Title
Ordinary language philosophy : a critical re-examination

Permalink
https://escholarship.org/uc/item/3wd2k80x

Author
Thibodeau, Jason Bruce

Publication Date
2006

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA, SAN DIEGO

Ordinary Language Philosophy: A Critical Re-examination

A Dissertation submitted in partial satisfaction of the requirements for the degree

Doctor of Philosophy

in

Philosophy

by

Jason Bruce Thibodeau

Committee in charge:

      Professor Rick Grush, Chair
      Professor Paul Churchland
      Professor Anthony Edwards
      Professor Ronald Langacker
      Professor Gila Sher

2006

The Dissertation of Jason Bruce Thibodeau is approved, and it is

acceptable in quality and form for publication on microfilm

_____

_____

_____

_____

_____
Chair

University of California, San Diego
2006

iii

For Heidi and Orion

TABLE OF CONTENTS

ACKNOWLEDGMENTS

I would like to thank Rick Grush for serving as my advisor for the past five years. When I came to him with the idea for this project about a topic that was anything but fashionable, though he voiced important concerns, he was enthusiastic and supportive and has remained so ever since.

The remaining members of my committee, Paul Churchland, Gila Sher, Tony Edwards, and Ron Langacker have also been generous with their time and effort and I also wish to thank them. In addition, I thank Kelly Jolley and Eric Marcus for reading and commenting on drafts of various chapters

Not long after I selected Ordinary Language Philosophy as my topic, Avrum Stroll gave an informal seminar on Wittgenstein's *On Certainty*. This seminar proved vital to my understanding of Wittgenstein's response to skepticism and directly informed my arguments in Chapter 3 of this dissertation. Though not an official committee member, Avrum has continued to take an active interest in my project, reading and commenting on drafts of nearly all of my chapters in addition to other projects that I was working on. He is a consistent source of encouragement and inspiration. I am deeply indebted to him.

I have benefited from discussions with several fellow graduate students including Nellie Wieland, Cory Wright, and Aaron Schiller (all of whom attended Avrum Stroll's *On Certainty* seminar). While working on this project, the graduate student's at UCSD were kind enough to invite me to speak twice, on which occasions

life, he has provided insights and wisdom that are invaluable to an aspiring academic. More than that, Henry has become an archetype for me of a devoted and thoughtful academic, teacher, husband, father, and friend. Thank you, Henry.

My parents have encouraged my intellectual development for as long as I can remember. I'm sure many philosophers can remember puzzled and concerned looks on their parent's faces when they told them of their plans to pursue a graduate education in philosophy. With my parents, there was only pride, commitment to help in any way possible, and constant encouragement. This was true when, as an eight-year-old boy, I informed them of my ardent desire to become an astrophysicist, and it remains true to this day. There is nothing so important to a child, especially one who has chosen an unorthodox career path, as the knowledge that he makes his parents proud. Mom and Dad, I love you.

My sisters, Kristin and Sarah, along with their husbands, Matt and Luke, and daughters, Allyson and Anna, are another vital source of love and support. Kristin and Sarah have also always been proud and supportive of their brother. Having Kristin and Matt close at hand during Heidi's and my stay in San Diego was extremely comforting and helpful. Family members who are so willing to listen to each other's concerns, stroke a bruised ego, shoot the breeze, or spend days braving the mosquitoes in the Sierras are not a dime a dozen. You guys are great; I love all of you.

Without the unwavering love and support of my wife Heidi, my graduate career would have ended years ago, long before I had even chosen a dissertation topic. We, too, met in 1998 as new graduate students. Having a wife who shares my interest

in philosophy has been invaluable. She has read everything I have ever written and provided useful stylistic and philosophical advice. As someone who is outside the world of professional philosophy, but who retains a deep interest in the field, she provides a unique perspective and offers important insights into the issues that concern me. But far more important, she has provided consistent emotional support and inspiration, refusing to let me wallow in self-doubt, picking me up after discouraging set-backs, and, during times in which I was ready to give up the entire enterprise as pointless, letting me know that she, at least, felt that what I was doing was important. I cannot imagine a partner who could display more patience, commitment, affection, and devotion, in short, a partner who is more completely a wife than the love of my life, Heidi.

Finally, though he has only been around for the last few months of the dissertation writing process, my son, Orion, along with Heidi, is my chief source of inspiration. He has taught me new meanings of the words love, patience and commitment. When I look at his face, I experience an unbounded sense of devotion and responsibility to him; and when I think about the world he is likely to experience in the future, I am very glad that I chose the career that I did. My hope is that my commitment to rational inquiry will be a model upon which my son will base his life and decisions. Orion, I hope that I will live up to the love you have for me.

CURRICULUM VITAE

Areas of Specialization: Philosophy of Language, Metaphysics

Areas of Competence: Epistemology, History of Analytic Philosophy, Philosophy of
Mind, Applied Ethics, Wittgenstein

EDUCATION

University of California San Diego

Ph.D. in Philosophy, 2006
Dissertation: *Ordinary Language Philosophy: A Critical Re-examination*

M.A. in Philosophy, 2002

University of California, Berkeley

B.A. in Philosophy, 1997

PROFESSIONAL EMPLOYMENT

Auburn University (Fall 2005 – Present)
Visiting Assistant Professor

UNDER REVIEW

"The Standard Meter Bar and Performative Utterances"
"Natural Kind Terms: Linguistic Intuitions and Presuppositions"

PROFESSIONAL TEACHING EXPERIENCE

Visiting Assistant Professor, Auburn University, Department of Philosophy
Courses: Fall 2005—Introduction to Ethics
Spring 2006—Introduction to Ethics
Fall 2005—Introduction to Ethics

PRESENTATIONS

"The Standard Meter and Performative Utterances"
Alabama Philosophical Society, October 2005

"Natural Kind Terms and Natural Substance Terms"
UCSD Graduate Philosophy Colloquium, February 2005

"Teaching an Old Dogma New Tricks"
UCSD Graduate Philosophy Colloquium, April 2004

"Moore and Wittgenstein On Certainty and Skepticism"
UCSD Graduate Philosophy Colloquium, February 2003

REFERENCES

Rick Grush, UCSD

Paul Churchland, UCSD

Gila Sher, UCSD

Avrum Stroll, UCSD

Kelly Jolley, Auburn University

ABSTRACT OF THE DISSERTATION

Ordinary Language Philosophy: A Critical Re-examination

by

Jason Bruce Thibodeau

Doctor of Philosophy in Philosophy

University of California, San Diego, 2006

Professor Rick Grush, Chair

This dissertation argues that ordinary language analysis offers a means of critically evaluating substantive philosophical claims as well as providing a tool for constructing answers to philosophical conundrums. It has been neglected due to some powerful, though ultimately misguided, objections. Those who called themselves ordinary language philosophers did not effectively explain or defend their

methodology and thus these objections were never adequately met. This dissertation articulates and defends the methodology and its core theoretical commitments.

Ordinary language philosophy is essentially an empirical investigation into how words are used in ordinary (non-philosophical) contexts. Some believed that we could eliminate philosophical conundrums via such analysis. Wittgenstein, for example, thought that most philosophical problems were caused by linguistic confusion. I argue that the methodology does have an important critical role to play, but I also believe that it can be an important tool in constructing genuine solutions to problems.

There are different sets of objections to ordinary language analysis corresponding to two different parts of the analysis. One set of objects concern the empirical inquiry into how words are used in ordinary contexts. Primary among them is the concern that, because it relies upon the linguistic intuitions of a single philosopher, the process is too idiosyncratic. My methodology (derived from J. L. Austin's work) achieves objectivity by expanding the pool of people who are consulted and by actually looking at real cases (by, for example, reading the relevant literature) in which the terms and idioms under consideration are used. Native speakers are experts in using these terms and idioms, so, if we expand the pool of speakers enough, we should achieve an adequate picture of how and under what circumstances the relevant terms are used.

Another set of objections concern what I call the philosophical part of ordinary language analysis. Among them are concerns that it is committed to the allegedly

repudiated notion of the analytic/synthetic distinction; concerns that since meanings are constantly changing, information about meaning is irrelevant to philosophical investigation; and concerns connected to Grice's work on conversational implicature. I discuss each objection and demonstrate how it can be overcome.

CHAPTER 1
An Introduction to Ordinary Language Philosophy

Those philosophers who we regard as prototypical ordinary language philosophers seemed to be much more interested in practicing ordinary language philosophy than in articulating and defending the method and its theoretical commitments. Perhaps a philosopher or a method of philosophizing is best judged by its results but in the case of ordinary language philosophy it is difficult to judge the results in isolation from the method.

Ordinary language philosophy (OLP) is a species of conceptual analysis. I say this with some reservations because, in a couple of important respects, it is very different from the type of conceptual analysis that is often disparaged by its detractors as "armchair metaphysics." First, it is crucial to note that the results of ordinary language analysis can be fruitfully applied to philosophical problems even with little investigation into the relevant concepts (I'll discuss some examples below). Second, the aim of OLP is not usually to determine the necessary and sufficient conditions of a given concept. In fact, it is near a tenet of OLP that most ordinary concepts do not admit of this sort of analysis. Wittgenstein was very explicit in holding such a position and famously asserted this of the concept of game.

Nevertheless, ordinary language analysis does aim at a clear understanding of concepts. The philosophical aims of OLP are dependent on claims about the meaning of words and the content of concepts. And OLP posits a vital role for the analysis of

concepts and thus a role for *a priori* philosophy. This, I suspect, ruffles the feathers of many philosophers who have roughly Quinian intuitions. Ordinary language philosophers suggest that we can resolve fundamental metaphysical questions just by examining how we talk in ordinary contexts. This is, on the face of it, a fairly controversial claim; at least one that many contemporary philosophers would not be too quick to accept.

Ordinary Language Analysis—A Non-controversial Case

As a way of motivating the view that ordinary language analysis is both fruitful and often a necessary antidote to premature theorizing on the part of philosophers, I will offer a brief description of an argument that is based upon a claim about the ordinary use of an expression. The argument is, I hope most readers will agree, relatively uncontroversial. Sometimes philosophers butcher ordinary language to such an extent and reach such profoundly ridiculous conclusions that it is easy to refute them.

Norman Malcolm, in his paper "Moore and Ordinary Language," characterized the arguments of G.E. Moore in such papers as "Proof of an External World" and "A Defense of Common Sense" as defenses of ordinary speech against philosophical perversion. In his paper, Malcolm gives a list of twelve propositions which have been made by philosophers at various times and provides what he takes to be a response characteristic to one that Moore would give. It is worth quoting this passage at some length.

(1) Philosopher: "There are no material things."
Moore: "You are certainly wrong, for here's one hand and here's another; and so there are at least two material things."

(2) Philosopher: "Time is unreal"
Moore: "If you mean that no event ever follows or precedes another event, you are certainly wrong; for *after* lunch I went for a walk, and after that I took a bath, and after that I had tea."

(3) Philosopher: "Space is unreal"
Moore: "If you mean that nothing ever is to the right of, or to the left of, or behind, or above, anything else, then you are certainly wrong; for the inkwell is to the left of this pen and my head is above them both."

(4) Philosopher: "No one ever perceives a material thing"
Moore: "If by 'perceive' you mean 'hear', 'see', 'feel', etc., then nothing could be more false; for I now both see and feel this piece of chalk."

(5) Philosopher: "No material thing exists unperceived."
Moore: "What you say is absurd, for no one perceived my bedroom while I slept last night and yet it certainly did not cease to exist."

(6) Philosopher: "All that one ever sees when one looks at a thing is part of one's own brain."
Moore: "This desk which both of us now see is most certainly not part of my brain, and, in fact, I have never seen a part of my own brain."

(7) Philosopher: "How would you prove that the statement that your own sensation, feelings, experiences are the only ones that exist is false?"
Moore: "In this way: I know that *you* now see me and hear me, and furthermore I know that my wife has a toothache, and therefore it follows that sensations, feelings, experiences other than my own exist.

(8) Philosopher: "You do not know for *certain* that there are any feelings or experiences other than your own."
Moore: "On the contrary, I know it to be *absolutely* certain that you now see me and hear what I say, and it is absolutely certain that my wife has a toothache. Therefore, I do know it to be absolutely certain that there exist feelings and experiences other than my own.

(9) Philosopher: "We do not know for certain that the world was not created five minutes ago, complete with the fossils."
Moore: "I know for certain that I and many other people have lived for many years, and that many other people lived many years before us; and it would be absurd to deny it."

(Malcolm 1942, 346-347)

Some of the responses Malcolm puts in the mouth of Moore are almost taken verbatim from papers published by Moore (e.g. the response in (1)) and all of the others are almost certainly answers that Moore would approve of. Many of the positions put forward in the mouths of the Philosopher are positions that were held by some philosopher or other during Moore's career. Some are positions held by philosophers for centuries and in others we can recognize the views of famous 20[th] century philosophers.

The argument that I am interested in is, I think, the most absurd of the bunch: argument (6). According to Malcolm, Bertrand Russell made an assertion very similar to that made in (6)[1]. What are we to make of the claim that all that we ever see is a part of our own brains?

There are a couple of important things that need to be said about this statement. First, it is obviously false, at least upon sober reflection. But it is not just false. It has the appearance of an empirical discovery, but, unlike normal empirical claims, we do not need an investigation to determine that it is false. We can know that it is false prior to any investigation.

Now, being a philosopher, I must admit that it is possible that there could be creatures that only seen a part of their own brain. We can imagine such creatures in some detail. These poor souls are born with a two-inch by two-inch area of their scalp and skull at the back of their heads missing, thus exposing a four square inch area of

---

[1] I don't know that Malcolm is correct, i.e., I don't know that Russell ever asserted precisely this view. But this is not my concern, as we'll see this is a position that, because of familiar worries, a normally reasonable philosopher might be led to assert.

the brain. To continue with the absurd example, a further malformation has given them two long stalks emanating from their faces at the place where the eyes would normally be. These stalks curve over the top of the head and terminate in eyes that are focused on the exposed area of brain. The eyestalks are very rigid and are not movable. The eyes themselves are not in sockets and thus cannot rotate as our eyes do. Thus, such a creature must endure his entire life enjoying a view of only a very small part of his own brain.

About these unfortunates it would be true to say, "they only ever see a part of their own brains."   Such creatures are surely conceivable and thus, in one sense, it is possible that Russell's statement might have been true. It would certainly be true if we were such creatures. But of course Russell was not trying to assert that humans possess such an odd anatomy. Indeed, if we were in the circumstances described, we would hardly need a philosopher to point out the fact; we wouldn't want to be reminded of it.

Nor was Russell revealing the fact that he has never opened his eyes. Russell did not take himself to be making an empirical claim that is proven false every time he looks around. No, Russell was in fact saying something very different. He was asserting that we ordinary people, with our ordinary anatomy, only ever see a part of our own brains, even though it seems to us the we see lots of things. It is this assertion that cannot possibly be true.

We know that it isn't true because we know what the verb 'see' means and we know what we imply when we say, e.g., "I see a cat" or "I see my own brain." Any

descriptive sentence describes a certain situation (or situations) and not others. "The cat is on the mat" describes a situation in which a specified feline is resting atop some manner of small rug. It does not describe a situation in which a dog is standing on the highest branch of a sycamore tree. Nor does it describe a situation in which a scientist is peering through a microscope at a slide containing neural tissue. If anyone tried to use the sentence "the cat is on the mat" to describe a dog lying on a rug, we would rightfully assume something is amiss with this person. Perhaps he is a small child who does not understand the difference between a dog and a cat; perhaps he has some type of aphasia. In any event, the point is clear: for any given sentence, there are countless situations that the sentence does not correctly describe.

"I see a part of my brain" is a sentence that can be used to describe a certain situation. If I am undergoing brain surgery and am conscious, the doctors might arrange some mirrors so that I can see what they are doing. In this case it would be correct to say that I see a part of my own brain. But when I am looking into the backyard at the sycamore tree, I cannot accurately describe this situation by saying that I see a part of my brain.

Why would a philosopher assert that a person can only ever see a part of his own brain? Though the position is absurd, the motivation for it is straightforward and based on a familiar epistemic worry. There are familiar arguments, which I will not rehearse here, for the conclusion that we do not see objects. These arguments imply that when I look, for example at my cat, I do not literally see my cat, I see some 'sense-data' or a 'sense-impression' of my cat. If we are convinced, for whatever

reason, that we can never actually see objects, that there is a veil of sense-data between me and an object, it is just a short step to conclude that what I really see is a part of my brain. If sense-data are mental phenomena and mental phenomena are realized in, or in some way reducible to activity in the brain, then it is logical to conclude that sense-data are brain states of some kind. Thus, if all I see when I look at my cat is sense-data, then all I see is a part of my own brain.

The argument is clear and perhaps even well motivated; but the conclusion is clearly non-sense, as the response Malcolm attributes to Moore makes obvious. Most of us, like Moore, have never seen any part of our brain. Russell believed, I suppose (and assuming Malcolm is correct that Russell really did assert such a claim), that the statement "I see a cat" would be false if the sense-data theory is true. And he also believed that "I see a sense-datum" would be true and thus, assuming sense-data are realized in brain processes, that "I see a part of my own brain" would also be true. But in this he is clearly wrong. Even if the sense-data theory were true "I see a part of my own brain" would not describe a situation in which I am looking at my cat (a situation that I would normally describe as seeing my cat). In fact, assuming that there are sense-data, "I see a part of my own brain" would always be false because even in circumstances in which I am really looking at my brain (in the operating room, e.g.) I wouldn't see my brain, I would see sense-data.

So Russell's claim is just false, no matter how you slice it. It is important to remember, as I have already pointed out, that we know this *a priori*. We don't need to perform any experiment to determine whether we can see anything other than bits of

our own brains. All that we need to know is how to use the word 'see' and we know that Russell is wrong. (Russell, of course, knew how to use the word, but his theories about vision got in the way of this understanding.)

Moore's arguments were not as sophisticated in their appeal to ordinary language as those who followed him, but we can see in his arguments the understanding that ordinary usage trumped philosophical theory. I will discuss Moore's response to external world skepticism in detail in chapter 3 and show how his approach differed from that of Wittgenstein. But for now I only intend to explain why I believe that Moore's response to Russell's position provides good reason to investigate ordinary language analysis in more detail

This is a very interesting argument if you think about it. We have shown, just on the basis of linguistic analysis, that a substantial metaphysical position is false. We reached our conclusion just by focusing on what a sentence means and what we imply when we say it. This observation leads to a whole host of questions: How is this possible? How do such arguments work? What is the status of claims about how words are used? Is this really *a priori* knowledge? Is it empirical in some sense? What do we mean by ordinary? Since meanings are constantly changing, can there really be any value in arguments that rely upon the meanings of words?

What is Ordinary Language Philosophy?

To begin to answer these questions we need to understand in more detail what ordinary language philosophy is and what it is committed to. The simplest answer is

that it involves a commitment to the belief that ordinary language is relevant to philosophical inquiry. To move beyond this less than informative characterization, it is important to distinguish two essential parts of ordinary language philosophy. First, OLP contains what we might call an empirical aspect: the investigation into the content of concepts via an examination of how words and expressions are used in ordinary contexts. This investigation is conducted because the information gathered is believed relevant to the resolution of philosophical problems. Thus the second aspect is the claim that what is said in ordinary contexts has philosophical relevance.

It is important to draw our attention to the distinction between these two parts because, as I will show, some of the objections to OLP involve its empirical aspect while others question the relevance of the data gathered. We might have reason to doubt that there is a reliable way of gathering the data about linguistic usage. We may have particular concerns about particular methods of collecting the data. I will spend some time discussing different conceptions of the empirical aspect of OLP and the objections to these conceptions. With any particular conception about what, for lack of a better term I will call the experimental method, we need a clear account of what it meant by 'ordinary contexts.' We might wonder why some contexts, labeled 'ordinary' are given priority over others. These concerns all have to do with the first part of OLP, the empirical part.

But even if we believe that there is nothing inherently suspect about the project of investigating linguistic facts in ordinary contexts, we may still find reasons to be suspicious of the second part. We might wonder about the value of such an enterprise

given that concepts themselves are quite fluid and meanings change over time. Are data about ordinary usage nothing more than historical snapshots of a conceptual framework at a particular time; pictures of interest only to historical linguists? Conceptual change is a real phenomena and it is not immediately obvious why capturing ordinary concepts at a particular time is relevant to philosophical issues. Combine this with the Quinian observation that conceptual frameworks embody theoretical commitments and the worry becomes even more forceful. The Quinian intuition is that there is no way, in principle, to distinguish questions about the content of concepts (or the meaning of words) from theoretical claims about the world. Thus any information gathered about a particular concept, say voluntary action, will necessarily embody theoretical beliefs about the subject matter, in this case, freedom and action. I will spend some time addressing these objections later in this chapter.

In order to evaluate the claims made by ordinary language philosophy, its empirical nature must be properly characterized. It is, essentially, an investigation into how words are used and what they mean. How are claims about the use of a word justified? Do we take a poll? Or do we just consult our linguistic intuitions?

In essence the question we must address is "How do ordinary language philosophers know the things that they claim?" Austin, Ryle, and others made many claims about what we say under certain circumstances, what we mean when we say X, what is implied when we say X; claims that we don't say X unless C is the case, that to say X in certain circumstances is inappropriate or meaningless. Consider, for example, just some of the claims made by Austin in his paper "Other Minds"

"Whenever I say I know, I am always liable to be taken to claim that, in a certain sense appropriate to the kind of statements (and to present intents and purposes), I am able to *prove* it." (Austin 1946, 85)

"The doubt or question 'But is it a *real* one?' has always (*must* have) a special basis, there must be some 'reason for suggesting' that it isn't real, in the sense of some specific way or limited number of ways, in which it is suggested that this experience or item may be phoney." (87)

"If you say 'That's not enough' [i.e., 'That's not enough evidence for your claim that you know], then you must have in mind some more or less definite lack. . . . If there is no definite lack, which you are at least prepared to specify on being pressed, then it's silly (outrageous) just to go on saying 'That's not enough'. (84)

"If we have made sure it's a goldfinch, and a real goldfinch, and in the future it does something outrageous (explodes, quotes Mrs. Woolf, or what not), we don't say we were wrong to say it was a goldfinch, *we don't know what to say*." (88)

These are all claims about what we would say, what we would imply, or what we would mean. How does Austin know these things? How can we be sure that Austin's claims are correct?

To begin to answer this question it is worth taking a closer look at Austin's method, the method that he believed allowed him license to make and use such claims. Austin wrote relatively little about his method but, luckily, many of his contemporaries who saw him at work have commented on it. In K. T. Fann's

collection, "Symposium on J. L. Austin, J. O. Urmson described in detail the method used by Austin to generate and justify claims such as those above.

According to Urmson, Austin employed a four-stage process. In this first stage we identify an area of discourse (about, e.g., responsibility, knowledge, etc) and then collect as many of the terms, idioms, and expressions, speakers employ in their discourse about this particular area. For example, if we were interested in the issue of responsibility, Urmson says, we would collect such linguistic resources as, " 'willingly', 'inadvertently', 'negligently', 'clumsily', and 'accidentally', idioms like 'he negligently did X' and 'he did X negligently'." (Fann 78). We would like the list to be exhaustive (or as nearly so as possible), and Urmson offers some tools we might use to generate the list including free association and consulting the dictionary.

The second stage involves creating stories (or describing situations) in which the terms or idioms are (or would be) used. It is important that the imagined circumstances are similar to actual normal circumstances in which the idioms occur (Austin suggested we use case studies, such as documented court proceedings and judgments). These stories, "give as clear and detailed examples as possible of circumstances under which this idiom is to be preferred to that, and that to this, and of where we should (do) use this term and where that." (79). We also need to note when the use of a certain term or idiom is inappropriate in addition to those in which it is natural.

In stage three we move beyond the collection of evidence and begin to look for generalizations and propose explanations of the data generated in parts one and two.

"At this stage we attempt to give general accounts of the various expressions (words, sentences, grammatical forms) under consideration; they will be correct and adequate if they make it clear why what is said in our various stories is or is not felicitous, is possible or impossible." (80). It is at this stage that such claims as "When I say I know I imply that I am able to prove it" are produced. Urmson points out that before stage three premature theorizing is strictly forbidden:

> Austin always insisted that during the work so far described [stages one and two] all theorizing should be rigidly excluded. We must make up detailed stories embodying the felicitous and the infelicitous, but carefully abstain from too early an attempt to explain why. Premature theorizing can blind us to the linguistic facts; premature theorizers bend their idiom to suit the theory, as is shown all too often by the barbarous idiom found in the writings of philosophers who outside of philosophy speak with complete felicity. (80)

In this quote, Urmson succinctly sums up the reason 'ordinary' appears in 'ordinary language philosophy.' What struck Austin, Wittgenstein, Ryle, and others, as so problematic about so many philosophical theories was precisely the fact that they are based upon premature theorizing. Suppose we want a theory of the good, that is, a theory of what quality 'good' refers to (assuming it refers to one at all). Would it not be important to look carefully at all of the different circumstances in which 'good' (and related terms, such as 'bad', 'proper', 'decent', etc.) is used? It was one of the guiding beliefs shared by these philosophers that when we look carefully at how words (such as good) are actually used in non-philosophical contexts (before the theorizing has begun) we will see that many philosophical explanations are misguided or unfounded. Ordinary contexts are precisely those that occur prior to theories. Once we begin to offer explanations and theories, this will affect what we believe it is

appropriate or inappropriate to say. We have already seen and instance of this, in which Russell believed that it made perfect sense to say, "all I ever see is a part of my own brain." Russell accepted this claim because of the particular theory of vision that he had adopted (and he adopted this theory of vision in response to certain philosophical worries) without stopping to reflect on how odd his use of the verb 'see' really was.

So at the very least we should acknowledge that it can often be useful to generate the expressions and idioms (step 1) and the stories in which the expressions are used (step 2) before we begin any sort of theory. In step three we try to look for explanations that make sense of the data generated in the first two steps. Step four involves using our accounts and explanations to address philosophical questions. This may involve comparing our results to what various philosophers have said about the subject matter under consideration. Or it may involve a fresh approach to an old problem, or a new problem for that matter. It is true to say that most ordinary language philosophy is critical in nature. Step four would normally involve critically evaluating previous attempts to solve particular problems. But there is no reason that the analysis arrived at using this method cannot be put to positive use.

We return then, to the issue of how philosophers could know the things they claim to know about how words are used; how does Austin know all of those things that I listed above? As an initial reply I will quote Urmson's remark concerning the validity of the accounts reached in stage four of the method: "it is an empirical question whether the accounts given are correct and adequate, for they can be checked

against the data collected. Of course, if we have rushed the earlier stages new linguistic facts may be later adduced that invalidate the accounts, this is the universal predicament of empirical accounts" (80).

The method is undeniably empirical and this fact leads to some important worries. First, we may object that the process is too subjective or idiosyncratic. Why should we trust the linguistic intuitions of one philosopher? The appropriate response to this worry is that, at least according to Austin's ideal method, the process does not rely on a single person's idiosyncrasies. Austin believed that the method should be done by a group: "Austin always insisted that the technique was at all stages best employed by a team of a dozen or so working together; the members supplemented each other and corrected each other's oversights and errors." (79). The validity of the conclusions reached at stage three, then, does not depend on the intuitions of one person.

It is important to point out that not all ordinary language philosophers practiced the method as Urmson describes it. Wittgenstein certainly did not convene seminars in which everyone generated a list of terms and idioms related to the issue of rule following. Nor is it necessary to engage in this method to be confident about any claim about how an expression is ordinarily used. It is also useful to remember that every native speaker of English is an expert in producing the sorts of expressions collected in stages one and two of Austin's method. As Stanley Cavell puts it, "Such speakers do not, in *general*, need evidence for what is said in the language; they are the source of such evidence." (Cavell 1958, 4). Every native speaker of English is

capable, in principle, of generating the evidence needed to justify the accounts given in stage three.

In any event, conclusions reached at stage three (such as the examples from "Other Minds" given above) are empirical and subject to disconfirmation. They are claims about how words are used, what they mean, and what they imply. Even though all native speakers may be experts at generating appropriate English expressions, this knowledge is only implicit. It is notoriously very difficult to capture, in an explicit account, implicit knowledge. That is why Austin wanted the technique to be employed by a group. At the very least, the larger the group, the less idiosyncratic will be the results. And of course there is always the check of peer-review. Nothing guarantees that the results will be accurate, but such is the lot of all empirical investigations.

These considerations answer the general objection that the process may be irreducibly idiosyncratic. Of course we do not need to engage in the elaborate investigation suggested by Austin every time we want to make a claim about ordinary language. If we want a detailed account of our concept of 'responsibility' we probably will want to employ something close to Austin's method. But if our aims are less ambitious, if we just want to show what is wrong with Russell's view that all we ever see is a part of our own brains, we need nothing so fancy. Sometimes we need very detailed and precise accounts of how a given expression is used, other times the claims we rely on are quite general, or they may not require a complete account of the

concepts involved and then we can just rely on our knowledge as native speakers of the language.

Using Data About Ordinary Usage

So much for the empirical side of OLP. I turn now to objections that concern the second part, the philosophical part.

One of the more demanding objections to ordinary language philosophy stems from the Quinian rejection of the analytic/synthetic distinction. Quine held that no statement is immune to empirical revision and conversely, that any statement can be held true, come what may. The net effect of this thesis is that all statements are on the same epistemological footing. No statement can be accepted, or rejected, *a priori* because all statements equally confront the world of facts. Ordinary language philosophy, and indeed all forms of conceptual analysis, privileges some statements over others. Some statements are immune to revision because they concern the meanings of words.

Consider again the statement that all a person ever sees is a part of his own brain. Moore rejected this view because it implies something that Moore knows to be false, namely that Moore himself has seen a part of his brain. How does Moore know this to be false? Well, he certainly does not remember ever seeing a part of his brain and if he had seen it, he would surely remember having seen it. But this would not convince Russell since his point is that those events we would normally describe as seeing a sunset or seeing the ocean are actually not cases of seeing external objects but

rather cases of seeing one's own brain. As Malcolm pointed out, the dispute is about how the word 'see' is correctly applied. The dispute is, essentially, about the meaning of the word 'see' or, if you like, about our concept of "seeing". As I have argued, Moore has the right analysis; the meaning of the verb 'see' implies that it is not logically possible for a person to only be able to see a part of his own brain. Or rather, to eliminate a possible misunderstanding, we are not logically incapable of seeing anything but our own brains. This is a claim about what the meaning of the word 'see' requires. It is the meaning of this word that rules out Russell's theory of vision.

It is difficult to accept this argument from a Quinian standpoint. All statements are subject to empirical disconfirmation. The statement "We are not logically incapable of seeing anything but our own brains" is, according to Quine's analysis, subject to being proven false by empirical evidence. Because of this it lacks the epistemological status that is required for it to serve as a basis for rejecting Russell's claim. On the other hand, the claim "All I ever see is a part of my own brain" can be held true come what may. And no amount of linguistic evidence to the effect that it violates the meaning of 'sees' can force us to reject it.

Quine insisted that there was, in principle, no distinction between statements that are about the meaning of words and those that are about the world. However, ordinary language philosophy relies upon data concerning the use of expressions and, consequently, the meaning of words. Thus we need some account of the distinction between statements that are about meanings and statements that are factual.

In addition to claims that are explicitly about the meaning of words there are the more ubiquitous claims about what is implied by certain expressions (claims such as those I quoted above from Austin's "Other Minds"). In chapter three, I will argue that we are entitled to say "I know . . ." only when we are in a position to offer evidence for what is known; saying "I know" implies that I can offer evidence. Statements of this form are similar to so-called analytic statements in that they are taken to be a priori (at least by the ordinary language philosophers who make them). But it is difficult to understand them as analytic statements. "I know . . ." does not mean "I have evidence for . . ." In his paper "Must We Mean What We Say" Stanely Cavell considers the following statement (call it *S*) "When we ask whether an action is voluntary we imply that the action is fishy." About *S* Cavell says, "When (if) you feel that *S* is necessarily true, that it is *a priori*, you will have to explain how a statement which is not analytic *can* be true *a priori*. . . . it is perfectly true that 'voluntary' does not mean (you will not find set beside it in a dictionary) 'fishy.'" (Cavell 1958, 13). Cavell insightfully points out the similarity that such statements as *S* have to rules, for example rules in a game (e.g. "when a player's king is in check, that player is required to move his king out of check;" call this statement *C*). He says that such statements (statements like *S*, as well as rules like *C*) have a kind of double life. They are, on the one hand, descriptions (*C* is part of a description of the game of chess; *S*, of how English speakers use the word 'voluntary') and, on the other, imperatives, implying that something ought to be done (the king must be moved out of check, 'voluntary' should be used only when there is something fishy about the action in question).

This is an important insight, but Cavell offers little by way of explanation of how statements such as *S* can be *a priori*. " 'Statements about ordinary language' [such as *S*] are not analytic, and they are not (it would be misleading to call them) synthetic (just like that). Nor do we know whether to say that they are *a priori*, or whether to account for their air of necessity as a dialectical illusion, due more to the motion of our argument than to their own nature. Given our current alternatives, there is no way to classify such statements; we don't know what they are." (16). Suffice it to say that I am not content to leave the discussion there. In chapter 2, I will explain why statements like S, along with so-called analytic statements, have an air of necessity. I will argue that such statements are, in an important sense, immune from empirical revision, and I will offer a new classification scheme that makes clear how such statements differ from run-of-the-mill empirical claims.

Here is just a brief preview of my argument: We need to distinguish two types of statements that are immune to revision. First, statements that establish criteria for using a word; second, statements that report those criteria. Statements of the type that have traditionally been classified as analytic are actually reports about how we are to use the words involved (statements about the criteria of use for the words). "All bachelors are unmarried" reports the fact that 'bachelor' is used to refer to unmarried individuals. This statement is, epistemologically, no different from a statement of the form "We don't say 'I know' unless we are in a position to offer evidence." The latter is also a report about how a particular word is used. Statements of this type (which I'll call 'reporters') are not about the world in the sense that they are not about the objects

to which the words refer. Thus they are not subject to empirical refutation via an inquiry into the nature of those objects (e.g., we won't refute the claim that all bachelors are unmarried by investigating bachelors). They are, however, reports about how words are used (or, if you like, about the meanings of words) and, as such, are subject to revision in the sense that it is possible to misreport the meaning of a word.

In addition to the issue of explaining the presumed *a priori* status of statements like *S*, such statements raise an additional concern. This involves the distinction between semantics and pragmatics. It might be thought that statements about what is implied by an utterance have to do with the pragmatics of the utterance, not necessarily its meaning. After all, as Cavell says, 'voluntary' doesn't mean fishy. It may be true that we assume that someone who says that he knows is able to prove it; but this is not because 'know' contains, as part of its meaning, the concept of proof. Clearly it does not.

Thus, yet another type of objection to ordinary language analysis stems from the work of Grice and others in the field of pragmatics. In chapter 4 I'll present a pragmatics-based objection to ordinary-language-type responses to external world skepticism given by Barry Stroud in his influential book *The Significance of Philosophical Scepticism*. Stroud's objection suggests a general kind of response to many ordinary language analyses. In essence the response is that statements like *S* above concern conversational maxims that speakers follow and have nothing to do with the meaning of the terms involved (e.g. 'voluntary' or 'know'). If this sort of argument is successful in general, it would greatly undermine the value of ordinary

language philosophy since much of the data referred to would not concern meanings at all and hence would be unhelpful to any proposed conceptual analysis

Questions about meaning are not easily handled and have been often mishandled. What is the meaning of a word? One of the things that all the ordinary language philosophers shared was a suspicion of this question. Wittgenstein began the Blue Book with this question and famously said, "if we had to name anything that is the life of the sign, we should have to say that it was its use." (Wittgenstein 1958b, 4). In "The Meaning of a Word", Austin said the following: "This supposed general question is really just a spurious question of a type which commonly arises in philosophy. We may call it the fallacy of asking about 'Nothing-in-particular' which is a practice decried by the plain man, but by the philosopher called 'generalizing' and regarded with some complacency." (Austin 1961, 58). And Austin concluded, "there is *no* simple and handy appendage of a word called 'the meaning of (the word) 'x.'''

In "Ordinary Language" Ryle says,

They [philosophers] construed the verb 'to mean' as standing for a relation between an expression and some other entity. The meaning of an expression was taken to be an entity which had that expression for its name. So studying the meaning of the phrase 'the solar system' was supposed or half-supposed to be the same thing as studying the solar system. It was partly in reaction against this erroneous view that philosophers came to prefer the idiom 'the use of the expression '. . . caused . . .' and '. . . the solar system.' We are accustomed to talking of the use of safety pins, banisters, table-knives, badges and gestures; and this familiar idiom neither connotes nor seems to connote any queer relations to any queer entities. It draws our attention to the teachable procedures and techniques of handling or employing things, without suggesting unwanted correlates. Learning how to manage a canoe-paddle, a traveller's cheque or a postage-stamp, is not being introduced to an extra entity. Nor is learning how to manage the words 'if', 'ought' and 'limit'. (Ryle 1953b, 113).

Here Ryle echoes the sentiment of Wittgenstein in both the *Blue Book* and in the *Philosophical Investigations*. All three of these philosophers were very suspicious of the phrase 'the meaning of a word.' They preferred instead to talk about the use of a word. For them knowing the meaning of a word means knowing how to use the word. If you want to know the meaning of the word 'loquacious' you can look it up in the dictionary; this will tell you what the word means. But being able to use a word obviously involves much more than being able to recite its dictionary definition. It is worth pointing out that dictionary definitions can only be understood by people who already speak the language.

The dictionary definition of 'voluntary' will not contain the information that to say that an action was voluntary is to imply that there is something fishy about it; but that doesn't mean that it isn't true or that it is only a pragmatic consideration. Dictionaries rely upon a general familiarity with the language, but the use of a word is obviously much more complicated than can be captured in a short description. Wittgenstein often spoke of the "grammar" of an expression. What he meant by this was that a word plays a certain role in the language. The role is very complicated, not exhausted by the dictionary definition, which can be captured only by a very complicated description of the rules that govern its use. When we learn the meaning of a word, we lean how to use the word.

If we understand meaning in this way, we will be less impressed by the distinction between semantics and pragmatics. Though this conception of meaning is, I believe, an assumption common to many ordinary language philosophers, it is no

simple task to articulate and defend it. Defending such an account of meaning is indeed a dissertation project unto itself. In this dissertation, rather than undertaking this task, which is hardly my primary concern, I will argue for a much more conservative conclusion. Namely, I'll argue that the general argument based upon the pragmatics/semantics distinction (and that used by Stroud) is not the universal ordinary-language-argument killer that Stroud, at least, makes it out to be. Though some who have used ordinary language analysis in the past have too quickly and easily moved past the concern expressed by Stroud (and indeed Grice before him), there is no general argument to the effect that the distinction between pragmatics and semantics renders hopeless all (or even most) claims to the effect that statements like *S* really do derive from the meanings of the terms under discussion.

By way of preview, here is a brief sketch of my argument: Austin, most famously in his paper "Other Minds", identified certain general conditions attaching to the use of words such as 'know' and its cognates; e.g. that when a speaker doubts or calls into question some knowledge claim, the speaker must be able to articulate the specific defect or lack which renders the knowledge claim subject to doubt (see chapter 4 for more detailed analysis of Austin's views). It is clear that Austin takes these as conditions on the meaningful utterance of these terms. Stroud argues that the conditions identified by Austin do not derive from the meaning of the terms but rather are the result of certain pragmatic requirements on conversation. Thus, for example, Austin believes that it would be unacceptable (in ordinary, non-philosophical discourse) to demand the elimination of a Cartesian-style dream hypothesis, and

further believes that it is unacceptable precisely because the meaning of the word 'know' entails that a possibility only undermines a piece of knowledge when there is some special reason to suppose that the possibility actually obtains. Stroud, on the other hand, acknowledges the impropriety of demanding the elimination of such doubts but insists that this impropriety is best accounted for by the hypothesis that the demand for the removal of Cartesian-style worries violates a conversational maxim. Thus, according to Stroud, the condition mentioned by Austin (that we have some special reason to suppose that a possibility actually obtains) is merely a pragmatic condition.

My response to Stroud will be that even if he is correct about Austin's condition, there is a different, though related, condition for which his pragamtics-based analysis does not work. Stroud himself admits that we do not (in ordinary contexts) demand the removal of doubts that are logically impossible to remove. In chapter 4, I argue that this linguistic phenomenon cannot be accounted for by the supposition that we are following a conversational maxim such as that proffered by Stroud. Rather, it is best accounted for by the alternative thesis that the condition derived from the meaning of the word 'know.' My general conclusion, as it relates to the project of Ordinary Language Analysis, is that it is always possible that the conditions identified by practitioners of the method are merely pragmatic and do not thus tell us anything of importance about meanings; however, there is no general reason to suppose that this must be the case, each case will be different. And, importantly, we should acknowledge that any philosophical project employing

Ordinary Language Analysis must include the additional step of considering whether the results obtained might best be explained by appeal to pragmatics.

Plan of the Dissertation

The remainder of this dissertation is devoted to developing, defending and using what I have referred to as the second part of Ordinary Language Analysis (or, as I have also been calling it, the philosophical part). In chapter 2, I defend the method against the powerful Quinian worries I mentioned earlier. Chapters 3 and 4 are devoted to an Ordinary Language Analysis of the problem of external world skepticism. In chapter 3 I develop a response to this sort of skepticism, a response which draws inspiration from Wittgenstein's remarks in *On Certainty*. This chapter thus provides a model for how information garnered from the empirical investigation into how words are used can be used to address philosophical issues. Chapter 4 is a response to the sort of pragmatics-based objection to Ordinary Language Analysis I mentioned above. Finally, in chapter 5, I turn the tools I have developed and defended toward a different project, namely the mind-body problem and specifically mind/brain reduction. For much of the time in this dissertation, I will be pre-occupied with external world skepticism and various responses to it. I think it is useful to see the method I have developed and defended put to use in a different context. This is my aim with chapter 5.

CHAPTER 2
The Analytic/Synthetic Distinction

It is hardly a secret that, in the milieu of contemporary philosophy, ordinary language analysis is passé. To the extent that anything like the method I described in chapter 1 is still practiced today it is usually done very much in the background and without the rather untroubled self-assuredness characteristic of the work of Austin, Ryle, and Wittgenstein. There are multiple reasons why this is the case. Much of it has to do with a failure on the part of many who used the techniques of ordinary language philosophy to do so in a careful and consistent manner. Undoubtedly the method can and has been misused. Another concern stems from the work of Paul Grice on pragmatics (a concern I address in chapter 4). But I think that foremost among the reasons that ordinary language philosophy fell on hard times stems from Quine's rejection of the analytic/synthetic distinction and the ensuing suspicion of arguments based on conceptual analysis. If Quine is correct that all statements equally confront the world of facts and no statement is immune from empirical revision, then philosophers can hardly privilege some statements (e.g. statements about the meaning of some term or set of terms) as inviolate propositions from which we can build (or criticize) solutions to philosophical problems.

It is not too far a step from the rejection of the analytic/synthetic distinction to a rather dogmatic scientism according to which there is no role whatsoever for

philosophy (understood as conceptual analysis) in the acquisition of knowledge.[2] While it is true that this kind of radical scientism has hardly captured the hearts and minds of the entire field, at the very least one gets the sense that most contemporary philosophers have rather ambivalent feelings about the role conceptual analysis might play in their field. Since ordinary language analysis is one type of conceptual analysis, it too must be suspect.

I think such concerns are misplaced. As I'll show in this chapter, Quine is quite wrong that all statements are immune to revision. Statements about meanings of terms (such as, " 'bachelor' refers to only unmarried men") do imply statements that are immune from revision (e.g. "all bachelors are unmarried"). Quinian concerns about the epistemic status of statements of either sort are misplaced.

For the purposes of this chapter, I'll be connecting this issue about the analytic/synthetic distinction with Wittgenstein's famous comments about the standard meter bar. "This bar is one meter long" will thus serve as my paradigmatic example of a statement that is immune from revision.

Introductory

That there are no statements that are true in virtue of meaning, or are immune to empirical revision, is a position widely held among contemporary philosophers. Since the publication of Austin's *How To Do Things With Words*, we have also been aware that those utterances Austin calls performatives do not have a truth-value.

---

[2] But see Gila Sher's "Is There a Place for Philosophy in Quine's Theory?" *The Journal of Philosophy* 96 (1999): 491-524.

About such statements as "I bid you good morning" and "I promise to take out the trash" the issue of truth does not arise because they are instances of performing an action (bidding you good morning and promising, respectively) and do not represent a pre-existing state of affairs. On this basis it is plausible to conclude that performatives are not subject to empirical revision. Nor are they true in virtue of meaning for the simple reason that they are not true. Most philosophers also will recognize that it is possible to use a performative (or, to be precise, a declaration) to define words (as in, "I hereby define 'bachelor' as 'an unmarried man'"). Whether and to what extent this is actually done in any natural language is open for debate but that it can be done shouldn't be. This suggests, contra Quine, that there are utterances that are about the meanings of words but are not subject to empirical revision (though not true by meaning).

Quine responded to Austin's views about performatives in "Symposium on Austin's Method", a collection of comments on Austin's philosophical career.[3] Quine says, " 'I bid you good morning' is true of us on a given occasion if and only if, on that occasion, I bid you good morning. A performative is a notable sort of utterance, I grant; it makes itself true; but then it is true." (Fann, 90). In this Quine is certainly mistaken. Austin's point was that performatives are, in essence, utterances whereby the speaker performs some action. Saying, "I bid you good morning," sincerely and in the appropriate circumstances, is the act of bidding you good morning; it is not a report of what I am doing, it is the act of doing it. The issue of truth no more arises for

---

[3] Published in K. T. Fann's "Symposium on J. L. Austin" (London: Routledge and Kegan Paul. 1969).

this utterance than it does for a wave of the hand. Of course there is an ambiguity created by the fact that 'bid' is both the present tense and past tense form of the verb. "I bid you good morning" can be used to report the fact that I have bid you good morning, in the past: "I bid you good morning (yesterday, but I'm not going to today)." But this doesn't prevent its use, in the present tense, as a performative. A genuine performative use of 'I bid you good morning' does make the past tense version true, but the performative itself is neither true nor false. The performative no more makes itself true than waving hello makes itself true.

This ambiguity suggests another. Suppose that the scientific authorities of the late eighteenth century uttered the statement "We hereby define 'meter' as the length of this platinum bar." Unambiguously this is a declaration. They could also have said simply, "This bar is one meter long," dispensing with the theatrics. We can certainly argue about whether this declaration would establish that 'meter' actually does mean this. I submit that it would so long as those who use the new term agree to abide by the declaration. In these circumstances it would be a fact about our usage of 'meter' that it refers to the length of the bar. But even if it did not succeed, so long as people used the term in some other (roughly uniform) way, there would still be facts about how the word is used. And of course it would be possible to make claims about what 'meter' means. An utterance of "This bar is one meter long" subsequent to the initial, successful performative, utterance, would be a report about the fact (about 'meter') that was established via the declaration.

What, then, is the status of a subsequent utterance of "This bar is one meter long" in regards to its potential for empirical revisability? As a report about the meaning of the word 'meter' it is revisable since it is potentially an incorrect report about that meaning. If the declaration by the authorities of the eighteenth century was efficacious, then it is a true report. If it is a true report, though, what are we to make of its status as a claim about the length of the bar? Well, in this case, it cannot possibly be true to say "This bar is not one meter long" since it has been established that 'meter' refers to the bar's length.

Or consider the empirical status of the statement "All bachelors are unmarried." It is unlikely that anyone established by fiat that 'bachelor' meant 'unmarried man.' Still, there are facts about what it does mean, about how we actually use the word. According to the old dogma rejected by Quine, we would understand this statement to be true just in virtue of the fact that 'bachelor' is defined as 'an unmarried man.' But its quite ambiguous to say that "All bachelors are unmarried" is true by definition. Is it a true claim about how 'bachelor' is used; or is it a true claim about bachelors? Quine rejected the notion that it is an empirically non-revisable statement about bachelors. But as a claim about the meaning of 'bachelor' it isn't non-revisable since, just as in the 'meter' case, it is possible to be wrong about what the word means. "All bachelors are unmarried", understood as this type of meaning report, is indeed subject to empirical revision. But the truth of this meaning report implies something very interesting about the report about bachelors (not 'bachelors'). The meaning report says that it is a criterion for being called a 'bachelor' that a person

be unmarried. If this is a true report about the criterion, then any object properly called a bachelor will be unmarried. Thus it will not be possible for a bachelor to be married and "All bachelors are unmarried", understood as a claim about bachelors (as opposed to a claim about the meaning of 'bachelor'), cannot be false.

The notion of truth by definition is indeed problematic, not (or not only) for the reasons that Quine suggested but because analytic statements are systematically ambiguous: they can be seen to be both reports about the meanings of the words and about the referents of those words.

These considerations suggest the following *reductio*:

(1) There are no statements that are immune to empirical revision.
(2) There are (empirically revisable) statements about how words are used.
(3) The statement "All bachelors are unmarried" is a true report of the type described in (2).
Therefore,
(4) Since it is a true report, in order for an object to properly be called a bachelor, it must be unmarried.
Therefore,
(5) It is not possible for a bachelor to be unmarried
Therefore,
(6) "All bachelors are unmarried" as a report about the extension of 'bachelor' (not about the meaning) cannot be false.
Therefore,
(7) There is a statement that is immune to empirical revision.

Perhaps we can resolve this by refusing to accept (2). Unfortunately there is no principled reason for doing so. These statements are not non-revisable, they aren't true by definition, and they aren't insulated from empirical inquiry. We might think of them as things that we learn when we learn the language. As such they are

discoverable by linguistic investigation. We could discover, for example, that every object that is called a bachelor is also called an unmarried man.

The Quinean would be suspicious of the move from (3) to (4). In "Two Dogmas of Empiricism" Quine says "There is no assurance here that the extensional agreement of 'bachelor' and 'unmarried man' rests on meaning rather than merely accidental matters of fact, as does the extensional agreement of 'creature with a heart' and 'creature with kidneys.'" (Quine 31). But this requires an unrealistically narrow view of the type of information that a linguistic inquiry could discover. We need not be limited to simply watching verbal behavior, identifying the objects that are called 'bachelor' and 'unmarried man.' Given that we speak the language, we can ask speakers such questions as "On what basis do you decide that a person is a bachelor?" And if the answer is "Well, if he is an unmarried man, that suffices," we can be pretty certain that the co-extensionality is no mere accident.

If we wanted to conduct an empirical inquiry into the question of whether there are any married bachelors, we would first have to have some way of identifying the bachelors. If it is true that being an unmarried male is the criterion for being properly termed 'bachelor,' then as we gather our sample of bachelors, we must reject any person who is married. Thus our sample will, of necessity, be comprised only of unmarried people. Again, if "all bachelors are unmarried" is an accurate report of how we use 'bachelor,' it is impossible to find a married bachelor. So the problem is real and cannot be solved by adhering to the belief that the only discoverable facts about

meanings will not suffice to show that agreement in extension is not simply an accident.

Notice how this case differs from the 'creature with a heart'/'creature with kidneys' case. That a creature has a heart is a good indication that it has kidneys. But that a person is an unmarried man is no mere indicator that he is a bachelor, it is constitutive of his bachelorhood

My own view is that we should reject (6) and (7); that is, we should maintain that "All bachelors are unmarried" (as a claim about bachelors) does not express a content. It is neither a true nor a false report about bachelors. The only way to give content to the sentence is to interpret it as a claim about the term 'bachelor.' It is a report about meaning disguised as a universal generalization. To extend the point, all so-called analytic statements are actually reports about the meaning of the terms involved and do not express a content about the referents of those words.

What are the independent reasons for thinking this? Consider what Robert Stalnaker says in the introduction of his book, *Ways a World Might Be*, "what it is to represent the world—to say how things are—is to locate the world in a space of possibilities. One understands what someone else is saying by understanding how that person is distinguishing between the possibilities, as one takes those possibilities to be." (Stalnaker 8). If we accept this reasonable view, then we should be suspicious of statements that cannot possibly be false. "All bachelors are unmarried," I have argued, is necessarily true. Thus it does not distinguish among possibilities since the

only possibility is that represented by the statement. So it is plausible, on this ground, to claim that "All bachelors are unmarried" does not express a content.[4]

In any event, in this paper I will argue that there are sentences that report the meanings of words but do not express content about the referents of the terms involved. Those statements traditionally thought of as analytic fall into this category.[5] As an example, I'll focus my efforts on the sentence "The standard meter bar is one meter long." I'll show that (in the relevant historical context) this sentence cannot be used to express a fact about the standard meter bar. First, though, I need to introduce some terminology.

Statements about meaning—Producers and Reporters

There is a distinction between statements that are about the meaning of words and those that are about the world. But even within the former category, those about meanings, there is a distinction to be made. One type of sentence establishes facts about the meanings of words and another reports these facts. Because of this, there are at least two distinct types of statements whose semantic and epistemological properties differ from run-of-the-mill empirical statements. The first of these types establish the criteria of use for a term (and thereby establish at least part of the word's meaning)

---

[4] At least not one about bachelors. It is, as I've argued, a contentful claim about 'bachelor.'

[5] I avoid discussing mathematical statements in this chapter. However, some readers may wonder whether they are empty of content in my sense since they too can't possibly be false. 'Indeed, '2 + 2 = 4' is not a statement about objects, according to the analysis that I defend. But this doesn't imply that it isn't true anymore than the fact that "All bachelors are unmarried" is not a claim about bachelors implies that it isn't true. "All bachelors are unmarried" is a true claim about the meaning of 'bachelor' and ' 2 + 2 = 4' is a true claim about a particular mathematical function. Though this view may be controversial, I don't think it is incoherent. Unfortunately, defending this position on mathematical statements is far beyond the scope of this dissertation.

and as such cannot accurately be called empirical. These statements are neither true nor false; they do not report facts. Consequently they are also known neither *a priori* nor *a posteriori* for they do not express knowledge claims at all.

The second type of statement reports the criteria of use of terms that have already been established. As I argued above, "all bachelors are unmarried" can be used to report the criteria (or part of the criteria) of use for the word 'bachelor'. As such, a statement of this type does not report a fact about the referents of the terms used and thus is not a claim that is true or false about those referents. However, unlike the first type of statement described above, such a report has a determinate truth-value; it is a report about how certain words are used (though not about material objects). Since any such report could be wrong it is open to disconfirmation but only in the sense that it could be an incorrect report about the meanings of the words.

Statements of the first type I will call 'meaning-producing statements' (producers for short), those of the second type, 'meaning-reporting statements' (reporters). This distinction can be somewhat confusing due to the fact that one and the same sentence can be used to make both a statement of the first type and one of the second. (I will shortly provide an illustration of this). For this reason I will sometimes speak of meaning-producing uses of sentences as well as meaning-reporting uses. I will be contrasting these two types of statements with run-of-the-mill empirical statements that express factual information about objects.

Statements of the first type (meaning *producing* statements) are not revisable in the way that common empirical statements such as "Pluto has two moons" are.

Since they do not report facts, there is nothing to revise. Those of the second type (meaning *reporting* statements) *are* revisable; but when they are revised what we learn is not some new fact about material objects, we learn something about the meaning of words. Also they are revised not by investigating the actual properties of e.g. bachelors, but by investigating how words are actually used. So there is also a sense in which reporters are not revisable. When they are discovered to be false, we do not revise the list of facts that we believe to be true about the referents. In other words, reporters do not express claims about material objects that might be revised. They are reports about the meaning of words. It is these claims that can be revised.

Consider, again, the sentence,

(1) "This bar is one meter"

Sentence (1) can be used to make a producer, a reporter and a normal empirical claim. That is, it can be used to establish the meaning of 'meter,' report the meaning of 'meter' and to report the length of an object in meters.

In the eighteenth century, when scientists decided to codify a new system of measurement standards, it was decided that the basic unit of length was to be called a 'meter.' The meter was to be equivalent to one ten-millionth of the distance from the North Pole to the equator. The French government financed an ambitious project to measure the section of the meridian running from Dunkerque to Barcelona.[6] When the measurements were completed, a bar of platinum was forged of the exact length and

---

[6] See Ken Alder's *The Measure of all Things* (New York: The Free Press, 2002) for a detailed history of this expedition.

thus became the standard meter bar.[7] I am not familiar with the precise history of this particular bar but we can assume (or not so unrealistically imagine) that after it was forged there was some form of ceremony during which the scientific and political powers of the day declared that the length of this bar should be one meter. This is our first instance of the sentence "This bar is one meter." The utterance established, once and for all, what length the term 'meter' stands for. Before the codification process, there was no fact of the matter about how the word 'meter' should be used; there was not an established meaning.[8]

Since the initial utterance of (1) established the meaning, it is not revisable; that is, it could not have gone wrong in the usual way. Let me say something about what I mean by 'the usual way.' Now-a-days, more than 200 years after the adoption of a standard, when I measure something, e.g. the width of my desk, and declare that it is one meter across; there is always the possibility that I have made an error. Perhaps upon closer and more careful inspection, it will turn out that the desk was actually 1.01 meters or only 0.99 m. My claim that my desk is 1m wide thus is revisable in the sense that I could be wrong. The statement made by the eighteenth century scientists is not revisable in that sense.[9] In order to show that my claim about the width of my

---

[7] Actually, as Alder recounts, things were a bit more complicated. The measurements were slightly off and as a result the bar was approximately 0.2 millimeters short. The distance from the North Pole to the Equator, therefore, is 10,002,290 meters, not 10,000,000 m. (see Alder, pp 5-9). See my footnote 18 for more about the implications that this has for my account.

[8] Or if there was, it was based upon a different practice and a different standard. The adoption of the new bar superceded any pre-existing practice and established a precise and universal standard for the entire world to employ.

[9] There is, perhaps, a sense of 'revisable' according to which it is revisable. We can certainly change the standard; we can adopt another bar or choose some other criteria (something that has been done at least twice). In a sense this would amount to a revision of the statement " The standard meter bar is one

desk is wrong, you must re-measure it or use a meter rule that is more accurate. But the idea of re-measuring the SMB (in meters) or of using a more accurate measuring tool to measure it is simply incoherent. To be sure, it does make sense to speak of re-measuring a bar that is a candidate for becoming the Standard Meter Bar to make sure that it is of the desired length. But once this particular candidate has been chosen, it is the Standard and as such it is the final and ultimate authority. No more accurate tool can be found because it is, by explicit decree, the most accurate measuring device in existence. We can't say, "Wait just a second. I've got my own meter stick and I'll see what is and what isn't precisely one meter long." There is no measuring device we could use to measure the Standard Meter bar that would show us that it isn't a meter long for the simple reason that, as the SMB, it trumps all claims to authority.

In addition to this first use of (1) there is another use, a meaning-reporting use. Imagine that you are visiting the museum in which the SMB is kept. Since the bar is in Paris and the signs are written in French, some fellow American tourists can't figure out what the bar is supposed to be. After deciding that it isn't an Andy Warhol piece, the group realizes that it must be the standard for a particular unit of measurement. However, they are split as to which unit it is: is it the yard or the meter? You take the opportunity to reveal your knowledge of the history of measurement and declare, "This bar is one meter long." There are many ways that this statement could go wrong. Your French may not be as good as you had believed: the bar might really be a Warhol piece; or it may be the standard unit for a defunct system of measurement

---

meter long"; but what we have done is change the meaning of 'meter.' We have not discovered evidence that the SMB isn't one meter long.

etc. In short, you could be wrong about what bar it is. But it is important to see that the aim of your statement is not to give the length of the particular bar, but to identify what the bar is by explaining what role it plays. And assuming that you are correct, you are not so much telling your friends the length of the bar as telling them the role that the bar plays in the system of metric measurement. This statement is revised, then, by investigating the role of the object in question, not by measuring its length. This is because the statement is not about the bar's length but about its use and status.

The third statement that can be made with the sentence at hand is an ordinary, run-of-the-mill empirical claim about the length of an object. In this case, someone is measuring a bar at a construction site, for example, and declares, "This bar is a meter long." This statement is open to empirical disconfirmation in the usual way; it can be falsified by a more accurate measurement.

I need to make one point of clarification here. I have said that (1) may be used to make a meaning-producing statement, a meaning-reporting statement, and an empirical statement. This is true but notice a subtle complication: In the producer and reporter statements the phrase 'this bar' refers to the same bar—the standard meter bar—but in the empirical statement it does not. It is a feature of my position that we cannot make an empirical use of (1) about the SMB. Thus, as we consider different possible uses of (1) it will be important to keep in mind which bar the 'this bar' in (1) refers to. As an expedient I will introduce a new statement just like (1) but with the caveat that 'this bar' will always designate the SMB:

(1a) This bar is one meter long.

(*Where 'this bar' refers to the standard meter bar.*)

How Producers Work

Earlier I claimed that statement (1a) can be used as a declaration establishing the meaning of 'meter.' I also pointed out that declarations and performatives are unique in that that they have the form of assertive statements but are not subject to the same standards of assessment as normal assertives. Peformatives, such as "I promise to pay you five dollars" and "I order you to cease and desist" and declarations such as "Court is adjourned," and "I now pronounce you man and wife" have the form of assertive sentences but are not subject to disconfirmation in the usual way. When a judge says, in the appropriate circumstances, "Court is adjourned," his statement has what Searle has called a "self-guaranteeing" character.[10]

Performatives and declarations have been discussed extensively so I will here only provide a brief overview of the features that are key to the present discussion.[11] The differences between declarations and assertives are best explained via the notion of direction of fit. Speech acts are said to have either the word-to-world or world-to-word direction of fit depending on whether the aim of the act is for the content to match the world (e.g. assertions) or whether the world is to match the content (e.g. requests).

---

[10] "How Performatives Work" in his *Consciousness and Language* (Cambridge: Cambridge Universtity Press, 2002)

[11] I will be drawing primarily on John Searle's work in *Expression and Meaning*, (Cambridge: Cambridge University Press, 1979).

Declarations, like other speech acts, can be successfully or unsuccessfully performed. But, unlike assertives, they do not aim at representing a pre-existing reality and, unlike directives, they are not evaluated on the basis of whether the world eventually changes to match them. This is because declarations change the world just in virtue of being performed. "It is the defining characteristic of this class that the successful performance of one of its members brings about the correspondence between the propositional content and reality, successful performance guarantees that the propositional content corresponds to the world."[12] Since a declaration both tries to create a change, and, when successful, its content matches the state of affairs consisting of that very change, its direction of fit is both world-to-word and word-to-world. As Searle makes clear, "declarations do attempt to get language to match the world. But they do not attempt to do it either by describing an existing state of affairs . . . nor by trying to get someone to bring about a future state of affairs."[13]

It is not at all surprising that words can be subjects for declarations; this is done quite often. Scientists, lawyers, and philosophers readily engage in this sort of activity: defining new words, tweaking and making more precise the definitions of pre-existing ones. The scientists of nineteenth century Europe were simply performing the not uncommon act of defining a new word. It should also be clear that the act could have been performed with more than one utterance. Though I have no idea what actually was said in 1792, it doesn't really matter for their purposes whether the scientists said, "This bar is one meter long" or "We hereby define the 'meter' as

---

[12] *Expression and Meaning,* p.16
[13] *Expression and Meaning,* p19

the length of this bar." Both are equally effective declarations. Similarly, "You are now man and wife" is as effective means of marrying a couple as is "I now pronounce you man and wife."[14]

Let's briefly summarize the consequences that follow from its status as a declaration: First, this statement was not an attempt to describe an independently existing reality; it was an attempt to make something the case merely by saying that it is. It has both the word-to-world and world-to-word direction of fit. The goal of the statement, then, is not to indicate the length of the bar but to make it the case that the word 'meter' refers to that length. The statement makes it the case that 'meter' has that meaning (world-to-word direction) and also states that it has that meaning (word-to-world direction). Thus, second, the statement is about the meaning of the word 'meter' and not about the length of the bar. Finally, since it is a declaration and not an assertion, the statement is not subject to empirical revision. It cannot be evaluated for truth or falsity as assertions can because it creates the state of affairs that it represents.[15] In other words, it has all of the properties that I listed for producers: it is about meaning, is neither true nor false, and is not open to empirical revision.

---

[14] At this point some may be wondering who has the authority to define words. Can anyone just decide to define any word any which way he chooses? The short answer is yes, anyone can, but you cannot expect that the rest of the world will accept and use your definition. Many, if not most, declarations require that certain extra-linguistic institutions be in place for them to be successful. Not just anyone can adjourn court. The power is reserved for the presiding judge; only his words have efficacy on that matter.

[15] Since, as I indicated, declarations have a self-guaranteeing character that insures that their content matches the world, one might think that, in a perfectly normal sense, they are true. Strictly speaking this may be correct; at least I am not prepared to dispute it here. Nonetheless, declarations do not attempt to describe a pre-existing reality and the point I am making here is that "This bar is one meter long" (understood as a declaration) is not a true description of the length of the standard meter bar. It establishes the meaning of 'meter' and so is, at best, a true description of the meaning of 'meter'. This point will become clear when I compare my position to that of Boghossian in what follows.

Paul Boghossian has argued[16] that the fact that a sentence is used in an ostensive definition of a word does not preclude it from asserting a fact. He puts the point very nicely. About the sentence,

(2) Stick S is a meter long at t.[17]

Boghossian says, "Suppose that stick S exists and is a certain length at t. Then it follows that 'meter' names that length and hence that [2] says that stick S is that length at t, and since it is that length at t, [2] is true."[18]

According to my analysis, however, (2) is not true and does not express a fact. If I am correct about its status as a declaration, then I believe that this conclusion is inevitable. However, I won't simply leave the issue at that. It is obviously controversial and deserves more explication and defense.

There *is* a class of declarations that both create a state of affairs by asserting that it exists and yet also attempt to describe a pre-existing reality. Examples include an umpire's "You're out" and "Strike three" or a jury's "We find the defendant guilty". These do attempt to correctly describe a state of affairs that exists independently of the utterance but also have the force of creating a state of affairs (at least in the eyes of the law, so to speak). When an umpire says "you're out" he is trying to state a fact –that the ball reached the base before the runner, e.g. –but the runner really is out regardless of whether the umpire got it right. The runner can

---

[16] "Analyticity" in B. Hale, C. Wright, eds., *A Companion to the Philosophy of Language*, (Oxford: Blackwell, 1997)

[17] Just to avoid possible confusion I should say that, for the purposes of the following discussion I will assume that (2) is uttered in the context I described earlier for "This bar is one meter." That is, in the eighteenth century, by the political and scientific authorities of the day.

[18] "Analyticity" p. 350

appeal, but so long as the decision is not overturned, the runner really is out (he must leave the playing field) even if the umpire got the call wrong. Similar things can be said for a jury's declaration that the defendant is guilty. Searle calls this type of statement 'assertive declarations' because, like assertions, they are assessable as true or false and, like declarations, they bring about a state of affairs by being uttered in the appropriate circumstances. After looking at the replay, we can say that the umpire made the wrong call even though, for the purposes of the game, the runner really is out. It is senseless, by contrast, to say that the judge got it wrong when he said "Court is adjourned."[19]

The key feature of assertive-declarations is that there are independent criteria (that is, independent of the utterance) by which the statement can be assessed, and the speaker is trying to say that those criteria have been satisfied. The reason we give declarative power to certain assertions is that we need a final authority to give an answer once and for all so that, e.g., the game can proceed. We empower a jury or an umpire to make difficult decisions that need to be made. However, the umpires and jury members are aware that their claims must ideally match the facts. What counts as an out in baseball is defined independently of any umpire's utterance; and when the ump calls an out he must use those criteria. So it is possible for him to get it wrong. For normal, non-assertive, declarations, there are no such independent criteria. Court is adjourned when the judge says it is, end of story.

---

[19] He may say it in inappropriate circumstance, of course, but, assuming the circumstances are appropriate, his utterance cannot be called into question. An umpire can also say, "you're out" in the wrong situation; but even when the context is entirely appropriate, we can always question his calls.

The moral of this story is that, for (2) to assert a fact it must be an assertive-declarative. Some might say that it obviously is: (2) defines 'meter' as the length of stick S and also asserts that S is a meter long. The problem with this view should be clear. In order for (2) to be an assertive-declarative, there must be independent criteria for what counts as a meter. But prior to the initial utterance by the scientists in the 1790's, there weren't such criteria.[20] That was the purpose of the utterance, to establish criteria for the use of the word 'meter.' To put the point as simply as possible, the initial utterance of a sentence like (2) could not state a fact about the length of the bar because there were no criteria for what counted as a meter.

Another way to see this point is as follows: If (2) did express a fact then it would make sense to ask, "how do you know" and it would also make sense to say that it could have been wrong. But of course the scientists could not have been wrong since there were no criteria prior to the utterance. Again, the umpire can be wrong about whether the pitch was a strike because there are criteria that are independent of the ump's utterance for what counts as a strike (there is a strike zone). But the judge cannot be wrong about whether court is adjourned because the only criterion for the adjournment of court is that the judge says that it is. Similarly, the existence of

---

[20] Again, since I am glossing over certain elements in the actual history, this is not entirely accurate. As stated above, a meter was supposed to be equivalent to one ten-millionth of the distance from the North Pole to the Equator. Because of an error in calculation, the actual bar is .2 mm too short. Thus there did exist initial criteria that were independent of the utterance that occurred later. But this just pushes the problem back one step. There were no independent criteria for precisely how long the meter should be. The length chosen was entirely contingent; it could have been one millionth of the distance between Paris and Cairo. The upshot is that it would have been senseless to say that the choice was the wrong choice; that somehow the scientists were wrong about how long a meter was. In addition, the fact is that, even when the error was discovered, the world did not change all of the meter sticks. This shows that, in practice, the bar did serve as the final authority and that, despite whatever earlier considerations there might have been about the length of the meter, the initial utterance of "This bar is one meter" was indeed a declaration superseding all other criteria.

independent criteria enables us to ask the umpire how he knows, e.g., whether the ball was really fair ("Did you really see it?" "Did you have the best angle?"); and the absence of such criteria make it nonsensical to ask the judge how he knows that court is adjourned.

Perhaps, one might argue, the utterance of (2) creates the fact while simultaneously asserting it. The idea here would be that (2), by defining 'meter' makes it the case that S is a meter long and, at the same time asserts that it is. Unfortunately this proposal is no better. It is true that all declaratives create a state of affairs by saying that that state of affairs exists. But the state of affairs created by an utterance of (2) is not that the bar is a meter long but that the word 'meter' refers to the length of the bar. The utterance does not establish that the bar has the length that it has. How could it? It would be very odd indeed if saying certain words could change the length of an object. The bar is the same length regardless of the utterance. A sincere and appropriate utterance of a sentence like (2) established that the word 'meter' would be used according to certain criteria; it did not establish that the bar was a certain length. [21]

One more point before we leave this issue. To repeat, Boghossian has said, "'meter' names that length and hence that (2) says that stick S is that length at t, and since it is that length at t, (2) is true." His view is, as I understand it, that since 'meter' refers to the length of S, it is obviously true that S is a meter long. But there is

---

[21] This is an important point that deserves to be elaborated upon. A declarative does establish a fact; it alters the world, in other words. "Court is adjourned" makes it the case that court is adjourned. Once the judge has said so, it is then a fact that court is adjourned. My point is that the fact created by (2) is not that S is a meter long but that the word 'meter' refers to the length of S. Following the utterance of (2) it is then a fact that 'meter' has that meaning.

something odd about this view. The phrase "that length" here refers to the length of S. Thus, we could rephrase Boghossian's point as follows:

(3) Statement (2) says that stick S is the length of stick S at t, and since it is the length of stick S at t, (2) is true.

Or, more succinctly, stick S is the length of stick S at t. This is, of course a tautology and, as such, is a necessary statement. But other claims about the length of objects are contingent. It is a contingent fact that an object has the length that it does. Specifically it is contingent that stick S is the length that it is. But "Stick S is the length of stick S" is not contingent.

The reason (2) is necessary is that stick S plays a special role in the practice of metric measurement. Since it is the standard, there can be no question about how long S is in meters. To say that an object is a meter is to say that it is the same length as S. When I know that my desk is a meter long, I know that it is the same length as S. This is a fact about my desk. But it is not a fact about an object that it is the same length as itself.

The passage from *Philosophical Investigations* that inspired this discussion occurs within the larger context of the need for external criteria. In order for a word to be used meaningfully there must be criteria for its use. In another passage Wittgenstein asks us to imagine the following: Someone is asked if she knows how tall she is. In reply she says, "Of course I know how tall I am, I'm this tall" while holding her hand on top of her head. Does this person know how tall she is? Not really. She hasn't placed herself within any sort of context; she hasn't said how tall

she is in meters or in inches or feet. All she has done is state the tautology that she as tall as herself. She hasn't said anything of content about her height. In order to say meaningful things about the length or height of an object, we need to relate the objects to set standards. In other words, we need to employ criteria.

In the sense in which we can use the criteria to say, about an object, that it satisfies the criteria, we cannot use the criteria to say that the criteria satisfy themselves. A claim that an object is a meter long is a claim that it is the same length as the standard meter bar. It means one thing to say that some random object X is a meter long (it means that it satisfies the criteria for being a meter long), quite another to say that the standard meter bar is a meter long. If the latter were a claim about the bar's length, it would be a claim that the bar is the same length as itself. But then we cannot speak of satisfaction here; we can only speak of satisfaction when there is the possibility of not satisfying the criteria. The standard meter bar cannot fail to satisfy the criteria for being a meter.

The Reporting use of (1)

A reporting use of a sentence such as,

(1a) This bar is one meter long.

is distinguished from a producing version by the mere fact that it is uttered after the initial, producing utterance. There are many examples of similar statements that report facts that have been previously established via utterances. Here are some:

(4) A touchdown is worth six points.

(5) The rook moves along ranks and files.

(6) When the king is in check and cannot move out of check without thereby moving into check again, the king is check mated and the match is over.

(7) This note is legal tender for all debts public and private

Just as there is a meaning-producing version of (1), we can imagine utterances of (4) through (7) which perform an analogous function. When the first football leagues were being established there was presumably a rules committee that met to establish such things as how many points a particular type of score is worth. The rules committee's use of (4) would have established that a touchdown is worth six points. Any subsequent utterance of (4) (spoken within the context of the football league) would be a report of this fact that was initially established by the rules committee.

Similarly, (5) and (6) report facts about the rules of chess which were established at some point in the past. It is perhaps less plausible to imagine a rules committee brought together for the purposes of ratifying a system of rules governing the game of chess. Perhaps chess was invented all at once by one person or a group of people but more likely the modern game is the most recent descendent of a very old game whose rules have probably evolved through the ages. The rules may have been implicit for many centuries and never officially, or even unofficially, codified. Nonetheless, (5) and (6) would still have been accurate reports about the rules.

A meaning-reporting utterance of (1a) can be understood on analogy with rule-reporting utterance of sentences (4), (5), or (6). Just as there is a use of (4) that creates

a fact (where (4) is a declaration), there is a use of (1a) that creates a fact. And just as a subsequent utterance of (4) can report a fact about the rules of football, so too there is a use of (1a) that reports the meaning of 'meter.' If this analogy stands up then the reporting uses of (1a) and (4) (and (5), (6) and (7)) should have similar epistemological properties.

Note that, in one sense, none of the statements (4) through (7) are empirical claims. The Federal Reserve did not do any research before it was decided that (7) would be written on all U.S. bills; nor did the inventors of chess conduct a series of scientific experiments to discover how a rook can move on the chessboard. In addition none of them can be empirically refuted in the way that we might refute the claim that Iraq possessed weapons of mass destruction in 2003. As Searle is apt to say, congress isn't going to form a commission to investigate the game of football in order to decide, once and for all, precisely how many points a touchdown is worth.

However, these statements are unlike declarations in that all of them report facts. Statement (5), for example, reports a fact about the role of the rook. It is a fact that the rook moves in that way. (7) reports the fact that the dollar bill can be used to pay debts. And it is a fact that a touchdown is worth six points. So it turns out that these statements are empirical after all. However, all of the statements including (1a) and (4) through (7) are unlike traditional empirical statements such as,

(8) Humans are the closest living relatives of Chimpanzees.

Once you know the rules of football, you know that a touchdown is worth six points, and once you know the rules of chess and what role the rook plays in the game, you

know that it moves along ranks and files. There is no need for further investigation. But, of course, even if you know what chimpanzees and humans are you do not know that the species are more closely related to each other than either is to any other species.

A reporter such as (1a) reports the meaning of the word 'meter.' Just like (4), (5) and (6) once you know the rules governing the usage of 'meter' and understand the role played by the standard meter bar, you know that (1a) is a true report. There is no need to measure the bar to make sure. Reporters are empirical because they do report facts. But these are not facts about the properties of the referents of the words; they are fact about the meanings of the words. Thus they are not revisable via any empirical information concerning the nature of the referents. Nothing that is learned about the standard meter bar could falsify the statement "The standard meter bar is one meter long."

This claim does not rest solely upon the analogies to the statements above. It can be shown that no matter how many investigations we launch to discover the "actual length" of the standard meter bar, nothing we discover in these investigation will demonstrate that the bar is not one meter long.

So far I have explained how it is possible for sentence (1) to perform a producing function. Now I will argue that it cannot be used to report a fact about the SMB's length. I will divide this task into two parts. First, I will argue that it is incoherent to speak of measuring the standard meter in meters. The argument for this will involve showing that no matter how many measuring ceremonies we engage in,

any evidence we gather from such ceremonies cannot count as evidence that the standard is not one meter long. This argument will sound verificationist to some but I do not want to rest my position upon such an argument. As we shall see, it is possible for one to accept my argument that it is meaningless to speak of measuring the SMB in meters but still hold that it is still quite meaningful to assert that it is (or is not) one meter long. One might argue that since it is possible for the bar to fluctuate to some minute degree, it is perfectly coherent to say that it is one meter long even if we cannot measure the change. Thus I will also consider this possibility.

Before we begin with these arguments, I need to say a few things about the purpose of the standard meter bar. Much of what I have to say will be obvious but it is important that we lay it all out explicitly for the arguments that follow. In order to have a universal system of measurement that can be used for various purposes, we need to make sure that everyone is using identical measuring devices. It will be of no use to have a system in which some people use a ruler that is shorter or longer than others. This would not be a coherent system. To make certain that everyone is adhering to the same standard, we need an external, objective method of checking various measuring devices. New measuring devices will be calibrated to the standard and old devices can be checked against it. The standard meter bar plays just this role. It is an object that can be used to check measuring devices to make certain that they agree in their measurements. This insures that when two people living in different parts of the world measure two different objects and each discovers that the object is 2.8 meters long, the two objects will be of the same length.

In order to adequately fulfill this role, the object chosen as the standard must meet certain criteria. First, it must be available to scientific authorities from all over the world so that whenever someone wishes to check a measuring device, the standard is accessible. Second, ideally, it must be the sort of object whose length does not fluctuate. If the bar expanded and contracted on a regular basis in response to, e.g., changes in temperature, it would not be a useful standard; it could very easily give conflicting results such that measuring sticks across the world do not agree with one another or with the standard. Thus the bar should be made of a material that is not susceptible to such fluctuations or, if it is, it must be kept someplace where it is not exposed to the sort of environmental shifts that could cause the fluctuations. In practice, of course, this sort of quality control may be more or less easy to achieve; and I would be remiss if I ignored the possibility of a standard bar whose length fluctuated through time.

To begin, however, we will indulge in a little bit of science fiction. First of all, let's imagine that the bar that was chosen perfectly satisfies the criteria I listed. Specifically, it is such that it is not susceptible to fluctuations in mass or length due to environmental conditions. In such circumstances, once the bar is cut, it will remain at precisely the same length so long as it is not twisted out of shape or in some other way physically damaged. We can further imagine that it is under lock and key, guarded twenty-four hours a day so that such an event is unlikely. Thus the bar cannot change physically and will forever remain the same length.

It is easy to see that measuring this bar in such conditions would be an entirely pointless exercise. First of all, we know that the bar's length cannot change so what would be the point of measuring it? Even if we did decide to periodically go through a ceremony of measuring the bar, this ceremony could serve no purpose. If we hold a meter rule or other measuring device up to the bar and discover that the bar is shorter the rule, this can mean only one thing: the rule itself is too long. The only thing we could discover by such measurements is that our measuring devices are inaccurate. Any discrepancy between the measuring device and the meter bar can only be evidence that the device is flawed.

So, any possible evidence gathered in such a ceremony cannot count as evidence that the bar is not a meter long. No evidence we could gather would affect our evaluation of the bar's length and so it really is a stretch to call such a ceremony a measurement. We measure an object precisely in order to determine its length; or to make sure we are correct about its length. But we can learn nothing from the ceremony of holding a meter rule up to the standard meter because none of the evidence gathered is at all relevant to our assessment of the bar's length. Thus it is really nonsense to speak of measuring such a standard meter bar.

Remember that this conclusion is based on the assumption of a perfect bar, one whose length never fluctuates by even the smallest degree. In reality, of course, there is no way to ensure that the SMB will never expand or contract. All we can do is limit the fluctuations to a tolerable degree. Even so, measuring such a bar would be as pointless an activity as measuring a perfect bar.

Again, suppose we suspect that the bar has shrunk to some minute degree. To determine if our supposition is correct, we bring out an alternative measuring device and hold it up to the SMB. Now assuming that the two bars do not agree, what can we conclude? Specifically, can we conclude that the standard bar is not one meter long? Unfortunately this conclusion is only warranted so long as we assume that the alternative bar against which we are measuring the SMB has not also undergone a change in length. But what would warrant this assumption? The SMB has been chosen so as to minimize any fluctuations in length. We would, therefore, have more reason to suspect that the result of the measurement proves that the alternative has changed its length. Unless we have reason to think that the alternative bar is less susceptible to variations in length, the so-called "measurement" cannot prove that the SMB is shorter (or longer) than one meter long. But, as I stated above, the SMB is supposed to be precisely that bar that we have the most confidence in; that is least likely to undergo a change in length. Any alternative measuring device will be at least as vulnerable to alterations in length as the SMB.

The standard meter bar is the final arbiter. As such it is impossible to measure it with another bar. If we accepted the outcome of such a "measurement" as proof that the SMB had shortened, this would only show that we no longer accept the SMB as the final authority. But that would mean only that the particular bar we had chosen as the SMB would no longer be the SMB; the alternative bar would be, in such a case, the *de facto* SMB. Of course, the scientific community can change the standard (and, indeed, it has done so), and then the old bar that served as the standard can be

measured with the new standard. But this type of change would amount to altering (if only slightly) the meaning of 'meter.' 'Meter' would no longer refer to the standard meter bar but to some new standard; the SMB would have been stripped of its status. Thus we would be permitted say that it is (or is not) a meter long but this would be a different assertion than it would have been in Wittgenstein's time (when the SMB was the standard).[22]

So far I have shown that it doesn't make sense to speak of measuring the standard meter bar. Thus the utterance "the SMB is one meter long" cannot be taken to mean that the bar has been determined by measurement to be one meter long. I would be remiss, however, if I did not take into account a further complication. Consider the following argument: Surely it is possible, even probable, that the meter bar may shrink and/or expand through time, even if only to a minute degree. So suppose that this did happen. Even if it is true, as I have argued, that we would be unable to measure this change, the change has occurred and thus, one might argue, the bar is no longer a meter long.

I do not know how likely such a scenario is, but for the sake of this argument, we can take for granted that at some point in its history the standard meter bar expanded or contracted to some small degree. Does even the possibility that this might occur prove that my position is mistaken? Since we cannot dismiss the possibility, isn't it simply an empirical question whether or not the bar has or has not shrunk? If so, (1) could surely be used to state a fact about the SMB. To be sure, we

---

[22] See p.61ff below

may never know that it is a fact but nonetheless the sentence "The standard meter bar is a meter long" is a claim about the bar's length.

Despite the seeming plausibility of this argument, I think that it is ultimately mistaken. First off, we should remember that the standard meter bar has to be relatively resistant to fluctuations in its environment. To further this goal, the bar was to be kept in conditions that would minimize fluctuations in its length. Because of this we can be confident that any variations in the bar's length would be minimal.[23]

So let's make this scenario a little more precise and assume that the bar has expanded by .02 microns. According to the argument sketched above it would be true to say that the SMB is no longer one meter long but is now 1.00000002 meters long. But suppose that this alteration took place in the early part of the nineteenth century. At this time, scientists most likely could not measure to the accuracy of one-millionth of a meter. A difference of two hundred-millionths of a meter would have been a difference that did not make a difference. Even today such a difference is irrelevant for most purposes.

If two bars differ in length by a couple hundred-millionths of a meter, I am inclined to say that they are the same length. I suspect many people would agree and I do not believe that we are being sloppy or imprecise. Suppose I say that the distance between San Diego and Las Vegas is exactly 535.3 kilometers. Would I be proven

---

[23] Barring, of course, some sort of catastrophe. If, for example, the room that the bar is kept in became super-heated and the bar expanded, melted and cooled, it would be obvious that it was no longer the same length. Thus, we would no longer be able to use that particular bar as our standard. It would cease to be the SMB. Note that in such a case we would be unable to precisely determine the bar's original length. We would have to choose a new bar but would never know whether or not the new bar was exactly the same length as the old one.

incorrect if it was shown that my measurement was off by even several millimeters? For some purposes, small differences in length are simply irrelevant. If I say that Mars is currently 227,900,000 km from the sun, someone who says that the "real" distance is 227,900,000.01 km will not prove me wrong.

A small difference in length is only relevant depending upon the purpose of the measurement and the accuracy needed. In the eighteenth century a difference of .02 microns was undetectable and did not matter. Scientists had not yet developed the tools to study phenomena in which such distances become relevant. For this reason, if an object's length changed by .02 microns from time t1 to t2, it would have been completely accurate to say (in the 1700's) that its length had not changed.

The degree of accuracy that is relevant depends upon the domain of inquiry. Even in the twenty-first century, a difference of a few microns is simply irrelevant for most purposes. In constructing a building, for example, the builders may use several steel beams that are 3.5 meters long. These beams may differ from one another by as much as a couple millimeters, but for the purposes of the task at hand, they are the same length. However, there are purposes for which a difference of 3 millimeters is relevant. In such a case, two objects that differ by 3mm are not the same length.

What I am saying, then, is that the meaning of the phrase "same length" depends upon the particular task we are engaged in. When our task is such that a difference of a few microns is irrelevant, then such a difference is irrelevant to whether two objects are the same length. This point becomes even more forceful when we are dealing with differences that are impossible to detect. Since distances

that are relevant at the atomic level were not measurable in the eighteenth century these distances did not matter to a determination of an objects length. It was impossible for anyone to know that two bars differed in length by .02 microns. Therefore, everyone at that time would have agreed that such bars were the same length. They were not wrong or even imprecise; they were completely accurate.

Let's return to the scenario in which the SMB has shrunk by .02 microns. Whether or not this means that it is no longer the same length depends upon the accuracy needed for our measurements. If all of our measurements do not need to be accurate to the hundred-millionth of a meter, or if we cannot even measure to the hundred-millionth of a meter, then this distance is irrelevant. A change in length of .02 microns, would, in this case, not indicate that the bar is no longer a meter long. It is still the same length it has always been, regardless of the microscopic change.

More generally, whether or not any difference in length, no matter how small or large, is relevant depends upon how accurate we need the measurement to be in a given domain. So the mere fact that the SMB has expanded or contracted by some miniscule degree does not yet tell us that it is not the same length as it used to be. Of course this point is too important and complex to defend adequately here. I have done enough if I have merely suggested its plausibility.[24]

---

[24] Many reject the view about contextual dependence I have been advancing. In his book *Ignorance* (Oxford: Oxford University Press, 1975), for example, Peter Unger says that our language is fraught with what he calls "absolute terms." As examples of such terms he offers 'flat' Unger argues that it is a condition of something being flat "that nothing could ever possibly be even the least bit flatter" and thus concludes, "in the most ordinary meaning of the words, that *at most hardly any* physical objects are *flat*." (*Ignorance,* p. 49). Limitations of length prevent me from fully responding to Unger's argument here. However, one might wonder why, if Unger is correct, we even have the expressions 'flatter' and 'flattest.' In order for one object to be flatter than another, isn't it a condition that both objects be flat to begin with? If it really is the case that if an object is flat then no other object

A standard meter whose length fluctuates to some small degree is a perfectly good standard so long as measurements do not have to be made to that degree of exactness. To put the point another way, if the SMB fluctuated by .5 microns this would only lead to imprecise measurements if we could measure .5 microns. If we couldn't measure on this microscopic level then our measurements would be as precise as needed (and as possible). A standard the fluctuated by this small amount would be, in this scenario, completely adequate.

As technology advances and there is an accompanying need for more and more precise measurements, a standard embodied in a bar that is vulnerable to even minute fluctuations in length will become inadequate. In such a case the standard meter bar may have to be abandoned for a more accurate standard. Indeed the standard for the meter has changed at least twice since its initial definition in the late eighteenth century. Currently 'meter' is defined as 1/299,792,458 the distance light travels in a vacuum in one second. Since the speed of light in a vacuum is a constant, this is a much more exacting standard. However, we shouldn't think that this implies that there was something inherently imprecise about the old standard. Relative to the measurements that were needed in the late nineteenth century, the standard meter bar was completely precise.

---

is in the least bit flatter, then flatness is not a matter of degree; either you have it or you don't. But surely flatness *is* a matter of degree: we say of some objects that they are 'pretty flat' or 'flat enough' of others that they are 'really flat' or 'flatter than the Kansas countryside.' One tabletop may be flatter than another even though both are flatter than Kansas and even though both exhibit peaks and valleys at the microscopic level. The surface of my desk really is flat even if on the microscopic level it isn't.

The existence of the new standard implies, of course, that we can now measure the bar that had served as the standard meter bar. And if we do so, we may discover that it is not exactly one meter long. This, however, would not be a refutation of Wittgenstein's point because when he asserted that we cannot say that the SMB is one meter long, the SMB was the standard. The word 'meter' meant something different in Wittgenstein's time than it does today.

Conclusion

Regardless of the merits of my specific claims about the standard meter example, the analysis of reporters that I have offered here can be expanded to include the so-called analytic statements. In as much as analytic statements are defined as those that are true in virtue of meaning, analytic statements do not exist. However, all (or almost all) of those statements that we have called analytic are actually reports about the meaning of words. They are neither true nor false reports about the referents of those words.

Take, for example, "All vixens are foxes." Traditionally we would have regarded this as an analytic truth and would cash this out as follows: the statement expresses a fact about vixens (that they are foxes) and it is true just in virtue of the meanings of the words 'vixen' and 'fox.' I have argued that this conception is inadequate. First saying that the statement is true in virtue of meaning is simplistic: it is a product of not recognizing the distinctions I have been making. Second, as a

claim about vixens, the statement lacks content because, since there cannot possibly be a vixen that is not a fox, it fails to distinguish among possibilities.

"All vixens are foxes" is a true report about the meaning of 'vixen.' It is a claim about meaning and can be revised if it is discovered that people no longer use the term according to these criteria. This would involve a change in meaning and thus would not amount to an empirical refutation of a claim about vixens and foxes.

Given that we have bought the special epistemic status for reporters by denying that they say anything about the world, one may wonder if they have any real epistemic value. What use are these statements if they don't tell us anything about the referents of the words? I am prepared to accept that reporters are of little value in providing us with substantive knowledge about the world. But this doesn't mean that they are of no value. They do, after all, express knowledge about how words are used. It is possible to violate linguistic conventions and Wittgenstein believed that philosophical conundrums resulted precisely because of the linguistic confusion that results from such violations. We don't have to agree with his negative view of philosophy to recognize that philosophical positions sometimes really are based on linguistic mistakes. As I indicated in chapter 1, Norman Malcolm showed us that we could reject the view that all a person sees when he looks at an object is a part of his own brain. We don't have to do any investigation to know that this view is false. We know it's false because we know that the meaning of 'see' implies that if we see anything when looking at, e.g. a tree, we don't see a part of our own brain. In fact

most of us have never seen any part of our own brain. So we can refute Russell's position just by considering the meaning of the verb 'see.'

Facts about meanings have at least this kind of value. At the very least they offer a tool, a means of assessing solutions to philosophical worries. To borrow a quote from Wittgenstein, we can evaluate wild conjectures with the quiet weighing of linguistic facts.

CHAPTER 3
An Analysis of External World Skepticism

Consider the following quote from Moore's "Proof of an External World,"

> I certainly did at the moment *know* that which I expressed by the combination of certain gestures with saying the words 'There is one hand and here is another.' I *knew* that there was one hand in the place indicated by combining a certain gesture with my first utterance of 'here' and that there was another in the different place indicated by combining a certain gesture with my second utterance of 'here.' How absurd it would be to suggest that I did not know it, but only believed it, and that perhaps it was not the case! You might as well suggest that I do not know that I am standing up and talking -- that perhaps after all I'm not, and that it's not quite certain that I am! (Moore 1959, 146-147)

In this passage Moore seems to be saying something quite banal and un-extraordinary. He is merely asserting that when he says "Here is one hand" he knows that there is a hand at the place indicated. Why is this interesting? It is philosophically provocative for two quite different reasons. First, it involves a complete rejection of philosophical skepticism that was, at the time Moore was writing (as well as at the present time), commonly thought to be a legitimate and open issue. If Moore is correct and he does know what he claims, then there must be something fundamentally flawed with the skeptical arguments that have dominated philosophy since at least Descartes. What Moore is in effect saying is that any argument that concludes that we cannot have knowledge of the external world is just a bad argument.

The second reason this passage is interesting is far less obvious. In his *On Certainty*, Wittgenstein undertook an extensive investigation of Moore's claims and

the issue of skepticism in general. Wittgenstein had much to say about the subject, and it is not easy to decipher all of his comments. However, it is clear that he felt that Moore had made a fundamental mistake when he declared that he *knew* that he had two hands. The mistake was not the one which someone defending skepticism would point out; namely that Moore cannot know that he has two hands unless he can also prove that he is not dreaming, having an hallucination, or is otherwise deceived. The mistake Wittgenstein points out involves the claim that Moore's assertion that he knew that there were two hands, involves a misuse of the word 'know.' The word 'know' simply does not mean what Moore wants it to mean

Moore's attempt to put to rest philosophical claims to the effect that we cannot have certain knowledge of the external world was a failure. However, the failure indirectly revealed the key to deflating philosophical skepticism

I will identify and focus on two key aspects to Moore's argument: (1) If you can prove that any object exists (e.g., a dog) exists, then you have proven that the external world exists. This amounts to a rejection of the verificationist theory of meaning espoused by the positivists and is a manifestation of Moore's commitment to common sense and ordinary language. (2) I can know something I cannot prove to be true. I know that I am not dreaming even though I cannot prove it. This is a rejection of the traditional conception of knowledge prevalent since at least Descartes. There is, of course, much more that is important in the papers I shall discuss, but for the purposes of this paper these are the relevant issues.

Moore begins "Proof of an External World" with the following quotation from the preface of the *Critique of Pure Reason*:

> It still remains a scandal to philosophy . . . that the existence of things outside of us . . . must be accepted merely on faith, and that, if anyone thinks good to doubt their existence, we are unable to counter his doubts by any satisfactory proof.[25]

It is interesting to note that Kant did not say that it is a scandal that we must accept the heliocentric theory on faith, that we are unable to provide a satisfactory proof that the sun is at the center of the solar system. But if we cannot prove that there are things outside of us, then we can hardly prove that the sun is at the center of the solar system. Since the sun is certainly an object that is outside of us, if there are no objects outside of us, then the sun does not exist and so can hardly be said to be at the center of the universe. It seems to follow that if we cannot prove that things outside of us exist, then we cannot prove the heliocentric theory of the solar system.

Certainly it would be a scandal for astronomy if, even at Kant's time, people had to accept the heliocentric theory on faith. It would be an equal scandal for astronomy in our time, if we had to accept the existence of planets beyond the orbit of Saturn on faith. And an even greater scandal for medicine if we were unable to prove that penicillin kills bacteria. Yet, if Kant is correct, and we cannot prove the existence of things outside of us, then it seems we can prove none of these scientific theories. So why did Kant not also mention the scandals in astronomy and the other sciences? Well, perhaps Kant was only interested in discussing the scandals in philosophy. But I think that there is another reason why Kant doesn't mention these other scandals. I

---

[25] B xxxix, note: Kemp Smith, p. 34

think that if you had asked him, in 1781, whether science had proved the heliocentric theory, he would have answered that he thought that science had proven it. Similarly, if you ask most philosophers today whether or not the theory of evolution has been proven or whether we must accept it on faith, they would say that they think that it has been proven. So why is it, then, that if philosophers agree that the heliocentric theory and the theory of evolution have been proven, they do not agree that we have proven the existence of the external world? I think that this question is important to keep in mind as we examine Moore's proof of an external world and Wittgenstein's response to this proof.

Of course the Logical Positivists had an answer to this question. Before I turn to a discussion of Moore's proof of the external world, I think it will be informative to look at their solution to the problem. This will be helpful because it will both help to illuminate what is ingenious in Moore's argument and also shed light on Wittgenstein's rejection of skepticism.

Rudolf Carnap thought that the question "Is there an external world?" is not legitimate, at least not as the skeptic understands it. Both possible answers, "Yes, there are external objects," and "No, external objects do not exist" are meaningless assertions. Carnap sharply distinguishes between what he calls 'internal' and 'external' questions. An internal question is one which is asked from within a system, an external one attempts to ask a question about the system from outside of the system. An internal question is one such as "Does Carnap have a wife?" It can be answered in the usual way, e.g., finding Carnap's home and looking for a wife, or by

perhaps simply asking him. Finding his wife or receiving an affirmative answer to the query would count as a confirming experience. The question "Does the external world exist?" cannot be answered in the usual way. There is no possible experience or range of experiences which could either confirm or deny the assertion that there is an external world. This is because any experience is fully compatible with the truth of both idealism and realism. The issue, as understood by the skeptic, simply cannot be decided.

Carnap says: "To recognize something as a real thing or event means to succeed in incorporating it into the system of things at a particular space-time position so that it fits together with the other things recognized as real, according to the rules of the framework." (Carnap 207). A system provides a framework within which questions have meaning; it gives rules for incorporating elements into the system. Carnap's assertion is that the language of physical objects is such a system. This language provides standards and rules for the incorporation of elements into its system, or, in other words, for deciding whether or not there is, e.g., a tenth planet (these are the internal questions). But if this is correct, we cannot ask, about the system, whether it exists (this is an external question). "To be real in the scientific sense is to be an element of the system; hence this concept cannot be meaningfully applied to the system itself" (Carnap 207). So if the skeptic's question is to have meaning, it cannot be understood as asking whether the external world exists. It would be meaningless to ask whether the system can be incorporated as an element of the system. Carnap concludes that if the skeptic's question is to have any meaning it

must be understood as a practical question: one which asks whether the system of physical objects is the one which we ought to employ. The only legitimate external questions are practical questions.

It is important to understand why Carnap thought that the question "Does the external world exist?" is meaningless. The rejection of the question ultimately rests upon the verificationist theory of meaning[26]. This theory is based upon the verificationist principle which is as follows:

(VP) In order for any proposition, P, to be meaningful, it must be possible for us to determine whether P is true or false.

The statement "The external world exists," is supposed to be unverifiable (understood as an empirical claim) for the reasons given by traditional empiricists such as Hume. Any possible experience is consistent with both this statement and its negation. There is no possible experience or range of experiences that could prove either that the external world exists or that it does not. So, according to VP, both the claim that it does exist and the claim that it does not, are meaningless.

Barry Stroud makes the point that the skeptic has as much reason to accept the truth of skepticism -- that the claim that there is an external world is both meaningful and unknowable -- as Carnap does for rejecting the claims as meaningless: "the skeptical philosopher has precisely the same reason for declaring the truth of scepticism-- all possible experience is equally compatible with the existence and with

---

[26] For an extended discussion of the verificationist principle and its use in anti-skeptical arguments see Barry Stroud's "Transcendatal Arguments" *The Journal of Philosophy* LXV (1968), no. 9: 241-256. My understanding of Carnap's motivations relies heavily on this paper.

the non-existence of the external world." (Stroud 1984, 179). Therefore, Carnap's position rests essentially upon the verificationist principle.

With this background we can now more fully comprehend one of the important aspects of Moore's argument. Moore's argument amounts to a rejection of the logical positivist's solution. He very carefully and deliberately showed that if you can prove that a dog, a rose, a house, a hippopotamus, or any object whatsoever, exists, then you have proven that the external world exists.

This amounts to a full rejection of Carnap's solution. Carnap's solution implies that we can have proven that abominable snowmen exist and yet have not proven that the external world exists. But Moore said that if we can prove that any objects exist, we have proven that the external world exists. Why does he think this? Consider the planets Neptune and Pluto. Neptune and Pluto are objects which are external to our minds. (Moore spends much time evaluating the phrase 'external to our minds' and decides that its meaning is best captured as follows: an object is external to our minds iff its existence is logically independent of the existence of mental states.) So if we have proven that Neptune and Pluto exist, we have proven that objects that are external to our minds exist. Since the external world is just the collection of objects (and their relations) which are external to our minds, if we prove that Neptune and Pluto exist, we have proven that there is an external world.

I suppose there are a couple of different reactions one might have to this argument and in fact many of us will be of two minds about it. Part of us might think that it is quite reasonable. As we saw earlier, there is at least something odd about

maintaining that we have proven the existence of the eighth and ninth planets and yet also claiming that we have not proven the existence of the external world. So we might be inclined to just accept Moore's proof and chuckle at philosophers who insist that we must accept the existence of objects external to our minds on faith or who claim that the question of the existence of the external world is a meaningless question. On the other hand, we also want to say that there is something deeply flawed with Moore's argument; it seems he has missed the whole point. In fact I believe there is some truth in both of these responses. Moore's argument is brilliant in its own way but it is also problematic. I think that the best way to see the importance of what Moore was saying is to go back and take a closer look at Carnap's solution. If Moore is correct, Carnap made an important but subtle mistake. We will wait to discuss the flaw in Moore's argument until we consider Wittgenstein's response to it.

Carnap's solution amounts to the claim that there is a sense of 'verify' according to which we can verify that Neptune exists and yet not verify that the external world exists. The problem is that by the standards of verification according to which the existence of the external world cannot be verified, the existence of Neptune cannot be verified either. The problem with the statement "the external world exists" is that is susceptible to all kinds of skeptical worries. I can only know that the external world exists if I can know that I am not dreaming, hallucinating, etc. Since I cannot prove that I am not thus deceived, I cannot verify that the external world exists. But the same is true of the statements "Neptune exists," "Pluto exists," "Alligators exist," and every empirical claim whatsoever. According to the skeptic, it is not possible for

me to verify any of these claims precisely because I cannot prove that I am not somehow deceived. So, if we accept the skeptic's standard of verification, there is no difference between the statement "The external world exists" and "The planet Neptune exists" as far as our ability to verify them is concerned. Carnap, however, claimed that we can verify the second but not the first.

We might put Carnap's point as follows: statements like "Neptune exists" can be verified according to our normal everyday standards of empirical investigation -- there are predictions and observations that can be made which confirm the existence of the eighth planet. "The external world exists" cannot be verified because every possible range of experiences is consistent with both its confirmation and its denial. Moore's point is that if we have verified that Neptune exists then we have thereby verified that an external world exists. I believe that Moore is quite right about this and his insight shows that Carnap is actually using a double standard of verification. We need to look at this more closely.

The interplay between Carnap's solution and Moore's argument reveals that there are two different standards of verification at play in philosophical discussions about skepticism and knowledge of the external world. (To distinguish the two senses, in the next few paragraphs, I will call the first sense "verify1" and use 'verify' when I am employing this sense. The second sense will be called "verify2" and I will use 'veriphy' when I am employing that sense). I will not go into precisely what the standards are according to which a statement is verified (verify1) but we can assume that they are the common everyday standards we employ when making claims like

"Bush's victory in Florida has been verified." A statement is veriphied (verify2), on the other hand, when all possible skeptical doubts surrounding it have been removed. So, using these new terms, we can understand the skeptic's claim to be that it is not possible to veriphy any empirical claim.

With this new terminology, we can also see that Carnap's solution is based upon a double standard. Suppose we grant the skeptic his claim and agree that no statement can be veriphied. Carnap is right to say that the existence of the external world cannot be veriphied. And if we understand the word 'verify' in VP to mean the same as 'veriphy,' it is also true that VP implies that "The external world exists" is meaningless. But the same is true of "Neptune exists." If the skeptic is correct, this statement is no more veriphiable than the first. So, according to VP, it should be meaningless as well.

Carnap's solution depends upon reading 'verify' as meaning 'verify1' when it is applied to most everyday empirical claims, and 'verify2' when it is applied to the statement "the external world exists." In normal, everyday speech, we do, of course, use 'verify' according to standards that are weaker than those which I gave for 'veriphy' above. We speak of a claim about an event in the war being independently verified, or of a candidate's victory in an election as being verified and we do not even blink. But all of these things that can be verified cannot be veriphied

So long as we don't distinguish these two senses, we can slip into thinking that there is something fundamentally different between a statement like "alligators exist" and "the external world exists." The difference is supposed to be this: the first can be

verified while the second cannot. But once we have distinguished the two senses we can see that in the sense in which "the external world exists" cannot be veriphied (verify2), "alligators exists" can also not be veriphied. If every possible range of experiences is consistent with both the truth and falsity of "The external world exists," then it also the case that every possible experience is consistent with the truth and with the falsity of "Alligators exist."

The upshot of all of this is as follows: if there is a standard of verification according to which "Neptune exists" can be verified, then, according to that same standard, it is possible to verify "The external word exists." And if there is a standard of verification according to which, "The external world exists" is unverifiable, then, according to this standard, "Neptune exists" is also unverifiable.

Moore did not make the argument I am making. He was not interested in responding directly to the logical positivists. I am here only using Carnap's solution as a foil to bring out one of the key features of Moore's argument, which I think is an important insight. His insight is that if we can prove or verify that any object (like a human hand or a planet) exists, then we have proven that the external world exists. So when Moore held up his hands and said, "Here is one hand and here's another," he had proven that the external world exists.

This leads to the second important aspect of Moore's proof. Moore claimed that he knew that the external world exists because he can provide a proof of its existence merely by showing that he has two hands. And he can be certain of this proof because he knows that his hand exists. However, he says, he cannot prove that

his hand exists. He readily admits that the skeptic will not be satisfied by his proof because the proof depends upon statements which themselves cannot be proven. Moore admits that he cannot provide a proof that his hand exists but that, nonetheless, he knows that it does. He cannot prove that they exist because he cannot prove that he is not sleeping or in some other way deceived. But even though he cannot prove that he has two hands, Moore insists that he knows that he does: "I can know things which I cannot prove; and among the things which I certainly did know . . . were the premises of my two proofs. I should say, therefore, that those, if any, who are dissatisfied with these proofs merely on the ground that I did not know their premises, have no good reason for their dissatisfaction." (Moore 1959, 150).

We now turn to Wittgenstein's response to Moore. *On Certainty* is a very dense work and there are a multitude of issues that Wittgenstein addresses. For the purposes of this paper I will focus on two of his claims which I think are the most relevant to his response to Moore and to skepticism. The most fundamental aspect of this response involves the claim that Moore has misused the word 'know' (and other such terms) in his discussions of skepticism. The second aspect I will discuss is very closely related to the first and in fact provides the basis for it. This is the claim that words and sentences get their meaning from the contexts in which they are employed. A statement only has meaning in certain contexts (it does not have meaning in all contexts) and we are led astray if we think that every sentence can be used in any context whatsoever. The key focus of the argument (as I will reconstruct it) is that Moore, as well as most philosophers, have not been appropriately aware of the

particular circumstances in which the word 'know' is meaningful. Thus Moore misuses 'know' because he attempts to use it in circumstances in which it does not have an established meaning. The statement "I know that here are two hands" does not have a determinate meaning in the context in which Moore uses it.

Very early in *On Certainty*, he says, "If e.g. someone says 'I don't know if there's a hand here' he might be told 'Look closer'. -- This possibility of satisfying oneself is part of the language-game. Is one of its essential features." (Wittgenstein 1969[27], 3). Someone might be told to look closer because we assume that if someone says that he doesn't know if there's a hand here, then there must be some reason, something preventing his seeing the hand or seeing that it is a hand. If he looks closer, he may finally see the hand and thus satisfy himself that there is a hand here. We might put Wittgenstein's point as follows: if someone says he doesn't know whether a proposition '*P*' is true, it is because he has not yet satisfied himself that it is true. The ability to satisfy oneself is an essential feature of the language-game involving the word 'know' and in fact, if the possibility of satisfying oneself is removed from a certain situation, then this language-game cannot be played.

"Now, can one enumerate what one knows (like Moore)? Straight off like that, I believe not. --For otherwise the expression 'I know' gets misused." (OC 6). In his "Certainty," Moore begins by listing several propositions he knows with certainty. Among them are, "I am standing up," "I have clothes on," "I am speaking in a loud voice." (Moore 1959, 227). So Moore has attempted to do just what Wittgenstein has

---

[27] Henceforth, I will use the abbreviation, OC, to refer to *On Certainty*.

said involves a misuse of the expression 'I know." Why should we think that this involves a misuse of this expression? We have already seen that Wittgenstein thinks that it is an essential feature of the language game involving 'I know' that the person using the expression be able to satisfy oneself. And we have also seen that Moore admits that he cannot prove, when he is standing in front of his audience, the statement "Here's one hand and here's another." If Moore cannot prove it, then can he satisfy himself that here is one hand? Can he, and does he, satisfy himself that, as he is lecturing, he is standing up, wearing clothes, and speaking in a loud voice?

"Now do I, in the course of my life, make sure I know that here is a hand -- my own hand, that is?" (OC 9). There seems to be something odd about claiming that, in the course of my life that I make sure (that I satisfy myself) that I know that I have two hands. I have never even doubted that I have two hands. And, as Moore pointed out in "Proof of an External World," I cannot prove that I have two hands. But there is even something odd about saying that Moore cannot prove that he has two hands. To say that he cannot prove it implies that it is a proposition that stands in need of a proof. But it is not clear that Moore's proposition are propositions which require proof. It doesn't really make sense to say that Moore could satisfy himself of these propositions because he doesn't even need to be satisfied. Nothing Moore could do would make him any more certain that he is standing etc. But if he cannot satisfy himself that he has two hands, then it does not make sense to say that he knows that he has two hands.

Wittgenstein makes it clear that he thinks it is nonsense to say that I know things which I don't need to be satisfied of and for which there is no means of

satisfying myself. "I know that a sick man is lying here? Nonsense! I am sitting at his bedside, I am looking attentively into his face. --So I don't know, then, that there is a sick man lying here? Neither the question not the assertion makes sense." (OC 10). When I am standing next to a sick man's bed, nothing could possibly make me more certain that there is a sick man there. I don't need to be satisfied that there is because I see him and, in fact, it is not clear what could count as satisfying myself that there is a sick man. Since the possibility of satisfying oneself is an essential feature of the language game, then the language game just does not operate in this circumstance.

Moore claims that I can know something even if I cannot prove it. I can know that there are two hands but I cannot prove it: "How am I to prove 'Here's one hand and here's another'? I do not believe I can do it." (Moore 1959, 149). All the same, Moore says, he does know that there are two hands. Wittgenstein's point is that in order to know some proposition, I must be able to show that it is true, to demonstrate it, to satisfy myself and others that it is true. But can I show, or satisfy myself that "Here's one hand" is true? Moore said that I cannot do it. But we can imagine certain circumstances in which we might say that I have shown that the proposition is true: I am standing in a room, in front of an audience, holding an opaque box. I hold up the box and say, "Here is one hand." Since the box is opaque and not attached to a human being, my audience might doubt my claim. In this case, opening the box and showing the audience the hand would constitute showing them that 'here is one hand' is true.

Moore admits that such circumstances exist, in which we can prove the proposition that here is one hand. But, he says, "I do not believe that any proof is

possible in all cases." (Moore 1959, 149). In a normal case, such as the one which Moore describes, in which a person is standing in front of an audience lecturing, it is not clear what could count as showing that here is one hand. Moore is standing, speaking, gesturing with his hands. In this instance, it is not at all clear what could possibly count as showing that Moore has two hands, is standing, speaking, etc.

Moore is wrong to say that he cannot prove that he has two hands. But he is not wrong because he *can* prove it but because it doesn't make sense to say either that he could or could not prove it. Similarly, he is wrong to say that he knows that here is one hand. He is wrong because (in normal circumstances) it doesn't make sense to say "I know that I have two hands." It doesn't make sense because an essential part of the language game, the possibility of satisfying oneself, is absent in these circumstances. In order for one to be in a position to say that he knows a proposition, he must be able to determine, show, demonstrate, etc. that the proposition is true.

It is very important to realize that Wittgenstein is not saying that "I know that I have two hands" is always meaningless. We can easily imagine circumstances in which it might have a perfectly clear meaning (e.g., a wounded soldier in a hospital bed might not be sure that he hasn't lost his hands -- see p.19 below). Wittgenstein is also not saying that, barring the sort of special circumstances of the soldier case, "I know that I have two hands" can never be used in a meaningful way. All he is claiming is that the particular use that Moore intends for 'know' in his argument is a misuse. Moore wants the "I know" to mean the same in "I know that I have two hands" as it does in "I know that Bush won the election" but it cannot because there is

an important difference between the contexts in which either would be uttered. In the circumstances surrounding Moore's utterance of the first statement, there is no possibility of satisfaction, while in the case of Al Gore's utterance of the second statement, such a possibility obviously exists. Wittgenstein is not asserting that "I know" can only be used in certain contexts or that it is meaningless in contexts that lack the possibility of satisfaction, he is only saying that in the statement "I know that I have two hands," 'I know' cannot mean what it does in "I know that Bush won."

In Moore's circumstances, "I know that I have two hands" might mean something, perhaps something like "The proposition that I have two hands stands fast for me" or "That I have two hands is just something that I take for granted." But it cannot mean "I have performed a check and satisfied myself that I have two hands." As such, simply asserting it cannot adequately answer the skeptics challenge.

The second key feature of Wittgenstein's discussion in *On Certainty* is the claim that the circumstances surrounding a proposition, in some sense determine the meaning of that proposition. (At this point this will sound very vague and unclear. It is in fact very closely connected to the previous point about the applicability of the word 'know.' Later, in the final section of the paper I will explain exactly how the two points connect). In 464 he says, "My difficulty can be shewn like this: I am sitting, talking to a friend, Suddenly I say: 'I knew all along that you were so-and-so.' Is that really just a superfluous, though true remark? I feel as if these words were like 'Good Morning' said to someone in the middle of a conversation." Suppose a father dresses up like Santa Claus on Christmas Eve and makes certain that his

children see him.   Perhaps, in this case, his daughter might say to him "I knew all along that you were in the Santa outfit."   However, in the case Wittgenstein describes, the assertion seems out of place; in the context, we have no clear idea of what is being asserted.   The child is asserting that she was not fooled, that she knew that it was her dad.   But if I am talking about baseball to a friend and then say, "You can't fool me.   I knew it was you I was talking to all along," it is not clear what I am asserting.

Earlier Wittgenstein imagined a different example:

I go to the doctor, shew him my hand and say "This is a hand, not . . .; I've injured it, etc., etc."   Am I only giving him a piece of superfluous information?   For example, mightn't one say: supposing the words 'This is a hand' were a piece of information - - how could you bank on his understanding this information.   Indeed, if it is open to doubt 'whether this is a hand', why isn't it also open to doubt whether I am a human being who is informing the doctor of this? -- But on the other hand one can imagine cases -- even if they are very rare ones -- where this declaration is not superfluous, or is only superfluous but not absurd. (OC 460).

If I were to say such a thing to a doctor, the doctor might wonder if I am completely sane.   The problem is that it is not clear what I would be asserting if I were to show my doctor a hand and say "This is a hand."   Wittgenstein argues that, in fact, this proposition, said in this context, simply does not have a determinate meaning: "Just as the words 'I am here' have a meaning only in certain contexts, and not when I say them to someone who is sitting in front of me and sees me clearly, -- and not because they are superfluous, but because their meaning is not *determined* by the situation, yet stands in need of such determination."   (OC 348).   Propositions like "This is a hand" and "I am here" do have a determinate meaning in certain contexts

but it is a mistake to think that they must have the same meaning, or *any* meaning, in any context whatsoever.

This point is closely connected to another that Wittgenstein makes repeatedly. If a proposition does not make sense in isolation, then a doubt cannot make sense in isolation either. If, when asking a doctor to treat my hand I feel it necessary to remove his doubt that it is a hand, then why shouldn't I also remove his doubt that I am a human being or that I understand the meaning of the words I am uttering? In other words, if what I am trying to do when I say "This is a hand," is convince the doctor that it really is a hand and that he is not hallucinating or dreaming, etc., then why should I not also have to convince him that I am a human being? It just doesn't seem to make sense that someone might doubt that the object he is shown is a hand and yet not doubt that the object connected to the hand is a human being and that this being understands the words he utters.

"What would it be like to doubt now whether I have two hands? Why can't I imagine it at all? What would I believe if I didn't believe that? So far I have no system at all within which this doubt might exist." (OC 247). If I were to doubt that I have two hands, would I also doubt that I have two arms? Why would I be certain that I have two arms but not that I have two hands? What would make me certain of the one but not of the other? A doubt cannot exist in isolation; you cannot isolate one proposition and doubt it but accept all others. Even if what we are doubting is the resurrection of Jesus Christ, we cannot simply doubt the proposition "Jesus rose from the dead." We also must doubt the truthfulness of people who have written about it,

we must doubt the veracity of Mark, Matthew, Luke, etc. We also doubt whether such a thing is possible; whether a person who is dead can come back to life; whether Jesus was a god etc., etc.

In the case of doubting the resurrection of Jesus, we can make sense of this doubt because, even if we do doubt all of these things, there are still things which we do not doubt (which are taken for granted). We don't have to doubt that a man named Jesus lived, that there was a Roman Empire, that people die, that people tell stories, etc. But in the case of doubting whether I have two hands, it is not clear what I would not doubt if I really did doubt that I had two hands. It is equally unclear what might possibly remove this doubt. Why would someone who doubted whether he had two hands be satisfied, for example, if he looked at his hands? Why, if he doubted the existence of his hands, would he not doubt his senses? Would he not also doubt that he knew the meaning of the word 'hand'? But if he doubts the meanings of his words can he really be said to doubt anything? If he doubts what 'hand' means, how can he be sure what he is doubting? Because, in normal circumstances, we are as sure that we have two hands as we are of anything, we cannot be certain what it would mean to doubt that we have two hands.

In normal circumstances people just do not doubt that they have two hands. If someone were to tell a doctor, "this is a hand" we should not regard it as superfluous information but we should say that we don't understand what is being asserted. We don't know what is being asserted and (because?) we don't know what doubt is being removed. Now, we can certainly imagine a case in which we would understand

someone who doubted that he had two hands. A soldier is sitting in a hospital in a body cast, he is paralyzed and cannot feel his body. Looking around he sees other soldiers who are missing hands, legs, feet. In this instance, the soldier might genuinely doubt whether he has two hands. He has grounds for his doubt: he remembers being in an explosion and feeling great pain in his right hand and arm, he sees other soldiers who have lost hands, etc. It is also clear what would satisfy him that he does indeed have two hands: his doubt will be removed when his bandages come off or when he is able to ask a doctor. Because of his special circumstances, there are things about which the soldier can be more certain than that he has two hands. However, in the normal case, there does not seem to be anything we can be more certain of than that we have two hands. We can imagine circumstances in which it makes sense for someone to say that he doubts that he has two hands, but this doesn't imply that it makes sense in every possible circumstance for someone to express this doubt. In normal circumstances, we would not understand what a person meant if he said "I doubt that I have two hands."

"My having two hands is, in normal circumstances, as certain as anything that I could produce in evidence for it." (OC 250). The skeptic is correct to point out that I do not have grounds for my belief that I have two hands. However, this does not imply that I can doubt it. There are some propositions that I can give no grounds for that are more certain than the proposition in question. "What reason have I, now, when I cannot see my toes, to assume that I have five toes on each foot? Is it right to say that my reason is that previous experience has always taught me so? Am I more

certain of previous experience than that I have ten toes? That previous experience may well be the *cause* of my present certitude; but is it its ground?" (OC 429). In order to say that some proposition, *P1*, is the ground of another, *P2* (such that knowledge of the former guarantees knowledge of the latter) we must know *P1* with more certainty than we know *P2*. But there is no such proposition that grounds my knowledge that I have ten toes. I am not more certain that I have had past experiences than I am that I have ten toes. The propositions that I have had past experience and that I have ten toes are on the same level, so to speak; I am as certain of the one as I am of the other. And so the fact that I have had experiences in the past of seeing my ten toes cannot serve as the ground of my certainty that I have ten toes.

The Last Word on Skepticism

So what, if anything, do these considerations imply for the validity of skepticism? Do Wittgenstein's arguments validate skepticism against the sort of attack launched by Moore, or do they undermine skepticism, perhaps providing ammunition for a successful attack from a different angle? An appropriate response to this question is somewhat complicated. The best way to begin is to take a look at a couple of different answers and evaluate them.

One possible response is to say that while Wittgenstein has shown that Moore uses the word 'know' inappropriately when he says that he knows that he has two hands, the rest of Moore's argument is untouched by the considerations Wittgenstein raises. We might think that Moore's more general claim is perfectly valid; if we can

prove that so much as two objects exist, we have thereby proven that there are objects external to our mind. The problem with Moore's proof, on this conception, is that he chose the wrong sort of proposition to take as the basis of his proof. It is nonsense for Moore to say that he knows that he has two hands, but it is not nonsense for someone to say that he knows that Neptune exists. Moore's general claim is correct, we can prove that the external world exists by proving that any given object exist. Thus a proof that Neptune exists is, *ipso facto*, a proof that the external world exists.

This response involves taking Wittgenstein to be expressing a very specific worry about Moore's use of the word 'know.' But this is a very narrow view of Wittgenstein's insight. Wittgenstein was not just admonishing Moore for his use of a particular word. He felt that Moore's misuse revealed something important about the nature of skepticism. Before we get to this, we need to understand why the response outlined in the previous paragraph does not answer the skeptic. When the skeptic demands proof of an external world, he is demanding more than just the standard of scientific proof; he is demanding that all possible doubt be removed. The skeptic should be perfectly willing to admit that by some standard of proof, we can prove that Neptune exists. What the skeptic is doing, in demanding a proof that the external world exists, is setting a higher standard. We might think of a sort of continuum of standards of proof; everyday non-scientific standards, rigorous scientific standards, etc. The skeptic is setting an absolute standard according to which something is only proven true when all possible doubts about its falsity have been eliminated. The proof that Neptune exists doesn't meet this standard and thus cannot count as a proof that the

external world exists. To return, for a moment, to the discussion of Carnap's solution from earlier, we can see that just as there are different senses of 'verify' corresponding to different standards, there are also different senses of 'proof.' According to some standard of proof, it is true that proof of the existence of extra terrestrials does constitute proof of an external world. But this is not the standard that the skeptic is concerned about.

A second way of looking at the implications of Wittgenstein's arguments on skepticism is that Wittgenstein has defended skepticism against a Moore-like attack. We might be inclined to conclude that Wittgenstein has shown that the skeptic is correct -- we cannot prove the existence of the external world hence we do not know that it exists. We cannot prove it because we cannot meet the standard described above -- we cannot remove all possible doubts. Simply holding up your hands and saying "here are some hands" does not meet this standard. It is possible to doubt that there are two hands just as it is possible to doubt any empirical claim whatsoever.

There are a couple of problems with this response: (1) It ignores everything that Wittgenstein has to say about doubt, and (2) The skeptic is in exactly the same situation as Moore concerning his use of 'know.' I will discuss each of these issues separately.

First, the issue of doubt. Wittgenstein shows that a doubt cannot exist in isolation. In order for me to doubt something, I must be certain of many other things. To see this we need to look at how the word 'doubt' is used in everyday speech. Imagine I said, "I doubt that George Bush won the election in Florida." This doubt

has a specific content. I am doubting whether Bush got as many votes in the state of Florida as Al Gore. When I express this doubt I am not also calling into question the existence of the state of Florida or of George W. Bush. In fact, in order to doubt whether Bush won, I must take these and many facts for granted; I must assume that an election was held, that people cast ballots, that there are a countable number of ballots cast, that George Bush was one of the candidates, etc. Now suppose I were a skeptic and believed that everything can be doubted. In this case I would not mean the same thing by "I doubt that Bush won the election." The skeptic doubts everything, including the very existence of the election. But if I doubt that there was an election, the content of my doubt that Bush won is not that I think Gore might have gotten more votes. Rather I doubt that Bush won because I doubt that there was an election. Therefore, in order to have the doubt with the specific content that involves suspecting that Bush might have gotten fewer votes, I must take for granted that an election was held.

The conclusion of this argument is that in order for me to doubt something there must be other things that I take for granted. So in order for skeptical doubts to get off the ground, some facts must be taken for granted. Well, the skeptic should certainly accept all of this and say that there are things that he takes for granted. Specifically he takes for granted that he exists and that he is having experiences. What is doubted is that any of these experiences are veridical. So, it may seem that the skeptic has answered this objection. However, Wittgenstein makes another important point about doubting: in order to doubt I must have grounds for doubting.

"But what about such a proposition as 'I know I have a brain'? Can I doubt it? Grounds for *doubt* are lacking! Everything speaks in its favor, nothing against it. Nevertheless, it is imaginable that my skull should turn out empty when it was operated on." (OC 2). Just because I can imagine that I don't have a brain, this does not mean that I can doubt it (or that it makes sense to say that I doubt it).

If I tell you that I have a tiger in my backyard you would have reason to doubt my truthfulness (knowing what you do about the size of tigers, the illegality of owning one as a pet, and the lack of wild tiger populations in Southern California). If I then take you to my house and show you my tiger and you still insist that you doubt that I have a tiger, I would either be dumbfounded or tell you that I do not know what you mean. Here is the tiger, its in my backyard, end of story. Once you have seen the tiger, all grounds for doubting its existence have vanished.

Wittgenstein is claiming that the criteria of application of the word 'doubt' require that when a person says "I doubt that x," he must have grounds for saying so. Additionally he is implicitly arguing that simply being able to imagine that some state of affairs might not obtain is not grounds to doubt that it does obtains. I think that this coincides with how we actually speak (at least in non-philosophical contexts). We count some things as evidence and grant some things as ground for doubt; but not all things. The fact that you can imagine that you are hallucinating when I show you the tiger in my backyard is not grounds to doubt the existence of my tiger; but before I show you the tiger, you do have ground for doubt.

This point can be made in a different way. Consider again the doubt that Bush won the election in Florida. By the criteria of use outlined above, in order for me to have this doubt, I must have grounds for believing that it may be false that Bush won. I might, perhaps, have reason to think that if all of the ballots were counted by hand, we would find that Gore had more votes. I might suspect that some republican operatives trashed a bunch of ballots in predominately democratic precincts. There are many factors that might give me ground for doubting the legitimacy of the outcome of the election. Suppose, however, that I claim that my reason for doubting that Bush won is that I suspect that I have been hallucinating for the past year or that I think that I might just be a brain in a vat. In this case, I think it is reasonable to say that my so-called doubt is not really about the election but about, perhaps, my conscious state. I don't really doubt that Bush won because I can give no specific grounds for thinking that the outcome of the election was illegitimate. Or, to go back to my tiger, if you say that your grounds for doubting the existence of him (once you have actually seen him) is that you think that you are having an elaborate hallucination, then your doubt is not really a doubt about whether I have a tiger. It is a doubt about whether this particular experience is veridical. Certainly a part of that doubt includes the claim that my tiger does not exist but it also includes the claim that you are not really in my backyard etc., and so it is not the same doubt as you might have had before you saw the tiger. So, the grounds that I give for my doubt are directly related to the content of the doubt. Without specific grounds, it just doesn't make sense to say that I have the specific doubt.

An important corollary of this claim is that when I have grounds for doubt, I have in mind some test that could be made which would either legitimize or negate my doubt. (I know that there is some experience I could have that would show whether or not the proposition that I doubted is true.) When I doubt the outcome of an election I know exactly what is needed to determine whether or not the outcome really is legitimate: the votes need to be recounted in a thorough manner. When you doubt the existence of my tiger, you know what it would take to alleviate your doubt.

This can be put as a challenge to the skeptic: What are the grounds for doubting the validity of my experiences and what possible experience could you have that would show whether or not the experiences are valid? Depending on how the skeptic wants to portray the nature of his doubt, he may have a response to this challenge. If the claim is that I can doubt the validity of every possible experience I could ever have, I think that the skeptic is in trouble. He might say that the grounds for the doubt are that we are sometimes deceived by our senses and so it is possible that we are always deceived. But this sort of doubt certainly undermines itself since if I really doubt whether any of my experiences are veridical, I must doubt the validity of the experiences which have shown that sometimes my senses are deceived. But if I doubt this, then I end up doubting the grounds for the doubt and then it looks as if the doubt can never get off of the ground. Furthermore, there cannot be any experience which could settle this doubt since any possible experience would itself be doubted. But without grounds for doubt or the possibility of resolving the doubt, there is no possibility for doubt.

The skeptic need not formulate his doubt in this way, however. He could say that all that he is claiming is that the proposition "The objects around me (the computer, desk, etc.) exist" can be doubted because there are possible experiences which would show that the proposition is not true. Imagine circumstances such as that in the film "The Matrix." A person takes a pill and wakes up in a vat and discovers that the entire world which he thought he lived in does not exist; he has been a brain in a vat his entire life. Here we are given conditions which might settle the doubt; an experience is described that would, if it occurred, determine that the world in which I think I live in does not exist.

The first thing to say about this formulation of the skeptical doubt is that the grounds for the doubt are extremely weak if not altogether non-existent. As we saw earlier, Wittgenstein claimed that the mere fact that it is possible that some proposition is false does not give us ground to doubt the truth of the proposition. But even supposing that there are grounds, this alleged doubt still runs into difficulty. Imagine that you really had a "Matrix" type experience. Would you really now believe that the world you have lived in for you entire life does not exist, or would you, rather, doubt whether this present experience is veridical? I think that anyone who had this sort of experience would naturally doubt the validity of the experience before accepting that the world does not exist. You would probably wonder who slipped you the pill or try to remember what you ate that is giving you such weird dreams. This sort of response is not only natural but completely logical. The experience described is not really grounds for doubting the existence of the world, but actually grounds for doubting the

validity of the experience. If you had such an experience then it would seem that anything is open to doubt; If I can't believe what I have seen my whole life, then what can I believe? The grounds given for the doubt actually undermine the doubt rather than really providing grounds.

Let's return now to the problem with the skeptic's use of 'know.' I have tried to argue that in saying that he knows that he has two hands, Moore cannot be using the word 'know' in its standard sense. The context in which he uses it is importantly different from the context in which it is normally employed. The skeptic's use of 'know' suffers from exactly the same problem. As the skeptic has posed the problem there simply is no way of satisfying ourselves that the world exists; there is no test that can be performed which would render a verdict on the question; any possible observation is compatible with both the existence and the non-existence of the external world. Since the possibility of performing some kind of check is absent in that context, the skeptic's use of the word 'know' cannot have the same significance as a normal everyday use of it.

In its most forceful form, philosophical skepticism claims that we cannot know anything because there is no way to check to make certain that all of our standards of evidence are ever actually satisfied. But it is this very aspect of skepticism, that which makes it most powerful and compelling, that ultimately undermines it. Imagine the following: a friend and I are walking through a field looking for monarch caterpillars on milkweed plants. I am looking at plants on one side of the path and she is looking on the other. After thirty minutes of fruitless searching my friend looks at me and,

pointing to a plant behind me, asks "Are there any on that plant there?" I say that there are not and she asks me if I am completely certain. "I am very certain, I just checked," I reply. At this point a normal person would be satisfied (if somewhat disappointed) but suppose my friend understands philosophical skepticism and decides to press the issue: "How can you be certain?" she asks "you don't really know, isn't it possible that you are hallucinating?" In response to this I might say that though it is quite hot and I am rather thirsty, I am certain that I am not hallucinating and was not hallucinating when I was looking at the plant. "But you can't really *know* that," she replies, "you might be manipulated by some evil demon, there might really be three caterpillars on that plant but you are systematically fooled and so cannot see them. All of your normal ways of acquiring information about the world could be radically misleading. There is no way for you to know that they are not, thus you do not know that there aren't any caterpillars on the plant!"

The problem with my friends reasoning is that, as she has posed the problem, there really is no way for me to check to make sure that my means of acquiring information about the world don't radically mislead me. But since there is no possibility for a check, it doesn't make sense to say either that I know or do not know whether I am systematically deceived. To repeat the point which is central to Wittgenstein's examination of skepticism, the possibility of satisfying oneself is an essential feature of the language game. Without that feature the word just does not have the same meaning. So when my skeptical friend says that I don't know that there aren't any bugs on the leaf, what she means should not be understood on analogy with

the claim that we don't know whether there is intelligent life on other planets. The second claim asserts that we do not have the requisite information, the first simply asserts the tautology that we cannot check that which we cannot check.[28]

So what should we say about skepticism then? Can we have knowledge about the external world? The response to this question (which Wittgenstein provides the basis for) is an excellent example of the type of response given by ordinary language philosophy analyses. When we ask the question, "Can we prove that the external world exists?" we need to consider why the question is posed to begin with. What makes philosophers ask such a question?

It doesn't make sense to say that I know that I have two hands, or that I know that Neptune exists (in the sense desired by the skeptic) because there is no possibility for satisfaction. Neither the proposition "I know that the external world exists" nor "I know that the external world does not exist" make sense because as philosophers have formulated the problem, there is no possibility for satisfaction (no check can be performed to determine the truth or falsity of the claims). So the OLP response to skepticism is neither a positive endorsement nor a rebuttal. The response is that the entire problem is ill-posed and that the only answer can be the dissolution of the problem.

---

[28] In this paper I have been speaking of different notions of proof, verification, etc. and I have suggested that we can think of a continuum of proof such that, as we move along the continuum we get higher and higher standards of proof until we ultimately get to the skeptics standard. While I think that this idea of a continuum of standards is a very useful way of representing the issue and may accurately describe the various ways that words like 'know', 'proof', and 'verify' are used in various areas of discourse, the picture is misleading in one important respect. It is not the case that the skeptic's standard really belongs at the far end of the continuum. Since it does not embody any criteria which actually might be satisfied it does not belong on the continuum at all.

Basic Assumptions

It is now time to identify the basic assumptions regarding Wittgenstein's response to skepticism and begin to outline a defense of these assumptions. Earlier I identified two important aspects of Wittgenstein's discussion of skepticism. They were, first, the claim that Moore had misused the word 'know.' The second feature was the claim that a statement gets its meaning from the circumstances of its utterance and that one and the same sentence can mean very different things depending on the circumstances surrounding each. In what follows I will briefly outline what is involved in defending these claims. I will also show how they are related.

The most fundamental aspect of Wittgenstein's discussion of skepticism is really the notion of misuse. In order to understand his position, we need to understand exactly what it means to say that a particular use of a word is a misuse. How, exactly do we identify a use as a misuse and what does it mean to say that it is a misuse? This problem can be seen very readily as a challenge to the ordinary language philosopher: "Look, if we take the whole of language use, all we have is a multitude of uses. Sometimes people say things like "I know that I have two hands." This is just one use of the word 'know.' Why should we set it aside from all of the other uses and claim that it is a misuse? (maybe it is the correct use and the others are abuses)?" To answer this question, we need be able to explain how the use picked out as a misuse differs in a fundamental way from most of the other uses.

The claim that a particular use is a misuse, then, depends upon identifying some essential feature of the meaning of the term and then claiming that, in this particular case, this feature is not present. Since this essential feature is not present, the statement which involves the misuse is meaningless (or lacks a determinate meaning). It is easy to misunderstand this claim. I am not saying that "I know that I have two hands" is, strictly speaking, senseless. We have seen that there are cases in which it might, in fact, mean that a check has been performed (the soldier case). Earlier, I also pointed out that the sentence might mean something, even when Moore says it; it just cannot mean that Moore has satisfied himself that he has two hands.

As Moore uses it, the word 'know' in "I know that I have two hands" is intended to have the same meaning as it does in "I know that Neptune exists." But in the second statement what is asserted is that I have checked and decided that there is adequate evidence that Neptune exists. In the circumstances surrounding Moore's utterance of the first statement, the possibility of a check is not present (we don't even know what it would mean for Moore to check) and so 'know' cannot signify the same thing that it does in the first sentence. So the claim is not that "I know that I have two hands" is just a meaningless string of words like "small green flobites stummple down the gumple garden" or "colorless green ideas sleep furiously." The claim we need to defend is that it is meaningless to assert that I know that I have two hands *in the same sense* in which I know that Neptune exists. The defense for this claim involves discussing how 'know' is used in various cases, pulling out the essential feature of

'know' common to all or most of the uses of the term and then showing that in the particular case under discussion, the feature is absent.

The second feature of Wittgenstein's analysis that I identified was what we might call the context-dependence of meaning: that the meaning of a sentence is dependent on the context of its utterance. We saw two different types of case in which this is apparent. First, we saw that "I know that I have two hands" will mean something different when uttered by the soldier than when uttered by Moore. In this first case the soldier is saying that he has checked to make sure while in the second, Moore can only be asserting that the proposition that he has hands stands fast for him (if, in fact, he is asserting anything at all). We also saw that the statement "I am here" is meaningless in many circumstances but in others it has a definite meaning.

This observation is a key aspect of much of Wittgenstein's later philosophy and it is in fact closely connected to the issue of misuse. Wittgenstein claims that the meaning of a statement is dependent on the circumstances surrounding its utterance. Philosophers are (or at least have been in the past) inclined to think that meaning is something that accompanies the word or statement. A word is just a series of lifeless marks but it is meaningful because there is something --the word's meaning -- that is associated with the word. This meaning is supposed to be context-independent (the meaning of a word doesn't change from context to context). Wittgenstein rejects this view and insists that it is a mistake to think that the meaning of a word (or statement) is something that accompanies the word. The meaning of a sentence depends upon the context in which it is uttered.

This is an important point and one that deserves more elaboration than I am able to give it here. The crucial connection to the claims of this chapter is that when Wittgenstein says that a use of a word is a misuse, he is really just making this point about the context-dependence of meaning. His point is that if the meaning of a word or sentence is context-dependent, then if some feature of the context surrounding most uses of a word ('know') is absent in a particular case, then the word cannot have the same meaning in this degenerate case as it does in the other cases when the relevant feature is present. I don't want to defend this rather strong version of meaning-holism (which could obviously be a book-length defense). Nor do I think that defending an elaborate account of meaning is necessary to establish the central claim of this dissertation: that paying careful attention to how philosophically significant terms are used in ordinary contexts can often yield important philosophical dividends. I do, of course, want to claim that context matters but debates about the relative importance of context and whether a context insensitive semantics can be used to attack the kind of ordinary language methodology I am defending will have to be set aside here. (My hunch, and it is only a hunch, is that for certain projects the context-sensitivity claim will have relatively more importance than others. Perhaps for the sort of response to skepticism I have outlined in this chapter, the view is essential. But I doubt that context-dependence will be an issue in every case of an ordinary-language-based argument. I'll have a bit more to say about this in the concluding chapter.)

Luckily, however, much can be said that is relevant and important that does not depend on the outcome of such debates. In the next chapter I will be taking a closer

look at one aspect of the context-dependence claim and one way in which it has been criticized. The influential and important work by H. P. Grice on conversational implicature brings to the fore a difficult problem for the any philosopher who aims to use information about contexts of utterance of a term (or terms) to derive conclusions regarding the meaning of that term(s). Grice demonstrated that certain general features of the context of utterance[29] of a term are the result of what he calls 'conversational maxims': rules that govern, in a very general way, what is appropriate or inappropriate to say in the course of conversation. The upshot of Grice's work is that some of the conditions attached to the utterance of a term (including, of course, those conditions that ordinary language philosophers are apt to identify) are, perhaps, not conditions on the *meaningful* utterance on the term, but on its *appropriate* utterance. What this would imply concerning the argument of this chapter is that, assuming the condition I identified (that a check be possible) is not a condition on the meaningful utterance of the word 'know' but rather an assertability condition, then it is perfectly meaningful, though inappropriate, to say "I know" or assert "You don't know" even in cases in which the condition is absent.

The reader will not be surprised to learn I do not think this a successful argument (at least not as it concerns the term 'know' and the conditions I have identified). And so, to this issue I will now turn.

---

[29] I don't mean to be adopting or using a technical term here. By 'context of utterance of a term' I simply refer to the context surrounding a run-of-the-mill or ordinary use of the term.

CHAPTER 4
Ordinary Language Analysis and Pragmatics

In the previous chapter I developed a response to external world skepticism derived from Wittgenstein's observations in *On Certainty*. I now want to consider an important and serious objection to the methodology that I employed in the argument of that chapter. My argument centered on the claim that the word 'know' can only be meaningfully employed in contexts in which some sort of check is possible. In contexts where no such check is possible (such as the possibility that I am manipulated by an evil demon), the word cannot be meaningfully applied. This claim was in turn justified by appealing to various (alleged) facts about how words like 'know,' 'certain,' and 'prove' are used in normal, everyday contexts. For example, I elaborated and defended Wittgenstein's claim that, in most normal contexts (such as two friends enjoying a chess match), it would be odd and inappropriate to say "I know you're here." Similarly it would be strange for me to say, while I am standing at the bed of a sick man in the hospital, "I know there is a sick man lying here." I claimed that these sorts of statements are not just superfluous but literally meaningless. The fact is that we do not say such things and the fact that we don't use the word 'know' in contexts in which the possibility of performing a check is absent is an indication that this feature (the possibility of satisfying) is essential to the meaning of the word. My analysis then, was that the odd assertions are meaningless because, in the contexts

described, the possibility of satisfaction is not present and, since the word 'know' is not applicable in such situations, the utterances cannot assert anything.[30]

In response to the sort of argument sketched above, someone might say the following: "I certainly agree that, in the contexts you described, it would be odd for someone to utter the statements you consider. Indeed it would be wildly inappropriate to do so. However, this may not be because the statements are themselves meaningless but because the utterance of such a statement would violate an important conversational maxim such as 'Do not say things that are obvious' or 'Do not contribute information that is not informative.'" I have argued that the linguistic data I cited above are a consequence of the meaning of the word 'know' and its cognates but this response asserts that meaning has nothing to do with it

This response reveals an important gap in my argument as I presented it in the last chapter. I have tried to argue that the fact that we do not normally use the word 'know' in certain contexts (contexts in which the possibility of performing a check is absent) is an indication that the word is not meaningful in those contexts. Given the above rebuttal, we can now see that this move is too hasty. If there are alternate explanations for the relevant linguistic data, then I must explain why my preferred explanation is best.

This counterargument to my position is based upon a distinction first described by Paul Grice. In his work "Logic and Conversation," Grice articulated the notion of

---

[30] This is too strong. As I pointed out in the previous chapter, such utterances might be used to assert something, for example, "the claim that you are here stands fast for me." My point here, as it was there, is that these utterances (in the described contexts) cannot assert that the speaker has knowledge in the sense of having a justified true belief that rests on the speaker's possession of evidence.

conversational implicature or what has come to be called pragmatics. According to Grice speakers follow certain general rules, which he calls conversation maxims, that determine, in a very general way, what is or is not appropriate to say in various contexts.

In chapter 1 of *Logic and Conversation*, Grice frames his project as a means of analyzing a certain type of move characteristic of ordinary language philosophers. He begins by offering a schematic of the move he is interested in:

> [O]ne begins with the observation that a certain range of expressions E, in each of which is embedded a subordinate expression α --let us call this range $E(\alpha)$ –is such that its members would not be used in application to certain specimen situations, that their use would be odd or inappropriate or even would make no sense; one then suggests that the relevant feature of such situations is that they fail to satisfy some condition C (which may be negative in character); and one concludes that it is a characteristic of the concept expressed by α, a feature of the meaning or use of α, that $E(\alpha)$ is applicable only if C is satisfied." (Grice, 3).

This move should be very familiar by now and is of course, precisely the move that I employed in my argument against philosophical skepticism. Grice gives several more examples, a couple of which I will provide:

> Ryle maintained: "In their most ordinary employment 'voluntary' and 'involuntary' are used, with a few minor elasticities, as adjectives applying to actions which ought not to be done. We discuss whether someone's action was voluntary or not only when the action seems to have been his fault." From this he draws the conclusion that 'in ordinary use, then, it is absurd to discuss whether satisfactory, correct or admirable performances are voluntary or involuntary. (5).

> Malcolm accused Moore of having misused the word 'know' when he said that he knew that this was one human hand and that this was another human hand; Malcolm claimed, I think, that an essential part of the concept 'know' is the implication that an inquiry is under way. Wittgenstein made a similar protest against the philosopher's

application of the word 'know' to supposedly paradigmatic situations.[31] (5)

In his book, *The Significance of Philosophical Scepticism*, Barry Stroud employs a strategy derived from Grice's work on conversational maxims to criticize the anti-skeptical arguments of J. L. Austin. Stroud thoroughly demonstrates that Descartes' skeptical conclusion is founded upon the claim that the Dream Hypothesis must be eliminated in order for one to have genuine knowledge. In his second chapter Stroud confronts the inconvenient fact that, in our everyday lives, we do not usually insist that someone prove that he is not dreaming before we grant that he knows something. This observation of the divergence in the criteria for knowledge in everyday circumstances from Descartes' criteria forms the basis for the type of anti-skeptical argument I presented in the previous chapter (related arguments can be found, as I have indicated, in the work of Norman Malcolm and J.L. Austin). Drawing upon Austin's paper "Other Minds," Stroud extracts what he takes to be an anti-skeptical argument that has the same form as the arguments described by Grice. Austin, Stroud claims, argues that the Dream Hypothesis can only undermine my knowledge claim if there is some special reason to believe that I am dreaming. In arguing for this conclusion Austin appeals to facts about when and in what circumstances it is appropriate to question a purported piece of knowledge. When, Austin asks, is it appropriate to demand more justification for some purported piece of knowledge? And when is it inappropriate to make such demands? One of Austin's

---

[31] Malcolm's views about the use of 'know' were, as he acknowledges, based upon Wittgenstein's remarks in *On Certainty*. His argument is very similar to my own, but, to my knowledge, he did not offer an adequate reply to the sort of objection that I am considering in this chapter.

key points is that there is a limit to when it is appropriate to demand the removal of doubts. Demanding the elimination of a possibility is only appropriate if there is some reason to believe that the possibility might actually obtain

The move Stroud is interested in examining, and which we will be examining in this chapter, is the move from the inappropriateness of using certain expressions to the claim that the inappropriateness is somehow due to the meaning of the terms involved. That is, Stroud is concerned with Austin's implicit assumption that facts about how expressions are appropriately and inappropriately *used* imply facts about what these expressions *mean*. This is the same worry articulated by Grice in "Logic and Conversation". For simplicity's sake, I'll call this move the OLP-Maneuver.

The OLP-Maneuver is that step in an argument in which we take information about how and in what circumstances speakers use a word, series of words, phrase(s), etc. and use it to establish conclusions about the meaning of the word(s) or phrase, etc. In this chapter I will focus on the OLP-Maneuver as it is used in anti-skeptical arguments such as that of Austin, as well as my own from the previous chapter. The underlying question guiding the investigation is whether the fact that, in ordinary contexts, we don't require the elimination of the Dream Hypothesis is a consequence of the meaning of the word 'know' (my view) or a consequence of our following a conversational maxim to the effect that such a demand is, in most normal circumstances, is inappropriate. In answering this question we will be adressing the issue of whether the requirement that the dream hypothesis be eliminated is dictated by the concept of knowledge or whether it is not; whether it is, in fact the result of

careless philosophers ignoring what is actually said in normal contexts. Stroud argues that the requirement is genuine while acknowledging that, in normal contexts, it would be inappropriate to insist on its fulfillment.

Stroud attempts to diffuse Austin's criticisms (and indeed any anti-skeptical argument employing the OLP-Maneuver) by arguing that the inappropriateness of the question "Were you dreaming?" or the assertion, "I wasn't dreaming?" derives not from the fact that knowing that one is not dreaming is not a requirement of knowledge, but from that fact that, in most cases, there is no special reason to suppose that the person might be dreaming. Thus asking whether someone is dreaming is odd and inappropriate not because of the concept of knowledge but because of certain general pragmatic conditions concerning when it is acceptable to assert that someone has knowledge. We don't normally require that one eliminate the dream hypothesis before one asserts that one has knowledge, even though it *is* a condition of knowledge that the hypothesis be eliminated, because, in most cases, it is not expedient to do so.

Stroud thus insists that we must distinguish between conditions regarding when it is appropriate to assert that one has knowledge (assertability conditions) and conditions regarding when knowledge claims are true (truth conditions). Thus we might be inclined to think, as Stroud does, that the elimination of the Dream Hypothesis is not an assertability condition; it is perfectly appropriate to assert knowledge, in most cases, even when we haven't eliminated the Dream Hypothesis. Nonetheless, its elimination is a truth condition on every knowledge claim; no knowledge claim is true unless the Dream Hypothesis has been eliminated.

Though I will be focused on the Dream Hypothesis (and the issue of why speakers regard demanding its elimination as inappropriate), the importance of this discussion goes beyond the narrow issue of the truth of external-world skepticism. It extends to the issue of whether the ordinary language methodology I have described and implemented in this dissertation is justified. As Grice pointed out, the general form of the anti-skeptical argument I have articulated is one that underlies many Ordinary Language arguments. This current chapter provides a model for how the Pragmatics-based response to such OLP arguments can be overcome. Grice himself continued to assert the value of ordinary language analysis.[32] His point, as I understand it, was not to undermine any argument of the form he described (and certainly not ordinary language philosophy as a whole), but rather to point out a particular flaw which arguments of this form are prone to. There is no reason to suppose that any attempt at such an argument will result in error. Rather, we should take from his observations the lesson that the move from the linguistic data to the philosophical conclusion is not nearly as straightforward as some (including, perhaps, Austin) have made it seem. In the case of the argument against external-world

---

[32] In his 1958 paper, "Postwar Oxford Philosophy" (reprinted as chapter 10 of *Studies in the Ways of Words*), Grice said, "it is almost certainly (perhaps quite certainly) wrong to reject as false, absurd, or linguistically incorrect some class of ordinary statements if this rejection is based merely on philosophical grounds" (172). This view was echoed in his "Retrospective Epilogue" to *Studies in the Ways of Words*. In a section entitled "Philosophical Method and Ordinary Language," he compares the meta-philosophical views of Moore and Austin, submitting Moore (who, according to Grice, believed that common sense propositions are immune from error) to some harsh criticism and generally expressing sympathy for Austin: "Austin plainly viewed ordinary language as a wonderfully subtle and well-contrived instrument, one which is fashioned not for idle display but for serious (and nonserious) use. So while there is no guarantee of immunity from error, if one is minded to find error embedded in ordinary modes of speech, one had better have a solid reason behind one. That which must be assumed to hold (other things being equal) can be legitimately rejected only if there are grounds for saying that other things are not, or may not be, equal." (384)

skepticism that I presented in the last chapter, I believe the argument does indeed work and in this chapter I will explain why the Gricean move employed by Stroud is unsuccessful.

Stroud, on the other hand, expresses tremendous pessimism for the prospects of the OLP maneuver (especially as it is used to attack skepticism):

> How is it to be shown that that weaker requirement [Austin's special reasons requirement] or any other description of the way we actually speak and respond to the assertions of others, does in fact state a condition of knowledge, as opposed to a condition of appropriately or justifiably saying that one knows? As long as it is even intelligible to suppose that there is a logical gap between the fulfillment of the conditions for the appropriately making and assessing assertions of knowledge on the one hand, and the fulfillment of the conditions for truth of those assertions on the other, evidence from usage or from our practice will not establish a conclusion about the conditions of knowledge. (Stroud 1984, 64)

The condition at issue here is Austin's special reasons requirement. Austin's position, as Stroud describes it, is that a person's knowledge claim is not undermined by a doubt unless there is some special reason to suppose that the possibility raised by the doubt might actually obtain. (I'll return to Austin's view below).

I think we need not be so pessimistic as Stroud suggests. Suppose I claim that it is a condition of a person's having any kind of knowledge that the person be over 60 years old. We know immediately that this is neither a truth condition nor an assertability condition. Similarly we know immediately that the condition that a person must be a Frenchman in order to have knowledge is not a genuine truth condition. And we know immediately that the condition that p be true *is* a truth condition for knowledge that p (and not just an assertability condition). But how do

we know these things? How do we know that *p*'s being true is a condition on my knowing that *p*? We know these things because we know how to use the word 'know' and we know that considerations about the age and nationality of a person have no bearing on the applicability of the word 'know.' And we know that whether in fact the proposition allegedly known is true is directly relevant to whether someone can be said to know it.

Of course the fact that we know how to use 'know' (i.e., that we have a tacit understanding of when and how to use it) does not mean that we will be able to verbalize the meaning conditions of the word (or the truth conditions of knowledge claims). Recall Moore's distinction between knowing what an expression means versus knowing its analysis (a distinction I discussed in some detail in chapter 1). This is where ordinary language analysis (the empirical part) comes in. It is the task of the ordinary language philosopher to gather such linguistic evidence that will enable us to formulate a careful description of (perhaps just some of) the meaning conditions of expressions. The fact, if it is one (and again I must emphasize that such linguistic facts are falsifiable), that, in non-philosophical contexts, we do not require the elimination of the Dream Hypothesis is evidence that it's elimination is not a truth condition on knowledge claims. Similarly, the fact, if it were one, that we sometimes do require the elimination of the DH is evidence for the contrary conclusion that it's elimination is a truth condition. So I think evidence about usage is evidence for conclusions about truth conditions (and meaning conditions) though, of course, sometimes it may be mere *evidence*, insufficient to *establish* the preferred conclusion.

But I think we can often go a step further; in the case at hand, for example, we can evaluate competing explanations of the linguistic data. I will argue that the *best* explanation for the inappropriateness of demanding the removal Dream Hypothesis (DH) is that its removal is not required by the truth conditions of knowledge claims, which themselves are a product of the meaning of the word 'know.' My point will be that Stroud's explanation, to the effect that we regard demanding the DH's removal as inappropriate because of our adherence to a conversational maxim, cannot explain the linguistic data that Stroud himself agrees to. My alternative explanation, since it does account for all of the data, is a better explanation.

This, I believe, is a good model for any response to any Pragmatics-based criticism of an argument employing the OLP-Maneuver. In general the project should be to determine which explanation, the Gricean one or the OLP-Maneuver-based one, does a better job of accounting for the observed linguistic data. Of course we should not expect that the OLP-Maneuver will always win. But nor do we have any general presumption against the OLP-Maneuver.

Stroud On Conversational Maxims

To begin the analysis of Stroud's argument, we should first note that he accepts the linguistic data described by Austin. He also agrees with my observation that the elimination of the Dream Hypothesis is a condition that can never be met (more on this below). Presumably he would also agree with my further observation that, in ordinary contexts, we only use the word know (and semantically related terms)

when it is possible to provide evidence.[33] His argument, to repeat, is that this linguistic data is the result of our following a conversational maxim rather than the result of the meaning of the word 'know.'

Though I think that the question of whether Austin's analysis of knowledge is adequate (as well as the question of whether Stroud has accurately interpreted Austin) is completely irrelevant to my task in this chapter, it is nonetheless essential to have an understanding of Austin's views (at least as they are understood by Stroud). In his paper "Other Minds" Austin was interested in our use of words like 'know,' the question of when and in what circumstances it is appropriate to challenge someone's knowledge claim, and the nature of the demands we make when we ask someone to demonstrate that he knows what he claims to know. He also discussed terms like 'real' (as in the difference between a real tiger and an hallucination of a tiger or a stuffed toy tiger) and, correspondingly, our demonstrations, when such are demanded, to the effect that something is, for example, a real tiger. Here are some examples of Austin's observations:

> If you say 'That's not enough', then you must have in mind some more or less definite lack . . . If there is no definite lack which you are at least prepared to specify on being pressed, then it's silly (outrageous) just to go on saying 'That's not enough.' (Austin 1946, 84)

> Whenever I say I know, I am always liable to be taken to claim that, in a certain sense appropriate to the kind of statement (and to present intent and purposes), I am able to *prove* it. (85)

> The doubt or question 'But is it a *real* one?' has always (*must* have) a special basis, there must be some 'reason for suggesting' that it isn't

---

[33] Perhaps he wouldn't, I have no way of knowing. But whether or not he does, and whether or not my observation is accurate is irrelevant to the issue at hand in this chapter, namely whether the pragmatics/semantics distinction undermines any and all applications of the OLP Maneuver.

real, in the sense of some specific way, or limited number of ways, in which it is suggested that this experience or item may be phony. Sometimes (usually) the context makes it clear what the suggestion is: the goldfinch might be stuffed but there's no suggestion that it's a mirage, the oasis might be a mirage but there's no suggestion it might be stuffed. If the context doesn't make it clear, then I am entitled to ask 'How do you mean? Do you mean it may be stuffed or what? *What are you suggesting*?' The wile of the metaphysician consists in asking 'Is it a real table?' (a kind of object which has no obvious way of being phoney) and not specifying or limiting what may be wrong with it, so that I feel at a loss 'how to prove' it *is* a real one? (87)

If you are aware you may be mistaken, you ought not to say you know, just as, if you are aware that you may break your word, you have no business to promise. But of course, being aware that you may be mistaken doesn't mean being aware that you are a fallible human being: it means that you have some concrete reason to suppose that you may be mistaken in this case . . . It is naturally *always* possible ('humanly' possible') that I may be mistaken or may break my word, but that by itself is no bar against using the expressions 'I know' and I promise' as we do in fact use them. (98).

These special cases where doubts arise and require resolving, are contrasted with the normal cases which hold the field *unless* there is some special suggestion that deceit, &c., is involved, and deceit, moreover, of an intelligible kind in the circumstances, that is, of a kind that can be looked into because motive &c., is specially suggested. There is no suggestion that I never know what other people's emotions are, nor yet that in particular cases I might be wrong for no special reason or in no special way. (113).

From quotes such as these, Stroud extracts a very general thesis concerning when it is acceptable to raise doubts against some piece of purported knowledge. Austin insists that a doubt can be legitimately raised only when there is some special reason to raise it:

Austin is arguing that even if the way the experience or item might be 'phoney' has been specified, the doubt or question 'But is it a *real* one?' is relevant to the original knowledge-claim and must be answered only if there is some special reason for suggesting that the specified possibility might obtain. It is not simply that the critic of the

knowledge-claim must specify some way in which knowledge would not be present on the occasion in question; he must also have some reason for thinking or suggesting that the possible deficiency he has in mind might be present on that occasion. In the absence of such a reason—that is, in the normal or non-special case—knowing it is a real goldfinch, for example, is not in question. (Stroud 1984, 52).

Certainly there is a thread running through the quotes I provided which suggests the reading Stroud offers: Austin's concerns that a doubt must have "a special basis", his claim that being aware that you might be mistaken means have some "concrete reason", and his suggestion that doubts arise only when there is "some special suggestion" all support Stroud's reading. As I said, whether Stroud has interpreted Austin correctly is not my concern here and so I will stipulate that Austin did hold roughly the sort of view Stroud describes.

If Austin's view is correct, then the elimination of the DH cannot be a general requirement on knowledge. According to Austin, for the DH to be relevant to any of my knowledge claims, there must be some special reason to suppose that I am dreaming. Now clearly there may be instances in which there may be reason to suppose that I might actually be dreaming (e.g., early in the morning after being awoken suddenly). But Descartes' skeptical argument asks me to consider the possibility that I am dreaming even in the most mundane contexts: while I am reading my book, or taking a walk in the cool spring sunshine. Clearly in such circumstances as these (or any normal circumstance) there is no concrete reason to suppose that I might be dreaming. Stroud concludes, "If there must be some special reason for suggesting or suspecting that one is dreaming before that reason for doubt is even allowed as relevant in everyday life, the most that is true of the dream-possibility with

respect to our knowledge of the world is that it must be known not to obtain whenever there is some special reason to think it might obtain." (Stroud 1984, 53).

I think the best way to understand what is at issue when we ask whether I must know that I am not dreaming in order to know anything is to think about the ways in which my knowledge claims can be undermined.   Now obviously any knowledge claim is defeated when the proposition that I claim to know to be true is, in fact, not true. (My purported knowledge that Anchorage is the capital of Alaska is defeated by the truth of the proposition "Anchorage is not the capitol of Alaska.")   Let's call any such proposition, which undermines a knowledge claim merely by being true, a knowledge claim 'defeater.'   The way in which knowledge is undermined by doubts about the truth or falsity of some proposition (such as "I am dreaming") is importantly different in that it is the possibility that the proposition *migh*t be true (not that it actually *is* true) that undermines the claim.   Hence, in the discussion of whether skeptical doubts undermine knowledge claims we are dealing with possibilities.   For clarity, then, we'll call such possibilities 'underminers' of knowledge claims.

Importantly, there are two different types of underminers. First a possibility might undermine my knowledge claim when, if it the possibility were actual, my belief would be false.   For example, the possibility that my cat is an elaborate hallucination undermines my purported knowledge that my cat weighs twenty pounds since, if the possibility were true, I would not have a cat and thus would not know that my cat weighed twenty pounds. If I cannot eliminate the possibility, I cannot know that the belief is not false. Hence I would lack knowledge.

The second type of underminer involves situations in which even if the possibility is actual, my belief might still be true but I would not be justified in believing it. Suppose that I believe that it is raining and that it actually is raining. However, it is possible that I am merely asleep, dreaming that it is raining, and my belief that it is raining is based upon my dreamt experience. Here, though my belief is true, I lack knowledge because, if the possibility were actual (in other words, if I were actually dreaming), I would lack appropriate justification.

There seem to be many positions one might take concerning when a possibility is an underminer of a knowledge claim. Here is an extremely conservative account:

(C) Possibility P is an underminer of B's knowledge claim K (of the form 'B knows that q') iff either (i) if P were actually true, then q would be false or (ii) if P were actually true, then B would lack justification for his belief that q, and (in either case) B is aware that P may be true.

On this account, a possibility does not undermine my knowledge if I am unaware of the possibility. So, if my knowledge claim, "I know that Iraq has weapons of mass destruction" is not undermined by the possibility that Iraq does not have weapons of mass destruction if I am not aware that it is possible that Iraq does not.

An extremely liberal account of underminers would be as follows:

(D) Possibility P is an underminer of B's knowledge claim K (of the form 'B

knows that q') iff either (i) were P true, q would be false or (ii) were P true, B would lack justification for his belief that q.

On this account the possibility that I am dreaming is an underminer of nearly any knowledge claim since, if it is true that I am dreaming, I lack justification for all of my beliefs (except, perhaps, some of those about my own internal mental states).

Each of the above accounts implies a corresponding account of the (partial) truth conditions for knowledge claims. Each of the theories acknowledges that a knowledge claim is true only if there are no underminers of the claim. The accounts of underminers, then, embody truth conditions for knowledge claims. The conservative account of underminers corresponds to a liberal account of truth conditions in the sense that it allows for much more knowledge than many of the other accounts:

(C^{TC}) B knows that q only if there is no possibility P such that either

(i) were P true, q would be false or (ii) were P true, B would lack

justification for his belief that q, and B is aware that P is a possibility.

The liberal account of underminers, on the other hand corresponds to a set of conservative truth conditions for knowledge claims:

(D^{TC}) B knows that q only if there is no possibility P such that either

(i) were *P* true, *q* would be false or (ii) were *P* true, *B* would lack

justification for his belief that *q*.

I believe that D and D^TC capture the Cartesian account of knowledge that Stroud examines. On the Cartesian account, a doubt about the truth of some statement S undermines my knowledge claim in the case that if S is false then I lack justification for my belief.

Here is my own account of underminers:

(T) Possibility *P* is an underminer of *B*'s knowledge claim *K* (of the form '*B* knows that *q*') iff either (i) if *P* were true, *q* would be false or (ii) if *P* is true, *B* would lack justification for his belief that *q*, and (in either case) *P* is something about which it is in principle possible for *B* to have evidence for or against.

Based on the above reading of Austin, we can capture his implicit account of underminers as follows:

(A) Possibility *P* is an underminer of *B*'s knowledge claim *K* (of the form '*B* knows that *q*') iff, either (i) if *P* were true, *q* would be false or (ii) if *P*

were true, B would lack justification for his belief that q, and (in either case) there is some reason to believe that P might really obtain.[34]

On Austin's account a doubt about the truth of some statement S can only undermine a knowledge claim if there is a special reason to believe that S is in fact false. On my account a doubt undermines my knowledge claim only when there is, in principle, some way of resolving the doubt (i.e., some way of acquiring evidence as to whether that which is doubted is in fact true). Thus, Austin's position implies

(A^{tc}) B knows that q only if there is no possibility P such that if P

were true B would lack justification for his belief that q and there is some

special reason to believe that P is true.

The overall issue, then, is whether DH is an underminer.

Stroud begins the explication of his position by describing the following context: While at a party, you are asked by the host if you know if your friend John will be there. You reply that you know that John will be there because you just spoke to him and he said he was coming because a person who he is interested in speaking to will be there. You and your host know John to be someone who keeps his word, is a careful, sober driver, etc. Nevertheless, John does not show up at the party. As you are leaving the host says, "You should be more careful about what you claim you know. You said John would be here and he isn't. You didn't know any such thing!"

---

[34] The phrase 'there is some reason" makes this account ambiguous. Does the reason have to be something that A is aware of or just something that a reasonable person could be expected to be aware of?

Stroud says that when you said that you knew that John was coming, "There could hardly be more favourable grounds for claiming knowledge about something not currently under my direct observation." (Stroud 1984, 59). And he concedes that the hosts' admonition was completely outrageous and inappropriate. Nonetheless he maintains that what the host said was literally true. Stroud's analysis of the situation is worth quoting at some length:

> My response when asked whether I knew John would be at the party was justified, reasonable, appropriate, and perfectly proper. It is not open to the kind of attack the host tries to subject it to. But what is invulnerable to those absurd attacks is my act of saying something, and also perhaps my coming to believe or to accept something. My asserting it is beyond criticism even if what I assert is (of course unknown to me) not true. And the host's remark about the state of my knowledge is true even if his making it is outrageous, unreasonable, and unjustified. So even if we claim that a certain attempt to criticize a knowledge-claim is outrageous or unreasonable or would not be listened to in everyday life, we cannot immediately infer that the knowledge-claim does not suffer from the deficiency stated in the criticism, or that the person does nevertheless know what he claims to know. Whether that is so or not will depend on the nature and source of the outrageousness or inappropriateness in question. The inappropriately-asserted objection to the knowledge-claim might not be an outrageous violation of the conditions of knowledge, but rather an outrageous violation of the conditions for the appropriate assessment and acceptance of *assertions* of knowledge. (60)

Stroud thus maintains that the meteorite hypothesis is an underminer and, since Austin's account wouldn't count it as an underminer, Austin is wrong. Stroud's point here is a good one and I will admit that, with some caveats I will discuss below, I find his analysis of the inappropriateness making features of the context very reasonable.

The general conclusion Stroud is interested in is that facts as to inappropriateness do not establish anything about meaning conditions. Specifically,

he concludes that such linguistic data do not justify the view that the removal of the DH is not a general requirement on knowledge. (More generally: the removal of skeptical doubts, which would be inappropriate to bring up in most ordinary circumstances, is not a general requirement on knowledge.) But even if his argument against Austin's position is a good one (and I'll state my reservations about this below), this conclusion is, as yet, far too strong since there are other less restrictive accounts of underminers. My own account, for example, would allow that the meteor hypothesis is an underminer (assuming that there is, in principle, some way to determine whether or not a meteor will strike John). So this single example cannot establish the general conclusion that the inappropriateness of raising skeptical doubts is a result of conversational maxims rather than the meaning of 'know.'

To be fair to Stroud, I think he would acknowledge this. He would be happy with a weaker conclusion to the effect that information about usage cannot demonstrate that raising skeptical doubts violates the meaning conditions of 'know.' In the passage I quoted earlier he makes this clear: "As long as it is even intelligible to suppose that there is a logical gap between the fulfillment of the conditions for appropriately making and assessing assertions of knowledge on the one hand, and the fulfillment of the conditions for truth of those assertions on the other, evidence from usage or from our practice will not establish a conclusion about the conditions of knowledge." (64) His position is that since there is this gap, the facts about usage cannot establish that requirements involving the removal of skeptical hypotheses are not genuine requirements on knowledge. This conclusion rests, though, upon the

claim that there is no method for distinguishing conditions for the appropriate making and assessing of knowledge claims (I'll call these assertability conditions for short) and conditions for assessing the truth of those assertions (truth conditions). And, as I indicated (in the introduction to this chapter), I think that this claim is too strong. We can sometimes know that some condition is a mere assertability condition, we can sometimes know that a condition is a truth condition, and sometimes, though we do not know for sure, we can have a good idea that some condition is a truth condition rather than a mere assertability condition.

The meteorite/party example only concerned the meteorite hypothesis and, even if Stroud's analysis of the example is correct, our conclusion should be that the elimination of the meteorite hypothesis really is a truth condition, but it is not an assertability condition. Now this conclusion certainly doesn't imply that what goes for this skeptical hypothesis goes for any skeptical hypothesis. It does so only if the more general thesis that it is impossible to distinguish assertability conditions from truth conditions is true. But, as I'll show, not all skeptical hypotheses are created equal; there is very good reason to suppose that the elimination of the dream hypothesis is not a truth condition (in addition to not being an assertability condition). I suggest we take things more slowly and consider the different skeptical hypotheses separately.

First a few words about the meteorite hypothesis:

I will agree with Stroud that I must know that John has not been struck by a meteor in order to know that John will be at the party. In other words, I agree that

knowing that John has not been struck by a meteor is a condition on knowing that he will be at the party. And thus, if I don't know that John has not been struck by a meteor, then I don't know that he will be at the party. But the kind of doubt that has been raised in this example is insufficient to justify radical skepticism of the type that Stroud is concerned to justify. Perhaps it justifies a local skepticism concerning whether we ever really know that our friends will actually arrive at parties. A more generalized meteor hypothesis (to the effect that for all I know a meteor may soon strike anyone and thus prevent them from performing any expected action) might justify, at most, a skepticism concerning our knowledge of future events. But the claim that we never really know what might happen in the future, far from being a radical skeptical position, sounds an awful lot like a standard cliché.

It seems to me that I don't really know that I won't be struck by a meteor the next time I leave my house and it seems to me that Stroud's analysis to the effect that the host's question is inappropriate because it violated a conversational maxim is, in its essentials, correct. His explanation is correct except for the fact that he misidentifies the maxim that is violated. In saying that I did not know that John would be coming since I did not know that he had not been struck by a meteor, my host has violated a maxim requiring charitable interpretation rather than a maxim enjoining speakers not to demand the removal of scenarios that are difficult to remove. Though, as Stroud describes the example, I do say, "I know he will be coming," I think it is clear that what is meant is something like, "For all I know, he is coming" or "I expect him to come given that he told me he would be coming." The host's comment is

inappropriate because he has given a too-literal interpretation of what I had said. When I said, "I know he is coming" I of course did not intend to indicate that, beyond a shadow of a doubt, John would be there. After all, both I and my host are aware that any number of unfortunate events can befall an expected partygoer: His car might break down, he might get an urgent call from the office; he might have to resolve a family emergency; and yes, however unlikely the possibility, he might be struck by lightning or even a meteor. When I said, "I know that John will be here" I was merely indicating the fact that we had every reason to expect that he would arrive given the fact that he had earlier told me his intentions. By insisting that, because I didn't know that John hadn't been struck by lightning, I didn't know whether he would be here, my host was guilty of giving a literal interpretation of a comment that was clearly not meant in this literal way.

So I think that the general structure of Stroud's analysis is correct: the elimination of the meteorite hypothesis (I'll concede, for now) is a condition on knowledge of future events. It's also true that it is not an assertability condition for roughly the reasons given by Stroud (we follow a maxim that says that we should not demand the removal of hypotheses that we have no reason to believe might be actual). However, I would add that often when people say they know some future event will occur, they assume that they are not to be taken exceedingly literally; they are merely indicating their belief that the events will occur.

We must also admit, though, that Stroud has shown Austin's thesis to be unsuccessful. The requirement that there be a special reason is an assertability

condition but not a truth condition (and thus the account of underminers that corresponds to Austin's view is incorrect). Even though there is no reason to suppose that the meteorite hypothesis might actually obtain, it is nonetheless an underminer of my claim that I know that John will be here (assuming I mean this is the strict and literal way that Stroud does).

My account (T) of underminers, on the other hand, does allow that the meteorite hypothesis is an underminer. Since, at least in principle, it is possible to have evidence for or against the claim that a meteorite will strike John, (T) implies that it is an underminer of my knowledge claim. While it would obviously be difficult to acquire such evidence, I could, at least in principle, track the orbits of all near-earth objects and determine if any are likely to strike John as he drives to the party.

Extending Stroud's Analysis to The Dream Hypothesis

Though, given the caveat I mentioned, I find Stroud's analysis of the meteorite example quite insightful and reasonable, I don't think the analysis can be carried over to the Dream Hypothesis (or any other general hypothesis that would have the effect of undermining all knowledge). Even if we grant that the host's claim, while inappropriate, is nonetheless literally true, this still does not imply that the inappropriateness of every skeptical doubt can be dealt with in the same way. In particular it does not show that denying a knowledge claim due to a failure to eliminate the Dream Hypothesis is merely inappropriate (and yet literally true).

Of course the analysis of the meteorite example does suggest, and I think Stroud is inclined to accept, the following very general hypothesis, which attempts to account for all of the linguistic data concerning the inappropriateness of raising certain skeptical doubts:

Inappropriateness as the Result of Conversational Maxims Hypothesis (IRCM):

> Any linguistic data, to the effect that the raising of any of various skeptical doubts as a means of undermining a knowledge claim is inappropriate, is the result not of a violation of the meaning of the word 'know' or any of its cognates but rather of our following one or another of various conversational maxims that make the raising of such doubts inappropriate.

Stroud suggests that the same Gricean analysis applied to the meteorite example can be extended to all skeptical doubt, including the Dream Hypothesis. To be fair to Stroud, I don't think that he insists that the IRCM is always true for every doubt and every bit of linguistic evidence. However, as is indicated in the quote above, he does seem to believe that so long as we can distinguish between assertability conditions and truth conditions, the IRCM is at least always a possible explanation and, since it is, evidence based on linguistic usage will not establish facts about meaning or truth conditions.

But even this weaker claim is false. The IRCM is not a very good explanation of the relevant linguistic data as it concerns the Dream Hypothesis; it is not even a potential explanation. Let's slow down at take a more careful look at what the IRCM can and cannot explain.

We have a set of linguistic data, some of which I mentioned in the previous chapter, some of which has been provided by Austin, the upshot of which, as it

concerns the DH, is that, in ordinary contexts, we do not demand the removal of the DH and we would regard any such demand as entirely inappropriate. We have two explanations for this on the table: on one hand we have the IRCM and, on the other, my suggestion that the word 'know' cannot be meaningfully applied in contexts in which there is no potential for evidence one way or the other. The question we need now consider is which explanation best accounts for the data. I will argue that the IRCM, far from being a good candidate, does not even account for the data.

First, though, we need to be a bit more specific about how the IRCM accounts for the inappropriateness of raising the DH as a means of undermining knowledge. What conversational maxim do we follow that makes the raising of the DH inappropriate? Here is what Stroud says concerning this:

> On the conception I have in mind, the requirement that there must be some 'special reason' for thinking a certain possibility might obtain in order for that possibility to be relevant to a particular knowledge-claim would be seen as a requirement on the appropriate or reasonable assertion of knowledge, but not necessarily as a requirement on knowledge itself. In the absence of such a 'special reason', one might perhaps be fully justified in saying 'I know that *p*' even thought it is not true that one knows that *p*. (Stroud 1984, 63)

> We do not ordinarily insist on the dream-possibility's being ruled out unless there is some special reason to think it might obtain; the philosopher insists that it must always be known not to obtain in order to know anything about the world around us. (70).

This suggests that the maxim we follow is this:

Conversational maxim concerning the Appropriate Raising of Doubts (CARD):

> A hypothesis should not be raised as a doubt intended to undermine a knowledge claim unless there is some special reason to suppose that the hypothesis might actually obtain.

So, unless there is some special reason to believe that the Dream Hypothesis might actually obtain, it is inappropriate to raise it as a doubt. The CARD might explain, claims Stroud, why we do not normally raise the DH and why we would regard its being raised as inappropriate. As I indicated, this explanation will not suffice.

Before we proceed I want to make one point regarding this maxim. Unfortunately, as it is written, the maxim is somewhat vague. What, exactly, counts as reason to suppose that a possibility might obtain? Is there mere suggestion that it *is* a possibility enough? Probably not—this is too weak a requirement. Does having reason amount to having evidence (inconclusive thought it may be) that the possibility does obtain? Again, probably not—this requirement is far too strong.

The vagueness is not something that Stroud worries about, nor does Austin, and, for the most part, I won't either. But we should be aware of it and be prepared to offer a more careful formulation. In addition I want to make it explicit that the requirement that we have reason to suppose that the possibility might obtain should not be read as indicating that we must have evidence that it obtains. In the normal case, there is no reason to suppose that the meteorite hypothesis obtains. But suppose that I have recently watched a news report that warned of thousands of asteroids in the path of earth's orbit, potentially falling as meteorites. This information, while hardly representing evidence that my friend *will* be hit by a meteorite, is nonetheless reason enough to suppose that he *might* be. Raising the meteorite hypothesis and acknowledging that there is reason to think that it might obtain does not imply having

evidence that it actually does obtain. When my host raises the hypothesis we do not, as yet, have any evidence one way or the other, but, given the news report, we do have reason to suppose it might obtain.

Returning now to the Dream Hypothesis (DH), we should first note that we do sometimes demand, entirely appropriately, that a person prove that he was not dreaming. Stroud provides us with an example of such a context:

> Of course it is sometimes relevant to ask how or whether we know we are not dreaming . . . If I am lying half-awake in bed early in the morning after a late night and seem to hear someone calling my name from outside the window, I might not be sure whether there really is someone out there or I am only dreaming I hear the call. I do not know whether there is someone out there or not. (50).

In such a case, raising the possibility that I am dreaming would be entirely appropriate. In most cases, however, there is no special reason to demand that I prove that I am not dreaming.

Now I certainly agree with Stroud that the case described is one in which demanding proof that I am not dreaming is entirely appropriate. It is obvious that I sometimes believe that I am awake when I am merely dreaming. This obvious fact is in fact the source of the force of DH. However, it is also obvious that I sometimes realize that I am no longer asleep. We have all experienced times in which we have just awoken, often suddenly in the middle of the night, and we feel that we are not quite sure whether we are really awake. You remember hearing a strange noise, perhaps, and wonder, *did I really hear that or was it a dream?* Sometimes it can take some time for the feeling to pass, but after getting up, walking around, perhaps having a glass of water, maybe even finding the source of the offending sound, the feeling

fades and you realize, with certainty, that you are awake. Indeed, once the feeling passes the issue quickly leaves your mind and the question, "Am I just dreaming all of this?" fades away and ceases to be relevant. Stroud himself ends the discussion of his early morning scenario in which he is unsure whether he is dreaming with the following:

> When the alarm clock has sounded and I have reached out and turned it off and got out of bed and gone over to the window and opened the curtains and found my friend calling and gesticulating in the garden, there is no question at that point that I might be dreaming or that I should check to see whether I am dreaming before I can know that he really is there—even though I can truly say to him that I didn't know he was there a few minutes ago because I didn't know whether or not I was dreaming. (50).

The use of the past tense, "I *didn't* know he was there," and the suggestion that once I have gone to the window I don't have to check whether I am dreaming, together imply that now I really do know that my friend is there. And if knowing that I am not dreaming is a condition of my knowing that he is there, as Stroud wants to insist it surely is, then I also know that I am not dreaming. I certainly agree with this suggestion; it seems to me that, after getting up and going to the window, I do know that my friend is there. And this implies that I have eliminated the possibility that I am dreaming. But have I really? Have I really eliminated the Dream Hypothesis?

Even though, on the one hand, we are inclined to say that in this case I have eliminated the DH (and even Stroud himself appears to be pulled in this direction), on the other hand, it is also apt to appear, on a bit of reflection, that DH could not have been eliminated. After all, I might have dreamt the entire episode, including hearing the alarm go off and seeing my friend.

Upon reflection it becomes obvious that it is impossible to eliminate the DH. This impossibility is the entire foundation of radical external-world skepticism. If, in eliminating the possibility that I was dreaming by turning off the alarm, going to the window, throwing open the curtain, and recognizing my friend in the garden, I had thereby eliminating the Dream Hypothesis, Cartesian skepticism would be a much less attractive position than it is. Indeed Stroud spends a good portion of his book arguing just this point: "If we agree that he [Descartes] must know that he is not dreaming if he is to know in his particular case that he is sitting by the fire with a piece of paper in his hand, we must also agree that we can know nothing about the world around us." (23).

Stroud's reasoning is impeccable. He argues that even if I had a test that would be an unfailing guarantee that I am not dreaming, I still could never use the test, in any circumstance, to prove that I am awake. Again, it is worth quoting Stroud at length:

> The test would have to be something he could perform successfully, the state of affairs would have to be something he could know obtains . . . But how is he to know that the test has been performed successfully or that the state of affairs in question does in fact obtain? Anything one can experience in one's waking life can also be dreamt about; it is possible to dream that one has performed a certain test or that a certain state of affairs obtains . . . In order to know that this test has been performed or that the state of affairs in question obtains Descartes would therefore have to establish that he is not merely dreaming that he performed the test successfully . . . Obviously the particular test or state of affairs already in question cannot serve as a guarantee of its own authenticity, since it might have been merely dreamt, so some further test or state of affairs would be needed to indicate that the original test was actually performed and not merely dreamt, or that the state of affairs in question was actually ascertained to obtain and not just

dreamt to obtain. But this further test or state of affairs is subject to the same general condition in turn. (22)

This general condition thus causes an infinite regress of tests and so we will never reach a point at which we can know that our test has actually been performed and not merely dreamt. The implications for the present argument are obvious. Every experience one could have is subject to this condition: it is possible that the experienced events were merely dreamt and did not actually happen. Therefore, no matter what the context, no matter what the experience, if Stroud is correct, it is impossible to eliminate the DH.

But if this is the case, what has been eliminated in Stroud's late night example? If not the Dream Hypothesis, is there some other possibility that has been eliminated by my experience of turning off the clock and going to the window and seeing my friend? But what could this be? Stroud does not confront this question; he seems happy enough to assert that, "there is no question at that point that I might be dreaming". But if, as Stroud would have us believe, we cannot eliminate DH, the question demands an answer.

I think that the most honest thing for Stroud to say at this point would be that DH really has not been eliminated in this case. And of course Stroud's assertion that there is no question that I might be dreaming might be interpreted as something short of an admission that DH is eliminated. Perhaps Stroud only meant to convey that, once I have turned off the alarm, etc., it is no longer reasonable to raise DH since there is not longer any reason to suppose that it obtains. But is *this* the case? Have we really gone from a situation in which it is reasonable to raise DH to one in which it is

not? And if so, what is it about the experiences described—turning off the clock, going to the window, seeing the friend—that makes raising DH unreasonable?

If we are to take this route we must believe that the experiences of hearing the alarm, turning off the clock, and seeing my friend (call this set of experiences *E*) makes it the case that it is not longer reasonable to raise DH. [When a hypothesis is such that it is reasonable to raise it as a means of undermining a knowledge claim, we'll say that it is a live hypothesis.] But somehow *E* does so without actually eliminating DH. So, if DH cannot be eliminated by *E*, can it be obviated in some way? How would we know that some hypothesis has been obviated, in virtue of what has it been obviated? The set of experiences *E* is entirely consistent with DH: "anything one can experience in one's waking life can also be dreamt about." So, if DH was live prior to *E* and *E*, being entirely consistent with DH, cannot be relevant to determining whether DH is true or false, how can *E* make it the case that DH is no longer live?

So the problem with claiming that after *E* it is no longer reasonable to raise DH is that since, according to Stroud, any possible waking experience can also be dreamt, if it is reasonable to raise DH at one point in time (in the course of a given set of experiences), it must be reasonable to raise DH at any point in time (in the course of any possible set of experiences). So Stroud's options here are to either declare that it is never reasonable to raise DH (a choice he would definitely reject) or to conclude that it is always reasonable to raise DH. But the second choice prevents him from claiming that, "When the alarm clock has sounded and I have reached out and turned it

off and got out of bed and gone over to the window and opened the curtains and found my friend calling and gesticulating in the garden, there is no question at that point that I might be dreaming or that I should check to see whether I am dreaming before I can know that he really is there".

Nonetheless, as I indicated, there is the nagging suspicion that Stroud has said something that is true here. After all, once I've gone to the window and seen my friend, it would be rather loony of me to go on questioning whether I am actually dreaming. The set of experiences, *E*, has eliminated something, hasn't it?

I think it reasonable to believe that the Dream Hypothesis has a kind of double life. On the one hand it is a hypothesis that we've all confronted at times, often late at night, when we're in that eerie transitional state between sleep and full wakefulness. In this guise, DH is a possibility often confronted and eliminated, it would seem. On the other hand, it is a hypothesis that philosopher's have been telling us for centuries that we cannot know for certainty does not actually obtain. It seems there are two different forms the DH can take; or perhaps not two forms but rather two different modes by which the hypothesize can be raised:

DH$_I$: The Dream Hypothesis raised in such a way that it is logically impossible for there to be evidence sufficient to eliminate it.

DH$_P$: The Dream Hypothesis raised in such a way that it is granted that there may be evidence sufficient to eliminate it.[35]

---

[35] I have been tempted to think that these are actually two different hypotheses; that one hypothesis is such that it cannot be eliminated and the other is such that it can. But I am not entirely convinced that this is the right way to think about it and I think it is ultimately unimportant, for my purposes, whether we do so or not.

Consider the following example: Suppose someone is giving testimony in a murder trial concerning certain events that she witnessed in the early morning hours on the night of the murder. She claims that she awoke at around 4:30 am and heard, over the course of a several minutes, a series of loud thumping noises coming from the apartment in which the victim was later found. When she looked out her window, she saw the defendant running from the apartment. The witness admits, under cross-examination, that when she heard these noises and looked out her window was quite groggy since she had just been jolted awake.

Clearly in this case it would be appropriate to raise the possibility that the witness might be dreaming. There is a special reason to suppose that, in the circumstances described, it might actually obtain. But would it be appropriate to raise DH in such a way that it is impossible to eliminate, that is, to raise DH_I?

Suppose further, then, that the defense attorney raises the DH with the witness and that, when asked whether she was certain that she wasn't dreaming when she heard the noises, the witness replied that immediately after she started hearing the noises her telephone rang; it was her mother and she told her mother about the noises and that her mother herself heard a loud banging noise over the phone and that her mother's testimony about the conversation was included in a police report. This would be enough, I submit, to satisfactorily eliminate the possibility that the witness had been dreaming; the witness did not just dream the noises and seeing the defendant. But, though she has eliminated DH_P, she has not eliminated DH_I. All of the corroborating testimony, though seemingly relevant information, would not be enough

to eliminate DH_I and, assuming the lawyer really was raising DH_I as a means of discrediting the witness, he ought to continue to press her as to the ground of her certainty that she hadn't just dreamt the entire episode including the conversation with her mother and her mother's subsequent reports concerning that conversation. Of course, such line of questioning would not be tolerated by any sane judge.

The judge, if he was very tolerant, would presumably remind the attorney that to the extent that DH_I undermines this witness's testimony, it would undermine every witness's testimony. To the extent that one might have legitimate doubts as to whether the witness really was awake and thus did not simply dream the noises she reported hearing, those doubts have been put to rest by the corroboration that her conversation with her mother and her mother's testimony provided. A question arose as to whether she had been dreaming and evidence was provided to the effect that she was not. It is therefore inappropriate to continue to raise the possibility that she might have been dreaming.

And this is the key point: Though there is a special reason to suppose that the witness might have been dreaming, it is still inappropriate to raise DH_I as a means of undermining her claim that she knows that she heard the noises and saw the defendant. The IRCM and CARD suggest that the inappropriateness of raising DH_I, in most circumstances, is due to the fact that, in most circumstances, there is no special reason to suppose that it might obtain. But here we have a circumstance in which there is such a special reason and yet it is still inappropriate to raise DH_I. Therefore IRCM cannot be the correct explanation for this linguistic evidence.

On the other hand, my suggestion that demanding the elimination of DH_I involves a violation of the meaning of 'know' does explain the linguistic evidence. It is impossible to have evidence either for or against DH_I and thus it would be a misuse of the word 'know' to say either that I know that DH_I is false or that I don't know that it is false. Saying I don't know whether I am dreaming suggests that I stand in the same cognitive relation to the proposition "I am dreaming" as I do to the proposition "there is intelligent life on other planets." But this is not the case; I don't know that there is life on other planets because I lack evidence, but there is no evidence I could gather that would indicate one way or the other whether the Dream Hypothesis is true of me. So it is a violation of the meaning of the word 'know' to claim either that one knows that the Dream Hypothesis is true or to claim that one does not know that it is true. And this explains why it is always inappropriate to raise DH_I as a means of undermining someone's knowledge claim.

My argument concludes that it is never appropriate to raise DH_I and I believe that the linguistic evidence would support this contention. Any instance in which a person's knowledge claim is criticized via the appropriate raising of the possibility that he might be dreaming is an instance in which the possibility that is raised in the manner of DH_P rather than DH_I. I could, of course, be wrong about this. But this is just to emphasize that the ordinary language methodology that I have been defending has an in-eliminable empirical element. But assuming that my analysis of the linguistic data is correct, the only way for Stroud to maintain that IRCM and CARD explain the linguistic data would be to insist that there is never any special reason to

raise DH_I. (If there is never any special reason to suppose that DH_I might obtain, then our use of CARD would predict that we would always regard the raising of DH_I as inappropriate). But this, I think, is implausible. Every instance in which there is reason raise DH_P is an instance in which we have reason to suppose that I am dreaming. But this is exactly what DH_I supposes as well and thus, in all such instances, we also have reason to raise DH_I as well. This is perhaps something that I need to make clear: I am not saying that we can never have reason to suppose that DH might actually obtain in such a way that it would be impossible to know that it did and impossible to eliminate (and, of course I am certainly not saying that we know that DH does not obtain in this way). What I am saying is that, since we can have no evidence (none whatsoever, not even inconclusive evidence) that would serve to eliminate (or confirm, for that matter) the possibility, it is meaningless to say either that we know or that we don't know that DH_I is not true.

There is something else that Stroud could say in response to my argument. He could argue that not only is DH_I never eliminated but DH_P is never eliminated either. This, of course would conflict with what he says about his own example (see above). But, in any event, is this actually the case: is DH_P never eliminated? In as much as we can eliminate the meteorite hypothesis—and, at least in principle, we certainly can—I think there is every reason to suppose that we can eliminate DH_P. Once I check the orbits of all near-earth objects (a task that may be extremely difficult but is at least logically possible to complete), I have thereby eliminated the meteorite hypothesis. But I haven't eliminated the hypothesis that I have only dreamed that I completed this

task. Similarly, once I have silenced the alarm, gone to the window, and asked my friend what he wants, as Stroud says, there is no longer any question that I might be dreaming; I have eliminated this hypothesis. But, once again, I have not eliminated DH_I.

The fact that Stroud says that in this example there is no longer any question that I might be dreaming is significant. Just a few pages earlier he had taken Austin to task for saying the following: "There are recognized ways of distinguishing between dreaming and waking (how otherwise should we know how to use and to contrast the words?" About this suggestion Stroud says, "[Austin] seems content with the idea that there must be some such procedures or else we would not be able to use and to contrast the words 'dreaming' and 'waking' as we do. I find that particular claim dubious, or at least difficult to establish . . ." (Stroud 1984, 47). And, as I've indicated, Stroud spends a good amount of energy showing that the Dream Hypothesis cannot be eliminated. So why would Stroud profess skepticism about the possibility of having some means of distinguishing between dreaming and waking at one moment and then, in the next moment, apparently declare that it *is* possible for someone to eliminate the possibility that he was dreaming.

First we should note that Stroud does not distinguish between DH_I and DH_P. When he talks about the Cartesian hypothesis he talks of "the possibility that I am dreaming" and he gives us no reason to think that the possibility that is active in the early-morning waking-to-my-friend's-calls example is in any way different than the possibility that functions in Descartes' skeptical argument and that Stroud argues is

impossible to eliminate. So maybe Stroud just made a mistake in his description of the early-morning example or maybe I am reading too much into his suggestion that "there is no question at that point that I might be dreaming." If it is a mistake, it is an instructive one. I think what's going on is that Stroud is moving back and forth between different sets of conditions on knowledge claims. In describing the early-morning example, Stroud has subtly shifted from the Cartesian conditions to what I would call the everyday conditions (where eliminating the DH_I is not required). Only when using the everyday conditions and not the Cartesian conditions can Stroud imply that I know, now (after turning off the clock etc.), that I am not dreaming. And only when using the Cartesian conditions can he insist that the Dream Hypothesis can never be eliminated. This mistake is easy to make precisely because the everyday conditions do not contain the Cartesian requirement that DH_I be eliminated.

Conclusion

The purpose of this chapter, as it relates to the project of this dissertation as a whole, is to demonstrate that the Pragmatics-based arguments do not undermine every use of the OL-Maneuver. As I have indicated, Grice has shown (as has Stroud) that Ordinary Language Philosophers need to be careful; they ought to consider whether the evidence they gather might best be explained via some conversational maxim rather than some postulation about the meaning of the terms involved. But Grice's observations in no way provide a general injunction against the OLP-Maneuver.

The reason for this is very obvious: words are meaningful, they have meaning-conditions, and statements that use them have truth-conditions that are partly a product of these meaning-conditions. It is always open to critics of Ordinary Language Analysis to claim that a purported meaning-condition is actually only an assertability condition, or that some condition thought not to be a truth condition actually is but is not an assertabiltity condition. But first, as I said, we know that words have meanings and that some conditions really are a product of meaning. And second, fluent speakers will be aware of many meaning conditions. We know that being 60 years old is not a condition on the possession of knowledge. But how do we justify this claim except by pointing to linguistic evidence and our own proficiency with the term 'know'? However, the fact that we never demand that a potential knower prove that he is at least 60 is consistent with the following explanation: that a person is 60 *is* a truth condition on knowledge claims but is not an assertability condition. This supposition predicts the linguistic behavior that we observe, namely that we never disqualify a person's knowledge due to the fact that the person is under 60. Nevertheless it is an entirely hopeless hypothesis since we know that the condition has nothing whatsoever to do with knowledge.

This topic suggests one final response to the argument of this chapter: Stroud, or someone defending Stroud's position, might say that though I have shown that the suggested maxim cannot explain the linguistic data, the real conclusion is not that the general strategy (embodied in the IRCM) of explaining the linguistic data by virtue of conversational maxims is invalid but that we were wrong about that the relevant

maxim. Perhaps the maxim that explains the data is 'Don't raise an hypothesis as a means of undermining a knowledge claim unless it is possible (at least in principle) to gather evidence for or against the truth of the hypothesis.'

The hypothesis that we conform to such a maxim does predict the linguistic data that I have relied on in this chapter and in the previous one; so this explanation is not subject to the criticism that forced us to reject the explanation that referenced the CARD maxim. But it is a suspicious maneuver nonetheless. Never mind the fact that the maxim is, quite transparently, an ad hoc formulation, used merely to respond to the criticisms I have leveled in this chapter. It is just much more likely that the requirement we be in a position (at least potentially) to gather evidence for or against a proposition is a product of the meaning of the word 'know' rather than of a conversational maxim. The distinction between having evidence and being in the position to offer evidence, on the one hand, and not having evidence and not being in a position to mention relevant evidence, on the other, is precisely the distinction that the word 'know' and its cognates is supposed to capture. Suggesting that this requirement (being in a position to get evidence) is only a pragmatic consideration, and not a meaning requirement, does not allow us to capture this important distinction.

Once again I'll reiterate a point I have made elsewhere in this dissertation: I do not stand in the same cognitive relation to the proposition 'There is life on Mars' as I do to the proposition 'The Dream Hypothesis is true of me.' The distinction between my cognitive relation to the former and my relation to the latter is that, with the former, I currently lack evidence but may one day come to have evidence that decides

the issue, while with the latter I lack evidence because nothing could possible count as evidence. It seems to me that the distinction between my cognitive relation to the statements 'Mt. Everest is the highest peak on the planet' and 'There is life on Mars' is a real distinction that the phrases 'I know X' and 'I don't know X' are supposed to capture. But if we regard the elimination of the Dream Hypothesis as a genuine requirement on knowledge-claims, we lose this distinction (since, by this criteria, I know neither claim).

CHAPTER 5
Ordinary Language Analysis in Action:
An Examination of (some aspects of) the Mind/Body Problem

Thus far I have endeavored to explain how ordinary language analysis is supposed to work and to defend the method against some important and difficult objections. I do not claim to have defended it against every objection that has been raised, to do so would be an absurdly uninteresting and most likely pointless effort. However, I do believe that I have responded to the most serious objections and, in so doing, shown that the type of analysis typified by the work of Austin, Wittgenstein, Ryle and others is a well-founded and legitimate way to conduct philosophical inquiry and criticism. At this point, I think it would be helpful to see how the tools I have been discussing can be applied in a realm other than that of the epistemological problems that I have been pre-occupied with thus far. For this task I have chosen the mind/body problem; more specifically, the problem that qualia pose for psycho-physical reduction.

This chapter is divided into two parts. In the first part I attempt to identify the nature of the problem that qualia are alleged to pose for mind/brain reduction. I discuss the work of many anti-reductionists (such as Joseph Levine, Frank Jackson, and David Chalmers) much of which is dependent upon the existence of an explanatory gap held to exist between physical phenomena on the one hand and qualitative phenomena on the other. That this explanatory gap poses a special problem for the philosophy of mind depends on a certain view of what qualia are. According to

this view, qualia exist only in the mental realm, i.e., without minds there would be no qualia. But this kind of restricted understanding of qualia cannot hold up under scrutiny. A cursory analysis of just a limited range of our qualia talk (i.e., the way we use terms associated with qualia, terms like 'sensation' and 'feels') reveals both that many anti-reductionist arguments are based on a confused conception of qualitative phenomena and also that it is far from obvious that qualitative phenomena are restricted to the mental realm.

However, even if anti-reductionist philosophers have ill-advisedly limited the range of qualitative phenomena, the explanatory gap arguably does exist. In the second part of the paper, I consider how far towards an anti-materialistic conclusion the recognition of this gap might bring us. Much of what I say here is based upon Wittgenstein's observations about our use of sensation terms.

I don't want to leave the impression that the arguments of this chapter entirely rest upon an extended analysis of ordinary speech. As I have indicated, I don't believe that ordinary language analysis should be our sole means of philosophical inquiry. This chapter is intended to demonstrate how the ordinary language methodology can be used to augment and compliment traditional modes of argument.

What is the Problem? / What are Qualia?

So much has been written about qualia, and the problem that qualia pose for a materialist account of mind, that it is easy to get the impression that the mind/body problem just is the problem of explaining qualia. Many philosophers write as if the

mere existence of qualia poses a serious and insurmountable problem for materialism. Qualia are supposed to be mysterious and any proposed materialist explanation of them is doomed to failure. I think that much of the thinking about qualia is muddled and in this paper I will try to disentangle the issues so that we may think clearly about qualia and the consequences that qualia have for mind/brain reduction.

It has been pointed out by many philosophers, such as Saul Kripke, Joseph Levine, David Chalmers, and others, that we can imagine pain to exist without the correlated brain state. Depending on the version of the argument, this is supposed to lead to one problematic conclusion or the other. For Kripke, it shows that pain cannot be identical to a brain state because all true identity statements involving rigid designators are necessary. The fact that it is possible that there can be pain without the correlated brain state (C-fiber stimulation, for example) shows that the statement 'pain is identical to C-fiber stimulation' is not necessary and therefore, since 'pain' and 'C-fiber stimulation' are rigid designators, it cannot be true. Levine rejects this metaphysical conclusion as too strong but argues that the observations do show that there is an explanatory gap between brain states and the correlated phenomenal state.

I agree with the observation that we can always imagine pain without the correlated brain state and I even agree that this implies that there is an explanatory gap. However, I reject the move from here to the conclusion that materialism (or physicalism) must be false. We would be warranted in declaring materialism false only if we believed that it is the job of materialism to fill in the explanatory gap. I do not believe it is plausible to think that any materialist theory of the universe must fill

in every explanatory gap. In this paper I offer two reasons for thinking this. First, there are many examples of explanatory gaps concerning phenomena that are decidedly physical. Second, it is quite plausible to believe that the non-mental world is full of qualitative phenomena that leave their own gaps. This fact is often obscured by the acceptance of the view that, borrowing from Peter Unger, I will call the Restriction of Qualia to the Mental (the Restriction, for short): the view that all phenomenal properties are mental in nature. If we reject the Restriction but continue to hold to the view that phenomenal properties are non-physical, we will be faced with the prospect of a world full of non-physical properties.

The second major aim of this chapter (in a sense the most important one) is to offer a compelling explanation for the existence of the explanatory gap and to state precisely what sort of problem it poses for materialist conceptions of mind. When we understand precisely why the gap exists, we will be in better position to evaluate the consequences it has for materialism. To anticipate a bit, I conclude that the identity theory does not offer the kind of explanation for qualia that many philosophers of mind have been after. But this does not mean that we need to reject materialism.[36] I will explain why it is that we feel compelled to ask these perplexing questions about phenomenal experience; why we feel compelled to find qualia so mysterious.[37] The answer will have nothing to do with whether or not these properties are physical.

---

[36] I should also say that this is not because I accept a functionalist account of qualia. For reasons I won't go into here, I believe functionalism about qualia is just as mistaken as the identity theory. In fact I offer no positive theory explaining qualia. My view is that materialism does not need to offer the sort of explanation for qualia demanded by Levine, Chalmers, Jackson and others.

[37] Since this is a paper in the philosophy of mind and not the philosophy of science, I will not offer much of an explanation for the corresponding questions concerning the fundamental laws of physics

When you read the relevant philosophical literature, you find philosophers asking a number of different questions concerning qualia. Thus it is important to identify what questions qualia pose for the philosophy of mind. I think we can identify at least five basic questions surrounding the issue of qualia. I bring them up not because I intend to answer them all but simply to clarify the issues involved. They are: First, and perhaps most central to the mind-body problem,

(1) How does the brain give rise to qualia? Or, more basically, what is the relationship between qualia and the brain?

This is the basic question that frames the entire analysis. One way of posing it is as follows: We know that there is a correlation between brain processes and mental states. What we want to know is what is the reason for the correlation.

The second issue is harder to state. It stems from the fact that the qualitative character of a particular sensation does not at all appear to be a consequence of the physical phenomena with which the sensation is associated. The best way to feel the force of this issue is to consider an example: Why should this particular electromagnetic reflectance profile be blue and not green? So the second question is,

(2) Why do the qualia have to be like this?

Why should these physical features be associated with these particular qualitative features? David Chalmers puts the question very well, "But why should that waveform, or even these neural firings, have given rise to a sound quality like that?" (Chalmers 7).

---

except to say that I suspect that it has something to do with Wittgenstein's famous assertion that all explanation must come to an end.

The third question concerns the traditional way of dividing up the features of objects into primary and secondary properties, where primary properties are those that objects genuinely possess and secondary properties are not genuinely possessed by objects but are somehow a consequence of our sensory apparatus.

(3) Are qualia features of the mind or are they features of objects? Does the sky genuinely possess the property of blueness or is color a property of experiences?

There are other questions, including epistemological worries such as (4) Can I know what another person's qualia are like? And (5) Can we know that such and such a qualitative state is always correlated with such and such a brain state? But I will be for the most part ignoring these.

My argument in this first section will be that much of the reason philosophers have seen qualia as mysterious stems from their stances toward questions (2) and (3).

As I have already indicated, Chalmers provides a very compelling formulation of question (2) in his book, *The Conscious Mind.*

> Given that conscious experience exists, why do individual experiences have their particular nature? When I open my eyes and look around my office, why do I have *this* sort of complex experience? At a more basic level, why is seeing red like *this* rather than like *that*? . . . Why, for that matter, do we experience the reddish sensation that we do, rather than some entirely different kind of sensation, like the sound of a trumpet? (Chalmers 5).

It is not obvious to me that any theory about the relationship between the mind and brain must answer such questions. Chalmers is quite right to point out that we haven't the beginnings of an explanation for why qualitative phenomena should have

the particular qualitative features that they do. He is right that it is a total mystery. But that does not mean that any solution to the mind/body problem must answer this question; or even that the question has an answer. It would in fact be not at all surprising if we failed ever to achieve any kind of explanation. Nor do qualia stand alone as phenomena for which such explanations are not forthcoming.

Against the kind of reasoning behind Chalmers question, I want to claim that this sort of question is not confined to the qualitative realm. We can equally ask, and be at great pains to answer, analogous questions concerning the nature of reality at a very fundamental level. Why, for example, does every particle in the universe attract every other particle with a force that varies inversely with the square of the distance between them? Why does the force vary inversely with the *square* of the distance, rather than the *cube* of the distance? Why does every electron repel every other electron and attract every proton? Why shouldn't it be that electron's attract other electrons and repel protons? Why is the world like *this* rather than *that*?

I don't want to deny that deeper explanations can be given for these phenomena. My point isn't that physicists have no explanation for why certain particles have the charge that they do. I certainly don't know enough physics to make that claim. My point is that we don't demand that physics provide these answers. If a physical theory had no explanation for why electrons have *this* charge rather than *that*, we would not regard the theory as a complete failure. Nor would we, I believe, want to claim that electric charge must be non-physical.

Suppose that we did have an explanation for electric charge (for all I know there is some such explanation). Suppose that this explanation was to the effect that electrons and protons were composed of different, more fundamental, sub-atomic particles and that these particles had special properties that give electrons and protons their respective charge. This explanation would invite the further question of why these explanatory fundamental particles have the particular physical nature that they do. And of course if we were offered a further explanation, we could equally well ask a yet further question as to why things have to be *this* way, and so on ad infinitum. The point is that we can ask the sort of question posed by Chalmers about any physical phenomenon; thus we must acknowledge the fact that we will never get a fully adequate and complete answer to this sort of 'why' question.

The view that qualia create an impenetrable mystery is also bolstered by the attitude taken toward question (3). Most philosophers who engage in the mind/body debate write as if they believe that qualitative features are found only within the mind. They are committed to a view that I will call, borrowing from Peter Unger, the Restriction of Qualia to the Mental, or the Restriction Thesis (RT) for short.

(RT) Every qualitative feature is a feature of some conscious mental state.

We can see this assumption at work in Frank Jackson's thought experiment about the color-deprived scientist named Mary. Mary spends all of her life in a room in which all color has been carefully removed. She has never seen any colors except

for varying shades of white, black and grey. Her life is not completely impoverished however, and she develops into a brilliant scientist who learns, over the course of her life in this drab environment, all of the physical facts about human neurobiology. (In some versions of the story Mary is so brilliant she knows every physical fact about the universe, not just those of neurobiology.) Still, Jackson claims, there are things that Mary doesn't know. One thing Mary does not know is what it is like to see a ripe tomato. Thus, though she knows all of the physical facts about my neuro-anatomy, she does not know what I experience when I look at a ripe tomato. When she leaves her room and sees a tomato for the first time, she will learn what this experience is like. Thus the physical facts cannot capture all of the facts about conscious experience.

> Jackson formalizes his argument as follows:
>
> (1) Mary (before her release) knows everything physical there is to know about other people.
>
> (2) Mary (before her release) does not know everything there is to know about other people (because she learns something about them upon her release).
>
> Therfore,
>
> (3) There are truths about other people (and herself) which escape the physicalist
>
> story.
>
> (Jackson, 568)

Much paper has been devoted to discussion of this example and there is little consensus about whether it is a successful criticism of physicalism. I don't intend to rehash the debates here. However, I do wish to point out that surprisingly little attention has been paid to the fact that it is far from obvious that what Mary doesn't know has anything to do with mental life.

It is undeniable that, before her release, Mary doesn't know what a ripe tomato looks like. One question we might ask is whether this is a mental fact. Is this fact that Mary learns a fact about mental life? Isn't it, rather, a fact about tomatoes that they turn red as they ripen? When Mary leaves her room one of the things she will learn is that ripe tomatoes are red. Of course, in a sense, she already knew this. Among the physical facts that she has learned is the fact that a beefsteak tomato turns red when it ripens. But Mary can't identify which color red is. If, immediately upon her release from her achromatic room, she is handed a color palette and asked to pick out the color of a ripe tomato, she won't be able to do it, having never seen a tomato in color. What Mary doesn't know then, is that a tomato has this color. Thus, in a very real sense, even though she has been told it is so, Mary doesn't know that tomatoes are red.

If this is so, then using reasoning analogous to Jackson's, we can reach the conclusion that there are truths about tomatoes that escape the physicalist story:

(1) Mary knows everything physical there is to know about tomatoes.

(2) Mary does not know that ripe tomatoes are red.

(3) Therefore, there are truths about tomatoes that escape the physicalist story.

   (Specifically, the fact that tomatoes are red escapes the physicalist story.)

Also note that, barring the Restriction Thesis, Chalmer's question (2) becomes one about all sorts of non-mental phenomena (that is, phenomena that are external to and independent of anyone's mind). "Why is red like *this* rather than *that*?" is a question about the color red, a color possessed by objects, such as many ripe tomatoes, some fire trucks, and the Anaheim Angels baseball cap. It is not a question about the mind. The upshot is that if we reject RT, question (2) cannot plausibly be thought to be a question that is exclusive to the mental realm. It is, rather, a question for metaphysics more broadly. There are, of course, various qualitative phenomena, such as pain and pleasure that are obviously mental in nature. But colors, smells, and sounds, barring RT, are not mind-dependent.

Of course this argument only works if we assume that the Restriction Thesis is false. Only if we allow that objects external to the mind can genuinely possess colors can we conclude that the fact that tomatoes are red is a fact about tomatoes rather than a fact about human consciousness. But why shouldn't we assume this? One clue that RT is false is the way we talk about qualitative phenomena. Consider again Chalmer's formulation of question (2): "Why, for that matter, do we experience the reddish sensation that we do, rather than some entirely different kind of sensation, like the sound of a trumpet?" Why does he use the word 'sensation' here? 'Sensation' is a term more properly used to refer to states such as pains, or itches, or pleasures. A pain is a sensation, an itch, etc, but it sounds at least a bit odd to say that red is a sensation. Red is a color and objects are colorful but sensations are not. I doubt that many

people would be inclined to say that when they looked at a ripe apple they were experiencing a sensation, at least in non-philosophical contexts.[38]

Using the phrase "reddish sensation" when talking about the experience of seeing a ripe tomato encourages the reader to identify the relevant qualitative feature (the redness) as something mental. Consider this reformulation of the question: "Why do we experience the reddish color that we do?" This question, without the loaded terminology, does not invite us to assume that what we experience is a property in the mind.[39]

I don't want to be dogmatic and simply assert that it is wrong to use the word 'sensation' to refer to visual experiences. However, if we pay attention to how we talk about the qualitative phenomena associated with visual and auditory experiences compared to how we describe sensations such as pleasure and pain, I believe that we are forced to recognize while some instances of qualia are features of mental life, many qualitative features are not (or at least that we do not conceive of them as such). For example, it is natural to use the expression 'feel' when talking about pain. We say things like, "I feel bad", "My head feels like its caught in a vise." But is very unnatural to say something like "When I look at the sunset I feel a beautiful redness" or "I feel a very high note." It's far from clear what these last two sentences might mean. We say, "I feel pain," "I feel an itch," " I feel pleasure," but never, "I feel redness," "I feel Middle C," or even "I feel a stench."

---

[38] Someone might say that he was experiencing a sensation of great pleasure while looking at the apple. But this would either be a very odd situation or a very odd person. We certainly wouldn't normally say that we experienced a sensation *of color* while looking at the apple.

[39] In fact, this question is rather ambiguous. It could be a question about the human visual system, or a question about the tomato (why is it red?). I will discuss this question and its ambiguity further below.

Despite this, some philosophers are comfortable using such tortured expressions: For example, Chalmers says, "any color can be awe-provoking if we reflect upon its nature. Why should it feel like that? Why should it feel like anything?" (Chalmers 7). To speak of a color as feeling a certain way is to stretch the meaning of 'feel.' Perhaps I am missing out on a whole range of qualitative experience but I wasn't aware that colors felt like anything. Again, using "feels" when talking about colors obscures the fact that, while some qualia are obviously mental, other types of qualia are, arguably, external to the mind.

When we start using expressions like this, it is probably an indication that we need to slow down and look more closely at how we actually talk, in ordinary contexts, about our sensations and experiences. Part of the confusion stems from the fact that 'feel' is used both as a noun and as a verb. When talking about tactile experiences we use the verb form of feel, as in "The blanket feels soft." Corresponding to this use of 'feel' there is a noun form, as in 'the feel of silk.' It seems likely that this noun form is derivative of the verb form used to report tactile experiences.

Adding to the confusion is the fact that we also use the verb form when talking about sensations such as pleasure and pain, about our emotions and sometimes about our beliefs ("I feel that there must be a god"). But it is pretty easy to distinguish the cases in which we are talking about a tactile experience or impression from those in which we are talking about some aspect of our mental life (i.e., when talking about sensations, emotions, beliefs etc.). It is interesting to note that the noun form

corresponding to this use is much more rare, if not non-existent. We talk of a pleasurable *feeling*, not a *feel*; we say that our feelings (not our feels) are hurt; we speak of a feeling that god exists, etc.

Despite the fact that 'feel' has a multitude of proper uses, it is, as I have said, relatively easy to distinguish those instances in which we use the word to talk about the tactile features of external objects from instances in which we use it to talk about our internal psychological states. That is, it's easy when we think about it carefully. When I say that a cat is soft, I am saying something about the cat's fur, not about my mental life. And when I say that I feel bad, I am talking about my internal psychological experiences. Nevertheless, the ambiguity of the word does cause confusion when philosophers start talking about the feel of the color red. Are we talking about a property that an object has or a property of mental states. If we aren't careful we are apt to interpret the use of 'feel' here as an indication that the qualitative phenomena in question is an internal phenomena, similar to emotions and sensations. Philosophers who use expressions like 'qualitative feel' and 'red sensation' invite themselves and their readers to understand the relevant phenomena as intimately connected to their own mental lives. But, as I have argued, this obscures the very real issue of whether some qualitative features are external to our minds.

In addition to our use of 'feels' there is other telling linguistic evidence. We say things like, "I have a pain" or "I've had an itch that I can't scratch." But we wouldn't say "I have a redness" or "I have a smell" (though this last sentence has a pretty clear alternative use). In other words we attribute pains and other sensations to

ourselves, or to parts of our bodies but we attribute colors to external objects. To be sure, we don't normally say that objects *have* colors; we are more apt to say, e.g., that the tomato *is* red. But we do not speak of our visual experiences as either having or being, e.g. red.

These facts, in addition to common sense, suggest that we need to distinguish qualitative phenomena that are properties of mental states (or *are* mental states) from those that are not. For convenience sake we can speak of 'internal qualia' to refer to those qualitative features that are a part of mental life, and 'external qualia' for the qualitative phenomena that are not aspects of mental life. Pain, pleasure, and itches, are types of internal qualia. They are events in the mental life of a conscious subject. Colors, sounds and smells are among the external qualia; they are not mental events, nor are they properties of mental events. They are properties of external objects.

If there really are external qualia then the Restriction Thesis is false. If RT is false, then questions such as "Why do we experience the reddish sensation that we do?" are not questions in the philosophy of mind. This question, if it makes sense at all, is not one that we should expect any theory of the relationship between mental states and brain states to answer. It is a question for metaphysics more broadly.

Some philosophers might deny that they hold the Restriction Thesis. Chalmers has the following interesting footnote to his use of the expression "red sensation": "I use expressions such as 'red sensation,' 'green experience,' and the like throughout this book. Of course by doing this I do not mean to imply that experiences instantiate the same sort of color properties that are instantiated by objects (apples, trees) in the

external world. This sort of talk can always be rephrased as 'experience of the type that I have (in the actual world) when looking at red objects,' and so on, but the briefer locution is more natural." (Chalmers 359).

In this passage Chalmers seems to be suggesting that when I look at a ripe tomato, there are two relevant phenomena: the tomato and my experience of the tomato. Only the tomato has the color property. This seems to be the cash-value of the suggestion that the experience does not instantiate the same property as the tomato. But this implies that it is wildly inappropriate (in addition to being unnatural) to speak of red sensations. In the experience of the tomato only one thing is red, and it isn't the experience. If this is right, then color is by no means a mental phenomenon. Sure, there are characteristic experiences people have when they look at tomatoes but, as Chalmers seems to recognize, the experiences themselves are not red. And if we want an answer to the question "Why does the experience have to be like *this*?" the only applicable answer is "Because tomatoes are red."

We may want to know, in addition, why tomatoes have this particular color, but this is a question about an external object, not a state of consciousness. But Chalmers, as we have seen, is very persistent in claiming that the question "Why does the color red (or 'red sensations') have the particular character that they do" is a question that must be answered by an adequate theory of the mind. The only way to

do this is to maintain that colors are mental in nature; that is to hold the Restriction Thesis.[40]

I think that Chalmers has gotten himself into a muddle. But it is not one that he created, nor is he alone among philosophers who write about qualia. I think many philosophers are guilty of it to some degree; we already saw an example in Jackson's colorblind Mary argument. John Searle, in his paper "The Problem of Consciousness" says the following:

> Indeed, until quite recently many workers in cognitive science and neurobiology regarded the study of consciousness as somehow out of bounds for their disciplines. They thought that it was beyond the reach of science to explain why warm things feel warm to us or why red things look red to us. I think, on the contrary, that it is precisely the task of neurobiology to explain these and other questions about consciousness. (Searle 2002, 15)

To be fair to Searle, the question "why do red things look red to us?" is not quite the same as "why does a red experience have to feel like *that*?" In fact it is quite ambiguous between this interpretation and another. We might understand Searle's question to be part of a larger concern about how human beings are able to see colors. It is a demonstrable fact that not all animals (not even all humans) see the color red. And neurobiology can certainly help us understand why most people see red whereas a total achromatic, for example, would see only a shade of grey. We are not puzzling at the mysteriousness of redness here, but wondering what is different about beings that

---

[40] Strictly speaking, there is another option. We might hold that experiences *do* instantiate the same sort of color properties as external objects. In other words, experiences have color and objects have color. Thus we could reject the Restriction Thesis and still think that the philosophy of mind has a job to do in explaining color experiences. First, though, Chalmers appears to explicitly reject this option when he denies that he is implying that "experiences instantiate the same color properties that are instantiated by objects." And second, and more to the point, if we accept this option, we must also accept that whatever mysteries are posed by qualia are not limited to the mental realm.

can see colors (or a specific color) from those that cannot. An adequate answer to this question would be of the form: Humans see red, while dogs do not, because humans have the right sort of visual and neurobiological apparatus.[41]

Why do philosophers conflate external and internal qualia? Searle says something interesting and revealing in connection with this in his book, *The Rediscovery of the Mind*. He discusses the notion of reduction and analyzes historical examples of successful reductions. He says that in cases where the reduced phenomena is a secondary property, "the point of the reduction was to carve off the surface features and redefine the original notion in terms of the causes that produce those surface features." (Searle 1992, 119). He offers the reduction of heat to mean kinetic energy and colors to reflectance profiles as examples:

> We then redefine heat and color in terms of the underlying causes of both the subjective experiences and the other surface phenomena. And in the redefinition we eliminate any reference to the subjective appearances and other surface effects of the underlying causes. 'Real' heat is now defined in terms of the kinetic energy of the molecular movements, and the subjective feel of heat that we get when we touch a hot object is now treated as just a subjective appearance caused by real heat, as an effect of heat. It is not longer a part of real heat. A similar distinction is made between real color and the subjective experience of color. (Searle 1992, 119)

The idea is that we distinguish between the physical reality of heat and color and the subjective appearance caused by the physical reality. The appearance, I suppose, exists only within the minds of conscious beings who interact with the physical reality. He goes on to say that pain cannot be reduced to physical states of

---

[41] Obviously a truly adequate answer will have much more to say about the precise nature of the apparatus required to see red.

the brain because, where pain is concerned, we cannot make the requisite distinction between appearance and reality. With pain, what we are interested in is the appearance so it will not work to carve it off and redefine 'pain' in terms of the physical reality of brain states. Searle explains: "the reduction of pain to its physical reality still leaves the subjective experience of pain unreduced, just as the reduction of heat left the subjective experience of heat unreduced." (121).

Here Searle seems to be implicitly accepting the Restriction Thesis. All qualitative features, such as colors and the feeling of heat, are subjective appearances, not any part of physical reality. According to his analysis of the structure of ontological reduction, qualia are never reduced. Why should we think this? The reason is quite straightforward: any proposed reduction of a qualitative feature such as heat (I should say, the feeling of heat to make it clear that I am talking about a qualitative feature) would be subject to the same sort of arguments as now confront the proposed reduction of conscious states. Specifically, it is imaginable that the feeling of heat could exist without kinetic energy; it is imaginable that the color blue could exist without the reflectance profile specific to blue, etc. To avoid this problem, successful reductions, according to Searle, must carve off the "subjective feature" (by which he means the qualitative feature) and redefine the phenomena, heat, in terms of objective reality. This process results, if carried to its logical consequence, in a conception of a world of objects that do not possess qualitative features. Qualia have been carved off external objects and pushed into the mind.

I will return to this point below when I discuss my own reasons for rejecting the identity theory for sensations. For now I only want to point out how Searle's analysis of reduction supports the view that all concerns about qualia are concerns for the philosophy of mind. If we believed that some qualia are external in the way that I have described, we would reject Searle's analysis. This would have the interesting consequence that, perhaps, the proposed reductions of heat, colors, etc. are not as successful as we assume.

Despite all that I have said to the effect that colors, sounds and smells are properly thought of as external qualia, I do admit the possibility that they are mind-dependent in a crucial way. Perhaps the answer to my question (3) is that qualia are only features of mental states. But I don't really want to take a stand on this issue here. I only want to point out that it is far from obvious that the problem of qualitative features is a problem exclusively for the philosophy of mind. It is enough if I have shown that the onus is on the philosopher of mind to show that qualia are excluded from the non-mental realm. If this cannot be shown, then we should conclude that question such as (2) are not confined to the philosophy of mind.

Before I conclude this section I need to make one more observation. I began this paper by noting that if one takes a look at recent work in the philosophy of mind, it is easy to get the impression that the mind/body problem consists of the problem of explaining qualia. By rejecting the Restriction Thesis (or at least calling it into question), we can now see that this is indeed a mischaracterization of the problem of consciousness. Consider again the conscious experience of seeing a ripe tomato.

Chalmers would have us believe that the hard problem is explaining why this experience has the peculiar character that it does. But if the Restriction Thesis is false, and external objects genuinely possess qualitative features, then Chalmers question is incoherent. The question is wrongly put because it is not the experience that has the qualitative feature, but the tomato itself. Thus, if the presence of the color red demands explanation, we cannot demand the explanation of the philosopher of mind. To repeat, the problem is not one in philosophy of mind, but in what we might call the metaphysics of qualia.

Nonetheless, the experience under discussion does pose a genuine problem for the philosopher of mind. Namely the question of how awareness is possible to begin with. This question is obscured by those who focus on qualia, but it is a fundamentally different issue from that of explaining qualitative features. The mind/body problem, then, is not to explain why a tomato is the color that it is, but to explain how it is possible for physical matter to be aware of a tomato. This, to my mind, is the classical formulation of the mind/body problem.

The Identity Theory

So far I have argued that some of the worries that arise about qualia and their relationship to the brain are not exclusive to that realm. Since similar concerns arise for decidedly physical phenomena as well as what I have termed non-mental qualitative features, it is at least questionable whether these worries are properly

thought of as exclusive to the philosophy of mind. Nevertheless, sensations[42] are obviously mental phenomena and thus the issue of explaining sensations false within the scope of philosophy of mind. And, as I have indicated, I think that the explanatory gap does arise for sensations. Philosophers have pointed to the gap and drawn many conclusions. Some, like Jackson declare that qualia must be non-physical. Colin McGinn suggests that we are constitutionally incapable of understanding an explanation that would bridge the gap. Before we draw conclusions, though, we ought to spend some time trying to understand why the gap exists. I believe that once we understand why there is an explanatory gap we will see that it is not much of a threat to materialism after all.

I will begin, then, with a discussion of sensations, since they are the most obvious place where the problems posed by qualia arise in the philosophy of mind. My argument will be that because of the special epistemic status that sensations possess, the identification of sensations with brain-processes is a non-starter. Much of what I have to say about sensations can also be said of external qualia and I will extend my analysis to cover qualitative features as a whole.

If we look closely at the criteria of application for sensation words like 'pain', we will begin to see why it is that identifying pains with brain states is problematic. As I said near the beginning of this paper, sensations have a special epistemic status that our knowledge of the external world of objects lacks. In the *Philosophical Investigations*, Wittgenstein says, "What I do is not, of course, to identify my

---

[42] Here, and in what follows, I am using the word 'sensation' to refer to things like pains, itches, and tickles; in other words what some might call 'bodily sensations.'

sensation by criteria" (Wittgenstein 1958a[43], 290). He is drawing our attention to the fact that, when I am in pain, I don't think "Oh there's a sensation. Now which one is it? Well it has such and such features so it must be . . . pain." In an important sense pains do not have features in the way that trees and houses and faces have features; thus we cannot identify them by recognizing features they have. Since we do not identify our sensations in virtue of criteria they satisfy, but simply experience them, it does not make sense to think of introspection as a faculty that gives us access to our sensations. It is true that my sensations give me access to processes in my body and brain (e.g. the sensations I get in my stomach after I eat a large meal might give me access to digestion) but these processes that I have access to are not the qualitative experiences themselves.

When Wittgenstein said that what I do is not to identify my sensation by criteria, he was discussing the criteria of application of sensation words. I don't identify my sensations by their properties in the way that I identify objects, like my dog, by her properties. When I see my dog after many months, I know that it is her because I recognize her size, her graying, black coat, her short stubby legs, etc. I cannot do the same thing to identify a pain. What property could I rely on to identify a pain except its painfulness. Wittgenstein's point was that I don't use properties to identify my sensations, so the property of a sensation cannot serve as a criteria of use of the word 'pain.' As I said earlier, introspection is not a faculty of the mind whereby

---

[43] Henceforth, I will refer to Wittgenstein 1958a as PI. Numbers cited in *Philosophical Investigations* are section numbers.

I sort sensations by their properties; I don't say, "Hmm, interesting sensation, it has this strange painful character so it must be a pain." I just feel the pain.

The upshot of this discussion is that sensations are very direct and basic. Sensations have an epistemic status that our knowledge of the external world lacks. Wittgenstein expresses this by saying, "it makes sense to say about other people that they doubt whether I am in pain; but not to say it about myself" (PI 246). When I am in pain, it does not even make sense to say that I could doubt that I am in pain. I can doubt that you are in pain because I do not experience your sensations directly.

The fact that sensations have this epistemic priority implies that we will think of them as different and separate phenomena from the states of the brain with which they are correlated. This feature also explains why it is that we will always be able to imagine that pain might exist without the underlying brain state. Since sensations are so direct while brain states are not, our knowledge of the former will always be more intimate than that of the latter.

Much of what I have said about sensations can also be said about other types of qualia. First, we do not identify a color, or a smell, or a taste, based upon properties that the color, smell or taste has. When I see the color blue, I do not recognize it on the basis of features that it has. What property could I use except its color? Furthermore, our knowledge of external qualia, while perhaps not as direct as our knowledge of sensations, is more direct and basic than our knowledge of objects. Our sensory faculties allow us to immediately experience the colors, sounds, and smells of the world. We encounter the world by experiencing its qualitative features.

We can see a tree, smell it, hear its leaves rustling in a breeze, even taste it if we wanted; but you cannot smell or taste a color and you cannot see a sound. Just as a sensation of pain is not an object that we can have access to, a color is not an object with features. Objects like trees have many features and we can experience those features through our different sensory modalities. But qualitative features themselves are featureless, and we cannot access a color, e.g. through other sensory modalities.

The point is that my experiences of qualia have a sort of directness that my knowledge of external objects lack. I can know that there is a computer in front of me but this is because I have a direct awareness of the computers qualitative features. Since I don't directly experience the computer in the way that I directly experience the color of the computer (gray), I can doubt that there is a computer out there, but I cannot doubt that I had an experience of gray. Qualia are the means by which we have access to the world and as such enjoy a special epistemic status. I will try to explain what I mean.

I have access to tables, chairs, books, people and I have access to these things because I can experience their qualitative features; but I do not have access to the qualitative features themselves. Qualia are not the sort of things that we have access to, they are the means by which we have access to the world. Or, to make the same point slightly differently, we can say that my sensory apparatus allows me to have access to objects because, in virtue of being able to see colors, feel shapes, hear sounds, I have access to the world. I have access to the world because I experience its qualitative features. If we understand qualia in this way then it is at best misleading to

speak of having access to qualia since qualitative features are the means by which I have access to the world and I do not have access to qualia in virtue of their qualitative features. Objects are things I have access to; qualia are not.

Our qualitative experiences give us access to the world and so we do not have access to qualia in the way that we have access to the world. It is extremely misleading to speak of introspection as a means of access to our mental lives. Arguments like Jackson's and Chalmer's correctly point to an epistemic asymmetry between conscious states and knowledge of the external world but it is misleading to say that I have direct access or direct knowledge of my mental states. To put it simply, my knowledge of my mental states is more direct than that. The epistemic status of qualia is so direct that it prompted Wittgenstein to say "It can't be said of me at all (except perhaps as a joke) that I *know* I am in pain. What is it supposed to mean -- except that I *am* in pain?" (PI 246). Qualia are direct and basic, we do not experience qualia on the basis of features that they have.

Some philosophers who are inclined toward a reductionist account of qualia are unmoved by the special epistemic status that sensations enjoy. Paul Churchland, for example, says this: "whatever else it may be, introspection is an epistemic modality, or perhaps a family of them. And while it may have its own quirks and distinguishing profile, it is entirely unclear whether it, alone among all of our epistemic modalities, constitutes a window onto a unique *ontological* domain of nonphysical properties. None of our other epistemic modalities has any such

distinction: they all access some aspect or other of the purely physical world. Why should introspection be any different?" (Churchland 1998, 133).[44]

Churchland admits this much about the nature of qualia: "Given any person at any time, there must be some set of features whose spontaneous or non-inferential discrimination is currently basic for that person, a set of features whose discrimination does not depend on the conscious discrimination of any more elemental perceptual features." (Churchland 1998, 138). But he says that too much has been made of this fact since it is completely inevitable. The epistemic status of qualia is inevitable but this does not undermine the claim that qualia are physical features of the brain, "If ultimately they are physical, then inner qualia ought to be accessible from more than just the first-person or "subjective" point of view; they ought to be accessible as well from one or more "objective" points of view, via some appropriate instruments that scan brain activity." (139).

If I am right about the epistemic status of qualia, the mistake here is thinking that qualia are accessible in the first place. If qualia were the sorts of things that were accessible, then we would have to have access to them in virtue of some feature that they have. But as I have shown, this makes no sense. Those phenomena that allow us to have access to the world are not things that we have access to in the way that we

---

[44] In the following discussion, I will use the word qualia without distinguishing between internal and external. I do this because Churchland, as with most philosophers, does not explicitly make this distinction and so to expedite the debate it is simpler to follow his terminology. This will cause some confusion since Churchland speaks of qualia being reduced to brain states. But if RT is false this is a non-starter to begin with (at least for external qualia); colors will never be reduced to brain states since they aren't even properties of mental states to begin with. I will try to avoid excessive confusion by using examples about sensations rather than external qualia. Nonetheless, everything *I* say about qualia is true, I believe, of both internal and external qualia.

have access to the world. For this reason we ought not to say that qualia are *accessible* from the first person perspective, but rather that they are *experienced* from the first person perspective.

But Churchland continues to make the mistake, describing introspection as an epistemic modality. He notes that the different sensory modalities of vision and touch can allow us access to the same phenomena. I have access to the shape of a box through the modality of vision and of tactile sensation. He says that just because we have access to qualia through the subjective perspective, this does not mean that we can't have access to the same phenomena from the objective perspective. He claims that introspection is just one of many epistemic modalities and that just as I can have access to the box through multiple sensory modalities, we may be able to have access to qualia from epistemic modalities other than introspection. Indeed, he says, "In at least some cases, one and the same (physical) state can be known both subjectively and objectively" (140). "One can tell by introspection that one's own bladder is full, but an ultrasound image will tell you the same thing." (139). This is a case in which a qualitative sensation, the pain associated with needing to empty one's bladder, gives me access to the fact that my bladder is full. The ultrasound image can give one access to the same fact but it cannot give one the experience that I am feeling.

The example he describes is a case of having *access* to physical properties but, as I have repeatedly said, qualia are not things that we have access to. I can experience the box in virtue of its visual appearance and its texture; I have access to it through different sensory modalities. But the only way to experience a qualitative

state (such as pain) is to have that state, and the only way to experience the color of an object is to look at it (you can't see the color of a tree by smelling the tree). Qualia are not something it even makes sense to talk about having access to through different modalities. The shape of an object is something that we can have access to through multiple sensory modalities, but the experiences we have (of, e.g. the color and texture of the box) in virtue of which we have such access are not accessible in that way.

Introspection is not an epistemic modality and qualia are not things or properties, or properties of things that we have access to. Touch and vision are epistemic modalities and we can have epistemic access to the same phenomena in virtue of its visible and feel-able qualities. We can see the shape of a box and feel the shape of a box. But it is what the box actually feels like or what it actually looks like that is the quale and we cannot feel what a box looks like. A blind person cannot see the shape of a box, though he can feel its shape. Similarly, we should not expect that a person whose nervous system is not connected to the injured part of a body could feel the pain. He might know that the person whose body is injured is in pain and he may be able to see the brain state that is associated with the pain (or is the pain), but he cannot feel the pain.

If Churchland were right, and qualia were brain states (internal qualia, in any event) and were accessible from multiple perspectives (not just the subjective, first person perspective), the picture would be this: When I have a particular qualitative experience, a pain of a needle prick, say, I should be able to have access to the pain from a different perspective (through a different epistemic modality) and others should

be able to have access to it as well. But when we stop and think about it, this idea is confused at best. When brain scientists look into my brain with a brain-o-scope they will have some experiences in virtue of which they have access to the brain state. But the qualitative features of the brain state that they would experience would not be the quale of pain because the quale of pain is not something you get just by looking at something (you don't feel pain in your finger just from looking at a brain). Since they do not have the experience of pain what sense does it make to say that they have access to the experience? They can't have access to some other properties of the experience because qualia are not objects with properties; they are the qualitative features (properties) in virtue of which we have access to objects. Churchland's solution might work if qualia were, like objects, things that we had access to in virtue of their properties. But this is precisely what qualia are not.

The epistemic status of qualia ensures that we think of them as distinct phenomena from physical states (brain states in the case of internal qualia). Qualia are those features that we do not have access to and in virtue of which we have access to the world. Pain and other bodily sensations give me access to my body, damage to it, its position etc.; colors and sounds give us access to the objects around us. When we reflect on the fundamental difference between qualia and the phenomena in the world that we have access to in virtue of qualia, the identity theory seems less plausible. We don't have access to pains in the way that we have access to this piece of paper, or brain state XYZ. The only way to experience something that is the rock-bottom

epistemic basis for getting at reality (to borrow a phrase from Searle) is to experience it.[45]

In my argument I have been discussing qualia as a whole (internal and external), but, strictly speaking, I could formulate the argument in terms of internal qualia alone. Sensations are not objects that we have access to but are qualitative phenomena that give us access to our bodies. Thus it would be a mistake to identify sensations with brain processes or state since brain states are the sorts of things that we have access to. This argument works regardless of whether the Restriction Thesis is false.

In any event, what I have to say about colors, sounds, smell etc. (all of those qualitative features that I have called external qualia), is not essential to my point about mind/brain reduction since, as I have argued, I am not convinced that these external qualia are properly thought of as mental phenomena. Perhaps my argument about sensations can be carried over to cover qualitative features like colors. This would imply, as noted earlier, that any proposed identification of, e.g., the color blue with reflectance profile *P*, would be unsuccessful in as much as a qualitative feature is

---

[45] Some of what Searle says in *Rediscovery of the Mind* is very similar to the claims that I have been making. He agrees, for example, that it is misleading to speak of introspection as an epistemic modality. He holds an anti-reductionist position according to which mental states have an irreducible first-person ontology. However, as I have said, I believe that he does hold something like the Restriction Thesis, as is demonstrated by his analysis of the reduction of heat which I quoted earlier. I understand him to hold that qualitative phenomena are mental in nature and that in typical cases of ontological reduction, the qualitative phenomenon (e.g., the feeling of heat) are carved off of the physical phenomena (e.g., mean molecular kinetic energy) and explained as subjective appearances that the physical phenomena cause in the mental life of conscious beings. All qualitative phenomena, according to Searle, have the first-person ontological status. According to my view, the first-person ontological status has nothing to do with irreducibility or with the explanatory gap. It is, rather, that qualitative features (whether they are features of objects, like color, or mental states, such as pain) are experienced directly as featureless features that give us access to objects and, as such, we tend to think of them as distinct phenomena from their correlates. (E.g., the experience we have when we feel heat is thought of as distinct from molecular kinetic energy.)

being identified with an object. I am not prepared to expand and defend that argument here. But it is interesting to note that, if I am correct, the explanatory gap exists not just within the mental realm but in the realm of external objects as well; and my analysis of qualitative features (taken as a whole) explains why the gap should exist.

Is this really an argument against Reduction?

One possible response to my argument is to claim that all that I have shown is that we are forced to think of, for example, pain as a distinct phenomena from the brain state correlate. I have not shown that pain is not a brain state; I have simply explained why it is that we think of pain as a distinct phenomena. But just because we make the distinction between pain and C-fiber stimulation (for example) does not mean that they really are different phenomena. At most what I may have shown is that it is difficult to understand the theoretical identification of pain and C-fiber stimulation, not that that the identity does not hold.

I think that this is a plausible response to the arguments I have made in this chapter. My primary aim, as I stated at the beginning, was to explain why there is an explanatory gap in proposed mental/physical reduction. I believe that I have offered a plausible explanation for why this gap will always exist. But I have not demonstrated the stronger metaphysical claim that the existence of the gap implies that mental states cannot be brain states. In order to establish this stronger claim I would have to offer an argument to the effect that the special epistemic status that qualitative phenomena

have gives them a distinct ontological status. I do not know if such an argument is possible; in any event I am not prepared to offer one here.

It is useful to consider, however, what effect the existence of the explanatory gap has on proposed identifications between qualitative phenomena and their physical correlates.[46] As I have said, because of the special epistemic status that qualitative phenomena have, identifying qualia with non-qualitative phenomena will be problematic. And this will be true of external as well as mental qualia. Take, for example, the identification of the color green with reflectance profile *P*. Since green is a qualitative phenomena, since it is directly experienced in the way that I have described, it is natural to think of it as a distinct phenomena from reflectance profile *P*. Because of this, there will always be a perceived gap; we will always be able to imagine that the color could exist without *P* and vice versa. The identity statement 'green is reflectance profile P' does not explain why this color must be associated with this reflectance profile. That is the explanatory gap. This, by the way, is one of the reasons philosophers have held the Restriction Thesis: the identification of colors with reflectance profiles does not offer the kind of explanation of color that we desire. Searle, for example, takes this as evidence that the color (or what he calls the subjective experience of color) is not actually reduced; it is explained as an effect on our nervous systems.

Perhaps, then, we should be perfectly happy to say that the identifications of colors with reflectance profiles (or heat with mean molecular kinetic energy) are

---

[46] I use the term 'physical correlates' for lack of a better term. I do not want to claim that qualitative phenomena are not physical (thereby begging the question).

genuine reductions of qualitative features. True, they do not offer explanations of the sort that Chalmers want; we don't get explanations for why this reflectance profile should look like this (should be green). But a full explanation of this sort will not be forthcoming and it is an error to conclude that the reduction must be false unless some such explanation is offered. Just as in the case of electric charge, we may ultimately have to say "That's just the way the world is."

Analogously, despite what I have said about the reasons for the inevitable existence of an explanatory gap, we might accept mind-brain reductions (as in 'pain = C-fiber stimulation') as genuine reductions of qualitative phenomena. We would accept this if we become convinced that the fact that we naturally think of pain as a distinct phenomenon from C-fiber stimulation does not imply that, metaphysically, they are distinct phenomena. But once we recognize that there will always be an explanatory gap, we also recognize that the proposed reduction is not doing all of the work that we wish it would. Specifically, the identity theory does not explain why this particular type of mental state should be pain (or be conscious for that matter). Ultimately, then, my argument may not be a demonstration that the mind-brain reduction is not possible, only that it is not a solution that we can readily understand nor will it answer all of the question philosophers want to ask.

Conclusion

Many philosophers working in the field of mind/brain reduction see the existence of an explanatory gap as a kind of bogeyman. Joseph Levine says the following:

> If there is nothing we can determine about C-fibers firing that explains why having one's C-fibers fire has the qualitative character that it does – or, to put it another way, if what it's particularly like to have one's C-fibers fire is not explained, or made intelligible, by understanding the physical of functional properties of C-fiber firings – it immediately becomes imaginable that there be C-fiber firings without the feeling of pain, and *vice versa*. (Levine 433)

I propose to turn Levine's reasoning on its head. It is precisely because it will always be imaginable that there can be C-fiber firing without the feeling of pain (a fact that is readily understandable once we understand the nature of qualia) that there will never be a satisfactory explanation for why having one's C-fibers stimulated has the qualitative character that it does. The nature of qualia guarantees that we must think of the sensation of pain as separate from whatever neuronal state underlies that sensation. But this inevitable fact should not lead us down the path to non-physicalism. Explanatory gaps are not bogeymen, they exist in other fields, and arguably the same sort of gap that exists between brain states and sensations exists between colors and reflectance profiles. We understand why the gap exists and it has nothing to do with the ontological status of qualitative states.

To this point this chapter has been almost entirely critical in nature. I have tried to argue that certain conceptions about what is required of materialism and about the relationship between qualia and the brain are mistaken. I have not offered anything positive, for example, to replace the identity theory as an account of the mind-brain connection. Because of this, some readers may get the impression that I am at best pessimistic about the prospects of saying anything substantive about that connection. This couldn't be further from the truth.

It is true that, in this paper, I do not want to offer a positive account of the relationship between qualitative states like sensations and the brain. This is primarily due to the fact that I am just not sure how best to speak about the relationship. Sometimes I am inclined to say that sensations such as pain and pleasure are properties of brain states, though just as often I think that this way of talking stretches the notion of property.[47] Ultimately I am not sure that it is very important how we characterize the connection. In other words, I am skeptical about the importance of deciding whether we should talk of sensations as properties, or whether we should talk of supervenience, or whether we should say that the brain causes qualitative states. Coming up with the right sort of metaphysical doctrine for explaining the ontological status of qualitative states is less important than answering certain other questions about the connection between qualia and the brain. And I absolutely believe that there are important questions to be answered. While I have argued that it is illegitimate to ask, *a la* Chalmers, why the experience of pain "feels like this?" there are legitimate questions that remain.

One question that it does make sense to ask is, "In virtue of what are some brain states associated with (to use the least loaded language) pain while others are associated with pleasure?" This question would be analogous to the questions, "Why are some objects red and others yellow?" or "Why do some fundamental particles have a negative charge while others have a positive charge?" We can ask, of a particular type of brain state, "why is it painful?" (or if we don't like that way of

---

[47] As I have argued, sensations are known directly and at least some properties, such as solidity, shape, size, are not known in the same way.

putting it, "why does the brain state have the property of being painful" or "why is it/does it cause pain?" or simply, "why is it associated with pain?") just as we can ask "Why is a tomato red?" [48]

Nothing I have said implies that such questions are unanswerable. In the case of the tomato, we have something of an answer to this question. The answer involves the fact that the surface of a tomato reflects certain wavelengths of light and absorbs others. There is every reason to think that neurobiology will be able to explain what brain states that are associated with painful sensations have in common and how they differ from brain states that are associated with pleasurable sensations. This will be the beginning (at least) of an explanation for why some brain states are associated with pain while others are associated with pleasure. Of course, as my argument in the first part of this paper shows, such explanations must come to an end. We know that red objects reflect certain wavelengths and absorb others. But we don't know why some wavelengths are red (or are associated with red) and others are blue. And, in fact, it is wrong to think that science must answer this question. At some point all we will be able to say is "That's just the way the world is." Similarly, at some point the neurobiological explanations will peter out and we will reach the end of explanation. I have no idea where that end will be but I have every confidence that neurobiologists

---

[48] Answering such questions, by the way, does not require us to commit ourselves to one or another metaphysical doctrine concerning the relationship between mind and brain. It is also important to note that some so-called solutions to the mind/body problem do not even address such questions. The identity theory may explain the ontological status of mental states, but it does not explain why some states are painful and others are pleasurable.

will soon come to understand a great deal about the differences and similarities among

various brain states and that this understanding will yield fruitful explanations.

CONCLUDING REMARKS

While working on this project I have had many conversations with philosophers many of whom are a bit more than mildly skeptical about the prospects of resurrecting ordinary language philosophy. I think many philosophers working today have the sense that this school of philosophy is unjustifiably conservative. Given the obvious fact of semantic change, what is the value (to philosophical enquiry) of investigating our ordinary concepts? In fact, they might say, isn't it part of the job of philosophy to hone, correct, and revamp our concepts in the light of inherent contradictions revealed by philosophical analysis and, perhaps more importantly, shouldn't we expect our ordinary conceptual scheme to undergo dramatic changes as the result of scientific discoveries? One particular argument stands out in this connection:

Suppose that the Catholic Church, in response to Galileo's theory that the Sun, not the Earth, is the center of the Solar system and that the Earth rotates around the Sun rather than vice-versa, had said the following: "Look Mr. Galileo, we've been through this before. You want to say that the Earth moves. But in uttering these words, you're guilty of abusing language. In fact, if you pay careful attention to how the word 'moves' is used in ordinary contexts, as we have done, you'll soon see that 'moves' is only used to speak about objects that change position relative to the Earth. This linguistic data suggests that 'x moves' can only be coherently applied in contexts

in which x is moving relative to the Earth. Now you come along with your Sun-centered picture of the Solar system and you want to claim that Earth moves. But given our linguistic analysis, this can only mean 'Earth changes position relative to the Earth' and obviously this is nonsense. We see no option, then, but to reject your Copernican views as meaningless."

Clearly this kind of argument cannot be justified and if this were a valid use of the methodology that I have been defending, it would serve as the basis of a very good objection to that methodology. Fortunately for us, our fictional religious friends have failed to fully appreciate the foundations of this form of philosophy. But before I explain where they have gone wrong, I want to flesh out the bogeyman that this example is supposed to represent.

The fear captured in this example is that of philosophy as an unjustified impediment shackling the progress of science. Galileo's claim is straightforwardly empirical and can only be proven false by empirical investigation, not liturgical lounge-chair linguistic legislation. The religious leaders of our fanciful example are guilty of closing their minds to important truths, obstructing their cognitive capacities with rhetorical absurdities. Thus, if we desire science and philosophy to advance unfettered, we should avoid armchair analysis, even when under the guise of fancy philosophical methodology. Since, it seems, ordinary language philosophy represents precisely this sort of un-empirical *a priori* conceptual impediment, we must reject it.

The first thing we must say in response to this concern is that, as I have described it, ordinary language analysis is most assuredly not un-empirical. It has a

crucial empirical component as I laid out in chapter 1. The underlying insight of ordinary language philosophy is that words have criteria of use; if a word is meaningful, it cannot be applied indiscriminately. Ordinary language philosophy is, in part, the attempt to discover what those criteria are for given terms. I think it will be granted by even the most ardent opponent that attending to how people actually use a term is at least one method for doing so.[49]

So how should a Galileo who is committed to ordinary language philosophy respond to his religious interrogators? Galileo ought, if he is smart, to accept the linguistic analysis offered by the defenders of the faith. It seems accurate after all. But, more to the point, it is neither his prerogative nor his interest to dispute it. What Galileo is suggesting is not a violation of the meaning of 'moves' but a slight adjustment of it. Galileo ought to say, "Indeed, up until now, I would grant that your account of the meaning of 'moves' has been accurate and that it has adequately served our needs for centuries. But recent astronomical discoveries have made it clear that, if that is what 'moves' does mean, we must alter it. I am suggesting that we make it into a two-place relation. Rather than 'x moves' being 'x changes position relative to the Earth,' I am proposing that we take 'x moves' to mean 'x changes position relative to y' where y is any other object whatsoever. Thus, when I say that the Earth moves, I am not asserting the absurd 'Earth changes position relative to the Earth' but the much more comprehensible, not to mention physically accurate, 'Earth changes position relative to the Sun.' This alteration has the dual advantages of being one that leaves us

---

[49] Again, the restriction to ordinary (non-philosophical) contexts is a methodological check preventing the results from being biased by philosophical theories which themselves make assumptions about how the term(s) under investigation can be applied.

with a word whose meaning is not different in any significant way and also, thereby, being one that will be readily understood even by those unversed in the field of astronomical science. ”

What Galileo has done (in our fictitious example) is simply altered (and really only slightly) the criteria of application of the word 'moves.' The key to his success lies in the fact that he provides us with readily accessible criteria with which we can judge new uses of the term as correct or incorrect. Thus the term will not be used nonsensically and also will not be used indiscriminately, thus preserving its utility.

Why have I belabored this point? Well, first because I do not think that ordinary language philosophy can be found guilty on the charge of being irrationally and destructively conservative. We must always be open to semantic changes to terms of philosophical import. But when those changes occur they must be readily describable. It is precisely when we begin to use a term in a manner that stretches its ordinary criteria of application that they get themselves into trouble. Case in point:

In their recent book, *Philosophical Foundations of Neuroscience*, M.R. Bennett and P.M.S. Hacker describe (as one example in which philosophers, neuroscientists, psychologists, and cognitive scientists are philosophically confused) a kind of mereological fallacy. They describe multiple instances in which such scholars attribute to the brain properties or processes that are only correctly applied to the person as a whole. They quote Fancis Crick, in his *Astonishing Hypothesis*, saying, "What you see is not what is *really* there; it is what your brain *believes* is there. . . . Your brain makes the best interpretation it can according to its previous experience

and the limited and ambiguous information provided by your eyes."[50]. They quote

Colin Blakemore, "We seem driven to say that such neurons have knowledge. They

have intelligence, for they are able to estimate the probability of outside events."[51]

David Marr, the pioneer in research into visual processing: "our brains must somehow

be capable of representing . . . information."[52] They quote many other scientists and

philosophers who attribute to the brain other psychological phenomena such as

thinking, hypothesizing, deciding, hearing, etc.

Here is what Bennett and Hacker have to say about such attributions:

> It is our contention that this application of psychological predicates to
> the brain *makes no sense*. It is not that, as a matter of fact brains do not
> think, hypothesize and decide, see and hear, ask and answer questions;
> rather it makes no sense to ascribe such predicates *or their negations* to
> the brain. The brain neither sees, *nor is it blind*—just as sticks and
> stones are not awake, *but they are not asleep either. . . . The brain is
> not a logically appropriate subject for psychological predicates*. Only
> a human being and what *behaves* like one can intelligibly and literally
> be said to see or be blind, hear or be deaf, ask questions or refrain from
> asking." (Bennett and Hacker 72)

I believe that their analysis is essentially correct and their point is very similar

to the point that I made in chapter 1 concerning the theory that all that we can ever see

is a part of our own brain. When taken literally, all of these claims are absurd at best.

When Crick says that the brain interprets the data provided by the eyes, he is not

making the astounding claim that he has met brains and that they have reported having

difficulty finding the right interpretation of the ambiguous data provided by the eyes.

---

[50] Quoted in Bennett and Hacker, p. 68.
[51] ibid, p.69
[52] ibid, p. 70.

He is trying to describe processes he and other scientists have observed in the brain and making the philosophical assumption that 'interpretation' can be accurately applied to what the brain is doing. Similarly, Blakemore has not interviewed or interrogated neurons to find out what knowledge they have. He is, at best, groping for a means to describing certain properties that neurons display. But such groping is hardly helpful and can actually be harmful if we take what are most likely (at least originally) metaphorical statements literally. They provide the illusion of understanding a process that in reality remains unexplained. When I know that the brain interprets data, I don't really know anything because I do not know how it can be that a brain interprets. I wonder if it has similar trouble interpreting the poetry of e.e. cummings. When I know that a neuron has knowledge, I don't know anything because I wonder how we can test whether it has knowledge, I wonder where it learned this knowledge. Who were its teachers? Does it forget? Does it have short-term and long-term memory? Does it have temporal lobes? I understand what it means for a person to have knowledge, but I do not understand what it is for a neuron to have knowledge. Words like 'interprets', 'believes', and 'thinks' have a range of application and this range does not readily admit application to objects like brains and neurons.

Bennett and Hacker confront the charge with which I began these remarks. They have this to say:

> It is not semantic inertia that motivates out claim that neuroscientists are involved in various forms of conceptual incoherence. It is, rather, the acknowledgement of the requirements of the logic of psychological expressions. Psychological predicates are predicable only of a whole

animal, not of its part. No conventions have been laid down to determine what is to be meant by the ascription of such predicates to a part of an animal, in particular to its brain. So the application of such predicates to the brain or the hemispheres of the brain transgresses the bounds of sense. (78)

To be fair to the scientists I mentioned above, I should point out that I do not know that they have not, à la our fictional Galileo, offered a compelling reason for using these psychological predicates and an adequate description of the altered criteria of application. But there can be no doubt (to use an expression famously misused by Richard Cheney) that, in their ordinary usage, psychological predicates are attributable only to the whole animal. We no more understand what it would mean for the brain to hear information, or for the brain to interpret data, or for the brain to make a decision, than we understand what it would mean for the brain to go on vacation, or look for a job, or write a dissertation. If there is some compelling reason to use such expressions to adequately and accurately describe the workings of the brain and its parts, then it is incumbent upon those who would suggest such usage to fully articulate the rules governing the novel application.

I conclude that ordinary language philosophy is not inherently conservative. Semantic change is a very real phenomenon. Ordinary language analysis can not only reveal it but also provide a necessary check on those who would extend the range of application for some term(s) without providing us with an understanding of how that range is to be delimited.

I'll close with the following more recent story of astronomical discoveries affecting our conceptual scheme: As I write over 2,500 astronomers are meeting in

Prague to debate changes to the definition of 'planet.' Recent discoveries of large planet-like objects beyond the orbit of Pluto, of similar size to, or even larger than, Pluto, have stretched our current definition of 'planet.' Such objects include an object known as Sedna (discovered at the Palomar observatory in November 2003, about 3/4 the size of Pluto, and three times as distant as Pluto), and the less memorably named 2003 UB313 (also discovered at Palomar, larger than Pluto, and even more distant than Sedna) among others. And of course we can expect similar discoveries in the years ahead. There is much internal debate among astronomers as to how these objects are to be classified. Are they planets, planetoids, or members of some as yet undefined class of astronomical body? These objects pose a challenge to our concept of a planet and the meeting in Prague is meant to answer this challenge.

Caltech astronomer Michael Brown, a member of the team that discovered Sedna and 2003 UB313, has a website (www.gps.caltech.edu/~mbrown) devoted to describing these objects, explaining how they were found, and addressing common questions about them, including the question of what counts as a planet. About this question, he says, "Astoundingly, no precise scientific definition of the word "planet" currently exists. It is rare for scientists to have to define a word that is already in common usage and that everybody from school children on up already understands. How does one then go about constructing a scientific definition of such a word after the fact?"[53] He considers four possible definitions of 'planet.' They are: first, what he calls the "purely historical" version, according to which all and only the nine

---

[53] http://www.gps.caltech.edu/~mbrown/sedna/index.html#planets

traditional planets (Mercury, Venus, Earth, Mars, Jupiter, Saturn, Uranus, Neptune, and Pluto) are planets; second, the "historical plus" version: all of the above traditional planets plus any object larger than Pluto; third, "gravitational rounding": any object that is round because of its own gravitational pull and orbits the sun is a planet; and finally a definition that takes advantage of the fact that objects in the solar system are (at least very often) either solitary individuals or members of some larger population (Brown calls this definition, "population classification").[54]

Brown also has some very interesting things to say about the relative worth of each of these definitions. The first two options (the historical ones), he claims, have the advantage of coinciding with how we have spoken historically, but neither is acceptable from a scientific point of view. The lines drawn by these definitions are totally arbitrary. If we maintain the purely historical definition, we must reject objects like Sedna and 2003 UB313 even though Sedna is nearly the same size as Pluto and 2003 UB313 is larger. The historical plus option is no better. Why, Brown asks, is Pluto the cutoff size? "Is there really a big enough difference in size between Pluto and Sedna and Quaoar [another Kuiper Belt object] that one should be called a planet while the others are not? The scientific answer remains a resounding no"[55]

The gravitational rounding version looks appealing at first but when you consider that there are a multitude of such objects in the Solar System that are not considered planets, it is less attractive. Brown says,

---

[54] ibid.

[55] ibid.

Unfortunately, this definition completely fails the historical sanity check. Historically, where does the criterion to be round come from, except for the near coincidence between the historical definition of planet and the transition size from round to not round? At no time in previous history has any discussion of whether or not an object is round been part of the discussion of whether or not it should be called a planet. Ceres [a member of the asteroid belt] *was* initially considered to be a planet, but not because it is round (which was unknown at the time), but because it was the only object known to exist between Mars and Jupiter. When other asteroids of similar sizes were found at nearly the same location it was decided to call them all members of the asteroid belt, rather than planets.

Roundness is an important physical property, and gravity is the dominant force in the solar system, so perhaps it *is* important to have a special word which describes the class of objects in the solar system which are round. But simply because all historical planets are round does not at all mean that it is good science to define all round objects to be planets. A much better idea is to use a different word to describe these objects. Spheroids? Gravispheres? Actually, we prefer the word "planetoid" as a new word to describe round objects orbiting the sun. All planets are planetoids. Not all planetoids are planets.[56]

The final option, population classification, is, claims Brown, more scientifically motivated. There is a clear distinction between solitary objects and objects that are members of a larger group. Most of the traditional planets (except Pluto) are solitary individuals and there are clearly objects that are part of large populations (such as the asteroid belt). Pluto is a member of what is known as the Kuiper Belt (a group of objects beyond the orbit of Neptune). Here is what Brown says in favor of this definition:

Let's examine this definition in more details. First, it is certainly scientifically motivated and well-founded. But so was the "gravisphere" definition above. Is there any historical basis for saying that a planet is a solitary individual that is not a member of a large population? Yes! As mentioned earlier, historically Ceres and the first few asteroids were

---

[56] ibid

initially classified as planets. Only when it became known that there were many many asteroids in similar orbits was it decided that they should no longer be classified as planets. Historically, there is a clear distinction between planets and populations. Any definition which fails to make this distinction is in strong trouble on historical grounds. This simple look at history shows that Pluto is completely analogous to Ceres. Pluto was initially thought to be a solitary individual. Over time we found more objects in the vicinity and realized instead that it is a member of a large population. Historically, then, Pluto, too, should no longer be considered a planet.[57]

The problem, of course, is this exclusion of Pluto. The fact that Pluto is a planet is so well entrenched in our minds that it would be difficult for people to accept a definition that excludes it. More recently Brown wrote the following,

In my view scientists should not be trying to legislate an entirely new definition of the word "planet." They should be trying to determine what it means. To the vast majority of society, "planet" means those large objects we call Mercury through Pluto. We are then left with two cultural choices. (1) Draw the line at Pluto and say there are no more planets; or (2) Draw the line at Pluto and say only things bigger are planets. Both would be culturally acceptable, but to me only the second makes sense for what I think we mean when we say the word planet. In addition, the second continues to allow the possibility that exploration will find a few more planets, which is a much more exciting prospect than that suggested by the first possibility. We don't think the number of planets found by the current generation of researchers will be large. Maybe one or two more. But we think that letting future generations still have a shot at planet-finding is nice.[58]

I have spent this time reporting Brown's views and quoted him so extensively because this issue directly relates to the claims I have been making in this dissertation in two important respects. First, the process of gathering scientific authorities together to decide on a universal definition of a planet is very similar to that which I described

---

[57] ibid

[58] www.gps.caltech.edu/~mbrown/planetlila/

in relation to the term 'meter' in chapter 2. In this case, of course, astronomers are not introducing a new term, but they will, nevertheless, be legislating the future use of the term 'planet.' And just as there was no particular decision that was forced on eighteenth century authorities as to what length 'meter' should refer to; there is today no particular decision that is forced on the astronomers gathered in Prague as to what the definition of 'planet' should be (though, obviously, as Brown's thoughts indicate, there will be many factors, scientific and cultural, that will constrain the range of possible choices). Until a decision is reached, there really is no fact of the matter whether Sedna or 2003 UB313 are planets, since on some proposed definitions they are planets and on some they are not. Their status ultimately hinges on a decision. It is not metaphysically odd or philosophically naive to say this; it is an accurate description of our current state of affairs. The final vote on the new definition will not take place until the end of the current month (August 2006) and we will have to wait until then to know the fate of these objects (and also of Pluto). What is interesting is that all of the relevant scientific facts are known (we know a great deal about the sizes, distances, and orbital paths of these objects); the purpose of the meeting in Prague is not to debate the veracity of scientific findings. The purpose is to adopt a new definition and to thereby articulate new, more rigorous criteria for the application of the term 'planet.'

Lest one think that Brown's opinions are those of a philosophically unsophisticated minority, remember that he is one of the lead members of the team responsible for the discovery of many of these objects. He is indisputably a scientific

authority. In addition, other astronomers gathering at Prague can be heard making similar statements. Here are a few relevant quotes taken from media reports about the event:

"My feeling is that they would like to re-classify Pluto. One idea would be to make Pluto an "ice dwarf". This would be a new subdivision of the planet family, rather like the asteroids," said Peter Bond of the UK's Royal Astronomical Society.[59]

"It could mean the number of planets leaps to 20 or more, or it drops to eight. But I think most people would prefer not to drop Pluto altogether," Mr Bond told the BBC News website.[60]

Among the possibilities at the 12-day meeting of the International Astronomical Union in the Czech Republic capital: Subtract Pluto or christen one more planet, and possibly dozens more. But the decision won't be an easy one. Scientists attending the conference are split over whether Pluto should be excluded from the list of planets, Pavel Suchan of the meeting's local organization committee said. "So far it looks like a stalemate," Suchan said. "One half wants Pluto to remain a planet, the other half says Pluto is not worth being called a planet."[61]

A decision on whether Pluto should be excluded or if "Xena" should be included on the list of planets will not be known before the end of the conference, Suchan said. "We of course need the definition of a planet first."[62]

Again, it is worth repeating that the decision on the new definition will not be based upon new scientific discoveries. Evidence about the composition, location, shape, size, distance, orbital trajectory or any other facet of these objects will not be presented that will force one definition over another. If this were the case, then the meeting in Prague would certainly be too hasty and should be cancelled until all the

---

[59] news.bbc.co.uk/2/hi/science/nature/4789531.stm

[60] ibid

[61] www.cnn.com/2006/TECH/space/08/14/planet.meeting.ap/index.html

[62] ibid

relevant Solar system objects are more thoroughly and accurately catalogued and described.

The second way in which this debate touches on the issues surrounding this dissertation is evident in Brown's claim that astronomers should be trying to discover what 'planet' means rather than trying to legislate an entirely new definition. He displays a great deal of concern for how the word 'planet' has historically been used and not just by astronomers but by non-scientists as well. In his view the way the term has ordinarily been used has at least as much importance as the scientific concerns he addresses. In fact, I suspect that astronomers wouldn't be in this mess if they were not interested in how the non-scientific community uses the term; we could just demote Pluto and be done with it. At the very least the decision would be much more straightforward if such concerns about historical usage held no sway. But I think it is both commendable and realistic for astronomers to voice such worries. If science is to be comprehensible by non-experts, the terms used in scientific discourse must have strict and compelling criteria of application and, importantly, must, to as large an extent as possible (and especially when using a term that has broad usage across the population, both scientific and non-scientific), connect with the common, ordinary understanding of what the terms mean.

BIBLIOGRAPHY

Alder (2002) *The Measure of all Things*. New York: The Free Press.

Austin, J. L. (1946) "Other Minds" in his *Philosophical Papers,* (1961), Oxford: Clarendon.

Austin, J.L. (1966) "Three Ways of Spilling Ink" *The Philosophical Review, vol. 75 No. 4,* reprinted in his *Philosophical Papers.*

Austin, J.L. (1961) *Philosophical Papers*. Oxford: Clarendon.

Austin, J.L. (1961) "The Meaning of a Word" in his *Philosophical Papers*, Oxford: Clarendon.

Austin, J.L. (1962) *Sense and Sensibilia*. Oxford: Oxford University Press.

Austin, J. L. (1962) *How To Do Things With Words*. Cambridge, Mass: Harvard University Press

Ayer, A. (1940) *The Foundations of Empirical Knowledge*. London: Macmillan and Co.

Ayer, A. (1950) *Language, Truth and Logic*. New York: Dover.

Baker, G.P. and Hacker, P.M.S (1985) *Wittgenstein: Rules Grammar and Necessity, Volume 2 of an Analytic Commentary on the* Philosophical Investigations. Oxford: Blackwell.

Bennet, M.R. and Hacker, P.M.S (2003) *Philosophical Foundations of Neuroscience*. Oxford: Blackwell.

Boghossian, P. (1997) "Analyticity" in *A Companion to the Philosophy of Language*, B. Hale, C. Wright eds. Oxford: Blackwell.

Carnap, R. (1947) "Empiricism, Semantics, and Ontology" in *Meaning and Necessity: A Study in Semantics and Modal Logic*. Chicago: The University of Chicago Press.

Caton, C. ed. (1970) *Philosophy and Ordinary Language*. Chicago: University of Illinois Press.

Cavell, S. (1958) "Must We Mean What We Say" *Inquiry*, I 172-212.

Cavell, S. (1965) "Austin at Criticism," *The Philosophical Review*, LXXIV: 204-219.

Cavell, S. (1979) *The Claim of Reason*. New York: Oxford University Press.

Chalmers, D. (1996). *The Conscious Mind: In Search of a Fundamental Theory*. New York: Oxford University Press.

Chisolm, R. (1951) "Philosophers and Ordinary Language" *The Philosophical Review* LX, 317-28.

Churchland, P.M. (1979) *Scientific Realism and the Plasticity of Mind*. Cambridge: Cambridge University Press.

Churchland, P.M. (1985) "Reduction, Qualia, and the Direct Introspection of Brain States," *Journal of Philosophy*. 82, no.1, (1-22).

Churchland, P. M. (1998) "The Rediscovery of Light." In *On the Contrary*, Churchland and Churchland, Cambridge: The MIT Press

Cook, J. (1965) "Wittgenstein on Privacy" *Philosophical Review* volume 74.

Dummett, M. (1979) "What does the Appeal to Use Do for the Theory of Meaning" in his *The Seas of Language* (1993) Oxford: Oxford University Press.

Fann, K.T. (1969) *Symposium on J. L. Austin*. London: Routledge and Kegan Paul.

Fodor, J. (1961) "Of Words and Uses" *Inquiry*, IV (1961), 190-208

Fodor, J. and Katz, J. (1962) "What's Wrong with the Philosophy of Language," *Inquiry*, V 197-237.

Frege, G. (1892). "On Sinn and Bedeutung." In *The Frege Reade*r (1997), Michael Beaney, ed. Oxford: Blackwell Publishers Ltd.

Gert, H. (2002) "The Standard Meter by Any Name is Still a Meter Long" in *Philosophy and Phenomenological Research* Vol. LXV, No. 1, pp. 50-68.

Grice, H.P. (1957) "Meaning," *Philosophical Review*, LXVI 377-388.

Grice, H.P (1989) *Studies in the Ways of Words*. Cambridge, Mass.: Harvard University Press

Jackson, F. (1997) "What Mary Didn't Know" in N.Block, O. Flanagan, G. Guzeldere (eds.) *The Nature of Consciousness: Philosophical Debates* (1997) Cambridge, Mass. : The MIT Press.

Jolley, K. (1998) "On Common Sense, Moore and Wittgenstein" *Enthographic Studies*, No.3 (Autumn 1998), pp. 41-58.

Kant, I. (1929) *Critique of Pure Reason*. translated by Norman Kemp Smith. New York: St. Martin's Press.

Kober, M. (1996) "Certainties of a World-Picture: the Epistemological Investigations of *On Certainty*" in *The Cambridge Companion to Wittgenstein*, H. Sluga and D. Stern eds. Cambridge: Cambridge University Press.

Kripke, S. (1971) "Identity and Necessity." New York University.

Kripke, S. (1972) *Naming and Necessity*. Cambridge, Mass: Harvard University Press.

Levine, J. (1983) "Materialism and Qualia: The Explanatory Gap," reprinted in T. O'Connor, D. Robb (eds.) *Philosophy of Mind: Contemporary Readings* (2003). London: Routlegde.

Luckhardt, C. G. (1978) "Beyond Knowledge: Paradigms in Wittgenstein's Later Philosophy" *Philosophy and Phenomenological Research*, Vol. 39, No. 2., pp. 240-252.

Malcolm, N. (1942) "Certainty and Empirical Statements," *Mind*, LI

Malcolm, N. (1942) "Moore and Ordinary Language" in *The Philosophy of G. E. Moore*, P.A. Schilpp, ed. La Salle, IL: Open Court.

Malcolm, N. (1949) "Defending Common Sense," *Philosophical Review*, LVII 201-220.

Malcolm, N. (1952) "Knowledge and Belief," *Mind*, LXI 178-189.

Malcolm, N. (1958) "Knowledge of Other Minds," *Journal of Philosophy*, LV 969-978.

Malcolm, N. (1972) "Thoughtless Brutes," *Proceedings and Addresses of the American Philosophical Association*, XLVI 5-20.

Malcolm, N. (1995) *Wittgensteinian Themes*. Ithaca: Cornell University Press.

Maxwell, G. and Feigl, H. (1961) "Why Ordinary Language Needs Reforming," *Journal of Philosophy*, LVII 488-498.

Moore, G.E. (1959) *Philosophical Papers*. London: Allen and Unwin.

Moyal-Sharrock (2001) "Logic in Action: Wittgenstein's *Logical Pragmatism* and the Impotence of Scepticism.

Nagel, T. (1974) "What is it Like to Be a Bat?" reprinted in his *Mortal Questions* (1979), Cambridge: Cambridge University Press.

Putnam, H. (1962) "The Analytic and the Synthetic," in his *Mind Language and Reality, Philosophical Papers Volume 2* (1975), Cambridge: Cambridge University Press.

Putnam, H. (1999) *The Threefold Cord: Mind, Body, and World*. New York: Columbia University Press.

Quine, W. V. O. (1953). "Two Dogmas of Empiricism" in his *From a Logical Point of View* (1953), Cambridge, Mass: Harvard University Press.

Russell, B. (1905) "On Denoting," in Martinich, ed. *The Philosophy of Language*, Oxford: Oxford Univeristy Press.

Ryle, G. (1949) *The Concept of Mind*. Chicago: The University of Chicago Press.

Ryle, G. (1953a) *Dilemmas*. Cambridge: Cambridge University Press.

Ryle, G. (1953b) "Ordinary Language" *The Philosophical Review*. LXII, 167-186. Reprinted in Caton (1970) *Philosophy and Ordinary Language*. Chicago: University of Illinois Press.

Ryle, G. (1957) "The Theory of Meaning" in Caton. (1970) *Philosophy and Ordinary Language*. Chicago: University of Illinois Press.

Salmon, N. (1993) "Analyticity and Apriority" *Philosophical Perspectives*, Vol. 7, Language and Logic, 15-133.

Searle, J. (1969) *Speech Acts*. Cambridge: Cambridge University Press.

Searle, J. (1979) *Expression and Meaning*. Cambridge: Cambridge University Press

Searle, J. (1989) "How Performatives Work" in his *Consciousness and Language* (2002), Cambridge: Cambridge Universtity Press.

Searle, J.   (1992) *The Rediscovery of the Mind.* Cambridge, Massachusetts: The MIT Press.

Searle, J. (1995) *The Construction of Social Reality*. New York: The Free Press.

Searle, J. (2002) *Consciousness and Language*. Cambridge: Cambridge University Press.

Sher, Gila, (1999) "Is There a Place for Philosophy in Quine's Theory?" *Journal of Philosophy*, 96: 491-524.

Stalnaker, R. (2003) *Ways a World Might Be*. Oxford: Oxford University Press.

Stern, D. (1995) *Wittgenstein on Mind and Language*. New York: Oxford University Press.

Strawson, P.F. (1950) "On Referring," *Mind* 59: 320-344.

Stroll, A. (1988) "Wittgenstein and Folk Psychology," in *Philosophy, Law, Politics, and Society: Proceedings of the 12th International Wittgenstein Symposium, August 1987,* 264-270. Vienna: Holder-Pichler-Tempsky

Stroll, A. (1988) *Surfaces*. Minneapolis: University of Minnesota Press.

Stroll, A. (1994) *Moore and Wittgenstein on Certainty*. New York: Oxford University Press.

Stroll, A. (1998) *Sketches of Landscapes: Philosophy by Example*. Cambridge, Mass: MIT Press.

Stroud, B. (1968) "Transcendental Arguments." *The Journal of Philosophy* LXV, no.9: 241-256.

Stroud, B. (1984) *The Significance of Philosophical Scepticism*. Oxford: Oxford University Press.

Unger, P. (1999) "The Mystery of the Physical and the Matter of Qualities: A Paper for Professor Shaffer" in Tim O'Conner and David Robb (eds.), *Philosophy of Mind: Contemporary Readings* (2003), London: Routledge.

Unger, P. (1975) *Ignorance: A Case for Skepticism*. Oxford: Oxford University Press.

Wittgenstein, L. (1922) *Tractatus Logico-Philosophicus*. London: Routledge.

Wittgenstein, L. (1958a). *Philosophical Investigations*. Englewood Cliffs, NJ: Prentice Hall.

Wittgenstein, L. (1958b) *The Blue and Brown Books*. New York: Harper and Row.

Wittgenstein, L. (1967). *Zettel*. Berkeley: University of California Press.

Wittgenstein, L. (1969). *On Certainty*. New York: Harper and Row Publishers.

Wittgenstein, L. (1993). *Philosophical Occasions 1912-1951*. Cambridge: Hacket.