Statistical Epistemic Logic [⋆]

Yusuke Kawamoto$^{1\,[0000-0002-2151-9560]}$

National Institute of Advanced Industrial Science
and Technology (AIST), Tsukuba, Japan

Abstract. We introduce a modal logic for describing statistical knowledge, which we call *statistical epistemic logic*. We propose a Kripke model dealing with probability distributions and stochastic assignments, and show a stochastic semantics for the logic. To our knowledge, this is the first semantics for modal logic that can express the statistical knowledge dependent on non-deterministic inputs and the statistical significance of observed results. By using statistical epistemic logic, we express a notion of statistical secrecy with a confidence level. We also show that this logic is useful to formalize statistical hypothesis testing and differential privacy in a simple and abstract manner.

Keywords: Epistemic logic · Possible world semantics · Divergence · Statistical hypothesis testing · Differential privacy

1 Introduction

Knowledge representation and reasoning have been studied in two research areas: *logic* and *statistics*. Broadly speaking, logic describes our knowledge using formal languages and reasons about it using symbolic techniques, while statistics interprets collected data having random variation and infers properties of their underlying probability models. As research advances demonstrate, logical and statistical approaches are respectively successful in many applications, including artificial intelligence, software engineering, and information security.

The techniques of these two approaches are basically orthogonal and could be integrated to get the best of both worlds. For example, in a large system with artificial intelligence (e.g., an autonomous car), both rule-based knowledge and statistical machine learning models may be used, and the way of combining them would be crucial to the performance and security of the whole system. However, even in theoretical research on knowledge models, there still remains much to be done to integrate techniques from the two approaches. For a very basic example, *epistemic logic* [39], a formal logic for representing and reasoning about knowledge, has not yet been able to model "statistical knowledge" with sampling and statistical significance, although a lot of epistemic models [14, 20, 21] have been proposed so far.

[⋆] This work was supported by JSPS KAKENHI Grant Number JP17K12667, and by Inria under the project LOGIS.

One of the important challenges in integrating logical and statistical knowledge is to design a logical model for statistical knowledge, which can be updated by a limited number of sampling of probabilistic events and by the non-deterministic inputs from an external environment. Here we note that non-deterministic inputs are essential to model the security of the system, because we usually do not have a prior knowledge of the probability distribution of adversarial inputs and need to reason about the worst scenarios caused by the attack. Nevertheless, to the best of our knowledge, no previous work on epistemic logic has proposed an abstract model for the statistical knowledge that involves non-deterministic inputs and the statistical significance of observed results.

In the present paper, we propose an epistemic logic for describing statistical knowledge. To define its semantics, we introduce a variant of a Kripke model [29] in which each possible world is defined as a probability distribution of states and each variable is probabilistically assigned a value. In this model, the stochastic behaviour of a system is modeled as a distribution of states at each world, and each non-deterministic input to the system corresponds to a distinct possible world. As for applications of this model, we define an accessibility relation as a statistical distance between distributions of observations, and show that our logic is useful to formalize statistical hypothesis testing and differential privacy [11] of statistical data.

Our contributions. The main contributions of this work are as follows:

- We introduce a modal logic, called *statistical epistemic logic* (StatEL), to describe statistical knowledge.
- We propose a Kripke model incorporating probability distributions and stochastic assignments by regarding each possible world as a distribution of states and by defining an accessibility relation using a metric/divergence between distributions.
- We introduce a stochastic semantics for StatEL based on the above models. As far as we know, this is the first semantics for modal logic that can express the statistical knowledge dependent on non-deterministic inputs and the statistical significance of observed results.
- We present basic properties of the probability quantification and epistemic modality in StatEL. In particular, we show that the transitivity and Euclidean axioms rely on the agent's capability of observation in our model.
- By using StatEL we introduce a notion of statistical secrecy with a significance level α. We also show that StatEL is useful to formalize statistical hypothesis testing and differential privacy in a simple and abstract manner.

Paper organization. The rest of this paper is organized as follows. Section 2 introduces background and notations used in this paper. Section 3 presents an example of coin flipping to explain the motivation for a logic of statistical knowledge. Section 4 shows the syntax and semantics of the statistical epistemic logic StatEL. Section 5 presents basic properties of the logic. As for applications, Sections 6 and 7 respectively model statistical hypothesis testing and statistical data privacy using StatEL. Section 8 presents related work and Section 9 concludes.

2 Preliminaries

In this section we recall the definitions of divergence and metrics, which are used in later sections to quantitatively model an agent's capability of distinguishing possible worlds.

2.1 Notations

Let $\mathbb{R}^{\geq 0}$ be the set of non-negative real numbers, and $[0,1] = \{r \in \mathbb{R}^{\geq 0} \mid r \leq 1\}$. We denote by $\mathbb{D}\mathcal{O}$ the set of all probability distributions over a set $\mathcal{O}$. For a finite set $\mathcal{O}$ and a distribution $\mu \in \mathbb{D}\mathcal{O}$, the probability of sampling a value y from μ is denoted by $\mu[y]$. For a subset $R \subseteq \mathcal{O}$, let $\mu[R] = \sum_{y \in R} \mu[y]$. The *support* of a distribution μ over a finite set $\mathcal{O}$ is $\mathtt{supp}(\mu) = \{v \in \mathcal{O} \colon \mu[v] > 0\}$. For a set $\mathcal{D}$, a randomized algorithm $A : \mathcal{D} \to \mathbb{D}\mathcal{O}$ and a set $R \subseteq \mathcal{O}$ we denote by $A(d)[R]$ the probability that given input $d \in \mathcal{D}$, A outputs one of the elements of R.

2.2 Metric and Divergence

A *metric* over a non-empty set $\mathcal{O}$ is a function $d : \mathcal{O} \times \mathcal{O} \to \mathbb{R}^{\geq 0}$ such that for all $y, y', y'' \in \mathcal{O}$, (i) $d(y, y') \geq 0$; (ii) $d(y, y') = 0$ iff $y = y'$; (iii) $d(y, y') = d(y', y)$; (iv) $d(y, y'') \leq d(y, y') + d(y', y'')$. Recall that (iii) and (iv) are respectively referred to as symmetry and subadditivity.

A *divergence* over a non-empty set $\mathcal{O}$ is a function $D(\cdot \parallel \cdot) : \mathbb{D}\mathcal{O} \times \mathbb{D}\mathcal{O} \to \mathbb{R}^{\geq 0}$ such that for all $\mu, \mu' \in \mathbb{D}\mathcal{O}$, (i) $D(\mu \parallel \mu') \geq 0$ and (ii) $D(\mu \parallel \mu') = 0$ iff $\mu = \mu'$. Note that a divergence may not be symmetric or subadditive.

To describe a statistical hypothesis testing in Section 6, we recall the definition of χ^2 divergence due to Pearson [16] as follows:

Definition 1 (Pearson's χ^2 divergence). Given two distributions μ, μ' over a finite set $\mathcal{O}$, the χ^2*-divergence* $D_{\chi^2}(\mu \parallel \mu')$ of μ from μ' is defined by:

$$D_{\chi^2}(\mu \parallel \mu') = \sum_{y \in \mathtt{supp}(\mu)} \frac{(\mu'[y] - \mu[y])^2}{\mu[y]}.$$

χ^2 *statistics* is the multiplication of χ^2-divergence with a sample size n.

To introduce a notion of statistical data privacy in Section 7, we recall the definition of the max-divergence D_∞ as follows.

Definition 2 (Max divergence). For two distributions μ, μ' over a finite set $\mathcal{O}$, the *max divergence* $D_\infty(\mu \parallel \mu')$ of μ from μ' is defined by:

$$D_\infty(\mu \parallel \mu') = \max_{R \subseteq \mathtt{supp}(\mu)} \ln \frac{\mu[R]}{\mu'[R]}.$$

Note that neither D_{χ^2} nor D_∞ is symmetric.

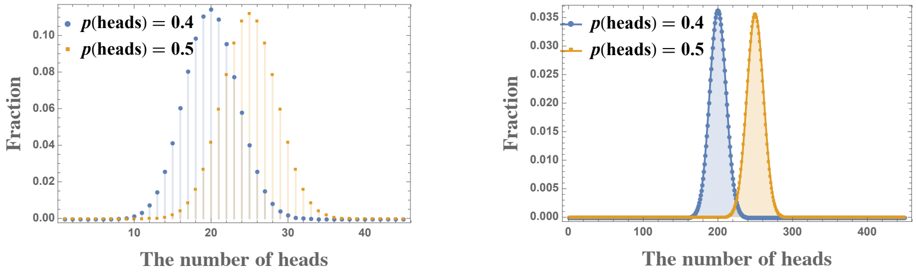

(a) Given 50 coin flips, the two distributions overlap much.

(b) Given 500 coin flips, the two distributions are distinguished more clearly.

Fig. 1: The frequency distributions of the numbers of heads in coin flipping.

3 Motivating Example

In this section we present a motivating example to explain why we need to introduce a new model for epistemic logic to describe statistical knowledge.

Example 1 (Coin flipping). Let us consider a simple running example of flipping a coin in two possible worlds w_0 and w_1 respectively. We assume that in the world w_0 the coin is fair (represented by $p(\mathit{heads}) = 0.5$), whereas in w_1 the probability of getting a heads is 0.4 (represented by $p(\mathit{heads}) = 0.4$). Here we do not have any prior belief on the probabilities of the worlds w_0 and w_1. This does not mean $p(w_0) = p(w_1) = 0.5$, but means we have no idea on the values of $p(w_0)$ and $p(w_1)$ at all, i.e., either w_0 or w_1 is chosen non-deterministically.

When we flip a coin just once and observe its outcome (heads or tails), we do not know whether the coin is fair or biased, that is, we cannot tell whether we are located in the world w_0 or w_1.

As shown in Fig. 1, however, when we increase the number n of coin flips, we can more clearly see the difference between the numbers of getting heads in w_0 and in w_1. If the fraction of observing heads goes to 0.5 (resp. 0.4), then we learn we are located in the world w_0 (resp. w_1) with a stronger confidence, namely, we have a stronger belief that the coin is fair (resp. biased). This implies that a larger number of observing the outcome enables us to distinguish two possible worlds more clearly, hence to obtain a stronger belief.

To model such statistical beliefs, we regard each possible world as a probability distribution of two states *heads* and *tails* as shown in Fig. 2 (e.g., $w_1[\mathit{heads}] = 0.4$ and $w_1[\mathit{tails}] = 0.6$). Then for a divergence D between two distributions, we define an accessibility relation $\mathcal{R}_\varepsilon$ between worlds such that for any worlds w and w', $(w, w') \in \mathcal{R}_\varepsilon$ iff $D(w \,\|\, w') \leq \varepsilon$. Then $(w_0, w_1) \in \mathcal{R}_\varepsilon$ for a smaller threshold ε represents that a larger number of sampling is required to distinguish w_0 from w_1.

This relation $\mathcal{R}_\varepsilon$ is used to formalize statistical knowledge in a model of epistemic logic in Section 4. Intuitively, given a threshold ε determining a confidence

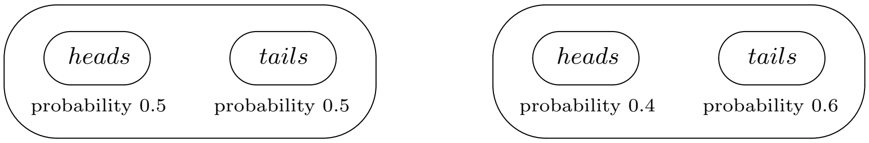

(a) The world w_0 with the fair coin.

(b) The world w_1 with the biased coin.

Fig. 2: One of the possible worlds (i.e., w_0 or w_1) is chosen non-deterministically. Then one of the states (i.e., *heads* or *tails*) is chosen probabilistically.

level, we say that we know a proposition φ in a world w if φ is satisfied in all possible worlds that are indistinguishable from w in terms of $\mathcal{R}_\varepsilon$. In Section 6 we will revisit the coin flipping example to see how we formalize it using our logic.

To our knowledge, no previous work on epistemic logic has modeled a statistical knowledge that depends on the agent's capability of observing events. In fact, in most of the Kripke models used in previous work, a possible world represents a single state instead of a probability distribution of states, hence the relation between possible worlds does not involve the probability of distinguishing them. Therefore, no prior work on epistemic logic has proposed an abstract model for the statistical knowledge that involves the sample size of observing random variables and the statistical significance of the observed results.

4 Statistical Epistemic Logic (StatEL)

In this section we introduce the syntax and semantics of the *statistical epistemic logic* (StatEL).

4.1 Syntax

We first present the syntax of the statistical epistemic logic as follows. To express both deterministic and probabilistic properties, we introduce two levels of formulas: *static formulas* and *epistemic formulas*. Intuitively, a static formula represents a proposition that can be satisfied at a state with probability 1, while an epistemic formula represents a proposition that can be satisfied at a probability distribution of states with some probability.

Formally, let `Mes` be a set of symbols called *measurement variables*, and $\mathit{\Gamma}$ be a set of atomic formulas of the form $\gamma(x_1, x_2, \ldots, x_n)$ for a predicate symbol γ and $x_1, x_2, \ldots, x_n \in \mathtt{Mes}$ $(n \geq 0)$. Let $I \subseteq [0, 1]$ be a finite union of intervals, and $\mathcal{A}$ be a finite set of indices (typically associated with the names of agents and/or statistical tests). Then the static and epistemic formulas are defined by:

$$\text{Static formulas:}\quad \psi ::= \gamma(x_1, x_2, \ldots, x_n) \mid \neg\psi \mid \psi \wedge \psi$$
$$\text{Epistemic formulas:}\quad \varphi ::= \mathbb{P}_I\, \psi \mid \neg\varphi \mid \varphi \wedge \varphi \mid \psi \supset \varphi \mid \mathsf{K}_a\, \varphi$$

where $a \in \mathcal{A}$. Let $\mathcal{F}$ be the set of all epistemic formulas. Note that we have no quantifiers over measurement variables. (See Section 4.5.)

The *probability quantification* $\mathbb{P}_I\,\psi$ represents that a static formula ψ is satisfied with a probability belonging to a set I. For instance, $\mathbb{P}_{(0.5,1]}\,\psi$ represents that ψ holds with a probability greater than 0.5. The *non-classical implication* $\supset$ is used to represent conditional probabilities. For example, by $\psi_0 \supset \mathbb{P}_I\,\psi_1$ we represent that the conditional probability of ψ_1 given ψ_0 is included in a set I. The *epistemic knowledge* $\mathsf{K}_a\,\varphi$ expresses that an agent a knows φ. The formal meaning of these operators will be shown in the definition of semantics.

As syntax sugar, we use *disjunction* $\vee$, *classical implication* $\rightarrow$, and *epistemic possibility operator* P_a, defined by: $\varphi_0 \vee \varphi_1 ::= \neg(\neg\varphi_0 \wedge \neg\varphi_1)$, $\varphi_0 \rightarrow \varphi_1 ::= \neg\varphi_0 \vee \varphi_1$, and $\mathsf{P}_a\,\varphi ::= \neg\,\mathsf{K}_a\,\neg\varphi$. When I is a singleton $\{i\}$, we abbreviate $\mathbb{P}_{[i,i]}$ as $\mathbb{P}_i$.

4.2 Modeling of Systems

In this work we deal with a simple stochastic system with measurement variables. Let $\mathcal{O}$ be the finite set of all data that can be assigned to the measurement variables in $\mathtt{Mes}$. We assume that all possible worlds share the same domain $\mathcal{O}$. We define a *stochastic system* as a pair (S, σ) consisting of:

- a stochastic program S that deals with input and output data through measurement variables in $\mathtt{Mes}$, behaves deterministically or probabilistically (by using some randomly generated data), and terminates with probability 1;
- a *stochastic assignment* $\sigma : \mathtt{Mes} \rightarrow \mathbb{D}\mathcal{O}$ representing that each measurement variable x has an observed value v with probability $\sigma(x)[v]$.

Here we present only a general model and do not specify the data type of those measurement variables, which can be (sequences of) bit strings, floating point numbers, texts, or other types of data. Thanks to the assumption on the program termination and on the finite range of data, the program S can reach finitely many states. For the sake of simplicity, our model does not take timing into account. Extension to time and temporal modality is left for future work.

4.3 Distributional Kripke Model

To define a semantics for StatEL, we recall the notion of a Kripke model [29]:

Definition 3 (Kripke model). Given a set Γ of atomic formulas, a *Kripke model* is defined as a triple $(\mathcal{W}, \mathcal{R}, V)$ consisting of a non-empty set $\mathcal{W}$, a binary relation $\mathcal{R}$ on $\mathcal{W}$, and a function V that maps each atomic formula $\gamma \in \Gamma$ to a subset $V(\gamma)$ of $\mathcal{W}$. The set $\mathcal{W}$ is called a *universe*, its elements are called *possible worlds*, $\mathcal{R}$ is called an *accessibility relation*, and V is called a *valuation*.

Now we introduce a Kripke model called a "distributional" Kripke model where each possible world is a probability distribution of states over $\mathcal{S}$ and each world w is associated with a stochastic assignment σ_w to measurement variables.

Definition 4 (Distributional Kripke model). Let $\mathcal{A}$ be a finite set of indices (typically associated with the names of agents and/or statistical tests), $\mathcal{S}$ be a finite set of states[1], and $\mathcal{O}$ be a finite set of data. A *distributional Kripke model* is a tuple $\mathfrak{M} = (\mathcal{W}, (\mathcal{R}_a)_{a\in\mathcal{A}}, (V_s)_{s\in\mathcal{S}})$ consisting of:

- a non-empty set[2] $\mathcal{W}$ of probability distributions of states over $\mathcal{S}$;
- for each $a \in \mathcal{A}$, an accessibility relation $\mathcal{R}_a \subseteq \mathcal{W} \times \mathcal{W}$;
- for each $s \in \mathcal{S}$, a valuation V_s that maps each k-ary predicate γ to a set $V_s(\gamma) \subseteq \mathcal{O}^k$.

We assume that each $w \in \mathcal{W}$ is associated with a function $\rho_w : \mathtt{Mes} \times \mathcal{S} \to \mathcal{O}$ that maps each measurement variable x to its value $\rho_w(x, s)$ observed at a state s. We also assume that each state s in a world w is associated with the assignment $\sigma_s : \mathtt{Mes} \to \mathcal{O}$ defined by $\sigma_s(x) = \rho_w(x, s)$.

Note that this model assumes a constant domain $\mathcal{O}$; i.e., all measurement variables range over the same set $\mathcal{O}$ in every world. Since each world w is a probability distribution of states, we denote by $w[s]$ the probability that a state s is sampled from w. Then the probability that a variable x has a value v in a world w is given by:

$$\sigma_w(x)[v] = \sum_{s\in\mathtt{supp}(w),\, \sigma_s(x)=v} w[s].$$

This means that when a state s is drawn from the distribution w, an input value $\sigma_s(x)$ is sampled from the distribution $\sigma_w(x)$.

4.4 Divergence-based Accessibility Relation

Next we introduce a family of accessibility relations used in typical statistical inferences. Since many notions of statistical distance are not metrics but divergences, we introduce an accessibility relation based on a divergence as follows.

Suppose that an agent a observes some data through a single measurement variable x. Then the distribution of the observed data at a world w is represented by $\sigma_w(x)$. Assume that the agent a distinguishes distributions in terms of a divergence $D(\cdot \,\|\, \cdot) : \mathbb{D}\mathcal{O} \times \mathbb{D}\mathcal{O} \to \mathbb{R}^{\geq 0}$. Then given a threshold $\varepsilon \geq 0$, we define a *divergence-based accessibility relation* $\mathcal{R}_{a,\varepsilon}$ by:

$$\mathcal{R}_{a,\varepsilon} \stackrel{\text{def}}{=} \left\{(w, w') \in \mathcal{W} \times \mathcal{W} \mid D(\sigma_w(x) \,\|\, \sigma_{w'}(x)) \leq \varepsilon\right\}.$$

For a smaller value of ε, the capability of distinguishing worlds is stronger.

If D is a metric instead, we call $\mathcal{R}_{a,\varepsilon}$ a *metric-based accessibility relation.* We often omit a to write $\mathcal{R}_\varepsilon$ when we do not compare different agents' knowledge.

Intuitively, $(w, w') \in \mathcal{R}_{a,\varepsilon}$ represents that the distribution of the data observed in w is indistinguishable from that in w' in terms of D. By the definition of a divergence/metric D, $D(\sigma_w(x) \,\|\, \sigma_{w'}(x)) = 0$ implies $\sigma_w(x) = \sigma_{w'}(x)$. Therefore, the relation $\mathcal{R}_{a,0}$ expresses that the agent a has an unlimited capability of observing the distributions $\sigma_w(x)$ and $\sigma_{w'}(x)$. In Sections 6 and 7 we will show examples of divergence-based accessibility relations.

[1] It is left for future work to investigate the case of infinite numbers of states.

[2] Since $\mathcal{W}$ is not a multiset, each world in $\mathcal{W}$ is a different distribution of states. However, this is still expressive enough when we take $\mathcal{S}$ to be sufficiently large.

4.5 Stochastic Semantics

In this section we define the *stochastic semantics* for the StatEL formulas over a distributional Kripke model $\mathfrak{M} = (\mathcal{W}, (\mathcal{R}_a)_{a\in\mathcal{A}}, (V_s)_{s\in\mathcal{S}})$ with $\mathcal{W} = \mathbb{D}\mathcal{S}$.

The interpretation of static formulas ψ at a state s is given by:

$$\begin{aligned} s \models \gamma(x_1, x_2, \ldots, x_k) &\text{ iff } (\sigma_s(x_1), \sigma_s(x_2), \ldots, \sigma_s(x_k)) \in V_s(\gamma) \\ s \models \neg\psi &\text{ iff } s \not\models \psi \\ s \models \psi \wedge \psi' &\text{ iff } s \models \psi \text{ and } s \models \psi'. \end{aligned}$$

Note that the satisfaction of the static formulas does not involve probability.

To interpret the non-classical implication $\supset$, we define the *restriction* $w|_\psi$ of a world w to a state formula ψ as follows. If there exists a state s such that $w[s] > 0$ and $s \models \psi$, then $w|_\psi$ can be defined as the distribution over the finite set $\mathcal{S}$ of states such that:

$$w|_\psi[s] = \begin{cases} \frac{w[s]}{\sum_{s':s'\models\psi} w[s']} & \text{if } s \models \psi \\ 0 & \text{otherwise.} \end{cases}$$

Then $\sum_s w|_\psi[s] = 1$. Note that $w|_\psi$ is undefined if w does not have a state s that satisfies ψ and has a non-zero probability in w.

Now we define the interpretation of epistemic formulas at a world w in $\mathfrak{M}$ by:

$$\begin{aligned} \mathfrak{M}, w \models \mathbb{P}_I\, \psi &\text{ iff } \Pr\Big[s \overset{\$}{\leftarrow} w : \; s \models \psi\Big] \in I \\ \mathfrak{M}, w \models \neg\varphi &\text{ iff } \mathfrak{M}, w \not\models \varphi \\ \mathfrak{M}, w \models \varphi \wedge \varphi' &\text{ iff } \mathfrak{M}, w \models \varphi \text{ and } \mathfrak{M}, w \models \varphi' \\ \mathfrak{M}, w \models \psi \supset \varphi &\text{ iff } w|_\psi \text{ is defined and } \mathfrak{M}, w|_\psi \models \varphi \\ \mathfrak{M}, w \models \mathsf{K}_a\, \varphi &\text{ iff } \text{for every } w' \text{ s.t. } (w, w') \in \mathcal{R}_a, \; \mathfrak{M}, w' \models \varphi, \end{aligned}$$

where $s \overset{\$}{\leftarrow} w$ represents that a state s is sampled from the distribution w.

Finally, the interpretation of an epistemic formula φ in $\mathfrak{M}$ is given by:

$$\mathfrak{M} \models \varphi \text{ iff for every world } w \text{ in } \mathfrak{M}, \; \mathfrak{M}, w \models \varphi.$$

We remark that in each world w, measurement variables can be interpreted using σ_w, as shown in Section 4.3. This allows one to assign different values to distinct occurrences of a variable in a formula; E.g., in $\varphi(x) \rightarrow \mathsf{K}_a\, \varphi'(x)$, the measurement variable x occurring in $\varphi(x)$ can be interpreted using σ_w in a world w, while x in $\varphi'(x)$ can be interpreted using $\sigma_{w'}$ in another w' s.t. $(w, w') \in \mathcal{R}_a$.

Note that our semantics for probability quantification is different from that in the previous work. Halpern [19] shows two approaches to defining semantics: giving probabilities (1) on the domain and (2) on possible worlds. However, our semantics is different from both. It defines probabilities on the states belonging to a possible world, while each world is not assigned a probability. Hence, unlike Halpern's approaches, our model can deal with both probabilistic behaviours of systems and non-deterministic inputs from an external environment.

We also remark that StatEL can be used to formalize conditional probabilities. If the conditional probability of satisfying a static formula ψ_1 given another static formula ψ_0 is included in a set I at a world w, then we have $\Pr\Big[s \overset{\$}{\leftarrow} w|_{\psi_0} : s \models \psi_1\Big] \in I$, hence we obtain $\mathfrak{M}, w \models \psi_0 \supset \mathbb{P}_I\, \psi_1$.

5 Basic Properties of StatEL

In this section we present basic properties of StatEL. In particular, we show the transitivity and Euclidean axioms rely on the agent's capability of observation.

5.1 Properties of Probability Quantification

We can define a dual operator of $\mathbb{P}_I$ as follows. Given a finite union $I \subseteq [0,1]$ of disjoint intervals, let $I^c \overset{\text{def}}{=} [0,1] \setminus I$ and $\overline{I} \overset{\text{def}}{=} \{1-p \mid p \in I\}$. Then $\overline{I^c} = \overline{I}^c$. Negation with $\mathbb{P}_I$ has the following properties.

Proposition 1 (Negation with probability quantification) *For any world w in a model $\mathfrak{M}$ and any static formula ψ, we have:*

1. $\mathfrak{M}, w \models \neg\,\mathbb{P}_I\, \psi$ *iff* $\mathfrak{M}, w \models \mathbb{P}_{I^c}\, \psi$
2. $\mathfrak{M}, w \models \mathbb{P}_I\, \neg\psi$ *iff* $\mathfrak{M}, w \models \mathbb{P}_{\overline{I}}\, \psi$.

By Proposition 1, $\neg\,\mathbb{P}_I\,\neg\psi$ is logically equivalent to $\mathbb{P}_{\overline{I^c}}\,\psi$. For instance, $\neg\,\mathbb{P}_{(0,1]}\,\neg\psi$ is equivalent to $\mathbb{P}_1\,\psi$, and $\neg\,\mathbb{P}_{[0,1)}\,\neg\psi$ is equivalent to $\mathbb{P}_0\,\psi$.

5.2 Properties of Epistemic Modality

Next we show some properties of epistemic modality. As with the standard modal logic, StatEL satisfies the necessitation rule and distribution axiom.

Proposition 2 (Minimal properties) *For any distributional Kripke model $\mathfrak{M}$, any $\varphi, \varphi_0, \varphi_1 \in \mathcal{F}$, and any $a \in \mathcal{A}$, we have:*

(**N**) *necessitation:* $\mathfrak{M} \models \varphi$ *implies* $\mathfrak{M} \models \mathsf{K}_a\, \varphi$
(**K**) *distribution:* $\mathfrak{M} \models \mathsf{K}_a(\varphi_0 \rightarrow \varphi_1) \rightarrow (\mathsf{K}_a\, \varphi_0 \rightarrow \mathsf{K}_a\, \varphi_1)$.

The satisfaction of other properties depends on the definition of the accessibility relation. Since many notions of statistical distance are not metrics but divergences, we present some basic properties when $\mathfrak{M}$ has a divergence-based accessibility relation: $\mathcal{R}_{a,\varepsilon} = \{(w, w') \in \mathcal{W} \times \mathcal{W} \mid D(\sigma_w(x) \parallel \sigma_{w'}(x)) \leq \varepsilon\}$.

Proposition 3 (Properties with divergence-based accessibility) *Let $a \in \mathcal{A}$ and $\varepsilon \geq \varepsilon' \geq 0$. For any distributional Kripke model $\mathfrak{M}$ with a divergence-based accessibility relation $\mathcal{R}_{a,\varepsilon}$ and any $\varphi \in \mathcal{F}$, we have:*

(**T**) *reflexivity:* $\mathfrak{M} \models \mathsf{K}_{a,\varepsilon}\, \varphi \rightarrow \varphi$
($\geq$) *comparison of observability:* $\mathfrak{M} \models \mathsf{K}_{a,\varepsilon}\varphi \rightarrow \mathsf{K}_{a,\varepsilon'}\varphi$.

If $\mathcal{R}_{a,\varepsilon}$ is symmetric (e.g., based on the Jensen-Shannon divergence [33]) then:

(**B**) *symmetry:* $\mathfrak{M} \models \varphi \rightarrow \mathsf{K}_{a,\varepsilon}\, \mathsf{P}_{a,\varepsilon}\, \varphi$.

Here the axiom ($\geq$) represents that an agent having a stronger capability of distinguishing worlds may have more beliefs.

Finally, we show some properties when $\mathcal{R}_{a,\varepsilon}$ is based on a metric (e.g. the p-Wasserstein metric [38], including the Earth mover's distance).

Proposition 4 (Properties with metric-based accessibility) *Let $a \in \mathcal{A}$ and $\varepsilon, \varepsilon' \geq 0$. For any distributional Kripke model $\mathfrak{M}$ with a metric-based accessibility relation $\mathcal{R}_{a,\varepsilon}$ and any $\varphi \in \mathcal{F}$, we have* (**T**)*reflexivity,* (**B**)*symmetry, and:*

(**4q**) *quantitative transitivity:* $\mathfrak{M} \models \mathsf{K}_{a,\varepsilon+\varepsilon'}\varphi \rightarrow \mathsf{K}_{a,\varepsilon}\, \mathsf{K}_{a,\varepsilon'}\varphi$
(**5q**) *relaxed Euclidean:* $\mathfrak{M} \models \mathsf{P}_{a,\varepsilon}\varphi \rightarrow \mathsf{K}_{a,\varepsilon'}\mathsf{P}_{a,\varepsilon+\varepsilon'}\varphi$.

If the agent has an unlimited capability of observation (i.e., $\varepsilon = \varepsilon' = 0$), then:

(**4**) *transitivity:* $\mathfrak{M} \models \mathsf{K}_{a,0}\, \varphi \rightarrow \mathsf{K}_{a,0}\, \mathsf{K}_{a,0}\, \varphi$
(**5**) *Euclidean:* $\mathfrak{M} \models \mathsf{P}_{a,0}\, \varphi \rightarrow \mathsf{K}_{a,0}\, \mathsf{P}_{a,0}\, \varphi$.

By this proposition, for $\varepsilon = 0$, StatEL has the axioms of **S5**, hence the epistemic operator $\mathsf{K}_{a,0}$ represents knowledge rather than beleif.

However, if the agent has a limited observability (i.e., $\varepsilon > 0$), then neither transitivity nor Euclidean may hold. This means that, even when he know whether φ holds or not with some confidence, he may not be perfectly confident that he knows it.

6 Modeling Statistical Hypothesis Testing Using StatEL

In this section we formalize statistical hypothesis testing by using StatEL formulas, and introduce a notion of statistical secrecy with a confidence level.

6.1 Statistical Hypothesis Testing

A *statistical hypothesis testing* is a method of statistical inference to check whether given datasets provide sufficient evidence to support some hypothesis. Typically, given two datasets, a *null hypothesis* H_0 is defined to claim that there is no statistical relationship between the two datasets (e.g., no difference between the result of a medical treatment and the placebo effect), while an *alternative hypothesis* H_1 represents that there is some relationship between them (e.g., the result of a medical treatment is better than the placebo effect).

Before performing a hypothesis test, we specify a *significance level* α, i.e., the probability that the test might reject the null hypothesis H_0, given that H_0 is true. Typically, α is 0.05 or 0.01. $1 - \alpha$ is called a *confidence level.*

6.2 Formalization of Statistical Hypothesis Testing

Now we define a distributional Kripke model $\mathfrak{M}$ with a universe $\mathcal{W}$ that includes at least two worlds w_{real} and w_{ideal} corresponding to the two datasets we compare:

- the real world w_{real} where we have a dataset sampled from actual experiments (e.g., from a medical treatment whose effectiveness we want to know);
- the ideal world w_{ideal} where we have a dataset that is synthesized from the null hypothesis setting (e.g., the dataset obtained from the placebo effect).

Note that $\mathcal{W}$ may include other worlds corresponding to different possible datasets.

Let n be the size of the dataset, and x be a measurement variable denoting a single data value chosen from the dataset we have. We assume that each world w has a state s corresponding to each single data value $\sigma_s(x)$ in the dataset. Then $\sigma_{w_{\mathsf{real}}}(x)$ is the empirical distribution (histogram) calculated from the dataset observed in the actual experiments in w_{real}, while $\sigma_{w_{\mathsf{ideal}}}(x)$ is the distribution calculated from the synthetic dataset in w_{ideal}. Then the number of data having a value v in the dataset in a world w is given by $n \cdot \sigma_w(x)[v]$.

Assume that $\mathfrak{M}$ has an accessibility relation $\mathcal{R}_{c_\alpha/n}$ that is specific to the sample size n, the statistical hypothesis test, and the *critical value* c_α for a significance level α we use. For brevity let $\varepsilon_{\alpha,n} = c_\alpha/n$. Intuitively, $(w_{\mathsf{real}}, w_{\mathsf{ideal}}) \in \mathcal{R}_{\varepsilon_{\alpha,n}}$ represents that the hypothesis test cannot distinguish the actual dataset from the synthetic one. For instance, when we use Pearson's χ^2-test as the hypothesis test, then $\mathcal{R}_{\varepsilon_{\alpha,n}}$ is defined by:

$$\mathcal{R}_{\varepsilon_{\alpha,n}} \overset{\text{def}}{=} \big\{(w, w') \in \mathcal{W} \times \mathcal{W} \mid D_{\chi^2}(\sigma_w(x) \parallel \sigma_{w'}(x)) \leq \varepsilon_{\alpha,n}\big\},$$

where D_{χ^2} is Pearson's χ^2 divergence (Definition 1).

Observe that when the confidence level $1-\alpha$ increases, then c_α decreases, hence $\varepsilon_{\alpha,n} = c_\alpha/n$ is smaller, i.e., the capability of distinguishing possible worlds is stronger.

Let φ_{syn} be a formula representing that the dataset is synthesized from the null hypothesis setting (e.g., representing the placebo effect). Then $\mathfrak{M}, w_{\mathsf{ideal}} \models \varphi_{\text{syn}}$. Since each world in $\mathcal{W}$ corresponds to a different dataset, it holds for any $w' \neq w_{\mathsf{ideal}}$ that $\mathfrak{M}, w' \models \neg\varphi_{\text{syn}}$. For instance, $\mathfrak{M}, w_{\mathsf{real}} \models \neg\varphi_{\text{syn}}$, since the actual dataset is used in w_{real} even when it looks indistinguishable from the synthetic dataset by the hypothesis test.

When the null hypothesis is rejected with a confidence level $1-\alpha$, then $(w_{\mathsf{real}}, w_{\mathsf{ideal}}) \notin \mathcal{R}_{\varepsilon_{\alpha,n}}$. Since $\mathfrak{M}, w' \models \neg\varphi_{\text{syn}}$ holds for any $w' \neq w_{\mathsf{ideal}}$, this rejection of the null hypothesis implies:

$$\mathfrak{M}, w_{\mathsf{real}} \models \mathsf{K}_{\varepsilon_{\alpha,n}} \neg\varphi_{\text{syn}},$$

which is logically equivalent to $\mathfrak{M}, w_{\mathsf{real}} \models \neg\, \mathsf{P}_{\varepsilon_{\alpha,n}}\, \varphi_{\text{syn}}$. This means that with the confidence level $1-\alpha$, we know we are not located in the world w_{ideal}, hence do not have a synthetic dataset.

On the other hand, when the null hypothesis is not rejected with a confidence level $1-\alpha$, then $(w_{\mathsf{real}}, w_{\mathsf{ideal}}) \in \mathcal{R}_{\varepsilon_{\alpha,n}}$. Thus we obtain:

$$\mathfrak{M}, w_{\mathsf{real}} \models \mathsf{P}_{\varepsilon_{\alpha,n}} \varphi_{\mathrm{syn}}. \tag{1}$$

This means that we cannot recognize whether we are located in the world w_{real} or w_{ideal}, i.e., we are not sure which database we have. To see this in details, let φ' be a formula representing that we have a third database (different from those in w_{real} and w_{ideal}). Suppose that another null hypothesis of satisfying φ' is not rejected with a confidence level $1-\alpha$. Then we have $\mathfrak{M}, w_{\mathsf{real}} \models \mathsf{P}_{\varepsilon_{\alpha,n}} \varphi'$. Since each world in $\mathcal{W}$ corresponds to a different database, we obtain $\mathfrak{M}, w_{\mathsf{real}} \models \mathsf{P}_{\varepsilon_{\alpha,n}} \neg\varphi_{\mathrm{syn}}$, which implies $\mathfrak{M}, w_{\mathsf{real}} \models \neg\, \mathsf{K}_{\varepsilon_{\alpha,n}} \varphi_{\mathrm{syn}}$. This represents that, when the null hypothesis is not rejected, we are not sure whether the null hypothesis is true or false.

6.3 Formalization of Statistical Secrecy

Now let us formalize the coin flipping in Example 1 in Section 3 by using StatEL as follows. Recall that $p(\mathit{heads}) = 0.5$ in w_0 and $p(\mathit{heads}) = 0.4$ in w_1. Let ψ be a static formula representing that the coin is a heads. Then $\mathfrak{M}, w_0 \models \mathbb{P}_{0.5}\, \psi$ and $\mathfrak{M}, w_1 \models \mathbb{P}_{0.4}\, \psi$. Assume that either $p(\mathit{heads}) = 0.5$ or $p(\mathit{heads}) = 0.4$ holds, i.e., $\mathfrak{M} \models \mathbb{P}_{0.5}\, \psi \vee \mathbb{P}_{0.4}\, \psi$.

When we have a sufficient number n of coin flips (e.g., $n = 500$), we can distinguish $p(\mathit{heads}) = 0.5$ from $p(\mathit{heads}) = 0.4$ (i.e., w_0 from w_1) by a hypothesis test. Hence we learn the probability $p(\mathit{heads})$ with some confidence level $1-\alpha$, i.e., $\mathfrak{M}, w_0 \models \mathsf{K}_{\varepsilon_{\alpha,n}} \mathbb{P}_{0.5}\, \psi$ and $\mathfrak{M}, w_1 \models \mathsf{K}_{\varepsilon_{\alpha,n}} \mathbb{P}_{0.4}\, \psi$. Therefore we obtain:

$$\mathfrak{M} \models \big(\mathbb{P}_{0.5}\, \psi \rightarrow \mathsf{K}_{\varepsilon_{\alpha,n}} \mathbb{P}_{0.5}\, \psi\big) \wedge \big(\mathbb{P}_{0.4}\, \psi \rightarrow \mathsf{K}_{\varepsilon_{\alpha,n}} \mathbb{P}_{0.4}\, \psi\big).$$

Note that for a larger sample size $n' > n$, we have $\varepsilon_{\alpha,n'} = c_\alpha/n' < c_\alpha/n = \varepsilon_{\alpha,n}$, hence it follows from the axiom $(\geq)$ in Proposition 3 that:

$$\mathfrak{M} \models \big(\mathbb{P}_{0.5}\, \psi \rightarrow \mathsf{K}_{\varepsilon_{\alpha,n'}} \mathbb{P}_{0.5}\, \psi\big) \wedge \big(\mathbb{P}_{0.4}\, \psi \rightarrow \mathsf{K}_{\varepsilon_{\alpha,n'}} \mathbb{P}_{0.4}\, \psi\big).$$

This means that if our knowledge derived from a smaller sample is statistically significant, then we derive the same conclusion from a larger sample.

On the other hand, when we have a very small number n'' of coin flips, we cannot distinguish w_0 from w_1. Then we are not sure about $p(\mathit{heads})$ with a confidence level $1-\alpha$, i.e., $\mathfrak{M}, w_0 \models \mathsf{P}_{\varepsilon_{\alpha,n''}} \mathbb{P}_{0.5}\, \psi$ and $\mathfrak{M}, w_1 \models \mathsf{P}_{\varepsilon_{\alpha,n''}} \mathbb{P}_{0.4}\, \psi$. Hence:

$$\mathfrak{M} \models (\mathbb{P}_{0.5}\, \psi \vee \mathbb{P}_{0.4}\, \psi) \rightarrow (\mathsf{P}_{\varepsilon_{\alpha,n''}} \mathbb{P}_{0.5}\, \psi \wedge \mathsf{P}_{\varepsilon_{\alpha,n''}} \mathbb{P}_{0.4}\, \psi).$$

This expresses a secrecy of $p(\mathit{heads})$. We generalize this to introduce the following definition of secrecy.

Definition 5 ((α, n)-statistical secrecy). Let Φ be a finite set of formulas, $\alpha \in [0, 1]$ be a significance level, and n be a sample size. We say that Φ is (α, n)-*statistically secret* if we have:

$$\mathfrak{M} \models \bigvee_{\varphi \in \Phi} \varphi \rightarrow \bigwedge_{\varphi \in \Phi} \mathsf{P}_{\varepsilon_{\alpha,n}} \varphi.$$

In the above coin flipping example, $\{\mathbb{P}_{0.5}\,\psi,\, \mathbb{P}_{0.4}\,\psi\}$ is (α, n)-statistically secret for some significance level α and sample size n. Syntactically, (α, n)-statistical secrecy resembles the notion of *total anonymity* [21], whereas in our definition, the epistemic operator $\mathsf{P}_{\varepsilon_{\alpha,n}}$ deals with the statistical significance and φ is not limited to a formula representing an agent's action.

7 Modeling Statistical Data Privacy Using StatEL

In this section we formalize a notion of statistical data privacy by using StatEL.

7.1 Differential Privacy

Differential privacy [11, 12] is a popular measure of data privacy guaranteeing that by observing a statistics about a database d, we cannot learn whether an individual user's record is included in d or not.

As a toy example, let us assume that the body weight of individuals is sensitive information, and we publish the average weight of all users recorded in a database d. Then we denote by d' the database obtained by adding to d a single record of a new user u's weight. If we also disclose the average weight of all users in d', then you learn u's weight from the difference between these two averages.

To mitigate such privacy leaks, many studies have proposed *obfuscation mechanisms*, i.e., randomized algorithms that add random noise to the statistics calculated from databases. In the above example, an obfuscation mechanism receives a database d and outputs a statistics of average weight to which some random noise is added. Then you cannot learn much information on u's weight from the perturbed statistics of average weight.

The privacy achieved by such obfuscation is often formalized as differential privacy. Intuitively, an ε-differential privacy mechanism makes every two "adjacent" (i.e., close) database d and d' indistinguishable with a degree of ε.

Definition 6 (Differential privacy). Let e be the base of natural logarithm, $\varepsilon \geq 0$, $\mathcal{D}$ be the set of all databases, and $\Psi \subseteq \mathcal{D} \times \mathcal{D}$ be an adjacency relation between two databases. A randomized algorithm $A : \mathcal{D} \to \mathbb{D}\mathcal{O}$ provides *ε-differential privacy* w.r.t. Ψ if for any $(d, d') \in \Psi$ and any $R \subseteq \mathcal{O}$,

$$\Pr[A(d) \in R] \leq e^{\varepsilon} \Pr[A(d') \in R]$$

where the probability is taken over the randomness in A.

For a smaller ε, the protection of differential privacy is stronger. It is known that differential privacy can be defined using the max-divergence D_∞ (Definition 2) as follows [12].

Proposition 5 *An obfuscation mechanism $A : \mathcal{D} \to \mathbb{D}\mathcal{O}$ provides ε-differential privacy w.r.t. $\Psi \subseteq \mathcal{D} \times \mathcal{D}$ iff for any $(d, d') \in \Psi$, $D_\infty(A(d) \parallel A(d')) \leq \varepsilon$ and $D_\infty(A(d') \parallel A(d)) \leq\ \varepsilon$.*

7.2 Formalization of Differential Privacy

Next we define a distributional Kripke model $\mathfrak{M} = (\mathcal{W}, \mathcal{R}_\varepsilon, (V_s)_{s\in\mathcal{S}})$ where there is a possible world corresponding to each database in $\mathcal{D}$. We assume that each world is a probability distribution of states in each of which an obfuscation mechanism A uses a different value of random seed for providing a probabilistically perturbed output. Let x (resp. y) be a measurement variable denoting the input (resp. output) of the obfuscation mechanism A. In each world w, $\sigma_w(x)$ is the database that A receives as input, and $\sigma_w(y)$ is the distribution of statistics that A outputs. Then the set of all databases is denoted by $\mathcal{D} = \{\sigma_w(x) \mid w \in \mathcal{W}\}$.

Now we define the accessibility relation $\mathcal{R}_\varepsilon$ in $\mathfrak{M}$ by using the max divergence D_∞ as follows[3]:

$$\mathcal{R}_\varepsilon \overset{\text{def}}{=} \{(w, w') \in \mathcal{W} \times \mathcal{W} \mid D_\infty(\sigma_w(y) \,\|\, \sigma_{w'}(y)) \leq \varepsilon,\ D_\infty(\sigma_{w'}(y) \,\|\, \sigma_w(y)) \leq \varepsilon\}\,.$$

Intuitively, $(w, w') \in \mathcal{R}_\varepsilon$ represents that, when we observe an output y of the obfuscation mechanism A, we do not know which of the two worlds w and w' we are located at. Hence we do not see which of the two databases $\sigma_w(x)$ and $\sigma_{w'}(x)$ was the input to A.

For each $d \in \mathcal{D}$, let φ_d be a formula representing that we have a database d. Then the ε-differential privacy of A w.r.t. an adjacency relation Ψ is expressed as:

$$\mathfrak{M} \models \bigwedge_{d\in\mathcal{D}} \Big(\varphi_d \rightarrow \bigwedge_{d'\in\Psi(d)} \mathsf{P}_\varepsilon\, \varphi_{d'}\Big).$$

Note that the privacy of user attributes defined as distribution privacy [27] can also be expressed using StatEL, since it is defined as the differential privacy w.r.t. a relation between the probability distributions that represent user attributes. We will elaborate on this in future work.

8 Related Work

In this section, we overview related work, including the integration of logical and statistical techniques, epistemic logic, and logical formalization of privacy.

Integration of logical and statistical techniques. There have been various studies on integrating logical and statistical techniques in software engineering. Notable examples are *probabilistic programming* [18], which has sampling from distributions and conditioning by observations, and *statistical model checking* [36, 40, 31], which checks the satisfiability of logical formulas by simulations and statistical hypothesis tests. In research of privacy, a few papers (e.g., [5]) present hybrid methods combining symbolic and statistical analyses to quantify privacy leaks. In future work, our logic may be used to define specifications of these techniques and characterize their properties.

[3] Since the relation $\mathcal{R}_\varepsilon$ is symmetric, the symmetry axiom (**B**) also holds.

Non-determinism and probability in Kripke models. Although many epistemic models have been proposed [14, 20, 21], they often assume that each possible world is a single deterministic state. To formalize the behaviours of stochastic systems in their model, they assume that every world is assigned a probability (e.g., [28]), which means the non-determinism needs to be resolved in advance.

However, not only probability but also non-deterministic inputs are essential to reason about security and many applications in statistics. In the context of security, we usually do not have a prior knowledge of the probability distribution of adversarial inputs. Also in the statistical hypothesis testing (Section 6.1), we do not assume the prior probabilities of the null/alternative hypotheses. The notion of differential privacy (Definition 6) is also independent of the prior distribution on the databases. Therefore, unlike ours, the Kripke models in previous work cannot be used for the purpose of formalizing such statistical knowledge.

Kripke model for some aspects of statistics. The *random worlds model* [20] is an epistemic model that tries to formalize some aspects of statistics. In that model, they assume that each possible world has an identical probability at the initial time, although this causes problems as mentioned in Chapter 10 of [20]. Unlike our distributional model, their model employs neither distributions of states nor statistical significance. They assume only finite intervals of errors, and analyze only the ideal situation that corresponds to an infinite sample size. Therefore, the random worlds model cannot formalize statistical knowledge in our sense.

In research of philosophical logic, [32, 2] formalize the idea that when a random value has various possible probability distributions, those distributions should be represented on different possible worlds. Unlike our work, however, they do not model statistical significance or explore accessibility relations.

Independently of our work, French et al. [15] propose a probability model for a dynamic epistemic logic where each world is associated with a (subjective) probability distribution *over the universe* and may have a different probability for a propositional variable to be true. This is different from our distributional Kripke model in that their model does not associate each world with a probability distribution of observable variables, hence deals with neither non-deterministic inputs, divergence-based accessibility relations, nor statistical significance.

Epistemic logic for privacy properties. Epistemic logic has been used to formalize and reason about privacy properties, including anonymity [37, 21, 35, 17, 25, 13, 4, 6], role-interchangeability [34], receipt-freeness of electronic voting protocols [23, 4], and its extension called coercion-resistance [30]. Unlike our formalization in Section 7, however, these do not regard possible worlds as probability distributions and cannot formalize privacy properties with a statistical significance.

Logical approaches to differential privacy. There have been studies that formalize differential privacy using logics, such as Hoare logic [3] and HyperPCTL [1]. Compared to StatEL, these formalizations need to explicitly describe inequalities of probabilities without much abstraction, hence the formulas are more complicated. In addition, none of them formalizes the situation with finite sample sizes or statistical significance.

9 Conclusion

We introduced statistical epistemic logic (StatEL) to describe statistical knowledge, and showed its stochastic semantics based on the distributional Kripke model. By using StatEL we introduced (α, n)-statistical secrecy with a significance level α and a sample size n, and showed that StatEL is useful to formalize hypothesis testing and differential privacy in a simple way. As shown in [24], StatEL can also express certain properties of statistical machine learning.

In our ongoing work, we extend StatEL to deal with the security of cryptography based on computational complexity theory. As for future work, we will extend this logic with temporal modality and give its axiomatization. Our future work includes an extension of StatEL to formalize the quantitative notions of anonymity [9] and asymptotic anonymity [26]. We are also interested in clarifying the relationships between our distributional Kripke model and the main stream probabilistic epistemic logic assigning probabilities to worlds. Furthermore, we plan to develop statistical epistemic logic for process calculi in an analogous way to [6, 22, 10, 8], and to investigate the relationships between statistical epistemic logic and bisimulation metrics analogously to [7].

Acknowledgments

I would like to thank the reviewers for their helpful and insightful comments. I am also grateful to Ken Mano, Gergei Bana, and Ryuta Arisaka for their useful comments on preliminary manuscripts.

References

1. Ábrahám, E., Bonakdarpour, B.: Hyperpctl: A temporal logic for probabilistic hyperproperties. In: Proc. QEST. pp. 20–35 (2018)
2. Bana, G.: Models of objective chance: An analysis through examples. In: Making it Formally Explicit. pp. 43–60. Springer International Publishing (2017). https://doi.org/10.1007/978-3-319-55486-0_3
3. Barthe, G., Gaboardi, M., Arias, E.J.G., Hsu, J., Kunz, C., Strub, P.: Proving differential privacy in hoare logic. In: Proc. CSF. pp. 411–424 (2014)
4. Baskar, A., Ramanujam, R., Suresh, S.P.: Knowledge-based modelling of voting protocols. In: Proc. TARK. pp. 62–71 (2007)
5. Biondi, F., Kawamoto, Y., Legay, A., Traonouez, L.: Hybrid statistical estimation of mutual information and its application to information flow. Formal Asp. Comput. **31**(2), 165–206 (2019). https://doi.org/10.1007/s00165-018-0469-z
6. Chadha, R., Delaune, S., Kremer, S.: Epistemic logic for the applied pi calculus. In: Proc. FMOODS/FORTE. pp. 182–197 (2009). https://doi.org/10.1007/978-3-642-02138-1_12

7. Chatzikokolakis, K., Gebler, D., Palamidessi, C., Xu, L.: Generalized bisimulation metrics. In: Proc. CONCUR. pp. 32–46 (2014). https://doi.org/10.1007/978-3-662-44584-6_4
8. Chatzikokolakis, K., Knight, S., Palamidessi, C., Panangaden, P.: Epistemic strategies and games on concurrent processes. ACM Trans. Comput. Logic **13**(4), 28:1–28:35 (2012). https://doi.org/10.1145/2362355.2362356
9. Chatzikokolakis, K., Palamidessi, C., Panangaden, P.: Anonymity protocols as noisy channels. Inf. Comput. **206**(2–4), 378–401 (2008). https://doi.org/10.1016/j.ic.2007.07.003
10. Dechesne, F., Mousavi, M., Orzan, S.: Operational and epistemic approaches to protocol analysis: Bridging the gap. In: Proc. LPAR. pp. 226–241 (2007)
11. Dwork, C.: Differential privacy. In: Proc. ICALP. pp. 1–12 (2006)
12. Dwork, C., Roth, A., et al.: The algorithmic foundations of differential privacy. Foundations and Trends® in Theoretical Computer Science **9**(3–4), 211–407 (2014)
13. van Eijck, J., Orzan, S.: Epistemic verification of anonymity. Electr. Notes Theor. Comput. Sci. **168**, 159–174 (2007). https://doi.org/10.1016/j.entcs.2006.08.026
14. Fagin, R., Halpern, J., Moses, Y., Vardi, M.: Reasoning about Knowledge. The MIT Press (1995)
15. French, T., Gozzard, A., Reynolds, M.: Dynamic aleatoric reasoning in games of bluffing and chance. In: Proc. AAMAS. pp. 1964–1966 (2019)
16. F.R.S., K.P.: On the criterion that a given system of deviations from the probable in the case of a correlated system of variables is such that it can be reasonably supposed to have arisen from random sampling. The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science **50**(302), 157–175 (1900)
17. Garcia, F.D., Hasuo, I., Pieters, W., van Rossum, P.: Provable anonymity. In: Proc. FMSE. pp. 63–72 (2005). https://doi.org/10.1145/1103576.1103585
18. Gordon, A.D., Henzinger, T.A., Nori, A.V., Rajamani, S.K.: Probabilistic programming. In: Proc. FOSE. pp. 167–181 (2014). https://doi.org/10.1145/2593882.2593900
19. Halpern, J.Y.: An analysis of first-order logics of probability. Artif. Intell. **46**(3), 311–350 (1990). https://doi.org/10.1016/0004-3702(90)90019-V
20. Halpern, J.Y.: Reasoning about uncertainty. The MIT press (2003)
21. Halpern, J.Y., O'Neill, K.R.: Anonymity and information hiding in multiagent systems. In: Proc. CSFW. pp. 75–88 (2003)
22. Hughes, D., Shmatikov, V.: Information hiding, anonymity and privacy: a modular approach. J. of Comp. Security **12**(1), 3–36 (2004)
23. Jonker, H.L., Pieters, W.: Receipt-freeness as a special case of anonymity in epistemic logic. In: Proc. Workshop On Trustworthy Elections (WOTE'06) (June 2006)
24. Kawamoto, Y.: Towards logical specification of statistical machine learning. In: Proc. SEFM (2019), to appear
25. Kawamoto, Y., Mano, K., Sakurada, H., Hagiya, M.: Partial knowledge of functions and verification of anonymity (in Japanese). Transactions of the Japan Society for Industrial and Applied Mathematics **17**(4), 559–576 (2007). https://doi.org/10.11540/jsiamt.17.4_559
26. Kawamoto, Y., Murakami, T.: On the anonymization of differentially private location obfuscation. In: Proc. ISITA. pp. 159–163 (2018)
27. Kawamoto, Y., Murakami, T.: Local obfuscation mechanisms for hiding probability distributions. In: Proc. ESORICS (2019), to appear
28. Kooi, B.P.: Probabilistic dynamic epistemic logic. Journal of Logic, Language and Information **12**(4), 381–408 (2003). https://doi.org/10.1023/A:1025050800836

29. Kripke, S.A.: Semantical analysis of modal logic i normal modal propositional calculi. Mathematical Logic Quarterly **9**(5-6), 67–96 (1963)
30. Küsters, R., Truderung, T.: An epistemic approach to coercion-resistance for electronic voting protocols. In: Proc. S&P. pp. 251–266 (2009). https://doi.org/10.1109/SP.2009.13
31. Legay, A., Delahaye, B., Bensalem, S.: Statistical model checking: An overview. In: Proc. RV. pp. 122–135 (2010). https://doi.org/10.1007/978-3-642-16612-9_11
32. Lewis, D.: A subjectivist's guide to objective chance. In: Studies in Inductive Logic and Probability, Volume II, pp. 263–293. Berkeley: University of California Press (1980)
33. Lin, J.: Divergence measures based on the shannon entropy. IEEE Transactions on Information Theory **37**(1), 145–151 (1991). https://doi.org/10.1109/18.61115
34. Mano, K., Kawabe, Y., Sakurada, H., Tsukada, Y.: Role interchange for anonymity and privacy of voting. J. Log. Comput. **20**(6), 1251–1288 (2010). https://doi.org/10.1093/logcom/exq013
35. van der Meyden, R., Su, K.: Symbolic model checking the knowledge of the dining cryptographers. In: Proc. CSFW. p. 280 (2004). https://doi.org/10.1109/CSFW.2004.19
36. Sen, K., Viswanathan, M., Agha, G.: Statistical model checking of black-box probabilistic systems. In: Proc. CAV. pp. 202–215 (2004). https://doi.org/10.1007/978-3-540-27813-9_16
37. Syverson, P.F., Stubblebine, S.G.: Group principals and the formalization of anonymity. In: World Congress on Formal Methods (1). pp. 814–833 (1999). https://doi.org/10.1007/3-540-48119-2_45
38. Vaserstein, L.: Markovian processes on countable space product describing large systems of automata. Probl. Peredachi Inf. **5**(3), 64–72 (1969)
39. von Wright, G.H.: An Essay in Modal Logic. Amsterdam: North-Holland Pub. Co. (1951)
40. Younes, H.L.: Verification and planning for stochastic processes with asynchronous events. Ph.D. thesis, Carnegie Mellon University (2005)

A Properties of Probability Quantification

In this section we present the proofs for properties of probability quantification.

Proposition 1 (Negation with probability quantification) *For any world w in a model $\mathfrak{M}$ and any static formula ψ, we have:*

1. $\mathfrak{M}, w \models \neg\, \mathbb{P}_I\, \psi$ *iff* $\mathfrak{M}, w \models \mathbb{P}_{I^c}\, \psi$
2. $\mathfrak{M}, w \models \mathbb{P}_I \neg\psi$ *iff* $\mathfrak{M}, w \models \mathbb{P}_{\overline{I}}\, \psi$.

Proof. We show the first claim as follows. By the definition of semantics, $\mathfrak{M}, w \models \neg\, \mathbb{P}_I\, \psi$ is logically equivalent to $\Pr\left[s \xleftarrow{\$} w : \ s \models \psi\right] \not\in I$, which is equivalent to $\Pr\left[s \xleftarrow{\$} w : \ s \models \psi\right] \in \ I^c$, namely, $\mathfrak{M}, w \models \mathbb{P}_{I^c}\, \psi$.

Next we show the second claim as follows. By the definition of semantics, $\mathfrak{M}, w \models \mathbb{P}_I \neg\psi$ is logically equivalent to $1 - \Pr\left[s \xleftarrow{\$} w : \ s \models \psi\right] \in I$, i.e., $\Pr\left[s \xleftarrow{\$} w : \ s \models \psi\right] \in \overline{I}$. This is equivalent to $\mathfrak{M}, w \models \mathbb{P}_{\overline{I}}\, \psi$. □

B Properties of the Epistemic Operators

In this section we present properties of our epistemic operators and their proofs.

Proposition 2 (Minimal properties) *For any distributional Kripke model $\mathfrak{M}$, any $\varphi, \varphi_0, \varphi_1 \in \mathcal{F}$, and any $a \in \mathcal{A}$, we have:*

(**N**) *necessitation:* $\mathfrak{M} \models \varphi$ *implies* $\mathfrak{M} \models \mathsf{K}_a\, \varphi$
(**K**) *distribution:* $\mathfrak{M} \models \mathsf{K}_a(\varphi_0 \rightarrow \varphi_1) \rightarrow (\mathsf{K}_a\, \varphi_0 \rightarrow \mathsf{K}_a\, \varphi_1)$.

Proof. We first show (**N**) necessitation rule as follows. Assume that $\mathfrak{M} \models \varphi$. Then for any world w in $\mathfrak{M}$, we have $\mathfrak{M}, w \models \varphi$. Hence $\mathfrak{M}, w \models \mathsf{K}_a\, \varphi$. Therefore the necessitation rule holds.

Next we show (**K**) distribution axiom as follows. Let w be a possible world in $\mathfrak{M}$. Assume that $\mathfrak{M}, w \models \mathsf{K}_a(\varphi_0 \rightarrow \varphi_1)$, and that $\mathfrak{M}, w \models \mathsf{K}_a\, \varphi_0$. Let w' be any world such that $(w, w') \in \mathcal{R}_{a,\varepsilon}$. Then we have $\mathfrak{M}, w' \models \varphi_0 \rightarrow \varphi_1$ and $\mathfrak{M}, w' \models \varphi_0$, hence $\mathfrak{M}, w' \models \varphi_1$. Thus we have $\mathfrak{M}, w \models \mathsf{K}_a\, \varphi_1$. Therefore we obtain $\mathfrak{M} \models \mathsf{K}_a(\varphi_0 \rightarrow \varphi_1) \rightarrow (\mathsf{K}_a\, \varphi_0 \rightarrow \mathsf{K}_a\, \varphi_1)$. □

Proposition 3 (Properties with divergence-based accessibility) *Let $a \in \mathcal{A}$ and $\varepsilon \geq \varepsilon' \geq 0$. For any distributional Kripke model $\mathfrak{M}$ with a divergence-based accessibility relation $\mathcal{R}_{a,\varepsilon}$ and any $\varphi \in \mathcal{F}$, we have:*

(**T**) *reflexivity:* $\mathfrak{M} \models \mathsf{K}_{a,\varepsilon}\, \varphi \rightarrow \varphi$
($\geq$) *comparison of observability:* $\mathfrak{M} \models \mathsf{K}_{a,\varepsilon}\varphi \rightarrow \mathsf{K}_{a,\varepsilon'}\varphi$.

If $\mathcal{R}_{a,\varepsilon}$ is symmetric (e.g., based on the Jensen-Shannon divergence [33]) then:

(**B**) *symmetry:* $\mathfrak{M} \models \varphi \rightarrow \mathsf{K}_{a,\varepsilon}\, \mathsf{P}_{a,\varepsilon}\, \varphi$.

Proof. Let w be a possible world in $\mathfrak{M}$.

We first show (**T**) reflexivity as follows. Assume that $\mathfrak{M}, w \models \mathsf{K}_{a,0}\, \varphi$. By $(w, w) \in \mathcal{R}_{a,0}$, we have $\mathfrak{M}, w \models \varphi$. Therefore, we obtain $\mathfrak{M} \models \mathsf{K}_{a,\varepsilon}\, \varphi \rightarrow \varphi$.

Next we show ($\geq$) comparison of observability as follows. Assume that $\mathfrak{M}, w \models \mathsf{K}_{a,\varepsilon}\, \varphi$. Let w' be any world such that $(w, w') \in \mathcal{R}_{a,\varepsilon}$. Then $\mathfrak{M}, w' \models \varphi$. By $\varepsilon' \leq \varepsilon$ and the definition of $\mathcal{R}_{a,\varepsilon}$, we have $\mathcal{R}_{a,\varepsilon'} \subseteq \mathcal{R}_{a,\varepsilon}$, hence $(w, w') \in \mathcal{R}_{a,\varepsilon'}$. Then $\mathfrak{M}, w \models \mathsf{K}_{a,\varepsilon'}\varphi$. Therefore we obtain $\mathfrak{M} \models \mathsf{K}_{a,\varepsilon}\varphi \rightarrow \mathsf{K}_{a,\varepsilon'}\varphi$.

Finally, we show (**B**) symmetry when $\mathcal{R}_{a,\varepsilon}$ is symmetric. Assume that $\mathfrak{M}, w \models \varphi$. Let w' be any world such that $(w, w') \in \mathcal{R}_{a,\varepsilon}$. Since $\mathcal{R}_{a,\varepsilon}$ is symmetric, we have $(w', w) \in \mathcal{R}_{a,\varepsilon}$. By $\mathfrak{M}, w \models \varphi$, we obtain $\mathfrak{M}, w' \models \mathsf{P}_{a,\varepsilon}\, \varphi$. Hence $\mathfrak{M}, w \models \mathsf{K}_{a,\varepsilon}\, \mathsf{P}_{a,\varepsilon}\, \varphi$. Therefore we obtain $\mathfrak{M} \models \varphi \rightarrow \mathsf{K}_{a,\varepsilon}\, \mathsf{P}_{a,\varepsilon}\, \varphi$. □

Proposition 4 (Properties with metric-based accessibility) *Let $a \in \mathcal{A}$ and $\varepsilon, \varepsilon' \geq 0$. For any distributional Kripke model $\mathfrak{M}$ with a metric-based accessibility relation $\mathcal{R}_{a,\varepsilon}$ and any $\varphi \in \mathcal{F}$, we have* (**T**)*reflexivity,* (**B**)*symmetry, and:*

(**4q**) *quantitative transitivity:* $\mathfrak{M} \models \mathsf{K}_{a,\varepsilon+\varepsilon'}\varphi \rightarrow \mathsf{K}_{a,\varepsilon}\, \mathsf{K}_{a,\varepsilon'}\varphi$
(**5q**) *relaxed Euclidean:* $\mathfrak{M} \models \mathsf{P}_{a,\varepsilon}\varphi \rightarrow \mathsf{K}_{a,\varepsilon'}\mathsf{P}_{a,\varepsilon+\varepsilon'}\varphi$.

If the agent has an unlimited capability of observation (i.e., $\varepsilon = \varepsilon' = 0$), then:

(**4**) *transitivity:* $\mathfrak{M} \models \mathsf{K}_{a,0}\,\varphi \rightarrow \mathsf{K}_{a,0}\,\mathsf{K}_{a,0}\,\varphi$
(**5**) *Euclidean:* $\mathfrak{M} \models \mathsf{P}_{a,0}\,\varphi \rightarrow \mathsf{K}_{a,0}\,\mathsf{P}_{a,0}\,\varphi$.

Proof. Since a metric satisfies the definition of a divergence (in Section 2), a metric-based accessibility relation is also a divergence-based accessibility relation. Therefore we obtain (**T**) reflexivity and (**B**) symmetry from Proposition 3.

Next we show (**4q**) quantitative transitivity as follows. Let w be a possible world in $\mathfrak{M}$. Assume that $\mathfrak{M}, w \models \mathsf{K}_{a,\varepsilon+\varepsilon'}\varphi$. Let w' be any world such that $(w, w') \in \mathcal{R}_{a,\varepsilon}$, and w'' be any world such that $(w', w'') \in \mathcal{R}_{a,\varepsilon'}$. By definition, we have $D(\sigma_w(x) \parallel \sigma_{w'}(x)) \leq \varepsilon$ and $D(\sigma_{w'}(x) \parallel \sigma_{w''}(x)) \leq \varepsilon'$. By the subadditivity of the divergence D, we have $D(\sigma_w(x) \parallel \sigma_{w''}(x)) \leq D(\sigma_w(x) \parallel \sigma_{w'}(x)) + D(\sigma_{w'}(x) \parallel \sigma_{w''}(x)) \leq \varepsilon + \varepsilon'$, hence $(w, w'') \in \mathcal{R}_{a,\varepsilon+\varepsilon'}$. Then it follows from $\mathfrak{M}, w \models \mathsf{K}_{a,\varepsilon+\varepsilon'}\varphi$ that $\mathfrak{M}, w'' \models \varphi$. By the definition of w'', we obtain $\mathfrak{M}, w \models \mathsf{K}_{a,\varepsilon}\,\mathsf{K}_{a,\varepsilon'}\varphi$. Therefore we have $\mathfrak{M} \models \mathsf{K}_{a,\varepsilon+\varepsilon'}\varphi \rightarrow \mathsf{K}_{a,\varepsilon}\,\mathsf{K}_{a,\varepsilon'}\varphi$.

We next show (**5q**) relaxed Euclidean as follows. Let w be a possible world in $\mathfrak{M}$. Assume that $\mathfrak{M}, w \models \mathsf{P}_{a,\varepsilon}\varphi$. Then there exists a world w' such that $\mathfrak{M}, w' \models \varphi$ and $(w, w') \in \mathcal{R}_{a,\varepsilon}$. Let w'' be any world such that $(w, w'') \in \mathcal{R}_{a,\varepsilon'}$. Since $\mathcal{R}_{a,\varepsilon'}$ is a metric-based accessibility relation, it is symmetric, hence $(w'', w) \in \mathcal{R}_{a,\varepsilon'}$. Then by $(w, w') \in \mathcal{R}_{a,\varepsilon}$, we obtain $(w'', w') \in \mathcal{R}_{a,\varepsilon+\varepsilon'}$. Hence $\mathfrak{M}, w \models \mathsf{K}_{a,\varepsilon'}\mathsf{P}_{a,\varepsilon+\varepsilon'}\varphi$. Therefore we obtain $\mathfrak{M} \models \mathsf{P}_{a,\varepsilon}\varphi \rightarrow \mathsf{K}_{a,\varepsilon'}\mathsf{P}_{a,\varepsilon+\varepsilon'}\varphi$.

Finally, for $\varepsilon = \varepsilon' = 0$, (**4**) transitivity and (**5**) Euclidean respectively follow from (**4q**) quantitative transitivity and (**5q**) relaxed Euclidean. □