Journal of Artificial Intelligence Research 26 (2006) 1–34 Submitted 11/05; published 05/06

A Logic for Reasoning about Evidence

Joseph Y. Halpern HALPERN@CS.CORNELL.EDU
Cornell University, Ithaca, NY 14853 USA

Riccardo Pucella RICCARDO@CCS.NEU.EDU
Northeastern University, Boston, MA 02115 USA

Abstract

We introduce a logic for reasoning about evidence that essentially views evidence as a function from prior beliefs (before making an observation) to posterior beliefs (after making the observation). We provide a sound and complete axiomatization for the logic, and consider the complexity of the decision problem. Although the reasoning in the logic is mainly propositional, we allow variables representing numbers and quantification over them. This expressive power seems necessary to capture important properties of evidence.

1. Introduction

Consider the following situation, essentially taken from Halpern and Tuttle (1993) and Fagin and Halpern (1994). A coin is tossed, which is either fair or double-headed. The coin lands heads. How likely is it that the coin is double-headed? What if the coin is tossed 20 times and it lands heads each time? Intuitively, it is much more likely that the coin is double-headed in the latter case than in the former. But how should the likelihood be measured? We cannot simply compute the probability of the coin being double-headed; assigning a probability to that event requires that we have a prior probability on the coin being double-headed. For example, if the coin was chosen at random from a barrel with one billion fair coins and one double-headed coin, it is still overwhelmingly likely that the coin is fair, and that the sequence of 20 heads is just unlucky. However, in the problem statement, the prior probability is not given. We can show than any given prior probability on the coin being double-headed increases significantly as a result of seeing 20 heads. But, intuitively, it seems that we should be able to say that seeing 20 heads in a row provides a great deal of *evidence* in favor of the coin being double-headed without invoking a prior. There has been a great deal of work in trying to make this intuition precise, which we now review.

The main feature of the coin example is that it involves a combination of probabilistic outcomes (e.g., the coin tosses) and nonprobabilistic outcomes (e.g., the choice of the coin). There has been a great deal of work on reasoning about systems that combine both probabilistic and nondeterministic choices; see, for example, Vardi (1985), Fischer and Zuck (1988), Halpern, Moses, and Tuttle (1988), Halpern and Tuttle (1993), de Alfaro (1998), He, Seidel, and McIver (1997). However, the observations above suggest that if we attempt to formally analyze this situation in one of those frameworks, which essentially permit only the modeling of probabilities, we will not be able to directly capture this intuition about increasing likelihood. To see how this plays out, consider a formal analysis of the situation in the Halpern-Tuttle (1993) framework. Suppose that Alice nonprobabilistically chooses one of two coins: a fair coin with probability 1/2 of landing heads, or a double-headed coin with probability 1 of landing heads. Alice tosses this coin repeatedly. Let φ_k be a formula stating: "the kth coin toss lands heads". What is the probability of φ_k according to Bob, who does not know which coin Alice chose, or even the probability of Alice's choice?

According to the Halpern-Tuttle framework, this can be modeled by considering the set of runs describing the states of the system at each point in time, and partitioning this set into two subsets, one for each coin used. In the set of runs where the fair coin is used, the probability of φ_k is 1/2; in the set of runs where the double-headed coin is used, the probability of φ_k is 1. In this setting, the only conclusion that can be drawn is $(\Pr_B(\varphi_k) = 1/2) \vee (\Pr_B(\varphi_k) = 1)$. (This is of course the probability from Bob's point of view; Alice presumably knows which coin she is using.) Intuitively, this seems reasonable: if the fair coin is chosen, the probability that the kth coin toss lands heads, according to Bob, is 1/2; if the double-headed coin is chosen, the probability is 1. Since Bob does not know which of the coins is being used, that is all that can be said.

But now suppose that, before the 101st coin toss, Bob learns the result of the first 100 tosses. Suppose, moreover, that all of these landed heads. What is the probability that the 101st coin toss lands heads? By the same analysis, it is still either 1/2 or 1, depending on which coin is used.

This is hardly useful. To make matters worse, no matter how many coin tosses Bob witnesses, the probability that the next toss lands heads remains unchanged. But this answer misses out on some important information. The fact that all of the first 100 coin tosses are heads is very strong *evidence* that the coin is in fact double-headed. Indeed, a straightforward computation using Bayes' Rule shows that if the prior probability of the coin being double-headed is α, then after observing that all of the 100 tosses land heads, the probability of the coin being double-headed becomes

$$\frac{\alpha}{\alpha + 2^{-100}(1-\alpha)} = \frac{2^{100}\alpha}{2^{100}\alpha + (1-\alpha)}.$$

However, note that it is not possible to determine the posterior probability that the coin is double-headed (or that the 101st coin toss is heads) without the prior probability α. After all, if Alice chooses the double-headed coin with probability only 10^{-100}, then it is still overwhelmingly likely that the coin used is in fact fair, and that Bob was just very unlucky to see such an unrepresentative sequence of coin tosses.

None of the frameworks described above for reasoning about nondeterminism and probability takes the issue of evidence into account. On the other hand, evidence has been discussed extensively in the philosophical literature. Much of this discussion occurs in the philosophy of science, specifically *confirmation theory*, where the concern has been historically to assess the support that evidence obtained through experimentation lends to various scientific theories (Carnap, 1962; Popper, 1959; Good, 1950; Milne, 1996). (Kyburg (1983) provides a good overview of the literature.)

In this paper, we introduce a logic for reasoning about evidence. Our logic extends a logic defined by Fagin, Halpern and Megiddo (1990) (FHM from now on) for reasoning about likelihood expressed as either probability or belief. The logic has first-order quantification over the reals (so includes the theory of real closed fields), as does the FHM logic, for reasons that will shortly become clear. We add observations to the states, and provide an additional operator to talk about the evidence provided by particular observations. We also refine the language to talk about both the prior probability of hypotheses and the posterior probability of hypotheses, taking into account the observation at the states. This lets us write formulas that talk about the relationship between the prior probabilities, the posterior probabilities, and the evidence provided by the observations.

We then provide a sound and complete axiomatization for the logic. To obtain such an axiomatization, we seem to need first-order quantification in a fundamental way. Roughly speaking, this is because ensuring that the evidence operator has the appropriate properties requires us to assert the existence of suitable probability measures. It does not seem possible to do this without existential quantification. Finally, we consider the complexity of the satisfiability problem. The complexity problem for the full language requires exponential space, since it incorporates the theory of real closed fields, for which an exponential-space lower bound is known (Ben-Or, Kozen, & Reif, 1986). However, we show that the satisfiability problem for a propositional fragment of the language, which is still strong enough to allow us to express many properties of interest, is decidable in polynomial space.

It is reasonable to ask at this point why we should bother with a logic of evidence. Our claim is that many decisions in practical applications are made on the basis of evidence. To take an example from security, consider an enforcement mechanism used to detect and react to intrusions in a computer system. Such an enforcement mechanism analyzes the behavior of users and attempts to recognize intruders. Clearly the mechanism wants to make sensible decisions based on observations of user behaviors. How should it do this? One way is to think of an enforcement mechanism as accumulating evidence for or against the hypothesis that the user is an intruder. The accumulated evidence can then be used as the basis for a decision to quarantine a user. In this context, it is not clear that there is a reasonable way to assign a prior probability on whether a user is an intruder. If we want to specify the behavior of such systems and prove that they meet their specifications, it is helpful to have a logic that allows us to do this. We believe that the logic we propose here is the first to do so.

The rest of the paper is organized as follows. In the next section, we formalize a notion of evidence that captures the intuitions outlined above. In Section 3, we introduce our logic for reasoning about evidence. In Section 4, we present an axiomatization for the logic and show that it is sound and complete with respect to the intended models. In Section 5, we discuss the complexity of the decision problem of our logic. In Section 6, we examine some alternatives to the definition of weight of evidence we use. For ease of exposition, in most of the paper, we consider a system where there are only two time points: before and after the observation. In Section 7, we extend our work to dynamic systems, where there can be multiple pieces of evidence, obtained at different points in time. The proofs of our technical results can be found in the appendix.

2. Measures of Confirmation and Evidence

In order to develop a logic for reasoning about evidence, we need to first formalize an appropriate notion of evidence. In this section, we review various formalizations from the literature, and discuss the formalization we use. Evidence has been studied in depth in the philosophical literature, under the name of *confirmation theory*. Confirmation theory aims at determining and measuring the support a piece of evidence provides an hypothesis. As we mentioned in the introduction, many different measures of confirmation have been proposed in the literature. Typically, a proposal has been judged on the degree to which it satisfies various properties that are considered appropriate for confirmation. For example, it may be required that a piece of evidence e confirms an hypothesis h if and only if e makes h more probable. We have no desire to enter the debate as to which class of measures of confirmation is more appropriate. For our purposes, most confirmation functions are inappropriate: they assume that we are given a prior on the set of hypotheses and observations. By marginalization, we also have a prior on hypotheses, which is exactly the information we do not have and do not want to assume. One exception is measures of evidence that use the *log-likelihood ratio*. In this case, rather than having a prior on hypotheses and observations, it suffices that there be a probability μ_h on observations for each hypothesis h: intuitively, $\mu_h(ob)$ is the probability of observing ob when h holds. Given an observation ob, the degree of confirmation that it provides for an hypothesis h is

$$l(ob, h) = \log\left(\frac{\mu_h(ob)}{\mu_{\overline{h}}(ob)}\right),$$

where $\overline{h}$ represents the hypothesis other than h (recall that this approach applies only if there are two hypotheses). Thus, the degree of confirmation is the ratio between these two probabilities. The use of the logarithm is not critical here. Using it ensures that the likelihood is positive if and only if the observation confirms the hypothesis. This approach has been advocated by Good (1950, 1960), among others.[1]

One problem with the log-likelihood ratio measure l as we have defined it is that it can be used only to reason about evidence discriminating between two competing hypotheses, namely between an hypothesis h holding and the hypothesis h not holding. We would like a measure of confirmation along the lines of the log-likelihood ratio measure, but that can handle multiple competing hypotheses. There have been a number of such generalizations, for example, by Pearl (1988) and Chan and Darwiche (2005). We focus here on the generalization given by Shafer (1982) in the context of the Dempster-Shafer theory of evidence based on belief functions (Shafer, 1976); it was further studied by Walley (1987). The description here is taken mostly from Halpern and Fagin (1992). While this measure of confirmation has a number of nice properties of which we take advantage, much of the work presented in this paper can be adapted to different measures of confirmation.

We start with a finite set $\mathcal{H}$ of mutually exclusive and exhaustive hypotheses; thus, exactly one hypothesis holds at any given time. Let $\mathcal{O}$ be the set of possible observations (or pieces of evidence). For simplicity, we assume that $\mathcal{O}$ is finite. Just as in the case of log-likelihood, we also assume that, for each hypotheses $h \in \mathcal{H}$, there is a probability measure μ_h on $\mathcal{O}$ such that $\mu_h(ob)$ is the probability of ob if hypothesis h holds. Furthermore, we assume that the observations in $\mathcal{O}$ are relevant to the hypotheses: for every observation $ob \in \mathcal{O}$, there must be an hypothesis h such that $\mu_h(ob) > 0$. (The measures μ_h are often called *likelihood functions* in the literature.) We define an *evidence space* (over $\mathcal{H}$ and $\mathcal{O}$)

1. Another related approach, the *Bayes factor* approach, is based on taking the ratio of *odds* rather than likelihoods (Good, 1950; Jeffrey, 1992). We remark that in the literature, confirmation is usually taken with respect to some background knowledge. For ease of exposition, we ignore background knowledge here, although it can easily be incorporated into the framework we present.

to be a tuple $\mathcal{E} = (\mathcal{H}, \mathcal{O}, \boldsymbol{\mu})$, where $\boldsymbol{\mu}$ is a function that assigns to every hypothesis $h \in \mathcal{H}$ the likelihood function $\boldsymbol{\mu}(h) = \mu_h$. (For simplicity, we usually write μ_h for $\boldsymbol{\mu}(h)$, when the the function $\boldsymbol{\mu}$ is clear from context.)

Given an evidence space $\mathcal{E}$, we define the weight that the observation ob lends to hypothesis h, written $w_{\mathcal{E}}(ob, h)$, as

$$w_{\mathcal{E}}(ob, h) = \frac{\mu_h(ob)}{\sum_{h' \in \mathcal{H}} \mu_{h'}(ob)}. \tag{1}$$

The measure $w_{\mathcal{E}}$ always lies between 0 and 1; intuitively, if $w_{\mathcal{E}}(ob, h) = 1$, then ob fully confirms h (i.e., h is certainly true if ob is observed), while if $w_{\mathcal{E}}(ob, h) = 0$, then ob disconfirms h (i.e., h is certainly false if ob is observed). Moreover, for each fixed observation ob for which $\sum_{h \in \mathcal{H}} \mu_h(ob) > 0$, $\sum_{h \in \mathcal{H}} w_{\mathcal{E}}(ob, h) = 1$, and thus the weight of evidence $w_{\mathcal{E}}$ looks like a probability measure for each ob. While this has some useful technical consequences, one should not interpret $w_{\mathcal{E}}$ as a probability measure. Roughly speaking, the weight $w_{\mathcal{E}}(ob, h)$ is the likelihood that h is the right hypothesis in the light of observation ob.[2] The advantages of $w_{\mathcal{E}}$ over other known measures of confirmation are that (a) it is applicable when we are not given a prior probability distribution on the hypotheses, (b) it is applicable when there are more than two competing hypotheses, and (c) it has a fairly intuitive probabilistic interpretation.

An important problem in *statistical inference* (Casella & Berger, 2001) is that of choosing the best parameter (i.e., hypothesis) that explains observed data. When there is no prior on the parameters, the "best" parameter is typically taken to be the one that maximizes the likelihood of the data given that parameter. Since $w_{\mathcal{E}}$ is just a normalized likelihood function, the parameter that maximizes the likelihood will also maximize $w_{\mathcal{E}}$. Thus, if all we are interested in is maximizing likelihood, there is no need to normalize the evidence as we do. We return to the issue of normalization in Section 6.[3]

Note that if $\mathcal{H} = \{h_1, h_2\}$, then $w_{\mathcal{E}}$ in some sense generalizes the log-likelihood ratio measure. More precisely, for a fixed observation ob, $w_{\mathcal{E}}(ob, \cdot)$ induces the same relative order on hypotheses as $l(ob, \cdot)$, and for a fixed hypothesis h, $w_{\mathcal{E}}(\cdot, h)$ induces the same relative order on observations as $l(\cdot, h)$.

Proposition 2.1: *For all ob, we have $w_{\mathcal{E}}(ob, h_i) \geq w_{\mathcal{E}}(ob, h_{3-i})$ if and only if $l(ob, h_i) \geq l(ob, h_{3-i})$, for $i = 1, 2$, and for all h, ob, and ob', we have $w_{\mathcal{E}}(ob, h) \geq w_{\mathcal{E}}(ob', h)$ if and only if $l(ob, h) \geq l(ob', h)$.*

2. We could have taken the log of the ratio to make $w_{\mathcal{E}}$ parallel the log-likelihood ratio l defined earlier, but there are technical advantages in having the weight of evidence be a number between 0 and 1.
3. Another representation of evidence that has similar characteristics to $w_{\mathcal{E}}$ is Shafer's original representation of evidence via belief functions (Shafer, 1976), defined as

$$w_{\mathcal{E}}^S(ob, h) = \frac{\mu_h(ob)}{\max_{h \in \mathcal{H}} \mu_h(ob)}.$$

This measure is known in statistical hypothesis testing as the *generalized likelihood-ratio statistic*. It is another generalization of the log-likelihood ratio measure l. The main difference between $w_{\mathcal{E}}$ and $w_{\mathcal{E}}^S$ is how they behave when one considers the combination of evidence, which we discuss later in this section. As Walley (1987) and Halpern and Fagin (1992) point out, $w_{\mathcal{E}}$ gives more intuitive results in this case. We remark that the parameter (hypothesis) that maximized likelihood also maximizes $w_{\mathcal{E}}^S$, so $w_{\mathcal{E}}^S$ can also be used in statistical inference.

Although $w_{\mathcal{E}}(ob, \cdot)$ behaves like a probability measure on hypotheses for every observation ob, one should not think of it as a probability; the weight of evidence of a combined hypothesis, for instance, is not generally the sum of the weights of the individual hypotheses (Halpern & Pucella, 2005a). Rather, $w_{\mathcal{E}}(ob, \cdot)$ is an encoding of evidence. But what is evidence? Halpern and Fagin (1992) have suggested that evidence can be thought of as a *function* mapping a prior probability on the hypotheses to a posterior probability, based on the observation made. There is a precise sense in which $w_{\mathcal{E}}$ can be viewed as a function that maps a prior probability μ_0 on the hypotheses $\mathcal{H}$ to a posterior probability μ_{ob} based on observing ob, by applying Dempster's Rule of Combination (Shafer, 1976). That is,

$$\mu_{ob} = \mu_0 \oplus w_{\mathcal{E}}(ob, \cdot), \tag{2}$$

where $\oplus$ combines two probability distributions on $\mathcal{H}$ to get a new probability distribution on $\mathcal{H}$ defined as follows:

$$(\mu_1 \oplus \mu_2)(H) = \frac{\sum_{h \in H} \mu_1(h)\mu_2(h)}{\sum_{h \in \mathcal{H}} \mu_1(h)\mu_2(h)}.$$

(Dempster's Rule of Combination is used to combine belief functions. The definition of $\oplus$ is more complicated when considering arbitrary belief functions, but in the special case where the belief functions are in fact probability measures, it takes the form we give here.)

Bayes' Rule is the standard way of updating a prior probability based on an observation, but it is only applicable when we have a joint probability distribution on both the hypotheses and the observations (or, equivalently, a prior on hypotheses together with the likelihood functions μ_h for $h \in \mathcal{H}$), something which we do not want to assume we are given. In particular, while we are willing to assume that we are given the likelihood functions, we are not willing to assume that we are given a prior on hypotheses. Dempster's Rule of Combination essentially "simulates" the effects of Bayes' Rule. The relationship between Dempster's Rule and Bayes' Rule is made precise by the following well-known theorem.

Proposition 2.2: (Halpern & Fagin, 1992) *Let $\mathcal{E} = (\mathcal{H}, \mathcal{O}, \boldsymbol{\mu})$ be an evidence space. Suppose that P is a probability on $\mathcal{H} \times \mathcal{O}$ such that $P(\mathcal{H} \times \{ob\} \mid \{h\} \times \mathcal{O}) = \mu_h(ob)$ for all $h \in \mathcal{H}$ and all $ob \in \mathcal{O}$. Let μ_0 be the probability on $\mathcal{H}$ induced by marginalizing P; that is, $\mu_0(h) = P(\{h\} \times \mathcal{O})$. For $ob \in \mathcal{O}$, let $\mu_{ob} = \mu_0 \oplus w_{\mathcal{E}}(ob, \cdot)$. Then $\mu_{ob}(h) = P(\{h\} \times \mathcal{O} \mid \mathcal{H} \times \{ob\})$.*

In other words, when we do have a joint probability on the hypotheses and observations, then Dempster's Rule of Combination gives us the same result as a straightforward application of Bayes' Rule.

Example 2.3: To get a feel for how this measure of evidence can be used, consider a variation of the two-coins example in the introduction. Assume that the coin chosen by Alice is either double-headed or fair, and consider sequences of a hundred tosses of that coin. Let $\mathcal{O} = \{m : 0 \leq m \leq 100\}$ (the number of heads observed), and let $\mathcal{H} = \{F, D\}$, where F is "the coin is fair", and D is "the coin is double-headed". The probability spaces associated with the hypotheses are generated by the following probabilities for simple observations m:

$$\mu_F(m) = \frac{1}{2^{100}}\binom{100}{m} \qquad \mu_D(m) = \begin{cases} 1 & \text{if } m = 100 \\ 0 & \text{otherwise.} \end{cases}$$

(We extend by additivity to the whole set $\mathcal{O}$.) Take $\mathcal{E} = (\mathcal{H}, \mathcal{O}, \boldsymbol{\mu})$, where $\boldsymbol{\mu}(F) = \mu_F$ and $\boldsymbol{\mu}(D) = \mu_D$. For any observation $m \neq 100$, the weight in favor of F is given by

$$w_{\mathcal{E}}(m, F) = \frac{\frac{1}{2^{100}}\binom{100}{m}}{0 + \frac{1}{2^{100}}\binom{100}{m}} = 1,$$

which means that the support of m is unconditionally provided to F; indeed, any such sequence of tosses cannot appear with the double-headed coin. Thus, if $m \neq 100$, we get that

$$w_{\mathcal{E}}(m, D) = \frac{0}{0 + \frac{1}{2^{100}}\binom{100}{m}} = 0.$$

What happens when the hundred coin tosses are all heads? It is straightforward to check that

$$w_{\mathcal{E}}(100, F) = \frac{\frac{1}{2^{100}}}{1 + \frac{1}{2^{100}}} = \frac{1}{1 + 2^{100}} \qquad w_{\mathcal{E}}(100, D) = \frac{1}{1 + \frac{1}{2^{100}}} = \frac{2^{100}}{1 + 2^{100}};$$

this time there is overwhelmingly more evidence in favor of D than F.

Note that we have not assumed any prior probability. Thus, we cannot talk about the probability that the coin is fair or double-headed. What we have is a quantitative assessment of the evidence in favor of one of the hypotheses. However, if we assume a prior probability α on the coin being fair and m heads are observed after 100 tosses, then the probability that the coin is fair is 1 if $m \neq 100$; if $m = 100$ then, applying the rule of combination, the posterior probability of the coin being fair is $\alpha/(\alpha + (1-\alpha)2^{100})$. $\square$

Can we characterize weight functions using a small number of properties? More precisely, given sets $\mathcal{H}$ and $\mathcal{O}$, and a function f from $\mathcal{O} \times \mathcal{H}$ to $[0,1]$, are there properties of f that ensure that there are likelihood functions $\boldsymbol{\mu}$ such that $f = w_{\mathcal{E}}$ for $\mathcal{E} = (\mathcal{H}, \mathcal{O}, \boldsymbol{\mu})$? As we saw earlier, for a fixed observation ob, f essentially acts like a probability measure on $\mathcal{H}$. However, this is not sufficient to guarantee that f is a weight function. Consider the following example, with $\mathcal{O} = \{ob_1, ob_2\}$ and $\mathcal{H} = \{h_1, h_2, h_3\}$:

$$\begin{array}{rclcrcl} f(ob_1, h_1) & = & 1/4 & \qquad & f(ob_2, h_1) & = & 1/4 \\ f(ob_1, h_2) & = & 1/4 & & f(ob_2, h_2) & = & 1/2 \\ f(ob_1, h_3) & = & 1/2 & & f(ob_2, h_3) & = & 1/4. \end{array}$$

It is straightforward to check that $f(ob_1, \cdot)$ and $f(ob_2, \cdot)$ are probability measures on $\mathcal{H}$, but that there is no evidence space $\mathcal{E} = (\mathcal{H}, \mathcal{O}, \boldsymbol{\mu})$ such that $f = w_{\mathcal{E}}$. Indeed, assume that we do have such $\mu_{h_1}, \mu_{h_2}, \mu_{h_3}$. By the definition of weight of evidence, and the fact that f is that weight of evidence, we get the following system of equations:

$$\begin{array}{ll} \frac{\mu_{h_1}(ob_1)}{\mu_{h_1}(ob_1)+\mu_{h_2}(ob_1)+\mu_{h_3}(ob_1)} = 1/4 & \qquad \frac{\mu_{h_1}(ob_2)}{\mu_{h_1}(ob_2)+\mu_{h_2}(ob_2)+\mu_{h_3}(ob_2)} = 1/4 \\ \frac{\mu_{h_2}(ob_1)}{\mu_{h_1}(ob_1)+\mu_{h_2}(ob_1)+\mu_{h_3}(ob_1)} = 1/4 & \qquad \frac{\mu_{h_2}(ob_2)}{\mu_{h_1}(ob_2)+\mu_{h_2}(ob_2)+\mu_{h_3}(ob_2)} = 1/2 \\ \frac{\mu_{h_3}(ob_1)}{\mu_{h_1}(ob_1)+\mu_{h_2}(ob_1)+\mu_{h_3}(ob_1)} = 1/2 & \qquad \frac{\mu_{h_3}(ob_2)}{\mu_{h_1}(ob_2)+\mu_{h_2}(ob_2)+\mu_{h_3}(ob_2)} = 1/4. \end{array}$$

It is now immediate that there exist α_1 and α_2 such that $\mu_{h_i}(ob_j) = \alpha_j f(ob_j, h_i)$, for $i = 1, 2, 3$. Indeed, $\alpha_j = \mu_{h_1}(ob_j) + \mu_{h_2}(ob_j) + \mu_{h_3}(ob_j)$, for $j = 1, 2$. Moreover, since μ_{h_i} is a probability measure, we must have that

$$\mu_{h_i}(ob_1) + \mu_{h_i}(ob_2) = \alpha_1 f(ob_1, h_i) + \alpha_2 f(ob_2, h_i) = 1,$$

for $i = 1, 2, 3$. Thus,

$$\alpha_1/4 + \alpha_2/4 = \alpha_1/4 + \alpha_2/2 = \alpha_1/2 + \alpha_4/4 = 1.$$

These constraints are easily seen to be unsatisfiable.

This argument generalizes to arbitrary functions f; thus, a necessary condition for f to be a weight function is that there exists α_i for each observation ob_i such that $\mu_h(ob_i) = \alpha_i f(ob_i, h)$ for each hypothesis h is a probability measure, that is, $\alpha_1 f(ob_1, h) + \cdots + \alpha_k f(ob_k, h) = 1$. In fact, when combined with the constraint that $f(ob, \cdot)$ is a probability measure for a fixed ob, this condition turns out to be sufficient, as the following theorem establishes.

Theorem 2.4: *Let $\mathcal{H} = \{h_1, \ldots, h_m\}$ and $\mathcal{O} = \{ob_1, \ldots, ob_n\}$, and let f be a real-valued function with domain $\mathcal{O} \times \mathcal{H}$ such that $f(ob, h) \in [0, 1]$. Then there exists an evidence space $\mathcal{E} = (\mathcal{H}, \mathcal{O}, \boldsymbol{\mu})$ such that $f = w_{\mathcal{E}}$ if and only if f satisfies the following properties:*

WF1. *For every $ob \in \mathcal{O}$, $f(ob, \cdot)$ is a probability measure on $\mathcal{H}$.*

WF2. *There exists $x_1, \ldots, x_n > 0$ such that, for all $h \in \mathcal{H}$, $\sum_{i=1}^{n} f(ob_i, h) x_i = 1$.*

This characterization is fundamental to the completeness of the axiomatization of the logic we introduce in the next section. The characterization is complicated by the fact that the weight of evidence is essentially a normalized likelihood: the likelihood of an observation given a particular hypothesis is normalized using the sum of all the likelihoods of that observation, for all possible hypotheses. One consequence of this, as we already mentioned above, is that the weight of evidence is always between 0 and 1, and superficially behaves like a probability measure. In Section 6, we examine the issue of normalization more carefully, and describe the changes to our framework that would occur were we to take unnormalized likelihoods as weight of evidence.

Let $\mathcal{E} = (\mathcal{H}, \mathcal{O}, \boldsymbol{\mu})$ be an evidence space. Let $\mathcal{O}^*$ be the set of sequences of observations $\langle ob^1, \ldots, ob^k \rangle$ over $\mathcal{O}$.[4] Assume that the observations are independent, that is, for each basic hypothesis h, take $\mu_h^*(\langle ob^1, \ldots, ob^k \rangle)$, the probability of observing a particular sequence of observations given h, to be $\mu_h(ob^1) \cdots \mu_h(ob^k)$, the product of the probability of making each observation in the sequence. Let $\mathcal{E}^* = (\mathcal{H}, \mathcal{O}^*, \boldsymbol{\mu}^*)$. With this assumption, it is well known that Dempster's Rule of Combination can be used to combine evidence in this setting; that is,

$$w_{\mathcal{E}^*}(\langle ob^1, \ldots, ob^k \rangle, \cdot) = w_{\mathcal{E}}(ob^1, \cdot) \oplus \cdots \oplus w_{\mathcal{E}}(ob^k, \cdot)$$

(Halpern & Fagin, 1992, Theorem 4.3). It is an easy exercise to check that the weight provided by the sequence of observations $\langle ob^1, \ldots, ob^k \rangle$ can be expressed in terms of the weight of the individual observations:

$$w_{\mathcal{E}^*}(\langle ob^1, \ldots, ob^k \rangle, h) = \frac{w_{\mathcal{E}^*}(ob^1, h) \cdots w_{\mathcal{E}^*}(ob^k, h)}{\sum_{h' \in \mathcal{H}} w_{\mathcal{E}^*}(ob^1, h') \cdots w_{\mathcal{E}^*}(ob^k, h')}. \tag{3}$$

4. We use superscript rather than subscripts to index observations in a sequence so that these observations will not be confused with the basic observations $ob_1, \ldots, ob_n$ in $\mathcal{O}$.

If we let μ_0 be a prior probability on the hypotheses, and $\mu_{\langle ob^1,\ldots,ob^k\rangle}$ be the probability on the hypotheses after observing $ob^1,\ldots,ob^k$, we can verify that

$$\mu_{\langle ob^1,\ldots,ob^k\rangle} = \mu_0 \oplus w_{\mathcal{E}^*}(\langle ob^1,\ldots,ob^k\rangle,\cdot).$$

Example 2.5: Consider a variant of Example 2.3, where we take the coin tosses as individual observations, rather than the number of heads that turn up in one hundred coin tosses. As before, assume that the coin chosen by Alice is either double-headed or fair. Let $\mathcal{O} = \{H, T\}$, the result of an individual coin toss, where H is "the coin landed heads" and T is "the coin landed tails". Let $\mathcal{H} = \{F, D\}$, where F is "the coin is fair", and D is "the coin is double-headed". Let $\mathcal{E}^* = (\mathcal{H}, \mathcal{O}^*, \boldsymbol{\mu}^*)$. The probability measure μ_h^* associated with the hypothesis h are generated by the following probabilities for simple observations:

$$\mu_F(H) = \frac{1}{2} \qquad \mu_D(H) = 1.$$

Thus, for example, $\mu_F^*(\langle H,H,T,H\rangle) = 1/16$, $\mu_D^*(\langle H,H,H\rangle) = 1$, and $\mu_H^*(\langle H,H,T,H\rangle) = 0$.

We can now easily verify results similar to those that were obtained in Example 2.3. For instance, the weight of observing T in favor of F is given by

$$w_{\mathcal{E}^*}(T,F) = \frac{\frac{1}{2}}{0+\frac{1}{2}} = 1,$$

which again indicates that observing T provides unconditional support to F; a double-headed coin cannot land tails.

How about sequences of observations? The weight provided by the sequence $\langle ob^1,\ldots,ob^k\rangle$ for hypothesis h is given by Equation (3). Thus, if $\underline{H} = \langle H,\ldots,H\rangle$, a sequence of a hundred coin tosses, we can check that

$$w_{\mathcal{E}^*}(\underline{H},F) = \frac{\frac{1}{2^{100}}}{1+\frac{1}{2^{100}}} = \frac{1}{1+2^{100}} \qquad w_{\mathcal{E}^*}(\underline{H},D) = \frac{1}{1+\frac{1}{2^{100}}} = \frac{2^{100}}{1+2^{100}}.$$

Unsurprisingly, this is the same result as in Example 2.3. □

3. Reasoning about Evidence

We introduce a logic $\mathcal{L}^{fo\text{-}ev}$ for reasoning about evidence, inspired by a logic introduced in FHM for reasoning about probability. The logic lets us reason about the weight of evidence of observations for hypotheses; moreover, to be able to talk about the relationship between prior probabilities, evidence, and posterior probabilities, we provide operators to reason about the prior and posterior probabilities of hypotheses. We remark that up to now we have been somewhat agnostic about whether the priors exist but are not given (or not known) or whether the prior does not exist at all. It is beyond the scope of this paper to enter the debate about whether it always appropriate to assume the existence of a prior. Although the definition of evidence makes sense even if the priors does not exist, our logic implicitly assumes that there are priors (although they may not be known), since we provide operators for reasoning about the prior. We make use of these operators in some of the examples below. However, the fragment of the logic that does not use these operators is appropriate for prior-free reasoning.

The logic has both propositional features and first-order features. We take the probability of propositions and the weight of evidence of observations for hypotheses, and view probability and evidence as propositions, but we allow first-order quantification over numerical quantities, such as probabilities and evidence. The logic essentially considers two time periods, which can be thought of as the time before an observation is made and the time after an observation is made. In this section, we assume that exactly one observation is made. (We consider sequences of observations in Section 7.) Thus, we can talk of the probability of a formula φ before an observation is made, denoted $\Pr^0(\varphi)$, the probability of φ after the observation, denoted $\Pr(\varphi)$, and the evidence provided by the observation ob for an hypothesis h, denoted $\mathrm{w}(ob, h)$. Of course, we want to be able to use the logic to relate all these quantities.

Formally, we start with two finite sets of primitive propositions, $\Phi_{\mathsf{h}} = \{h_1, \ldots, h_{n_{\mathsf{h}}}\}$ representing the hypotheses, and $\Phi_{\mathsf{o}} = \{ob_1, \ldots, ob_{n_{\mathsf{o}}}\}$ representing the observations. Let $\mathcal{L}_{\mathsf{h}}(\Phi_{\mathsf{h}})$ be the propositional sublanguage of *hypothesis formulas* obtained by taking primitive propositions in Φ_{h} and closing off under negation and conjunction; we use ρ to range over formulas of that sublanguage.

A *basic term* has the form $\Pr^0(\rho)$, $\Pr(\rho)$, or $\mathrm{w}(ob, h)$, where ρ is an hypothesis formula, ob is an observation, and h is an hypothesis. As we said, we interpret $\Pr^0(\rho)$ as the prior probability of ρ, $\Pr(\rho)$ as the posterior probability of ρ, and $\mathrm{w}(ob, h)$ as the weight of evidence of observation ob for hypothesis h. It may seem strange that we allow the language to talk about the prior probability of hypotheses, although we have said that we do not want to assume that the prior is known. We could, of course, simplify the syntax so that it did not include formulas of the form $\Pr^0(\rho)$ or $\Pr(\rho)$. The advantage of having them is that, even if the prior is not known, given our view of evidence as a function from priors to posteriors, we can make statements such as "if the prior probability of h is 2/3, ob is observed, and the weight of evidence of ob for h is 3/4, then the posterior probability of h is 6/7; this is just

$$\Pr^0(h) = 1/2 \wedge ob \wedge \mathrm{w}(ob, h) = 3/4 \Rightarrow \Pr(h) = 6/7.$$

A *polynomial term* has the form $t_1 + \cdots + t_n$, where each term t_i is a product of integers, basic terms, and variables (which range over the reals). A *polynomial inequality formula* has the form $p \geq c$, where p is a polynomial term and c is an integer. Let $\mathcal{L}^{fo\text{-}ev}(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$ be the language obtained by starting out with the primitive propositions in Φ_{h} and Φ_{o} and polynomial inequality formulas, and closing off under conjunction, negation, and first-order quantification. Let *true* be an abbreviation for an arbitrary propositional tautology involving only hypotheses, such as $h_1 \vee \neg h_1$; let *false* be an abbreviation for $\neg$*true*. With this definition, *true* and *false* can be considered as part of the sublanguage $\mathcal{L}_{\mathsf{h}}(\Phi_{\mathsf{h}})$.

It should be clear that while we allow only integer coefficients to appear in polynomial terms, we can in fact express polynomial terms with rational coefficients by crossmultiplying. For instance, $\frac{1}{3}\Pr(\rho) + \frac{1}{2}\Pr(\rho') \geq 1$ can be represented by the polynomial inequality formula $2\Pr(\rho) + 3\Pr(\rho') \geq 6$. While there is no difficulty in giving a semantics to polynomial terms that use arbitrary real coefficients, we need the restriction to integers in order to make use of results from the theory of real closed fields in both the axiomatization of Section 4 and the complexity results of Section 5.

We use obvious abbreviations where needed, such as $\varphi \vee \psi$ for $\neg(\neg\varphi \wedge \neg\psi)$, $\varphi \Rightarrow \psi$ for $\neg\varphi \vee \psi$, $\exists x\varphi$ for $\neg\forall x(\neg\varphi)$, $\Pr(\varphi) - \Pr(\psi) \geq c$ for $\Pr(\varphi) + (-1)\Pr(\psi) \geq c$, $\Pr(\varphi) \geq \Pr(\psi)$ for $\Pr(\varphi) - \Pr(\psi) \geq 0$, $\Pr(\varphi) \leq c$ for $-\Pr(\varphi) \geq -c$, $\Pr(\varphi) < c$ for $\neg(\Pr(\varphi) \geq c)$, and $\Pr(\varphi) = c$ for $(\Pr(\varphi) \geq c) \wedge (\Pr(\varphi) \leq c)$ (and analogous abbreviations for inequalities involving $\Pr^0$ and w).

Example 3.1: Consider again the situation given in Example 2.3. Let Φ_{o}, the observations, consist of primitive propositions of the form $\mathsf{heads}[m]$, where m is an integer with $0 \leq m \leq 100$, indicating that m heads out of 100 tosses have appeared. Let Φ_{h} consist of the two primitive propositions fair and $\mathsf{doubleheaded}$. The computations in Example 2.3 can be written as follows:

$$\mathsf{w}(\mathsf{heads}[100], \mathsf{fair}) = 1/(1 + 2^{100}) \wedge \mathsf{w}(\mathsf{heads}[100], \mathsf{doubleheaded}) = 2^{100}/(1 + 2^{100}).$$

We can also capture the fact that the weight of evidence of an observation maps a prior probability into a posterior probability by Dempster's Rule of Combination. For example, the following formula captures the update of the prior probability α of the hypothesis fair upon observation of a hundred coin tosses landing heads:

$$\Pr^0(\mathsf{fair}) = \alpha \wedge \mathsf{w}(\mathsf{heads}[100], \mathsf{fair}) = 1/(1 + 2^{100}) \Rightarrow \Pr(\mathsf{fair}) = \alpha/(\alpha + (1 - \alpha)2^{100}).$$

We develop a deductive system to derive such conclusions in the next section. □

Now we consider the semantics. A formula is interpreted in a world that specifies which hypothesis is true and which observation was made, as well as an evidence space to interpret the weight of evidence of observations and a probability distribution on the hypotheses to interpret prior probabilities and talk about updating based on evidence. (We do not need to include a posterior probability distribution, since it can be computed from the prior and the weights of evidence using Equation (2).) An *evidential world* is a tuple $w = (h, ob, \mu, \mathcal{E})$, where h is a hypothesis, ob is an observation, μ is a probability distribution on Φ_{h}, and $\mathcal{E}$ is an evidence space over Φ_{h} and Φ_{o}.

To interpret propositional formulas in $\mathcal{L}_{\mathsf{h}}(\Phi_{\mathsf{h}})$, we associate with each hypothesis formula ρ a set $[\![\rho]\!]$ of hypotheses, by induction on the structure of ρ:

$$\begin{aligned} [\![h]\!] &= \{h\} \\ [\![\neg\rho]\!] &= \Phi_{\mathsf{h}} - [\![\rho]\!] \\ [\![\rho_1 \wedge \rho_2]\!] &= [\![\rho_1]\!] \cap [\![\rho_2]\!]. \end{aligned}$$

To interpret first-order formulas that may contain variables, we need a valuation v that assigns a real number to every variable. Given an evidential world $w = (h, ob, \mu, \mathcal{E})$ and a valuation v, we assign to a polynomial term p a real number $[p]^{w,v}$ in a straightforward way:

$$\begin{aligned} [x]^{w,v} &= v(x) \\ [a]^{w,v} &= a \\ [\Pr^0(\rho)]^{w,v} &= \mu([\![\rho]\!]) \end{aligned}$$

$$\begin{aligned}
[\Pr(\rho)]^{w,v} &= (\mu \oplus w_{\mathcal{E}}(ob,\cdot))([\![\rho]\!]) \\
[\mathrm{w}(ob',h')]^{w,v} &= w_{\mathcal{E}}(ob',h') \\
[t_1 t_2]^{w,v} &= [t_1]^{w,v} \times [t_2]^{w,v} \\
[p_1 + p_2]^{w,v} &= [p_1]^{w,v} + [p_2]^{w,v}.
\end{aligned}$$

Note that, to interpret $\Pr(\rho)$, the posterior probability of ρ after having observed ob (the observation at world w), we use Equation (2), which says that the posterior is obtained by combining the prior probability μ with $w_{\mathcal{E}}(ob,\cdot)$.

We define what it means for a formula φ to be true (or satisfied) at an evidential world w under valuation v, written $(w,v) \models \varphi$, as follows:

$(w,v) \models h$ if $w = (h, ob, \mu, \mathcal{E})$ for some ob, μ, $\mathcal{E}$

$(w,v) \models ob$ if $w = (h, ob, \mu, \mathcal{E})$ for some h, μ, $\mathcal{E}$

$(w,v) \models \neg\varphi$ if $(w,v) \not\models \varphi$

$(w,v) \models \varphi \wedge \psi$ if $(w,v) \models \varphi$ and $(w,v) \models \psi$

$(w,v) \models p \geq c$ if $[p]^{w,v} \geq c$

$(w,v) \models \forall x \varphi$ if $(w,v') \models \varphi$ for all v' that agree with v on all variables but x.

If $(w,v) \models \varphi$ is true for all v, we write simply $w \models \varphi$. It is easy to check that if φ is a closed formula (that is, one with no free variables), then $(w,v) \models \varphi$ if and only if $(w,v') \models \varphi$, for all v, v'. Therefore, given a closed formula φ, if $(M,w,v) \models \varphi$, then in fact $w \models \varphi$. We will typically be concerned only with closed formulas. Finally, if $w \models \varphi$ for all evidential worlds w, we write $\models \varphi$ and say that φ is valid. In the next section, we will characterize axiomatically all the valid formulas of the logic.

Example 3.2: The following formula is valid, that is, true in all evidential worlds:

$$\models (\mathrm{w}(ob,h_1) = 2/3 \wedge \mathrm{w}(ob,h_2) = 1/3) \Rightarrow (\Pr^0(h_1) \geq 1/100 \wedge ob) \Rightarrow \Pr(h_1) \geq 2/101.$$

In other words, at all evidential worlds where the weight of evidence of observation ob for hypothesis h_1 is $2/3$ and the weight of evidence of observation ob for hypothesis h_2 is $1/3$, it must be the case that if the prior probability of h_1 is at least $1/100$ and ob is actually observed, then the posterior probability of h_1 is at least $2/101$. This shows the extent to which we can reason about the evidence independently of the prior probabilities. □

The logic imposes no restriction on the prior probabilities to be used in the models. This implies, for instance, that the formula

$$\mathsf{fair} \Rightarrow \Pr^0(\mathsf{fair}) = 0$$

is satisfiable: there exists an evidential world w such that the formula is true at w. In other words, it is consistent for an hypothesis to be true, despite the prior probability of it being true being 0. It is a simple matter to impose a restriction on the models that they be such that if h is true at a world, then $\mu(h) > 0$ for the prior μ at that world.

We conclude this section with some remarks concerning the semantic model. Our semantic model implicitly assumes that the prior probability is known and that the likelihood functions (i.e., the measures μ_h) are known. Of course, in many situations there will be uncertainty about both. Indeed, our motivation for focusing on evidence is precisely to deal with situations where the prior is not known. Handling uncertainty about the prior is easy in our framework, since our notion of evidence is independent of the prior on hypotheses. It is straightforward to extend our model by allowing a set of possible worlds, with a different prior in each, but using the same evidence space for all of them. We can then extend the logic with a knowledge operator, where a statement is known to be true if it is true in all the worlds. This allows us to make statements like "I know that the prior on hypothesis h is between α and β. Since observation ob provides evidence $3/4$ for h, I know that the posterior on h given ob is between $(3\alpha)/(2\alpha+1)$ and $(3\beta)/(2\beta+1)$."

Dealing with uncertainty about the likelihood functions is somewhat more subtle. To understand the issue, suppose that one of two coins will be chosen and tossed. The bias of coin 1 (i.e., the probability that coin 1 lands heads) is between $2/3$ and $3/4$; the bias of coin 2 is between $1/4$ and $1/3$. Here there is uncertainty about the probability that coin 1 will be picked (this is uncertainty about the prior) and there is uncertainty about the bias of each coin (this is uncertainty about the likelihood functions). The problem here is that, to deal with this, we must consider possible worlds where there is a possibly different evidence space in each world. It is then not obvious how to define weight of evidence. We explore this issue in more detail in a companion paper (Halpern & Pucella, 2005a).

4. Axiomatizing Evidence

In this section we present a sound and complete axiomatization $\mathbf{AX}(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$ for our logic.

The axiomatization can be divided into four parts. The first part, consisting of the following axiom and inference rule, accounts for first-order reasoning:

Taut. All substitution instances of valid formulas of first-order logic with equality.

MP. From φ and $\varphi \Rightarrow \psi$ infer ψ.

Instances of **Taut** include, for example, all formulas of the form $\varphi \vee \neg\varphi$, where φ is an arbitrary formula of the logic. It also includes formulas such as $(\forall x \varphi) \Leftrightarrow \varphi$ if x is not free in φ. In particular, $(\forall x(h)) \Leftrightarrow h$ for hypotheses in Φ_{h}, and similarly for observations in Φ_{o}. Note that **Taut** includes all *substitution* instances of valid formulas of first-order logic with equality; in other words, any valid formula of first-order logic with equality where free variables are replaced with arbitrary terms of our language (including $\Pr^0(\rho)$, $\Pr(\rho)$, $\mathrm{w}(ob, h)$) is an instance of **Taut**. Axiom **Taut** can be replaced by a sound and complete axiomatization for first-order logic with equality, as given, for instance, in Shoenfield (1967) or Enderton (1972).

The second set of axioms accounts for reasoning about polynomial inequalities, by relying on the theory of real closed fields:

RCF. All instances of formulas valid in real closed fields (and, thus, true about the reals), with nonlogical symbols $+$, $\cdot$, $<$, 0, 1, -1, 2, -2, 3, -3, $\ldots$.

Formulas that are valid in real closed fields include, for example, the fact that addition on the reals is associative, $\forall x \forall y \forall z((x+y)+z = x+(y+z))$, that 1 is the identity for multiplication, $\forall x(x \cdot 1 = x)$, and formulas relating the constant symbols, such as $k = 1+\cdots+1$ (k times) and $-1+1=0$. As for **Taut**, we could replace **RCF** by a sound and complete axiomatization for real closed fields (cf. Fagin et al., 1990; Shoenfield, 1967; Tarski, 1951).

The third set of axioms essentially captures the fact that there is a single hypothesis and a single observation that holds per state.

H1. $h_1 \vee \cdots \vee h_{n_\mathsf{h}}$.

H2. $h_i \Rightarrow \neg h_j$ if $i \neq j$.

O1. $ob_1 \vee \cdots \vee ob_{n_\mathsf{o}}$.

O2. $ob_i \Rightarrow \neg ob_j$ if $i \neq j$.

These axioms illustrate a subtlety of our logic. Like most propositional logics, ours is parameterized by primitive propositions, in our case, Φ_h and Φ_o. However, while axiomatizations for propositional logics typically do not depend on the exact set of primitive propositions, ours does. Clearly, axiom **H1** is sound only if the hypothesis primitives are exactly $h_1, \ldots, h_{n_\mathsf{h}}$. Similarly, axiom **O1** is sound only if the observation primitives are exactly $ob_1, \ldots, ob_{n_\mathsf{o}}$. It is therefore important for us to identify the primitive propositions when talking about the axiomatization $\mathbf{AX}(\Phi_\mathsf{h}, \Phi_\mathsf{o})$.

The last set of axioms concerns reasoning about probabilities and evidence proper. The axioms for probability are taken from FHM.

Pr1. $\Pr^0(\mathit{true}) = 1$.

Pr2. $\Pr^0(\rho) \geq 0$.

Pr3. $\Pr^0(\rho_1 \wedge \rho_2) + \Pr^0(\rho_1 \wedge \neg\rho_2) = \Pr^0(\rho_1)$.

Pr4. $\Pr^0(\rho_1) = \Pr^0(\rho_2)$ if $\rho_1 \Leftrightarrow \rho_2$ is a propositional tautology.

Axiom **Pr1** simply says that the event *true* has probability 1. Axiom **Pr2** says that probability is nonnegative. Axiom **Pr3** captures finite additivity. It is not possible to express countable additivity in our logic. On the other hand, just as in FHM, we do not need an axiom for countable additivity. Roughly speaking, as we establish in the next section, if a formula is satisfiable at all, it is satisfiable in a finite structure. Similar axioms capture posterior probability formulas:

Po1. $\Pr(\mathit{true}) = 1$.

Po2. $\Pr(\rho) \geq 0$.

Po3. $\Pr(\rho_1 \wedge \rho_2) + \Pr(\rho_1 \wedge \neg\rho_2) = \Pr(\rho_1)$.

Po4. $\Pr(\rho_1) = \Pr(\rho_2)$ if $\rho_1 \Leftrightarrow \rho_2$ is a propositional tautology.

Finally, we need axioms to account for the behavior of the evidence operator w. What are these properties? For one thing, the weight function acts essentially like a probability on hypotheses, for each fixed observation, except that we are restricted to taking the weight of evidence of basic hypotheses only. This gives the following axioms:

E1. $\mathrm{w}(ob, h) \geq 0$.

E2. $\mathrm{w}(ob, h_1) + \cdots + \mathrm{w}(ob, h_{n_\mathsf{h}}) = 1$.

Second, evidence connects the prior and posterior beliefs via Dempster's Rule of Combination, as in (2). This is captured by the following axiom. (Note that, since we do not have division in the language, we crossmultiply to clear the denominator.)

E3. $ob \Rightarrow (\Pr^0(h)\mathrm{w}(ob, h) = \Pr(h)\Pr^0(h_1)\mathrm{w}(ob, h_1) + \cdots + \Pr(h)\Pr^0(h_{n_\mathsf{h}})\mathrm{w}(ob, h_{n_\mathsf{h}}))$.

This is not quite enough. As we saw in Section 2, property **WF2** in Theorem 2.4 is required for a function to be an evidence function. The following axiom captures **WF2** in our logic:

E4. $\exists x_1 \ldots \exists x_{n_\mathsf{o}}(x_1 > 0 \wedge \cdots \wedge x_{n_\mathsf{o}} > 0 \wedge \mathrm{w}(ob_1, h_1)x_1 + \cdots + \mathrm{w}(ob_{n_\mathsf{o}}, h_1)x_{n_\mathsf{o}} = 1 \wedge$
$\cdots \wedge \mathrm{w}(ob_1, h_{n_\mathsf{h}})x_1 + \cdots + \mathrm{w}(ob_{n_\mathsf{o}}, h_{n_\mathsf{h}})x_{n_\mathsf{o}} = 1)$.

Note that axiom **E4** is the only axiom that requires quantification. Moreover, axioms **E3** and **E4** both depend on Φ_h and Φ_o.

As an example, we show that if h and h' are distinct hypotheses in Φ_h, then the formula

$$\neg(\mathrm{w}(ob, h) = 2/3 \wedge \mathrm{w}(ob, h') = 2/3)$$

is provable. First, by **RCF**, the following valid formula of the theory of real closed fields is provable:

$$\forall x \forall y (x = 2/3 \wedge y = 2/3 \Rightarrow x + y > 1).$$

Moreover, if $\varphi(x, y)$ is any first-order logic formula with two free variables x and y, then

$$(\forall x \forall y (\varphi(x, y))) \Rightarrow \varphi(\mathrm{w}(ob, h), \mathrm{w}(ob, h'))$$

is a substitution instance of a valid formula of first-order logic with equality, and hence is an instance of **Taut**. Thus, by **MP**, we can prove that

$$\mathrm{w}(ob, h) = 2/3 \wedge \mathrm{w}(ob, h') = 2/3 \Rightarrow \mathrm{w}(ob, h) + \mathrm{w}(ob, h') > 1,$$

which is provably equivalent (by **Taut** and **MP**) to its contrapositive

$$\mathrm{w}(ob, h) + \mathrm{w}(ob, h') \leq 1 \Rightarrow \neg(\mathrm{w}(ob, h) = 2/3 \wedge \mathrm{w}(ob, h') = 2/3).$$

By an argument similar to that above, using **RCF**, **Taut**, **MP**, **E1**, and **E2**, we can derive

$$\mathrm{w}(ob, h) + \mathrm{w}(ob, h') \leq 1,$$

and by **MP**, we obtain the desired conclusion: $\neg(\mathrm{w}(ob, h) = 2/3 \wedge \mathrm{w}(ob, h') = 2/3)$.

Theorem 4.1: *$\boldsymbol{AX}(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$ is a sound and complete axiomatization for $\mathcal{L}^{fo\text{-}ev}(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$ with respect to evidential worlds.*

As usual, soundness is straightforward, and to prove completeness, it suffices to show that if a formula φ is consistent with $\mathbf{AX}(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$, it is satisfiable in an evidential structure. However, the usual approach for proving completeness in modal logic, which involves considering maximal consistent sets and canonical structures does not work. The problem is that there are maximal consistent sets of formulas that are not satisfiable. For example, there is a maximal consistent set of formulas that includes $\Pr(\rho) > 0$ and $\Pr(\rho) \leq 1/n$ for $n = 1, 2, \ldots$. This is clearly unsatisfiable. Our proof follows the techniques developed in FHM.

To express axiom **E4**, we needed to have quantification in the logic. This is where the fact that our representation of evidence is normalized has a nontrivial effect on the logic: **E4** corresponds to property **WF2**, which essentially says that a function is a weight of evidence function if one can find such a normalization factor. An interesting question is whether it is possible to find a sound and complete axiomatization for the propositional fragment of our logic (without quantification or variables). To do this, we need to give quantifier-free axioms to replace axiom **E4**. This amounts to asking whether there is a simpler property than **WF2** in Theorem 2.4 that characterizes weight of evidence functions. This remains an open question.

5. Decision Procedures

In this section, we consider the decision problem for our logic, that is, the problem of deciding whether a given formula φ is satisfiable. In order to state the problem precisely, however, we need to deal carefully with the fact that the logic is parameterized by the sets Φ_{h} and Φ_{o} of primitive propositions representing hypotheses and observations. In most logics, the choice of underlying primitive propositions is essentially irrelevant. For example, if a propositional formula φ that contains only primitive propositions in some set Φ is true with respect to all truth assignments to Φ, then it remains true with respect to all truth assignments to any set $\Phi' \supseteq \Phi$. This monotonicity property does not hold here. For example, as we have already observed, axiom **H1** clearly depends on the set of hypotheses and observations; it is no longer valid if the set is changed. The same is true for **O1**, **E3**, and **E4**.

This means that we have to be careful, when stating decision problems, about the role of Φ_{h} and Φ_{o} in the algorithm. A straightforward way to deal with this is to assume that the satisfiability algorithm gets as input Φ_{h}, Φ_{o}, and a formula $\varphi \in \mathcal{L}^{fo\text{-}ev}(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$. Because $\mathcal{L}^{fo\text{-}ev}(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$ contains the full theory of real closed fields, it is unsurprisingly difficult to decide. For our decision procedure, we can use the exponential-space algorithm of Ben-Or, Kozen, and Reif (1986) to decide the satisfiability of real closed field formulas. We define the length $|\varphi|$ of φ to be the number of symbols required to write φ, where we count the length of each coefficient as 1. Similarly, we define $\|\varphi\|$ to be the length of the longest coefficient appearing in f, when written in binary.

Theorem 5.1: *There is a procedure that runs in space exponential in $|\varphi|\ \|\varphi\|$ for deciding, given Φ_{h} and Φ_{o}, whether a formula φ of $\mathcal{L}^{fo\text{-}ev}(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$ is satisfiable in an evidential world.*

This is essentially the best we can do, since Ben-Or, Kozen, and Reif (1986) prove that the decision problem for real closed fields is complete for exponential space, and our logic contains the full language of real closed fields.

While we assumed that the algorithm takes as input the set of primitive propositions Φ_{h} and Φ_{o}, this does not really affect the complexity of the algorithm. More precisely, if we are given a formula φ in $\mathcal{L}^{fo\text{-}ev}$ over some set of hypotheses and observations, we can still decide whether φ is satisfiable, that is, whether there are sets Φ_{h} and Φ_{o} of primitive propositions containing all the primitive propositions in φ and an evidential world w that satisfies φ.

Theorem 5.2: *There is a procedure that runs in space exponential in $|\varphi|\ \|\varphi\|$ for deciding whether there exists sets of primitive propositions Φ_{h} and Φ_{o} such that $\varphi \in \mathcal{L}^{fo\text{-}ev}(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$ and φ is satisfiable in an evidential world.*

The main culprit for the exponential-space complexity is the theory of real closed fields, which we had to add to the logic to be able to even write down axiom **E4** of the axiomatization $\mathbf{AX}(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$.[5] However, if we are not interested in axiomatizations, but simply in verifying properties of probabilities and weights of evidence, we can consider the following propositional (quantifier-free) fragment of our logic. As before, we start with sets Φ_{h} and Φ_{o} of hypothesis and observation primitives, and form the sublanguage $\mathcal{L}_{\mathsf{h}}$ of hypothesis formulas. Basic terms have the form $\Pr^0(\rho)$, $\Pr(\rho)$, and $\mathrm{w}(ob, h)$, where ρ is an hypothesis formula, ob is an observation, and h is an hypothesis. A quantifier-free polynomial term has the form $a_1 t_1 + \cdots + a_n t_n$, where each a_i is an integer and each t_i is a product of basic terms. A quantifier-free polynomial inequality formula has the form $p \geq c$, where p is a quantifier-free polynomial term, and c is an integer. For instance, a quantifier-free polynomial inequality formula takes the form $\Pr^0(\rho) + 3\mathrm{w}(ob, h) + 5\Pr^0(\rho)\Pr(\rho') \geq 7$.

Let $\mathcal{L}^{ev}(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$ be the language obtained by starting out with the primitive propositions in Φ_{h} and Φ_{o} and quantifier-free polynomial inequality formulas, and closing off under conjunction and negation. Since quantifier-free polynomial inequality formulas are polynomial inequality formulas, $\mathcal{L}^{ev}(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$ is a sublanguage of $\mathcal{L}^{fo\text{-}ev}(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$. The logic $\mathcal{L}^{ev}(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$ is sufficiently expressive to express many properties of interest; for instance, it can certainly express the general connection between priors, posteriors, and evidence captured by axiom **E3**, as well as specific relationships between prior probability and posterior probability through the weight of evidence of a particular observation, as in Example 3.1. Reasoning about the propositional fragment of our logic $\mathcal{L}^{ev}(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$ is easier than the full language.[6]

5. Recall that axiom **E4** requires existential quantification. Thus, we can restrict to the sublanguage consisting of formulas with a single block of existential quantifiers in prefix position. The satisfiability problem for this sublanguage can be shown to be decidable in time exponential in the size of the formula (Renegar, 1992).

6. In a preliminary version of this paper (Halpern & Pucella, 2003), we examined the quantifier-free fragment of $\mathcal{L}^{fo\text{-}ev}(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$ that uses only linear inequality formulas, of the form $a_1 t_1 + \cdots + a_n t_n \geq c$, where each t_i is a basic term. We claimed that the problem of deciding, given Φ_{h} and Φ_{o}, whether a formula φ of this fragment is satisfiable in an evidential world is NP-complete. We further claimed that this result followed from a small-model theorem: if φ is satisfiable, then it is satisfiable in an evidential world over a small number of hypotheses and observations. While this small-model theorem is true, our argument that the satisfiability problem is in NP also implicitly assumed that the numbers associated with the probability measure and the evidence space in the evidential world were small. But this is not true

Theorem 5.3: *There is a procedure that runs in space polynomial in $|\varphi|\ \|\varphi\|$ for deciding, given Φ_{h} and Φ_{o}, whether a formula φ of $\mathcal{L}^{ev}(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$ is satisfiable in an evidential world.*

Theorem 5.3 relies on Canny's (1988) procedure for deciding the validity of quantifier-free formulas in the theory of real closed fields. As in the general case, the complexity is unaffected by whether or not the decision problem takes as input the sets Φ_{h} and Φ_{o} of primitive propositions.

Theorem 5.4: *There is a procedure that runs in space polynomial in $|\varphi|\ \|\varphi\|$ for deciding whether there exists sets of primitive propositions Φ_{h} and Φ_{o} such that $\varphi \in \mathcal{L}^{ev}(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$ and φ is satisfiable in an evidential world.*

6. Normalized Versus Unnormalized Likelihoods

The weight of evidence we used throughout this paper is a generalization of the log-likelihood ratio advocated by Good (1950, 1960). As we pointed out earlier, this measure of confirmation is essentially a normalized likelihood: the likelihood of an observation given a particular hypothesis is normalized by the sum of all the likelihoods of that observation, for all possible hypotheses. What would change if we were to take the (unnormalized) likelihoods μ_h themselves as weight of evidence? Some things would simplify. For example, **WF2** is a consequence of normalization, as is the corresponding axiom **E4**, which is the only axiom that requires quantification.

The main argument for normalizing likelihood is the same as that for normalizing probability measures. Just like probability, when using normalized likelihood, the weight of evidence is always between 0 and 1, and provides an absolute scale against which to judge all reports of evidence. The impact here is psychological—it permits one to use the same rules of thumb in all situations, since the numbers obtained are independent from the context of their use. Thus, for instance, a weight of evidence of 0.95 in one situation corresponds to the "same amount" of evidence as a weight of evidence of 0.95 in a different situation; any acceptable decision based on this weight of evidence in the first situation ought to be acceptable in the other situation as well. The importance of having such a uniform scale depends, of course, on the intended applications.

For the sake of completeness, we now describe the changes to our framework required to use unnormalized likelihoods as a weight of evidence. Define $w^u_{\mathcal{E}}(ob, h) = \mu_h(ob)$.

in general. Even though the formula φ involves only linear inequality formulas, every evidential world satisfies axiom **E3**. This constraint enables us to write formulas for which there exist no models where the probabilities and weights of evidence are rational. For example, consider the formula

$$\Pr^0(h_1) = \mathrm{w}(ob_1, h_1) \land \Pr^0(h_2) = 1 - \Pr^0(h_1) \land \Pr(h_1) = 1/2 \land \mathrm{w}(ob_1, h_2) = 1/4$$

Any evidential world satisfying the formula must satisfy

$$\Pr^0(h_1) = \mathrm{w}(ob_1, h_1) = -1/8(1 - \sqrt{17})$$

which is irrational. The exact complexity of this fragment remains open. We can use our techniques to show that it is in PSPACE, but we have no matching lower bound. (In particular, it may indeed be in NP.) We re-examine this fragment of the logic in Section 6, under a different interpretation of weights of evidence.

First, note that we can update a prior probability μ_0 via a set of likelihood functions μ_h using a form of Dempster's Rule of Combination. More precisely, we can define $\mu_0 \oplus w^u_{\mathcal{E}}(ob, \cdot)$ to be the probability measure defined by

$$(\mu_0 \oplus w^u_{\mathcal{E}}(ob, \cdot))(h) = \frac{\mu_0(h)\mu_h(ob)}{\sum_{h' \in \mathcal{H}} \mu_0(h')\mu_{h'}(ob)}.$$

The logic we introduced in Section 3 applies just as well to this new interpretation of weights of evidence. The syntax remains unchanged, the models remain evidential worlds, and the semantics of formulas simply take the new interpretation of weight of evidence into account. In particular, the assignment $[p]^{w,v}$ now uses the above definition of $w^u_{\mathcal{E}}$, and becomes

$$\begin{array}{rcl} [\Pr(\rho)]^{w,v} & = & (\mu \oplus w^u_{\mathcal{E}}(ob, \cdot))([\![\rho]\!]) \\ [\mathrm{w}(ob', h')]^{w,v} & = & w^u_{\mathcal{E}}(ob', h'). \end{array}$$

The axiomatization of this new logic is slightly different and somewhat simpler than the one in Section 3. In particular, **E1** and **E2**, which say that $\mathrm{w}(ob, h)$ acts as a probability measure for each fixed ob, are replaced by axioms that say that $\mathrm{w}(ob, h)$ acts as a probability measure for each fixed h:

E1′. $\mathrm{w}(ob, h) \geq 0$.

E2′. $\mathrm{w}(ob_1, h) + \cdots + \mathrm{w}(ob_{n_{\mathsf{o}}}, h) = 1$.

Axiom **E3** is unchanged, since $w^u_{\mathcal{E}}$ is updated in essentially the same way as $w_{\mathcal{E}}$. Axiom **E4** becomes unnecessary.

What about the complexity of the decision procedure? As in Section 5, the complexity of the decision problem for the full logic $\mathcal{L}^{fo\text{-}ev}(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$ remains dominated by the complexity of reasoning in real closed fields. Of course, now, we can express the full axiomatization for the unnormalized likelihood interpretation of weight of evidence in the $\mathcal{L}^{ev}(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$ fragment, which can be decided in polynomial space. A further advantage of the unnormalized likelihood interpretation of weight of evidence, however, is that it leads to a useful fragment of $\mathcal{L}^{ev}(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$ that is perhaps easier to decide.

Suppose that we are interested in reasoning exclusively about weights of evidence, with no prior or posterior probability. This is the kind of reasoning that actually underlies many computer science applications involving randomized algorithms (Halpern & Pucella, 2005b). As before, we start with sets Φ_{h} and Φ_{o} of hypothesis and observation primitives, and form the sublanguage $\mathcal{L}_{\mathsf{h}}$ of hypothesis formulas. A quantifier-free linear term has the form $a_1\mathrm{w}(ob^1, h^1) + \cdots + a_n\mathrm{w}(ob^n, h^n)$, where each a_i is an integer, each ob^i is an observation, and each h^i is an hypothesis. A quantifier-free linear inequality formula has the form $p \geq c$, where p is a quantifier-free linear term and c is an integer. For example, $\mathrm{w}(ob', h) + 3\mathrm{w}(ob, h) \geq 7$ is a quantifier-free linear inequality formula.

Let $\mathcal{L}^w(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$ be the language obtained by starting out with the primitive propositions in Φ_{h} and Φ_{o} and quantifier-free linear inequality formulas, and closing off under conjunction and negation. Since quantifier-free linear inequality formulas are polynomial inequality formulas, $\mathcal{L}^w(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$ is a sublanguage of $\mathcal{L}^{fo\text{-}ev}(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$. Reasoning about $\mathcal{L}^w(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$ is easier than the full language, and possibly easier than the $\mathcal{L}^{ev}(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$ fragment.

Theorem 6.1: *The problem of deciding, given Φ_h and Φ_o, whether a formula φ of $\mathcal{L}^w(\Phi_h, \Phi_o)$ is satisfiable in an evidential world is NP-complete.*

As in the general case, the complexity is unaffected by whether or not the decision problem takes as input the sets Φ_h and Φ_o of primitive propositions.

Theorem 6.2: *The problem of deciding, for a formula φ, whether there exists sets of primitive propositions Φ_h and Φ_o such that $\varphi \in \mathcal{L}^w(\Phi_h, \Phi_o)$ and φ is satisfiable in an evidential world is NP-complete.*

7. Evidence in Dynamic Systems

The evidential worlds we have considered until now are essentially static, in that they model only the situation where a single observation is made. Considering such static worlds lets us focus on the relationship between the prior and posterior probabilities on hypotheses and the weight of evidence of a single observation. In a related paper (Halpern & Pucella, 2005b), we consider evidence in the context of randomized algorithms; we use evidence to characterize the information provided by, for example, a randomized algorithm for primality when it says that a number is prime. The framework in that work is dynamic; sequences of observations are made over time. In this section, we extend our logic to reason about the evidence of sequences of observations, using the approach to combining evidence described in Section 2.

There are subtleties involved in trying to find an appropriate logic for reasoning about situations like that in Example 2.5. The most important one is the relationship between observations and time. By way of illustration, consider the following example. Bob is expecting an email from Alice stating where a rendezvous is to take place. Calm under pressure, Bob is reading while he waits. We assume that Bob is not concerned with the time. For the purposes of this example, one of three things can occur at any given point in time:

(1) Bob does not check if he has received email;

(2) Bob checks if he has received email, and notices he has not received an email from Alice;

(3) Bob checks if he has received email, and notices he has received an email from Alice.

How is his view of the world affected by these events? In (1), it should be clear that, all things being equal, Bob's view of the world does not change: no observation is made. Contrast this with (2) and (3). In (2), Bob does make an observation, namely that he has not yet received Alice's email. The fact that he checks indicates that he wants to observe a result. In (3), he also makes an observation, namely that he received an email from Alice. In both of these cases, the check yields an observation, that he can use to update his view of the world. In case (2), he essentially observed that nothing happened, but we emphasize again that this is an observation, to be distinguished from the case where Bob does not even check whether email has arrived, and should be explicit in the set $\mathcal{O}$ in the evidence space.

This discussion motivates the models that we use in this section. We characterize an agent's state by the observations that she has made, including possibly the "nothing happened" observation. Although we do not explicitly model time, it is easy to incorporate time in our framework, since the agent can observe times or clock ticks. The models in this section are admittedly simple, but they already highlight the issues involved in reasoning about evidence in dynamic systems. As long as agents do not forget observations, there is no loss of generality in associating an agent's state with a sequence of observations. We do, however, make the simplifying assumption that the same evidence space is used for all the observations in a sequence. In other words, we assume that the evidence space is fixed for the evolution of the system. In many situations of interest, the external world changes. The possible observations may depend on the state of the world, as may the likelihood functions. There are no intrinsic difficulties in extending the model to handle state changes, but the additional details would only obscure the presentation.

In some ways, considering a dynamic setting simplifies things. Rather than talking about the prior and posterior probability using different operators, we need only a single probability operator that represents the probability of an hypothesis at the current time. To express the analogue of axiom **E3** in this logic, we need to be able to talk about the probability at the next time step. This can be done by adding the "next-time" operator $\bigcirc$ to the logic, where $\bigcirc\varphi$ holds at the current time if φ holds at the next time step.[7] We further extend the logic to talk about the weight of evidence of a sequence of observations.

We define the logic $\mathcal{L}^{fo\text{-}ev}_{dyn}$ as follows. As in Section 3, we start with a set of primitive propositions Φ_{h} and Φ_{o}, respectively representing the hypotheses and the observations. Again, let $\mathcal{L}_{\mathsf{h}}(\Phi_{\mathsf{h}})$ be the propositional sublanguage of hypotheses formulas obtained by taking primitive propositions in Φ_{h} and closing off under negation and conjunction; we use ρ to range over formulas of that sublanguage.

A basic term now has the form $\Pr(\rho)$ or $\mathrm{w}(\underline{ob}, h)$, where ρ is an hypothesis formula, $\underline{ob} = \langle ob^1, \ldots, ob^k\rangle$ is a nonempty sequence of observations, and h is an hypothesis. If $\underline{ob} = \langle ob^1\rangle$, we write $\mathrm{w}(ob_1, h)$ rather than $\mathrm{w}(\langle ob^1\rangle, h)$. As before, a polynomial term has the form $t_1 + \cdots + t_n$, where each term t_i is a product of integers, basic terms, and variables (which intuitively range over the reals). A polynomial inequality formula has the form $p \geq c$, where p is a polynomial term and c is an integer. Let $\mathcal{L}^{fo\text{-}ev}_{dyn}(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$ be the language obtained by starting out with the primitive propositions in Φ_{h} and Φ_{o} and polynomial inequality formulas, and closing off under conjunction, negation, first-order quantification, and application of the $\bigcirc$ operator. We use the same abbreviations as in Section 3.

The semantics of this logic now involves models that have dynamic behavior. Rather than just considering individual worlds, we now consider sequences of worlds, which we call *runs*, representing the evolution of the system over time. A model is now an infinite run, where a run describes a possible dynamic evolution of the system. As before, a run records the observations being made and the hypothesis that is true for the run, as well as a probability distribution describing the prior probability of the hypothesis at the initial state of the run, and an evidence space $\mathcal{E}^*$ over Φ_{h} and Φ^*_{o} to interpret w. We define an *evidential run* r to be a map from the natural numbers (representing time) to histories of

7. Following the discussion above, time steps are associated with new observations. Thus, $\bigcirc\varphi$ means that φ is true at the next time step, that is, after the next observation. This simplifies the presentation of the logic.

the system up to that time. A history at time m records the relevant information about the run—the hypothesis that is true, the prior probability on the hypotheses, and the evidence space $\mathcal{E}^*$—and the observations that have been made up to time m. Hence, a history has the form $\langle (h, \mu, \mathcal{E}^*), ob^1, \ldots, ob^k \rangle$. We assume that $r(0) = \langle (h, \mu, \mathcal{E}^*) \rangle$ for some h, μ, and $\mathcal{E}^*$, while $r(m) = \langle (h, \mu, \mathcal{E}^*), ob^1, \ldots, ob^m \rangle$ for $m > 0$. We define a point of the run to be a pair (r, m) consisting of a run r and time m.

We associate with each propositional formula ρ in $\mathcal{L}_{\mathsf{h}}(\Phi_{\mathsf{h}})$ a set $[\![\rho]\!]$ of hypotheses, just as we did in Section 3.

In order to ascribe a semantics to first-order formulas that may contain variables, we need a valuation v that assigns a real number to every variable. Given a valuation v, an evidential run r, and a point (r, m), where $r(m) = \langle (h, \mu, \mathcal{E}^*), ob^1, \ldots, ob^m \rangle$, we can assign to a polynomial term p a real number $[p]^{r,m,v}$ using essentially the same approach as in Section 3:

$$\begin{array}{rcl}
[x]^{r,m,v} & = & v(x) \\
[a]^{r,m,v} & = & a \\
[\Pr(\rho)]^{r,m,v} & = & (\mu \oplus w_{\mathcal{E}^*}(\langle ob^1, \ldots, ob^m \rangle, \cdot)))([\![\rho]\!]) \\
 & & \quad \text{where } r(m) = \langle (h, \mu, \mathcal{E}^*), ob^1, \ldots, ob^m \rangle \\
[\mathrm{w}(\underline{ob}, h')]^{r,m,v} & = & w_{\mathcal{E}^*}(\underline{ob}, h') \\
 & & \quad \text{where } r(m) = \langle (h, \mu, \mathcal{E}^*), ob^1, \ldots, ob^m \rangle \\
[t_1 t_2]^{r,m,v} & = & [t_1]^{r,m,v} \times [t_2]^{r,m,v} \\
[p_1 + p_2]^{r,m,,v} & = & [p_1]^{r,m,v} + [p_2]^{r,m,v}.
\end{array}$$

We define what it means for a formula φ to be true (or satisfied) at a point (r, m) of an evidential run r under valuation v, written $(r, m, v) \models \varphi$, using essentially the same approach as in Section 3:

$(r, m, v) \models h$ if $r(m) = \langle (h, \mu, \mathcal{E}^*), \ldots \rangle$

$(r, m, v) \models ob$ if $r(m) = \langle (h, \mu, \mathcal{E}^*), \ldots, ob \rangle$

$(r, m, v) \models \neg\varphi$ if $(r, m, v) \not\models \varphi$

$(r, m, v) \models \varphi \wedge \psi$ if $(r, m, v) \models \varphi$ and $(r, m, v) \models \psi$

$(r, m, v) \models p \geq c$ if $[p]^{r,m,v} \geq c$

$(r, m, v) \models \bigcirc\varphi$ if $(r, m+1, v) \models \varphi$

$(r, m, v) \models \forall x \varphi$ if $(r, m, v') \models \varphi$ for all valuations v' that agree with v on all variables but x.

If $(r, m, v) \models \varphi$ is true for all v, we simply write $(r, m) \models \varphi$. If $(r, m) \models \varphi$ for all points (r, m) of r, then we write $r \models \varphi$ and say that φ is valid in r. Finally, if $r \models \varphi$ for all evidential runs r, we write $\models \varphi$ and say that φ is valid.

It is straightforward to axiomatize this new logic. The axiomatization shows that we can capture the combination of evidence directly in the logic, a pleasant property. Most of the axioms from Section 3 carry over immediately. Let the axiomatization $\mathbf{AX}_{dyn}(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$ consists of the following axioms and inference rules: first-order reasoning (**Taut**, **MP**), reasoning about polynomial inequalities (**RCF**), reasoning about hypotheses and observations (**H1**,**H2**,**O1**,**O2**), reasoning about probabilities (**Po1**–**4** only, since we do not have $\Pr^0$ in the language), and reasoning about weights of evidence (**E1**, **E2**, **E4**), as well as new axioms we now present.

Basically, the only axiom that needs replacing is **E3**, which links prior and posterior probabilities, since this now needs to be expressed using the $\bigcirc$ operator. Moreover, we need an axiom to relate the weight of evidence of a sequence of observation to the weight of evidence of the individual observations, as given by Equation (3).

E5. $ob \Rightarrow \forall x(\bigcirc(\Pr(h) = x) \Rightarrow$
$$\Pr(h)\mathrm{w}(ob, h) = x\Pr(h_1)\mathrm{w}(ob, h_1) + \cdots + x\Pr(h_{n_{\mathsf{h}}})\mathrm{w}(ob, h_{n_{\mathsf{h}}})).$$

E6. $\mathrm{w}(ob^1, h)\cdots\mathrm{w}(ob^k, h) = \mathrm{w}(\langle ob^1, \ldots, ob^k\rangle, h)\mathrm{w}(ob^1, h_1)\cdots\mathrm{w}(ob^k, h_1) + \cdots +$
$$\mathrm{w}(\langle ob^1, \ldots, ob^k\rangle, h)\mathrm{w}(ob^1, h_{n_{\mathsf{h}}})\cdots\mathrm{w}(ob^k, h_{n_{\mathsf{h}}}).$$

To get a complete axiomatization, we also need axioms and inference rules that capture the properties of the temporal operator $\bigcirc$.

T1. $\bigcirc\varphi \wedge \bigcirc(\varphi \Rightarrow \psi) \Rightarrow \bigcirc\psi$.

T2. $\bigcirc\neg\varphi \Leftrightarrow \neg\bigcirc\varphi$.

T3. From φ infer $\bigcirc\varphi$.

Finally, we need axioms to say that the truth of hypotheses as well as the value of polynomial terms not containing occurrences of Pr is time-independent:

T4. $\bigcirc\rho \Leftrightarrow \rho$.

T5. $\bigcirc(p \geq c) \Leftrightarrow p \geq c$ if p does not contain an occurrence of Pr.

T6. $\bigcirc(\forall x\varphi) \Leftrightarrow \forall x(\bigcirc\varphi)$.

Theorem 7.1: *$\boldsymbol{AX}_{dyn}(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$ is a sound and complete axiomatization for $\mathcal{L}^{fo\text{-}ev}_{dyn}(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$ with respect to evidential runs.*

8. Conclusion

In the literature, reasoning about the effect of observations is typically done in a context where we have a prior probability on a set of hypotheses which we can condition on the observations made to obtain a new probability on the hypotheses that reflects the effect of the observations. In this paper, we have presented a logic of evidence that lets us reason about the weight of evidence of observations, independently of any prior probability on the hypotheses. The logic is expressive enough to capture in a logical form the relationship between a prior probability on hypotheses, the weight of evidence of observations, and the result posterior probability on hypotheses. But we can also capture reasoning that does not involve prior probabilities.

While the logic is essentially propositional, obtaining a sound and complete axiomatization seems to require quantification over the reals. This adds to the complexity of the logic—the decision problem for the full logic is in exponential space. However, an interesting and potentially useful fragment, the propositional fragment, is decidable in polynomial space.

Acknowledgments. A preliminary version of this paper appeared in the *Proceedings of the Nineteenth Conference on Uncertainty in Artificial Intelligence*, pp. 297–304, 2003. This work was mainly done while the second author was at Cornell University. We thank Dexter Kozen and Nimrod Megiddo for useful discussions. Special thanks to Manfred Jaeger for his careful reading of the paper and subsequent comments. Manfred found the bug in our proof that the satisfiability problem for the quantifier-free fragment of $\mathcal{L}^{fo\text{-}ev}(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$ that uses only linear inequality formulas is NP-complete. His comments also led us to discuss the issue of normalization. We also thank the reviewers, whose comments greatly improved the paper. This work was supported in part by NSF under grants CTC-0208535, ITR-0325453, and IIS-0534064, by ONR under grant N00014-01-10-511, by the DoD Multidisciplinary University Research Initiative (MURI) program administered by the ONR under grants N00014-01-1-0795 and N00014-04-1-0725, and by AFOSR under grant F49620-02-1-0101.

Appendix A. Proofs

Proposition 2.1: *For all ob, we have $w_{\mathcal{E}}(ob, h_i) \geq w_{\mathcal{E}}(ob, h_{3-i})$ if and only if $l(ob, h_i) \geq l(ob, h_{3-i})$, for $i = 1, 2$, and for all h, ob, and ob', we have $w_{\mathcal{E}}(ob, h) \geq w_{\mathcal{E}}(ob', h)$ if and only if $l(ob, h) \geq l(ob', h)$.*

Proof. Let ob be an arbitrary observation. The result follows from the following argument:

$$\begin{aligned}
&w_{\mathcal{E}}(ob, h_i) \geq w_{\mathcal{E}}(ob, h_{3-i})\\
&\quad \text{iff } \mu_{h_i}(ob)/(\mu_{h_i}(ob) + \mu_{h_{3-i}}(ob)) \geq \mu_{h_{3-i}}(ob)/(\mu_{h_i}(ob) + \mu_{h_{3-i}}(ob))\\
&\quad \text{iff } \mu_{h_i}(ob)\mu_{h_i}(ob) \geq \mu_{h_{3-i}}(ob)\mu_{h_{3-i}}(ob)\\
&\quad \text{iff } \mu_{h_i}(ob)/\mu_{h_{3-i}}(ob) \geq \mu_{h_{3-i}}(ob)/\mu_{h_i}(ob)\\
&\quad \text{iff } l(ob, h_i) \geq l(ob, h_{3-i}).
\end{aligned}$$

A similar argument establishes the result for hypotheses. □

Theorem 2.4: *Let $\mathcal{H} = \{h_1, \ldots, h_m\}$ and $\mathcal{O} = \{ob_1, \ldots, ob_n\}$, and let f be a real-valued function with domain $\mathcal{O} \times \mathcal{H}$ such that $f(ob, h) \in [0, 1]$. Then there exists an evidence space $\mathcal{E} = (\mathcal{H}, \mathcal{O}, \mu_{h_1}, \ldots, \mu_{h_m})$ such that $f = w_{\mathcal{E}}$ if and only if f satisfies the following properties:*

WF1. *For every $ob \in \mathcal{O}$, $f(ob, \cdot)$ is a probability measure on $\mathcal{H}$.*

WF2. *There exists $x_1, \ldots, x_n > 0$ such that, for all $h \in \mathcal{H}$, $\sum_{i=1}^{n} f(ob_i, h)x_i = 1$.*

Proof. ($\Rightarrow$) Assume that $f = w_{\mathcal{E}}$ for some evidence space $\mathcal{E} = (\mathcal{H}, \mathcal{O}, \mu_{h_1}, \ldots, \mu_{h_m})$. It is routine to verify **WF1**, that for a fixed $ob \in \mathcal{O}$, $w_{\mathcal{E}}(ob, \cdot)$ is a probability measure on $\mathcal{H}$. To verify **WF2**, note that we can simply take $x_i = \sum_{h' \in \mathcal{H}} \mu_{h'}(ob_i)$.

($\Leftarrow$) Let f be a function from $\mathcal{O} \times \mathcal{H}$ to $[0, 1]$ that satisfies **WF1** and **WF2**. Let $x_1^*, \ldots, x_{n_{\mathsf{h}}}^*$ be the positive reals guaranteed by **WF2**. It is straightforward to verify that taking $\mu_h(ob_i) = f(ob_i, h)/x_i^*$ for each $h \in \mathcal{H}$ yields an evidence space $\mathcal{E}$ such that $f = w_{\mathcal{E}}$. □

The following lemmas are useful to prove the completeness of the axiomatizations in this paper. These results depend on the soundness of the axiomatization $\mathbf{AX}(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$.

Lemma A.1: *$\boldsymbol{AX}(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$ is a sound axiomatization for the logic $\mathcal{L}^{fo\text{-}ev}(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$ with respect to evidential worlds.*

Proof. It is easy to see that each axiom is valid in evidential worlds. □

Lemma A.2: *For all hypothesis formulas ρ, $\rho \Leftrightarrow h_1 \vee \cdots \vee h_k$ is provable in $\boldsymbol{AX}(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$, when $[\![\rho]\!] = \{h_1, \ldots, h_k\}$.*

Proof. Using **Taut**, we can show that ρ is provably equivalent to a formula ρ' in disjunctive normal form. Moreover, by axiom **H2**, we can assume without loss of generality that each of the disjuncts in ρ' consists of a single hypothesis. Thus, ρ is $h_1 \vee \cdots \vee h_k$. An easy induction on structure shows that for an hypothesis formula ρ and evidential world w, we have that $w \models \rho$ iff $w \models h$ for some $h \in [\![\rho]\!]$. Moreover, it follows immediately from the soundness of the axiomatization (Lemma A.1) that $\rho \Leftrightarrow h_1 \vee \ldots \vee h_k$ is provable iff for all evidential worlds w, $w \models \rho$ iff $w \models h_i$ for some $i \in \{1, \ldots, k\}$. Thus, $\rho \Leftrightarrow h_1 \vee \ldots \vee h_k$ is provable iff $[\![\rho]\!] = \{h_1, \ldots, h_k\}$. □

An easy consequence of Lemma A.2 is that ρ_1 is provably equivalent to ρ_2 if and only if $[\![\rho_1]\!] = [\![\rho_2]\!]$.

Lemma A.3: *Let ρ be an hypothesis formula. The formulas*

$$\Pr(\rho) = \sum_{h \in [\![\rho]\!]} \Pr(h) \text{ and}$$

$$\Pr^0(\rho) = \sum_{h \in [\![\rho]\!]} \Pr^0(h)$$

are provable in $\boldsymbol{AX}(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$.

Proof. Let $\Phi_{\mathsf{h}} = \{h_1, \ldots, h_{n_{\mathsf{h}}}\}$ and $\Phi_{\mathsf{o}} = \{ob_1, \ldots, ob_{n_{\mathsf{o}}}\}$. We prove the result for Pr. We proceed by induction on the size of $[\![\rho]\!]$. For the base case, assume that $|[\![\rho]\!]| = 0$. By Lemma A.2, this implies that ρ is provably equivalent to *false*. By **Po4**, $\Pr(\rho) = \Pr(\mathit{false})$, and it is easy to check that $\Pr(\mathit{false}) = 0$ is provable using **Po1**, **Po3**, and **Po4**, thus $\Pr(\rho) = 0$, as required. If $|[\![\rho]\!]| = n + 1 > 0$, then $[\![\rho]\!] = \{h_{i_1}, \ldots, h_{i_{n+1}}\}$, and by Lemma A.2, ρ is provably equivalent to $h_{i_1} \vee \cdots \vee h_{i_{n+1}}$. By **Po4**, $\Pr(\rho) = \Pr(\rho \wedge h_{i_{n+1}}) + \Pr(\rho \wedge \neg h_{i_{n+1}})$. It is easy to check that $\rho \wedge h_{i_{n+1}}$ is provably equivalent to $h_{i_{n+1}}$ (using **H2**), and similarly $\rho \wedge \neg h_{i_{n+1}}$ is provably equivalent to $h_{i_1} \vee \cdots \vee h_{i_n}$. Thus, $\Pr(\rho) = \Pr(h_{i_{n+1}}) + \Pr(h_{i_1} \vee \cdots \vee h_{i_n})$ is provable. Since $|[\![h_{i_1} \vee \cdots \vee h_{i_n}]\!]| = n$, by the induction hypothesis, $\Pr(h_{i_1} \vee \cdots \vee h_{i_n}) = \sum_{h \in \{h_{i_1}, \ldots, h_{i_n}\}} \Pr(h) = \sum_{h \in [\![\rho]\!] - \{h_{i_{n+1}}\}} \Pr(h)$. Thus, $\Pr(\rho) = \Pr(h_{i_{n+1}}) + \sum_{h \in [\![\rho]\!] - \{h_{i_{n+1}}\}} \Pr(h)$, that is, $\Pr(\rho) = \sum_{h \in [\![\rho]\!]} \Pr(h)$, as required.

The same argument applies *mutatis mutandis* for $\Pr^0$, using axioms **Pr1**–**4** instead of **Po1**–**4**. □

Theorem 4.1: $\boldsymbol{AX}(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$ *is a sound and complete axiomatization for the logic with respect to evidential worlds.*

Proof. Soundness was established in Lemma A.1. To prove completeness, recall the following definitions. A formula φ is *consistent* with the axiom system $\mathbf{AX}(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$ if $\neg\varphi$ is not provable from $\mathbf{AX}(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$. To prove completeness, it is sufficient to show that if φ is consistent, then it is satisfiable, that is, there exists an evidential world w and valuation v such that $(w, v) \models \varphi$.

As in the body of the paper, let $\Phi_{\mathsf{h}} = \{h_1, \ldots, h_{n_{\mathsf{h}}}\}$ and $\Phi_{\mathsf{o}} = \{ob_1, \ldots, ob_{n_{\mathsf{o}}}\}$. Let φ be a consistent formula. By way of contradiction, assume that φ is unsatisfiable. We reduce the formula φ to an equivalent formula in the language of real closed fields. Let $u_1, \ldots, u_{n_{\mathsf{h}}}$, $v_1, \ldots, v_{n_{\mathsf{o}}}$, $x_1, \ldots, x_{n_{\mathsf{h}}}$, $y_1, \ldots, y_{n_{\mathsf{o}}}$, and $z_1^1, \ldots, z_{n_{\mathsf{h}}}^1, \ldots, z_1^{n_{\mathsf{o}}}, \ldots, z_{n_{\mathsf{h}}}^{n_{\mathsf{o}}}$ be new variables, where, intuitively,

- u_i gets value 1 if hypothesis h_i holds, 0 otherwise;
- v_i gets value 1 if observation ob_i holds, 0 otherwise;
- x_i represents $\Pr^0(h_i)$;
- y_i represents $\Pr(h_i)$;
- $z_{i,j}$ represents $\mathrm{w}(ob_i, h_j)$.

Let $\mathbf{v}$ represent that list of new variables. Consider the following formulas. Let φ_h be the formula saying that exactly one hypothesis holds:

$$(u_1 = 0 \vee u_1 = 1) \wedge \cdots \wedge (u_{n_{\mathsf{h}}} = 0 \vee u_{n_{\mathsf{h}}} = 1) \wedge u_1 + \cdots + u_{n_{\mathsf{h}}} = 1.$$

Similarly, let φ_o be the formula saying that exactly one observation holds:

$$(v_1 = 0 \vee v_1 = 1) \wedge \cdots \wedge (v_{n_{\mathsf{o}}} = 0 \vee v_{n_{\mathsf{h}}} = 1) \wedge v_1 + \cdots + v_{n_{\mathsf{h}}} = 1.$$

Let φ_{pr} be the formula that expresses that $\Pr^0$ is a probability measure:

$$\varphi_{pr} = x_1 \geq 0 \wedge \cdots \wedge x_{n_{\mathsf{h}}} \geq 0 \wedge x_1 + \cdots + x_{n_{\mathsf{h}}} = 1.$$

Similarly, let φ_{po} be the formula that expresses that Pr is a probability measure:

$$\varphi_{po} = y_1 \geq 0 \wedge \cdots \wedge y_{n_{\mathsf{h}}} \geq 0 \wedge y_1 + \cdots + y_{n_{\mathsf{h}}} = 1.$$

Finally, we need formulas saying that w is a weight of evidence function. The formula $\varphi_{w,p}$ simply says that w satisfies **WF1**, that is, it acts as a probability measure for a fixed observation:

$$\begin{gathered} z_{1,1} \geq 0 \wedge \cdots \wedge z_{1,n_{\mathsf{h}}} \geq 0 \wedge z_{n_{\mathsf{o}},1} \geq 0 \wedge \cdots \wedge z_{n_{\mathsf{o}},n_{\mathsf{h}}} \geq 0 \wedge \\ z_{1,1} + \cdots + z_{1,n_{\mathsf{h}}} = 1 \wedge \cdots \wedge z_{n_{\mathsf{o}},1} + \cdots + z_{n_{\mathsf{o}},n_{\mathsf{h}}} = 1. \end{gathered}$$

The formula $\varphi_{w,f}$ says that w satisfies **WF2**:

$$\begin{gathered} \exists w_1, \ldots, w_{n_{\mathsf{o}}}(w_1 > 0 \wedge \cdots \wedge w_{n_{\mathsf{o}}} > 0 \wedge z_{1,1}w_1 + \cdots + z_{n_{\mathsf{o}},1}w_{n_{\mathsf{o}}} = 1 \wedge \\ \cdots \wedge z_{1,n_{\mathsf{h}}}w_1 + \cdots + z_{n_{\mathsf{o}},n_{\mathsf{h}}}w_{n_{\mathsf{o}}} = 1) \end{gathered}$$

where $w_1, \dots, w_{n_o}$ are new variables.

Finally, the formula $\varphi_{w,up}$ captures the fact that weights of evidence can be viewed as updating a prior probability into a posterior probability, via Dempster's Rule of Combination:

$$\begin{aligned}
&(v_1 = 1 \Rightarrow (x_1 z_{1,1} = y_1 x_1 z_{1,1} + \cdots + y_1 x_{n_h} z_{1,n_h} \wedge \\
&\qquad\qquad \cdots \wedge x_{n_h} z_{1,n_h} = y_{n_h} x_1 z_{1,1} + \cdots + y_{n_h} x_{n_h} z_{1,n_h})) \wedge \\
&\cdots \wedge \\
&(v_{n_o} = 1 \Rightarrow (x_1 z_{n_o,1} = y_1 x_1 z_{n_o,1} + \cdots + y_1 x_{n_h} z_{n_o,n_h} \wedge \\
&\qquad\qquad \cdots \wedge x_{n_h} z_{n_o,n_h} = y_{n_h} x_1 z_{n_o,1} + \dots y_{n_h} x_{n_h} z_{n_o,nh})).
\end{aligned}$$

Let $\hat{\varphi}$ be the formula in the language of real closed fields obtained from φ by replacing each occurrence of the primitive proposition h_i by $u_i = 1$, each occurrence of ob_i by $v_i = 1$, each occurrence of $\Pr^0(\rho)$ by $\sum_{h_i \in [\![\rho]\!]} x_i$, each occurrence of $\Pr(\rho)$ by $\sum_{h_i \in [\![\rho]\!]} y_i$, each occurrence of $\mathrm{w}(ob_i, h_j)$ by $z_{i,j}$, and each occurrence of an integer coefficient k by $1 + \cdots + 1$ (k times). Finally, let φ' be the formula $\exists \mathbf{v}(\varphi_h \wedge \varphi_o \wedge \varphi_{pr} \wedge \varphi_{po} \wedge \varphi_{w,p} \wedge \varphi_{w,f} \wedge \varphi_{w,up} \wedge \hat{\varphi})$.

It is easy to see that if φ is unsatisfiable over evidential worlds, then φ' is false when interpreted over the real numbers. Therefore, $\neg\varphi'$ must be a formula valid in real closed fields, and hence an instance of **RCF**. Thus, $\neg\varphi'$ is provable. It is straightforward to show, using Lemma A.3, that $\neg\varphi$ itself is provable, contradicting the fact that φ is consistent. Thus, φ must be satisfiable, establishing completeness. □

As we mentioned at the beginning of Section 5, $\mathcal{L}^{fo\text{-}ev}$ is not monotone with respect to validity: axiom **H1** depends on the set of hypotheses and observations, and will in general no longer be valid if the set is changed. The same is true for **O1**, **E3**, and **E4**. We do, however, have a form of monotonicity with respect to satisfiability, as the following lemma shows.

Lemma A.4: *Given Φ_h and Φ_o, let φ be a formula of $\mathcal{L}^{fo\text{-}ev}(\Phi_h, \Phi_o)$, and let $\mathcal{H} \subseteq \Phi_h$ and $\mathcal{O} \subseteq \Phi_o$ be the hypotheses and observations that occur in φ. If φ is satisfiable in an evidential world over Φ_h and Φ_o, then φ is satisfiable in an evidential world over Φ'_h and Φ'_o, where $|\Phi'_h| = |\mathcal{H}| + 1$ and $|\Phi'_o| = |\mathcal{O}| + 1$.*

Proof. We do this in two steps, to clarify the presentation. First, we show that we can add a single hypothesis and observation to Φ_h and Φ_o and preserve satisfiability of φ. This means that the second step below can assume that $\Phi_h \neq \mathcal{H}$ and $\Phi_o \neq \mathcal{O}$. Assume that φ is satisfied in an evidential world $w = (h, ob, \mu, \mathcal{E})$ over Φ_h and Φ_o, so that there exists v such that $(w, v) \models \varphi$. Let $\Phi'_h = \Phi_h \cup \{h^*\}$, where h^* is a new hypothesis not in Φ_h, and let $\Phi'_o = \Phi_o \cup \{ob^*\}$, where ob^* is a new observation not in Φ_o. Define the evidential world $w' = (h, ob, \mu', \mathcal{E}')$ over Φ'_h and Φ'_o, where $\mathcal{E}'$ and μ' are defined as follows. Define the probability measure μ' by taking:

$$\mu'(h) = \begin{cases} \mu(h) & \text{if } h \in \Phi_h \\ 0 & \text{if } h = h^*. \end{cases}$$

Similarly, define the evidence space $\mathcal{E}' = (\Phi'_{\mathsf{h}}, \Phi'_{\mathsf{o}}, \boldsymbol{\mu}')$ derived from $\mathcal{E} = (\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}}, \boldsymbol{\mu})$ by taking:

$$\mu'_h(ob) = \begin{cases} \mu_h(ob) & \text{if } h \in \Phi_{\mathsf{h}} \text{ and } ob \in \Phi_{\mathsf{o}} \\ 0 & \text{if } h \in \Phi_{\mathsf{h}} \text{ and } ob = ob^* \\ 0 & \text{if } h = h^* \text{ and } ob \in \Phi_{\mathsf{o}} \\ 1 & \text{if } h = h^* \text{ and } ob \in ob^*. \end{cases}$$

Thus, μ'_h extends the existing μ_h by assigning a probability of 0 to the new observation ob^*; in contrast, the new probability μ'_{h^*} assigns probability 1 to the new observation ob^*. We can check that $(w', v) \models \varphi$.

The second step is to "collapse" all the hypotheses and observations that do not appear in φ into one of the hypotheses that do not appear in $\mathcal{H}$ and $\mathcal{O}$, which by the previous step are guaranteed to exist. By the previous step, we can assume that $\Phi_{\mathsf{h}} \neq \mathcal{H}$ and $\Phi_{\mathsf{o}} \neq \mathcal{O}$. Assume φ is satisfiable in an evidential world $w = (h, ob, \mu, \mathcal{E})$ over Φ_{h} and Φ_{o}, that is, there exists v such that $(w, v) \models \varphi$. Pick an hypothesis and an observation from Φ_{h} and Φ_{o} as follows, depending on the hypothesis h and observation ob in w. Let $h^\dagger$ be h if $h \notin \mathcal{H}$, otherwise, let $h^\dagger$ be an arbitrary element of $\Phi_{\mathsf{h}} - \mathcal{H}$; let $\Phi'_{\mathsf{h}} = \mathcal{H} \cup \{h^\dagger\}$. Similarly, let $ob^\dagger$ be ob if $ob \notin \mathcal{O}$, otherwise, let $ob^\dagger$ be an arbitrary element of $\Phi_{\mathsf{o}} - \mathcal{O}$; let $\Phi'_{\mathsf{o}} = \mathcal{O} \cup \{ob^\dagger\}$. Let $w' = (h, ob, \mu', \mathcal{E}')$ be an evidential world over Φ'_{h} and Φ'_{o} obtained from w as follows. Define the probability measure μ' by taking:

$$\mu'(h) = \begin{cases} \mu(h) & \text{if } h \in \mathcal{H} \\ \sum_{h' \in \Phi_{\mathsf{h}} - \mathcal{H}} \mu(h') & \text{if } h = h^\dagger. \end{cases}$$

Define $\mathcal{E}' = (\Phi'_{\mathsf{h}}, \Phi'_{\mathsf{o}}, \boldsymbol{\mu}')$ derived from $\mathcal{E} = (\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}}, \boldsymbol{\mu})$ by taking:

$$\mu'_h(ob) = \begin{cases} \mu_h(ob) & \text{if } h \in \mathcal{H} \text{ and } ob \in \mathcal{O} \\ \sum_{ob' \in \Phi_{\mathsf{o}} - \mathcal{O}} \mu_h(ob') & \text{if } h \in \mathcal{H} \text{ and } ob = ob^\dagger \\ \sum_{h' \in \Phi_{\mathsf{h}} - \mathcal{H}} \mu_{h'}(ob) & \text{if } h = h^\dagger \text{ and } ob \in \mathcal{O} \\ \sum_{h' \in \Phi_{\mathsf{h}} - \mathcal{H}} \sum_{ob' \in \Phi_{\mathsf{o}} - \mathcal{O}} \mu_{h'}(ob') & \text{if } h = h^\dagger \text{ and } ob = ob^\dagger. \end{cases}$$

We can check by induction that $(w', v) \models \varphi$. □

Theorem 5.1: *There is a procedure that runs in space exponential in $|\varphi|\ \|\varphi\|$ for deciding, given Φ_{h} and Φ_{o}, whether a formula φ of $\mathcal{L}^{fo\text{-}ev}(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$ is satisfiable in an evidential world.*

Proof. Let φ be a formula of $\mathcal{L}^{fo\text{-}ev}(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$. By Lemma A.4, φ is satisfiable if we can construct a probability measure μ on $\Phi'_{\mathsf{h}} = \mathcal{H} \cup \{h^*\}$ (where $\mathcal{H}$ is the set of hypotheses appearing in φ, and $h^* \notin \mathcal{H}$) and probability measures $\mu_{h_1}, \ldots, \mu_{h_m}$ on $\Phi'_{\mathsf{o}} = \mathcal{O} \cup \{ob^*\}$ (where $\mathcal{O}$ is the set of observations appearing in φ and $ob^* \notin \mathcal{O}$) such that $\mathcal{E} = (\Phi'_{\mathsf{h}}, \Phi'_{\mathsf{o}}, \boldsymbol{\mu})$, $w = (h, ob, \mu, \mathcal{E})$ with $(w, v) \models \varphi$ for some h, ob, and v.

The aim now is to derive a formula φ' in the language of real closed fields that asserts the existence of these probability measures. More precisely, we can adapt the construction of the formula φ' from φ in the proof of Theorem 4.1. The one change we need to make is ensure that φ' is polynomial in the size of φ, which the construction in the proof of

Theorem 4.1 does not guarantee. The culprit is the fact that we encode integer constants k as $1+\cdots+1$. It is straightforward to modify the construction so that we use a more efficient representation of integer constants, namely, a binary representation. For example, we can write 42 as $2(1+2^2(1+2^2))$, which can be expressed in the language of real closed fields as $(1+1)(1+(1+1)(1+1)(1+(1+1)(1+1)))$. We can check that if k is a coefficient of length k (when written in binary), it can be written as a term of length $O(k)$ in the language of real closed fields. Thus, we modify the construction of φ' in the proof of Theorem 4.1 so that integer constants k are represented using the above binary encoding. It is easy to see that $|\varphi'|$ is polynomial in $|\varphi|\ \|\varphi\|$ (since $|\Phi'_{\mathsf{h}}|$ and $|\Phi'_{\mathsf{o}}|$ are both polynomial in $|\varphi|$). We can now use the exponential-space algorithm of Ben-Or, Kozen, and Reif (1986) on φ': if φ' is satisfiable, then we can construct the required probability measures, and φ is satisfiable; otherwise, no such probability measures exist, and φ is unsatisfiable. □

Theorem 5.2: *There is a procedure that runs in space exponential in $|\varphi|\ \|\varphi\|$ for deciding whether there exist sets of primitive propositions Φ_{h} and Φ_{o} such that $\varphi \in \mathcal{L}^{fo\text{-}ev}(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$ and φ is satisfiable in an evidential world.*

Proof. Let $h_1, \ldots, h_m$ be the hypotheses appearing in φ, and $ob_1, \ldots, ob_n$ be the hypotheses appearing in φ. Let $\Phi_{\mathsf{h}} = \{h_1, \ldots, h_m, h^*\}$ and $\Phi_{\mathsf{o}} = \{ob_1, \ldots, ob_n, ob^*\}$, where h^* and ob^* are an hypothesis and observation not appearing in φ. Clearly, $|\Phi_{\mathsf{h}}|$ and $|\Phi_{\mathsf{o}}|$ are polynomial in $|\varphi|$. By Lemma A.4, if φ is satisfiable in an evidential world, it is satisfiable in an evidential world over Φ_{h} and Φ_{o}. By Theorem 5.1, we have an algorithm to determine if φ is satisfied in an evidential world over Φ_{h} and Φ_{o} that runs in space exponential in $|\varphi|\ \|\varphi\|$. □

Theorem 5.3: *There is a procedure that runs in space polynomial in $|\varphi|\ \|\varphi\|$ for deciding, given Φ_{h} and Φ_{o}, whether a formula φ of $\mathcal{L}^{ev}(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$ is satisfiable in an evidential world.*

Proof. The proof of this result is very similar to that of Theorem 5.1. Let φ be a formula of $\mathcal{L}^{ev}(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$. By Lemma A.4, φ is satisfiable if there exists a probability measure μ on $\Phi'_{\mathsf{h}} = \mathcal{H} \cup \{h^*\}$ (where $\mathcal{H}$ is the set of hypotheses appearing in φ, and $h^* \notin \mathcal{H}$), probability measures $\mu_{h_1}, \ldots, \mu_{h_m}$ on $\Phi'_{\mathsf{o}} = \mathcal{O} \cup \{ob^*\}$ (where $\mathcal{O}$ is the set of observations appearing in φ and $ob^* \notin \mathcal{O}$), a hypothesis h, observation o, and valuation v such that $(w, v) \models \varphi$, where $w = (h, ob, \mu, \mathcal{E})$ and $\mathcal{E} = (\Phi'_{\mathsf{h}}, \Phi'_{\mathsf{o}}, \boldsymbol{\mu})$.

We derive a formula φ' in the language of real closed fields that asserts the existence of these probability measures by adapting the construction of the formula φ' from φ in the proof of Theorem 4.1. As in the proof of Theorem 5.1, we need to make sure that φ' is polynomial in the size of φ, which the construction in the proof of Theorem 4.1 does not guarantee. We modify the construction so that we use a more efficient representation of integer constants, namely, a binary representation. For example, we can write 42 as $2(1+2^2(1+2^2))$, which can be expressed in the language of real closed fields as $(1+1)(1+(1+1)(1+1)(1+(1+1)(1+1)))$. We can check that if k is a coefficient of length k (when written in binary), it can be written as a term of length $O(k)$ in the language of real closed fields. We modify the construction of φ' in the proof of Theorem 4.1 so that integer constants k are represented using this binary encoding. It is easy to see that $|\varphi'|$ is polynomial in $|\varphi|\ \|\varphi\|$ (since $|\Phi'_{\mathsf{h}}|$ and $|\Phi'_{\mathsf{o}}|$ are both polynomial in $|\varphi|$). The key now is to notice that the resulting formula φ' can be written as $\exists x_1 \ldots \exists x_n(\varphi'')$ for some quantifier-free formula φ''. In this form, we can apply the polynomial space algorithm of Canny (1988) to φ'': if φ'' is satisfiable, then we can construct the required probability measures, and φ is satisfiable; otherwise, no such probability measures exist, and φ is unsatisfiable. $\square$

Theorem 5.4: *There is a procedure that runs in space polynomial in $|\varphi|\ \|\varphi\|$ for deciding whether there exists sets of primitive propositions Φ_{h} and Φ_{o} such that $\varphi \in \mathcal{L}^{ev}(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$ and φ is satisfiable in an evidential world.*

Proof. Let $h_1, \ldots, h_m$ be the hypotheses appearing in φ, and $ob_1, \ldots, ob_n$ be the hypotheses appearing in φ. Let $\Phi_{\mathsf{h}} = \{h_1, \ldots, h_m, h^*\}$ and $\Phi_{\mathsf{o}} = \{ob_1, \ldots, ob_n, ob^*\}$, where h^* and ob^* are an hypothesis and observation not appearing in φ. Clearly, $|\Phi_{\mathsf{h}}|$ and $|\Phi_{\mathsf{o}}|$ are polynomial in $|\varphi|$. By Lemma A.4, if φ is satisfiable in an evidential world, it is satisfiable in an evidential world over Φ_{h} and Φ_{o}. By Theorem 5.3, we have an algorithm to determine if φ is satisfied in an evidential world over Φ_{h} and Φ_{o} that runs in space polynomial in $|\varphi|\ \|\varphi\|$. $\square$

The proofs of Theorem 6.1 and 6.2 rely on the following small model result, a variation on Lemma A.4.

Lemma A.5: *Given Φ_{h} and Φ_{o}, let φ be a formula of $\mathcal{L}^{fo\text{-}ev}(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$, and let $\mathcal{H} \subseteq \Phi_{\mathsf{h}}$ and $\mathcal{O} \subseteq \Phi_{\mathsf{o}}$ be the hypotheses and observations that occur in φ. If φ is satisfiable in an evidential world over Φ_{h} and Φ_{o}, then φ is satisfiable in an evidential world over Φ'_{h} and Φ'_{o} where $|\Phi'_{\mathsf{h}}| = |\mathcal{H}| + 1$ and $|\Phi'_{\mathsf{o}}| = |\mathcal{O}| + 1$, and where, for each $h \in \Phi'_{\mathsf{h}}$ and $ob \in \Phi'_{\mathsf{o}}$, the likelihood $\mu_h(ob)$ is a rational number with size $O(|\varphi|\ \|\varphi\| + |\varphi| \log(|\varphi|))$.*

Proof. Let φ be a formula satisfiable in an evidential world over Φ_{h} and Φ_{o}. By Lemma A.4, φ is satisfiable in an evidential world over Φ'_{h} and Φ'_{o}, where $|\Phi'_{\mathsf{h}}| = |\mathcal{H}|+1$ and $|\Phi'_{\mathsf{o}}| = |\mathcal{O}|+1$. To force the likelihoods to be small, we adapt Theorem 2.6 in FHM, which says that if a formula f in the FHM logic is satisfiable, it is satisfiable in a structure where the probability assigned to each state of the structure is a rational number with size $O(|f|\ \|f\| + |f| \log(|f|))$. The formulas in $\mathcal{L}^{w}(\Phi'_{\mathsf{h}}, \Phi'_{\mathsf{o}})$ are just formulas in the FHM logic. The result adapts immediately, and yields the required bounds for the size of the likelihoods. $\square$

Theorem 6.1: *The problem of deciding, given Φ_{h} and Φ_{o}, whether a formula φ of $\mathcal{L}^{w}(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$ is satisfiable in an evidential world is NP-complete.*

Proof. To establish the lower bound, observe that we can reduce propositional satisfiability to satisfiability in $\mathcal{L}^{w}(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$. More precisely, let f be a propositional formula, where $p_1, \ldots, p_n$ are the primitive propositions appearing in f. Let $\Phi_{\mathsf{o}} = \{ob_1, \ldots, ob_n, ob^*\}$ be a set of observations, where observation ob_i corresponds to the primitive proposition p_i, and ob^* is another (distinct) observation; let Φ_{h} be an arbitrary set of hypotheses, and let h be an arbitrary hypothesis in Φ_{h}. Consider the formula $\hat{f}$ obtained by replacing every occurrence of p_i in f by $\mathrm{w}(ob_i, h) > 0$. It is straightforward to verify that f is satisfiable if and only if $\hat{f}$ is satisfiable in $\mathcal{L}^{w}(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$. (We need the extra observation ob^* to take care of the case f is satisfiable in a a model where each of $p_1, \ldots, p_n$ is false. In that case, $\mathrm{w}(ob_1, h) = \cdots \mathrm{w}(ob_n, h) = 0$, but we can take $\mathrm{w}(ob^*, h) = 1$.) This establishes the lower bound,

The upper bound is straightforward. By Lemma A.5, an evidential world over Φ_{h} and Φ_{o} can be guessed in time polynomial in $|\Phi_{\mathsf{h}}| + |\Phi_{\mathsf{o}}| + |\varphi|\ \|\varphi\|$, since the prior probability in the world requires assigning a value to $|\Phi_{\mathsf{h}}|$ hypotheses, and the evidence space requires

$|\Phi_{\mathsf{h}}|$ likelihood functions, each assigning a value to $|\Phi_{\mathsf{o}}|$ observations, of size polynomial in $|\varphi|\ \|\varphi\|$. We can verify that a world satisfies φ in time polynomial in $|\varphi|\ \|\varphi\| + |\Phi_{\mathsf{h}}| + |\Phi_{\mathsf{h}}|$. This establishes that the problem is in NP. $\square$

Theorem 6.2: *The problem of deciding, for a formula φ, whether there exists sets of primitive propositions Φ_{h} and Φ_{o} such that $\varphi \in \mathcal{L}^w(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$ and φ is satisfiable in an evidential world is NP-complete.*

Proof. For the lower bound, we reduce from the decision problem of $\mathcal{L}^w(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$ over fixed Φ_{h} and Φ_{o}. Let $\Phi_{\mathsf{h}} = \{h_1, \ldots, h_m\}$ and $\Phi_{\mathsf{o}} = \{ob_1, \ldots, ob_n\}$, and let φ be a formula in $\mathcal{L}^w(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$. We can check that φ is satisfiable in evidential world over Φ_{h} and Φ_{o} if and only if $\varphi \wedge (h_1 \vee \cdots \vee h_m) \wedge (ob_1 \vee \cdots \vee ob_n)$ is satisfiable in an evidential world over arbitrary Φ'_{h} and Φ'_{o}. Thus, by Theorem 6.1, we get our lower bound.

For the upper bound, by Lemma A.5, if φ is satisfiable, it is satisfiable in an evidential world over Φ_{h} and Φ_{o}, where $\Phi_{\mathsf{h}} = \mathcal{H} \cup \{h^*\}$, $\mathcal{H}$ consists of the hypotheses appearing in φ, $\Phi_{\mathsf{o}} = \mathcal{O} \cup \{ob^*\}$, $\mathcal{O}$ consists of the observations appearing in φ, and h^* and ob^* are new hypotheses and observations. Thus, $|\Phi_{\mathsf{h}}| \leq |\varphi| + 1$, and $|\Phi_{\mathsf{o}}| \leq |\varphi| + 1$. As in the proof of Theorem 6.1, such a world can be guessed in time polynomial in $|\varphi|\ \|\varphi\| + |\Phi_{\mathsf{h}}| + |\Phi_{\mathsf{o}}|$, and therefore in time polynomial in $|\varphi|\ \|\varphi\|$. We can verify that this world satisfies φ in time polynomial in $|\varphi|\ \|\varphi\|$, establishing that the problem is in NP. $\square$

Theorem 7.1: *$\boldsymbol{AX}_{dyn}(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$ is a sound and complete axiomatization for $\mathcal{L}^{fo\text{-}ev}_{dyn}(\Phi_{\mathsf{h}}, \Phi_{\mathsf{o}})$ with respect to evidential runs.*

Proof. It is easy to see that each axiom is valid in evidential runs. To prove completeness, we follow the same procedure as in the proof of Theorem 4.1, showing that if φ is consistent, then it is satisfiable, that is, there exists an evidential run r and valuation v such that $(r, m, v) \models \varphi$ for some point (r, m) of r.

As in the body of the paper, let $\Phi_{\mathsf{h}} = \{h_1, \ldots, h_{n_{\mathsf{h}}}\}$ and $\Phi_{\mathsf{o}} = \{ob_1, \ldots, ob_{n_{\mathsf{o}}}\}$. Let φ be a consistent formula. The first step of the process is to reduce the formula φ to a canonical form with respect to the $\bigcirc$ operator. Intuitively, we push down every occurrence of a $\bigcirc$ to the polynomial inequality formulas present in the formula. It is easy to see that axioms and inference rules **T1**–**T6** can be used to establish that φ is provably equivalent to a formula φ' where every occurrence of $\bigcirc$ is in the form of subformulas $\bigcirc^n(ob)$ and $\bigcirc^n(p \geq c)$, where p is a polynomial term that contains at least one occurrence of the Pr operator. We use the notation $\bigcirc^n\varphi$ for $\bigcirc \ldots \bigcirc\varphi$, the n-fold application of $\bigcirc$ to φ. We write $\bigcirc^0\varphi$ for φ. Let N be the maximum coefficient of $\bigcirc$ in φ'.

By way of contradiction, assume that φ' (and hence φ) is unsatisfiable. As in the proof of Theorem 4.1, we reduce the formula φ' to an equivalent formula in the language of real closed fields. Let $u_1, \ldots, u_{n_{\mathsf{h}}}$, $v_1^0, \ldots, v_{n_{\mathsf{o}}}^0, \ldots, v_1^N, \ldots, v_{n_{\mathsf{o}}}^N$, $y_1^0, \ldots, y_{n_{\mathsf{o}}}^0, \ldots, y_1^N, \ldots, y_{n_{\mathsf{o}}}^N$, and $z_{\langle i_1, \ldots, i_k \rangle, 1}, \ldots, z_{\langle i_1, \ldots, i_k \rangle, n_{\mathsf{h}}}$ (for every sequence $\langle i_1, \ldots, i_k \rangle$) be new variables, where, intuitively,

- u_i gets value 1 if hypothesis h_i holds, 0 otherwise;
- v_i^n gets value 1 if observation ob_i holds at time n, 0 otherwise;
- y_i^n represents $\Pr(h_i)$ at time n;

- $z_{\langle i_1,\ldots,i_k\rangle,j}$ represents $\mathrm{w}(\langle ob^{i_1},\ldots,ob^{i_k}\rangle,h_j)$.

The main difference with the construction in the proof of Theorem 4.1 is that we have variables v_i^n representing the observations at every time step n, rather than variables representing observations at the only time step, variables y_i^n representing each hypothesis probability at every time step, rather than variables representing prior and posterior probabilities, and variables $z_{\langle i_1,\ldots,i_k\rangle,j}$ representing the weight of evidence of sequences of observations, rather than variables representing the weight of evidence of single observations. Let $\mathbf{v}$ represent that list of new variables. We consider the same formulas as in the proof of Theorem 4.1, modified to account for the new variables, and the fact that we are reasoning over multiple time steps. More specifically, the formula φ_h is unchanged. Instead of φ_o, we consider formulas $\varphi_o^1,\ldots,\varphi_o^N$ saying that exactly one observation holds at each time time step, where φ_o^n is given by:

$$(v_1^n = 0 \vee v_1^n = 1) \wedge \cdots \wedge (v_{n_\mathsf{o}}^n = 0 \vee v_{n_\mathsf{h}}^n = 1) \wedge v_1^n + \cdots + v_{n_\mathsf{h}}^n = 1.$$

Let $\varphi_o' = \varphi_o^1 \wedge \cdots \wedge \varphi_o^N$.

Similarly, instead of φ_{pr} and φ_{po}, we consider formulas $\varphi_p^1,\ldots,\varphi_p^N$ expressing that Pr is a probability measure at each time step, where φ_p^n is given by:

$$y_1^n \geq 0 \wedge \cdots \wedge y_{n_\mathsf{h}}^n \geq 0 \wedge y_1^n + \cdots + y_{n_\mathsf{h}}^n = 1.$$

Let $\varphi_p = \varphi_p^1 \wedge \cdots \wedge \varphi_p^N$.

Similarly, we consider $\varphi_{w,p}$ and $\varphi_{w,f}$, except where we replace variables $z_{i,j}$ by $z_{\langle i\rangle,j}$, to reflect the fact that we now consider sequences of observations. The formula $\varphi_{w,up}$, capturing the update of a prior probability into a posterior probability given by **E5**, is replaced by the formulas $\varphi_{w,up}^1,\ldots,\varphi_{w,up}^N$ representing the update of the probability at each time step, where $\varphi_{w,up}^n$ is given by the obvious generalization of $\varphi_{w,up}$:

$$\begin{aligned}
&(v_1^n = 1 \Rightarrow (y_1^{n-1} z_{1,1} = y_1^n y_1^{n-1} z_{1,1} + \cdots + y_1^n y_{n_\mathsf{h}}^{n-1} z_{1,n_\mathsf{h}} \wedge \\
&\qquad \cdots \wedge y_{n_\mathsf{h}}^{n-1} z_{1,n_\mathsf{h}} = y_{n_\mathsf{h}}^n y_1^{n-1} z_{1,1} + \cdots + y_{n_\mathsf{h}}^n y_{n_\mathsf{h}}^{n-1} z_{1,n_\mathsf{h}})) \wedge \\
&\cdots \wedge \\
&(v_{n_\mathsf{o}}^n = 1 \Rightarrow (y_1^{n-1} z_{n_\mathsf{o},1} = y_1^n y_1^{n-1} z_{n_\mathsf{o},1} + \cdots + y_1^n y_{n_\mathsf{h}}^{n-1} z_{n_\mathsf{o},n_\mathsf{h}} \wedge \\
&\qquad \cdots \wedge y_{n_\mathsf{h}}^{n-1} z_{n_\mathsf{o},n_\mathsf{h}} = y_{n_\mathsf{h}}^n y_1^{n-1} z_{n_\mathsf{o},1} + \ldots y_{n_\mathsf{h}}^n y_{n_\mathsf{h}}^{n-1} z_{n_\mathsf{o},n_\mathsf{h}})).
\end{aligned}$$

Let $\varphi_{w,up}' = \varphi_{w,up}^1 \wedge \cdots \wedge \varphi_{w,up}^N$.

Finally, we need a new formula $\varphi_{w,c}$ capturing the relationship between the weight of evidence of a sequence of observations, and the weight of evidence of the individual observations, to capture axiom **E6**:

$$\begin{aligned}
&\bigwedge_{\substack{1\leq k\leq N\\ 1\leq i_1,\ldots,i_k\leq n_\mathsf{o}}} z_{\langle i_1\rangle,h_1} \cdots z_{\langle i_k\rangle,h_1} = z_{\langle i_1,\ldots,i_k\rangle,h_1} z_{\langle i_1\rangle,h_1} \cdots z_{\langle i_k\rangle,h_1} \\
&\qquad\qquad + \cdots + z_{\langle i_1,\ldots,i_k\rangle,h_1} z_{\langle i_1\rangle,h_{n_\mathsf{h}}} \cdots z_{\langle i_k\rangle,h_{n_\mathsf{h}}} \wedge \\
&\cdots \wedge \bigwedge_{\substack{1\leq k\leq N\\ 1\leq i_1,\ldots,i_k\leq n_\mathsf{o}}} z_{\langle i_1\rangle,h_{n_\mathsf{h}}} \cdots z_{\langle i_k\rangle,h_{n_\mathsf{h}}} = z_{\langle i_1,\ldots,i_k\rangle,h_{n_\mathsf{h}}} z_{\langle i_1\rangle,h_1} \cdots z_{\langle i_k\rangle,h_1} \\
&\qquad\qquad + \cdots + z_{\langle i_1,\ldots,i_k\rangle,h_{n_\mathsf{h}}} z_{\langle i_1\rangle,h_{n_\mathsf{h}}} \cdots z_{\langle i_k\rangle,h_{n_\mathsf{h}}}.
\end{aligned}$$

Let $\hat{\varphi}$ be the formula in the language of real closed fields obtained from φ by replacing each occurrence of the primitive proposition h_i by $u_i = 1$, each occurrence of $\bigcirc^n ob_i$ by $v_i^n = 1$, and within each polynomial inequality formula $\bigcirc^n(p \geq c)$, replacing each occurrence of $\Pr(\rho)$ by $\sum_{h_i \in [\![\rho]\!]} y_i^n$, each occurrence of $\mathrm{w}(\langle ob^{i_1}, \ldots, ob^{i_k}\rangle, h_j)$ by $z_{\langle i_1,\ldots,i_k\rangle,j}$, and each occurrence of an integer coefficient k by $1 + \cdots + 1$ (k times). Finally, let φ' be the formula $\exists \mathbf{v}(\varphi_h \wedge \varphi'_o \wedge \varphi_p \wedge \varphi_{w,p} \wedge \varphi_{w,f} \wedge \varphi'_{w,up} \wedge \varphi_{w,c} \wedge \hat{\varphi})$.

It is easy to see that if φ is unsatisfiable over evidential systems, then φ' is false about the real numbers. Therefore, $\neg\varphi'$ must be a formula valid in real closed fields, and hence an instance of **RCF**. Thus, $\neg\varphi'$ is provable. It is straightforward to show, using the obvious variant of Lemma A.3 that $\neg\varphi$ itself is provable, contradicting the fact that φ is consistent. Thus, φ must be satisfiable, establishing completeness. □

References

Ben-Or, M., Kozen, D., & Reif, J. H. (1986). The complexity of elementary algebra and geometry. *Journal of Computer and System Sciences*, *32*(1), 251–264.

Canny, J. F. (1988). Some algebraic and geometric computations in PSPACE. In *Proc. 20th Annual ACM Symposium on the Theory of Computing (STOC'88)*, pp. 460–467.

Carnap, R. (1962). *Logical Foundations of Probability* (Second edition). University of Chicago Press.

Casella, G., & Berger, R. L. (2001). *Statistical Inference* (Second edition). Duxbury.

Chan, H., & Darwiche, A. (2005). On the revision of probabilistic beliefs using uncertain evidence. *Artificial Intelligence*, *163*, 67–90.

de Alfaro, L. (1998). *Formal Verification of Probabilistic Systems*. Ph.D. thesis, Stanford University. Available as Technical Report STAN-CS-TR-98-1601.

Enderton, H. B. (1972). *A Mathematical Introduction to Logic*. Academic Press.

Fagin, R., & Halpern, J. Y. (1994). Reasoning about knowledge and probability. *Journal of the ACM*, *41*(2), 340–367.

Fagin, R., Halpern, J. Y., & Megiddo, N. (1990). A logic for reasoning about probabilities. *Information and Computation*, *87*(1/2), 78–128.

Fischer, M. J., & Zuck, L. D. (1988). Reasoning about uncertainty in fault-tolerant distributed systems. Technical report YALEU/DCS/TR–643, Yale University.

Good, I. J. (1950). *Probability and the Weighing of Evidence*. Charles Griffin & Co. Ltd.

Good, I. J. (1960). Weights of evidence, corroboration, explanatory power, information and the utility of experiments. *Journal of the Royal Statistical Society, Series B*, *22*, 319–331.

Halpern, J. Y., & Fagin, R. (1992). Two views of belief: belief as generalized probability and belief as evidence. *Artificial Intelligence*, *54*, 275–317.

Halpern, J. Y., Moses, Y., & Tuttle, M. R. (1988). A knowledge-based analysis of zero knowledge. In *Proc. 20th Annual ACM Symposium on the Theory of Computing (STOC'88)*, pp. 132–147.

Halpern, J. Y., & Pucella, R. (2003). A logic for reasoning about evidence. In *Proc. 19th Conference on Uncertainty in Artificial Intelligence (UAI'03)*, pp. 297–304.

Halpern, J. Y., & Pucella, R. (2005a). Evidence with uncertain likelihoods. In *Proc. 21th Conference on Uncertainty in Artificial Intelligence (UAI'05)*, pp. 243–250.

Halpern, J. Y., & Pucella, R. (2005b). Probabilistic algorithmic knowledge. *Logical Methods in Computer Science*, *1*(3:1).

Halpern, J. Y., & Tuttle, M. R. (1993). Knowledge, probability, and adversaries. *Journal of the ACM*, *40*(4), 917–962.

He, J., Seidel, K., & McIver, A. (1997). Probabilistic models for the guarded command language. *Science of Computer Programming*, *28*(2–3), 171–192.

Jeffrey, R. C. (1992). *Probability and the Art of Judgement.* Cambridge University Press.

Kyburg, Jr., H. E. (1983). Recent work in inductive logic. In Machan, T., & Lucey, K. (Eds.), *Recent Work in Philosophy*, pp. 87–150. Rowman & Allanheld.

Milne, P. (1996). $\log[p(h|eb)/p(h|b)]$ is the one true measure of confirmation. *Philosophy of Science*, *63*, 21–26.

Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference.* Morgan Kaufmann.

Popper, K. R. (1959). *The Logic of Scientific Discovery.* Hutchinson.

Renegar, J. (1992). On the computational complexity and geometry of the first order theory of the reals. *Journal of Symbolic Computation*, *13*(3), 255–352.

Shafer, G. (1976). *A Mathematical Theory of Evidence.* Princeton University Press.

Shafer, G. (1982). Belief functions and parametric models (with commentary). *Journal of the Royal Statistical Society, Series B*, *44*, 322–352.

Shoenfield, J. R. (1967). *Mathematical Logic.* Addison-Wesley, Reading, Mass.

Tarski, A. (1951). *A Decision Method for Elementary Algebra and Geometry* (2nd edition). Univ. of California Press.

Vardi, M. Y. (1985). Automatic verification of probabilistic concurrent finite-state programs. In *Proc. 26th IEEE Symposium on the Foundations of Computer Science (FOCS'85)*, pp. 327–338.

Walley, P. (1987). Belief function representations of statistical evidence. *Annals of Statistics*, *18*(4), 1439–1465.