



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

A First-Order Logic of Probability and Only Knowing in Unbounded Domains

Citation for published version:

Belle, V, Lakemeyer, G & Levesque, HJ 2016, A First-Order Logic of Probability and Only Knowing in Unbounded Domains. in *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA.* AAAI Press, pp. 893-899, Thirtieth AAAI Conference on Artificial Intelligence, Phoenix, Arizona, United States, 12/02/16.
<<http://www.aaai.org/ocs/index.php/AAAI/AAAI16/paper/view/12271>>

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Publisher's PDF, also known as Version of record

Published In:

Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA.

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



A First-Order Logic of Probability and Only Knowing in Unbounded Domains

Vaishak Belle*
 Dept. of Computer Science
 KU Leuven
 Belgium
 vaishak@cs.kuleuven.be

Gerhard Lakemeyer†
 Dept. of Computer Science
 RWTH Aachen University
 Germany
 gerhard@cs.rwth-aachen.de

Hector J. Levesque
 Dept. of Computer Science
 University of Toronto
 Canada
 hector@cs.toronto.edu

Abstract

Only knowing captures the intuitive notion that the beliefs of an agent are precisely those that follow from its knowledge base. It has previously been shown to be useful in characterizing knowledge-based reasoners, especially in a quantified setting. While this allows us to reason about incomplete knowledge in the sense of not knowing whether a formula is true or not, there are many applications where one would like to reason about the degree of belief in a formula. In this work, we propose a new general first-order account of probability and only knowing that admits knowledge bases with incomplete and probabilistic specifications. Beliefs and non-beliefs are then shown to emerge as a direct logical consequence of the sentences of the knowledge base at a corresponding level of specificity.

Somewhat surprisingly, \mathcal{OL} has a particularly simple semantics. A model, or an epistemic state, is simply a set of worlds e , which satisfies only knowing a sentence α just in case e is maximal, that is, adding any other world to e would lead to not believing α any more. Here worlds are simply truth assignments to atoms of a first-order language and believing is interpreted in the usual way as truth in all worlds in e . Over the years, only knowing has been studied and extended in various ways, such as capturing defaults (Lakemeyer and Levesque 2005), modeling dynamical worlds (Lakemeyer and Levesque 2004), and reasoning about multiple agents (Halpern and Lakemeyer 2001; Waaler and Solhaug 2005; Belle and Lakemeyer 2010).

Be that as it may, first-order accounts of only knowing have limited themselves to categorical knowledge, where a formula ϕ is either known or not known. In many real-world applications, it is not enough to deal with incomplete knowledge in that sense. Here, one must also know which of ϕ or $\neg\phi$ is the more likely, and by how much. Similarly, in a robotic setting (Bacchus, Halpern, and Levesque 1999), some outcomes are more likely than others, perhaps due to error prone actuators. Reasoning about probabilistic beliefs is then a critical feature. To that end, it is natural to ask the following: if all that is known is p , what is the *degree of belief* in q ? More generally, if a knowledge base includes partial, incomplete or probabilistic specifications, what can the agent be expected to believe?

In this work, we attempt a response to these questions by proposing a *first-order* account of *subjective probability* and *only knowing*. The reasonableness of our proposal can be evaluated in different ways. At one extreme, belief is shown to exhibit the usual properties of subjective probabilities (Fagin and Halpern 1994). At the other extreme, the logic is also shown to exhibit the properties of classical only knowing as desired. Moreover, like \mathcal{OL} , our account admits a substitutional interpretation of quantifiers over an infinite domain, giving us a simple way to address issues such as quantifying-in (Kaplan 1968). Perhaps most significantly, the properties of belief emerge as a logical consequence of the sentences of the knowledge base at a corresponding level of specificity. In particular, this will mean that for partial and incomplete specifications, where a single probability distribution would often not suffice, suitable beliefs are nonetheless entailed.

1 Introduction

When considering knowledge-based reasoners, it seems intuitive that the beliefs of the agent are those that follow from the assumption that its knowledge base is all that is known.¹ Levesque (Levesque 1990) was among the first to capture this idea in the logic of *only knowing* \mathcal{OL} , where a modality \mathcal{O} is introduced in addition to the classical epistemic operator \mathbf{K} . For example, from $\mathcal{O}p$ it follows that $\mathbf{K}p$, but more interestingly, it also follows that $\neg\mathbf{K}q$. This is different from classical epistemic logic where $\mathbf{K}p$ does not preclude $\mathbf{K}(p \wedge q)$. Similarly, in a quantified setting, we get the following valid sentence:

$$\mathcal{O}(P(a) \vee P(b)) \supset \mathbf{K}(\exists x[P(x) \wedge \neg\mathbf{K}P(x)])$$

which says that “if all I know is $P(a)$ or $P(b)$ then I know that there is an instance of P but not what.” In other words, a precise characterization of the *beliefs* and the *non-beliefs* of a knowledge-based reasoner can be given in a succinct manner using only knowing. With this meta-knowledge, the agent is then in a position to do something when its knowledge is incomplete, ask a question for example.

*Supported by the Research Foundation-Flanders (FWO-Vlaanderen).

†Also affiliated with the University of Toronto. Copyright © 2016, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹In this paper, we do not distinguish between “knowledge” and “belief” and freely use the terms interchangeably.

We structure our work as follows. We begin by introducing the logic OBL (= OL + *degrees of belief* modality). Then, we turn to properties and example specifications, OBL 's relation to OL , discuss related work, and conclude.

2 The Logic OBL

Syntax

The non-modal fragment of OBL consists of standard first-order logic with = (that is, connectives $\{\wedge, \vee, \neg\}$, syntactic abbreviations $\{\exists, \equiv, \supset\}$) and a countably infinite set of *standard names* \mathcal{N} , which includes the set of rationals as a sub-sort. As we shall see, these standard names will serve as a fixed domain of discourse, permitting a substitutional interpretation for quantifiers. To keep matters simple, function symbols are not considered in this language. We call a predicate other than =, applied to first-order variables or standard names, an *atomic* formula. We write α_n^x to mean that the variable x is substituted in α by a standard name n . If all the variables in an atom are substituted by standard names, then we call it a *ground* atom. Let \mathcal{P} be the set of ground atoms in OBL .

OBL has two epistemic operators: $B(\alpha: r)$ is to be read as “ α is believed with a probability r ,” where r is a rational number. Next, the modality $O(\alpha_1: r_1, \dots, \alpha_k: r_k)$, where α_i does not mention modalities and r_i is a rational, is to be read as “all that is believed is: α_1 with probability r_1, \dots , and α_k with probability r_k .” We also use $K\alpha$, to be read as “ α is known,” as an abbreviation for $B(\alpha: 1)$. We write $O\alpha$, to be read as “ α is all that is known,” to mean $O(\alpha: 1)$.

A formula not mentioning modalities is called *objective*, and a formula where all predicate symbols appear within the scope of a modal operator is called *subjective*.

Semantics

The semantics is given in terms of possible worlds. Here, a world w maps the ground atoms in \mathcal{P} to $\{0, 1\}$, and let \mathcal{W} be the set of all worlds. By a *distribution* d we mean a function from \mathcal{W} to the set of positive reals $\mathbb{R}^{\geq 0}$. An *epistemic state* e is defined as a set of distributions. By a *model*, we mean a pair (e, w) .

As can be surmised, the definition of a distribution here does not immediately lead to a probability distribution in the usual sense. To obtain probability distributions, we appeal to three simple conditions that will be used in the semantics:

Definition 1: Let d be any distribution, $\mathcal{V} \subseteq \mathcal{W}$ any set of worlds, and r any real number. We define conditions BOUND, EQUAL and NORM as follows:

- $\text{NORM}(d, \mathcal{V}, r)$ iff there is a number $b \neq 0$ such that $\text{EQUAL}(d, \mathcal{V}, b \times r)$ and $\text{EQUAL}(d, \mathcal{W}, b)$;
- $\text{EQUAL}(d, \mathcal{V}, r)$ iff $\text{BOUND}(d, \mathcal{V}, r)$ and there is no $r' < r$ such that $\text{BOUND}(d, \mathcal{V}, r')$;
- $\text{BOUND}(d, \mathcal{V}, r)$ iff there is no $k, w_1, \dots, w_k \in \mathcal{V}$ such that

$$\sum_{i=1}^k d(w_i) > r.$$

Intuitively, given $\text{NORM}(d, \mathcal{V}, r)$, r can be seen as the normalization of the weights of worlds in \mathcal{V} in relation to the set of all worlds \mathcal{W} as accorded by d . Here, $\text{EQUAL}(d, \mathcal{V}, b)$ expresses that the weight accorded to the worlds in \mathcal{V} is b , and finally $\text{BOUND}(d, \mathcal{V}, b)$ ensures the weight of worlds in \mathcal{V} is bounded by b . In essence, although \mathcal{W} is uncountable, the conditions BOUND and EQUAL admit a well-defined summation of the weights on worlds:

Theorem 2 : *Suppose d is a distribution, and let $\mathcal{V} = \{w \in \mathcal{W} \mid d(w) \neq 0\}$. For any number $b \geq 0$, if $\text{BOUND}(d, \mathcal{W}, b)$ then \mathcal{V} is countable.*

Proof: Consider the sequence of sets of worlds $\mathcal{V}_i = \{w \in \mathcal{V} \mid d(w) \geq 1/i\}$ for positive integers i . It is easy to see that $\mathcal{V} = \cup \mathcal{V}_i$, a property we denote by (*). Now suppose \mathcal{V} is uncountable. Then there is some $\epsilon > 0$ such that $\mathcal{V}_\epsilon = \{w \in \mathcal{V} \mid d(w) \geq \epsilon\}$ is infinite. (For otherwise we could enumerate \mathcal{V} by enumerating \mathcal{V}_i starting at $i = 1$ using (*).) Now, consider any countably infinite sequence of worlds w_j taken from \mathcal{V}_ϵ . Since $d(w_j) \geq \epsilon$ for all w_j , the sum $\sum_{j=1}^{\infty} d(w_j)$ is clearly unbounded, contradicting the assumption that \mathcal{W} is bounded by b . ■

Corollary 3: *Suppose d is as above, $\mathcal{V} \subseteq \mathcal{W}$ is infinite, and $\text{BOUND}(d, \mathcal{V}, b)$. Then $\{w \in \mathcal{V} \mid d(w) \neq 0\}$ is countable.*

Theorem 4 *Suppose $\mathcal{V} \subseteq \mathcal{W}$ is infinite, d is a distribution, b a number and $\text{EQUAL}(d, \mathcal{V}, b)$ holds. Then there is an infinite sequence w_1, \dots , with $w_i \in \mathcal{V}$ such that $\sum_{i=1}^{\infty} d(w_i) = b$.*

Proof: Since $\text{EQUAL}(d, \mathcal{V}, b)$, by definition, $\text{BOUND}(d, \mathcal{V}, b)$. By Corollary 3, $\mathcal{U} = \{w \in \mathcal{V} \mid d(w) \neq 0\}$ is countable. Consider any sequence consisting of all worlds $w_i \in \mathcal{U}$. We recall the property that for any series $\sum_{i=1}^{\infty} a_i$, where a_i is nonnegative, the series converges iff it is bounded. Then, we have that $\sum_{i=1}^{\infty} d(w_i)$ is bounded by assumption and so it converges, and suppose $\sum_{i=1}^{\infty} d(w_i) = b'$. It is easy to see that if $b' \neq b$, then $\text{EQUAL}(d, \mathcal{V}, b)$ is a contradiction. ■

We are now prepared for a semantics. Given $\alpha \in OBL$ and a model (e, w) , the semantic rules are as follows:

- $e, w \models p$ iff $w[p] = 1$;
- $e, w \models n_1 = n_2$ iff n_1 and n_2 are the same standard names;
- $e, w \models \neg\alpha$ iff $e, w \not\models \alpha$;
- $e, w \models \alpha \wedge \beta$ iff $e, w \models \alpha$ and $e, w \models \beta$;
- $e, w \models \forall x\alpha$ iff $e, w \models \alpha_n^x$ for all names n ;
- $e, w \models B(\alpha: r)$ iff
for all $d \in e$, $\text{NORM}(d, \{w' \mid e, w' \models \alpha\}, r)$;
- $e, w \models O(\alpha_1: r_1, \dots, \alpha_k: r_k)$ iff
for all d , $d \in e$ iff $\text{NORM}(d, \{w' \mid e, w' \models \alpha_1\}, r_1), \dots, \text{NORM}(d, \{w' \mid e, w' \models \alpha_k\}, r_k)$.

We often write $e \models \alpha$ when the world is irrelevant (that is, when α is a subjective formula), and write $w \models \alpha$ when the epistemic state is irrelevant (that is, when α is an objective formula). Given any set of sentences Σ , we write $\Sigma \models \alpha$ (read: “ Σ entails α ”) to mean that if $e, w \models \beta$ for every $\beta \in \Sigma$ then $e, w \models \alpha$. Finally, we write $\models \alpha$ (read: “ α is valid”) to mean $\{\} \models \alpha$. We let TRUE denote a tautologous sentence, such as $\forall x(x = x)$, and let FALSE denote its negation.

3 Properties of Belief

In this section, we argue for the reasonableness of our semantic rules for the B modality. (A discussion on only knowing is deferred to Section 5.) To that end, it will be convenient to identify epistemic states that can interpret every atomic event:

Definition 5: We say e is measurable iff for every $d \in e$, for every atom p , there is a $r \geq 0$ such that $\{d\} \models B(p: r)$.

As we now show, measurability² is intimately related to the notion of a “reasonable” epistemic state, in that valid sentences are always known. In the sequel, given α and e , we write \mathcal{W}_α^e to mean the set $\{w \mid e, w \models \alpha\}$, often dropping the superscript when the context is clear.

Theorem 6: e is measurable iff $e \models B(\text{TRUE}: 1)$.

Proof: Suppose e is measurable. Then for every $d \in e$, for every p , there is some r , such that $\{d\} \models B(p: r)$. That is, $\text{NORM}(d, \mathcal{W}_p, r)$, and so for some $b \neq 0$, $\text{EQUAL}(d, \mathcal{W}_p, b \times r)$ and $\text{EQUAL}(d, \mathcal{W}, b)$. Thus, for every $d \in e$, $\text{NORM}(d, \mathcal{W}, 1)$ and so $e \models B(\text{TRUE}: 1)$.

Conversely, suppose $e \models B(\text{TRUE}: 1)$ and so, for every $d \in e$, $\text{NORM}(d, \mathcal{W}, 1)$, that is, $\text{EQUAL}(d, \mathcal{W}, b)$ for some positive b . Then for any ground atom p , $\text{BOUND}(d, \mathcal{W}_p, b)$ since $\mathcal{W}_p \subseteq \mathcal{W}$. By Corollary 3, $\mathcal{V} = \{w \in \mathcal{W}_p \mid d(w) \neq 0\}$ is countable, and since bounded countable series of nonnegative terms converge (see argument for Theorem 4), let b^* be such that $\text{EQUAL}(d, \mathcal{W}_p, b^*)$. So, $\text{NORM}(d, \mathcal{W}_p, b^*/b)$. Since this applies to all $d \in e$, clearly e is measurable. ■

Consequently, we obtain the following properties for B regarding truth:³

- $\neg K_{\text{TRUE}}$ is satisfiable

Proof: Suppose d maps every $w \in \mathcal{W}$ to an integer > 0 . Then there is no b such that $\text{BOUND}(d, \mathcal{W}, b)$ holds, and so also not $\text{EQUAL}(d, \mathcal{W}, b)$. Thus, for any e such that $d \in e$, $e \not\models B(\text{TRUE}: 1)$. ■

- $e \models B(\text{TRUE}: 1) \supset B(\text{FALSE}: 0)$

Proof: Suppose $e \models B(\text{TRUE}: 1)$. Then for every $d \in e$, $\text{NORM}(d, \mathcal{W}, 1)$ and so $\text{NORM}(d, \{\}, 0)$, that is, $e \models B(\text{FALSE}: 0)$. ■

Additivity and Equivalence

Beyond truth and falsehood, B can be shown to exhibit the following reasonable properties (as in, for example, (Fagin and Halpern 1994)):

- $e \models B(\alpha \wedge \beta: r) \wedge B(\alpha \wedge \neg\beta: r') \supset B(\alpha: r + r')$

Proof: Suppose $e \models B(\alpha \wedge \beta: r)$ and $e \models B(\alpha \wedge \neg\beta: r')$. This means that for all $d \in e$, $\text{NORM}(d, \mathcal{W}_{\alpha \wedge \beta}, r)$ and

²Our terminology is such because this property can be seen to assign a probability measure to $\mathcal{W}_p = \{w \mid w \models p\}$ for every atom p . Roughly, this is a definition to realize the *measurable case* in the sense of (Fagin and Halpern 1994).

³Of course, if needed, we can simply restrict ourselves to measurable epistemic states by means of which truth is always believed with probability 1.

$\text{NORM}(d, \mathcal{W}_{\alpha \wedge \neg\beta}, r')$. So for some $b \neq 0$, $\text{EQUAL}(d, \mathcal{W}_{\alpha \wedge \beta}, b \times r)$, $\text{EQUAL}(d, \mathcal{W}_{\alpha \wedge \neg\beta}, b \times r')$ and $\text{EQUAL}(d, \mathcal{W}, b)$ hold. Since $\mathcal{W}_{\alpha \wedge \beta} \cap \mathcal{W}_{\alpha \wedge \neg\beta}$ is empty, and $\mathcal{W}_{\alpha \wedge \beta} \cup \mathcal{W}_{\alpha \wedge \neg\beta} - \mathcal{W}_\alpha$ is empty, we have $\text{EQUAL}(d, \mathcal{W}_\alpha, b \times r + b \times r')$. Hence $\text{NORM}(d, \mathcal{W}_\alpha, r + r')$ for every $d \in e$, and so $e \models B(\alpha: r + r')$. ■

- if r is a negative number then $B(\alpha: r)$ is unsatisfiable

Proof: If $e \models B(\alpha: -r)$ where r is a positive real, then for every $d \in e$, $\text{NORM}(d, \mathcal{W}_\alpha, -r)$ which is impossible by the definition of d that maps worlds to positive reals. ■

- if $e \models \alpha \equiv \beta$, then $e \models B(\alpha: r) \equiv B(\beta: r)$

Proof: For any e , and for every $d \in e$, since $\mathcal{W}_\alpha = \mathcal{W}_\beta$, if $\text{NORM}(d, \mathcal{W}_\alpha, r)$ then $\text{NORM}(d, \mathcal{W}_\beta, r)$ and vice versa. Thus, $e \models B(\alpha: r) \supset B(\beta: r)$ and vice versa. ■

- $e \models B(\alpha: r) \wedge B(\beta: r') \wedge B(\alpha \wedge \beta: r'') \supset B(\alpha \vee \beta: r + r' - r'')$
 $OB\mathcal{L}$ satisfies the addition law of probability.

Knowledge

As in \mathcal{OL} (Levesque and Lakemeyer 2001), the following properties of knowledge hold, which include the universal and existential versions of the Barcan formula:

- $e \models K\alpha \wedge K(\alpha \supset \beta) \supset K\beta$

Proof: Suppose $e \models K\alpha$ and $e \models K(\alpha \supset \beta)$. Then for all $d \in e$, $\text{NORM}(d, \mathcal{W}_\alpha, 1)$ and $\text{NORM}(d, \mathcal{W}_{\alpha \supset \beta}, 1)$. So for every $d \in e$, there is a $b \neq 0$ such that $\text{EQUAL}(d, \mathcal{W}_\alpha, b)$, $\text{EQUAL}(d, \mathcal{W}_{\alpha \supset \beta}, b)$ and $\text{EQUAL}(d, \mathcal{W}, b)$. In particular, because $\text{EQUAL}(d, \mathcal{W}_\alpha, b)$ and $\text{EQUAL}(d, \mathcal{W}, b)$, there is an infinite sequence w_1, \dots with $w_i \in \mathcal{W}_\alpha$ such that⁴ $\sum d(w_i) = b$ and for $w \notin \mathcal{W}_\alpha$, $d(w) = 0$. (Otherwise, $\text{EQUAL}(d, \mathcal{W}, b)$ cannot be true.) Moreover, there is no finite sequence w_1, \dots, w_k with $w_i \in \mathcal{W}_\alpha$ such that $d(w_1) + \dots + d(w_k) > b$. An analogous observation can be made for $\text{EQUAL}(d, \mathcal{W}_{\alpha \supset \beta}, b)$, and so there is an infinite sequence w_1, \dots with $w_i \in \mathcal{W}_\alpha \cap \mathcal{W}_{\alpha \supset \beta}$ such that $\sum d(w_i) = b$, and for $w \notin \mathcal{W}_\alpha \cap \mathcal{W}_{\alpha \supset \beta}$, $d(w) = 0$. Since $\mathcal{W}_\beta \supseteq \mathcal{W}_\alpha \cap \mathcal{W}_{\alpha \supset \beta}$ and $d(w) = 0$ for $w \in \mathcal{W}_\beta - \mathcal{W}_\alpha \cap \mathcal{W}_{\alpha \supset \beta}$, there is an infinite sequence w_1, \dots with $w_i \in \mathcal{W}_\beta$ such that $\sum d(w_i) = b$ but no finite sequence w_1, \dots, w_k such that $d(w_1) + \dots + d(w_k) > b$. That is, $\text{EQUAL}(d, \mathcal{W}_\beta, b)$ and because $\text{EQUAL}(d, \mathcal{W}, b)$, we get $\text{NORM}(d, \mathcal{W}_\beta, 1)$. This holds for every $d \in e$. So $e \models K\beta$. ■

- $e \models \forall x K\alpha \supset K\forall x\alpha$

Proof: Suppose $e, w \models \forall x K\alpha$. Then, $e, w \models K\alpha_n^x$ for all names n . Thus, for all $d \in e$, $\text{NORM}(d, \mathcal{W}_{\alpha_n^x}, 1)$ for all names n , and so $\text{NORM}(d, \mathcal{W}_{\forall x\alpha}, 1)$. So $e, w \models K\forall x\alpha$. ■

- $e \models K\forall x\alpha \supset \forall x K\alpha$

- $e \models \exists x K\alpha \supset K\exists x\alpha$ but $e \not\models K\exists x\alpha \supset \exists x K\alpha$

The converse of the Barcan does not hold for the existential: knowing that α holds for someone does not imply knowing that individual.

⁴We simply write $\sum d(w_i)$ for $\sum_{i=1}^{\infty} d(w_i)$.

Introspection

For introspection, we prove general versions involving $B(\alpha: r)$ for arbitrary numbers r , and so the usual K -properties (Fagin et al. 1995) come out as special cases.

- $\models B(\alpha: r) \supset KB(\alpha: r)$

Proof: Suppose $e \models B(\alpha: r)$, and so for all $d \in e$, $\text{NORM}(d, \mathcal{W}, 1)$. By assumption, also note that $\mathcal{W}_{B(\alpha: r)} = \mathcal{W}$. (If $e \not\models B(\alpha: r)$, $\mathcal{W}_{B(\alpha: r)}$ would be empty.) Therefore, for every $d \in e$, $\text{NORM}(d, \mathcal{W}_{B(\alpha: r)}, 1)$. So, by definition, $e \models KB(\alpha: r)$. ■

- $\models K\alpha \supset KK\alpha$
- $K_{\text{TRUE}} \models \neg B(\alpha: r) \supset K\neg B(\alpha: r)$

Proof: Suppose $e \models B(\text{TRUE}: 1)$. Then, for every $d \in e$, $\text{NORM}(d, \mathcal{W}, 1)$. Suppose $e \not\models B(\alpha: r)$. Then $\mathcal{W}_{\neg B(\alpha: r)} = \mathcal{W}$. (That is, only from $e \not\models B(\alpha: r)$ it does not follow that for every $d \in e$, $\text{NORM}(d, \mathcal{W}, 1)$ holds.) So, for all $d \in e$, $\text{NORM}(d, \mathcal{W}_{\neg B(\alpha: r)}, 1)$ and thus, by definition, $e \models K\neg B(\alpha: r)$. ■

- $K_{\text{TRUE}} \models \neg K\alpha \supset K\neg K\alpha$

More generally, we have the following regarding *weak S5*:

Theorem 7: *Suppose $\alpha \in \text{OBL}$ is propositional and only mentions the modality K . Then α is valid in weak S5 iff $K_{\text{TRUE}} \supset \alpha$ is valid in OBL .*

4 Extensions

The formula $B(\alpha: r)$ expresses that the probability of α is precisely r . It may be desirable to also allow arbitrary inequalities, which we introduce here by building on our previous discussions. Let OBL be extended with the following modalities: $B\alpha \circ r$, where $\circ \in \{\leq, <, \geq, >\}$, to be read as “ α is believed with a probability $\circ r$.” (Analogous O -based modalities are also possible along the same lines, which we omit here.)

In the semantic rules, we add the following definition:

- $e, w \models B\alpha \circ r$ iff $e, w \models B(\alpha: r')$ for some $r' \circ r$.

Obvious properties then follow from this inequality operators, e.g., regarding equality and inequalities, we have:

- $\models B(\alpha: r) \equiv B\alpha \leq r \wedge B\alpha \geq r$

Proof: Here, $e \models B\alpha \leq r \wedge B\alpha \geq r$ iff (by definition) $e \models B(\alpha: r')$ for some $r' \leq r$ and $e \models B(\alpha: r'')$ for some $r'' \geq r$ iff (by definition) for all $d \in e$, $\text{NORM}(d, \mathcal{W}_\alpha, r')$ and $\text{NORM}(d, \mathcal{W}_\alpha, r'')$, that is, $r' = r''$ (because for any $\mathcal{V} \subseteq \mathcal{W}$, $\text{NORM}(d, \mathcal{V}, r^*)$ can only be true for a unique r^*) and by the linear constraints $r' = r$ iff for all $d \in e$, $\text{NORM}(d, \mathcal{W}_\alpha, r)$ iff (by definition) $e \models B(\alpha: r)$. ■

- $\models B(p: r) \vee B(p: r') \supset Bp \geq \min(r, r')$

which says that if the agent believes p with a probability of r or r' , without being able to say which, it can be said to believe p with a probability of at least $\min(r, r')$.

- $\models Bp \geq r \supset Bp \geq r'$ for every $0 \leq r' \leq r$

where r is the upper bound for the \geq relation.

- $\models Kp \wedge B(q: r) \supset B(p \wedge q) \geq r$

That is, beliefs and knowledge can be combined freely.

A less obvious property in the context of a universally quantified formula is the following:

Proposition 8: *Suppose n is any name, and $r \notin \{0, 1\}$. Then $\not\models B(\forall xP(x): r) \supset B(P(n): r)$ but $\models B(\forall xP(x): r) \supset BP(n) \geq r$.*

Proof: Suppose $e \models B(\forall xP(x): r)$. Suppose w is a world such that $w \models P(n')$ for every name n' , and w' is a world such that $w' \models P(n) \wedge \neg P(n')$ for every name $n' \neq n$. Let d map w to r and w' to $1 - r$, and 0 to all other worlds. Clearly $d \in e$, but $\text{NORM}(d, \mathcal{W}_{P(n)}, r)$ is false, proving the falsifiability of $B(\forall xP(x): r) \supset B(P(n): r)$. In contrast, it is easy to see that $B(\forall xP(x): r) \supset BP(n) \geq r$ is not falsifiable. ■

5 Only Knowing and Probabilities

We take a closer look at the specification of knowledge bases and what they entail, that is, we inspect the O modality. In particular, given a knowledge base Σ that is possibly incomplete or partial, we provide examples as to what the agent can be expected to believe. In what follows, α is an arbitrary formula, ϕ, ψ are objective, and p, q are distinct atoms.

For starters, as in OL (Levesque and Lakemeyer 2001), only knowing implies knowing, but also not believing what does not logically follow from the knowledge base:

- $\models O\alpha \supset K\alpha$
- $\models O\phi \supset K\psi$ iff $\models \phi \supset \psi$
- $\models O(\forall xP(x)) \supset KP(n)$ for any name n
- $\models Op \supset \neg Kq$

Proof: Suppose $e \models Op$ but $e \models Kq$. Then for every $d \in e$, $\text{NORM}(d, \mathcal{W}_q, 1)$. Consider d' that assigns $w \in \mathcal{W}_p$ a value 1, where $w \models p \wedge \neg q$. By definition $d' \in e$, but $\text{NORM}(d', \mathcal{W}_q, 1)$ cannot be true. This is a contradiction. ■

Clearly, the last property is not true in classical epistemic logic if we were to replace O by K ; in general:

- $\models O\phi \supset \neg K\psi$ iff $\not\models \phi \supset \psi$

Similarly, the example involving quantifying-in from OL turns out to be valid also in OBL :

- $\models O(P(a) \vee P(b)) \supset K(\exists x[P(x) \wedge \neg KP(x)])$

More generally, as far as disjunctions are concerned, we get:

- $\models O(p \vee q) \supset \neg K\neg p$

This validity also does not hold in classical epistemic logic if we were to replace O by K .

Nonetheless, $O(p \vee q)$ does not mean that p is believed with a strictly positive probability:

- $\not\models O(p \vee q) \supset Bp > 0$

Proof: Let d assign 1 to some $w \in \mathcal{W}_q$, where $w \models q \wedge \neg p$, and 0 to all other worlds. Suppose $e \models O(p \vee q)$. Clearly $d \in e$, and so $e \not\models Bp > 0$. ■

In a similar vein, consider the case of existentially quantified knowledge bases. It follows that for any name n :

- $\not\models O(\exists xP(x)) \supset B(P(n): 0)$
- $\not\models O(\exists xP(x)) \supset B(P(n)) > 0$

Proof: If $e \models O(\exists xP(x))$, there is a $d \in e$ that assigns a strictly positive probability to a world w where $w \models P(n)$, and so $O(\exists xP(x)) \supset B(P(n): 0)$ cannot be valid. But there is also a $d \in e$ that assigns 0 to any world where $P(n)$ holds, and so $O(\exists xP(x)) \supset B(P(n)) > 0$ is not valid. ■

The examples above demonstrate incomplete specifications, in that the probabilities of atoms are left open. The behavior of only knowing is analogous when we turn to probabilistic knowledge base specifications, in the sense that we obtain appropriate generalizations of the properties identified earlier. For example, only knowing implies believing:

- $\models O(\alpha_1: r_1, \dots, \alpha_k: r_k) \supset \bigwedge B(\alpha_i: r_i)$

In the probabilistic context, we can also speak of the negations of sentences in the knowledge base; for example:

- $\models O(p: r) \supset B(p: r) \wedge B(\neg p: 1 - r)$

Proof: Suppose $e \models O(p: r)$. Then for all $d, d \in e$ iff $\text{NORM}(d, \mathcal{W}_p, r)$, and so $\text{NORM}(d, \mathcal{W}, 1)$, leading to $e \models B(p: r)$. But $\text{NORM}(d, \mathcal{W}_{\neg p}, 1 - r)$; so $e \models B(\neg p: 1 - r)$. ■

In general, if ϕ is objective, then:

- $\models O(\phi: r) \supset B(\phi: r) \wedge B(\neg\phi: 1 - r)$

In case only $p \vee q$ is believed for ground atoms p and q , we obtain, for example:

- $\models O(p \vee q: r) \supset B(\neg p \wedge \neg q: 1 - r)$

Of course, partial and probabilistic specifications can be paired, as in the following example:

- $\models O(p \vee q: 1, p: r) \supset B(q \wedge \neg p: 1 - r)$

Proof: Suppose $e \models O(p \vee q: 1, p: r)$. Then by definition for all $d \in e$, $\text{NORM}(d, \mathcal{W}_{p \vee q}, 1)$ and $\text{NORM}(d, \mathcal{W}_p, r)$. Since $\mathcal{W}_{q \wedge \neg p} = \mathcal{W}_{p \vee q} - \mathcal{W}_p$, we have $\text{NORM}(d, \mathcal{W}_{q \wedge \neg p}, 1 - r)$ for all $d \in e$. Thus, we obtain $e \models B(q \wedge \neg p: 1 - r)$. ■

Let us also reiterate that even in the presence of probabilistic specifications about an atom p , as far as atoms other than p are concerned, these comes out as not being believed:

- $\models O(p: r) \supset \neg Kq$

This need not be the case if O is replaced by B :

- $\not\models B(p: r) \supset \neg Kq$

6 OL is part of $OB\mathcal{L}$

In this section, we show that $OB\mathcal{L}$ is downward compatible with OL . As noted, an epistemic state in OL is simply a set of worlds. Using \models' to denote the satisfaction relation in OL , and using e' to denote epistemic states in OL , the semantic rules differ only in how the modalities are interpreted:

- $e', w \models' K\alpha$ iff for all $w' \in e'$, $w' \models' \alpha$;
- $e', w \models' O\alpha$ iff for all $w', w' \in e'$ iff $e', w' \models' \alpha$.

Satisfaction and validity are defined as usual.

On the relationship between OL and $OB\mathcal{L}$, we obtain:⁵

⁵We need to consider measurable epistemic states because $\models' K_{\text{TRUE}}$ but as observed earlier K_{TRUE} is falsifiable in $OB\mathcal{L}$.

Theorem 9: For any $\alpha \in OL \cap OB\mathcal{L}$, $\models' \alpha$ iff $\models K_{\text{TRUE}} \supset \alpha$.

Its proof rests on two intermediate lemmas:

Lemma 10 : Let e' be any set of worlds and $e = \{d \mid \text{NORM}(d, e', 1)\}$. Then (a) $e \models K_{\text{TRUE}}$ and (b) for all w and $\alpha \in OL \cap OB\mathcal{L}$, $e', w \models' \alpha$ iff $e, w \models \alpha$.

Lemma 11: Let e be any measurable epistemic state, and $e' = \{w' \mid \text{there is a } d \in e \text{ such that } d(w') > 0\}$. Then for all w and $\alpha \in OL \cap OB\mathcal{L}$, $e, w \models \alpha$ iff $e', w \models' \alpha$.

The argument for Theorem 9 is then as follows:

Proof: Suppose $\models K_{\text{TRUE}} \supset \alpha$. Let e' be any set of worlds, $w \in \mathcal{W}$, and $e = \{d \mid \text{NORM}(d, e', 1)\}$. By Lemma 10 (a), $e \models K_{\text{TRUE}}$ and hence $e, w \models \alpha$. By Lemma 10 (b), $e', w \models' \alpha$.

Suppose $\models' \alpha$ and let w be a world and e be any set of distributions such that $e \models K_{\text{TRUE}}$, that is, e is measurable by Theorem 6. Let e' be as in Lemma 11. Then $e', w \models' \alpha$ by assumption and hence, by Lemma 11, $e, w \models \alpha$. ■

7 A Specification in Action

To see the logic in action in a dynamic world, we treat a quantitative variant of the well-known litmus test example (Moore 1985a). In the original version, an agent identifies the acidity of a solution by dipping litmus paper in the solution; the paper turning red implies acidity. In our version, we assume the sensor is noisy, as is usual in robotic applications (Thrun, Burgard, and Fox 2005).

To reason about dynamics, let $OB\mathcal{L}^+$ be $OB\mathcal{L}$ augmented with an *observation modality* $[\alpha_1: r_1, \dots, \alpha_k: r_k]$, where α_i is objective and r_i is a rational, to be read as saying “the agent observes that α_1 has a likelihood r_1, \dots, α_k has likelihood r_k .” (The modality can be seen as a simple kind of *action likelihood axioms* with no preconditions (Bacchus, Halpern, and Levesque 1999); we implicitly assume $\models \alpha_i \supset \neg\alpha_j$ for every $i \neq j$.) For example, after dipping the paper, suppose a indicates that the solution is acidic, then $[a: 2, \neg a: 1]$ is a noisy observation where seeing a is twice as likely.

We define the meaning of $\alpha \in OB\mathcal{L}^+$ inductively as before, with the following new rule:

- $e, w \models [\alpha_1: r_1, \dots, \alpha_k: r_k]\phi$ iff $e', w \models \phi$

where $e' = \{\text{prog}(d) \mid d \in e\}$, and $\text{prog}(d)$ is a distribution d' such that for all $w \in \mathcal{W}$:

- if $w \models \alpha_i$ then $d'(w) = d(w) \times r_i$;
- otherwise, $d'(w) = d(w)$.

We can model the litmus in the presence of probabilistic and incomplete knowledge bases as follows:

- $\models O(a: .6) \supset [a: 2, \neg a: 1]B(a: .75)$

The prior is updated after the (noisy) observation to yield a posterior belief of .75 for the solution being acidic.

- $\models O_{\text{TRUE}} \supset [a: 2, \neg a: 1]\forall x(B(\neg a: x) \supset B(a: 2 \cdot (1 - x)))$

If all that the agent knows are tautologies and it senses, then for any (posterior) probability x accorded to $\neg a$, a probability of $2 \cdot (1 - x)$ is accorded to a .

8 Related Work

There is a fairly extensive literature on reasoning about probabilities (see, for example, (Gaifman 1964; Nilsson 1986; Halpern 2003)), and we briefly survey the major camps below. At the outset, we remark that the focus of our work is only knowing,⁶ which has not been considered for probabilistic specifications (save for one exception; see below).

The inspiration for our work, and perhaps the one closest in spirit, is the work of Bacchus, Halpern and Levesque (BHL) (1999) on degrees of belief in the situation calculus. Although an axiomatic proposal (whose semantic formulations are much in need (Lakemeyer and Levesque 2004)), BHL propose a conceptually attractive definition of belief in a first-order setting. Roughly speaking, the belief in ϕ for BHL is simply the normalized sum of the weights of worlds where ϕ holds. But because this sum is not always well-defined in a first-order setting, they require an axiom stipulating that the belief in true (that is, sum of weights of all worlds) is some number. Our conditions on measurable epistemic states are essentially a modal reworking of this intuition. BHL do not consider only knowing, and we only considered a simple sort of actions.

Also inspired by BHL, Gabaldon and Lakemeyer (Gabaldon and Lakemeyer 2007) consider a logic of only knowing and probability. However, to ensure well-defined sums in the definition of belief, they need to make artificial meta-linguistic assumptions. For example, among other things, they need to assume that the agent only ever encounters a finite number of standard names by means of which they can construct probability distributions on finitely many (equivalence classes of) worlds. Consequently, quantification also ranges over this finite set. Our approach overcomes these problematic limitations in a general manner, especially in handling quantification precisely as in OL . Outside of this, in a game theory context, Halpern and Pass (Halpern and Pass 2009) have considered a (propositional) version of only knowing in probability structures. It would be interesting to see how this latter notion relates to (propositional) OBL .

Reasoning about knowledge and probability has appeared in a number of works prior to BHL, of course, in computer science (Nilsson 1986; Fagin and Halpern 1994), game theory (Monderer and Samet 1989; Heifetz and Mongin 2001), among others (Halpern 2003). Properties discussed in this paper, such as introspection and additivity, are also well studied (Aumann 1999). Notably, the work of Fagin and Halpern (Fagin and Halpern 1994) can be seen to be at the heart of BHL (and our work). The Fagin-Halpern scheme is a general one formulated for Kripke frames (Fagin et al. 1995) where in addition to the worlds an agent considers possible, the agent may consider some worlds more likely than others. In terms of expressivity of the logical

⁶Only knowing is related to analogous notions such as *minimal knowledge* (Halpern and Moses 1984) and *total knowledge* (Pratt-Hartmann 2000). There are significant differences, however. For example, in the proposal of minimal knowledge, the notion of “all I know” is a meta-linguistic notion (and, surprisingly, harder to reason with (Rosati 2000)), and total knowledge requires knowledge to be true. See (Levesque and Lakemeyer 2001) for discussions.

language, they allow sums of linear inequalities, but are propositional. It would not be hard to extend our language to also allow such sums. However, theirs is a multiagent framework, and ours is not (yet). We also consider the simple case where a set of global distributions apply to W as seen at every world; in theirs, the probability spaces can differ arbitrarily across the worlds. The Fagin-Halpern scheme shares some similarity with probabilistic logics for programs (Kozen 1981) and variants thereof (Halpern and Tuttle 1993; Van Benthem, Gerbrandy, and Kooi 2009); see (Fagin and Halpern 1994) for discussions.

There are many previous first-order accounts of probabilities, such as Bacchus (1990) and Halpern (1990); see (Ognjanovic and Raškovic 2000) for a recent list. As discussed by Halpern, probabilities could be accorded to the domain of discourse or on possible worlds (or both), leading to different sorts of properties. Not surprisingly, our B modality shares properties with the possible-worlds version (*e.g.*, additivity). See also (Abadi and Halpern 1994) who show that the set of valid formulas of the logic with probabilities on worlds is not recursively enumerable. (Incidentally, the same goes for OL (Halpern and Lakemeyer 1995).) Recently in AI, limited versions of probabilistic logics have been discussed (Poole 2003; Domingos and Webb 2012), with things like a finite domain assumption built-in.

Finally, as far as the O modality is concerned, we considered a set of probability distributions, which was argued to have properties identical to OL , which is what we desired. There are other ways to interpret $O(p \vee q)$ in a probabilistic context, say using the principle of maximum entropy (Halpern 2003), but this, we feel, differs from the intuitive reading of the logical consequences of $O(p \vee q)$ in OL . There are possibly other ways to realize the sets of probability distributions admitted by $O(p \vee q)$, perhaps by means of Dempster-Shafer belief functions (Shafer 1990) or others (Grove and Halpern 1998). We leave this for the future.

9 Conclusions

In a categorical setting, only knowing admits a perspicuous account of the beliefs and non-beliefs of an incomplete knowledge base. To date, however, how this natural notion can be extended for partial or incomplete probabilistic specifications was left unaddressed. This paper investigates that concern, treating quantification in a general manner, and not only do we show that our proposal OBL is downward compatible with OL in terms of what follows from only knowing, we also show that beliefs emerge as a logical consequence at a corresponding level of specificity.

There are two main directions for the future. As we remarked, our inspiration is the BHL framework, and we would like to investigate a full dynamic version of OBL , as a semantic basis to the axiomatic work of BHL. Second, Levesque showed that OL captures *autoepistemic defaults* (Moore 1985b) when the knowledge base includes beliefs about itself. What these defaults would be like in a probabilistic context, and how they would relate to the family of *statistical defaults* considered in (Bacchus 1990; Halpern 1990) would bring only knowing and the latter ideas closer together.

References

- Abadi, M., and Halpern, J. 1994. Decidability and expressiveness for first-order logics of probability. *Information and Computation* 112(1):1 – 36.
- Aumann, R. J. 1999. Interactive epistemology II: probability. *Int. J. Game Theory* 28(3):301–314.
- Bacchus, F.; Halpern, J. Y.; and Levesque, H. J. 1999. Reasoning about noisy sensors and effectors in the situation calculus. *Artificial Intelligence* 111(1–2):171 – 208.
- Bacchus, F. 1990. *Representing and Reasoning with Probabilistic Knowledge*. MIT Press.
- Belle, V., and Lakemeyer, G. 2010. Multi-agent only-knowing revisited. In *Proc. KR*, 49–60.
- Domingos, P., and Webb, W. A. 2012. A tractable first-order probabilistic logic. In *Proc. AAAI*.
- Fagin, R., and Halpern, J. Y. 1994. Reasoning about knowledge and probability. *J. ACM* 41(2):340–367.
- Fagin, R.; Halpern, J. Y.; Moses, Y.; and Vardi, M. Y. 1995. *Reasoning About Knowledge*. MIT Press.
- Gabalton, A., and Lakemeyer, G. 2007. ESP: A logic of only-knowing, noisy sensing and acting. In *Proc. AAAI*, 974–979.
- Gaifman, H. 1964. Concerning measures in first order calculi. *Israel Journal of Mathematics* 2(1):1–18.
- Grove, A. J., and Halpern, J. Y. 1998. Updating sets of probabilities. In *Proc. UAI*, 173–182.
- Halpern, J. Y., and Lakemeyer, G. 1995. Levesque’s axiomatization of only knowing is incomplete. *Artificial Intelligence* 74(2):381–387.
- Halpern, J., and Lakemeyer, G. 2001. Multi-agent only knowing. *Journal of Logic and Computation* 11(1):251–265.
- Halpern, J. Y., and Moses, Y. 1984. Towards a theory of knowledge and ignorance: Preliminary report. In *Proc. NMR*, 125–143.
- Halpern, J. Y., and Pass, R. 2009. A logical characterization of iterated admissibility. In *Proc. TARK*, 146–155.
- Halpern, J. Y., and Tuttle, M. R. 1993. Knowledge, probability, and adversaries. *J. ACM* 40:917–960.
- Halpern, J. 1990. An analysis of first-order logics of probability. *Artificial Intelligence* 46(3):311–350.
- Halpern, J. Y. 2003. *Reasoning about Uncertainty*. MIT Press.
- Heifetz, A., and Mongin, P. 2001. Probability logic for type spaces. *Games and Economic Behavior* 35(1-2):31 – 53.
- Kaplan, D. 1968. Quantifying in. *Synthese* 19(1):178–214.
- Kozen, D. 1981. Semantics of probabilistic programs. *Journal of Computer and System Sciences* 22(3):328 – 350.
- Lakemeyer, G., and Levesque, H. J. 2004. Situations, Si! situation terms, No! In *Proc. KR*, 516–526.
- Lakemeyer, G., and Levesque, H. J. 2005. Only-knowing: Taking it beyond autoepistemic reasoning. In *AAAI*, 633–638.
- Levesque, H., and Lakemeyer, G. 2001. *The logic of knowledge bases*. The MIT Press.
- Levesque, H. J. 1990. All I know: a study in autoepistemic logic. *Artificial Intelligence* 42(2-3):263–309.
- Monderer, D., and Samet, D. 1989. Approximating common knowledge with common beliefs. *Games and Economic Behavior* 1(2):170–190.
- Moore, R. C. 1985a. A Formal Theory of Knowledge and Action. In *Formal Theories of the Commonsense World*. Norwood, NJ: Ablex. 319–358.
- Moore, R. C. 1985b. Semantical considerations on non-monotonic logic. *Artificial Intelligence* 25(1):75–94.
- Nilsson, N. 1986. Probabilistic logic. *Artificial intelligence* 28(1):71–87.
- Ognjanovic, Z., and Raškovic, M. 2000. Some first-order probability logics. *Theoretical Computer Science* 247(1–2):191 – 212.
- Poole, D. 2003. First-order probabilistic inference. In *Proc. IJCAI*, 985–991.
- Pratt-Hartmann, I. 2000. Total knowledge. In *Proc. AAAI*, 423–428.
- Rosati, R. 2000. On the decidability and complexity of reasoning about only knowing. *Artificial Intelligence* 116(1-2):193–215.
- Shafer, G. 1990. Perspectives on the theory and practice of belief functions. *International Journal of Approximate Reasoning* 4(5–6):323 – 362.
- Thrun, S.; Burgard, W.; and Fox, D. 2005. *Probabilistic Robotics*. MIT Press.
- Van Benthem, J.; Gerbrandy, J.; and Kooi, B. 2009. Dynamic update with probabilities. *Studia Logica* 93(1):67–96.
- Waalder, A., and Solhaug, B. 2005. Semantics for multi-agent only knowing: extended abstract. In *Proc. TARK*, 109–125.