

Radu J. Bogdan

WHAT IS EPISTEMIC DISCOURSE ABOUT?

1. INTRODUCTION

In the 1960's and 70's Jaakko Hintikka has written extensively about epistemic logic, epistemic concepts and ordinary epistemic discourse. As a (graduate) student of Jaakko's toward the end of that period, I was somewhat familiar with that body of work and even discussed some fragments of it in my dissertation on the pragmatics of knowledge. Since then my interests developed in different directions, toward philosophy of mind and cognitive science in general and commonsense or naive psychology in particular. This paper looks back at Jaakko's work on epistemic discourse from the vantage point of my later work on naive psychology.

The title question, about the subject matter of epistemic discourse, is not an easy one to answer, for several reasons, surveyed in the next section. These reasons bear on the tight yet not fully identified and understood links between epistemic discourse, on the one hand, and naive

psychology, ordinary language, and epistemological expectations, on the other hand. Section 3 focuses on the relation between naive psychology and epistemic discourse, surveys some empirical data about the development of epistemic discourse in the context of a more general mental development, and suggests the need for a top-down explanatory approach to our competence for epistemic discourse. In section 4 I propose to read Jaakko Hintikka's work on epistemic discourse in cognitive-scientific terms and view it as contributing to such an explanatory approach. A concluding section 5 suggests deeper evolutionary reasons why epistemic discourse would likely work along the lines such as those suggested by Hintikka's account.

2. PROBLEMS

The notion of epistemic discourse is meant here to include epistemic terms, such as those for perception, belief, memory or knowledge, and epistemic locutions, such as 'she knows (believes, perceives) that p.' Epistemic discourse is part of a broader mentalist discourse that also contains terms and locutions for various other sorts of mental states and intentional relations, such as desire, intention, fear,

regret, and so on. The mentalist discourse and its epistemic subset express in ordinary language a body of knowledge or a competence variously known as naive or folk or commonsense psychology or, somewhat more technically, theory of mind. Naive psychology is credited with enabling us to conceptualize, explain and predict what is going on in other minds and our own, and how these mental goings on translate into behaviors. Since naive psychology is the home base of our epistemic concepts and discourse, our understanding of the former is bound to affect our understanding of the latter. Whence the first problem: in different ways and terms, philosophers and psychologists disagree rather sharply over the nature, job and modus operandi of our naive psychology, and this disagreement is inevitably echoed by disagreements over what epistemic discourse is all about and how it works. I will return to this problem in the next section.

But even if this first problem were solved, there is a second problem to be faced. Ordinary epistemic discourse appears to have duties that go much beyond the basic job description of naive psychology -- or, if we choose to look at it somewhat differently, duties that would expand the job description of naive psychology in directions that are even more elusive and controversial than those of its basic job description. These

other duties, emanating from the nature of linguistic communication and social interactions, bring a host of contextual and pragmatic parameters into our epistemic discourse and weave them -- in ways still poorly understood -- into the fabric of naive-psychological concepts and attributions.

The problems just surveyed are problems of substance that will be in focus throughout this paper. There are also methodological problems of how to approach and study epistemic discourse. Philosophers have done more work on epistemic discourse, and ventured more explanations of it, than either psychologists or linguists. Since I will review some psychological data and hypotheses in the next section, I should say a few words now about the philosophical approaches.

Among the ways in which philosophers and logicians have gone about studying ordinary discourse in general, one stands out. It consists in making sense of and explicating the linguistic intuitions of the speakers of a natural language. In the case of epistemic discourse, we are talking about the epistemic intuitions that speakers have about knowledge or belief or memory claims and attributions. I use the notion of intuition here in a loose analogy to Noam Chomsky's notion of grammatical intuitions that speakers of a natural language have in distinguishing grammatically

correct from incorrect uses of words and sentences in that language.

They can do so without necessarily knowing how they do it, by what rules -- at least not until properly schooled. The same is thought to be true of the intuitions speakers have about epistemic discourse.

One serious problem with basing philosophical reconstructions on epistemic intuitions is that it is not always clear whether the intuitions in question are used to support an epistemological analysis or, alternatively, an epistemic-discourse analysis -- that is, whether the intuitions bear on what the concepts of knowledge or belief ought to be or, alternatively, on elucidating our ordinary discursive practices involving knowledge or belief claims. To see what the distinction is and why it matters, suppose we ask, for example, what the concept of knowledge is and when it is attributed. On the epistemological reading, we are asking what knowledge is or what it takes to have knowledge in general and in what conditions knowledge can be said to be instantiated in some organism or system. This is the question that philosophers have asked forever but whose answer is still rather elusive. Since Plato, most answers have taken the form of conceptual analyses that unpack a typically idealized or normative concept. The standard analysis is that of knowledge as justified true belief, but there are other analyses as well, including naturalist accounts

that replace justification with causation or reliable information processing. In the heyday of Gettier-like games that epistemologists used to play with gusto (but mercifully, no more), the examples of and counterexamples to some definition of knowledge were checking epistemological intuitions with an eye to this conceptual project, even though many of these intuitions may have originated in the use of epistemic discourse.

The epistemic-discourse reading, on the other hand, is concerned solely with the epistemic terms and locutions used normally, rather descriptively and usually contextually by language users. The analytic project here is to figure out and explain the rules of and constraints on the ordinary use of such terms and locutions. The distinction between the two readings allows us to say, without contradiction, that one can make a knowledge claim to the effect that he knows that *p* even though, on some epistemological analysis, he does not really have that knowledge because, for example, he lacks appropriate justification. Perhaps the best known and most insightful analyses of epistemic discourse can be found in Wittgenstein's later writings, in some of Gilbert Ryle's, John Austin's and those of other philosophers of ordinary language (Urmson, Malcolm, etc.).

Yet even when the territory is clearly demarcated, and epistemic discourse is seen for what it is, independently of epistemology, there is no guarantee that the linguistic intuitions associated with it would or could reveal how it works and why. As the main source of insights about and analyses of epistemic discourse, ordinary language philosophy is programmatically descriptive and rather uninterested in explanation. Explanation requires taking linguistic intuitions and practices at best as data or clues pointing to deeper causes or reasons why they work as they do. If epistemic discourse is handled by a psycholinguistic competence or expertise, then explanation requires a theory of that competence or expertise. The first order of business, in the next section, is to establish the antecedent of this conditional and get a sense of the competence in question. In the section following it, I propose to look at Jaakko Hintikka's work on epistemic discourse as contributing to an explanatory theory of that competence.

3. NAIVE PSYCHOLOGY AND EPISTEMIC DISCOURSE

Making epistemic attributions in ordinary contexts by employing an ordinary epistemic discourse requires the resources of naive psychology, which is our competence to recognize and represent how other minds and our own relate to the world in perception, desire, intention, thinking, or memory. The most systematic empirical study of naive psychology and its language has been so far undertaken by developmental psychologists.

Despite inevitable and often sharp disagreements over the nature of naive psychology as a mindreading competence, most researchers agree on some innate and prelinguistic basis for the competence and also agree on several age-related milestones in the development of the competence. The earliest abilities to detect and represent basic intentional relations, such as looking at, seeing, trying to, and wanting, begin to be exercised by children before the age of 1 and thus before the onset of language. It is likely that these early and prelinguistic naive-psychological and epistemic insights might influence and perhaps constrain the later development of epistemic discourse and of the concepts based on it. Thus, some psychologists think that the concept of belief may be modeled on the earliest and prelinguistic concept of perception or gaze. In close analogy, an early version of the concept of knowledge may be modeled on the prelinguistic concept of seeing, in the

sense of successful perception or, more generally, successful access to information (see Perner 1991 for a general survey; also Bogdan 1997).

The main implication for our discussion is that the earliest and most central concepts of naive psychology are prelinguistic and thus owe nothing to the rules and practices of linguistic communication in general and mentalist and epistemic discourses in particular -- although they may owe much to prelinguistic interpersonal interactions (Bogdan 2000, chapter 3). Epistemic discourse itself seems to have its own developmental schedule, although it has been less investigated than the developmental schedule of naive psychology. Still, there some pertinent data, which I will report telegraphically and then weave into our discussion. The main sources are Bartsch and Wellman (1995) and Nelson (1996).

The former authors distinguish several phases in the child's acquisition of epistemic terms. The first terms, acquired a few months after the age of 2, are WANT and LIKES. Around 3 or soon afterwards emerge DESIRE, BELIEF and THOUGHT. Around 4 the term for BELIEF is used for explaining actions, first those of others, before those of self. Only around 4 do children begin to distinguish between KNOW and THINK or GUESS, although they do not seem to distinguish between THINK and

GUESS before the age of 8. What do these data mean? They clearly show an emerging mastery of epistemic terms and attributions. But what exactly do these terms and attributions represent at each developmental stage? This is a difficult question that psychological research has not yet answered fully. There are two main reasons for the difficulty. The first and most important is that there is no agreement among philosophers or psychologists over what epistemic discourse in general is about. The nature and function of this discourse are still to be plausibly defined. Do epistemic terms and locutions represent what is going on in the minds of those targeted by them? Or do such terms and locutions represent something entirely different, such as conversational appropriateness, evaluation of information and evidence, prediction of behavior, and the like? Or are epistemic representations targeting a mixture of mind and outside factors in combinations still to be figured out?

This indeterminacy reflects a parallel but deeper uncertainty and hence disagreement about the nature and function of the more basic and broader competence for naive psychology. Again, is naive psychology directed at minds, our own and those of others, or at something else, of which minds may be only a part? What is clearer already is that the first intentional relations to be represented by very young children (and

possibly great apes), such as gaze, attention, and behavioral postures and motions indicating goals or simple desires, are the most visible, informative and causally manipulable relations. There is also a rather wide consensus that the naive-psychological categories of these relations are likely to be innate and modularized or at least to reflect innate predispositions. These properties make evolutionary sense (Bogdan 1997). It is also likely that the meanings of the first mentalist words, such as SEE, WANT or DESIRE, track closely the relevant perceptual experiences that activate these primordial categories.

The more difficult problem is determining what happens in the second and later phases of the acquisition of epistemic and, more generally, mentalist vocabularies, when the higher-level and more abstract categories of belief, opinion, thought, and knowledge no longer track perceptual experiences and depend increasingly on linguistic descriptions and other social and contextual factors. The child's naive psychology now becomes inextricably linked -- indeed woven into -- the mentalist and specifically epistemic discourses. As a result, it is during these later phases that the child assimilates most of the adult epistemic meanings and the language games in which they are involved. It appears that this assimilation process is complex, difficult and protracted, with

comprehension emerging earlier than production. It is symptomatic that the child's meaning of KNOW remains different from the adult meaning, and keeps changing, until about the age of 12.

So what is going on? Hard to tell, but a few developmental facts point to a dramatic mental revolution that affects naive psychology during the second, nonperceptual phase that begins around the age of 4, when it gradually moves beyond representing here-and-now situations and tracks more abstract attitudes, such as nonperceptual belief, intention or knowledge. Two contrasting metaphors may help clarify this transformation. Until around the age of 3 to 4 the young mind operates on a single screen, where perceptual inputs of current events are displayed and constantly updated by new inputs. It is a mind largely confined to current motivation and perception, controlled by a unique focus of attention, and representing things on a single mental screen. After 4, the young mind (mostly its prefrontal cortex) is shaken by several mental commotions, executive as well as cognitive, and revolutionary in their cumulative impact. Chief among them are the inhibition of current perception, the linguistic recoding and representational explicitation of earlier procedural competencies, such as counting, mental imagery, and naive theories of various domains,

including naive psychology, and the ability to deploy multi-layered clusters of thoughts and to embed thoughts into other thoughts. These changes liberate the young mind from the captivity of single-screen mentation and enable it to entertain simultaneously, on different and interconnected mental screens, nested sets of alternative and often conflicting representations of actual and nonactual, current, past and counterfactual situations. The single screen of early childhood is replaced by a multi-screen or multiplex mentation.

Among the new mental activities made possible by the emerging multiplex mind, two are relevant to our discussion. One is the imaginative and often counterfactual access to nonactual, future as well possible situations or worlds. As a result, it becomes now possible to reconfigure earlier naive-psychological categories and to conceptualize new propositional attitudes, such as thought, intention or knowledge, in terms of possible worlds -- at least when the default attributions fail. Hintikka's possible worlds semantics for epistemic attributions thus has some psychological bite. The point is not that the young or adult epistemic attributor envisages possible worlds whenever she makes a belief or knowledge claim for herself or others. The point is that she could do that when the need arises -- for example, in contexts of uncertainty, doubt,

criticism, incomplete evidence, high stakes, rigorous inquiry, and so on. And the further related point is that she would not recognize epistemic ascriptions and the concepts behind them for what they are, if, for some reason, thinking and talking in terms of possible worlds would be always unavailable. The other new mental activity made possible by multiplex mentation is integrating many sources of information across several modes of representation -- language, memory, perception, imagination, inference, and so on. Epistemic attributions require such integration -- for example, in iterating attributions involving different attitudes, such as belief, perception and memory, as in 'he believes that she remembered seeing him going home.'

The point of these remarks is that the naive-psychological and epistemic concepts and attributions available to and employed by a multiplex mind are vastly different and more complex from those of the younger uniplex mind. I think that neither observations nor experiments nor analyses of linguistic intuitions are sufficient to reveal what the former are all about. What is needed is a theory of the cognitive tasks that the new naive-psychological and epistemic concepts and attribution abilities are designed to carry out -- in other words, a theory in the spirit of recent cognitive science. From a logical and semantic angle, this is how

I read Jaakko Hintikka's work on epistemic logic and how I think it can contribute to a better understanding of epistemic discourse.

4. A THEORY OF EPISTEMIC MEANINGS

I begin with a familiar picture of cognitive-scientific explanation, in whose terms I want to frame my discussion. To understand the mind as a system of mechanisms that act on information in pursuing its goals, the theorist must figure out the programs run by the mechanisms and thus the competencies involved in processing information and acting on it; but to figure out the programs, the theorist must antecedently identify the tasks the programs execute and the problems encountered and solved in the execution. The analysis and explanation thus proceed top-down: from Tasks to Programs to Mechanisms to Ware (hard, wet, whatever). I abbreviate it as the T->P->M->W method. Thanks to the influential work of Noam Chomsky on language, David Marr on vision, and Allen Newell on artificial intelligence (to cite pioneers), the T->P->M->W method has achieved classic status in cognitive science.

One indication that Hintikka's angle on epistemic discourse fits the explanatory methodology of cognitive science is that he is no friend of intuitions as the primary basis for philosophical or formal analysis. According to Hintikka, when not wrong, which they may often be, intuitions tend to lead not to the concepts or abilities they purport to illuminate but to some collateral relations or indirect associations. This is why a deeper analysis is needed. It is in this spirit that Hintikka conceives of epistemic logic as an explanatory model of the workings of ordinary language. It brings out the "deep logic" (which I read as: core tasks) underlying epistemic discourse (Hintikka 1969, 3-5). It should be noted that Hintikka's work in many other domains -- such as inductive logic, the logic of questions, mathematical reasoning, and game-theoretical semantics -- is also intuition-free, theory-driven and task-oriented.

Hintikka writes that, "as the case is with theoretical models in general, it [the explanatory model] does not seem to be derivable from any number of observations concerning ordinary language. It has to be invented rather than discovered" (1969, 5; with, significantly, a footnote reference to Chomsky). Hintikka is thinking in the same spirit as Chomsky about the tasks of epistemic attributions. Hintikka points out that the explanatory model embodied in epistemic logic does not reproduce what

is found observationally or intuitively in ordinary discourse as surface phenomena. The latter at best point to the tasks of the “deep logic” of our epistemic-language competence, just as the surface grammars of English or Chinese point to the computational tasks of the “deep grammar” that characterizes our grammatical competence in general. In the case of epistemic discourse, not only are the surface phenomena distinct from the deep-logic tasks but they are constantly influenced by all sorts of collateral interests and pressures, such as conflicting goals, pragmatic considerations, cognitive limitations, and contextual factors.

Given all these considerations, Hintikka’s suggestion is to treat as basic the meaning of an epistemic expression captured by the explanatory model and then view its modifications and variations by the collateral factors just cited as residual meanings (Hintikka 1969, 6-7). The actual use of the expression reflects the specific relation between the basic and the residual meanings, which is the relation between what the expression (through its terms and operations) is designed to convey according to its deep logic (i.e., its core tasks) and the collateral conditions of its use.

This, quite roughly, is the line of metaepistemic analysis pursued by Hintikka in his classic *Knowledge and Belief* (1962) and many other works,

including the collection of historical essays, *Knowledge and the Known* (1974). The distinction between basic and residual meanings is best revealed in his equivalence claim that for one to know is to know that one knows. Call this the KK equivalence. Its critics, according to Hintikka, failed to see that the equivalence concerns the basic meaning of a knowledge claim (what its deep logic conveys) and not its surface variations in ordinary discourse, due to collateral interests. In its different surface manifestations the KK equivalence is bound to break down most of the time, precisely because of collateral interferences. 'Being certain' or 'being aware' or 'having enough evidence' are expressions of residual epistemic meanings that often defeat the KK equivalence for contextual and pragmatic reasons. But these are not the reasons why the KK equivalence holds fundamentally. This is why, according to Hintikka, the ordinary language analyses of epistemic terms and locutions highlight variety and diversity but fail to address their deep logic or basic meanings or core tasks. Hintikka is an essentialist realist about the deep logic of ordinary epistemic discourse whereas most ordinary-language analysts are postmodernist impressionists.

Having sketched the broader picture, we can now ask the key question: What is the deep logic or basic meaning of a knowledge claim

and hence of the KK equivalence? Recall that the question is not, epistemologically speaking, about having knowledge or instantiating it in some form. The question is about an item of epistemic discourse, in particular a knowledge (or some other epistemic) claim or description made explicitly by a speaker of a language. The answer is that if one knows something, one ipso facto knows that one knows, because the same circumstances that would justify one in saying 'I know that I know' would justify one in saying 'I know' simpliciter (Hintikka 1962, chapter 5). There are demonstrably no situations or possible worlds in which one claim would be true and the other false. The reason is that one always knows what one is thinking when saying something, such as making an epistemic claim; for not knowing it would be epistemically indefensible or inconsistent. This, then, seems to be the core task of a knowledge claim: to ascertain that the one making the claim has as good a justification as there can be (in terms of all possible eventualities) and that further doubt or criticism are beyond the point. It is the task of discussion (criticism, inquiry)-stopper (Hintikka 1962, 111; 1969, 13).

Notice that this account fits the main joints of the standard epistemological analysis. When one makes a knowledge claim, one presupposes that one has a true belief that is justified. The difference is

that in the case of epistemic discourse and its knowledge claims, the question of evidence and justification (which frustratingly eludes most epistemological analyses) is settled, as it were, by the deep-logic design of the enterprise. That is the very point of making a knowledge claim, its core task -- to indicate an end to inquiry and to the pursuit of evidence and justification.

It may appear that in the first-person case, the KK equivalence entails that mental states are transparent to self or self-intimating or introspectable. But Hintikka does not take self knowledge to be a mental state and therefore one's self knowledge claim does not say anything substantive about one's own mind -- except that it made up its mind, so to speak, to conclude an inquiry or the search for evidence, and so declares publicly. For Hintikka, the deep logic of epistemic discourse has no room for privileged access or incorrigible authority. Although the focus here has been on knowledge -- perhaps the epistemic-discourse notion most systematically investigated by Hintikka -- I expect similar conclusions to be drawn, *mutatis mutandis*, about the deep logic of other epistemic terms and attributions. Contrary to many historically and recently fashionable views, Hintikka's analysis of its deep logic or core tasks suggests that epistemic discourse in general is not about the mind,

nor about the vagaries of context and conversation. My reading of his analysis is that the basic job of epistemic discourse is to inform publicly about the range of actual and possible situations or worlds compatible with a given intentional attitude (or a sequence of attitudes) of the person discoursed about. I find this reading congenial to a larger picture I draw of naive psychology and mentalist discourse, as I explain in the next and concluding section.

5. A DEEPER WHY

Why would the deep logic of epistemic discourse work the way Hintikka proposes, at least according to my reading? To answer this question, we need to go beyond the confines of the classic top-down method of explanation in cognitive science, the T->P->M->W method mentioned at the beginning of the previous section. To see what I mean, consider a methodological distinction that is familiar in evolutionary biology. It is the distinction between proximate and ultimate explanations of biological traits, in particular competencies. A proximate explanation tells us how a competence works, according to what program, executing which tasks. But it does not explain the reason for the tasks themselves

and hence for the design of the program and the modus operandi of the mechanism running the program. It is the business of an ultimate or evolutionary explanation to identify the deeper reason. The implication, then, is that the T->P->M->W method organizes and provides an order to the proximate explanations of mental capacities but does not tell us why these capacities exist and why they evolved. The why question must be answered at a higher level of theorizing, that of the evolutionary function (E) of the tasks and programs under scrutiny. The classic method must be augmented to take the E->T->M->P->W form.

In proposing this addition, I argued in an earlier work that, unlike the more transparent tasks of our competencies for vision or grammar, the tasks of naive psychology are not obvious without a careful look at their evolutionary function (Bogdan 1997, chapters 3 and 5). When their evolutionary function is factored in, it becomes apparent and plausible that the tasks of naive psychology are to detect, represent and categorize those relational (not intrinsic) properties of other individuals, which are mentally intentional as well as behavioral, and which the naive psychologist can use causally to engage, influence or otherwise exploit in order to pursue her goals in a variety of social and communicational contexts. In other words, naive psychology is an evolved mental tool kit

that services the active goals of the naive psychologist when interacting with conspecifics.

This analysis works best in the case of simple, most visible and informative as well as causally manipulable intentional relations, such as gazing, noticing, seeing, or wanting. The categories of such relations may indeed have evolved by natural selection. To get a flavor of the analysis, consider the competence to represent gaze -- a basic pillar of primate naive psychology (Bogdan 1997, 137-138). The metaintentional category of gaze, underlying this competence, contains instructions and procedures to represent those aspects of someone's gaze that predict behavior and allow causal interventions that meet the goals of the naive psychologist. This idea can be unpacked and illustrated in the following analysis of the gaze category:

(a) eyes open --> alertness and propensity for behavior --> involvement

(b) eyes open + line of regard --> interest and its general direction --> involvement

(c) eyes open + line of regard tracked --> goal to be identified or the direction of a behavior to be initiated or something happening somewhere --> involvement

(d) eyes open + line of regard tracked + the target of the line of regard identified --> specific goal --> involvement

A human child or a chimpanzee can use these different components of the gaze category to find out about others and the situations they are in or will be in, and also to interfere with those situations or the actions of others. For example, the (b) instructions enable a naive psychologist to obstruct the line of regard of a gazer and prevent him from seeing something of interest to the naive psychologist. Apes are known to do this trick quite often.

Consider now the much more sophisticated naive psychology of human adults tracking complex, invisible and linguistically expressed propositional attitudes, such as opinions, memories, intentions, thoughts, and claims to knowledge. Epistemic discourse becomes the main avenue not only to the identity of such attitudes but also and crucially to what the attitudes inform about -- mainly people and situations, actual, past or possible -- and to the opportunities of interference, manipulation or utilization afforded by this informativeness of the attitudes. Looked at from the perspective of its initial evolutionary rationale, the naive-psychological game played with the epistemic discourse about

propositional attitudes is not that different in its basic tasks from the much primitive game played with the much simpler metaintentional categories, such as gaze. As I read it, Jaakko Hintikka's analysis portrays the basic epistemic meanings and their deep logic in the same instrumental light, as providing information on attitudes and the possible situations in which they hold -- information that can be exploited in some way or put to some use.

Department of Philosophy

Tulane University

REFERENCES

- Bartsch, K. and Wellman, H. (1995) *Children Talk about the Mind*, Oxford: Oxford University Press.
- Bogdan, R. (1997) *Interpreting Minds*, Cambridge, MA: MIT Press.
- Bogdan, R. (2000) *Minding Minds*, Cambridge: MA, MIT Press.
- Hintikka, J. (1962) *Knowledge and Belief*, Ithaca, NY: Cornell University Press.

Hintikka, J. (1969) *Models for Modalities*, Dordrecht: Reidel.

Hintikka, J. (1974) *Knowledge and the Known*, Dordrecht: Reidel.

Hintikka, J. (1975) *The Intentions of Intentionality*, Dordrecht: Reidel.

Nelson, K. (1996) *Language in Cognitive Development*, Cambridge:
Cambridge University Press.

Perner, J. (1991) *Understanding the Representational Minds*, Cambridge,
MA: MIT Press.

Published in

D.Kolak and J. Symons (eds), *Quantifiers, Questions and Quantum Physics*
Springer, Dordrecht, 2004