Likelihood, probability, and knowledge[1]

JOSEPH Y. HALPERN
IBM Research Division, Almaden Research Center, San Jose, CA 95120, U.S.A.
AND
DAVID A. MCALLESTER
MIT Artificial Intelligence Laboratory, Cambridge, MA 02139, U.S.A.

Received August 13, 1986
Revision accepted September 25, 1989

The modal logic LL was introduced by Halpern and Rabin as a means of doing qualitative reasoning about likelihood. Here the relationship between LL and probability theory is examined. It is shown that there is a way of translating probability assertions into LL in a sound manner, so that LL in some sense can capture the probabilistic interpretation of likelihood. However, the translation is subtle; several more obvious attempts are shown to lead to inconsistencies. We also extend LL by adding modal operators for knowledge. This allows us to reason about the interaction between knowledge and likelihood. The propositional version of the resulting logic LLK is shown to have a complete axiomatization and to be decidable in exponential time, provably the best possible.

Key words: qualitative reasoning about likelihood, relating probability and likelihood, combining knowledge and likelihood, modal logic.

La logique modale LL a été proposée par Halpern et Rabin comme moyen de procéder à un raisonnement qualitatif à propos de la vraisemblance. Dans cet article, la relation entre la logique modale LL et la théorie des probabilités est examinée. Les auteurs démontrent qu'il existe une façon de bien traduire des assertions probabilistiques en logique modale LL de façon à ce que cette dernière puisse saisir l'interprétation probabilistique de la vraisemblance. Cependant, cette traduction est subtile; plusieurs tentatives plus évidentes ont entraîné des incohérences. Des opérateurs modaux ont été ajoutés à la logique modale LL afin de permettre un raisonnement sur l'interaction de la connaissance et de la vraisemblance. On a constaté que la version propositionnelle de la logique résultante possédait une axiomatisation complète et s'avérait un facteur décisif en temps exponentiel.

Mots clés : raisonnement qualitatif à propos de la vraisemblance, lien probabilité et vraisemblance, combinaison connaissance et vraisemblance, logique modale.

[Traduit par la revue]

Comput. Intell. 5, 151–160 (1989)

1. Introduction

Reasoning in the presence of incomplete knowledge plays an important role in many AI expert systems. One way of representing partially constrained situations is with sentences of first-order logic (cf. McCarthy and Hayes 1969; Lipski 1977; Reiter 1984). A set A of first-order sentences specifies a set of possible worlds (first-order models). Intuitively, this is the set of all possible worlds that satisfy the constraints in the first-order statements.

While such assertions can deal with partial knowledge, they cannot adequately represent knowledge about relative likelihood. This would suggest that we should attach probabilities to sentences and/or possible worlds. Indeed, that approach has been taken by many authors, going back to Carnap (1950). More recently, this approach has seen renewed study (Bacchus 1988; Fagin *et al.* 1988; Halpern 1989; Nilsson 1986).

However, the use of probability is not always appropriate. Philosophers have spent years debating the situation (see Nutter (1987) and Shafer (1976) for some interesting discussion of this issue, and Cheeseman (1985) for a spirited defence of probability). The epistemological problems with the use of probability in AI were first noted by McCarthy and Hayes (1969), who made the following comments:

We agree that the formalism will eventually have to allow statements about the probabilities of events, but attaching probabilities to all statements has the following objections:

1. It is not clear how to attach probabilities to statements containing quantifiers in such a way that corresponds to the amount of conviction that people have.
2. The information necessary to assign numerical probabilities is not ordinarily available. Therefore, a formalism that required numerical probabilities would be epistemologically inadequate.

There have been a number of proposals for numerical representations of likelihood where a numerical estimate, or certainty factor, is assigned to each bit of information and to each conclusion drawn from that information (see Davis *et al.* (1977), Shafer (1976), and Zadeh (1978) for examples). But none of these proposals have been able to adequately satisfy the objections raised by McCarthy and Hayes. It is never quite clear where the numerical estimates are coming from; nor do these proposals seem to capture how people approach such reasoning. While people seem quite prepared to give qualitative estimates of likelihood, they are often notoriously unwilling to give precise numerical estimates to outcomes (cf. Szolovits and Pauker 1978).

Halpern and Rabin (1987) introduce a logic LL that is designed to allow qualitative reasoning about likelihood without the requirement of assigning precise numerical probabilities to outcomes. Indeed, numerical estimates and probability do not enter anywhere in the syntax or semantics of LL.

[1]This is an expanded version of a paper with the same title that appears in the Proceedings of the National Conference on Artificial Intelligence, Austin, TX, 1984.

Printed in Canada / Imprimé au Canada

Despite the fact that no use is made of numbers, LL is able to capture many properties of likelihood in an intuitively appealing way by using a modal operator L to capture the notion of likelihood. For example, consider the following chain of reasoning: if P_1 holds, then it is reasonably likely that P_2 holds, and if P_2 holds, it is reasonably likely that P_3 holds. Hence, if P_1 holds, it is somewhat likely that P_3 holds. (Clearly, the longer the chain, the less confidence we have in the likelihood of the conclusion.) In LL, this essentially becomes "from $P_1 \Rightarrow LP_2$ and $P_2 \Rightarrow LP_3$, conclude $P_1 \Rightarrow L^2P_3$." Note that the powers of L denote dilution of likelihood.

While LL is meant to capture a nonprobabilistic approach to reasoning about likelihood, given the prevalent usage of probability theory, it is important to understand relationship between LL and probability theory. This is especially so, since one reasonable way of understanding likelihood is via probability theory. To quote Halpern and Rabin (1987), "We can think of likely [the modal operator L] as meaning 'with probability greater than α' (for some user-defined α)." The exact relationship between LL and probability theory is not studied in Halpern and Rabin (1987). However, a close examination shows that it is not completely straightforward. Indeed, as show below, if we simply translate "P holds with probability greater than α" by LP, we quickly run into inconsistencies. Nevertheless, we confirm the sentiment in the quote above by showing that there is a way of translating numerical probability statements into LL in such a way that inferences made in LL are *sound* with respect to this interpretation of likelihood. Roughly speaking, this means that if we have a set of probability assertions about a certain domain, translate them (using the suggested translation) into LL, and then reason in LL, any conclusions we draw will be true when interpreted as probability assertions about the domain. However, our translation is somewhat subtle, as is the proof of its soundness; several more obvious attempts fail. These subtleties also shed some light on nonmonotonic reasoning.

In many situations, it does not suffice to reason about likelihood alone. We also have to reason about the subtle interplay between knowledge, belief, and likelihood. Work in the modal logic of knowledge and belief goes back to Hintikka (1962); more recent work can be found, for example, in Moore (1985), Fagin *et al.* (1984), and Halpern and Moses (1985) (see Halpern (1986) for an overview). It is clearly important to be able to reason simultaneously about knowledge and likelihood; there are many cases in which knowledge is heuristic or probabilistic. For example, suppose I know that Mary is a woman, but I have never met her and therefore do not know how tall she is. Under such circumstances, I consider it unlikely that she is over six feet tall. However, suppose that I am told that she is on the Stanford women's basketball team. My knowledge about her height has now changed, although I still don't know how tall she is. I now consider it reasonably likely that she is over six feet tall. Recently, a logic for reasoning simultaneously about knowledge and probability has been proposed (Fagin and Halpern 1988). Here we provide a logic for reasoning simultaneously about knowledge and likelihood, by enriching LL with modal operators for knowledge to get the modal logic LLK. LLK is shown to have a complete axiomatization, which is essentially obtained by combining the complete axiomatization of LL with that of the modal logic of knowledge. In addition, we show that there is a procedure for deciding validity of LLK formulas which runs in deterministic exponential time, the same as that for LL. This is provably the best possible.

The rest of the paper is organized as follows. In the next section, we review the syntax and semantics of LL. In Sect. 3, we discuss the translation of English sentences into LL and show that there is a translation which is sound with respect to the probabilistic interpretation of L. In Sect. 4, we add knowledge to the system to get the logic LLK, and state some technical results on decision procedures and axiomatizations for LLK. We conclude in Sect. 5 with some further directions for research.

2. Syntax and semantics

We briefly review the syntax and semantics of LL here, referring the reader to Halpern and Rabin (1987) for more details.

LL is a logic which extends standard propositional logic by the addition of two modal operators, L and G.[2] Roughly speaking, a formula of the form Lp should be viewed as saying "p is likely," while Gp should be viewed as saying "necessarily p." The syntax of LL is quite straightforward. Starting with a set, $\Phi = \{P, Q, R, \ldots\}$, of *primitive propositions*, we build more complicated LL formulas using the propositional connectives $\neg$ and $\wedge$ and the modal operators G and L. Thus, if p and q are formulas, then so are $\neg p$, $(p \wedge q)$, Gp (necessarily p), and Lp. We omit parentheses if they are clear from context. We also use the abbreviations $p \vee q$ for $\neg(\neg p \wedge \neg q)$, $p \Rightarrow q$ for $\neg p \vee q$, $p \equiv q$ for $(p \Rightarrow q) \wedge (q \Rightarrow p)$, Fp (possible p) for $\neg G \neg p$, and $L^i p$ for $L \ldots Lp$ (i Ls). Thus a typical LL formula is L^2GQ, which can be read "it is somewhat likely that G is necessarily the case." Note that the syntax allows arbitrary nestings and alternations of Ls and Gs.

We give semantics to LL formulas by means of Kripke structures. An LL *model* is a triple $M = (S, \mathcal{L}, \pi)$, where S is a set of *states,* $\mathcal{L}$ is a reflexive binary relation on S (i.e., for all $s \in S$, we have $(s, s) \in \mathcal{L}$), and $\pi : \Phi \times S \rightarrow$ {**true, false**}. Thus π tells us for each primitive proposition $P \in \Phi$ and each state $s \in S$ whether P is true in s.

Intuitively, a state is a complete and consistent set of "working hypotheses" concerning the situation under consideration, which we take to be "true for now." The *likely successors of a state* s (i.e., those states t such as $(s, t) \in \mathcal{L}$) are those states that describe a set of hypothesis that is reasonably likely, given our current hypotheses.[3]

We can think of $(S, \mathcal{L})$ as a graph with vertices S and edges $\mathcal{L}$. If $(s, t) \in \mathcal{L}$ then we say that t is an $\mathcal{L}$-successor of s. We will say t is *reachable* (*in k steps*) from s if, for

[2]The choice of the modal operator L is perhaps somewhat unfortunate, since it has been used in other papers (e.g., McDermott 1982) to denote necessity, and still others (e.g., Levesque 1984) to denote implicit belief. Nevertheless, we stick with L for consistency with Halpern and Rabin (1987) and because it suggests likelihood.

[3]In Halpern and Rabin (1987), besides likely successors there were also conceivable successors. For ease of exposition, we have omitted "conceivable" relation here, thus identifying the operator L^* of Halpern and Rabin (1987) with the F operator, which is the dual of G. We leave it to the reader to check that all our results also hold if we reinstate the conceivable relation.

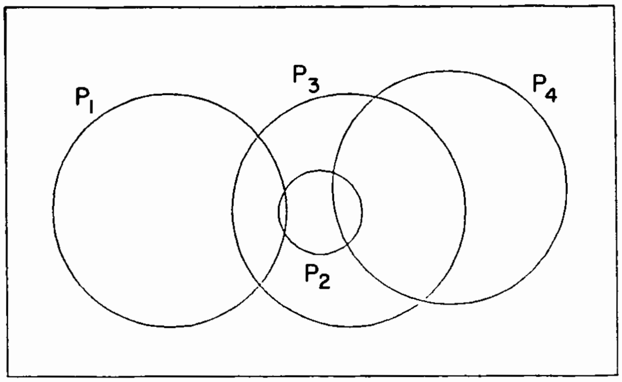

FIG. 1. A representation of a doctor's view.

some finite sequence $s_0, \ldots, s_k$, we have $s_0 = s$, $s_k = t$, and $(s_i, s_{i+1}) \in \mathcal{L}$ for $i < k$.

We define $M, s \models p$, read *p is satisfied in state s of model M*, by induction on the structure of p:

- $M, s \models P$ for $P \in \Phi$ iff $\pi(P, s) =$ **true**,
- $M, s \models \neg p$ iff $(M, s \not\models p)$,
- $M, s \models p \wedge q$ iff $M, s \models p$ and $M, s \models q$,
- $M, s \models Gp$ iff $M, t \models p$ for all t reachable from s,
- $M, s \models LP$ iff $M, t \models p$ for some t with $(s, t) \in \mathcal{L}$.

Definitions

A formula p is *valid* (resp. *satisfiable*) iff for all (resp. some) $M = (S, \mathcal{L}, \pi)$ and all (resp. some) $s \in S$ we have $M, s \models p$. It is easy to check that p is valid iff $\neg p$ is not satisfiable. If Σ is a set of LL formulas, we write $M, s \models \Sigma$ iff $M, s \models p$ for every formula $p \in \Sigma$.

Note that while we require L be reflexive, we do not take it to be transitive or symmetric. There are good reasons for this. The fact that $\mathcal{L}$ is reflexive guarantees us that $p \Rightarrow Lp$ will be valid. This says that if p is true, then it is likely to be true. On the other hand, if $\mathcal{L}$ were transitive, then it is easy to see that $L^i p \Leftrightarrow Lp$ would be valid for all $i > 0$. This would not be desirable, since we would like to capture the diminution of likelihood using powers of L. If $\mathcal{L}$ were symmetric, then $L \neg Lp \Rightarrow \neg Lp$ would be valid; again, this does not seem to be an appropriate property for likelihood. We remark that in Halpern and Rabin (1987) a complete axiomatization is provided for LL, which completely characterizes the properties of the L operator. (The axioms are described in Sect. 4, where we give a complete axiomatization for LLK.)

3. The probabilistic interpretation of likelihood

Lp is supposed to represent the notion that "p is reasonably likely." Certainly one way of interpreting this statement is "p holds with probability greater than or equal to α." However, as already noted in Halpern and Rabin (1987), there are problems with this interpretation of Lp. Suppose we take $\alpha = \frac{1}{2}$, and consider a situation where we toss a fair coin twice. If P represents "the coin will land heads both times" and Q represents "the coin will land tails both times," then we clearly have $L(P \vee Q)$, as well as $\neg LP$ and $\neg LQ$. But, for any LL model, $L(P \vee Q)$ is true iff $LP \vee LQ$ is true, giving us a contradiction.

We solve this problem by changing the way we translate statements of the form "p is reasonably likely" into LL. Note that if a state s satisfies the formula p (i.e., $M, s \models p$), this does *not* imply that p is necessarily true at s, but simply that p is one of the hypotheses that we are taking to be true at this state. We must use Gp to capture the fact that p is *necessarily* true at s, since $M, s \models Gp$ iff $M, t \models p$ for all t reachable from s, and thus in *no* state reachable from s is it the case that $\neg p$ is true. The English statement "The coin is likely to land heads twice in a row" should be interpreted as "It is likely to be (necessarily) the case that the coin lands heads twice in a row" and thus should be translated as LGP rather than LP. Similarly, "the coin is likely to land tails twice in a row" is LGQ, while "it is likely that the coin lands either heads or tails" is $LG(P \vee Q)$. With these translations, we do not run into the problem described above, for $LG(P \vee Q)$ is not equivalent to $LGP \vee LGQ$. These observations suggest that the only LL formulas that describe real-world situations are (Boolean combinations of) formulas of the form L^iGC, where C is a Boolean combination of primitive propositions. We will return to this point later.

Having successfully dealt with that problem, we next turn our attention to translating statements of conditional probability: "if P, then it is reasonably likely that Q" or "Q is reasonably likely given P." Again, the obvious translation, $GP \Rightarrow LGQ$, runs into trouble.

Consider a doctor making a medical diagnosis. His view of the world can be described by primitive propositions which stand for diseases, symptoms, and test results. The relationship between these formulas can be represented by a joint probability distribution, or a Venn diagram where the area of each region indicates its probability, and the basic regions correspond to the primitive propositions.

For example, the Venn diagram shown in Fig. 1 might represent part of the doctor's view, where P_1 and P_2 represent diseases and P_3 and P_4 represent symptoms. The diagram shows (among other things) that (1) disease P_1 is reasonably likely given symptom P_3; (2) P_3 is always a symptom of P_2; (3) if a patient has P_2, then it is not reasonably likely that he also has P_1; and (4) P_1 and P_4 never occur simultaneously.

The second statement is clearly $G(P_2 \Rightarrow P_3)$, from which we can deduce $GP_2 \Rightarrow GP_3$. Now suppose that we represented the first and third statements, as suggested above, by $GP_3 \Rightarrow LGP_1$ and $GP_2 \Rightarrow \neg LGP_1$, respectively. Then simply using propositional reasoning, we could deduce that $GP_2 \Rightarrow LGP_1 \wedge \neg LGP_1$, surely a contradiction.

The problem is that when we make such English statements as "P_1 is reasonably likely given P_3" or "the conditional probability of P_3 given P_1 is greater than one half," we are implicitly saying "given P_1 and all else being equal" or "given P_1 and no other information," P_3 is likely. We cannot quite say "given P_1 and *no* other information" in LL. Indeed, it is not quite clear precisely what this statement means (cf. Halpern and Moses 1984). However, we can say "in the absence of any information about the formulas $P_1, \ldots, P_k$ which would cause us to conclude otherwise," and this suffices for our applications. In our present example, P_1 is reasonably likely given P_3, as long as we are not given $\neg P_1$ or P_2 or P_4. Thus, a better translation of "P_1 is reasonably likely given P_3" is

$$\neg G \neg P_1 \wedge \neg GP_2 \wedge \neg GP_4 \wedge GP_3 \Rightarrow LGP_1$$

Similarly, "if a patient has P_2, then it is unlikely that he has P_1" can be expressed by

$\neg GP_1 \wedge GP_2 \Rightarrow \neg LGP_1$

In general, we must put all the necessary caveats into the precondition to avoid contradictions.

This translation seems to avoid the problem mentioned above, but how can we be sure that there are no further problems lurking in the bushes? We now show that, in a precise sence, there are not.

Fix a finite set of primitive propositions $\Phi = \{P_1, ..., P_n\}$. An *atom* of Φ is any conjunction $Q_1 \wedge ... \wedge Q_n$, where each Q_i is either P_i or $\neg P_i$. Note that there are 2^n such atoms. Let AT(Φ) be the set of atoms of Φ, and let LIT(Φ) $= \{P, \neg P | P \in \Phi\}$ be the set of *literals* of Φ; thus LIT(Φ) consists of all the primitive propositions and their negations. Let CON(Φ) be the set of all possible conjunctions of literals in LIT(Φ). We identify the empty conjunction with the formula *true*; thus, *true* is always a member of LIT(Φ). Finally, let PROP(Φ) be the set of all propositional formulas that can be formed using the propositions in Φ. If C, C' is an element of PROP(Φ), we write $C \leq C'$ if $C \Rightarrow C'$ is a propositional tautology.

We say a function $\Pr : \text{AT}(\Phi) \rightarrow [0, 1]$ is a *probability assignment on* Φ if $\Sigma_{A \epsilon \text{AT}(\Phi)} \Pr(A) = 1$. Intuitively, Pr assigns a probability to all the atoms in At(Φ) in such a way that the total probability is 1. We can extend Pr to PROP(Φ) by defining $\Pr(C) = \Sigma_{A \leq C, A \in \text{AT}(\Phi)} \Pr(A)$. (Since $A \leq true$ for every atom A, this means that $\Pr(true) = 1$, as expected.) If $\Pr(C') \neq 0$, we define the *conditional probability of C given D*, written $\Pr(C|C')$, as $\Pr(C \wedge C')/\Pr(C')$. Note that $\Pr(C|true) = \Pr(C)$.

Define a *propositional probability space W* to be a pair $(\Phi, \Pr)$, where Pr is a probability assignment on Φ. We now consider a restricted class of probability statements about W. Fix α with $0 < \alpha < 1$. A *probability assertion about W* is a formula of the form $\Pr(C|D) \geq \alpha^i$ or $\Pr(C|D) < \alpha^i$, where $i \geq 0$, $C \in$ PROP(Φ), $D \in$ CON(Φ), and $\Pr(D) > 0$. (Closure under negation is built into these formulas since, for example, $\neg \Pr(C|D) \geq \alpha^i$ iff $\Pr(C|D) < \alpha^i$.) The language is powerful enough to express assertions such as $\Pr(C) \geq \alpha^i$, since we can take $D = true$ in the formula $\Pr(C|D) \geq \alpha^i$. By taking $i = 0$, we can assert that a certain statement holds with probability 1.

It may seem that taking D to be in CON(Φ) in a statement such as $\Pr(C|D) \geq \alpha^i$ is a rather powerful restriction, but this is not so. If we wish to talk about the conditional probability of C with respect to an arbitrary formula D, we can simply extend Φ by adding one more primitive proposition, say P_D, extend Pr so that $\Pr(P_D \equiv D) = 1$ (this can be done easily), and write $\Pr(C|P_D)$ instead of $\Pr(C|D)$. In fact, this observation shows that we could have restricted to conditional probability statements of the form $\Pr(C|D)$, where D is a primitive proposition. There are two reasons not to do so. The first is pragmatic: we would like to keep Φ, the set of primitive propositions, small since, as we shall see, this keeps the set of possible "caveats" small. Second, we think of D as representing the collection of facts that the agent has learned or observations that the agent has made so far; in practice, this can often be represented as a conjunction of literals.

Corresponding to these probability assertions about W, we consider *standard LL formulas over* Φ. These are formed by taking formulas of the form L^iGC and $\neg L^iGC$, $i \geq 0$, where $C \in$ PROP(Φ), and closing off under conjunction and disjunction. By the observations above, these are, in some sense, exactly those LL formulas that describe a "real world" situation involving the primitive propositions of Φ.

We want to translate probability assertions about W into standard LL formulas over Φ. As discussed above, a conditional probability assertion of the form $\Pr(C|D) \geq a^i$ will be translated into a formula of the form $\neg GQ_1 \wedge ... \wedge \neg GQ_k \wedge GD \Rightarrow L^iGC$, where $Q_1, ..., Q_k$ are the "necessary caveats." We now make the notion of a "necessary caveat" precise. Given $C \in$ PROP(Φ), $D \in$ CON(Φ), and $Q \in$ LIT(Φ), we say Q has *negative* (resp. *positive*) *impact on C given D with respect to* (*w.r.t.*) Pr if

$$\Pr(D \wedge Q) > 0 \quad \text{and} \quad \Pr(C|D \wedge Q) < \Pr(C|D)$$
$$(\text{resp. } \Pr(C|D \wedge Q) > \Pr(C|D))$$

Thus Q has negative (resp. positive) impact on C given D w.r.t. Pr if discovering Q lowers (resp. increases) the probability of C given D. We say Q has *potential negative* (resp. *positive*) *impact on C given D w.r.t. Pr* if for some $D' \leq D$ with $\Pr(D') > 0$, Q has negative (resp. positive) impact on C given D' w.r.t. Pr. (In the sequel, we omit the phrase "w.r.t. Pr" if Pr is clear from context.)

Intuitively, if $D' \leq D$, then D' represents more information than D. Thus, if Q does not have potential negative impact on C given D, then once we know D, no matter what extra information we get, finding out Q will not lower the probability that C is true. Similar remarks hold for potential positive impact. We define

$$\text{PNI}(C, D) = \{Q \in \text{LIT}(\Phi) | Q \text{ has potential negative impact on } C \text{ given } D\}$$
$$\text{PPI}(C, D) = \{Q \in \text{LIT}(\Phi) | Q \text{ has potential positive impact on } C \text{ given } D\}$$

Now using the idea of potential positive and negative impact, we give a translation $q \rightarrow q^t$ from probability assertions about W to standard formulas over Φ. We define

$$[\Pr(C|D) \geq \alpha^i]^t = ((\wedge_{Q \in \text{PNI}(C,D)} \neg GQ) \wedge GD) \Rightarrow L^iGC$$
$$[\Pr(C|D) < \alpha^i]^t = ((\wedge_{q \in \text{PPI}(C,D)} \neg GQ) \wedge GD) \Rightarrow \neg L^iGC$$

Again we note that the term $\wedge_{Q \in \text{PNI}(C,D)} Q$ (resp. $\wedge_{Q \in \text{PPI}(C,D)} Q$) in the translation of $\Pr(C|D) \geq \alpha^i$ (resp. $\Pr(C|D) < \alpha^i$) is intended to capture the idea of "putting in all the necessary caveats in order to avoid contradictions."

With these definitions in hand, we can now state the theorem which asserts that there is a translation from probability assertions about W into LL which is sound.

Theorem 1

Let Σ be a finite set of probability assertions true in W, and Σ^t the conjunction of the standard LL formulas that result from translating the formula in Σ into LL (via $p \rightarrow p^t$). If q is a probability assertion such as $\Sigma^t \Rightarrow q^t$ is valid, then q is true about W.

We prove the theorem by constructing, for every propositional probability space W, an LL model $M_W = (S, \mathcal{L}, \pi)$ which we call the *canonical model corresponding to W*. The set S of states consists of countably many copies of each $C \in$ PROP(Φ) with $\Pr(C) > 0$. Successive copies are connected by $\mathcal{L}$, as well as a state you are likely to move to as your knowledge increases. More formally,

$$S = \{(C_i | i \geq 0,\ C \in \text{PROP}(\Phi),\ \Pr(C) > 0\}$$
$$\mathcal{L} = \{(C_i, C_i), (C_i, C_{i+1}) \mid i \geq 0\} \cup \{(C_i, C'_0) | C' \leq C,\ \Pr(C'|C) \geq \alpha^{i+1}\}$$

We define π as follows. For each $C \in \text{PROP}(\Phi)$ such that $\Pr(C) > 0$, choose some atom $A \in \text{AT}(\Phi)$ such that $A \leq C$ and $\Pr(A) > 0$ (such an atom must exist, since $\Pr(C) = \Sigma_{A \leq C, A \in \text{AT}(\Phi)} \Pr(A)$). Call this atom AT($C$). We then define $\pi(P, C_i) =$ **true** iff P is one of the conjuncts in AT(C). Note that this definition guarantees that $M_W, C_i \models C$. In fact, for any propositional formula C', we have $M_w, C_i \models C'$ iff $\text{AT}(C) \leq C'$.

The result now follows from three lemmas. We just state the lemmas here, leaving their proof to the appendix.

Lemma 1

Suppose $C, C' \in \text{PROP}(\Phi)$ and $\Pr(C) > 0$. Then

(a) If $k \geq 1$, then C'_0 is reachable from C_i in k steps iff $C' \leq C$ and $\Pr(C'|C) \geq \alpha^{k+i}$;

(b) $M_W, C_i \models GC'$ iff $\Pr(C'|C) = 1$ in W;

(c) $M_W, C_0 \models L^k GC'$ iff $\Pr(C'|C) \geq \alpha^k$ in W.

Lemma 2

If q is a probability assertion true about W, then $M_W, D_0 \models q^t$ for all $D \in \text{CON}(\Phi)$.

Lemma 3

If q is of the form $\Pr(C|D) \geq \alpha^i$ or $\Pr(C|D) < \alpha^i$ and $M_W, D_0 \models q^t$, then q is true in W.

Proof of Theorem 1

Suppose Σ is a finite set of probability assertions true in W, M_W is the canonical model for W constructed above, and q is a probability assertion about W such that $\Sigma^t \Rightarrow q^t$ is valid. Assume that q is of the form $\Pr(C|D) \geq \alpha^i$. (The case that q is of the form $\Pr(C|D) < \alpha^i$ is similar.) By Lemma 2, since each formula $q' \in \Sigma$ is true in W, we have $M_W, D_0 \models \Sigma^t$. Since $\Sigma^t \Rightarrow q^t$ is valid, we have $M_W, D_0 \models q^t$. Now by Lemma 3, it follows that q is true in W. ■

Roughly speaking, Theorem 1 says that if an agent reasoning about a situation has some probabilistic information about how likely certain events are and translates this information into LL using the translation described above, then every conclusion he can draw in LL that can be given a probabilistic interpretation will be true about the underlying situation. Thus it can be viewed as a soundness result, in that it indicates that doing probabilistic reasoning in LL will not lead to any contradictory results.

The following example should give the reader a feel of how the translation might work.

Example

Suppose that we are told that a randomly selected Stanford student is likely (with probability at least α) to be both intelligent and athletic. This statement implies that a randomly selected jock is likely to be smart. More specifically, if S stands for the proposition "x is a Stanford student," I stands for the proposition "x is intelligent," and A stands for the proposition "x is athletic." We are told that $\Pr(I \wedge A|S) \geq \alpha$. We leave it to the reader to check that the laws of probability allow us to conclude $\Pr(I|A \wedge S) \geq \alpha$. We now show that this conclusion can be derived in LL.

We must first translate the assertion $\Pr(I \wedge A|S) \geq \alpha$. To translate this into LL we need to consider a particular probability space. Assume that Φ includes the primitive propositions I, A, and S, and fix some propositional probability space $W = (\Phi, \Pr)$. We assume that $\Pr(A \wedge S) > 0$, otherwise the assertions we are interested in are not legitimate probability assertions about W. We want to show that

[1] $(\Pr(I \wedge A|S) \geq \alpha)^t \Rightarrow (\Pr(I|A \wedge S) \geq \alpha)^t$

is valid.

To compute these translates, we first need to compute $\text{PNI}(I \wedge A, S)$ and $\text{PNI}(I, A \wedge S)$, i.e., the literals which have potential negative impact on $I \wedge A$ given S and those which have potential negative impact on I given $A \wedge S$, respectively. In the special case where $\Phi = \{I, A, S\}$ and W assigns positive probability to all atoms, then it is easy to see that $\text{PNI}(I, A \wedge S) = \{\neg I\}$ and $\text{PNI}(I \wedge A, S) = \{\neg I, \neg A\}$. Note that in this case we have $\text{PNI}(I, A \wedge S) \subseteq \text{PNI}(I \wedge A, S)$. It turns out that this relationship holds in general. Suppose $Q \in \text{PNI}(I, A \wedge S)$. Then there exists some $D \in \text{CON}(\Phi)$ such that $\Pr(A \wedge S \wedge D \wedge Q) > 0$ and $\Pr(I|A \wedge S \wedge D \wedge Q) < \Pr(I|A \wedge S \wedge D)$. Note that for any conjunction $D' \in \text{CON}(\Phi)$, we have $\Pr(I|A \wedge D') = \Pr(I \wedge A|A \wedge D')$. Thus, $\Pr(I \wedge A|S \wedge A \wedge D \wedge Q) < \Pr(I \wedge A|S \wedge A \wedge D)$. It immediately follows that $Q \in \text{PNI}(I \wedge A, S)$. This shows that $\text{PNI}(I, A \wedge S) \subseteq \text{PNI}(I \wedge A, S)$, as desired.

Now

$$(\Pr(I \wedge A|S) \geq \alpha)^t = ((\wedge_{Q \in \text{PNI}(I \wedge A, S)} \neg GQ) \wedge GS) \Rightarrow LG(I \wedge A)$$

and

$$(\Pr(I|A \wedge S) \geq \alpha)^t = ((\wedge_{Q \in \text{PNI}(I, A \wedge S)} \neg GQ) \wedge G(A \wedge S)) \Rightarrow LGI$$

Since $\text{PNI}(I, A \wedge S) \subseteq \text{PNI}(I \wedge A, S)$ and $G(A \wedge S) \Rightarrow GS$ is valid, if the precondition of $(\Pr(I|A \wedge S) \geq \alpha)^t$ holds, then so does the precondition of $(\Pr(I \wedge A|S) \geq \alpha)^t$. Since $LG(I \wedge A) \Rightarrow LGI$ is clearly valid, it is easy to see that [1] is valid as well.

Remarks

1. Note that the translation given in Theorem 1 depends on W, the underlying propositional probability space. Thus, we should really write $q^{t(W)}$ rather than q^t to denote the translate of the probability of assertion q. Intuitively, the reason for this dependence on W is that the translation can thereby take into account background knowledge about the situation the agent finds himself in. This background knowledge comes up in the computation of PNI and PPI.

In practice, it is of course not always possible to compute $\text{PNI}(C, C')$ or $\text{PPI}(C, C')$. The probabilistic information required may not be available, or it may be available but hard to compute, perhaps in part because the set of primitive propositions Φ may be large. In the examples discussed in McCarthy (1980), Φ is viewed as being essentially infinite. If we take P to be "Tweety is a bird" and Q to be "Tweety can fly," then Q is likely given P as long as Tweety is not an ostrich, Tweety in not a penguin, Tweety is not dead, Tweety's wings are not clipped, ... The list of possible disclaimers is endless. However, our assumption of having only a finite (and reasonably small) number of primitive propositions does seem to be both epistemologically and practically reasonable in many natural applications. It amounts to restricting oneself to considering only a small set of relevant events, which is something that seems to be frequently done in practice. For example, in medical diagnosis, we could take Φ to consist of relevant symptoms, diseases, and possible treatments, where the symptoms are qualitative (his temperature is very high) rather than quantitative (his temperature is 104°F.).

In any case, if we cannot compute PNI or PPI, we can always "play it safe", by replacing PNI(C, C') (resp. PPI(C, C')) wherever it occurs in the translation by a superset. It is straightforward to modify the proof of Theorem 1 to show that the resulting translation is still sound. More precisely, given a finite set Σ of probability assertions true about W, let Σ' be the result of replacing each LL formula p^t in Σ^t by a formula p', where instead of PNI or PPI, we use some superset of caveats. Then it is still the case that if $\Sigma' \Rightarrow q^t$, then q is true about W.[4] However, since putting extra caveats limits the applicability of a rule, this is not always a good strategy to follow.

If, on the other hand, we use a subset of PNI or PPI in the translation, then our reasoning may be unsound (in the sense of Theorem 1). This may help to explain where the nonmonotonicity comes from in certain natural language situations. People often use a type of informal default reasoning, saying "P is likely given Q," without specifying the situations where the default Q may not obtain. Of course, this means that the conclusion Q may occasionally have to be withdrawn in the light of further evidence.

2. Our translation is somewhat sensitive to the set of allowable probability assertions. We could easily extend the set of probability assertions about W to allow conjunctions and still prove an analogue of Theorem 1. However, if we further extend the set of probability assertions about W so that it is closed under disjunction (so that $\Pr(C|D) \geq \alpha \vee \Pr(C'|D') < \alpha^2$ would be a typical assertion), and extend the translation $q \rightarrow q^t$ to deal with disjunctions by taking $(p \vee q)^t = p^t \vee q^t$, then Theorem 1 no longer holds. For an example, let $\Phi = \{P\}$, and take $\Sigma = \{\Pr(P) \geq \alpha, \Pr(\neg P) \geq \alpha\}$ and $q = \Pr(P|\neg P) \geq \alpha \vee \Pr(\neg P|P) \geq \alpha$. Let W be any probability space such that the assertions in Σ are true in W. Since $\Pr(P|\neg P)$ and $\Pr(\neg P|P)$ must both be 0, it is clear that q is false in W. However, it is easy to see that $\Sigma^t = (\neg G\neg P \Rightarrow LGP) \wedge (\neg GP \Rightarrow LG\neg P)$ and $q^t = (G\neg P \Rightarrow LGP) \vee (GP \Rightarrow LG\neg P)$, from which it follows by propositional reasoning that $\Sigma^t \Rightarrow q^t$ is valid.

3. We have viewed Theorem 1 as a soundness result. It is natural to ask if there is a complementary completeness result. Ideally, we would like to prove that if Σ is a collection of probability assertions true about W and q is a probability assertion that follows from Σ, then $\Sigma^{t(W)} \Rightarrow q^{t(W)}$ is valid.[5] However, this is too much to expect.

For example, suppose that $\Phi = \{P, Q\}$, Σ consists of the assertions $\Pr(P|true) \geq \alpha$, $\Pr(\neg P|true) \geq \alpha$, and $\Pr(Q|true) < \alpha$, and q is the assertion $\Pr(\neg Q|true) \geq \alpha$. The first two assertions in Σ are true precisely in situations where $\alpha \leq 1/2$. Thus, if W is a space where all the assertions in Σ are true, then q must also be true in W, no matter what the interpretation of α. But it is easy to construct a W such that $\Sigma^{t(W)} \wedge \neg q^{t(W)}$ is consistent.

There are other examples of probabilistic information that is lost as a result of our translation. For one thing, there are some valid deductions in probability theory, such as the axiom $\Pr(p \vee q) = \Pr(p) + \Pr(q) - \Pr(p \wedge q)$, that cannot even be expressed in LL (since LL has no means of expressing addition). Another example is provided by the following "chain rule." It is easy to see that $\Pr(R|P) > \alpha^2$ follows from $\Pr(Q|P) > \alpha$ and $\Pr(R|P \wedge Q) > \alpha$.[6] However, we leave it to the reader to check that we cannot deduce $[\Pr(R|P) > \alpha^2]^t$ from $[\Pr(Q|P) > \alpha]^t \wedge [\Pr(R|P \wedge Q) > \alpha]^t$.

So where does this leave us? We have shown that some probabilistic reasoning can be done within LL, while some cannot. We do not have a precise characterization of how much probabilistic reasoning can be done in LL; providing such a characterization seems to be a difficult task (although the decision procedure stated in Theorem 3 in the next section at least gives us a procedure for checking whether a fact is deducible). In fact, it seems that the issue of how much probabilistic reasoning can be captured within LL is perhaps not the right issue to be worrying about. LL is not meant as a replacement for probability theory in cases where probability theory is applicable. In such cases, it is best to just stick to probability theory, without bothering to translate into LL.

However, as has been pointed out by many authors (see the discussion in the introduction, as well as the papers cited there and in Halpern and Rabin (1987)), there are many situations where probability theory is not immediately applicable or precise numerical probabilities are not available. In such situations, an agent may still have some qualitative information (or subjective feelings) about the likelihood of various events. In fact, often this information is of the form "P is reasonably likely given Q." What Theorem 1 shows is that care must be taken in order to capture this information in LL in a way that is reasonably consistent with probabilistic intuitions and that, if such care is taken, there is a way of capturing it in LL.

4. Reasoning about knowledge and likelihood

We can augment LL in a straigthforward way in order to accommodate reasoning about knowledge. The syntax of the resulting language, which we call LLK, is the same as that of LL except that we add unary operators $K_1, \ldots, K_n$, one for each of the "players" or "agents" 1, ..., n, and allow formulas of the form K_ip (which is intended to mean "agent i knows p"). Thus, a typical formula of LLK might be $K_i(GQ \wedge LGP)$: agent i knows that Q is actually the case and it is likely that P is the case. Again we allow arbitrary nesting and alternation of Ks, Ls, and Gs. It is very different for agent i to know that P is likely (i.e., K_iLGP) and for it to be likely that agent i knows P (LK_iGP).

We assume that our knowledge operator satisfies the axioms of the classifical modal logic S5 (cf. Halpern and Moses 1985). In particular, we know only true things and

[4]Note that in order to ensure soundness, we cannot replace PNI or PPI by a superset in the conclusion q^t of the implication. In fact, soundness is preserved only if we use a subset, rather than a superset, of the caveats specified in the definition of q^t.

[5]We say *q follows from* Σ if, for any interpretation of α with $0 < \alpha < 1$, if Σ is true of a propositional probability space W under this interpretation of α, then so is q.

[6]On the other hand, $\Pr(R|P) > \alpha^2$ does not follow $\Pr(Q|P) > \alpha$ and $\Pr(R|Q) > \alpha$. It is easy to construct examples to show that this is not a valid probabilistic inference.

we know exactly what we know and do not know. Thus we have

$$K_ip \Rightarrow p$$
$$K_ip \Rightarrow K_iK_ip$$
$$\neg K_ip \Rightarrow K_i\neg K_ip$$

We remark that it is easy to change the semantics for knowledge given below so that each of these axioms can be dropped (and possibly replaced by others). We omit further details here (see Halpern and Moses (1985) for more discussion of this point).

A model for LLK is a tuple $M = (S, \mathcal{L}, \mathcal{K}_1, \ldots, \mathcal{K}_n, \pi)$, where, just as before, S is a set of states, $\mathcal{L}$ is a reflexive binary relation describing likely successors, and π assigns truth values to the primitive propositions at every state. For each i, $\mathcal{K}_i$ is an equivalence relation on S: it is reflexive ($(s, s) \in \mathcal{K}_i$ for each $s \in S$), symmetric (if $(s, t) \in \mathcal{K}_i$, then so is (t, s)), and transitive (if (s, t), $(t, u) \in \mathcal{K}_i$, then so is (s, u)). As we shall see, the $\mathcal{K}_i$ will be used to give semantics to the K_ip formulas; making it an equivalence relation guarantees that it satisfies the three axioms above.

We can think of a state and all the states reachable from it via the $\mathcal{L}$ relation as describing a "likelihood distribution." Two states are joined via the $\mathcal{K}_i$ relation iff agent i views them as possible likelihood distributions (rather than just possible worlds in Hintikka (1962), McCarthy *et al.* (1979), and Moore (1985)) given his/her current knowledge. Thus, once agent i knows that Mary is on the women's basketball team, he/she would presumably consider it likely that Mary is over six feet tall, so that in all the likelihood distributions joined by $\mathcal{K}_i$ where there is no information to the contrary, LG(Mary is over six feet tall) would hold.

As before, we define $M, s \models p$ by induction on the structure of p. The only new clause is

$$M, s \models K_ip \text{ iff } M, t \models p \text{ for all } t \text{ such that } (s, t) \in \mathcal{K}_i$$

We now state some technical results for LLK. They essentially parallel the corresponding results for each of LL and S5 alone, showing that combining likelihood and knowledge does not lead to any special complications. The proofs use standard techniques of modal logic, so are omitted here. The reader is referred to Emerson and Halpern (1985), Fischer and Ladner (1979), Halpern and Moses (1985), Halpern and Rabin (1987), and Kozen and Parikh (1981) for further details.

Definition

The *size of a formula p*, written $|p|$, is its length as a string over the alphabet $\Phi \cup \{\neg, \wedge, G, L, K_1, \ldots, K_n,), (\}$. The *size of a model M* is the number of states in S (and thus could be infinite).

Theorem 2

An LLK formula p is satisfiable iff it is satisfiable in a model of size $\leq 2^{|p|}$.

Theorem 3

There is a procedure for deciding if a formula p is satisfiable (respectively, valid) which runs in deterministic time $O(2^{c|p|})$ for some $c > 0$.

This is the best we can do, as the following shows:

Theorem 4

The problem of deciding satisfiability (validity) of LL formulas is complete for deterministic exponential time.

Theorem 5

The following axiom system is sound and complete for LLK:

Axion schemes

AX1. All (substitution instances of) tautologies of propositional logic.
AX2. $Gp \Rightarrow p$
AX3. $Gp \Rightarrow GGp$
AX4. $Gp \Rightarrow \neg L\neg p$
AX5. $p \Rightarrow Lp$
AX6. $L(p \vee q) \equiv (Lp \vee Lq)$
AX7. $G(p \Rightarrow q) \Rightarrow (Gp \Rightarrow Gq)$
AX8. $G(p \Rightarrow q) \Rightarrow (Lp \Rightarrow Lq)$
AX9. $G(p \Rightarrow \neg L\neg p) \Rightarrow (p \Rightarrow Gp)$
AX10. $K_ip \Rightarrow p$
AX11. $K_ip \Rightarrow K_iK_ip$
AX12. $\neg K_ip \Rightarrow K_i\neg K_ip$
AX13. $K_i(p \Rightarrow q) \Rightarrow (K_ip \Rightarrow K_iq)$

Rules of inference

R1. From p and $p \Rightarrow q$ infer q (*modus ponens*)
R2. From p infer Gp (generalization)
R3. From p infer K_ip (knowledge generalization)

The axiom system given here is simply a combination of the axiom systems for LL (AX1–9, R1, R2) and S5 (AX1, AX10–13, R1, R3). The axioms for LL presented in Halpern and Rabin (1987) differ slightly due to the presence of the "conceivable" relation in the semantics. Had we reinstated the conceivable relation, we would have to modify AX9 accordingly (cf. Halpern and Rabin 1987).

5. Conclusions

We have examined the relationship between the logic LL and probability theory. We have shown that there is a precise sense in which a restricted class of probabilistic assertions about a domain can be captured by LL formulas. However, the translation from probabilistic assertions to LL is subtle; translations more naive than the one we use turn out not to be sound. In particular, in order to correctly deal with statements of conditional probability, we must specifically list all the situations in which the conclusion may not hold. The failure to do so in informal human reasoning is frequently the cause of the nonmonotonicity so often observed in such reasoning. (However, we note here in passing that a number of the problems that McDermott (1982) suggests can be dealt with by nonmonotonic logic can also be dealt with by LL, in a completely monotonic fashion. See Halpern and Rabin (1987) for further discussion on this point). The need to explicitly list all caveats can be viewed as a discipline which forces a practitioner to list explicitly all the exceptions to his rules. Of course, this method does not guarantee correctness. If an exception is omitted, then any conclusion made using that rule may be invalid. But, whenever a conclusion is retracted, it should be possible to find the missing exception and correct the rule appropriately.

The fact that we have a translation from probability assertions to LL does not necessarily make LL a good language for probabilistic assertions. Not all probabilistic assertions can be captured in LL, and even for the ones that can, some information can be lost in the process of translating into LL. This suggests that LL may not be an appropriate language to use when we really have probabilistic data.

In retrospect, this conclusion is perhaps not too surprising. LL is designed to deal with situations where likelihood is interpreted as being something other than just probability. While a given LL formula may be true of any situation where L is interpreted as meaning "with probability $\geq \alpha$," it may not be true for some other interpretation of L. We could, for example, take LGp to mean "the agent has belief (in the sense of the Dempster–Shafer notion of belief (Shafer 1976) at least α in p." With this interpretation, the formula $q^{t(W)}$ in the third remark after Theorem 1 (which we used to show why we could not get a completeness result) would not follow from $\Sigma^{t(W)}$. This may provide some intuition as to why we cannot get a completeness result of analogues to Theorem 1.

On the other hand, it is also possible that the lack of a completeness result is an artifact of the particular translation we used from probability assertions to LL. It may be worth investigating whether there is a different translation from probability theory to LL for which a soundness *and* completeness result in the spirit of Theorem 1 is provable. Perhaps this translation would give a reasonable probabilistic interpretation to nonstandard LL formulas (which do not correspond to probabilistic assertions under our present translation).

Since LL can express some notions of likelihood other than probability, this may make it applicable in contexts where probability theory is not. It would be interesting to know whether LL is able to capture other notions of reasoning about uncertainty, such as *possibility theory* (Zadeh 1978, 1981) or belief functions (Shafer 1976). (See the survey paper by Prade (1984) for a thorough discussion of various approaches to modelling reasoning about uncertainty.) It would be interesting to investigate for these other approaches whether an analogue of Theorem 1 is provable. Perhaps when we consider these other approaches, there will be a need to use nonstandard LL formulas.

It would also be interesting to understand the relationship between LL and the logic QP (Qualitative Probability) introduced by Gärdenfors (1975). In QP, we can say "p is more likely than q", but not "p is likely." The axioms of QP seem more complicated than those of LL, and although QP is decidable, it seems that the decision procedure would be quite complex. It is not clear how easy it would be to combine the approaches taken in LL and QP to construct a logic where one can both make absolute and relative statements about likelihood.

Finally, a rich area for further work is simultaneous reasoning about knowledge and likelihood. LLK provides a first step, but clearly there is more work to be done in order to find a truly appropriate logic that is both formally and epistemologically adequate.

Acknowledgements

We would like to thank Daniel Hunter for his useful criticisms of both our explanation of the soundness result and our translation into LL. We also thank the referees for the constructive criticism of an earlier version of this paper.

BACCHUS, F. 1988. Representing and reasoning with probabilistic knowledge. Ph.D. thesis, University of Alberta, Edmonton, Alta. Also issued as Technical Report CS-88-31, University of Waterloo, Waterloo, Ont.

CARNAP, R. 1950. Logical foundations of probability. University of Chicago Press, Chicago, IL.

CHEESEMAN, P. 1985. In defense of probability. Proceedings of the Ninth International Joint Conference on Artificial Intelligence, Los Angeles, CA, pp. 1002–1009.

DAVIS, R., BUCHANAN, B., and SHORTLIFFE, E. 1977. Production rules as a representation for a knowledge-based consultation system. Artificial Intelligence, **8**: 15–45.

EMERSON, E.A., and HALPERN, J.Y. 1985. Decision procedures and expressiveness in the temporal logic of branching time. Journal of Computer and System Sciences, **30**(1): 1–24.

FAGIN, R., and HALPERN J.Y. 1988. Reasoning about knowledge and probability: preliminary report. Proceedings of the Second Conference on Theoretical Aspects of Reasoning about Knowledge, Monterey, CA, pp. 277–293. *Edited by* M.Y. Vardi. Morgan Kaufmann, San Mateo, CA.

FAGIN, R., HALPERN, J.Y., and VARDI, M.Y. 1984. A model-theoretic analysis of knowledge: preliminary report. Proceedings of the 25th IEEE Symposium on Foundations of Computer Science, Singer Island, FL, pp. 268–278. A revised and expanded version appears as IBM Research Report RJ 6641, IBM Research Division, Almaden Research Center, San Jose, CA. 1988.

FAGIN, R., HALPERN, J.Y., MEGIDDO, N. 1988. A logic for reasoning about probabilities. Proceedings of the Third Symposium on Logic in Computer Science, Ithaca, NY, pp. 277–291. A revised and expanded version appears as IBM Research Report RJ 6190, IBM Research Division, Almaden Research Center, San Jose, CA. April 1988.

FISCHER, M.J., and LADNER, R.E. 1979. Propositional dynamic logic of regular programs. Journal of Computer and System Sciences, **18**(2): 194–211.

GÄRDENFORS, P. 1975. Qualitative probability as an intensional logic. Journal of Philosophical Logic, **4**: 171–185.

HALPERN J.Y. 1986. Reasoning about knowledge: an overview. *In* Theoretical aspects of reasoning about knowledge: Proceedings of the 1986 Conference. *Edited by* J.Y. Halpern. Morgan Kaufmann, San Mateo, CA, pp. 1–17. *Reprinted in* Proceedings of the National Computer Conference, Las Vegas, NV, pp. 219–228.

——— 1989. An analysis of first-order logics of probability. Proceedings of the Eleventh International Joint Conference on Artificial Intelligence, Detroit, MI.

HALPERN, J.Y., and MOSES, Y. 1984. Towards a theory of knowledge and ignorance. Proceedings of the American Association for Artificial Intelligence Workshop on Non-monotonic Logic, Mohonk, NY, pp. 125–143. *Reprinted in* Logics and models of concurrent systems. *Edited by* K. Apt. Springer-Verlag, New York, NY, pp. 459–476. 1985.

——— 1985. A guide to the modal logics of knowledge and belief. Proceedings of the Ninth International Joint Conference on Artificial Intelligence, Los Angeles, CA, pp. 480–490.

HALPERN, J.Y., and RABIN, M.O. 1987. A logic to reason about likelihood. Artificial Intelligence, **32**(3): 379–405.

HINTIKKA, J. 1962. Knowledge and belief. Cornell University Press, Ithaca, NY.

KOZEN, D., and PARIKH, R. 1981. An elementary proof of the completeness of PDL. Theoretical Computer Science, **14**(1): 113–118.

LEVESQUE, H. 1984. A logic of implicit and explicit belief. Proceedings of the National Conference on Artificial Intelligence, Austin, TX, pp. 198–202.

LIPSKI, W. 1977. On the logic of incomplete information. Proceedings of the 6th International Symposium on Mathematical Foundations of Computer Science, Lecture Notes in Computer Science, Vol. 53, pp. 374–381. Springer-Verlag, New York, NY.

MCCARTHY, J. 1980. Circumscription — a form of non-monotonic reasoning. Artificial Intelligence, **13**: 1,2.
MCCARTHY, J.M., and HAYES, P.J. 1969. Some philosophical problems from the standpoint of artificial intelligence. Machine intelligence 4. *Edited by* D. Michie. American Elsevier, New York, NY. pp. 463–502.
MCCARTHY, J., SATO, M., and HAYASHI, T., and IGARISHI, S. 1979. On the model theory of knowledge. Technical Report STAN-CS-78-657, Stanford, CA.
MCDERMOTT, D.V. 1982. Nonmonotonic logic II: nonmonotonic model theories. Journal of the ACM, **29**(1): 33–57.
NILSSON, N. 1986. Probabilistic logic. Artificial Intelligence, **28**: 71–87.
NUTTER, J.T. 1987. Uncertainty and probability. Proceedings of the Tenth International Joint Conference on Artificial Intelligence, Milan, Italy, pp. 373–379.
PRADE, H. 1984. Quantitative methods in approximate and plausible reasoning: the state of the art. Technical Report, Université Paul Sabatier, Toulouse, France.
REITER, R. 1984. Towards a logical reconstruction of relational database theory. *In* Conceptual modelling: perspectives from artificial intelligence, databases, and programming languages. *Edited by* M.L. Brodie, J. Mylopoulos, and J. Schmidt. Springer-Verlag, New York, NY. pp. 191–233.
SHAFER, G. 1976. A mathematical theory of evidence. Princeton University Press, Princeton, NJ.
SZOLOVITS, P., and PAUKER, G. 1978. Categorical and probablistic reasoning in medical diagnosis. Artificial Intelligence, **11**: 115–144.
ZADEH, L.A. 1978. Fuzzy sets as a basis for a theory of possibility. Fuzzy Sets and Systems, **1**: 3–28.
——— 1981. Possibility theory and soft data analysis. *In* Mathematical frontiers of the social and policy sciences. *Edited by* L.M. Cobb and R.M. Thrall. American Association for the Advancement of Science Selected Symposium, Vol. 54. Westview Press, Boulder, CO. pp. 69–129.

Appendix. Proof of lemmas 1, 2, and 3

For easy reference, we repeat the statements of the lemmas here.

Lemma 1

Suppose $C, C' \in \text{PROP}(\Phi)$ and $\Pr(C) > 0$. Then

(a) If $k \geq 1$, then C'_0 is reachable from C_i in k steps iff $C' \leq C$ and $\Pr(C'|C) \geq \alpha^{k+i}$,

(b) $M_W, C_i \vDash GC'$ iff $\Pr(C'|C) = 1$ in W,

(c) $M_W, C_0 \vDash L^k GC'$ iff $\Pr(C'|C) \geq \alpha^k$ in W.

Proof

For part (a) we proceed by induction on k. The case $k = 1$ is immediate from the definition of $\mathcal{L}$. Suppose $k > 1$. If C'_0 is reachable from C_i in k steps, then there exists a state D such that D is reachable from C_i in one step and C'_0 is reachable from D in $k - 1$ steps. By the definition of $\mathcal{L}$, D is of the form C_i, C_{i+1}, or E_0, where $E \leq C$ and $\Pr(E|C) \geq \alpha^{i+1}$. If one of the first two possibilities holds, then we get $C' \leq C$ and $\Pr(C'|C) \geq \alpha^{k+i}$ immediately from the inductive hypothesis. If the third case holds, by the inductive hypothesis we have $C' \leq E$ and $\Pr(C'|E) \geq \alpha^{k-1}$. Since $C' \leq E$ and $E \leq C$, we have $C' \leq C$. This also means that for any atom A, if $A \leq C'$ then $A \leq E$ and if $A \leq E$ then $A \leq C$. As a consequence, we must have $\Pr(C') = \Pr(C' \wedge E) = \Pr(C' \wedge C)$ and $\Pr(E \wedge C) = \Pr(E)$. Easy manipulations now show that $\Pr(C'|C) \geq \alpha^{k+i}$. For the converse, if $C' \leq C$ and $\Pr(C'|C) \geq \alpha^{k+i}$, then $(C_{k+i-1}, C'_0) \in \mathcal{L}$, so it follows immediately from the definition of $\mathcal{L}$ that C'_0 is reachable from C_i in k steps.

For part (b), first suppose $\Pr(C'|C) < 1$ in W. Thus $\Pr(\neg C'|C) > 0$, so there must be some atom A such that $A \leq \neg C'$ and $\Pr(A|C) > 0$. From part (a) it follows that A_0 is reachable from C_i. From the definition of π, it follows that $M_W, A_0 \vDash A$ and hence that $M_W, A_0 \vDash \neg C'$. Thus, $M_W, D_i \vDash \neg GC'$. For the converse, suppose that $\Pr(C'|C) = 1$. We want to show that $M_W, C_i \vDash GC'$. Suppose not. Then there must be some state E reachable from C_i such that $M_W, E \vDash \neg C'$. But if E is reachable from C_i, then E is of the form C''_j where $C'' \leq C$ and $\Pr(C'') > 0$. Thus, from the definition of π, we must have $\text{AT}(C'') \leq \neg C'$. But, again by definition, $\text{AT}(C'') \leq C''$, so $\text{AT}(C'') \leq C$. It follows that $\text{AT}(C'') \leq C \wedge \neg C'$. Moreover, since $\Pr(\text{AT}(C'')) > 0$, we have $\Pr(\neg C'|C) > 0$. But this contradicts our assumption that $\Pr(C'|C) = 1$. Hence we have $M_W, C_i \vDash GC'$, as desired.

For part (c), note that if $C_0 \vDash L^k GC'$, then there exists a state E which is reachable from C_0 in k steps such that $M_W, E \vDash GC'$. By part (b), it follows that $\Pr(C'|E) = 1$ and, by part (a), we have that $E \leq C$ and $\Pr(E|C) \geq \alpha^k$. Using standard probabilistic reasoning, we get that $\Pr(C'|C) \geq \alpha^k$. For the converse, note that the case $k = 0$ follows immediately from part (b). If $k > 0$, since $\Pr(C'|C) \geq \alpha^k$, clearly $\Pr(C' \wedge C|C) \geq \alpha^k$. But $(C' \wedge C) \leq C$, so $(C' \wedge C)_0$ is reachable from C_0 in k steps. Since $M_W, (C' \wedge C)_0 \vDash GC'$ by part (b), it follows that $M_W, C_0 \vDash L^k GC'$. ■

Lemma 2

If q is a probability assertion true about W, then $M_W, D_0 \vDash q^t$ for all $D \in \text{CON}(\Phi)$.

Proof

Suppose q is of the form $\Pr(C|D') \geq \alpha^i$. (The proof if q is of the form $\Pr(C|D') < \alpha^i$ is similar and left to the reader.) By definition, we have

$$q^t = ((\wedge_{Q \in \text{PNI}(C,D')} \neg GQ) \wedge GD') \Rightarrow L^i GC$$

Suppose $M_W, D_0 \vDash (\wedge_{Q \in \text{PNI}(C,D')} \neg GQ) \wedge GD'$. Our goal is to show that $\Pr(C|D) \geq \alpha^i$ in W, for then by part (c) of Lemma 1, we have $M_W, D_0 \vDash L^i GC$, as desired. To see this, first observe that by part (b) of Lemma 1, we have $\Pr(D'|D) = 1$. Easy calculations now show that $\Pr(D' \wedge D) = \Pr(D)$ and $\Pr(C|D' \wedge D) = \Pr(C|D)$. Thus it suffices to show that $\Pr(C|D' \wedge D) \geq \alpha^i$. Next observe that none of the conjuncts that make up D can be in $\text{PNI}(C, D')$ (for if $Q \in \text{PNI}(C, D')$ were one of the conjuncts in D, the definition of π would guarantee that $M_W, D_0 \vDash GQ$, contradicting our assumption). It follows that $D' \wedge D$ is of the form $D' \wedge Q_1 \wedge \ldots Q_k$, where $Q_i \notin \text{PNI}(C, D')$. Now using the definition of $\text{PNI}(C, D')$, we can show, by induction on j,

$$\Pr(C|D' \wedge Q_1 \wedge \ldots \wedge Q_j) \geq \Pr(C|D' \wedge Q_1 \wedge \ldots \wedge Q_{j-1}) \geq \ldots \geq \Pr(C|D')$$

Since $\Pr(C|D) = \Pr(C|D' \wedge D) = \Pr(C|D' \wedge Q_1 \wedge \ldots Q_j)$, and $\Pr(C|D') \geq \alpha^i$ in W (since q is true in W, by assumption), we have that $\Pr(C|D) \geq \alpha^i$, as desired. ■

We remark that Lemma 2 does not hold for arbitrary $C \in \text{PROP}(\Phi)$. That is, there exists a probability assertion true q about W and a $C \in \text{PROP}(\Phi)$ such that $M_W, C_0 \nvDash q^t$. (We can find a counterexample by taking

$\Phi = \{P, Q, R\}$, $C = (P \wedge Q) \vee (P \wedge R)$, and $q = \Pr(\neg Q|P) \geq \alpha$. We omit details here.)

Lemma 3

If q is of the form $\Pr(C|D) \geq \alpha^i$ or $\Pr(C|D) < \alpha^i$ and $M_W, D_0 \models q^t$, then q is true in W.

Proof

Assume that q is of the form $\Pr(C|D) \geq \alpha^i$ and $M_W, D_0 \models q^t$. (Again, the case where q is of the form $\Pr(C|D) < \alpha^i$ is similar and left to the reader). It is easy to check that if $Q \in \mathrm{PNI}(C, D)$ then we must have $M_W, D_0 \models \neg GQ$. (For if $M_W, D_0 \models GQ$, then by part (c) of Lemma 1, we must have $\Pr(Q|D) = 1$, from which it follows that if $D' \leq D$, we also have $\Pr(Q|D') = 1$, and $\Pr(C|D' \wedge Q) = \Pr(C|D')$, so that $Q \notin \mathrm{PNI}(C, D)$.) Thus we have $M_W, D_0 \models (\wedge_{Q \in \mathrm{PNI}(C, D)} \neg GQ) \wedge GD$. Since, by assumption, $M_W, D_0 \models q^t$, it follows that $M_W, D_0 \models L^i GC$. By part (c) of Lemma 1, we have that $q = \Pr(C|D) \geq \alpha^i$ in W, as desired. ■