

REASONING ABOUT KNOWLEDGE IN PHILOSOPHY:
THE PARADIGM OF EPISTEMIC LOGIC

Jaakko Hintikka
Department of Philosophy
Florida State University
Tallahassee, FL 32306

ABSTRACT

Theories of knowledge representation and reasoning about knowledge in philosophy are considered from the vantage point of epistemic logic. This logic is primarily a logic of knowing that, and its semantics can be considered an explication of the well-known idea that "information means elimination of uncertainty". A number of other theories are discussed as further developments of epistemic logic. They include:

- (1) theory of questions and answers.
- (2) interplay of quantifiers and epistemic concepts.
- (3) representations of other kinds of knowledge than knowing that, especially those expressed by knows + an indirect wh-question and by knows + direct grammatical object.
- (4) the problem of cross-identification; the coexistence of different cross-identification methods.
- (5) the problem of logical omniscience.
- (6) informational independence in epistemic logic and its manifestations, including the de dicto - de re contrast and wh-questions with outside quantifiers.
- (7) an interrogative model of inquiry and its applications, especially the conceptualization of tacit knowledge and of range of attention.

Epistemic logic as a vehicle of knowledge representation

The main vehicle of speaking and reasoning about knowledge in philosophy has recently been epistemic logic.¹ Even though epistemic logic is not the only relevant language-game in town, it offers a useful perspective here, for certain other approaches can be thought of as improvements on epistemic logic. In its axiomatic-deductive forms, epistemic logic is normally considered a branch of modal logic, and its semantics is usually subsumed under the misleading heading of "possible-worlds semantics".² I will not attempt here a survey of the existing literature on epistemic logic.² Most of this literature is focused on syntactical (e.g., deductive and axiomatic) methods of dealing with knowledge representation and reasoning about knowledge. This is in my view a serious defect in much of the current work on epistemic logic. For typically the most interesting problems and solutions are found by considering the model-theoretical (semantical) situation. For this reason, I will not attempt here a survey of existing literature, but a review of some of the central conceptual problems arising in epistemic logic.

The basic laws of epistemic logic are in fact easily obtained on a basis of a simple semantical idea. It is that all talk about knowledge (in the sense of knowing that) presupposes a set (space) W of models (a.k.a. scenarios, worlds or situations) and a knower b who has so much information that he or she can restrict his or her attention to a subset W_1 of the space W . Since W_1 is relative not only to b but also to the scenario $w_0 \in W$ in which b 's knowledge is being considered, the obvious implementation of this intuitive idea, which of course is but a form of the old adage "information is elimination of uncertainty", is to assume that a two-place relation R is defined on W for each b . The members of W_1 are the worlds compatible with what b knows in w_0 . Then W_1 is the set of all scenarios to which w_0 bears this relation. The relation will be called an epistemic alternativeness relation, and the members of W_1 are called the epistemic b -alternatives to w_0 . Thus " b knows that S " is true in w_0 iff S is true in all the epistemic b -alternatives to w_0 .

Each such alternativeness relation must be assumed to be reflexive (what is known is true) and transitive. (If b is in a position to rule out all scenarios in $W - W_1$, b is ipso facto in a position to rule out the claim that he or

¹The idea of epistemic logic goes back at least to G.H. von Wright, An Essay in Modal Logic, North-Holland, Amsterdam, 1951. The first book-length treatment was my Knowledge and Belief: An Introduction to the Logic of the Two Notions, Cornell U.P., Ithaca, 1962.

²For a partial survey of earlier work, see Wolfgang Lenzen, Recent Work in Epistemic Logic (Acta Philosophica Fennica vol. 30, no. 1) Societas Philosophica Fennica, Helsinki, 1978.

she is not in such a position.) These definitions and stipulations (combined with a suitable semantics for the usual quantification theory) specify the semantics of a system of epistemic logic, and hence its deductive-axiomatic treatment, subject to the qualifications to be discussed below.

The resulting logic turns out not to be devoid of interest. Its propositional part (restricted to one knower) is the logic of the topological closure operation. Hence epistemic logic is related to the logic of topology. Its laws are in effect those of intuitionistic logic. There are also close relations³ between the semantics of epistemic logic and the technique of forcing. However, in order to reach this connection, the semantics of negation and conditional have to be modified somewhat.

Many of the further developments in epistemic logic can be thought of as solutions to problems concerning the epistemic logic so far set up. One of the first problems is to represent the other kinds of knowledge for instance, the kinds of knowledge expressed by knows + indirect wh-questions and by knows + a direct grammatical object, by starting from the knows that construction.⁴ This basic construction (b knows that S) in my notation " { b } K S".⁴

Two comments are in order here.

(a) The propositional alternatives I have called "scenarios" or "models" can be states of affairs, situations, courses of events, or entire world histories. The last of these applications is highlighted by philosophers' misleading term "possible-worlds semantics". This term is misleading, because applications to entire universes are scarcely found outside philosophers' speculations.⁵ The primary intended applications are to scenarios covering relatively small pieces of space-time. Thus the label "situation semantics",⁶ which has recently been applied to a study of additional relations between what I have called scenarios, does not mark any sharp contrast to rightly understood possible-worlds semantics.

³See here, e.g., Melvin Fitting, Intuitionistic Logic, Model Theory, and Forcing, North-Holland, Amsterdam, 1969; Kenneth A. Bowen, Model Theory for Modal Logic, D. Reidel, Dordrecht, 1979. These treatises are not addressed to the specific problems of epistemic logic, however.

⁴In the earlier literature, the knower used to be indicated by a subscript. This is misleading, however, for the term referring to the knower is not within the scope of the epistemic operator.

⁵Some philosophers have tried to find a difference in principle between the two kinds of applications. It is nevertheless clearer in epistemic logic than in some of the parallel theories that the intended applications have always been to "small worlds", to use L.J. Savage's phrase.

⁶See Jon Barwise and John Perry, Situations and Attitudes, MIT Press, Cambridge, Mass., 1983.

(b) The most important application of epistemic logic is to the theory of questions and answers. No separate treatment is needed, however, for a direct question like

(1.1) Who's living here?

can be construed as a request of information which might as well be expressed by

(1.2) Bring it about that I know who is living here.

Here the subordinate clause is an indirect question with "knows" as its main verb. I have called it the desideratum of the question (1.1). It falls within the purview of epistemic logic (see sec. 2 below). And obviously the logical study of direct questions like (1.1) reduces largely to the study of their desiderata.

The first and foremost problem is the theory of questions and answers concerns the relation of a question to its (conclusive) answers. When does a reply, say "d" to a wh-question like (1.1) do its job? Obviously when it makes the desideratum

(1.3) I know who lives here

true. But what does the reply "d" in fact accomplish? Obviously, the truth of

(1.4) I know that d lives here.

Hence the problem of answerhood is the question as to when (1.4) implies (1.3). Now the logical forms of (1.3) and (1.4) are, fairly obviously (but see sec. 4 below),

(1.5) $(\exists x) \{1\} K (x \text{ lives here})$

and

(1.6) $\{1\} K (d \text{ lives here})$

⁷See here Jaakko Hintikka, The Semantics of Questions and the Questions of Semantics (Acta Philosophica Fennica, vol. 28, no. 4), Societas Philosophica Fennica, Helsinki, 1976.

Hence the operative problem is when (1.6) implies (1.5). This is a question concerning the interplay of quantifiers and epistemic operators. This interplay will be discussed in section 2 below.

Quantifiers in epistemic logic. Knowing + indirect wh-questions

The first conceptual problem I shall analyze is the representation of other kinds of knowledge than knowing that. They include the kind of knowledge expressed linguistically by the constructions knows + indirect wh-question and knows + a direct grammatical object. Here philosophers' preoccupation with the surface phenomena of ordinary usage has seriously hampered their theorizing. In fact, the right treatment is nevertheless not hard to find. It can be presented as a succession of steps.

(i) In order to use quantifiers in context which, like epistemic contexts involve a multitude of scenarios, it must be assumed that criteria of identity for individuals across worlds have been given. In order to have a vivid vocabulary, I shall speak of the imaginary lines connecting the counterparts of the same individual in different model or scenarios as "world lines".

Once a warp of world lines connecting the members of W is given (for each relevant knower), truth conditions for quantified sentences in a (first order) epistemic language. Such truth-conditions solve all the conceptual problems Quine and others have raised or, rather, transform them into problems concerning the way world lines are drawn.

(ii) It cannot be assumed that the same individuals exist in all models. Then the basic laws of quantification theory have to be revised by changing some of the instantiation rules. For instance, the law of universal instantiation might be changed so as to read:

(UI) If "x" occurs in $S[x]$ outside the scopes of all epistemic operators,

$$(2.1) \quad \frac{(\forall x)S[x] \quad \varepsilon \quad (\exists y)(z=y)}{S[z]}$$

In other words, when we speak of z as a member of the actual world, we have to assume that it exists in that world in order for it to be a bona fide value of quantifiers pertaining inter alia to the actual world.¹⁰

⁸The question here is under what conditions existential generalization is valid in epistemic logic. The conditions are of course the same as the conditions on valid universal instantiation dealt with in sec. 2, part ii, below.

⁹Cf. here chapter 1 of my book, The Intentions of Intentionality, D. Reidel, Dordrecht, 1974.

¹⁰I am assuming that a distinction is made between those name-like free singular terms which pick out the same individual from different worlds and

(iii) The obvious formal counterpart of a knows + an indirect wh-question, e.g., of

(2.2) \underline{b} knows who (say, x) is such that $S[x]$

is

(2.3) $(\exists x) \{ \underline{b} \} K S[x]$.

This paraphrase amounts to saying that \underline{b} knows who (x) is such that $S[x]$ is there to be a world line which in all of \underline{b} 's knowledge worlds picks out an individual x satisfying in that world the condition $S[x]$.

The best proof of the aptness of this rational reconstruction of "knowing + wh-construction" sentences is that it leads into an elegant and powerful analysis of the relation of a (direct) wh-question to its (conclusive) answers.

(iv) In order for this idea to work, we nevertheless must allow world lines to break down in a more radical sense than the failure of an individual to exist in a given "world". We must allow a world line of x to break down in a world w_1 in the strong₁ sense that it does not even make sense to ask whether x exists in w_1 or not.¹¹ For otherwise (if all world lines could be extended ad libitum) everybody would know the identity of every individual (under some guise or other, so to speak), on the basis of the paraphrase of constructions of the form knowing + indirect wh-question agreed on in (iii).

In other words, the natural model-theoretical counterpart of \underline{b} 's knowing the identity of x is that some world line passing through x (considered as a member of the actual world) spans all of \underline{b} 's knowledge worlds.

How is the well-definedness of x in a world w_1 (i.e., the extendibility of the world line of x to w_1) to be expressed in a formal language? The obvious candidate is the truth of " $x=x$ " in w_1 . This simple idea yields the first fully satisfactory treatment of quantification in epistemic contexts. This treatment has not been worked out in the literature, but the main points are nevertheless clear. Since the truth of $(\exists z)(x=z)$ implies that of $x=x$, no changes are needed in the quantifier rules. Instead, we need a three-value logic with different kinds of negations, forced on us by the idea that if x is undefined in a world w_1 , any atomic sentence containing " x " does not have either of the two usual truth-values "true" and "false" in w_1 .

(v) Essentially the same treatment can be extended to higher-order logic. (Such a treatment is needed, among other purposes, for applications of epistemic logic to the theory of questions and answers.) There is one

those that might refer to different individuals in different worlds. Here " z " is assumed to be of the former kind.

¹¹This matter will be dealt with in a greater detail in a projected monograph of mine.

difference, however. In the case of higher-order entities, existence is not needed as a condition of being a value of a quantified variable. (It is not clear what existence might mean here.) But well-definedness is still required. Hence the counterpart to (UI) for, e.g., one-place second-order quantifier is

$$(UI)' \quad \frac{(\forall X)S[X] \quad \varepsilon \quad (Y = Y)}{S[Y]}$$

where $(Y = Y)$ is to be taken to same as

$$(2.4) \quad (\forall x)(Y(x) \leftrightarrow Y(x))$$

(vi) There is another crucially important feature of the conceptual situation here which has been obscured by the surface phenomena of ordinary language and therefore neglected by philosophers. This is the fact that in certain situations there are two systems of world lines in operation.¹² A knower's (say, \underline{b} 's) cognitive relations to his or her environment (including past situations in which \underline{b} was directly involved) span a framework which can be used for the purpose of drawing world lines. Such a world line connects such (scenario-bound) individuals as play the same role in these first-hand cognitive relations to \underline{b} from \underline{b} 's perspective. The simplest example is visual perception (visual knowledge). There the relevant framework is \underline{b} 's visual space and the world lines perspectively drawn connect the individuals occupying the same slot in \underline{b} 's visual geometry. (If \underline{b} does not see who or what they are, they are not the same absolutely or descriptively identified entities). In other words, \underline{b} 's perspectively identified objects are his or her visual objects. This can be extended to other kinds of knowledge in a fairly straightforward way.

Because of the presence of two systems of world lines, we must have two pairs of quantifiers corresponding to (relying on) them. Success in the perspectival cross-identification will then be expressed in the same way as with the other (public, descriptive) mode of identification, but with a different kind of quantifiers, say $(\underline{a}x)$ and $(\underline{e}x)$, instead of $(\forall x)$ and $(\exists x)$. Then

$$(2.5) \quad (\underline{e}x) \{ \underline{b} \} K(\underline{d} = x)$$

will say, assuming a case of visual knowledge situation, that \underline{b} can find a niche for \underline{d} among \underline{b} 's visual objects. In other words, (2.5) says that \underline{b} sees (and recognizes) \underline{d} . More generally speaking, (2.5) says that \underline{b} is acquainted with \underline{d} , i.e., knows \underline{d} . Here we thus have an analysis of the knows + a direct grammatical object construction. This construction can be analyzed in terms of knowing that plus perspectival (contextual) quantifiers.

In intuitive terms, the contrast between the two different cross-identification methods has variously been described as putting a name to

¹²See here chapters 3-4 of my book The Intentions of Intentionality, D. Reidel, Dordrecht, 1974.

a face vs. putting a face to a name, answering what-questions vs. answering a where-question, etc. These explanations are only partial ones, however, and do not bring out fully the underlying model-theoretical situation. Hence they (like the terms "visual object" and "object of acquaintance") have to be taken with a pinch of salt.

Even though most philosophical logicians have chosen to disregard this duality of quantifiers and modes of cross-identification, it is one of the most important phenomena in the field of knowledge representation, important both epistemologically and psychologically. In epistemology, the distinction has in effect figured as Bertrand Russell's¹³ contrast between knowledge by description and knowledge by acquaintance. In psychology, the same distinction is manifested as the contrast between what are called semantical and¹⁴ episodic memory¹⁵ as well as a distinction between two kinds of visual systems.

The problem of "logical omniscience"

The model-theoretical treatment of epistemic logic so far outlined leads to a paradoxical result. Suppose that $\vdash (S_1 \supset S_2)$, i.e., that all models of S_1 are models of S_2 . Then all the epistemic alternatives in which S_1 are true are also alternatives in which S_2 is true. From this it follows that the following will be true for any b and in any scenario:

$$(3.1) \quad \{b\}K S_1 \supset \{b\}K S_2$$

In brief, everybody always knows all the logical consequences of what he or she knows. This is obviously an unacceptable result, and in certain quarters it is still considered a¹⁶ sufficient reason for rejecting a model-theoretic analysis of epistemic concepts.

This follows only if the paradox, known naturally as the paradox of logical omniscience, is unavoidable. And it has been known for quite a while that it is not. In fact, we have have one of the several unmistakable but unheralded triumphs of epistemic logic. There are in fact two equivalent ways of

¹³Bertrand Russell, "Knowledge by Acquaintance and Knowledge by Description", in *Mysticism and Logic*, George Allen & Unwin, London, 1917; chapter 5 of *The Problems of Philosophy*, Home University Library, London, 1912; and cf. Jaakko Hintikka, *Knowledge by Acquaintance - Individuation by Acquaintance*, in D.F. Pears, editor, *Bertrand Russell (Modern Studies in Philosophy)*, Doubleday, Garden City, N.J., 1972, pp. 52-79.

¹⁴See Endel Tulving. *Elements of Episodic Memory*, Clarendon Press, Oxford, 1983.

¹⁵See Lucia Vaina, *From Vision to Cognition: A Computational Theory of Higher-Level Visual Functions*, D. Reidel, Dordrecht, 1986.

¹⁶Cf., e.g., Noam Chomsky, *The Generative Enterprise*, Foris, Dordrecht, 1982. From what is reported in the rest of this section, this objection against possible-worlds analysis of knowledge was effectively disposed of more than ten years ago.

delineating the subclass of logical consequences $\vdash (S_1 \supset S_2)$ for which (3.1) holds.

(i) This set can be defined by putting syntactical restrictions on the deductive argument which leads from S_1 to S_2 . This argument can of course be of many different kinds. It turns out, however, that for all the half-way natural ones, the same heuristic idea works and gives the same result. In an easily appreciated sense, the number of free individual symbols together with the number of layers of quantifiers determine how many individuals are considered in a given sentence S (or in a given argument). The natural restriction on the argument from S_1 to S_2 now is that this parameter should never be larger at any stage of the argument than it is in S_1 or S_2 .

Even though this basic idea is thus easy to understand and to implement, no simple axiomatic-deductive system codifying it has been presented in the literature.

This way of defining knowledge-preserving logical inferences is connected, via the idea of so many individuals considered in their relation to each other in an argument with a wealth of traditional philosophical issues in the philosophy of logic and mathematics.¹⁸ The same idea also promises connections with the psychology of deductive reasoning.¹⁹

(ii) This syntactical (deductive-axiomatic) restriction on "logical omniscience" yield the same result as a different and apparently completely unlike line of thought. This line begins with an interesting generalization of the concept of model (world, scenario). Unlike its rivals as a candidate for the role of a logically₂₀ nonstandard world (semantical bases of so-called paraconsistent logics),²⁰ this generalization is completely realistic. Indeed, this generalization is but a variant of the notion₂₁ of urn model in probability theory, and is referred to by the same term.²¹ In an obvious sense, nested quantifiers can be thought of as representing successive "draws" of individuals from an "urn", i.e. from the domain of a model, or (perhaps a little more

17See Jaakko Hintikka, "Knowledge, Belief, and Logical Consequence" in Hintikka, The Intentions of Intentionality, D. Reidel, Dordrecht, 1975.

18See Jaakko Hintikka, Logic, Language-Games, and Information, Clarendon Press, Oxford, 1973.

19See Jaakko Hintikka, "Mental Models, Semantical Games, and Varieties of Intelligence", in Lucia Vaina, ed., Varieties of Intelligence, D. Reidel, Dordrecht, 1986.

20The so-called paraconsistent logics have never been any realistic model-theoretical and pragmatic interpretation, and hence have in their present form little interest. Cf. here Nicholas Rescher and Robert Brandom The Logic of Inconsistency, Basil Blackwell, Oxford, 1979.

21See Veikko Rantala, "Urn Models: A New Kind of Non-Standard Model for First-Order Logic", Journal of Philosophical Logic, vol. 4 (1975), pp. 455-474, reprinted in Esa Saarinen, Game-Theoretical Semantics, D. Reidel, Dordrecht, 1979.

vividly) as a series successful searches of individuals from the model. The concept of urn model is obtained by letting the set of available individuals change between successive "draws". (The world, in other words, is run by a malicious demon who can restrict the set of available individuals in tandem with our examination of the world (via successive searches or "draws").

Actually, not all and sundry urn models are natural candidates for the role of epistemically possible but (classically speaking) logically impossible worlds which are the model-theoretical codification of the failure of logical omniscience. For that role, only those urn models (changing models) are acceptable which vary so subtly that they cannot be told apart from the invariant (classical) models by means of sequences of draws as long as those involved in a given sentence.²² It turns out that the conditionals $(S_1 \supset S_2)$ which are true in all such "almost invariant" urn models (at the length of sequence of draws envisaged in the conditional) are precisely the same as those for which the step from $\vdash (S_1 \supset S_2)$ to (3.1) is authorized by the syntactical restriction.

Epistemic logic and informational independence

One of the most characteristic features of epistemic logic has barely been mentioned in the existing literature.²³ In order to see what is involved, let us consider the familiar distinction between

(4.1) \underline{b} knows that there is an individual x such that $S[x]$

and

(4.2) \underline{b} knows of some individual x that $S[x]$.

Usually, it is said that these two are to be represented in the language of epistemic logic by

(4.3) $\{ \underline{b} \} K (\exists x) S[x]$

and

(4.4) $(\exists x) \{ \underline{b} \} K S[x]$,

respectively. Here the latter has roughly the force

²²See Jaakko Hintikka, "Impossible Possible Worlds Vindicated", Journal of Philosophical Logic vol. 4 (1975), pp. 475-484, reprinted ibid.

²³The first time this interesting phenomenon was pointed out in the literature is in Lauri Carlson and Alice ter Meulen, "Informational Independence in Intensional Context", in Esa Saarinen et al., eds., Essays in Honour of Jaakko Hintikka, D. Reidel, Dordrecht, 1979, pp. 61-72.

(4.5) \underline{b} knows who (what), say x , is such that $S[x]$.

Hence the contrast in question is roughly that between "knowing that there is" and "knowing who or what".

I am not questioning the status of (4.4) as being logically equivalent with (4.2), i.e., as a possible translation of (4.2). What I am asking is how this translation comes about, i.e., what the mechanism is that leads us from (4.2) to (4.4). It is usually assumed, as the rendering (4.4) of (4.2) shows, that the mechanism in question is the relative order (relative scopes) of K and $(\exists x)$. The linguistic evidence for this idea is unconvincing, however. It is much more natural to assume that in epistemic context like (4.2), the choice of the individual we are talking about is independent of epistemic considerations, i.e., that the quantifier somehow ranges over just the actually existing individuals. This independence can be captured by the two-dimensional expression

(4.6) $(\exists x) \begin{matrix} \diagdown \\ S[x] \\ \diagup \\ \{ \underline{b} \} K \end{matrix}$

Even though the meaning of (4.6) is intuitively obvious, further explanations are needed here to incorporate expressions like (4.6) in our formal language. In order to spell out the semantics of expressions like (4.6) we must combine possible-worlds analysis of epistemic concepts (as sketched ever so briefly in section 1 above) with what has been called game-theoretical semantics (GTS).²⁴ Very briefly, in GTS the truth of a sentence S in a model M is explicated as the existence of a winning strategy in a certain verification game (semantical game $G(S)$ played with S on M for one of the players, called Myself (or the Verifier) against an opponent called Nature (or the Falsifier). Most of the rules of these games can be anticipated on the basis of the verification idea. The following are cases in point:

(G.E) If the game has reached the sentence $(\exists x) S[x]$

and M , Myself chooses an individual from the domain $do(M)$ of M . Then the game is continued with respect to $S[b]$ and M .

(G.U) Similarly except that Nature chooses b .

(G.v) $G(S_1 \vee S_2)$ (played on M) begins with Myself's choice of S_i ($i = 1$ or 2). The rest of the game is $G(S_i)$ (played on the same model M).

²⁴See here Esa Saarinen, ed., Game-Theoretical Semantics, D. Reidel, Dordrecht, 1979, and Jaakko Hintikka and Jack Kulas, The Game of Language, D. Reidel, Dordrecht, 1983.

(G. ε) Similarly, except that Nature chooses S_i .

(G. \sim) $G(\sim S)$ is like $G(S)$ except that the roles of the two players are reversed.

(G.K) If the game has reached the sentence $\{b\} K S$ and the model (world) M_0 , Nature chooses an epistemic b -alternative M_1 to M_0 . The game is continued with respect to S and M_1 .

In terms of GTS, the semantics of branching formulas like (4.6) can be dealt with explicitly.²⁵ What they instantiate is the well-known game-theoretical phenomenon of informational independence. In (4.6), the moves connected with " $(\exists x)$ " and " $\{b\} K$ " are each made without knowledge of the other move. More generally, each move is associated with an information set including those other moves which are known to the player making the move. Hence the operator structure of a sentence need not always be partly even partially ordered.

We can sometimes linearize the branching notation used in (4.6) by attaching to each quantifier an indication which shows which of the earlier epistemic operators (if any) it is informationally independent of. Thus (4.6) can be written, in a self-explanatory notation, as

(4.7) $\{b\} K (\exists x / \{b\} K) S[x]$

which is of course logically equivalent with (4.6) and (4.4).

It is a most important general fact about the logic of epistemic concepts that when they mix with quantifiers, these quantifiers frequently have to be taken to be independent of some of the epistemic operators present. For instance, the so-called de re reading of quantifiers is in reality precisely the reading obtained by taking the quantifier in question to be informationally independent of an epistemic (or other intentional) operator.

GTS shows that (and how) other primitives of one's language, not just the quantifiers, can be independent of epistemic operators. For instance, an atomic predicate $A(x)$ or a proper name a may be evaluated in M_0 independently of an epistemic operator; say " $\{b\} K$ ". This may be indicated by writing $A(x / \{b\} K)$ or $a / \{b\} K$. Since the epistemic operator governs the choice of an alternative world M_1 , this means evaluating these primitive in the pre-epistemic-move-model M_0 . And since in the winning strategy they clearly have to be evaluated so as to assign to them their actual references, such

²⁵For branching quantifier structures, there exists a growing body of studies. For references, see the bibliography of Jaakko Hintikka and Jack Kulas, The Game of Language, D. Reidel, Dordrecht, 1983. Independences between other kinds of concepts have scarcely been studied, except for the papers referred to here.

expressions as $A(x/\{b\}K)$ and $a/\{b\}K$ in effect pick out the actual references in M_0 .

As the equivalence of (4.6) and (4.4) illustrates, in the simplest cases sentences containing informationally independent quantifiers have non-independent equivalents. But even in such cases, a notation which spells out the independence can bring out the intended logical form of our epistemic statements more clearly than the dependent (linear) notation. For instance, natural-language statements of the form

(4.8) \underline{b} knows who (say, x) is such that $A(x)$

(where A is an atomic predicate) have normally two readings, which in the linear traditional notation are expressed as follows:

(4.9) $(\exists x) \{b\}K A(x)$

(4.10) $(\forall x)(A(x) \supset (\exists z)(y=z \ \& \ \{b\}K A(x)))$.

The parallelism of the two is not obvious in (4.9)-(4.10) but is brought out much more clearly in an independence-friendly notation as follows:

(4.11) $\{b\}K (\exists x/\{b\}K) (A(x/\{b\}K) \ \& \ A(x))$

(4.12) $\{b\}K (\forall x/\{b\}K)(A(x/\{b\}K) \supset A(x))$.

However, in other cases the independence notation can be indispensable.²⁶ For instance,

(4.13) \underline{b} knows whom everybody adores

cannot be expressed (on the reading according to which different persons may have different idols) by

(4.14) $(\forall x)(\exists y)(\exists z) (x = z \ \& \ \{b\}K (z \text{ admires } y))$,

for (4.14) implies (as you can easily see) that \underline{b} knows the identity of each person and of his or her admirer. In fact, the force of (4.13) can only be expressed (unless we resort to higher-order quantifiers) by something like

(4.15) $\{b\}K(\forall x)(\exists y/\{b\}K) (x \text{ admires } y)$.

Indeed, (4.15) (unlike (4.14)) is logically equivalent with

²⁶See here Jaakko Hintikka, "Questions with Outside Quantifiers", in R. Schneider, K. Tuite and R. Chametzky, eds., Papers from the Parasession on Nondeclaratives, Chicago Linguistics Society, Chicago, 1982, pp. 83-92.

$$(4.16) \quad (\exists f) \{ \underline{b} \} K (\forall x) (x \text{ admires } f(x))$$

which neatly brings out the obviously intended force of (4.13). For what (4.13) says is, obviously, that \underline{b} knows how to find, for a given x , whom x admires, i.e., knows a function \underline{f} which takes in from any person to someone she or he admires. And this is precisely what (4.16) says.

Notice that the informational dependence of the different operators $\{ \underline{b} \} K$, $(\forall x)$, and $(\exists y)$ in (4.15) is not even partially ordered, but exhibit a loop structure:

$$(4.17) \quad \begin{array}{ccc} & \{ \underline{b} \} K & \\ \nearrow & & \searrow \\ (\exists y) & \leftarrow & (\forall x) \end{array} \quad (x \text{ admires } y)$$

This is logically equivalent with (4.15). A semantical game connected with (4.17) is hard to implement in the usual move-by-move form. It is easy to "play" in what game theorists call the normal form of a game: Both players choose a strategy, which jointly determine the course of the game, including its outcome. This observation is connected with the equivalence of (4.16) with (4.17). It has some general interest as an illustration of what different kinds of games (in the sense of game theory) can be like.

Thus the independence notation sometimes lends added power to epistemic logic. How much more? It is known that the force of quantification theory with branching quantifiers²⁷ is extremely strong, coming close to that of the entire second-order logic. Hence no complete axiomatization of quantified epistemic logic with unlimited independence is possible. However, for a variety of simpler cases, for instance, when all the cases of independence pertain to ignorance of a move connected with a single unnegated epistemic operator $\{ \underline{b} \} K$, an explicit formal treatment may very well be possible, even though it has not been presented in the literature.

Knowledge acquisition by questioning

So far, I have been concerned with problems of knowledge representation and only indirectly with reasoning about knowledge. Of course, the representation problem has to be solved before problems about reasoning are tackled. But what is represented is already acquired and already available knowledge, whereas much of the actual reasoning about knowledge is also concerned with the step-by-step processes of knowledge acquisition. For instance, in the well-known puzzle variously known as the case of cheating

²⁷See here Jaakko Hintikka, "Quantifiers vs. Quantification Theory", Linguistic Inquiry, vol. 5 (1974), pp. 153-177, reprinted in Esa Saarinen, editor, Game-Theoretical Semantics, D. Reidel, Dordrecht, 1979.

husbands or of the wise men,²⁸ the reasoning of the participants depends essentially on their knowing what the others knew or did not know at the preceding stage of their synchronized reasoning processes.

One way of modelling knowledge acquisition is to conceptualize it as a series of questions a reasoner, here termed the Inquirer, addresses to a source of information, to be called the Oracle (in some applications, more naturally called Nature).²⁹ The answers, when available, may be used by the Inquirer as premises for the purpose of deriving a given conclusion C (or, in an alternative version of the model) for the purpose of answering the question "C or not-C"? In this process, steps of deduction may alternate with each other, and the Inquirer may have a fixed initial premise T (called the theoretical premise) available for the purpose.³⁰ The deductive rules to be used are restricted to those that satisfy the subformula principle. Before a question may be asked by the Inquirer, its presupposition must have been established.³¹

In this way we obtain what I have called the interrogative game or the interrogative model of inquiry. Since the logic of questions and answers is in effect (as was pointed in section 1 above, part (b)) a branch of epistemic logic, the interrogative model can also be thought of as another outgrowth of epistemic logic. One of the main advantages of the interrogative model is that

28See, e.g., Danny Dolev, Joseph Y. Halpern and Yoram Moses, "Cheating Husbands and Other Stories: A Case Study of Knowledge, Action and Communication", preprint, 1985.

29The model sketched here has been studied in a number of papers of mine. See, e.g., Jaakko Hintikka and Merrill B. Hintikka, "Sherlock Holmes Encounters Modern Logic: Towards a Theory of Information-Seeking by Questioning", in E.M. Barth and J.L. Martens, Argumentation: Approaches to Theory Formation, John Benjamins, Amsterdam, 1982, pp. 55-76; "The Logic of Science as a Model-Oriented Logic", in P.D. Asquith and P. Kitcher, eds., PSA 1984, vol. 1, Philosophy of Science of Association, East Lansing, MI, 1984, pp. 177-185.

30As a book-keeping device we can use a Beth-type semantical tableau. (For them, see W.W. Beth, "Semantic Entailment and Formal Derivability", Mededelingen van de Koninklijke Nederlandse Akademie van Wetenschappen, Afd. Letterkunde, N.R. vol. 18, no. 13, Amsterdam, 1955, pp. 309-342.) Then we can use all the usual terminology of the tableau method, and the deductive "moves" will be simply tableau-building rules. (We shall minimize movements between the left and the right column, however, and restrict the rules to those in keeping with the subformula principle.) Each application of the game rules is then relative to a given stage of some one subtableau. As is well known, the tableau method is simply the mirror image of a Gentzen-type sequent calculus. The only novelty here is that Nature's answers are entered into the left column of a subtableau as additional premises.

31For the concept of presupposition presupposed here, see Jaakko Hintikka The Semantics of Questions and the Questions of Semantics (Acta Philosophica Fennica, vol. 28, no. 4), Societas Philosophica Fennica, Helsinki, 1976.

it enables us to discuss cognitive strategies and not only static cognitive situations. This possibility can be realized in many different ways. The Oracle can be literally nature, and nature's answers can then be scientific experiments. In a simpler case, the source of information is one's environment, and the answers the Inquirer can hope to obtain are perceptual observations. But in other cases, the available answers can be items of information stored in the database of a computer, which will then play the role of the Oracle. In still other applications, the computer is the Inquirer's own brain, and the totality of available answers define's the Inquirer's tacit knowledge. The most natural application is undoubtedly one in which the answerer (the Oracle) is another human being, with whom the Inquirer is engaged in a dialogue. We can also allow for the interrogation "game" to be a symmetrical n-person game in which at each interrogative move each player can address a question to each of the other players. Some of the most intriguing types of reasoning about knowledge can be dealt with by means of such games, for instance, the "case of the cheating husbands". Such applications and extensions can be thought of as belonging to the logical theory of dialogues (discourse).³²

I shall not discuss the details of any of these applications detail. It is in order, however, to locate some of the crucial parameters which play a role in these different interrogative games and distinguish them from each other.

(i) In different applications, different kinds of questions are answered by the Oracle. The most clear-cut restraints here are those that depend on the logical complexity, especially the quantifier prefix structure of the available answers. For instance, sense-perception can only answer yer-or-no questions concerning particular matters of fact of one's environment. For a logician, these are yes-or-no questions concerning atomic sentences.

(b) In contrast, controlled experiments can yield answers which codify functional dependencies. Such answers³³ must have an AE quantifier prefix (i.e. a prefix of the form $(\forall x)(\exists y)$).

(c) Again the information stored in the memory of a human being or of a computer can logically speaking be of arbitrarily high complexity.

There are of course but special cases of a long spectrum of different kinds of interrogative procedures, distinguished from each other by the quantificational complexity of available answers. This spectrum ranges from

³²An excellent example of what can be done in this direction is Lauri Carlson, *Dialogue Games*, D. Reidel, Dordrecht, 1982.

³³This observation has important consequences for the contemporary philosophy of science, where it has generally been assumed that only questions concerning the truth or falsity of atomic sentences are answerable by Nature. In reality, the logic of experimental inquiry is an AE logic, not the logic of the atomistic case.

case (a) via the different AEA... prefixes to the unlimited case (c). This hierarchy turns out to be highly important for many purposes, especially in the philosophy of science.

There of course are normally other kinds of limitations on available answers. In all these cases, it obviously makes a difference to the strategy selection of the Inquirer what (partial or total) knowledge he or she has of the limitations on available answers. This is illustrated by the knowledge which the user of a database may have as to what information is or is not stored in it.

(ii) In many applications of the interrogative model and indeed in many applications of epistemic logic, it makes a crucial difference what kind of knowledge we are dealing with. For instance, tacit knowledge must fairly obviously be modelled by a sub-oracle. The list of propositions stored in the "memory" of this sub-oracle defines the extent of the Inquirer's tacit knowledge by delimiting the set of available answers this sub-oracle can answer. (I am speaking of sub-oracle here because in realistic uses of my model the Inquirer can of course consult other sources of information than her or his tacit knowledge.)

At the other extreme there is the completely activated knowledge the Inquirer has. This is naturally modelled by the set of those sentences which have been put forward by the Inquirer as outcomes of interrogative or deductive moves (or which have the status of explicit initial premises of the interrogative process). This might be called the Inquirer's active knowledge.

Active knowledge, unlike tacit knowledge, is relative to a stage of the interrogative game. Dealing with such vital for reasoning about knowledge, for it is a speaker's active knowledge that he or she is aware of and can report to others.

But neither tacit knowledge nor active knowledge obeys the laws of epistemic logic. For instance, neither is closed with respect to logical consequence, not even if relations of logical consequence are restrained as indicated in section 3 above. In order to be able to develop satisfactory ways of reasoning about knowledge, we consider other kinds of knowledge. Among the most important ones there are the following:

The Inquirer's potential knowledge consists of all the conclusions C the Inquirer can establish by means of the interrogative process.

The Inquirer's virtual knowledge consists of all the conclusions C the Inquirer can establish by means of the interrogative process without introducing new "auxiliary" individuals into the argument in the sense of sec. 2 (i).

By limiting the Inquirer's moves to deductive ones we can similarly define potential deductive knowledge and virtual deductive knowledge.

(iii) The purely logical properties of the interrogative games are also of considerable interest. Let us denote the interrogative derivability of C from T in a model M by $T \Vdash_M C$. This relation depends of course on whatever restrictions there may be on available answers. It turns out that this relation depends also on the set RA of available tautological premises of the form

$$(5.1) \quad (S_i \vee \sim S_i)$$

where $i \in ra$. The reason is that in interrogative processes one cannot always restrict one's methods to those satisfying the subformula principle.³⁴ In other words, metalogical results analogous to Gentzen's first Hauptsatz, which implies the eliminability of premises of the form (5.1), of the cut rule, of unrestricted modus ponens, etc.

There is a sense in which the notion of interrogative derivability is between the relations of logical consequence $T \vdash C$ and the truth of C in M , i.e., $M \vDash C$. For if no questions are allowed, $T \Vdash_M C$ obviously reduces to $T \vdash C$. On the other hand, if no restrictions are imposed on available answers or on available tautological premises RA , it can be shown that $\emptyset \Vdash_M C$ iff $M \vDash C$ (\emptyset = the empty set).

The set RA has an intuitive interpretation which is worth noting here. What the set RA codifies is essentially the totality of yes-or-no questions which the Inquirer is prepared to ask (independently of the initial premise T). A restriction on RA is therefore very much like a restriction on the Inquirer's range of attention, primarily in the sense of a restriction on the range of questions the Inquirer is prepared to raise, secondarily in the sense of a restriction on the items of tacit knowledge the Inquirer can activate. This is because the activation of such knowledge can only happen by means of questions whose presuppositions have to be available to the Inquirer. Thus the concept of range of attention is not purely subjective and psychological but has an objective logical and epistemological counterpart.

This is but an example of the many possibilities of analyzing - and synthesizing - interesting epistemic concepts by means of the interrogative model. Most of the work in utilizing these possibilities still remains to be done.

Acknowledgement: The research reported here was made possible by NSF Grant # IST-8310936 (Information Science and Technology).

³⁴The notions of subformula principle, cut elimination, Gentzen's Hauptsatz, etc. are explained in any decent introduction to proof theory. For Gentzen's classical papers, see M.E.Szabo, ed., The Collected Papers of Gerhard Gentzen, North-Holland, Amsterdam, 1969.