Extending Probabilistic Dynamic Epistemic Logic*

Joshua Sack†

March 25, 2009

Abstract

This paper aims to extend in two directions the probabilistic dynamic epistemic logic provided in Kooi's paper [7] and to relate these extensions to ones made in [10]. Kooi's probabilistic dynamic epistemic logic adds to probabilistic epistemic logic sentences that express consequences of public announcements. The paper [10] extends [7] to using action models, but in both papers, the probabilities are discrete, and are defined on trivial σ-algebras over finite sample spaces. The first extension offered in this paper is to add a previous-time operator to a probabilistic dynamic epistemic logic similar to Kooi's in [7]. The other is to involve non-trivial σ-algebras and continuous probabilities in probabilistic dynamic epistemic logic.

1 Introduction

Probabilistic epistemic logic has been developed to express interaction between both qualitative and quantitative beliefs. This logic lets us formally express statements such as "Bob believes the probability of φ to be at least 1/2" or "Ann considers the probability of ψ to be 1/4". As we are often concerned about how beliefs and probabilities change over time, there have been papers written that mix probability, belief, and time. Examples include [6] and [3], which use probabilistic systems of runs, and [7] and [10], which combine probability with public announcement logic and dynamic epistemic logic (DEL). The approach using probabilistic systems of runs provides a natural way to view time, both past and future, but conditions need to be imposed in order to ensure that agents' probability measures change in a realistic way. Now public announcement logic (PAL), and more generally DEL, provides a mechanical procedure for changes in belief upon receipt of public information, and [7] extends this mechanical procedure to show how a probability measure may change given public information. But DEL has limitations in its ability to express features of the past and future. It cannot express features of the past at all, and

*Many thanks to Larry Moss, Wiebe van der Hoek, Johan van Benthem, and Eric Pacuit.

†The author was partly supported by the project "New Developments in Operational Semantics" (nr. 080039021) of The Icelandic Research Fund and a grant from Reykjavík University's Development Fund.

it can only express future situations as a result of specific actions. By adding temporal logic to DEL in a non-probabilistic setting, the paper [9] captures both some of the temporal flexibility of the system of runs as well as the mechanical method offered by DEL of going from one stage in time to the next.

One goal of this paper, the focus of Section 3, is to involve probability in the combination of temporal logic and DEL, focusing on the inclusion of a previous-time operator and exploring the possibility of completeness. A previous-time operator will give us the ability to specify what time the current state is by counting how many steps into the past we have to go before reaching the initial time. The interaction between the previous-time operator and probability will allow us to express agents' probabilities about past situations as well as probabilities agents have previously assigned to situations. We may even express agents' probabilities and beliefs about when in the past something was true. A previous-time operator will also enable us to express when in the past an agent became more or less certain of a situation holding true. For example, we may express that an agent considers the probability of φ to be 1 now but .5 the time before, thus indicating that it was the previous event that made the agent confident in φ holding. If the truth of φ never changes, we can describe this change in probability as a form of learning, and the previous-time operator lets us keep track of when such learning took place. Example 3.4 provides a formula that expresses what exact time it is as well as formulas that provide interaction between probability and the previous-time operators.

Another goal of this paper, which will be the focus of Section 4, is to involve non-trivial σ-algebras in the probability spaces. A σ-algebra is a set of sets on which a probability measure is defined, and a trivial σ-algebra is just the power set of the sample space. Work on probabilistic dynamic epistemic logic, such as in both [7] and [10], has so far restricted probability measures to finite spaces defined over trivial σ-algebras. As the addition of the dynamic component to probabilistic epistemic logic greatly increases the complexity of the system, shedding the σ-algebra has been a helpful means of simplifying the resulting system and focusing on the new dynamic contribution. But the involvement of unmeasurable sets, that is sets not in the σ-algebra and hence not in the domain of the probability measure, is important if we want to consider infinite sample spaces; and it is necessary if we want a uniform probability distribution over a continuous interval to avoid Vitali sets disrupting the laws of probability. Involving continuous probability distributions would make the probabilistic semantics more mathematically complete, and would allow probabilistic dynamic epistemic logic to be used in modeling more realistic situations. One example may be robotics; although robots make discrete measurements, they are modeled using continuous probabilities.

But even in the finite case, unmeasurable sets can be a convenient way of representing that an agent does not assigning a probability to an event in which there is just not enough known to formulate a probability. Such an event could be another person choosing between two numbers. Section 4.4 gives a discussion about how the use of non-trivial σ-algebras lets us model in a very simple way that an agent does not assign a probability to such an event, by excluding the set corresponding to the event from the σ-algebra, and hence from the domain of the probability measure. This section also makes use of a language defined in Section 4.3, which lets us express that a formula corresponds to an unmeasurable set, and hence has an unknown probability. An alternative typically used for expressing that an agent does not assign a probability to a formula is to express that the agent considers all probability values possible, and this may be more cumbersome than having a set not in the σ-algebra. In addition, we would still have to commit to an actual probability.

2 Probabilistic Epistemic Models and Variations

As the underlying structure for dynamic epistemic logic is the epistemic model, the underlying structure for probabilistic dynamic epistemic logic is ideally the probabilistic epistemic model, which adds probability spaces to an epistemic model.

Definition 2.1 [Probabilistic Epistemic Model] Let Φ be a set of proposition letters, and $\mathbf{I}$ be a set of agents. A probabilistic epistemic model is a tuple $\mathbf{M} = (X, \{\xrightarrow{i}\}_{i\in\mathbf{I}}, \|\cdot\|, \{\mathbf{P}_{i,x}\})$, where

- X is a set of "states" or "possible worlds".
- $\xrightarrow{i} \subseteq X^2$ is a binary relation for each agent $i \in \mathbf{I}$.
- $\|\cdot\| : \Phi \to \mathcal{P}X$ is a valuation function.
- $\mathbf{P}_{i,x}$ is a probability space for each agent i and state x, that is $P_{i,x} = (S_{i,x}, \mathcal{A}_{i,x}, \mu_{i,x})$, where
  - $S_{i,x} \subseteq X$ is a set called the sample space.
  - $\mathcal{A}_{i,x}$ is a σ-algebra over $S_{i,x}$ (that is, a collection of subsets of $S_{i,x}$ that is closed under complements and countable unions). We call the sets in the σ-algebra "measurable sets".
  - $\mu_{i,x} : \mathcal{A}_{i,x} \to [0,1]$ is a probability measure over $S_{i,x}$ (that is, $\mu_{i,x}(S_{i,x}) = 1$ and for each countable collection $A_1, A_2, \ldots$ of pairwise disjoint sets in $\mathcal{A}$, $\mu(\bigcup_{k=1}^{\infty} A_k) = \sum_{k=1}^{\infty} \mu(A_k)$).

$\triangleleft$

The set X is called the *carrier set* of the model. For the binary relations, we interpret $x \xrightarrow{i} y$ to mean that agent i considers y possible from x. Given this reading, we call it an *epistemic relation.* We will at times find it useful to consider structures that replace the probability spaces with a different or less restrictive space. For simplicity, we may want to assume that the set X is finite and the set $\mathcal{A}_{i,x}$ is the power set $\mathcal{P}(S_{i,x})$. We will impose such a restriction throughout Section 3, so as to highlight the contribution of a previous-time operator to the system in a less technically burdened environment.

In general, the logics in this paper involve transforming one model into another, and in Section 3, we will find it useful to let probability measures be transformed into functions that do not necessarily obey all of the properties of a probability measure. In particular, we may allow in Section 3 probability measures to be transformed into the 0 measure, a function that maps every set to 0. Such a transformation will be triggered by an agent receiving information that she assigned 0 probability to. Bayesian updating would result in division by 0. We will discuss a variety of methods for dealing with such updates, and we will employ the 0 measure for its technical convenience. Simply defining the resulting measure to be the 0 measure is an effective way of marking that there was an update upon a set of measure 0.

3 Temporal Probabilistic Public Announcement Logic

In this section, we add a previous-time operator to probabilistic public announcement logic. We will view time like the natural numbers, with an initial time, and discrete steps from one time to the next. We will first go over the semantics of probabilistic public announcement logic that we will use. For simplicity, we will focus on single agent models in this section. Rather than transforming probabilistic epistemic models into other probabilistic epistemic models, we will use a slight modification of these probabilistic epistemic models, called the *simple lax probabilistic epistemic models*, defined as follows:

Definition 3.1 [Simple Lax Probabilistic Epistemic Models] Let Φ be a set of proposition letters. A simple lax probabilistic epistemic model is a tuple $\mathbf{M} = (X, \to, \|\cdot\|, \{\mathbf{P}_x\}_{x\in X})$, where

- X is a finite set of states,
- $\to$ and $\|\cdot\|$ are defined the same way as $\xrightarrow{i}$ and $\|\cdot\|$ are in probabilistic epistemic model (Definition 2.1), and
- $P_x = (S_x, \mu_x)$, where $S_x \subseteq X$ and $\mu_x : \mathcal{P}(S_x) \to [0, 1]$ is either a probability measure or the 0 function.

◁

It is recommended that $S_x \subseteq \{z : x \to z\}$, as every outcome in the sample space is a state the agent considers possible. For technical convenience in Definition 3.2, we will not impose such a restriction. One might assume that the converse of the recommendation should hold too, thus making the agent's sample space S_x equal to the set $\{z : x \to z\}$ of states the agent considers possible, but the example in the beginning of the next section motivates why we prefer not to make this restriction either. The example presents a situation in which an agent does not know enough to assign a probability to everything she considers possible, and although there are different ways of handling the uncertainty about the probability, omitting some states from the sample space is an attractive solution. As sample spaces are defined for each state, the agent may still be uncertain about which sample space is correct.

Public announcement logic is concerned with how an agent revises his/her beliefs given new information, knowing that this information is received by all other agents. Of greatest interest is new information that is consistent with the agent's beliefs, and PAL provides an enlightening mechanical procedure called an *update* for producing a new epistemic model from an old one given the new information. But although the updates provide a reasonable method of revising beliefs upon consistent information, they do not upon inconsistent information. The goal of probabilistic public announcement logic is to provide an update procedure that shows how to produce a new simple lax probabilistic epistemic model from an old one, given information that is not only consistent with the agents' beliefs, but is also given positive probability. The case where the probability of the new information is 0 poses difficulties, and the goal of the definition of such a case is more to provide technical convenience than to offer a realistic result.

Definition 3.2 [Updates] Given a probabilistic epistemic model $\mathbf{M} = (X, \{\rightarrow\}, \|\cdot\|, \{\mathbf{P}_x\})$ and a subset Y of X, the update of $\mathbf{M}$ given Y is written $\mathbf{M} \otimes Y$ and is the model $(X', \{\rightarrow\}', \|\cdot\|', \{\mathbf{P}'_x\}_{x \in X'})$, where

- $X' = Y$
- $x \rightarrow' y$ iff $x, y \in Y$ and $x \rightarrow y$
- $\|p\|' = \|p\| \cap Y$
- If $\mu_x(Y) = 0$, then let $\mathbf{P}'_x(S'_x, \mu'_x)$ be defined by:
  - $S'_x = S_x \cap Y$
  - $\mu'_x(X) = 0$ for every $X \subseteq S'_x$

  Otherwise, let $\mathbf{P}'_x$ be defined by
  - $S'_x = S_x \cap Y$
  - For each subset $Z \subseteq Y$, $\mu'_x(Z) = \mu_x(Z)/\mu_x(Y)$

$\triangleleft$

This definition differs from the one in [7] in that here updating removes states while in [7] it removes relational connections but not states. Probability is updated in the same way as long as the set Y has positive probability.

There are a number of choices for how a measure should be updated upon information that has probability 0. The one used here is most similar to one by [10]. The method used in [7] allows the probabilities to remain the same, after updating, and has a number of advantages, such as ensuring that the updated measure is still a probability measure and the new measure retains information about the old measure; but the definition of this update requires states not to be eliminated. Another method could be to have a default probability measure defined to give the actual state probability 1. This method would have the advantage of ensuring that the updated measure is a probability measure, and will also be compatible with the elimination of states, but a disadvantage is that there is no way to reconstruct any information about the previous measure, thus making it nearly impossible to have axioms relating probability with past. So with the method we have adopted, we will at least know when there is a major disconnect between the current measure and some previous measure. A possibly desirable definition for updating a probability measure upon a set of probability 0 may be to give the set upon which we are updating a probability 1, but make all subsets of this set unmeasurable. We do not adopt this method here, for we wish to benefit from the simplicity of avoiding non-measurable sets. But we will use it in the next section, as we will focus on the contribution of using non-measurable sets.

A natural choice for semantics that includes a previous-time operator for this language is to involve structures that consist of a list of all past and present models. This is what was done in [9].

Definition 3.3 [History] A history H is a list of models $(\mathbf{M}_0, \mathbf{M}_1, \ldots, \mathbf{M}_n)$, where for each k, $\mathbf{M}_k = (X_k, \{\xrightarrow{i}_k\}, \|\cdot\|_k, \{\mathbf{P}_{kx}\})$, and $M_{k+1} = M_k \otimes X_{k+1}$. ◁

Given a history $H = (\mathbf{M}_0, \mathbf{M}_1, \ldots, \mathbf{M}_n)$, let $\widetilde{P}(H) = (\mathbf{M}_0, \mathbf{M}_1, \ldots \mathbf{M}_{n-1})$ be the previous history and $\widetilde{\mathbf{M}}(H) = \mathbf{M}_n$ be the last (most recent) model in the list, and let $\widetilde{X}(H) = X_n$. We may write $x \in H$ for $x \in \widetilde{X}(H)$. To make the function $\widetilde{P}$ total, define $\widetilde{P}((\mathbf{M})) = \emptyset$.

Language

Let Φ be a set of proposition letters. We define by mutual recursion a multi-sorted language $\mathcal{L}$ with sentences and terms for each agent. The sentences (also called formulas) are given by

$$\varphi ::= \mathsf{true} \,|\, p \,|\, \neg\varphi \,|\, \varphi_1 \wedge \varphi_2 \,|\, \Box\varphi \,|\, [\varphi_1]\varphi_2 \,|\, t \geq q \,|\, \bar{Y}\varphi$$

where t is a term and $p \in \Phi$.

The terms are given by

$$t ::= qP(\varphi) \,|\, t_1 + t_2$$

where $q \in \mathbb{Q}$ is a rational number and φ is a sentence.

We could have defined a simpler language that does not have the symbol $+$, a language where the probability terms are much simpler. But we will benefit from the ability to add probabilities when expressing axioms.

The semantics is defined by a function $[\![\cdot]\!]$ from formulas to functions f that map each history H to a subset of $\widetilde{X}(H)$, the carrier set of the most recent model in H. Then

$$
\begin{array}{lll}
x \in [\![\mathsf{true}]\!](H) & \text{iff} & x \in \widetilde{X}(H) \\
x \in [\![p]\!](H) & \text{iff} & x \in \|p\|_{\widetilde{\mathbf{M}}(H)} \\
x \in [\![\neg\varphi]\!](H) & \text{iff} & x \in [\![\mathsf{true}]\!](H) - [\![\varphi]\!](H) \\
x \in [\![\varphi \wedge \psi]\!](H) & \text{iff} & x \in [\![\varphi]\!](H) \cap [\![\psi]\!](H) \\
x \in [\![\Box\varphi]\!](H) & \text{iff} & z \in [\![\varphi]\!](H) \text{ whenever } x \to_{\widetilde{\mathbf{M}}(H)} z \\
x \in [\![[\psi]\varphi]\!](H) & \text{iff} & x \in [\![\varphi]\!](H \otimes [\![\psi]\!](H)) \text{ whenever } s \in [\![\psi]\!](H) \\
x \in [\![t \geq q]\!](H) & \text{iff} & \left\{ \begin{array}{l} \sum_{j=1}^{n} q_j \mu_x([\![\varphi_j]\!](H) \cap S_x) \geq q \\ \text{whenever } t = q_1 P(\varphi_1) + \cdots + q_n P(\varphi_n) \end{array} \right. \\
x \in [\![\bar{Y}\varphi]\!](H) & \text{iff} & x \in [\![\varphi]\!](\widetilde{P}(H)) \text{ whenever } \widetilde{P}(H) \neq \emptyset
\end{array}
$$

We have the usual modal abbreviations, such as $\Diamond\varphi \equiv \neg\Box\neg\varphi$ and $\langle\psi\rangle\varphi \equiv \neg[\psi]\neg\varphi$, and we let $\widehat{Y}\varphi \equiv \neg\bar{Y}\neg\varphi$, which asserts that there is a previous time and φ is true then. Here are some abbreviations in the language that express a variety of inequalities and equality.

- $t \leq q \equiv -t \geq -q$
- $t < q \equiv \neg(t \geq q)$
- $t > q \equiv \neg(t \leq q)$
- $t = q \equiv t \leq q \wedge t \geq q$
- $t \geq s \equiv t - s \geq 0$
- $t = s \equiv t - s \geq 0 \wedge s - t \geq 0$

Example 3.4 (Formulas involving past) Here are some examples of (non-primitive) formulas that we can express:

$\bar{Y}\bar{Y}\neg\,\mathsf{true} \wedge \widehat{Y}\,\mathsf{true}$, which states that now is the second stage in time since the beginning of the scenario being modeled.

$P(\widehat{Y}\widehat{Y}\varphi) \geq P(\widehat{Y}\varphi)$, which states that the agent considers the probability of φ holding true two stages ago to be greater than the probability of φ to be true one stage ago. These probabilities are given by the agent's current probability measure.

$P(q) = 1 \wedge \widehat{Y}(P(q) = .5)$, which states that the agent now is certain of atomic proposition q holding true, but one step ago only gave q a 50% chance of holding. Note that atomic propositions do not change their truth value over time.

$P(q) = 1 \wedge \widehat{Y}[\psi](P(q) = .5)$, which says that the agent assigns 1 to the probability of q, but had ψ been announced before, then the agent would have assigned .5 to the probability of q

$\widehat{Y}\varphi \to \Box\widehat{Y}\varphi \wedge P(\widehat{Y}\varphi) = 1$, which states that if φ was just true before, then the agent knows that it was just true and assigns 1 to the probability that it was just true. We can view this as omniscience about the previous stage in time.

$\widehat{Y}(P(\varphi) = a) \to P(\widehat{(Y)}\varphi) = a$, which states that if the agent assigned a to the probability of φ, then the agent will remember that he assigned a to the probability of φ. This turns out to be a special case of the *probability yesterday 1 axiom*, which will be define next.

Proof system

Include axioms of proposition logic together with those in Tables 1 and 2.

$\Box$-normality	$\Box(\varphi \to \psi) \to (\Box\varphi \to \Box\psi)$
$[\varphi]$-normality	$[\varphi](\psi_1 \to \psi_2) \to ([\varphi]\psi_1 \to [\varphi]\psi_2)$
$\bar{Y}$-normality	$\bar{Y}(\varphi \to \psi) \to (\bar{Y}\varphi \to \bar{Y}\psi)$
Update partial functionality	$[\varphi]\neg\psi \leftrightarrow (\varphi \to \neg[\varphi]\psi)$
$\bar{Y}$-partial functionality	$\widehat{Y}\psi \leftrightarrow (\widehat{Y}\,\mathsf{true} \to \bar{Y}\psi)$
Future atomic permanence	$(\varphi \to p) \leftrightarrow [\varphi]p$
Past atomic permanence	$\bar{Y}p \leftrightarrow (\widehat{Y}\,\mathsf{true} \to p)$
Update yesterday	$[\varphi]\bar{Y}\psi \leftrightarrow (\varphi \to \psi)$
Probability yesterday 0	$\widehat{Y}(\sum_{k=1}^{n} q_k P(\varphi_k) = 0) \to (\sum_{k=1}^{n} q_k P(\widehat{Y}\varphi_k) = 0)$
Probability yesterday 1	$(\widehat{Y}(\sum_{k=1}^{n} q_k P(\varphi_k) = \sum_{k=1}^{n} q_k P(\mathsf{true}))$ $\to (\sum_{k=1}^{n} q_k P(\widehat{Y}\varphi_k) = \sum_{k=1}^{n} q_k P(\mathsf{true}))$
Epistemic-yesterday mix	$\bar{Y}\Box\varphi \to \Box\bar{Y}\varphi$
Epistemic update	$[\varphi]\Box_i\psi \leftrightarrow (\varphi \to \Box_i[\varphi]\psi)$
Probability update	$P(\varphi) > 0 \to ([\varphi]\sum_{k=1}^{n} q_k P(\varphi_k) \geq q$ $\leftrightarrow (\varphi \to \sum_{k=1}^{n} q_k P(\varphi \wedge [\varphi]\varphi_k) \geq qP(\varphi)))$
Probability 0 update	$P(\varphi) = 0 \to ([\varphi](\sum_{k=1}^{n} q_k P(\varphi_k)) \geq q$ $\leftrightarrow (\varphi \to \sum_{k=1}^{n} q_k P(\mathsf{false}) \geq q)$
Non-initial time	$\widehat{Y}\,\mathsf{true} \to \Box\widehat{Y}\,\mathsf{true} \wedge P(\widehat{Y}\,\mathsf{true}) = 1$
Initial time	$\bar{Y}\,\mathsf{false} \to \Box\bar{Y}\,\mathsf{false} \wedge P(\bar{Y}\,\mathsf{false}) = 1$

Table 1: Axioms for public announcement logic and interaction with Probability

Include rules in Table 3.

3.1 Soundness

Theorem 3.5 (Soundness) *If $\vdash \varphi$, then $[\![\varphi]\!] = [\![\mathsf{true}]\!]$, for all $\varphi \in \mathcal{L}$.*

Proof. The soundness of many axioms follow arguments given in other papers. Soundness for the axioms used in dynamic epistemic logic can be found in [1], those that add to dynamic epistemic logic a previous time operator $\bar{Y}$ can be found in [8], and the probabilistic epistemic logic axioms can be found in [3].

0 terms	
	$\sum_{k=1}^{n} q_k P(\varphi_k) \geq q$
	$\leftrightarrow (\sum_{k=1}^{n} q_k P(\varphi_k)) + 0P(\varphi_{k+1}) \geq q$
Permutation	
	$\sum_{k=1}^{n} q_k P(\varphi_k) \geq q \rightarrow \sum_{k=1}^{n} q_{j_k} P(\varphi_{j_k}) \geq q$
	where$j_1, \ldots, j_n$ is a permutation of $1, \ldots, n$
Addition	
	$\sum_{k=1}^{n} q_k P(\varphi_k) \geq q \wedge \sum_{k=1}^{n} q'_k P(\varphi_k) \geq q'$
	$\rightarrow \sum_{k=1}^{n} (q_k + q'_k) P(\varphi_k) \geq (q + q')$
Multiplication	
	$(\sum_{k=1}^{n} q_k P(\varphi_k) \geq q)$
	$\leftrightarrow (\sum_{k=1}^{n} dq_k P(\varphi_k) \geq dq)$ where $d > 0$
Dichotomy	$(t \geq q) \vee (t \leq q)$
Monotonicity	$(t \geq q) \rightarrow (t > q')$ where $q > q'$
Nonnegativity	$P(\varphi) \geq 0$
Probability of truth	
$(P(\mathsf{true}) = 1 \vee P(\mathsf{true}) = 0)) \wedge (\bar{Y}\,\mathsf{false} \rightarrow P(\mathsf{true}) = 1)$	
Additivity	$P(\varphi \wedge \psi) + P(\varphi \wedge \neg\psi) = P(\varphi)$

Table 2: Axioms for Probability

$\Box_i$-necessitation	From $\vdash \varphi$ infer $\vdash \Box\varphi$
$[\varphi]$-necessitation	From $\vdash \varphi$ infer $\vdash [\psi]\varphi$
$\bar{Y}$-necessitation	From $\vdash \varphi$ infer $\vdash \bar{Y}\varphi$
Equivalence	
From $\vdash \varphi \leftrightarrow \psi,$ infer $\vdash P(\varphi) = P(\psi)$	

Table 3: Rules

Of the probability axioms, only the *probability of truth* axiom is significantly different, and its soundness follows from the fact that our initial model is a probabilistic epistemic model, and the ones that follow allow for 0 measures. Slightly different axioms than the ones used here for showing the interaction between updates and probability are given in [7]. As our axioms and semantics differ significantly enough from those, we will address the soundness of them here, as well as the soundness of axioms for the interaction between probability and past.

Probability-yesterday 0 $\widehat{Y}(\sum_{k=1}^{n} q_k P(\varphi_k) = 0) \rightarrow (\sum_{k=1}^{n} q_k P(\widehat{Y}\varphi_k) = 0)$: Suppose $s \in [\![\widehat{Y}(\sum_{k=1}^{n} q_k P(\varphi_k) = 0)]\!](H)$. Then H is not an initial history and $s \in [\![\sum_{k=1}^{n} q_k P(\varphi_k) = 0]\!](\widetilde{P}(H))$. Then using propositional logic, the definition of the abbreviation of $t = 0$, and basic knowledge of inequalities, we have that $\sum_{k=1}^{n} q_k \mu'_s([\![\varphi_k]\!](\widetilde{P}(H))) = 0$, where μ'_s is the measure at s in $\widetilde{P}(H)$. Let μ_s be the measure at s in H. Either $\mu_s(S_s) = 1$ and was obtained from

μ'_s through Bayesian conditioning or $\mu_s(S_s) = 0$. Either way, we have that $\sum_{k=1}^n q_k \mu_s([\![\widehat{Y}\varphi_k]\!](H)) = 0$, whence $s \in [\![\sum_{k=1}^n q_k P(\widehat{Y}\varphi_k) = 0]\!](H)$.

Probability-yesterday 1 $\widehat{Y}(\sum_{k=1}^n q_k P(\varphi_k) = \sum_{k=1}^n q_k P(\mathsf{true}))$ $\rightarrow (\sum_{k=1}^n q_k P(\widehat{Y}\varphi_k) = \sum_{k=1}^n q_k P(\mathsf{true}))$: Suppose $s \in [\![\widehat{Y}(\sum_{k=1}^n q_k P(\varphi_k) = \sum_{k=1}^n q_k P(\mathsf{true}))]\!](H)$. Using similar reasoning as that given for the soundness for probability-yesterday 0, we obtain $s \in [\![\sum_{k=1}^n q_k P(\widehat{Y}\varphi_k) = \sum_{k=1}^n q_k P(\widehat{Y}\,\mathsf{true}))]\!](H)$. But as H is not an initial history, $\widehat{Y}\,\mathsf{true}$ and true are satisfied in the exact same states in H. Thus we obtain our desired consequent.

Probability update $P(\varphi) > 0 \rightarrow ([\varphi]\sum_{k=1}^n q_k P(\varphi_k) \geq q$ $\leftrightarrow (\varphi \rightarrow \sum_{k=1}^n q_k P(\varphi \wedge [\varphi]\varphi_k) \geq q P(\varphi)))$: Suppose $s \in [\![P(\varphi) > 0]\!](H)$. Then the following are equivalent (we use μ_s to be a measure in H and μ'_s to be a measure in $H \otimes [\![\varphi]\!](H)$):

1. $s \in [\![[\varphi]\sum_{k=1}^n q_k P(\varphi_k) \geq q]\!](H)$
2. if $s \in [\![\varphi]\!](H)$, then $s \in [\![[\sum_{k=1}^n q_k P(\varphi_k) \geq q]\!](H \otimes [\![\varphi]\!](H))$
3. if $s \in [\![\varphi]\!](H)$, then $\sum_{k=1}^n q_k \mu'_s([\![\varphi_k]\!](H \otimes [\![\varphi]\!](H))) \geq q$
4. if $s \in [\![\varphi]\!](H)$, then $\sum_{k=1}^n q_k \mu_s q_k([\![\varphi \wedge [\varphi]\varphi_k]\!](H))/\mu_s q_k([\![\varphi]\!](H)) \geq q$
5. if $s \in [\![\varphi]\!](H)$, then $\sum_{k=1}^n q_k \mu_s q_k([\![\varphi \wedge [\varphi]\varphi_k]\!](H)) \geq q\mu_s([\![\varphi]\!](H))$
6. $s \in [\![\varphi \rightarrow \sum_{k=1}^n q_k P(\varphi \wedge [\varphi]\varphi_k) \geq q P(\varphi))]\!](H)$

Probability 0 update $P(\varphi) = 0 \rightarrow ([\varphi](\sum_{k=1}^n q_k P(\varphi_k)) \geq q$ $\leftrightarrow (\varphi \rightarrow \sum_{k=1}^n q_k P(\mathsf{false}) \geq q)$: Suppose $s \in [\![P(\varphi) = 0]\!](H)$. Then the following are equivalent (we use μ_s to be a measure in H and μ'_s to be a measure in $H \otimes [\![\varphi]\!](H)$, and note here that μ'_s is a trivial measure, assigning probability 1 to s):

1. $s \in [\![[\varphi](\sum_{k=1}^n q_k P(\varphi_k)) \geq q]\!](H)$
2. if $s \in [\![\varphi]\!]$, then $s \in [\![\sum_{k=1}^n q_k P(\varphi_k)) \geq q]\!](H \otimes [\![\varphi]\!](H))$
3. if $s \in [\![\varphi]\!]$, then $\sum_{k=1}^n q_k \mu'_s([\![\varphi_k]\!](H \otimes [\![\varphi]\!](H))) \geq q$
4. if $s \in [\![\varphi]\!]$, then $0 \geq q$
5. if $s \in [\![\varphi]\!]$, then $\sum_{k=1}^n q_k \mu_s([\![\mathsf{false}]\!](H)) \geq q$
6. $s \in [\![\varphi \rightarrow \sum_{k=1}^n q_k P(\mathsf{false}) \geq q]\!](H)$

QED

The proofs for the soundness of the rules are standard propositional and modal arguments.

3.2 Derived Formulas and Rules

Definition 3.6 [Provable Equivalence] We say two formulas are provably equivalent, and write $\varphi \equiv \psi$, if $\vdash \varphi \leftrightarrow \psi$. $\triangleleft$

Proposition 3.7 $\vdash [\varphi](\psi \to \chi) \leftrightarrow ([\varphi]\psi \to [\varphi]\chi)$ *and* $\vdash \bar{Y}(\psi \to \chi) \leftrightarrow (\bar{Y}\psi \to \bar{Y}\chi)$

Proof. Of course one direction of each of these is the normality axiom scheme for $[\varphi]$ and $\bar{Y}$ respectively. But the other direction comes from the partial functionality axioms together with basic modal logic. QED

Proposition 3.8 $\vdash [\varphi]\sum_{k=1}^{n} q_k P(\psi_k) \geq q \leftrightarrow$
$(-P(\varphi) \geq 0 \wedge (\varphi \to \sum_{k=1}^{n} q_k P(\mathsf{false}) \geq q))) \vee$
$(\neg(-P(\varphi) \geq 0) \wedge (\varphi \to \sum_{k=1}^{n} q_k P(\varphi \wedge [\varphi]\psi_k) \geq qP(\varphi)))$

Proof. This can be proved from the probability update axiom and probability 0 update axiom using propositional logic alone. QED

One may ask whether we can replace the probability update and probability 0 update axiom schema with Proposition 3.8. It turns out that the axiom schema guarantees the following that might not be guaranteed by Proposition 3.8: $[\varphi]\sum_{k=1}^{n} q_k P(\psi_k) \wedge (\varphi \to \sum_{k=1}^{n} q_k P(\mathsf{false}) \geq q)) \wedge (\varphi \to \sum_{k=1}^{n} q_k P(\varphi \wedge [\varphi]\psi_k) \geq qP(\varphi))$.

Proposition 3.9 *From* $\vdash \varphi_k \leftrightarrow \psi_k$ *for* $k = 1, \ldots, n$, *infer*
$\vdash (\sum_{k=1}^{n} q_k P(\varphi_k) \geq q) \leftrightarrow (\sum_{k=1}^{n} q_k P(\psi_k) \geq q)$

Proof. We will prove the left to right direction. The converse is entirely symmetric. From the rule equivalence, we have for each j that $\vdash P(\varphi_j) = P(\psi_j)$ and in particular $\vdash P(\psi_j) - P(\varphi_j) \geq 0$. From the multiplication axiom and modus ponens, we have for each j that $\vdash q_j P(\psi_j) - q_j P(\varphi_j) \geq 0$. We then apply modus ponens to these and several applications of the zero terms axiom (to produce terms with coefficient 0), then several applications of the permutation axiom, and finally the addition axiom to obtain,

$$\vdash \sum_{j=1}^{n} q_j P(\psi_j) - q_j P(\varphi_j) \geq 0. \tag{1}$$

Finally we have the following sequence, where each provably implies the next:

(a) $\sum_{j=1}^{n} q_j P(\varphi_j) \geq q$

(b) $\sum_{j=1}^{n} 0P(\psi_j) + q_j P(\varphi_j) \geq q$

(c) $\sum_{j=1}^{n} q_j P(\psi_j) + 0P(\varphi_j) \geq q$

(d) $\sum_{j=1}^{n} q_j P(\psi_j) \geq q$

The implication from (a) to (b) is provable from the multiple applications of the zero term and permutation axioms. The implication from (b) to (c) is provable from the addition axiom (adding formula (b) to the provable formula 1). The implication from (c) to (d) is provable using multiple applications of the permutation and zero term axioms. QED

4 Involving Non-Trivial σ-Algebras

In this section, we relax the restriction that every σ-algebra be trivial.

4.1 Measure Theoretic Background

The following measure theoretic background will be useful in defining our updating framework. The book [5] provides a similar discussion, but in a more general setting that does not involve probability; the probabilistic discussion that follows just involves measures for which the measure of the whole set is normalized to 1.

Definition 4.1 [Outer Probability Measure] A function $\mu : \mathcal{P}(S) \to [0, 1]$ is called an outer probability measure if the following conditions hold:

1. $\mu(\emptyset) = 0$ and $\mu(S) = 1$
2. $\mu(A) \leq \mu(B)$ whenever $A \subseteq B$ (monotonicity)
3. $\mu\left(\bigcup_{j=1}^{\infty} A_j\right) \leq \sum_{j=1}^{\infty} \mu(A_j)$ (countable sub-additivity)

$\triangleleft$

Given a probability space $(S, \mathcal{A}, \mu)$, the measure μ can be extended to an outer measure μ^* on $\mathcal{P}(S)$ by

$$\mu^*(T) = \inf\{\mu(A) : A \in \mathcal{A}, T \subseteq A\}$$

for each $T \subseteq S$.

But the relationship between probability measures and outer measures can go the other way, for given an outer probability measure $\mu : \mathcal{P}(X) \to [0, 1]$, one can construct a probability space by restricting μ to the set $\mathcal{A}(\mu)$, the "measureable sets of outer measure μ", where we define

$$\mathcal{A}(\mu) = \{A : \mu(B) = \mu(B \cap A) + \mu(B - A) \text{ for all } B \subseteq S\}$$

Definition 4.2 [Product Space] Given probability spaces $\mathcal{P}_1 = (S_1, \mathcal{A}_1, \mu_1)$ and $\mathcal{P}_2 = (S_2, \mathcal{A}_2, \mu_2)$, we define the product measure written $\mathcal{P}_1 \times \mathcal{P}_2$ to be the space $(S, \mathcal{A}, \mu)$, where

- $S = S_1 \times S_2$ is the Cartesian product of S_1 and S_2
- $\mathcal{A}$ is the smallest σ-algebra containing $\{A_1 \times A_2 : A_1 \in \mathcal{A}_1, A_2 \in \mathcal{A}_2\}$

- $\mu(A) = \inf\{\sum_{j=1}^{\infty} \mu_1(A_1^i)\mu_2(A_2^i) : A_j^i \in \mathcal{A}_j, j = 1, 2 \text{ and } A \subset \cup(A_1^i \times A_2^i)\}$

◁

It turns out that the product probability space is a probability space.

Remark 4.3 *We use the word "measurability" in two ways: with respect to a probability space and with respect to an outer probability measure. With respect to a probability space, the set of measurable sets is just the σ-algebra in the definition of the space. With respect to a outer probability measure μ, the set of measurable sets is the set $\mathcal{A}(\mu)$.*

4.2 Action Models and Update Products

The paper [10] involves (in probabilistic dynamic epistemic logic) action models and three probabilities: prior, occurrence, and observation probabilities. These probabilities, however, are discrete with finite sample spaces and trivial or unspecified σ-algebras. The prior probabilities are the ones in the probabilistic epistemic models, and they represent the initial probabilities agents have. The occurrence probabilities provide objective likelihoods for certain actions to occur in each of a collection of situations. For example, an action could be a positive test result to an illness; this action may be very likely (though not necessarily certain) to occur if the patient being tested has the illness. The observation probabilities provide the likelihoods that each agent assigns to the event that just occurred. For example, the agent may have been told something through an announcement, but is so hard of hearing, he is not sure which of several possibilities the announcement was. The observation probabilities may be particularly important in robotics, where sensors may not be precise.

This paper will define a semantics that involves prior probabilities and occurrence probabilities that are not necessarily discrete, but that can also be continuous. It will, however, not involve occurrence probabilities, leaving for future work the involvement of non-trivial σ-algebras in occurrence probabilities.

Definition 4.4 [Action Model] An *action model* $(\Sigma, \{\stackrel{i}{\to}\}, \{\mathcal{P}_{i,\sigma}\}, \mathsf{pre})$ is a probabilistic epistemic model with the valuation function $\|\cdot\|$ replaced by a function pre which assigns to each $\sigma \in \Sigma$ a function that assigns to each probabilistic epistemic model a subset of the carrier set of that model. Each element $\sigma \in \Sigma$ is called an *action type.* ◁

We define the update product between a probabilistic epistemic model and an action model in two stages. We first define the product between the original probabilistic epistemic model and an action frame (which is the action model without the function pre), and then relativize the result according to the pre function. The first product is called the unrestricted product. The second is called the relativization.

Definition 4.5 [Unrestricted Product] The unrestricted product between a probabilistic epistemic model $\mathbf{M}$ and an action model $\mathbf{\Sigma}$ is $\mathbf{M} \otimes_U \mathbf{\Sigma}$ with the following components:

1. $X_\otimes = X \times \Sigma$
2. $(x, \sigma) \stackrel{i}{\hookrightarrow} (z, \tau)$ iff $x \stackrel{i}{\hookrightarrow} z$ and $\sigma \stackrel{i}{\hookrightarrow} t$
3. $\|p\|_\otimes = \|p\| \times \Sigma$
4. We define $\mathcal{P}_{i,(x,\sigma)}$ to be $\mathcal{P}_{i,x} \times \mathcal{P}_{i,\sigma}$, the product space

◁

As the product probability measure is itself a probability measure, this product is a probabilistic epistemic model.

Definition 4.6 [Relativization] The relativization of a probabilistic epistemic model $\mathbf{M}$ to $Y \subseteq X$ is given by $\mathbf{M} \otimes_R Y$ with the following components:

1. $X_Y = Y$
2. $x \stackrel{i}{\hookrightarrow}_Y z$ iff $x \stackrel{i}{\hookrightarrow} z$ and $x, z \in Y$
3. $\|p\|_Y = \|p\| \cap Y$
4. For $x \in Y$, if $\mu^*_{i,x}(Y) = 0$, then define $\mathcal{P}_{i,x}$ to be the only probability space on the σ-algebra $\{\emptyset, Y\}$. Otherwise let $\widehat{\mu}_{i,x} : \mathcal{P}(Y) \to [0,1]$ be the outer measure defined by
$$\widehat{\mu_Y}_{i,x}(B) = \frac{\mu^*_{i,x}(B)}{\mu^*_{i,x}(Y)}$$
for each $B \subseteq Y$. Then let
   (a) ${S_Y}_{i,x} = S_{i,x} \cap Y$
   (b) ${\mathcal{A}_Y}_{i,x} = \mathcal{A}(\widehat{\mu_Y}_{i,x}) \cap \{A \cap Y : A \in \mathcal{A}_{i,x}\}$
   (c) ${\mu_Y}_{i,x}$ is the restriction of $\widehat{\mu_Y}_{i,x}$ to ${\mathcal{A}_Y}_{i,x}$

◁

There are a couple of points to make about this definition. The first is with updating by a set with probability 0. This has the advantage that the updated function remains a probability function. But it is debatable whether it is desirable for the agent to give up probabilities of everything not characterizable as the set Y. The second point is with the updated σ-algebra. It is defined to be the intersection of two σ-algebras, one of which is generated by the outer measure and the other as a relativized version of the previous σ-algebra (that is the collection of all sets in the previous σ-algebra intersected by the relativizing set). One can easily verify that the intersection of two σ-algebras is itself a σ-algebra. From measure theory, we know that the outer measure satisfies the countable additivity condition on the set of its measurable sets, and hence it is a measure when restricted to those sets. Furthermore, an outer measure is a measure on any σ-algebra contained in the σ-algebra of its measurable sets. We choose to restrict the outer measure beyond its measurable sets to the relativized σ-algebra, as this ensures that agents remain uncertain about probabilities. Thus the first σ-algebra is to sufficiently guarantee that the restricted outer probability measure is a probability measure, and the second is to capture our use of sets not in a σ-algebra as indicating uncertainty about probabilities.

Interestingly, when the prior σ-algebra is finite, then the relativized σ-algebra will suffice as the domain of a probability measure. When the σ-algebra $\mathcal{A}$ of a space $(S, \mathcal{A}, \mu)$ is finite, the outer measure of μ applied to a set $T \subseteq S$ becomes

$$\mu^*(T) = \bigcap\{\mu(A) : A \in \mathcal{A}, T \subseteq A\} = \mu(\bigcap\{A : A \in \mathcal{A}, T \subseteq A\})$$

Thus the outer measure of a (not necessarily measurable) set is equal to the measure of an appropriate measurable set. This is not guaranteed in the infinite case. But this property helps us guarantee that the updated function is indeed a measure. The most difficult case is the additivity condition. If $A_1, \ldots, A_n$ is a set of pairwise disjoint sets measurable in the relativized model, let $\hat{A}_i = \bigcap\{B : A_i \subseteq B, B \in \mathcal{A}_{i,x}\}$, where $\mathcal{A}_{i,x}$ is the σ-algebra for the original model. Unlike the A_i, the $\hat{A}_i$ are necessarily measurable in the first space. Also $\hat{A}_1, \ldots, \hat{A}_n$ is pairwise disjoint, for if Y is the set with which we relativized, then $B = \hat{A}_j \cap \hat{A}_k \subseteq \bar{Y}$ (otherwise A_j and A_k would not be disjoint). But $A_j \subseteq \hat{A}_j - B$ and $\hat{A}_j - B \in \mathcal{A}_{i,x}$, thus $\hat{A}_j = \hat{A}_j - B$, and so we conclude that $B = \emptyset$. Also observe that $\widehat{\bigcup A_i} = \bigcup \hat{A}_i$. We can then make use of this and the fact that $\mu^*(C) = \mu(\hat{C})$ for any set C in order to establish the additivity property of the new measure.

Definition 4.7 [Update Product] Let $\mathbf{\Sigma} = (\Sigma, \{\xrightarrow{i}\}_{i \in \mathbf{I}}, \{\mathcal{P}_{i,x}\}, \mathsf{pre})$ be an action model and $\mathbf{M} = (X, \{\xrightarrow{i}\}, \| \cdot \|, \{\mathcal{P}_{i,x}\})$. Let $Y = \{(x, \sigma) : x \in \mathsf{pre}(\sigma)(\mathbf{M})\}$. The update product between $\mathbf{M}$ and $\mathbf{\Sigma}$ is written $\mathcal{M} \otimes \mathbf{\Sigma}$ and is defined as $(\mathbf{M} \otimes_U \mathbf{\Sigma}) \otimes_R Y$. ◁

4.3 Language

Let Φ be a set of atomic propositions and $\mathbf{I}$ a set of agents. Define a language with sentences given by:

$$\varphi ::= \mathsf{true} \mid p \mid \neg\varphi \mid \varphi_1 \wedge \varphi_2 \mid \Box_i \varphi \mid [\mathbf{\Sigma}, \sigma]\varphi \mid t_i \leq a \mid \mathsf{meas}_i(\varphi)$$

where $p \in \Phi$ is an atomic proposition, $i \in \mathbf{I}$ is an agent, $(\mathbf{\Sigma}, \sigma)$ is a pointed action model[1] (an action model together with an action point in it), t is a term and $a \in \mathbb{R}$ is a real number.

The terms are given by

$$t_i ::= aP_i(\varphi) \mid t_j + t_k$$

[1]The language represents action models by using a primitive symbol for each model.

where $a \in \mathbb{R}$ is a real number, t_i, t_j, and t_k are terms for possibly but not necessarily distinct agents, and φ is a sentence. Define abbreviations for such formulas such as $t = a$ and $\Diamond_i \varphi$ similarly to those given in Section 3.

The semantics is defined by a function $[\![\cdot]\!]$ from formulas to functions f that map each history $\mathbf{M}$ to a subset of $\widetilde{X}(\mathbf{M})$, the carrier set of $\mathbf{M}$. Then

$$
\begin{array}{lll}
x \in [\![\,\mathsf{true}]\!](\mathbf{M}) & \text{iff} & x \in \widetilde{X}(\mathbf{M}) \\
x \in [\![p]\!](\mathbf{M}) & \text{iff} & x \in \|p\|_{\mathbf{M}} \\
x \in [\![\neg\varphi]\!](\mathbf{M}) & \text{iff} & x \in [\![\,\mathsf{true}]\!](\mathbf{M}) - [\![\varphi]\!](\mathbf{M}) \\
x \in [\![\varphi \wedge \psi]\!](\mathbf{M}) & \text{iff} & x \in [\![\varphi]\!](\mathbf{M}) \cap [\![\psi]\!](\mathbf{M}) \\
x \in [\![\Box_i \varphi]\!](\mathbf{M}) & \text{iff} & z \in [\![\varphi]\!](\mathbf{M}) \text{ whenever } x \xrightarrow{i}_{\widetilde{\mathbf{M}}(\mathbf{M})} z \\
x \in [\![[\mathbf{\Sigma}, \sigma]\varphi]\!](\mathbf{M}) & \text{iff} & (x, \sigma) \in [\![\varphi]\!](\mathbf{M} \otimes \mathbf{\Sigma}) \text{ whenever } \sigma \in \mathsf{pre}_{\mathbf{\Sigma}}(x) \\
x \in [\![t_i \geq q]\!](\mathbf{M}) & \text{iff} & \left\{ \begin{array}{l} \sum_{j=1}^{n} q_j \mu^*_{i,x}([\![\varphi_j]\!](\mathbf{M}) \cap S_{i,x}) \geq q \\ \text{whenever } t_i = q_1 P_i(\varphi_1) + \cdots + q_n P_i(\varphi_n) \end{array} \right. \\
x \in [\![\mathsf{meas}_i(\varphi)]\!](\mathbf{M}) & \text{iff} & [\![\varphi]\!](\mathbf{M}) \in \mathcal{A}_{i,x}
\end{array}
$$

4.4 Discussion about Unmeasurability

Suppose agent j says to agent k, "I am about to write down on a piece of paper a letter that is either a or b." Agent k knows that b is agent j's favorite letter, but agent j could easily write a. If asked to give a probability for the letter a being written down, it might be reasonable for k to respond "I cannot give a specific probability to what will be written down." Although probability reflects degrees of uncertainty about an outcome, it does reflect an awareness of likelihood, something that is not always guaranteed. It would thus be helpful to be able to express in the language that probabilities are not available.

One method for doing this is to mix epistemic and probabilistic components of the language, as given in [3]. There, we would capture this uncertainty by having uncertainty between or among probability spaces. For example, the probabilistic epistemic model could consist of two states a and b, where a is written down at a and b is written down at b. From a, agent k's probability space consists of just $\{a\}$ as the sample space (where the measure assigns the value 1 to its only non-empty set), and from b, the probability space consists of just $\{b\}$ as the sample space. But the epistemic relation consists of all pairs, capturing complete uncertainty between a and b. Suppose that the actual state is a (agent j writes down a). Then the following (non-primitive) formulas are true in this model at a:

1. $P_k(\mathsf{a}) = 1$
2. $\neg\Box_k P_k(\mathsf{a}) = 1$
3. $\Box_k(P_k(\mathsf{a}) = 1 \vee P_k(\mathsf{b}) = 1)$.

If P_k were to represent k's subjective probability, then the first formula would suggest that k assigns probability 1 to the event that agent j writes down a. This is not quite what we want to capture, as we want to reflect that agent k is unsure about the probability. So it might be better to view these as objective probabilities. In this case the first formula states that the actual probability of a being written down is 1, the second formula states that you do not know that the actual probability of a being written down is 1, but the third formula says that agent k is certain that either the probability of a being written down is 1 or the probability of b being written down is 1 (and hence the probability of a being written down is 0). The second of these formulas may be appropriate, but the first and third could be debated. Regarding the first, we may not want to say the probability that agent j writes down the letter a is 1, until he has actually written it down. Regarding the second, we may want to capture that agent k considers any probability for j to write down a possible; this could involve an infinite model with infinitely many probability distributions, and an epistemic equivalence class over all of the distributions.

With unmeasurable sets, we can capture this situation with a model that includes the same two states and the same epistemic structure. But this time, the probability space from both states will be the space with σ-algebra $\{\emptyset, \{a, b\}\}$, assigning probability 1 to the set $\{a, b\}$ and providing no probability for the singletons. We interpret the formula $\mathsf{meas}(\mathsf{a})$ to mean that agent k assigns a probability to formula a, and we interpret $P(\mathsf{a}) = 1$ to mean that agent k assigns probability 1 to a being written down in the event that agent k assigns any probability to a at all. Thus we interpret this as a subjective probability. Now at a in the model described above, the formula $\mathsf{meas}(\mathsf{a})$ is false, reflecting that agent k does not assign probability to a being written down. The truth of its negation, $\neg\mathsf{meas}(\mathsf{a})$, signifies that the formula $P(\mathsf{a}) = 1$ is not about probabilities, but rather about outer probabilities.

4.5 Protecting formulas from unmeasurable sets

We may wish that every formula correspond to a measurable set. One way to do this is to place a restriction on the initial model that the σ-algebra be large enough for all formulas to correspond to measurable sets. But it is desirable, particularly when extending the language to one that involves past, to ensure that this property remains after updating. To complicate matters, we might not easily guarantee that the set Y in Definition 4.6 is itself measurable. Consider an action model for which k's probability space includes two action types: σ and τ; and suppose there are only two measurable sets: the whole set $\{\sigma, \tau\}$ and the empty set. Suppose a probabilistic epistemic model $\mathbf{M}$ has two states: x and y, and k's probability sample space is $\{x, y\}$ and all subsets are measurable. Then in the product measure, the measurable sets are

$$\{\emptyset, \{(x, \sigma), (x, \tau)\}, \{(y, \sigma), (y, \tau)\}, \{(x, \sigma), (y, \sigma), (x, \tau), (y, \tau)\}\}.$$

Suppose there were a formula φ for which only x is true, and another formula for which only y is true. Then these formulas correspond to measurable sets. Let the function pre reflect these two formulas, by defining $\mathsf{pre}(\sigma)(\mathbf{M}) = x$ and $\mathsf{pre}(\tau)(\mathbf{M}) = y$. Then when taking the full update product, we would be relativizing with respect to the set $Y = \{(x, \sigma), (y, \tau)\}$, which is not measurable. In general, if an action model has only discrete probability spaces (probability spaces where the σ-algebras are power sets of the sample spaces), then the measurability of the sets $\mathsf{pre}(\sigma)(\mathbf{M})$ for each $\sigma \in \Sigma$, does guarantee that the set Y in Definition 4.6 is measurable. It remains to be seen that in an updated model, every formula still corresponds to a measurable set.

5 Conclusion

This paper explores ways of extending probabilistic dynamic epistemic logic on two fronts. One is to add a previous-time operator to a probabilistic dynamic epistemic logic similar to the one given in [7], and the other is to involve non-trivial σ-algebras in probabilistic dynamic epistemic logic. The inclusion of a previous-time operator allows us to express agents' beliefs in the past or about past events. Allowing the involvement of non-trivial σ-algebras in probabilistic dynamic epistemic logic can be helpful in letting us express that an agent does not assign a probability to a formula being true, and in allowing us to involve continuous probabilities.

References

[1] Baltag, A., Moss, L., and Solecki, S. Logics for epistemic actions: Completeness, decidability, expressivity. ms. Indiana University, 2003.

[2] Baltag, A.and Moss, L. S. Logics for epistemic programs. *Synthese*, 139(2, Knowledge, Rationality & Action):165–224, 2004.

[3] Fagin, R. and Halpern, J. Reasoning about knowledge and probability. *JACM*, 41(2):340–367, 1994.

[4] Fagin, R., Halpern, J., Moses, Y., and Vardi, M. *Reasoning About Knowledge*. The MIT Press, 1995.

[5] Gariepy, R.F., Ziemer, W.P. *Modern Real Analysis*. PWS Publishing Company, 1995.

[6] Halpern, J. and Tuttle, M. Knowledge, probability, and adversaries. *JACM*, 40(4):917–962, 1993.

[7] Kooi, B. P. Probabilistic dynamic epistemic logic. *J. Logic Lang. Inform.*, 12(4):381–408, 2003.

[8] Sack, J. Logic for update products and steps into the past. ms., 2007.

[9] Sack, J. Temporal languages for epistemic programs. *Journal of Logic, Language and Information*, 17(2):183–216, 2008.

[10] van Benthem, J., Gerbrandy, J., and Kooi, B. Dynamic update with probabilities. In W. van der Hoek and M. Wooldridge, editors, *Proceedings of LOFT'06*, Liverpool, 2006