

Reasoning About Knowledge and Probability

RONALD FAGIN AND JOSEPH Y. HALPERN

IBM Almaden Research Center, San Jose, California

Abstract. We provide a model for reasoning about knowledge and probability together. We allow explicit mention of probabilities in formulas, so that our language has formulas that essentially say “according to agent i , formula φ holds with probability at least b .” The language is powerful enough to allow reasoning about higher-order probabilities, as well as allowing explicit comparisons of the probabilities an agent places on distinct events. We present a general framework for interpreting such formulas, and consider various properties that might hold of the interrelationship between agents’ probability assignments at different states. We provide a complete axiomatization for reasoning about knowledge and probability, prove a small model property, and obtain decision procedures. We then consider the effects of adding common knowledge and a probabilistic variant of common knowledge to the language.

Categories and Subject Descriptors: B.4.4 [Input / Output and Data Communications]: Performance Analysis and Design Aids—*formal models; verification*; C.2.4 [Computer-Communication Networks]: Network Operations, D.2.4 [Software Engineering]: Program Verification; F.2 [Analysis of Algorithms and Problem Complexity]; F.3.1 [Logics and Meanings of Programs]: Specifying and Verifying and Reasoning about Programs; F.4.1 [Mathematical Logic and Formal Languages]. Mathematical Logic—*model theory*; F.4.m [Mathematical Logic and Formal Languages]: Miscellaneous; G.3 [Probability and Statistics]; I.2.4 [Artificial Intelligence]: Knowledge Representation Formalisms and Methods

General Terms: Theory, Verification

Additional Key Words and Phrases. Knowledge, modal logic, nondeterminism vs. probability, possible worlds, probabilistic common knowledge, probabilistic knowledge, reasoning about knowledge and probability

1. Introduction

Reasoning about knowledge has become an active topic of investigation for researchers in such diverse fields as philosophy [Hintikka, 1962], economics [Aumann, 1976], and artificial intelligence [Moore, 1985]. Recently the interest of theoretical computer scientists has been sparked, since reasoning about knowledge has been shown to be a useful tool in analyzing distributed systems (see Halpern [1987] for an overview and further references).

A preliminary version of this paper appeared in *Proceedings of the 2nd Conference on Theoretical Aspects of Reasoning about Knowledge*, M. Y. Vardi, ed. Morgan-Kaufmann, San Mateo, Calif., 1988, pp. 277–293.

Authors’ address: IBM Almaden Research Center, 650 Harry Road, San Jose, CA 95120.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

© 1994 ACM 0004-5411/94/0300-0340 \$03.50

In many of the application areas for reasoning about knowledge, it is important to be able to reason about the probability of certain events as well as the knowledge of agents. In particular, this arises in distributed systems applications when we want to analyze randomized or probabilistic programs. Not surprisingly, researchers have considered knowledge and probability before. Indeed, all the publications in economics on reasoning about knowledge, going back to Aumann's seminal paper [Aumann, 1976], have probability built into the model. However, they do not consider a logical language that explicitly allows reasoning about probability. In this paper, we consider a language that extends the traditional logic of knowledge by allowing explicit reasoning about probability along the lines discussed in a companion paper [Fagin et al., 1990].

In the standard possible-worlds model of knowledge (which we briefly review in the next section), agent i knows a fact φ , written $K_i\varphi$, in a world or state s if φ is true in all the worlds the agent considers possible in world s . We want to reason not only about an agent's knowledge, but also about the probability he places on certain events. In order to do this, we extend the language considered in [Fagin et al., 1990], which is essentially a formalization of Nilsson's probability logic [Nilsson, 1986]. Typical formulas in the logic of Fagin et al. [1990] include $w(\varphi) \geq 2w(\psi)$ and $w(\varphi) < 1/3$, where φ and ψ are propositional formulas. These formulas can be viewed as saying " φ is twice as probable as ψ " and " φ has probability less than $1/3$ ", respectively. Since we want to reason about the probability that agent i places on events, we modify their language to allow formulas such as $w_i(\varphi) \geq 2w_i(\psi)$. We also allow φ and ψ here to be arbitrary formulas (which may themselves contain nested occurrences of the model operators w_j and K_j) rather than just propositional formulas. This gives us the power to reason about higher-order probabilities (see Gaifman [1986] for more discussion on this subject, as well as added references) and to reason about the probability that an agent knows a certain fact.

In order to give semantics to such a language in the possible-worlds framework, we assume that, roughly speaking, at each state, each agent has a probability on the worlds he considers possible. Then a formula such as $w_i(\varphi) \geq 2w_i(\psi)$ is true at state s if, according to agent i 's probability assignment at state s , the event φ is twice as probable as ψ . For technical and philosophical reasons, we find it convenient to view the probability as being placed on an arbitrary set of worlds, rather than the set of all worlds that the agent considers possible in a given state. As we shall show by example, different choices for the probability space seem to correspond to different assumptions about the background context.

Despite the richness of the resulting language, we can combine the well-known techniques for reasoning about knowledge with the techniques for reasoning about probability introduced in [Fagin et al., 1990] to obtain an elegant complete axiomatization for the resulting language. Just as there are different assumptions we can make about the relationship between the worlds that agent i considers possible, leading to different axioms for knowledge (see Halpern and Moses [1992] for an overview), there are also different assumptions about the interrelationships between agents' probability assignment spaces at different states, which also can be captured axiomatically. We discuss these assumptions and their appropriateness, and show how these assumptions can effect the complexity of the decision procedure for the language.

This paper is related to a number of other works. We give a brief overview of the related literature here. Propositional probabilistic variants of temporal logic [Hart and Sharir, 1984; Lehman and Shelah, 1982] and dynamic logic [Feldman, 1984; Kozen, 1985] have also been studied, with the goal of analyzing probabilistic programs. Probabilistic temporal logic papers have traditionally limited the language so that the only probabilistic statements that can be made are Boolean combinations of formulas of the form “ φ occurs with probability one.” The logics studied in [Feldman, 1984; Kozen, 1985] do bear some superficial resemblance to ours in that explicit probability statements are allowed, as well as linear combinations of statements. Indeed, the probability logic considered in [Fagin et al., 1990], where the only formulas in the scope of the modal operator w are propositional formulas, is a fragment of Feldman’s logic. However, there are some fundamental differences as well, which arise from the fact that the main object of interest in these other logics are programs. As a result, our language and those used in [Feldman, 1984] and [Kozen, 1985] are incomparable. The languages used in [Feldman, 1984] and [Kozen, 1985] are richer than the one we consider here in that they allow explicit reasoning about programs, but poorer in that they can talk about the probability of only a restricted class of formulas. Moreover, there are significant technical differences in the semantics of knowledge operators (our K_i ’s) and the program operators of [Feldman, 1984] and [Kozen, 1985].

As we mentioned above, probabilistic knowledge has been an issue of great interest in the economics community. Although they have not considered formal languages containing knowledge and probability, their models can be viewed as a special case of the models we discuss in this paper. In a recent paper [Monderer and Samet, 1989] in the economics literature, Monderer and Samet investigate *probabilistic common knowledge*, a topic that shall also concern us here. We compare our framework to theirs in more detail when we discuss probabilistic common knowledge.

The framework developed in this paper has also been applied to distributed systems and cryptography in some recent papers [Fisher and Zuck 1987, 1988; Halpern et al., 1988; Halpern and Tuttle, 1993], where the issues raised here have been examined more carefully in the context of these applications areas.

Finally, we should mention two other papers that consider reasoning about knowledge and uncertainty in a possible worlds framework somewhat similar to our own. Halpern and McAllester [1989] consider a language that allows reasoning about knowledge and likelihood, but their notion of likelihood, based on the logic of likelihood of [Halpern and Rabin, 1987] considers only a qualitative notion of likelihood, rather than explicit probabilities. Although this may be appropriate for some applications, it is not useful for an analysis of protocols. Ruspini [1987] discusses certain relations that hold between knowledge and probability in the one-agent case, and relates this in turn to Dempster–Shafer *belief functions* [Shafer, 1976].

The rest of this paper is organized as follows: The next section contains a brief review of the classical possible-worlds semantics for knowledge and a discussion of how knowledge can be ascribed to processes in a distributed system. In Section 3, we describe the extended language for knowledge and probability and discuss some assumptions that can be placed on the interrelationships between agents’ probability assignments at different states. In Section 4, we give results on complete axiomatizations and decision procedures. In

Section 5, we extend the language to allow common knowledge and probabilistic common knowledge. In Section 6, we give our conclusions.

2. The Standard Kripke Model for Knowledge

In this section, we briefly review the standard S5 possible-worlds semantics for knowledge. The reader is referred to Halpern and Moses [1992] for more details.

In order to reason formally about knowledge, we need a language. Suppose we consider a system with n agents, say $1, \dots, n$, and we have a nonempty set Φ of primitive propositions about which we wish to reason. (For distributed systems applications, these will typically represent statements, such as “The value of variable x is 0”; in natural language situations, they might represent statements of the form “It is raining in San Francisco.”) For convenience, we define *true* to be an abbreviation for the formula $p \vee \neg p$, where p is a fixed primitive proposition. We abbreviate $\neg \text{true}$ by *false*. We construct more complicated formulas by closing off Φ under conjunction, negation, and the modal operators K_i , for $i = 1, \dots, n$ (where $K_i \varphi$ is read “agent i knows φ ”).

We give semantics to these formulas by means of *Kripke structures* [Kripke, 1963], which formalize the intuitions behind possible worlds. A *Kripke structure for knowledge* (for n agents) is a tuple $(S, \pi, \mathcal{R}_1, \dots, \mathcal{R}_n)$, where S is a set of *states* (thought of as states of affairs or possible worlds), $\pi(s)$ is a truth assignment to the primitive propositions of Φ for each state $s \in S$ (i.e., $\pi(s)(p) \in \{\mathbf{true}, \mathbf{false}\}$ for each primitive proposition $p \in \Phi$ and state $s \in S$), and \mathcal{R}_i is an equivalence relation on the states of S , for $i = 1, \dots, n$. The \mathcal{R}_i relation is intended to capture the possibility relation according to agent i : $(s, t) \in \mathcal{R}_i$ if in world s agent i considers t a possible world.¹ We define $\mathcal{R}_i(s) = \{s' \mid (s, s') \in \mathcal{R}_i\}$.

We now assign truth values to formulas at a state in a structure. We write $(M, s) \models \varphi$ if the formula φ is true at state s in Kripke structure M .

$$\begin{aligned} (M, s) \models p \quad (\text{for } p \in \Phi) & \text{ iff } \pi(s)(p) = \mathbf{true} \\ (M, s) \models \varphi \wedge \psi & \text{ iff } (M, s) \models \varphi \text{ and } (M, s) \models \psi \\ (M, s) \models \neg \varphi & \text{ iff } (M, s) \not\models \varphi \\ (M, s) \models K_i \varphi & \text{ iff } (M, t) \models \varphi \text{ for all } t \in \mathcal{R}_i(s). \end{aligned}$$

The last clause in this definition captures the intuition that agent i knows φ in world (M, s) exactly if φ is true in all worlds that i considers possible.

Given a structure $M = (S, \pi, \mathcal{R}_1, \dots, \mathcal{R}_n)$, we say that a formula φ is *valid* in M , and write $M \models \varphi$, if $(M, s) \models \varphi$ for every state s in S , and say that φ is *satisfiable* in M if $(M, s) \models \varphi$ for some state s in S . We say that a formula φ is *valid* if it is valid in all structures, and it is *satisfiable* if it is satisfiable in some structure. It is easy to check that a formula φ is valid in M (respectively, valid) if and only if $\neg \varphi$ is not satisfiable in M (respectively, not satisfiable).

¹We could take \mathcal{R}_i to be an arbitrary binary relation, but for distributed systems applications, taking it to be an equivalence relation seems most appropriate (see Halpern [1987] for further discussion of this point). Our results could easily be modified to deal with the general case where \mathcal{R}_i is an arbitrary binary relation.

We are often interested in characterizing by an axiom system the set of formulas that are valid. An axiom system AX is said to be *sound* for a language \mathcal{L} with respect to a class \mathcal{M} of structures if every formula in \mathcal{L} provable in AX is valid with respect to every structure in \mathcal{M} . The system AX is *complete* for \mathcal{L} with respect to \mathcal{M} if every formula in \mathcal{L} that is valid with respect to every structure in \mathcal{M} is provable in AX. We think of AX as characterizing the class \mathcal{M} if it provides a sound and complete axiomatization of that class. Soundness and completeness provide a connection between the *syntactic* notion of provability and the *semantic* notion of validity.

It is well known that the following set of axioms and inference rules, which goes back to Hintikka [1962], provides a sound and complete axiomatization for the logic of knowledge just defined with respect to the class of Kripke structures for knowledge (see Halpern and Moses [1992] for a proof).

K1. All instances of propositional tautologies.

K2. $(K_i \varphi \wedge K_i(\varphi \Rightarrow \psi)) \Rightarrow K_i \psi$.

K3. $K_i \varphi \Rightarrow \varphi$.

K4. $K_i \varphi \Rightarrow K_i K_i \varphi$.

K5. $\neg K_i \varphi \Rightarrow K_i \neg K_i \varphi$.

R1. From φ and $\varphi \Rightarrow \psi$ infer ψ (modus ponens).

R2. From φ infer $K_i \varphi$ (knowledge generalization).

We remark that this axiom system for the case of one agent has traditionally been called S5. Philosophers have spent years debating the appropriateness of this set of axioms and, indeed, of this whole approach for capturing the notion of knowledge as applied to human reasoning (see Lenzen [1978] for a review of the pertinent literature). Other axiom systems for knowledge have been considered. We mention two here, since they will arise in our later discussion: the axiom system K, consisting of K1, K2, R1, and R2, and the axiom system KD45, consisting of K1, K2, K4, K5, R1, R2 and the axiom $\neg K_i(\text{false})$. The system S5 has proved particularly useful in distributed systems applications. We now briefly review how knowledge is ascribed to processes in distributed systems. More discussion and details on the model can be found in Halpern [1987].

A distributed system consists of a collection of processes, say $1, \dots, n$, connected by a communication network. We think of these processes as running some protocol. At any time in the execution of such a protocol, the system is in some *global state*, which is a tuple of the form $\langle s_e, s_1, \dots, s_n \rangle$, where s_i is the local state of process i , and s_e is the state of the *environment*. We think of the global state as providing a “snapshot” of the state of the system at any time. The environment includes everything that we consider relevant to the system that is not described in the state of the processes. A *run* of a system is just a function from the natural numbers to global states. Intuitively, a run describes a possible execution of a system over time (where we think of the time as ranging over natural numbers). We identify a system with a set of runs (these can be thought of as the possible runs of the system when running a particular protocol). We often speak of a pair (r, m) , consisting of a run r and a time m , as a *point*. Associated with any point (r, m) we have $r(m)$, the global state of the system at this point. We can define equivalence relations \sim_i , for $i = 1, \dots, n$, on points via $(r, m) \sim_i (r', m')$ iff process i has the same local state at the global states $r(m)$ and $r'(m')$.

Suppose we fix a set Φ of primitive propositions. In distributed systems applications, we can think of these propositions as saying things like “the value of variable x is 0”, “process 1’s initial input is 3”, and so on. We define an interpreted system \mathcal{I} to be a pair (\mathcal{R}, π) , where \mathcal{R} is a system (set of runs), and π is a truth assignment to the primitive propositions of Φ at every point in \mathcal{R} . With this definition, it is easy to view an interpreted system as a Kripke structure, where the points play the role of states and the \mathcal{R}_i relation is given by \sim_i . Truth is now defined with respect to a point (r, m) in an interpreted system \mathcal{I} . In particular, we have

$$(\mathcal{I}, r, m) \models K_i \varphi \quad \text{iff} \quad (\mathcal{I}, r', m') \models \varphi \quad \text{for all} \quad (r', m') \\ \text{such that } (r', m') \sim_i (r, m).$$

Since \sim_i is an equivalence relation, it is easy to check that all the axioms of S5 hold for this interpretation of knowledge.

3. Adding Probability

The formula $K_i \varphi$ says that φ is true at all the worlds that agent i considers possible. We want to extend our language to allow formulas such as $w_i(\varphi) \geq b$, which intuitively says that “according to agent i , formula φ holds with probability at least b .” In fact, it turns out to be convenient to extend the language even further. Specifically, if $\varphi_1, \dots, \varphi_k$ are formulas, then so is $a_1 w_i(\varphi_1) + \dots + a_k w_i(\varphi_k) \geq b$, where a_1, \dots, a_k, b are arbitrary rational numbers, and $k \geq 1$. We call such a formula an *i -probability formula* (or simply a *probability formula*, if we do not wish to specify i). An expression of the form $a_1 w_i(\varphi_1) + \dots + a_k w_i(\varphi_k)$ is called a *term*. Allowing arbitrary terms in i -probability formulas, rather than just formulas of the form $w_i(\varphi) \geq \alpha$, gives us a great deal of flexibility in expressing relationships between probabilities of events. Notice we do not allow mixed formulas such as $w_i(\varphi) + w_j(\psi) \geq b$.²

We use a number of abbreviations throughout the paper for readability. For example, we use $w_i(\varphi) \geq w_i(\psi)$ as an abbreviation for $w_i(\varphi) - w_i(\psi) \geq 0$, $w_i(\varphi) \leq b$ for $-w_i(\varphi) \geq -b$, $w_i(\varphi) < b$ for $\neg(w_i(\varphi) \geq b)$, and $w_i(\varphi) = b$ for $(w_i(\varphi) \geq b) \wedge (w_i(\varphi) \leq b)$. We also use $K_i^b(\varphi)$ as an abbreviation for $K_i(w_i(\varphi) \geq b)$. Intuitively, this says that “agent i knows that the probability of φ is greater than or equal to b .” It might seem that the formula $w_i(\varphi) \geq b$ should already say that “agent i knows that the probability of φ is greater than or equal to b ”, even without the K_i operator. This is not the case under our semantics. In a given state s , the formula $w_i(\varphi)$ denotes the probability of φ according to agent i ’s probability distribution in state s . Although it may at first seem counterintuitive, it is useful to allow agent i not to know what probability distribution is being used to calculate $w_i(\varphi)$. For example, if agent i knows that one of two distributions governs φ , and according to one, the probability of φ is $1/2$ and according to the other, the probability of φ is $3/4$, then we can model this by saying that there are two states of the world that i cannot distinguish, such that $w_i(\varphi) = 1/2$ in one, and $w_i(\varphi) = 3/4$ in the other. In such a situation, it would be the case that $K_i(w_i(\varphi) \geq 1/2)$ holds.

²There would be no difficulty giving semantics to such formulas, but some of our results on decision procedures and axiomatizations seem to require that we not allow such mixed formulas. We return to this point in the next section.

The language used here extends that considered in [Fagin et al., 1990] in two ways. First, rather than having just one “probability modality” w , we have a modality w_i for each agent i , in order to allow us to reason about the probability assigned by different agents to the same event. Secondly, rather than restricting the formulas that appear in the scope of the probability modality to be propositional, we allow them to be arbitrary. In particular, we allow higher-order probability formulas such as $w_i(w_j(\varphi) \geq b) \geq c$.

Before we give formal semantics to this language, we briefly review some material from probability theory (see Halmos [1950] for more details). A *probability space* is a tuple $(\Omega, \mathcal{X}, \mu)$ where Ω is a set, called the *sample space*, \mathcal{X} is a σ -algebra of subsets of Ω (i.e., a set of subsets containing Ω and closed under complementation and countable union), whose elements are called the *measurable sets*, and a probability measure μ defined on the elements of \mathcal{X} . Note that μ does not assign a probability to all subsets of Ω , but only to the measurable sets. One natural way of attaching a weight to every subset of Ω is by considering the *inner measure* μ_* induced by μ ; if $A \subseteq \Omega$, we have

$$\mu_*(A) = \sup\{\mu(B) \mid B \subseteq A \text{ and } B \in \mathcal{X}\}.$$

Thus, the inner measure of A is essentially the measure of the largest measurable set contained in A . The properties of probability spaces guarantee that μ_* is well defined, and that if A is measurable, then $\mu_*(A) = \mu(A)$.³ Given a structure $M = (S, \pi, \mathcal{X}_1, \dots, \mathcal{X}_n)$, in order to decide whether a probability formula is true at a state s in M , we need to associate with each state s a probability space. Thus, we take a *Kripke structure for knowledge and probability* (for n agents) to be a tuple $(S, \pi, \mathcal{X}_1, \dots, \mathcal{X}_n, \mathcal{P})$, where \mathcal{P} is a *probability assignment*, which assigns to each agent $i \in \{1, \dots, n\}$ and state $s \in S$ a probability space $\mathcal{P}(i, s) = (S_{i,s}, \mathcal{X}_{i,s}, \mu_{i,s})$, where $S_{i,s} \subseteq S$. We shall usually write $\mathcal{P}(i, s)$ as $\mathcal{P}_{i,s}$. Intuitively, the probability space $\mathcal{P}_{i,s}$ describes agent i 's probabilities on events, given that the state is s . We allow $S_{i,s}$ to be an arbitrary subset of S . It might seem reasonable to take $S_{i,s} = \mathcal{X}_i(s)$, thus requiring that the agent places probability on precisely on the set of worlds he considers possible; however, as we shall see below, there are good technical and philosophical reasons to allow $S_{i,s}$ to be distinct from $\mathcal{X}_i(s)$. It is often natural to require at least that $S_{i,s}$ be a subset of $\mathcal{X}_i(s)$; we consider the consequences of this assumption below.

We can give semantics to formulas not involving probability just as before. To give semantics to i -probability formulas, assume inductively that we have defined $(M, s) \models \varphi$ for each state $s \in S$. Define $S_{i,s}(\varphi) = \{s' \in S_{i,s} \mid (M, s') \models \varphi\}$. Then, the obvious way to define the semantics of a formula such as $w_i(\varphi) \geq b$ is

$$(M, s) \models w_i(\varphi) \geq b \quad \text{iff} \quad \mu_{i,s}(S_{i,s}(\varphi)) \geq b.$$

The only problem with this definition is that the set $S_{i,s}(\varphi)$ might not be measurable (i.e., not in $\mathcal{X}_{i,s}$), so that $\mu_{i,s}(S_{i,s}(\varphi))$ might not be well defined. We discuss this issue in more detail below (and, in fact, provide sufficient conditions to guarantee that this set is measurable), but in order to deal with

³We remark that there is also a dual notion of *outer measure*; the outer measure of A , denoted $\mu^*(A)$, is essentially the measure of the smallest measurable set containing A . It is easy to see that $\mu^*(A) = 1 - \mu_*(\bar{A})$, so that the outer measure is expressible in terms of the inner measure.

this problem in general, we use the inner measures $(\mu_{i,s})_*$ rather than $\mu_{i,s}$. Thus, $w_i(\varphi) \geq b$ is true at the state s if there is some measurable set (according to agent i) contained in $S_{i,s}(\varphi)$ whose measure is at least b . More generally, we have

$$(M, s) \models a_1 w_i(\varphi_1) + \cdots + a_k w_i(\varphi_k) \geq b$$

$$\text{iff } a_1 (\mu_{i,s})_*(S_{i,s}(\varphi_1)) + \cdots + a_k (\mu_{i,s})_*(S_{i,s}(\varphi_k)) \geq b.$$

This completes the semantic definition for the whole language.

Before we discuss the properties of this language, it is helpful to consider a detailed example. This example illustrates some of the subtleties involved in choosing the probability spaces at each state.

Suppose we have two agents. Agent 2 has an input bit, either 0 or 1. He then tosses a fair coin, and performs an action a if the coin toss agrees with the input bit, that is, if the coin toss lands heads and the input bit is 1, or if the coin lands tails and the input bit is 0. We assume that agent 1 never learns agent 2's input bit or the outcome of his coin toss. An easy argument shows that according to agent 2, who knows the input bit, the probability (before he tosses the coin) of performing action a is $1/2$. There is also a reasonable argument to show that, even according to agent 1 (who does not know the input bit), the probability that the action will be performed is $1/2$. Clearly from agent 1's viewpoint, if agent 2's input bit is 0, then the probability that agent 2 performs action a is $1/2$ (since the probability of the coin landing heads is $1/2$); similarly, if agent 2's input bit is 1, then the probability of agent 2 performing action a is $1/2$. Thus, no matter what agent 2's input bit, the probability according to agent 1 that agent 2 will perform action a is $1/2$. Thus, it seems reasonable to say that agent 1 knows that the *a priori* probability of agent 2 performing action a is $1/2$. Note that we do not need to assume a probability distribution on the input bit for this argument to hold. This is a good thing: We do not want to assume that there is an input on the probability distribution, since none is provided by the problem statement. Of course, if there were a probability distribution, then this argument would hold independent of the actual probability distribution.

Now suppose we want to capture this argument in our formal system. From agent 1's point of view, there are four possibilities: $(1, h)$, $(1, t)$, $(0, h)$, $(0, t)$ (the input bit was 1 and the coin landed heads, the input bit was 1 and the coin landed tails, etc.). We can view these as the possible worlds or states in a Kripke structure. Call them s_1 , s_2 , s_3 , and s_4 , respectively; let S be the set consisting of all four states. Assume that we have primitive propositions A , H , T , B_0 , and B_1 in the language, denoting the events that action a is performed, the coin landed heads, the coin landed tails, agent 2's input bit is 0, and agent 2's input bit is 1. Thus, H is true at states s_1 and s_3 , A is true at states s_1 and s_4 , and so on. What should agent 1's probability assignment be? We now describe three plausible answers to this question.

- (1) We can associate with each state the sample space consisting of all four states, that is, all the possible worlds. This might seem to be the most natural choice, since we are taking the probability space at each state s to be $\mathcal{X}(s)$, so that at each state, agent 1 is putting the probability on the set

of states that he considers possible. Because we assume that there is no probability on the event “the input bit is 0” (respectively, “the input bit is 1”), the only candidates for measurable sets (besides the whole space and the empty set) are $\{s_1, s_3\}$ (which corresponds to the event “the coin landed heads”) and $\{s_2, s_4\}$ (“the coin landed tails”). Each of these sets has probability $1/2$. Call the resulting Kripke structure M_0 . Notice that the events $\{s_1\}$ and $\{s_2\}$ cannot both be measurable, for then the event $\{s_1, s_2\}$, which corresponds to “the input bit is 1”, would also be measurable. Similarly, we cannot take $\{s_1, s_4\}$, which corresponds to the event “action a is performed”, to be measurable. This is because if it were measurable, then, since the set of measurable sets is closed under finite intersection and complementation, each of $\{s_1\}$, $\{s_2\}$, $\{s_3\}$, and $\{s_4\}$ would have to be measurable.

- (2) We can associate with states s_1 and s_2 , where the input bit is 1, the sample space consisting only of s_1 and s_2 , with $\{s_1\}$ and $\{s_2\}$ both being measurable and having measure $1/2$. Similarly, we can associate with states s_3 and s_4 the sample space consisting only of s_3 and s_4 , with $\{s_3\}$ and $\{s_4\}$ each having measure $1/2$. Thus, when the input bit is 1, we take the sample space to consist of only those states where the input bit is 1, with the obvious probability on that space, and similarly for when the input bit is 0. Call this Kripke structure M_1 .
- (3) Finally, we can make the trivial choice of associating with each state the sample space consisting of that state alone, and giving it measure 1. Call the resulting Kripke structure M_2 .

Of the three Kripke structures above, it is easy to see that only M_1 supports the informal reasoning above. It is easy to check that we have $(M_1, s) \models K_1^{1/2}A$, for every state $s \in S$. On the other hand, in every state of M_2 , we have either $w_1(A) = 1$ (in states s_1 and s_4) or $w_1(A) = 0$ (in states s_2 and s_3). Thus, for every state $s \in S$, we have $(M_2, s) \models K_1((w_1(A) = 1) \vee (w_1(A) = 0))$ and $(M_2, s) \models \neg K_1^{1/2}A$. Finally, in M_0 , the event A is not measurable, nor does it contain any nonempty measurable sets. Thus, we have $(M_0, s) \models K_1(w_1(A) = 0)$ (where now w_1 represents the inner measure, since A is not measurable).

Does this mean that M_1 is somehow the “right” Kripke structure for this situation? Not necessarily. A better understanding can be attained if we think of this as a two-step process developing over time. At the first step, “nature” (nondeterministically) selects agent 2’s input bit. Then agent 2 tosses the coin. We can think of M_2 as describing the situation after the coin has landed. It does not make sense to say that the probability of heads is $1/2$ at this time (although it does make sense to say that the *a priori* probability of heads is $1/2$), nor does it make sense to say that the probability of performing action a is $1/2$. After the coin has landed, either it landed heads or it didn’t; either a was performed or it wasn’t. This is the intuitive explanation for why the formula $K_1((w_1(A) = 1) \vee (w_1(A) = 0))$ is valid in M_2 . M_1 describes the situation after nature has made her decision, but before the coin is tossed. Thus, agent 1 knows that either the input bit is 1 or the input bit is 0 (although he doesn’t know which one). As expected, the formula $K_1((w_1(B_0) = 1) \vee (w_1(B_1) = 0))$ holds in this situation. An (admittedly weak) argument can be made that M_0 describes the initial situation, before nature has made her decision. At this point, the event “the input bit is 0” is not measurable and we cannot attach a probability to it.

We can capture these intuitions nicely using runs. There are four runs, say r_1, r_2, r_3, r_4 , corresponding to the four states above. There are three relevant times: 0 (before nature has decided on the input bit), 1 (after nature has decided, but before the coin is tossed), and 2 (after the coin is tossed). Agent 1's local state contains only the time (since agent 1 never learns anything about the coin or the input bit); agent 2's local state contains the time, the input bit (at times 1 and 2), and the outcome of the coin toss (at time 2). We can omit the environment from the global state; everything relevant is already captured by the states of the agents. Thus, at time 1 in run r_3 , agent 1's local state is 1 (since the time is 1), while agent 2's local state is the pair $(1, 0)$, since the time is 1 and the input bit is 0. Thus, $r_3(1) = \langle 1, (1, 0) \rangle$. Similarly, we have that $r_3(2) = \langle 2, (2, 0, h) \rangle$. We now interpret the propositions A, H , etc. to mean that the action a has been or eventually will be performed, heads has been or eventually will be tossed, etc. Thus, proposition A is true at the point (r_j, k) if the action a is performed at $(r_j, 3)$. Similarly, H is true at (r_j, k) if heads is tossed in run r_j , and so on.

Clearly at each time $k = 0, 1, 2$, agent 1 considers the four points (r_j, k) , $j = 1, 2, 3, 4$, possible. At time 0, we can define the probability space at each state to make this look like M_0 . At time 1, defining the probability spaces so that we get Kripke structure M_1 seems to be appropriate, while at time 2, Kripke structure M_2 seems appropriate. Thus, although it seems that, in some sense, agent 1's knowledge about the input bit and the outcome of the coin toss does not change over time, the probability assignments used by agent 1 may change. For example, after the coin has been tossed, the probability assignment should change to reflect the fact that, although agent 1 has not learned anything about the outcome of the coin flip, he does know that the coin has been tossed.

But why and how should the fact that the coin has been tossed affect the probability assignment used by agent 1? This question is perhaps best answered in the framework discussed in [Halpern and Tuttle, 1989], where the point of view is taken that the choice of probability assignment should reflect the agent's view of the adversary it is playing against or, more accurately, the knowledge of the adversary it is playing against. Different choices of probability assignment correspond to playing adversaries with different knowledge. Suppose we play an adversary with complete information about all that has happened in the past, but who does not know the outcome of probabilistic events that will take place in the future. Thus, at time 1, the adversary does not know the outcome of the coin toss, while at time 2, he does. As shown in [Halpern and Tuttle, 1989], when agent i is playing against such an adversary, the probability assignment used by agent i should reflect what the adversary knows as well as what agent i knows. Technically, this amounts to taking the intersection of the set of possible worlds describing agent i 's knowledge with the set of possible worlds describing the adversary's knowledge. Thus, when playing against an adversary with complete information about the past, the assignment described by M_1 is appropriate at time 1, while the assignment described by M_2 is appropriate at time 2. (See Halpern and Tuttle [1989] for details of the arguments regarding appropriateness.) Interestingly, the probability assignment described by M_0 —which initially may have seemed to be the most reasonable choice of probability assignment—does not correspond to playing against an adversary in the framework of Halpern and Tuttle [1989]. In retrospect, it is the hardest probability assignment to justify.

Even in this simple example, we can already see that the decision of how to assign the probability spaces is not completely straightforward. In general, it seems that it will depend in more detail on the form of the analysis. This example already shows that in general at a state s , we do not want to take $S_{i,s} = \mathcal{K}_i(s)$. Note that $S_{i,s} = \mathcal{K}_i(s)$ only in M_0 above; in particular, in M_1 , where we can carry out the informal reasoning which says that action a occurs with probability $1/2$, we have $S_{1,s}$ as a strict subset of $\mathcal{K}_1(s)$.⁴ Although in this example, we do not want $S_{i,s} = \mathcal{K}_i(s)$, we do want $S_{i,s} \subseteq \mathcal{K}_i(s)$. This is quite a natural condition. Without it, it is possible that an agent can place positive probability on a fact that he knows to be false; for example, the formula $K_i \neg p \wedge w_i(p) > 0$ is consistent. We would view an agent who places positive probability on an event he knows to be false as inconsistent. Thus, we term the following condition CONS (for *consistent*).

CONS. For all i and s , if $\mathcal{P}_{i,s} = (S_{i,s}, \mathcal{K}_{i,s}, \mu_{i,s})$, then $S_{i,s} \subseteq \mathcal{K}_i(s)$.

Note that CONS does not imply that $s \in S_{i,s}$; an agent is not required to view the state that he is in as one of the set of states in his probability space. Although it may seem unusual, there are times (in particular, when analyzing *asynchronous* distributed systems), when it turns out to be appropriate not to require that $s \in S_{i,s}$ [Halpern and Tuttle, 1989].

In some applications, although the agents have different sets of points they consider possible, it is useful to model them as agreeing on what the probability space is at each point. In this case, we say that the probability assignment is *objective*. This is a quite natural assumption in contexts where all the probabilistic events are common knowledge, for example, if there is a global coin. Alternatively, in the framework of Halpern and Tuttle [1989], it is appropriate if the agents are viewed as all playing the same adversary, who has at least as much knowledge as each of the agents individually. Note that, under this assumption, the intersection of the set of states describing agent i 's knowledge with the set of states describing the adversary's knowledge is the same for all i . This means that, according to the framework of Halpern and Tuttle [1989], the agents should all use the same probability assignment. Note that this assumption is appropriate, for example, if the agents all play an adversary who has complete information about the global state of the system, they would agree on what the appropriate probability space should be.⁵

In the context of a Kripke structure for knowledge and probability, having an objective probability assignment corresponds to the following condition:

OBJ. $\mathcal{P}_{i,s} = \mathcal{P}_{j,s}$ for all i, j , and s .

Note that if we had required that $S_{i,s} = \mathcal{K}_i(s)$ for each agent i and each state s , then OBJ could hold only in Kripke structures where $\mathcal{K}_i(s) = \mathcal{K}_j(s)$ for all agents i and j and all states s .

⁴The example presented here is a simplification of one given by Mark Tuttle. It was Mark who first pointed out to us that it is not always appropriate to take $S_{i,s} = \mathcal{K}_i(s)$.

⁵Mark Tuttle and Yoram Moses first pointed out to us that in distributed systems applications, an appropriate choice is often an objective probability with the probability space consisting of all the points with the same global state. This approach was first taken in [Halpern, 1988]. See Fisher and Zuck [1988] and Halpern and Tuttle [1989] for further discussion on the appropriate choice of probability assignment in distributed systems.

We now consider some other assumptions about the interrelationship between an agent's probability assignments at different states. A rather natural assumption to make on the choice of probability space is that it is the same in all worlds the agent considers possible. In the context of distributed systems, this would mean that an agent's probability space is determined by his local state. We call this property SDP (*state-determined probability*). Formally, we have:

SDP. For all i , s , and t , if $t \in \mathcal{X}_i(s)$, then $\mathcal{P}_{i,s} = \mathcal{P}_{i,t}$.

Of the three Kripke structures we considered above, only M_0 satisfies SDP. As the discussion in [Halpern and Tuttle, 1989] shows, SDP is most natural in situations where no nondeterministic (or, perhaps better, *nonprobabilistic*) choices are made by "nature". (In our example, the choice of the agent's input bit is nonprobabilistic; the outcome of the coin toss is probabilistic.) SDP is an assumption that has often been made. Indeed, it is implicitly assumed in much of the economists' work (e.g., [Aumann, 1976; Cave, 1983]). In these papers, it is assumed that each agent initially defines a probability space over the sample space of all worlds. Thus, for each agent i , we have a probability space $\mathcal{P}_i = (S, \mathcal{X}_i, \mu_i)$, where S is the set of all worlds.⁶ Agent i 's probability of an event e at a state s is taken to be the conditional probability of e given agent i 's set of possible worlds. This means that $\mathcal{P}_{i,s} = (\mathcal{X}_i(s), \mathcal{X}_{i,s}, \mu_{i,s})$, where $\mathcal{X}_{i,s} = \{A \cap \mathcal{X}_i(s) \mid A \in \mathcal{X}_i\}$, and $\mu_{i,s}(A \cap \mathcal{X}_i(s)) = \mu_i(A) / \mu_i(\mathcal{X}_i(s))$.⁷ Note that the resulting Kripke structure has the SDP property.

Although M_1 and M_2 in our example above do not satisfy SDP, they do satisfy a weaker property that we call *uniformity*. Roughly speaking, uniformity holds if we can partition $\mathcal{X}_i(s)$ into subsets such that at every point in a given subset T , the probability space is the same. Formally, we say uniformity holds if:

UNIF. For all i , s , and t , if $\mathcal{P}_{i,s} = (S_{i,s}, \mathcal{X}_{i,s}, \mu_{i,s})$ and $t \in S_{i,s}$, then $\mathcal{P}_{i,t} = \mathcal{P}_{i,s}$.

Notice that UNIF does not require that $S_{i,s} \subseteq \mathcal{X}_i(s)$; thus, in order to be able to partition $\mathcal{X}_i(s)$ into subsets such that at every point in a given subset T , the probability space is the same, we require both UNIF and CONS. Uniformity arises in a natural way when considering appropriate probability assignments in distributed systems. Each subset of $S_{i,s}$ turns out to correspond to the result of "nature" making a particular nonprobabilistic choice, just as is the case in the structure M_1 in our example (see Halpern and Tuttle [1989] for details). Uniformity also has some interesting connections with a well-studied principle regarding higher-order probabilities called *Miller's principle* [Miller, 1966; Skyrms, 1980]; we comment on this in a little more detail below. Note that CONS and SDP together imply UNIF, and that all the structures in our example above satisfy UNIF.

There is one last property of interest to us, which seems to have been assumed in all previous papers involving reasoning about probability, and that

⁶Aumann actually assumes that there is an objective probability on the whole space, so that $\mathcal{P}_i = \mathcal{P}_j$ for all agents i and j . This corresponds to the agents having a common prior distribution.

⁷This approach runs into slight technical difficulties if $\mathcal{X}_i(s)$ is not measurable, or has measure 0. However, it is always assumed that this is not the case.

is that all formulas define measurable sets. As shown in [Fagin et al., 1990] (and as we shall see again below), reasoning about probability is simplified if we assume that all formulas define measurable sets. More precisely, we say that formulas define measurable sets in M if

MEAS. For all i and s and for every formula φ , the set $S_{i,s}(\varphi) \in \mathcal{P}_{i,s}$.

(Recall that $S_{i,s}(\varphi) = \{s' \in S_{i,s} \mid (M, s') \models \varphi\}$.)

Clearly if primitive propositions define measurable sets, then all propositional formulas define measurable sets. However, there is no particular reason to expect that a probability formula such as $w_i(p) + w_i(q) \geq 1/2$ will define a measurable set (in fact, it is easy to show that in general it will not). Let PMEAS be the property which says that all primitive propositions define measurable sets. (Note that PMEAS does not hold in M_0 , but does hold in M_1 and M_2 .) The following lemma describes sufficient conditions for MEAS to hold.

LEMMA 3.1. *If M is a structure satisfying CONS, OBJ, UNIF, and PMEAS, then M satisfies MEAS.*

PROOF. A straightforward induction on the structure of formulas φ shows that $S_{i,s}(\varphi)$ is measurable for all formulas φ . The assumptions CONS and OBJ together imply that for all agents i and j , we have $S_{i,s} \subseteq \mathcal{K}_j(s)$, so it is easy to see that $S_{i,s}(K_j\varphi)$ is either $S_{i,s}$ or \emptyset . In either case, it is measurable. Similarly, we can show that OBJ and UNIF together imply that for any probability formula φ , we have that $S_{i,s}(\varphi)$ is either $S_{i,s}$ or \emptyset . \square

It seems that OBJ, UNIF, and PMEAS are often reasonable assumptions in distributed systems applications, so this lemma is of more than just pure technical interest.

We close this section by briefly considering one more property of probabilities that has appeared in the literature. *Miller's principle* is an axiom that connects higher-order probabilities (that is, probabilities on probabilities) with probabilities on formulas [Miller, 1966; Skyrms, 1980]. It says:

$$w_i(\varphi) \geq bw_i(w_i(\varphi) \geq b).$$

It is easy to see that, in general, this axiom does not hold in structures for knowledge and probability. However, it is not hard to show that our condition UNIF implies this axiom. In systems satisfying UNIF, we have either (a) $(w_i(\varphi) \geq b)$ is false at state s , in which case UNIF implies that $w_i(w_i(\varphi) \geq b) = 0$ at state s , or (b) $(w_i(\varphi) \geq b)$ is true at state s , in which case UNIF implies that $w_i(w_i(\varphi) \geq b) = 1$ at state s . In either case, it is easy to see that Miller's principle holds. It turns out that there is a precise sense in which Miller's principle completely characterizes uniform structures; see Halpern [1991] for details.

4. Complete Axiomatizations and Decision Procedures

We now describe a natural complete axiomatization for the logic of probability and knowledge. The axiom system can be modularized into four components. The first component allows us to do propositional reasoning, the second allows us to reason about knowledge, the third allows us to reason about inequalities (so it contains axioms that allow us to deduce, for example, that $2x \geq 2y$

follows from $x \geq y$), while the fourth is the only one that has axioms and inference rules for reasoning about probability.

I. Axiom and rule for propositional reasoning

Axiom K1 and rule R1 from Section 2

II. Axioms and rule for reasoning about knowledge

Axioms K2–K5 and rule R2 from Section 2

For reasoning about inequalities, we need a system that allows us to prove all valid formulas about linear inequalities; one particular system that will do the trick is given in [Fagin et al., 1990]. We repeat it here.

III. Axioms for reasoning about linear inequalities

- 11.** $(a_1w_i(\varphi_1) + \cdots + a_kw_i(\varphi_k) \geq b) \Leftrightarrow (a_1w_i(\varphi_1) + \cdots + a_kw_i(\varphi_k) + 0w_i(\varphi_{k+1}) \geq b)$ (adding and deleting 0 terms).
- 12.** $(a_1w_i(\varphi_1) + \cdots + a_kw_i(\varphi_k) \geq b) \Rightarrow (a_{j_1}w_i(\varphi_{j_1}) + \cdots + a_{j_k}w_i(\varphi_{j_k}) \geq b)$, if j_1, \dots, j_k is a permutation of $1, \dots, k$ (permutation).
- 13.** $(a_1w_i(\varphi_1) + \cdots + a_kw_i(\varphi_k) \geq b) \wedge (a'_1w_i(\varphi_1) + \cdots + a'_kw_i(\varphi_k) \geq b') \Rightarrow (a_1 + a'_1)w_i(\varphi_1) + \cdots + (a_k + a'_k)w_i(\varphi_k) \geq (b + b')$ (addition of coefficients).
- 14.** $(a_1w_i(\varphi_1) + \cdots + a_kw_i(\varphi_k) \geq b) \Leftrightarrow (da_1w_i(\varphi_1) + \cdots + da_kw_i(\varphi_k) \geq db)$ if $d > 0$ (multiplication of nonzero coefficients).
- 15.** $(t \geq b) \vee (t \leq b)$ if t is a term (dichotomy).
- 16.** $(t \geq b) \Rightarrow (t > c)$ if t is a term and $b > c$ (monotonicity).

Finally, we need axioms for reasoning about probability. The axioms we take are also given in [Fagin et al., 1990]; they are simply a translation of the standard axioms for probability in finite domains to our language.

IV. Axioms for reasoning about probabilities

- W1.** $w_i(\varphi) \geq 0$ (nonnegativity).
- W2.** $w_i(\text{true}) = 1$ (the probability of the event *true* is 1).
- W3.** $w_i(\varphi \wedge \psi) + w_i(\varphi \wedge \neg \psi) = w_i(\varphi)$ (additivity).
- W4.** $w_i(\varphi) = w_i(\psi)$ if $\varphi \Leftrightarrow \psi$ is a proposition tautology (distributivity).
- W5.** $w_i(\text{false}) = 0$ (the probability of the event *false* is 0).⁸

Axiom W3 corresponds to finite additivity. Although we allow infinite domains, as noted in [Fagin et al., 1990], we do not need an axiom that corresponds to countable additivity. Indeed, we could not even express such an axiom in our language. Roughly speaking, we can get away with finite additivity because we can show that if a formula is satisfiable at all, it is satisfiable in a finite domain.

Things get more complicated if we drop the measurability assumption. It is easy to check that in this case, W3 is no longer sound. As shown in [Fagin et al., 1990], there is another axiom with which we can replace W3 to get a complete axiomatization. This axiom is also the key axiom that characterizes

⁸Axiom W5 is actually redundant. It is included, since it will be needed later when we replace axiom W3 by another axiom in the nonmeasurable case.

belief functions in the Dempster–Shafer approach to reasoning about uncertainty.

$$\mathbf{W6.} \quad w_i(\varphi_1 \vee \cdots \vee \varphi_k) \geq \sum_{I \subseteq \{1, \dots, k\}, I \neq \emptyset} (-1)^{|I|+1} w_i(\bigwedge_{i \in I} \varphi_i).$$

Although this axiom may appear somewhat mysterious, note that if we replace \geq by $=$, then in the measurable case, this becomes an instance of the well-known *inclusion-exclusion rule* for probabilities [Feldman, 1984].

It turns out that if we replace W3 by W6, we get a complete axiomatization for i -probability formulas in the nonmeasurable case. (See Fagin et al. [1990] for more details, as well as proofs of soundness and completeness.)

Let AX_{MEAS} consist of K1–K5, I1–I6, W1–W5, and R1–R2. Let AX be the result of replacing W3 in AX_{MEAS} by W6. The following theorem says that these axiomatizations are sound and complete.⁹

THEOREM 4.1. *AX (respectively, AX_{MEAS}) is a sound and complete axiomatization for the logic of knowledge and probability (respectively, for structures satisfying MEAS).*

PROOF. Soundness is straightforward, as usual, so we focus on completeness. We sketch the proof for the measurable case; the nonmeasurable case follows the same lines.

In order to prove completeness, we need only show that if the formula φ is consistent with AX_{MEAS} , then φ is satisfiable in a Kripke structure for knowledge and probability satisfying MEAS. Let $Sub(\varphi)$ be the set of all subformulas of φ , and let $Sub^+(\varphi)$ be the set of subformulas of φ and their negations.

Let s be a finite set of formulas, and let φ_s be the conjunction of the formulas in s . We say that s is *consistent* if it is not the case that $AX_{MEAS} \vdash \neg \varphi_s$, where as usual, we write $AX_{MEAS} \vdash \psi$ if the formula ψ is provable in the axiom system AX_{MEAS} . The set s is a *maximal* consistent subset of $Sub^+(\varphi)$ if it is consistent, a subset of $Sub^+(\varphi)$, and for every subformula ψ of φ , includes one of ψ or $\neg \psi$. (Note that it cannot include both, for then it would not be consistent.) Following Makinson [1966] (see also Halpern and Moses [1992]), we first construct a Kripke structure for knowledge (but not probability) $(S, \pi, \mathcal{K}_1, \dots, \mathcal{K}_n)$ as follows: We take S , the set of states, to consist of all maximal consistent subsets of $Sub^+(\varphi)$. If s and t are states, then $(s, t) \in \mathcal{K}_i$ precisely if s and t contain the same formulas of the form $K_i \psi$. We define π so that for a primitive proposition p , we have $\pi(s)(p) = \mathbf{true}$ iff p is one of the formulas in the set s . Our goal is to define a probability assignment \mathcal{P} such that if we consider the Kripke structure for knowledge and probability $M = (S, \pi, \mathcal{K}_1, \dots, \mathcal{K}_n, \mathcal{P})$, then for every state $s \in S$ and every formula $\psi \in Sub^+(\varphi)$, we have $(M, s) \models \psi$ iff $\psi \in s$.

We now sketch the techniques from Fagin et al. [1990] required to do this. It is easy to see that the formulas φ_s are provably mutually exclusive for $s \in S$; that is, $AX_{MEAS} \vdash \varphi_s \Rightarrow \neg \varphi_t$ for $s \neq t$. Indeed, the proof uses only propositional reasoning, namely K1 and R1. Moreover, again using only propositional reasoning, we can show that $AX_{MEAS} \vdash \psi \Leftrightarrow \bigvee_{\{s \in S \mid \psi \in s\}} \varphi_s$, for all $\psi \in$

⁹The proofs of the technical results in this section presume familiarity with the results of Fagin et al. [1990], and with standard proofs of completeness and complexity for modal logics (cf. [Halpern and Moses, 1992; Ladner, 1977]).

$\text{Sub}^+(\varphi)$. Using these observations, we can show, using W1–W5, that $\mu_i(\psi) = \sum_{\{s \in S \mid \psi \in s\}} \mu_i(\varphi_s)$ is provable in AX_{MEAS} (cf. [Fagin et al., 1990, Lemma 2.3]).¹⁰ Using this fact together with I1 and I3, we can show that an i -probability formula $\psi \in \text{Sub}^+(\varphi)$ is provably equivalent to a formula of the form $\sum_{s \in S} c_s \mu_i(\varphi_s) \geq b$, for some appropriate coefficients c_s .

Fix an agent i and a state $s \in S$. We now describe a set of linear equalities and inequalities corresponding to i and s , over variables of the form $x_{iss'}$, for $s' \in S$. We can think of $x_{iss'}$ as representing $\mu_{i,s}(s')$, that is, the probability of state s' under agent i 's probability distribution at state s . We have one inequality corresponding to every i -probability formula ψ in $\text{Sub}^+(\varphi)$. Assume that ψ is equivalent to $\sum_{s' \in S} c_{s'} \mu_i(\varphi_{s'}) \geq b$. Observe that exactly one of ψ and $\neg\psi$ is in s . If $\psi \in s$, then the corresponding inequality is

$$\sum_{s' \in S} c_{s'} x_{iss'} \geq b.$$

If $\neg\psi \in s$, then the corresponding inequality is

$$\sum_{s' \in S} c_{s'} x_{iss'} < b.$$

Finally, we have the equality

$$\sum_{s' \in S} x_{iss'} = 1.$$

As shown in Theorem 2.2 in Fagin et al. [1990], since φ_s is consistent, this set of linear equalities and inequalities has a solution $x_{iss'}^*, s' \in S$.

For each i and s , we solve the corresponding set of inequalities separately. We now define \mathcal{P} so that $\mathcal{P}_{i,s} = (S, 2^S, \mu_{i,s})$, where if $A \subseteq S$, then $\mu_{i,s}(A) = \sum_{s' \in A} x_{iss'}^*$. Since $\sum_{s' \in S} x_{iss'}^* = 1$, it is easy to see that $\mu_{i,s}$ is indeed a probability measure. Note that, in the probability space $\mathcal{P}_{i,s}$, every set is measurable. This probability assignment does not necessarily satisfy CONS; it may well be the case that there are sets disjoint from $\mathcal{N}_i(s)$ that are assigned positive measure under $\mu_{i,s}$.

As we said above, we now want to show that for every formula $\psi \in \text{Sub}^+(\varphi)$ and every state in s , we have $(M, s) \models \psi$ iff $\psi \in s$. The proof proceeds by induction on ψ . If ψ is a primitive proposition, the result is immediate from the definition of π . The cases where ψ is a negation or a conjunction are straightforward and left to the reader. The case where ψ is an i -probability formula follows immediately from the arguments above, since the appropriate inequality corresponding to ψ is satisfied by $\mu_{i,s}$. Finally, if ψ is of the form $K_i\psi'$, the proof proceeds using well-known arguments from modal logic (cf. [Hughes and Cresswell, 1968; Halpern and Moses, 1992]). We sketch a few of the details here. If $K_i\psi' \in s$, then, by construction of \mathcal{N}_i , for all $t \in \mathcal{N}_i(s)$, we have $K_i\psi' \in t$. Since t is a maximal consistent subset of $\text{Sub}^+(\varphi)$, it must be the case that one of ψ' or $\neg\psi'$ is in t . From Axiom K3, it follows that it must in fact be ψ' . By the induction hypothesis, we have that $(M, t) \models \psi'$. Since this argument holds for all $t \in \mathcal{N}_i(s)$, we have that $(M, s) \models K_i\psi'$.

¹⁰Note that this proof makes crucial use of W3; this formula is not provable using the axiom system AX.

Now suppose that $(M, s) \models K_i \psi'$. We want to show that $K_i \psi' \in s$. Let s' be the subset of s consisting of all formulas in s of the form $K_i \psi''$ or $\neg K_i \psi''$. Notice that, in particular, s' includes one of $K_i \psi'$ or $\neg K_i \psi'$; we plan to show that in fact it must include $K_i \psi'$. We claim that

$$\text{AX}_{\text{MEAS}} \vdash \varphi_{s'} \Rightarrow \psi'. \quad (1)$$

For suppose not. Then $\varphi_{s'} \wedge \neg \psi'$ is consistent. Thus, there is a maximal consistent subset of $\text{Sub}^+(\varphi)$, say t , that includes $s' \cup \{\neg \psi'\}$. But then, by construction of \mathcal{N}_i , we have $(s, t) \in \mathcal{N}_i$, and by the induction hypothesis, $(M, t) \models \neg \psi'$. But this contradicts our assumption that $(M, s) \models K_i \psi'$. Thus, (1) holds.

By R2, from (1), we have

$$\text{AX}_{\text{MEAS}} \vdash K_i(\varphi_{s'} \Rightarrow \psi'). \quad (2)$$

Using A2 and propositional reasoning, it follows that

$$\text{AX}_{\text{MEAS}} \vdash K_i \varphi_{s'} \Rightarrow K_i \psi'. \quad (3)$$

Every conjunct of $\varphi_{s'}$ is of the form $K_i \psi''$ or $\neg K_i \psi''$. Thus, if σ is one of the conjuncts of $\varphi_{s'}$, using either Axiom A4 or A5, it follows that

$$\text{AX}_{\text{MEAS}} \vdash \sigma \Rightarrow K_i \sigma. \quad (4)$$

It is well known (and can be proved using K1, K2, R1, and R2) that for any formulas σ_1 and σ_2 , we have

$$\text{AX}_{\text{MEAS}} \vdash K_i(\sigma_1 \wedge \sigma_2) \Leftrightarrow K_i \sigma_1 \wedge K_i \sigma_2.$$

Thus, from (4), it follows that

$$\text{AX}_{\text{MEAS}} \vdash \varphi_{s'} \Rightarrow K_i \varphi_{s'}. \quad (5)$$

From (3) and (5), we now get

$$\text{AX}_{\text{MEAS}} \vdash \varphi_{s'} \Rightarrow K_i \psi'.$$

Since $\varphi_{s'}$, and hence $\varphi_{s'}$, is consistent, it now follows that $\neg K_i \psi'$ cannot be one of the conjuncts of $\varphi_{s'}$, and hence that $K_i \psi' \in s$, as desired.

If φ is consistent, it must be in one of the maximal consistent subsets of $\text{Sub}^+(\varphi)$. Thus, it follows that if φ is consistent, then it is satisfiable in the structure M . This complete the proof in the measurable case.

Note that the proof shows the modularity of the axiom system. In order to deal with i -probability formulas, we just need the axioms for reasoning about probability and inequalities (together with propositional reasoning); the axioms for reasoning about knowledge play no role. Similarly, in order to deal with knowledge formulas, we just used the axioms for reasoning about knowledge.

This modularity is important when it comes to dealing with the nonmeasurable case. We must now replace the arguments above for constructing \mathcal{P} by analogous arguments from Theorem 3.8 of Fagin et al. [1990] for the nonmeasurable case. As these arguments show, it is not quite the case that we can construct a Kripke structure satisfying φ whose set of states is the set S above consisting of maximal consistent subsets of $\text{Sub}^+(\varphi)$. Rather, we need to make copies of each of the maximal consistent sets. Thus, for each maximal consistent set s , there will be states s_1, \dots, s_n (as shown in [Fagin et al., 1990], we can take $n \leq |\text{Sub}(\varphi)|$). We can now define a probability assignment \mathcal{P} on this set; it will no longer be the case that in the probability space $\mathcal{P}_{i,S}$, all sets are

measurable. Modulo this change to \mathcal{P} , we can construct a Kripke structure M for knowledge and probability such that for each state s_k corresponding to a maximal consistent set s and each formula $\psi \in \text{Sub}^+(\varphi)$, we have $(M, s_k) \models \psi$ iff $\psi \in s$. The proof follows identical lines to that of the measurable case. The only change comes in dealing with i -probability formulas. Again, this is done by constructing a collection of linear equalities and inequalities that $\mu_{i,s}$ must satisfy, in the proof of Theorem 3.7 of Fagin et al. [1990]. We omit further details here. \square

We can also capture some of the assumptions we made about systems axiomatically. In a precise sense (as we shall see), CONS corresponds to the axiom

W7. $K_i \varphi \Rightarrow (w_i(\varphi) = 1)$.

This axiom essentially tells us that the set of states that agent i considers possible has measure 1 (according to agent i).

OBJ corresponds to the axiom

W8. $(a_1 w_i(\varphi_1) + \dots + a_k w_i(\varphi_k) \geq b) \Rightarrow (a_1 w_j(\varphi_1) + \dots + a_k w_j(\varphi_k) \geq b)$.

Axiom W8 says that each i -probability formula implies the corresponding j -probability formula. This is clearly sound if we have an objective probability distribution.

UNIF corresponds to the axiom

W9. $\varphi \Rightarrow (w_i(\varphi) = 1)$ if φ is an i -probability formula or the negation of an i -probability formula.

Since a given i -probability formula has the same truth value at all states where agent i 's probability assignments is the same, the soundness of W9 in structures satisfying UNIF is easy to verify.

SDP corresponds to the axiom:

W10. $\varphi \Rightarrow K_i \varphi$ if φ is an i -probability formula or the negation of an i -probability formula.

Since SDP says that agent i knows the probability space (in that it is the same for all states in $\mathcal{S}_i(s)$), it is easy to see that SDP implies that in a given state, agent i knows all i -probability formulas that are true in that state. Axioms W7 and W10 together imply W9, which is reasonable since CONS and SDP together imply UNIF.

The next theorem proves our claims about correspondence between various properties and various axioms.

THEOREM 4.2. *Let \mathcal{A} be a subset of $\{\text{CONS}, \text{OBJ}, \text{UNIF}, \text{SDP}\}$ and let A be the corresponding subset of $\{\text{W7}, \text{W8}, \text{W9}, \text{W10}\}$. Then $AX \cup A$ (respectively, $AX_{\text{MEAS}} \cup A$) is a sound and complete axiomatization for the logic of knowledge and probability for structures satisfying \mathcal{A} (respectively, $\text{MEAS} \cup \mathcal{A}$).¹¹*

PROOF. Again, soundness is straightforward, so we focus on completeness. We obtain completeness in each case by a relatively straightforward modification of the proof of Theorem 4.1. We just sketch the details here.

¹¹Although it is straightforward to extend Theorem 4.1 to the case where we have mixed formulas of the form $w_i(\varphi) + w_j(\psi) \geq b$ (with appropriate modifications to axioms I1, I2, I3, and I4), the situation seems much more complicated in the presence of the properties UNIF and SDP. It is due to these complexities that we do not allow such mixed formulas in our language.

First, consider the relationship between CONS and axiom W7. Assume that W7 is included as an axiom. In this case, it is easy to see that we can modify our construction in the proof of Theorem 4.1 so that we can take $\mathcal{P}_{i,s} = (S_{i,s}, \mathcal{X}_{i,s}, \mu_{i,s})$ such that $S_{i,s} \subseteq \mathcal{X}_i(s)$. We sketch the details in the measurable case. Recall that in this case, in the proof of Theorem 4.1, we took $\mathcal{P}_{i,s} = (S, 2^S, \mu_{i,s})$, so that all sets were measurable, and $\mu_{i,s}$ was defined in terms of a solution to a set of linear equalities and inequalities. Now we claim that in the presence of W7, we can show that if $s \in S$ and $s' \notin \mathcal{X}_i(s)$, then $\varphi_s \Rightarrow (\mu_i(\varphi_{s'}) = 0)$ is provable. To see this, observe that $\varphi_s \Rightarrow K_i(\neg \varphi_{s'})$ is provable using K4 or K5. Thus, applying W7, we have that $\varphi_s \Rightarrow (\mu_i(\neg \varphi_{s'}) = 1)$ is provable. Finally, by using W2, W3, and W4, it is not hard to show that $\varphi_s \Rightarrow (\mu_i(\varphi_{s'}) = 0)$ is provable. As a consequence, we can augment the linear system of equalities and inequalities defining $\mu_{i,s}$ by adding $x_{i,s'} = 0$ for $s' \notin \mathcal{X}_i(s)$. The proof that shows that the consistency of φ_s implies that the original linear system was satisfiable can easily be extended to show that the augmented system must now be satisfiable in the presence of W7. Again, we can use the solution to this system to define $\mu_{i,s}$. Since $x_{i,s'} = 0$ for $s' \notin \mathcal{X}_i(s)$, we can take $S_{i,s} \subseteq \mathcal{X}_i(s)$, so that CONS is satisfied. Note that if W7 is the only additional axiom, then we can take $S_{i,s}$ to be $\mathcal{X}_i(s)$; as we shall see, some of the other axioms may force $S_{i,s}$ to be a strict subset of $\mathcal{X}_i(s)$.

Now consider the relationship between OBJ and axiom W8. Assume that W8 is included as an axiom. It is easy to see that the subscripts in i -probability formulas can then be ignored (i.e., we can think of $w_i(\varphi)$ as simply $w(\varphi)$). Thus, we can easily modify the construction of Theorem 4.1 so as to take $\mathcal{P}_{i,s} = \mathcal{P}_{i,s}$. This guarantees that the Kripke structure for knowledge and probability that we construct satisfies OBJ.

Next, consider the relationship between UNIF and axiom W9. Assume that W9 is included as an axiom. Let $T_i(s)$ be the set of states that contain precisely the same i -probability formulas and negations of i -probability formulas as s . Just as we showed that the presence of W7 allowed us to assume without loss of generality that $S_{i,s}$ is a subset of $\mathcal{X}_i(s)$, we show that the presence of W9 allows us to assume that $S_{i,s}$ is a subset of $T_i(s)$. Again, we consider only the measurable case here. Suppose that $s' \notin T_i(s)$. Then s and s' disagree on some i -probability formula, say ψ . Without loss of generality, $\psi \in s$ and $\psi \notin s'$. Thus, $\varphi_s \Rightarrow \psi$ and $\psi \Rightarrow \neg \varphi_{s'}$ are both provable. Since, by W9, $\psi \Rightarrow (\mu_i(\psi) = 1)$ is provable, it easily follows using the axioms of probability and propositional reasoning that $\varphi_s \Rightarrow (\mu_i(\varphi_{s'}) = 0)$ is provable. Thus, we can augment the linear system of equalities and inequalities defining $\mu_{i,s}$ by adding $x_{i,s'} = 0$ for $s' \notin T_i(s)$. Just as in our proof of the relationship between W7 and CONS, we can now show that we can take $T_i(s)$ to be a subset of $S_{i,s}$. Now note that for each $t \in S_{i,s}$, we must have $T_i(s) = T_i(t)$. Since, as the proof of Theorem 4.1 shows, the definition of $\mu_{i,t}$ depends only on the i -probability formulas and negations of i -probability at state t , it follows that we can take $\mathcal{P}_{i,t} = \mathcal{P}_{i,s}$ for all $t \in T_i(s)$. Thus, UNIF holds. We remark that if our only additional axiom is W9, then we can actually take $S_{i,s} = T_i(s)$; however, if we have both W7 and W9, then by combining the arguments used in each case, we can show that we can take $S_{i,s}$ to be $\mathcal{X}_i(s) \cap T_i(s)$.

Finally, consider the relationship between SDP and axiom W10. Assume that W10 is included as an axiom. We want to show that we can slightly modify the construction of Theorem 4.1 so that if $t \in \mathcal{X}_i(s)$, then $t \in T_i(s)$. This is trivial

to do: We just change the definition of the \mathcal{R}_i relation so that $t \in \mathcal{R}_i(s)$ iff it is the case both that $t \in T_i(s)$ and that s and t contain all the same subformulas of the form $K_i\psi$. With this change, we can assume without loss of generality that if $t \in \mathcal{R}_i(s)$, then $\mathcal{P}_{i,s} = \mathcal{P}_{i,t}$, since, as we have already noted, the definition of $\mu_{i,t}$ depends only on the i -probability formulas and negations of i -probability formulas at state t . Now for the structure M constructed in this way, we still want to show (in the measurable case) that $(M, s) \models \psi$ iff $\psi \in S$. The proof is almost identical to that given for Theorem 4.1. There is only one case where we must be a little careful: when proving that if $(M, s) \models K_i\psi'$, then $K_i\psi' \in s$. Rather than taking s' to be the subset of s consisting of all formulas in s of the form $K_i\psi''$ or $\neg K_i\psi''$, we now extend it to consist of all these formulas together with all i -probability formulas or negations of i -probability formulas in s . With this change, the proof now proceeds as before. It is still the case that for every formula $\sigma \in s'$, we have that $\sigma \Rightarrow K_i\sigma$ is provable; for formulas of the form $K_i\psi''$ we use K4, for formulas of the form $\neg K_i\psi''$ we use K5, and if σ is an i -probability formula or the negation of an i -probability formula, we use W10.

Let \mathcal{A} be a subset of {CONS, OBJ, UNIF, SDP}, and let A be the corresponding subset of {W7, W8, W9, W10}. If A is included among the axioms, then our discussion shows that given a consistent formula φ , we can modify our original construction of a Kripke structure for knowledge and probability satisfying φ to get a Kripke structure that not only satisfies φ , but also the conditions in \mathcal{A} . This proves completeness. \square

As is often the case in modal logics, the ideas in our completeness proof can be extended to get a small model property and a decision procedure. In order to state our results here, we need a few definitions. Recall that $\text{Sub}(\varphi)$ is the set of all subformulas of φ . It is easy to see that an upper bound on the size $|\text{Sub}(\varphi)|$ of $\text{Sub}(\varphi)$ is the number of symbols in φ , where we treat a rational number as a single symbol. We also define the size of a Kripke structure $(S, \pi, \mathcal{R}_1, \dots, \mathcal{R}_n, \mathcal{P})$ to be the number of states in S . (Note that the size of a Kripke structure may be infinite.)

THEOREM 4.3. *Let \mathcal{A} be a subset of {MEAS, CONS, OBJ, UNIF, SDP}. The formula φ is satisfiable in a Kripke structure satisfying \mathcal{A} iff it is satisfiable in a Kripke structure satisfying \mathcal{A} of size at most $|\text{Sub}(\varphi)|2^{|\text{Sub}(\varphi)|}$ (or just $2^{|\text{Sub}(\varphi)|}$ if $\text{MEAS} \in \mathcal{A}$).*

PROOF. We need only show that the Kripke structure for knowledge and probability constructed in the proof of Theorem 4.2 is no bigger than the size given in the statement of this theorem. If $\text{MEAS} \in \mathcal{A}$, then the set of states is simply the set of maximal consistent subsets of $\text{Sub}^+(\varphi)$. Now a subformula of φ and its negation cannot both be in a maximal consistent subset, so the cardinality of a maximal consistent subset is at most equal to $|\text{Sub}(\varphi)|$. Hence, the number of states in the Kripke structure for knowledge and probability constructed in the proof of Theorem 4.2 is at most $2^{|\text{Sub}(\varphi)|}$.

If $\text{MEAS} \notin \mathcal{A}$, then, as we mentioned in the proof of Theorem 4.1, we cannot take the states of the Kripke structure satisfying φ to be the maximal consistent subsets of $\text{Sub}^+(\varphi)$. Rather, we must make copies of the sets. As shown in [Fagin et al., 1990], we need to make at most $|\text{Sub}(\varphi)|$ copies of each state, so the size of the resulting structure is at most $|\text{Sub}(\varphi)|2^{|\text{Sub}(\varphi)|}$. \square

It can be shown that this result is essentially optimal, in that there is a sequence of formulas $\varphi_1, \varphi_2, \dots$ and a constant $c > 0$ such that (1) $|\text{Sub}(\varphi_k)| \leq ck$, (2) φ_k is satisfiable, and (3) φ_k is satisfiable only in a structure of size at least 2^k .¹² Indeed, this exponential lower bound holds even when there is only one agent. However, if we assume that CONS and either UNIF or SDP hold, then we can get polynomial-sized models in the case of one agent.

THEOREM 4.4. *If the formula φ talks about the knowledge and probabilities of only one agent and \mathcal{A} is a subset of $\{\text{MEAS}, \text{CONS}, \text{OBJ}, \text{UNIF}, \text{SDP}\}$ containing CONS and either UNIF or SDP, then φ is satisfiable in a structure satisfying \mathcal{A} iff φ is satisfiable in a structure of size polynomial in $|\text{Sub}(\varphi)|$ satisfying \mathcal{A} .*

PROOF. Let $M = (S, \pi, \mathcal{K}_1, \mathcal{P})$ be a structure satisfying \mathcal{A} where φ is satisfiable. For each $s \in S$, let $\mathcal{P}_{1,s} = (S_{1,s}, \mathcal{X}_{1,s}, \mu_{1,s})$. Since CONS is in \mathcal{A} , we know that $S_{1,s} \subseteq \mathcal{K}_1(s)$ for each $s \in S$. Without loss of generality, we can assume that \mathcal{K}_1 is a single equivalence class, that is, that $\mathcal{K}_1 = S \times S$. (PROOF. Suppose that $(M, s) \models \varphi$. Let $M' = (S', \pi', \mathcal{K}'_1, \mathcal{P}')$, where S' is the equivalence class of \mathcal{K}_1 that includes s , and let π', \mathcal{K}'_1 , and \mathcal{P}' be the restrictions of π, \mathcal{K}_1 , and \mathcal{P} , respectively, to S' . It is easy to see that $(M', s) \models \varphi$, and \mathcal{K}'_1 is a single equivalence class by construction.) Since CONS and SDP together imply UNIF, and since \mathcal{A} contains CONS and either UNIF or SDP, it follows that M satisfies UNIF. Observe that it follows that no two distinct probability spaces $\mathcal{P}_{1,s}$ have overlapping sample spaces; that is, if $\mathcal{P}_{1,s} = (S_{1,s}, \mathcal{X}_{1,s}, \mu_{1,s})$ and $\mathcal{P}_{1,t} = (S_{1,t}, \mathcal{X}_{1,t}, \mu_{1,t})$, and if $\mathcal{P}_{1,s} \neq \mathcal{P}_{1,t}$, then $S_{1,s} \cap S_{1,t} = \emptyset$. This is because if $u \in S_{1,s} \cap S_{1,t}$, then by UNIF we have $\mathcal{P}_{1,u} = \mathcal{P}_{1,s}$ and $\mathcal{P}_{1,u} = \mathcal{P}_{1,t}$, so $\mathcal{P}_{1,s} = \mathcal{P}_{1,t}$.

We now describe a small probability space $\mathcal{P}'_{1,s}$ (one whose sample space has cardinality at most $|\text{Sub}(\varphi)|$) that we shall later “replace” $\mathcal{P}_{1,s}$ with. For each state s , let Σ_s be the set of formulas $\sigma \in \text{Sub}^+(\varphi)$ such that $(M, s) \models \sigma$. By techniques of Fagin et al. [1990] (see Theorem 2.4 for the measurable case, and Theorem 3.4 for the general case), for each state s , there is a probability space $\mathcal{P}'_{1,s} = (S'_{1,s}, \mathcal{X}'_{1,s}, \mu'_{1,s})$ where

- (1) $S'_{1,s} \subseteq S_{1,s}$,
- (2) the cardinality of $S'_{1,s}$ is at most $|\text{Sub}(\varphi)|^2$ (in the measurable case, $|S'_{1,s}| \leq |\text{Sub}(\varphi)|$)
- (3) if we interpret $w_i(\psi)$ to mean the inner measure of the set of states s where $\psi \in \Sigma_s$, for each $\psi \in \text{Sub}^+(\varphi)$, then each of the i -probability formulas and negations of i -probability formulas of Σ_s is satisfied, and
- (4) if MEAS is in \mathcal{A} , then every subset of $S'_{1,s}$ is measurable (i.e., a member of $\mathcal{X}'_{1,s}$).

Let $s_0 \in S$ be a state such that $(M, s_0) \models \varphi$. For each formula $\sigma \in \Sigma_{s_0}$ of the form $\neg K_1 \psi$, we select some state t_σ such that $(M, t_\sigma) \models \neg \psi$ (there is such a state t_σ , since $(M, s_0) \models \neg K_1 \psi$). Let T consist of s_0 , along with each of these states t_σ . Note that the cardinality of T is at most $1 + |\text{Sub}(\varphi)|$. Define $M' = (S', \pi', \mathcal{K}'_1, \mathcal{P}')$ by letting S' be the union of the sample spaces of $\mathcal{P}'_{1,s}$

¹²The idea is that φ_k forces a structure to contain a binary tree of depth k . In fact, the result follows from the corresponding result for the modal logic K (cf. [Halpern and Moses, 1992; Ladner, 1977]). Without any assumptions on the probability assignment, $w_i(\varphi) = 1$ acts like the \square operator. We omit details here.

for each $s \in T$, by letting π' be π restricted to S' , by letting $\mathcal{K}'_1 = S' \times S'$, and letting $\mathcal{P}'(1, s) = \mathcal{P}'_{1,s}$. It is straightforward to show that $(M', s_0) \models \varphi$, that M' satisfies \mathcal{A} , and that M' is of size polynomial in $|\text{Sub}(\varphi)|$. \square

We now consider the complexity of decision procedures for the validity problem. The difficulty of deciding whether φ is valid will be a function of the length of φ , written $|\varphi|$. In computing this length, we also include the length of the coefficients in probability terms. (Since all coefficients are rational, the length of the coefficient is just the sum of the lengths of the numerator and denominator, written in binary.)

THEOREM 4.5. *Let \mathcal{A} be a subset of $\{\text{MEAS}, \text{CONS}, \text{OBJ}, \text{UNIF}, \text{SDP}\}$. If $\text{CONS} \in \mathcal{A}$, but it is not the case that UNIF or SDP is in \mathcal{A} , then the validity problem with respect to structures satisfying \mathcal{A} is complete for exponential time (i.e., that is an algorithm that decides if a formula φ is valid in all structures satisfying \mathcal{A} that runs in time exponential in $|\varphi|$, and every exponential time problem can be reduced to the validity problem). If $\text{CONS} \notin \mathcal{A}$ or UNIF or SDP is in \mathcal{A} , then the validity problem with respect to structures satisfying \mathcal{A} is complete for polynomial space.*

PROOF. The proof requires combining techniques for proving upper and lower bounds on the complexity of the validity problem for logics of knowledge and logics of probability, as discussed in [Halpern and Moses, 1992] and [Fagin et al., 1990], respectively. We briefly sketch the main ideas here, referring the reader to [Halpern and Moses, 1992] and [Fagin et al., 1990] for further details.

The polynomial space lower bound follows from the polynomial space lower bound for logics of knowledge alone [Halpern and Moses, 1992]. For the exponential time lower bounds, let $B_i\varphi$ be an abbreviation of $w_i(\varphi) = 1$. We can view B_i as a modal operator, just like K_i . If UNIF or SDP is in \mathcal{A} , then it can be shown that B_i satisfies the axioms of the modal system KD45, but without these assumptions, B_i is unconstrained (in particular, it satisfies only the axioms of the modal system K). If CONS is in \mathcal{A} , then everything “reachable probabilistically” is also considered possible. More formally, suppose we have a sequence of states s_0, s_1, \dots, s_k such that s_k is reachable probabilistically from s_0 , as far as agent 1 is concerned; that is, s_{j+1} is in S_{1,s_j} and $\mu_{1,s_j}(\{s_{j+1}\}) > 0$, for $0 \leq j \leq k-1$. Then CONS implies that $(s_0, s_k) \in \mathcal{K}'_1$. As a consequence, it is not hard to show that B_1 and K_1 can be used to essentially simulate the $[a]$ and $[a^*]$ operators in Propositional Dynamic Logic (PDL). Since the validity problem for PDL is exponential-time complete [Fischer and Ladner, 1979], we get the exponential time lower bound if CONS is in \mathcal{A} , but neither UNIF nor SDP is. Note that the lower bound holds even with only one agent.

In the cases, where we claim a polynomial space upper bound, this is shown by proving that if a formula φ is satisfiable at all, it is satisfiable in a structure that looks like a tree, with polynomial branching and depth no greater than the depth of nesting of K_i and w_i operators in φ . The result now follows along similar lines to corresponding results for logics of knowledge.

Finally, the exponential time upper bound follows by showing that if a formula is satisfiable at all, it is satisfiable in an exponential-size model, that can be constructed in deterministic exponential time; the technique is similar to that used to show that logics of knowledge with common knowledge are

decidable in deterministic exponential time [Halpern and Moses, 1992], or that PDL is decidable in deterministic exponential time [Pratt, 1979]. \square

Again, if we restrict attention to the case of one agent and structures satisfying CONS and either UNIF or SDP, then we can do better. In fact, the complexity of the validity problem is no worse than that for propositional logic.

THEOREM 4.6. *Let \mathcal{A} be a subset of $\{MEAS, CONS, OBJ, UNIF, SDP\}$ containing CONS and either UNIF or SDP. For the case of one agent, the validity problem with respect to structures satisfying \mathcal{A} is co-NP-complete.*

PROOF. We show that the satisfiability problem is NP-complete. It follows that the validity problem is co-NP-complete. The lower bound is immediate, since clearly the logic is at least as hard as propositional logic. For the upper bound, by Theorem 4.4, φ is satisfiable in a structure satisfying \mathcal{A} iff φ is satisfiable in a structure M of size polynomial in $|\text{Sub}(\varphi)|$ satisfying \mathcal{A} . It might seem that this suffices to complete the proof: We simply guess a polynomial-sized structure satisfying φ . However, there is one additional subtlety: In order to describe the polynomial-sized structure, we have to describe the probabilities of all of its subsets. *A priori*, this might take us far more than polynomial space.

By results of Fagin et al. [1990], we can assume without loss of generality that M has the property that for each state s in M and agent i , the probability assigned to every measurable subset of $S_{i,s}$ in the probability space $\mathcal{P}_{i,s}$ is a rational number a/b , such that the length of a and b is linear in $|\varphi|$. If MEAS is in \mathcal{A} , we can assume even more, namely that every subset of $S_{i,s}$ is measurable. This means we can describe the probability space $\mathcal{P}_{i,s}$ by describing the probability of each point. If MEAS $\notin \mathcal{A}$, then we cannot assume that every subset of $S_{i,s}$ is measurable. Instead, we describe the probability space $\mathcal{P}_{i,s}$ by describing the probabilities of the *basis sets*, that is, the nonempty measurable sets none of whose proper nonempty subsets are a measurable set. Since every measurable set is the disjoint union of basis sets, this again completely describes $\mathcal{P}_{i,s}$. In either case, we get a polynomial-sized description of the structure M . Thus, in order to check if φ is satisfiable, we just guess a structure M with a probability-sized description that satisfies it. This gives us an NP procedure for checking satisfiability. \square

5. Adding Common Knowledge

For many of our applications, we need to reason not only about what an individual process knows, but about what everyone in a group knows, or what everyone in a group knows that everyone else in the group knows. *Common knowledge* can be viewed as the state of knowledge where everyone knows, everyone knows that everyone knows, everyone knows that everyone knows that everyone knows, etc.

It is easy to extend our language so that we can reason about common knowledge. We add modal operators E_G (where G is a subset of $\{1, \dots, n\}$) and C_G , where $E_G \varphi$ and $C_G \varphi$ are read “everyone in the group G knows φ ” and “ φ is common knowledge among the group G ”, respectively.

$$(M, s) \models E_G \varphi \quad \text{iff} \quad (M, s) \models K_i \varphi \quad \text{for all } i \in G,$$

$$(M, s) \models C_G \varphi \quad \text{iff} \quad (M, s) \models E_G^k \varphi \quad \text{for all } k \geq 1,$$

where $E_G^1\varphi$ is an abbreviation for $E_G\varphi$, and $E_G^{k+1}\varphi$ is an abbreviation for $E_G E_G^k\varphi$.

It is well known (again, see Halpern and Moses [1992]) that we can get a complete axiomatization for the language of knowledge and common knowledge by adding the following axioms and rule of inference to the axiom system described in Section 2:

- C1.** $E_G\varphi \Leftrightarrow \bigwedge_{i \in G} K_i\varphi$
- C2.** $(C_G\varphi \wedge C_G(\varphi \Rightarrow \psi)) \Rightarrow C_G\psi$
- C3.** $C_G\varphi \Leftrightarrow E_G(\varphi \wedge C_G\varphi)$
- RC1.** From $\varphi \Rightarrow E_G\varphi$ infer $\varphi \Rightarrow C_G\varphi$.


Axiom C3, called the *fixed-point axiom*, says that $C_G\varphi$ can be viewed as a fixed point of the equation $X \Leftrightarrow E_G(\varphi \wedge X)$. In fact, with a little work it can be shown to be the greatest fixed point of this equation, that is, it is implied by all other fixed points. For most of our applications, it is the fixed-point characterization of common knowledge that is essential to us (see Halpern and Moses [1990] for a discussion of fixed points). The rule of inference RC1 is called the induction rule. The reason is that from the fact that $\varphi \Rightarrow E_G\varphi$ is valid, we can easily show by induction on k that $\varphi \Rightarrow E_G^k\varphi$ is valid for all k . It follows that $\varphi \Rightarrow C_G\varphi$ is valid. In fact, the same proof can be used to show that for any structure M , if $\varphi \Rightarrow E_G\varphi$ is valid in M , then $\varphi \Rightarrow C_G\varphi$ is valid in M (see Halpern and Moses [1990] for further discussion of these points).

It is perhaps not surprising that if we augment AX_{MEAS} with the axioms for common knowledge, we get a complete axiomatization for the language of knowledge, common knowledge, and probability for structures satisfying MEAS. If we want to deal with nonmeasurable structures, we must use the axiom system AX rather than AX_{MEAS} . And again we get small model theorems and an exponential-time complete decision procedure (regardless of what additional assumptions among MEAS, OBJ, UNIF, and SDP we make). The proofs involve a combination of the techniques for dealing with common knowledge, and the techniques for probability introduced in [Fagin et al., 1990] and the previous section. We omit details here.

In [Halpern and Moses, 1990], it was observed that common knowledge is often not attainable in practical distributed systems, but weaker variants of it are. One obvious variant to consider is a probabilistic variant (indeed, this was already mentioned as something to consider in [Halpern and Moses, 1990]). Recall that we defined $K_i^b\varphi$ to be an abbreviation for $K_i(w_i(\varphi) \geq b)$. We now extend our syntax to allow modal operators of the form E_G^b and C_G^b . We define

$$(M, s) \models E_G^b\varphi \quad \text{iff} \quad (M, s) \models K_i^b\varphi \quad \text{for all } i \in G.$$

By analogy to $C_G\varphi$, we want $C_G^b\varphi$ to be the greatest fixed point of the equation $X \Leftrightarrow E_G^b(\varphi \wedge X)$. The obvious analogue to the definition of $C_G\varphi$, namely, $E_G^b\varphi \wedge (E_G^b)^2\varphi \wedge \dots$ does not work. For example, consider a structure M for knowledge and probability defined as follows: There are four states in M , say s_1, s_2, s_3, s_4 . Agent 1 cannot distinguish s_1 from s_2 and cannot distinguish s_3 from s_4 , while agent 2 cannot distinguish s_1 from s_3 and cannot distinguish s_2 from s_4 . Thus, \mathcal{R}_1 is the reflexive symmetric closure of $\{(s_1, s_2), (s_3, s_4)\}$ (i.e., \mathcal{R}_1 is the least relation containing (s_1, s_2) and (s_3, s_4) that is reflexive and symmetric) while \mathcal{R}_2 is the reflexive symmetric closure of $\{(s_1, s_3), (s_2, s_4)\}$. We take the primitive proposition p to be true at s_2 and s_3 ,

FIG. 1. The Kripke structure M .

and false at s_1 and s_4 (so that $\pi(s_1)(p) = \text{false}$, etc.). The Kripke structure M is sketched in Figure 1 (where reflexive loops between states are ignored).

We assume that M is an SDP structure (so that $S_{i,y} = \mathcal{X}_i(s)$). We take $\mathcal{P}_{1,s_1} = \mathcal{P}_{1,s_2}$ to be the probability space that assigns probability $1/2$ to both s_1 and s_2 . Similarly, $\mathcal{P}_{2,s_1} = \mathcal{P}_{2,s_3}$ is a probability space where both s_1 and s_3 have probability $1/2$. On the other hand, we take $\mathcal{P}_{1,s_3} = \mathcal{P}_{1,s_4}$ and $\mathcal{P}_{2,s_2} = \mathcal{P}_{2,s_4}$ to be such that the probability of s_4 is 1. Take $G = \{1, 2\}$, and let ψ be the infinite conjunction $E_G^{1/2}p \wedge E_G^{1/2}E_G^{1/2}p \wedge \dots$. It is now easy to check that (a) $(M, s_1) \models \psi$, (b) $(M, s_2) \models \neg E_G^{1/2}p$, and (c) $(M, s_3) \models \neg E_G^{1/2}p$. Since $(M, s_1) \not\models p$, it follows that none of s_1 , s_2 , or s_3 satisfy $p \wedge E_G^{1/2}p$. Thus, $(M, s_1) \not\models E_G^{1/2}(p \wedge E_G^{1/2}p)$, so $(M, s_1) \not\models E_G^{1/2}(p \wedge \psi)$. In particular, this means that ψ does not satisfy the fixed-point equation $X \Leftrightarrow E_G^{1/2}(p \wedge X)$.

However, a slight variation does work. Define $(F_G^b)^0\varphi = \text{true}$ and $(F_G^b)^{k+1}\varphi = E_G^b(\varphi \wedge (F_G^b)^k\varphi)$. Then, we take

$$(M, s) \models C_G^b\varphi \quad \text{iff} \quad (M, s) \models (F_G^b)^k\varphi \quad \text{for all } k \geq 1.$$

We remark that this actually is a generalization of the nonprobabilistic case. The reason is that if we define $F_G^0\varphi = \text{true}$ and $F_G^{k+1}\varphi = E_G(\varphi \wedge F_G^k\varphi)$, then we get $F_G^k\varphi \Leftrightarrow E_G^k\varphi$ (since both $E_G(\varphi \wedge \psi) \Leftrightarrow E_G\varphi \wedge E_G\psi$ and $E_G\varphi \Rightarrow \varphi$ are valid). The analogous facts do not hold once we add probabilities, as we have already observed.¹³

The following lemma shows that this definition indeed does have the right properties:

LEMMA 5.1. $C_G^b\varphi$ is the greatest fixed-point solution of the equation $X \Leftrightarrow E_G^b(\varphi \wedge X)$.

PROOF. We first show that $C_G^b\varphi$ is a fixed-point solution of the equation, that is, that $C_G^b\varphi \Leftrightarrow E_G^b(\varphi \wedge C_G^b\varphi)$ is valid. One implication is straightforward: if $E_G^b(\varphi \wedge C_G^b\varphi)$ holds at (M, s) , then so does $E_G^b(\varphi \wedge (F_G^b)^k\varphi)$ for each k , since $C_G^b\varphi \Rightarrow (F_G^b)^k\varphi$ is valid. That is, $(F_G^b)^{k+1}\varphi$ holds for each k at (M, s) , and so $C_G^b\varphi$ holds at (M, s) . As for the other implication, assume that $C_G^b\varphi$ holds at (M, s) . Hence, $(F_G^b)^{k+1}\varphi$, that is, $E_G^b(\varphi \wedge (F_G^b)^k\varphi)$, holds at (M, s) for each k . For each agent i , let $A_{i,k}$ be the set of states in $S_{i,y}$ where $\varphi \wedge (F_G^b)^k\varphi$ holds, for $k = 1, 2, \dots$. Since $(M, s) \models E_G^b(\varphi \wedge (F_G^b)^k\varphi)$, it follows that $(\mu_{i,s})_*(A_{i,k})$

¹³It is interesting to note that the infinite conjunction $E_G^b\varphi \wedge (E_G^b)^2 \wedge \dots$ is a solution to a slightly different fixed-point equation, namely $X \Leftrightarrow E_G^b\varphi \wedge E_G^bX$. This is the definition taken by Monderer and Samet [1980]. Both definitions are generalizations of the nonprobabilistic case, since, as we observed above, $E_G(\varphi \wedge X)$ is equivalent to $E_G\varphi \wedge E_GX$, so $C_G\varphi$ is also a solution to the fixed-point equation $E_G\varphi \wedge E_GX$. The two definitions are quite similar. Which is the right one to use seems to depend on the application. Our definition seems to be somewhat more appropriate in analyzing probabilistic coordinated attack and Byzantine agreement protocols [Halpern and Tuttle, 1993].

$\geq b$, for each agent i and for all k . It is a standard result of probability theory that there exists $B_{i,k} \subseteq A_{i,k}$ such that $B_{i,k}$ is measurable and $\mu_{i,s}(B_{i,k}) \geq b$ [Halmos, 1950; Neveu, 1964]. It is straightforward to verify, by induction on k , that $(F_G^b)^{k+1}\varphi \Rightarrow (F_G^b)^k\varphi$ is valid. (PROOF. The case $k = 0$ is easy, since $(F_G^b)^0\varphi =_{\text{def}} \text{true}$. For the inductive step, note that the validity of $(F_G^b)^{k+1}\varphi \Rightarrow (F_G^b)^k\varphi$ implies the validity of $E_G^b(\varphi \wedge (F_G^b)^{k+1}\varphi) \Rightarrow E_G^b(\varphi \wedge (F_G^b)^k\varphi)$. But this last formula is precisely $(F_G^b)^{k+2}\varphi \Rightarrow (F_G^b)^{k+1}\varphi$.) Thus, we have $A_{i,1} \supseteq A_{i,2} \supseteq A_{i,3} \supseteq \dots$. Without loss of generality, we can assume also that $B_{i,1} \supseteq B_{i,2} \supseteq B_{i,3} \supseteq \dots$ (since we can always replace $B_{i,k}$ by the union of $B_{i,k'}$ for $k' \geq k$). The set $B_{i,\infty} = \bigcap_{i=1}^{\infty} B_{i,k}$ is a measurable set; it is easy to see that it must have measure at least b . By construction, $\varphi \wedge C_G^b(\varphi)$ holds at $B_{i,\infty}$. It thus follows that $E_G^b(\varphi \wedge C_G^b(\varphi))$ holds at (M, s) , as desired.

We now show that $C_G^b\varphi$ is the greatest fixed point. Assume that ψ is a fixed point in a structure M , that is, that $M \models \psi \Leftrightarrow E_G^b(\varphi \wedge \psi)$. We want to show that $M \models \psi \Rightarrow C_G^b\varphi$. We first show, by induction on k , that $M \models \psi \Rightarrow (F_G^b)^k\varphi$. Since $(F_G^b)^0\varphi =_{\text{def}} \text{true}$ by definition, the result is immediate in the case of $k = 0$. For the induction step, suppose $M \models \psi \Rightarrow (F_G^b)^m\varphi$. It follows easily that $M \models E_G^b(\varphi \wedge \psi) \Rightarrow E_G^b(\varphi \wedge (F_G^b)^m\varphi)$. Hence, since $M \models \psi \Leftrightarrow E_G^b(\varphi \wedge \psi)$, we must also have $M \models \psi \Rightarrow E_G^b(\varphi \wedge (F_G^b)^m\varphi)$. But $(F_G^b)^{m+1}\varphi =_{\text{def}} E_G^b(\varphi \wedge (F_G^b)^m\varphi)$. So $M \models \psi \Rightarrow (F_G^b)^{m+1}\varphi$. This completes the inductive step. It now follows that if $(M, s) \models \psi$, then $(M, s) \models (F_G^b)^k\varphi$ for all k , and hence that $(M, s) \models C_G^b\varphi$. Thus, $M \models \psi \Rightarrow C_G^b\varphi$. This proves that $C_G^b\varphi$ is the greatest fixed point of the equation $X \Leftrightarrow E_G^b(\varphi \wedge X)$. \square

It is now easy to check that we have the following analogues to the axioms for E_G and C_G .

CP1. $E_G^b\varphi \Leftrightarrow \bigwedge_{i \in G} K_i^b\varphi$.

CP2. $C_G^b\varphi \Leftrightarrow E_G^b(\varphi \wedge C_G^b\varphi)$.

RCP1. From $\psi \Rightarrow E_G^b(\psi \wedge \varphi)$ infer $\psi \Rightarrow C_G^b\varphi$.

We remark that these axioms and rule of inference are sound in all structures for knowledge and probability. And again, we can actually show the following strengthening of RCP1: For any structure M , if $\psi \Rightarrow E_G^b(\psi \wedge \varphi)$ is valid, in M then $\psi \Rightarrow C_G^b\varphi$ is valid in M .

It can be shown that these axioms and inference rule, together with the axioms and inference rules C1–C3 and RC1 for common knowledge discussed above and AX_{MEAS} (respectively, AX) gives us a sound and complete axiomatization for this extended language in the measurable case (respectively, in the general case). Moreover, we can prove a small model theorem, and show that the validity problem for all variants of the logic is complete for exponential time. These proofs are quite difficult; we hope to provide the details in a later paper.

6. Conclusions

We have investigated a logic of knowledge and probability that allows explicit reasoning about probability. We have been able to obtain complete axiomatizations and decision procedures for our logic. We have also identified some important properties that might hold of the interrelationship between agents' probability assignments at different states.

It seems to us that the most important area for further research lies in having a better understanding of what the appropriate choice of probability

space is. Some discussion of this issue appears in [Fischer and Zuck, 1988]; a more general treatment appears in [Halpern and Tuttle, 1993]. Using the ideas in this paper together with Moses' recent work [1988] on resource-bounded reasoning, Moses, Tuttle, and Halpern have made progress on capturing *interactive proofs* and *zero knowledge* [Goldwasser et al., 1989] in the framework of knowledge and probability discussed in this paper. These results appear in [Halpern et al., 1988]. The analysis in [Halpern et al., 1988] is done using the same style of probability assignment as in our examples in Section 3, that is, they take $S_{i,(r,m)}$ to consist of all points with the same global state as (r, m) . This probability assignment satisfies OBJ and UNIF, but not necessarily SDP. Although this is not the only choice of probability assignment that is reasonable in this context, there are good grounds for believing that no reasonable choice will satisfy SDP. If there are nonprobabilistic events in a system as well as probabilistic events, SDP seems inappropriate. As we said earlier, a general framework for deciding which choice of probability assignment is appropriate, presented in terms of adversaries, appears in [Halpern and Tuttle, 1993].

As this discussion may suggest, although our understanding of the subtle interaction between knowledge and probability is increasing, more work needs to be done in this area. It would be especially useful to have a larger body of examples on which to test our ideas. The economics and game theory literature may be a good source for such examples. We expect that further progress can be made by combining the intuitions from both computer science and game theory.

ACKNOWLEDGMENTS. The foundations of this paper were greatly influenced by discussions the second author had with Yoram Moses and Mark Tuttle in the context of their joint work on capturing interactive proofs [Halpern et al., 1988]. In particular, their observation that it was necessary to allow $S_{i,s}$ to be a subset of $\mathcal{A}_i(s)$ caused us to rethink many of our ideas. They also suggested taking $K_i^b\varphi$ to be an abbreviation for $K_i(w_i(\varphi) \geq b)$ rather than $w_i(\varphi) \geq b$, as was done in an early draft of this paper. As usual, Moshe Vardi's comments helped improve both the style and content of the paper. Finally, we would like to thank an anonymous referee for a close reading of the paper and many useful comments and suggestions that improved the presentation of the results.

REFERENCES

- AUMANN, R. J., 1976. Agreeing to disagree. *Ann. Stat.* 4, 6, 1236–1239.
- CAVE, J., 1983. Learning to agree. *Econ. Lett.* 12, 147–152.
- FAGIN, R., HALPERN, J. Y., AND MEGIDDO, N., 1990. A logic for reasoning about probabilities. *Inf. Comput.* 87, 1/2, 78–128.
- FELDMAN, Y., 1984. A decidable propositional probabilistic dynamic logic with explicit probabilities. *Inf. Control*, 63, 11–38.
- FELLER, W., 1957. *An Introduction to Probability Theory and its Applications*, volume 1. Wiley, New York.
- FISCHER, M. J. AND LADNER, R. E., 1979. Propositional dynamic logic of regular programs. *J. Comput. Syst. Sci.* 18, 2, 194–211.
- FISCHER, M. J. AND ZUCK, L. D., 1987. Relative knowledge and belief (extended abstract). Tech. Rep. YALEU/DCS/TR-589. Yale Univ.
- FISCHER, M. J. AND ZUCK, L. D., 1988. Reasoning about uncertainty in fault-tolerant distributed systems. Tech. Rep. YALEU/DCS/TR-643. Yale Univ.
- GAIFMAN, H., 1986. A theory of higher order probabilities. In *Theoretical Aspects of Reasoning about Knowledge: Proceedings of the 1986 Conference*, J. Y. Halpern, ed., Morgan-Kaufmann, San Mateo, Calif., pp. 275–292.

- GOLDWASSER, S., MICALI, S., AND RACKOFF, C., 1989. The knowledge complexity of interactive proof systems. *SIAM J. Comput.* 18, 1 (Feb.), 186–208.
- HALMOS, P., 1950. *Measure Theory*. Van Nostrand, New York.
- HALPERN, J. Y., 1987. Using reasoning about knowledge to analyze distributed systems. In *Annual Review of Computer Science*, vol. 2. J. F. Traub, B. J. Grosz, B. W. Lampson, and N. J. Nilsson, eds. Annual Reviews Inc., Palo Alto, Calif., pp. 37–68.
- HALPERN, J. Y., 1991. The relationship between knowledge, belief, and certainty. *Ann. Math. Artif. Int.* 4, 301–322.
- HALPERN, J. Y. AND MCALLESTER, D. A., 1989. Knowledge, likelihood, and probability. *Computat. Int.* 5, 151–160.
- HALPERN, J. Y. AND MOSES, Y., 1990. Knowledge and common knowledge in a distributed environment. *J. ACM* 37, 3 (July), 549–587.
- HALPERN, J. Y. AND MOSES, Y., 1992. A guide to completeness and complexity for modal logics of knowledge and belief. *Artif. Int.* 54, 319–379.
- HALPERN, J. Y., MOSES, Y., AND TUTTLE, M. R., 1988. A knowledge-based analysis of zero knowledge. In *Proceedings of the 20th ACM Symposium on Theory of Computing* (Chicago, Ill., May 2–4). ACM, New York, pp. 132–147.
- HALPERN, J. Y. AND RABIN, M. O., 1987. A logic to reason about likelihood. *Artif. Int.* 32, 3, 379–405.
- HALPERN, J. Y. AND TUTTLE, M. R., 1993. Knowledge, probability, and adversaries. *J. ACM* 40, 4 (Sept.), 917–962.
- HART, S. AND SHARIR, M., 1984. Probabilistic temporal logics for finite and bounded models. In *Proceedings of the 16th ACM Symposium on Theory of Computing* (Washington, DC, Apr. 30–May 2). ACM, New York, pp. 1–13.
- HINTIKKA, J., 1962. *Knowledge and Belief*. Cornell University Press, Ithaca, NY.
- HUGHES, G. E. AND CRESSWELL, M. J., 1968. *An Introduction to Modal Logic*. Methuen, London.
- KOZEN, D., 1985. Probabilistic PDL. *J. Comput. Syst. Sci.* 30, 162–178.
- KRIPKE, S., 1963. A semantical analysis of modal logic. I: Normal modal propositional calculi. *Z. Math. Logik Grundl. Math.* 9, 67–96. (Announced in *J. Symb. Logic* 24, 1959, p. 323.)
- LADNER, R. E., 1977. The computational complexity of provability in systems of modal propositional logic. *SIAM J. Comput.* 6, 3, 467–480.
- LEHMANN, D. AND SHELAH, S., 1982. Reasoning about time and change. *Inf. Control* 53, 165–198.
- LENZEN, W., 1978. Recent work in epistemic logic. *Acta Phil. Fen.* 30, 1–219.
- MAKINSON, D., 1966. On some completeness theorems in modal logic. *Z. Math. Logik Grundl. Math.* 12, 379–384.
- MILLER, D., 1966. A paradox of information. *Brit. J. Phil. Sci.*, 17.
- MONDERER, D. AND SAMET, D., 1989. Approximating common knowledge with common beliefs. *Games and Economic Behavior* 1, 170–190.
- MOORE, R. C., 1985. A formal theory of knowledge and action. In *Formal Theories of the Commonsense World*, J. Hobbs and R. C. Moore, eds. Ablex Publishing Corp., Norwood, N.J., pp. 319–358.
- MOSES, Y., 1988. Resource-bounded knowledge. In *Proceedings of the 2nd conference on Theoretical Aspects of Reasoning about Knowledge*. M. Y. Vardi, ed., Morgan-Kaufmann, San Mateo, Calif., pp. 261–276.
- NEVEU, J., 1964. *Bases Mathematiques du Calcul des Probabilités*. Mason.
- NILSSON, N., 1986. Probabilistic logic. *Artif. Int.* 28, 71–87.
- PRATT, V. R., 1979. Models of program logics. In *Proceedings of the 20th IEEE Symposium on Foundations of Computer Science*. IEEE, New York, pp. 115–122.
- RUSPINI, E. H., 1987. Epistemic logics, probability, and the calculus of evidence. In *Proceedings of the 10th International Joint Conference on Artificial Intelligence (IJCAI-87)*, pp. 924–931.
- SHAFFER, G., 1976. *A Mathematical Theory of Evidence*. Princeton University Press, Princeton, N.J.
- SKYRMS, B., 1980. Higher order degrees of belief. In *Prospects for Pragmatism: Essays in Honor of F. P. Ramsey*. D. H. Mellor, ed., Cambridge University Press, Cambridge, U.K.

RECEIVED APRIL 1990; REVISED AUGUST 1992; ACCEPTED OCTOBER 1992