

Chapter 26

Epistemology Without Knowledge and Without Belief

Jaakko Hintikka

Knowledge and Decision-Making

Epistemology seems to enjoy an unexpectedly glamorous reputation in these days. A few years ago, William Safire (1995) wrote a popular novel called *The Sleeper Spy*. It depicts a distinctly post-Cold War world in which it is no longer easy to tell the good guys—including the good spies—from the bad ones. To emphasize this sea change, Safire tells us that his Russian protagonist has not been trained in the military or in the police, as he would have been in the old days, but as an epistemologist.

But is this with-it image deserved? Would the theory of knowledge that contemporary academic epistemologists cultivate be of any help to a sleeper spy? This question prompts a critical survey of the state of the art or, rather, the state of the theory of knowledge. I submit that the up-to-date image is not accurate and that most of the current epistemological literature deals with unproductive and antiquated questions. This failure is reflected in the concepts that are employed by contemporary epistemologists.

What are those concepts? It is usually thought and said that the most central concepts of epistemology are knowledge and belief. The prominence of these two notions is reflected in the existing literature on epistemology. A large chunk of it consists in discussions of how the concept of knowledge is to be defined or is not to be defined. Are those discussions on the target? An adequate analysis of such concepts as knowledge and belief, whether it is calculated to lead us to a formal

Jaakko Hintikka was deceased at the time of publication.

J. Hintikka (deceased)
Boston University, Helsinki, Finland

definition or not, should start from the role that they play in real life. Now in real life we are both producers and consumers of knowledge. We acquire knowledge in whatever ways we do so, and we then put it to use in our actions and decision-making. I will here start from the latter role, which takes us to the question: What is the role that the notion of knowledge plays in that decision-making?

To take a simple example, let us suppose that I am getting ready to face a new day in the morning. How, then, does it affect my actions if I know that it will not rain today? You will not be surprised if I say that what it means is that I am entitled to behave as if it will not rain—for instance to leave my umbrella home. However, you may be surprised if I claim that most of the important features of the logical behavior of the notion of knowledge can be teased out of such simple examples. Yet this is the case. My modest example can be generalized. The role of knowledge in decision-making is to rule out certain possibilities. In order to use my knowledge, I must know which possibilities it rules out. In other words, any one scenario must therefore be either incompatible or compatible with what I know, for I am either entitled or not entitled to disregard it. Thus the totality of incompatible scenarios determines what I know and what I do not know, and vice versa. In principle, all that there is to logic of knowledge is this dichotomy between epistemically impossible and epistemically possible scenarios.

It is also clear how this dichotomy serves the purposes of decision-making, just as it does in my mini-example of deciding whether or not to take an umbrella with me. But the connection with overt behavior is indirect, for what the dichotomy merely demarcates are the limits of what I am *entitled to* disregard. And being entitled to do something does not always mean that I do it. It does not always show up in the overt ways one actually or even potentially acts. For other considerations may very well enter into my decision-making. Maybe I just want to sport an umbrella even though I know that it need not serve its function of shielding myself from rain. Maybe I am an epistemological *akrates* and act against what I know. The connection is nevertheless real, even though it is a subtle one. There is a link between my knowledge and my decisions, but it is, so to speak, a *de jure* connection and not a *de facto* connection. I think that this is a part of what John Austin (1961a) was getting at when he compared “I know” with “I promise.” To know something does not mean simply to have evidence of a superior degree for it, nor does it mean to have a superior kind of confidence in it. If my first names were George Edward, I might use the open-question argument to defend these distinctions. By saying “I promise,” I entitle you to expect that I fulfill my promise. By saying “I know,” I claim that I am entitled to disregard those possibilities that do not agree with what I know. There is an evaluative element involved in the concept of knowledge that does not reduce to the observable facts of the case. Hence, it is already seen to be unlikely that you could define what it means to know by reference to matters of fact, such as the evidence that the putative knower possesses or the state of the knower’s mind.

This evaluative element is due to the role of knowledge in guiding our life in that it plays a role in the justification of our decisions. This role determines in the last analysis the logic and in some sense the meaning of knowledge. A Wittgensteinian might put this point by saying that decision-making is one of the language-games

that constitute the logical home of the concept of knowledge. You can remove knowledge from the contexts of decision-making, but you cannot remove a relation to decision-making from the concept of knowledge. For this reason, it is among other things misguided in a fundamental way to try to separate epistemic possibility from actual (natural) possibility. Of course, the two are different notions, but the notion of epistemic possibility has conceptual links to the kind of possibility that we have to heed in our decision-making. For one thing, the set of scenarios involved in the two notions must be the same.

But the main point here is not that there is an evaluative component to the notion of knowledge. The basic insight is that there is a link between the concept of knowledge and human action. The evaluative element is merely a complicating factor in the equation. The existence of a link between the two is not peculiar to the notion of knowledge. There is a link, albeit of a different kind, also in the case of belief. In fact, the conceptual connection is even more obvious in the case of belief. Behavioral scientists have studied extensively decision principles where belief constitutes one component, as, for instance, in the principle of maximizing expected utility. It usually comes in the form of degrees of belief. (They are often identified with probabilities.) Typically, utilities constitute another component. Whether or not such explicit decision principles capture the precise links between belief and behavior, they illustrate the existence of the link and yield clues to its nature.

Indeed, from a systematic point of view, the relative roles assigned to knowledge and to belief in recent epistemology and recent decision theory cannot but appear paradoxical. Belief is in such studies generally thought of as a direct determinant of our decisions, whereas knowledge is related to action only indirectly, if at all. Yet common sense tells us that one of the main reasons for looking for more knowledge is to put us in a better position in our decision-making, whereas philosophers often consider belief—especially when it is contrasted with knowledge—as being initially undetermined by our factual information and therefore being a much worse guide to decision-making. Probability is sometimes said to be a guide to life, but surely knowledge is a better one. Or, if we cannot use black-or-white concepts here, shouldn't rational decision-making be guided by degrees of knowledge rather than degrees of mere belief?

The same point can perhaps be made by noting that in many studies of decision-making, a rational agent is supposed to base his or her decisions on the agent's beliefs (plus, of course, utilities) and then by asking: Would it not be even more rational for the agent to base his or her decisions on what the agent *knows*?

In order for a rational agent to act on his or her belief, this belief clearly must be backed up by some evidence. Otherwise, current decision theory makes little sense. The difference is that the criteria of what entities are to act are different in the case of belief from what they are in the case of knowledge. If I act on a belief, that belief must satisfy my personal requirements for that role. They may vary from person to person. In contrast, the criteria of knowing are impersonal and not dependent on the agent in question. In order to define knowledge as distinguished from beliefs, we would have to spell out those impersonal criteria. This is obviously an extremely difficult task at best.

Another fact that complicates the connection between knowledge and behavior—that is, between what I know and what I do—is that in principle, this link is holistic. What matters to my decisions in the last analysis is the connection between the totality of my knowledge. There is not always any hard-and-fast connection between particular items of knowledge and my behavior. In principle, the connection is via my entire store of knowledge. This is reflected by the fact emphasized earlier that the dichotomy that determines the logic of knowledge is a distinction between scenarios that are ruled out by the *totality* of what I know and scenarios that are compatible with the *totality* of my knowledge and that I therefore must be prepared for. The same feature of the concept of knowledge also shows up in the requirement of total evidence that is needed in Bayesian inference and which has prompted discussion and criticism there. (See, e.g., Earman 1992.)

To spell out the criteria of the justification involved in the applications of the concept of knowledge is to define what knowledge is as distinguished from other propositional attitudes. Characterizing these conditions is obviously a complicated task. I will return to these criteria later in this chapter.

The Logic of Knowledge and Information

Meanwhile, another dimension of the concept of knowledge is brought out by homely examples of the kind I am indulging in. By this time it should be clear—I hope—that it is extremely hard to specify the kind of entitlement or justification that knowing something amounts to. This difficulty is perhaps sufficiently attested to by the inconclusiveness of the extensive discussions about how to define knowledge that one can find in the literature. (See, e.g., Shope 1983.) But another aspect of this notion is in principle as clear as anything one can hope to find in philosophical analysis (or synthesis). It may be difficult to tell whether a certain propositional attitude amounts to knowledge, belief, opinion or whatnot, but there is typically no difficulty in spelling out the *content* of any one of these attitudes on some particular occasion. Here, the lesson drawn from my rain-and-umbrella example is applicable. It was seen that what someone knows specifies, and is specified by, the class of possible scenarios that are compatible with what he or she knows. And such classes of scenarios or of “possible worlds” can be captured linguistically as the classes of scenarios (alias possible worlds) in which a certain sentence is true. Indeed, for Montague (1974, p. 153) such classes of possible worlds (or, strictly speaking, the characteristic functions of these classes, in the sense of functions from possible worlds to truth-values) *are* propositions. In this way, the content of a propositional attitude can normally be captured verbally. For another instance, for Husserl (1983, sec. 124), the task would be to capture the noematic *Sinn* of an act, which he says can in principle always be accomplished linguistically—that is, in Husserl’s terminology, through *Bedeutungen*.

Let us now call the members of the class of scenarios admitted by someone’s knowledge that someone’s epistemic alternatives. That I know that it will not rain

today means that none of the scenarios under which the wet stuff falls down are among my epistemic alternatives, and likewise for all *knowing that* statements. What the concept of knowledge involves in a purely logical perspective is thus a dichotomy of the space of all possible scenarios into those that are compatible with what I know and those that are incompatible with my knowledge. What was just seen is that this dichotomy is directly conditioned by the role of the notion of knowledge in real life. Now this very dichotomy is virtually all we need in developing an explicit logic of knowledge, better known as epistemic logic. This conceptual parentage is reflected by the usual notation of epistemic logic. In it, the epistemic operator K_a (“a knows that”) receives its meaning from the dichotomy between excluded and admitted scenarios, while the sentence within its scope specifies the content of the item of knowledge in question.

Basing epistemic logic on such a dichotomy has been the guiding idea of my work in epistemic logic right from the beginning. I have seen this idea being credited to David Lewis, but I have not seen any uses of it that predate my work.

But here we seem to run into a serious problem in interpreting epistemic logic from the vantage point of a dichotomy of excluded and admitted scenarios. Such an interpretation might seem to exclude “quantifying in”—that is to say, to exclude applications of the knowledge operator to open formulas for them, it would not make any sense to speak of scenarios in which the content of one’s knowledge is true or false. Such “quantifying in” is apparently indispensable for the purpose of analyzing the all-important *wh*-constructions with *knows*. For instance, “John *knows* who murdered Roger Ackroyd” apparently must be expressed by

$$(\exists x) K_{\text{John}} (x \text{ murdered Roger Ackroyd}) \quad (26.1)$$

as distinguished from

$$K_{\text{John}} (\exists x) (x \text{ murdered Roger Ackroyd}) \quad (26.2)$$

which says that John knows that someone murdered the victim and hence can serve as the presupposition of the question, “Who murdered Roger Ackroyd?”

But in (26.1), the notion of knowledge apparently cannot be interpreted by reference to a distinction between admitted and excluded scenarios. The reason is that the knowledge operator in (26.1) is prefixed to an open formula. Such an open formula cannot be said to be true or false in a given scenario, for its truth depends on the value of the variable x . Hence it cannot implement the required dichotomy.

In order for our epistemic discourse to express the *wh*-constructions, the knowledge operator must apparently be allowed to occur also internally, prefixed to open formulas rather than sentences (formulas without free variables). This prompts a serious interpretational problem. Indeed we can see here the reason for the deep theoretical interest of the problem of “quantifying in,” which otherwise might strike one as being merely the logicians’ technical problem. Fortunately, this apparent problem can be solved by means of suitable analysis of the relations between different logical operators (see section “[Information acquisition as a questioning procedure](#)”).

An epistemic logic of this kind can obviously be developed within the framework of possible worlds semantics. (For a sketch of how this can be done, see Hintikka 2003b.) In fact, the truth condition for *knows that* is little more than a translation of what was just said: “b knows that S” is true in a world W if and only if S is true in all the epistemic b-alternatives to W. These alternatives are all the scenarios or “worlds” compatible with everything b knows in W. In certain important ways, this truth condition for knowledge statements is clearer than its counterpart in the ordinary (alethic) modal semantics, in that in epistemic logic the interpretation of the alternativeness relation (alias accessibility relation) is much clearer than in the logic of physical or metaphysical modalities.

Here we have already reached a major conclusion. Epistemic logic presupposes essentially only the dichotomy between epistemically possible and epistemically excluded scenarios. How this dichotomy is drawn is a question pertaining to the definition of knowledge. However, we do not need to know this definition in doing epistemic logic. Thus the logic and the semantics of knowledge can be understood independently of any explicit definition of knowledge. Hence it should not be surprising to see that a similar semantics and a similar logic can be developed for other epistemic notions—for instance, belief, information, memory, and even perception. This is an instance of a general law holding for propositional attitudes. This law says that the content of a propositional attitude can be specified independently of differences between different attitudes. This law has been widely recognized, even if it has not always been formulated as a separate assumption. For instance, in Husserl (1983, e.g., sec. 133) it takes the form of separating the noematic *Sinn* from the thetic component of a noema. As a consequence, the respective logics of different epistemic notions do not differ much from each other. In particular, they do not differ at all in those aspects of their logic that depend merely on the dichotomical character of their semantics. These aspects include prominently the laws that hold for quantifiers and identity, especially the modifications that are needed in epistemic contexts in the laws of the substitutivity of identity and existential generalization.

The fact that different epistemic notions, such as knowledge, belief, and information, share the same dichotomic logic should not be surprising in the light of what has been said. The reason is that they can all serve the same purpose of guiding our decisions, albeit in different ways. Hence the same line of thought can be applied to them as was applied earlier to the concept of knowledge, ending up with the conclusion that their logic is a dichotomic logic not unlike the logic that governs the notion of knowledge. The common ingredient in all these different logics is then the true epistemic logic. But it turns out to be a logic of information rather than a logic of knowledge.

This distinction between what pertains to the mere dichotomy between admitted and excluded scenarios and what pertains to the criteria relied on in this dichotomy is not a novelty. It is at bottom only a restatement in structural terms of familiar contrast, which in the hands of different thinkers has received apparently different formulations. The dichotomy defines the content of a propositional attitude, while the criteria of drawing it determine which propositional attitude we are dealing with.

Hence we are naturally led to the project of developing a generic logic of contents of attitudes, independent of the differences between different attitudes.

This generic logic of epistemology can be thought of as the logic of information. Indeed, what the content of a propositional attitude amounts to can be thought of as a certain item of information. In attributing different attitudes to agents, different things are said about this information—for instance, that it is known, believed, remembered, and so on. This fits in well with the fact that the same content can be known by one person, believed by another, remembered by a third one, and so on. This idea that one and the same objective content may be the target of different people's different attitudes is part of what Frege (see, e.g., 1984) was highlighting by his notion of *the thought*. Thus it might even be happier to talk about the logic of information than about epistemic logic. John Austin (1961b) once excused his use of the term “performative” by saying that even though it is a foreign word and an ugly word that perhaps does not mean very much, it has one good thing about it: It is not a deep word. It seems to me that epistemology would be in much better shape if instead of the deep word “knowledge,” philosophers cultivated more the ugly foreign word “information,” even though it perhaps does not capture philosophers' profound sense of knowing. In any case, in the generic logic of epistemology here envisaged, philosophers' strong sense of knowledge plays no role.

Information Acquisition as a Questioning Procedure

But what about the other context in which we encounter knowledge in real life—the context of knowledge acquisition? As was noted, what the concept of knowledge amounts to is revealed by two questions: What is it that we are searching for in the process of knowledge acquisition? What purpose can the product of such an inquiry serve? The second question has now been discussed. It remains to examine the crucial first question. Surely the first order of business of any genuine theory of knowledge—the most important task both theoretically and practically—is how new acquired, not merely how previously obtained information can be evaluated. A theory of information (knowledge) acquisition is both philosophically and humanly much more important than a theory of whether or not already achieved information amounts to knowledge. Discovery is more important than the defense of what you already know. In epistemology, as in warfare, offense frequently is the best defense.

This point can be illustrated in a variety of ways. For instance, a thinker who does not acquire any information cannot even be a skeptic, for he or she would not have anything to be skeptical about. And a skeptic's doubts must be grounded on some grasp as to how that information is obtained, unless these doubts are totally irrational. Epistemology cannot start from the experience of wonder or doubt. It should start from recognition of where the item of information that we are wondering about or doubting came from in the first place. Any rational justification or rational distinction of such wonder or doubt must be based on its ancestry.

Fortunately we now have available to us a framework in which to discuss the logic and epistemology of knowledge acquisition or, rather, if I have the terminological courage of my epistemological convictions, information acquisition. The framework is what is referred to as the interrogative model of inquiry or interrogative approach to inquiry. (See Hintikka 1999.) Its basic idea is the same as that of the oldest explicit form of reasoning in philosophy, the Socratic method of questioning or *elenchus*. In it, all new information enters into an argument or a line of reasoning in the form of answers to questions that the inquirer addresses to a suitable source of information.

It might at first seem implausible that this approach might yield a viable theory of ampliative reasoning in general, for several different reasons. Fortunately all these objections can be overcome. First, it might not seem likely that this model can be developed into a form explicit and detailed enough to allow for precise conclusions. This objection would have been eminently appropriate as recently as a decade or two ago. For it is only in the last several years that there has existed a general and explicit logical theory of all the relevant kinds of questions. This logic of questions and answers is the backbone of the interrogative model. This theory has not yet been presented in a monographic or textbook form, but its basic ideas are explained in recent and forthcoming papers of mine. (See, e.g., Hintikka 2003a.) This logic of questions and answers is an extension and application of epistemic logic (logic of knowledge). It has been made possible by a quiet revolution in epistemic logic. One of the main problems in representing questions is to specify which ingredients of the aimed-at information are the questioned elements—that is to say, are supposed to be made known by the answer. It turns out that their specification can sometimes be accomplished only by means of the independence indicators whose logic is only now being explored, even though it cannot be done in the earlier “first-generation” epistemic logic. The details of the new “second-generation” epistemic logic that makes use of the notion of independence need not concern us here. It may nevertheless be noted that this new logic solves the problem of “quantifying in” in that in it, the epistemic operator *K* always occurs sentence-initially. There is no problem of quantifying in, one might say here, only quantifying (binding variables) independently of an epistemic operator.

Another main requirement that can be addressed to the interrogative approach—and indeed to the theory of any goal-directed activity—is that it must do justice to the strategic aspects of inquiry. Among other things, it ought to be possible to distinguish the definitory rules of the activity in question from its strategic rules. The former spell out what is possible at each stage of the process. The latter express what actions are better and worse for the purpose of reaching the goals of the activity. This requirement can be handled most naturally by doing what Plato already did to the Socratic *elenchus* and by construing knowledge-seeking by questioning as a game that pits the questioner against the answerer. Then the study of the strategies of knowledge acquisition becomes another application of the mathematical theory of games, which perhaps ought to be called “strategy theory” rather than “game theory” in the first place. The distinction between the definitory rules—usually called simply the rules of the game—and strategic principles is built right into the structure of such games.

The greatest obstacle to generality might seem to be the apparently restricted range of applicability of the interrogative model. Some of the resistance to this approach, which I have referred to as the idea of “inquiry as inquiry,” can be dispelled by pointing out that questions and answers can be understood in a wide sense, and have to be so understood if the generality claim is to be acceptable. Sources of answers to explicit or implicit questions have to include not only human witnesses and other informants or databases in a computer, but observation and experimentation as well as memory and tacit knowledge. One of the leading ideas of the interrogative approach is that all information used in an argument must be brought in as an answer to a question. In claiming such generality for the interrogative model, I can appeal to such precedents as Collingwood’s (1940) and Gadamer’s (1975) “logic of questions and answers,” even though what they called logic really was not. My claims of generality on behalf of the interrogative approach are not even as sweeping as Collingwood’s thesis that every proposition may be considered as an answer to a question. Likewise in construing experiments as questions to nature, I can cite Kant (Kant 1787) and Bacon.

Interrogation and Justification

But the context of knowledge acquisition is vital even if the aim of your game is justification and not discovery. Suppose that a scientist has a reason to think that one of his or her conclusions is not beyond doubt. What is he or she to do? Will the scientist try to mine his or her data so as to extract from them grounds for a decision? Sometimes, perhaps, but in an overwhelming majority of actual scientific situations, the scientist will ask what further information one should in such circumstances try to obtain in order to confirm or disconfirm the suspect proposition—for instance, what experiments it would be advisable to perform or what kinds of observation one should try to make in order to throw light on the subject matter. Unfortunately such contexts—or should I say, such language-games—of verification by means of new information have not received much attention from recent philosophers. They have been preoccupied with the justification of already acquired knowledge rather than with the strategies of reaching new knowledge.

Thus we must extend the scope of the interrogative model in such a way that it enables us to cope with justification and not just pure discovery. What we need is a rule or rules that authorize the rejection—which is tentative and may be only temporary—of some of the answers that an inquirer receives. The *terminus technicus* for such rejection is *bracketing*. The possibility of bracketing widens the scope of epistemological and logical methods tremendously. After this generalization has been carried out, the logic of interrogative inquiry can serve many of the same purposes as the different variants of non-monotonic reasoning, and serve them without the tacit assumptions that often make nonmonotonic reasoning epistemologically restricted or even philosophically dubious. A telling example is offered by what is known as circumscriptive reasoning. (See McCarthy 1990.)

It relies on the assumption that the premises present the reasoner with all the relevant information, so that the reasoner can assume that they are made true in the intended models in the simplest possible way. This is an assumption that in fact can often be made, but it is not always available on other occasions. As every puzzle fan knows, often a key to the clever reasoning needed to solve a puzzle lies precisely in being able to imagine circumstances in which the normal expectations evoked by the specification of the puzzle are not realized. Suppose a puzzle goes as follows: “Evelyn survived George by more than 80 years, even though she was born many decades before him. How come?” The explanation is easy if you disregard the presumption that “George” is a man’s name and “Evelyn” a woman’s. Evelyn Waugh in fact survived George Eliot by 86 years. Here the solution of the puzzle depends entirely on going beyond the *prima facie* information provided by the putative—in other words, on violating the presuppositions of a circumscriptive inference. Reasoning by circumscription is enthymemic reasoning. It involves tacit premises that may be false.

Thus by introducing the idea of bracketing, we can dispense with all modes of ampliative reasoning. The only rules besides rules of logical inference are the rules for questioning and the rule allowing bracketing. This may at first look like a cheap trick serving merely to sweep all the difficulties of epistemic justification under the rug of bracketing. In reality, what is involved is an important insight. What is involved is not a denial of the difficulties of justification, but an insight into their nature as problems. Once a distinction is made between strategic and definitory rules, it is realized that the definitory rules can only be permissive, telling what one may do in order to reach knowledge and to justify it. The problem of justification is a strategic problem. It pertains to what one ought to do in order to make sure that the results of one’s inquiry are secure. This is to be done by the double process of disregarding dubious results and confirming the survivors through further inquiry. The only new permissive rule needed for the purpose is the rule that allows bracketing.

Thus the question as to which answers to bracket is always at bottom a strategic problem. It is therefore futile in principle to try to capture the justificatory process by means of definitory rules of this or that kind. To attempt to do so is a fallacy that in the last analysis vitiates all the usual “logics” of ampliative reasoning. This mistake is committed not only by non-monotonic logics but also by inductive logic and by the current theories of belief revision. Ampliative logics can be of considerable practical interest and value, but in the ultimate epistemological perspective, they are but types of enthymemic reasoning, relying on tacit premises quite as much as circumscriptive reasoning. An epistemologist’s primary task here is not to study the technicalities of such modes of reasoning, fascinating though they are in their own right. It is to uncover the tacit premises on which such enthymemic reasoning is in reality predicated.

Allowing bracketing is among other things important because it makes it possible to conceive of interrogative inquiry as a model also of the confirmation of hypotheses and other propositions in the teeth of evidence. The interrogative model can thus also serve as a general model of the justification of hypotheses. It should

in fact be obvious that the processes of discovery and justification cannot be sharply separated from each other in the practice or in the theory of science. Normally, a new discovery in science is justified by the very same process—for instance, by the same experiments—by means of which it was made, or could have been made. And this double duty service of questioning is not due only to the practical exigencies of “normal science.” It has a firm conceptual basis. This basis is the fact that information (unlike many Federal appropriations) does not come to an inquirer earmarked for a special purpose—for instance, for the purpose of discovery rather than justification. The inquirer may ask a question for this or that proximate purpose in mind, but there is nothing in the answer that rules out its being used for other purposes as well.

And such an answer can only be evaluated in terms of its service for both causes. This is because from game theory we know that in the last analysis, game-like goal-directed processes can be evaluated only in terms of their strategies, not in terms of what one can say of particular moves—for instance, what kinds of “warrants” they might have. As a sports-minded logician might explain the point, evaluating a player’s skills in a strategic game is in principle like judging a figure-skating performance rather than keeping score in a football game. In less playful terms, one can in general associate utilities (payoffs) only with strategies, not with particular moves. But since discovery and justification are aspects of the same process, they have to be evaluated in terms of the different possible strategies that are calculated to serve both purposes.

When we realize this strategic inseparability of the two processes, we can in fact gain a better understanding of certain otherwise puzzling features of epistemic enterprise. For instance, we can now see why it sometimes is appropriate to jump to a conclusion on the basis of relatively thin evidence. The reason is that finding what the truth is can help us mightily in our next order of business of finding evidence for that very truth. Sherlock Holmes has abductively “inferred” that the stablemaster has stolen the famous racing horse “Silver Blaze” (see the Conan Doyle story with this title) in order to lame it partially. He still has to confirm this conclusion, however, and in that process he is guided by the very content of that abductive conclusion—for instance, in directing his attention to the possibility that the stablemaster had practiced his laming operation on the innocent sheep grazing nearby. He puts a question to the shepherd as to whether anything had been amiss with them of late. “Well, sir, not of much account, but three of them have gone lame, sir.” Without having already hit on the truth, Holmes could not have thought of asking this particular question.

If you disregard the strategic angle, the frequent practice of such “jumps to a conclusion” by scientists may easily lead one to believe that scientific discovery is not subject to epistemological rules. The result will then be the hypothetico-deductive model of scientific reasoning, which is hence seen to rest on a fallacious dismissal of the strategic angle.

Thus we reach a result that is neatly contrary to what were once prevalent views. It used to be held that discovery cannot be subject to explicit epistemological theory, whereas justification can. We have found out that not only can discovery be

approached epistemologically, but that justification cannot in the long run be done justice to by a theory that does not also cover discovery.

A critical reader might initially have been wondering why contexts of verification and of other forms of justification do not constitute a third logical home of the notion of knowledge, besides the contexts of decision-making and information-acquisition. The answer is that processes of justification can only be considered as aspects of processes of information-acquisition.

The Generality of the Interrogative Model

The most general argument for the generality of the interrogative approach relies only on the assumption that the inquirer's line of thought can be rationally evaluated. What is needed for such an evaluation? If no new information is introduced into an argument by a certain step, then the outcome of that step is a logical consequence of earlier statements reached in the argument. Hence we are dealing with a logical inference step that has to be evaluated by the criteria of logical validity. It follows that interrogative steps are the ones in which new information enters into the argument. In order to evaluate the step, we must know what the source of this information is, for the reliability of the information may depend on its source. We must also know what else might have resulted from the inquirer's approaching this particular source in this particular way and with what probabilities. If so, what the inquirer did can be thought of as a question addressed to that source of information. Likewise, we must know what other sources of information the inquirer could have consulted and what the different results might have been. This amounts to knowing what other sources of answers the inquirer might have consulted. But if all of this is known, we might as well consider what the inquirer did as a step in interrogative inquiry.

In an earlier work (Hintikka 1998), I have likened such tacit interrogative steps to Peircean abductions, which Peirce insists are inferences even though they have interrogative and conjectural aspects.

The interrogative model can be thought of as having also another kind of generality—namely, generality with respect to the different kinds of questions. Earlier epistemic logic was incapable of handling questions more complicated than simple *wh*-questions. In particular, it could not specify the logical form of questions in which the questioned ingredient was apparently within the scope of a universal quantifier, which in turn was in the scope of a *knows that* operator. This defect was eliminated by means of the independence indicator (slash) /. (See Hintikka 2003b.) What characterizes the questioned ingredient is its independence of the epistemic operator, and such independence is perfectly compatible with its being dependent on a universal quantifier, which is in turn dependent on the universal quantifier. In symbols we can now write, for instance, $K(\forall x)(\exists y/K)$ without having to face the impossible task of capturing the threefold dependence structure by means of scopes—that is, by ordering K , $(\forall x)$, and $(\exists y)$ linearly so as to capture their dependence relations.

In this way, we can treat all *wh*-questions and all propositional questions (involving questions where the two kinds of question ingredients are intermingled). The question ingredient of propositional questions turns out to be of the form (v/K) and the question ingredient of *wh*-questions of the form $(\exists x/K)$. We can also close a major gap in our argument so far. The connection between knowledge and decision-making discussed in section “[Knowledge and decision-making](#)” is apparently subject to the serious objection mentioned in section “[The logic of knowledge and information](#)”. It helps to understand a knowledge operator K only when it occurs clause-initially, prefixed to a closed sentence. For it is only such sentences, not all and sundry formulas, that express a proposition that can serve as a justification of an action. Occurrences of K inside a sentence prefixed to an open formula cannot be interpreted in the same way. Now we can restrict K to a sentence-initial position, which eliminates this objection. This also helps to fulfill the promise made in section “[The logic of knowledge and information](#)” of constructing a general logic for the epistemic operator. Here we are witnessing a major triumph of second-generation epistemic logic, which relies on the notion of independence. It solves once and for all the problem of “quantifying in.” It turns out that we do not at bottom *quantify into* a context governed by the epistemic operator K . What we in effect do is to *quantify independently* of this operator.

Why-questions and *how*-questions require a special treatment, which nevertheless is not hard to do. (See, e.g., Hintikka and Halonen 1995.)

The most persuasive argument for the interrogative model nevertheless comes from the applications of the interrogative viewpoint to different problems in epistemology. An important role in such applications is played by the presuppositions of questions and by the presuppositions of answers, better known as their conclusiveness conditions. Examples of such application are offered in Chaps. 4 and 5 of this volume.

The Place of Knowledge in Inquiry

It would take me too far afield here to essay a full-fledged description of the interrogative model. It is nevertheless easy to make an inventory of the concepts that are employed in it. In an explicit model, question-answer steps are interspersed with logical inference steps. Hence the concepts of ordinary deductive logic are needed. As long as the inquirer can trust all the answers, the concepts that are needed are the presuppositions of a question, the conclusiveness condition of an answer (which might be called the “presupposition” of the answer), and the notion of information. To describe an interrogative argument with uncertain answers (responses), we need the notion of tentative rejection of an answer, also known as *bracketing*, and hence also the converse operation of unbracketing, plus ultimately also the notion of probability needed to judge the conditions of bracketing and unbracketing.

What is remarkable about this inventory is that it does not include the concept of knowledge. One can construct a full epistemological theory of inquiry as inquiry

without ever using the k-word. This observation is made especially significant by the generality of the interrogative model. As was indicated, not only is it by means of an interrogative argument that all new information can be thought of as having been discovered, it is by the same questioning method that its credibility must be established in principle.

What this means is that by constructing a theory of interrogative inquiry we apparently can build up a complete theory of epistemology without using the concept of knowledge. We do not need the notion of knowledge in our theory of knowledge—or so it seems. We do not need it either in the theory of discovery or in the theory of justification.

This conclusion might seem to be too strange to be halfway plausible. It is not, but it needs explanations to be seen in the right perspective.

It might perhaps seem that the concept of knowledge is smuggled into interrogative argumentation by the epistemic logic that has to be used in it. This objection is in fact a shrewd one. I said earlier that the logic of questions and answers, which is the backbone of the interrogative model, is part of the logic of knowledge. And this need to resort to epistemic notions is grounded deeply in the facts of the case. It might at first seem that in an interrogative inquiry, no epistemic notions are needed. The presuppositions of questions, questions themselves, and replies to them can apparently be formulated without using epistemic notions.

However, this first impression turns out to be misleading. The structure of and the rules governing it cannot be specified without using some suitable epistemic logic. For one thing, many of the properties of questions and answers are best explained by reference to what is known as the desideratum of a question. This desideratum specifies the epistemic state that the questioner wants to be brought about (in the normal use of questions). For instance, the desideratum of “Who murdered Roger Ackroyd?” is “I know who murdered Roger Ackroyd.” But the desideratum with its *prima facie* knowledge operator is not only a part of a theory of question-answer sequences, it is a vital ingredient of the very interrogative process.

In particular, it is needed to solve Meno’s problem (Plato 1924) applied to interrogative inquiry. In the initial formulation of the rules for interrogative inquiry, it is apparently required that we must know not only the initial premises of inquiry but also their ultimate conclusion. This seems to mean that we can use interrogative inquiry only to explain conclusions we have already reached but not to solve problems—in other words, answer questions by means of questions. But in trying to answer a question by means of interrogative inquiry, we apparently do not know what the ultimate conclusion is. We are instead looking for it. How, then, can we use interrogative inquiry for the purpose of answering questions? The answer is that we must formulate the logic of inquiry in terms of what the inquirer knows (in the sense of being informed about) at each stage. Then we can solve Meno’s problem merely by using the desideratum of the overall question as the ultimate conclusion. But then we seem to need the notion of knowledge with vengeance.

What is true is that a viable theory of questions and answers will inevitably involve an intensional operator, and in particular an epistemic operator in a wide

sense of the word. However, the epistemic attitude this operator expresses is not knowledge in any reasonable sense of the word, not just not in the philosopher's solemn sense. Here, the results reached in section "[The logic of knowledge and information](#)" are applicable. Before an interrogative inquiry has reached its aim—that is, knowledge—we are dealing with information that has not yet hardened into knowledge. It was seen earlier that the logic of such unfinished epistemological business is indeed a kind of epistemic logic, but a logic of information rather than of knowledge.

This point is worth elaborating. Indeed the real refutation of the accusation of having smuggled the concept of knowledge into interrogative inquiry in the form of the epistemic operator used in questions and answers lies in pointing out the behavior of this operator in epistemic inquiry. It may sound natural to say that after having received what is known as a conclusive answer to a question, the inquirer now knows it. But the notion of knowledge employed here is a far cry from the notion of knowledge that philosophers have tried to define. It looks much more like the ugly foreign notion of information. It does not even carry the implication of truth, for the answer might very well have to be bracketed later in the same inquiry. By the same token, it does not even presuppose any kind of stable belief in what is "known." Instead of saying that after having received a conclusive answer, the inquirer knows it, it would be more accurate to say that he or she has been informed about it. Here the advantages of the less deep notion of information are amply in evidence. Unlike knowledge, information need not be true. If an item of information offered to me turns out to be false, I can borrow a line from *Casablanca* and ruefully say, "I was misinformed." The epistemic operator needed in the logic of questions and answers is therefore not a knowledge operator in the usual sense of the term. My emphasis on this point is a penance, for I now realize that my statements in the past might have conveyed to my readers a different impression. What is involved in the semantics of questions and answers is the logic of information, not the logic of knowledge. This role of the notion of information in interrogative inquiry is indeed crucial, but it does not involve epistemologists' usual concept of knowledge at all.

This point is so important as to be worth spelling out even more fully. Each answer presents the inquirer with a certain item of information, and the distinction between question-answer steps and logical inferences steps hinges on the question of whether this information must be old or whether it can be new information. But it is important to realize that such information does not amount to knowledge. In an ongoing interrogative inquiry, there are no propositions concerning which question is ever raised, whether they are known or not. There may be a provisional presumption that, barring further evidence, the answers that an inquirer receives are true, but there is not even a whiff of a presumption that they are known. Conversely, when an answer is bracketed, it does not mean that it is definitively declared not to be known, for further answers may lead the inquirer to unbracket it. In sum, it is true in the strictest possible sense that the concept of knowledge in anything like philosophers' sense is not used in the course of interrogative inquiry.

These observations show the place of knowledge in the world of actual inquiry, and it also shows the only context in which questions about the definition of knowledge can legitimately be asked. The notion of knowledge may or may not be a discussion-stopper, but it is certainly an inquiry-stopper.

It might be suspected that this is due to the particular way the interrogative model is set up. Such a suspicion is unfounded, however. The absence of the concept of knowledge from ampliative inquiry is grounded in the very nature of the concept of knowledge. Questions of knowledge do not play any role in the questioning process itself, only in evaluating its results. For what role was it seen to play in human life? It was seen as what justifies us to act in a certain way. The concept of knowledge is therefore related to interrogative inquiry by asking: When has an interrogative inquiry reached far enough to justify the inquirer's acting on the basis of the conclusions it has so far reached? Or, to align this question with the locutions used earlier, when has the inquiry entitled the inquirer to dismiss the scenarios that are incompatible with the propositions accepted in the inquiry at the time? This is a genuine question, and it might seem to bring the concept of knowledge to the center of the theory of interrogative inquiry.

In a sense it does that. But this sense does not bring the notion of knowledge back as a concept that can possibly figure in the definitory rules of inquiry. It brings knowledge back to the sphere of strategic aspects of inquiry. The question as to whether a conclusion of inquiry has been justified strongly enough for it to qualify as knowledge is on a par with the question as to whether or not a step in an inquiry (typically an answer to a question) should perhaps be bracketed (however tentatively). Both are strategic questions. It is hopeless to try to model knowledge acquisition in a way that turns these decisions into questions of definitory correctness.

Any context-free definition of knowledge would amount to a definitory rule in the game of inquiry—namely, a definitory rule for stopping an inquiry. And once one realizes that this is what a definition of knowledge would have to do in the light of the conception of inquiry as inquiry, one realizes that the pursuit of such a definition is a wild goose chase.

It is important to realize that this conclusion does not only apply to attempted definitions of knowledge that refer only to the epistemic situation that has been reached at the putative end stage of the “game” of inquiry. In other words, it does not apply only to the state of an inquirer's evidence at the end of an inquiry. It also applies to definitions in which the entire history of inquiry so far is taken into account.

This conclusion is worth spelling out more fully. What the conclusion says is that no matter how we measure the credence of the output of interrogative inquiry, there is no reason to believe that an answer to the question as to when an inquirer is justified to act on his or her presumed knowledge depends only on the process of inquiry through which the inquirer's information has been obtained independently of the subject matter of the inquiry. In an old terminology, the criteria of justification cannot be purely *ad argumentum*, but must also be *ad hoc*. Neither the amount of

information nor the amount of justification that authorizes an agent to stop his or her inquiry and act on its results can always be specified independently of the subject matter—for instance, independently of the seriousness of the consequences of being wrong about the particular question at hand. And if the justification depends on the subject matter, then so does the concept of knowledge, because of the roots of our concept of knowledge in action.

But since the notion of knowledge was seen to be tied to the justification of acting on the basis of what one knows, the concept of knowledge depends on the subject matter and not only on the epistemological situation. Accordingly, no general definition of knowledge in purely epistemological terms is possible.

This point is not a relativistic one as far as the possibility of *a priori* epistemology is concerned. If anything, the divorce of knowledge from inquiry underlines the objectivity of inquiry and its independence of the value aspects of the subject matter. The fashionable recent emphasis on the alleged value-ladenness of science is misleading in that it is typically predicated on forgetting or overlooking that the question as to when the results of scientific inquiry authorize acting on them is different from questions concerning the methodology of scientific inquiry itself. The dependence of the criteria of knowledge on subject matter ought to be a platitude. It is one thing for Einstein to claim that he knew that the special theory of relativity was true notwithstanding *prima facie* contrary experimental evidence, and another thing for a medical researcher to be in a position to claim to know that a new vaccine is safe enough to be administered to sixty million people. But some relativists mistakenly take this platitude to be a deep truth about scientific methodology and its dependence on subject matter. This is a mistake in the light of the fact that the allegedly value-laden concept of knowledge does not play any role in the actual process of inquiry.

Here, a comparison with such decision principles as the maximization of expected utility is instructive. What an inquiry can provide is only the expectations (probabilities). But they do not alone determine the decision, which depends also on the decider's utilities. Hence the criteria of knowing cannot be defined by any topic-neutral general epistemology alone. But this dependence does not mean that the probabilities used—misleadingly called “subjective” probabilities—should in rational decision-making depend on one's utilities. Decision-making based on such probability estimates would be paradigmatically irrational.

The influence of subject matter on the notion of knowledge does not imply that the interrogative process through which putative knowledge has been obtained is irrelevant for the evaluation of its status. Here lies, in fact, a promising field of work for applied epistemologists. Material for such work is available in, among many other places, different kinds of studies of risk-taking. Even though considerations of strategies do not help us to formulate a topic-neutral definition of knowledge, in such a topic-sensitive epistemology they are bound to play a crucial role. This is a consequence of the general fact that in game-like processes, only strategies, not individual moves, can in the last analysis be evaluated.

Comparisons with Other Epistemologists

Relativizing our humanly relevant concept of knowledge to some particular subject matter also provides a strategy of answering a philosophical skeptic. If knowledge claims depend for their very meaning on the criteria governing some particular walk of human action, then so also must reasonable doubts. It is only unspecific “philosophical” doubts that do not have built into their own logic standards that show how they can be surmounted.

One philosopher who would have agreed with my thesis concerning the dependence of the criteria of knowledge on the subject matter, and who in fact supplied reasons for it, is Ludwig Wittgenstein. In Hintikka [forthcoming](#), I have shown that according to Wittgenstein’s mature views, the concept of knowledge cannot be used in what I have called “primary language-games.” These language-games are for Wittgenstein the direct links between language and reality. In them, we cannot, in Wittgenstein’s metaphor, drive a wedge between language and what it expresses. Such a primary language-game does not operate by means of criteria, but by means of spontaneous responses. If I try to say in such a primary language-game “I know that I am in pain,” all that I can express is the same as “I am in pain.” And in a primary language-game, to utter “I am in pain” is but a form of pain-behavior.

In Wittgenstein’s view, epistemic concepts can be used only in what I have called “secondary language-games.” These secondary language-games presuppose primary ones. They do not operate through spontaneous responses, verbal or behavioral, and hence they must involve criteria. For this reason, epistemic vocabulary can be used in them. But those criteria are different in different secondary games. Hence the force of epistemic terms depends on the particular secondary game in which they are being used. Saying this is very nearly nothing but Ludwigspeak for saying that the criteria of knowing depend on the subject matter.

Other epistemologists have not been unaware, either, of connections between the justifiability of knowledge claims and the subject matter involved. (See, e.g., DeRose 1995; Cohen 1998; Williams 2001, ch. 14; Bonjour 2002, pp. 267–271.) They seem to have ascribed the dependence in question to the context of inquiry rather than to its subject matter, however. Unless and until the notion of context used here is clarified, I remain doubtful of such claims of context-dependence. For instance, criteria of knowing that a vaccine is safe depend on the life-or-death character of the subject matter, but they presumably should not depend on the context, which may be an administrative decision to initiate compulsory vaccination or a pharmaceutical company’s promise to produce the requisite vaccine. However, if the notion of context is interpreted in such a way that it includes first and foremost the subject matter of inquiry, contextualist epistemology might very well converge with the views expressed here. In this work, contextualism is not examined further, however.

Moreover, contextual epistemologists seem to have assimilated the insight into the context-dependence of knowledge to another insight—namely, to the insight that every epistemological inquiry concerns some particular model, a “system”

as physicists would call it, which typically is not an entire world. (See here Hintikka 2003a.) All epistemological inquiry is therefore contextual in this sense of being relative to a model (scenario or “possible world”). But this does not make epistemology itself contextual or relative as a scientific theory is made contextual or relative by the fact that it is inevitably applied to reality system by system. Hence the impact of the line of thought pursued here is diametrically opposed to the most common form of contextualism. This form of contextualism aims at the rejection of global epistemological questions. (See Bonjour 2002, p. 267). For us, global epistemological questions concern in the first place the nature of interrogative inquiry, and they are in no sense context-dependent or even dependent on the subject matter.

The Folly of Trying to Define Knowledge

The concept of knowledge thus belongs to applied epistemology, not to general epistemology. The criteria of knowledge concern the conditions on which the results of epistemological inquiry can be relied as a basis of action. It follows that it is an exercise in futility to try to define knowledge in any general epistemological theory. Such a definition could never help Satire’s sleeper spy. But my point is not only about what is not useful in practice. The extensive discussions about how to define knowledge are not only useless for applications, they are theoretically misguided. Here the true relations of the concepts knowledge and truth to definability are almost precisely opposite to what they have been taken to be recently. Tarski (1956) proved certain results concerning the undefinability of truth. Philosophers and other thinkers have taken Tarski’s results at their apparent face value, without realizing how restrictive the assumptions are on which these impossibility results are predicated. (See Hintikka 2002.) They have even let Tarski’s results discourage them to the extent of giving up attempts to define truth. Tarski notwithstanding, a truth predicate can be formulated for sufficiently rich languages in a philosophically relevant sense in the same language. In contrast, no major philosopher has to the best of my knowledge openly maintained it to be a folly to try to define knowledge. Yet if someone has done so, that philosopher would have been much closer to truth than a philosopher who argues that it is foolish to try to define truth. (See Davidson 1996.)

Belief as a Product of Inquiry

The notion of knowledge belongs to applied epistemology because it is connected conceptually with the notions of acting and decision-making. The particular connection is not crucial. But if it does not matter, similar conclusions must hold also for those other epistemic concepts that are connected conceptually with behavior,

especially with decision-making. The concept of belief is a case in point. And conclusions similar to the ones that have been reached here concerning the notion of knowledge can in fact be drawn concerning the notion of belief. If you are inspired by this line of thought to review the structure of interrogative inquiry with a view to finding a role for the notion of belief there, you will not find such a role. Receiving an answer and incorporating it into one's interrogative inquiry is not the same as adopting a new belief. Acceptance is not the same thing as belief-formation. (For a discussion of their relationship, see Cohen 1992.) For one thing, at no stage of an interrogative inquiry are there any indications whether or not the inquirer is prepared to act on the truth of the propositions that the inquirer has at that stage accepted (and not bracketed). Hence the entire theory of knowledge acquisition can—and must—be developed without using the notion of belief. This notion does not play any role in an interrogative inquiry, only in the evaluation of its putative end-point. If one thinks about it, the notion of belief does not play much of a role in the methodology of science. What I am suggesting is that it should not play any more of a role in general epistemology either.

There is thus a certain partial epistemological parallelism between belief and knowledge. This parallelism has not been appreciated by epistemologists. Ever since Plato, the two notions are habitually contrasted to each other. This contrast is nevertheless seriously misleading, as far as the epistemology of belief is concerned.

It seems to me that the same point is unwittingly attested to by all the decision theorists who are using beliefs as an ingredient in rational decision-making. Such a use would be pointless unless there were some previous reasons to think that the beliefs in question can rationally be expected to be true. And such reasons must somehow come from the inquirer's previous experience, if one is a good empiricist.

Belief, too, is connected with criteria as to when I am ready to act on a certain item of information I have received. But whereas the criteria of knowing are impersonal (even though they can be relative to the subject matter), the criteria of belief can be personal and dependent on an even wider selection of the aspects of the subject matter. In claiming to know, I am making a commitment to others, but in forming a belief, I am usually responsible only to myself.

There are also intermediate cases. For instance, a scientist's beliefs *qua* scientist are subject to the standards of acceptance in his or her scientific community. The crucial point is that those beliefs are, in the case of a scientist, formed as a result of an inquiry, rather than, so to speak, as a response to the question, "What do you think about it?" One may very well catch a physicist asking whether he or she should believe a certain hypothesis in the light of available evidence. But one is even likelier to find a scientific inquirer asking what new information he or she should try to acquire—for instance, what experiments to carry out—in order to be in a position to entertain a certain belief.

In general, the same things can thus be said of belief and its standards as were said earlier of knowledge. Belief statements, like knowledge statements, express entitlement of a certain sort. In neither case does an agent have to avail himself or herself of such entitlement. Beliefs need not manifest themselves in overt behavior any more than knowledge. Hence, decision theorists' frequent assumption that an

agent's beliefs (or degrees of belief) together with utilities determine his, her, or its behavior is in need of scrutiny. Above all, beliefs, too, must be thought of as being formed by means of inquiry.

What I take to be a related point has been expressed by Timothy Williamson by pointing out that a "reason is needed for thinking that beliefs tend to be true." (Quoted from the abstract of his contribution to the conference on "Modalism and Mentalism in Modern Epistemology," Copenhagen, January 29–31, 2004.) The relationship is mediated by the fact that, if I am right, interrogative inquiry is, in the last analysis, the only way of arriving at true beliefs.

The conclusions reached here have repercussions for the entire research strategies that should be pursued in epistemology. For instance, there is a major school of thought that conceives of inquiry as a series of belief revisions. But is this at all realistic as a description of what good reasoners actually do? Georges Simenon's Inspector Maigret is sometimes asked what he believes about the case he is investigating. His typical answer is: "I don't believe anything." And this does not mean, contrary to what one might first suspect, that Maigret wants only to know and not to believe and that he has not yet reached that state of knowledge. No—in one story he says, "The moment for believing or not believing hasn't come yet." (Georges Simenon, *Maigret and the Pickpocket*, Harcourt Brace Jovanovich, San Diego, 1985.) It is not that Maigret has not carried his investigation far enough to be in a position to know something. He has not reached far enough to form a belief. (The mere possibility of using the locution "belief formation" is instructive.) In serious inquiry, belief too is a matter whether an inquiry has reached far enough.

Belief, too, concerns the question of when to stop an inquiry. That is the place of this concept in the framework of the interrogative approach. The difference between belief and knowledge does not lie merely in the degree of justification the believer has reached. It does not mean that there is an evaluative component in knowledge but not in belief. The difference lies in the kind of evaluation involved. It is much more like the difference between satisfying an agent's own freely chosen standards of epistemic confidence and satisfying certain impersonal standards that are appropriate to the subject matter.

In linguists' terminology, knowing is an achievement verb. In a way, although not in a literal sense, believing is in the context of interrogative inquiry likewise an achievement notion. What should be studied in epistemology is belief-formation and not only belief change. The notion of belief cannot serve the role as a determinant of human action that is assigned to it in decision theory if it is not influenced by what the agent knows. But such influence is usually not studied in decision theory.

One corollary to the results we have reached concerns philosophers' research strategies. What we can see now is that the interrogative model is not only a rational reconstruction of knowledge acquisition, it can also be used as a model of belief formation. The insight that belief, too, is typically a product of inquiry lends some renewed interest to the "true belief" type of attempted definitions of knowledge. What they perhaps succeed in capturing is admittedly not philosophers' strong sense of knowledge. But there may be other uses (senses?) of the words *knowledge* and *knowing* that can be approached by means of such characterizations.

Philosophers tend to downplay the role of certainty, especially of experienced certainty, in explicating the notion of knowledge. There is nevertheless a third-person use of knowledge attributions in which the meaning of knowing is very close to true conviction reached through inquiry. In such cases, the inquirer has convinced himself or herself by means of inquiry of the truth of some proposition or other even when, by some standards, the inquirer has not yet reached sufficient justification.

A typical context is when an investigator has reached a correct conclusion—for instance, identified the perpetrator—through inquiry and has become convinced of this conclusion even though his or her reasons would not satisfy the standards of evidence in a court of law. It is interesting to note that in such a usage, the true conclusion must have been reached through a viable strategy. A mere guess would not amount to knowledge even in such cases. (Notice that one could not attribute knowledge in this sense to an automaton or to a database.) This observation may be related to Frank Ramsey's (1978) attempt to characterize knowledge as true belief obtained through a reliable method. This sense of knowing seems to be much closer to colloquial usage than the one philosophers have in vain been trying to define.

Repercussions for Other Approaches

From the point of view we have reached, we can also see some serious problems about the Bayesian approach to inquiry. (See, e.g., Earman 1992.) This approach deals with belief change rather than belief-formation. Insofar as we can find any slot for belief-formation within the Bayesian framework from the point of view of any simple application of, it is pushed back to the selection of priors. In other words, it is made entirely *a priori*, at least locally. This is by itself difficult to implement in the case of theory-formation (belief-formation) in science. Is it, for instance, realistic to assume that a scientist can associate an *a priori* probability with each and every possible law of nature? And these doubts are reinforced by general conceptual considerations. Assignments of priors amount to assumptions concerning the world. What is more, prior probabilities pertain to the entire system (model, “world”) that the inquirer is investigating bit by bit. How can the inquirer choose such priors on the basis of his or her limited knowledge of the world? These difficulties might not be crucial if there existed a Bayesian theory of belief-change that included a study of changes of priors. Even though such changes have been studied, it seems to me that their theory has not been developed far enough in the Bayesian framework to cover all possibilities.

All sorts of difficult questions face us here. For instance, in order to use Bayesian inference, we need to know the prior probabilities. It seems to be thought generally that this does not amount to asking very much. This may be true in situations in which the primary data is reasonably reliable, as in typical scientific contexts. However, if our evidence is likely to be relatively unreliable, the situation may be different—for instance, when we are dealing with testimony as our basic form of

evidence. I may easily end up asking: Do I really have enough information to make the guess concerning the world that was seem to be involved in the choice of the priors?

For one thing, even though the matter is highly controversial, fascinating evidence to this effect comes from the theory of so-called cognitive fallacies studied by mathematical psychologists such as Amos Tversky and Daniel Kahneman. (See, e.g., Kahneman et al. 1982; Piatelli-Palmarini 1994.) These alleged fallacies include the conjunction fallacy and the base-rate fallacy. As I have suggested in Chap. 9 of this volume (and in Hintikka 2004), at least in certain “crucial experiment” cases, the alleged mistakes are not fallacious at all, but rather point to certain subtle but very real ways in which one’s prior probabilities can (and must) be changed in the light of new evidence. They do not show that certain fallacious ways of thinking are hardwired into human beings. Rather, what they show is that Bayesians have so far failed to master certain subtle modes of ampliative reasoning. Tversky’s and Kahneman’s Nobel Prize notwithstanding, epistemologists should take a long critical look at the entire theory of cognitive fallacies.

Here I can only give indications of how to view the cognitive fallacies conundrum. Very briefly, in the kind of situation that is at issue in the alleged conjunctive fallacy, the prior probabilities that one in effect relies on include the degrees of probability (credibility) assigned to the reports one receives. But that credibility can not only be affected by suitable new evidence, it can be affected by the very report itself. If the report shows that the reporter is likely to know more about the subject matter than another one, it is not fallacious to assign a higher prior probability to his or her report, even though it is a conjunction of a less credible report and further information.

In the case of an alleged base-rate fallacy, there is no conceivable mistake present if the intended sample space consists simply of the different possible courses of events concerning the crucial event—for example, a traffic accident. Base rates enter into the picture only when a wider class of courses of events is considered—for example, all possible courses of events that might have led to the accident. This means considering a larger sample space. Either sample space can of course be considered entirely consistently, depending on one’s purposes. A fallacy would inevitably be committed only if the only legitimate application of our language and our epistemological methods was to the entire world—in this case, the larger sample space. But such an exclusive preference of the larger sample space is but an instance of the one-world assumption, which I have criticized elsewhere. (See Hintikka 2003a.)

Whither Epistemology?

The moral of the insights we have thus reached is not merely to avoid certain words in our epistemological theorizing. It calls for rethinking our overall research strategies in epistemology. And the spirit in which we should do so is perhaps

illustrated by the first epistemologist in the Western philosophical tradition. Socrates did not claim that he knew anything. In the manner of a practitioner of my interrogative method, what he did was to ask questions. I suspect that it is only in Plato's dialogues that he was looking for a definition of knowledge. And Plato put this question (and other questions of definition) into Socrates's mouth because Plato shared the widespread Greek assumption that the definition of X gives us the "blueprint" that enables us to bring about X. (See Hintikka 1974, ch. 1–2.) This applies both to the generic search for knowledge and to the quest of particular items of knowledge. Thus, insofar as Plato contemplated knowledge-seeking (information-seeking) by questioning in our sense, he would have had to say that we must know what we are looking for there and that it is this knowledge alone that can guide our search. (No wonder he was worried about Meno's problem.) By the same token, all search for knowledge would have had to be guided by our knowledge of what knowledge is.

Hence it is seen that Plato had in one important respect the same focus as we: the quest for knowledge rather than the justification of beliefs. The definition of knowledge was thought of by Plato as a means for this quest. If so, the pursuit of the definition of knowledge would indeed have been the alpha and omega of epistemology. But we do not think in that way. The training that Satire's spymaster is supposed to have received did not aim exclusively at learning the definition of knowledge. For us, the fact that knowledge can be considered the end product of inquiry shows on the contrary that it cannot play any role in the process of inquiry. Hence the wild goose chase of the definition of knowledge only shows that too many contemporary epistemologists are still bewitched by Plato's assumptions. This is one of the reasons why at the beginning of this chapter, I called contemporary academic epistemology antiquated. Maybe it is time for its practitioners to take up some more up-to-date problems.

References

- Austin, J. (1961a) (original 1946). Other minds. In *Philosophical papers*. Oxford: Clarendon Press, especially pp. 67–68.
- Austin, J. (1961b). Performative utterances. In *Philosophical papers*. Oxford: Clarendon Press, ch. 10, especially p. 220.
- Bonjour, L. (2002). *Epistemology: Classical problems and contemporary responses*. Lanham: Rowman & Littlefield.
- Cohen, L. J. (1992). *Essay on belief and acceptance*. Oxford: Clarendon.
- Cohen, S. (1998). Contextual solutions to epistemological problems. *Australasian Journal of Philosophy*, 76, 289–306.
- Collingwood, R. G. (1940). *An essay on metaphysics*. Oxford: Clarendon Press, especially ch. I, sec. 5.
- Davidson, D. (1996). The folly of trying to define truth. *Journal of Philosophy*, 93, 263–278.
- DeRose, K. (1995). Solving the skeptical problem. *Philosophical Review*, 104, 1–52.
- Earman, J. (1992). *Bayes or bust? A critical examination of Bayesian confirmation theory*. Cambridge: MIT Press.

- Frege, G. (1984) (original 1918–19). Thoughts. In *Collected papers*. Oxford: Basil Blackwell.
- Gadamer, H. -G. (1975) (original 1960). *Truth and method*. Continuum, New York, especially the section “Logic of Questions and Answers,” pp. 333–341.
- Hintikka, J. (1974). *Knowledge and the known*. Dordrecht: Reidel.
- Hintikka, J. (1998). What is abduction? The fundamental problem of contemporary epistemology. *Transactions of the Charles Peirce Society*, 34, 503–533. A revised version, “Abduction—Inference, Conjecture, or an Answer to a Question,” appears as Chapter 2 in this volume.
- Hintikka, J. (1999). *Inquiry as inquiry: A logic of scientific discovery*. Dordrecht: Kluwer.
- Hintikka, J. (2002). Post-tarskian truth. *Synthese*, 126, 17–36.
- Hintikka, J. (2003a). A distinction too few or too many: A vindication of the analytic vs. synthetic distinction. In Carol C. Gould (Ed.), *Constructivism and practice: Toward a historical epistemology* (pp. 47–74). Lanham: Rowman & Littlefield.
- Hintikka, J. (2003b). A second-generation epistemic logic and its general significance. In Vincent F. Hendricks, et al. (Eds.), *Knowledge contributors* (pp. 33–56). Dordrecht: Kluwer Academic. And as Chap. 3 in this volume.
- Hintikka, J. (2004). A fallacious fallacy? *Synthese*, 140, 25–35. And as Chap. 9 in this volume.
- Hintikka, J. (forthcoming). *Wittgenstein on knowledge and skepticism* (Working paper).
- Hintikka, J., & Halonen, I. (1995). Semantics and pragmatics for why-questions. *Journal of Philosophy*, 92, 636–657.
- Husserl, E. (1983) (original 1913). *Ideas pertaining to a pure phenomenology. First book*. The Hague: Martinus Nijhoff.
- Kahneman, D., Slovic, P., & Tversky, A. (Eds.). (1982). *Judgment under uncertainty: Heuristics and biases*. Cambridge: Cambridge University Press.
- Kant, I. (1787). *Kritik der reinen Vernunft*, 2nd ed. (see Preface, p. xiii).
- McCarthy, J. (1990). Circumscription—A form of non-monotonic reasoning. *Artificial Intelligence*, 13, 27–39 and 171–172.
- Montague, R. (1974). *Formal philosophy*. New Haven: Yale University Press.
- Piatelli-Palmarini, M. (1994). *Inevitable illusions*. New York: Wiley.
- Plato. (1924). *Meno, Plato: With an English translation* (Loeb classical library, vol. IV). Cambridge: Harvard University Press.
- Ramsey, F. (1978) (original 1929). Knowledge. In *Foundations: Essays in philosophy, logic, mathematics, and economics* (pp. 126–127). London: Routledge & Kegan Paul.
- Safire, W. (1995). *The sleeper spy*. New York: Random House.
- Shope, R. K. (1983). *The analysis of knowing: A decade of research*. Princeton: Princeton University Press.
- Tarski, A. (1956) (original 1935). The concept of truth in formalized languages. In *Logic, semantics, metamathematics* (pp. 152–278). Oxford: Clarendon Press.
- Williams, M. (2001). *Problems of knowledge: A critical introduction to epistemology*. Oxford: Oxford University Press.