

Chapter Two

THE CONCEPT OF KNOWLEDGE

When, therefore, anyone forms the true opinion of anything without rational explanation, you may say that his mind is truly exercised, but he has no knowledge; for he who cannot give and receive a reason for a thing, has no knowledge of that thing; but when he adds rational explanation, then, he is perfected in knowledge.

--Plato

1. Definitions and Word Games

Suppose that we are concerned with the question of economic justice -- the fact that a few are ridiculously wealthy, while many are pitifully poor. We might convene an academic conference to discuss the issue and suggest some sort of coherent social policy. Economists might tell us about how income distribution is empirically related to national productivity. Political scientists might tell something about relative tax rates and the amount of government services. Sociologists could address the social effects of long-term poverty. Historians could give us some sense of whether the problem is better or worse than it was a hundred years ago. It would not be at all surprising if a philosopher contributed a paper on the meaning of economic justice. In one way such a contribution seems necessary and foundational. After all, how can we reasonably construct some social policy aimed at greater economic justice, if we are not crystal clear as to what we mean by this

concept? In another light, however, the philosopher's contribution seems frivolous, and even counter-productive. If there is wide agreement that there is a problem that needs to be solved, the philosopher's concern with long dead thinkers like Plato, Adam Smith and Marx may strike us as an irresponsible waste of time and intellectual energy. To carry this example just a bit further, suppose the philosopher's paper offers a definition of economic justice that suggests some kind of tension with other widely held values and social policies, and goes so far as to suggest that we will never have a concept of economic justice that everyone will feel comfortable with. Now the philosopher's concern with theory and the definition of terms may strike us as subversive. It may be difficult and controversial to articulate a theory about the nature of economic justice that everyone will agree to. Nevertheless, we know injustice when we see it. And to suggest that we spend our time defining terms and teasing out subtle philosophical arguments, rather than offering constructive solutions to the obvious problems that plague our society is both dangerous and immoral. But all of this is quite unfair. No sane philosopher is going to suggest that we spend all of our time and energy in academic theoretical pursuits. Obviously, there are crises that call for immediate action, and we all recognize the need to make decisions on less than perfect information. But there is also a need for abstract theoretical work. It does seem crazy to propose significant social changes that will effect all of us without some kind of clear understanding of what we are trying to bring about. Pausing to reflect on the nature of economic justice -- defining our terms, as they say -- may be worthwhile even in a time of some urgency.

Please excuse the above digression. I have included it because I believe that many beginning students see much of traditional epistemology in the same uncharitable light that our philosopher was portrayed. Every reader of this book is a mature speaker of English. The verb, 'to know', and the abstract noun, 'knowledge', are fairly normal words within the English language. Obviously, we must know what they mean. We will discover, however, that it proves exceedingly difficult to articulate a clear and coherent definition, or theory, of knowledge.

2. The Myth of Definition

This chapter discusses the prospects for offering a helpful analysis, or definition, of the concept of knowledge. We need to take a little time dispelling a common misunderstanding about the importance of

definition in everyday contexts, as well as philosophical contexts. It is widely believed that people do not know the meaning of the words they use -- they do not know what they are talking about -- unless they can provide adequate definitions for all of those words. This is simply a mistaken view of meaning.

Someone can be an excellent athlete, a hitter in baseball, for example, yet be a very poor coach or teacher of how to hit. Surprisingly, perhaps, others can be mediocre hitters, but turn into outstanding hitting coaches. The reason these things are possible is that there is all the difference in the world between doing something, and describing or explaining how to do something. Think for a moment about those things that you are most skilled at doing -- shooting free throws, playing a musical instrument, riding a bicycle, etc. How confident would you be that you could teach someone else how to be skillful at this? Could you write a manual for them on how to do it?

Speaking a language is much more like hitting a baseball than being a good little league coach. Language is skillful activity that human beings master with remarkable facility, in ways that philosophers, psychologists, and linguists are only beginning to appreciate. I can safely assume that any reader of this book is an accomplished enough user of English that he or she knows full well the meaning of almost every word that philosophers have spent a great deal of time and energy trying to analyze or define. You all know the meaning of terms like 'beauty', 'justice', and 'knowledge', because you can use sentences like the following to communicate with other English speakers.

1. That's a beautiful painting.
2. Not to let Sarah play would be unjust.
3. You don't really know that the Dodgers will win the pennant, you just hope they will.

All of this is important because it is so easy to forget in the middle of philosophical battles. We are going to analyze the concept of knowledge in this chapter. We will see that this task is difficult, controversial, and perhaps in the end impossible to complete satisfactorily. This doesn't mean for a second that you, or the great minds of western philosophy, do not know how to use words like 'know' and 'knowledge' for the purposes of clear communication.

3. *The Need for Conceptual Clarity*

Although I stand one-hundred percent behind what I said above, it would be a mistake to infer that careful conceptual analysis is not important. People sometimes make remarkable claims about knowledge. We have just seen how the skeptic can put together plausible and disturbing arguments that we know next to nothing. The arguments of the last chapter are classical examples of the sorts of intellectual concerns that occupy the attention of professional philosophers. Intellectual disputes about knowledge are not limited to philosophers, however. We often hear that modern scientists do not know that evolution by natural selection is true. Many claim that it is only a "theory." This is sometimes backed up with an argument. Science, so this line of thinking goes, is only concerned with what can be directly observed or proved with laboratory experiments. But evolution cannot be directly observed, both because it is too slow of a process, and because the most interesting observations would have needed to take place in a time before there were human observers. Furthermore, creationists claim that no controlled laboratory experiment can prove that evolution is true.

If we are to make any progress in understanding, let alone resolving, these kinds of intellectual disputes, we are going to need to be much clearer in our own minds as to what counts as knowledge. I claim to know that I am at my computer composing this chapter. The skeptic tells me I don't know this after all, it might only be a dream. I believe that modern biology knows that natural selection is true. Creationists claim that they don't, and that their "faith" in the theory is no different than religious belief. How can we possibly hope to make progress toward resolving these disputes without some fairly specific agreement as to what counts as genuine knowledge?

For some, the kind of conceptual analysis we are engaged in in this chapter can be fun and exciting in its own right. Most of you, however, should see it as a necessary means to an end. I assume most of you care about whether scientists know what they are talking about. If you are like I am, you think they probably do. But to really feel confident about this you need to have some answers to the philosophical skeptic who says it might all be a dream, and the methodological skeptic who argues from a specific model of scientific knowledge. To answer either of these skeptics productively, you need some agreement about the nature of knowledge.

4. Knowledge and Belief

Human beings seem to be a very credulous species, we believe an amazing verity of things. Our ancestors believed in witches, that the earth was flat, and in the divine right of kings. People today believe that their futures are foretold in horoscopes, that good writing can be accomplished in first drafts, and a few diehard fans believe that the Cubs will finally get it together. From the perspective of history it is easy to find countless beliefs that we sincerely held that strike us as foolish, dangerous, and immoral. But, of course, not all beliefs fit into this category.

Other things we don't *merely* believe, we know. I, of course, believe that I am a philosophy professor, an aging softball player, and a movie freak. But, I don't just believe these things, I know them. The distinction between belief and knowledge is not like the one between being a sibling and being an only child -- it is not an exclusive, either/or difference. It is rather like the distinction between an automobile and a convertible. To be a convertible is to be a special kind of automobile. As logicians put it, being an automobile is a *necessary* condition of being a convertible. Not all automobiles are convertibles, but all convertibles are automobiles.

Traditional models, or definitions, of knowledge have attempted to articulate a list of necessary conditions that are jointly *sufficient* for having genuine knowledge. Since the abstract noun, 'knowledge,' is kind of artificial, these models have usually been articulated using the more familiar verb. Our observations about knowledge and belief suggest the first entry on our list of necessary conditions.

***J* knows *P* only if:**

i. *J* believes *P*.

There is a fairly common expression that seems to call this into question. Suppose we have a friend who is headed for heartache partly because he refuses to take seriously the obvious evidence of his lover's infidelity. We might say, "Jake knows that she's untrue, but he can't bring himself to believe it." Or, perhaps, we know someone who is foolishly refusing to take heed of medical symptoms. "Sarah knows something is wrong, but just won't believe it." How seriously should we take the claim that both Jake and Sarah have knowledge, but lack belief? Not very.

Jake sees the obvious signs, and has his moments of doubt. Sarah, too. If they didn't we wouldn't be inclined to say they knew. It is, of course, possible for people to be perversely dense. People can be totally oblivious to things that are perfectly obvious to others. Jessica may genuinely believe that her lover is totally faithful, despite the lame excuses and the lipstick on the collar. But we would never be tempted to say Jessica knows, though perhaps she should. When we use the above idiom we are getting at something interesting about Jake and Sarah. They seem to be engaging in what philosophers call self-deception. This is an important issue in both philosophy and psychology, but really says nothing about how to define knowledge.

I take it to be settled that knowledge implies some kind of genuine conviction or intellectual confidence. Thus, the first necessary condition of knowledge turns out to be relatively secure, uncontroversial, and philosophically straightforward. Would that we could say the same about the conditions to follow.

5. In Search for the Truth

You are the District Attorney and you've got a great case. The defendant is the kind of low-life that society needs to do something about. You've got the goods on him too. Lots of physical evidence, a clear motive, and witnesses. The case will be an easy one to try, and it will be a feather in your cap to be the one who put him away. You are convinced the slime ball's guilty. There's only one problem with this scenario; the guy didn't do it. It does not matter how sincere your belief is, nor how good the evidence seems to be -- if what you thought you knew turns out to be false, it's back to the drawing board. Truth is an absolute precondition for knowledge. Unfortunately, truth is a philosophical mess.

Here's the truth about truth, and it's not a pretty picture. Contemporary philosophy is about as far from consensus about the nature of truth as any issue in the field. Some believe that truth is *correspondence* with reality. Others that it is *coherence* with other widely held beliefs. Yet others claim that the assertion that "snow is white is true" is just a fancy way of saying that "snow is white." All of these theories of truth have plausible arguments in their defense, and all suffer from serious conceptual problems.

In spite of all of the confusion about the nature of truth, however, the relationship between truth and knowledge is as clear as could be. The only beliefs that we have that are viable candidates for being knowledge are those that are true. The surest way to defeat someone's

claim that they know something is to show that what they claim to know is false. This suggests a workable epistemological definition of truth.

truth =_{df} not not-false

Admittedly, this is a trivial definition. It does, however, have the advantage of separating philosophical disputes about the contentious nature of truth from the non-controversial connection between truth and knowledge.

Truth supplies a second necessary condition for knowledge. We can expand our evolving model of knowledge as follows.

***J* knows *P* only if:**

i. *J* believes *P*.

ii. *P* is true.

6. Epistemic Justification

Perhaps we already have all that we need. The concept of knowledge seems both subjective and objective -- it is something that individual subjects have or fail to have, but it is objective in the sense that the way the world is, independent of particular subjects, also plays a huge part. Condition (i) takes care of the subjective element, and (ii) covers the objective. What more do we need?

I have been hoping for a raise. Unfortunately, my latest evaluation left a lot to be desired, and the state's budget looks pretty bleak. Forever the optimist, I continue to think the best. I woke up yesterday and as I was having my morning coffee I glanced at my horoscope. The entry for Pisces was way cool.

You will receive something long overdue and well deserved. All the signs are positive.

My raise! What could be clearer? I went to work with a smile on my face absolutely confident that I would get the good news. And I did! The Governor decided that all state employees should get a modest salary adjustment, and that in the afternoon we were all formally notified.

The two conditions for knowledge are satisfied. Johnson believes that he will get a raise, and it is true that he will get a raise. He knows

he will get a raise. Most of us would be very reluctant to say this is knowledge. What he believes turns out to be true, but merely by coincidence or good luck. The subjective element of belief, and the objective element of truth seem much too tenuously connected. What seems to be missing is some reason or evidence in support of my belief. Sure, the horoscope is a reason in the sense of providing a psychological explanation for why I happen to have this belief. But, it's such a poor reason -- it's unreliable -- that we attribute the belief's truth to good fortune and not the strength of the reason.

Epistemologists have adopted the idiom of normative obligation to get at the stronger connection between belief and truth that is required for genuine knowledge. You are entitled to claim knowledge, according to this way of thinking about things, only if your belief is *justified*. That is, just in case you have *very good reason* for thinking it is true. Thus, on the so-called *standard analysis of knowledge* a third necessary condition of knowledge, one that completes the package and makes it *jointly sufficient*, is the justification condition.

***J* knows *P* if and only if:**

i. *J* believes *P*.

ii. *P* is true.

iii. *J* is justified in believing *P*.

7. What Does It Take to Be Justified?

We have seen how skeptics can produce a formidable battery of arguments designed to show that we are never completely justified in believing anything. The problem concerns the connection between truth and justification. The only standard that completely eliminates the possibility of our beliefs being held in error is one of self-evidence or certainty. But as the Cartesian project has convinced most of us, epistemological certainty is unattainable. This means that whatever model of knowledge is finally endorsed will be committed to some sort of epistemic fallibility. This is not that serious a worry for most natural or social scientists, but does run counter to the dominant tradition in western epistemology.

Self-evidence and certainty may have set unrealistically high standards for knowledge, but these epistemic standards had the superficial appearance of being clear and identifiable. Models of knowledge that substitute criteria for epistemic justification must be prepared to state some new criterion for distinguishing unfounded belief, from a promising theory, from established knowledge. The contemporary literature offers many intriguing possibilities -- some highly formal, and some quite commonsensical, but, none that have won anything approaching consensus.

I suggest that we understand the idea of epistemic justification in terms of evidence. The things that we *know* are those true beliefs for which we have very, very, very good evidence. Good evidence is something that we are all familiar with, and something that we can learn to reliably spot. I will be offering in the chapters to follow a model of -- or a kind of formula for testing for -- good evidence. I hope to convince you that this model captures everything we care about when we assess the quality of a person's evidence, or for that matter, their claims to knowledge.

Let's transform the *standard analysis of knowledge* in light of all of this into the following.

***J* knows *P* only if:**

i. *J* believes *P*.

ii. *P* is true.

iii. *J* has exceedingly good evidence for *P*.

8. An Unsolved Problem

If you were reading very carefully, you may have noticed a slight difference in the way I stated the standard analysis of knowledge at the end of Section 6, and at the end of Section 7. You are all smart enough to see the obvious change in condition (iii.), but can you find the other difference? The way the philosophic tradition has defined knowledge is to articulate necessary and sufficient conditions for knowing something. The standard analysis of knowledge claims that the three necessary conditions are, taken together, sufficient for knowing something. In my statement of a "transformed" analysis, I wimped-out a bit. I claimed

that my three conditions were all necessary -- that's what the *only if* signifies -- but, I left it open whether the three conditions were sufficient. Here's why.

Consider the following little thought-experiment. My wife and I have spent the last hour collaborating on our special spaghetti sauce. Just as we are getting ready to serve dinner we discover that we are out of Parmesan cheese. We divide responsibilities -- she will toss the salad and serve dinner; I'll make the emergency run to the store. While at the store I meet a colleague doing research in contemporary epistemology -- she wants an example of knowledge. I suggest that I know there is a spaghetti dinner sitting on our dining room table right now. And as luck would have it, it's true that a spaghetti dinner is on the table. I believe it; it's true; and I'm justified in believing it. All is well. Unfortunately not. After I left our German shepherd, Guido, got rambunctious and knocked the pot of simmering spaghetti sauce on the dirty kitchen floor. My wife considered violence to the dog, but before anything could happen a neighbor arrived with a pot of left over spaghetti sauce, announcing that she was leaving on vacation and it would surely spoil before she returned. Thus, the spaghetti sauce that made my knowledge claim true is unconnected to the spaghetti sauce that provided the justification for my belief. It is odd in the extreme to claim that I had knowledge of the pot of spaghetti sitting on my table. It is pure serendipity that my belief turned out to be true.

A lot of contemporary epistemology has been concerned with ruling out these kind of "Guido" cases. Many philosophers have suggested that some fourth, or fifth, or sixth, etc., condition must be added to our analysis of knowledge. I am not sure whether I personally agree, or not. To be on the safe side, however, I will be content with the above transformed analysis. The epistemic action in this little book will focus on condition (iii). What the heck is it to have evidence, or good evidence, or exceedingly good evidence, for something?