The penultimate chapter, *Real machines*, is the major exposition on AI techniques and programs found in this book. It is here that heuristic search is discussed and classic programs such as SHRDLU and GPS are described. It is here that a sampling of AI material and its flavor as research is presented. Some of the material here is repeated without real analysis. For example, the author repeats the standard textbook mistake on the size of the chess space. On page 178, he states that 10^{120} is the size of this space, and uses this to suggest that no computer will ever play perfect chess. Actually, an estimate of 10^{40} is more realistic. If one considers that no chess board can have more than sixteen pieces of each color and there are many configurations that are illegal or equivalent, then the state space is reduced considerably. It is likely that chess is in some fundamental way related to NP-hard problems and as such is probably not amenable to polynomial time solution.

There is an appropriate focusing on the problems of programming common sense. As is pointed out, relatively limited progress has been made on this problem and the related "frame" problem. These limits become the lead in to the final chapter and its summing up. This last chapter is titled *Real people.* It emphasizes the author's belief that feeling and sensation are important aspects of understanding. In this regard, he differs with those in the AI community who believe these features of understanding are also readily abstracted symbolically. He also recalls Herb Simon's famous 1957 predictions of success for AI programs in four major areas by 1967—namely the computer as world chess champion, the computer as important mathematician, the computer as composer of critically acclaimed music, and finally, that most psychological theory will be in the form of programs. He rightly characterizes these predictions as overeager, but does not view the failure to accomplish them as demonstrating the incorrectness of the GOFAI hypotheses. Throughout this book his account is a balanced and judicious one. One might even suggest tame. His final pronouncement is "the hard questions remain open" (p. 254).

This is a useful addition to the joint inquiry of philosophers and AI researchers into the adequacy of programming theories of the mind. It is a well-written and readily accessible work. Technical jargon is avoided or carefully explained. Its weakness from an AI perspective is its lack of scope. Much of importance in AI is omitted, such as expert systems, heuristic search theory, and formal methods in theorem proving. They may not be directly related to GOFAI, but they do not constitute a minor part of the work in AI, as the author states on page 116. This book is a worthwhile and highly readable supplement to the AI and philosophical literature. IRA POHL

JOSEPH Y. HALPERN. *Reasoning about knowledge: an overview.* ***Theoretical aspects of reasoning about knowledge, Proceedings of the 1986 conference***, edited by Joseph Y. Halpern, Morgan Kaufmann Publishers, Los Altos 1986, pp. 1–17.

There has been a great deal of work of interest to logicians and philosophers going on in computer science and artificial intelligence (AI). This has not merely been work in some particular discipline that is of interest to logicians or philosophers, as, for example, work in physics is of interest to philosophers of physics. Rather, this is logic and philosophy *per se*. As Joseph Halpern says in his foreword to the volume containing the paper under review—the proceedings of a conference on epistemic logic—"This conference represents the first attempt to bring together researchers from [philosophy, linguistics, AI, economics, and theoretical computer science] to discuss issues of mutual interest" (p. vii). And, as he says at the beginning of his own paper, "The commonality of concerns of researchers in all these areas has been quite remarkable. Unfortunately, lack of communication between researchers in the various fields, while perhaps not as remarkable, has also been rather noticeable" (p. 2). This volume—and future such conferences—should do a great deal to remedy that lack. Logicians and philosophers interested in epistemic logic who do not become familiar with the work described in these articles will soon be left behind. This review and the following nineteen reviews will briefly summarize those articles from the volume that are most relevant to readers of this JOURNAL, to give a flavor of the kind of *philosophical* work going on in computer science and AI.

Halpern's article is a useful survey with which to begin the collection. After introducing the syntax and Kripke-style possible-worlds semantics for a propositional epistemic logic for m agents, he notes that "while it is not clear whether [this] model... is appropriate for human reasoning, it can capture quite well much of the reasoning that goes on in analyzing distributed systems" (p. 5). The analogy is this: The abstract notion of a possible world can be interpreted as a global state of a distributed system (i.e., as a description of each processor's state), and the accessibility relation for agent i can be interpreted as the relation between two global states s and t such that processor i has the same state in s and t. Thus, processor i "knows" proposition φ iff φ is true in all global states consistent with i's current state, where φ expresses information about processors' states or the values of their variables, for example. (In the opinion of this reviewer, computational interpretations such as this of the abstract paraphernalia of possible-worlds semantics for modal logics are among the clearest, most revealing, and least metaphysically suspect.) Another major topic discussed in Halpern's survey is the problem of "logical omniscience"—that all agents "know all valid formulas and all logical consequences of their knowledge" (p. 7). Halpern discusses three approaches to the solution of this problem. First, there is Kurt Konolige's syntactic approach, which employs incomplete sets of deduction rules. Secondly, there is Hector Levesque's semantic approach, in which a distinction is made between "explicit" knowledge (the propositions that the agent is explicitly aware of) and "implicit" knowledge (the logical consequences of the explicit knowledge). Of interest to philosophical logicians are, first, Levesque's explicit-belief operator B which is such that $B\varphi \Rightarrow B\psi$ iff φ relevantly entails ψ, and, secondly, Levesque's semantics, which are akin to Barwise and Perry's situation semantics. Finally, there is the combined syntactic-semantic approach of Ronald Fagin and Halpern's "logic of general awareness," which "adds to each state [of a Kripke structure] a set of formulas that the agent is 'aware' of at that state" (p. 8). On this view, implicit knowledge is the same as the standard epistemic-logic concept of knowledge, and an agent a explicitly knows φ iff a implicitly knows φ and φ is in a's awareness set. (It is of some, perhaps sociological, interest that the most serious attention to the problem of logical omniscience has been paid, not by pure philosophers of mind or of language, but by computer scientists; but compare the discussion of Hintikka's article below.) Other topics Halpern discusses include "common" belief and knowledge (variously called "shared" or "mutual" belief and knowledge), with applications to economics and distributed systems; and relations between knowledge and action for planning, querying databases, and communication systems. Each of these topics is dealt with in detail in subsequent articles in the collection.

WILLIAM J. RAPAPORT

BRIAN CANTWELL SMITH. *Varieties of self-reference.* Ibid., pp. 19–43.

Brian Cantwell Smith's philosophically sophisticated contribution discusses a theory of self-reference and examines Levesque's and Fagin and Halpern's epistemic logics in light of that theory. He begins by identifying two problems: the lack of a "clear, single concept of the self" (p. 20) and the relation between self-reference and an AI system that can reason about the world. In good analytic philosophical fashion, Smith "chisholms" away at these problems, beginning with certain assumptions: A "representational system" (an agent—person or computer) has parts, called "representational structures," which themselves have parts, called "aspects." He usefully distinguishes "*im*pressions," which are "internal structures that are causally responsible for an agent's ... actions," from "*ex*pressions," which are "tokens or utterances, external to an agent, in a consensual language" (p. 22). Finally, each part has (1) a "meaning," which "indicates ... what and how it contributes to the content of the composite wholes in which it participates," (2) a "content" or "what a representation ... is about," and (3) a "significance," its content together with the "full conceptual or functional role that the representational structure can play in and for the agent" (p. 23). Smith argues for "circumstantial relativity": the significance of a representational system goes beyond what is represented by its representational structures, i.e., its significance is relative to its circumstances. For example, the referent of 'I' is not part of its meaning, and that Lisp is dynamically scoped is not represented in Lisp. And a representation must be "efficient," allowing different uses of parts of an agent to have different consequences, depending on circumstances. The self makes its appearance in Smith's discussion of three limitations of circumstantial relativity: (1) Circumstantial relativity makes communication difficult if speaker and hearer are in different circumstances; (2) a representation "can empower a system with respect to situations remote in space or time," but then the system must be able to "represent its own relativity"; and (3) a representation is "partially disconnected"—it can represent things "in ways other than how they are" (p. 26). To overcome these limitations, the representation of circumstantial relativity requires the representation of self, which is its source, and, since representation is relative to a theory or "conceptual scheme," self-reference that has "causal connection" between action and reasoning requires a theory of the self, encoded within the system, and a "mechanism of connection" between (1) thinking about and acting in the world and (2) reasoning about oneself and one's situation. Smith notes that representational structures or impressions have "immediate" properties, which are directly causally related to computations over them. Some definitions follow: x *is explicitly represented by* a structure or impression y iff x is represented by an

immediate aspect of *y*, whereas *x is implicit with respect to* a representation *y* iff *x* is not explicitly represented by *y* and *x* is part of the circumstances that determine the content or significance of *y*. For example, "if I say 'there's a bear to the right', I am implicitly involved, but not explicitly represented" (p. 29). If *x* plays some role, then *x* is *explicit* iff *x* is explicitly represented and *x* plays its role in virtue of the explicit representation. Finally, two representational structures are *self-relative* iff different occurrences of them are part of the circumstances that determine their content. Circumstances are "external" if they are parts of the world in which the system does not participate; "indexical" if they are parts of the world in which the system is a constituent; and "internal" if they are impressions, processes defined over them, or relations among them. Representations are similarly categorized according to the category of circumstances that their content depends on. Smith goes on to discuss two varieties of self-reference: "autonymy" and "introspection." A system is defined to be *autonymic* iff it can use a name for itself in a causally connected way; this requires (1) "a mechanism to convert K-ary impressions [e.g. RIGHT (SOMEONE)] to K + 1-ary impressions," e.g. RIGHT (SOMEONE, ____) (p. 31), and (2) an efficient name (e.g. 'I') such that the explicit (K + 1-ary) version has the same content as the implicit (K-ary) version. An *introspective* system has "causally connected self-referential mechanisms that render explicit ... some of their ... implicit internal structure" (p. 31). Smith is one of the few computer scientists to note the need to distinguish *our* theoretical commitments from those of the agents we study, e.g. in a formula such as $\boldsymbol{B}\varphi \Rightarrow \boldsymbol{BB}\varphi$, "the inner ***B***'s represent the agents' views; the outer ones the theorists'" (p. 32). This is the point first made over twenty years ago in Hector-Neri Castañeda's *'He': a study in the logic of self-consciousness*, ***Ratio***, vol. 8 (1966), pp. 130–157 (cf. William J. Rapaport, *Logical foundations for belief representation*, ***Cognitive science***, vol. 10 (1986), pp. 371–422, for a discussion and computer implementation).
WILLIAM J. RAPAPORT

FRED LANDMAN. *Pegs and alecs.* An abridged version of LIII 656. Ibid., pp. 45–61.

Fred Landman's article deals with problems of identity and "partial objects." The objects that we talk about, which are objects in a discourse semantics, are neither objects in the real world nor "symbols in some representation" (p. 46), and they are incomplete, in a Meinongian sense. He rejects Hans Kamp's discourse representation theory and Terence Parsons's interpretation of Meinong, arguing that what is needed is a notion of partial objects that can grow (change) yet be the same object. Landman's analysis is based on a theory of "data semantics," consisting of a set D of partial objects, a system $\mathscr{I}$ of propositions, a set $\mathscr{P}$ of properties (functions from n-tuples of objects to propositions), an interpretation function ι mapping constants to objects and predicates to properties, and a recursive specification of truth and falsity based on the notion of an "information state" (a set of propositions). The first of Landman's key notions, a "peg," is an "informational object" on which compatible properties can be hung by information states. Since "properties map pegs on facts" (p. 48), pegs must be the partial objects. The reference of a peg d in a total information state $\mathbf{w}$ is the equivalence class of all pegs "identical" to d, where d' is identical to d if d and d' are indiscernible (in a technical, but obvious, sense) in all extensions of $\mathbf{w}$. Finally, the real objects of $\mathbf{w}$ are the referents. To this reviewer, pegs seem very like nodes in a propositional semantic network (such as SNePS nodes, given a Meinongian interpretation; cf. Stuart C. Shapiro and William J. Rapaport, *SNePS considered as a fully intensional propositional semantic network*, ***The knowledge frontier: essays in the representation of knowledge***, Springer-Verlag, 1987, pp. 262–315; and cf. William J. Rapaport, *Meinongian theories and a Russellian paradox*, ***Noûs***, vol. 12 (1978), pp. 153–180, on the notion of change) or—Landman's protestations notwithstanding—to the very similar discourse markers of Kamp's theory. They also bear a very close resemblance to bare particulars, since pegs are indiscernible (in the technical sense) if they have no properties hanging on them. Landman gives no indication of being familiar with this well-investigated ontological notion. Pegs are used to analyze three puzzles: Kripke's Pierre-puzzle, the paradox of the hooded man (a paradox about knowing-who—you don't know who the man in the hood is, you do know who your brother is, yet the man in the hood is your brother), and the Hesperus-Phosphorus puzzle. The latter is analyzed as follows. Suppose two Babylonians (B_1 and B_2) talk about Hesperus (h) and Phosphorus (p), agreeing that $h \neq p$ and that B_1's $h = B_2$'s h. Then there is one "real object," four "objects in their heads," and two pegs, namely h and p (p. 51). Landman's second key notion, an "alec," is part of his proposal for a theory of semantics for variables similar to (but different from) Skolem constants or Kit Fine's arbitrary objects (see LIII 305). Inspired by Alec Guinness's having played the roles of all the victims in the film *Kind hearts and coronets*, an alec is a peg that can play several roles. More precisely, let $\mathbf{s}$ be an information state (part of a conversation), let $\mathbf{s}_0$ be its beginning point, let

$\mathscr{B}_s$ be the extensions of s_0, and let $b, b' \in \mathscr{B}_s$ be "alternatives" if they have the same endpoints; then peg a is *an alec with respect to* set of properties $\mathscr{P}$ *in* s iff $(\forall b \in \mathscr{B}_s)$ ($\forall$ peg d with properties in $\mathscr{P}$) ($\exists b'$ alternative to b) [a has properties in $\mathscr{P}$ and a is an indiscernible approximation of d]. As an example, Landman shows how to use alecs to analyze problems of anaphora arising from donkey sentences and witch sentences.

WILLIAM J. RAPAPORT

JAAKKO HINTIKKA. *Reasoning about knowledge in philosophy: the paradigm of epistemic logic.* Ibid., pp. 63–80.

It is only fitting that a conference on reasoning about knowledge should have a contribution by Jaakko Hintikka. Unfortunately, his article has been poorly proofread and seems to be little more than a summary of his earlier work. He begins by discussing epistemic logic and "knowledge representation," taking the latter term in what this reviewer would call the "strong" sense: the representation of *knowledge* as opposed to belief. (The "weak" sense would refer to the representation of information in general, whether true or false.) Hintikka reviews the basic semantic intuitions underlying epistemic logic, namely, that 'knower b knows that S' is true in the "scenario" w_0 (in which b's knowledge is being considered) $\in W$ (a presupposed set of scenarios) iff S is true in all "epistemic b-alternatives" to w_0, where $w \in W_1$ is said to be an epistemic b-alternative to w_0 provided that $w \in W_1 \subseteq W$ iff w is compatible with what b knows in w_0; it is a reflexive and transitive relation. Scenarios are like Barwise and Perry's situations, "covering relatively small pieces of space-time" (p. 65). Hintikka prefers the notation '$\{b\}KS$' to the more familiar 'K_bS', since 'b' should not be in the scope of 'K'; but this avoids an interesting issue: Does 'K' now represent a unique, two-place relation between knower and sentence, or are there different K-operators, one per knower (as suggested by the older notation)? In any event, the aim of Hintikka's paper is to discuss extensions of this intuition about *knowing that* to *knows* + an indirect wh-question (e.g. *knowing who*) and to *knows* + a direct grammatical object. To handle the former, he assumes that there are "criteria of identity for individuals across worlds," i.e., "world lines" "connecting the counterparts of the same individual in different... scenarios" (p. 67). Knowing-who can then be analyzed in terms of knowing-that as follows: b knows who is such that $S[x]$ iff $(\exists x)\{b\}KS[x]$, i.e., iff there is "a world line which in all of b's knowledge worlds picks out an individual x satisfying in that world the condition $S[x]$" (p. 68). To handle knowing an object, Hintikka observes that sometimes "there are two systems of world lines," the second one being "A knower's... cognitive relations to his or her environment" (p. 69). By using a second pair of quantifiers—(ax), (ex)—corresponding to the second world-line system, knowing an object is analyzed as follows: b knows d iff (ex)$\{b\}K(d = x)$, i.e., iff b is acquainted with d, i.e., iff there is among the objects to which b is cognitively related something that is d. Hintikka also sketches two solutions to the problem of logical omniscience, which he characterizes as follows: If $\vdash(S_1 \supset S_2)$, then, for any b and any scenario, (3.1) $\{b\}KS_1 \supset \{b\}KS_2$. He asserts this to be paradoxical but avoidable since "there are... two equivalent ways of delineating the subclass of logical consequences $\vdash(S_1 \supset S_2)$ for which (3.1) holds" (pp. 70–71). The first way is to put syntactic restrictions on the argument from S_1 to S_2, in particular, to require the number of individuals "considered" in any sentence in the argument not to exceed the number in S_1 or S_2. The second way is to use the game-theoretical semantic notion of "urn models" (which Hintikka assumes the reader is familiar with). Game-theoretical semantics is also brought to bear on the interpretation of 'b knows of some individual x that $S[x]$.' Its analysis as '$(\exists x)\{b\}KS[x]$' is rejected in favor of one with branching quantifiers:

$$\left.\begin{array}{l}(\exists x)\\ \\ \{b\}K\end{array}\right> S[x]$$

for which a game-theoretical semantics is appropriate—the moves associated with '$(\exists x)$' and with '$\{b\}K$' are made independently of each other. Hintikka then presents a linearization of this "in a self-explanatory notation" (p. 74): $\{b\}K(\exists x/\{b\}K)S[x]$. Perhaps '/' is to be read "independent of." This notation is used in the analysis of 'b knows whom each person admires.' Rather than '$(\forall x)(\exists y)(\exists z)$ $(x = z \,\&\, \{b\}K(z$ admires $y))$' Hintikka prefers '$\{b\}K(\forall x)(\exists y/\{b\}K)(x$ admires $y)$,' which, he notes, is logically equivalent to '$(\exists f)\{b\}K(\forall x)(x$ admires $f(x))$.' Finally, Hintikka discusses "the most important application of epistemic logic," which "is to the theory of questions and answers" (p. 66). Here, he views knowledge acquisition as a series of questions asked of an information source, whose answers can be used by the questioner in an inference to determine whether some conclusion is true. This is of clear relevance to issues in AI, but it seems to this reviewer that here 'knowledge' is being used in a weaker sense, perhaps akin to (merely) justified belief. (But cf. Hector-Neri Castañeda, *The theory of questions, epistemic powers, and the indexical theory of knowledge*, ***Studies in epistemology***, Midwest studies in philosophy, vol. 5, University of Minnesota Press, 1980, pp. 193–237, for a similar theory.)

WILLIAM J. RAPAPORT

STANLEY J. ROSENSCHEIN and LESLIE PACK KAELBLING. *The synthesis of digital machines with provable epistemic properties.* Ibid., pp. 83–98.

Stanley Rosenschein's "situated-automata approach" to epistemic logic replaces the interpretation of a system x knowing a proposition φ as describing "the *propositional content* of information encoded in x's state without specifying the details of the encoding" (p. 84) by the following: Process x knows φ "in a situation where its internal state is v if φ holds in all possible situations in which x is in state v" (p. 84), a definition that satisfies the S5 axioms. The paper, written with Leslie Pack Kaelbling, extends this approach to "hierarchically constructed machines... by viewing their components as elements of a multi-agent system and reasoning about the flow of information among these components" (p. 84). Rosenschein and Kaelbling present—a bit too compactly for ease of comprehension—the syntax and semantics of a very rich formal propositional language for dealing "with processes, their states, time, and knowledge" (p. 85); the semantics is given in terms of instants of time, locations, and possible worlds. The language contains, *inter alia*, (1) a knowledge operator, K, that operates on pairs of processes and formulas and that satisfies principles of truth ($K(x,\varphi) \to \varphi$), consequential closure ($K(x,\varphi \to \psi) \to (K(x,\varphi) \to K(x,\psi))$), positive introspection ($K(x,\varphi) \to K(x,K(x,\varphi))$), and negative introspection ($\neg K(x,\varphi) \to K(x,\neg K(x,\varphi))$); (2) an operator, *, that takes processes into values and that satisfies a principle of self-awareness ($*X = v \to K(X, *X = v)$); and (3) time and world modal operators, $\Box_T$, $\Box_W$, together with sets of possible worlds and time instants. The language is used to model complex machines built from, for example, logic gates and delay components. The authors go on to describe Rex, an extension of Lisp, for specifying such complex machines described in the language. Later sections give examples of such machines and their epistemic properties. In particular, the approach has been applied to a mobile robot under development at SRI International—now *there's* applied philosophy for you!

WILLIAM J. RAPAPORT

LEORA MORGENSTERN. *A first order theory of planning, knowledge, and action.* Ibid., pp. 99–114.

The earliest and still most important application of epistemic logic to AI is the work of Robert Moore (*Reasoning about knowledge and action*, ***5th International Joint Conference on Artificial Intelligence–1977*** (IJCAI-77), volume 1, pp. 223–227; *A formal theory of knowledge and action*, ***Formal theories of the commonsense world***, Ablex, 1985, pp. 319–358). Moore's system added concepts of action to an S4 epistemic logic with a possible-worlds semantics. Leora Morgenstern's article extends Moore's work to situations in which an agent lacks complete knowledge. She critiques Moore's system on four grounds: (1) An executable description of an action should not be a rigid designator, as Moore has it, but a world-specific property of the action; (2) in Moore's system, "all general procedures for action are known" by all agents, but sometimes "an agent cannot do an action because he has no idea how to do the general procedure" (p. 102); (3) Moore's first-order epistemic logics are not expressive enough—for example, they can't express 'John knows that Bill knows something that he doesn't know'; and (4) "a complete theory of knowledge and action" should deal with learning and with the delegation of authority (p. 103). Morgenstern's system uses a first-order predicate logic with quotation and a syntactic predicate 'Know' that ranges over pairs of agents and names of sentences. To avoid the knower paradox (Richard Montague, *Syntactical treatments of modality, with corollaries on reflexion principles and finite axiomatizability*, XL 600; David Kaplan and Richard Montague, *A paradox regained*, XXX 102) concerning a sentence S such that S is true iff Know(a, '$\sim S$'), she adapts Kripke's theory of truth (Saul Kripke, *Outline of a theory of truth*, L 1068). Let L be a classical first-order language containing the relation Believe(a,p), and let L_0 be the sentences of L without truth- or knowledge-predicates (where knowledge is true belief). Then extend L_0 to a sequence of languages L_i such that Know(a,p) is (1) true in L_i if p has a positive truth value in L_{i-1} and Believe(a,p), (2) false in L_i if either p has a negative truth value in L_{i-1} or else p is true in L_{i-1} and $\neg$Believe(a, p), and (3) undefined otherwise. A sentence is *grounded* if it has a truth value at a fixed point L_σ. The knower paradox sentence S is ungrounded, hence no paradox. Morgenstern's unified theory of planning, knowledge, and action is expressed in a first-order language L, in which 'Know' is a three-place predicate over agents, (names of?) propositions, and "situations," with a temporal logic based on that of Drew McDermott, *A temporal logic for reasoning about processes and plans*, ***Cognitive science***, vol. 6 (1982), pp. 101–155. The axiom schemata for grounded sentences of L include (K1) the axioms of predicate logic, an axiom schema (K4) saying "that agents can and do reason with the rules of inference of predicate logic" (p. 108), and axiom schemata (K5) saying that "all agents know all axioms of logic and the basic axioms of knowledge" (p. 109). An agent, a, is said to *know how to do* action A iff either A is a "primitive" action and a knows the parameters of A, or a knows how A is constructed from primitive actions and a knows how to do those primitive actions. Concepts of an agent *knowing how to perform* and *being able to perform* an action and *knowing how to achieve* a situation are introduced. Planning is introduced as follows. Saying that agent *a can execute* plan *p in* situation *s* means that a knows in s that he will be able to perform all actions in p for which he is an agent and moreover a can predict that all other events in p occur in proper order. Note—as Morgenstern does not—the essential use of quasi-indicators here: for a to execute p in s, a must know in s that he* (i.e., he himself) will be able to perform all actions in p for which he* is an agent. Finally, delegation of authority is represented by a "controls" predicate of L: if agent a controls agent b with respect to task t, then a can delegate t to b.

WILLIAM J. RAPAPORT

JIM DES RIVIÈRES and HECTOR J. LEVESQUE. *The consistency of syntactical treatments of knowledge.* Ibid., pp. 115–130.

Jim des Rivières and Hector Levesque's article is a fine example of the kinds of *philosophical* contributions being made by AI researchers. The authors present an argument against Richard Montague's and Richmond Thomason's claims that the $\Box$ *operator* (e.g. necessity, knowledge, belief) on sentences cannot be replaced by a *predicate* L on sentences. Montague 1963 (op. cit.) claimed, roughly, that a first-order syntactic theory containing arithmetic and the basic axioms of an alethic or epistemic modal logic is inconsistent. Again roughly, Richmond Thomason extended this result to a modal logic for idealized belief (*A note on syntactical treatments of modality*, ***Synthese***, vol. 44 (1980), pp. 391–395). Rivières and Levesque's straightforwardly philosophical result, however, is motivated by computational issues: although Quine was the major philosopher who wanted such a result, it has been AI researchers (such as John McCarthy, Lewis Creary, Moore, Konolige, and Donald Perlis) who have felt a need for it. Rivières and Levesque's approach begins by considering the sentences over which α should range in modal axiom schemata such as $\Box\alpha \supset \alpha$, and their non-modal analogues $L(\ulcorner\alpha\urcorner) \supset \alpha$. They determine these by finding the set of sentences described by the modal schema, de-modalizing each of these, and then finding a "classical" (i.e., non-modal) schema for the resulting set. This cannot be done easily; "there are some sentences in the classical language that have no shorthand equivalent in the modal language," e.g. $\exists x L(x)$. Therefore, "our re-reading of sentences in the modal language only yields a [proper] subset of the classical language," called the set of "regular" sentences (p. 117). So $\Box\alpha \supset \alpha$ for all sentences α in the modal language becomes $L(\ulcorner\alpha\urcorner) \supset \alpha$ for all regular sentences α in the classical language. This allows a syntactical treatment of modal operators. More formally, let $\mathscr{L}$ be a first-order language with $\neg$, $\supset$, $\forall$, and predicate, function, and constant symbols; let $\mathscr{L}(\Box)$ be $\mathscr{L}$ + the modal sentential operator $\Box$; and let $\mathscr{L}(L)$ be $\mathscr{L}$ + all $(n+1)$-ary predicate symbols L_n (n a natural number). For any term or formula α of $\mathscr{L}(L)$, there is called an *encoding term* $\ulcorner\alpha\urcorner$ of $\mathscr{L}(L)$, viz., its name. Next, $\diamond : \mathscr{L}_1 \to \mathscr{L}_2$ is called an *embedding of* one first-order language $\mathscr{L}_1$ *in* another, $\mathscr{L}_2$, iff (i) $\alpha^\diamond = \alpha$ for atomic α, (ii) $\diamond$ distributes over $\neg$, $\supset$, $\forall$, and (iii) $\alpha^\diamond$ and α have the same free variables. The sentence φ is called an *extended theorem of* $\mathscr{L}_1$ iff $\varphi^\diamond$ is a theorem of first-order logic for all $\diamond : \mathscr{L}_1 \to \mathscr{L}(L)$. Sentence φ *is derivable from* a set of sentences S ($S \vdash \varphi$) iff φ follows from S and the extended theorems of $\mathscr{L}_1$ by modus ponens alone. The set S *is inconsistent* iff $S \vdash \varphi$ for all φ in $\mathscr{L}_1$. Next, define $* : \mathscr{L}(\Box) \to \mathscr{L}(L)$ such that $(\Box\alpha)^* = L_n(\ulcorner\alpha^*\urcorner, x_1, \ldots, x_n)$, where the x_i are the free variables of α. Sentence $\alpha \in \mathscr{L}(L)$ is called *regular* iff there is $\gamma \in \mathscr{L}(\Box)$ such that $\alpha = \gamma^*$. Then $\diamond : \mathscr{L}_1 \to \mathscr{L}_2$ is said to be a *reduction of* $S \subseteq \mathscr{L}_1$ *to* $T \subseteq \mathscr{L}_2$ iff for all $\varphi \in \mathscr{L}_1$, $S \vdash \varphi$ iff $T \vdash \varphi^\diamond$, and it *is a general reduction of* $\mathscr{L}_1$ *to* $\mathscr{L}_2$ iff for all $T \subseteq \mathscr{L}_1$, $\diamond$ is a reduction of T to $\{\varphi^\diamond | \varphi \in T\}$. The main result is Theorem 6: $*$ is a general reduction of $\mathscr{L}(\Box)$ to $\mathscr{L}(L)$. Thus, $\Box$ can be taken "as shorthand for the multigrade predicate L and a pair of quotation marks that Quine has countenanced. Applied to a genuine modal theory such as **S5**, one obtains a classical, first-order system that rightfully deserves to be considered as *a syntactical treatment of modality*" (p. 125). The authors conclude their paper with an interpretation of Montague's and Thomason's results in light of the notion of a regular sentence.

WILLIAM J. RAPAPORT

NICHOLAS M. ASHER and JOHAN A. W. KAMP. *The knower's paradox and representational theories of attitudes.* Ibid., pp. 131–147.

Nicholas Asher and Hans Kamp's article is a nice companion to the previous one, since it is by two philosophers and is also on the Montague–Thomason paradoxes. The "knower's paradox" of the article's title refers to Montague's result for "representational" theories, e.g. ones in which the objects of propositional attitudes are sentences. Asher and Kamp begin by showing how a similar paradox can be derived in Montague's non-representational intensional logic (where propositions are sets of possible worlds) augmented by "enough arithmetic to permit goedelization" (p. 133), and they do the same for a variety of other theories, with and without sentence-forming operators and with and without a recursive representational structure. The solution they offer to these paradoxes is an extension of the techniques developed by H. Herzberger (*Naive semantics and the liar paradox*, ***The journal of philosophy***, vol. 79 (1982), pp. 479–497) and A. Gupta (*Truth and paradox*, L 1068) for a theory of truth immune to the liar paradox. In particular, given a Hintikka-style analysis of $\mathbf{a}$'s knowledge in world $\mathbf{w}$ in terms of a set $\mathbf{W}_{\mathbf{K},\mathbf{a}}(\mathbf{w})$ of possible worlds, they wish "to determine, for any possible world structure W with alternativeness relations for knowledge and belief, what in each world of W are the extensions for the knowledge and belief predicates, K and B," especially for "self-referential belief reports about a" (p. 137). More formally (for the case of belief; knowledge is treated analogously): Let L be a first-order language with a two-place predicate $B(x, y)$ (x believes that y). A model $\mathcal{M}$ for L is a structure $\langle \mathbf{W}, \mathbf{D}, [\ ], \{\mathbf{R}_{\mathbf{B},\mathbf{a}} \mid \mathbf{a} \in \mathbf{A}\}\rangle$, where $\mathbf{w}\mathbf{R}_{\mathbf{B},\mathbf{a}}\mathbf{w}'$ iff (0) $\mathbf{w}'$ is compatible with the totality of $\mathbf{a}$'s beliefs in $\mathbf{w}$, (i) $\mathbf{W}$ is a set of worlds, (ii) $\mathbf{D}\colon \mathbf{w} \in \mathbf{W} \mapsto \mathbf{D}_{\mathbf{w}} \neq \varnothing$ (the *universe* of $\mathbf{w}$), (iii) $\mathbf{D}_{\mathbf{w}} = \mathbf{D}_{\mathbf{w}'}$ for all $\mathbf{w}, \mathbf{w}' \in \mathbf{W}$, (iv) the set of agents $\mathbf{A} \subseteq \mathbf{D}_{\mathbf{w}}$, (v) $[\ ]$: non-logical constant of $L \mapsto$ its classical extension at each $\mathbf{w}$, (vi) for all $\mathbf{w}, \mathbf{w}' \in \mathbf{W}$ in $\mathcal{M}$, $[\mathbf{c}]_{\mathbf{w}} = [\mathbf{c}]_{\mathbf{w}'}$ for each individual constant $\mathbf{c}$, and (vii) every sentence of L is included in the universe of $\mathcal{M}$. $\mathcal{M}$ is *doxastically coherent* iff for every sentence ψ and for every world $\mathbf{w}$, $\langle \mathbf{a}, \psi\rangle \in \mathbf{B}_{\mathcal{M},\mathbf{w}}$ iff ψ is true at all $\mathbf{w}' \in \{\mathbf{w}'' \mid \mathbf{w}\mathbf{R}_{\mathcal{M},\mathbf{B},\mathbf{a}}\mathbf{w}''\}$. Then, given $\mathcal{M}$, define for each ordinal α the model $\mathcal{M}^{\alpha} = \langle \mathbf{W}_{\mathcal{M}}, \mathbf{D}_{\mathcal{M}}, \mathbf{R}_{\mathcal{M}}, [\ ]^{\alpha}\rangle$ where $[\mathbf{Q}]^{\alpha} = [\mathbf{Q}]_{\mathcal{M}}$ for any non-logical constant $\mathbf{Q} \neq \mathbf{B}$, and $[\mathbf{B}]^{\alpha}$ is defined by the conditions: (i) $[\mathbf{B}]^{0}_{\mathbf{w}} = [\mathbf{B}]_{\mathbf{w}}$, (ii) $[\mathbf{B}]^{\alpha+1}_{\mathbf{w}} = \{\varphi \mid \forall \mathbf{w}'(\mathbf{w}\mathbf{R}\mathbf{w}' \to [\varphi]_{\mathcal{M}^{\alpha},\mathbf{w}'} = 1)\}$, and (iii) if α is a limit ordinal, then $[\mathbf{B}]^{\alpha}_{\mathbf{w}} = \{\varphi \mid (\exists \beta < \alpha)(\forall \gamma)(\beta \leqq \gamma < \alpha \to \varphi \in [\mathbf{B}]^{\gamma}_{\mathbf{w}})\}$. Asher and Kamp conclude by proving a number of theorems about such issues as these: Which models become coherent after such revisions? When do they become coherent? Which sentences get "settled"? What epistemic and doxastic logics are determined by the (in)coherent models? The necessary background for a statement of these theorems is, unfortunately, beyond the scope of this brief review. There is clearly, however, a need to compare the several approaches discussed in these two papers, as well as the approaches mentioned in Halpern's paper.

WILLIAM J. RAPAPORT

CYNTHIA DWORK and YORAM MOSES. *Knowledge and common knowledge in a Byzantine environment I: crash failures.* Ibid., pp. 149–169.

There are four articles on applications of epistemic logic to distributed processing. The first of these, by Cynthia Dwork and Yoram Moses, is, perhaps, of less interest to some readers of this JOURNAL, having to do with "the problem of designing effective protocols for distributed systems whose components are unreliable" (p. 150), a so-called "Byzantine event." It is, nonetheless, of interest as an application of epistemic logic to a non-cognitive domain. The crucial formal notion is the following: A processor p_i in a synchronous distributed system S *knows* a fact ψ in S at an "execution" (ρ, k), denoted $(S, \rho, k) \models K_i\psi$ (where ρ is a set of messages sent and received by S—its "history"—and k is the number of "rounds" of message-passing) iff for all $(\rho', k) \in S \times \{k\}$ such that $v(p_i, \rho, k) = v(p_i, \rho', k)$ we have $(S, \rho', k) \models \psi$, where $v(p, \rho, k)$ is processor p's "view" of S's configuration at (ρ, k), and $(S, \rho, k) \models \psi$ iff $(\rho, k) \in \tau(\psi)$, a set of executions (= a "ground fact"). That is, p_i knows ψ if ψ holds given p_i's view. In particular, K_i satisfies the usual epistemic axioms, including $K_i\varphi \supset \varphi$. This notion is extended to a group of processors knowing that φ at (ρ, k). Although the authors are only interested in applications of their theory to technical problems in distributed systems, it would be of interest to see how such notions might be adapted to issues of mutual (or common) knowledge among cognitive agents in a natural-language discourse.

WILLIAM J. RAPAPORT

RICHARD E. LADNER and JOHN H. REIF. *The logic of distributed protocols (preliminary report).* Ibid., pp. 207–222.

The notion of a protocol is somewhat more intuitively presented in the article by Richard Ladner and John Reif. "A protocol is an algorithm whose execution is shared by a number of independent participants or ... *players.* Each player may be unaware of what the other players are exactly doing. Therefore, a key ingredient found in the behavior of many protocols (and not found in sequential or parallel algorithms) is the *lack of knowledge* about the complete state of the protocol by each of its players" (p. 208). Their own theory is stated in terms of a temporal epistemic logic.

WILLIAM J. RAPAPORT

RONALD FAGIN and MOSHE Y. VARDI. *Knowledge and implicit knowledge in a distributed environment: preliminary report.* Ibid., pp. 187–206.

The article on distributed systems by Ronald Fagin and Moshe Vardi is perhaps the most accessible to those who are not computer scientists. Their main result is that "S5 is not complete for reasoning about knowledge in distributed systems" (p. 189)—an additional axiom is needed, for *implicit knowledge*, understood here as follows: "The implicit knowledge of a group G is what someone could infer given complete knowledge of what each member of G knows. For example, if Alice knows φ_1 and Bob knows $\varphi_1 \Rightarrow \varphi_2$, then together they have implicit knowledge of φ_2, even though neither of them might individually know φ_2" (p. 189). (This has also been noted by Nicolas Goodman, *Egocentric and sociocentric epistemic logic* (abstract), this JOURNAL, vol. 50 (1985), p. 1096.) The new axiom is $I \sim \alpha \Rightarrow (K_1 \sim \alpha \vee \cdots \vee K_n \sim \alpha)$, where $1, \ldots, n$ are all the members of the group.

WILLIAM J. RAPAPORT

RICHMOND H. THOMASON. *Paradoxes and semantic representation.* Ibid., pp. 225–239.

Richmond Thomason's brief and rather programmatic paper belongs thematically with the earlier papers in this volume on the knower's paradox. He reviews various solutions to the liar paradox, classified under three strategies (p. 226). (1) "Hold certain seemingly intelligible notions to be inexpressible." (2) "Give up plausible schematic principles on truthlike predicates, such as Convention T." (3) "Impose limits on the extent to which language can be used to talk about its own syntax." He then discusses two "conflicting assumptions" concerning intensional versions of the liar paradox (such as the knower's paradox). (I′) Semantic representations of sentences can be calculated. (II′) Semantic representations of sentences "yield ... simple explanations of boolean connectives, modal operators, and the like" (p. 231). The programmatic portion of the paper consists of a brief discussion of four strategies for resolving this conflict. (1′) "Impose limits on the extent to which language can be used to talk about its propositions, while ... allowing it the full expressive power of quotation" (p. 233). (2′) "Do not require every sentence to express a proposition, and ... limit Convention T to sentences that express propositions" (p. 233). (4) Treat propositional attitudes as not satisfying alethic modal conditions. The fourth, unnumbered, is to use ramified type theory.

WILLIAM J. RAPAPORT

KURT KONOLIGE. *What awareness isn't: a sentential view of implicit and explicit belief.* Ibid., pp. 241–250.

Propositional-attitude approaches to epistemic and doxastic logics treat knowledge and belief as relations between agents and propositions; sentential approaches treat them as relations between agents and sentences that express propositions. The former approach, using Hintikka- and Kripke-style possible-worlds semantics, have the property of logical omniscience or "consequential closure." As Kurt Konolige expresses it in his article, "an agent's beliefs [or knowledge] are closed under logical consequence" (p. 242). Fagin and Halpern's logic of general awareness, following Levesque's, distinguishes between "explicit" and "implicit" beliefs (R. Fagin and J. Y. Halpern, *Belief, awareness, and limited reasoning: preliminary report*, ***Proceedings of the ninth International Joint Conference on Artificial Intelligence*** (IJCAI-85), volume 1, Morgan Kaufmann, 1985, pp. 491–501; Hector J. Levesque, *A logic of implicit and explicit belief*, ***Proceedings of the National Conference on Artificial Intelligence*** (AAAI-84), William Kaufmann, 1984, pp. 198–202). Kurt Konolige's article argues against their approach. He begins with a brief review of Fagin and Halpern's theory (a propositional language with B, L, and A operators for explicit belief, implicit belief, and awareness, respectively), in particular (and roughly; cf. the earlier discussion of Halpern's contribution to this volume): (1) an agent implicitly believes sentence φ in state s iff φ is true in all states accessible from s (where accessibility is a transitive, Euclidean, serial binary relation), and (2) an agent explicitly believes φ in s iff he implicitly believes φ in s and is aware of φ in s. An agent's being aware of φ amounts to the agent's being "able to determine whether or not φ follows from his initial premises in time T" (p. 245). Konolige's first objection is that "in the case of awareness, the formal correspondence between accessibility conditions and sets of awareness sentences breaks down; hence the connection between accessibility conditions and belief is ruptured" (p. 246). The second objection is that "the logic of general awareness represents agents as perfect reasoners, restricted to considering some syntactic class of sentences." But there are no "clear intuitions that this is the case for human or computer agents" (p. 247). WILLIAM J. RAPAPORT

MOSHE Y. VARDI. *On epistemic logic and logical omniscience.* Ibid., pp. 293–305.

Moshe Vardi's article also treats the problem of logical awareness. He rejects Montague's intensional logic on the grounds that the notion of a possible world is primitive, hence unexplained and non-constructive. In its place, he offers a constructive theory of belief worlds. Roughly, a depth-0 world is a truth assignment to atomic propositions, and a depth-k world is a collection of sets of depth-$(k - 1)$ worlds. The formula $B_a\varphi$ is satisfied in an $(r + 1)$-ary world if $r \geqq 1$ and $\{w \mid w \in W_r \,\&\, \varphi \text{ is satisfied in } w\} \in f_r(a)$, where W_r = the set of all r-ary worlds and $f_r(a)$ is a set of propositions (i.e., sets of r-ary worlds). This is then extended to infinitary worlds. Vardi proves three main theorems. Theorem 3: The validity problem for belief worlds is decidable. Theorem 4: Validity in belief worlds can be soundly and completely axiomatized by (A1) all substitution instances of propositional tautologies and (R1) from $\varphi \equiv \psi$, infer $B_a\varphi \equiv B_a\psi$. Theorem 5: There is a belief structure (*á la* Montague) that models the collection of all belief worlds. Vardi then extends his system to allow for agents to reason (e.g. to allow $B_a\varphi \wedge B_a\psi \supset B_a(\varphi \wedge \psi)$ to be valid), to obtain *knowledge* worlds, and to allow for non-standard belief structures in the fashion of Nicholas Rescher and Robert Brandom (***The logic of inconsistency***, XLVII 233).

WILLIAM J. RAPAPORT

CHRISTOPHE GEISSLER and KURT KONOLIGE. *A resolution method for quantified modal logics of knowledge and belief.* Ibid., pp. 309–324.

Christophe Geissler and Kurt Konolige's article offers a resolution inference rule for quantified epistemic and doxastic logics. Let L be a first-order modal language with function symbols and modal operators $[S]$ for each agent S. The semantics for L are given in terms of Kripke possible-worlds models with accessibility relations for each agent, such that the domain of each possible world is a subset of the domain of any accessible world. To get a version of Herbrand's theorem that holds for L, they revise the definition of "substitution" and of "instance" using a rigid-designator operator, $\bullet$, such that $\bullet t$ refers to what t denotes in the actual world. Now consider the modal logic K, in which accessibility has no restrictions. Let Γ be a set of formulas of L, and let $\Gamma^{\bullet} = \Gamma$ with the $\bullet$-operator uniformly replaced by a new unary function. Then $[S]\varphi_1 \vee A_1, \ldots, [S]\varphi_n \vee A_n, \neg[S]\delta \vee A \vdash A_1 \vee \cdots \vee A_n \vee A$, when $\{\varphi_1, \ldots, \varphi_n, \neg\delta\}^{\bullet}$ is K-unsatisfiable, is a sound and complete resolution rule for K.

WILLIAM J. RAPAPORT

GERHARD LAKEMEYER. *Steps towards a first-order logic of explicit and implicit belief.* Ibid., pp. 325–340.

Gerhard Lakemeyer's article extends Levesque's propositional theory of implicit and explicit belief to include quantification. He argues that the semantics for the theory of explicit belief should satisfy the following sentences (where L is the explicit-belief operator, P is a unary predicate, x is a variable, a is a non-rigid designator, and n is a standard name): (1) $\vDash \exists xLPx \supset L\exists xPx$, (2) $\nvDash L\exists xPx \supset \exists xLPx$, (3) $\vDash LPa \supset L\exists xPx$, (4) $\nvDash LPa \supset \exists xLPx$, (5) $\vDash LPn \supset \exists xLPx$, and (6) $\vDash L\forall xPx \supset \forall xLPx$, but not (7) $\vDash \forall xLPx \supset L\forall xPx$. This is accomplished by means of a first-order language containing B and L operators and *parameters*, i.e., rigid designators, which allow the language to distinguish between knowing-what and knowing-who. The semantics for the language is an extension of a variant of tautological entailment, called *t-entailment*, due to P. F. Patel-Schneider (*A decidable first-order logic for knowledge representation*, ***Proceedings of the ninth International Joint Conference on Artificial Intelligence*** (IJCAI-85), volume 1, Morgan Kaufmann, 1985, pp. 455–458). WILLIAM J. RAPAPORT

RAYMOND M. SMULLYAN. *Logicians who reason about themselves.* Ibid., pp. 341–352.

Raymond Smullyan's article, written in his typical style, uses "knights" (who only make true statements) and "knaves" (who only make false ones), to investigate "some curious epistemic counterparts of undecidability results in metamathematics" (p. 341). Since it is difficult to summarize Smullyan's results—and doing so would make them lose most of their appeal—I shall merely give an example: A reasoner with (1) a complete knowledge of propositional logic, who (2) "also *knows* that his beliefs are closed under modus ponens," who is such that (3) "if whenever he believes *p*, he also believes that he believes *p*," and who (4) believes that he has property (3) is said to be a type-4 reasoner (p. 345). Here is Smullyan's analogue of Gödel's second theorem. "*Problem 3*... A logician L of type 4 visits a knight-knave island (or at least he *believes* it to be one) and meets N who says: 'You will never believe that I'm a knight.' Prove that if L is consistent, he can never know that he is—or put another way, if L ever believes that he is consistent, he will become inconsistent" (p. 346). (The solution is provided.)

WILLIAM J. RAPAPORT

JOÃO P. MARTINS and STUART C. SHAPIRO. *Theoretical foundations for belief revision.* Ibid., pp. 383–398.

It has been known since the mid-1970's that relevance logic is an appropriate logic for a database or knowledge base that uses automated reasoning (cf. Nuel Belnap, *How a computer should think*, ***Contemporary aspects of philosophy***, Oriel Press, 1975, pp. 30–56). After all, if one user tells the knowledge base that *p*, another tells it that $\neg p$, and a third queries it whether *q*, the system should not (necessarily) respond with 'yes.' The article by João Martins and Stuart Shapiro discusses a formal relevance logic underlying a belief revision system, i.e., a system for detecting and eliminating contradictions in a knowledge base. Such a system must record the inferential source of every proposition in its knowledge base. There are two ways to do this. In a *justification-based* system, each proposition's support record consists of all propositions from which it was inferred; in an *assumption-based* system, each proposition's support record consists only of the hypotheses (i.e., non-derived propositions) from which it was inferred. Martins and Shapiro choose the latter because it is easier (1) to find the source of the contradiction, (2) to change beliefs (i.e., to eliminate the source), and (3) to compare belief sets. They introduce SWM, a quantified relevance logic that allows the hypotheses supporting a given proposition to be computed and that "remembers" contradictions that have been derived. SWM deals with *supported wffs*, i.e., wffs accompanied by an origin tag (OT), an origin set (OS), and a restriction set (RS). The OS is the set of all and only the "hypotheses... *actually used* in the derivation of that wff" (p. 390). The OT indicates whether the wff is a hypothesis, a derived wff, or a special "extended" wff. As for the RS, "A wff, say *A*, whose RS is $\{R1, \ldots, Rn\}$ means that the hypotheses in [its OS] added to any of the sets $R1, \ldots, Rn$ produce an *inconsistent set*. The RS of an extended wff will contain *every* set which unioned with the wff's OS will produce a set that is known to be inconsistent. Our rules of inference guarantee that the information contained in the RS is carried over to the new wffs whenever a new proposition is derived. Furthermore, the rules of inference guarantee that RSs do not contain any redundant information" (p. 390). The rules of inference (generally, introduction-elimination rules for a Fitch-style natural-deduction system, together with an obligatory rule of "updating restriction sets") specify how the OS, OT, and—especially—the RS must be updated. They satisfy the following properties. "1. The OS of a supported wff contains *every* hypothesis that was used in its derivation. 2. The OS of a supported wff *only* contains the hypotheses that were used in its derivation. 3. The RS of a supported wff records *every* set known to be inconsistent with the wff's OS. 4. The application of rules of inference is blocked if the resulting wff would have an OS known to be inconsistent" (pp. 391–392). The belief revision system based on SWM is called the Multiple Belief Reasoner and has been implemented in the SNePS semantic-network processing system (Stuart C. Shapiro, *The SNePS semantic network processing system*, ***Associative networks***, Academic Press, 1979, pp. 179–203).

WILLIAM J. RAPAPORT

GORDON PLOTKIN and COLIN STIRLING. *A framework for intuitionistic modal logics.* Ibid., pp. 399–406.

Gordon Plotkin and Colin Stirling's article provides "a Kripkean analysis of intuitionistic modal logic" (p. 399). Consider a set *W* of worlds and two relations on it: $\sqsubseteq$, an intuitionistic partial order, and *R*, a modal accessibility relation, such that (1) if $w \sqsubseteq w'$ and wRv, then $\exists v'[w'Rv' \,\&\, v \sqsubseteq v']$, and (2) if $v \neq v'$ and $v \sqsubseteq v'$ and wRv, then $\exists v'[w'Rv' \,\&\, w \sqsubseteq w']$. These "guarantee" (1) $\Diamond A \rightarrow \neg\Box\neg A$ and (2) $\neg\Diamond A \rightarrow \Box\neg A$, respectively. Plotkin and Stirling develop a sentential modal logic corresponding to this and extend it to other systems.

WILLIAM J. RAPAPORT

NIMROD MEGIDDO and AVI WIGDERSON. *On play by means of computing machines (preliminary version).* Ibid., pp. 259–274.

HAIM GAIFMAN. *A theory of higher order probabilities.* Ibid., pp. 275–292.

SILVIO MICALI. *Knowledge and efficient computation.* Ibid., pp. 353–362.

JOHN C. MITCHELL and MICHAEL J. O'DONNELL. *Realizability semantics for error-tolerant logics (preliminary version).* Ibid., pp. 363–381.

The paper by Nimrod Megiddo and Avi Wigderson is a discussion of the prisoner's dilemma game, using computers to play the game. Haim Gaifman's contribution discusses a system in which a probability-assignment operator can be iterated in the way modal or epistemic operators can be. Silvio Micali's paper discusses the quantity of knowledge, measured by "knowledge complexity"— a notion related to, but distinct from, notions from information theory—that needs to "be communicated for proving a theorem" (p. 353) and for correctness proofs for cryptographic protocols. Finally, the article by John Mitchell and Michael O'Donnell proposes a new semantics for relevance logic "based on the intuitionistic concept of realizability" (p. 366).

The best summary one can give for this volume and the research programs described in it comes from Thomason's contribution, with which I shall conclude this sequence of reviews. "The more closely I have become acquainted with the theories of reasoning that are being presently developed in Computer Science, the more urgently I have felt the need for philosophers to become acquainted with these theories. Philosophers, for instance, seem to know little about knowledge representation, whereas Computer Scientists have learned the relevant philosophy. But in the case of the paradoxes, it seems to me that Computer Scientists still have as much to learn from Philosophers as Philosophers have to learn from Computer Scientists. One short-range consequence for Computer Science, then, is that familiarity with the philosophical literature on the paradoxes is important for research on paradox-related issues concerning semantic representation.... Meetings such as this one provide a pleasant way to speed up the interdisciplinary interactions" (p. 236). And so do volumes such as this.

WILLIAM J. RAPAPORT