The Theory of Questions, Epistemic Powers, and the Indexical Theory of Knowledge

HECTOR-NERI CASTAÑEDA

The real danger of oversimplified models is not that they are over-simple, but that we may be satisfied with them, and fail to compare them with regions of experience other than those which suggested them.

> Wilfred Sellars, *The Structure of Knowledge*

A main cause of philosophical illness—one-sided diet: one nourishes one's thinking with only one kind of example.

> Ludwig Wittgenstein, *Philosophical Investigations*

Philosophical method has the anti-Augustinian property. When somebody asks me about philosophical method I know what it is. But when nobody asks me, and I am philosophizing, I often do not know what it is.

> Oscar Thend, *Philosophical Method*

INTRODUCTION

To have knowledge is to have a network of beliefs that stand in a most important isomorphism with some facts in the world. What is the nature of that isomorphism? This is the question of *basic epistemology*. My general purpose here is twofold: protophilosophically, to heighten our understanding of the methodology of the answer to the question; theoretically, to deepen our insights into the nature of epistemic isomorphisms and the contextual structure of justified belief and of knowledge. This is, therefore, a study in philosophical method and a contribution to basic epistemology and to the philosophy of cognitive language.

This essay is dedicated to:
RODERICK CHISHOLM, *the great philosophical analyst*
        and
EDMUND GETTIER, *the great philosophical iconoclast*

My plan involves several intertwined lines of development. *Two* general lines deserve foremention. *First*, I propose to look at the phenomenon of knowledge directly, not through the existing theories or definitions.[1,2] By exegizing our cognitive experiences we are bound to gather some insights into the role of knowledge in our normal transactions with the world. Such insights will be a positive gain, whatever the destiny of the theories we build upon them. In any case, they crystalize in criteria of adequacy for any future theory of knowledge. *Second*, I propose to subject those criteria to theoretical criticism so as to distill from *them* valuable hints for one or more initially plausible theories, even definitions, of knowledge. The more hints for alternative theories we can place in our theoretical cellar, the better able we are to build a most comprehensive and beautiful theory. Unfortunately, we do not have enough space here to consider all the useful theoretical hints. We must hurry our exegesis and criticism of the phenomenon of knowledge to outline a tentative analysis of knowledge. Although this analysis is only a temporary way station, it may not be amiss to provide at this juncture a general idea of its nature.

The exegesis of data reveals that our concept of knowledge is actually that of a *family* of particular species of knowledge determined by rich contexts of epistemic assessment. These epistemic contexts are normally pinpointed by the context of thought and speech. Hence the epistemic words, particularly 'know' and 'knowledge', turn out to be *indexical*: they either denote a generic state or have a semantic incompleteness of denotation analogous to, although somewhat different from, the denotational incompleteness present in the indicator 'here' and the one present in the color words in their perceptual uses. We shall explain these contrasts.

Our analysandum is the indexed 'X knows$_i$ that p' (for a species of knowledge determined by context i). Our analysans differs from all the standard analysantia because of our exegesis of data, it goes deliberately outside the circle of just truth and belief. It adopts the Plato-Powers principle that knowledge involves essentially the non-doxastic component of a power to answer questions. It also includes the non-doxastic component of some mechanisms of inference. In short, our preliminary analysis includes six conditions about: (i) belief; (ii) relevance; (iii) an appropriate epistemic power; (iv) the normality or limited abnormality of the truth circumstances; (v) some appropriate inferential powers, and (vi) evidence.[3] The standard truth condition is implied by these conditions.

There are other non-standard features of our analysis, e.g., it requires no conception of, let alone belief in, probability; it is non-Fichtean since it does not require self-knowledge, indeed it does not even require the knower to have the concepts of self, belief, or of any other mental state. We shall see all of this and other things in detail and in their proper order later on.

I. SOME UNAVOIDABLE CRITERIA OF ADEQUACY FOR ANY THEORY OF KNOWLEDGE

Gettier has taught us that knowledge is not equivalent to, nor is it identical with, justified true belief. What else is needed? What is justified belief? What is truth?

Here we shall not go into the theory of truth. It will suffice to fix this param-
eter by simply assuming that what one believes is either a truth or a falsehood and
that truths correspond to facts. We may even adopt the *fact/truth equational view*
and take truths to be identical with facts. Naturally, no analysis of knowledge is
fully illuminating until it is placed in the context of a theory of truth and facts.

What other conditions are required for knowledge besides justified true belief?
This most important question demands a careful examination of what is involved in
our attributions of knowledge. We must scrutinize the ways our concept of knowl-
edge functions in experience in order to glean crucial points about the function and
nature of that concept. Then we must posit a concept of knowledge characterized by
a pattern that threads those crucial points. Such points are our criteria of adequacy
for the posited conceptual pattern. The formulation of criteria through the exegsis
of data I call *protophilosophy*, and the positing of theories or patterns *symphilos-
ophy*.[4] Protophilosophy must be pursued systematically.

Undoubtedly, our use of the verb 'know' is vague and imprecise. There is,
therefore, no good reason for arguing whether a certain proposed analysis of knowl-
edge captures exactly our ordinary use or meaning of 'know'. The issue should be,
rather, whether a certain proposal characterizes the pattern of a fruitful concept that
coincides with a large segment of our ordinary use of 'know' and its inflections, thus
illuminating our actual cognitive dealings with others and with the world. The cri-
teria of adequacy fix the crucial points of the illumination and hence both demarcate
the field of theorization and provide tests for all proposed analysis or theories of
knowledge.

Let us proceed to our protophilosophical disquisition.

1. The Criterion of Knower Reliability

The most obvious and general feature of our concept of knowledge is that to attri-
bute to a person, say, Sharon, knowledge that p is to attribute to Sharon a maximal
degree of reliability concerning p. This reliability has several dimensions, and it
grounds other criteria. If Sharon knows that p, then Sharon not only has a true be-
lief but has beliefs that both guarantee that p and make her an unimpeachable source
of information about that p.[5] Let us develop these points in some detail.

2. The Gestalt Character of Knowledge

Both the points made above about the knower's maximal reliability involve a non-
atomist structure of knowledge. Each belief, like each piece of knowledge, is indeed
associated with a proposition, or a state of affairs (or a fact, if you wish). Each true
belief that p and each piece of knowledge that p is, on the equational view of truth
mentioned above, a fact. But note our language! Here we have a clue to the non-
atomist character of knowledge. We speak naturally of each belief, but we need a
fragmentation operation in the case of knowledge, and then we speak of each *piece*
of knowledge. This linguistic datum suggests that knowledge, unlike belief, is a non-
individuated massive whole. Indeed, going beyond the linguistic datum, it seems that
at a certain surface level of analysis believing is atomistic in that a person could con-

ceivably believe each proposition he believes in isolation of his other beliefs, without taking into account its implications relationships. At that surface level of analysis, on the equational view of truth and fact, the isomorphism between true beliefs and facts is an identity isomorphism of a set of beliefs.

Knowledge is, on the contrary, non-atomistic. For a person to know some truth, she must have evidence for the truth in question, or she must be at least justified in believing that truth because it agrees or coheres with other beliefs she has. A person who knows a truth must believe a whole battery of truths, and even know other truths. To have knowledge is to have beliefs about propositions and about their implication relationships. To know is to have a network of beliefs with a certain *Gestalt*. Thus the belief-fact isomorphism constitutive of knowledge connects networks of beliefs with networks of facts, as wholes, not piecemeal through the association of the individual beliefs with their corresponding facts. On the fact/truth equational view, the knowledge isomorphism may still be an identity isomorphism, but it must be the identity isomorphism of a set of sets of beliefs (or truths).[6]

3. The Knower's Beliefs Must Guarantee Truths

Let us return to the knower's maximal reliability. Patently, the knower cannot be maximally reliable if he is not reliable without the faintest possible correction to his reliability. Since we are dealing with thinking sources of information, the only relevant reliability in basic epistemology is the reliability of what our thinkers believe. In basic epistemology we can assume that if a person attempts to communicate, he or she will, if successful, communicate candidly. We can, that is, assume that the person's utterances will convey the truths, or falsehoods, he or she believes—at least to the extent that language is propositionally transparent.[7] Furthermore, in basic epistemology we do not concern ourselves with the information provided by the circumstances of the speaker, by the speaker's body, movements, or vestment, or by the circumstances of the speaker's speech acts. All of these and other sources of information are valuable and must be studied by later, more sophisticated branches of epistemology. In basic epistemology we deal with fundamental structures to be assumed by those later branches.

In short, the maximal reliability of the knower requires that the knower's beliefs (whether we call them evidence or not) guarantee in a very strong way the truth of what he believes and knows.

4. The Cognitive Irrelevance of Probability

Patently, the criterion that the knower's beliefs guarantee the truth of what he believes leads immediately to the irrelevance of probability in the central structure of knowledge. It is a triviality that to believe that it is very highly probable that p is *not* to know that p, regardless of how high the probability one believes to obtain may turn out to be, regardless of how correct this belief may be, and regardless of how deeply and thoroughly justified one may be in believing that it is very probable that p. To illustrate, consider

The Supercautious Gambler. There are 1000 tickets in a lottery. Ann, our cautious gambler, has bought 999 tickets. The prize will be won by the ticket with a number, between 1 and 1000, which also appears on a wooden disk to be picked out of an urn, completely at random.

Obviously, given the knower's reliability criteria, our supercautious gambler does *not* know that she will win the prize. To be sure, she has a right to believe with confidence that she will win; but neither she nor we know that this will happen, even if in fact she will win.

The Supercautious Gambler is the gambler who does not know that she will win when she will in fact win. This is an unavoidable datum for any theory or analysis of knowledge. Interestingly enough, it rebuts the very sophisticated analysis of knowledge proposed by Keith Lehrer in [42]. Indeed, Lehrer's analysis is rebutted by the weaker datum contained in:

The Minimally Cautious Gambler. Like the Supercautious Gambler, except that this gambler buys only 501 lottery tickets.[8]

The moral of these examples, when they are properly exegized, is that the beliefs of the knower must guarantee the truth of what he knows.[9]

5. The Paradox of Empirical Knowledge

Clearly, the reason the supercautious gambler does not know that she will win is because she is not as supercautious as she may be: she could have been maximally cautious if she had bought all 1000 tickets. Then she would have had, assuming the circumstances to be normal, a full guarantee that one of her tickets would be the winner. Requiring probability 1 for knowledge would guarantee truth. But probability 1, as *The Maximally Cautious Gambler* illustrates, requires logical implication. Lehrer knows this very well, and he mentions [3], [39], and [59] as having shown that a hypothesis *h* has probability 1 on evidence E only if E logically implies *h* ([42], p. 145).

This seems to raise a paradox. Our contingent claims about the world, including claims about physical objects in our neighborhood, are not logically implied by our evidence for them. For instance, our sensory experiences do not imply our perceptual claims; our perceptual claims do not logically imply our claims about other persons' mental states; our perceptual claims, including the testimony of other persons, do not logically imply our claims about theoretical entities, and so on. Thus, if we require knowledge to include a guarantee of truth, we seem to be ruling out empirical knowledge altogether, except perhaps for the solipsistic knowledge of our present contents of consciousness. But perhaps not even this. Thus Lehrer claims:

Nevertheless, for any strictly coherent probability function, restricting justification to a probability of one will lead us deeply into the den of skepticism. All of the contingent and non-general statements we naturally assume we know would turn out to be statements we are not completely justified in believing and could not possibly know ([42], p. 151).

Apparently we have a serious dilemma: either we insist on knowledge beliefs guaranteeing truth, or we do not. If we choose the first horn of the dilemma, we are apparently lost irretrievably in some deep skepticism. If we opt for the second horn of the dilemma, we know even, as in *The Minimally Cautious Gambler*, when we believe correctly that the chances of truth are barely more than 50 percent, just in case what we believe happens to be true. On the first horn we know too little; on the second horn we know too much. Lehrer has chosen the second horn.

There should be a way out of the dilemma without having to endorse either horn. Clearly, our efficient and useful concept of knowledge functions in our experience in a way such that neither horn is a viable alternative. Understanding the nature of our experience and the structure of the world we find ourselves in requires that we dissolve the dilemma, rather than choose one of its horns. Since the dilemma rests on the use of probability, perhaps we can do without probability—or perhaps we can require probability 1 for knowledge and yet salvage empirical knowledge.[10] After all, empirical truths do imply other empirical truths.

The presence of the paradox and the need of its dissolution are themselves a most important datum. A satisfactory theory of knowledge, and even perhaps a satisfactory analysis of a fruitful and illuminating concept of knowledge, should shed light on the paradox, its source, and its dissolution.

6. Belief in Probability Not Required for Knowledge

In our world we often encounter beings with the power to think, have beliefs, and even know (in our ordinary use of this word), even though they lack the concept of probability. Small children, not to mention animals or robots, do seem to have knowledge, perceptual knowledge, if nothing else, even though they lack the power to think of some states of affairs, or propositions, as probable. A proposition of the form "Probably p" or "It is (highly) probably that p" is a modal complex proposition. To think a modalized complex proposition Mp, a thinker must have the ability to think p, but the converse is not generally necessary. The more complex the modality M, the larger the gap between the ability to think the arguments p of the modality and the power to think of the modalized proposition Mp. Naturally, given the syncategorematic character of modalities they cannot be thought by themselves —except as abstractions from modalized propositions.

Clearly, a probabilitized concept of knowledge, like the one Lehrer has constructed, lacks the fruitfulness of the ordinary concept of knowledge. First, the fruitful ordinary concept of knowledge is the concept of a maximally reliable source of information, where, as noted, the constructed probabilitized concept of knowledge can be very unreliable. Second, the ordinary concept of knowledge dispenses altogether with probability as the content of some networks of beliefs that constitute knowledge.

7. Knowledge Does Not Require Knowledge, or Belief, that One Knows

Similar to the previous criterion about the concept of probability is another important criterion for elucidating our ordinary concept of knowledge. A person can know

without being able to *think* that he (himself) or she (herself) knows: the person may lack the concept of the first-person. This is a crucial point I raised in detail in [7] and assumed as far back as 1960, when I began writing [4].[11]

In [7] I contrasted the consciousness and mind of a thinking being lacking the power of self-reference, being which I called *Externus*, with the consciousness and mind of a being endowed with the power of self-reference who is such that: for him consciousness is self-consciousness, and whenever he believes, knows, imagines, etc., he knows (believes) that he knows, believes, imagines, and so on. This reflexive property I called Fichtean, and I spoke of *Fichtean knowledge*, *Fichtean belief*, and *Fichtean consciousness*—in honor of Fichte who forcefully claimed that "without self-consciousness there is no consciousness"([26], p. 41).

It is an empirical matter whether there exist, or not, examples of Externus-type of mind or of consciousness. David Schwayder has assured me, in correspondence commenting on [7], that, of course, we are surrounded by Externi, namely: animals, especially pets, and very young babies. He may be right about this.

It is also an empirical matter whether there are, or there are no, Fichtean minds in our midst. Hintikka, as is well known, has produced in [36] formal theories for species of Fichtean knowledge and Fichtean belief. It is not clear, however, that he is committed to the view that ordinary human beings have Fichtean knowledge and Fichtean belief. Lehrer has also proposed in [42] an account of Fichtean knowledge.[12] Perhaps he also holds that ordinary human beings are Fichtean in this sense. Chisholm in [19] and [20] does seem to hold the strong Fichtean view that consciousness is self-consciousness.

It seems to me that we, ordinary human beings, do not possess Fichtean minds. When I read in Sartre's [54] about episodes of unreflective consciousness, I find myself experiencing similar things. Many a time I am aware of something, I think of something, without being aware or thinking that I am aware or thinking of it. I have often been surprised to learn that I have known some important truths all along.

We now have the concepts of Fichtean knowledge and Fichtean belief. They are important concepts. They can be of great value in a general theory of knowledge in which they are used to characterize an epistemic or doxastic ideal limiting case. I submit, however, that our ordinary concepts of knowledge and belief are more general, and have the Fichtean cases as possible limit instances. The generality of our ordinary concepts of knowledge and belief infuses them with great flexibility and usefulness. We must, in order to understand *our world*, produce accounts of general knowledge and belief, which need not be Fichtean.

8. The Concept of Belief Is Not Required to Have Knowledge

Just as an ordinary knower need have neither the concept of probability nor the concept of self nor the concept of the first or second person, so a knower need not have any psychological concept at all. An Externus knower may simply know the physical world, having no idea of other minds or of the contrast between object and subject. In [4] I describe briefly the structure of the contents of consciousness and of belief of the Externus stage of a being called Privatus. Privatus thinks out loud, so

that the attributions of knowledge to him can be made on the standard basis both of perceptions of his body and circumstances, and of interpretations of his utterances. As in our normal attributions of knowledge, Privatus may be correctly attributed the maximal reliability characteristic of knowledge.

Briefly, an Externus type of consciousness and mind will think and believe, not *cogito* propositions of the form "I see (hear, feel) that such and such," but perceptual *cogitatum* propositions of the form "There is such and such here (there)," which are states of affairs in perceptual fields. Obviously, an Externus knowledge of physical objects cannot be built upon *cogito* propositions. It has to be built, if 'built' is the appropriate word, from *cogitatum* propositions—if these are not already propositons about physical objects. This is a crucial non-Cartesian aspect of knowledge set aside by philosophers who hold that knowledge has a foundation on experience and who then go on to interpret experience as what *cogito* propositions are about.[13]

9. The Basic Epistemic Paradox of Other Persons' Information

We need other persons to furnish us with information. This raises a problem. I do *not* mean the traditional problem of other minds. This problem, about how we know that there are other beings with mental states, certainly has a place in the theory of knowledge. See, e.g., [4]. But here, in basic epistemology, we can set that problem aside. The problem of other persons that belongs here is different. We assume here that we are in fact surrounded by other persons, that we deal with them, and that they have mental states to which we have access somehow. Yet within these assumptions we have a problem. We consider X, say, as a person. Thus X is endowed with both the power to know and the power to act. Because of X's knowledge, X is a source of information and we can rely on him to the extent that we claim that he knows. But X is also an agent, and, as such, we must think of X as capable of choosing freely at least some of his courses of action. To the extent that he is a free agent, X is naturally unreliable. Thus, whenever we find a free agent in a chain of evidence, we have there a point of unreliability, a point of non-knowledge. This is illustrated by:

> *The First Cheater.* Lottery L is composed of 1000 tickets. The Official Drawer OD will pick out a numbered chip from a revolving urn. Our Cheater has bought ticket no. 5. He has paid OD an appropriate fee, in return for OD's promise that he will pick out chip no. 5.

Our cheater clearly does not know that he will win. Yet he has the relevant information both about the Official Drawer's knowledge of the lottery and about OD's character. OD is known to keep his promises. But the more First Cheater thinks of OD as a free agent, who will be free to choose or not to choose the winning chip at the moment of drawing, the less First Cheater is in a position to know that he will win. Regardless of how many promises he has kept in the past, free agent OD will have to act in the light of all the reasons he considers *at the time* he decides what to do then.

By contrast, the more Official Drawer's action is determined and the more it is known how this action is determined, i.e., the more OD is like a *mechanism* rather than a *genuine agent*, the more one can claim to know that he will do action A.

An expert agent who is both knowledgeable and free to inform, or not to inform, is a perplexing beast. He is reliable to the extent that he knows, yet he is unreliable to the extent that he is free to inform or not to inform.

This internal conflict characteristic of each free informant presents problems that we cannot even start to consider here. All we can do is note that in practical life we solve the problem by adopting an *attitude of trust*. This attitude grounds beliefs, is itself grounded on some beliefs, but it cannot be the logical consequence of the beliefs grounding it. Having knowledge entails more than having beliefs: it also entails having *propensities* to have certain beliefs and make certain inferences.

10. Propositions Are Evidentially Organic Wholes

We need the testimony of others to secure the certainty that comes with the essential intersubjectivity of physical reality. This need of others' testimony includes a paradox of its own. We need others with their true beliefs, but we must have them with their freedom to deceive, mislead, and lie. Thus the potential for false evidence and true counterevidence are essential components of our dependence on others.

Yet our proneness to error lies deeper than our social dimension. Our finitude creates an enormous potential for error. Our need to believe and to act on our beliefs leads us to plunge into errors and into partial views of things. We must build our knowledge on beliefs that may be falsehoods or misleading truths, simply because we do not know enough about their *connections* with other truths. Falsehoods may sometimes be the best vehicles for reaching truths and even knowledge. Counterevidence at a certain level may turn out to be part and parcel of valuable evidence.

Propositions, or states of affairs, are, with respect to evidentiality, *organic wholes*. They form compound propositions that have evidential values which are not merely the sum of the evidential values of the components. Thus some set E of propositions may provide very strong evidence for a proposition p, whereas a superset E' of E may provide counterevidence for p, and a superset E'' of both E and E' may provide again strong evidence for p. Therefore, it may happen that a superset E'' may justify belief in p much better than E, thanks to the fact that the set E'-E provides, by its subsumption into E', negative evidence.

11. False Evidence and Counterevidence May Be Needed for Knowledge

The preceding feature is nearly a commonplace when it is discussed in its naked generality. However, it has not enjoyed the full appreciation it deserves in the mainstream dialectic about the analysis of knowledge. The dominant tendency among the most distinguished practitioners is a form of what has been called the *indefeasibility approach*.[14] This approach is *puritanical*: it requires the evidence or justification of belief that constitutes knowledge to be pure true belief, or belief purified of falsehoods.

The easiest indefeasibility formula is simply to require both (*a*) that there is no counterevidence anywhere, even outside the knower's realm of beliefs, and (*b*) that the knower does not have counterevidence for what he knows.[15] Each of these two

conditions is patently too strong. Each runs against the grain of the principle of organic unities for evidence and justification of belief. Thus the most energetic epistemologists have been engaged in a search for methodologies that modify (a) and (b) in the appropriate way—recognizing the fact that there are no ways of suppressing counterevidence from the world, or even from the knower's mind.

The best work on defeasibility does not attempt pure justification as in (a)-(b). The best proposals aim, rather, at the purification of evidence and justification in order to secure knowledge. The idea is to purge the knower's evidence of the bad features and insist that knowledge obtains only if the purged evidence retains its justificatory power. This core idea receives different concrete specifications. The most prevailing view is, roughly, that the knower has a *wholly* veridical evidential path (pyramid or tree, says Sosa) that links a sufficient set of pieces of evidence with what he knows. This is, however, too puritanical.

The purification techniques that have been developed, with ever increasing insight into the nature of justification, are of great value. I am, however, not entirely happy with two implicit *assumptions* that typically underlie the development of those techniques:

(A) counterevidence always has to do with falsehood;

(B) falsehoods must be purged: a set E of propositions that justifies belief in p does not yield knowledge if the believer does not have evidence that would justify his believing p if the falsehoods in E are purged somehow.[16]

Of these, (B) is the more widespread. The mainstream discussions seem to agree that at least (B) is necessary for knowledge. The issue has been whether it is also sufficient.

Undoubtedly, counterevidence sometimes involves falsehoods—as with false testimony; and evidence containing falsehoods is usually incapable of yielding knowledge. Gettier's classic counterexamples illustrate this very well, and so do many others. But not always. *Sometimes the one who fails to know need not have any false belief.*[17]

I want to urge that (B) is *not necessary* for knowledge. It may occasionally happen that the one who knows needs some false belief in order to come to know. Consider, for example,

The Cross-Wired Rememberer. Crispin has an interesting memory mechanism. Whenever he perceives a date d written on a book, if he perceives it very clearly, he will remember it as d + 10. (This causes him some problems concerning exchanges of information about times, of course.) He has read only one history book about Columbus's preparations for his first voyage to America. In that book there was a misprint: it gave Columbus's date of departure from Palos as August 13, 1482. He repeated the date in order to memorize it. Naturally for him, his memory mechanism corrected his correct perception. Crispin knows, just like anyone of us, that Columbus left Palos on August 3, 1492. Just ask him!

In sections 14–17 below we shall examine the varying contextual assumptions of different species of knowledge. Then we shall better see in what sense both Crispin

and we know that Columbus left Palos on August 3, 1492. For the moment let us simply say that Crispin passes the regular tests we pass: mentioning books where we can find the right information, answering the question "When did Columbus depart from Palos (Spain)?" as we do. We may even suppose that Crispin *knows* of his cross-wired memory for dates, so that he knows that he needs false perceptions, or false initial beliefs, of the appropriate kind to know the dates of events and to lead others to know what he knows.

The general point is this: *within a standard correlation between perception, memory, and belief, there can be systematic exceptions that correlate false perceptions with true, fully justified beliefs*, which are fully justified *because they are grounded on those false perceptions*.[18] The connections among the different operations of the mind are capable of combining in more ways than we are used to and still yield knowledge. This matter deserves thorough study, without which a full theory of knowledge will not be attained. For basic epistemology we need note only that we cannot assume an easy correlation between perception, memory, and belief. In fact, even within our ordinary scheme we can have:

> *The Twitching Color Blind.* Norman is color blind, not capable of visually discriminating between green and blue. He has found, fortunately, that often, when he sees pairs of blue-green objects, one member of the pair causes his eyelids to twitch. To make the story short, it happens that Norman learns to discriminate some blue-green objects from other blue-green objects by his twitching, namely, the ones others see as blue when he sees them paired with green objects.

Given that Norman must see a contrasting pair to have his twitching indicative of a blue object, he must, then, have the initial false perceptual judgment that the two objects he sees are of the same color. This judgment grounds his further judgment about the difference in color. The requirement that eliminates false judgment from knowledge-making evidence would deprive Norman of the proper perceptual basis for coming to know that the object he has seen, or is seeing, is blue.

It is easy to conceive of correcting methods for acquiring knowledge that depend crucially on there being some false steps to be corrected. Indeed, I wonder whether this is not more common than purificational epistemologists have thought. In general we can imagine:

> *The False-Believing Knower.* In certain types of situations Louise reacts by acquiring a belief, which turns out to be false. Louise learns the systematic connection between her circumstances and her acquiring a false belief of the requisite type, and then learns to use her own false beliefs as an inductive basis for positing the corresponding truth. She reasons as follows: "In situations of type s I always acquire a false belief of type t; I am now in situations of type s and I have acquired the belief that p, which belief is of type t; hence, it is not the case that p."

Obviously, the knower can more easily, and naturally, use essentially the false beliefs of others to acquire his knowledge.

Nothing in the preceding discussion contravenes the chief intuition of the analytic epistemologists, namely, that the belief constitutive of knowledge be undefeated and that the evidence for it be undefeated. These intuitions are palpably sound. The discussion shows that the puritanical view, which requires for the undefeasibility of the evidence a wholly veridical path (or pyramid, as Sosa says), is *too restrictive*. The evidence must be undefeated in the sense of beating up the opposition, but there may very well be strong opposition that, so to speak, leaves indelible scars on the victor.

12. The Objective-Subjective Paradox of Epistemic Justification

The requirement of a wholly veridical evidential path seems like the most natural one to complete the classical list of conditions for knowledge. Yet even full evidence without a taint of falsehood may *fail to be sufficient* for knowledge. We shall show this presently through some easy useful examples. But first, let us reflect about it to find our bearings. The situation is perplexing. Nothing seems to be missing from the list: truth; belief; complete justification; wholly veridical justification. What else can be added?

A look at the standard characterizations of complete justification reveals that the completeness of justification in question is a matter of the knower's beliefs overriding all the counterevidence available to him. What we feel is the need to ask for an objective completeness. This is precisely the idea behind the naive condition (*a*) discussed in section 11 above. There is, however, the crucial fact that a merely objective condition that we cannot take into account is of little value—except as an ideal. (See Tienson [71].) Hence we must require a subjective completeness, which is precisely what Sosa, Lehrer, Harman, Dretske, and the others do. But then we are back to the position where wholly veridical (subjectively) complete justification may fail to yield knowledge. We face here one of the most fundamental problems: *the subjective-objective paradox of epistemic justification*. It is a most pressing one. It is a special case of our egocentric predicament: We encounter it just because we cannot get outside our circle of beliefs.

Perhaps it may not be amiss to see the paradox illustrated in Harman's very fine perception of the problem in [31]. He devises some examples to show that completely justified true belief may fail to constitute knowledge: some missing evidence keeps knowledge away. I am not entirely sure that his examples are all effective. But the point is clear even on general grounds. In any case, there are the examples in *The Cheater Series* in section 13 below. Harman proposes a solution in terms of a subjective requirement for justified inference:

Q. One may infer a conclusion *only if* one also infers that there is no undermining evidence one does not possess ([31], p. 151; my italics).

This principle *cannot* be part of an account of *our* concept of knowledge. According to Harman's Q, the conclusion the prospective knower would infer is of the form

(1) There is no evidence undermining [. . .] that I do not possess.

The indefinite personal pronoun 'one' in the subordinate clause 'there is no under-mining evidence ONE does not possess' in Q is what I call a *quasi-indicator*. (See [6] and [10].) Propositions of this form can be thought only by entities that have: (i) the concept of, or the mechanism of reference in, the first person; (ii) the concept of evidence; (iii) the concept of belief, which is what one possesses as this word is used in (1) and Q. Thus, as explained in section 7, were Harman's Q a universal condition of knowledge, it would at best govern *Fichetean knowledge*. For non-Fichtean, Sartrean creatures like ourselves, Q would impose an unbearable amount of self-consciousness.

Can Harman's Q be a principal characteristic of the knowledge of those thinkers capable of self-reference (not merely capable of referring to themselves but capable of referring to themselves in the first person)? To decide this, we need to know what 'undermining' means. Harman tells us that "the label 'undermining evidence one does not possess' has been explained in terms of knowledge" ([31], p. 151). A perusal of the text preceding this quotation shows the phrase 'undermining knowledge'. Hence the full version of form (1) is:

> (2) There is no evidence undermining my knowledge that *p* that I do not pos-
> sess.

Clearly, then, adding to our list of conditions for knowledge (truth of that *p*, belief that *p*, complete justification, wholly veridical justification of one's belief that *p*) the fifth condition that the knower infers in accordance with Harman's Q does not accomplish much. Any believer can so infer. Should we require, instead, as the *fifth condition* guaranteeing sufficiency (not necessity, remember) that the knower infers *justifiedly and correctly* in accordance with Q? This new condition would give us a desirable combination of subjectivity (justified inference) and objectivity (correct inference). Is this the end of the post-Gettier search?

The above fifth condition seems to be successful if 'undermining' means falsifying. Clearly, if there is no evidence, whether I possess it or not, that falsifies that I know that *p*, then on the assumption of bivalence, I know that *p*. The trouble with this fifth condition is its unhelpfulness. The five conditions cannot provide an analysis of knowledge because of the circularity of the fifth condition. The conjunction of the five conditions is not equivalent to "X knows that *p*"; it implies, but it is not implied by, "X knows that *p*."

To sum up, Harman's Q cannot help us understand the structure of knowledge. He suspects that this is so, for he merely claims that Q "is a principle concerning inference." I am sure that some scientific inferences conform to Q; but Q certainly cannot account for most of the valid and justified inferences we make in daily life—not to mention the inferences made by children who have not yet acquired the concept of evidence, or of undermining knowledge, let alone the concept of self.

The objective-subjective paradox of epistemic justification is a most serious matter. It has to be solved, we can be sure, only by a requirement that combines the subjective element of some belief with the objective element of the truth of that belief.

13. Wholly Veridical, Completely Justified True Belief Need Not Be Knowledge: The Cheater Series

Let us know discuss some examples that establish this criterion. Exegizing them should help us find valuable clues for solving the subjective-objective paradox of epistemic justification. Consider:

> *The Second Cheater.* Lottery L is composed of 1000 tickets. The winning ticket will be the one with a number matching the number on a disk in a revolving urn picked out by the Official Drawer (OD). Ann bought ticket no. 5. She has ensured that OD picks out disk no. 5 in some way, the mechanics of which do not matter. (On one version Ann arranges to have disk no. 5 magnetized and a fine iron filing inserted in OD's right hand, his lottery hand.)

In an example like this, with more obvious details added, Ann knows that she will win the lottery prize. But that is so only because of a major detail that is seldom discussed, yet it is of the greatest importance in Ann's knowledge that she will win the prize. The general silence on this detail, it seems to me, accounts for a good deal of the difficulties normally encountered in discussions on the analysis of knowledge. That detail is this: *the total circumstances surrounding and linking Ann and the lottery are assumed to be normal.* The normality of the circumstances has been the crucial neglected factor. Such normality is implicitly assumed in the many ingenious counterexamples which show the inadequacy of proposed analyses of knowledge by destroying the normality of the circumstances.[19]

How must we deal with normality? This is the nuclear question. To find a guide to the answer we must exegize the epistemic roles of alterations of normality. Consider:

> *The Third Cheater.* The situation is the same as in *The Second Cheater*, except that without anybody having any idea (knowledge, if you wish) of it, there is a pair of mechanisms M_1 and N_1, say, in the Lottery Hall. Mechanism M_1 replaces the number n on a disk picked up by OD with the number $n + h$. Mechanism N_1 maps a number n on a disk (either in the hand of OD or right upon the disk touching the official lottery tray) into the number $n - h$. Thus, when OD picks up disk no. 5, mechanism M_1 causes the number $5 + h$ to be on the disk. Then mechanism N_1 replaces $5 + h$ with 5 again. M_1 and N_1 are causally independent of each other. Thus, when OD has placed the winning disk on the official tray the disk shows ticket no. 5 to be the winner. Ann collects her prize.

Perhaps it is useful to add a word about the mechanisms. Their physical embodiment is immaterial. What is important is that they operate as noted and that N_1 acts after M_1 during the appropriate time and location. M_1 could be embodied in a chemical process acting on a disk in OD's hand because of chemicals he used to cleanse his hands. N_1 could be embodied in an optical process acting through mirrors reflecting light on the official tray. But this is a problem of engineering, not of philosophy.

Does Ann know before OD picks up disk no. 5, or before she is declared the winner, that she will be the winner? This question cannot be answered immediately,

one way or the other, just on the description of the situation contained in *The Third Cheater*. But it is of the utmost importance to emphasize one crucial point: *We are not here concerned with falsehoods*, whether in Ann's beliefs or in the relevant reports made by others. We are deliberately limiting ourselves to wholly veridical evidence.

Does Ann, in *The Third Cheater*, know that she will win, if she will in fact win, the lottery prize? *It depends on the context*. By hypothesis Ann has no idea of the balancing mechanisms M_1 and N_1. If these mechanisms are *normal fixtures of the lottery situation*, Ann knows that she will win the prize—in spite of her ignorance about them! Let the mechanisms be standard elements of the normal circumstances. Then they function normally in the total process that both leads causally to the truth of the proposition "Ann [I] will win the prize" and furnishes Ann's evidence for this proposition. That is, those mechanisms function in exactly the same way as other normal causal segments of the standard situation about which nobody has any idea, e.g., the processes inside OD's lottery hand, or inside the disks in the lottery urn, whose chemical constitution has not been studied up to the time of drawing. Yet such processes and those involved in other objects in the situation have an epistemically relevant place in the proceedings leading up to the declaration of the winning ticket.

The fundamental idea is, then, that the *normality* of the circumstances pertaining to the truth of the known proposition plays a crucial role in the knower's knowing that proposition: that normality is part of the *objective requirement* for knowledge to come about. The *subjective requirement* is this: the knower's assuming, taking it for granted, rather than believing it in a more substantial sense, that the circumstances are normal. Here is, I submit, the first stage in the solution of the subjective-objective paradox of epistemic justification.

Suppose now that, alternatively, neither mechanism M_1 nor mechanism N_1 is a normal part of the standard circumstances; they are instantiated for the first time in the history of Lottery L. Then Ann does *not* know that she will win the prize. To be sure, Ann, as described in *The Third Cheater*, still takes it for granted that the circumstances are normal. This is precisely what creates the discrepancy between her beliefs (and evidence) and the external reality. In short:

Crucial Datum I. In *The Normal Third Cheater*, Ann knows, whereas in *The Non-Normal Third Cheater* Ann does not know, that she will win the prize.

The Non-Normal Third Cheater establishes that the typical undefeatedness analyses of knowledge do not provide universally valid sufficient conditions for knowledge. Ann does not know that she will win the prize. Yet she has wholly veridical justification in her belief that she will win the prize, lacks the slightest whit of counterevidence, and in fact she will win the ticket. Furthermore, we may suppose that she has drawn, on her own initiative, the conclusion that no evidence that undermines her knowing that she will win is not known to her.[20] She has inferred that the best hypothesis that explains the facts as she believes (knows, if you wish) them is that she will win. In *The Normal Cheater* she knows in spite of defeasibility.[21]

The Non-Normal Third Cheater brings out a possible ambiguity in Harman's principle Q. Does mechanism M_1 undermine Ann's, or anybody else's, belief that Ann will win the prize? In the context in question *as a whole*, which contains also the reversing mechanism N_1, the answer seems to be an emphatic "NO!" On the other hand, we may consider, not the whole context, but only the context known to Ann, her evidence, and then the extended evidential context that includes M_1 but not N_1. In this *piecemeal* view it is not out of order to hold that there is in reality a piece of evidence that when added to Ann's evidence undermines her knowledge that she will win. I cannot find in Harman's book a text that decides this issue, at least in a way clear to me. But he does include an insightful discussion of how one's evidence may change from time to time. This suggests that *perhaps* Harman at least at some moments adopts the piecemeal view.

If 'undermines one's knowledge that p' is taken in the holistic way, then Ann is correct in inferring that there is no evidence that undermines her knowledge. If the expression is taken in the piecemeal sense, then Ann is mistaken in that inference. But then her being able to make the inference correctly is irrelevant to her knowing that she will win the prize. We can, if we wish, have Ann not engage in irrelevant inferences.

Let us continue our exegesis of the data. Consider *The Modified Non-Normal Third Cheater*. This is like *The Non-Normal Third Cheater*, except that Ann knows somehow, or, more weakly, has strong evidence, about the generally unthought of mechanisms M_1 and N_1; everything else remains the same. Clearly:

Crucial Datum II. In *The Modified Non-Normal Third Cheater* Ann does know that she will win the lottery prize.

Palpably, there are in principle infinitely many pairs of mechanisms, like M_1 and N_1, that can intervene anywhere in the standard causal circuit leading from the setting up of the lottery to the delivery of the prize to the owner of the winning ticket. Furthermore, pairs of mechanisms like M_1 and N_1 can zero in on new members of the expanded causal circuit. For instance, a mechanism $M_{1.1}$ may stop the action of mechanism M_1, or it may operate after M_1, say, mapping $n + h$ onto $n + h/k$, but then another mechanism $N_{1.1}$ reverses the action of $M_{1.1}$.

The requirements are that these mechanisms work in pairs in the appropriate time and order. Naturally, there are also pairs of mechanisms that epistemically cancel each other out, so to speak, even though they do not operate contiguously, converging on one link in the causal chain. There may be all sorts of compensatory pairs of changes.

We shall speak of *The Kth Cheater* to refer to a *Cheater* situuation built on *The Second Cheater* by $K - 2$ pairs of compensatory mechanisms. Once again, we distinguish between *Normal* and *Non-Normal Cheaters*. We also distinguish, as before, between *The Non-Normal Kth Cheater* in general and *The Modified Non-Normal Kth Cheater*. *The Fully Modified Non-Normal Kth Cheater* is the variant of *The Kth Cheater* in which there are $K - 2$ pairs of compensatory mechanisms such that: (i) they are exceptions to the normality of the circumstances of *The Second Cheater*;

erties we consider in daily life, however, do not seem to be truly analyzable in a finite number of steps into ultimately primitive ones. Thus there seems to be no ultimate analysis of many propositions we know to be true. Hence the hierarchy of logical forms (and guises) of most propositions is not composed of (definitely) a finite number of rungs. Thus the hierarchy of questions determined by most propositions is not a finite system. Hence there is in principle an indefinite, and even infinitely large, number of species of knowledge.

The preceding result seems to explain why we lack in ordinary language a systematic mechanism for denoting, or signaling, the species of knowledge with which we are concerned. Obviously, we are not concerned in practical life with the whole hierarchy of logical forms (and guises) of a proposition, at least not as such. The interrogative locutions we have mark points of interrogation. By placing them in an interrogative sentence, we reveal both the level of logical form (and guise) and within this form the interrogation point in which we are interested. Compare, e.g., in the case above, the questions: "*Is it* the case that (2)?," "If all flights were on time today, *where* is John?," and "If all the flights were on time today, then, since he is neither in the men's room nor in a telephone booth, *where* is he?"

Now the question appears: If in ordinary language we lack a systematic mechanism for specifying which questions and, *a fortiori*, which species of knowledge we are interested in, how do we manage? For we do seem to manage to communicate about what certain persons know or, for that matter, fail to know. The answer seems to be this: *We manage to communicate about species of knowledge in exactly the same way we manage to communicate about the species of other properties, when we need a specification from a genus with indefinitely and perhaps infinitely many species*, namely: the context of communication pinpoints the relevant species.

What I am claiming for 'know' and 'knowledge' is in detail somewhat different from what happens to demonstrative words like 'here' and 'there' and from what happens to color words. Of course, there are fundamental reasons for the differences. The common claim is that in their primary uses all these types of expressions are at bottom indexical. (See [6].)

With 'here' and 'there' the expressions themselves, and the sentences in which they occur, need to be complemented by an association with a place of *possible* (*not* actual) utterance. In perceptul uses color words and the sentences containing them need an association with actual shades of color; but they can be used in a generic way. One reason, I suppose, is that the actual determinate color content of experience is normally not important in communication. Thus, in some uses, especially non-perceptual uses, color words do not denote a determinate shade determined by context, including pointing. They must, then, be understood as either denoting a generic, a determinable quality, or as being implicitly quantified. For instance, 'The book on the table in the chairman's office is red' may be taken as of the form 'There is a shade ϕness of the red family such that the book on the table in the chairman's office is ϕ'.

The epistemic words 'know' and 'knowledge' may perhaps be used in a generic sense, as the color words are. As in the case of these words, the generic use rests on

simpliciter but must distinguish *species of knowledge*: at least one for each context of belief justification. Nothing at this stage of exegesis requires that within each context there be no other parameters to which knowledge is relativized.

One consequence of the diversity of normal contexts impinging upon one and the same person at one and the same place and time is this: for some proposition that *p* the person so impinged may know that *p* in one respect or species, without knowing that *p*, in another respect or species. This is an important result. *First*, it is testable in experience. *Second*, it will furnish some aid in understanding the historical puzzle, merely accentuated in the last decades, of the stubbornness of the disputes about the analysis of knowledge: they continue unabated in spite of the tremendous cleverness of the arguments and the ingenuity of the examples taken as data; the debate has provided penetrating and satisfying insights, yet mastering those disputes produces a deep sense of unfulfillment. The stubbornness of such ingenious disputes suggests that perhaps underlying them all there is often no true joining of issues: that sometimes one species of knowledge may be under consideration and sometimes another species may be examined, even perhaps by the same epistemologist—under the implicit, or explicit, assumption that there is just one species of knowledge.

Let us return to our daily cognitive experiences. Let us investigate whether there are species of knowledge. For this we need cases in which for some person, e.g. Andrea, and some proposition that *p*:

(1) At time *t* Andrea knows that P;

(2) At time *t* Andrea does not know that P;

(3) There is no contradiction in (1) and (2);

(4) The common words in sentences (1) and (2) seem to have, on all fours, the same sense or meaning.

There is unavoidable vagueness in condition (4). The reason is that (1)-(4) are descriptions of data, and the data have certain vaguenesses that represent the junctures at which the puzzles they represent are to be solved by the theories to be developed. Naturally, we may say that

(5) The sentence 'Andrea knows that *p*' does not express in isolation, in (1), the same proposition, or truth, that it expresses in (2).

Hence we may say that the sentence 'Andrea knows that *p*' occurs ambiguously in the pair (1)-(2). But we must *not* conclude from this that the word 'know' is ambiguous. In general, the (propositional) ambiguity of a sentence does not imply that any of its words or constituent expressions is ambiguous. Let us consider examples before we theorize about the semantics of the verb 'know'.

Is there any situation in which for some proposition that *p* (1)-(5) obtain? Yes; there is. Powers in [52] has discussed a beautiful simple example, which shows this.[22] (He, however, exegizes it somewhat differently from the way I do.) He considers the following question:

(Q) Is there (in English) a four-letter word that ends in EE, ENN, and WHY?

The reader may want to investigate the answer to (Q) before reading more of this paper. A reflective pause may not be unsalutary.

I assume that the reader has gone through the dictionary and has found that there is, at least initially, a difficulty in locating an affirmative answer to (Q). I am sure that the reader has also found that an affirmative answer to (Q) is delivered by the word 'deny'. Thus we have for proposition that *p* in (1)-(4):

p: There is (in English) a four-letter word ending in EE, ENN, and WHY.

Now, Andrea (like Powers himself, me, and many others, when first confronted with question (Q) has run through the alphabet mentally and has not located the word 'deny'. That occurred at time *t*. Hence (2) is true. But, obviously, Andrea would have emphatically answered "Yes" to the question:

(R) Is the word 'deny' a four-letter (English) word ending in EE, ENN, and WHY?

To the extent that Andrea can answer "Yes" to (R) she may correctly be said to know that there is an English four-letter word ending in EE, ENN, and WHY. Hence to this extent (1) is also true.

Powers used his example to conclude that we must distinguish between *propositional knowledge*, which is, *roughly*, the power to answer a question of the form "Is P the case?" from *cognitive knowledge*, which is, roughly again, the power to answer some questions [52], pp. 347 ff). This stipulation formalizes an important distinction. More important, the distinction builds on Powers's Platonic idea of construing knowledge as the power to answer questions. The distinction must not be construed, however, as dealing with two different senses of 'know' and 'knowledge'. *The unity of knowledge must be maintained*. That unity is the one that underlies certain constructions that puzzle Powers:

> The situation [concerning the distinction between propositional and cognitive knowledge] is *obscured* by our tendency in ordinary discourse to use "knowing-that" locutions *even where* no merely propositional knowledge is in question ([52], p. 347; my italics).

This fact of ordinary discourse is another *crucial datum*. One must consider it seriously. It seems to me premature to regard that tendency Powers mentions as obscuring a distinction. At least my *desideratum* is to develop an account of knowledge that accepts Powers's datum and illuminates it by showing the reason why ordinary discourse has it. Clearly, the drift of the desideratum is this: we want an account of knowledge that shows the fundamental unity of knowledge through Powers's constructions.[23]

It seems to me that the word 'know' does not have a different meaning in (1), related to question (R), from the one it has in (2), related to question (Q). It must be carefully observed that the unity of sense of 'knows' in (1) and (2), on which I am insisting, is perfectly compatible with the thesis that in some sense we are dealing in (1) with a different kind of knowledge from the one dealt with in (2). This harks

back to the exegetical thesis that the ambiguity of a sentence must *not* be taken automatically as, i.e., does not imply, the ambiguity of a component of such a sentence.

Assuming a sentence S to be ambiguous creates the problem of disambiguation. If we want to adopt a canonical notation that exhibits the ambiguity in question, we must have in the ambiguous sentence two (or more) readings that exhibit the diversity of interpretations. If there is no ambiguous expression in S, then we must suppose that some implicit sentential element has to be introduced. This additional exegetical principle does not automatically determine a path for disambiguation, only a framework.

In our Powersian example we want to recognize the Platonic-Powersian approach, which equates types of knowledge with powers to answer questions. Thus we can introduce the following canonical expressions for the disambiguation of the sentence 'Andrea knows that there is (in English) a four-letter word ending in EE, ENN, and WHY:

> (1a) At time t Andrea knows$_R$ that there is (in English) a four-letter word ending in EE, ENN, and WHY.
>
> (2a) At time t Andrea does not know$_Q$ that there is (in English) a four-letter word ending in EE, ENN, and WHY.

Here the subscripts refer to questions (Q) and (R) determining the species of knowledge involved. They denote operators mapping generic knowledge on species of it.

The tentatively proposed analyses (1a) and (2a) of (1) and (2), respectively, maintain the common meaning of 'knows' in (1) and (2). Through the subscripts attached to 'knows' they signal a differentiation of the *one genus* knowledge into *distinct species*. Naturally, there are other alternatives. One can take the needed elements to be, not operators specifying species of the relation *knows* between a person and a proposition or propositional function, but operators on propositional functions. In the latter case we could represent, e.g., (1), as of the form:

> (1b) R (at time t Andrea knows that there is [in English] a four-letter word ending in EE, ENN, and WHY).

Powers's example of questions (Q) and (R) partially confirms that, taking questions as determiners of epistemic contexts, *different normal contexts* may apply to a person at one and the same time. Thus we better acknowledge different determinate species of knowledge.

16. One Dimension of the Semantics of 'Know' and 'Knowledge': The Theory of Questions and Propositional Guises

We must generalize upon Powers's example. Both Plato and Powers start with a *what*-question, like (Q), rather than with a *whether*-question. Undoubtedly, there are many other questions: *Where? When? How? Why? What with? Whenceforth?* and so on. All these questions can be asked with the same resources of the sentence that formulates the correct answer. They are all questions that select a component of the correct answer as the interrogation point, so to speak. But such questions *do not* ex-

haust *the interrogation points of a proposition*. For instance, we can ask about the logical connection between two propositions, or propositional functions, when these form a compound proposition or propositional function. Yet ordinary language does not have easy mechanisms that allow one to ask such questions without describing, rather than merely presenting, the question-generating proposition.

Consider the proposition expressed by this sentence:

> (2) If all the flights were on time today, John is either in the men's room or in a telephone booth, or he did not come.

Logically, even if not practically, we could ask questions that take the main connective 'if' as the interrogation point, or the disjunction 'or' in the consequent, or the universal quantifier 'all' of the antecedent, or the time denoted by the locution 'on time', and so on. In sentence (2), it seems that each word, locution, and clause determines an interrogation point. This includes grouping questions like "If all the flights were on time today, where is John?" and the overall question "Is it the case that (2)?"

Perhaps we can gather the whole family of questions as follows:

> (*a*) There is the hierarchy of logical forms of a proposition, which is the correct answer to each member of a corresponding family of questions, e.g., in (2) we have the forms: p [of a proposition whatever], $p \supset q$ [of any conditional whatever], $p \supset (qvr)$ [of any conditional with disjunctive consequent], $(x)(Fx) \supset (qvr)$ [of a conditional with a universal antecedent and a disjunctive consequent], and so on. There are the deeper and deeper forms of (2), for instance, which result from analyzing any of its component concepts. If we analyze *flight, men's room, telephone booth, come,* etc., we find a very complex logical form underlying the ostensible form that sentence (2) reveals.[24]

> (*b*) We take each component, or part of a proposition considered as having a certain logical form, as an interrogation point, whether the part is a proper part or not. If we analyze any component in (2), e.g., 'flight' or 'telephone', we find more interrogation points.

The result is a hierarchy of the family of questions determined by a given proposition. Each level of questions corresponds to what in [11], part II, I have called a *propositional guise*. A propositional guise is, roughly, a proposition conceived as having a certain logical form. The theory of propositional guises permits a unified account of Moore's paradox of analysis, the discrimination of attention, and the increase of knowledge and belief through the exegesis of propositions.[25] Here I mention it because it has additional unifying power in bringing under the same account the diversity of questions and the diversity of epistemic contexts.

Following the Plato-Powers line, we take a person's powers to answer certain questions as determining species of knowledge. Perhaps some properties, whether qualities or relations, of objects are ultimate, truly unanalyzable. Perhaps some properties can be fully analyzed in terms of primitive properties. Most of the prop-

erties we consider in daily life, however, do not seem to be truly analyzable in a finite number of steps into ultimately primitive ones. Thus there seems to be no ultimate analysis of many propositions we know to be true. Hence the hierarchy of logical forms (and guises) of most propositions is not composed of (definitely) a finite number of rungs. Thus the hierarchy of questions determined by most propositions is not a finite system. Hence there is in principle an indefinite, and even infinitely large, number of species of knowledge.

The preceding result seems to explain why we lack in ordinary language a systematic mechanism for denoting, or signaling, the species of knowledge with which we are concerned. Obviously, we are not concerned in practical life with the whole hierarchy of logical forms (and guises) of a proposition, at least not as such. The interrogative locutions we have mark points of interrogation. By placing them in an interrogative sentence, we reveal both the level of logical form (and guise) and within this form the interrogation point in which we are interested. Compare, e.g., in the case above, the questions: "*Is it* the case that (2)?," "If all flights were on time today, *where* is John?," and "If all the flights were on time today, then, since he is neither in the men's room nor in a telephone booth, *where* is he?"

Now the question appears: If in ordinary language we lack a systematic mechanism for specifying which questions and, *a fortiori*, which species of knowledge we are interested in, how do we manage? For we do seem to manage to communicate about what certain persons know or, for that matter, fail to know. The answer seems to be this: *We manage to communicate about species of knowledge in exactly the same way we manage to communicate about the species of other properties, when we need a specification from a genus with indefinitely and perhaps infinitely many species*, namely: the context of communication pinpoints the relevant species.

What I am claiming for 'know' and 'knowledge' is in detail somewhat different from what happens to demonstrative words like 'here' and 'there' and from what happens to color words. Of course, there are fundamental reasons for the differences. The common claim is that in their primary uses all these types of expressions are at bottom indexical. (See [6].)

With 'here' and 'there' the expressions themselves, and the sentences in which they occur, need to be complemented by an association with a place of *possible* (*not* actual) utterance. In perceptul uses color words and the sentences containing them need an association with actual shades of color; but they can be used in a generic way. One reason, I suppose, is that the actual determinate color content of experience is normally not important in communication. Thus, in some uses, especially non-perceptual uses, color words do not denote a determinate shade determined by context, including pointing. They must, then, be understood as either denoting a generic, a determinable quality, or as being implicitly quantified. For instance, 'The book on the table in the chairman's office is red' may be taken as of the form 'There is a shade ϕness of the red family such that the book on the table in the chairman's office is ϕ'.

The epistemic words 'know' and 'knowledge' may perhaps be used in a generic sense, as the color words are. As in the case of these words, the generic use rests on

their use in which they denote determinate species contextually pinpointed. With 'know' and 'knowledge' it is not perceptual context, but the *assumed context of inquiry*, that pinpoints the determinate species relevant to episodes of assertion or of thought.

In each context of communication we are concerned with certain questions. Often there is no need to formulate them, because the courses of action we are engaged in, or are planning to engage in, demand certain answers. Thus the implicit set of relevant questions determines by itself the species of knowledge we are interested in. The context of communication, and the context of inquiry, whether the inquiry is carried out by one person or another, determines, therefore, a species operator, like the subscripts 'Q' and 'R' in section 15, which maps the generic knowledge into a relevant species of knowledge.

Aside from inquiry, a person has at a given time the epistemic powers to answer certain questions pertaining to a certain proposition p. Hence, independently of any inquiry, or of any context of communication, a person has *knowledge$_a$* that p for several classes a of questions that include questions pertaining to that p. The linguistic point we are making is that in a given context of speech, even if it is not communicational, a speaker will, in choosing his topic of discourse, select at least one class a of questions, with respect to which he attributes knowledge to persons. Patently, the selection of a is here chiefly the setting in readiness, even if not the actual triggering, of a bundle of dispositions to think of (and formulate) certain questions. Hence the selection of a class a of questions need not be exhibited in a listing of the members of a or in the formulation of a description of a.

17. Another Dimension of the Semantics of 'Know' and of the Theory of Questions: Methodological and Contextual Constraints on Interrogation Ranges

Each proposition, or state of affairs, we claim, determines a hierarchy of questions. This hierarchy is determined simply by the hierarchy of logical forms of the given proposition (or state of affairs). But the total logical hierarchy of questions is often too broad for practical concerns. We are normally interested in a modest segment of the interrogative hierarchy pertaining to a proposition. The segment in question is cut off from the total hierarchy by means of a set of constraints. In a given context of inquiry we are interested in questions that comply with certain conditions embodying the relevant constraints. Such conditions are part of the circumstances determining the truth of the proposition to which the question-*cum*-constraints, as well as the hierarchy of questions containing it, belongs. Obviously, a person may have the power to find a proposition P upon thinking a question $Q^c(P)$ with constraint c belonging to P, and yet that person may lack the power to think of P upon thinking a question $R^d(P)$ with constraint d also belonging to P.

The constraints on the relevant questions for a certain inquiry are part of the context of the inquiry. To the extent that a person may be simultaneously involved in different inquiries, that person may be involved in different sets of *normal circumstances* for the truth of the propositions of the inquires in question. More generally,

a person's power to offer a proposition P as the true answer to a question Q^c(P) may not be matched by a power to offer P as the true answer to a different question R^d(P). Yet the set of circumstances involved in the truth of P by constraint c and question Q may be just as normal as the set of circumstances involved in the truth of P by question R and constraint d.

In Powers's example of questions (Q) and (R) in section 15, there are no methodological constraints on the questions; the two normal contexts include simply the capacity of thinking the questions and of English words. In that sense the examples are oversimple.

Let us take a closer look at questions with constraints.

Each point of interrogation in a propositional guise (i.e., a proposition considered as having a certain logical form) determines both *a blank in the proposition*, so to speak, and a range of possible *fillers of that blank*. For example, the proposition (2) of section 16 above, analyzed as far as sentence (2) shows, has the point of articulation represented by the expression 'flights', which yields, among others, the following two questions:

(Q1) What events are such that if all of *them* were on time today, John is either in the men's room or in a telephone booth, or he did not come?

(Q2) What commercial events scheduled to take place here are such that if all of *them* were (occurred) on time today, John is either in the men's room or in a telephone booth, or did not come?

Patently, an analysis of the concept *flight* will open up more interrogation points within this concept in the proposition P expressed by sentence (2). Such interrogation points will yield further questions belonging to P with their proper ranges of fillers.

Question (Q1) determines a large number of fillers for the blank represented by the italicized pronoun *'them'*. However, the context of inquiry in which the question is raised may make it clear that certain fillers are not relevant. Question (Q2) cuts down the set of fillers by introducing some important constraints. Yet it may still be too broad as the formulation of the problem of the actual context of inquiry. Further constraints may make the set of relevant fillers a yet smaller subset of the total range of logically viable fillers.

There may also be methodological requirements on the relevant answers. These are external requirements that need not have anything to do with the logical structure of a proposition, or of each of the questions in the hierarchy of questions the proposition determines. The interrogation context may establish that a question is not answered by simply finding the proposition to which the question belongs. It may require that the proposition be found in accordance with certain processes or techniques. In order to know$_\beta$ that p, for some set β of questions with methodological constraints, not only must one have the power to think that p as the answer to mere questions in β: one must have the power to think that p as also conforming to the constraints attached to the questions in β.

The preceding becomes apparent through a simple reflection on: (*a*) what counts as knowing that Columbus discovered America on October 12, 1492, in a television quiz show, with (*b*) what counts as knowing that in a high school student's essay on Columbus's discovery of America, with (*c*) what counts as knowing it when a historian defends the traditional date of the discovery from some ingenious and famous Harvard historian's claim that Columbus discovered America on October 11, 1492.

18. The Hierarchy of Epistemic Powers and the Fundamental Question-Proposition Complexes

The Plato-Powers view, as discussed and endorsed above, makes knowing not just a matter of believing a truth with a certain type of justification but a matter of having the power to find a (true) answer to certain questions. This needs to be clarified in terms of the concept of power or ability used here; the conditions that activate such powers or abilities; the answering relation between question and answer; what it means for a person to believe that a certain proposition is a (true) answer to a question, and so on.

We have already provided the basic theory of questions by developing the conception of a hierarchy of constrained questions, determined first by the logical forms of a proposition, and then reduced both by the stipulation on conditions on the fillers for the blanks at the interrogation points and by methodological constraints. Naturally this theory of questions needs additional development. But what we have presented above suffices for the elementary sections of basic epistemology. We opt here for the simple terminology: *an answer P to a question Q is a true proposition P such that P satisfies the constraints on Q and Q belongs to P in the sense explained above*, namely: Q arises from P by the occurrence of a blank in a propositional guise of P, with the blank determining a subset of all the logically possible fillers for that blank. Thus 'true answer' and 'answer' mean the same thing.

Consider the following:

(1) At time *t* John believed that 'deny' was (is) a four-letter word (in English) ending in EE, ENN, and WHY.

(2) At time *t* John believed that 'deny' was (is) an (incomplete) answer to the question "What is a four-letter word (in English) ending in EE, ENN, and WHY?"

These sentences can be taken *literally* at face value, that is, as having all the expressions of the subordinate clause *internally* (*de dicto* many writers would say, in a terminology I do not like). If so taken, then the John talked about in (1) and (2) is claimed to believe the propositions formulated in the subordinate that-clauses. On that interpretation, (1) and (2) would express true propositions if John can understand the question and think of 'deny' at the same time. Thus in the internal syntactical construal neither sentence (1) nor sentence (2) describes John's ability to offer 'deny' as an example of a word ending in EE, ENN, and WHY when asked to do so.

Perhaps we can express that John has that ability by taking the expression 'deny' in (1) and in (2) as occurring *externally* (or *de re*, again in a terminology I do not like[26]). On this construal, (1) and (2) are not perspicuous, for they represent what is better put as follows:

(1a) At time *t* John believed 'deny' to be a four-letter word (in English) ending in EE, ENN, and WHY.

(2a) At time *t* John believed 'deny' to be an (incomplete) answer to the question "What is a four-letter word (in English) ending in EE, ENN, and WHY?"

These sentences, in their perspicuous use, assign the reference to the verb 'deny' to the speaker of the whole sentence (1a) or (2a). Sentences (1) and (2), on the other hand, in their internal construal represent a cumulation of references both by the speaker and by the John spoken of. Therefore, (1a) and (2a) leave it unspecified how John refers to the verb 'deny'; they leave it unspecified what exactly is the proposition John is said to believe, and thus they are *propositionally opaque* with respect to the subject of the proposition believed by John. Hence sentences (1a) and (2a) may express truths when (1) and (2), construed internally, express truths; but they may also express truths when (1) and (2) do not express truths, because John refers to the verb 'deny' in other ways than by using the expression '(the verb) 'deny''.

Yet neither (1a) nor (2a) describes John's ability to *offer* 'deny' as an example of a word ending in EE, ENN, and WHY when asked to do so.[27] This ability is not identical with, or part of, John's belief that 'deny' is composed of DEE, EE, ENN, and WHY and his ability to *understand* the question "What is a four-letter word (in English) ending in EE, ENN, and WHY?"

Whatever abilities a state of believing that *p* may consist in, it is clear from the discussion of the Powers example above, and from the generalization to all types of constrained questions, that the ability to offer an answer to a question does not consist in, but it does presuppose, both understanding the question and believing the proposition that answers the question. The ability to answer a question has to do with one's ability to *marshall evidence*, and this, in its turn, has to do with the organization of one's beliefs in the unconscious depths of one's mind. Knowledge has to do with structures of evidence. Clearly, a passive account of knowledge, in terms of beliefs and truths *only*, as we have already intimated, cannot be adequate; the analysis of knowledge must connect a person's beliefs with his abilities to marshal evidence.

The Plato-Powers Principle, that knowledge is power, particularly the power to answer questions, is, therefore, an important criterion of adequacy for any analysis of knowledge.

Socrates believed that knowledge is power in a much stronger sense than that enthroned in the Plato-Powers principle about knowledge. For Socrates, knowledge is a power to act, to produce physical movement other than the cerebral and vocal activities required, causally, for thinking and asserting. Socrates may have been right, and a comprehensive theory of knowledge should include the subtheory about the general connections between knowledge and action. These general connections in-

clude not only the most fundamental ones, e.g., that every proposition a person believes, *a fortiori*, every proposition a person knows, is a possible premise both for the acquisition of further beliefs and for the determination of what he or she ought to do. The general connections between knowledge and action include causal relationships between the beliefs and the powers composing knowledge, on the one hand, and the production of action, on the other hand. These causal connections include, but do not exhaust, the phenomenon of volition which welds together a person's action with his knowledge of what to do.[28]

Basic epistemology, our concern here, must deal with the most pervasive and fundamental connection between knowledge and action. This is the connection embedded in our powers to think answers to questions. These are the epistemic powers *par excellence*. They depend on the organization of our beliefs; their exercise, in its turn, reorganizes our beliefs. But this organization and reorganization of beliefs, which feeds and results from, respectively, our thinking questions and looking for their answers, is *not* itself a belief, not even the additional belief that the other beliefs stand in a certain hierarchical arrangement. The evidential organization of beliefs, which the engaging in questioning and answering brings about, is a *causal organization* of degrees of readiness one has to think (assertively, for the most part) this or that proposition.

In brief, *our most basic epistemic powers have to do with the causal arrangement of our degrees of readiness to bring forth premises we may need in our reasonings*. Because our reasonings may lead to diverse courses of action, our most basic epistemic powers are at the foundation of the connection between knowledge (and belief) and action.

The basic epistemic powers consist of powers to think questions, to think corresponding answers, and to be caused to think the answers by the very thinking of the questions. One need *not* have the additional linguistic power to use the second-order words 'answer' and 'question'. Consider, for instance, the case of a very young child, Amy. A visitor asks Amy: "Where is your baby doll, Amy, the one Santa brought you?" And Amy, unable to articulate the answer, takes the visitor's hand and leads her to her bedroom upstairs and says, "Here Baby."[29] Undoubtedly, Amy can think propositions and can think questions and can think propositions as answers to questions. These thinking powers have, of course, a foundation in her linguistic powers. But thinking powers, however grounded in linguistic power and control they may be, do run ahead of linguistic development.

Now, aside from the powers of articulation, we must distinguish between first- and second-order thinking. The fundamental thinking ability one needs is simply the ability to think first-order propositions and first-order questions. Thus one must be able to think, say, (a) *John is happy*, (b) *Is John happy?* and (c) *Who is happy?* in order to "see" that (a), if true, stands in a special relationship to (b) and (c). One can, to put it differently, think (A) or (B) below:

(A) Is John happy? . . . [Examination of beliefs and perceptions] Yes, John is happy.

(B) Who is happy? . . . [Examination of beliefs and perceptions] Ah! John
is happy.

One can do so *without* being able to think either of the following:

(A+) Question: "Is John happy?" . . . [Examination of beliefs and percep-
tions]

Answer: "John is happy."

(B+) Question: "Who is happy?" . . . [Examination of beliefs and percep-
tions]

Answer: "John is happy."

What we need at the foundation of our epistemic powers is merely the power of
making an *Operational connection* between answers to questions as in (A) and (B).
These illustrate the primary dialectic of thinking. We can, I suppose, speak of this
operational connection as the *Ah!-Yes* or, simply, YES connection between a ques-
tion and a proposition taken to be one of the question's answers. This operational
connection must be carefully distinguished from the *predicative relation*, which one
can passively contemplate as holding between a question and each of its answers.
This relation is the one being thought of in (A+) and (B+). This relation is expressed
by the predicative locution *is an answer to* and, naturally, the sentences expressing
the second-order propositions involving this relation require nouns to flank 'is an
answer to'. For an example consider:

(3) "John is happy" is an answer to "Who is happy?"

Here the quotation marks form nouns. On the other hand, the expression 'YES',
which we have just stipulated to be a signal of the operational connection in under-
lying (3), needs no nouns, but applies directly to sentences, and perhaps forms with
them a sentence expressing a compound proposition. Thus we need no quotation
marks or any other nominalizing device and may simply write:

(3*) Who is happy? *YES* John is happy

to represent the operational connection underlying the relation *answering* denoted
in (3).

In general, I am proposing here to use the word 'YES' as a technical symbol·
for the most basic connective linking a question and a proposition that may be
thought to be an answer to that question. Thus we have *question-proposition com-
plexes* of the form:

(4) Q? *YES p*.

Once again, complexes of form (4) are *not* second-order relational statements like
(3); they are first-order complexes. They are not questions. They may be regarded
as propositions, since they seemed to be believed by a person who reasons as in (A)
and (B) above. The main principles of truth values for such complexes are these:

(Q*p*.T*g*) A question-proposition complex of form "*q YES p*" is true, just in
case the question *q* belongs to the hierarchy of questions determined
by the proposition that *p*; otherwise it is false.

$(Q_p.T_i)$ A question-proposition complex of the form "q YES_ip," where i denotes, or signals, a set a_i of constraints, is true, just in case the question q belongs to the subhierarchy of questions determined by that p and is governed by the constraints in a_i; otherwise the complex "q YES_ip" is false.

We have no room here to develop the fundamental logic of the question-proposition complexes. Obviously, no analysis of knowledge along the Platonic-Powersian lines followed here can be fully illuminating, until that logic is fully developed, too.

Evidently, a thinking being may be able to think question-proposition complexes, even if it is not able to think the corresponding second-order classificatory relational propositions about questions and their answers. A small child, for instance, may be able to think complexes like (3*), but not propositions like (3). Thus there is a hierarchy of basic *epistemic powers*: at the very bottom, the raw epistemic power is the power to be caused to think that p by the very thinking of the question q. This is the power that children, and perhaps sophisticated primates and pets, acquire first. Then there is the power to think complexes of the form "q YES p" and "q YES_i p" for constraints that merely demand the normality of the assumed circumstances. Then comes the power to think complexes of the form "q YES_ip" for abnormal constraints, and later on the powers to think the most sophisticated constraints. Somewhere along the development between normality and abnormality there comes the power to raise to the meta-language: to have thoughts of the relational kind of the form "The question q has as an answer (the proposition) that p."

This hierarchy of epistemic powers provides room for broad and narrow uses of the epistemic vocabulary, especially the words 'know' and 'knowledge'. Do house cats or domestic dogs know, e.g., that their owner is in the house, is angry, is loving? This all depends on what powers of thinking we attribute to cats and which ones to dogs. And small children?

19. The Truth-Circumstances of the Known Truth Involved in the Constraints of Questions Need Not Connect the Knower Causally with That Truth

It is widely taken for granted that for a person to know that p there need be neither a common cause of the person's believing that p and of that p nor a causation path from that p to the person's believing that p.[30] I believe this to be correct.

We cannot review the arguments in favor, or against, the causal account of knowledge. To the attentive reader, it will be obvious that some of the examples offered above, especially *The Non-Normal Cheater* subseries, raise problems for most causal accounts so far proposed. Yet our present interest lies in revealing some general features of epistemic contexts of justification.

The Blind Tiresias. Tiresias, blind seer, has the extraordinary power of seeing very vivid visions in his imagination. This happens after he has eaten cucumbers marinated in a mixture of tequila and some very secret concoctions. Tiresias's visions come, he says, with a date at the bottom. His visions have been found

to correspond point by point to events in the past, or events in the future, or events simultaneous with his visions, in accordance with the dates he sees. Yet there is no way those events can be causally connected with Tiresias's visions. This power came to Tiresias, as his memory goes, on his twenty-fifth birthday. He paid no attention to his visions for about three years. But then some reports on the newspapers described what he had seen, and then be began to believe his dated visions correspond to reality.

Does Tiresias know that his current vision is true? Do we know—we who have tested him thousands of times and have found no single failure? Here again we need the notion of stable, normal circumstances. Tiresias's early beliefs were not grounded at all. His later ones are grounded on the generalization that *caetaris paribus* his visions depict reality, that is, given the assumption of a *normal context of circumstances*, connecting the occurrences of his visions with what they depict. If the generalization, relative to the normal circumstances, only part of which are described in *The Blind Tiresias*, hold, then his belief corresponds to reality and, it seems to me, it constitutes knowledge.

Perhaps Tiresias's normal circumstances include elements we have no idea of. If so, regardless of how bizarre those elements may be, the implicit assumption that the circumstances are normal remains valid. If some bizarre circumstances affect the normality of the context, then Tiresias will fail to know, even if those circumstances come in balancing pairs, just as, the reader may recall, it happens in *The Cheater Series*. Then, if Tiresias knows of (or has reasonable belief about) those circumstances and considers them in his claims about the truth of his visions, Tiresias again knows that his revised generalization is, by hypothesis, true.

There is *no* need for Tiresias to know that his beliefs are causally connected with what he believes. This is more clearly so, the more detailed his visions and beliefs are about the events depicted in his visions. At this juncture serious and profound problems appear regarding how he can think of events he cannot be causally connected with. Undoubtedly, if the contents of consciousness are essentially universals, then Tiresias, like anybody else, thinks of all sorts of particulars in terms of traits or characteristics, and their relations to himself and to the particulars present in his perceptual fields.[31] These issues are difficult, yet they must be dealt with in a comprehensive epistemological theory.

20. Normal Truth Circumstances and World Order

We have returned to the most crucial and deep-seated assumption underlying our cognitive claims, namely: that the truth circumstances of the known propositions are normal or have identifiable abnormalities. We have already formulated in sections 13-17 some of the most fundamental principles governing the epistemic role of normality. Let us say something about what normality consists of. To begin with, the most general and fundamental principle is this:

MNP*. *The master presupposition about normality.* The changes in our immediate environment have a reliable order, determined by: (1) a hier-

archy of general principles and laws, which we cannot for the most part (at least at the moment) specify, and (2) a set of particular relevant regularities, which we can specify, and both the general order and the particular regularities remain in operation at least during the period with which we are concerned.

The idea that there is an underlying world order that abides but that we cannot specify at a given moment is the permanent assumption that not only unifies each personal life but also unifies all of us as members of one epistemic community. The assumption of a deep-seated world order that sustains our actions, but an order we cannot specify beyond some particularly relevant regularities, is needed to plan actions and count with our being able to carry out our plans. Given our thorough ignorance of the ways of the world, we need the master presupposition MPN*. This is a most economical assumption: we can carry on our tasks without tarrying to formulate the principles that govern the order of the world. It is a vital assumption, since we must live and do our things, regardless of how much knowledge of the order of the world we have and regardless of how much time we have to find out what the structure of that order is.

We cannot know that the world is fully ordered. Some of us, scientists and philosophers, are discovering more and more principles of world order. But we lack the power to answer *all* the questions about the order of the world, especially if we place on our questions some strenuous methodological constraints, like those now applicable to the natural sciences. Thus the fundamental principle MPN* is part of the unspecifiable framework within which we determine that a certain person knows certain truths. On the other hand, the fundamental presupposition and the other structural presuppositions constituting frameworks within which beliefs turned out to be justified, or not, cannot themselves be justified within those frameworks. Whether there are more abstract, or more comprehensive, frameworks within which we can ask whether MPN* is justified, or not, is a most important *topic—but it goes* beyond basic epistemology. In any case, we have the principle:

> NSSJ. Structural assumptions or presuppositions about the world, or a part thereof, which determine epistemic frameworks, are not self-justifiable and are not justifiable in the frameworks they determine.

This principle ties the *Gestalt* character of knowledge, introduced in section 2, and the contextual determination of species of knowledge, discussed in section 15. To know something, a person must have a battery of beliefs of sorts, but these beliefs, which by converging together on other beliefs confer upon these the status of knowledge$_a$, are not themselves known$_a$.

Several pervasive assumptions govern the epistemic role of truth circumstances. They are of different sorts: some stipulate general features of the world order, others stipulate connections between the world order and the mind. The following are just *some* of the most widely discussed:

> EP. *Experience principle*. Every proposition that a person is justified in believing must be available to the person through experience.

EBP. *Empirical-basis principle.* The strongest availability through experience of a proposition P for a person S at a time *t* is for P to be fully present in a specious present experience of S at *t*, i.e., for P to be at *t* either the total content of an episode of consciousness of S or the total content of a segment of one such episode.

RP. *Rationalist principle.* Propositions available to a person S need not be exhaustive contents of S's episodes of consciousness (as in EBP), at a time *t*, or during a set of times, but then such propositions must be connected to the experiential contents of S at some time or other and to S's beliefs, through a network of principles that compose in part the order of the world.

LKP. *Leibniz-Kant's principle.* Both the principles of deductive logic and the principles governing structural relationships between concepts (or properties) are principles of world order.

PHOW. *Principle of the Hierarchial Order of the World.* The different categories of entities or states of affairs composing the world are connected to one another by principles of ranking of the categories, escalation from one category to another higher up, descension from one category to another lower down, and commutation across categories not linearly ordered. Principles of these different types link: (a) sensory experiences to physical objects; (b) physical states of affairs to mental states of affairs; (c) perceived objects to unperceived ones; (d) experienceable states of affairs to theoretical ones; (e) particular propositions to generalizations; (f) witnesses' testimony to attested matters; (g) facts to values; (h) values to obligations and rights; and so on.

PIAWO. *Principle of the inferential nature of our attitude toward the world order.* We tend [for economy of thinking and action] to treat the principles of world order we recognize as built-in mechanisms of inference, rather than as formulable beliefs and explicit premises in reasonings. The most abstract and pervasive principles of world order [like the ones on this list] are seldom, if ever, thought of, although they structure most of our reasoning and thinking.[32]

HHP. *Heidegger-Hahn's principle.* We always find ourselves in the midst of a world, which has an assumed general order, and we have a conception of some aspects of that order.

DHP. *Duhem-Hahn's principle.* The world order is not fixed for all times; it may be changed at will, but only piecemeal.

QP. *Quine's principle.* The hierarchy of the world order for a person S at time *t* is a hierarchy of X's degrees of willingness to resist a renunciation of a given proposition.[33]

DA. *Descartes's axiom.* The *cogito* propositions about a person S and a

time t to the effect that S is at t having *himself* such and such experiences *then* are maximally resistant to renunciation by S at t [Here *'himself'* and *'then'* are quasi-indicators.[34]]

CP. *Cogitatum principle.* The *cogitatum* propositions pertaining to a person S and a time t to the effect that, within S's total field of consciousness, including S's perceptual fields, such and such occurs, are maximally resistant to renunciation by S at t.

These general principles do not deliver particular regularities that a person can use as premises in building up his knowledge of the world. They, with other similar generalities, merely constitute the broadest schema within which we posit less pervasive, more definite, and more rigorous principles and laws. All such principles and laws, then, constitute the framework within which a person's beliefs are to cohere.

The abnormalities of the truth circumstances of a given proposition about particular matters of fact always deal with specific regularities within the general framework. More general propositions may involve abnormalities revolving around more general regularities. More general laws may be involved in the abnormalities of the truth circumstances in which laws are considered. Finally, the very innocuous-seeming principles listed above and their ilk may yield abnormalities, at least merely thinkable ones, in the truth contexts in which the most general laws of nature, or other fundamental principles, are examined. This hierarchical arrangement of real and merely thinkable abnormalities of the relevant truth circumstances for the justification of belief is simply a consequence of the hierarchical structure of the world order.

We have barely touched the surface of the huge topic of normal relevant truth circumstances. But this must suffice for our present reflections in basic epistemology. We must bring these reflections to an end by tying the several strands of the fabric of knowledge, which we have deployed, in order to take a preliminary look at the weave of knowledge.

II. A PRELIMINARY ANALYSIS OF KNOWLEDGE THAT *P*

1. The Elements of Epistemic Contexts

In our exegesis of data in part I, we saw how species of knowledge that p come about, and how those species are determined by different parameters. We noted how there is a framework of beliefs taken for granted, which scan the surface of an assumed normal context, and beliefs explicitly formulated or formulable, which constitute specific pieces of evidence and justification for other beliefs. These latter beliefs are the ones whose epistemic status is at issue. We noted how a set of questions with constraints, often assumed and left to the context of communication to reveal, determines both a set of epistemic goals and a set of possible epistemic powers. We also noted that our beliefs depend on, and include, the beliefs of others, in spite of the inherent unreliability of the reports of others.

In brief, an epistemic context is an ordered sextuple:

$C_i = \langle F_i, A_i, S_i, R_i, Q_i, P_i \rangle$, where:

1. F_i is the set of presupposed "facts": the relevant beliefs about the issues pertaining to P_i, which the members of A_i have, including beliefs about what other members of A_i believe and say, if the testimony of others is relevant.

2. A_i is the set of agents in the relevant epistemic-community involved in deciding the issues pertaining to P_i.

3. S_i is the hierarchical network of structural presuppositions and assumptions, held by at least some members of A_i, which are presumed relevant to the determination of the truth, or falsehood, of the propositions in P_i by the members of F_i and of R_i: these are the presuppositions and assumptions that function in the mental economy of the persons in A_i as mechanisms of inference, rather than as major premises.

R_i is the set of relevant respects of abnormality applicable to S_i.

Q_i is the set of questions, governed, perhaps, by some constraints on the range of the blanks in the questions and by some methodological constraints.

P_i is the set of propositions whose epistemic status is at issue: they are the possible answers to the questions in Q_i.[35]

2. A Tentative Analysis of Knowledge that P

The proposal that ensues is only tentative. I am not even sure that an explicit definition of 'X knows that p' is the accomplishment we should expect in basic epistemology. As I have argued elsewhere, the definitional method of classical analytic philosophy cannot, in spite of the illuminating power it may attain, provide the full philosophical illumination some of us desire. This is so, partly, because definitions are subject to the strict rule of deploying sets of conditions that are necessary and sufficient. Thus we can easily miss the illumination that the discovery of conditions that are necessary, but not sufficient, or vice versa, may deliver.[36]

Naturally, definitions of certain concepts are sometimes feasible. Furthermore, we can offer here a tentative analysis of knowledge that p, which brakes the standard approach at several junctures. *First*, we provide a development of the Plato-Powers principle that cognitive knowledge has to do, not just with beliefs and truths, but also essentially with powers to answer questions. *Second*, we assume the theory of questions sketched out in sections 16-18. *Third*, we take into account the multiplicity of species of knowledge and their indexical specification by epistemic contexts. *Fourth*, we weave into our analysis the crucial feature of the normality of a context and its respects of abnormality. *Fifth*, we are not, besides, providing a Fichtean analysis of knowledge. *Sixth*, we are also not requiring probability computations; but probability may be included in the constraints on the questions in Q_i. Recall, nevertheless, that the analysis schema is only the tip of the iceberg described in part I.

Now the plunge: Recall that '$q\text{YES}_i p$' means, as explained in section 18, the first-order non-classificatory proposition that (question) q has (proposition) p for a (true) answer; also recall the truth-condition for such propositions given in section 18.

The Contextually Indexed Analysis of Knowledge that p is tentatively deployed in the following schema, where the subscript '*i*' is a schematic letter signaling epistemic contexts, its proper substituents are expressions denoting epistemic contexts, each of which expressions, by being subscripted to 'know', represents a contextual operator having generic knowledge as operand and having a (contextually identified) species of knowledge as value. We assume, thus, an epistemic context:

$$C_i = \langle F_i, A_i, S_i, R_i, Q_i, P_i \rangle$$

as characterized above. Then:

C*i*K. At time t X knows*i* that p, if and only if:

    (i) At time t X believes that p;

    (ii) that p is in P_i;

    (iii) There is at least one question q in Q_i such that:

        1) q YES*i* p,

        2) if at t X considers q and reflects, searching within his beliefs for an answer to q, X can psychologically think, believingly, at t' that qYES*i*p, where the interval (t, t') is a retrieval and computation time that conforms with the constraints governing question q, and

        3) at t X believes that qYES*i*p;

    (iv) There are truth circumstances Z for at least some members of P_i, there is a positive number b, and there are respects of abnormality r_i, . . . , r_b in R_i, such that:

        1) Z is a set of normal circumstances for that p, except for including r_i and . . . and r_b,

        2) Z obtains, and

        3) at t X believes that the truth circumstances for that p are normal except for including respects r_i and . . . r_b;

    (v) There is a subset s_i of S_i such that at t X has the propensity to make inferences in accordance with the members of s_i; [it is not ruled out that there may be unconscious, or subconscious, processes of inference and that some of them may occur in accordance with the members of s_i];

    (vi) There is a conjunction E*i* of members of F_i such that:

        1) given Z and s_i, if E*i*, then p,

        2) both s_i and E*i* are true, and

        3) at t X believes that both E*i* and, *ceteris paribus* [that is, given his believing-inferential attitude toward s_i and his believing that the truth circumstances for that p are normal except for r_i, . . . , and r_b], if E*i*, then p.

NOTES ON THE SCHEMA

    1. The locutions within square brackets in clauses (v) and (vi) indicate remarks and references by us, or whoever uses schema C*i*K to attribute knowledge that p to

someone; they do *not* represent references or parts of the thought content of the person X who is said to know.

2. Particular exemplifications of schema C_iK may have to be carefully formulated because they may require quasi-indicators, so as to attribute to knower X demonstrative references to time (by his using, in English, mainly 'now'), space (by his using 'here'), self (by his using 'I'), and so on. See [6] and [10].

3. Schema C_iK applies to mathematical knowlege. The stringent standards for mathematical *proof* prevailing now at the end of the twentieth century are much higher than those prevailing during the seventeenth century, and these were higher than those prevailing two centuries before. All those differences belong to different epistemic contexts, and are generically represented in the constraints built into the questions in Q_i.

4. Schema C_iK has the following pattern:

     (i): the belief condition;
     (ii): the relevance condition;
     (iii): the epistemic-power condition;
     (iv): the normality-abnormality condition;
     (v): the inferential-power condition;
     (vi): the evidence condition.

3. Skepticism, Scientific Knowledge, and Other Matters

Schema C_iK conforms, I believe, to *all* the criteria of adequacy encountered in part I. The reader will, undoubtedly, investigate whether this is so or not.

C_iK has the virtue of indicating where different developments belong. Each of the different parameters of an epistemic context gives rise to a branch of general epistemology. In particular, different problems in the methodology of science concern different sets of constraints to be placed on the questions determining unified sets of epistemic goals.

The issue of skepticism, its types and variants, must be reopened. Obviously, different epistemic contexts give rise to different types of skepticism. The most radical forms of skepticism involve contexts with the slimmest possible parameters. A radical version of skepticism, for instance, works with the following epistemic context:

⟨ϕ, {X}, {deduction rules}, ϕ, {all logically possible questions}, {all propositions}⟩, where ϕ is the null set.

Clearly, some extreme forms of skepticism are indefensible, whereas others, especially local ones about some definite small group of propositions, are wholly defensible. Determining where to draw the line between defensible and indefensible skepticisms requires careful investigation. Things concerning skepticism are, it seems to me, more complex than they are usually taken to be. One thing is certain: however indefensible a form of skepticism may be, this should be shown by the nature of the epistemic context involved in that form of skepticism. All simple appeals to the nonsensicality of skepticism are better shunned. No skeptic need use 'know' in a nonordinary way.

4. Colophon

This essay has naturally culminated in analysis schema CiK—not so much because it unfolds a new, Platonic-Powersian inferential, and indexical approach in basic epistemology but because it represents a sustained methodological effort at complying with both admonitions, Sellars's and Wittgenstein's, quoted under the title.[37]

Notes

1. The similarity between this sentence and Husserl's slogan about returning to the things themselves is *not* coincidental. What I am calling protophilosophy, i.e., the collection and exegesis of data about certain types of experience is, in general terms, similar to Husserl's phenomenological description. One crucial difference is this: protophilosophy is more outspokenly linguistic than is Husserl's description of essences, because it regards syntactical contrasts in ordinary languages as fundamental philosophical data. For additional discussions of philosophical method see [9] chap. 6, [13], [14], and the complementary references mentioned in them. In the first reference, the contrast between local and comprehensive theories is illustrated and discussed.

2. We will not survey here the huge post-Gettier literature. This literature is impressive not only because of its bulk but also because of the complexity and the ingenuity of a large portion of it. We will only be able to make some contrasting references to a few of the best or most influential pieces. But others not mentioned are, if not better, equally ingenious and intriguing. *One question for the sociology of philosophical knowledge*: the search for a definition of knowledge during the last fifteen years has been typically an American quest—why?

3. This tentative analysis of contextually determined species of knowledge is inspired by the analysis of the justification of actions in terms of what I call the Legitimacy of actions practically considered, or practitions. (See [9], chaps. 5 and 6.) The pivotal analogy is that of 'X knows$_i$ that p' to 'X ought$_i$ to A', the subscript 'i' signaling in the two cases a context of justification: in the case of ought$_i$ we have the context i of Legitimacy of the action (or practition) X ⋯ to A; in the case of knows$_i$, the context i of the justification of the state of affairs *it being the case that p*, rather than that of the psychological state of *believing that p*. See [9], chap. 8 for the justification of norms, or ought-statements, in the sense of the establishment of the truth of a norm. In spite of the analogy between ought$_i$ and knows$_i$, it is worth noting that the "ethics" of belief, to use Chisholm's expression in [15], is not really an ethics or a genuine system of norms. This is so for many reasons. One is this: a person does not have the crucial freedom, characteristic of actions, whether to believe or not to believe. Another reason, grounded on the finitude of the mind is this: whereas one ought$_i$ to do the actions (practitions) implied by what one ought$_i$ to do (see [9], chap. 7 for relevant distinctions), it is just not the case that one knows$_i$, or believes, or ought epistemically to believe, a proposition implied by what one knows$_i$ or believes. One may even be unable to think such a proposition, let alone believe it. There are no obligations, epistemic or otherwise, to believe a proposition. (Of course, it may be *better* to believe than not. But this is something else.) On the other hand, whether one can, or cannot, think of an action implied by what one ought$_i$ to do is immaterial; one still ought$_i$, derivately, of course, to do it.

My original plan was to write a series of essays, or one with a section devoted to the justification of belief, which section could show in detail the connections depicted in the following chart:

	reasons for:	*reasons for:*
Y {believing, supposing}	the proposition that X is P	X's being P
Y {endorsing, wanting}	the command X to A	X's A'ing
X {deciding, intending}	to do A	My (X's) A'ing

A very brief preliminary discussion of this chart appears in my note "There Are Command Sh-Inferences," *Analysis* 32 (1971):13–19. The issue is yet of the utmost importance. It has to do with the fact that reasons for believing that *p* (*a fortiori*, grounds for knowing) are founded on reasons for *it being the case that*, just as reasons for wanting agents to do action A are grounded on reasons for *agents' doing A*. Here *agents' doing A* is not a propositional matrix. (See [9], chaps. 6 and 7.) This connects with the sense in which propositions, rather than states of belief, are warranted. But all this must be left for another study. See note 17 of this paper for another influence of my theory of action on my theory of knowledge.

4. See the essays mentioned in note 1.

5. This crucial and pervasive feature of knowing is, probably, part of what moved J. L. Austin in [2] to claim, outrageously, that one's assertion "I know that *p*" is like one's assertion "I promise to A": one gives one's word that *p*. I say "outrageously" because, for one thing, whenever one asserts something one gives one's word, so that the likeness is too general to be informative. For another thing, the contrast between "I know" and "He knows" does not correspond to the contrast between "I promise" and "He promises": My "He promises" about Karl is parasitic on Karl's "I promise"; my "He knows" is not parasitic on Karl's "I know." Furthermore, to the extent that my "I know" gives my word non-trivially, my "Karl knows" compromises me just as much as it compromises Karl, in a way in which my "He promises" does not compromise Karl or even me.

6. The elementary *Gestalt* character of knowing has been emphatically affirmed by Sellars (see [56]–[58]), Leibniz (see [43]), Kant (in [38]), and others. It appears in the coherentist view of knowledge and in the coherentist view of truth. It is a datum for epistemological coherentism, but not for truth coherentism. In any case, the fundamental *Gestalt* character of knowledge must be carefully distinguished from all coherentisms. Thus it must not be confused with the issue between coherentism and foundationalism in epistemology. On this issue there is a huge literature, but a large amount of ground is covered by [1], [18], [31], [42], [48]–[50], [53], [56]–[58], and [64]–[67], and [74].

7. For a fundamental theory of communication and for a discussion of the propositional transparence of language, some of its limits, and its contrasts with Quine's referential transparence, see [10], [12].

8. The original plan for this essay included the writing of an appendix devoted to an assessment of Lehrer's definition of knowledge in [42]. That definition is the culmination of many years of reflection on the topic, and it is mounted on the most complex and illuminating examination of alternatives and consideration of rich data. But there is no space here for that. It must suffice to note that Lehrer's analysis suffers seriously because it hinges on probability comparisons. (It is also too restrictive to illuminate fully our ordinary cognitive experiences. It is too drastic in its ruling falsehoods out of the path to knowledge, as we note in section 11; that ruling is also not sufficient, falling prey to the data labeled *The Non-Normal Cheater Series* in section 13. It is Fichtean, as noted in section 7.) Lehrer's main clause for belief justification is literally as follows:

(iv) *h** is the strongest competitor of *h* for S if and only if *h** competes with *h* for S and, for any *k*, if *k* competes with *h* for S, then *p(h*)* is at least as great as *p(h)*. . . .

[Here, '*p(h)*' means (as Lehrer explains on p. 201): 'the chance S believes *h* to have of being true within his corrected doxastic system'.]

. . . Thus we obtain the following *final result*:

(v) S is completely justified in believing that *h* if and only if *p(h)* is greater than *p(h*)*. (See [42], p. 207; my italics.)

The hero in *The Minimally Cautious Gambler* believes correctly that *p* (he will win the prize) = .501, and that *p* (the strongest competitor) = *p* (he will not win) = .499. Hence, by Lehrer's (iv) above, our minimally cautious gambler is completely justified in believing that he will win. That belief conforms to Lehrer's other conditions for knowledge: it is true that *p*; S believes that *p*, and "S is completely justified in believing that *p* in some way [in every way, in our case] that

does not depend on any false statement" ([42], p. 21). Therefore, the minimally cautious gambler knows in Lehrer's analysans sense of knowledge that he will win the prize. Patently, as far as normal experience goes, he does *not* know that he will win the prize.

Pastin, in [50], has raised some serious difficulties against Lehrer's analysis. A very tough one, also arising from Lehrer's stipulation of a deep, thorough, and essential involvement of knowledge with probability, pertains to our perceptual knowledge. Remembering that one has often experienced illusions, and sometimes hallucinations, decreases the probability one believes one's perceptual judgments to have.

9. This idea that the evidence must suffice for the truth of known propositions has been appreciated by most writers on basic epistemology. A particularly interesting discussion is Tienson's [71] and Tomberlin's [73]. [71] argues that most of our attributions of knowledge are literally false—just as most of our attributions of flatness to the ordinary surfaces we encounter in experiences are literally false. This is an important thesis that needs serious reflection. It accounts very well for our use of the word 'flat', and it illuminates the difficulties encountered by basic epistemologists in their search for an analysis of 'knows'.

Most basic epistemologists simply follow a line like Lehrer's: allow the evidence to remain insufficient for the truth of what is known. But this raises difficulties. Two apparent exceptions are Skyrms in [60] and Dretske in [25]. These writers hold that the reasons for what one knows must be *conclusive*. I said that they are "apparent exceptions" because the conclusiveness of their reasons is not logical, but causal or subjunctive, depending on the circumstances. This is actually a step in the right direction, but it needs to be complemented with a study of normal and abnormal circumstances. (See sections 13 and 20 below, where we provide the beginning of that study.) Thus Skyrms's and Dretske's definitions of knowledge fall prey to the data contained in *The Non-Normal Third Cheater Series*, discussed in section 13. Valuable and ingenious attacks on Dretske's [25] appear in Sosa's [64] and Pappas and Swain's [47].

10. For a sustained and impressive argument for the view that induction does not include, or rest on, the calculus of probability, see Cohen's [24].

11. Other reasons against the rule that knowledge implies one knows appear in Hilpinen [34], Powers [52], and Tomberlin [72]-[73]. The original defense of the rule, which gave rise to much discussion, appears in Hintikka [36], pp. 17-22, 24-28, 104-13, *et al*. Most critics of Hintikka's defense of the rule have not noticed that Hintikka insisted from the very start that he accepted the rule with one qualification, namely: "only if the person referred to by *a* (the subject expression in 'a knows') knows that he is referred to by it" ([36], p. 106). He first formulated that qualification on pp. 158f. Hintikka's qualification conforms to what he characterized as his data, which includes the situation a person is in when he or she is in a position to say "I know" ([36], p. 33). Thus the main issue is whether Hintikka's [36] studies, and illuminates, a concept of knowledge that is widely used in life, *not* whether the rule, which he takes the concept of knowledge he is discussing to obey, holds; it does hold for that concept. A secondary issue is whether Hintikka's formulations of the main rule, and of the additional condition for it, in his system, are satisfactory. This issue is taken up in [7], sections V-VII.

12. This was first pointed out by Pastin in [50] without calling it Fichtean.

13. Given his Fichtean view of the mind, Chisholm is a notable example of a *cogito-foundationalist*. For his building knowledge on *cogito*-propositions and assimilating *cogitatum*-propositions that are not about perceived physical objects to adverbial modalities of the self, see, e.g., [19], chap. 1; and see [14] for a discussion of these views of Chisholm's.

14. See, for instance, the works, mentioned in the bibliography, by Clark, Dretske, Ginet, Harman, Hilpinen, Klein, Lehrer, Paxson, Sosa, Swain, and Thalberg. A nice survey of the main branch of the defeasibility approach appears in Swain [69].

15. Condition (*a*), strengthened with a causal stricture, appears in Swain [68]. Condition (*b*) has been demanded by Clark in [21], Dretske in [25], Ginet in [28], and Thalberg in [70]. For examples of critical responses, sometimes with very ingenious, even brilliant counterexamples, see Coder [22]-[23] (against Thalberg), Pappas [46] (against Dretske), Harman [31], and Lehrer [42]. *The Non-Normal Cheaters* described in section 13 below are useful counterexamples. See note 21 below.

16. Lehrer has formulated several versions of condition (B) in a string of essays that preceded [42]. This book contains a record of those earlier proposals of his as well as of others.

17. This important datum has been appreciated by a good number of epistemologists. Yet it has not been appreciated fully that, as we show in this section and next, with *The Non-Normal Cheater Series*, that datum condemns the puritanical view of the evidence constitutive of knowledge. An excellent protophilosophical paper presenting examples of this datum is Coder [23]. Coder makes clear the role of the normalcy of the truth circumstances of the known proposition by exegizing a three-piece datum. I must confess, however, that I did not appreciate the force of Coder's unified set of examples, nor the depth of his suggestion, until I had reflected on the topic on my own and was surveying the post-Gettier literature. His suggestion is this: "A comparison of the three cases with one another suggests that what is needed for knowledge in addition to justified true belief is *simply* that *one's total picture of events, from first evidence to last belief, be not too skewed*" ([23], p. 116; my italics). He is essenially right: this is simply the condition that has to be added to the classical conditions. On the other hand, it is not a simple matter to spell out this condition. The present essay is, in fact, a preliminary effort at giving an account of what is needed for a person's picture not to be too skewed for him or her to have knowledge.

I diverge from Coder when he requires that the knower's picture cover all the ground from first evidence to last belief, if 'first evidence' is taken in a strong foundationalist sense. I doubt very much, however, that Coder had this sense in mind. I hold very firmly the view that there is an important analogy between knowledge and intentional action. For some foundationalists of intentional action, an action is intentional only if *every* segment of it is intentional, including some basic actions. I believe that this is incorrect. Whether they are just bodily movements or not, basic actions are intersubstitutable. See [13] for five arguments against this view in the sophisticated form Alvin Goldman has given it. Similarly, I believe that for John to know that *p* it is not necessary that someone, whether John or not, be able to trace his belief that *p* to evidence going back to basic knowledge, e.g., John's perceptions. *Perceptions are also intersubstitutable*. At any rate, our exegesis of the phenomenon of knowledge in this essay does not reveal any such foundationalist character of knowledge.

18. This principle became very clear to me in 1963, when I was composing [5], which on p. 511 describes the kernel of all those examples in this section 11 illustrating it.

19. Most basic epistemologists have made references to the circumstances of the knower and of what he knows. Yet it has been a troublesome, often unacknowledged factor. For some of the difficulties it has created, see the papers mentioned in note 15. As far as I can detect, the constructive awareness of the normality of the circumstances appears more clearly first in Coder [23] (see note 17 above), and, even more definitely, in Sosa [64]-[65]. In [64] Sosa presents a valuable discussion; he attempts to elucidate the reference to normal circumstances in terms of nomological connections and discusses a series of interesting cases. He distinguishes knowledge from a human context from knowledge from a layman's perspective and from an expert's point of view. His relativized analysis of knowledge that *p* is this: "S knows (from the K point of view) that *p* iff: (a) it is true that *p*; (b) S believes that *p*; and (c) there is a non-defective epistemic pyramid (from the K point of view) for S and the proposition that *p*" ([64], p. 118).

My source for the role of the normality of the circumstances is twofold Leibnizian-Kantian. On the one hand, I have been impressed by the role the parameter of normal circumstances plays in Sellars's theory of perception. (See [56]-[58].) I learned my first Kant from Sellars. On the other hand, the reference to the normality of the circumstances as a whole is precisely the central point of the lawfulness of the phenomena, characteristic of empirical knowledge, as argued in Kant's transcendental deduction of the categories. This idea, I learned later on, is Leibnizian.

20. Of course, putting Harman's inferential principle Q as a defining condition in a definition of knowledge introduces a devastating circularity—unless some (incomplete) recursive schema is adopted.

21. *The Normal Third Cheater* affects, thus, the analysis of knowledge proposed by Lehrer, Skyrms, Dretske, Ginet, and others. The reader can examine this in the case of the de-

fining passages that have been quoted, e.g., Lehrer's. One more illustration: Ginet defines thus: "S knows that p if and only if: S is confident that p, this confidence is supported by a disinterested justification that S has for it, and *there is no truth r such that*, were S to be justified in believing that r and to retain all his properties that are compatible with his having justification for believing that r, then he would be very far from justified in being confident that p" ([28], p. 80; my italics). Clearly, this is a view of type (*b*) as characterized above in section 11. Evidently, *The Normal Cheater* refutes this; the requisite statement r present in the example is simply: "There is a mechanism M_1 operating as described." Ann knows in spite of such r.

22. Powers's [52] is one of the most insightful studies in basic epistemology. In a field filled with ingenious, penetrating, and even brilliant papers, Powers's [52] stands out for its deep insights. Villoro in [74] connects 'conocer', but not 'saber', with the power to answer questions.

23. We maintain the unity of the propositions about, and the unity of the content of, propositional attitudes. Instead of multiplying senses of 'believe', 'suppose', 'think', 'know', etc., we recognize that these verbs appear in constructions canonically put as follows: "X believes (knows, etc.) of --- that . . . ," where the dashed blanks are to be occupied by a list of expressions, or a description of a set of entities, and the dotted blank is occupied by a sentence or a clause. We consider as a limiting case, although it is the fundamental case, that in which the dashed blank is occupied by a null list or null description of a set of entities. In that case, the whole sentence occupying the dotted blank expresses a truth or a falsehood. In short, there are no senses of 'know', 'believe', etc., but different constructions, in which the *oratio obliqua* is sometimes a proposition and sometimes a propositional function. See [10] for additional discussion.

24. See [9], chap. 3 for a discussion of the hierarchy of logical forms of a proposition and for some of the basic laws governing such hierarchies. See Appendix.

25. See [11], part II, for a detailed discussion of propositional guises and their application to the analysis of perceptual judgments, sensory fields, and so on.

26. See [10] for some reasons for prefering the terminology 'internal'-'external' over the terminology '*de dicto*'-'*de re*' to refer to occurrences of expressions in *oratio obliqua* constructions, and for a discussion of propositional transparence and propositional opacity.

27. Powers has, it seems, a different view of sentences (1) and (2), if we can construe him as holding for 'believes' what he claims for 'knows', namely: "Again, we use '*x* knows that the answer to the question Q is P' *not* to mean that if *x* were asked 'Is the answer to Q P? he would answer yes, but rather to mean that if *x* were asked Q he would answer P. Thus, what we call 'knowing that the answer to Q is P' is more than merely knowing (correctly accepting) the proposition that the answer to Q is P" ([52], p. 348; his italics).

28. See [9], chap. 10 for an account of volition and chap. 12 for a discussion of several types of intentional action.

29. I owe this example to Miriam M. Castañeda. See also [12] for a double generality of language.

30. See e.g. Skyrms [60], Sosa [63], Paxson [51], Pappas [46], Lehrer [42]. One of the earliest causal accounts of knowledge is the one proposed by Goldman in [29]. Skyrms attacked it with an example that was regarded as conclusive by later writers. Loeb in [44] makes an important clarification to Skyrms's argument and shows that it is not as devastating as it was thought to be. *The Non-Normal Third Cheater Series* shows the insufficiency of Goldman's causal principle.

31. See [12] and [11] for accounts of the nature of our consciousness of particulars.

32. For a fascinating discussion of the different roles of principles of inference, especially their role in the structuring of concepts, see Sellars [55]. An insightful discussion of inference appears in Harman [31].

33. It may seem improper to formulate Quine's axiom in terms of propositions rather than sentences. But we are using the word 'proposition' to mean either a truth or a falsehood, and obviously propositions in this sense are not sentences. In fact, they are not even classes of equivalent sentences, under some equivalence or other, as is shown in [9], pp. 34f. Yet we still leave it open that the truths (and falsehoods) of the world may be reducible to sentences, or, better,

classes of sentences (in order to allow for intra- and inter-language equivalences), together with something else. If the student of Quine is not yet comfortable with the word 'proposition', he may accept the axiom above as a Quinelike axiom, rather than as one of Quine's very own.

34. For the characteristics and crucial referential roles of quasi-indicators, see [6], [10], and [11], section II.9. For an anti-Cartesian interpretation of the *cogito ergo sum*, see [11], section II.4.

35. As is explained in note 2 above, the main inspiration for the account of contexts of epistemic justification developed here is the account of contexts of justification for actions, intentions, and norms proposed in [9], chaps. 5, 6, and 8. I must note an error in the latter account, kindly pointed out by Michael Bratman. This error consists in not having maintained in an actional context the important parameter consisting of a set of *prime actions*, i.e., actions which are for the context in question like atoms of action—even though in other contexts they may be analyzed, into either pure compound actions, whose components are nothing but actions, or mixed actions, having some propositional components. Bratman proved the important theorem that a context that has as prime actions actions that are compounds of prime actions, may lead to contradictions. The counterpart epistemological error is avoided in this essay by having each epistemic context C_i of belief justification contain the crucial parameter P_i. In the epistemic case the rationale for this parameter is, fortunately, even more obvious: P_i is, as noted, the set of answers, whether true or false, to the questions in Q_i. Conversely, given the theory of questions sketched out here, the members of P_i are the propositions that yield hierarchies of questions containing the questions in Q_i. Thus the relevance of P_i to Q_i is obvious and deep.

36. See [13] and [14] for complementary discussions of the role of definitions in philosophy and for a critique of the classical method of philosophical analysis.

37. I am very grateful to Peter French for having put me in the position of having to finish, at last, this essay, even if it is only the first part of a series conceived and planned in 1970 on the model of theory of practical thinking. See note 3 above. The delay was, actually, beneficiary. It allowed me to learn from and to refer to the rich post-Gettier literature. And it also allowed me both to utilize Powers's data and to incorporate his approach. I am grateful to Victoria Haire of the University of Minnesota Press and to Howard Wittstein for having thoroughly overhauled the grammar and style of this paper.

References

[1] Alston, William, "Two Types of Foundationalism," *The Journal of Philosophy* 73 (1976): 165–85.

[2] Austin, John, "Other Minds," *Aristotelian Society Supplementary Volume* 20 (1946):148–87.

[3] Carnap, Rudolf, *Logical Foundations of Probability* (Berkeley, 1949).

[4] Castañeda, Hector-Neri, "Consciousness and Behavior: Their Basic Connections," in *Intentionality, Minds, and Perception*, ed. H.-N. Castañeda (Detroit, 1963).

[5] Castañeda, Hector-Neri, "Knowledge and Certainty," *The Review of Metaphysics* 18 (1965): 508–47.

[6] Castañeda, Hector-Neri, "Indicators and Quasi-Indicators," *American Philosophical Quarterly* 4 (1967):85–100.

[7] Castañeda, Hector-Neri, "On Knowing (or Believing) that One Knows (or Believes)," *Synthese* 21 (1970):187–203.

[8] Castañeda, Hector-Neri, "There Are Command Sh-Inferences," *Analysis* 32 (1971):13–19.

[9] Castañeda, Hector-Neri, *Thinking and Doing: The Philosophical Foundations of Institutions* (Dordrecht, 1975).

[10] Castañeda, Hector-Neri, "On the Philosophical Foundations of the Theory of Communication: Reference," *Midwest Studies in Philosophy* 2 (1977):165–86. Reprinted in *Contemporary Perspectives in the Philosophy of Language*, ed. P. French, T. Uehling, and H. Wettstein (Minneapolis, 1979).

[11] Castañeda, Hector-Neri, "Perception, Belief, and the Structure of Physical Objects and of Consciousness," *Synthese* 35 (1977):285–351.

[12] Castañeda, Hector-Neri, "The Causal and Epistemic Roles of Proper Names in Our Thinking of Particulars," in *Contemporary Perspectives in the Philosophy of Language*, ed. P. French, T. Uehling, and H. Wettstein (Minneapolis, 1979).

[13] Castañeda, Hector-Neri, "Intentionality and Identity in Human Action and Philosophical Method," *Nous* 13 (1979):235–60.

[14] Castañeda, Hector-Neri, "Philosophical Method and Chisholm's Direct Awareness of Self," *Grazer philosophische Studien* (forthcoming).

[15] Chisholm, Roderick M., "Epistemic Statements and the Ethics of Belief," *Philosophy and Phenomenological Research* 16 (1956):281–312.

[16] Chisholm, Roderick M., "Evidence as Justification," *The Journal of Philosophy* 58 (1961): 739–48.

[17] Chisholm, Roderick M., "Contrary-to-Duty Imperatives and Deontic Logic," *Analysis* 14 (1963):33–36.

[18] Chisholm, Roderick M., *Theory of Knowledge* (Englewood Cliffs, N.J., 1966).

[19] Chisholm, Roderick M., *Person and Object: A Metaphysical Study* (Chicago, 1976).

[20] Chisholm, Roderick M., "Castañeda's Thinking and Doing," *Nous* 13 (1979):385-96.

[21] Clark, Michael, "Knowledge and Grounds: A Comment on Mr. Gettier's Paper," *Analysis* 24 (1963):46–48.

[22] Coder, David, "Thalberg's Defense of Justified True Belief," *The Journal of Philosophy* 67 (1970):424–25.

[23] Coder, David, "Naturalizing the Gettier Argument," *Philosophical Studies* 26 (1974):111–18.

[24] Cohen, L. Jonathan, *The Implications of Induction* (London, 1970).

[25] Dretske, Fred, "Conclusive Reasons," *Australasian Journal of Philosophy* 49 (1971):1–22.

[26] Fichte, Johann Gottlieb, *Science of Knowledge*, trans. P. Heath and J. Lachs (New York, 1970).

[27] Gettier, Edmund L., "Is Justified True Belief Knowledge?" *Analysis* 23 (1963):121–23.

[28] Ginet, Carl, *Knowledge, Perception, and Memory* (Dordrecht, 1975).

[29] Goldman, Alvin, "A Causal Theory of Knowing," *The Journal of Philosophy* 64 (1967): 355–72.

[30] Goldman, Alvin., "Discrimination and Perceptual Knowledge," *The Journal of Philosophy* 73 (1976):771–91.

[31] Harman, Gilbert, *Thought* (Princeton, N.J., 1973).

[32] Heidegger, Martin, *El ser y el tiempo*, trans. José Gaos (Mexico, 1951).

[33] Heidelberg, Herbert, "Chisholm's Epistemic Principles," *Nous* 13 (1979):73–82.

[34] Hilpinen, Risto, "Knowing that One Knows and the Classical Definition of Knowledge," *Synthese* 21 (1970):109–32.

[35] Hilpinen, Risto, "Knowledge and Justification," *Ajatus* 33 (1971):7–39.

[36] Hintikka, Jaakko, *Knowledge and Belief* (Ithaca, N.Y., 1962).

[37] Johnsen, Bredo, "Knowledge," *Philosophical Studies* 25 (1974):273–82.

[38] Kant, Immanuel, *Critique of Pure Reason*.

[39] Kemeny, John G., "Fair Bets and Inductive Probabilities," *The Journal of Symbolic Logic* 20 (1955):263–73.

[40] Klein, Peter D., "A Proposed Definition of Propositional Knowledge," *The Journal of Philosophy* 68 (1971):471–82.

[41] Klein, Peter D., "Knowledge, Causality, and Defeasibility," *The Journal of Philosophy* 73 (1976):792–812.

[42] Lehrer, Keith, *Knowledge* (Oxford, 1974). This book supersedes a dozen or so previous essays by Lehrer.

[43] Leibniz, Gottfried Wilhelm, "On the Method of Distinguishing Real from Imaginary Phenomena," in *Gottfried Wilhelm Leibniz: Philosophical Papers and Letters*, ed. L. L. Loemker (Dordrecht, 1969).

[44] Loeb, Louis, "On a Heady Attempt to Befiend Causal Theories of Knowledge," *Philosophical Studies* 29 (1976):331-36.

[45] Lucey, Kenneth, "Scales of Epistemic Appraisal," *Philosophical Studies* 25 (1974):423-28.

[46] Pappas, George, "Knowledge and Reasons," *Philosophical Studies* 25 (1974):423-28.

[47] Pappas, George, and Swain, Marshall, "Some Conclusive Reasons against 'Conclusive Reasons'," *The Australasian Journal of Philosophy* 51 (1973):72-76.

[48] Pastin, Mark, "Foundationalism Reux," *The Journal of Philosophy* 71 (1974):709-10.

[49] Pastin, Mark, "Modest Foundationalism and Self-Warrant," in *Studies in Epistemology*, ed. N. Rescher (Oxford, 1975).

[50] Pastin, Mark, Review of Lehrer's *Knowledge, Nous* 11 (1977):431-37.

[51] Paxson, Thomas D., "Prof. Swain's Account of Knowledge," *Philosophical Studies* 25 (1974):57-61.

[52] Powers, Lawrence, "Knowledge by Deduction," *The Philosophical Review* 87 (1978):337-71.

[53] Rescher, Nicholas, "Foundationalism, Coherentism, and the Idea of Cognitive Systematization," *The Journal of Philosophy* 71 (1974):695-708.

[54] Sartre, Jean-Paul, *Being and Nothingness*, trans. Hazel E. Barnes (New York, 1956).

[55] Sellars, Wilfrid, "Meaning and Inference," *Mind* 62 (1953):313-38.

[56] Sellars, Wilfrid, "Empiricism and the Philosophy of Mind," vol. I in *Minnesota Studies in the Philosophy of Science*, ed. H. Feigl and M. Scriven (Minneapolis, 1956).

[57] Sellars, Wilfrid, "Phenomenalism," in *Intentionality, Minds, and Perception*, ed. H.-N. Castañeda (Detroit, 1967).

[58] Sellars, Wilfrid, "The Structure of Knowledge," in *Action, Knowledge and Reality: Studies in Honor of Wilfrid Sellars*, ed. H.-N. Castañeda (Indianapolis, 1975).

[59] Shimony, Abner, "Coherence and the Axioms of Confirmation," *The Journal of Symbolic Logic* 20 (1955):1-28.

[60] Skyrms, Brian, "The Explication of 'X knows that p'," *The Journal of Philosophy* 64 (1967):333-89.

[61] Sosa, Ernest, "The Analysis of 'Knowledge that p'," *Analysis* 25 (1964):1-8.

[62] Sosa, Ernest, "Propositional Knowledge," *Philosophical Studies* 20 (1969):33-43.

[63] Sosa, Ernest, "Two Conceptions of Knowledge," *The Journal of Philosophy* 67 (1970):59-66.

[64] Sosa, Ernest, "The Concept of Knowledge. How Do You Know?" *American Philosophical Quarterly* 11 (1974):113-22.

[65] Sosa, Ernest, "On Our Knowledge of Matters of Fact," *Mind* 83 (1974):388-405.

[66] Sosa, Ernest, Review of Harman's *Knowledge, Nous* 11 (1977):421-30.

[67] Sosa, Ernest, "The Foundations of Foundationalism," *Nous* (forthcoming).

[68] Swain, Marshall, "Knowledge, Causality, and Justification," *The Journal of Philosophy* 69 (1972):291-300.

[69] Swain, Marshall, "Epistemic Defeasibility," *American Philosophical Quarterly* 11 (1974):15-25.

[70] Thalberg, Irving, "In Defense of Justified True Belief," *The Journal of Philosophy* 66 (1969):795-803.

[71] Tienson, John, "On Analyzing Knowledge," *Philosophical Studies* 25 (1974):289-93.

[72] Tomberlin, James, "Knowing without Knowing that One Knows," *Philosophia* 2 (1972):239-46.

[73] Tomberlin, James, Review of Carl Ginet's *Knowledge, Perception, and Memory, Nous* (forthcoming).

[74] Villoro, Luis, "Conocer y saber," *Critica* 10 (1970):75-91.

[75] Wittgenstein, Ludwig, *Philosophical Investigations* (Oxford, 1952).

APPENDIX

1. After the preceding essay was copy edited, Nuel Belnap showed me that the hierarchy of questions described in section I.16 does *not* include all the questions to which a given proposition is a true answer. For instance, the following question escapes that hierarchy: Can you formulate a proposition that is a member of the set S? In general, the questions that are left out seem to me to be questions that arise from propositions containing P as a component. If this is so, then the hierarchy described in section I.16 is only the *proper* or *characteristic hierarchy* of questions determined by P. I wonder now whether what are called constraints on questions in the essay can, at least in some cases, be considered parts of propositions that have as a component a given proposition in which we are interested.

There is much more to the theory of questions than we have touched upon in our discussion. Given our concern with knowledge and the state of mind constitutive of knowledge, it was crucial, however, that we did not consider a question as a set of propositions, but as a special thought content, which stands in the relations YES_i to its answers. I am not sure that the logic of questions that treats questions as sets of propositions offers us a fully serviceable object of interrogative attitudes. But this is a large issue that we cannot decide here.

In any case, the tentative analysis of part II remains the same, I believe, if the parameter Q_i of questions is understood to include more questions than our discussion in part I concentrated on. See the rich bibliographies in [1] and [2] below.

2. On October 30, 1979 George Nakhnikian gave me the following quotation:

If by 'omniscience' we mean the ability to answer with certainty every conceivable question, including questions concerning the future . . . ([3], p. 121).

Thus, instead of speaking of the Plato-Powers principle as I did above, I should have spoken of the Plato-Popper-Powers principle.

Bibliography for the Appendix

[1] Belnap, Nuel, and Steel, Thomas B., Jr., *The Logic of Questions and Answers* (New Haven and London, Yale University Press, 1976).

[2] Hintikka, Jaakko, *The Semantics of Questions and the Questions of Semantics* (Amsterdam: North Holland Publishing Co., 1976).

[3] Popper, Karl, "Indeterminism in Quantum Physics and in Classical Physics," *The British Journal for the Philosophy of Science* 1 (1950-51).