This is to certify that the

dissertation entitled

MEANINGS, MANNERS, AND SCEPTICISM

presented by

Sarah Black Jones

has been accepted towards fulfillment
of the requirements for

Ph.D. degree in Philosophy

Albert C Cafagna
Major professor

Date 12/14/01

DATE DUE	DATE DUE	DATE DUE

MEANINGS, MANNERS, AND SCEPTICISM

By

Sarah Black Jones

A DISSERTATION

Submitted to

Michigan State University

in partial fulfillment of the requirements

for the degree of

DOCTOR OF PHILOSOPHY

Department of Philosophy

2001

ABSTRACT

MEANINGS, MANNERS, AND SCEPTICISM

By

Sarah Black Jones

In Meanings, Manners, and Scepticism, I argue that the ordinary language philosopher, the relevant alternativist, and the contextualist fail to refute the sceptic. I argue that they fail to refute the sceptic because their arguments critically rest on a number of both dogmatic and erroneous assumptions. In particular they rest on the dogmatic and erroneous assumptions that common sense is threatened by skepticism, that ordinary linguistic behaviour is of epistemic significance, and that the sceptic's departure from the linguistic norm suffices to prove his sceptical alternatives irrelevant to the truth of ordinary claims to know. I argue that the first assumption is mistaken because it wrongly assumes that common sense consists not only of first-order beliefs about the world, but also second-order knowledge beliefs about those beliefs. I argue that the second assumption is mistaken because it conflates two logically distinct sets of standards: the standards of appropriate and the standards of true assertion. Finally, I argue that the third assumption is mistaken because it rests on a confusion of the kind of standard operative in ordinary knowledge-discourse with the kind of standard the sceptic appeals to when he reasons that we can never know. I conclude that appealing to ordinary linguistic behaviour and habits cannot suffice to prove the sceptic's hypothesis irrelevant to the truth of our knowledge claims, though insofar as common sense does not include second-

order knowledge beliefs the truth of even the most radical skepticism cannot threaten our ordinary claims to know.

For Julia and James, my mother and father
and
Katharine, my daughter

ACKNOWLEDGMENTS

There a number of people who I wish to thank for having made the completion of this work possible. Most especially, I would like to thank certain of my teachers, each of whom it was a great privilege to know: First, and most especially, professor Herbert J. Garelick whose call to the vocation of philosophy must surely have been Divine, and of whom no words of praise or appreciation could possibly do justice. To him, I am most grateful for having been his student, and for the unique and unforgettable privilege of having been in his presence. Professor Charles J. McKracken whose classical style and commitment to tradition will forever influence me. Professor Hal Walsch for his absolute refusal to acquiesce to senseless dogmatism, and, whose penchant to lapse into Latin and Greek made for the most wonderful of learned lectures. Professor Richard J. Hall, for the years of his always intellectually stimulating courses and seminars, and for his encouragement. The members of my dissertation committee, especially Debrah Nails and Winston Wilkinson for their detailed comments. Finally, professors Stephen Esquith and Albert Cafagna, the kindest of philosophers, without whose support I surely could not have completed this work.

TABLE OF CONTENTS

CHAPTER I:
Introduction

In his *Outlines of Pyrrhonism*, Sextus Empiricus wrote: "Scepticism is an ability, or mental attitude, which opposes appearances to judgments in any way whatsoever . . . ."[1] The goal of this sceptical attitude was initially to reach a state of quietude; a state initially thought to be achieved by discovering the truth. But upon subjecting every belief to criticism, the Pyrrhonian discovered that he arrived not at certainty and quietude, but rather, at uncertainty and disquietude.[2] This uncertainty was a consequence of what the Pyrrhonian inquiries did uncover: not truth, but a number of competing possible truths. For every belief that P, that is, they found alternative beliefs that – P.[3] Moreover, for each of those alternative beliefs that – P, they discovered reasons and evidence. The reasons and evidence for – P, they asserted, were "equipollent" to those asserted for P. These reasons, that is, were equal with respect to their probability.

---

[1] Sextus Empiricus, *Outlines Of Pyrrhonism*, trans. R.G. Bury, (Cambridge: Harvard University Press, 1933), 7.

[2] What they also found was that they arrived at a state of quietude, not through knowledge of the truth, however, but through the recognition that the search for knowledge of the true nature of reality is futile, and can end only in frustration and disquietude. Thus did they replace this search for knowledge with the suspension of judgment about all matters "not evident."

[3] One here, I think, begins to see the origins of not only the relevant alternativist's position, but equally, of each sceptical and anti-sceptical effort to deal with the relevance of alternatives to one's belief that P. Indeed, as we shall see in chapters II, III, and IV of this work, the question of the relevance of such alternatives is at the heart of much of the twentieth century attempt to answer the sceptic.

The Pyrrhonian sceptic's recognition of this state of evidentiary "equipollence"[4] led the Pyrrhonian to "suspend judgment" about P altogether:

> The phrase "I suspend judgment" we adopt in place of "I am unable to say which of the objects presented I ought to believe and which I ought to disbelieve," Indicating that the objects appear equal to us as regards credibility and incredibility.[5]

They argued, therefore, that to assent either to P or to – P would be wholly dogmatic, given P's non-evident nature. Such assent is dogmatic, the Pyrrhonian argued, because there is no criterion sufficiently justified by which we may separate the true from the false;[6] no criterion, that is, according to which we can *know* with certainty which of our beliefs are true. Why is there no criterion sufficient for distinguishing the true from the false? The Pyrrhonian summed it up this way:

> In order to decide the dispute which has arisen about the criterion, we must possess an accepted criterion by which we shall be able to judge the dispute; and in order to possess an accepted criterion, the dispute about the criterion must first be decided. And when the argument reduces itself to this form of circular reasoning, the discovery of the criterion becomes impracticable, since we do not allow them [our opponent] to adopt a criterion by assumption, while if they offer to judge the criterion by a criterion, we are forced to a regress *ad infinitum*.[7]

---

[4] *Outlines*, 111.

[5] *Outlines*, 115.

[6] The importance of this ancient skeptical argument is evidenced throughout the history of philosophy. Indeed, it was just such an objection which Gassendi offered against Cartesianism in the 17th century and which Descartes regarded as the "objection of objections." CF both "The Fifth Set of Objections," *in The Philosophical Works of Descartes*, trans. Elizabeth Haldane and G.T. Ross, (Cambridge: Cambridge University Press, 1955), and Richard Popkin, "The High Road to Pyrrhonism," *American Philosophical Quarterly 2*, (January 1965): 1-15.

[7] *Outlines*, 165.

Both the spirit of this sceptical attitude, as well as the consequences which follow from its application, have haunted knowledge claims ever since. Indeed, it is precisely the question of such a criterion which is at the heart both of Descartes' *Meditations* as well as the criticisms leveled against it.[8] Wherever there is uncritical assumption, there is also an arsenal of alternatives, and the question of the criterion of the criterion. This plethora of alternatives to P, together with the vexing problem of the criterion, is sufficient to expose the dogmatic nature of any claim to certainty.[9][10] So unrivaled, in fact, is its power both to disrupt and to throw into chaos all varieties of dogmatic thought, that after its long absence during the medieval period, and upon its rediscovery (as well, of course, as the translation of it into Latin and French) at the dawning of the renaissance, the *Outlines of Pyrrhonism* (a work which catalogues all of the means of arriving at alternatives, as well as the insoluble problems for distinguishing amongst them) gave rise to a crisis for all presumed knowledge; a crisis for dogmatic thought unrivaled in all of history.[11]

Upon the *Outlines'* rediscovery, there ensued an unparalleled struggle between reason

---

[8] See especially, *Meditations* I and II, Trans. Elizabeth S. Haldane, and G.R.T. Ross (Cambridge: Cambridge University Press , 1911).

[9] Indeed, the mere possibility of such equipollent alternatives, combined, of course, with the absence of any defensible criterion, has constituted the chief agnoilogical weapon of the sceptic from the Pyrrhonians through Hume.

[10] It is interesting to note, that in spite of the now ancient origins of the regress argument (regarding the criterion), it was precisely to deal with this problem (as inherited by Descartes) that contemporary theories of knowledge such as coherentism, and contextualism were developed. And, as we shall see, throughout the contemporary period it is dispute over the criterion which provides at least a major part of the impetus for the development of relevant alternativism, whose central concern is just the delineation of the relevant from the non-relevant alternatives to knowing.

[11] For a very thorough presentation of the Pyrrhonian impact in the Modern period, see Richard Popkin's "The High Road to Pyrrhonism," and "Scepticism and Anti-scepticism," in Popkin, *The High Road to Pyrrhonism*, Ed. Richard A. Watson and James E. Force, Indianapolis: Hacket Publishing Company, 1980).

and belief. This struggle is everywhere evident: from the attack of the French intellectuals such as Pierre Gassendi, Simon Foucher, and Samuel Sorbiere, on the "new Dogmatism" of Descartes, to the complete unraveling of "the warp and woof of man's intellectual world"[12] effected by Pierre Bayle's sceptical *Dictionary*. [13] That struggle's centuries long exposition and disputation put an end (at least for the French and Germans) to the former satisfaction achieved by resting the veracity of either one's religious or one's scientific knowledge in the lap of speculative reason.[14] Indeed for the first time in nearly a millennium, dogmatic belief was the subject of relentless and philosophical scrutiny; a Pyrrhonian scrutiny, the effect of which shook the foundations of western culture to a degree from which it has not recovered.

What made possible the seismic impact of these ancient arguments? What made possible the resulting destruction of man's faith in reason? The modern response to the arguments of the sceptic was by and large to concede that the sceptic established beyond any doubt the lack of any foundation *in reason* for both religious and scientific knowledge claims. It was widely acknowledged that by proving that choosing one alternative description of reality's nature rather than another has no *rational* basis, (owing to the lack of any justified criterion) the sceptic established the impossibility of knowledge.[15]

---

[12] Popkin, "The High Road to Pyrrhonism," 25.

[13] Perhaps the most significant of which was the impact of Hume's Pyrrhonism on the rest of the history of philosophy through its influence on Kant.

[14] Again, Popkin identifies the insipid reaction of the British – especially amongst the common sense philosophers. See, for example, his "Scepticism and Anti-scepticism in the Latter Part of the Eighteenth Century," in Popkin, *The High Road to Pyrrhonism.*

[15] Such reasoning is echoed throughout the history of scepticism.

Furthermore, she presented what seemed an irrefutable case for the lack of any criterion sufficiently justified to determine what is the case and to distinguish it from what is not. What reason's application to belief revealed, then, was not truth, but a plethora of uncertainty. Reason, that is, had established the utter indefensibility of *judging* reality to be one way rather than another.

As a consequence, the dogma of a millennium was either (though atypically) rejected in total or was said to rest (necessarily) on faith. The sceptic's ruthless and Pyrrhonian criticism of the claims to *knowledge* of belief's content resulted, therefore, not only, in the overt recognition of reason's inability to defend belief, but also in a proliferation of the defense of convention. Notable, however, was the absence of convention's defense by reason. Equally notable was the concomitant proliferation of claims that only faith could save belief. Throughout this neo-Pyrrhonian period the primary response to scepticism was to separate what could be known through reason, from what could be known only through faith. This is evident, of course, in Pascal, Montaigne, and the other great French intellectuals of the time. Furthermore, though this strategy of defending the norm (the common beliefs of the time) preserved to a certain extent much of what was dogmatically assumed, it did so without any pretense to a rational foundation.[16] The Modern defense of belief acknowledged the agnoilogical power of the sceptical arguments, and the consequent futility of giving a *reasoned* defense of belief. Faith, after all, not reason,

---

[16] Here, of course, I am excepting Descartes' fantastic dogmatism, as well as that of other Modern rationalists.

could alone save belief.[17]

In fact, although nearly every sceptic of the Modern period was -- in spite of her scepticism -- a devout believer, none pretended to justify or defend those beliefs through human reason.[18] David Hume, of course, represents the culmination of this Pyrrhonian critique of judgment. Hume went beyond the sceptical conclusion that reason could not ground belief to argue that belief itself is grounded in something other than reason.[19] Instinct, Hume argued, not reason, is the source and foundation of belief. Believing that there is an external world, that there are other minds, that the future will be like the past, is a matter of *natural*, not rational, necessity.

> Nature by an absolute and uncontrollable necessity has required us to judge as well as to breathe and feel.[20]

A fundamental consequence of belief's *natural* necessity is that no sceptical argument, however powerful, can effect its abandonment. "Philosophy" Hume acknowledged, "would render us completely Pyrrhonian, were not nature too strong for it."[21] Hume's insight into the *necessity* of belief is in fact recognized repeatedly in the history of scepticism. Indeed, the *necessity* of belief in spite of reason's opposition is evidenced in

---

[17] There were, of course, exceptions. One might think Descartes among the most notable of such exceptions. Yet even he, in his response to Gassendi, for example, did not pretend to be able to overcome the Pyrrhonian attack on the criterion; rather, he simply pointed to the devastating consequences this attack would have for knowledge if taken seriously.

[18] Montaigne, Pascal, Foucher, Gassendi, Huet, and Sorbieri.

[19] A point which leads Hume to disagree with the Pyrrhonians about the possibility that we *can* suspend belief.

[20] David Hume, *A Treatise of Human Nature, ed.*, L. A. Selby-Bigge (Oxford: Oxford University Press, 1988).

[21] Hume, *Treatise*,183-4

ancient Pyrrhonism, and pronounced unequivocally by Sextus.

> And though the both the Academics and the Sceptics say they believe some things, yet here too the difference between the two philosophies is quite plain. For the word "believe" has different meanings: it means not to resist but simply to follow . . . but sometimes it means to assent to a thing of deliberate choice . . . . our belief is a matter of simple yielding without any consent . . . .[22]

> Adhering, then, to appearances we live in accordance with the normal rules of life, undogmatically, seeing that we cannot remain wholly inactive.[23]

In spite of the apparent conflict, then, between reason and belief; between the examined and the unexamined; between scepticism and common sense, the sceptic (from Pyrrho to Hume) has long recognized the impossibility of suspending belief *altogether*. Belief, after all, is necessary to existence. (Without the belief that there is an external world, for example, we would each surely perish. Without the belief in other minds, we would have no reason for communication. Without the belief in the order of nature we would make no plans.) The sceptic, therefore, has thematically and consistently maintained a distinction between the natural, ineradicable, and indispensable, propensity and necessity to believe (in spite of reason's inability to confer certainty upon that belief), and the eradicable and unjustified judgment regarding the knowability of those beliefs. Accordingly, she has maintained that we must acquiesce to the natural necessity of belief, but without the addition of indefensible and unnecessary *judgments* concerning our knowledge of belief's truth or falsity.

The sceptic, then, does not recommend suspension of mere belief regarding ordinary

---

[22]*Outlines*, 141.

[23]*Outlines*, 17.

matters (a suspension she would have regarded as impossible), but rather the suspension only of *judgment* about the veracity of those beliefs. For mere belief – of the sort that is *not* accompanied by assent – is itself neither dogmatic nor suspendable. What *is* both dogmatic and suspendable, however, (and, therefore, objectionable to the sceptic) is the kind of "second-order" belief that makes pretensions to *know* what is believed. It is *judgment*, that is, about whether we know the truth of those beliefs that has the sceptic's attention. For it is only in judgment, not mere belief, that one goes beyond what is evident. We must distinguish, therefore, between our direct beliefs about the world (what henceforth I shall refer to as *first-order* beliefs) and our judgments *about* those beliefs (what I call *second-order* beliefs, indicating thereby a kind of belief which is not about the world, but rather, *about beliefs* about the world). The sceptic accepts both what she *must* believe as well as the (universal) natural necessity of our unwillingness to give up those (*first-order* beliefs), but maintains all the while that no matter how unwilling (or incapable) we may be to give up those beliefs, any claim to *know* what we believe cannot enjoy rational defense, and never, therefore, can constitute *knowledge*. Why can our second-order knowledge beliefs enjoy no rational defense? There is, again, no rationally justifiable criterion for choosing our belief that P over its alternative. There is, that is to say, no rational justification for supposing that P is true while –P is false.

In the twentieth century the natural necessity (and the consequent universality and ineradicability) of belief has, however, been widely confused with belief's epistemic justification. The distinction between *judgments about* those necessary beliefs (what I am calling *second-order* beliefs), and beliefs of the first-order has also been largely

overlooked. The necessity of believing that there is a well-ordered external world, for example, which causes our perceptual experience has been confused with the unnecessary (and unordinary) belief that we can *know* the world is thus and so. As a consequence, contemporary anti-sceptics have not only misunderstood the true object of radical philosophical scepticism, but also, therefore, have misunderstood the kind of argument needed to refute such scepticism. This conflation and error is clearly evident in the manner in which three of the most important anti-sceptical movements of the twentieth century argue against the sceptic: it is evident in the ordinary language philosopher, the relevant alternativist, and the contextualist's attempt to refute the sceptic. Indeed, it is evidenced not only in the particulars of their anti-sceptical arguments, but also in their conviction that common sense is threatened by scepticism.

It is only on the basis of their failure to give appropriate consideration to the sceptic's insistence that mere first-order belief about the world is necessary, not in itself objectionable, and to be distinguished from the unnecessary and objectionable practice of judging the veracity of those beliefs, that these anti-sceptical philosophers feel compelled to defend common sense against the specific threat of philosophical scepticism. For, correctly understood, these distinctions (which must be correctly understood if the very object of the sceptic's agnoilogical moorings is to be identified) reveal that the sceptic's arguments do not threaten common sense, any more than they threaten the 'googoos' and 'gagas' of newborns. These basic *first-order* beliefs are not *in themselves* of any interest to the sceptic; for such beliefs are based not on reason but nature. Moreover, they are not threatened, and cannot be threatened by scepticism, because they are not judgments. It is

only judgments whose epistemic credentials can be impugned. The object of scepticism is not the natural first order-beliefs which comprise the content of common- sense, but rather, scepticism's object is judgments; judgements *about* the accuracy of those (and other) beliefs.[24] An accurate understanding of what scepticism does and does not imply, as well as an accurate understanding of the nature and contents of common sense, reveals, therefore, the complete inanity and futility of each of these contemporary anti-sceptical refutations.

In the following chapters, I argue that neither these distinctions nor their consequences have been clearly understood. Had they been understood, we should not in the last century have had such a plethora of defenses of common sense, defenses erected against the particular (and supposed) threat of philosophical scepticism. Equally, we should not have had the particular *variety* of sceptical refutation so widely embraced in the twentieth century; that which rests its entire refutation on mere facts of ordinary behaviour, ordinary belief, and ordinary language. For the sceptic, I argue, is not concerned with how we *happen ordinarily* to behave, speak, or believe; at least not directly. Rather, she is concerned with how we *ought* to behave, speak and believe; and then only with respect to the second-order, for it is only here that epistemological criticism makes sense. Consequently, how we speak and think is relevant *only* to the extent that it is revelatory of the second-order.

Of course, the ordinary language philosopher, as well as her heirs, have *assumed* that

---

[24] I can, for example, (not to mention, I most assuredly do) believe that there is an external world, a world that it is the kind of world in which there are objects, people, and planets, and I can do so without thereby believing *that I know* there is such a world.

how we happen to speak *is* relevant to the question of whether we have knowledge. So too they have assumed that the ordinary is relevant to the question of scepticism's truth or falsity. Equally, therefore, they have assumed that the irrelevance of the sceptic's alternatives is a *fait accompli*. All that remained to be done, they thought, was to explain how exactly the sceptic manages to get our attention, to seem so *nearly* right. The answer, they argued, lay in the manner in which the sceptic raises, manipulates, or otherwise distorts the standards of knowing ordinarily in place.

The present work is in part a critical inquiry into the conditions which make possible this particular variety of anti-sceptical refutation. I argue that these conditions reside not in the nature of things, but rather in a number of philosophical (and dogmatic) assumptions. If what follows achieves nothing else, I hope that it at least makes clear that there are a number of clearly identifiable *assumptions*, and that these are necessary to support the variety of anti-sceptical arguments I discuss. Moreover, because each of these assumptions is neither necessary nor otherwise compelling, I argue they should be rejected.

My inquiry reveals that the kind of anti-sceptical inference made by the ordinary language philosopher, the relevant alternativist, and the contextualist – from the ordinary to the epistemic – rests on a number of misunderstandings. Among those misunderstandings is the already mentioned assumption that the ordinary carries epistemic and semantic relevance while the extraordinary does not (or, at least not much). It is this assumption, for example, which allows not only the ordinary language philosopher, but also the relevant alternativist, and the contextualist to conclude that the sceptic's departure from the ordinary suffices to refute (for all practical purposes) her agnoiological

conclusions. This conclusion rests, however, on an additional misunderstanding of the nature and contents of common sense. It rests on a misunderstanding of the *kind* of standards that govern common sense. It rests also, therefore, on a misunderstanding of the nature and significance of the disparity between the sceptic's queries and our ordinary linguistic habits, the standards of knowledge she cites, and the relation of both to the claims of common sense.

A further, and perhaps not less significant, result of my inquiries is the discovery that the relevance of the Cartesian alternatives to knowledge *may* be much greater than these anti-sceptical views have acknowledged. Like the ordinary language philosopher, the contextualist, and the relevant alternativist, however, I too conclude that the Cartesian alternatives to *knowledge* are of *little significance* for common sense, though not at all for the reasons these others have given. Rather, it is because the presumed correlation between defeating the sceptic and saving common sense is fundamentally mistaken. This correlation is mistaken because it falsely assumes that the standards that govern common sense are of the same kind as the standards that govern knowledge. It is mistaken, moreover, because it falsely attributes to the contents of common sense second-order knowledge beliefs which are, I argue, anything but common. It is mistaken because it confuses the kinds of beliefs we do ordinarily hold with the kind of beliefs that we do not; while it is the latter alone which the sceptic means to confute. And because it is the latter kind of beliefs the sceptic means to refute, common sense beliefs remain unscathed by scepticism, even if the sceptic is correct about the kinds of justification necessary for knowledge.

There are, of course, different varieties of sceptical argumentation. Equally, there are different varieties of anti-sceptical argumentation. In this work I am interested only in that variety of scepticism which argues that because there are always ineliminable alternatives to what we claim to know all knowledge is impossible. According to this sceptical argument, the logical possibility of these alternatives, coupled with their ineliminability, suffices to make them relevant. What kinds of alternatives? Alternatives like that I may now be dreaming; or, that unbeknownst to me, I may be the unwitting victim of an evil and cunning genius; or that, again unbeknownst to me, I may be an envatted brain having the experiences I do have not because of an orderly external world, but rather, because of the stimulations my envatted brain receives from the super-computer to which it is connected.

Of course, what each of these alternatives has in common is the fact that each describes a logically possible alternative to what we ordinarily believe is the cause of our experience. More importantly, each of these alternatives describes a state of affairs that we cannot eliminate. Thus the alternatives of philosophical significance are those that are not only *logically* possible, but also, ineliminable. What makes these alternatives ineliminable is, as the Pyrrhonian discovered long ago, the fact that we lack any criterion by which we may *judge* the one to be the true cause of our experience rather than the other. Lacking any such criterion, the sceptic concludes we must suspend our judgement regarding the actual nature of the cause of our experience.

Just as there are a variety of kinds of scepticism, so also, of course, there are any number of kinds of *anti*-scepticism. Again, however, I am not here concerned (nor would

it be practical to be concerned) to refute *every* sort of anti-sceptical argument. Rather, in what follows, I confine my attention to just one variety of anti-sceptical argument – that which infers from the sceptic's mere linguistic oddity her epistemological irrelevance – and to three of its particularly important representatives: the ordinary language philosopher, the relevant alternativist, and the contextualist.[25]

There are three primary reasons I confine my attention these kinds of anti-sceptical argument: first, each argues specifically against the kind of sceptical position I am concerned with; second, each confuses the first-order beliefs which are both necessary and constitutive of common sense for the second-order beliefs which are neither; and third, each appeals to the same kind of argument – the argument from ordinary language – in its attempt to refute the sceptic and to save common sense.

Like the defenders of the religious norm during the renaissance and modern period, these contemporary anti-sceptics felt compelled to safeguard the doxastic norm; they felt compelled, that is, to safeguard the contents of common sense. Novel, however, to their twentieth century approach, is their defense of what we ordinarily believe *simply* on the grounds that it *is* ordinary. Of course, they do not argue explicitly from the mere fact that our ordinary beliefs are ordinary that they must therefore be correct. Nonetheless, ultimately the argument of each – surface complexities and dissimilarities aside – does

---

[25]It might be thought that I should address Malcolm's or Putnam's radical anti-scepticism which claims that scepticism is self-defeating, that the "dreaming" and "brain-in-the-vat" alternatives can't get off the ground. There are at least two reasons why I am not addressing this radical anti-scepticism: first, while I am confident their arguments, such as Putnam's which assumes a causal theory of reference, fall prey to the Pyrrhonian reply that they presuppose question-begging criteria, it would take another thesis to do justice to their challenge. Second, the anti-sceptical adversaries I am addressing do not regard these standard sceptical alternatives as self-defeating or self-refuting. In fact, they take them seriously enough to advance extensive arguments against their relevance.

reduce to this form.

Chapters II, III, and IV of this work are largely devoted to establishing that in spite of the obvious surface differences in their arguments, the ordinary language philosopher, the contextualist, and the relevant alternativist's arguments do in fact reduce to the same kind of argument. To establish this, I argue that each appeals to the same fundamental principle in their anti-sceptical refutations. The principle to which they appeal states that any challenge to a knowledge claim which violates ordinary linguistic practices is irrelevant to the truth of that knowledge claim. In chapter II, I lay the foundations for the critique of this mode of argumentation. I there argue not only that this principle is mistaken, but also, that to argue from mere linguistic oddity to epistemic irrelevance is transparently fallacious. Why? Because to the extent that one assumes the ordinary is revelatory of the epistemic, one conflates two *kinds* of standards which are in fact *logically* distinct: the standards governing appropriate assertion, and the standards governing true assertion.

In the second chapter I also give numerous illustrations that demonstrate that these kinds of standards are in fact distinct, and must be distinguished if either common sense or scepticism is to be clearly understood. I argue that to the extent that the standards are distinct *in kind*, the proof that the sceptic violates the one (that of ordinary assertion) is *not* sufficient to prove she violates the other (of true assertion). And insofar as each of these philosophers establishes only that she violates the standards of ordinary assertion, there arguments do nothing to settle either the question of the truth or falsity of the sceptic's conclusions, nor do they settle the truth or falsity of the contents of common

sense. I conclude that one cannot refute the sceptic merely by pointing to the conventionality of everyday belief, nor by pointing merely to the absence (or violation) of conventionality in the alternatives to which the sceptic appeals.

With respect, then, not only to the relevant alternativists, but also to the contextualists and to the ordinary language philosopher, the truth of the sceptic's claim that knowledge is unattainable comes down to a question of the *relevance* of her alternatives. Moreover, the relevance of those alternatives comes down to how well they fit in with our ordinary linguistic practices. But why should their relevance be determined by our ordinary practices? Why should a question of truth be determined by a matter of convention?

In my fifth and final chapter I argue that the entire project of protecting the ordinary rests on a fundamental misunderstanding of what the ordinary entails. In particular, I argue, it rests on the confused, but widely embraced, belief that the contents of common sense include not merely first-order beliefs about the world, but also second-order beliefs about those beliefs. That is, it rests on the conviction (largely unexamined) that common sense includes not only the belief that there is an external world, that one has hands, and so on, but also the additional set of corresponding beliefs that each of these first-order beliefs is known to be true. For it is only if common sense does include such second-order beliefs, such judgements about the knowability of the first-order, that common sense is (or can be) threatened by scepticism.

I argue, however, that this view of common sense is in error. One does not in fact ordinarily believe, in addition to each of one's first-order beliefs about the world, that each of those beliefs are known. Rather, one simply believes. Moreover, what concerns

the sceptic is not what we believe about the world, but rather, whether we can know what we believe. To the extent, therefore that the ordinary man's common sense does not include such *second-order* beliefs, to the extent that it does not include beliefs about what we can know, to that same extent are the contents of common sense protected from the sceptic, and the refutations of radical scepticism (*on behalf of* common sense) unnecessary.

Why, one might ask, would the likes of Austin, Lewis, Dretske, Goldman, and DeRose be mistaken about the nature and contents of ordinary belief? How, after all, could one get the *ordinary* wrong? I believe this mistake is explained by a philosophical confusion over the kinds of standards which govern what we ordinarily say and do, and the kinds of standards which govern truth. It is a confusion, that is, of the standards of good manners with those of good truths. This confusion, I believe, is the result of a further confusion between the strength of our natural convictions and the kind of certainty required for knowing. Had due attention been given to the Humean (indeed the sceptical) insistence on the natural necessity of (first-order) belief, such confusion about both the object of scepticism, and the sceptic's conception of ordinary belief, might have been avoided. For, again, it is not mere unavoidable belief about the world that the sceptic attacks. Rather, it is judgement about belief. Only in judgement do we find the kind of dogmatism that the sceptic attacks. Indeed, from the earliest systematic scepticism, the object of scepticism has been explicitly defined as the dogmatic beliefs of the philosopher. Thus the object of the earliest scepticism was directed at the claims of the Eleatics, the Pythagoreans, the

stoic, the Epicurean, and the Platonist, for example.[26]

The goals of this work are, therefore, threefold: first, to establish that each of these major anti-sceptical movements is reducible to an argument from ordinary language; second, to establish that these arguments from ordinary language are fundamentally flawed; and third, to establish that because the ordinary language philosopher, the contextualist, and the relevant alternativist misconstrue both the nature of common sense and the object of scepticism, they fail to refute the sceptic and to save common sense.

In addition to these thesis, I conclude that even if we grant to the ordinary language philosopher, the contextualist, and the relevant alternativist that common sense does include second-order knowledge beliefs, the "knowledge" they allow us to claim is completely unsatisfactory. For the kind of knowledge their theories afford us rests essentially on ignoring what is difficult, presupposing what cannot be established by argument, and begging the very question at issue. Their kind of knowledge is indistinguishable from dogmatic assumption. It requires neither truth, nor justification. This philosopher finds such knowledge deeply unsatisfactory. For this philosopher still believes that knowledge is different from mere belief, that knowing requires more than mere inter-subjective agreement, and that winning an argument still requires the giving of good reasons. In so far as "knowledge" is (by their lights) a matter of presupposing, ignoring, and *ad hoc* excepting, I conclude then that even if we can "know" according to *these* contextually sensitive standards – such a victory for "knowledge" is cheap indeed. For to overcome uncertainty by presupposing, ignoring, and *ad hoc*

---

[26]See, for example, the *Outlines*. Sextus Empiricus makes clear throughout the *Outlines* that the subject of sceptical criticism is not those beliefs we are compelled to hold, but rather, those non-compulsory beliefs held by the philosopher, whom he would call "the dogmatist".

excepting is not to give a philosophical defense of knowledge, nor is it to refute the sceptic. Rather, it is to re-embrace the very dogmatism – that is, assent to the non-evident claims – that centuries of philosophical scepticism has been designed to end.

CHAPTER II:
Meanings, Manners, and Scepticism

In this chapter, I examine the relation between meanings, manners, and scepticism. In particular, I examine the effectiveness of the ordinary language[1] response to scepticism. I argue that the ordinary language argument against philosophical scepticism is fundamentally mistaken, and does not, therefore, refute scepticism. For the ordinary language refutation[2] of scepticism confuses good manners for good truths, and it is only to the extent that one overlooks this confusion that one is tempted by its anti-sceptical implications. Once, however, that confusion is made explicit, it is clear that the argument from ordinary language poses no serious threat to — let alone *refutation* of — the sceptic.

The kind of philosophical sceptic I examine (I suppose there can be many) in what follows is the one who (like Descartes at the end of his First, and the beginning of his Second, *Meditation,*) argues that there is nothing we can know: Not that the flowers in the garden are in bloom, not that there are birds in the garden, nor indeed that there is a garden.

---

[1]What I *intend* by 'ordinary language' is somewhat narrow, and shall become clear shortly. Suffice it for now to say that I intend the following: the position that ordinary language is correct and can be used to resolve philosophical problems; more specifically, with respect to scepticism, *any* proposed refutation of scepticism which argues from the mere linguistic oddity of the sceptic's remarks to their supposed epistemological irrelevance or falsity. (As will become clear, there are reasons to include both the relevant alternativist, and the contextualist within this genre.) My remarks throughout apply only to this variety of ordinary language.

In particular, I am interested in the philosophical sceptic who argues for this conclusion on the grounds of competing *possible* descriptions of the etiology of our experience. It is the sceptic who argues, for example, that my experience of the flowers in bloom *may be* the result of a garden in bloom, but *may also* be the result of the machinations of an evil and all powerful genius who delights in deceiving me. It is the sceptic who argues that my experience of the birds in the garden *may be* caused by birds in a garden but *may also* be caused by the stimulation received by my envatted brain from a super-computer to which it is attached. It is the sceptic who argues that my experience of the garden itself *may be* caused by a garden, but *may also* be caused by a dream I am now having.

> But in thinking over this I remind myself that on many occasions I have in sleep been deceived by similar illusions, and in dwelling carefully on this reflection, I see so manifestly that there are no certain indications by which we may clearly distinguish wakefulness from sleep that I am lost in astonishment. And my astonishment is such that it is almost capable of persuading me that I now dream.[1]

It is the sceptic who argues that because the content of my experience is equally compatible *with* these alternatives, yet incompatible with what I ordinarily believe, I can never come to know. *This* is the sceptic who interests me, and whose purported refutation by ordinary standards shall be the focus not only of this chapter, but of the chapters to come.

What is it about such competing possibilities that grounds the sceptic's systematic doubt? His conclusion that we cannot escape eternal ignorance? Why are these

---

[1]Descartes, *Meditation I*, 146.

competing possibilities grounds for concluding that I — and you — can never know? The sceptic argument is this: *knowing* has certain *necessary* conditions. Among these necessary conditions is the condition of certainty. Because certitude is a *necessary* condition of knowing,[2] it follows that if one is uncertain of what she believes, she fails to know.[3] Certainty is a condition of knowing that the flowers in the garden are in bloom. It is a condition of knowing that the birds are in the garden. It is a condition of knowing that there is a garden. Indeed, certainty, the sceptic argues, is a condition of knowing *anything* at all.

The sceptic, however, thinks that the *existence* of competing possibilities (the issue of all likelihood or probability aside) undermines certainty. Why? What has certainty to do with dreaming, vatting, and evil and cunning genius? The sceptic argues that certainty requires the *complete* absence of doubt. Here it is useful to distinguish two kinds of doubt: on the one hand, there is the doubt of the subject's state of mind, what we might call 'subjective doubt'; on the other hand, there is the fact of the matter – whether there exist factors which make something dubious – what we might call objective doubt. What kind of doubt does the sceptic require be absent? Not mere subjective doubt, for it is easy enough to fail to doubt what *seems* perfectly ordinary. Rather, it is objective doubt which cannot co-exist with certainty, and objective doubt follows from objective possibility. Where objective doubt exists, so also does (objective) uncertainty. And where uncertainty reigns,

---

[4]I will not, in this work, look into the further issue of which conditions are sufficient for knowledge. My present concern is simply the sceptic's claim that ruling out the Cartesian alternatives to one's knowledge claims is a *necessary* condition of knowing anything at all.

> I suppose, then, that all the things which I see are false; I persuade myself that nothing has ever existed of all that my fallacious memory represents to me; I consider that I possess no senses; I imagine that body, figure, extension, movement, and place are but fictions of my mind. What, then, can be esteemed as true? Perhaps nothing at all, unless that there is nothing in the world that is certain.[2]

In particular, I am interested in the philosophical sceptic who argues for this conclusion on the grounds of competing *possible* descriptions of the etiology of our experience. It is the sceptic who argues, for example, that my experience of the flowers in bloom *may be* the result of a garden in bloom, but *may also* be the result of the machinations of an evil and all powerful genius who delights in deceiving me. It is the sceptic who argues that my experience of the birds in the garden *may be* caused by birds in a garden but *may also* be caused by the stimulation received by my envatted brain from a super-computer to which it is attached. It is the sceptic who argues that my experience of the garden itself *may be* caused by a garden, but *may also* be caused by a dream I am now having.

> But in thinking over this I remind myself that on many occasions I have in sleep been deceived by similar illusions, and in dwelling carefully on this reflection, I see so manifestly that there are no certain indications by which we may clearly distinguish wakefulness from sleep that I am lost in astonishment. And my astonishment is such that it is almost capable of persuading me that I now dream.[3]

It is the sceptic who argues that because the content of my experience is equally

---

[2]Rene Descartes, *Meditations on First Philosophy,* Trans. Elizabeth Haldane and G.T. Ross, *The Philosophical Works of Descartes* (Cambridge: Cambridge University Press, 1911), 149.

[3]Descartes, *Meditation I,* 146.

compatible *with* these alternatives, yet incompatible with what I ordinarily believe, I can never come to know. *This* is the sceptic who interests me, and whose purported refutation by ordinary standards shall be the focus not only of this chapter, but of the chapters to come.

What is it about such competing possibilities that grounds the sceptic's systematic doubt? His conclusion that we cannot escape eternal ignorance? Why are these competing possibilities grounds for concluding that I — and you — can never know? The sceptic argument is this: *knowing* has certain *necessary* conditions. Among these necessary conditions is the condition of certainty. Because certitude is a *necessary* condition of knowing,[4] it follows that if one is uncertain of what she believes, she fails to know.[5] Certainty is a condition of knowing that the flowers in the garden are in bloom. It is a condition of knowing that the birds are in the garden. It is a condition of knowing that there is a garden. Indeed, certainty, the sceptic argues, is a condition of knowing *anything* at all.

The sceptic, however, thinks that the *existence* of competing possibilities (the issue of all likelihood or probability aside) undermines certainty. Why? What has certainty to do with dreaming, vatting, and evil and cunning genius? The sceptic argues that certainty requires the *complete* absence of doubt. Here it is useful to distinguish two kinds of doubt: on the one hand, there is the doubt of the subject's state of mind, what we might call 'subjective

---

[4]I will not, in this work, look into the further issue of which conditions are sufficient for knowledge. My present concern is simply the sceptic's claim that ruling out the Cartesian alternatives to one's knowledge claims is a *necessary* condition of knowing anything at all.

doubt'; on the other hand, there is the fact of the matter – whether there exist factors which make something dubious – what we might call objective doubt. What kind of doubt does the sceptic require be absent? Not mere subjective doubt, for it is easy enough to fail to doubt what *seems* perfectly ordinary. Rather, it is objective doubt which cannot co-exist with certainty, and objective doubt follows from objective possibility. Where objective doubt exists, so also does (objective) uncertainty. And where uncertainty reigns, so also does ignorance.

What, then, is 'objective possibility'? What is *objectively* possible? Anything, I should think, which does not involve a logical contradiction. It is, for example, objectively *possible* that I am now dreaming, however unlikely it may be either that I am, or that I will (subjectively) seriously entertain the possibility. Equally, it is objectively *possible* (that is, there is no contradiction in supposing) that I am an envatted brain, or a victim of an evil and cunning genius. Indeed, no matter what my peculiar subjective response to these possibilities is, they *are* possibilities. In order, then, for me to know what I believe, I must be able to establish that I am not merely dreaming, not an envatted brain, and not a victim of evil genius. Unless I can rule these out, I cannot be objectively certain. To the extent that I cannot rule them out, these alternatives undermine my certainty. They undermine my certainty because they are compatible with all my evidence and yet incompatible with what I believe. They cannot be ruled out because from the contents of my experience I cannot discriminate between the possibilities. For all I know, I am now dreaming, or a brain in a vat, or a victim of ingenious deception.

There are, however, those who reject this harsh evaluation of our cognitive powers.

Among the most strongly opposed are the defenders of ordinary language. These guardians of the linguistic norm regard the sceptic as more of a trickster than a serious epistemological rival. The sceptic, they think, is seduced into disbelief by estrangement from ordinary linguistic practices. They believe, moreover, that proving that it is (perfectly) *ordinary* to *say* we know suffices to refute the sceptic. For what is ordinary, they argue, is manifestly correct.

In what follows, I offer a critical examination of this argument from ordinary language. I argue that the ordinary language refutation of scepticism rests on the uncritical *Principle of Epistemic Conventionality:* the principle according to which we may infer the irrelevance of the sceptic's doubts simply from their unconventionality.[5] I argue that this principle which constitutes the foundation of this kind of argument is ambiguous. On one interpretation the principle is true, but ineffective; on the other interpretation, it is false. According to that interpretation which makes it true, it is trivial. On the other hand, according to that interpretation which makes it a weapon against scepticism, it is false. In fact the principle *seems* effective against the sceptic only when one overlooks its (illegitimate) conflation of two logically distinct sets of standards; the standards of appropriate and the standards of true assertion.[6]

In this chapter, I offer numerous illustrations of the logical separability of these standards, and conclude that the ordinary language advocate believes that mere facts about

---

[5]As I will argue in chapters III, IV, and V, the arguments of the relevant alternativist and the contextualist also rest on this principle.

[6]My understanding of this distinction as well as the work it can do in defending the sceptic was *very much* influenced by Barry Stroud's *The Significance of Philosophical Scepticism* (Oxford: Clarendon Press, 1984): in particular chapter II.

ordinary language can refute the sceptic only because the ordinary language advocate overlooks the logical independence of two kinds of standards. The ordinary language advocate believes falsely that the sceptic's challenges to our knowledge claims are irrelevant to their truth, but only because she believes — also falsely — that the standards of appropriate assertion (which underlie our knowledge claims) are identical with the standards of truth (which do not).

Section i: The Sceptic

There are three essential elements in the argument of the (at least my) philosophical sceptic: First, the sceptic claims that objective certainty is a *necessary condition* of knowing; second the sceptic claims that ruling out those alternatives explanations of the cause of our experience which are *equally compatible* with that experience is a necessary condition of objective certainty; third, and finally, the sceptic claims that because we cannot eliminate those alternatives (nor, therefore, distinguish the true from the false with respect to these causal explanations), we cannot meet the conditions necessary for knowledge. In other words, we cannot be objectively certain. The ineliminability of the Cartesian alternatives undermines the possibility of objective certainty, and therefore, the possibility of knowledge.

The sceptic, of course, at least the modern sceptic[7], does not expect you to give up your ordinary beliefs about the world (that the world which is the cause of your experience is

[7]As a guide to living well, the ancient sceptics approach the issue of suspending belief somewhat differently. See especially Sextus Empiricus, *Outlines of Pyrrhonism*, Book I, Trans. R.G. Bury.

an external world, that the objects of your experience exist, and that you are among those objects). The abandonment of ordinary belief is not the point of the sceptic's reasoning.[8] (And so, objecting to the sceptic on the grounds that it would be impossible to get by without these *beliefs* is quite without point.) Rather, the sceptic expects only that you abandon your *judgement* that you *know* what you believe; for that judgement, he argues, must, inevitably, prove false. This is not to say that *what* you believe (e.g., *that* there are trees outside your window) is false. Rather, it is only to say that your belief *that* you *know* (if you possess such belief) is false, false according to the objective standards of knowing.

Take, for example, my claim to *know* certain things about my dog Beau. Beau, at the moment, appears to be a real dog. Indeed, he appears to be a perfect specimen of the only breed which can boast of being "gentlemen in white fur"; a Pyrenees, that is, of Siberian ancestry and majestic appearance. And, as is characteristic of his breed, he has soft dark eyes, and appears to be drooling (somewhat unmajestically) as he lounges about. The question for the sceptic — the question the sceptic wishes us to entertain — is this: appearances and beliefs aside, what do I *know*? Do I *know* these things which *appear* to me to be so obviously the case? What would it take for me to *know*, for example, that I have a dog, of Pyrenean origin, with soft dark eyes, and incessant (and unmajestic) drool?

Of course, undeniably — indubitably — it *seems* perfectly evident that I have a dog, a real dog, that his name is Beau, that he lounges about, and so on. Indeed, at times I even

---

[8] One might think the sceptic does argue we should abandon belief. Strictly speaking, however, the sceptic argues only against judgement about the veracity of belief, not belief in itself. For further explanation of this often overlooked feature of the sceptic's actual intentions, see chapters I and V.

say of myself *that* I *know* these things. (As, for example, when I am asked how I know he is a real or pure-bred Pyrenees, I point to his double dew claws.) But what do I *know*? Do I *know*, for example, that what appears now to be a dog *is* a dog? Do I *know* that I have a dog? Do I *know* that I now sit in a room, and he in it with me? Do I *know* that the clear liquid which appears to run from his mouth is drool? Real drool? Do I *know* his double dew claws are real? Do I know this dog, this room, this drool, these dew claws are not merely dreamt, vatted, or the results of computer stimulation? The sceptic argues I do not — *and can not* — no matter how unwilling I may be to withdraw my beliefs. For knowing, the sceptic reminds me, requires certainty. And so long as I cannot prove I am not now dreaming, so long as I cannot prove that I am not now vatting, so long as I cannot not prove that I am not now the unwitting victim of a cunning and evil genius, certain is something I am not. Q.E.D.

Might it not be that the room I believe I am in is really only a figment of my, or someone else's, imagination? Might the room-appearance not be caused by a dream I am now having? Might not this hypothesis explain my experience equally well? (Descartes' familiar argument that there are no convincing signs distinguishing waking and dreaming proves powerful here.) Or, might not the appearance of my room be caused by a super computer which, in virtue of its connections to my envatted brain, stimulates this appearance in me? How would the experience I am now having *differ* from the experience I *would* have if either of these alternative possibilities obtained? The sceptic's answer is: *not at all*. Indeed, that is precisely the sceptic's point, not to mention the well-spring of the sceptic's power.

And what of my perception of Beau, my (apparent) dog? I see what I take to be his soft eyes, his clean white fur, and yes, his drool. But might not all of this too be caused by a dream I am now having? "How often in sleep have I *dreamt of Beau* . . ."? Might not these perceptions be the deliverances not of my eyes, but of my artificially stimulated and envatted brain? The answer, of course, is yes. They *might* be. After all, these alternatives are objectively possible.[9] And moreover, and perhaps most importantly, each element of my present experience would not differ, would not differ at all, in any detail, if it were caused not by a real room, dog, eyes and drool, but by a dream, a computer, or an evil (a *very* evil) genius. For, *ex hypothesi*, the content would — in each case — be *precisely* the same; the same as it would be if the cause of the content were the things I take to be real. I would perceive my familiar room, with all of its familiar furnishings, the same architectural ornaments. I would perceive the same white fur, dark eyes, and drool. And, importantly, I would believe the same things.

From such objectively possible, and yet competing, perceptual equivalents follows the doubt from which the sceptic derives her claim that we cannot know. Though, of course, it is *only if* these competing alternatives are also ineliminable that they threaten my claims to know; as noted, it is the fact that they are ineliminable that undermines my certainty, since it is the extent to which I cannot know I am n*o*t now dreaming, vatting, or deceived, that the certainty required for knowing cannot be mine.

---

[9]Because there is general agreement amongst the philosophers that I will discuss that — according to the sceptic's standards of elimination — these are both possible and ineliminable, I will not spend time arguing that they in fact are. Instead, because the real question becomes whether or not objective possibility and ineliminability stand in the way of knowing, my focus will be on this question, and how one might go about answering it.

One may, of course, object that (at least some) such alternatives really *are* eliminable, *can* be refuted by ordinary evidence, and hence cannot threaten what I (seem to) take myself to know. Indeed, within the ordinary language camp, there are a number of such protestations; Austin's is a paradigm example. Distinguishing waking states from dream states, he argued, is not mysterious at all. One need merely appeal to well established methods: "There are recognized ways of distinguishing between waking and dreaming".[10] Unfortunately, Austin does not tell us what these methods are, nor how they answer the sceptic. One can assume, however, that the methods in question are the methods one ordinarily employs when one checks to see if one's child is asleep, or even if one is oneself asleep (methods like seeing if one's eyes are open, if one is dressed, if the sun is up, and so on).

The problem, of course, is that one cannot assume the *veracity* of such methods without begging the question against the sceptic. The sceptic, no doubt, will protest that using any such methods to verify one's belief that one is awake assumes the veracity of the very thing in question, namely, one's perceptions. One cannot -- without circularity -- appeal to these *perceptions* to justify the conclusion that one is not asleep; these perceptions might themselves be caused by dreams, vatting, or evil and cunning genius. In the absence of a criterion by which one can prove one is not dreaming, vatting, or the victim of ingenious deception in the first place, it will be of no (non-question-begging) use to appeal to these ordinary methods to show that one isn't. It is, after all, precisely the ordinary that

---

[10]J.L. Austin, "Other Minds", *J.L. Austin Philosophical Papers*, ed. J.O. Urmson, and G.J. Warnock, (Oxford: Oxford University Press, 1979) 87.

is in question.

Moreover, long before the likes of Descartes, the ancient Greeks[11], recognized the problem of the criterion of the criterion.[12] Even if one had a criterion for ruling out having merely dreamt, or believed by deceit, how could one know *this* criterion is not itself the result of dreaming, or deception? The problem is quite obvious, and answering it lands us either in an infinite regress, a circular justification, or an arbitrary stopping point.[13] As Sextus Empiricus states in his *Outlines:*

> . . . he who prefers one impression to another, or one 'circumstance' to another, does so either uncritically and without proof, or critically and with proof; but he can do this neither without these means (for then he would be discredited) nor with them. For if he is to pass judgement on the impressions he must certainly judge them by a criterion; this criterion, then, he will declare to be true, or else false. But if false, he will be discredited; whereas if he shall declare it to be true, he will be stating that the criterion is true either without proof or with proof. But if without proof, he will be discredited; and if with proof, it will certainly be necessary for the proof also to be true, to avoid being discredited. Shall he, then, affirm the truth of the proof adopted to establish the criterion after having judged it or without having judged it? If without judging, he will be discredited; but if after judging, plainly he will say he has judged it by a criterion; and of that criterion we shall ask for a proof, and of that proof again a criterion. For the proof always requires a criterion to confirm it, and the criterion

---

[11]Sextus Empiricus, *Outlines*, Book IV, 67-7. Since Sextus was primarily a compiler of the sceptical arguments of others, this argument predates the *Outlines*.

[12]Stroud offers an effective discussion of essentially this problem in his *Significance of Philosophical Scepticism*, though he leaves its ancient origin out of the discussion.

[13]For further reading here, the debate between the foundationalists and coherentists is of course relevant. See, for example, Lawrence Bonjour, "A Critique of Foundationalism," in *The Theory of Knowledge: Classic and Contemporary Readings*. Ed., Louis P. Pojman (Belmont, CA., Wadsworth, 1993).

> also a proof to demonstrate its truth; and neither can a proof be sound without the previous existence of a true criterion nor can the criterion be true without the previous confirmation of the proof. So in this way both the criterion and the proof are involved in the circular process of reasoning, and thereby both are found to be untrustworthy; for since each of them is dependent on the credibility of the other, the one is lacking in credibility just as much as the other.[14]

Even if Austin (and Moore, and Bouwsma,[15] and countless others) were correct in thinking that to distinguish wakefulness from sleep is a simple, *ordinary*, matter, what about vatting ? What about an evil and cunning genius? Surely these are no *ordinary* matters. What kind of criterion could possibly enable me to distinguish these possible causes from the causes I think are responsible for my experience? The evidential states are identical. *For, ex hypothesi*, my vision of Beau can be cleverly produced in my envatted brain by a super computer. Even the calm look of his soft dark eyes, the clean white fur, even the drool. All of these perceptions *might* be caused by clever wiring which stimulates precisely those brain states (and thereby simulates the experiences) I would have if Beau were a real dog with real drool. Likewise, each of these perceptions might be brought about in me through the clever machinations of an evil genius, leaving me no perceptual or evidential difference; *might*, that is, in the sense that there is no logical contradiction involved in supposing it to be so.

---

[14]Sextus Empiricus, *Outlines* Book I, 68-9. (This argument is also found throughout the *Outlines*.)

[15]See, for example, G.E. Moore, "Proof of an External World," *G.E. Moore: Philosophical Papers* (New York, Collier Books, 1962); and O.K. Bouwsma, "Descartes' Evil Genius." *Bouwsma's Philosophical Essays* (Lincoln, University of Nebraska, 1965).

> Consequently, if a man can prefer one impression to another neither without a proof and a criterion nor with them, then the different impressions due to the differing conditions will admit of no settlement; so that as a result of this Mode also we are brought to suspend judgement regarding the nature of external realities[16]

Section ii: The Argument From Ordinary Language

Admittedly, the sceptic's claim that knowing that we are not dreaming, vatting, or victims of evil and cunning genius is a necessary condition of knowledge is odd. In fact, of course, it seems perfectly absurd. Is it not absurd to think that I fail to know of the existence of my dog because it is *possible* that I *might* be deceived in one of these ways? Because my perception of him *might* be caused by a super-computer, or an evil genius? How often, after all, do I thus scrutinize my evidence before I give Beau his daily treat? How often do I stop to consider the competing possible etiologies of my experience before I wipe away his drool? How often does anyone scrutinize the basis of her certainty before claiming to know? Absurd indeed! Indubitably so.

In what follows I examine the view that the sceptic's hypothesis is absurd; her hypothesis, that is, that because I *might* now be deceived in some undetectable way, I cannot know. Specifically, I examine the *inference* from the sceptic's mere absurdity to the conclusion that her assertions must also be irrelevant or false. The proponent of this view argues that because of the degree to which the sceptic violates ordinary language her doubts are irrelevant. Moreover, she claims, the sceptic's violation of ordinary language is sufficient to prove her claim that we fail to

---

[16]Sextus, *Outlines*, Book I, 69.

know is false, her doubts innocuous.[17]

J. L. Austin's "Other Minds" is paradigmatic of the kind of anti-sceptical argument I call the argument from ordinary language. I shall, therefore, refer to his 'refutation' throughout. Austin argues that doubting our ordinary claims to know, for Cartesian-like reasons, is inappropriate. More importantly, he argues that the sceptic's raising of the Cartesian alternative possibilities to what we know is itself inappropriate. Austin supports this contention with some observations about our *ordinary* linguistic *practices*. Of particular interest are his remarks about what *ordinarily* happens when one makes a claim to know.

Austin observes, for example, that ordinarily when a claim to know has been made, a number of rules or conventions are observed. The claim to know that "There is a goldfinch in the garden," for example, is one such knowledge claim, and — ordinarily — when it is made, a number of linguistic conventions are observed. For instance, unless one has some specific reason to doubt what has been claimed, one accepts it. And, moreover, one does not doubt unless one has in mind, "and can specify," some definite lack.[18]

If, for example, I have a reason to doubt your claim to see a bird in the garden, I might point out that the evidence you offer (e.g., that what you perceive has several

---

[17]Representatives of the view I have in mind here include Austin, Wittgenstein, and Malcolm. No one of them individually represents this view in whole. (Austin, for example, does not explicitly state that the sceptical thesis is false, at least not in the material we look at, though he strongly suggests, if not implies, this. And Wittgenstein too would hesitate to say that this view is false. Both he and Austin would likely argue it is senseless, rather than false. Malcolm, on the other hand, explicitly states this view is false. Yet the spirit of what I here call the ordinary language argument is an essential part of the work of each.

[18]Austin, "Other Minds," 84.

characteristics unique to the *Carduelis*) is compatible with the perceived things being a species of large moth, or some other winged *non-Carduelis*. Moreover, having specified such a *definite* lack in your evidence, if you go on to explain why the perceived entity could not be other than you said, I would not — ordinarily — go on to insist on more. Instead, I would recognize that "enough" evidence had been given when, "relative to current intents and purposes,"the evidence "le[aves] no room for an alternative competing description of it."[19] Hence, if in response to my query, for example, you cite your training in ornithology, your specialization in distinguishing small birds from other winged non-birds, and your extensive field experience, it would be unreasonable and inappropriate for me to protest: "That is not enough." As Austin observes, "Enough is enough. It doesn't mean everything. . . . It does not mean enough to show it isn't a stuffed goldfinch," at least, not with respect to the conventions governing *ordinary* linguistic situations.[20]

Suppose, with these conventions in mind, that you and I travel abroad to England. Suppose, moreover, that while there we attend a garden party. In the garden, the trees appear to be in bloom, and several small species of birds appear to be perched about the foliage. Austin's view is this: were you to exclaim, in this context, "Look, there is a goldfinch in the greenery," my most appropriate response — assuming my familiarity with the conventions governing ordinary discourse — in the absence of any specific reason to

---

[19]Austin, "Other Minds," 84.This appeal of Austin's to rules implicit in ordinary conversation, as well as his emphasis on context, anticipates the appeal of many later Contextualists. See, for example, my discussion of Lewis in chapter III.

[20]Austin, "Other Minds," 84. (To Rich Hall I owe the following: Austin appears here to anticipate the denial of deductive closure; a denial found in a number of later thinkers; Nozick and Dretske, for example.) Austin also anticipates the contextualists and the relevant alternativists insofar as he argues against scepticism on the grounds that some alternatives to our knowledge claims are not relevant.

doubt your knowledge claim, is to express my agreement with you, or at the very least, to refrain from antagonistic challenge. I might, for example, say "Indeed!" Or, not myself being familiar with this particular species of British bird, I might query, "How do you know?" Or, "How can you tell?" Or, "How do you know it is not a woodpecker?" In either case, my response is *perfectly ordinary*, and, therefore, in making it I violate no obvious linguistic norms. Austin would approve.

The sceptic, however, asserts that neither you, nor I, nor anyone else can know anything at all. The sceptic reasons that you, for example, cannot know any empirical proposition because certainty is a condition of knowing, a necessary condition, and so long as you do not know you are not now dreaming, vatting, or deceived by a strange and malevolent genius, certain is something you are (certainly) not. Certainty requires the ability to rule out such alternatives, but -- lacking a criterion -- you cannot: whatever your evidence that there is a finch in the foliage, that same evidence supports the hypothesis that you are now dreaming that there is a finch in the foliage, dreaming that you and I are in a garden, dreaming that you are an ornithologist. Equally, your evidence that there is a finch in the foliage supports the further hypothesis that your brain is envatted. So also it supports the hypothesis that you are merely the unwitting victim of an evil genius of the Cartesian variety.

Suppose that the sceptic appears at the garden gate. Suppose further that, upon approaching, the sceptic overhears our discussion concerning the goldfinch and your claim to know. How might the sceptic respond? It is quite improbable that the sceptic would

accept the literal content[21] of your claim to *know* there is a goldfinch in the greenery, since you imply, thereby, that you are certain it is so. Rather, one imagines the sceptic, Pyrrhonian opposer that she is, querying thus: "Are you *certain* it is a *real* goldfinch?" "After all, for all you know ('which', slyly she murmurs, 'is nothing'), you *might* be dreaming you see a goldfinch, you *might* be vatting you are an ornithologist, you *might* be the unwitting victim of evil and cunning genius who deceives you into believing you are in a garden, when all the while you are not." In articulating her view that you cannot be certain that what you perceive is a *real* goldfinch, the sceptic here plants the seeds of doubt. For, even though you may steadfastly adhere to your belief that you do perceive a goldfinch, a real goldfinch, and that you are an ornithologist, it is clear that you do not know it is so.[22] As the sceptic argues, clearly you are *not* certain it is so, and cannot be certain unless you can also rule out the alternative possibilities regarding the etiology of what you now experience. Lacking certainty, you also lack knowledge; necessarily so.

What should we make of the sceptic's contention that your claim to know is ill-made? It seems that the reservations which ground *her* uncertainty differ from the more ordinary doubts expressed, for example, by my suggesting that the bird is perhaps a woodpecker, rather than a goldfinch. It is quite unremarkable for someone to challenge your knowledge claim by suggesting that perhaps you have mistaken one variety of bird for another. Contrariwise, to challenge your claim to know on the grounds that you may be

---

[21]The *intended* content has rather different implications from the *literal* content. This distinction is covered to some extent in Chapter IV's discussion of Dretske, but will be taken up in greater depth in chapter V.

[22]Whether you ought to *claim* that you know this to be so is a different matter.

dreaming, or vatting, or deceived by an evil and cunning genius, is extraordinary. Indeed, Austin suggests (as do the later contextualists and relevant alternativists), the degree to which the sceptic's alternative hypotheses depart from the norm suffices to make them not merely unusual, but irrelevant and false.

Austin reminds us that "knowing it is a *real* goldfinch is not in question in the *ordinary* case."[23] In the ordinary case — in the context of an ordinary conversation — one does not doubt the reality of a purported perception. One does not raise the possibility that the perception might be unreal. One certainly does not raise the possibility that the purported knower might be dreaming, vatting, or the victim of an evil genius. But why not? Why does one not raise such doubts in contexts like the garden party? And, furthermore, *of what significance* is it that one *ordinarily* does not raise such doubts?

With respect to the first question, *why* one ordinarily does not raise such doubts, Austin maintains that one does not raise these kinds of doubts because doing so would violate a number of linguistic conventions. What are these conventions? Austin argues that doubts concerning the reality of a perception — your perception of the goldfinch, for example, or my perception of Beau — can appropriately be raised (and seemingly therefore relevant[24]) only when two conditions are met:

> First, the reason for doubting must describe a *possible* state of affairs.
> Second, there must be a *reason to think* the possibility

---

[23]Austin, "Other Minds," 88. (My emphases.)

[24]The question of the relation of epistemic relevance to appropriateness is somewhat ambiguous in Austin, hence I say "seemingly". I address the importance of this ambiguity to the ordinary language refutation in section iii of the present chapter.

described actually obtains.[25]

These conditions are dictated by the norms of ordinary linguistic behavior.

Importantly, meeting the first condition is not (always) sufficient for meeting the second. Thus, though it may be *possible* that your perception of the finch in the garden is a causal result of stimulations to your envatted brain, there may be no reason to think it is so. Likewise, though it may be *possible* that my perception of Beau's soft dark eyes is the causal result of a dream I am now having, there may be no reason to think it is so. In the absence of such reasons, however, convention decrees that any challenge to my knowledge claims based on these Cartesian possibilities is inappropriate. Hence when I claim to know that Beau is near, you ought not, unless you have reason to think it so, challenge my claim on the grounds that I may be deceived in one of these ways.

The sceptic's suggestion that you *might* be dreaming, vatting, or deceived by an evil genius, can be raised only if she has a *reason to think* it is so.[26] She cannot (*appropriately*) challenge your knowledge claims on the grounds of mere *possibility* alone. Thus, if she is *to say* that you do not know that it is a (real) goldfinch causing what you perceive, because it *might* be stuffed, she must have a reason to think the goldfinch is stuffed. Indeed, Austin seems to imply, the sceptic's raising of the more important and specifically Cartesian possibilities that you may be dreaming, vatting, or in some other way

---

[25]Austin, "Other Minds," 87. See also 98-99. I should note here that with these two conditions, Austin anticipates central elements of both the Contextualists as well as the Relevant Alternativists. See, for example, my discussion of Stine in chapter IV, and also of Dretske's "Pragmatic Dimension of Knowledge," also in chapter IV.

[26]Austin, "Other Minds," 98-99.

deceived, are *irrelevant* to your claim to know; irrelevant, that is, unless she can specify some reason to think that you *are* in some such way deceived.

What follows, it seems, is that unless the sceptic can produce *a reason for doubting* (beyond mere possibility) the genuineness of your perceptual experience, she *cannot challenge* your claims to know: not your claim to know that the flowers in the garden are in bloom, nor your claim to know that the birds are in the garden, nor indeed your claim to know that there is a garden, a garden in which all this "blooming buzzing confusion" takes place. She is, Austin thinks, effectively silenced; silenced by the conventions which govern appropriate assertion.[27]

Return for a moment to the garden. Suppose we are sitting amongst what appear to be diverse species of trees, other flora, and birds. The lighting appears to be adequate. Our vision is good. We've not had much to drink. What *reason* could the sceptic have to think we might be dreaming? What *reason* could she have to think we might be vatting? What *reason* could she have to think we might be victims of evil and cunning genius? It appears she has none. Surely, at least, that is Austin's opinion.[28] Yet, in the absence of some such reason, her peculiar doubts cannot be raised.

That the sceptic has, in the garden, no *reason* to think our experience is phoney in any of these ways seems true. I have, after all, defined the situation as *ordinary*, and if the

---

[27]The question of whether she is not only silenced, but refuted, is considered in the next section of this chapter.

[28]Of course, if the sceptic is correct in describing the sceptical alternatives as equally compatible with all the evidence from experience, then it may be that we *at least have as much reason* for believing we are dreaming, vatting, or deceived by an evil genius, as we do for believing the cause of our experience is an external world much like the one we ordinarily believe is the actual cause of our experience.

situation is truly ordinary, it will not (it seems) be the case that there is reason to think we are dreaming, vatting, or deceived. An *ordinary* situation, by definition, seems to be one where such concerns are absent. What follows, in Austin's opinion, is that *ordinarily* — which is *most of the time* — the sceptic will lack (conventional) justification for raising her peculiar — and most extraordinary — doubts.

Can she raise them elsewhere?[29] Can she, for example, raise her doubts where circumstances suggest she has reason to think these alternative possibilities might obtain? Suppose, for example, that you and I, having had a bit too much to drink, find ourselves slumbering on the garden benches. And suppose that upon waking in the garden, I exclaim, "That goldfinch is bothering me," implying, of course, that there is a real goldfinch nearby, and (seemingly) that I believe that I know this. The sceptic, returning to the garden and overhearing my exclamation, might query thus: "Are you sure you are fully awake? You had rather a lot to drink last night. Perhaps you have just been dreaming of a goldfinch?" And, she might continue, "I assure you, I heard no goldfinch as I approached, but only the faint whimper of what seemed to be a slumbering voice."

Applying Austin's conventional criteria of appropriate doubt in these circumstances reveals that the sceptic's doubt (about whether I might merely be dreaming) is — in the face of my recent intoxication and slumber — appropriate (and, seemingly, therefore relevant). For these latter conditions — my recent intoxication and slumber — give her *reason to think* I might have been dreaming. Indeed, in this case, her reason to think I

---

[29]One must be clear to distinguish here the question of whether she *can* raise such doubts from the question of whether she *should* raise such doubts. Austin does not do this, and this is a serious omission on his part. More on this forthwith.

might have just been dreaming is the reason she raises the possibility.[30] Likewise, if upon awaking in the garden, we learn that our host — heretofore unknown to us — is an avid taxidermist who collects and stuffs goldfinches to display in his greenery, the sceptic's query regarding our certainty that the thing we identified last evening as a (real) goldfinch was not a stuffed goldfinch may be appropriate. After all, our somewhat peculiar host *is* in the habit of scattering stuffed goldfinches about his garden. Here again, it is not merely possible that we may have been deceived in our perceptions, there is some reason to think it so. This being so, the sceptic falters not in doubting our claims to know.[31]

There exist, then, contexts in which (at least some of) the sceptic's doubts about our claims to know can be *appropriately* raised.[32] What Austin suggests, however, is that the contexts — such as those described — where it is *appropriate* for the sceptic to raise these specifically Cartesian doubts are never present in the ordinary case. In particular, although in the ordinary case it is always *possible* that I may be dreaming, vatting, or the

---

[30]So as not to mislead or misrepresent the sceptic, I should mention that though according to Austin's criteria, the sceptic's doubts are appropriate in this context, this is precisely the kind of context in which a global sceptic would not raise her doubts. At least, it is not the kind of context from which she would derive her universal doubt. The sceptic's power lies in finding reason to doubt even in the ideal cases, though the a-typical cases are important in so far as they play a role in revealing the objective conditions of knowing, a role Austin (as well as the contextualists and the relevant alternativists) fails to note.

[31]Though in the case of the stuffed finch I might, of course, be able to discover the finch is stuffed, unlike the case where I might be dreaming, there is an important relation between these kinds of cases; the conditions of the possibility of certain doubts *ever* being relevant to my claims to know.

[32]Of course, confining the sceptic's doubts to the atypical misses the point of the sceptic's doubt. For her power resides in finding typical cases, actually better than typical cases, where perceptual conditions and the like are ideal, and then showing that even under such ideal conditions, one cannot know. Nonetheless, even when confined to atypical cases, the sceptic's doubts have universal significance, for the fact that her doubts are ever relevant reveals something important about the *general* conditions of the applicability of our concept of knowledge.

victim of evil and cunning genius, *ordinarily* (which is to say most of the time) there is not any *reason to think* that I am. What follows is that it is *inappropriate* (most of the time) for the sceptic to challenge my claims to know. Hence the fact that there are always a number of possible competing descriptions of the etiology of our experiences does not — contrary to what the sceptic has argued — undermine our knowledge claims. At least, it cannot prevent us from *saying* that we know, *most* of the time.[33]

Section iii: Disentangling: a conflation of manners and truth

Austin argues that the sceptic's Cartesian doubts are, more often than not, *inappropriate*. He argues that the two conditions of appropriate doubt are not often met. Austin's conditions, again, are that the reason for doubt must express a *possible* state of affairs, and that there must be a *reason to think* the state of affairs expressed by that doubt obtains. The problem for the sceptic, then, is that although the first condition may be met, this second condition is not. In *ordinary* circumstances, there is no *reason to think* of you or me that we might be dreaming, vatting, or deceived by evil and cunning genius. What follows, Austin argues, is that it is *inappropriate* for the sceptic to raise her doubts against our ordinary claims to know. Her doubts are inappropriate in all but the most unusual circumstances (e.g., when the subject of the purported knowledge is inebriated, or half

[33]Austin's *satisfaction* with this silencing of the sceptic rests on (among other things) a misunderstanding, or overlooking, of the *transcendental* significance of the applicability of the sceptic's doubts to such cases. Though I must reserve any sufficient argument about this for another place, the mere fact that the sceptic's doubts are at times applicable suggests that the requirement of certainty may be built in to our concept of knowledge.

asleep). Such doubts cannot be raised appropriately unless conditions (either of the perceiver or the environment) render not only possible, but to some extent likely, that one or more of the competing etiological descriptions is true.

One must concede that the conditions described by Austin as necessary for the *appropriate* assertion of doubt (hereafter, *the conditions of appropriate assertion*) do represent how we *ordinarily* expect others to behave in the context of *ordinary* conversation. Therefore,one must also concede that the sceptic behaves inappropriately when she raises these Cartesian challenges in the midst of *ordinary* conversation. As I confessed earlier, her doubts are contextually absurd. What, after all, could be more absurd than to challenge my claim to know that my dog exists on the grounds that it is logically possible that I may be a mere envatted brain? Surely the unlucky individual who puts forth such an assertion is likely to wind up in a straight-jacket.

> And how could I deny that I these hands and this body are mine, were it not perhaps that I compare myself to certain persons, *devoid of sense*, whose *cerebella are* so *troubled and clouded* by the violent vapours of black bile, that they constantly assure us that they are kings when they are really quite poor, or that they are clothed in purple when they are really without covering, or who imagine that they have an earthenware head or are nothing but pumpkins or are made of glass. But *they are mad*, and I should not be any less insane were I to follow examples so extravagant.[34]

Nonetheless, one must be cautious. One must be cautious not to infer from the sceptic's absurdity more than one is entitled. One must be cautious , that is, not to infer too much from the fact that the sceptic's behavior is inappropriate. For, though the sceptic's behavior — her articulation of Cartesian doubt — may be absurd, may violate

---

[34]Descartes, *Meditation I,* 145. (My emphases)

the conventions which govern appropriate assertion, it *does not follow*, at least not straightaway and without further argument, that her conclusions are false, irrelevant, or epistemically misguided.

It is here — with respect to the *epistemic* relevance of the sceptic's doubts, and with respect to the *truth or falsity* of her conclusions — that one must consider the second question: what is the *significance* of the sceptic's absurdity? What is the significance of her (or anyone else's) violating the conventions governing *appropriate* assertion? And, in particular, what — if any — is the *epistemic* significance of so violating these standards?

Austin has argued only that the sceptic's doubts are inappropriate, since — when raised — they are not supported by any reason to think they obtain. Asking for the *epistemic* significance of the sceptic's violation is another way of asking how Austin's observations about ordinary linguistic practices are supposed to *refute* the sceptic. The sceptic, after all, is concerned with knowledge, not — it seems — appropriate (linguistic) behavior or manners. Moreover, Austin, it seems, wishes not merely to comment on the sceptic's manners, but to refute her.

Without further argument, argument establishing a link between what is linguistically inappropriate and what is false, however, the sceptic remains unrefuted; whether ill-mannered or not.[35] It is precisely here that Austin's anti-sceptical suggestions go too far, beyond the evidence of ordinary linguistic behavior and convention. For, though such ordinary convention may, in all but the most extraordinary contexts, effectively *silence* the

---

[35]Equally problematic is any inference from what is linguistically appropriate to what is true. I will discuss this at length in chapter V.

sceptic, they do not prove her false. They silence, but they do not refute.

Consider what reasons Austin has given for rejecting scepticism. Austin has stated what are the standards for *appropriately* raising doubts to another's claim to know. He has observed that in ordinary conversation two standards are conventionally observed. He has argued that the sceptic violates *these* standards of *appropriate* assertion when she challenges our ordinary knowledge claims. From this violation of conventional standards, however, Austin infers also the irrelevance of the sceptic's doubts. But, *irrelevance with respect to what?* Irrelevance *in what sense*? Are the sceptic's hypotheses irrelevant to the truth of our knowledge claims? Or to something else? In what follows, I argue, it is only the latter, not the former, which Austin can justifiedly infer.

Austin's conclusion — that the doubts raised by the sceptic are irrelevant — is somewhat ambiguous, and derives its ambiguity from the equally ambiguous principle on which it rests: *If a challenge to a knowledge claim violates, when raised, the conventional standards for appropriate assertion, it is irrelevant to that knowledge claim.*[36] Let us call this principle the *Principle of Epistemic Conventionality*. Straightaway, this principle begs for an answer to the question: irrelevant in what sense? Irrelevant to what? There are two reasonable interpretations of this principle. There are also, therefore, two reasonable interpretations of Austin's conclusion that the sceptic's doubts are irrelevant to our ordinary knowledge claims. What I find, however, is that each interpretation proves

---

[36]In this, we shall see, Austin anticipates the central arguments of both the later Contextualists as well as the Relevant Alternativists. Their use of this Principle is discussed in chapters III and IV of the present work.

incapable of refuting philosophical scepticism.

The first interpretation of the Principle of Epistemic Conventionality follows immediately from the evidence Austin gives for concluding that the sceptic's doubts are absurd. According to this interpretation, when it is asserted that the doubts of the sceptic are irrelevant, on the grounds that they violate *conventional* standards for *appropriate* assertion, their irrelevance is confined to the context of good manners. The sceptic's violation, that is, makes her comments embarrassing, annoying, and irrelevant to the goals of ordinary conversation (goals which are generally limited to following convention and effectively communicating), goals aided by observing the (conversational) manners each duly educated person learns as she acquires the English (and, presumably, any other) language.[37]

If this first interpretation of the principle is the correct interpretation, Austin's 'refutation' of the sceptic is seriously deficient. For such an interpretation of the principle does no more than to bar the sceptic from *speaking* at (and perhaps attending) parties. *It silences her, but it does not refute her*. Refutations concern truth and falsehood. Refuting someone requires showing that what she says is false, not merely odd, linguistically aberrant, or ill-mannered. In fact, because establishing that what the sceptic says is false requires showing that she has violated, not merely the standards of

---

[37]I do not intend to argue at any length that these are the goals of ordinary conversation. For now, however, perhaps it shall suffice to merely indicate what I mean by ordinary conversation. Suppose, for example, I take my car in for an oil change. The attendant tells me that my oil is a quart low. I ask her how she knows this. She shows me the dip stick. My interest in this ordinary conversation about my oil is not Cartesian certainty, but keeping my car on the road.

appropriate assertion, but also the standards of *true assertion*, appealing to her violation of the former cannot suffice to refute her, unless there is a necessary connection between speaking inappropriately and speaking falsely.

But does there exist such a connection between the standards governing appropriate and those governing true assertion? Only if there is such a connection can the Principle of Epistemic Conventionality, and, therefore, the argument from ordinary language, refute, and not merely silence, the sceptic. The question of the connection between these two kinds of standards leads to the second interpretation of the principle of Epistemic Conventionality. According to the second interpretation, the sceptic's linguistic transgressions suffice to refute her, not merely prove her ill-mannered. According to that interpretation, the kind of irrelevance the sceptic has is epistemic. But this second interpretation rests on the presumed existence of a necessary connection between those standards governing what it is appropriate to say and those governing what is true. I turn now to that presumed connection.

There are, in fact, compelling reasons to believe that the standards governing appropriate assertion are not only *not* strongly connected, but are in fact *logically independent*. If the standards are logically independent, then, rather than giving one compelling reason to reject scepticism, one will have compelling reason to conclude that showing the sceptic's expressions of doubt are inappropriate or ill-mannered — which is all the *first* interpretation of the principle can justify — does not suffice to refute her.

Moreover, and consequently, if there is no such strong connection, then only the first

interpretation of the Principle — according to which the sceptic's irrelevance is confined to the context of manners — is justified. But this interpretation, though it may justify shunning or ignoring the sceptic for pragmatic considerations, will not do so on semantic ones. That is, it will not do so on the basis of precisely the grounds necessary to refute, rather than merely silence.

Let me then establish the logical distance between the standards of appropriate and true assertion. I begin by considering the following: Suppose that having outdone ourselves at the garden party, we find ourselves resting up at the Opera. Suppose, moreover, that while taking in *La Boheme*, I notice not only that the gentleman next to me is wearing a *toupe*, but also that his dental hygiene is so poor that I can discern, without effort, the particular herbs which have seasoned his dinner. Should I mention to this fellow that his toupe has slipped and his teeth need brushing? It is perfectly obvious that my so doing would be completely inappropriate. But, what follows from my mere inappropriateness? Does it follow that the statements I make are false? Are they irrelevant to the truth about his toupe or his teeth? Surely not. His *toupe has*, after all, slipped, and his teeth *are* laden with basil.

One might argue that the sceptic's queries are much like my (proposed) undiplomatic and ill-mannered behavior toward the stranger at the opera. Indeed, Austin and his ordinary language companions would, no doubt, appreciate the analogy. Importantly, however, just as it does not follow from my ill-manneredness alone that *what* I assert is false, so also it does not follow from the sceptic's ill-manneredness alone that *what* she

asserts is false, no matter how many conventions (governing *appropriate* assertion) she may twist or violate.

Consider another example.[38] Suppose my good friend Rodger invites me to a party, and I accept. Suppose, moreover, that just prior to my arriving at the party I have parted company with my friend John. Suppose that I have just encountered John buying supplies for Rodger's party, to which he assured me he is going forthwith. Now, suppose further that, after I have arrived at the party, the host (Rodger) asks of me "Is John going to be at the party?" And suppose, finally, that I respond, "Yes. I know he will. I just saw him at the store and he assured me he would be here forthwith." In the case just described, I think, the conditions for *appropriately* asserting that I *know* are met. I have, after all, just seen John. John did in fact assure me he was on his way to the party. (Perhaps earlier in the day I even helped John pick out a jacket for the party.) But, from the appropriateness of my assertion alone, does it follow that I know John will be at the party? Does it follow that my knowledge claim is true? Surely not. For, in spite of all my appropriateness, consider what *could* happen. John, upon leaving the corner store, might be struck by a car. Indeed, for all I know, he might be struck by a meteorite. In fact, he may be so struck that he is unable to go to the party and finds instead he must go to the hospital where he lapses into a coma and dies.

Suppose that John is struck by a meteorite, does lapse into a coma, and dies. Suppose, moreover, that upon my leaving the party, the host — having been informed of John's

---

[38]I borrow the structure of this example (and much detail) from Barry Stroud's *Significance of Philosophical Scepticism*, 68-71.

misfortune — approaches me, not to thank me for having attended, nor to console me on the loss of my friend, but only to berate me for the false arrogance I expressed in saying of John that I *knew* he would be there. After all, he points out, I *knew* no such thing. A necessary condition of knowing John would be at the party is knowing he would not be struck down by a meteorite. And this, he exclaims, I clearly did not know.

In such circumstances, surely one must concede that the host's beration of me is wholly inappropriate and inexplicably rude. Yet, does it follow from his ill-manneredness *alone* that what he says is false? Did I in fact *know* that John would be at the party? Did I *know* John would not be struck down by a meteorite? Surely not. After all, John was not at the party, and he was struck down by a meteorite.[39]

What I wish to suggest here, of course, is that the party host is much like the sceptic. The sceptic too behaves rather inappropriately at parties. At the garden party, the sceptic claimed that you failed to know there was a goldfinch in the greenery because you did not know that you were not dreaming, vatting, or deceived by an evil and cunning genius. In so doing, the sceptic violated our ordinary linguistic conventions. Her suggestion, therefore, is rather like Rodger's suggestion that I failed to know that John would be at the party because I did not know that he would not be struck down by a meteorite. And, just as it did not follow from the inappropriateness of Rodger's comments that what he said was mistaken, so also it does not follow from the sceptic's inappropriateness that what she says is mistaken.

---

[39]I assume here that it is also a necessary condition of knowing that what is believed be *true*.

What surely does follow, however, from these examples, is that there is no necessary connection between fulfilling (or failing to fulfill) the conditions of appropriate assertion, on the one hand, and fulfilling (or failing to fulfill) the conditions of true assertion, on the other. These examples clearly establish the logical independence of these two sets of conditions: the host's remarks are *both* absurd (inappropriate) and true, and my remarks are *both* appropriate and false. What follows is that one cannot infer from the inappropriateness of the sceptic's behavior alone that what she says is false. In particular, one cannot infer from the inappropriateness of her objections that they are irrelevant *to the truth of* our knowledge claims. Merely establishing that she violates linguistic conventions which govern appropriate assertion cannot, therefore, establish that she violates conventions governing true assertion, and hence though she may be silenced, she most assuredly is not refuted. The *epistemic* relevance of the sceptic's poor manners is nil.

One might object to my examples on the grounds that in them the conditions of appropriate assertion are not in fact fulfilled. Failing to fulfill these, one might argue, they cannot establish the logical independence of this set of conditions from that of the conditions governing true assertion. Thus, for example, one might argue that the conditions of *appropriately* asserting that John would be at the party *include* knowing that John would not be struck down by a meteorite. And, the objection might continue, since clearly I did not know that John would not be struck down by a meteorite, I ought not to have asserted that I knew he would be at the party. At the very least, one might argue, it was not appropriate for me to do so. Thus, one might argue, the illustration was not one

where the conditions of appropriate assertion were fulfilled, and hence not one where the latter where fulfilled and the conditions of true assertion were not.

There are a number of things one can say about this objection. First, and perhaps most damaging to its proponent, unless the objection is made during a record-breaking meteor shower, it sets the standards of appropriate assertion incredibly high. Making knowledge that my friend will not be struck down by a meteorite a condition of *appropriately* asserting I know he will be at the party implies the inappropriateness, and irrelevance, of nearly every knowledge claim one might ordinarily make.[40] Moreover, setting the standards of appropriate assertion *this* high amounts to giving the sceptic her cake and eating it *with* her. (Surely, the standards of true assertion will not be lower or easier to fulfill than the standards of *appropriate* assertion.) For precisely this last reason, not to mention the general counter-intuitiveness of the objection, this is an objection no ordinary language philosopher would make. In fact, it is more likely an objection a sceptic would make, and hence serves no purpose in the *defense* of an ordinary language *refutation* of scepticism.

One must then recognize a logical distinction between the conditions of appropriate assertion and the conditions of truth. One must recognize that this distinction is such that the one set of conditions can be fulfilled while the other is not. Thus I may, while at the opera, make a true though inappropriate assertion, just as you may, at the party, make an

---

[40]Many, of course, have argued that the sceptic is at fault for doing precisely this. I shall discuss the error of this view in chapter V. Suffice it for now to point out that the sceptic's conclusion are not directed at what we may appropriately assert, but rather, at what we may justifiedly *believe*.

inappropriate but true assertion. The two sets of conditions are logically distinct. Because they are logically distinct, one cannot infer from the fulfillment (or lack thereof) of the one set, the fulfillment (or lack thereof) of the other.

What, then, should one make of Austin's claim that because the sceptic's doubts are inappropriate, they are also irrelevant to our claims to know? I have argued that this conclusion is based on the Principle of Epistemic Conventionality, the Principle that if raising Cartesian doubts violates the conventions governing appropriate assertion, then those doubts are irrelevant to the knowledge claim against which they are directed. I argued that the conclusion that the sceptic's doubts are irrelevant must be based on this principle since it is only by thus linking inappropriateness to irrelevance that the argument from ordinary language can move from facts about ordinary discourse alone to conclusions about relevance. In particular, it is only if these two sets of conditions are so linked that the ordinary language philosopher can infer the irrelevance of the sceptic's hypotheses from their mere inappropriateness

I am now in a better position to answer the question: what is the *epistemic* significance of the sceptic's violating the conventions of appropriate assertion? Is her violating these conditions sufficient reason to conclude that her doubts are *epistemically* irrelevant; irrelevant to the *truth* of our claims to know? Is the sceptic's inappropriateness grounds for concluding her hypotheses are *philosophically* insignificant or that her sceptical conclusions are *false*? Clearly not. From their inappropriateness alone, nothing follows about their truth.

Of the two possible interpretations of the Principle of Epistemic Conventionality, then, the first is true, but ineffective; the second false. According to the first,one may conclude from the inappropriateness of the sceptic's assertions only that they are ill-mannered; *not irrelevant*. But on this interpretation, the principle is philosophically inept. It offers no guidance on the issue of truth or falsehood, and, in particular, it justifies no inferences about the truth or falsehood of the sceptic's conclusions. It justifies no inference regarding the *epistemic* merits of her doubts. It is, however, to the *epistemic* merits of her doubts that any proposed refutation must be directed. Whether she is ill-mannered, whether her doubts seem absurd, or out of place, the question of her correctness, the question of whether one can in fact know anything, is a question of truth. It is a question about the truth of our knowledge claims, and hence a question about the objective conditions which must be met if one is to know. It is, therefore, to no avail that the argument from ordinary language lists the multitude of linguistic rules the sceptic violates. Only the second interpretation of the principle is sufficient to refute the sceptic. Yet this interpretation is neither justified nor true. Manners and truths stand independent. Hence, that interpretation that alone could justify the conclusion that the sceptic's doubts are *epistemically* irrelevant— is false; based as it is on the false assumption that ill-manneredness alone is *semantically* significant.

If, therefore, the ordinary language philosopher believes she has refuted the sceptic (as opposed to having merely silenced her), she has confused good manners for good truths. For it is only by illegitimately conflating the standards governing appropriate assertion

with those governing true assertion that she concludes from the mere inappropriateness of the sceptic's hypotheses their *epistemic* irrelevance. The sceptic makes no claim to share her linguistic habits with the ordinary; nor should she. Her concern is not pragmatic, but semantic. She is interested not in what is conventional, but in what is true.

In both her inference from linguistic conventions to epistemic truths, and in her underlying conflation of the standards governing appropriate and true assertion, the ordinary language philosopher anticipates the themes — as well as the fundamental errors — of the anti-sceptical philosophers of my next two chapters; the contextualists and the relevant alternativist. Like the ordinary language philosopher, both the contextualists and the relevant alternativists base their refutations of the sceptic on the contextual inappropriateness of the sceptic's doubts. To the extent to which they do base their refutations on mere facts of ordinary linguistic behaviour, they can be viewed as the contemporary representatives of the ordinary language movement. Unfortunately, to precisely this same extent, they too fail; both in their attempts to refute the sceptic, and in their attempts to save — from the specter of scepticism — what they take to be common sense. It is to these contemporary defenders, and their likenesses to the ordinary language philosopher I now turn.

CHAPTER III
Contextualism

In this chapter, I look at a relatively new position in epistemology called 'contextualism.'[1] The contextualist, like the ordinary language philosopher, seeks to defend the veracity of our common sense knowledge claims, and, thereby, to save our "ordinary" knowledge. The contextualist seeks also to defeat the ever-present threat of epistemological scepticism. For the contextualist thinks that scepticism will continue to threaten our knowledge claims unless we can locate the source of its mysterious power. In her attempt to explain this mysterious power, the contextualist goes further than the ordinary language philosopher. The contextualist argues that the power of scepticism is located in a kind of (contextual) truth about the sceptic's position. The contextualist finds it possible both to defeat the sceptic and to concede (a contextual) victory to the sceptic because, unlike the ordinary language philosopher's, the contextualist's refutation of scepticism emerges from a semantical theory of knowledge according to which there are a *number* of correct analyses of what it means 'to know'. In fact, it seems that on the contextualist's account of knowledge, there are as many correct definitions of knowledge as there are knowledge-contexts. For knowledge, on the contextualist's analysis, is a

---

[1]The origin of the term 'contextualism' is, I believe, Peter Unger's *Philosophical Relativity,* though, as we shall find, the position properly called 'contextualism' predates Unger's book. See, for example, David Lewis, "Scorekeeping in a Language Game," *Journal of Philosophical Logic* 8, (1979): 339-359. And, in contextualist theories of justification, David Annis, "A Contextualist Theory of Epistemic Justification," *American Philosophical Quarterly*, 15, (1978): 105-112.

contextually-sensitive term. Hence, according to the contextualist, because of this semantic diversity, the sceptic's denial of our claims to know *is compatible* with a victorious common sense.[2] The contextualist argues, therefore, that the mysterious hold the sceptic exercises on philosophy can be explained in a way which protects ordinary knowledge.

In the following pages, I argue that the contextualist's successes in both refuting the sceptic and saving common sense rest on the same fundamental error as did the ordinary language philosopher's of chapter II: For like the ordinary language philosopher, the contextualist appears to succeed *only to the extent that* the she fails to separate the standards of *appropriate* assertion from the standards of *true* assertion, and to the extent that this failure goes unnoticed. Once these standards are distinguished, however, her arguments succumb to the same difficulties. For both her refutation of the sceptic and her conclusion that we do know those things we ordinarily say that we know, rest on her conflation of these logically independent sets of standards. Morever, since the contextualist grounds her standards of ordinary knowing *only in* the logically independent standards of appropriate assertion, the failure of her victory for ordinary knowledge is a *fait accompli*. As a consequence, the objective standards of knowledge remain, in spite of the contextualist's efforts, undetermined.

---

[2]I find, however, that the nature and extent of this victory, though conceded by all contextualists (see, for example, David Annis, (1978), Stewart Cohen, "Knowledge, Context, and Social Standards," *Synthese*, 73, (1987): 3-26; and Keith DeRose, "Solving The Sceptical Problem," *The Philosophical Review*, 104, (1995): 1-52; and David Lewis "Scorekeeping"), is viewed somewhat differently by each.

I argue, moreover, that *even if* one ignores the contextualist's failure to distinguish the logically independent standards for appropriate and true assertion, and therefore limit one's scrutiny to other aspects of her argument, the contextualist remains unvictorious. For her victory for common-sense knowledge is accomplished *only through* a series of conversational rules; rules which sanction the presupposing, ignoring, and assuming of *whatever is necessary* to give our ordinary 'knowledge' contextual veracity. But, I argue, if our ordinary 'knowledge' is veracious only to the extent that we presuppose, ignore, and assume, the victory here is deeply unsatisfactory.

Section i: What is Contextualism?

What is Contextualism? As a clearly recognizable theory in epistemology,[3] contextualism emerged recently. Initially, the position was neither clearly defined nor understood.[4] This is due not only to the late appearance of the term 'contextualism,' but also to the difficulty of distinguishing the contextualist's position from that of another relatively new position, that of the relevant alternativist.[5] The Contextualist position does, however, emerge as a response to a very familiar tradition in epistemology. Indeed,

---

[3]Contextualism has long been familiar and discussed at much greater length in ethical theory.

[4]See, for example, Cohen, "Knowledge," which will be discussed at some length here.

[5]See, for example, Fred Drestke, "Epistemic Operators," *Journal of Philosophy*, 67, (1970): 1007-23; and Alvin Goldman, "Discrimination and Perceptual Knowledge," *Journal of Philosophy*, 73, (1976): 771-91; and Gail Stine, "Skepticism, Relevant Alternatives, and Deductive Closure," *Philosophical Studies*, 29, (1976): 249-61. Since much of chapter IV shall be devoted to the details of the relevant alternativist position, and to what extent various relevant alternativists are contextualist, I will not discuss that here.

as David Annis notes[6], it emerges as an alternative response to the regress argument;[7] a response that claims to avoid the difficulties of both coherentism and foundationalism.

The clearest and earliest attempt to define contextualism is Unger's *Philosophical Relativity*.[8] Unger defines contextualism by contrasting it with a position he calls 'invariantism'. Invariantism, in brief, is the familiar view that the necessary conditions of knowing remain constant, rigid, inflexible with regard to idiosyncracies of both subject and object. By contrast, the contextualist position is characteristically dynamic, pliable, and flexible regarding the necessary conditions of knowing, especially with respect to the *conversational* context in which a knowledge claim is made. 'Knowledge', the contextualist argues, is a *contextually sensitive* term: it changes both in meaning *and* in *truth* conditions as the environment -- both internal and external to the subject -- alters. Thus, when S is said *to know* that p, whether she *knows* that p or not is a function of the *context* in which she claims to know.

Unger correctly identifies a contextually sensitive semantics as the most fundamental feature of any contextualist position. But there are a number of other features which are particularly worth noting as well. The contextualists, for example, are united by their common goal to refute the sceptic. Indeed, this, I think, is the motivation behind their contextually-sensitive semantics. Moreover, like the ordinary language philosophers, the

---

[6]Annis, "Contextualist,"105.

[7]Though I will not attempt to adjudicate who has the better solution to the regress argument, see: Robert Audi, "Contemporary Foundationalism," *The Theory of Knowledge*, ed. Lewis Pojman, (Belmont, CA: Wadsworth Publishing Company, 1993), 206-214.

[8]Unger, *Philosophical Relativity* (Minneapolis: University of Minnesota Press, 1984).

contextualists wish to rescue the knowledge assertions of common sense from the perceived threat of global scepticism. They believe, moreover, that with their contextualist semantics they can *explain* both the (contextually circumscribed) victory of the sceptic, and the enigmatic sway this position has held. Finally, they believe this explanation -- in terms of a contextual semantics -- proves that the sceptic's success is limited in such a way that it cannot implicate our ordinary claims to know.

In the following page I look at a variety of contextualist positions. The first contextualist position I discuss is that of David Lewis.[9] Lewis is suitably first in my discussion both because his, to my knowledge, is the *earliest* contextualist thesis, and because it is a near paradigm of contextualist reasoning.

Section ii: David Lewis

In "Scorekeeping in a Language Game", Lewis argues that *in any conversation* a number of things are *presupposed*. Thus, for example, if in the course of an ordinary conversation, I assert *that* "The table-top is flat", it is presupposed (lest I be made incorrect) that the standard of flatness according to which I assert that the table top is flat, is *not* the absolute standard according to which nothing could ever be flatter than the table top I now describe.[10] Thus, owing to the necessary presuppositions, my claim is true. Moreover, Lewis observes, these conversational presuppositions are both dynamic

---

[9]Lewis, "Scorekeeping."

[10]This example is used throughout contextualist literature, though I believe its origin is Peter Unger's *Ignorance*.

and rule-governed: that is, they *change* in accordance with *rules*. Thus, were you to assert that "The table-top is not actually *flat*", the presupposed standard of flatness *would change*, moving more in the direction of the absolute standard, according to which that which is flat is not *at all* bumpy (under microscopic examination or otherwise), and according to which your assertion is true.

The course of these changing presuppositions is charted by the changing contextual needs of the conversation. And, again, this change is rule-governed. Of the rules governing the changing conversational presuppositions, the most important is what Lewis calls *The Rule of Accommodation*. The Rule of Accommodation is as follows:

> If at time t something is said that requires presupposition p and if p is not presupposed just before t, then *ceteris paribus* and within certain limits, that presupposition comes into existence at t.[11]

What does this mean? Lewis argues it *means* that when, for example, I claim to know that p (*that* the table is flat, for example), then, *if* my claim to know would be false in the *absence* of certain presuppositions (for example, the presupposition *that* the standard of flatness is relatively low), those truth-bringing presuppositions "spring into existence"[12] at t, *thereby making my claim to know true.*

Because, however, the Rule of Accommodation allows the presuppositions within a conversation *to* change *as needed*, it follows also from the Rule that if at t1 *you* assert that I *fail* to know p, then whatever presuppositions necessary to make *your* statement true

---

[11]Lewis, "Scorkeeping," 2.

[12]Lewis, "Scorekeeping," 1. The full quotation is, "say something that requires a missing presupposition and straightaway that presupposition springs into existence . . . . "

(for example, the presupposition that flat is an absolute term) spring also (at t1) into existence. And so it goes. When, and to what extent, is a sentence true? That, Lewis claims, "depends on the context":

> What is true enough on one occasion is not true enough on another. The standards of precision in force are different from one conversation to another and may change in the course of a single conversation. [13]

Thus Lewis maintains, "it is not as easy as one might think to say something that is unacceptable,"[14] or false: As Lewis states, the Rule of Accommodation is simply an instance of the following more general rule:

> If at time t something is said that requires component Sn . . . if what is said is to be true . . . then at t the score component Sn [comes into existence].[15]

One might think that having the relevant presuppositions in place which make your claim (*that* I fail to know that p) true, my claim (to *know* that p, at t) would *thereby* be rendered false. Not so, according to the contextualist: for the presuppositions operative at t and t1 are different: the context has changed. *In fact*, the contextually sensitive presuppositions allow *both* my claim to know at t and your claim at t1 that I fail to know, to be true. For "*ceteris paribus* and within certain limits" *whatever* presuppositions are necessary "spring into existence". And presuppositions, on a contextualist's account, do not discriminate. Indeed at the heart of the contextually sensitive semantics is the claim that seemingly contradictory knowledge-claims may *both* be true: *both* my claim to know *and* your claim

---

[13]Lewis, "Scorkeeping," 114.

[14]Lewis, "Scorekeeping," 1.

[15]Lewis, "Scorekeeping," 9.

that I fail to know, for example.

The above considerations have, of course, important implications for Lewis' attempt to reach the contextualist's goals: the contextually sensitive presuppositions allow that both the claims to know in ordinary conversations and the sceptical claim (in philosophical conversations) that we fail to know are true. Moreover, the sceptic's *persuasiveness* (her power) can be explained by reference to this contextually-sensitive semantics. And, of course, the ordinary man's claims to know can be explained by reference to the same.

Furthermore, the sceptical victory can be *contextually confined*: for though what may be truly said in one context (for example my claim to know that p in an ordinary context) may be only falsely said in another context (my claim to know that P in a philosophical context), since the contexts and correlative presuppositions are different, my failure to know in the sceptic's context of high standards, does not undermine my success in knowing in the average man's more ordinary context, where the standards of knowing are lower. Instances of this rule, which governs the change of conversational presuppositions, according to Lewis, are prevalent in ordinary conversation. [16]

Section iii: Stewart Cohen

Stewart Cohen is another philosopher of whom I think contextualism is the right classification.[17] The construction of his contextualist semantics differs from Lewis',

---

[16]Lewis, "Scorekeeping," 347.

[17]Although Cohen seems to align himself with the Relevant Alternativists (see, for example p. 583), my sense is that is so because he is familiar with the RA position (as he notes, p. 580) but lacks familiarity with the term 'contextualism'. Indeed, to my knowledge, contextualism in epistemology had just begun to emerge as a distinct category shortly prior to Cohen's paper (see Unger's 1984, and

however, in important ways. Like Lewis, Cohen wants to refute the sceptic. And, as is characteristic of contextualists, he wants to save the claims to know of the ordinary person. But, unlike Lewis, Cohen's 1986 paper, rather than looking to any specific conversational rules to find the well-springs of contextualism, finds them in the necessary conditions of the *concept* of knowledge *itself*.[18] Cohen argues that it is a *necessary* condition of our concept of knowledge that the putative knower possess "good reasons."[19] Having good reasons, Cohen claims, entails the ability to defeat those defeaters to one's good reasons which are contextually important. But there are several kinds of "good reasons," and Cohen argues, not all of them are required for knowledge.

Cohen distinguishes three kinds of good reasons: there are

a. Subjectively good reasons
b. Intersubjectively good reasons, and
c. Ideally good reasons.[20]

Cohen defines each *kind* of good reason in terms of the absence of different kinds of defeaters: One possesses *subjectively* good reasons for her belief that p, just in case she

---

my earlier comment on this), while the relevant alternativist position was fairly well known much earlier. In fact Cohen discusses the RA positions of Dretske, Harman, and Goldman at p. 583 in "Knowledge". Supposing for the sake of argument, however, that Cohen was cognizant of 'contextualism' and hence consciously chose, for whatever reason, not to call his position by that name, there is plenty of textual evidence in his "Knowledge" to support my claim that he falls squarely within the contextualist school. Moreover, Cohen clearly identifies his position as a contextualist one in later papers: see, for example, his "Contextualism, Skepticism, and the Structure of Reasons," *Philosophical Perspectives* 13, (1999).

[18]This is somewhat of an over simplification. The well-springs of contextualism for Cohen are located in both the concept, and the standards of reasoning of the relevant social group. But it is the concept which specifies what is necessary for knowing, and Cohen goes outside the concept only in defining what those necessary conditions entail.

[19]Cohen, "Knowledge," 574.

[20]Cohen, "Knowledge," 575.

possesses no *subjectively evident* defeaters to those reasons; one possesses *intersubjectively* good reasons just in case there are no *intersubjectively evident* defeaters to her reasons; and, finally, one possesses *ideally* good reasons just in case there are *no defeaters* (evident or otherwise) to her "good reasons." [21]

Cohen argues that knowledge "entails" having good reasons.[22] But *which kind* of good reasons does knowledge entail? Cohen argues that though subjectively good reasons are entailed by knowledge, they are not a sufficient condition of knowing.[23] His reasons here are as expected: Given that subjectively good reasons are present *whenever* one possesses no defeaters to those reasons which *she* finds evident, if subjectively good reasons were all that were required for knowledge, then even where there were numerous and widely recognized defeaters to what the subject claimed to know, so long as she did not recognize them as such, she could be said to know. Surely, on any adequate account of knowledge, this will not do.

What about *ideally* good reasons? Requiring of the putative knower that she possess ideally good reasons surely would solve the problem of the standards for knowledge being too low: if the possession of ideally good reasons were requisite for knowledge, no one could rightly claim to know whether she possessed even one defeater to her good reasons which she could not rule out. This, of course, is a condition of knowing most hospitable to the sceptic. But, Cohen rejects this very strong requirement that *ideally good* reasons

---

[21]Cohen, "Knowledge," 576.

[22]Cohen, "Knowledge," 575.

[23]Cohen, "Knowledge," 577.

are necessary for knowing. He argues instead that knowledge entails *only* the possession of *intersubjectively good* reasons.[24] But what grounds has Cohen for dismissing the possibility that knowledge entails ideally good reasons? And what reason has he for thinking that knowledge entails merely *intersubjectively good* reasons?[25]

Cohen argues that it is the *intersubjectively* good, and not the *ideally* good, reasons which are required for knowledge because whereas the former are dictated by the "intersubjective standards of reasoning" the latter are not.[26] He supports this claim by appealing to our *ordinary* sense of when we say of someone that she either knows, or fails to know. He argues that we withhold knowledge only up to a certain level of intersubjective opacity; that is, when a subject possesses a defeater which is *intersubjectively* opaque – not evident as a defeater to the group – then we do not withhold knowledge from the subject, but rather, we readily assent to her claim to know. Contrariwise, when a subject possesses an intersubjectively *evident* defeater, and -- in spite of the presence of this intersubjectively evident defeater, which she fails to recognize as such, claims to know that p, we readily acknowledge her failure to know. [27]

Having appealed to the intersubjective standards of reasoning in defining what is required for knowledge, Cohen turns to the question: How are these standards set? It is here that Cohen's contextualism is most apparent. Cohen argues that the standards of

---

[24]Cohen, "Knowledge," 578.

[25]There are a number of obvious objections to this requirement. I will discuss these in the latter part of this chapter.

[26]Cohen, "knowledge," 578.

[27]Cohen, "Knowledge," 578.

reasoning are affected by two factors: a. The "*normal* reasoning powers of the attributor's social group,"[28] and b. The *context* of the attribution. [29] He concludes, therefore, like David Lewis, that whether a given knowledge attribution is *true* depends upon both the *context* in which it is uttered and the social group relevant in the context of the attribution.

Once again, therefore, I find that the semantics of knowing is a contextually sensitive matter. Again, as in Lewis' contextualism, *both* I, in my claim to know *that* the table top is flat, and you, in your claim *that* I fail to know the table top is flat, are correct. I am correct because the acuity of my group is lower than yours, making any defeaters to my "good reasons" intersubjectively opaque; and you are correct because your group, being a little brighter than most, finds the defeaters to my good reasons quite evident and indefeasible. As Cohen asserts:

> The truth-value of a knowledge attribution will depend on the status of the defeater the subject possesses, relative to the standards that apply *in the context* of the attribution.[30]

We come, therefore, to a thoroughly contextualist semantics of 'knowledge': a semantics according to which the necessary conditions vary, not only from group to group, but from context to context as well. Once again, therefore, the contextualist — in this case Cohen — meets his contextualist goals by contextualizing the necessary conditions of knowledge. And again, because the conditions of knowing are contextually sensitive (sensitive to the

---

[28]Cohen, "Knowledge," 579

[29]Cohen, "Knowledge," 579.

[30]Cohen, "Knowledge," 579.

relative acuity of the subject's social group), both the sceptic and the ordinary person are correct: The sceptic is correct that the ordinary person fails to know, and the ordinary person is correct her claims to know. For what is intersubjectively opaque or evident differs from group to group, and so, therefore, do the conditions of knowing.

Section iv: Keith DeRose

Let us turn now to a more recent contextualist,[31] Keith DeRose. Of the DeRose papers I will discuss or refer to, each is more recent than those of Lewis and Cohen. Moreover, perhaps owing to the advantage of time, DeRose recognizes his view as a 'contextualist'one, and titles his 1992 paper accordingly: "Contextualism and Knowledge Attributions"[32] It is his 1995 paper "Solving the Sceptical Problem," however, that is most relevant for us here.[33] In this paper, DeRose, like Lewis and Cohen, attempts to defeat the sceptic. Like Lewis and Cohen, he defends a *contextualist* account of knowing. And again, like the latter, in an effort to both defeat scepticism and explain its enigmatic power, he defends the knowledge-claims of ordinary folk.. Further still, DeRose, like the other contextualists, argues that his solution establishes the "compatibility" of common sense knowledge claims and global scepticism.

Unlike Lewis and Cohen, however, DeRose has the advantage of a contextualist

---

[31]I leave out some other very fine examples of contextualists, due only to a limitation of space and a desire to avoid redundancy. See, for example, David Annis, "A Contextualist theory of Justification," *American Philosophical Quarterly* 15, (1978).

[32]Keith DeRose, "Contextualism and Knowledge Attributions," *Philosophy and Phenomenological Research*, LII (1992): 916-929. I shall refer to the 1992 paper infrequently. This decision is justified both by considerations of length and of the greater complexity of his 1992 paper.

[33]I will discuss his "Contextualism" at some length in chapter IV.

tradition on which to fall back. In fact, he argues that though the main lines of the "traditional" contextualist solution are correct, they lack the power to explain the sceptic's persuasiveness.[34] Though DeRose endorses their basic contextualist strategy of arguing that the sceptic succeeds only at different contextual standards, he argues that the account of Lewis, for example, offers no real *insight* into *why* it is we are so tempted by the sceptical claim that we fail to know, as for example, when the sceptic asserts we fail to know we are not mere Brains in vats.[35] It is in this area, of providing a real explanation of the persuasiveness of the sceptical argument (the argument from ignorance, in this case), that DeRose thinks his account is superior.

DeRose, however, does not make his efforts alone. In his superior attempt both to explain and defeat the sceptic he draws heavily on Nozick's *subjunctive conditional analysis* of knowledge. [36] DeRose believes that an adequate explanation of the power of the sceptic's argument must locate the "mechanism" or "conversational rule" by which the sceptic "*manipulates* the standards" of knowledge: This, he thinks, Lewis failed to do (in his 1979 paper). Indeed, DeRose believes it is only because the sceptic so *manipulates* the standards of knowing that she succeeds at all.[37] But what is the mechanism by which

---

[34]DeRose, "Solving," 7 ff

[35]Specifically, DeRose argues that Lewis' rule of Accommodation can afford no explanation of why it is that the argument of the complex sceptic is more convincing than the mere assertion of what he calls "the simple sceptic" that we fail to know.

[36]Robert Nozick, *Philosophical Explanation,* (Cambridge, Harvard University Press, 1981): *167-228*

[37]DeRose ("Solving") claims that the earlier contextualists also maintain that the sceptic succeeds only by "raising the standards" of knowledge. This is somewhat innaccurate and needs to be clarified. Neither Lewis or Cohen argue that the sceptic succeeds by manipulating the standards (unfairly or otherwise). Rather, as in the case of Lewis, it is by the rule of accommodation -- a rule which operates equally in the spheres of common and philosophical discourse) that the sceptic's

the sceptic does this? It is in locating this supposed mechanism by which the sceptic manipulates the standards of knowledge that DeRose appeals to Nozick's subjunctive conditional analysis of 'knowledge'. DeRose argues that Nozick's claim that a belief must be "sensitive" to count as knowledge provides the insight needed for a more advanced contextualism.

A belief is *sensitive,* on Nozick's account, iff we would not hold that belief if it were false. Equivalently, a belief is *insensitive* iff we would continue to believe it *even if* it were false.[38] DeRose argues that we in fact have a *tendency* to believe that a belief *must* be sensitive if it is to constitute knowledge. And, it is this tendency, he claims, that is responsible for the extent to which we are psychologically overpowered by the sceptical argument. However, though we are persuaded by our tendency to believe sensitivity is required for knowledge, sensitivity is not a part of our concept of knowledge, and, therefore, is *not* ordinarily required for knowing.

DeRose parts company here with Nozick when it comes to building the sensitivity requirement into the concept of knowledge itself: Nozick argues that sensitivity is a *necessary* condition of the concept of knowing, and that *this* explains *why* the sceptic's claim that *we don't know* that -H (where H is a sceptical hypothesis to the effect that we may be mere brains in vats, for example) is so powerful. DeRose rejects this explanation

---

claims are made true. And, in the case of Cohen, it is in virtue of the fact that what is intersubjectively *opaque* to the *average person* is not opaque, but evident, to the sceptic, that the sceptic correctly asserts that we fail to know. Perhaps these are not so different from DeRose's explanation, but the detail is worth noting, especially in light of the somewhat hostile tone DeRose takes toward the sceptic; which is absent in both Lewis and Cohen.

[38]DeRose, "Solving," 18.

in favour of another. He argues, in contradistinction to Nozick, that sensitivity is *not* a part of our concept of knowledge, and therefore cannot be the explanation for the persuasiveness of the sceptic's argument from ignorance. The sensitivity requirement ,he argues, is only brought into play by the sceptic's manipulation of a *conversational rule*, a rule that when manipulated *raises* the standards of knowing to a level at which sensitivity is required. And it is *only* at this level, in the *unusual* context of sceptical manipulation, that the sceptic correctly claims that we know nothing.

This rule, manipulated by the sceptic, to raise the standards of knowing is as follows:

> When it is asserted that some subject knows (or does not know) that p, the standards for knowledge (how good an epistemic position one must be in to count as knowing) tend to be raised, if need be, to such a level as to require S's belief in that particular p to be sensitive in order to count as knowledge. [39]

DeRose calls this "The Rule of Sensitivity."[40] And, he claims, it is *only* at the artificially high standards brought into play by the rule of sensitivity and the *philosophical context* of sceptical discourse, that the sceptic succeeds in arguing that we fail to know. For, according to DeRose, because sensitivity is *not* a part of our ordinary concept of knowing, the sceptic's success is limited to this kind of contextual bubble wherein the rule of sensitivity is brought into play. Thus DeRose concludes:

> At these high standards, the sceptic truthfully asserts that we don't know that -H [that we are not brains in vats, for example],[41]

though

---

[39]DeRose, "Solving," 36.

[40]DeRose, "Solving," 36.

[41]DeRose, "Solving," 37.

> to the extent that he does succeed, he does so only by raising the standards for knowledge and so the success of his argument has no tendency to show that our ordinary claims to know are in any way defective. [42]

This, he thinks is the correct explanation of the sceptic's (mitigated) success. Thus, he concludes, like Lewis and Cohen, our ordinary knowledge claims are true *and* compatible with scepticism. The sceptic succeeds only in the unusual *context* -- created by the sceptic -- where the rule of sensitivity drives up the standards for knowledge; while common sense succeeds in the more ordinary contexts of everyday conversation, a level where the standards for knowing are admittedly low.

Section v: Contextualism and Ordinary Language

I have stated both the way in which DeRose's contextualism attempts to defeat the sceptic and save our ordinary claims to know by describing a semantics that is contextually sensitive, and the extent to which the contextually sensitive semantics of Lewis and Cohen do the same. It is time now, however, to ask for some *justification* of this proposed semantic theory. One must ask for the justification not only of the contextualist's semantics, but also, to the extent that it differs, her justification for the standards she argues are the "normal" necessary conditions of knowing, and according to which the sceptic cannot defeat us. It is here, with respect to the ground of her position, that I find that the Contextualists fall victim to the same kind of problem as that suffered by the ordinary language philosopher of chapter II.

I argued in chapter II that Austin, for example, in his attempt to argue that the sceptical

---

[42]DeRose, "Solving," 6.

hypotheses are *not relevant* to our ordinary knowledge claims, relied on the *Principle of Epistemic Conventionality*. According to that principle, whenever a sceptical objection is made (to the effect that we cannot meet the objective conditions of knowing), if it violates the conditions of appropriate assertion, that challenge can be ignored. The conditions of appropriate assertion, I showed, are defined by facts about usage: the things we *ordinarily* say about knowledge, the kinds of objections we *ordinarily* make to knowledge claims, and the kinds of responses we *ordinarily* both give and accept (or disregard) to queries about our claims to know. This set of linguistic *facts* about what we *ordinarily do* was thought to make irrelevant any objection, or claim, that was non-ordinary. Yet, by showing that the conditions of appropriate assertion (claim, doubt, etc.) are logically independent of the conditions of true assertion, I showed also that it is *only* to the extent that she conflates these independent standards that she succeeds in refuting the sceptic and saves ordinary knowledge.

The contextualists -- in the course of arguing for their most important conclusions -- make nearly the same illegitimate conflation as the ordinary language philosopher. And though the details to some extent vary, the principles involved are much the same. The extent that the contextualist relies on the same *kind* of principles in her effort to save the linguistic norm, therefore, the failure of her purported successes is a *fait accompli*.

That such a formidable error is made by the contextualists is quite evident in David Lewis' discussion of the Rule of Accommodation. Lewis argued that, in accordance with the Rule of Accommodation, the presuppositions necessary to make a claim true "spring into existence" as needed. Thus, he claimed, though the sceptic's denials of our

claims to know are true, our individual ordinary knowledge claims are *not* thereby rendered false: in fact they remain true in spite of their correctly asserted falsehood!

Lewis' explanation for this exceedingly strange phenomenon is the Rule of Accommodation. He argued that this rule (which is merely an instance of a general schema operative in ordinary conversations) makes those *presupposition* needed to make the attributor's (or the denier's) knowledge claim true a part of the conversational presuppositions, *ceteris paribus* of course, and "within certain limits." And, moreover, the rule-governance of these changes in the contextually sensitive presuppositions is what makes possible both an explanation of the sceptic's success and the vindication of our ordinary knowledge-speak from her sceptical clutches. For, as I noted, our knowledge claims made at t remain true, *even if* the sceptic correctly claims at t1 that our knowledge claims are false: This is the magic of a contextually-sensitive semantics.

In what sense, however, are the changing presuppositions of our conversations "rule-governed?" In what sense is the Rule of Accommodation a rule. There are at least two possible answers here: 1. The rule is a *rule* in the sense that it *prescribes* what can and can not be assumed in a conversation, or 2. The rule is a *rule* because it *describes* the *procedures followed* in ordinary conversation; that is, it describes the presuppositions we happen to make. It is, however, *only if* it is a rule in the *first* sense that Lewis' description of how the presuppositions change in a conversation can save knowledge. For if the "rule" is merely descriptive, it offers no guidance, and certainly no standards, according to which a claim to know is either true or false. It offers no more guidance than would rules which describe how we happen to reason. (The gambler's fallacy, affirming

the consequent, and *ad hominens* come to mind.) Rather, it tells us merely what we happen to do, how presuppositions *happen* to get made in a conversation, *not* how they ought to get made. [43]

But is Lewis' rule a rule in the first sense? Clearly it is not. This is evident if we consider the origin of the rule: the origin is not the analysis of our concept of knowledge. Indeed, Lewis offers none.[44] Lewis' rule does not even purport to prescribe *any* conditions which must be followed for the changing of presuppositions to be legitimate.[45] He tells us merely that this rule – the Rule of Accomodation – can be *seen* to be operative in many instances in conversations; that it *happens* to operate in a way which favours neither scepticism nor common sense.

The only "support" Lewis offers for his proposed rule is that it can be seen to be operating in conversation. The rule, it seems, has its origin in our *ordinary* everyday practices, as this is the only origin Lewis describes. This would go a long way to explaining its non-prescriptive nature. But if the origin of the rule of accommodation just is our practices, it seems that the rule itself is not very rule-like at all: if it is "discovered" simply by paying attention to what we ordinarily do, to what happens when assertions and counter-assertions are made, and is merely descriptive thereof, it provides no basis for the *objective* prohibition or prescription of any conversational practice. For, as I noted in

---

[43]Indeed, it is a pattern throughout contextualist literature to ignore this important distinction; the distinction between what is the case and what ought to be the case.

[44]Lewis does, to some extent, offer such an analysis in his later paper "Elusive Knowledge." I shall address that shortly.

[45]Again, in his "Elusive Knowledge," he does attempt a list of both prohibitive and prescriptive rules. I discuss these shortly.

chapter II, examples abound of knowledge-claims and practices that violate the standards of true assertion, but not the standards of appropriate assertion, and also of knowledge-claims that do not violate the standards of appropriate assertion, but which are clearly false. And since we are here interested in the rules governing *knowledge*, the rule of Accommodation, since it is grounded merely in what we happen to do, brings us no closer to an *objective* criterion of knowing than at the end of chapter II. The extent to which this rule is *merely* descriptive, it offers no guidance, but merely narration. [46]

Having argued that Lewis' Rule of Accommodation fails to be anything but descriptive in character, in fairness, I should note that Lewis' more recent paper, his 1996 "Elusive Knowledge," at least appears not to make this mistake. In this later paper Lewis again argues for a contextualist position, but his approach seems much more refined, both with respect to his earlier paper and with respect to other contextualists. Lewis, here, attempts not only to provide a number of types of rules, he also goes so far as to argue that knowledge is infallible. He argues, however, that in spite of the infallibility of knowledge, we can save common sense knowledge claims.

How does Lewis propose to save ordinary knowledge? He begins by offering a definition of knowledge and then providing a number of rules that fully explain its meaning. His definition of knowledge is as follows:

S knows that P iff P holds in *every* possibility left uneliminated by S's evidence.

Or, equivalently,

---

[46]Again, I think my analogy with how we happen to reason is instructive. For surely a mere description of how we happen commonly to reason does not reveal how we ought to reason.

S knows that P iff S's evidence eliminates *every* possibility in which not-P.[47]

This definition of knowledge looks surprisingly congenial to the sceptic. After all, it is precisely because the sceptic holds that there are uneliminated possibilities in which not-P, that she holds S fails to know that P. Lewis, of course, is aware of this difficulty and proceeds to resolve it by a special definition of what it *means* to require of *every* possibility that it be eliminated.

He argues that the term 'every' is "normally" restricted to a certain domain. Thus, for example, when at a party one asserts "*all* the glasses are empty", it is understood that the assertion is confined to a limited domain, namely, the glasses *at the party*. The glasses outside of the party (in the rest of the world) are understood to be irrelevant to the truth of the claim, as they are outside the limited domain.[48]

Lewis argues, by analogy, that this is the case with knowledge. Hence, he asserts, "when we say that *every* possibility must be eliminated, we can ignore some of the possibilities, for they are irrelevant to the truth of what is said."[49] For, "what does it mean to say *every* possibility in which not-p is eliminated? An idiom of quantification like 'every' is normally restricted to some limited domain."[50] Feeling the contextualist need to keep the sceptic at a safe distance, Lewis proceeds straightaway to *modify* his initial definition of knowledge *dramatically* as follows:

---

[47]Lewis, "Elusive Knowledge," 551. (My emphases.)

[48]Lewis, "Elusive Knowledge," 553.

[49]Lewis, "Elusive Knowledge," 553. (My emphasis.) One would like to exclaim here "So much for the infallibility of knowledge!"

[50]Lewis, "Elusive Knowledge," 553. My emphasis.

> S *knows* that P iff S's evidence eliminates every possibility in which not-P
> — Psst! — except for those possibilities that we are properly ignoring.[51]

The remainder of Lewis' paper is devoted largely to clarifying this *sotto voce* proviso. That clarification takes the form of an enumeration of rules, rules which specify what may and may not be "properly ignored".

Among the more outstanding of these rules are the rules of Actuality, Resemblance, Reliability, Conversation, and Attention. The Rule of Actuality simply stipulates that we may not ignore the possibility that actuality exists. In other words, "actuality is always a relevant alternative".[52] The Rule of Resemblance is a prohibitive rule: it states that, if "one possibility saliently resembles another . . . then, if one of them may not be properly ignored, neither may the other." [53]

Unlike the Rule of Accommodation of his 1979 paper, these two rules do *offer guidance* concerning what may and may not be presupposed in the course of a knowledge conversation. Thus, to some extent, they appear more rule-like. We now know, for example, that we *may not* simply presuppose that actuality is false, nor of anything saliently resembling it, that it is so. This helps to *place restrictions* on both what the ordinary person may be said to know, and on what the sceptic may claim is not known. To this extent, the rules of "Elusive Knowledge" seem a great improvement on those of "Scorekeeping."

---

[51]Lewis, "Elusive Knowledge," 554. (One can nearly hear the sceptic gasp! "Infallibility indeed!")

[52]Lewis, "Elusive Knowledge," 554. (What Lewis means by 'actuality' is merely what we ordinarily take to be the case about the world.)

[53]Lewis, "Elusive Knowledge," 556

Ironically, however, these rules — taken together — seem to place the sceptic at something of an advantage. For clearly the sceptic's hypothesis that I may be a mere BIV, or that I may be dreaming, are compelling precisely because they saliently resemble what Lewis calls "actuality." Thus, it seems, according to the rule of Resemblance, the ordinary person *cannot* "properly ignore" the sceptical hypotheses, and furthermore, given Lewis' definition of knowledge, she cannot claim to know. (Though, of course, the Rule of Actuality ensures also that the sceptic may not properly ignore the possibility that actuality is true.)

Lewis, of course, recognizes the problem his Rule of Resemblance seems to pose,[54] and in the interest of avoiding scepticism and saving ordinary knowledge, he makes an admittedly *ad hoc* exception to the rule. In conceding that the rule *could* apply to the possibility in which the subject is "radically deceived by a demon" and, therefore, favour the sceptic, he remarks:

> That will be so even if W [the uneliminated possibility] is in other respects very dissimilar to actuality. . . . Plainly we dare not apply the Rules of Actuality and Resemblance to conclude that any such W is a relevant alternative — that would be capitulation to scepticism. The Rule of Resemblance was never meant to apply to *this* resemblance! We seem to have an *ad hoc* exception to the Rule . . . . What would be better, though, would be to find a way to reformulate the Rule so as to get the needed exception without *ad hoc*ery. I do not know how to do this.[55]

---

[54]Actually, though he addresses one problem this rule may pose, it is not the one I mention. He notes that the rule of resemblance may make relevant the sceptical possibilities simply because they resemble all possibilities which are uneliminated by the evidence. (See: 557.) In fact he nowhere addresses the problem I mention -- the salient resemblance of the sceptical possibilities to actuality, though he solves both it and the one he does discuss with his *ad hoc* exception to the rule of resemblance.

[55]Lewis, "Elusive Knowledge," 556-7.

The effort of Lewis to save knowledge with these two rules seems, then, to succeed only insofar as we accept his *ad hoc*ery.

What about his other rules? Do they offer some improvement over the earlier Rule of Accommodation? These too seem, at first glance, to hold out some hope for providing truly prescriptive ("real") rules according to which our claims to know can be not only made and judged, but vindicated. The Rules of Reliability and Conversation — like those of Actuality and Resemblance — seek to save our ordinary 'knowledge' from the sceptical hypotheses. And, unlike the Rule of Accommodation, (though like the Rules of Resemblance and Attention) these rules *offer guidance* regarding what is — and what is not — *acceptable* to presuppose when making claims to know.

The Rule of Reliability, for example, allows us to "assume that perception, memory, and testimony are . . . fairly reliable." Indeed, "we are entitled," Lewis says, "to take them for granted." "We may properly *presuppose* that they work without a glitch." Thus, he says, we may assume, albeit defeasibly, that the possibilities "in which they fail may be properly ignored."[56] The Rule of Conversation, moreover, specifies that we are "entitled to ignore — defeasibly — what those around us ignore." [57] Thus, for example, if the members of a group I happen to be conversing with ignore all but the most obvious defeaters to our knowledge claims, I too am entitled — without argument or further justification — to ignore them.

These "rules", however, are not argued for by Lewis, but merely assumed. And though

---

[56]Lewis, "Elusive Knowledge," 558.

[57]Lewis, "Elusive Knowledge," 559.

assuming perception is reliable, and *ignoring* all those sceptical possibilities those around us ignore, should surely help us to defeat the sceptic, one is struck by the fact that one wins (defeats the sceptic) only by begging the questions at issue. After all, is it not at the very heart of the sceptic's argument that we *may* be *misled* by our perceptions, that we *may* be *wrongly* ignoring certain uneliminated possibilities which undermine our claims to know? Moreover, are mere presumptions rule-like at all? It seems the hope held out for grounding our ordinary claims to know in real rules is dashed, even in this later paper of Lewis. For though these rules are indeed more rule-like, their basis — assuming and question-begging — is fundamentally flawed. And so, though not subject to the same criticisms as the Rule of Accommodation, they are particularly inept at moving contextualism into a more positive light.

There remains, however, one final "rule" to examine: The Rule of Attention. This rule, like the others of Lewis discussed above, is meant to fill out the *sotto proviso* of Lewis' definition of knowledge and, thereby, to clarify the conditions of knowing. The Rule of Attention states that what is and is not ignored is a feature of a particular *conversational* context.[58] It states, moreover, that

> No matter how far fetched a possibility may be . . . if we are not now ignoring it, then for us now it is a relevant alternative."[59] "Do some epistemology. Let your fantasies rip. Find uneliminated possibilities of error everywhere. Now that you are attending to them . . . you are no longer ignoring them, properly or otherwise.[60]

---

[58]Lewis, "Elusive Knowledge," 559.

[59]Lewis, "Elusive Knowledge," 559.

[60]Lewis, "Elusive Knowledge," 559.

What is the nature of this rule? Is it, for example, dictated by the necessary conditions of knowledge? Is it merely descriptive of what happens when a lonely sceptic seeks our conversation? It is not clear. Surely the rule is not of much help in *saving* our knowledge claims. For the more uneliminated possibilities we must attend to, the less likely it is we will rightly claim to know. Thus, at least if a verbose sceptic is around, our knowledge claims fall asunder. Though perhaps, in Lewis' view the other rules (of Actuality, Resemblance, Reliability, and Conversation) are enough to save our knowledge from the sceptical possibilities. (After all, what more could one who seeks to save knowledge hope for than to be entitled to ignore whatever uneliminated possibilities of error those around him ignore?) In fact, rather than functioning to safeguard our ordinary knowledge claims, the Rule of Attention's most important function appears to be *to explain* how it is that the sceptic has such power, or in Lewis' words, how it is that knowledge can be so "elusive."

It is this rule, Lewis argues, that enables epistemology to "destroy itself:" making immediately relevant any considered hypothesis, no matter how bizarre, and hence making improper the ignoring of the sceptical hypotheses, at least when they are before our attention. It is because of this rule – the Rule of Attention – that the sceptic's raising of her sceptical hypotheses -- that we may be mere BIVs, or idle dreamers -- *undermines* our claims to know. For given that knowing requires the ruling out of *every* uneliminated possibility we are not proper in ignoring, we will *never* fulfill the necessary conditions of knowledge, at least not when the sceptic is present.

Importantly, this last rule is a *contextualist* one: it makes the sceptical alternatives alternatives relevant *only* in those *contexts* in which we attend to them. This

contextualized semantics, therefore, once again allows not only that the sceptic has some degree of success, but at the same time, safeguards our ordinary knowledge claims. And it is for this reason — *the reason of contextualism* — that Lewis can say:

> That is how epistemology destroys itself. But it does so only temporarily. The pastime of epistemology [of looking more closely at these strange alternatives than we ordinarily do[61]] does not plunge us forevermore into *its special context*. We can still do a lot of proper ignoring, a lot of knowing, and a lot of true ascribing of knowledge, to ourselves and others, the rest of the time.[62]

Thus, "Elusive Knowledge," provides us with a *kind* of answer to how, in spite of the sceptic's victory, we can correctly claim to know.

But how satisfactory is this answer? And, has it improved upon the answer he gave in 1979? I turned to Lewis' 1996 paper after finding dissatisfaction with his earlier "Scorekeeping". I was dissatisfied there because, though there *are* two kinds of rules — the merely descriptive and the prescriptive — it seemed that Lewis' Rule of Accommodation fit only the former category. This was problematic because rules which are descriptive merely entail, in themselves, *nothing* about what is *objectively* correct: nothing regarding what we ought or ought not do. As I stated there, if such rules are narrative merely, based only on the descriptions of the way we happen to behave, they are not justificatory. For it never follows from the mere facts of verbal behaviour alone that one is behaving as she ought to behave. As I have argued, one can easily behave as others do, and in that sense meet the conditions of intersubjective appropriateness, while at the

---

[61]Lewis, "Elusive Knowledge," 565.

[62]Lewis, "Elusive Knowledge," 559. (My emphasis.)

same time violate the conditions of objective truth.[63]

Has Lewis managed in this more recent paper to overcome those difficulties? It is unclear. For, on the one hand, he gives no explicit reference to the source of these rules or to their justification. And, on the other hand, even to the extent which we accept them without justification, they seem somewhat *ad hoc* and unsatisfactory. Thus, though he does not locate the origin of these rules in facts about ordinary language (given that he does not locate them at all), and *hence* may escape any claim that he errs in the same manner as the ordinary language philosopher of chapter II, at the same time, he seems not to improve upon the problems of his earlier paper. I am left without *any* vision as to the purported justification of these rules, and hence somewhat bewildered as to their real nature, and what epistemological work they can do.

In addition to bewilderment, I am left in a state of dissatisfaction: for it is only to the extent that he either begs the question (as with the Rule of Reliability), makes assumptions (as with the Rule of Conversation), or makes *ad hoc* exceptions (as with the Rule of Resemblance) that his rules manage to save our "knowledge." In fact the difficulties I mention are so apparent and overwhelming, I almost forget to ask, "What happened to the *infallibility* of knowledge?"; an infallibility Lewis argued was essential to knowledge. Indeed, Lewis himself seems to recognize the unsatisfactory nature of his "solution" when, near the end of his paper, he asserts:

> I have to grant, in general, that knowledge just by presupposing and ignoring *is* knowledge; but it is an especially elusive sort of knowledge, and

---

[63]See the examples of chapter II, pp. 15-20.

consequently an unclaimable sort of knowledge.[64]

Both the early and later paper of Lewis are unsuccessful, then, in providing a solution to the contextualist's sceptical problem: the problem of explaining the success of the sceptic while guarding our ordinary knowledge. His first paper (1979) failed because, like the ordinary language philosopher, he succeeded in saving ordinary knowledge claims *only to the extent that* he failed to separate two logically independent sets of conditions, those of appropriate and true assertion. His later paper fails because he saves our knowledge there only by presupposing, ignoring, and question-begging.

Section vi: DeRose and Invariantism

Keith DeRose, I noted earlier, believes that his solution to the problem of scepticism and ordinary knowledge is superior to that of other contextualists. Is he more successful? Let's return briefly to his Rule of Sensitivity. Does his Rule of Sensitivity avoid the problems of the ordinary language philosopher of chapter II?

Whether DeRose does avoid those difficulties depends in part on the sense in which his Rule of Sensitivity is a 'rule'. But it also depends on the extent to which, if any, his Rule is used to define the standards according to which we may be said, correctly, to know. Here one must be a little cautious, because DeRose (unlike Lewis) does *not* use his Rule

---

[64]Lewis, "Elusive Knowledge," 562. (His emphasis.)

to argue our knowledge-claims are safe, at least *not directly*. Rather, he uses the Rule of Sensitivity to explain the success of the sceptic. Only indirectly, by arguing the standards put in place by this rule -- the standards, that is, that explain the sceptic's success -- *do not* apply in more ordinary contexts, does his Rule have consequences for our ordinary standards of knowledge. So it is only indirectly that DeRose argues through this rule that our ordinary knowledge claims are safe. Before I address the important question of the origin of the standards according to which we (ordinarily) *do* know, and according to which the standards brought about by the rule of sensitivity do *not* apply to our ordinary knowledge claims, we must further examine of the Rule of Sensitivity.

DeRose used his Rule of Sensitivity to *explain* how the sceptic "drives up" the standards for knowing to such a degree that we fail to know even those things we ordinarily take ourselves to know. What is the justification for this rule? What is *its* origin? Is DeRose's Rule of Sensitivity a "rule" in the prescriptive sense -- one that *prescribes* what ought to happen -- or merely a rule in the descriptive sense, describing what *happens* when we make and evaluate claims to know?

Like Lewis' Rule of Accommodation, DeRose's Rule of Sensitivity seems *merely* descriptive: it offers a *description* of what takes place when the sceptic either asserts that we fail to know, or brings forth one of her famous sceptical hypotheses which, she claims, we cannot know to be false. According to DeRose, it is only when the sceptic makes some such assertion, that the standards of knowing are raised to a level at which our beliefs must be sensitive to count as knowing. Thus, DeRose says, the success of the sceptic is restricted to the *context* created by her manipulation of the Rule of Sensitivity.

Again that Rule states that:

> When it is asserted that some subject knows (or does not know) that p, the standards for knowledge (how good an epistemic position one must be in to count as knowing) tend to be raised, if need be, to such a level as to require S's belief in that particular p to be sensitive in order to count as knowledge. [65]

Noting that this Rule of Sensitivity operates in sceptical contexts is *not enough*, however, to justify the claim that DeRose, like Lewis and the ordinary language philosopher violates the distinction between the standards of appropriate assertion and the standards of true assertion in defining the ordinary standards of knowledge. For, as I have noted, it is only in a negative sense that his Rule of Sensitivity has implications here for the ordinary standards of knowledge. The Rule does not purport to offer a positive account of the ordinary knowledge requirements. The Rule tells us only that at the "low" levels of ordinary conversation, since the sceptical hypotheses are *not* at issue, sensitivity is *not* required for knowledge, since it is only when the sceptic enters the discussion that the Rule of Sensitivity is brought into play. Thus, whether we would believe that we have hands, for example, even if we did not — that is, even if our belief that we have hands is *insensitive* — does not affect the truth-value of our claim to know this in an *ordinary* context.[66]

The Rule of Sensitivity seems, therefore, to limit the sceptic's success to the *specific*

---

[65]DeRose, "Solving," 36.

[66]DeRose assumes throughout that such beliefs are sensitive. But the problem the sceptic brings in manipulating the Rule of Sensitivity is that our ordinary beliefs become insensitive. For because at these high standards we fail to know that we are not, for example, brains in vats, and because moreover, we must be in at least as good an epistemic position to know that O as we are to know that -H, we cannot, therefore, at these levels know what we ordinarily claim to know.

context of her sceptical queries. So one must look elsewhere for a justification for DeRose's *positive* claim that we *do* in fact know those things we ordinarily take ourselves to know, and for the standards according to which this judgement is made.

What, then, *are* the requirements of knowing at these low levels? And *how* does DeRose arrive at a knowledge of them? By what standard does he know that these are the ordinary rules? That is hard to say; for DeRose is not very forthcoming here. In fact, DeRose nowhere speaks of their origin or their justification. The most I am able to do, therefore, is to infer the origin and justification of these rules from what DeRose *does* say. What DeRose does say leads me to conclude that the *ordinary* requirements of knowing, for DeRose, correlate precisely with the standards which seem to govern what we *let pass* as knowledge in everyday contexts. In other words, the conditions of *ordinary* knowing are, once again, to be deduced from our *ordinary* practices. Whether this is in fact DeRose's position, I cannot say for certain since DeRose does not give much to go on: he merely asserts, time and again, throughout his paper, that at the *ordinary* low levels of *ordinary* conversation what we *ordinarily* take ourselves to know (that we have hands, for example) is known.

Thus, it seems once again that the contextualist's standards of knowing are derived from mere *facts* of *ordinary* discourse: *what* we say, *what* we do, *how* we respond, and so forth.[67] But again, as with the ordinary language philosopher of chapter II, this is, at best, an incomplete argument for what the objective standards -- what is *truly* required -- of

---

[67]This I have argued is the method of derivation of the ordinary standards for knowledge in Lewis' 1979. And it is a method explicitly advocated in the contextualist David Annis' paper "A Contextual Theory of Justification". See, for example, 107.

knowledge are. For how can the mere facts of what we *happen* to say, what we *happen* to accept others saying, and what they *happen* to accept us saying, and so on, *establish* the objective conditions of knowing? It seems, as I argued in chapter II, they cannot. For, it is only if the standards of what we do allow (what I have called the standards of appropriate assertion) are logically related to the standards of what we ought to allow (what I have called the standards of true assertion), in such a way that allows us to deduce the latter from the former, that the mere descriptive facts can provide us with an objective list of the criteria of knowledge.

But, as my many examples from chapter II (of appropriate yet false assertion, of true yet inappropriate assertion, and so on) have shown, this is precisely what is *not* possible because the standards of appropriate assertion are logically independent of the standards of true assertion: it is entirely possible to speak in accordance with the former *while violating* the latter. Hence the former cannot function as a definitive guide to the latter. And, therefore, the mere fulfillment of the former cannot – in itself – support the claim that we know what we ordinarily take ourselves to know.

Before concluding my discussion of DeRose, however, I should note that he does consider a criticism to his position that is very much like my own. He argues that the invariantist can offer an argument against DeRose's contextualism, which although it offers a radically different analysis of ordinary behaviour, essentially parallels his own. According to the invariantist position DeRose describes, however, it is the standards of true assertion that remain constant while it is the standards of *useful* or *warranted* assertion that vary. And, this invariantist argues, in our common sense knowledge claims

we mistake the standards of true assertion — which we do not meet — for those of warranted or useful assertion, which we do meet. DeRose attributes at least the beginnings of such a view to both Peter Unger and Barry Stroud[68] (to whom, as I mentioned in chapter II, I am very much indebted.) And DeRose goes on to argue that though anyone can offer such a parallel argument, if she has no other more weighty considerations on her side, there is no reason to prefer her argument to DeRose's own.

For, according to DeRose, the invariantist's parallel argument involves the speakers of English in widespread and systematic falsehood. And this, he thinks, counts "decisively" against their position. I will address the extent to which my argument (that it is the standards of appropriate assertion, not true assertion, which govern our knowledge claims) is similar to and dissimilar from the argument he describes in chapter V. Therefore, at this point I would like simply to consider DeRose's objection to the invariantist position, namely, that it involves us in systematic and widespread falsehood. I should like to point out that DeRose's objection is accurate only if the invariantist takes *precisely* the position he describes: in particular, it is only if the invariantist assumes that it *is* part of our common sense that we *believe we know* the things we say we know in ordinary conversations, that the invariantist criticism of DeRose's position does indeed implicate us in systematic and widespread falsehood.

As I shall make clear in chapter V, however, this is *not* my view of common sense; for I do not believe it is the correct view of common sense. Rather, as I will argue there, our

---

[68]The reference is to Peter Unger's *Ignorance: A Case For Scepticism* (Oxford: Oxford University Press, 1975), and Barry Stroud's *The Significance of Philosophical Scepticism.*

common sense beliefs about our knowledge claims are quite different from what philosophers ordinarily take them to be. In fact, I shall argue that we do *not* believe that *we know* the things we *say* we know in ordinary conversations. And, if one's sceptical invariantism starts here, rather than from the assumption that we do believe we know such things, then the invariantist's argument that it is the standards of appropriate assertion, and not the standards of true assertion, which vary, will *not* have us all believing falsely. For, if we do not *believe* we know the things we loosely say we know, we cannot *falsely* believe we know these things. Thus, DeRose's *only* and purportedly "decisive" objection is, with respect to a different invariantist position, quite without foundation and ineffective. Moreover, since this is the only reason he gives for preferring his contextualism to invariantism, we have no reason to embrace his position.

Finally, let's look at the contextualism of Stewart Cohen. He, like Lewis and DeRose, argued that the sceptic's success is explained by, and limited to, contextually-sensitive standards of knowing. Moreover, like the other contextualists, he argued that we do know those things we ordinarily take ourselves to know. And, though Cohen offered neither "rules" for the making of ordinary knowledge claims nor for the explanation of the continued hold the sceptic exercises on us, he did, like Lewis and DeRose, argue, that the standards of knowing *change* from context to context: that the standards of knowing are semantically *contextual*.

Recall that for Cohen, knowledge requires only "intersubjectively good reasons". That is, reasons for which there are no *intersubjectively evident* defeaters. And for Cohen, which defeaters are evident — and therefore, what determines the standards of knowledge

-- is determined by the reasoning powers of the group which claims to know. Thus, if the group to which I happen to belong happens to be particularly dull, what counts as an "evident defeater" will differ greatly from the evident defeaters of a more accomplished group of reasoners. For, while my group may lack the acuity to discern that some piece of evidence *is* a defeater to our good reasons, another group may not.

What follows is that the standards *of knowledge* are only as high as the (logical) intelligence of one's group. As Cohen argued, "In general, the standards in effect in a particular context are determined by *the normal reasoning powers* of the attributor's social group."[69] It is for this reason that Cohen, like DeRose and Lewis, concluded that both we and the sceptic can be correct in our respective claims about knowledge: The ordinary person is (usually) correct in the knowledge-claims she makes, and the sceptic is correct in claiming the ordinary person knows nothing at all. This can be so because, by Cohen's contextualist semantics, what it *means* to know for the ordinary person is *determined* by the degree to which her group can penetrate the opacity of defeaters; while for the sceptic and her group, being of a less dull sort, the evidentness of defeaters is inescapable, and hence the requirements of knowing are unmeetable.

Here again, in seeking the *foundation, the justification,* of those standards which Cohen claims make both our ordinary claims to know true, and the claims of the sceptic irrelevant, we find only description. His initial promise to derive the standards of knowledge from the concept deliquesces into a discussion of the relative intellectual acuity of the groups and the standards which they dictate. There is no *analysis* according to

---

[69]Cohen, "Knowledge," 579.

which 'knowledge' requires of us the fulfillment of certain necessary conditions. Rather, to the extent that he offers an analysis at all, there is only talk of what would or would not violate the "ordinary standards of reasoning", which vary according to the *ordinary* I.Q. of the attributor's social group.

But why should the standards of knowing thus vary? Why should the requirements of myself and the group to which I belong be lower than the requirements of the sceptic or a society of geniuses? The only explanation seems to be that this is in fact how we talk in ordinary everyday contexts. But if this is the only explanation forthcoming, we once again have a confusion of what passes as appropriate in ordinary (and extraordinary) discourse with what is objectively required to make the claims made therein veracious. And, if we take these standards to extremes there can be no doubt that the two sets of conditions are logically distinct, and, therefore, no doubt that the one does not follow neatly from the other.[70] Once again, therefore, the contextualist fails to save ordinary knowledge because, in her failure to separate the logically distinct conditions of appropriate and true assertion, he gives no objective foundation or criteria against which any such claim to know may be judged.

---

[70]I have already offered a number of such illustrations in Chapter II. But suppose, for example, that I have a group of friends who are all rather dull. Suppose therefore that nearly any defeater which is evident to you or me is opaque to them. Then, what passes for knowledge in their group is propositions such as that the United States is the only country in the world, the moon is only just above, and so on. Clearly, even though the assertion to know these things in their group is appropriate, the things asserted are neither true nor, therefore, known. Precisely the same kind of reasoning applies to the group of average persons, when compared, for example, with learned scientists or philosophers. This, of course is not to say that either the scientists or the philosophers know anything. It is just to point out that what defeaters are intersubjectively opaque to the average person might be perfectly evident to the more learned.

Section vii: Conclusion

I have claimed that the contextualists make essentially the same critical error as the ordinary language philosophers. The error of the ordinary language philosopher, as I argued in chapter II, was that she rested her whole defense of the ordinary standards of knowing on mere *facts* of linguistic behaviour. She rehearsed how we respond, when asked, *if* we know, *how*, and *why*. She offered *only* these *descriptions* of how we ordinarily behave in defense of her claim that the objective criteria are one and the same as those standards seemingly embodied in our linguistic behaviour. In so doing, I argued, she failed to separate the standards of appropriate assertion from those of true assertion. For it was only on the basis of her conflation of these that she arrived at her claim to save the knowledge of the common person. This failure is critical since it is possible, as I showed, to behave verbally in accordance with the former while violating the latter. Without further argument (to the effect that the rules of true assertion are derivable from the rules of appropriate assertion), therefore, the ordinary language philosopher failed to save ordinary knowledge from the sceptic.

I have shown, I believe, that the contextualist makes precisely the same kind of error. Lewis, by deriving the Rule of Accommodation — and the corresponding standards according to which we know — from ordinary practices; DeRose by locating the supposed fact that we know the things we ordinarily let pass for knowledge in just those conversational facts which evidence this; and Cohen, by relativizing what is necessary for knowledge to what the group perceives as opaque or evident. Each of the contextualists, then, derives the standards of ordinary knowing from the mere facts of how we happen to

behave, reason, or speak. As a consequence, each fails to establish that our ordinary knowledge claims meet the conditions of truth. Each fails, that is, to prove that we can know what the sceptic argues we cannot.

Suppose, however, that we ignore the contextualist's failure to distinguish the now familiar categories of appropriate and true assertion. And suppose, therefore, that the standards of knowledge are as low as they seem to be in ordinary conversation. Suppose, then also that we admit they succeed in proving that we do know what we say we know. How satisfactory is the contextualist's conclusion that we know? Not very, I am afraid. For we *know*, on the contextualists account, *only* to the extent that we presuppose (as in Lewis' 1979), we assume, ignore, or beg the question (as in Lewis' 1996), or we rely on our own lack of acuity (as in Cohen 1986). But how satisfactory is it to be able to say that one knows only because one presupposes or assumes what one cannot argue for, or ignores or begs the question with issues that are intractable, or hides behind the average intelligence of her group? Not very satisfactory I think. Indeed, as a victory for common sense, it is cheap indeed!

In *knowing* that I am not a mere brain in a vat, or a mere idle dreamer, I want to know *because* I can rule out, *because* I can argue against, *because* I can provide definitive evidence; not *because* I am sneaky and dull. In fact it gives me no satisfaction to know that a contextualist would claim I know I am not a brain in a vat simply because I may presuppose what those around me (no matter how ignorant) presuppose, or ignore possibilities (no matter how compelling) simply because I cannot rule them out, or overlook defeaters to my "good reasons" simply because my group finds them opaque.

Indeed, on this definition of knowledge, it seems the very foundation of what distinguishes mere belief (even true belief) from knowledge is wholly lost. For a belief based on ignoring, presupposing, and question-begging is not epistemically different from a belief based on mere guess-work or prejudice. To call such belief knowledge, in fact, makes little sense at all.

CHAPTER IV:
Relevant Alternativism

In the previous chapter I found that the problems which riddle the contextualist theory of knowledge gave conclusive reasons for rejecting both the contextualist's claim to refute the sceptic and, therewith, her claim to save — from the specter of scepticism — common sense. In this chapter, I examine the relevant alternativist's epistemological theory. The relevant alternativist's theory of knowledge is very similar to the contextualist's theory; so similar in fact that Keith DeRose, a prominent contextualist, calls relevant alternativism a "prominent *form of* contextualism."[1] Of course, if DeRose is correct in describing relevant alternativism as a *form* of Contextualism, then we should expect relevant alternativism to fail in precisely the same way as its explicitly Contextualist counterpart.[2] In the present chapter, I examine a number of such "relevant alternativist" theories of knowledge. I argue that properly conceived the relevant alternativists do in fact fail in much the same way as the contextualists, though, properly conceived, the relevant alternativists are more like distant cousins of the contextualists than immediate family. For though there are in some of the relevant alternativists certain fundamental

---

[1]Keith DeRose, "Contextualism and Knowledge Attributions," (1992): 19. (My emphasis.)

[2]Hereafter, I shall sometimes refer to contextualist as CTs, to the contextualist theory of knowledge as CT, and so on. I shall likewise often refer to the relevant alternativists as RAs, to the relevant alternatives theory of knowledge as RA, and so on.

similarities to the contextualists[3] (they allow, for example, a kind of contextuality to the conditions of knowing and, thereby, provide the foundation for a kind of contextualist account of knowledge), the most prominent relevant alternativist theories of knowledge are *not* contextualist theories *in the sense that* DeRose has in mind. The failures, therefore, are not to be attributed to the kinds of similarities DeRose believes the relevant alternativists have to the contextualists, but rather to a different set of similarities – a set responsible also for the shortcomings of contextualism – which he overlooks.

On DeRose's account, a contextualist is one who stipulates (after chapter III, I dare not say "argues") that the conditions of knowing are altered by contextual features, features which include mere *conversational* elements. The relevant alternativists, however, argue that such conversational elements alone can never change the fundamental conditions of knowing.[4] In fact, the relevant alternativist argues that the fundamental conditions of knowing never change. To the extent that the relevant alternativists differ from the contextualists on such positive theoretical commitments, they have distinct advantages over them. For example, though the relevant alternativists – like the contextualists – do sanction a kind of flexibility in the conditions of knowing, the contextualists allow for much greater flexibility and thereby (arguably) undermine the semantics of our ordinary concept of knowledge that much more. Moreover, unlike the contextualists, the relevant alternativists claim that there *are* certain conditions of knowing which are *necessary*,

---

[3]This is particularly true of Gail Stine. See: "Skepticism, Relevant Alternatives, and Deductive Closure," *Philosophical Studies*, 29, (1976).

[4]Strictly speaking, I am not aware of any RA actually arguing this directly. Yet each does imply this.

irrespective of context. To this extent, the relevant alternativist concept of knowledge accords much better with our ordinary beliefs. Furthermore, and as a consequence of their commitment to at least one immutable condition of knowing, the relevant alternativists are better able to maintain the integrity of the belief/knowledge distinction, a distinction wholly lost on their contextualist cousins. These differences must be kept in mind one is to appreciate the uniqueness of the relevant alternativist position.

Despite these important differences between the contextualists and relevant alternativists, however, I find that there are number of crucial similarities, similarities not surprisingly passed over by DeRose. These similarities lie not only in their respective approaches to the problem of scepticism, but also in their consequent failures to deal adequately with that problem. Both the contextualists and the relevant alternativists, for example, share the common goal of establishing that the sceptic's Cartesian doubts are *irrelevant* to the truth of our *ordinary* claims to know. To that extent they are, of course, similar, though not in a way which uniquely defines their anti-sceptical arguments. More important is the similarity evidenced in their respective arguments (or lack thereof) *for* that irrelevance. (A lack of argument sometimes touted as a virtue, or at least a right. See, for example, Dretske's "Entitlement: Epistemic Rights Without Epistemic Duties?" *Philosophy and Phenomenological Researach*, Vol. 60, No. 3, May 2000; and DeRose's "Ought we to Follow the Evidence?", *Philosophy and Phenomenological Research*, Vol. 60, No.3, May 2000). For it is here that each appeals to the same fallacious principle,

namely, the Principle of Epistemic Conventionality.[5] (In this respect, of course, the relevant alternativists share a fundamental characteristic not only with the contextualists, but also with the ordinary language philosopher.) Arguably, it is this fallacious mode of reasoning, of refuting the sceptic, which is the *most* important characteristic shared by each of these positions. For it is *this* fallacious mode of reasoning which proves fatal to each of their anti-sceptical positions.

The relevant alternativist's appeal to this fallacious principle, like the contextualist's and the ordinary language philosopher's, evidences a failure to distinguish two logically distinct sets of standards now familiar to my readers; namely, the set of standards governing appropriate assertion and the set of standards governing true assertion. In the final analysis, therefore, I find that the RA's *err* in precisely the same way as their Contextualist cousins: in their zeal to defend ordinary knowledge, they too rest too much on our ordinary linguistic habits; too much that is to refute the sceptic and rescue our ordinary knowledge from her apparent threat.

Section i: Relevant Alternatives and Epistemic Operators

There is one notable exception to this simple failure amongst the relevant alternativists, and that is Dretske. Fred Dretske is probably the most prominent relevant alternativist,

---

[5]This is something of an oversimplification, though it will be cleared up later in my discussion. In brief, though the RAs do each evidence this mistake — in a way that is fundamental to the details of their respective positions – Dretske does not make this mistake in at least his early paper, "Epistemic Operators," *The Journal of Philosophy*: 24 (December, 1970). Though, even there, there is reason to believe that this mistake is behind his overall relevant alternativist strategy.

and arguably the most sophisticated.[6] His seminal paper "Epistemic Operators"[7] was the first to attack the problem of scepticism in terms of this notion of relevant alternatives, and did so in such a way that the sceptic seemed (almost) to have been definitively refuted. His arguments there, in this early paper (though not the arguments of his later papers), go beyond the simple-minded appeals of the other relevant alternativists to our linguistic habits and focus, rather, on the logical behaviour of sentential operators. I devote, therefore, a greater amount of space here to Dretske than the other relevant alternativists. I also, therefore, examine Dretske's anti-sceptical arguments first, and only later look to the arguments of two other prominent RA's, Alvin Goldman and Gail Stine.[8] In my discussion of Dretske, I find that despite his detailed and persuasive arguments, his attack on scepticism ultimately fails. It fails because in the course of arguing for his anti-sceptical conclusions, central premises prove false and fundamental analogies fall apart.

But first, what is relevant alternativism? Its distinguishing feature as a theory of knowledge is its doctrine that the capacity to rule out *relevant* alternatives (to what we claim to know) is a *necessary* condition of knowing.[9] It is the relevant alternativist's uniform emphasis on this singular necessary condition of knowing that distinguishes their approach to scepticism from other anti-sceptical theories, and distinguishes it in particular,

---

[6]His sophistication, compared to the other Ras, will become quite clear by the end of the chapter.

[7]Dretske, "Epistemic Operators."

[8]Gail Stine, "Skepticism, Relevant Alternatives and Deductive Closure," *Philosophical Studies*, (1976); Alvin Goldman, "Discrimination and Perceptual Knowledge," *The Journal of Philosophy*, 73.20 (1976).

[9]This feature is unlike the sceptic's insofar as 'relevant' here implies the existence of the non-relevant.

from the contextualist's.[10] Knowledge does *not* — as the philosophical sceptic has argued — require ruling out *every* alternative which is incompatible with what we claim to know. It does not, for example, require ruling out the possibility that we may be mere brains in vats (or BIV's), or unwitting victims of an evil and deceiving genius. Knowledge also does not — contrary to what the contextualist has argued — require ruling out *mentioned*, though inane, alternatives. Rather, the relevant alternativist argues, it requires only the ability to rule out those alternatives which are *both* incompatible with *and relevant* to the knowledge claim that is made.[11] What determines relevance, the relevant alternativist argues, is not mere logical possibility, nor is it conversational utterance or assumption.[12] Rather, it is the (environmental) context in which the knowledge claim is made.[13]

Thus, when S claims to know that P, for example, there may be any number of alternatives which are *both possible* and incompatible with what she claims to know. In claiming to know *that* she has hands, for example, the *possibility* that S is really a mere

---

[10]As I will discuss in section v of this chapter, though the CTs also think some alternatives are irrelevant, they do not make ruling out relevant alternatives *a general* or necessary condition of knowing, and for that reason, are not, I think, properly characterized as relevant alternativists.

[11]See, for example, Dretske's "Epistemic Operators," Stine's "Skepticism," and Goldman's "Discrimination."

[12]I mean here to hold distinct what is conversational from what is linguistic habit.

[13]At this point, the RA position sounds very much like the contextualist's. But, as will become clear later on, the positions are quite distinct in virtue of which features of the context they allow to affect the relevance of any incompatible alternative. Moreover, to say that the determinant of relevance for the RA is context is something of an oversimplification. As I will illustrate, there is actually a bit of divergence on this issue, and even where there is divergence the actual criteria of relevance are at times somewhat unclear. At this point, however, to say that the determinant of relevance is context is an innocuous oversimplification.

brain in a vat, or that S is really the unwitting victim of a deceiving and evil genius are examples of such incompatible and *possible* alternatives. The sceptic, of course, has argued that *each* of these alternatives must be ruled out, if S is to correctly claim to *know*. The contextualist argued that such alternatives must be ruled out *only* if they are raised in the context of a knowledge conversation or if they are part of what the group perceives as relevant. The relevant alternativist, however, argues that the knowing subject need rule out only those alternatives which *are relevant* to P, and such Cartesian alternatives as these — she argues — are *not*.

In "Epistemic Operators," Dretske's effort to establish that the sceptic is wrong in her claim that we must rule out such incompatible alternatives proceeds somewhat differently than the other relevant alternativists. The argument Dretske offers proceeds indirectly and by analogy. There is no simple-minded appeal to our ordinary linguistic practices which, I will show, is a central feature of later relevant alternativist reasoning. There is no appeal to the intersubjectively opaque, or the conversationally raised or assumed, as with the contextualists. Instead, the focus of Dretske, in this paper, is on the logical behaviour of operators. The singular purpose of his analysis is to establish that the *epistemic operators*, in particular, are not fully penetrating. For, from here, Dretske thinks, the refutation of the sceptic is a *fait accompli*.

Before we can fully appreciate or evaluate Dretske's argument, one must then first be fluent in the various ways operators are distinguished, one from another. One must, for instance, be privy to the distinction between an epistemic and a non-epistemic operator, and also to the distinction between fully penetrating, semi-penetrating, and non-

penetrating operators. To this end, let us consider the following presentation of some of the more pertinent examples Dretske offers.[14]

On Dretske's view, some operators have the quality of being *fully* penetrating.[15] If O stands for the operator, and P the proposition on which O operates, then if O is fully penetrating, then the consequences which follow logically from P follow also from O(P). The following are examples of such fully penetrating operators:

It is *a fact that*
It is *necessary that*
It is *true that*
It is *possible that*

Thus, If Q is a necessary consequence of P, it follows that: if it is a *fact that* P, then it is a *fact that* Q; if it is *necessary that* P, then it is *necessary that* Q; if it is *true that* P, then it is *true that* Q, and so on.[16] In other words, no matter what the consequences of P, if we apply a *fully* penetrating operator to P, the operator will penetrate or "carry over" to those consequence of P.

That some operators are non-penetrating is obvious. (Dretske's example is: though it is *strange that* P, and Q is a necessary consequences of P, it is not necessarily *strange that* Q.[17]) What about the *epistemic* operators? Are *they* fully penetrating? Is it the case that if I *know that* P, and I know that Q is a necessary consequence of P, I *know that* Q?

---

[14]To the reader already familiar with these distinctions, I apologize.

[15]Dretske, "Epistemic Operators," 1007.

[16]Dretske, "Epistemic Operators," 1007.

[17]Dretske, "Epistemic Operators," 1008.

Dretske argues it is not. He maintains that *all* of the epistemic operators are semi-penetrating only: they are such that the operator does not penetrate to *all* of the necessary consequences of the sentence it operates on, though they do penetrate to some. In particular, he argues, it is the case with the operator 'knows that'.[18]

Thus, for example, if S claims to *know that* P, and Q is a necessary consequence of P, and moreover, S *knows* that Q is a necessary consequence of P, it *does not follow* that S *knows* that Q, *even though* she *knows that* P *and knows* that Q is a necessary consequence of P.[19] [20]

To illustrate: Suppose that S claims to *know that* she has hands (P). A necessary consequence Q (admitted by most, and, in particular, by Dretske) of her having hands (of P) is that she not be a mere BIV (since, obviously, mere BIVs have no limbs or appendages at all). Suppose further that S *knows that* Q (she is not being a mere BIV) is a necessary consequence of P (her having hands). It does not follow, according to Dretske, that she *knows that* she is not a mere BIV. For according to Dretske, the operator *knows that* does not penetrate to such consequences. Therefore, even if S *knows that* P, and *knows that* Q is a consequence of P, she need not *know that* Q.

Clearly, if Dretske's analysis of the operator *knows that* is correct, and, in particular, if she is right that one may know that P, and know that P entails Q, without knowing that Q, this has important implications for the *relevance* of the challenges posed by the sceptic.

---

[18]Dretske, "Epistemic Operators," 1009.

[19]For the sake of argument and simplicity, for now I will assume that if *knows that* is only semi-penetrating then it does not penetrate to any of the known consequences of P.

[20]Dretske, "Epistemic Operators," 1010.

For the sceptic has argued, in effect, that if S is to know that P, and if P is known to entail Q, then she must know that Q (where the Q's in question, of course, include not being a BIV, not being the victim of an evil genius, and so on). In fact, according to Dretske, *almost all* sceptical objections trade on this pattern of argument.[21] He maintains that:

> Whether it be a question of dreams or demons, illusions or fakes, the same pattern of argument emerges. If you know this is a chair, you must know you are not dreaming (or being deceived by a cunning demon), since its being a (real) chair entails it is not simply a figment of your imagination.[22]

Whether Dretske is correct or not about how prevalent this pattern of reasoning is, the pattern is certainly familiar enough that Dretske's argument against it must be taken seriously.[23] For Dretske, if successful, will have shown the Cartesian kinds of doubts and alternatives that are thought to threaten our ordinary knowledge are actually *irrelevant* thereto. As he states: Whether or not a certain objection *is* relevant *depends* on whether the consequence is one to which the operators penetrate. [24] Moreover, he states:

> What [he] wish[es] to argue . . . is that the traditional sceptical arguments exploit precisely those consequences of a proposition to which the epistemic operators *do not* penetrate, precisely those consequences which distinguish the epistemic operators from the fully penetrating operators.[25]

---

[21]Dretske, "Epistemic Operators," 1011.

[22]Dretske, "Epistemic Operators," 1011-1012.

[23]Clearly many sceptic's have argued for the much stronger thesis that to know that P one must know all the consequences of P, whether these are known to follow from P or not. But since Dretske confines himself to refuting the weaker version of this thesis (knowing the known consequences of P) I will do the same. See Dretske, "Epistemic Operators," 1010.

[24]Dretske, "Epistemic Operators," 1012.

[25]Dretske, "Epistemic Operators," 1012. (My emphases.)

Insofar as these kinds of consequences are the primary ones which have proved both enigmatic and retractable to the history of epistemology, Dretske's anti-sceptical argument, if successful, will then have lain to rest a substantial source of philosophical doubt.

Dretske calls the thesis behind this pattern of reasoning the thesis of *penetrability*.[26] [27] For, in effect, the pattern rests on the assumption that the epistemic operator 'knows that' *penetrates to* all the known consequences of the object of knowledge: Such that if Q is a consequence of P, and S knows that P, and knows that Q is a consequence of P, then S *must* know that Q, else she is only incorrectly described as *knowing that* P.

Up to this point in his reasoning, however, Dretske, has merely stated how one *might* undermine the sceptical challenges to our knowledge claims. He has stated that *if* the epistemic operators do not penetrate to the kinds of alternatives the sceptic appeals, *then* those alternatives will prove both irrelevant to and, therefore, inefficacious against our ordinary claims to know. The question is: how does he proceed from here? How does Dretske establish that the epistemic operators do not penetrate to such consequences?

Dretske's argument is fairly complex (much more complex than that offered by any other relevant alternativist). As I mentioned earlier, he proceeds indirectly and by analogy; arguing first that the *epistemic* operators have two characteristics in common with what he

---

[26]Dretske, "Epistemic Operators," 1011.

[27]This same thesis is also, of course, discussed under the name *principle of deductive closure*, or simply deductive closure by both Stine and Nozick, among others. See, for example, Stine's "Skepticism," and Nozick's *Philosophical Explanations,* Cambridge: Harvard University Press, (1981).

calls the *explanatory* operators,[28] and arguing then that because the *explanatory* operators possess the third characteristic of not penetrating to the "contrast consequences" of explained propositions, the epistemic operators probably (and by analogy) also possess this characteristic.[29] Significantly so. For if they do share this characteristic with the explanatory operators, they will be quite inefficacious in the sceptical arguments. For, Dretske argues, it is precisely such contrast consequences which the sceptic argues we must know if we are to know anything at all.

But what characteristics do the epistemic operators share with the explanatory operators? The first, and simplest to prove, is the characteristic of semi-penetrability. Or, I should say, not non-penetrability. For, as Dretske argues the epistemic operators are at least more than non-penetrating. This can be seem from elementary logic. If, for example, S *knows that* P *and* Q, it follows, necessarily that S *knows* that Q, a logical consequence of what she knows, namely P and Q. [30] More importantly, however, and less obviously, Dretske argues that the epistemic operators share with the explanatory operators the further characteristic of not penetrating to the "presuppositional consequences" of a proposition.[31] What precisely a "presuppositional consequence" is (or perhaps I should say *why* a presupposition is considered to *be* a consequence), is not particularly clear.

---

[28]Explanatory operators include the operators is *a reason (justificatory) for*, is *a reason (explanatory)* for, and *would not be the case unless*. Dretske, "Epistemic Operators," 1017.

[29]The analogy is admittedly weak, though at times persuasive. There remains, however, a somewhat stronger analogy hidden within this rather weaker more explicit one. This will be detailed forthwith.

[30]Dretske, "Epistemic Operators," 1009.

[31]Dretske, "Epistemic Operators," 1012-1014.

Such a consequence, however, appears to be what most of us would call simply a *presupposition*. Dretske, however, views presuppositions as a "*class of consequences* to which the epistemic operators do not penetrate."[32] And because establishing that the epistemic operators do not penetrate to such presuppositional consequences is a central element in his analogy, one must examine with care how he claims to do this.

Dretske claims to establish that the epistemic operators do not penetrate to presuppositional consequences with two examples. Dretske's examples focus on the epistemic operator *realizes that* and purport to establish that a given proposition can be *realized*, even though a necessary presuppositional consequence of that proposition is not. The examples, therefore, serve to undermine the general claim that the epistemic operators penetrate to *all* the consequences of a proposition: If they do not penetrate to all of the *presuppositional* consequences, they do not penetrate to *all* consequences. Moreover, and more importantly, the examples serve to establish that the epistemic operators share this second characteristic – of not penetrating to the presuppositional consequences – with the explanatory ones.[33]

Dretske's first example: Dretske asks us to imagine that he is traveling on a train with his brother. And, as is typical of rude siblings (of which, I must say, I have none), his brother refuses to give up his seat to an older lady who boards the train. Dretske proposes that the correct description of the older lady's perception of his brother's

---

[32]Dretske, "Epistemic Operators," 1014. (My emphasis.)

[33]I do not address the issue of whether the *explanatory* operators penetrate to the presuppositional consequences because inasmuch as Dretske fails to establish the epistemic operators do, this is sufficient to undermine his claim that the two sets of operators *share* this characteristic.

behaviour is the following: she *realizes that*[34] his (Dretske's) brother will not move.[35]

Importantly, however, a (*presuppositional*) consequence of *what* the older lady realizes is *that* the man on the train *is* Dretske's brother. And though, according to Dretske, this is a necessary consequence of what she knows (*that* his brother will not move), the lady need not know *this* (*that* this is Dretske's brother) to know *what* she does know (that his brother will not move). Thus, Dretske concludes, the epistemic operator *realizes that* does not penetrate to the *presuppositional* consequences, and therefore, does not penetrate to *all* consequences.

Dretske's next example takes the same form as the last and is meant to establish the same thing. When I have a reason to believe *that* the Church is empty, he says, a *necessary consequence* of *what* I have a reason to believe is *that* the building (which I have a reason to believe is empty) *is* a church. But, he says, I *need not* have a reason to believe *this* (that the building is a church) to have *a reason to believe that* the church is empty. Therefore, again, the epistemic operator *reason to believe that* does not penetrate to *all* the consequences of a proposition. In particular, it does not penetrate to its presuppositions.

Dretske realizes, of course, that it is not the presuppositions of a proposition that the sceptic is interested in. So, one must ask, what is the point of proving of epistemic operators that they do not penetrate to the presuppositions of a proposition? Again, the answer lies in Dretske's desire to establish that at least some of the epistemic operators

---

[34]At times he says *knows that*. See, Dretske, "Epistemic Operators," 1013.

[35]Dretske, "Epistemic Operators," 1013.

share *some* characteristics with the explanatory operators; for it is only from here that he can argue *analogically* that they probably also share others. Establishing the latter is quite important since it is only if he can establish a *likeness* of the epistemic operators to the explanatory operators that he can proceed *analogically* to claim that the epistemic operators probably also share the characteristic of not penetrating to the contrast consequences of the proposition which they operate upon. It is this latter class of consequences which Dretske believes the sceptic exploits, and to which he argues the epistemic operators do not penetrate.[36] Hence establishing that the explanatory operators do not penetrate to the presuppositions is really only a secondary move; one which enables him to get to the class of consequences he is really interested in, namely, the contrast consequences.[37]

Why does Dretske think that such contrast consequences are exploited by the Sceptic? And how might this exploitation be related to the thesis of penetrability; a thesis he claims is also fundamental to the sceptic's challenge? "The exploitation," Dretske claims, "proceeds as follows":

> Someone purports to know that X is A, that the wall is red, say. The sceptic now finds a predicate 'B' that is incompatible with 'A'. In this particular example, we may let 'B' stand for the predicate 'is white'. Since 'x is red' entails that 'x is not white' it also entails that x is not-(white and Q) . . . . Therefore the sceptic selects a 'Q' that gives expression to a condition . . . under which a white wall would appear exactly the same as a red wall . . . let

---

[36]Dretske, "Epistemic Operators," 1015.

[37]A 'contrast consequence' is a consequence of what we claim to know which is incompatible with (or 'contrasts' with) what is known. Therefore, for example, it is a consequence of my knowing that I have hands that I am not a BIV (since, among other things, BIVs do not have hands).

'Q' stand for: cleverly illuminated to look red.[38]

The sceptic, of course, proceeds to argue that unless S knows that not (B and Q), she cannot know that A. For a necessary consequence of A is that -(B and Q). This reasoning, Dretske argues, rests on the thesis of penetrability: that if S knows that P, and S knows that P entails Q, then S must know that Q.[39] Or, stated differently, if P entails Q and S does not know whether Q is true or false, it follows that S does not know that P. The question, Dretske claims,[40] is whether we do in fact ever know any such thing?[41] Do we, when we claim to know that the wall is white, know that it is not merely cleverly illuminated to look white? A possibility the sceptic argues we must know the falsehood of to know that the wall is red. Dretske argues that though we habitually take such things for granted, in particular we habitually assume that the conditions under which we make our knowledge claims are normal, we do *not* have sufficient evidence to know any such thing.

Importantly, he concedes this much to the sceptic. Moreover, Dretske goes on to illustrate his agreement with the sceptic on our lack of such general knowledge with his notorious zoo example. Suppose, he says, you take your son to the zoo. Suppose, moreover, that while there you point to a group of zoo animals and profess to *know* that they are zebras. One consequence of the zoo animal's *being* zebras is that they are not

---

[38]Dretske, "Epistemic Operators,"1015. I think this description distorts the sceptic's motivation, though since this does not affect Dretske's overall argument, I won't pursue it here.

[39]Dretske, "Epistemic Operators," 1011.

[40]Though, as will become evident, this is not the question all relevant alternativists would ask.

[41]Dretske, "Epistemic Operators," 1015.

merely mules cleverly painted to look like zebras. Yet, Dretske concedes, even though their not being cleverly painted mules is a consequence of what you purport to know, it is not something you do know. You know, he argues, no such thing. At least, *ex hypothesi*, you have not tested the animal's hair for paint, and you know sufficiently little about zoology that you could easily be fooled by a painted mule into believing it is a zebra. Yet, in spite of your ignorance, Dretske claims, *you do know* it is a zebra.[42]

The point of Dretske's example is general: when we make a claim to know, normally we do not *know* that the conditions under which we make our claim are normal. We do not normally check, for example, whether we might be dreaming. We do not inspect for phony lighting when we enter a room. We cannot (and do not try to) prove we are not mere brains in vats. Dretske concedes all of this. Nonetheless, Dretske argues, we do know the things we claim to know: that the animals in the zoo are zebras, that the wall is white, and so on. If he is to support these contentions, Dretske must, however, find a way to reject closure. For unless it is not the case that a condition of knowing that P is knowing all the known logical consequences of P, he must, having conceded we fail to know these, admit that we fail to know generally. (In this Dretske differs markedly from the contextualist, whose strategy here is, as described by DeRose a "non-heroic Moorean" one. See DeRose, "How Can we Know we are not Brains in Vats?", *The Southern Journal of Philosophy*, Vol. 38, 2000.) It is here, of course, that Dretske's analogy is supposed to do its important work.

Because of the analogical form his argument takes, Dretske must first establish that the

---

[42]This, of course, is quite reminiscent of Austin's goldfinch example. Cf chapter II of this work, 8-11.

explanatory operators (*explains* that, is a *reason that*, *would not be the case unless*, and so on) do not *themselves* penetrate to contrast consequences in order to establish the further conclusion – by analogy – that the *epistemic* operatorss do not penetrate to those consequences. He must, that is, establish that the explanatory operators do not penetrate to the particular *type* of consequences which the sceptic "exploits." For from here, Dretske thinks, by analogy, we may conclude that the *epistemic* operators do not penetrate to contrast consequences either and, therefore, that the thesis of penetrability — upon which the edifice of scepticism is built — is false.

Dretske, to this end, presents us with another example.[43] Suppose, he says, that after having dinner with Brenda, I notice that Brenda orders no dessert. There are a number of contrast consequences (consequences, that is, that contrast with the fact that she did not order desert) which follow necessarily from the fact that Brenda did not order dessert. It is implied, for example, that Brenda did not order dessert *and* throw it at the waiter. It is also implied that Brenda did not order dessert *and* take it home to a sick friend. When, however, I *explain why* Brenda did not order dessert by remarking that she was not hungry, I do not explain either of the aforementioned consequences, even though they are both necessary (contrast) consequences of the *fact* that she did not order dessert; implied, that is, by the fact I *have* explained. We have then the following set of entailment relations:

Let P stand for the fact *that* Brenda did not order desert.
Let Q stand for the (contrast) consequence of P, *that* Brenda did not order desert *and* throw it at the waiter.

---

[43]Dretske, "Epistemic Operators," 1021.

Let R stand for any consequence at all.
Q is a (contrast) consequence of P, and
O(P) *does not* entail O(Q), nor O(Q&R).

That is to say, the *explanatory* operator *explains why* does not penetrate to the contrast consequences of the proposition upon which it operates. For, even though it is implied that she did not order desert *and* throw it at the waiter, explaining the fact that she did not order desert does not explain why she did not order *and* throw it at the waiter; it does not explain one of the consequences of what it does explain. Returning to his analogy, Dretske therefore concludes that the *epistemic* operators also do not penetrate to the contrast consequences of the proposition upon which they operate. As he reasons: the epistemic operators share each of two characteristics with the explanatory operators (they are not non-penetrating, and they do not penetrate to the presuppositional consequences of the propositions upon which they operate). Therefore, he reasons, it is *likely* that the epistemic operators share also this third characteristic possessed by the explanatory operators, viz., they also do not penetrate to the contrast consequences.[44]

At this point, I think, Dretske's analogical argument appears so weak as to stand in need of no refutation. For those, however, who find Dretske's analogy persuasive, I offer the following objections. First, even if Dretske has established his claim that the epistemic operators share the first characteristic (not non-penetrability) with the explanatory operators, the illustrations he uses to establish that they share the second characteristic (not penetrating to the presuppositional consequences) fail.[45] Recall his proof took the

---

[44]Dretske, "Epistemic Operators," 1022.

[45]In so far as I am right about this, it would be a waste of paper to discuss whether or not he succeeds in establishing this for the explanatory operators. For, if he fails to establish that the

following form: S *e's* that P (where 'e' is any epistemic operator); a presupposition of P is that Q; yet S need not 'e' that Q. Thus he argued that though the lady on the train *realized that* my brother would not move, she did not (need to) *realize that* the person who would not move was my brother.

But what grounds does Dretske offer for claiming that what she realized was that *his brother* would not move? He offers no argument. Merely assertion. Yet most of us, I think, would reject this description of *what* the lady realized. We would reject it because it is false. By hypothesis, the lady on the train had no idea what relation Dretske's brother stood in to him, if any. It is, therefore, simply incorrect to describe her as knowing or realizing *that* Dretske's *brother* would not move. The correct description, given her ignorance of Dretske's familial relations, is this: she realized that *the man* sitting to another man's side would not give up his seat. Without further information, she could not, I think, *realize* anything more.

Hence, even if the *explanatory* operators happen to possess the characteristic of not penetrating to the contrast consequences of the proposition upon which *they* operate, Dretske has not proven that the *epistemic* operators do. He has also, therefore, *not* established that the two sets of operators share two characteristics, but one. This, however, is surely not reason enough to accept his conclusion that they probably also share the characteristic of not penetrating to the contrast consequences. Yet the minimal likeness of the two sets of operators to one another is the only reason Dretske sufficiently

---

epistemic operators have this characteristic, whether the explanatory operators have it is moot: their doing so would contribute nothing to the analogy.

defends for concluding that the epistemic operators do not penetrate to the all important contrast consequences.

Though unnecessary, I think, in light of my first objection, one might offer the further objection that Dretske fails also to establish that the purported consequences to which the epistemic operators do not penetrate are *consequences*, and therefore, fails not only to show that they share this characteristic with the explanatory operators, but also to provide a counter-example to the thesis of penetrability. For, even if we allowed his description of what the lady realized on the train, his own description of such presuppositions *as* a class of *consequences* is quite unconvincing. Presuppositions are not consequences. They are presuppositions.

Finally, of course, there is the problem of the weakness of the form Dretske's argument takes. For even if he had established that the epistemic operators shared *two* characteristics with the explanatory operators, such minimal likeness is a very weak footing upon which to base his conclusion that they share the third characteristic, not penetrating to the contrast consequences. It assumes, among other things, that the two "sets of operators share the same logic in this respect."[46] Indeed, his entire refutation rests on this assumption. But what grounds are there for believing that these distinct sets of operators — the epistemic and the explanatory — share the same logic here? Dretske provides a partial, though somewhat opaque and unsatisfactory, answer. His answer is articulated in the context of his discussion of the *explanatory* operators and their failure to penetrate to the contrast consequences of the propositions upon which they operate.

---

[46]Dretske, "Epistemic Operators," 1019.

In discussing the explanatory operator *explains that*, for example, Dretske argues that when we give *an explanation* of a particular fact, *what* is being explained is a function of two things:

1. The fact
2. The Range of relevant alternatives[47]

Thus, for example, when I explain *why* Brenda did not order dessert, I explain why she did this (the fact) *as opposed to* ordering dessert *and* eating it (the *relevant* alternative). Hence, my explanation 'she wasn't hungry' succeeds *relative to* this fact and this relevant alternative. However, if we *change* the relevant alternatives – to include ordering dessert *and* throwing it at the waiter, for example – my former explanation is no longer *explanatory*. (This, of course, is just another way of saying that the explanatory operators do not penetrate to contrast consequences. For that she did not order dessert and throw it at the waiter *is* a contrast consequence of what I explained originally; namely, that she did not order dessert. Yet my having *explained* why she did not order dessert does not *explain* why she did not order dessert *and* throw it at the waiter.) For the fact that she was not hungry does not *explain why* she did not order desert and throw it at the waiter. *What* I explain, therefore, is rightly argued by Dretske to be a function of these two factors. I cannot change the relevant alternatives (or, of course, the fact) without thereby changing both *what* I explain and *whether* the explanation I have given is explanatory.

This much, then, is clear: *with respect to explanation*, when I change the range of

---

[47]Dretske, "Epistemic Operators," 1021. As a point of historical interest, this is the first occurrence of the phrase 'relevant alternatives.'

alternative facts, I change thereby the explanandum. In so doing, I change also what will count as adequate explanans. Dretske goes on, however, to conclude, without further argument, "So it is with our epistemic operators":[48]

> Within the context of explanation *and* within the context of our other operators, the proposition we operate on must be understood as embedded within a matrix of relevant alternatives.[49]

> *To know* that x is A is to know *that* x is A within a framework of relevant alternatives . . . . This set of contrasts [the *relevant* alternatives], together with the fact that x is A, serve to define *what* it is that is known when one knows *that* x is A. One cannot change this set of contrasts without changing *what* a person is said to know when he is said to know *that* x is A.[50]

Dretske – in the above cited passages – has argued only that when I change the range of relevant alternatives, I change *what* I am purported to know. Thus, for example, while at the zoo, if I claim to know that the animals are zebras *as opposed to* gazelles, giraffes, or bison, I claim to know something different from *what* I claim to know when I change the range of relevant alternatives to include cleverly painted mules. The change in alternatives causes a corresponding change in *what* I am purported to know.

To complete the analogy, however, and establish that the sceptical alternatives are not relevant to the truth of my knowledge claim – to what I am (intending) to claim to know (less they change the very thing that I claim to know), Dretske *needs* to argue that changing the range of relevant alternatives changes *whether* I know what I claim to know (*just as* changing the relevant alternatives in the case of explanation changes *whether* my

---

[48]Dretske, "Epistemic Operators," 1022.

[49]Dretske, "Epistemic Operators," 1022. (My emphasis.)

[50]Dretske, Epistemic Operators," 1022.

proposed explanation is explanatory); not just *what* I know. Without this further argument the analogy is incomplete. Dretske, however, does not provide further argument. He seems to think that it is obvious that this further conclusion follows.

But how, exactly, might the (purported) fact that changes in the relevant alternatives change *what* I claim to know *establish* that those alternatives are irrelevant to my (original) knowledge claim and hence to *whether* I know *that*? I think Dretske's reasoning is as follows: Suppose I claim to know that 'S knows Lefty killed Otto.'[51] *What* I claim S knows is a function of which alternatives are relevant to the fact S is purported to know.[52] Dretske illustrates this in terms of the variety of emphases we might place on the proposition that S is purported to know. Thus, for example, I might be claiming of S that he knows the *identity* of Otto's killer (that *Lefty* killed Otto). The relevant alternative is that *a non-lefty* killed Otto. Or, I might be claiming of S that he knows Otto was *killed*, rather than merely injured (S knows Lefty *killed* Otto). *What* I claim of S to know here is different from *what* I claimed S knew when I claimed he knew that *Lefty* killed Otto. And so, Dretske thinks, are the conditions of knowing. For if *what* I claim to know is changed by the changing relevant alternatives then so must be, he thinks, the conditions according to which my claim to know should be evaluated. If I don't (intend to) claim to know that Otto was not killed by someone other than Lefty, then I do not have to rule this alternative out. Relative to what (I intend to) claim to know, this alternative is not relevant, even though it is a necessary consequence of what I claim to know. At least this must be the

---

[51]Dretske's example, "Epistemic Operators," 1022.

[52]I think what Dretske means here is this: what I *intend* to claim to know is a function of which alternatives I *intend* to rule out/view as relevant.

case *if* the epistemic operator *knows that* functions in the same way as the explanatory operator *explains that*.

When I claim to know that the animals in the zoo are zebras, then, as opposed to gazelles, I need not rule out that they are merely painted mules cleverly disguised to look like zebras. Not being part of my (intended) set of relevant alternatives, such an alternative, Dretske thinks, is irrelevant to what I claim to know, and, therefore, to whether I know what I claim to know. For, Dretske maintains, in claiming that the animals are zebras as opposed to gazelles, I *am not (and normally neither are you) claiming* to know that they are not merely cleverly painted mules. Hence even though their not being cleverly painted mules *is* a necessary consequence of *what* I claim to know, it is not a relevant alternative and hence not something I must know in order to know *that* they are zebras. Hence the fact that I cannot know that the zoo conditions are normal does not undermine my claim to know that the zoo animals are zebras. So Dretske argues.

> What I am suggesting is that we simply admit that we do not know some of these contrasting 'sceptical alternatives' are not the case, but refuse to admit that we do not know what we originally said we knew.[53]

But is Dretske right about this? Is it the case that knowing functions like explaining when it comes to relevant alternatives? Is it the case that if a particular alternative to what I claim to know is not part of what I (intend to) claim to know, I need not know it to know what I do (intend to) claim to know? Is it possible for me to selectively intend which alternatives are relevant to my claim to know and which are not? Is it the case that

---

[53]Dretske, "Epistemic Operators," 1016.

the thesis of penetrability is false, and the alternatives which *are* (on all accounts) necessary consequences of what I claim to know need not be known for me to know what I claim to know? Do alternatives to what I (intend to) claim to know and changes therein function like alternatives to what I claim to explain? I think not. Let me explain.

Suppose that I claim to know *that* I have hands. And suppose further that I do not (intend thereby to) claim to know *that* I am not a brain in a vat, but only that I am not an amputee (a necessary consequence of *what* I claim to know). Is the *possibility* that I might be a mere brain in a vat *irrelevant* to what I claim to know? I think it is arguable that it is irrelevant to what I *intend* to claim to know, but not to *whether* I do know what I claim to know.[54] For unlike explanations, where what I claim to explain is a function *both* of the fact and the intended relevant alternatives, in the case of knowing, while *what* I claim to know is a function of the fact and the limited alternatives I have in mind, *whether* I know is not, at least not exclusively. Whether I know is admittedly a function of what I claim to know, but it is also a function of the objective criteria of knowledge which must be fulfilled *no matter what* particular I may (*intend* to) claim to have knowledge about. Unlike the explanatory operator, therefore, changes in the set of criteria for *what* I know do not affect changes in the criteria for *whether* I know.[55]

---

[54]For, if we allow it is possible for me to intend to claim only this much, then arguing that I do not know more than I intend, does not undermine what I claimed to know. As I will argue, however, it is not at all clear that 'know' functions in such a way as to allow such selective intending.

[55]Actually, although I do not have time to pursue this here, 'explaining' is arguably just like 'knowing' insofar as there are certain conditions which must be met if something is to count as an act of explanation (as opposed, for example, to exclamation, imitation, or denunciation) at all. And these remain so as to enable us to distinguish the kind of act that is explaining from other kinds of acts (like exclaiming, imitating, and denouncing).

To illustrate: when I claim to know that I have hands, no matter what alternatives *I* have in mind which govern *what* I claim to know, the set of criteria which must be met are constant, if I am to know anything at all. And, what the sceptic argues, of course, is that these criteria of knowing *in general* include the following: I must know that I am not a handless brain in a vat. I must know I am not constantly deceived by an evil genius. And, the sceptic argues, I must know these whether my claim *to know* is about hands, ships, sealing wax, cabbages or kings. For in each case, she argues, a necessary condition of its being so is that none of these Cartesian alternatives obtain. One does not refute the sceptic, therefore, unless one also shows that the general conditions of knowing anything at all (i.e., of correctly describing myself as knowing, as opposed to merely believing) do *not* include the kind of Cartesian alternatives that the sceptic argues are ineliminable.

Because the necessary conditions of knowing anything at all are always a constant and proper subset of the conditions of what I know, (just as the necessary condition of 'doing', 'singing' and 'thinking' anything at all are always a constant and proper subset of the conditions of what I do, sing, or think,) shifting what I claim to know *never* affects a shift in this (constant and proper) subset of conditions governing whether I know. And whether I claim to know that x is a tree, a hand, or a wall, I must, the sceptic argues, in each case know that the object is not phony in any respect. I must know it is not the result of a mere dream, a manipulative evil genius, or clever lighting. For these simply are the conditions of knowing. Thus, even with changes in fact (with *what* I claim to know) the range of relevant alternatives (with respect to knowing) remains constant and does not change whether I know, for these alternatives represent *general* conditions of knowing

anything at all.[56]

That this is so comes out readily when, rather than focusing on the intentions of the knower, we look at the following kinds of assertions the knowing subject might make, or at least be coaxed into making when the implications of his (and Dretske's) position are drawn out. [57] Take, for example, the father at the zoo. If queried whether he knows the zoo animals are zebras, he would, of course, say "Yes." If we go on to inquire whether he knows they are not painted mules, he will (at least *Dretske will*) say "No," he does not. We have, then, the following conjunctive statement "I know they are zebras, though they might, for all I know, be painted mules."

Or, to take another of Dretske's examples, suppose I place what appears to be an ordinary tomato before you. Suppose I ask you if you ***know*** it is a tomato. You assert that you do. But, when queried, you also assert that it might be a fake tomato, or merely an image of a tomato produced in your envatted brain by a deceiving and powerful super computer, for all you know. We then have the conjunctive statement that "I know it is a (real) tomato, though (for all I know) it might be a fake."

The sceptic argues (and I am inclined to agree) that each of these statements is absurd, and the reason they are absurd is because they are contradictory.[58] They are contradictory

---

[56]This is not to say, of course, that when I claim to know x is a tree I must know it is not a cleverly painted mule, or that when I claim to know that x is a zebra I must know it is not a cleverly painted tree. Rather, again, it is to say that no matter what I claim to know, if the sceptic is correct, then I must always know the object of my knowledge is not phony, an imposter. Whereas when I claim to have explained something, I need not explain every possible alternative, but only the one I mean to include in my set of relevant alternatives.

[57]I owe this insight to Peter Unger, *Ignorance*, Clarendon Press: Oxford, (1975), 140.

[58]I owe this also to Unger, *Ignorance*, 139-40.

because a condition of knowing anything at all is knowing that the object of knowledge is not fake in any way. And a condition of knowing the object is not a fake is knowing that you are not an envatted brain. A further condition is knowing you are not merely dreaming. For, by definition, the objects of your perceptions when your brain is envatted or dreaming are paradigmatically fake. The Moral? Dretske's analogy fails. Why does it fail? It fails because it does not take into account the fact that certain conditions of knowing are constant, and do not themselves change with the changing intentions of what I claim to know.

At this point, Dretske has at least two possible alternatives: first, he might try to fall back on the main analogy; second, he might return to the (purported) fact that I can intend to know something more limited. Let's consider the second possibility first. Can I, simply by intending to know only that the zebras are not gazelles, for example, rightly claim to know that they are zebras, as opposed to gazelles. I think that the sceptical arguments above make a strong case that I cannot. For, there are certain *general* conditions on knowing *anything*, and if these include knowing that the object of your purported knowledge is not phony in any way, then one cannot ignore these possibilities. Once you make a knowledge claim, therefore, no matter how limited you might *intend* it to be, you are bound by the general conditions of knowing. And, as I argued above, the alternatives you do not claim to know might nonetheless a necessary component of the conditions of what you do (intend to) claim to know, since they represent the conditions of 'knowing' anything at all. Thus, to refute the sceptic, one must show that these general conditions do not include those which the sceptic has argued they do.

Therefore, even if you intend only to claim to know that the zoo animals are not gazelles, your knowledge claim has a much wider logical implication: in claiming to know the zoo animals are not gazelles, you thereby — simply in virtue of making *a knowledge-claim* — claim to fulfill the general conditions of knowing. And, unless you establishe that these do not include knowing they are not phony in *any* way, you cannot correctly claim to know merely that the zoo animals are zebras (as opposed to gazelles) simply by selective intending. For, just as claiming to have *done* something implies claiming to have fulfilled the conditions of *doing*, and claiming to have *made* something implies claiming to have fulfilled the conditions of *making*, so also, claiming to *know* something implies claiming to have fulfilled the conditions of *knowing*, and these, the sceptic argues, go far beyond the mere ruling out that zoo animals are not gazelles.

What about Dretske's potential to fall back on his main analogy? Recall Dretske argued that the epistemic operators share two characteristics with the explanatory operators and, therefore, probably also share a third characteristic, namely, not penetrating to the contrast consequences. This analogy was meant to establish that the thesis of penetrability was false. For if the contrast consequences were not part of what was operated on when the claim to know was made, then it would be false to assert that if Q is a logical (contrast) consequence of P, and you know that Q is a consequence of P, you must know that Q (in order to know that P). This response, however, has been already shown to fail. I argued above that not only is the general form of the argument very weak, it was not established that the epistemic operators share this characteristic — not penetrating to presuppositional consequences — with the explanatory operators. Hence, even if the

weakness of the argument form were not a factor, the analogy would still fail, since Dretske does not establish the elements of the analogy.

Section ii: Does Henry Know? Stine's Reason to Believe and the Simple Approach to Relevant Alternatives

In the preceding pages, I have argued that Dretske's arguments in "Epistemic Operators," though somewhat beguiling and more complex than those offered either by his contextualist cousins or (as we shall see) his relevant alternativist colleagues, fail. They fail to substantiate the conclusion that the sceptical alternatives to what we claim to know need not be known to know what we claim to know, and they fail, therefore, to support the claim that the thesis of penetrability is false. I now turn, therefore, to the relevant alternativist positions of Dretske's early contemporaries, Gail Stine and Alvin Goldman. The arguments of Stine and Goldman, though ultimately relevant alternativist in nature, differ significantly from the arguments of Dretske's early paper. Stine, for example, is a relevant alternativist but maintains that the thesis of penetrability need *not* be shown to be false in order to establish that we know what we claim to know. And Goldman argues for the irrelevance of at least some alternatives without any reference to the logical behaviour of operators, epistemic or otherwise.

Insofar as these later relevant alternativist arguments differ quite markedly from Dretske's early ones, they hold out hope for the RA's to reach their anti-sceptical position. To the rue of many an anti-sceptic, however, I find that they too are unsuccessful in their efforts to refute the sceptic. Not only are they – like Dretske – unsuccessful in arguing for the major theses of relevant alternativism, but because they fall back on the simple-minded

kind of anti-sceptical argument offered by both the ordinary language philosophers as well as the contextualists,[59] their failures are only too familiar. Moreover, because of the greater simplicity of their arguments, their failures are both more obvious than Dretske's and, therefore, easier to establish.

Like Dretske, both Stine and Goldman argue that knowing is not a matter of ruling out *all* alternatives to what we claim to know, but rather, is a matter of ruling out only those alternatives which are *relevant* to what we claim to know. And, like Dretske, they argue that the relevant alternatives are not identical with the set of all possible alternatives, but instead are a proper subset thereof. Stine and Goldman's arguments for the irrelevance of some alternatives are, as I have indicated, much simpler. They proceed *not* from the perspective of epistemic logic, but rather, from certain assumptions about ordinary language and *its* function in defining what it means to know. It is from this springboard of ordinary language that both Stine and Goldman argue that at least some alternatives to what we claim to know are not relevant. Each fails, however, because each rests her/his conclusions about what it means to know on our ordinary linguistic practices. Each fails also, therefore, to defeat the sceptic and so also to vindicate (what she/he takes to be) ordinary knowledge.

Stine's paper, "Scepticism, Relevant Alternatives, and Deductive Closure," like Dretske's "Epistemic Operators," is a defense of the relevant alternativist view. She – like he – defends the view that knowing is a matter of ruling out not all alternatives, but all *relevant* alternatives. Unlike Dretske, however, Stine devotes no attention or space to

---

[59]The arguments I have in mind, of course, are those discussed in chapter II of the present work.

establishing the view that some alternatives to what we claim to know are irrelevant. Instead, she simply assumes that the relevant alternativist view is correct, and proceeds from this assumption to argue that Dretske's defense of our ordinary claims to know is wholly unnecessary.

In Stine's opinion, there is *no need* to argue that the principle of deductive closure (what Dretske calls the thesis of penetrability) is false. Because Dretske believed both that we must concede we do not know the falsity of the contrast consequences implied by what we claim to know, and that the principle of deductive closure — according to which we must know these — underlay all (or at least most) sceptical arguments, he had argued that if scepticism could be defeated, closure must be proven false. Stine, contrariwise, argues that the thesis of penetrability and relevant alternativism are wholly compatible. Stine believes we *do* know the falsity of the contrast consequences, whose falsity the sceptic says we cannot know. Hence, the question whether we need know their falsity is set aside. Hence her defense of relevant alternativism takes a much different turn.

In the course of arguing for the anti-sceptical conclusions of relevant alternativism, however, Stine makes a number of critical mistakes. Her argument for the criterion of relevance, for example, rests explicitly (and merely) on an appeal to our ordinary linguistic practices, and her claim that we know the falsity of the Cartesian alternatives is question begging at best.

What grounds does Stine give for claiming that the thesis of penetrability, the same thesis which proved so imposing and problematic to Dretske, poses no threat to relevant alternativism? To understand Stine's view, we must remind ourselves first of a number of

features of Dretske's. Dretske argued that the kinds of alternatives which the sceptic claims we must know do follow *necessarily* from the facts we claim to know. Dretske conceded, moreover, that we do *not* know those sceptical alternatives to be false. If, therefore (he thought), we are to know the things we claim to know, we must either show that we need not know the consequences which follow necessarily from what we do know, *or* we must show that we do know them to be false. Conceding, however, that we do not know those (Cartesian) consequences to be false, Dretske was compelled to argue that we *need not* know them to be false, in order to know what we do (take ourselves to) know. In other words, he argued that though Q (a Cartesian contrast consequence) is a necessary consequence of P, and S knows that Q is a necessary consequence of P, she need not know that Q in order to know that P. This, of course, is just to say that he argued that the thesis of penetrability was false, and hence that our ignorance of the falsity of the Cartesian alternatives was immaterial.

Stine, however, does not concede that we fail to know that the necessary contrast consequences of our knowledge claims are false. In fact, she argues, we do know those alternatives to be false. Moreover, because we know their falsity, she argues, the fact that they follow necessarily from what we claim to know can pose no threat to our knowledge claims. Therefore, she argues, insofar as we know these alternatives do not obtain (i.e., we know conditions are normal), we need not show the principle of closure is false. For if we know the falsity of the Cartesian alternatives whose falsity follows necessarily from what we claim to know, then there is no incompatibility between the principle of epistemic closure and our common sense knowledge claims. The incompatibility exists only so long

as we fail to know the known consequences of what we claim to know. Therefore, the veracity of the principle of closure — on her view — poses no threat to ordinary knowing.

If Stine is correct in her assertion that we *do* know that conditions are normal — that we know, contrary to what Dretske has argued, that the animals in the zoo are not merely cleverly painted mules, that we are not mere brains in vats, and so on — this truly is a revolutionary and far-reaching achievement. For insofar as scepticism about knowledge rests fundamentally on the claim that we cannot know that conditions are normal, but must, if we are to know anything at all, Stine will have (as Dretske intended) laid to rest most, if not all, sources of philosophical scepticism. Though intriguing, Stine's arguments for these philosophical conclusions are inherently disappointing. Indeed, they are not arguments at all.

Stine's "proof" that we know the Cartesian alternatives are false, for example, *consists* simply of two assumptions: 1) Evidence is not always needed for knowledge. In particular, it is not needed (as Dretske has argued) to know that the conditions under which we make our knowledge claims are normal; and 2) We know what we take for granted. Stine gives no argument for either of these claims. She simply asserts them. That is it: Stine's proof that we know the falsity of the Cartesian alternatives, and so need not know the falsity of the principle of deductive closure, in an unreduced nutshell.

When it comes to establishing that certain alternatives to what we know are irrelevant, Stine's paper is weaker still. To begin with, Stine again assumes without argument that the relevant alternativist position is correct. She offers no proof or evidence for the thesis

that some alternatives to what we claim to know are not relevant.[60] Her delineation of the relevant from the non-relevant alternatives is equally without merit. She distinguishes the relevant from the non-relevant alternatives in two ways: first, she appeals to ordinary language, and second, she rejects the only (from her standpoint) alternative relevant alternativist position of the determinants of relevance.

Stine's first approach is framed in the context of a discussion of how we might know if an alternative is relevant or not. With respect to the simple cases, Stine merely asserts the following:

> in ordinary life, we do exhibit rather strong agreement about what is relevant and what is not.[61]

Clearly, in the context of her discussion, this singular observation is considered sufficient to establish which alternatives are relevant and which not. But like every argument which draws conclusions about what is objectively the case from what happens as a matter of fact to be the case, Stine's 'argument' is subject to the familiar criticism that it conflates the standards of appropriate with the standards of true assertion. Stated differently, she, like the ordinary language philosophers and contextualists before her, appeals to the principle of epistemic conventionality, arguing (by implication) that the mere facts of linguistic habit alone are revelatory of the objective criteria of knowing. In so far, then, as

---

[60]The absence of any argument here strikes me as rather astounding given the relative newness of the RA position at the time Stine wrote this, and also given her refutation of Dretske's arguments, one of the only other papers on the subject at the time. Stine does, in a sense, appeal to both Austin and Armstrong here, but she does not argue that the position contained in the passages she quotes is correct. Rather, she merely asserts her agreement with it. See 254 and 258. In this respect, Stine's response to scepticism is like the self-described (by DeRose) "non-heroic Moorean" responses of the contextualists.

[61]Stine, "Skepticism," 252.

her only support for this criterion of relevance — our agreement in ordinary life — is her reliance on common linguistic practices, her criterion remains unsatisfactory and unconvincing.

The second way Stine gives for distinguishing the relevant from non-relevant alternatives takes as its subject the grey areas where there is less obvious agreement on which alternatives are relevant. Here she appeals to the example attributed by Goldman to Carl Ginet.[62] She argues that when there are barn facades in the immediate vicinity of Henry, we count the alternative that what Henry might be seeing is a mere facade as clearly relevant. When, by contrast, there are no barn facades in Henry's immediate vicinity, we reject the alternative that he may be viewing a mere facade as irrelevant. When, however, there are barn facades neither in the immediate vicinity of Henry nor yet too far off, she states that the issue of relevance is less clear.[63]

Stine argues that there are two possible relevant alternativist ways for deciding the question of relevance in such cases: the first, a view she defends as her own, states that an alternative is relevant only if there is *a reason to think* it is true. The second, a view she mistakenly attributes to Dretske, states that an alternative is relevant if *it could be true*. Stine, having mistakenly attributed the second view to Dretske, goes on to argue against it. She offers two reasons for rejecting the view that relevance is a function merely of what *could be* true: First, this criterion of relevance, she argues, would deprive the relevant alternativist view of its distinguishing feature, namely, the claim that some

---

[62]See Stine, "Skepticism," 252.

[63]Stine, "Skepticism," 252.

possible alternatives -- in particular, the Cartesian alternatives — are not relevant. Second, she claims (without argument) that determining relevance by mere possibility is not as easy as determining relevance by when there is a reason to think an alternative is true. In light of these two objections, Stine concludes the correct criterion of relevance is the first: an alternative is relevant only if there is a reason to think it is true.[64]

Again, however, Stine's defense of the relevant alternativist view is disappointing. For, to begin with, although Stine characterizes the second view of relevance as a possible *relevant alternativist* view (in fact, an actual one, insofar as she attributes it to Dretske), it is quite evident that the view she describes is paradigmatically sceptical, and not, therefore, one to which *any* relevant alternativist would subscribe. For the making relevant of *any* alternative so long as it is merely possible makes irrelevant only those alternatives which are logical falsehoods (or at least physical impossibilities).

Hence, because all of the Cartesian alternatives are possible (in the sense that none are logically contradictory), and we have no reason to think any are physical impossibilities, the criterion of relevance Stine attributes to Dretske would make relevant all those alternatives the relevant alternativist view is designed specifically to reject. Hence, even though Stine correctly points out that this criterion of relevance would deprive relevant alternativism of its distinguishing feature, it is something of a straw man to argue it is a possible *relevant alternativist* view. Attacking this view then, as a means of supporting her own, is quite unpersuasive. Moreover, in arguing against this paradigmatically sceptical

---

[64]One, of course, is reminded of Austin's conditions of appropriate assertion. See chapter II of this work.

view of relevance, her reason why it should be rejected – because it would allow sceptical alternatives to be relevant – is clearly question begging.

Stine's second reason for rejecting the logical possibility criterion of relevance (according to which an alternative is relevant so long as it is possible) is quite unmistakably false. For it is not, as she claims, difficult to decide if an alternative is logically possible: we simply assess whether it is logically false.[65] On the other hand, whether there is some reason to think that an alternative is true is a much more complicated matter. Moreover, even if it were more difficult to assess logical (or physical) possibility, this is hardly an *objective* basis for ruling out this – or any other – criterion of relevance. Clearly the fact that something is difficult is not proof that it is false or incorrect.

Section iii: Goldman and Perceptual Discrimination

In my discussion of Stine, I noted that it is odd that she leaves the relevant alternativist view essentially unargued for, especially odd in light of the dearth of extant arguments for the relevant alternativist position. It is possible, however, that Stine believes Goldman to have argued adequately for the relevant alternativist position, and relies therefore, implicitly, on his relevant alternativist achievements. I turn now, therefore, to Goldman's defense of relevant alternativism, as presented in "Discrimination and Perceptual Knowledge." Like Dretske and Stine, Goldman defends the view that to know that P one

---

[65]Perhaps proving physical possibility would be more difficult, but I don't see why this should be so, except insofar as we can't know the laws of nature, but this is a difficulty only if we are sceptics, and Stine clearly is not.

need rule out only those alternatives *relevant to* P. Also like Dretske and Stine, he argues that at least some alternatives are not relevant to our ordinary claims to know. In a vein similar to Dretske, Goldman argues that it is a *necessary* condition of knowing anything at all that we be able to rule out those alternatives which are relevant to what we claim to know.[66] Goldman, however, argues that this condition of knowing follows from the more general requirement that the processes or mechanisms which produce belief be reliable.

A reliable belief forming process, according to Goldman, is one which produces true beliefs rather than false ones.[67] Moreover,

> to be reliable, a cognitive mechanism must enable a person to *discriminate* or *differentiate* between incompatible states of affairs. It must operate in such a way that incompatible states of the world would generate different cognitive responses.[68]

With respect to a *perceptual* belief forming mechanism,

> [it] is reliable to the extent that contrary features of the environment (e.g., an object's being red, verses its being yellow) would produce contrary perceptual states of the organism, which would, in turn, produce suitably different beliefs about the environment.[69]

It is, therefore, an essential and necessary condition of knowing that we be able to discriminate the state of affairs we believe to obtain from other possible states of affairs that we do not believe to, but might, obtain. As Goldman points out, this view of

---

[66]I don't include Stine here because although she argues that ruling out relevant alternatives is a sufficient condition of knowing, she does not argue it is as necessary condition. In this respect, her position is much closer to a contextualist one. I will discuss this shortly.

[67]Goldman, "Discrimination," 163.

[68]Goldman, "Discrimination," 162.

[69]Goldman, "Discrimination," 162.

knowledge accords with a sense of the verb *to know* listed in the OED; a sense oft ignored by philosophers:

> "to distinguish (one thing) from (another,)" as in "I know a hawk from a handsaw" (Hamlet) and "We'll teach him to know Turtles from Jays" (Merry Wives of Windsor).[70]

At this point, Goldman's views are quite congenial to those of the sceptic.[71] For the view that a necessary condition of knowing is the capacity to distinguish the state of affairs we believe to obtain from other contrary states of affairs which might obtain is foundational to sceptical epistemologies. Goldman goes on, however, to distinguish his view from that of the sceptic, and in the course of so doing, his relevant alternativism is made clear:

> A knowledge attribution imputes to someone the discrimination of a given state of affairs from possible alternatives, *but not* necessarily *all logically possible alternatives*.[72]

In the span of a single sentence, then, Goldman commits himself to the central tenet of relevant alternativism: the view that we need know to be false *only* those alternatives which are relevant to our knowledge claim, not all those which are merely logical possible.

But what is Goldman's justification for so qualifying his original contention that knowledge requires the ability to discriminate between incompatible states of affairs? What justifies his claim that only some of the class of possible alternatives are relevant?

---

[70]Goldman, "Discrimination," 163.

[71]In fact, to some extent, his views — unlike those of any other RA — *remain* congenial to the sceptic. This he explicitly recognizes, although he states that he remains "officially neutral" on the question of scepticism. See, "Discrimination," 166.

[72]Goldman, "Discrimination," 163.

Like Stine, Goldman appeals to ordinary language. He justifies the circumscription on ruling out alternatives by appealing to what we "*normally* do" and "*ordinarily* require":

> In forming beliefs about the world we do not *normally* consider all logical possibilities.[73]

And,

> In deciding whether someone knows that P, we do not *ordinarily* require him to discriminate P from all logically possible alternatives.[74]

Clearly what Goldman here asserts is correct: we do not, in fact, ordinarily do or require either of these things. When, for example, in the company of friends, someone claims to know that P, we do not ask of her that she prove that she knows she is not dreaming, nor do we insist that she establish she knows she is not a mere brain in a vat. Equally, we do not impose such requirements on ourselves. But does it follow *straightaway* – as Goldman and Stine suggest – that these alternatives are *not* therefore *relevant* to the *objective truth* of what is claimed to be known? Does it follow that knowing the falsity of these alternatives is not — *objectively* — a condition of knowing? A condition of being correct in professing to know? Clearly it does not. The lesson of both Chapter's II and III is that the irrelevance of the sceptic's alternatives cannot follow — without further argument — from the mere facts of linguistic habit alone. It is entirely possible for an individual to *meet* the standards of ordinary behaviour *while failing to meet* the standards objectively required to make that behaviour veracious.

Let us return for a moment to my friend Rodger's party, a situation discussed in chapter

---

[73]Goldman, "Discrimination," 163-4.

[74]Goldman, "Discrimination," 164. Once again, the appeal to what we ordinarily do and require is reminiscent of Austin.

II.[75] Recall that upon arriving at the party, Rodger asked of me if John would be coming. I asserted that I *knew* he would. Rodger, of course, did not ask me to prove that I had not merely been dreaming that John was on his way, nor that I was not an envatted brain, incapable of real contact with John; to do so would have been inappropriate, requiring as it would more than we *ordinarily* require when someone makes a claim to know. Thus Rodger accepted my claim to know and went about his way.

Nonetheless, *what* I claimed to know proved false. John did not arrive at the party because, having been struck by a meteorite, he lapsed into a coma and died. Thus, though my claim to know was perfectly ordinary, meeting the standards of appropriate assertion, it was nonetheless objectively false. Obviously, though I met the conditions of appropriate assertion, I failed, nonetheless to meet the conditions of true assertion. Clearly, therefore, the conditions of appropriate (habitual) assertion and the conditions of true assertion are not one and the same, and meeting one does not amount to meeting the other.

The moral, of course, is this: one cannot infer from the mere fact that the conditions of ordinary assertion are met that the conditions of true assertion are met also. The illustration establishes the obvious truth that the conditions we ordinarily impose (what I have called the conditions of appropriate assertion) are logically separable from the conditions objectively required by the truth conditions of knowing. To infer, therefore, from the fact that we ordinarily do not require that certain alternatives be ruled out, the conclusion that we *need not* rule these alternatives out to have objective knowledge is to

---

[75]See Chapter II of this work, 16-17.

argue fallaciously. This, however, is precisely what both Stine and Goldman do.

That we do not normally consider all logically possible alternatives when we claim to know something, and that we do not normally require of others that they rule out such alternatives are the primary justifications Goldman offers for thus limiting the alternatives we must rule out so as to exclude at least some of the logically possible alternatives to what we claim to know. Goldman's defense, therefore, of his claim that some alternatives to our claims to know are irrelevant suffers from the same kinds of difficulties as those which riddled both the ordinary language approach to scepticism and the contextualist. To the extent that his argument does so suffer, it fails. It fails to establish that some alternatives are not relevant and, therefore, fails to establish that the sceptical alternatives in particular are irrelevant. [76]

Goldman has, however, an additional argument to support his relevant alternativist position, namely its explanatory power. He illustrates the explanatory power of relevant alernativism by reference to Carl Ginet's example of Henry and the issue of relevance concerning barn facades.[77] Like Stine, Goldman asserts that when there are no facades present, we acknowledge that Henry knows that what he sees is a barn. Also like Stine, he argues that if there are such facades present — unbeknownst to Henry — then Henry does not know that what he sees is a barn. But how, Goldman asks, can we explain the

---

[76]This latter, of course, is not strictly a problem for Goldman since he wishes to remain neutral on the issue of scepticism, but it is a problem for the general RA program which — apart from Goldman — is fundamentally anti-sceptical.

[77]Goldman attributes the example to Ginet, though he does not state where Ginet discusses the example. See Goldman, "Discrimination," 164.

difference in our reactions to the two alternative situations?[78] How do we explain that in the first case we say of Henry that he knows, while in the second we say he knows not?

Goldman considers a number of explanations which, he argues, will not do. Among the explanations he considers is the justified true belief analysis of knowledge. According to this analysis, of course, one knows only when she has a justified, true, belief. Such an analysis will not explain our different reactions to the alternative situations, however, because, Goldman argues, in each case Henry's belief that what he sees is a barn is both justified and true.[79] It is with respect to explaining our attribution of knowledge on the one hand, and our denial of the same on the other, that Goldman believes the relevant alternativist analysis of knowledge both succeeds and thereby gains further support. For the relevant alternativist analysis *can*, he claims, unlike other analyses of knowledge, explain both why we say of Henry that he knows (when there are no barn facades in his vicinity), and why we say of Henry that he fails to know (when there are, unbeknownst to him, barn facades in his vicinity).

The correct explanation of our different reactions, Goldman thinks, lies, therefore, in the respective irrelevance and relevance of the alternative that what Henry may be seeing is a mere facade. That is, it lies in the principles of relevant alternativism. In the first case, where there are no barn facades in the vicinity, the reason we say that Henry knows, even though it is of course *logically* possible that what he sees is a mere facade, is that the *alternative* that what he sees is a mere facade is simply *not relevant* and hence needn't be

---

[78]Goldman, "Discrimination," 164.

[79]Goldman, "Discrimination," 164.

ruled out by him in order to correctly claim to know. When, on the other hand, the environment changes such that there are facades in Henry's vicinity, the *alternative* that what he may be seeing is a mere facade *becomes relevant,* and thereby, insofar as Henry cannot rule it out, undermines his claim to know

Goldman's relevant alternativism, then, does indeed provide *an* explanation for our differing responses to the alternative instances where Henry is, in the one case, said to know, and, in the other, is not. Goldman, however, by my lights, still owes us an account of what affects the change in relevance from one case to the next. Goldman remains officially neutral on the question of whether there is a unique set of alternatives which must be ruled out for any particular claim to know.[80] But, on my count, there are three potential explanations of the *change* in relevance to be found in Goldman's paper.

First, there is the fact (agreed upon by both Goldman and Stine) that we do not *ordinarily* require such alternatives to be ruled out. Second, there is Goldman's discussion of perceptual equivalents, and the requirement that an alternative be a perceptual equivalent in order to be relevant.[81] Finally, there is the fact that the only apparent change from one case to the next is the objective presence or absence of an environmental feature.[82]

The first explanation — that changes in relevance are simply a result of our ordinary

---

[80]Goldman, "Discrimination," 166. This leads him also to remain officially neutral on the question of scepticism; a neutrality absent in all other RA accounts.

[81]Goldman, "Discrimination," 166.

[82]This, we shall see, is Dretske's criterion of relevance in his paper "The Pragmatic Dimension of Knowledge," *Philosophical Studies*, 40, (1981).

attitudes toward relevance — has already been shown to fail. In both Stine and Goldman, there is no argument beyond the fallacious inference from facts of ordinary practice to objective requirements. Hence to explain the apparent fact that the facades are not relevant in the one case but are in the other by reference to the mere fact that ordinarily we would not require their presence to be ruled out in the first case but would in the second, is not to give an explanation at all.

What about the second possibility? Explaining changes in relevance by reference to the notion of perceptual equivalents? Goldman argues that a necessary condition of an alternative's being relevant is that the alternative represent a perceptual equivalent of what we claim to know.[83] One might think, therefore, that Goldman could explain the purported difference in relevance by arguing that in the first case where no facades are present, the alternative that Henry is seeing a facade is not a perceptual equivalent, whereas in the case where facades are present the alternative is a perceptual equivalent. This explanation, however, makes no sense when we recognize that on Goldman's *final* analysis, the alternative in each case does count as a perceptual equivalent and, therefore, cannot be ruled out as irrelevant on these grounds.

What distinguishes the respective relevance and irrelevance of the facades, must, therefore, be something beyond the meeting of the condition of perceptual equivalents, and also beyond the mere facts of linguistic habit. Might the determinant of relevance then be that in the one case there is no objective feature in the environment corresponding to

---

[83]Goldman, "Discrimination," 167.

the alternative, but in the other there is? If so, how do we explain *why* that would affect a change in relevance? Working with the basic assumptions of the relevant alternativist view, and what Goldman has given us to go on, I think the only plausible explanation is that in ordinary life this is pretty much how we would decide whether or not Henry need know that what he is looking at is not a facade in order to know that what he sees is a barn. If there were actually circumstances (a practical joke, say) in which a number of facades had been erected in an area of Henry's travels, we would require him to rule out that what he was seeing was a mere facade before we would allow that he knows that what he sees is a barn. Likewise, normally, where no such facades are present, the mere logical possibility alone that there could be facades would not suffice to demand of Henry that he rule such an alternative out.

Yet, this last explanation also fails to make clear what are the objective determinants of relevance. For clearly to explain the *change* in relevance in terms of how we would *ordinarily* judge the relevance of alternatives is simply another version of the fallacious appeal to the principle of epistemic conventionality. Hence, the last potential explanation can fare no better than the first. For, if the presence or absence of environmental features affects relevance merely because this is how we *ordinarily* make judgements about relevant alternatives, we have once again an argument from mere facts of how we happen to behave to conclusions about the necessary conditions of truth. This, of course, is an argument form whose fallaciousness I have hitherto established.

In spite of the differences in Goldman's approach to relevant alternativism, I find no

greater reason for accepting the relevant alternativist picture of knowledge than I found in Stine. At this point, therefore there is no given reason (at least no good reason) to adopt the relevant alternativist view of knowledge, nor therefore to find refuge therein from the threat of scepticism.

IV: Dretske Revisited

Before closing my discussion of relevant alternativism, however, I owe Dretske one final look. "Epistemic Operators" does not represent his final word.[84] In "The Pragmatic Dimensions of Knowledge," Dretske offers another argument for relevant alternativism, an argument which makes no mention of epistemic operators, and which is quite independent of any arguments offered in that seminal paper. In this later paper, Dretske approaches the relevant alternativist position from a model of concepts which he borrows, in part, from Peter Unger.[85] Dretske argues that 'knowledge' is, as Unger argues, an 'absolute' concept.[86] An absolute term, Dretske notes, "depicts a situation as being completely devoid of a certain thing . . . . The fact that there can be nothing of this sort is what makes it an absolute concept."[87] As an absolute concept therefore, 'knowledge' cannot

---

[84]Dretske has at least one further paper relevant to the discussion here, but which for consideration of time I leave out. That is his "Conclusive Reasons," *Australasian Journal of Philosophy*, 49, May, (1971).

[85]Unger, "A Defense of Skepticism," *Philosophical Review* 80, (1971).

[86]Dretske, "Pragmatic Dimensions," 364.

[87]Dretske, "Pragmatic Dimensions," 366.

admit of degrees: one either knows or one fails to know.

Unlike Unger, however, Dretske argues that *knowledge* is not just an absolute concept, but a *relationally* absolute one; as such, it can be sensitive to the changing needs and interests of the individuals who employ it. It is, Dretske argues, from the relationally absolute character of the concept of knowledge that the irrelevance of the Cartesian alternatives follows. For insofar as knowledge has this relationally absolute characteristic, the alternatives which must be ruled out differ with changes in context. The purpose of "The Pragmatic Dimensions of Knowledge" is to explore how, and the extent to which, knowledge exhibits such contextual relativity.[88] Dretske states:

> In short, I want to explore the way, and extent to which, this absolute notion exhibits a degree of contextual relativity in its *ordinary* use.[89]

A paradigm example of an absolute term (perhaps made paradigmatic by Peter Unger) is the term 'flat.' Flat is an absolute term in the sense that if a particular object or thing *is* flat, it can have absolutely no bumps. Hence if a particular surface is flat, it follows that nothing can be flatter than it. 'Flat' is, as Unger says, an absolute limit term; defining, as it does, an absolute limit which must be met if a thing is to be flat at all. If the limit is not met *completely*, then the thing in question is not (really) flat. If a thing really is flat, then nothing can be flatter. Thus, Unger argues, when we *say* of one thing that it is flatter than another, we are really assessing the degree to which the respective surfaces approximate

---

[88]In the final section of this chapter, I will explore the relationship between Dretske's position here and that of the contextualists.

[89]Dretske, "Pragmatic Dimensions," 365. (My emphasis.)

flatness. Because flat is an absolute term, however, Unger argues it has no true application. Nothing, he says, is *really* such that it is (*absolutely*) flat. For to be absolutely flat is to have no bumps at all, while everything has at least some surface irregularities, even if only microscopically.

Knowledge, Unger argues, is also an absolute term. Like flat, therefore, it too defines an absolute limit. If I *know*, I must be *completely* justified. There must be nothing that undermines my justification. Because of the absolute nature of 'knowledge,' anything less than *complete* justification will not do. So, like flat, given the "epistemic bumpiness of our world,"[90] nothing is truly known, just as nothing is truly flat. Thus, on Unger's view, when we claim to know, we are either speaking loosely, incorrectly, or indirectly about the degree to which we *approximate* knowing. I cannot truly know, Unger argues, because there will always be competing possibilities to what I claim to know which, since I cannot rule them out, undermine my justification.[91]

Dretske agrees that 'knowledge' is an absolute term. But, unlike Unger, he does not draw from this any sceptical conclusions.[92] In fact, he argues, terms like flat and knowledge are *relationally* absolute and hence admit of application according to contextually relevant standards. Thus he argues that "Although nothing can be flat *if* it

---

[90]The phrase is Dretske's, "Pragmatic Dimensions," 367.

[91]Unger, of course, later (and regrettably, I think) gives up this view. See, for example, his "The Cone Model of Knowledge," and his *Philosophical Relativity.*

[92]Dretske, "Pragmatic Dimensions," 366.

has any bumps, what *counts* as a bump . . . depends on the surface being described."[93] Similarly, he argues: "Something is empty if it has *nothing* in it, but this does not mean that an abandoned warehouse isn't empty because it has light bulbs and molecules in it."[94] What *counts* as a 'thing,' Dretske maintains, depends upon the context and interests of those who wish to know if the warehouse is empty.

The contextual relativity, then, of these absolute terms, is in part a function of the terms we use to define them. It is also, however, a function of the interests of the speakers involved in the conversation. Whether or not the warehouse has things in it, for example, depends upon the interest we have in establishing whether or not it is empty. If we are scientists interested in creating a vacuum, then certain things will count *as things* in the warehouse which, if we were merely apartment dwellers looking for a place to store our goods, would not. Similarly, if we are expert golfers, looking for a *flat* green on which to practice our putting, certain minute bumps may disqualify a green from being or *counting* as flat which, for the amateur golfer, would go wholly unnoticed and not undermine the purported flatness of the surface.

What, then, has all of this to do with knowledge, scepticism, and relevant alternatives? Dretske wants to take this model of relationally absolute terms as a model for knowledge.

> In accordance with this approach . . . I propose to think of knowledge as an evidential state in which all *relevant* alternatives are eliminated.

---

[93]Dretske, "Pragmatic Dimensions," 366.

[94]Dretske, "Pragmatic Dimensions," 366.

Thus, he argues, just as what *counts* as a bump in defining what *counts* as a flat surface may vary with context, so also what *counts* as an 'alternative' to what we know — a *relevant* alternative — varies with context. An alternative may be relevant in one context, and hence in need of being ruled out, but not relevant in another context and hence not in need of being ruled out. Of course, it is precisely those Cartesian alternatives appealed to by the sceptic which Dretske wishes to argue are *not* relevant, do not *count* as relevant, in any *ordinary* contexts and hence cannot undermine our *ordinary* claims to know.

> This [requiring that we rule out *all* relevant alternatives] makes knowledge an absolute concept, but the restriction to *relevant* alternatives makes it, like empty and flat, applicable to this epistemically bumpy world we live in.[95]

What, then, does the relationally absolute model of concepts tell us about our *ordinary* claims to know? It does not tell us that they are justified. But it does, if Dretske is correct, tell us that *to be justified* it is unnecessary to first rule out the Cartesian kinds of alternatives appealed to by the sceptic. I need not, for example, know that I am not a brain in a vat in order to know that I have hands. I need not know that I am not the victim of an all-powerful, evil and deceiving genius. In fact, I need not know the falsity of *any* Cartesian alternatives to know what I ordinarily take myself to know. For, *ordinarily*, what *counts* as a relevant alternative are the less outrageous alternatives we are familiar with in ordinary life. And it is what *counts* as known, *ordinarily*, which determines relevance.

---

[95]Dretske, "Pragmatic Dimensions," 367.

On Dretske's view, then, the fact that I cannot know I am not a mere brain in a vat does not undermine my (ordinary) claim to know. Moreover, once again, the pattern of sceptical argumentation seems to be undermined. For the sceptic, on Dretske's view reasons thus: If Q (that I am not a brain in a vat, for example) follows from P (that I have hands, for example), and I know that Q follows from P, it follows that I cannot know that P is true unless I know that Q. (That is, I cannot know that I have hands unless I can know that I am not a handless brain in a vat.) Yet Dretske's argument appears to establish the contrary: I need not know that Q in order to know that P, even though Q may be a necessary consequence of P. For Q (and its kind) is an alternative not *normally counted as* relevant. In relation to most contexts and interests, not knowing the falsity of -Q does not *normally count* as a blow to certainty. Hence not knowing the falsity of the Cartesian alternatives cannot, normally, undermine our knowledge claims.

But what has Dretske actually established? Has he, for example, established that the Cartesian alternatives are irrelevant to the *truth* my knowledge claim? Has he, unlike Stine and Goldman, given us an *objective* basis from which we may dismiss the sceptical alternatives? Dismiss them as objectively irrelevant to ordinary claims to (objective) knowledge? Has he gone at all beyond the reasons offered by Austin, Stine, Goldman, or any of the contextualists (Lewis, DeRose, or Cohen, for example) in arguing that the alternatives raised by the sceptic are *objectively* irrelevant to the question of knowledge? I

think not.[96]

Consider his Proof. Dretske's "proof" thus far for the relational or relative nature of the correct application of the concept of knowledge has consisted merely in the fact that we often let certain things *count* as known or certain. He has argued that ordinarily we do not *count* the Cartesian type of alternatives to what we claim to know as relevant to whether we know. We do not — ordinarily — require that any such alternatives be ruled out for the purported knower to count as being completely certain. Such facts as these are, however, the only considerations Dretske offers for his conclusion that knowledge is only *relationally* absolute.

By my lights, Dretske still owes us *an argument*. He owes us an argument for the conclusion that such practices are not merely widespread, but correct. Merely pointing out what we *happen to let count* as known or certain, and concluding therefrom that knowledge does not require the Cartesian alternatives to be ruled out, is but one more fallacious appeal to the principle of epistemic conventionality: the principle that any doubts the introduction of which violates the conventions of appropriate assertion are irrelevant. As we saw with the other relevant alternativists, arguing from *the mere fact* that in ordinary conversations and contexts certain things *count* as known or as irrelevant to knowledge, to the conclusion that those things manifest the objective nature of the

---

[96]Actually, the absence of reasons in defense of his position that we need not rule out such alternatives is consistent with Dretske's defense of his externalism, according to which one need not be justified in order to have not only 'the epistemic right to believe,' but also, in order to know. See for example, Dretske's "Entitlement: Epistemic Rights Without Epistemic Duties."

concept of knowledge itself is paradigmatically fallacious. It leaves out the crucial step and hard work of *establishing* that what we let *count* as irrelevant *is* irrelevant. As I have illustrated time and again, the *mere* fact that we happen to behave verbally in a certain way, that we happen to have certain ordinary standards for appropriate assertion, does not justify any inference to the nature of the objective standards for true assertion.

In this case, it is of course true (and the sceptic happily concedes) that we *allow* many a knowledge claim to go unchallenged, to "count" as certain or known, even though the sceptical alternatives to it have not been ruled out. When, for example, my brother claims to *know* which stocks I should buy to double my money in the very near term, I do not challenge his knowledge claim. He was, after all, trained as a broker, he has studied market fundamentals, P/E ratios, and so on. Nonetheless, even though what he claims to know may *count* as knowledge, it does not follow that he *knows*. It does not follow that the *objective* conditions of knowing have been met. It does not follow that the knowledge claim is true. And as I have already argued, though there may (or may not) be a relation between meeting the standards of appropriate assertion (of making claims which we let count as known) and the meeting the standards of true assertion, the two sets of conditions are logically distinct. The one set having been fulfilled does not — straightaway — justify the inference that the other set has been filled as well.

In fact, the very examples which Dretske uses to prove the relationally absolute nature of absolute concepts themselves evidence clearly, I think, the distinction or difference between having fulfilled the conditions of appropriate assertion and having fulfilled the

conditions of true assertion. Take, for example, the admittedly absolute term *flat*. Recall that for something to be flat it must be *not at all* bumpy. When, therefore, the expert golfer is searching for a truly flat surface, though microscopic bumps in the surface may not "count" as bumps relative to her interests, they do undermine the *truth* of her claim that the surface is flat. For to be flat *is* just to be *not at all* bumpy: there is no allowance for degrees. Ordinarily, of course, we take her to be speaking somewhat loosely. We assume, for example, that he means that the surface is *very close to being flat* or *flat enough* to play well on. No one, I think, believes her to mean that the surface really is completely devoid of even minute bumps; and so Dretske's claim, that the term 'flat' is relationally absolute *because* what *counts* as a bump varies with context, is false. A bump is a bump, no matter how small. What changes with context is not whether a given thing *is* a bump, but whether or not *we care* if the bumps we encounter are big or small. If we are only concerned that there be no *perceivable* bumps we *call* a surface flat, though only loosely. To the extent that there must be no bump at all, no matter how small, we *say* of even the seemingly flattest of surfaces either that they are not flat or that they are very close to being flat.

The mere fact that we can speak differently of different surfaces in different contexts, or even speak differently of the same surface in different contexts, does not establish that our *terms* are flexible, but rather that *we* are. It establishes that we *use* our terms flexibly, not that the objective conditions of their true application are flexible. And, the fact that our flexible application of our terms is ordinarily — in ordinary conversation —

unobjectionable (or, at least, not objected to), establishes only that the conditions of *appropriate* assertion that we hold ourselves and others to are themselves somewhat flexible.

To the extent, therefore, that Dretske's proof here of the irrelevance of the Cartesian alternatives rests merely on what we let "count" as knowledge, it is inefficacious in establishing that knowledge *itself* is relationally absolute. So also, therefore, is it inefficacious in establishing that the sceptical alternatives to what we know are *irrelevant* to the *objective* truth (or falsity) of what we claim to know.

Admittedly, however, Dretske does not rest his entire case for the relevant alternativist view on the mere facts of what we happen to let count as knowledge. Instead, he goes on to offer additional reasons for viewing knowledge as relationally absolute. The reasons he offers are not, however, conclusive. Not even close. Rather, they are indirect and contingent, contingent upon the presumed correctness of factors for which he does not argue. He argues, for example, that his view, can accommodate the position of those — like the contextualists — who believe that knowledge is a function of pragmatic factors, factors such as the changing interests of the conversational participants.[97] His view also – unlike that of the contextualists – avoids "the problem of proliferating different meanings of knowledge."[98] Additionally, he thinks, his view enables us to handle "otherwise

---

97Dretske, "Pragmatic Dimensions," 368.

98Dretske, "Pragmatic Dimensions," 368.

puzzling examples."[99]

None of these reasons, however, for viewing knowledge as thus relationally absolute, are compelling. For, though it may be an advantage over other analyses of knowledge that he *can* accommodate pragmatic views of knowing, why should we want a theory that can accommodate such a view? Clearly the ability to accommodate a view like the contextualist's is not a material advantage *unless* the pragmatic view is objectively correct. Without an argument for their objective correctness, there is no reason to accommodate them, and hence no (proven) advantage in so doing. Moreover, the only arguments I have found in defense of such pragmatic views of knowing have all been based on the erroneous principle of epistemic conventionality.[100] The fact, therefore, that his model of knowledge *can* accommodate such views is not a reason to prefer it to any other.

What about Dretske's claim that his view of knowledge avoids the proliferation of different meanings of knowledge? To some extent this is true. For, unlike the contextualists, for example, who claim that the necessary or truth conditions of knowing change with context, Dretske's view is that the necessary conditions remain constant — we must , for example, *always* rule out *all* relevant alternatives. Yet, like the contextualists, this is a reason to prefer his view, I think, only if there are independently good reasons for thinking his relevant alternativism is correct. Yet the reasons he offers

---

99Dretske, "Pragmatic Dimensions," 368.

[100]In particular, this is so of the contextualist views, and the RA views which have a contextualist component.

in "The Pragmatic Dimension of Knowledge" are fundamentally the same as the reasons offered by the contextualists for preferring their own. Each rests their case ultimately on the fallacious and direct inference from our linguistic habits and the standards of appropriate assertion manifested therein to the objective conditions of knowing. I have already argued, however, that this inference fails. Hence, the mere fact that his view avoids proliferating multiple senses of 'knowledge' is not reason enough to prefer it.

Finally, what about Dretske's claim that his view enables us to handle otherwise puzzling examples? Dretske defends this claim by reference to the following: Suppose an amateur bird watcher is searching for an exclusively North American bird — the Godwall — in a North American pond. He discovers a bird with all of the markings of a Godwall floating restfully on the pond. In this case Dretske argues we would say that the amateur bird-watcher *knows* that the bird he sees is a Godwall.

Suppose however that an expert ornithologist is visiting in the area. Suppose, moreover, that he has *good reasons* to believe that the Siberian Grebe — not ordinarily found in North America — has been migrating to North America. Suppose, moreover, that the Siberian Grebe is indistinguishable from the American Godwall, except for markings on the stomach which, however, would not have been visible to the amateur bird-watcher, since the underside of the bird was hidden beneath the pond waters. Does the amateur bird-watcher *know* that what he identified as an American Godwall is not a Siberian Grebe? He does not, Dretske says, and therefore cannot know that it is an American Godwall. For its being a Siberian Grebe is — in this context — a relevant

alternative, and hence must be ruled out if the amateur bird watcher is to correctly claim to know.

Consider one final possibility: suppose that the expert ornithologist does *not* have good reasons for suspecting that the Siberian Grebes have been migrating to North America, and hence, does *not* have good reasons for thinking that the bird identified by the amateur bird-watcher may be a Siberian Grebe. Does the amateur bird-watcher still need to rule out that the bird he sees is a Siberian Grebe, in order to *know* that it is an American Godwall? It is, after all, *possible* that the bird he sees *is* a Siberian Grebe. Dretske argues that the amateur bird-watcher need not rule this possibility out because, in this instance, such an alternative is not relevant. Hence the amateur bird-watcher *knows* he sees a Godwall, even though it is *possible* that the bird he sees is a Siberian Grebe.

Dretske's explanation here seems to rest on Stine's criterion for the relevance or irrelevant of alternative possibilities. After all, one factor which changes from case to case is simply the presence or absence of good reasons for thinking the alternative might obtain. Clearly Dretske does not base his assessment of relevance on the mere possibility (as Stine had suggested he does) that the purported Godwall might be a Siberian Grebe. Yet, although Stine claimed there to be only two alternative relevant alternativist views, Dretske here in fact rejects both of these in favour of a third. For it is not, he argues the mere presence or absence of good reasons which determines the relevance of the alternative. Rather, he argues, such "good reasons" are actually irrelevant to the question of relevance. As he states:

> The Wisconsin lakes could be overloaded with migrant Siberian Grebes without the bird watcher having any reason to think that such look-alike birds actually existed. His lack of any reason to doubt, his ignorance of the possibility that what he sees is a Grebe and not a Godwall, is irrelevant.[101]

Hence, Dretske concludes that *having a reason to doubt* is *not* -- as Stine had claimed -- a necessary condition of an alternative's being relevant, though it may be sufficient. He argues that the real criterion for distinguishing relevant from irrelevant alternatives resides not "in what we happen to regard as a real possibility (whether reasonably or not), but in the kind of possibilities that actually exist in the objective situation." [102]

Dretske's relevant aternativist explanation of the above "otherwise puzzling" example is meant to serve as further support for his relevant alternativist view. But again, the fact that his view *can* explain the example is reason to prefer Dretske's view to that of others *only if* either his is the only view capable of explaining the example, or, of the alternatives available, his is the best. There are, however, other views at least equally capable of explaining the example.[103] Moreover, the fact that he can explain at least one (kind) of puzzling example is not reason enough to accept his view in the absence of other reasons. Yet, as we have seen, his other reasons (e.g., ordinary linguistic practices) are really non-productive in the absence of further, and better, arguments.

Finally, I should note that Dretske's explanation of the bird-watcher example amounts,

---

[101]Dretske, "Pragmatic Dimensions," 376.

[102]Dretske, "Pragmatic Dimensions," 337.

[103]I will defend this claim at length in chapter V.

in the end, to either a quite circular argument or a virtual capitulation to scepticism. Take the circularity first. The issue in question is what we need or need not know in order to know what we claim to know. Dretske has argued that what determines relevance is the objective features of the environment — whether there are or are not migrant Grebes in North America, for example. Yet, if the determinant of what it takes to know is the objective nature of the external environment, then either we know what the world is like and, therefore, which alternatives are relevant, or we know nothing at all. But to claim that we know what the external environment is like is to beg the very question at issue: whether we *can* know anything at all. While, to concede that we do not know what the environment is like is to concede that we can never know which alternatives we must rule out in order to know, since what determines relevance just is what is *objectively* possible in that environment. Thus, to concede that we do not know what the objective world is like is to concede we know nothing at all, or at least that if we do know anything, we cannot know that we know it.

In light of this latter objection, one might expect Dretske to fall back upon what we *ordinarily suppose* the objective environment to be like in determining which alternatives are relevant; since, after all, he defended his analysis of knowledge by appealing to what we *let count* as relevant. But, to do so, would — by his own argument — be arbitrary. For, he has already stated that what we happen to think, even what we happen to have reason to believe, is irrelevant to the question of what might objectively be the case, and therefore irrelevant to which alternatives might be relevant to our claims to know. Of

course, such a move would also suffer the further -- and familiar -- difficulty of appealing to the principle of epistemic conventionality, a principle already proven fallacious and incapable of establishing the objective conditions of knowledge.

Section v: Relevant Alternativism and Contextualism: relevant differences

In my discussion of relevant alternativism I argued that the relevant alternativists err — in a number of critical instances — in precisely the same way as their contextualist cousins. I would like now, therefore, to ask to what extent these positions not only are fundamentally the same, but also to what extent they differ. At the beginning of this chapter I quoted the noted contextualist Keith DeRose who views the relevant alternativist position as but a prominent form of contextualism. I suggested there that though the two views are quite importantly different, and, therefore, that the relevant alternativists should really be viewed as a distinct group of anti-sceptical epistemologists, they also share certain fundamental similarities, though not of the kind DeRose contends. It is time now for the details.

In DeRose's paper, "Contextualism and Knowledge Attributions,"[104] DeRose argues that there are essentially two types of factors which may affect the relevance of an alternative.

[104]DeRose, "Contextualism," 918.

DeRose calls these *subject* and *attributor* factors.[105] [106] DeRose argues that insofar as a relevant alternativist allows that the *attributor* factors affect the relevance of an alternative, he is a Contextualist. He states that nearly all relevant alternativists allow such attributor factors to affect relevance, though he gives no textual evidence of this, and goes on to conclude that most relevant alternativists are, therefore, Contextualists.[107]

In what follows, I shall argue that DeRose's categories of the different factors which may affect relevance are construed somewhat too narrowly; too narrowly to capture either what is essential to the contextualist position or to the relevant alternativist position. In fact, once the narrowness of his categories is augmented, I find that what distinguishes a relevant alternativist from a contextualist is not the extent to which he allows *attributor* factors to affect relevance (though this may be part of the story), but rather, the extent to which he allows *conversational* factors to do so. I then argue that because none of the relevant alternativists allow conversational factors alone to affect relevance — as do the contextualists — they are not properly viewed as contextualists. For it is the allowance of mere conversational factors which defines the contextualist's identity, and it precisely such factors[108] which none of the relevant alternativists allow to affect the relevance of an alternative to a claim to know.

---

[105]DeRose, "Contextualism," 918.

[106]DeRose develops this distinction also in his "Relevant Alternatives and the content of knowledge Attributions," *Philosophy and Phenomenological Research* Vol. LVI (1996).

[107]DeRose, "Contextualism," 918.

[108]Though ,of course, there are others; e.g., mere logical possibility.

Finally, I argue that though the RA's and contextualists are really, therefore, fundamentally different positions — in precisely that respect in which DeRose thinks they are the same, the manner in which they argue for the irrelevance of the Cartesian alternatives to our knowledge claims *is* fundamentally the same. Hence, the relevant alternativists and contextualists are alike, not — as DeRose contends — in positive theory, but rather, in the practical failure of their respective attempts to support their positive theoretical commitments.

DeRose defines 'subject' and 'attributor' factors as follows: Subject factors are factors that have to do with the subject of the purported knowledge; that is, the knower.[109] Attributor factors, on the other hand, are factors which have to do with the speaker or one who makes the knowledge attribution.[110] (Such factors might, of course, overlap if the purported knower and the attributor are one and the same, as when I say *of myself* that *I know* that P.) Examples of subject factors given by DeRose include there being a shortage of zebras, leading zoos to use painted mules, and there being a change in the hours of local banks. Examples of attributor factors, on the other hand, have always to do with the "speaker's linguistic and psychological context."[111] From contextualist literature, including DeRose's own, linguistic and psychological factors include what alternatives to the knowledge claim are *raised* in a conversation, the *importance* of being right about

---

[109]DeRose, "Contextualism," 918.

[110]DeRose, "Contextualism," 919.

[111]DeRose, "Contextualism," 919.

what one claims, and whether or not the alternative possibility *is considered* by the parties to the conversation.[112]

DeRose claims that it is the extent to which the relevant alternativist allows the latter — the linguistic and psychological factors — to be a determinant of relevance that the relevant alternativist is a contextualist. And, as noted above, he asserts that most relevant alternativists *are* contextualists. Yet, in my reading of the more prominent relevant alternativists I find none who allows things such as the mere mentioning of a possibility or the importance of being right to affect the *relevance* of an alternative. In Stine, for example, a *necessary* condition of a factor's being relevant is that we have some reason to think it is true. Hence the mere mentioning of a possibility does not, for Stine, make an alternative relevant. For Goldman, a *necessary* condition of a factor's being relevant is that it be a perceptual equivalent of that which is claimed to be known. Again, therefore, whether it is important to be right or whether a possibility is entertained is not sufficient to make an alternative relevant, and certainly is not necessary. In Dretske (as well as in Stine), one finds the *explicit* rejection of the view that mere mentioned possibilities can affect or determine relevance. Dretske argued, for example, that if the expert ornithologist has no basis for his hypothesis that there may be migrating Siberian Grebes in the area, then that possibility is not relevant, whether mentioned or not (unless of course the objective environment is such that there are Siberian Grebes in the area. But this has nothing to do with whether the possibility is mentioned, considered, important, or

[112]See chapter III of this work, and also DeRose, "Contextualism," 915.

discussed). In fact, then, contrary to DeRose's claim, *not one* of the prominent relevant alternativists does consider 'attributor' factors alone to be a criterion -- necessary or sufficient -- of relevance.[113]

Earlier I noted that I believe DeRose's construal of the factors which make one a contextualist is too narrow. This is evident I think because though the attributor proper is, of course, a *part* of the linguistic context, she is a part only, unless, of course, she is conversing with herself alone. Usually, as the examples of the contextualists normally evidence (including DeRose's own), the factors a contextualist considers to be relevance determining are those factors which contribute to the linguistic context *as a whole*. These factors include any statements made by the attributor, as well as any group assumptions, individual queries, interests, and so forth. They are not, then, limited to the contributions of the attributor alone. Thus, I would suggest that the relevance determinant category which best identifies the Contextualist position ought to be named "conversational factors," not attributor factors. Thus revised, I would agree that to the extent that a relevant alternativist allows such *conversational* factors to determine relevance, she a contextualist. Yet, as suggested above, each of the relevant alternativists rejects such factors as a determinant of relevance, and does so unequivocally. Why, therefore, DeRose asserts that *most* relevant alternativists are contextualists, or that the relevant alternativist

---

[113]Indeed, it is precisely this difference which leads to another difference in the respective views of the RAs and the CTs, namely the fact that the RAs do not (ordinarily) allow any room for the correctness of scepticism or the compatibility of scepticism with common sense, whereas this is another hallmark of the contextualist view. (Stine is something of an exception here, though the extent to which she is, she is a contextualist, not an RA.)

view is a prominent *Contextualist* one escapes me.

On the other hand, DeRose's categorization of the factors the non-contextualist-relevant- alternativist allows to affect relevance as 'subject factors' also seems inappropriate. For examples of such factors are invariably *not* something uniquely tied to the subject, but rather to the objective state of his environment.[114] Such factors as whether there *are* barn facades nearby, whether there *is* a shortage of zebras, whether there *are* any migrant Siberian Grebes are the kinds of factors the relevant alternativist allows to determine relevance. Thus I suggest we change DeRose's 'subject factors' to 'environmental factors.' This label, I think, captures better the range of factors allowed and disallowed by the relevant alternativist to be determinants of relevance.

We have then two categories, the 'conversational' and the 'environmental' which, I think, enable us to distinguish a relevant alternativist from a contextualist, much in the way that an expert ornithologist might distinguish a Godwall from a Siberian Grebe: if the philosopher in question claims that conversational factors can affect which alternatives to a knowledge claim must be ruled out in order for S to know that P, this is sufficient reason for identifying her as a contextualist. (It is also, as I argued above, a necessary condition. For a philosopher cannot be a contextualist unless she does consider such factors to be relevant.) If, on the other hand, she allows that only environmental factors can affect which factors must be ruled out in order for S to know that P, she is a non-contextualist

[114]Stine may be something of an exception here, but insofar as she is, she is again closer to being a CT than an RA.

relevant alternativist. If she allows both, she is not a relevant alternativist, but a contextualist.

Interestingly, DeRose in this (his 1992) paper argues also that it is contextual changes in the attributor factors (what I am calling conversational factors) which change the meaning of the term 'knowledge.' Equally interesting is his claim that mere subject factors alone are not capable of affecting such a change in meaning. This interests me because , though I do not agree that either factor can change the meaning of 'knowledge,' I do believe that another factor which distinguishes the relevant alternativist view from that of the contextualist is the RA's steadfast adherence to the view that there are *necessary* conditions of knowing anything at all, which as such do not change in accordance with changes in context. This is supportive of, or at least consistent with, my view that strictly speaking the relevant alternativists do not allow the mere conversational factors to affect the relevance of alternatives, or the meaning of 'knowledge.'[115] It is also supportive of, or again at least consistent with, my view that the contextualists do allow for such changes in meaning. Indeed as I asserted at the outset, this is one area where the RA's position is clearly preferable to the contextualist's. For, unlike the contextualists (and as noted by Dretske) the relevant alternativists do not allow for a proliferation of meanings with respect to the term 'knowledge,' and hence do not undermine the semantics of 'knowledge' to the extent to which the contextualists do.

---

[115] With respect to the question of meaning, Stine again seems to be something of a noteworthy exception here, but to the extent that she is, I would argue she is a contextualist, not an RA. This issue must be left for another time, however.

With respect to positive theoretical commitments, then, the relevant alternativists and contextualists are — contrary to DeRose's assessment — quite distinct. Though each tries to refute the sceptic by arguing that the Cartesian alternatives appealed to by the sceptic are irrelevant to knowledge, the theoretical basis on which they do so differs: For the Contextualists, the sceptical alternatives are irrelevant so long as they are not part of the conversational context; which, since *ordinarily* they are not, saves our *ordinary* knowledge claims. For the relevant alternativist, the sceptical alternatives are not relevant — whether part of the conversational context or not — unless there are objective environmental features which correspond to them; which since (it is presumed) *ordinarily* there are not, also saves our *ordinary* knowledge.[116]

Yet there are fundamental similarities too. These similarities have primarily to do with the extent to which representatives of both groups attempt to establish their respective positive theoretical commitments in the manner of Austin.[117] Each, I have argued, argues for her position by appealing either implicitly or explicitly to our ordinary linguistic practices. But, as I have also argued, this equally dooms each of their respective proofs to failure. For the inference from mere facts of linguistic behaviour to the objective conditions of knowing is fallacious. It involves an appeal to the fallacious principle of

---

[116]This may be something of an oversimplification with respect to Stine, since she requires only that there be a reason to think an alternative is true if it is to be relevant. But I do not think this oversimplification affects my main point which is simply that the RA's require more than mere mentioning and conversational attention for an alternative to be relevant.

[117]As detailed in chapter II of this work.

epistemic conventionality, a principle which supports such an inference only to the extent that it conflates two logically distinct sets of standards, the conditions of appropriate and the conditions of true assertion. Therefore, though DeRose is correct in attributing a fundamental similarity to the relevant alternativists and the contextualists, the source of that similarity lies not in positive theoretical commitments, but rather in the manner in which each tries to establish what are the objective conditions of knowing, and thereby to refute the sceptic.

CHAPTER V:
The Epistemic Relevance of Common Sense

Section i: Prefatorial remarks:

What follows is a matter not of the way in which things most assuredly are, but rather, of how things seem to me now, after reflecting on the positions of the ordinary language philosopher, the contextualist, and the relevant alternativist. I do not intend what is said be taken in the spirit of the dogmatist (nor in the tradition which follows therefrom), a tradition which asserts *positively* that things *are* thus and so, without appreciation for the potential equipollence of the arguments.[1] The following, rather, is intended in the spirit of the Pyrrhonian, who claims only that things *appear* thus and so.

> Regarding the phrase "I determine nothing" this is what we say. We hold that to determine is not simply to state a thing but to put forward something non-evident combined with assent. For in this sense it will be found that the sceptic determines nothing, not even the very proposition "I determine nothing," what he means is "I am now in such a state of mind as neither to affirm dogmatically nor deny any of the matters now in question." And this he says undogmatically by way of announcing undogmatically what appears to himself regarding the matters presented, not making any confident declaration, but just explaining his own state of mind.[2]

In this chapter, I examine the extent to which the contextualists, the ordinary

---

[1]I allude to the Pyrrhonian observation that for every argument there is a counter argument, and for every belief that P there are reasons and evidence also for not-P. See my discussion of this in chapter I.

[2]*Outlines*, 117.

language philosopher, and the relevant alternativist met their goals. Did they, for example, explain the force of scepticism? Did they explain why common sense is *compatible* with the truth of scepticism? Did they *save* common sense? And, moreover, to the extent that they did achieve these goals, did they do so without those dogmatic assumptions the sceptic argues undo even the cleverest of arguments? On the assumption that arguments and reasons still matter, I argue, they failed establish their anti-sceptical conclusions. For, not only do each of the theories we have looked at suffer from the difficulties already discussed, they also suffer from the making of many dogmatic assumptions, assumptions which by their very nature *as dogmatic* are both inherently objectionable and fallible. In fact, further exploration into the foundations of this genre of anti-scepticism (the genre embraced, that is, by the ordinary language philosopher, the contextualist, and the relevant alternativist) reveals that their multiplication of contextualized theories of knowing (together with their numerous *ad hoc* assumptions, and relativistic exceptions[3]) rests essentially on an insidious (though not uncommon) misunderstanding of both the nature of scepticism and the nature and contents of common sense. This misunderstanding rests on a number of not only uncritical and dogmatic assumptions, but assumptions that are erroneous. The erroneous nature of these dogmatic assumptions is clear once they are subjected to critical reflection.

One aspect of the ordinary language philosopher, the contextualist, and the relevant alternativist's misunderstanding of the nature and contents of common sense is the

[3]Exceptions and assumptions I have described in the foregoing chapters.

unstated, but crucial, assumption that common sense consists not just in first-order beliefs about the world, but also in a set of second-order beliefs *about those beliefs*. (By *first-order,* I mean beliefs about the world; by *second-order* I mean a kind of meta-belief about those beliefs.) This false and uncritical assumption leads to the further – and equally erroneous – conviction that common sense is threatened by scepticism. It is this latter conclusion that leads to the further (and also erroneous) conclusion that scepticism must be refuted in order to save common sense.

What follows upon these assumptions, therefore, is the proliferation of attempts to establish that the sceptic's doubts are irrelevant. Moreover, having made the assumption that the ordinary is revelatory of the true, each of these attempts rests on proving that by the ordinary standards of knowledge revealed in ordinary behavior the sceptic's standards are false in almost all (that is, excepting the extraordinary) contexts. For it is *context* which determines the standards one must meet in order to correctly claim knowledge. Hence both the sceptic's failure, as well as her mitigated success, are explained in terms of a contextualized semantics of knowing.

In this chapter I argue that because *this* anti-sceptical tradition rests on such uncritical assumptions about the nature and contents of common sense, rejecting these assumptions (as well, therefore, as the inferences made therefrom) leads to the end, not only of the proliferation of this particular *kind* of anti-sceptical refutation, but also, to the very need or foundation upon which many anti-sceptical theories are built; the perceived need to defend common sense. Moreover, I argue, though common sense and scepticism are indeed compatible (as the contextualist maintains), one can establish

this compatibility without making any of these mistaken assumptions, and furthermore, one can do so without relativizing, contextualizing, or otherwise arguing that the standards which govern knowledge are lower than the sceptic has claimed.

The assumption that common sense includes second-order knowledge beliefs is the assumption ultimately responsible for the impression that scepticism threatens the contents of common sense. The further assumption that the ordinary reveals the objective standards of knowledge (an assumption evidenced by the direct inference from ordinary linguistic behaviour to the standards of knowing) is the belief ultimately responsible for refuting the sceptic by proving her standards are not ordinary. Once, however, one rejects the assumption that common sense includes a corresponding set of meta-beliefs to the effect that each of the beliefs we hold can also be known, refutations, proofs, and the like, each prove unnecessary. For what immediately follows is the conclusion that scepticism *cannot* threaten common sense, even if scepticism is *objectively* correct in its most radical form.

In fact, I think, scepticism has not historically been directed at common sense. Rather, scepticism's object is second-order beliefs about those common sense (and other) beliefs.

> Adhering then to appearances, we live in accordance with the normal rules of life, undogmatically, seeing that we cannot remain wholly inactive, and it would seem that this regulation of life is fourfold, and that one part of it lies in the guidance of nature, another in the constraint of the passions, another in the tradition of laws and customs, another in the instruction of the arts. . . . But we make all of these statements undogmatically.[4]

---

[4]*Outlines*, 17.

The sceptic's argument is neither directed at nor threatening to *common* sense. Moreover, once the further assumption that the ordinary reveals the objective is subjected to critical reflection, I find that merely noting the sceptic's concerns are extraordinary cannot refute her.

The argument of this chapter requires, therefore, a careful re-examination of our (mostly unexamined) conception of the relation between scepticism and common sense. It requires careful reconsideration of our assumptions about both what is and what is not implied by our ordinary behavior. It requires an *attention* to the potential (and, I argue, *actual*) distinctness *in kind* of the standards which govern knowing behavior and those which govern knowledge. Finally, it requires that each of these reconsiderations be carried through, as much as possible, without the dogma of traditional assumption.[5] Truly re-examining both what common sense is and how it and scepticism stand related necessitates the removal of *dogmatic* (especially philosophical dogmatic) assumption.

The result of both becoming aware of and removing these dogmatic assumptions is an *ordinary* conception of the contents of common sense. According to this resulting conception of common sense, second-order beliefs about knowledge are *not* part of the contents of common sense. That is, though we may indeed *believe* that P, it is not the

---

[5]Of course, as all of the great epistemologists have pointed out for the last several millennia, the instinct to believe and to accept belief without evidence is everywhere. So ubiquitous, in fact, that it is impossible to be certain one is not under its influence without continuous and systematic attention. Even then, of course, we cannot transcend certain fundamental principles of thought. Nonetheless, within such limitations, one can still assess whether or not a position rests only on what are arguably dogmatic assumptions and nothing else. That is the task of this last chapter.

case that for each of our beliefs that P – even of our most intractable beliefs that P – we also ordinarily believe *that* we know that P. Rather, according to the resulting conception of common sense, though we believe many things, we do not *ordinarily* also believe – in addition to each of those beliefs – *that* we know them. That, of course, is not to say that it is part of common sense to believe that we do not know them. Rather, the point is that common sense is concerned with existence, not semantics, not epistemology. Its business is guiding action, not the conscious fulfillment of the objective conditions of knowledge or the analysis of knowledge. Neither has the least bit to do with *common* sense. And, insofar as common sense makes no such pretensions to fulfill the objective conditions of knowing, it cannot be undermined by arguing its pretensions to know are mistaken.

Supposing that common sense and ordinary belief are *not* governed by semantical or epistemological standards, it follows then that common sense is not, and *cannot be*, threatened by scepticism. Moreover, because the content of common sense does not include second-order beliefs about knowing, to the extent that our ordinary behavior is governed by common sense, that behavior is neither revelatory of the standards of knowing, nor is it governed by that kind of standard. To the extent that it is not governed by *epistemic* standards, common sense cannot be threatened by scepticism, even when scepticism's objective truth is most radically understood. For common sense consists not of second-order beliefs about knowing, but of first-order beliefs about the world.

I, therefore, urge the reader to take a suspensive, dubiative, and non-dogmatic

approach to the traditional assumptions herein discussed. I urge, that is, a kind of Pyrrhonian approach, an approach committed to rejecting assumptions for which there is not unequivocal evidence or reason. In particular, I urge the reader to reconsider the evidence for the position that common sense includes not only first but second-order beliefs. This approach reveals an important set of consequences for our understanding of the relation between epistemology, scepticism, and common sense. This set of consequences is uncovered as one leaves behind one's philosophical sophistication (until it becomes once again necessary, that is, and appropriate).

When one begins with a proclivity and commitment not to what is schooled and philosophical, but to what is ordinary and common, our non-dogmatic study of common sense leads to conclusions so contrary to those thought evident by the great *anti*-sceptics of this century (indeed, of others also) that one marvels at the divergence in conclusion, given the seeming equivalence in beginning; the ordinary, that is. The rub, of course, is that the beginnings are not really equivalent, nor is the kind of reasoning which follows therefrom. Their beginnings are fraught with unargued for and unnecessary assumptions, leaving the inferences drawn therefrom in a quite vulnerable, because ill-founded, position. It is *these latter* – their assumptions, and ill-founded inferences – which account for how the end can differ so greatly from beginnings which seem synonymous.[6]

One's instinct to judge without evidence is everywhere, and must -- along the way –

---

[6]Both I and they appeal to the ordinary, though I begin also with a commitment to resisting the dogmatic, while they begin with a proclivity to engage implicitly in it.

be intentionally resisted. To resist such instinct by speaking the language of the Pyrrhonian, however, can also not be done, at least without a great deal of unnatural attention to language. Since the language of the Pyrrhonian – relative to our current linguistic habits – is at once both unnatural and convoluted, it is distracting, and would serve, I think, only to undermine what we are after, namely, a clean sense of what is ordinary; a pre-philosophical understanding of common sense.

> If then while the dogmatizer posits the matter of his dogma as substantial truth, the sceptic enunciates his formulae so that they are virtually cancelled by themselves, he should not be said to dogmatize in his enunciation of them. And, most important of all, in his enunciation of these formulae he states what appears to himself and announces his own impression in an undogmatic way, without making any positive assertion regarding the external realities.[7]

My central aim in this chapter is to challenge the dogmatic assumptions of the clearest representatives of the kind of anti-sceptical approach I have considered throughout. My aim is not to offer an alternative theory of knowledge. Rather, it is to look critically at the theories I have discussed, and delineate the consequences of rejecting certain principles which they uncritically assume. I attempt to use only those principles which must be assumed (e.g., -(A & -A), (A v -A), etc)),by all parties, lest I be dealt the fate of Craytalus, unable to explore human discourse and reason, capable only of wiggling my little finger. To the extent, therefore, that I appear to make positive assertions that things are thus and so, let it be understood that I mean only that *it appears to be the case that* things are thus and so.

---

[7]*Outlines*, 11.

Again, the direct result of removing uncritical assumption is the discovery that the ordinary language philosopher, the relevant alternativist, and the contextualist's conception of common sense is erroneous. The assumption that knowledge beliefs can be derived directly from knowledge behavior is false, and, therefore, to defend the latter against the sceptic represents a fundamental error; an error regarding the kind of beliefs and the kind of standards at stake. When the sceptic claims that we cannot know that P, her assertion is not governed by the same kind of standard as that which governs our *saying* that we know that P. The argument from ordinary language (embraced by not only the ordinary language philosopher, but also, the contextualist, and the relevant alternativist) confuses the distinctness in kind of the standards governing what we say and the objective conditions which govern knowledge. The two are separable in kind: the one a question of pragmatics, manners, and habit; the other a question of truth.

The sceptic's criteria of knowing pose, therefore, no threat to ordinary common sense. In that sense, scepticism is, as the contextualists maintain, fully compatible with common sense. It must be emphasized, however, that scepticism's consequent compatibility with common sense and ordinary knowing is *not* (as the OLs, CTs, and RAs have supposed) owed to the sceptic's *raising* or *manipulating* 'ordinary' standards (she does not). It is, furthermore, not owed to the criteriological significance of her doubts lying outside (as they have oft been accused[8]) of the conditions of knowing (there is no reason to suppose they do). Equally, this compatibility between common sense and scepticism is not owed to the sometime or special relevance of the

[8]An argument variously made by each of the OLs, CTs, and RAs.

alternatives to P.

The irrelevance of scepticism and Cartesian doubt to the contents of common sense is owed, therefore, not to any error in the sceptic's analysis, but rather, to the separability *in kind* of the standards which govern the objective conditions of knowing from those which govern knowing-behavior, the first being epistemic and semantic, and the second not. It is owed to the fact that common sense consists not in the abstract and unordinary beliefs about those beliefs which concern philosophers, but rather, in *ordinary* first-order beliefs about the world.

In what follows, therefore, I offer a compatibilist defense of common sense which does not rest on *ad hoc* assumptions, contextualized standards of knowing, or relativized alternatives to P. In fact, my defense of common sense rests not on defeating the sceptic, nor on a particular theory of knowledge, but rather, on the recognition that the subject matter of ordinary common sense is distinct *in kind* from the subject matter of philosophical scepticism.

Section ii: The Dogmatic Foundations of Epistemic Conventionality

In chapters II, III, and IV, I argued that there are critical threads which bind the pages of these three major anti-sceptical epistemologies: the ordinary language, the relevant alternativist, and the contextualist. Salient among these is their common perception that in order to save common sense from ruin, the alternatives the sceptic claims we must rule out to know must be proven irrelevant to the accuracy of our knowledge claims. These anti-sceptics believed that unless the sceptic's alternatives

could be proven irrelevant to knowledge, our common sense beliefs would be suspect of systematic falsehood. Why would the sceptic implicate us in systematic falsehood? Because ordinarily we claim to know many things. Proving that we need not rule out the sceptical alternatives to our knowledge claims, they thought, was the *only* way to make safe our ordinary knowledge claims from the peril of scepticism.

In what follows I too argue that the Cartesian alternatives to P are irrelevant to our ordinary knowledge claims. I too argue that these alternatives need not be ruled out to save *ordinary* belief. Nor need they be ruled out to vindicate our ordinary claims to know. However, this irrelevance derives not from any error in the sceptic's reasoning. It derives not from the irrelevance of the sceptic's alternatives to actual conditions of knowledge. (In fact, I largely leave the question of the relevance the sceptic's alternatives to knowledge undetermined.) For the irrelevance of those alternatives derives from the fact that the object of scepticism, what the sceptic argues cannot be justified, is not the kind of belief evidenced by common sense (a belief the sceptic acknowledges is both necessary and natural), but instead, the kind of belief that is philosophical. In my argument for the compatibility of scepticism and common sense I find fault, then, not with scepticism, but with the *assumption* that scepticism's truth threatens the ordinary. An assumption which *creates* the need for such contextualized semantics.

Again, the method the OLs, CTs, and RAs chose to refute the sceptic and prove her irrelevance to knowledge was to appeal to the facts of ordinary language , what we might call 'ordinary speak.' This latter -- ordinary speak -- proved ,they thought, the

irrelevance of sceptical doubt. They argued that from the conditions which (appear to) govern ordinary language we can infer – without further argument – the criteria which govern the truth conditions of knowledge. Moreover, they believed that the conditions which appear to govern ordinary speak could be discerned from our ordinary behavior. Of course, were it the case that our ordinary behaviour had this kind of significance, the weaker standards evidenced there *would* insulate our knowledge claims from the sceptic's more demanding criteria. Yet, as I have already shown, the inference from the ordinary to the normative is fallacious.

What I have not, as yet, fully explained is *why* philosophers with such ability would chose *this* means, rather than another, to save common sense. I have not explained why this kind of approach to scepticism might *seem* correct. It is here, however, that the uncritical assumptions I spoke of above come into play. Consider the role played by the assumption that what we ordinarily do reveals the standards of what we ought to do. It is true, of course, (and obviously so) that *ordinarily* we neither observe nor require what the sceptic asserts is necessary for knowledge. We do not, for example, require of S that she rule out that she may now be dreaming; we do not require that she rule out that she may be vatting; nor do we require that she rule out that she may be a victim of evil and cunning genius. Such conditions are neither *ordinarily* observed, nor *ordinarily* are they expected to be fulfilled. Indeed, to the extent that we evaluate the *claim* of S to know that P, one set of considerations we *ordinarily* never consider are these.

But without first assuming that the standards we follow in ordinary contexts reveal

the standards we ought to follow in claiming to know, what can we actually infer from these practices? What follows from these facts about how we behave? Of what philosophical significance is the fact, for example, that *ordinarily* we do not rule out, nor expect others to rule out, these Cartesian alternatives before we claim to know? Does it follow that we need not – according to the objective conditions of knowing – rule out these Cartesian alternatives to P? Or, might there be an alternative and better explanation of what we ordinarily do?

The ordinary language philosopher, the contextualist, and the relevant alternativists assume that the only explanation is that the condition which govern our ordinary linguistic behaviour are the conditions of knowledge. They assume also, therefore, that we can save the ordinary by establishing that the sceptic of changes, manipulates, or otherwise distorts the true standards of knowledge. For, they also assume, the standards governing ordinary behavior are the same kind of standard as those governing knowledge. They conclude, therefore, that from our ordinary behavior the irrelevance of sceptical doubt is clearly implied. They argue, if we *do not* ordinarily rule out those alternatives to P, we *need not* rule them out either. What happens, however, when we deduct from their arguments the uncritical assumptions on which they rest?

As I argued earlier[9], without such assumptions in place, the matter of what we happen to do, is of little philosophical significance with respect to scepticism. At least, it does not have the kind of significance these philosophers have taken it to have. For

---

[9]In each of the previous chapters.

without the assumption, for example, that the ordinary reveals the objective we cannot infer from the fact that we *do not* ordinarily rule out Cartesian alternatives that we *ought not* to. We can not infer that these alternatives are not objectively *relevant* to what we claim to know. No such conclusion follows. Without these assumptions, what happens to be true of how we speak (or how we respond to the spoken word of others) does not imply the necessary conditions of the truth of what is so spoken. The relevance, of course, in the case of knowledge in particular, is that from mere facts detailing how we habitually or ordinarily use the term 'knowledge' we cannot derive the conditions of its true application.[10]

These anti-sceptical philosophers, however, do not pause to consider whether the standards we (seemingly) happen, ordinarily, to use in ordinary discourse and the standards which concern the sceptic might differ not in strength, but in kind. They do not take pause to consider the possibility that what we say in ordinary conversation might be governed by a different kind of standard altogether than the kind which governs our concept of knowledge and, therefore, the objective conditions of its correct application. They do not consider whether it might differ in kind from the kind of standard which concerns the sceptic. Each, therefore, is led to conclude that the standards the sceptic appeals to are much higher than those necessary for knowledge. They are lead also to conclude that it is only by her artificially high standards that we can fail to know.

But where is the proof? Where is the proof that the sceptic's standards are

---

[10]I established this in chapter II.

artificially high? Where is the proof that our common sense behavior represents *the* standards of knowing? What is the evidence that the standards which govern our ordinary knowledge assertions play any *epistemic* function? Beyond their uniform chorus of homage to what is ordinary, little of *rational* defense can be discerned, excepting perhaps the oft rehearsed – though inadequate – 'explanation' that we must avoid scepticism.

How important is this avoidance of scepticism? Important enough that Lewis allows assumptions to "spring into existence" when needed, "lest we be made false."[11] And Cohen requires only that we rule out only those defeaters which are *inter-subjectively* evident. This, of course, at the same time entitles us to ignore all those defeaters which are "inter-subjectively opaque."[12] Knowledge is, after all, only *relationally* absolute.[13] Stine too states directly that avoiding scepticism is in-itself a reason to embrace a contextualistic theory of knowing.[14] Lewis concedes he has no reason for contextualism other than the desire to avoid capitulation to scepticism, defining the exceptions necessary to avoid scepticism as decidedly "*ad hoc.*"[15] Indeed, DeRose goes so far as to argue that the mere avoidance of being implicated in systematic error is the central

---

[11]David Lewis, "Scorekeeping in a Language Game," *Journal of Philosophical Logic* 8, (1979).

[12]Stewart Cohen, "Knowledge and Context," *Synthese* 73, (1987).

[13]Cf Dretske's "The Pragmatic Dimension of Knowledge," *Philosophical Studies*, 40, (1981).

[14]Gail Stine, "Skepticism, Relevant Alternatives and Deductive Closure," *Philosophical Studies*, (1976).

[15]David Lewis, "Elusive Knowledge," *Australian Journal of Philosophy* 74 (1996).

reason for preferring contextualism.[16] And in a later paper, DeRose confesses that this is the motivation behind not only his, but also, Stine, Lewis, and Cohen's contextualism as well.[17] As DeRose states in his "How can we know we are not brains in vats?", " . . . the motivation – or at least *a* main motivation – for thinking we know we are not BIVs is to avoid scepticism . . . . "[18] Even Goldman and Dretske lay ultimately their conclusions before the tribunal of the ordinary (the latter, if not the former[19]) in the hope of avoiding scepticism.

But these reasons are quite incomplete for claiming that the standards of ordinary language refute the sceptic, or prove her doubts irrelevant to ordinary claims to know; incomplete, that is, unless you are naturally inclined to make the additional correlative assumptions. In light of both the role and the uncritical nature of these assumptions, I believe that not only do these anti-scepticisms fail, but that ultimately, the very need to refute the sceptic rests on an ill-gotten fear, a fear which rests on an erroneous understanding.

The drive to refute the sceptic is premised upon an erroneous understanding of common sense; an understanding of common sense bred by too much philosophy and a lack of attention to what we really, ordinarily, believe. Properly understood common

---

[16]I discuss this at some length in section iv of the present chapter.

[17]Keith DeRose, "How do we know we are not Brains in Vats?", The Southern Journal of Philosophy, 38, (2000).

[18]DeRose, "Brains in Vats," 122.

[19]Since Goldman maintains neutrality on the question of scepticism. I discussed this in Chapter IV of this work. See also Goldman's "Discrimination and Perceptual Knowledge," *The Journal of Philosophy*, 73.20, (1976).

sense is not threatened by the truth of scepticism, and the fear that scepticism might land us in systematic error is unwarranted. So too, therefore, is the need for the OL, CT, and RA 'refutation' of scepticism. Their fear that the sceptic poses a threat to ordinary belief is illusory; common sense is not in danger. That the sceptic's doubts and conclusions pose no threat to common sense does not, however, imply that the sceptic's philosophical conclusions are false or ill-derived. (In fact, though I do not have the space to argue it here, I believe they are neither.) Yet there is a marked assumption not uncommon amongst philosophers, embraced also by each of the philosophers I have addressed, that the vindication of common sense amounts to the refutation of scepticism, or, that the refutation of the sceptic is an event sufficient to vindicate common sense.

Yet this tendency to embrace the vindication of common sense as a positive refutation of the sceptic rests also on the singular, yet false, assumption that the standards which govern ordinary common sense discourse differ only in strength from the sceptic's standards of knowing. It is this singular assumption which leads not only to conflating the vindication of common sense with the refutation of the sceptic, but also, to a distorted conception of common sense now taken on as a virtual first principle, and perhaps most disappointingly, at the expense of objectivity.

Section iii: Common Sense: the separability of first and second-order belief

To establish that the fear behind the anti-sceptical philosopher's epistemic theories is ill- founded, we must then not only reacquaint ourselves with our (ordinary) concept of

knowledge, but also, with our (ordinary) conception of common sense. For so deeply entrenched is *the philosopher's* (and in particular, the anti-sceptical philosopher's) understanding of common sense, that to return us to our native beliefs, we must undo the confusion she has laden that conception with. We must retrieve that sense of what we actually, ordinarily, believe from the hands of the philosopher and what she *claims* we believe. I, therefore, again invite the reader to approach the following comments on common sense anew, without the *preconceived* (or perhaps post-philosophically conceived) notion that common sense is threatened by scepticism and, must, therefore, be defended therefrom. In particular, I invite her to re-examine the long held (and un-ordinary) *association* between the vindication common sense and the refutation of the sceptic. It is, after all, what the *ordinary* person believes which constitutes *common* sense, not the beliefs of the non-ordinary person whose brain is nourished by a diet rich in philosophy. For the association between the vindication of what is common and the refutation of what is not represents not the ordinary person's beliefs, but the philosopher's. In particular, the inclusion of second-order beliefs in common sense – in addition to the first-order – is an act of the philosopher, not a fact of the ordinary person.

That the truth of scepticism can pose no threat to our common sense follows not only from the fact that the common beliefs which comprise our 'common sense' do not include such second-order beliefs, but also from the falsity of the singular assumption that the standards which govern both ordinary (common sense) discourse and scepticism are of the same kind, differing merely in strength, but not in kind. The

resulting conclusion that these are but two alternative sets of standards for the applicability of one concept must be rejected forthwith. The assumption that they are otherwise is but one more instance of the conflation of what is logically distinct, the standards of appropriate and true assertion.

How then do I establish what our common sense behavior does or does not entail? How, in particular, do I establish that it is *not* part of common sense to believe that we know what we believe or what we say we know? I suggest we take a closer look at that behavior. And, as we consider that behavior, we ask not merely what *can* explain it, but also what minimal (non-extraneous) suppositions about that behavior we must assent to in order that it be adequately explained. Once we reflect on our ordinary behaviour, we find that none of the above assumptions is necessary to its adequate explanation.

In my efforts to explain our ordinary behaviour, and in particular, those aspects of our ordinary behaviour the OLs, CTs, and RAs believed could be adequately explained only by assuming it revealed second-order knowledge beliefs, I return to an example of my earlier chapter II. Now, recalling my earlier discussion of Beau, my (apparent) dog. What is it that I actually believe about him? What are my *common* sense beliefs about Beau? I, of course, *say* of him many things. I *say* that he is my dog, that he is a perfect specimen of the Great Pyrenees, that he is in the room with me, and alas, that he drools. These are things I most definitely *say* and believe about Beau. No sceptic, however talented, can convince me that I do not. (Equally, I think, no talented sceptic would

try.[20])

The sceptic asserts only that if I cannot *know* that I am not now dreaming, if I cannot *know* that I am not now vatting, if I cannot *know* that I am not now the victim of evil and cunning genius, then I cannot *know* these things that I believe. I do not (and cannot) *know* that Beau is a dog. I do not (and cannot) *know* that he is a Pyrenees. I do *not* (perhaps a good thing, for which I ought to be thankful!) *know* that he now drools, nor *that* his drool is real, nor *that* it is about to be shaken from his gentle face onto your newly cleaned clothes.

Such assertions from the sceptic tear at the heart of the anti-sceptic. For the anti-sceptic *hears* her beliefs being challenged. She *senses* a threat to her most fundamental beliefs. Beliefs which are dear not only to her but to all; the kind of beliefs which she thinks make up common sense. Thus does DeRose assert that the truth of scepticism lands us in systematic falsehood.[21]

The sceptic's challenge is not, however, whether what I (and you) believe is true (though, of course, one implication of what she says is that it *may* be false). Rather, she challenges most specifically the belief *that* one can *know;* know, that is the propositions that you and I believe.[22] *Know*, that is, what one ordinarily believes (a belief the unschooled in philosophy do not have). The sceptic's challenge is directed to knowledge

---

[20]Anyone who thinks she has should review, as she has clearly misunderstood her.

[21]I discuss this argument in section iv of this chapter.

[22]The point, of course, is to distinguish the first order-belief *that* P, from the second-order belief *that* I know that P.

beliefs.

Wherein, then, lies her threat? From the beliefs which I surely do have about Beau -- *that* he is a dog, of Pyrenean origin, with a penchant to drool, especially in the company of friends – it does, of course, *seem* as if I take myself to know them; and, moreover, it seems (to some) that these *kinds* of beliefs are at the very heart of the ordinary individual's common sense. I think, however, that the attribution of the *second*-order belief *that* I know what I believe is both unnecessary, and mistaken.

Suppose, for the sake of discussion, that you come for a visit. I do, after all, invite you in, introduce you to Beau, call to Beau when it is time for his treat, feed him and wipe away his drool. My behaviour seems to reveal that I believe many things, e.g., that you are real, that Beau is real, that Beau needs to be fed, and so on. But is it *necessary* to the explanation of my (or of your) behavior to suppose that I not only believe these things *but also* that I believe *that* I know them? I think it is not. The explanation that I believe these things – perhaps with great conviction – explains, *by itself*, why it is that I pet Beau and wipe away his drool. Equally, it explains, *by itself*, why I invite you in, engage you in discourse, and wave goodbye as you leave.

Supposing simply that I believe these things is quite sufficient for explaining why I act as I do. I *believe* that there is an external world. I *believe* that Beau is real. I *believe* that you have a mind. I, therefore, treat Beau, you, and the word as if you are real. But ordinarily I give no consideration to whether I have met a set of knowledge conditions which justify my saying that I know them. Why do I, at times, say that I know them? I do so to evidence the strong conviction with which I hold these beliefs, and my

unwillingness to give them up. I do not say that I know to thereby evidence that I believe I have met the objective conditions of knowledge. In fact, although it certainly seems to be the case that I do believe these things, I do not (ordinarily) believe *that* I know them. Ordinarily, I do not give any thought to whether I have met the conditions of knowing. In fact, ordinarily, I give no thought to what those conditions are, not to any distinction they might point to between mere strong conviction and knowledge.

Moreover, though you also believe these things, you do not (ordinarily) believe that you *know* them; at least, I have no reason to suppose that you do. You evidence your belief that there is a world, for example, that you are in it, that you are not now dreaming. But nowhere, ordinarily do you evidence your belief that you know them, in addition to strongly believing them. You, ordinarily, do not recite how you have defeated uncertainty, that you are conscious that knowing has a specific set of conditions, nor that you have scrutinized your evidence to make sure you meet them.

Though the respective beliefs we do have are, in fact, part of our 'common sense,' believing *that* we know them is not. Adding to the explanation of my, or anyone else's, ordinary behavior an entire additional set of beliefs does not add anything to the explanation of that behavior. There is, in fact, nothing that this supposed additional set of beliefs can explain that the original beliefs themselves cannot.

Putting one's psychological and philosophical associations aside, then, I think we can see that though it is the case that I believe these things, it is not the case (at least there is no reason to suppose it is the case) that I believe *that I know* them. It is neither necessary nor is it productive to suppose that, in addition to my first-order beliefs about Beau and his

world, I have also the *second*-order beliefs that I *know* each of these things.[23] That in addition to believing of Beau that he is a soft-eyed, gentle, Pyrrhonian drooler that I also believe – as part of *my* share of our common sense – *that* I *know* he is thus and so. In fact, I see no reason to suppose that the *second*-order beliefs – that *I know* my common sense beliefs – are really part of the content of my share of common-sense. In the absence of any reason for supposing I do hold such beliefs, the supposition that I do is an artificial addition to common sense, added on by a certain temperament, inattention, and confusion of philosophers.

Suppose that upon visiting, you and I (and Beau) proceed to the parlour. I sit across from you, and Beau, protective and cautious, near me. Suppose, moreover, that we begin to discuss Descartes, Sextus Empiricus, or Hume. We laugh. We exclaim how difficult it is to teach students (first year, especially) to take seriously the Cartesian challenge, especially when the sceptic's victory for the absurd seems to win hands down. I confess that even I, at times, have trouble with the seeming absurdity of her alternatives: I may now be dreaming, a mere envatted brain, or a victim of evil and cunning genius.[24] I concede the utter impossibility of suspending my ordinary beliefs about the world. Indeed, I confess, that in one of my classes I burst into laughter (Perhaps the long winter. Perhaps my common sense beliefs conflicting with the philosophical.) whilst trying to persuade my

---

[23]By first-order beliefs I mean the beliefs I hold about the world and the things in it. By second-order beliefs I mean the beliefs I may or may not hold about my first order beliefs.

[24]Note: our trouble here is correct since we do not believe we know this things, and the doubts apply only to knowing. It is therefore inappropriate to worry about removing these doubts from mere belief (of the first-order, that is).

students of the importance and centrality of the Cartesian *Meditations* on doubt; of the possibility that Descartes – philosophically – was right (before his immediate fall back into dogmatism); of the possibility that we do not *really* know anything at all.

We, (you and I, that is. Beau, though highly intelligent does not take part in the conversation) go on to discuss what reasons Descartes' had entertained which might have led him to conclude that such alternatives (to what we believe) are indeed relevant to knowing. Why such troubling doubts ought, to anyone, to be troubling at all. At once, however, my thoughts are interrupted: Perhaps I really do not know I am not now dreaming. I recall a dream I once had in which you (quite unexpectedly) came to my door and I, pleased to see you, invited you in. I remember that in my dream I introduced you to Beau, we discussed Descartes, we sat in the parlour, you across from me, and Beau -- cautious and protective -- near me. I recall also that in my dream I told you the story of my laughter in class, and of my students astonishment when my laughter continued. Of course, I believe that I am not now dreaming, but I also, after considering my dream, realize that I have no criterion within my present experience which would enable me to *know* that I am not now dreaming. I believe that I am not, but I have never conducted a study into making sure. How can I *know* that you are not now part of a dream in which I tell you of a dream I once had? This is something I most assuredly believe, but I realize I have never subjected that belief to anything like philosophical scrutiny. Initially, I search fiendishly for some sign, and you with me, to *prove* I am now awake, that you really are here with me, and that I am here with you.

It occurs to me, however, that whether Descartes' reasons to doubt are good or not,

whether his alternatives are relevant to the conditions of knowledge, is irrelevant to the fact of my belief. I realize the sceptic threatens me only to the extent that I believe *that* I know these things. Yet I do not believe that I know them. The fact that I have never subjected these and the like beliefs to any set of standards implies clearly that I do not believe that I know these things, no matter how tenaciously I cling to them. Hence, I realize, the sceptic can pose no threat to my common sense. Again I relax.

You assert that surely we both at least *believe* that the dog is real, that the drool he has now shaken onto your pants is also real, and that the conversation we've just had was not a mere dream. There is simply no mistake about it. Indeed you insist that these beliefs so are so strong and so fundamental that you cannot really bring yourself to doubt them. You cannot engage in real doubt about them, the kind that leaves you in a state of uncertainty, the kind that convinces you not to belief them until you have some further evidence.

But from the mere strength of our convictions need it follow that we believe that we know? Or, mightn't it perhaps be that these beliefs are so fundamental, so non-negotiable, to our common sense that we *feel* they are challenged by the sceptic when really they are not? Mightn't it be that we confuse the strength of these convictions for the conviction *that* we know them? After all, where in the context of common sense and ordinary discourse do the objective conditions of knowing come up? They do not. The objective conditions of knowledge are neither a part of the content of common sense nor of ordinary discourse. Thus it is quite evident that it is not necessary to suppose that we have such second-order beliefs; not necessary, that is, to the explanation of what we say or to the explanation of what we do. It is not necessary to suppose we believe that we know what

we believe because we do not. We do not hold those beliefs which are the object of sceptical doubt. At least, there seems to be no good reason for supposing that we do.

To the extent that you and I discuss knowledge at all, its objective conditions, the relevance of certain alternatives to P, our discussion is not *ordinary*, nor is it about what any of us would characterize as *common* sense. Rather, the discussion is *philosophical*, everything what is common is not. Its concern is epistemic and justificatory; again, what the common is not.

Moreover we can now see how the association between the vindication of common sense and the refutation of scepticism might get started. It rests on a misunderstanding of these relationships. It rests, in particular, on confusing the beliefs we do have for beliefs they are not. It rests on confusing the strength and degree of conviction with which we hold our beliefs for the strength of evidence required for knowledge. It rests on confusing what we have a right to believe and assert, with the fulfillment of those conditions necessary for true knowledge. It rests, that is, on confusing the fact of our fulfilling the conditions of appropriate with those of true assertion.

One might object that our linguistic habits reveal that we do have such knowledge beliefs. We do *say*, after all, that we know many things. Perhaps, however, our critic infers with too much haste. Perhaps she is misled into believing that we believe we know what we believes by too much philosophy. Perhaps she is again misled by the strength of our (or her own) convictions, the strength of her ordinary beliefs. After all, it is not on the basis of excluding uncertainty that I greet you, invite you in, or introduce you to Beau. Equally, it is not on the basis of your exclusion of uncertainty that you get thru the day.

Yet, the critic does seem right about something: our claims to know must have *some* explanation. Surely our behaviour which evidences our belief in an external world stands also in need of explanation.

The question, however, is one of the right explanation. The right explanation, I think, is that our conventions allow us to claim that we know, even when in fact we may not. After all, one violates no convention when she does not normally attempt to prove she is not now dreaming before brushing his teeth, dressing for the day, or walking his dog, any more than I violate linguistic (or other) convention before I stroke Beau's white fur, give him his daily treat, or wipe away his drool. Equally, I violate no convention when I greet you, invite you in and introduce you to Beau. But what kind of convention governs the acceptability of our behaviour? What kind of standards are the conventions which must be met for our behaviour to count as acceptable.

I think the correct answer is this: the kind of conventions governing ordinary behaviour are pragmatic, not semantic, not epistemological. As I argued in chapter II, it is because what governs ordinary discourse is standards of appropriate (not true) assertion that we can say with impunity that we know, when in fact we may not know (witness my false assertion that I know John will be at the party. A perfectly appropriate assertion, though false). When I say I know that P, in ordinary conversation, it is not on the basis of conscientiously applying a set of criteria I consider the objective criteria of knowledge. Indeed, most ordinary people, when asked, concede they do not know the meaning of this phrase 'objective criteria of knowledge,' nor what I have in mind when I ask them if they think they can meet them. Importantly, in fact, if I ask of the individual on the street if she

is certain that P, it is not at all difficult to get her to concede that she probably does not *know* that P, even though she is pretty darn sure that P. If I ask, "Are you absolutely certain that you know that P?" she concedes she is not. This concession, however, is not part of her ordinary or common sense beliefs. It is arrived at upon reflexion. What she does insist upon, however, is that she will continue to *believe* that the world is thus and so no matter what I say. She will continue to believe that my dog is real, that she has hands, and that she is not dreaming. These beliefs, she insists, *are* part of common sense. I agree.

Believing that these things are so, however, is quite different than believing, upon reflection and with intention, that one knows them absolutely, or that one has -- in her *common* sense -- *any* conception of the objective criteria of knowledge (whatever they may be) which must be applied before she can *correctly* (as opposed to merely appropriately) claim to know. It seems much more likely that what one has is a sense of what the *appropriate* use of the term is, and it is *that* – that sense of appropriate assertion — which explains her behavior, for it is that which gives her her linguistic direction.

What then is behind the ordinary claim to know, and the behavior ordinarily taken to imply that one believes oneself to know? I propose first of all, that it is one's sense of what it is appropriate to say, and secondly, that it is the intention of establishing the differing degrees with which one is subjectively certain. I say that I know when I am pretty darn sure, or when nothing could conceivably bring me to relinquish my belief.[25] This

[25]In fact, on Dretske's account of things, it is just this kind of irrelinquishability that gives me the 'epistemic right' to claim I know. See his 2000, "Ought we to follow the evidence?".

behavior (as well as the standards that govern it) is learned as we learn the way the term knowledge is appropriately used in our culture. If, for example, you ask me if I *know* that the sun warms the earth. Without any reflection I would, ordinarily, most assuredly say yes. If, by contrast, you ask if I know the building I see off in the distance is round, I would, ordinarily, say I do not. My intention, however, is not to indicate that in the former case I have mounds of evidence which I have scrutinized according to that set of conditions I believe constitute the objective conditions of knowing, while in the latter case I have not. It does not indicate that I have scrutinized the concept of knowledge, analyzed its objective conditions, and carefully adjusted my beliefs accordingly. Not, anyway, if I am an *ordinary* individual. Rather, I intend, I think, to convey to you that in the former case I have a kind of belief about which no one can convince me I am wrong. The kind of belief that I absolutely refuse to give up. In the latter case, by contrast, I indicate my willingness to give up the belief in light of other considerations.

To the extent that I'm not *that* sure (in an appropriate sense of 'sure') I will most likely, therefore, say I "think so," as opposed to, "absolutely," "without a doubt," and so on. I intend by my words to indicate to you that I am more subjectively certain of one than of the other. The fact that I *say* I know the one and not the other is but my loose way of letting you know that the one belief I will not relinquish whereas the other I (perhaps) will, that the one belief is more certain than the other.

Why then are so many claims to know made (or implied) by so many ordinary people? For precisely those reasons I have for claiming to know that the sun warms the earth, that the moon is above my head, and the earth under my feet. One indicates thereby one's

degree of conviction, one's unwillingness to go without, one's instinct. One does not, as it has been widely assumed, indicate her epistemology, her semantics, her solution to philosophical scepticism. The ordinary person habitually claims to know, but not because she habitually believes that she knows. I accept her claim, as an ordinary person, not because it is absolutely true that she knows. Rather, she claims to know, and I accept her claim, because it is acceptable, given the conventions governing our linguistic culture. In her case and in mine, the standards to which we appeal are the standards of *appropriate*, not true, assertion.

Section iv: Objections and replies

***First objection*:** One might, given too cursory a reading perhaps of this (and the previous) chapters, conclude that ultimately there is no real reason to prefer my means of reconciling common sense and scepticism to that of the contextualist, and perhaps even no reason to prefer it to the position of the relevant alternativist. For, one might think, I have offered but one more form of contextualism or relevant alternativism. One might conclude that my compatibilism saves knowledge only by claiming that the objective standards of knowing are different at different levels of discourse. On this reading of my argument our ordinary knowledge claims are true, and – in all but the most extraordinary contexts – the sceptic's claim that I fail to know is false.

Such an objection, however, misses what is at the heart of my defense of common sense; namely the separation in kind of two very different kind of standards. As I have shown, once we distinguish the appropriate from the epistemic, we find that common sense is

made up not of complex second order beliefs about knowing, but rather, is made up of simple ordinary beliefs about the world: I have a dog, his name is Beau, he is a drooler, and the like. And these are not judgements about the truth of my claims; they are beliefs about the world. On my understanding of common sense the question of whether my knowledge claims are true represents a confusion about the kinds of standards that my knowledge claim purports to meet with the kinds of claims it does not. Moreover, since on my understanding the claims of ordinary speak are governed by the conditions of *appropriate* assertion, they are really claims that it is appropriate to say that P; not that it is true – according to the conditions of knowledge – to say that I know that P.

What follows, as I have argued, is that our ordinary knowledge claims are not objectionable from the vantage point of one kind of standard; that which governs what is appropriate. It does not follow, however, that these claims are unobjectionable from the vantage point of another kind of standard; that which governs what is true. The importance of the separation of the second-order from the first (as well as the correlative separation in kind of the standards governing each), and the consequent separation of the subject matter of epistemology from the contents of common sense cannot be overstated. With these distinctions, the very fear which serves as the foundation of ordinary language, relevant alternative, and contextualist anti-scepticism proves irrational. Scepticism can pose no threat where there is no claim to fulfill the conditions of knowledge

So too, therefore, does the multiplication of contextualized theories of knowing come to an end. For it was only because common sense was thought to be threatened by scepticism that these theories of knowing were developed. Each was designed specifically

to show how our common sense knowledge claims could be compatible with the truth of scepticism. Each of the proponents of these contextualized theories of knowledge, moreover, because of the dogmatic and false assumptions she made, could establish this compatibility and save knowledge only by contextualizing the victory of the sceptic, arguing that her standards are never applicable to ordinary knowing.

What I have suggested here, however, is that it is possible to save common sense without contextualizing the sceptic's victory. Moreover, one can argue for the compatibility of common sense and scepticism without contextualizing the conditions of knowledge, and without accusing the sceptic of distorting those same conditions. One can achieve these results by establishing that the subject of common sense is distinct in kind from that of the sceptic's, by establishing that the kind of standards which govern ordinary assertion and discourse are different in kind than those which govern knowledge, and by establishing that common sense is not concerned with epistemology or meta-theory about belief, but rather, direct belief and action.

Moreover, by finding that second-order knowledge beliefs are not a part of the content of common sense, one can avoid a number of serious problems faced by the other theories I have discussed. One avoids, for example, the proliferating meanings for the singular term knowledge, solely for the reason of avoiding scepticism. Furthermore, one can respect the time-honored distinction between mere belief and knowledge. And, perhaps most importantly, arguing for the compatibility of scepticism and common sense on these grounds allows us to acknowledge the possible truth of even radical scepticism without thereby threatening common sense, and it does so without arbitrary or contextualized or

relativized standards of knowledge.

Second Objection:

There is, however, a further objection that I must consider. Keith DeRose constructs this objection in "Solving the Sceptical Problem."[26] Because his objection is directed at a position which appears very much like my own, I must not pass it by. Both the position DeRose considers and the position I defend agree that if our ordinary claims to know are interpreted as having an epistemic intent, they are no doubt false. (Importantly, the position he constructs and then refutes does not consider whether our knowledge-claims might be devoid of such intent. In this, the position described is significantly different from my own.) Additionally, the position which DeRose considers is an invariantist position. (Where, by *invariantism* I mean a position according to which the conditions of knowing are *unchanging, illiquid, absolute*.) Both the position I argue for and the invariantist position which DeRose addresses represent *alternative* descriptions of the nature of ordinary discourse. Each, in certain ways, parallels DeRose's contextualist description, but from an invariantist, not a contextualist, perspective.[27]

---

[26]*The Philosophical Review*, (1995) 104:1-52.

[27]I, of course, have not argued for such an invariantist position. What I have argued, however, leaves open the possibility that an invariantist position is correct.

Suppose, DeRose suggests, that the invariantist argues the following: when one claims to know in the context of ordinary discourse, she confuses the standards of *warranted* assertability with those of knowledge. In fact, the invariantist goes on, though she meets the conditions of warranted assertability, she does not meet the conditions of knowledge. It follows that when one claims to know, her claim is false.[28] Moreover, because any claim of any S to know that (any) P is false, all of S's claims to know are *systematically* false. This invariantist theory, of course, is precisely contrary to the theory DeRose would like us to embrace, a theory which is contextualistic and relativistic; a theory according to which we do meet the standards of knowing because those standards take into account our contextually determined interests.

Of particular interest in DeRose's response to invariantism is his contention that this singular fact – that invariantism implicates us in systematic falsehood – is *the* reason to prefer his own contextualist theory to the absolutist theory of the invariantist. One might point out scepticism's implicating us in systematic falsehood is hardly an objection to it. Yet in response to a theory, or genre of theories, whose self-proclaimed *raison d'etre* is the avoidance of such a threat, no such response will do. A much better response is to consider the form DeRose gives the invariantist position and ask if belief's performing an epistemic function is a necessary corollary of that position. We must ask, that is: Does invariantism necessarily implicate us in systematic falsehood? By my lights, it does not.

---

[28](Now this, of course, sounds very much like my claim that certain individuals confuse the conditions of assertability with those of truth. As I will argue shortly, however, because I argue it is the philosopher who makes this mistake, not the ordinary person, my and this variety of invariantism have radically different consequences.)

At least, it need not.

DeRose, of course, is correct that the invariantist position *he describes* does implicate us in systematic falsehood. But the invariantist theory he describes implicates us in systematic falsehood *only because* it too takes for granted that common sense is concerned with the same kind of standard as the epistemology. (The confusion of the two kinds of standards is evidenced in the claim that we meet the standards of *warranted* assertibility. It is also evidenced in the claim that our common sense knowledge claims are *false*.) This, however, is a view which is mistaken. Moreover, the view that our common sense knowledge claims should be judged according to the standards of knowledge is certainly not a necessary accompaniment to an invariantist position. The invariantist need only maintain that the objective standards of knowledge are unchanging, irrespective of context. He is not committed to *any* particular view about the nature or contents of common sense.

In fact the correct position, I suspect, is a combination of invariantism and the explicit rejection of the supposition that common sense includes second-order judgements about belief as a piece of unsupported dogma.[29] Once the dogmatic assumption that ordinary common sense includes second-order epistemological belief is rejected, what becomes of DeRose's objection to invariantism? What becomes of his *(apparently singular)* reason for preferring a contextualist semantics to an invariantist one?

Immediately, both DeRose's objection and his singular reason for embracing a contextualist semantics over an invariantist semantics loses all force. For it is *only given*

---

[29]Indeed I suspect (though I have not the space to defend my view here) that invariantism is the *only* theory which makes any sense of either the power of scepticism or the concept of knowledge.

the supposition that our ordinary knowledge claims are epistemic in nature that scepticism's truth can imply their falsehood. It is only if second-order beliefs are included in the contents of common sense that invariantism can (invariably) implicate our common sense beliefs in systematic falsehood. Once we remove these uncritical assumptions, the *raison d'etre* of such defenses of common sense against philosophical scepticism evaporates. And because, as I mentioned, there is no reason an invariantist cannot reject this erroneous view of common sense whilst still maintaining her invariantism, equally, therefore, there is no reason (by DeRose's own account) to prefer contextualism to invariantism. Invariantism ***need not*** implicate us in sytematic falsehood, then, any more than need scepticism, and, of course, niether does once we understand the erroneousness of that conception of common sense which places second-order epistemological beliefs in its contents.

Once we recognize that second-order beliefs are *not* part of the contents of common sense, then, we immediately see also that no systematic falsehood *can* result from the truth of even the most radical invariantistic scepticism. Hence, DeRose's objection to (at least the variety I suggest) of invariantism is quite without power. Invariantism need not implicate us in systematic falsehood, and does not, once we release it from the dogma of traditional assumptions.

Section v: Concluding Remarks

Let us suppose, however, for the sake of argument, that none of my conclusions thus far are correct. Second-order beliefs really *are* an ordinary component of common sense, and,

therefore, if the sceptic's conditions of knowledge are unobjectionable, we fail to know, systematically. Let us suppose, moreover, that we can -- as the ordinary language philosopher, the relevant alternativist, and the contextualist suppose -- discriminate the relevant from the non-relevant alternatives by appealing to ordinary linguistic practices. Let us suppose also, therefore, that the conditions of knowledge the sceptic proposes are not relevant, and that the real conditions of knoweledge change in accordance with changes in the contextually sensitive determinants of relevance. Finally, let us suppose, that because the conditions of knowing are thus contextually sensitive, we *can* correctly claim to know, at least most of the time.

I then, in effect (and for the sake of argument), concede that the ordinary language philosopher, the contextualist, and the relevant alternativist are correct: correct in supposing that in certain contexts the sceptic's criteria apply, but in most (because ordinary) contexts they do not. I concede also, therefore, they are correct in supposing that both we and the sceptic can be right (i.e., that common sense and radical scepticism are compatible). I concede also, therefore, they are correct in supposing that we know those things we ordinarily take for granted: that we have hands, that there is a finch in the garden, that there is an external world. For, if the ordinary conditions of knowledge (as opposed to the ordinary conditions of appropriate assertion) are determined by ordinary context, it follows that we do not need to rule out any of the Cartesian alternatives to our knowledge claims in order to correctly claim to know. We need not know we are not merely dreaming. We need not know we are not brains-in-vats. We need not know that we are not victims of evil and cunning genius. Why need we not know the falsity of these

alternatives to correctly claim to know? Because, ordinarily, no one rules them out. Ordinarily, no one ever thinks of such alternatives. Ordinarily such alternatives as these are, as Cohen says, "intersubjectively opaque."

The following assertions are, then, not only unobjectionable from an epistemological standpoint, they are true: "I know I am not merely dreaming, even though for all I know I might be dreaming." "I know I have hands, even though for all I know I may be a handless brain in a vat." "I know that two plus two is four, even though for all I know two plus two is not four, but merely seems to be four because I am systematically deceived by an evil and cunning genius." In fact, not only are these assertions true, but for the contextualist at least, so are *their negations.*[30] Equally, the follwing set of assertions are also both unobjectionable and true: "I know my dog is real, even though for all I know he is a fake." "I know there are other minds, even though for all I know there are not." "I know there is an external world, even though for all I know there is not."

Under these conditions, *what* then does it *mean* to *know*? What does our *knowledge* actually amount to? What criteria, if any, must we meet? How does knowledge differ from mere (even true) belief? First, because on each account, there is no one singular set of conditions which must be met for S to correctly claim to know that P, knowing does not require of you the same things it may require of me. Moreover, since the conditions I must meet to know may differ from the conditions you must meet, I can know that P while you

---

[30]See Cohen's "Knowledge," and his "Contextualism, Skepticism, and the Structure of Reasons," Philosophical Perspectives 13, (1999). See also DeRose, "Solving the Sceptical Probelm" Both DeRose and Cohen explicitly commit to the possiblity that both the ordinary claim to know and the (extraordinary) claim of the sceptic that I fail to know are true.

know that-P. Hence, truth, it seems, cannot be a condition of knowledge. Furthermore, because we need only rule out those alternatives to P which the context brings relevance, the justification of P becomes such a subjective matter that it is hardly a requirement at all. This leads to the further problem that mere belief seems to become indistinguishable from knoweldge.

How satisfactory, in light of its meaning and consequences, is this victory. How satisfactory, under *these* conditions, is it to *know* what we ordinarily believe? By my lights, this kind of victory -- in the face of scepticism -- is empty. This kind of victory yields no satisfaction. Why is it empty? Why does it yield no satisfaction? Wherefrom comes the honest philosopher's feeling of deep and unavoidable dissatisfaction with her true knowledge? It comes from first our intuitive understanding that knowing is different from mere belief. It comes from the fact that knoweldge is not of the false, but of the true. Moreover, it comes from the fact that the success our knowledge claims, on this theory, is contingent upon our willingness to make knowledge a function of inter-subjectivity: the intersubjectively evident, the inter-subjectively ignored, the inter-subjectively presupposed. In fact, insofar as saving knowledge from scepticism was one of the unifying goals of not only the contextualist, but the relevant alternativist and the ordinary language philosopher as well, the anti-sceptic's I have considered seem to fail. For surely this is not the 'knowledge' we want rescued from the threat of scepticism. In fact it does not look like knowledge at all.

In the last analysis, we are not obligated to fulfill *any* conditions *not* ordinarily imposed by the average speaker in an average conversation. As Lewis states in "Elusive

knowledge," "we are entitled to ignore whatever alternatives to P those around us ignore."[31] This same conviction, as we discovered in earlier chapters, is found also in Austin, DeRose, Cohen, Stine, and others. And, since ordinarily in ordinary conversation we ignore the Cartesian alternatives to our knowledge claims, it follows that we need not rule them out to correctly claim to know. Should the conversation change direction, away from the ordinary and toward the extra-ordinary *(ala* the sceptic, for example) then, either the pressupositions necessary to make our knowledge claims true "spring into existence,"[32] or the higher level of standards brought into play—no matter how compelling -- because they are not ordinary, need not be met.[33] All of this thanks to a contextualistic semantics.

Among the consequences, therefore, of embracing this genre of anti-scepticism, are the following: truth, certainty, and justification are not *(necessarily) necessary* conditions of knowledge.[34] Mere belief, even true belief, cannot be distinguished from knowledge. We know, *but only to the extent that* we ignore presuppose, assume, make *ad hoc* exceptions, or beg the question. That, after all, is how people *ordinarily* behave. As Lewis concedes in "Elusive Knowledge,"

> I have to grant, in general, that Knowledge just by presupposing and

---

[31]Lewis, "Elusive Knowledge," 559.

[32]An argument of Lewis from his earlier paper "Scorekeeping and the language Game," 2.

[33]I allude to DeRose's "Solving The Sceptical Problem," 36.

[34]Of course none of the philosopher's we've discussed states this explicitly. To do so would bring instant scrutiny, something a position without supporting argument cannot afford. Each, however, with the possible exception of Austin, does imply that P's being true is not a necessary condition of knowing that P.

ignoring *is* knowledge.[35]

This is a theory which, by Lewis' admission, avoids scepticism only by making *ad hoc* exceptions,[36] a theory *the only 'evidence'* for which (by the admission of Lewis, Stine, DeRose, and others)[37] is its capacity to both embrace what is ordinary and ignore what is not. A theory, that is, which is victorious over scepticism, but *only* to the extent that it throws out the heart of the issue, namely, justification, a theory in which our 'knowledge' is 'elusive'; vanishing (as Lewis remarks) the moment we *pay attention* to a higher standard of knowing. Surely this is a theory whose victory is cheap indeed!

To leave out of epistemology what is arguably its most fundamental concept, justification, is to leave out epistemology altogether. By my estimation, therefore, what each of the philosophers I have considered is doing is not epistemology, but something much different, psychology I suppose.[38] And, of course, changing the subject matter is not a means of defeating the sceptic, but only of avoiding her. Mere avoidance, however, is not victory.

Equally, as I have argued throughout, the ordinary language philosopher, the contextualist, and the relevant alternativist have given no *philosophical* grounds for rejecting any of these conditions on knowledge. Inasmuch, moreover, as the OLs, CTs,

---

[35]Lewis, "Elusive," 562.

[36]As I quoted in chapter III, Lewis concedes the *ad hoc* nature of a central assumption he makes to avoid scepticism. See my discussion of this in chapter III of this work, 23.

[37]See my discussion of these in chapters III and IV of this work.

[38]In this last point, I am very much in agreement with, as well as indebted to, the opinions of Jaegwon Kim, as argued for in his "What Is 'Naturalized Epistemology?" (1988)

and RAs are responding directly to the threat of the philosophical scepticism, they clearly have not met (nor, in fact, showed much interest in meeting) their burden of proof.[39] Rather than telling us what about knowledge makes possible its co-existence with question-begging, assuming, and ignoring, they have simply appealed to convention. Moreover, in so appealing to convention they have overlooked the distinction between two fundamentally different kinds of standards: the pragmatic and the justificatory. They have confused first for second-order belief, and only on the basis of the confusions, they have laid the objective conditions of knowing before the subjective conditions of group behavior.

What is *wrong* with their kind of knowledge? What makes the proclamation of it and of the consequent victory over philosophical scepticism so vacuous? As I have indicated, winning an argument still requires the giving of reasons. It requires also that we not beg the very question at issue. It requires that we have a better reason than mere dislike for rejecting our opponent's position. Moreover, any winning argument ought not to violate fundamental laws of thought (the law of non-contradiction, for example, not to mention excluded middle.)[40] Equally, it ought to *defend* (not merely assume)the foundations from which it derives.

---

[39]To my near astonishment, DeRose seems to consider this a virtue,or at least, nothing to be ashamed of. He states, for example, in his "BIVs," "I don't make this knowledge that I claim we have dependent on any argument I give -- or any argument whatsoever." And, "As a Moorean, I face the pointed question: *How* do we know that we are not BIV's? I join my contextualist Moorean cohorts [Stine , Cohen, and Lewis] in answering negatively: *not* by means of any argument, nor even by means of our having any effective evidence to that effect." 135.

[40]I allude to the consequence of this kind of contextualistic semantics that the same proposition can at the same time be both true and false.

In addition to these and the other philosophical criticisms I've made throughout this work, this philosopher still believes that what determines both the conditions of knowing, as well as the conditions of their fulfillment, is not the average IQ or linguistic habits of one's peers. Such a means of determining knowledge's truth conditions is no more rational than determining the conditions of sound argument, the morally good, or the beautiful by popular opinion polls. Furthermore, this philosopher still believes (especially in light of no argument to the contrary) that a necessary condition of knowledge is truth, that a corrollary of correctly claiming to know that P is that one's belief that P is justified, and that there is a legitimate distinction between mere belief and knowledge. This philosopher is *not* satisfied with *knowing* that her dog is real, nor that his Pyrronean origins are true, nor that he really is a drooler, without first knowing she is *not* merely an envatted brain, not the victim of evil and cunning genius, nor merely dreaming. This philosopher is not satisfied with knowing simply by ignoring, presupposing, or making *ad hoc* exceptions for what is difficult or impossible to prove.[41] Finally, this philosopher is not satisfied by arguments made -- at the expense of reason -- the sole purpose of which is to avoid the consequences of a position one does not like.

[41]The objective standards of knowledge; the nature of the external world; the reliability of perception, and the like.

WORKS CITED

Annis, David. 1978. "A Contextualist Theory of Epistemic Justification." *American Philosophical Quarterly 15*: 105-112.

Audi, Robert. 1993. "Contemporary Foundationalism." *The Theory of Knowledge*. Ed. Lewis Pojman. Belmont, CA: Wadsworth Publishing.

Austin, J.L. 1979. "Other Minds." *Austin: Philosophical Papers*. Ed. J.O. Urmson, and G.J. Warnock. Oxford: Oxford University Press.

Bonjour, Lawrence. 1993. "A Critique of Foundationalism." *The Theory of Knowledge: Classic and Contemporary readings*. Ed. Louis P. Pojman. Belmont, CA: Wadsworth Publishing.

Bouwsma, O.K. 1965. "Descartes' Evil Genius." *Bouwsma: Philosophical Essays*. Lincoln: University of Nebraska Press.

Cohen, Stewart. 1987. "Knowledge, Context, and Social Standards." *Synthese* 73: 3-26.

DeRose, Keith. 2000. "How Can We Know that We're Not Brains in Vats?" *The Southern Journal of Philosophy* XXXVIII: 121-149.

_____. 2000. "Ought we to Follow Our Evidence?" *Philosophy and Phenomenological Research* LX: 697-706.

_____. 1995. "Solving the Sceptical Problem." *The Philosophical Review* 104: 1-52.

_____. 1992. "Contextualism and Knowledge Attributions." *Philosophy and Phenomenological Research* LII: 916-929.

Descartes, Rene. 1911. *Meditations*. Trans. Elizabeth S. Haldane, and G.R.T. Ross. Cambridge: Cambridge University Press.

Dretske, Fred. 2000. "Entitlement: Epistemic Rights Without Epistemic Duties?" *Philosophy and Phenomenological Research* 60, No. 3, 591-611.

_____. 1970. "Epistemic Operators." *Journal of Philosophy* 67: 1007-23.

Empiricus, Sextus. 1933. *Outlines of Pyrrhonism, Book I.* Trans. R.G. Bury. Cambridge: Harvard University Press.

Goldman, Alvin. 1976. "Discrimination and Perceptual Knowledge." *Journal of Philosophy* 73: 771-91.

Hume, David. 1988. *A Treatise of Human Nature.* Ed., L. A. Selby-Bigge. Oxford: Oxford University Press.

Kim, Jaegwon. 1988. "What is Naturalized Epistemology?" *The Theory of Knowledge: Classic and Contemporary Readings.* Ed. Louis P. Pojman. Belmont, CA: Wadsworth Publishing.

Lewis, David. 1996. "Elusive Knowledge." *Australasian Journal of Philosophy* 74: 549-567.

_____. 1979. "Scorekeeping in a Language Game." *Journal of Philosophical Logic* 8: 339-335.

Moore, G.E. 1962. *G.E. Moore: Philosophical Papers.* New York: Collier Books

Nozick, Robert. 1981. *Philosophical Explanations.* Cambridge: Harvard University Press.

Popkin, Richard. 1980. "Skepticism and Anti-Skepticism in the latter part of the Eighteenth Century." *The High Road to Pyrrhonism.* Ed., Richard Watson and James E. Force. Cambridge: Hackett Publishing Company.

_____. 1965. "The High Road to Pyrrhonism." *American Philosophical Quarterly* 2, January: 1- 15.

Stine, Gail. 1976. "Skepticism, Relevant Alternatives, and Deductive Closure." *Philosophical Studies* 299: 249-61.

Stroud, Barry. *The Significance of Philosophical Scepticism.* Oxford: Clarendon Press.

Unger, Peter. 1975. *Ignorance: A Case for Scepticism.* Oxford: Oxford University Press.

_____. 1971. "A Defense of Skepticism." *Philosophical Review* 80.

WORKS CONSULTED

Cohen, Stewart. "Two Kinds of Skeptical Argument." *Philosophy and Phenomenological Research* 68, No.1, (1998): 143-159.

DeRose, Keith. "Knowledge Assertions and Lotteries." *Australasian Journal of Philosophy* 74, No.4, (1996): 568-580.

Feldman, Richard. "Contextualism and Skepticism." *Philosophical Perspectives* 13 (1999): 92-113.

Heller, Mark. "The Proper Role For Contextualism in an Anti-Luck Epistemology." *Philosphical Perspectives* 13 (1999): 115-129.

Putnam, Hillary. *Reason Truth and History*. Cambridge: Cambridge University Press, 1981.

Reiber, Steven. "Skepticism and Contrastive Explanation." *Nous* 32, No. 2, (1998): 189-204.

Schiffer. S. "Contextualist Solutions to Scepticism." In *The Proceedings of the Aristotlean Society held in the Senior Common Room* 20 May (1996): 317-334.