112.
[2] Polya 1957, p. xvi.

As important as understanding the problem is, however, it may still leave the solution radically underdetermined. We still need additional constraints to limit the solution search space to a manageable size. One source of these constraints, background beliefs and concepts, is essential if understanding a problem is even to be possible.

II. Background Beliefs and Concepts

The interrelated set of background beliefs and concepts possessed by a cognitive agent at a given time constitutes that agent's *noetic structure*. Having a noetic structure is essential for hypothesis generation because a problem cannot even be defined apart from concepts. This evident fact poses a difficulty: If concepts are learned, and if learning is the result of inquiry, and if inquiry can only make progress after the problem has been defined, and if defining a problem requires concepts, then how could we have acquired those concepts in the first place? The answer to this difficulty is obviously that not all concepts are learned, as Thomas Reid points out:

> [I]f mankind had not a natural language, they could never have invented an artificial one by their reason and ingenuity. For all artificial language supposes some compact or agreement to affix a certain meaning to certain signs; therefore, there must be compacts or agreements before the use of artificial signs; but there can be no compact or agreement without signs, nor without language; and, therefore, there must be a natural language before any artificial language can be invented.[3]

Reid argues that in order to have acquired our spoken language we must have already possessed an innate, natural language that allowed us to grasp the meanings implicit in things like facial expressions, tones of voice, body postures, etc. Similarly, since inquiry requires a defined problem and a defined problem requires concepts, our noetic structure cannot itself be wholly the outcome of inquiry; part of it must be innate or *a priori*.

---

[3] Reid 1983, p. 32.

Furthermore, if this *a priori* noetic structure is to guide us to *true* hypotheses—and the successes of science suggest that to a significant degree it has—then there must be some affinity between how our minds are "wired" and the intelligible structure of the world. This is the correspondence thesis that I discuss in chapter five.

So we have to suppose that hypothesis generation is guided in part by our possession of an innate noetic structure that corresponds significantly with the intelligible structure of reality. At the same time it is equally clear that most of our noetic structure, like most of our language, is acquired, not innate. Working out the specifics of this innate noetic structure is not important for our purposes. Indeed, that job is better suited to cognitive scientists than to philosophers. What is important is that having a noetic structure not only gives us the conceptual resources to define and formulate problems but also constrains what kinds of solutions are permissible. Just as we expect the piece of the puzzle that we are looking for to fit with the pieces already in place, we want the hypotheses we generate to be consistent, if possible, with our existing noetic structure. Furthermore, our existing noetic structure provides a database of concepts and schemata that can be searched for possible answers to the problem at hand. If we see that the current problem is analogous to a problem that we already know how to solve, then that suggests that the solution is analogous as well.[4]

As already stated, most of our noetic structure is acquired. But from where do we get this new conceptual content? The obvious answer is *experience*. Learning takes place through the interplay between experience and an existing noetic structure. Our noetic

---

[4] Thomas Kuhn points out that much of scientific training consists in acquiring a stock set of paradigmatic problem–solution pairs that provide models for how to solve the particular problems that a scientist is likely to encounter. Cf. "Second Thoughts on Paradigms," in Kuhn 1977, pp. 293-319.

structure helps us interpret experience, but experience also informs and shapes our noetic structure. So experience must provide in some sense an independent constraint on our interpretations of it.

III. Experience

In its broadest sense, "experience" refers to that which a subject undergoes—thus we might say that a hit baseball "experiences" the impact of the bat. Typically, however, we use the term to refer to that which *sentient* subjects undergo *insofar as they are conscious*. In this sense "experience" refers to the *presentational* content of a subject's conscious state. To have an experience is to be appeared to in a certain way. The term "perception," by way of contrast, refers to the *re-presentational* content of a subject's conscious state. In other words, perception differs from experience in the same way that *seeing as* differs from mere *seeing*. The key question is how something non-propositional (experience) gives rise to and informs something propositional (a perceptual judgment). As Quine put it, how do we get from the "meager input" of sensory stimulation to the "torrential output" of significant knowledge about the world?[5] This is a large and difficult question. That experience does inform our perceptual judgments seems obvious. How this occurs is not so clear. Fortunately, delving into details here is not necessary for my purposes.[6]

IV. Correction of Prior Hypotheses

This fourth guide to hypothesis generation won't help us generate an initial hypothesis, but it can make subsequent revisions of that hypothesis *much* more efficient

[5] Quine 1994, p. 25.
[6] A recommended source for the interested reader is Kelley 1986.

in progressing toward adequate answers to our questions, as James Blachowicz explains in his recent book.[7] He begins by distinguishing *correction* from *elimination*:

> To make a correction, you need to use some information from the previous error regarding how close you were to the desired result. If you know you missed by such-and-such amount in such-and-such a direction, you can intelligently vary your trials, "zeroing-in" on your goal. This must be distinguished from the sense of increased "closeness" that is a product of elimination. Obviously, you get closer to a final answer as you exhaust a finite number of alternatives; but you do not use positive information about the nature of this closeness to affect the relative probability of future trials.[8]

For example, suppose you are asked to guess a number between 1 and 100. If after each guess you are simply told either "yes" or "no," you receive no information that can be used to reassess the plausibilities of other guesses. You have to proceed by simple elimination, working through the numbers one at a time until you get the right one. Now suppose that after each guess you are told "higher" or "lower." If the number is 87 and your first guess is 50, you can now eliminate a whole range of guesses with one trial, and you will home in on the correct answer much more rapidly. In this case you know in what *direction* to adjust your trials. If, in addition, after each trial you were told not only "higher" or "lower" but also given an estimate of the size of the discrepancy (e.g., "you're within 20"), then the search process will be faster still. You now know not only what direction to move in but also approximately what *distance* to move. In such cases, it matters little what one's initial guess is—the corrective process will quickly overcome the initial error.

It should be obvious that when corrective procedures can be applied, we have a powerful guide to hypothesis generation. A major reason why science has flourished with

---

[7] Blachowicz 1998.
[8] Blachowicz 1998, p. 18.

the application of quantitative methods is that such methods allow us to measure in a fairly definite way the direction and distance that we need to travel to reach our goal. But correction has its limits. Correction is useful when considering a given *type* of hypothesis and trying to find a quantitative *variant* of that type that will yield more accurate predictions. It can enable one to quickly find the optimum variant of that type. But if the optimum variant of a given type is still inadequate, it is clear that we must look for another type of hypothesis, and for *that* sort of task correction is not very helpful. In the above examples, the type of hypothesis or guess is understood to be an integer. There is also a specified range of variants to consider: integers between 1 and 100. Because we can define direction and distance in the 'variant-space,' information that tells us in what direction or how far to adjust our trials can readily be put to use. But where neither directions nor distances can be defined in the variant-space, correction must give way to elimination.[9] Blachowicz provides a nice example of this type of situation:

> You get into your car in the evening and find that your headlights don't work. There are, you may estimate, four possible reasons: the battery, the fuses, the bulbs, or a loose wire. You start the engine successfully, thereby eliminating the battery as the cause. But this doesn't tell you anything about the probabilities of the remaining three alternatives.[10]

Here we have four distinct hypothesis types, and the only available procedure to decide among them is elimination, testing each hypothesis-type in turn until a positive result is obtained.

---

[9] On this point I differ from Blachowicz. He seems to think that *all* hypothesis generation is informed by correction: "The generation of a new hypothesis should be understood as the intelligent correction of an antecedent hypothesis in response to the precise nature of the discrepancies between the desired result and results generated from (predicted by) the antecedent hypothesis" (1998, p. 29). I think correction is only applicable *within* a given hypothesis type, so that it cannot guide one to the generation of a distinctly new type of hypothesis.

[10] Blachowicz, 1998, p. 17.

ABSTRACT

Alan Robert Rhoda

BA, University of Nevada, Las Vegas

MA, Fordham University

The Problem of Induction: An Epistemological and Methodological Response

Dissertation directed by John Greco, PhD

In this dissertation I respond to two related problems of induction. The *negative problem* is to rebut the Humean skeptical argument, which argues that induction cannot result in epistemically justified beliefs because the substantive assumptions upon which induction depends cannot be appropriately justified. The *positive problem* is to show that it is reasonable to rely on induction as a source of epistemically justified beliefs. With respect to the negative problem, I examine and reject several responses and then propose an original response of my own that combines a partially internalist / partially externalist account of epistemic justification with an internalist account of inference (i.e., one that requires having a conscious perspective on the relation between premises and conclusion). On this basis I argue that the substantive assumptions upon which induction depends can be *noninferentially* justified. With respect to the positive problem, I argue that inferences are essentially embedded within a process of inquiry, the aim of which is to solve epistemic problems. Two important consequences follow. First, we discover two distinct kinds of induction, which I call explanatory and evaluative

inference, respectively. A complete solution to the positive problem, therefore, needs to address both. Second, methodological considerations bearing on the conduct of inquiry are epistemically relevant to the inferential justification of belief. I exploit this result by showing that there are *methodologically necessary* assumptions undergirding the conduct of inquiry. Chief among these is the *correspondence thesis*, i.e., that reality is significantly *knowable*, and significantly knowable *by us*. From this I derive several important epistemological corollaries. I then employ Bayes' Theorem to express the logic of explanatory and evaluative inference in terms of probabilities, and use the correspondence thesis and its corollaries to argue that we can make epistemically responsible estimations of those probabilities. It follows that we can responsibly rely on induction as a source of epistemically justified beliefs.

VITA

Alan Robert Rhoda, son of R. Carl Rhoda and Ursula Limuti, was born on August 24, 1970, in Las Vegas, Nevada. After graduating as valedictorian of Las Vegas High School in 1989, he entered the University of Nevada, Las Vegas as the recipient of the four-year Elardi-Wynn Scholarship. In 1993, he graduated *summa cum laude* with a Bachelor of Arts degree in Philosophy.

In 1993, Alan entered Fordham University with the Bennett Scholarship and a four-year Presidential Fellowship. While at Fordham, he served for two years as staff assistant for the *International Philosophical Quarterly*. He earned a Master of Arts degree in Philosophy in 1996. While working toward his doctoral degree in Philosophy, under the mentorship of John Greco, he has taught undergraduate philosophy courses at both Fordham University and the University of Nevada, Las Vegas. In the summer of 2003 he married Heather Logsdon and currently resides in Las Vegas.