informs
http://pubsonline.informs.org/journal/opre

OPERATIONS RESEARCH
Vol. 69, No. 5, September–October 2021, pp. 1575–1590
ISSN 0030-364X (print), ISSN 1526-5463 (online)

Methods

Private Sequential Learning

John N. Tsitsiklis,[a] Kuang Xu,[b] Zhi Xu[a]

[a] Laboratory for Information and Decision Systems (LIDS), Massachusetts Institute of Technology, Cambridge, Massachusetts 02139; [b] Graduate School of Business, Stanford University, Stanford, California 94305

Contact: jnt@mit.edu, https://orcid.org/0000-0003-2658-8239 (JNT); kuangxu@stanford.edu, https://orcid.org/0000-0002-2221-1648 (KX); zhixu@mit.edu, https://orcid.org/0000-0002-1421-2309 (ZX)

Received: May 13, 2018
Revised: July 29, 2019; January 6, 2020
Accepted: February 13, 2020
Published Online in Articles in Advance: January 14, 2021

Subject Classification: decision analysis: sequential; statistics: estimation; probability: stochastic model applications
Area of Review: Stochastic Models

https://doi.org/10.1287/opre.2020.2021

Abstract. We formulate a private learning model to study an intrinsic tradeoff between privacy and query complexity in sequential learning. Our model involves a learner who aims to determine a scalar value v^* by sequentially querying an external database and receiving binary responses. In the meantime, an adversary observes the learner's queries, although not the responses, and tries to infer from them the value of v^*. The objective of the learner is to obtain an accurate estimate of v^* using only a small number of queries while simultaneously protecting his or her privacy by making v^* provably difficult to learn for the adversary. Our main results provide tight upper and lower bounds on the learner's query complexity as a function of desired levels of privacy and estimation accuracy. We also construct explicit query strategies whose complexity is optimal up to an additive constant.

Keywords: sequential learning • privacy • bisection algorithm

1. Introduction

Organizations and individuals often rely on relevant data to solve decision problems. Sometimes such data are beyond the immediate reach of a decision maker and must be acquired by interacting with an external entity or environment. However, these interactions may be monitored by a third-party adversary and subject the decision maker to potential privacy breaches, a possibility that has become increasingly prominent as information technologies and tools for data analytics advance.

This paper studies a decision maker, henceforth referred to as the *learner*, who acquires data from an external entity in an interactive fashion by submitting sequential queries. The *interactivity* benefits the learner by enabling him or her to tailor future queries based on past responses and thus reduce the number of queries needed and, at the same time, exposes the learner to substantial privacy risk: the more the learner's queries depend on past responses, the easier it might be for an adversary to use the observed queries to infer those past responses. Our main objective is to articulate and understand an intrinsic *privacy versus query complexity tradeoff* in the context of such a private sequential learning model.

We begin with an informal description of the model. A *learner* would like to determine the value of a scalar v^* referred to as the *true value* that lies in a bounded subset of $\mathbb{R}$, for example, the interval [0,1). To search for v^*, the learner must interact with an external database through sequentially submitted queries: at step k, the learner submits a query $q_k \in \mathbb{R}$ and receives a binary response r_k, where $r_k = 1$ if $v^* \geq q_k$ and $r_k = 0$ otherwise. The interaction is sequential in the sense that the learner may choose a query depending on the responses to all previous queries. Meanwhile there is an *adversary* who eavesdrops on the learner's actions: the adversary observes all the learner's queries q_k, but not the responses, and tries to use these queries to estimate the true value v^*. The learner's goal is to submit queries in such a way that he or she can learn v^* within a prescribed error tolerance while v^* cannot be accurately estimated by the adversary with high confidence. The learner's goal is easily attained by submitting an unlimited number of queries, in which case the queries need not depend on the past responses and hence reveal no information to the adversary. Our quest is, however, to understand the *least number of queries* that the learner needs to submit in order to successfully retain privacy. Is the query complexity significantly different from the case where privacy constraints are absent? How does it vary as a function of the levels of accuracy and privacy? Is there a simple and yet efficient query strategy that the learner can adopt? Our main results address these questions.

1.1. Motivating Examples

We discuss three examples that provide some context for our model. Although the examples are stylized, they are intended to motivate and provide insights on the general concept of privacy preservation in sequential learning.

Example 1 (Learning an Optimal Price). A firm is to release a new product and would like to identify a revenue-maximizing price p^* before the product launch. The firm believes that there is an underlying unknown demand function $g(p)$ and that the revenue function $f(p) = g(p) \times p$ is a strictly concave and differentiable function of the price p. A sequential learning process is used to identify p^* over a series of epochs: in epoch k, the firm assesses how the market responds to a test price p_k and receives a binary feedback as to whether $f'(p_k) \geq 0$ or $f'(p_k) < 0$. This may be achieved, for instance, by contracting a consulting firm to conduct market surveys on the price *sensitivity* around p_k. The survey could ask for a customer's willingness to purchase the product at price levels p and $p + \xi$, for some small ξ, from which the demand $g(p)$ and its gradient can be estimated. Using the chain rule, this information can then be converted to an estimate of the sign of $f'(p)$. The firm would like to be able to estimate p^* with reasonable accuracy after a small number of epochs but is wary that a competitor might be able to observe the surveys, either by purposely participating in them or by interviewing survey participants, and then deduce the value of p^* ahead of the product launch. In the context of private sequential learning, the firm is the learner, the competitor is the adversary, the revenue-maximizing price is the true value, and the test prices are the queries. The binary response on the revenue's price sensitivity indicates whether the revenue-maximizing price is less than the current test price.

Example 2 (Learning Consumer Preferences). e-Commence firms such as Amazon are often incentivized to learn consumer preferences that could later be used for devising personalized promotions. Consider a consumer who is looking for an item with an ideal scalar product feature (e.g., size) on an online merchant's platform. Although the consumer does not initially know the optimal feature value, when presented with a specific product, he or she will be able to assess whether the ideal value is greater than the current option (e.g., if the current option is too large or too small). The consumer browses different products in a sequential manner and hopes to eventually narrow in on the ideal feature value. Meanwhile, the platform sees all the products browsed by the consumer but does not directly observe his or her internal assessments. Our model can be used to investigate how a privacy-aware consumer can browse products in such a manner that will eventually allow him or her to identify the ideal product feature value while preventing the platform from confidently inferring that value from his or her browsing history.

Example 3 (Online Optimization with Private Weights). In Examples 1 and 2, the adversary is a third-party that does not observe the responses to the queries. We now provide a different example in which the adversary is the database to which queries are submitted and thus has partial knowledge of the responses. Consider a learner who wishes to identify the maximizer x^* of a function $f(x) = \sum_{i=1}^{m} \alpha_i f_i(x)$ over some bounded interval $\mathcal{X} \subset \mathbb{R}$, where $\{f_i(\cdot)\}_{1 \leq i \leq m}$ is a collection of strictly concave differentiable constituent functions, and $\{\alpha_i\}_{1 \leq i \leq m}$ are positive (private) weights representing the importance that the learner associates with each constituent function. The learner knows the weights but does not have information about the constituent functions; such knowledge is to be acquired by querying an external database. During epoch k, the learner submits a test value x_k and receives from the database the derivatives of all constituent functions at x_k, $\{f_i'(x_k)\}_{1 \leq i \leq m}$. Using the weights, the learner can then compute the derivative $f'(x_k)$, whose sign serves as a binary indicator of the position of the maximizer x^* relative to the current test value. The database, which possesses complete information about the constituent functions but does not know the weights, would like to infer from the learner's querying pattern the maximizing value x^* or possibly the weights themselves. Although the previously mentioned model appears to differ from the previous two examples, it turns out that the modeling methodology and query strategies that we develop can also be applied to this setting. The connection between the two settings is made precise in chapter 2 of Xu (2017).

1.2. Preview of the Main Result

We now preview our main result. Let us begin by introducing some additional notation. Recall that both the learner and the adversary aim to obtain estimates that are close to a true value $v^* \in [0, 1)$. We denote by $\epsilon/2$ and $\delta/2$ the estimation error that the learner and the adversary are willing to tolerate, respectively. We will use a privacy parameter $L \in \mathbb{N}$ to quantify the learner's level of privacy at the end of the learning process: the learner's privacy level is L if the adversary can successfully approximate the true value within an error of $\delta/2$ with probability at most $1/L$, so higher values of L correspond to enhanced privacy. (A precise, formal definition of *privacy level*, in terms of L, will be provided in Section 3.2.) A private query strategy for the learner, then, must be able to produce an estimate of the true value within an error of at most $\epsilon/2$ while simultaneously guaranteeing that the desired privacy level L holds against the adversary.

Our main objective is to quantify the *query complexity* of private sequential learning $N^*(\epsilon, \delta, L)$, defined as the minimum number of queries needed for a private learner strategy, under a given set of parameters ϵ, δ, and L. Specifically, we will focus on the regime where $2\epsilon < \delta \le 1/L$. The reason for this choice will become clear after a formal introduction of the model, and we will revisit it at the beginning of Section 4. In this regime, we have the following upper and lower bounds on query complexity:

1. We establish an upper bound[1] of $\log(1/L\epsilon) + 2L$ by explicitly constructing a private learner strategy, which applies for any δ in the range $(2\epsilon, 1/L]$.
2. We establish a lower bound of $\log(\delta/\epsilon) + 2L - 4$ by characterizing the amount of information available to the adversary.

We note that our bounds are tight in the sense that when the adversary's accuracy requirement is as loose as possible (i.e., $\delta = 1/L$), the upper bound matches the lower bound up to an additive constant equal to four. Furthermore, comparing with the elementary lower bound of $\log(1/\epsilon)$ for the case where privacy is not a concern, we see that the extra effort necessary to guarantee a privacy level L is at most an additive factor of $2L$.

To put our results in context, we examine in Section 5 two simple strategies situated at two extreme points of the privacy–efficiency tradeoff curve. On the one hand, the classical bisection search algorithm achieves an optimal query complexity of $\log(1/\epsilon)$, but it completely reveals the responses of past queries and is hence almost never private. On the other hand, the learner could use a nonadaptive *ϵ-dense* strategy that places $1/\epsilon$ equally spaced queries throughout the unit interval. The nonadaptive nature of this strategy allows the learner to always be private, but its query complexity is significantly worse. Our results and policies essentially aim to understand the optimal tradeoff between these two extremes for any given privacy level L.

1.3. Related Work

In the absence of a privacy constraint, the problem of identifying a value within a compact interval through (possibly noisy) binary feedback is a classical problem arising in domains such as coding theory (Horstein 1963) and root finding (Waeber et al. 2013). It is well known that the bisection algorithm achieves the optimal query complexity of $\log(1/\epsilon)$ (Waeber et al. 2013), where $\epsilon > 0$ is the error tolerance. In contrast, to the best of our knowledge, the question of how to preserve a learner's privacy when his or her actions are fully observed by an adversary and what the resulting query complexity would be has received relatively little attention in the literature.

Related to our work, in spirit, is the body of literature on differential privacy (Dwork et al. 2006, Dwork and Roth 2014), a concept that has been applied in statistics (Wasserman and Zhou 2010, Smith 2011, Duchi et al. 2016) and learning theory (Raskhodnikova et al. 2008, Chaudhuri and Hsu 2011, Blum et al. 2013, Feldman and Xiao 2014). Differential privacy mandates that the output distribution of an algorithm be insensitive under certain perturbations of the input data. For instance, Jain et al. (2012) study regret minimization in an online optimization problem while ensuring differential privacy, in the sense that the distribution of the sequence of solutions remains nearly identical when any one of the functions being optimized is perturbed. In Dwork et al. (2012), the authors study differential privacy in data analysis when multiple analysts query the same database. More recently, adaptive versions of differentially private data analysis have been studied (Dwork et al. 2015a,b,c; Cummings et al. 2016). Notably, our work departs from this literature by using a goal-oriented privacy framework: our definition of privacy measures the adversary's ability to perform a *specific* inferential goal. In contrast, differential privacy aims to prevent an adversary from performing *any* meaningful inference. As such, the goal-oriented privacy framework leads to substantially more efficient decision policies, whereas differential privacy offers stronger privacy protection for settings where the privacy goal is not clear a priori, at the cost of an increased efficiency loss. In this regard, our approach echoes a number of recent papers that also employ a goal-oriented privacy framework (Fanti et al. 2015, Liao et al. 2018, Tsitsiklis and Xu 2018).

In a different model, Tsitsiklis and Xu (2018) study the issue of privacy in a sequential decision problem, where an agent attempts to reach a particular node in a graph, traversing it in a way that obfuscates the agent's intended destination, against an adversary who observes the agent's past trajectory. The authors show that the probability of a correct prediction by the adversary is inversely proportional to the time it takes for the agent to reach his or her destination. Similar to the setting of Tsitsiklis and Xu (2018), the learner in our model also plays against a powerful adversary who observes all past actions. However, a major new element is that the learner in our model strives to *learn* a piece of information of which the learner himself or herself has no prior knowledge, in contrast to the agent in Tsitsiklis and Xu (2018), who tries to conceal private information already in his or her possession. In a way, the central conflict of trying to learn something while preventing others from learning the same information sets the present work apart from the extant literature.

Finally, our model is close in spirit to private information retrieval problems in the field of cryptography (Kushilevitz and Ostrovsky 1997, Chor et al. 1998, Gasarch 2004). In these problems, a learner wishes to retrieve an item from some location i in a database in such a manner that the database obtains no information on the value of i, where the latter requirement can be either information theoretic or based on computational hardness assumptions. Compared with this line of literature, our privacy requirement is substantially weaker: the adversary may still obtain *some* information on the true value. This relaxation of the privacy requirement allows the learner to deploy richer and more sample-efficient query strategies.

1.4. Organization

The remainder of this paper is organized as follows. We formally introduce the private sequential learning model in Section 2. In Section 3, we motivate and discuss private learner strategies. Our main results are stated in Section 4. Before delving into the proofs, we examine in Section 5 three examples of learner strategies that provide further insight into the structure of the problem. Sections 6 and 7 are devoted to the proof of the upper and lower bounds in our main theorem, respectively. We conclude in Section 8, where we also describe some interesting variations of our model, which are further elaborated on in Appendices A and B.

2. The Private Sequential Learning Model

Here we formally introduce our private sequential learning model. The model involves a *learner* who aims to determine a particular *true value* v^*. The true value is a scalar in some bounded subset of $\mathbb{R}$. Without loss of generality, we assume that v^* belongs to the interval[2] $[0,1)$ and that the learner knows that this is the case. The true value is stored in an external database. In order to learn the true value, the learner interacts with the database by submitting queries as follows: at each step k, the learner submits a *query* $q_k \in [0,1)$ and receives from the database a *response* r_k indicating whether v^* is greater than or equal to the query value; that is,

$$r_k = \mathbb{I}(v^* \geq q_k),$$

where $\mathbb{I}(\cdot)$ stands for the indicator function. Furthermore, each query is allowed to depend on the responses to previous queries through a learner strategy to be defined shortly.

Denote by N the total number of learner queries and by $\epsilon > 0$ the learner's desired accuracy. After receiving the responses to N queries, the learner aims to produce an estimate $\hat{x}$ for v^* that satisfies

$$|\hat{x} - v^*| \leq \frac{\epsilon}{2}.$$

In the meantime, there is an *adversary* who is also interested in learning the true value v^*. The adversary has no access to the database and hence seeks to estimate v^* by free-riding on observations of the learner queries. Let $\delta > 0$ be an accuracy parameter for the adversary. We assume that the adversary can observe the values of the queries but not the responses and knows the learner's query strategy. Based on this information, and after observing all the queries submitted by the learner, the adversary aims to generate an estimate $\hat{x}^a$ for v^* that satisfies

$$|\hat{x}^a - v^*| \leq \frac{\delta}{2}.$$

2.1. Learner Strategy

The queries that the learner submits to the database are generated by a (possibly randomized) *learner strategy* in a sequential manner: the query at step k depends on the queries and their responses up to step $k-1$, as well as on a discrete random variable Y. In particular, the random variable Y allows the learner to randomize, if needed, and we refer to Y as the *random seed*. Without loss of generality, we assume that Y is uniformly distributed over $\{1,2,\ldots,\mathcal{Y}\}$, where $\mathcal{Y}$ is a large integer. Formally, fixing $N \in \mathbb{N}$, a learner strategy ϕ of length N is comprised of two parts:

1. A finite sequence of N query functions $(\phi_1,\ldots,\phi_N)$, where each ϕ_k is a mapping that takes as input the values of the first $k-1$ queries submitted, the corresponding responses, and the realized value of Y and outputs the kth query q_k.
2. An estimation function ϕ^E, which takes as input the N queries submitted, the corresponding responses, and the realized value of Y and outputs the final estimate $\hat{x}$ for the true value v^*.

More precisely, we have

1. If $k=1$, then $\phi_1 : \{1,2,\ldots,\mathcal{Y}\} \to [0,1)$ and $q_1 = \phi_1(Y)$; if $k = 2,3,\ldots,N$, then $\phi_k : [0,1)^{k-1} \times \{0,1\}^{k-1} \times \{1,2,\ldots,\mathcal{Y}\} \to [0,1)$ and

$$q_k = \phi_k(q_1,q_2,\ldots,q_{k-1},r_1,r_2,\ldots,r_{k-1},Y);$$

2. Finally, $\phi^E : [0,1)^N \times \{0,1\}^N \times \{1,2,\ldots,\mathcal{Y}\} \to [0,1)$ and $\hat{x} = \phi^E(q_1,q_2,\ldots,q_N,r_1,r_2,\ldots,r_N,Y)$.

Observe that the preceding definition can be simplified: knowing the value of the random seed Y and the responses to the queries is sufficient for reconstructing the values of the queries. As an example, we have $q_2 = \phi_2(q_1,r_1,Y) = \phi_2(\phi_1(Y),r_1,Y) = \phi_2'(r_1,Y)$ for some new function ϕ_2'. Through induction, it then suffices to let the input to ϕ_k be just $r_1, \ldots, r_{k-1}$ and Y. This leads to an alternative, simpler definition of learner strategies:

1. If $k = 1$, then $\phi_1 : \{1, 2, \ldots, \mathcal{Y}\} \to [0, 1)$ and $q_1 = \phi_1(Y)$; if $k = 2, 3, \ldots, N$, then $\phi_k : \{0, 1\}^{k-1} \times \{1, 2, \ldots, \mathcal{Y}\} \to [0, 1)$ and $q_k = \phi_k(r_1, r_2, \ldots, r_{k-1}, Y)$;

2. Finally, $\phi^E : \{0,1\}^N \times \{1,2,\ldots,\mathcal{Y}\} \to [0,1)$ and $\hat{x} = \phi^E(r_1, r_2, \ldots, r_N, Y)$.

In what follows, we adopt the latter, simpler definition. In addition, we will consider learner strategies that submit distinct queries because repeated queries do not provide additional information to the learner. We will denote by Φ_N the set of all learner strategies of length N, defined previously.

Fix a learner strategy $\phi \in \Phi_N$. To clarify the dependence on the random seed, for any $x \in [0, 1)$ and $y \in \{1, 2, \ldots, \mathcal{Y}\}$, we will use $\overline{q}(x, y)$ to denote the realization of the sequence of queries $(q_1, q_2, \ldots, q_N)$ when the true value v^* is x and the learner's random seed Y is y. Similarly, we will denote by $\hat{x}(x, y)$ the learner's estimate of the true value when $v^* = x$ and $Y = y$.

2.2. Information Available to the Adversary

We summarize in this section the information available to the adversary. First, the adversary is aware that the true value v^* belongs to $[0, 1)$. Second, we assume that the adversary can observe the values of the queries but not the corresponding responses and that the learner strategy ϕ, including the distribution of the random seed Y, is known to the adversary. In particular, the adversary observes the value of each query q_k, for $k = 1, \ldots, N$, and knows the N mappings $\phi_1, \phi_2, \ldots, \phi_N$. This means that if the adversary had access to the values $r_1, r_2, \ldots, r_{k-1}$ and the realized value of Y, he or she would know exactly what q_k is for step k. These assumptions stem from a worst-case consideration: the privacy guarantees hold even when the adversary knows the learner's strategy and can automatically extend to more practical scenarios where such knowledge may not be exact. Although it may seem that an adversary who sees both the learner strategy and the learner's actions is too powerful to defend against, we will see in the ensuing analysis that the learner is still be able to implement effective and efficient obfuscation by exploiting the randomness in Y.

3. Private Learner Strategies

In this section, we introduce and formally define private learning strategies, the central concept of this paper. Although we will briefly discuss the underlying intuition, further interpretation is provided in Appendix A. As was mentioned in the Introduction, a private learner strategy must always make sure that its estimate is close to the true value v^* while keeping the adversary's probability of accurately estimating v^* sufficiently small. Our goal in this section is to formalize these ideas. To this end, we first introduce in Section 3.1 ways of quantifying the amount of information acquired by the adversary as a function of the learner's queries. This then leads to a precise privacy constraint, presented in Section 3.2.

3.1. Information Set

Recall from Section 2.2 that the adversary knows the values of the queries and the learner strategy. We now convert this knowledge into a succinct representation: the *information set* of the adversary. Fix a learner strategy ϕ. Denote by $\mathcal{Q}(x)$ the set of query sequences that have a positive probability of appearing under ϕ when the true value v^* is equal to x:

$$\mathcal{Q}(x) = \{\overline{q} \in [0, 1)^N : \mathbb{P}_\phi(Q_x = \overline{q}) > 0\}, \tag{1}$$

where Q_x is a vector-valued random variable representing the sequence of learner queries when the true value is equal to x and where $\overline{q}$ ranges over possible realized values; the probability is measured with respect to the randomness in the learner's random seed Y.

Definition 1. Fix $\phi \in \Phi_N$. The information set for the adversary $\mathcal{I}(\overline{q})$ is defined by

$$\mathcal{I}(\overline{q}) = \{x \in [0, 1) : \overline{q} \in \mathcal{Q}(x)\}, \quad \overline{q} \in [0, 1)^N. \tag{2}$$

From the viewpoint of the adversary, the information set represents all possible true values that are consistent with the queries observed. As such, it captures the amount of information that the learner reveals to the adversary.

3.2. The (ϵ, δ, L)-Private Strategies

A private learner strategy should achieve two aims: accuracy and privacy. Accuracy can be captured in a relatively straightforward manner by measuring the absolute distance between the learner's estimate and the true value. An effective measure of the learner's privacy, by contrast, is more subtle because it depends on what the adversary is able to infer. To this end, we develop in this subsection a privacy metric by quantifying the *effective size* of the information set $\mathcal{I}(\overline{q})$ described in Definition 1. Intuitively, because the information set contains all possible realizations of the true value v^*, the larger the information set, the more difficult it is for the adversary to pin down the true value.

The choice of such a metric requires care. As a first attempt, the diameter of the information set $\sup_{y_1, y_2 \in \mathcal{I}(\overline{q})} |y_1 - y_2|$ may appear to be a natural candidate.

Because the adversary has an accuracy parameter of δ, we could impose, as a privacy constraint, that the diameter of $\mathcal{I}(\bar{q})$ be greater than δ. The diameter, however, is not a good metric because it paints an overly optimistic picture for the learner. For example, consider the case where the information set is the union of two intervals of length δ each placed far apart from each other. By setting his or her estimate to be the center of one of the two intervals, chosen at random with equal probabilities, the adversary will have probability 1/2 of correctly predicting the true value, even though the diameter of the information set could be large. As a second attempt, we might consider the Lebesgue measure of the information set. However, it also fails to capture the intended meaning of learner privacy. For example, consider the case where the information set consists of many distantly placed but very small intervals. It is not difficult to see that the adversary would not be able to correctly estimate the true value with high certainty, even if the Lebesgue measure of the set is arbitrarily small.

The shortcomings of the preceding metrics motivate a more refined notion of *effective size* and, in particular, one that would be appropriate for disconnected information sets. To this end, we use set coverability to measure the size of the information set, defined as follows.

Definition 2. Fix $\delta > 0$, $L \in \mathbb{N}$, and a set $\mathcal{E} \subset \mathbb{R}$. We say that a collection of L closed intervals $[a_1, b_1]$, $[a_2, b_2]$, ..., $[a_L, b_L]$ is a (δ, L) *cover* for $\mathcal{E}$ if $\mathcal{E} \subset \bigcup_{1\le j\le L}[a_j, b_j]$ and $b_j - a_j \le \delta$ for all j.

We say that a set $\mathcal{E}$ is (δ, L)-*coverable* if it admits a (δ, L) cover. In addition, we define the δ-cover number of a set $\mathcal{E}$, $C_\delta(\mathcal{E})$, as

$$C_\delta(\mathcal{E}) \triangleq \min\{L \in \mathbb{N} : \mathcal{E} \text{ is } (\delta, L)\text{-coverable}\}. \quad (3)$$

We are now ready to define (ϵ, δ, L)-private learner strategies.

Definition 3 (Private Learner Strategy)**.** Fix $\epsilon > 0$, $\delta > 0$, and $L \ge 2$, with $L \in \mathbb{N}$. A learner strategy $\phi \in \Phi_N$ is (ϵ, δ, L)-*private* if it satisfies the following:

a. *Accuracy constraint:* The learner estimate accurately recovers the true value with probability one:

$$\mathbb{P}(|\hat{x}(x, Y) - x| \le \epsilon/2) = 1, \quad \forall x \in [0, 1),$$

where the probability is measured with respect to the randomness in Y.

b. *Privacy constraint:* For every $x \in [0, 1)$ and every possible sequence of queries $\bar{q} \in \mathcal{Q}(x)$, the δ-cover number of the information set for the adversary, $C_\delta(\mathcal{I}(\bar{q}))$, is at least L, that is,

$$C_\delta\big(\mathcal{I}(\bar{q})\big) \ge L, \quad \forall \bar{q} \in \mathcal{Q}(x). \quad (4)$$

The accuracy constraint requires that a private learner strategy always produce an accurate estimate within the error tolerance ϵ for any possible true value in $[0, 1)$. The privacy constraint controls the size of the information set induced by the sequence of queries generated, and the parameter L can be interpreted as the learner's privacy level: because the intervals used to cover the information set are of length at most δ, each interval can be thought of as representing a plausible guess for the adversary. Therefore, the probability of the adversary successfully estimating the location of v^* is essentially inversely proportional to the number of intervals needed to cover the information set, which is at most $1/L$. It turns out that this intuition can be made precise: in Appendix A, we formally establish the equivalence between $1/L$ and the adversary's probability of correct estimation.

3.3. Worst-Case vs. Bayesian Formulations

Our definition of learner privacy involves worst-case requirements for both the learner and the adversary. In a Bayesian formulation, these are replaced by requirements that only need to hold, on average, under a prior distribution for the value of the unknown target v^*. We formulate such a Bayesian variant in detail in Appendix B and argue that it is complementary, not directly comparable, to our main formulation. On the technical side, in Appendix B, we present a learner strategy that achieves privacy with a query complexity that depends *multiplicatively* on L; in more recent work that follows up on this paper, Xu (2018) establishes a lower bound that shows that such dependence is also tight. We note that not only are the two formulations not comparable, but their analysis is also different: the lower bounds for our base model rely on combinatorial arguments in contrast to information-theoretic arguments for the Bayesian variant (Xu 2018).

4. Main Result

The learner's overall objective is to use a minimal number of queries while satisfying the accuracy and privacy requirements. We state our main theorem in this section, which establishes lower and upper bounds for the query complexity of a private learner strategy as a function of the adversary accuracy δ, learner accuracy ϵ, and learner privacy level L. Recall that Φ_N is the set of learner strategies of length N. We define $N^*(\epsilon, \delta, L)$ as the minimum number of queries needed across all (ϵ, δ, L)-private learner strategies. Thus,

$$N^*(\epsilon, \delta, L) = \min\{N \in \mathbb{N} : \Phi_N \text{ contains at least one } \times(\epsilon, \delta, L)\text{-private strategy}\}. \quad (5)$$

Our result focuses on the regime of parameters where

$$0 < 2\epsilon < \delta \le 1/L. \tag{6}$$

1. Having $2\epsilon < \delta$ corresponds to a scenario where the learner wants to estimate the true value with high accuracy while the adversary is content with a coarse estimate. In contrast, the regime where $\delta < \epsilon$ is arguably much less interesting: it is not natural for the adversary, who is not engaged in the querying process, to aim at higher accuracy than the learner.

2. The requirement that $\delta \le 1/L$ stems from the following argument: if $\delta \ge 1/(L-1)$, then the entire interval $[0,1)$ is trivially $(\delta, L-1)$-coverable, and $C_\delta(\mathcal{I}(\overline{q})) \le C_\delta([0,1)) \le L-1 < L$. Thus, the privacy constraint is automatically violated, and no private learner strategy exists. To obtain a nontrivial problem, we therefore only need to consider the case where $\delta < 1/(L-1)$, which is only slightly broader than the regime $\delta \le 1/L$ that we consider.

The following theorem is the main result of this paper.

Theorem 1 (Query Complexity of Private Sequential Learning). *Fix $\epsilon > 0$, $\delta > 0$, and a positive integer $L \ge 2$ such that $2\epsilon < \delta \le 1/L$. Then*

$$\max\left\{\log\frac{1}{\epsilon}, \log\frac{\delta}{\epsilon} + 2L - 4\right\} \le N^*(\epsilon, \delta, L) \le \log\frac{1}{L\epsilon} + 2L. \tag{7}$$

The proof of the upper bound in Theorem 1 is constructive, providing a specific learner strategy that satisfies the bound. If we set $\delta = 1/L$, where the adversary's accuracy requirement is essentially as loose as possible, and thus corresponds to a worst case for the learner, then Theorem 1 leads to the following corollary, which yields upper and lower bounds that are tight up to an additive constant of four. In other words, the private learner strategy that we construct achieves essentially the optimal query complexity in this scenario.

Corollary 1. *Fix $\epsilon > 0$ and a positive integer $L \ge 2$ such that $2\epsilon < 1/L$. The following holds:*

a. *If $L = 2$, then*

$$\log\frac{1}{\epsilon} \le N^*\left(\epsilon, \frac{1}{L}, L\right) \le \log\frac{1}{\epsilon} + 4. \tag{8}$$

b. *If $L \ge 3$, then*

$$\log\frac{1}{L\epsilon} + 2L - 4 \le N^*\left(\epsilon, \frac{1}{L}, L\right) \le \log\frac{1}{L\epsilon} + 2L. \tag{9}$$

A key message from these results is about the price of privacy: it is not difficult to see that in the absence of a privacy constraint, the most efficient strategy, using a bisection search, can locate the true value with $\log(1/\epsilon)$ queries. Our results thus demonstrate that the price of privacy is at most an *additive* factor of $2L$.

We close by noting the following two important aspects of our upper bounds:

1. *Randomization:* The proof of our upper bounds involves a strategy that relies strongly on the availability of the randomization variable Y. It is not known whether a deterministic private learner strategy with comparable query complexity is possible.

2. *Practical relevance and limitation:* Although mathematically valid, the particular strategy we developed heavily exploits the underlying structure of the information set. When the learner's desired accuracy ϵ is very small, the resulting information set will contain separate, tiny guesses (intervals). From an applied perspective, the chance that these tiny guesses contain v^* is negligible, and a practical adversary might benefit by essentially ignoring them instead of having to cover them under the current worst-case formulation. We postpone a detailed discussion together with two potential remedies until Section 8, after developing a necessary understanding of the proof in the next sections.

5. Examples of Learner Strategies

Before delving into the proofs of our main result, we first provide some intuition and motivation by examining three representative learner strategies situated at different locations along the complexity–privacy tradeoff curve.

Strategy 1: Bisection. A most natural candidate is the classical bisection strategy, which is known to achieve the optimal query complexity in the absence of privacy constraints. Under this strategy, the learner first submits a query at the midpoint of $[0,1)$, that is, $q_1 = 0.5$. Then, based on the response, the learner identifies the half interval that contains the true value and subsequently submits its midpoint as the next

Figure 1. (Color online) Example of the Bisection Strategy

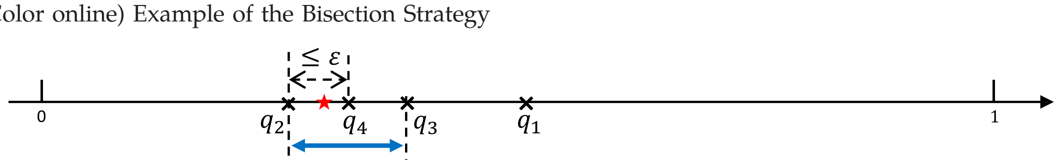

Notes. The star represents the true value v^*. The dashed line with arrows represents the learner's error tolerance. The solid line with arrows represents the information set of the adversary $\mathcal{I}(\overline{q})$.

query q_2. The process continues recursively until the learner finds an interval of length at most ϵ that contains the true value v^*. Figure 1 provides an illustration of this strategy.

Under the bisection strategy, the learner knows that the interval containing the true value is halved with each successive query. It follows that the number of queries needed under the bisection strategy is $N = \log(1/\epsilon)$. Unfortunately, the favorable query complexity afforded by the bisection strategy comes at the cost of the learner's privacy. In particular, at the end of the process, the adversary knows that the true value must be close (within ϵ) to the last query the learner submitted. Moreover, the information set is an interval of length at most 2ϵ. Hence, under our assumption that $\delta > 2\epsilon$, its δ-cover number is one, and the strategy is not private for any $L \geq 2$. The bisection strategy lies at one extreme end of the complexity–privacy tradeoff, with a minimal query complexity but no privacy.

Strategy 2: ϵ-Dense. At the opposite end of the spectrum is the ϵ-dense strategy, where the learner submits a predetermined sequence of $N = 1/\epsilon - 1$ queries, with $q_1 = \epsilon$, $q_2 = 2\epsilon, \ldots, q_N = N\epsilon$ (Figure 2). The strategy is accurate because the distance between any two adjacent queries is equal to the error tolerance ϵ. Moreover, because the sequence of queries is predetermined and does not depend on the location of the true value, the adversary obtains no information from the learner's query pattern, and the information set remains the interval $[0, 1)$ throughout. Thus, as long as $\delta \leq 1/L$, the strategy is (ϵ, δ, L)-private. Compared with the bisection strategy, the perfect privacy of the ϵ-dense strategy is achieved at the expense of an *exponential* increase in query complexity, from $\log(1/\epsilon)$ to $1/\epsilon$. The ϵ-dense strategy is therefore overly conservative and, as our proposed strategy will demonstrate, leads to unnecessarily high query complexity for moderate values of L.

Strategy 3: Replicated Bisection. The contrast between Strategies 1 and 2 highlights the tension between the learner's conflicting objectives: on the one hand, to maximally exploit the information learned from earlier queries and shorten the search and, on the other hand, to reduce adaptivity so that the queries are not too revealing. An efficient private learner strategy should therefore strike a balance between these two objectives. To start, it is natural to consider a learner strategy that combines Strategies 1 and 2 in an appropriate manner, which leads us to the replicated bisection strategy, which we describe next and which consists of two phases:

Phase 1: Deterministic Queries. The learner submits $L - 1$ queries, chosen deterministically:

$$q_1 = \frac{1}{L},\ q_2 = \frac{2}{L}, \ldots, q_{L-1} = \frac{L-1}{L}. \tag{10}$$

These queries partition the unit interval into L disjoint subintervals of length $1/L$ each, namely $[0, 1/L), [1/L, 2/L), \ldots, [1 - 1/L, 1)$. At this point, the learner can determine which one of the L subintervals contains the true value, whereas the adversary has gained no additional information about the true value. We refer to the subinterval that contains the true value as the *true subinterval* and all other subintervals as *false subintervals*. This phase uses $L - 1$ queries.

Phase 2: Replicated Bisection. In the second phase, the learner conducts a bisection strategy within the true subinterval until the true value has been located, while in the meantime submitting translated replicas of these queries in each false subinterval in parallel. The exact order in which these queries are submitted can be arranged in such a manner as to be independent from the identity of the true subinterval. This phase uses $L\log(1/L\epsilon)$ queries, where $\log(1/L\epsilon)$ is the number of queries needed to conduct a bisection strategy in a subinterval.

An example of this strategy is provided in Figure 3. When the process is completed, the learner will have identified the true value via the bisection strategy within the true subinterval, whereas the adversary will have seen L identical copies of the same bisection strategy, leading to an information set that consists of L disjoint length-2ϵ intervals separated from each other by a distance of $1/L - 2\epsilon$. It is not difficult to show that the replicated bisection strategy is (ϵ, δ, L)-private, with $L\log(1/L\epsilon) + L - 1$ queries. In particular, the replicated bisection strategy achieves privacy at the cost of an increase in query complexity that is a *multiplicative* factor of L compared with that of the bisection strategy ($N = \log(1/\epsilon)$).

The replicated bisection strategy thus appears to be a natural and successful combination of the bisection and ϵ-dense strategies: it ensures privacy while requiring substantially fewer than the $N = 1/\epsilon$

Figure 2. (Color online) Example of the ϵ-Dense Strategy

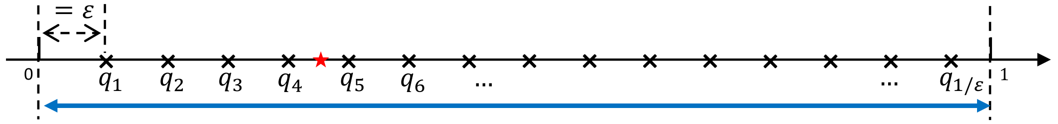

Notes. The dashed line with arrows represents the learner's error tolerance. The solid line with arrows represents the information set of the adversary $\mathcal{I}(\bar{q})$.

Figure 3. (Color online) Example of the Replicated Bisection Strategy with $L = 2$

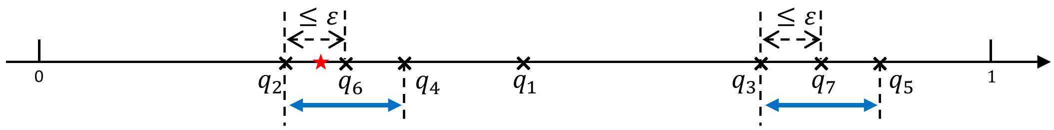

Notes. The dashed lines with arrows represent the learner's error tolerance. The solid lines with arrows represent the information set of the adversary $\mathcal{I}(\overline{q})$.

queries of the ϵ-dense strategy. Nevertheless, the upper bound in Theorem 1 indicates that the query complexity can be significantly improved, with an additive—rather than multiplicative—dependence on L.

6. Proof of the Upper Bound: Opportunistic Bisection Strategy

We prove in this section the upper bound on the query complexity in Theorem 1. This is achieved by constructing a specific learner strategy, which we will refer to as the *opportunistic bisection* (OB) strategy. We start with some terminology to facilitate the exposition.

Definition 4. Fix $M \in \mathbb{N}$ and an interval $\mathcal{J} \subset [0,1)$. Let $Z = (Z_1, Z_2, \ldots)$ be an infinite sequence of independent and identically distributed (i.i.d.) Bernoulli random variables with $\mathbb{P}(Z_i = 0) = 1/2$. Let $(q_1, q_2, \ldots, q_M)$ be a sequence of M queries, where q_1 is equal to the midpoint of $\mathcal{J}$, and let $(r_1, r_2, \ldots, r_M)$ be their corresponding responses.

a. We say that $(q_1, q_2, \ldots, q_M)$ is a *truthful bisection search* of $\mathcal{J}$ if it satisfies the following criteria, defined inductively. Let $\mathcal{J}_1 = \mathcal{J}$. For $i = 1, 2, \ldots, M$,

(a) The query q_i is set to

$$q_i = \text{midpoint of interval } \mathcal{J}_i. \tag{11}$$

(b) The interval $\mathcal{J}_{i+1}$ is set to

$$\mathcal{J}_{i+1} = \{[\inf \mathcal{J}_i, q_i), \quad \text{if } r_i = 0, [q_i, \sup \mathcal{J}_i), \quad \text{if } r_i = 1. \tag{12}$$

b. We say that $(q_1, q_2, \ldots, q_M)$ is a *fictitious bisection search* of $\mathcal{J}$ if it satisfies the following criteria, defined inductively. Let $\mathcal{J}_1 = \mathcal{J}$. For $i = 1, 2, \ldots, M$,

(a) The query q_i is set to

$$q_i = \text{midpoint of interval } \mathcal{J}_i. \tag{13}$$

(b) The interval $\mathcal{J}_{i+1}$ is set to

$$\mathcal{J}_{i+1} = \{[\inf \mathcal{J}_i, q_i), \quad \text{if } Z_i = 0, [q_i, \sup \mathcal{J}_i), \quad \text{if } Z_i = 1. \tag{14}$$

In words, whether a bisection search is truthful or fictitious depends on how the interval $\mathcal{J}_i$ is updated. In a truthful search, $\mathcal{J}_{i+1}$ is set to the half-interval within $\mathcal{J}_i$ that, according to the response r_i, contains the true value. In a fictitious search, this choice is made uniformly at random, according to Z.

We are now ready to define the OB strategy, which consists of two phases.

Phase 1: Opportunistic Guesses. The first $2L$ queries submitted by the strategy are deterministic and do not depend on responses from earlier queries, with

$$q_i = (i-1)\frac{1}{L}, \quad i = 1, \ldots, L, \tag{15}$$

and

$$(q_{L+1}, q_{L+2}, \ldots, q_{2L}) = (q_1 + \epsilon, q_2 + \epsilon, \ldots, q_L + \epsilon). \tag{16}$$

The two queries q_i and q_{i+L} determine an interval $[q_i, q_{i+L})$ of length ϵ. At the end of this phase, there will be L such intervals, evenly spaced across the unit interval. Each such interval $[q_i, q_{i+L})$ thus represents a guess on the true value v^*; if v^* lies in $[q_i, q_{i+L})$ for some $i \in \{1, \ldots, L\}$, then the learner learns the location of v^* within the desired level of accuracy. We refer to the interval $[q_i, q_{i+L})$ as the *ith guess.*

Phase 2: Local Bisection Search. The guesses submitted in Phase 1 are few and spaced apart, and it is possible that none of the L guesses contains v^*. The goal of Phase 2 is to ensure that the learner identifies v^* at the end, but the queries are to be executed in a fashion that conceals from the adversary whether v^* was identified during Phase 1 or Phase 2.

Define $\mathcal{J}^{(i)}$ as the interval between the ith and $(i+1)$th guesses:

$$\mathcal{J}^{(i)} = [q_{L+i}, q_{i+1}) = \left[(i-1)\frac{1}{L} + \epsilon, \frac{i}{L}\right), \quad i = 1, 2, \ldots, L. \tag{17}$$

We refer to $\mathcal{J}^{(i)}$ as the *ith subinterval*. Importantly, by the end of Phase 1, if none of the guesses contains the true value, then the learner knows which subinterval contains the true value, which we denote by $\mathcal{J}^*$. The queries in Phase 2 will be chosen according to the following rule:

1. If none of the guesses in Phase 1 contains v^*, then let $(q_{2L+1}, q_{2L+2}, \ldots, q_{2L+M})$ be a truthful bisection search of $\mathcal{J}^*$ with $M = \log(1/\epsilon L)$.

2. If one of the guesses in Phase 1 contains v^*, then let $\tilde{\mathcal{J}}$ be a subinterval chosen uniformly at random among all L subintervals, and let $(q_{2L+1}, q_{2L+2}, \ldots, q_{2L+M})$ be a fictitious bisection search of $\tilde{\mathcal{J}}$ with

Figure 4. (Color online) Example of the OB Strategy with $L = 3$

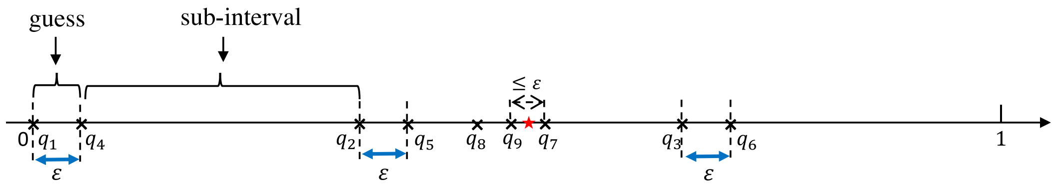

$M = \log(1/\epsilon L)$, using the randomization provided by Y (i.e., using Y to generate the sequence of i.i.d. Bernoulli random variables Z_i).

An example of this strategy is provided in Figure 4.

Remark. It is interesting to contrast OB with the replicated bisection strategy in Section 5. Both strategies use deterministic queries in the first phase, but instead of submitting L queries, the OB strategy incurs a slight overhead and submits L *guesses* ($2L$ queries). Crucially, the guesses make it possible to immediately discover the location of the true value in the first phase, although such discoveries might be unlikely. In the second stage, although the replicated bisection strategy conducts a bisection search in *each* of the L subintervals, the OB strategy does so in only *one* of the subintervals, hence drastically reducing the number of queries.

It follows directly from the definition that the number of queries submitted under the OB strategy is

$$N = 2L + \log\left(\frac{1}{\epsilon L}\right). \tag{18}$$

To complete the proof of the upper bound in Theorem 1, it thus suffices to show that the OB strategy satisfies both the accuracy and privacy constraints. This is accomplished in the following proposition, which is the main result of this section.

Proposition 1. *Fix $\epsilon > 0$, $\delta > 0$, and a positive integer $L \geq 2$ such that $2\epsilon < \delta \leq 1/L$. Then the OB strategy is (ϵ, δ, L)-private.*

Proof. We first show that the OB strategy is accurate and, specifically, that it allows the learner to produce an estimate of v^* with an absolute error of at most $\epsilon/2$. To this end, we consider two possible scenarios:

Case 1. Suppose that some guess in Phase 1, namely the interval $[q_{i'}, q_{i'+L})$, contains the true value, v^*. In this case, the learner can set $\hat{x}$ to be the midpoint of the guess, that is, $\hat{x} = (q_{i'} + q_{i'+L})/2$. Because the length of each guess is exactly ϵ, we have $|\hat{x} - v^*| \leq \epsilon/2$.

Case 2. Suppose that none of the guesses in Phase 1 contains v^*. This means that a *truthful* bisection search will be conducted in Phase 2 in the subinterval that contains v^*. Because the search is truthful, we know that one of the two intervals adjacent to q_N must contain v^*. Let this interval be denoted by H^*. Furthermore, because the length of each subinterval is less than $1/L$ and there are $\log(1/\epsilon L)$ steps in the bisection search, we know that the length of H^* is at most ϵ. Therefore, the learner can generate an accurate estimate by setting $\hat{x}$ to be the midpoint of H^*. Together with Case 1, this shows that the OB strategy leads to an accurate estimate of v^*.

We now show that the OB strategy is private and, in particular, that the δ-cover number of the information set of the adversary $C_\delta(\mathcal{I}(\overline{q}))$ is at least L.

Denote by $\mathcal{G}$ the union of the guesses, that is,

$$\mathcal{G} = \bigcup_{i=1}^{L} [q_i, q_{i+L}). \tag{19}$$

It is elementary to show that for two sets U and V, with $U \subset V$, if $C_\delta(U)$ is at least L, then so is $C_\delta(V)$. Therefore, it suffices to prove the following two claims:

Claim 1. The δ-cover number of $\mathcal{G}$, $C_\delta(\mathcal{G})$, is at least L.

Claim 2. The information set $\mathcal{I}(\overline{q})$ contains $\mathcal{G}$.

We first show Claim 1. Consider any interval $J \subset [0, 1)$ with length at most δ used in a cover for $\mathcal{G}$. By construction, each guess has length ϵ, and two adjacent guesses are separated by a distance of $1/L - \epsilon$. Because $\delta \leq 1/L$, this implies that the Lebesgue measure of $J \cap \mathcal{G}$ is at most ϵ. Because the Lebesgue measure of $\mathcal{G}$ is ϵL, we conclude that it will take at least L such intervals J (i.e., of size at most δ) to cover $\mathcal{G}$. Therefore, $C_\delta(\mathcal{G}) \geq L$. This proves Claim 1.

We next show Claim 2. Any particular query sequence $\overline{q}$ can arise in two different ways: (1) it may be that v^* is an arbitrary element of one of the guesses (i.e., $v^* \in \mathcal{G}$), and $\overline{q}$ is the result of a fictitious bisection search, or (2) it may be that v^* lies outside the guesses, and $\overline{q}$ is the result of a truthful bisection search. The adversary has no way of distinguishing between these two possibilities. Furthermore, there is no information available to the adversary that could distinguish between different elements of $\mathcal{G}$. As a consequence, all elements of $\mathcal{G}$ are included within the information set (i.e., $\mathcal{G} \subset \mathcal{I}(\overline{q})$), which proves Claim 2. $\square$

7. Proof of the Lower Bound

We now derive the two lower bounds in Theorem 1 on the query complexity. The query complexity of the OB strategy carries a $2L$ overhead compared with the (nonprivate) bisection strategy. The value $2L$ admits an intuitive justification: a private learner strategy must create L plausible locations of the true values, and each such location is associated with at least two queries. One may question, however, whether the $2L$ queries need to be *distinct* from the $\log(1/\epsilon)$ queries already used by the bisection search or whether the query complexity could be further reduced by *blending* the queries for obfuscation with those for identifying the true value in a more effective manner. The key to the proof of the lower bound in this section is to show that such blending is not possible: in order to successfully obfuscate the true value, one needs $2L$ queries that are *distinct* from those that participate in the bisection algorithm.

7.1. Information Sets

We first introduce some notation to facilitate our discussion. Recall that $(q_1, q_2, \ldots, q_N)$ is the sequence of learner queries. For the remainder of this section, we augment this sequence with two more queries, $q_0 \triangleq 0$ and $q_{N+1} \triangleq 1$, so that $\overline{q} = (0, q_1, q_2, \ldots, q_N, 1)$. This is inconsequential because $0 \le v^* < 1$, and hence adding q_0 and q_{N+1} does not provide additional information to either the learner or the adversary.

We start by examining the information provided to the learner through the queries and the responses. Let us fix an arbitrary $y \in \{1, \ldots, \mathcal{Y}\}$ that has positive probability and some $v^* \in [0, 1)$. Consider the resulting sequence of queries $\overline{q} = \overline{q}(v^*, y)$, and then let $\overline{q}^S = (q^0, q^1, \ldots, q^N, q^{N+1})$ be the sequence of queries in $\overline{q}$ arranged in increasing order (in particular, $q^0 = 0$ and $q^{N+1} = 1$). For each query q^i, the learner knows (through the response to the corresponding query) whether $v^* < q^i$ or $v^* \ge q^i$. In particular, at the end of the learning process, the learner has access to an interval of the form $H = [q^i, q^{i+1})$, for some $i \in \{0, 1, \ldots, N\}$, such that v^* is certain to belong to that interval. Furthermore, from the definition of learner strategies, all elements of that interval would have produced identical responses to the queries, and the learner has no information that distinguishes between such elements.

It is not hard to see that if $q^{i+1} - q^i > \epsilon$, then the learner has no way of producing an $(\epsilon/2)$-accurate estimate of v^*.[3] Because we are interested in learner strategies that satisfy the accuracy constraint in Definition 3, we conclude that the length of H is at most ϵ.

Let us now consider the situation from the point of view of the adversary. The adversary can look at the query sequence $\overline{q}$, form the intervals of the form $[q^i, q^{i+1})$, and select those intervals whose length is at most ϵ; we refer to these as *special intervals*. We have already argued that v^* must lie inside a special interval. Therefore, the adversary has enough information to conclude that v^* lies in the union of the special intervals. We denote that union by $\overline{\mathcal{I}}(\overline{q})$, and we have

$$\mathcal{I}(\overline{q}) \subset \overline{\mathcal{I}}(\overline{q}). \tag{20}$$

7.2. Completing the Proof of the Lower Bound

We are now ready to prove the lower bound. We begin with a lemma.

Lemma 1. *Fix a learner strategy ϕ that satisfies the accuracy constraint. For every $y \in \{1, 2, \ldots, \mathcal{Y}\}$, there exists $x \in (0, \delta)$ such that there are at least $\log(\delta/\epsilon)$ of the queries in $\overline{q}(x, y)$ that belong to $(0, \delta)$.*

Lemma 1 is essentially the classical result that $\log(1/\epsilon)$ query complexity of the bisection strategy is optimal for the unit interval (Waeber et al. 2013), which proves the first term of the lower bound in Theorem 1. We omit the proof of Lemma 1, which is fairly standard, but provide an intuitive argument. Fix $y \in \{1, 2, \ldots, \mathcal{Y}\}$. The interval $(0, \delta)$ consists of δ/ϵ disjoint subintervals of length ϵ each. An accurate learner strategy therefore must be able to distinguish in which one of these subintervals the true value resides. Distinguishing among δ/ϵ possibilities using binary feedback therefore implies that there will be some $v^* \in (0, \delta)$ whose accurate identification requires $\log(\delta/\epsilon)$ queries in $(0, \delta)$.

For the rest of this proof, fix $v^* = x_0$ and $Y = y_0$ for some x_0 and y_0 satisfying Lemma 1, and use $\overline{q}$ to denote $\overline{q}(x_0, y_0)$. We now consider the queries in the interval $[\delta, 1]$. Among them, we restrict attention to queries that are endpoints of special intervals. We call these *special* queries and let K be their number. We sort the special queries in ascending order and denote them by $s_1, s_2, \ldots, s_K$, where K is their number.

The outline of the rest of the argument is as follows. For a private learner strategy, the δ-cover number of $\mathcal{I}(\overline{q})$ is at least L. From Equation (20), it follows that the δ-cover number of $\overline{\mathcal{I}}(\overline{q})$ is also at least L. Because endpoints of special intervals are within ϵ of each other, every s_i must be within ϵ of a *neighboring* query (namely s_{i-1} or s_{i+2}), with the possible exception of s_1, which could be the right endpoint of a special interval whose left endpoint, denoted s_0, is in $[0, \delta)$. Using the assumption that $\delta > 2\epsilon$, each interval used in the cover can include two, and often three, queries s_i. In what follows, we will make this argument precise and show that the number K of special queries is at least $2L - 3$.

Let us consider first the case where the previously mentioned exception does not arise; that is, we assume that s_1 is *not* the right endpoint of a special interval. We decompose the set $\overline{\mathcal{I}}(\overline{q}) \cap [\delta, 1]$ as the

Figure 5. Illustration of the Proof

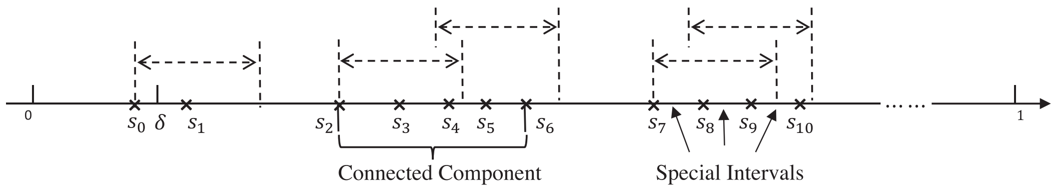

Notes. Here $s_1 - s_0 \leq \epsilon$ and hence $[s_0, s_1]$ is a special interval. There are three connected components: $[s_0, s_1]$, $[s_2, s_6]$, and $[s_7, s_{10}]$. The dashed lines with arrows represent a cover.

union of its (finitely many) connected components, which we call just *components* for short (see Figure 5 for an illustration). Each connected component is an interval whose endpoints are special queries (s_i for some i). Furthermore, within each such interval, special queries are separated by at most ϵ. Suppose that we have m components. For the jth component, let k_j be the number of special queries it contains, and let d_j be its length. We have $d_j \leq (k_j - 1)\epsilon$. In particular, the jth interval can be covered by at most

$$\left\lceil \frac{d_j}{\delta} \right\rceil \leq \left\lceil (k_j - 1)\frac{\epsilon}{\delta} \right\rceil \leq \left\lceil \frac{k_j - 1}{2} \right\rceil \leq \frac{k_j}{2}$$

intervals of length δ. (We have used here our standing assumption that $\delta > 2\epsilon$.) Summing over the different components, we conclude that $\overline{\mathcal{I}}(\overline{q}) \cap [\delta, 1]$ can be covered by at most $\sum_j k_j/2 = K/2$ intervals of length δ.

For the exceptional case where s_1 is the right endpoint of a special interval $[s_0, s_1]$, we just apply the same argument, now on the set $\overline{\mathcal{I}}(\overline{q}) \cap [s_0, 1]$, and for an augmented collection of special queries $(s_0, s_1, \ldots, s_K)$. Having effectively increased the number of points of interest, from K to $K+1$, we obtain an upper bound of $(K+1)/2$ on the number of intervals of length δ that are needed to cover $\overline{\mathcal{I}}(\overline{q}) \cap [s_0, 1]$.

By combining the results of the two cases and using one more interval of length δ to cover the set $[0, \delta]$, we conclude that the δ-cover number of $\overline{\mathcal{I}}(\overline{q})$, and therefore of $\mathcal{I}(\overline{q})$ as well, is at most $(K+3)/2$. By contrast, as long as we are dealing with an (ϵ, δ, L)-private strategy, this δ-cover number is at least L. Thus, $(K + 3)/2 \geq L$ or $K \geq 2L - 3$.

Recall that the argument is being carried out for the case of particular x_0 and y_0 with the properties specified earlier. For such x_0 and y_0, we have at least $\log(\delta/\epsilon)$ queries in the set $[0, \delta)$ and at least $2L - 3$ queries in the set $[\delta, 1]$. Recall that we have introduced an additional artificial query at $x = 1$ (i.e., $q^{N+1} = 1$), which is above and beyond the N queries used by the strategy, and this query may be included in the K special queries. For this reason, the lower bound on N has to be decremented by one, leading to the lower bound in Theorem 1.

8. Conclusions and Future Work

This paper studies an intrinsic privacy–complexity tradeoff faced by a learner in a sequential learning problem who tries to conceal his or her findings from an observant adversary. We use the notion of information set, the set of possible true values, to capture the information available to the adversary through the learner's learning process and focus on the coverability of the information set as the main metric for measuring a learner strategy's level of privacy. Our main result shows that to ensure privacy, that is, so that the resulting information set requires at least L intervals of size δ to be fully covered, it is necessary for the learner to use at least $\log(\delta/\epsilon) + 2L - 4$ queries. We further provide a constructive learner strategy that achieves privacy with $\log(1/L\epsilon) + 2L$ queries. Together the upper and lower bounds on the query complexity demonstrate that increasing the level of privacy L leads to a *linear* additive increase in the learner's query complexity.

Although mathematically valid and worst-case optimal, the OB strategy could be potentially undesirable in applications where the learner aims for an extremely accurate estimate, that is, $\epsilon \to 0$. Indeed, when ϵ is small, the chance that the guesses contain the true value becomes negligible. In practice, as long as the adversary is willing to tolerate a small chance of making an error, predicting according to the subsequent bisection phase would be much more beneficial: as ϵ goes to zero, the probability of accurately estimating v^* approaches one for the adversary. As was mentioned in Section 3.3, this naturally motivates an average-case Bayesian variant that essentially focuses on covering most, rather than all, of the information set (see Appendix B). There the OB strategy performs arbitrarily badly when ϵ is small enough, and the formulation warrants a new private strategy, which turns out to be the replicated bisection in Section 5. We refer readers to Appendix B for a much more elaborated discussion of the Bayesian variant. Relatedly, as another direction of resolving the preceding practical issue, one could attempt to refine the information set to contain only those x' values whose probability of producing the observed sequence of queries is greater than some confidence level ξ. This is to be contrasted with only requiring a positive probability in the current formulation. More precisely, in Equation (1), one could let the requirement be $\mathbb{P}_{\phi}(Q_x = \bar{q}) > \xi$ for some confidence level $\xi > 0$. It would be interesting to see whether we can use the OB strategy as the base case to construct a class of strategies with respect to different confidence levels ξ.

In addition, there are several interesting extensions and variations of the model that were left unaddressed. One may consider the binary query model in higher dimensions, where $v^* \in [0,1)^d$. A query in this setting will be a hyperplane in $\mathbb{R}^d$, where the response indicates whether the true value is to the right or to the left of the queried hyperplane. We further investigate this extension in Appendix C. Yet another interesting extension is to consider alternative response models, such as noisy binary responses, or real-valued responses that can convey a richer set of information. More broadly, there could be other interesting problem formulations for understanding the privacy implications in a number of sequential decision problems in learning theory, optimization, and decision theory. It is not difficult to see that standard algorithms, originally designed to optimize runtime or query complexity, often provide little or no protection for the learner's privacy. Can we identify a universal procedure to design sample-efficient private decision strategies? Is there a more general tradeoff between privacy and complexity in sequential decision making? We are optimistic that there are many fruitful inquiries along these directions.

Appendix A. Coverability and the Adversary's Probability of Correct Estimation

In this appendix, we provide an alternative, probabilistic interpretation of the definition of privacy in Definition A.2. In particular, we show that $1/L$ can be interpreted as a worst-case guarantee on the adversary's probability of correct detection.

Recall the definition of $\mathcal{Q}(x)$ as the set of all possible query sequences when $v^* = x$ (cf. Equation (1)), and let $\mathcal{Q} = \cup_{x\in[0,1)}\mathcal{Q}(x)$. We next define adversary estimators $\hat{x}^a$ as random variables whose values are determined by the observed query sequence $\bar{q}$ together with an independent randomization seed.

Definition A.1. Fix $\delta > 0$, $L \geq 2$, a learner strategy $\phi \in \Phi_N$, and a sequence of queries $\bar{q} \in \mathcal{Q}$. We say that an adversary estimator $\hat{x}^a$ is (δ, L)-*correct* with respect to $\bar{q}$ if

$$\mathbb{P}\left(\left|\hat{x}^a(\bar{q}) - x\right| \leq \delta/2\right) > \frac{1}{L}, \quad \forall x \in \mathcal{I}(\bar{q}), \tag{A.1}$$

where the probability is taken with respect to any randomization in the adversary's estimator $\hat{x}^a$.

In words, an adversary estimator is (δ, L)-correct with respect to $\bar{q}$ if as soon as the learner deploys the queries $\bar{q}$, the adversary will *know* that the resulting estimate will incur an error of at most $\delta/2$ with probability larger than $1/L$. In a sense, this means that $\bar{q}$ conceals the true value poorly.

Based on this probabilistic definition of the correctness of adversary estimators, we can then define a new notion of privacy for learner strategies in a manner similar to Definition 3. To distinguish the two, we use the term *secure learner strategies* for the new notion.

Definition A.2 (Secure Learner Strategy). Fix $\epsilon > 0$, $\delta > 0$, and $L \geq 2$, with $L \in \mathbb{N}$. A learner strategy $\phi \in \Phi_N$ is (ϵ, δ, L)-*secure* if it satisfies the following:

a. *Accuracy constraint:* The learner estimate accurately recovers the true value, with probability one:

$$\mathbb{P}(\,|\hat{x}(x, Y) - x| \leq \epsilon/2) = 1, \quad \forall x \in [0,1),$$

where the probability is measured with respect to the randomness in Y.

b. *Privacy constraint:* For every possible sequence of queries $\bar{q} \in \mathcal{Q}$, there does not exist an adversary estimator $\hat{x}^a$ that is (δ, L)-correct with respect to $\bar{q}$. That is, for every $\bar{q} \in \mathcal{Q}$, there exists $x \in \mathcal{I}(\bar{q})$ such that

$$\mathbb{P}\left(\left|\hat{x}^a(\bar{q}) - x\right| \leq \delta/2\right) \leq \frac{1}{L}, \tag{A.2}$$

where the probability is taken with respect to any randomization in $\hat{x}^a$.

The following proposition establishes an equivalence between the private learner strategies, defined in terms of the coverability of information set, and secure learner strategies, defined in terms of adversary's probability of correct estimation.

Proposition A.1. *Fix $\delta > 0$, $\epsilon > 0$, and $L \geq 2$, with $L \in \mathbb{N}$. A learner strategy $\phi \in \Phi$ is (ϵ, δ, L)-secure if and only if it is (ϵ, δ, L)-private.*

Proof. We first establish the forward direction that if ϕ is (ϵ, δ, L)-secure, then ϕ is (ϵ, δ, L)-private. We will actually establish the equivalent statement that if ϕ is not (ϵ, δ, L)-private, then it is not (ϵ, δ, L)-secure. For this, it suffices to show the claim that for a sequence of queries $\bar{q} \in \mathcal{Q}$, if the δ-cover number of the information set $\mathcal{I}(\bar{q})$, $C_\delta(\mathcal{I}(\bar{q}))$, is at most $L-1$ (thus violating (ϵ, δ, L)-privacy), then there exists an adversary estimator that is (δ, L)-correct with respect to $\bar{q}$. To show the claim, fix $\bar{q} \in \mathcal{Q}$ such that $\mathcal{I}(\bar{q})$ is $(\delta, L-1)$-coverable. Then there exist $L-1$ intervals $[a_1, b_1]$, $[a_2, b_2]$, …, $[a_{L-1}, b_{L-1}]$ each of length δ that cover $\mathcal{I}(\bar{q})$. Consider a randomized adversary estimator $\hat{x}^a$ that is distributed uniformly at random among the $L-1$ midpoints of the intervals. Then, with probability $1/(L-1)$, the resulting estimator $\hat{x}^a$ will lie in the same interval as the true value; when this event occurs, and because the length of each one of the intervals is at most δ, the estimate will be at a distance of at most $\delta/2$ from the true value, implying that

$$\mathbb{P}\left(\left|\hat{x}^a(\bar{q}) - x\right| \leq \delta/2\right) = \frac{1}{L-1} > \frac{1}{L}, \qquad \forall x \in \mathcal{I}(\bar{q}).$$

This shows that $\hat{x}^a$ is (δ, L)-correct given $\bar{q}$, which proves the claim.

Conversely, we now prove that if ϕ is (ϵ, δ, L)-private, then it is (ϵ, δ, L)-secure. It suffices to show the following claim: for any sequence of queries $\overline{q} \in \mathcal{Q}$, if $C_\delta(\mathcal{I}(\overline{q}))$ is at least L, then there does not exist an adversary estimator that is (δ, L)-correct with respect to $\overline{q}$. We make use of the following lemma.

Lemma A.1. *Fix $\delta \in (0,1)$ and $L \geq 2$. Let J be a subset of $[0,1)$ such that the δ-cover number of J, $C_\delta(J)$, is at least L. Then there exist points $\{x_1, x_2, \ldots, x_L\}$ in the closure of J such that*

$$|x_i - x_j| > \delta, \quad \forall i \neq j. \tag{A.3}$$

Proof. We prove the lemma by constructing the set of x_j values explicitly, with the aid of a helper sequence $\{z_j\}$. In particular, we construct $\{z_j\}$ so that the sequence of intervals $\{[z_j, z_j + \delta]\}$ forms a cover. The sequence $\{x_j\}$ will then be derived from a perturbed version of $\{z_j\}$, so the constraint in Equation (A.3) can be satisfied.

Let $\overline{J}$ be the closure of J. Consider the following procedure:

1. Let $z_1 = x_1 = \min \overline{J}$.
2. For $i = 2, 3, \ldots$, consider the following recursive process of constructing z_i and x_i: let

$$y_i \triangleq \min\{x \in \overline{J} : x \geq z_{i-1} + \delta\}.$$

Now consider two scenarios:

2.1. If $y_i > z_{i-1} + \delta$, then let $z_i = x_i = y_i$.

2.2. If $y_i = z_{i-1} + \delta$, then check whether y_i is a right endpoint of some interval in $\overline{J}$.

a. If there exists $\lambda_i > 0$ small enough such that $[y_i, y_i + \lambda_i] \subset \overline{J}$, then let $z_i = y_i$ and $x_i = y_i + \lambda_i'$, where $0 < \lambda_i' < \lambda_i$ and λ_i' is sufficiently small.

b. Otherwise, if such a λ_i does not exist, let

$$z_i = x_i = \min\{x \in \overline{J} : x > z_{i-1} + \delta\}.$$

The procedure terminates at some step T when $(z_T + \delta, 1) \cap \overline{J} = \emptyset$. By construction, all z_i and x_i values belong to the closure of J. Furthermore, the intervals

$$W_i := [z_i, z_i + \delta], \quad i = 1, 2, \ldots, T,$$

form a cover of $\overline{J}$. Because $C_\delta(J) \geq L$ by assumption, it follows that we must have $T \geq L$. Finally, it is easy to verify that the points $\{x_1, x_2, \ldots, x_L\}$ satisfy the conditions outlined in the lemma (Equation (A.3)). In particular, this is guaranteed by choosing λ_i' to be sufficiently small and $\lambda_i' < \lambda_j'$ for $i < j$, when Case 2(i) occurs. This completes the proof. □

Fix an adversary estimator $\hat{x}^a$ and some $\overline{q} \in \mathcal{Q}$ such that $C_\delta(\mathcal{I}(\overline{q})) \geq L$. Apply Lemma A.1 with $J = \mathcal{I}(\overline{q})$, and let $\{x_1, x_2, \ldots, x_L\}$ be as defined in the lemma. Because the x_i values belong to the closure of $\mathcal{I}(\overline{q})$ and $|x_i - x_j| > \delta$ for any $i \neq j$, by slightly perturbing them, we can obtain a set of points $\{\tilde{x}_1, \tilde{x}_2, \ldots, \tilde{x}_T\} \subset \mathcal{I}(\overline{q})$ such that

$$|\tilde{x}_i - \tilde{x}_j| > \delta, \quad \forall i \neq j,$$

still holds. Define intervals

$$U_i := [\tilde{x}_i - \delta/2, \tilde{x}_i + \delta/2], \quad i = 1, 2, \ldots, L.$$

Because the distance between any two distinct $\tilde{x}_i$ values is greater than δ, we know that the intervals U_i are disjoint, which implies that at least one of these L intervals will have probability of containing $\hat{x}^a(\overline{q})$ less than or equal to $1/L$. In particular, there exists $i^* \in \{1, 2, \ldots, L\}$ such that

$$\mathbb{P}\big(|\hat{x}^a(\overline{q}) - \tilde{x}_{i^*}| \leq \delta/2\big) = \mathbb{P}\big(\hat{x}^a(\overline{q}) \in U_{i^*}\big) \leq 1/L.$$

Because $\tilde{x}_{i^*} \in \mathcal{I}(\overline{q})$ by construction, we conclude that the adversary estimate $\hat{x}^a$ is not (δ, L)-correct with respect to $\overline{q}$. This completes the proof of the claim and hence the converse direction of the proposition. □

Appendix B. Bayesian Private Learning Model

The private sequential learning model we studied in this paper assumes that neither the learner nor the adversary has any prior information on the true value v^* and that they can obtain information only gradually, through queries. In this section, we discuss a Bayesian variant of the model where the true value v^* is generated according to a prior distribution, known to both parties. In particular, we assume that v^* is distributed according to a distribution P_{v^*}, where, for simplicity, we assume that the support of P_{v^*} is equal to $[0,1)$.

Naturally, we allow the learner strategy, defined in Section 2.1, to depend on P_{v^*}. In addition, because the adversary aims to produce an estimate $\hat{x}^a$ that is close to the true value, we define an adversary strategy to be a function ψ that maps the adversary's available information (i.e., the prior distribution P_{v^*}, the learner strategy ϕ, and the sequence of observed queries $\overline{q}$) to a probability distribution over $[0,1)$ or, equivalently, to a random variable $\hat{x}^a$ that takes values in $[0,1)$. Denote by Ψ the set of all such functions, that is, the set of all adversary strategies.

Because the true value in the Bayesian private learning model admits a prior distribution, instead of using the information set, it is sufficient for the adversary to keep track of the posterior distribution of v^* given the learner's queries. The Bayesian formulation also allows us to measure the probability that the adversary is able to provide an estimate of the true value that is within a given error tolerance. This leads to the following definition of Bayesian private learner strategies.

Definition B.1. Fix $\epsilon > 0, \delta > 0$, an integer $L \geq 2$, and a prior distribution P_{v^*}. A learner strategy $\phi \in \Phi_N$ is (ϵ, δ, L)-B-private if it satisfies the following:

a. *Accuracy constraint:* The strategy accurately recovers v^* with probability one:

$$\mathbb{P}\left(\left|\hat{x}(v^*, Y) - v^*\right| \leq \frac{\epsilon}{2}\right) = 1,$$

where the probability is taken with respect to the randomness in v^* and Y.

b. *Privacy constraint:* Under this learner strategy, and for every adversary strategy $\psi \in \Psi$, we have

$$\mathbb{P}\big(|\hat{x}^a - v^*| \leq \delta/2\big) \leq \frac{1}{L}, \tag{B.1}$$

where the probability is taken with respect to the randomness in v^*, Y, and $\hat{x}^a$.

Notice the resemblance of this definition with Definition 3. The parameters ϵ and δ have the same meaning as in the original model, and L mirrors the role of L in (δ, L)-coverability but now has a more concrete interpretation in terms of the adversary's probability of error.

Fix a prior distribution P_{v^*}. Denote by $N_B^*(\epsilon, \delta, L)$ the minimum number of queries needed for there to exist an (ϵ, δ, L)-B-private learner strategy:

$$N_B^*(\epsilon, \delta, L) = \min\{N \in \mathbb{N} : \exists \phi \in \Phi_N \text{ s.t. } \phi \text{ is } \times (\epsilon, \delta, L)\text{-B-private}\}.$$

Similar to the original model, we would like to obtain lower and upper bounds on $N_B^*(\epsilon, \delta, L)$.

B.1. Comparing the Two Formulations

In single-player optimization, Bayesian formulations are always less demanding than worst-case formulations. However, this is not the case in game-theoretic situations, and an automatic comparison between $N^*(\epsilon, \delta, L)$ and $N_B^*(\epsilon, \delta, L)$ is not available. This is because the Bayesian formulation can, in principle, work in the favor of either player, as discussed:

1. Under the worst-case formulation, once an information set is determined, and in order to break the learner's privacy, the adversary must be able to cover (with intervals of length δ) the *entire* information set. In contrast, in the Bayesian formulation, the adversary can break the learner's privacy by covering *most* of the information set. Because the Bayesian formulation seems to make the adversary's objective easier to achieve, one might expect that the situation has become more demanding for the learner, resulting in higher query complexity.
2. Under the worst-case formulation, part (b) of Definition 3 involves a privacy requirement that holds for every $x \in [0, 1)$. In a Bayesian formulation, this is essentially replaced by a requirement that holds, on average, over x. In principle, this is an easier requirement for the learner and might work in the direction of lower query complexity under a Bayesian formulation.

For a concrete illustration of the difference between the two formulations, consider the OB strategy in Section 6, which was private under the worst-case formulation. Under that strategy, privacy was ensured by having many small intervals in the information set (the guesses) that the adversary could not ignore. By contrast, under the Bayesian formulation, and as long as ϵ is small enough, the probability that these small intervals contain v^* is negligible, and privacy is completely lost in the Bayesian setting. Specifically, suppose that the prior distribution of v^* is uniform over $[0, 1)$, and consider an adversary strategy that always uses the last query of Phase 2 of the OB strategy as his or her estimator, that is, the last query of the local bisection search phase. Recall the regime of interest $2\epsilon < \delta \le 1/L$. With probability $1 - L\epsilon$, the L guesses will not contain v^*, and Phase 2 would be a truthful bisection search, under which the adversary strategy would succeed. That is, the probability that the adversary recovers v^* with error at most $\delta/2$ is at least $1 - L\epsilon$. These arguments demonstrate that when ϵ is too small, privacy cannot be guaranteed under the Bayesian formulation.

B.2. Results

For the case where P_{v^*} is a uniform distribution over $[0, 1)$, we can obtain the following result by adapting the proof for the original model.

Proposition B.1. *Fix $\epsilon > 0$, $\delta > 0$, and a positive integer $L \ge 2$ such that $2\epsilon < \delta \le 1/L$. Suppose that the prior distribution P_{v^*} is uniform over $[0, 1)$. Then*

$$\log \frac{1}{\epsilon} \le N_B^*(\epsilon, \delta, L) \le L \log \frac{1}{L\epsilon} + L - 1.$$

The lower bound $\log(1/\epsilon)$ follows directly from the accuracy constraint (see Section 7). The upper bound is achieved by the replicated bisection strategy in Section 5, with L replications. Because the replicated bisection strategy creates L identical copies of query patterns across L subintervals, the symmetry ensures that the posterior distribution of v^* will be evenly spread across all subintervals, forcing the adversary's probability of correct detection to be at most $1/L$.

In contrast to the bounds in Theorem 1, the leading terms in the upper and lower bounds in Proposition B.1 differ by a multiplicative factor of L, and closing this gap is nontrivial. In particular, the approach used in this paper to prove the lower bound is highly dependent on the structure of the information set and is unlikely to apply.[4]

Appendix C. High-Dimensional Private Sequential Learning

In this appendix, we provide a brief discussion of a high-dimensional variant of the private sequential learning model that allows us to apply the insights from the original model to a more general setting. Consider a learner who wants to identify a true value v^* situated in a d-dimensional unit cube $[0, 1)^d$. A query is now a hyperplane in $\mathbb{R}^d$, where the response indicates whether the true value is to the right or to the left of the queried hyperplane. Formally, a query is specified by some $q = (q^{(1)}, q^{(2)}, \ldots, q^{(d)}, c) \in \mathbb{R}^{d+1}$, and the corresponding response is $r = \mathbb{I}(\langle (q^{(1)}, \ldots, q^{(d)}), v^* \rangle \le c)$. The distance metric for measuring estimation errors is the l_∞ norm. In this setting, the adversary aims to cover the information set using less than L hypercubes of edge length δ (and hence volume δ^d).

Let us focus on the regime where $2\epsilon < \delta \le 1/L^{1/d}$ and assume, for the sake of this discussion, that $L^{1/d}$ is an integer. For an upper bound, it is not hard to see that we can extend the OB strategy to this case in a straightforward manner. In particular, we can use the following two-phase strategy:

Phase 1: Opportunistic Guesses. During the first phase, we first partition the space into L equal-sized large hypercubes, each with volume $1/L$, so that each side of a hypercube has length $L^{-1/d}$. This can be achieved by submitting $L^{1/d}$ queries for each dimension. For example, the queries for the first dimension are $\{(1, 0, \ldots, 0, (i-1)L^{-1/d})\}_{i=1}^{L^{1/d}}$ and the queries for the second dimension are $\{(0, 1, \ldots, 0, (i-1)L^{-1/d})\}_{i=1}^{L^{1/d}}$. To achieve successful obfuscation so that the adversary is not able to cover the information set with $L - 1$ hypercubes, the second step of this phase is to construct a smaller subcube with edge length ϵ inside each one of the L hypercubes to serve as the plausible guesses. Again, this is achieved by another $L^{1/d}$ query for each dimension. For instance, the queries for the first dimension are $\{(1,0,\ldots,0,(i-1)L^{-1/d} + \epsilon)\}_{i=1}^{L^{1/d}}$, and the queries for the second dimension are $\{(0, 1, \ldots, 0, (i-1)L^{-1/d} + \epsilon)\}_{i=1}^{L^{1/d}}$.

Phase 2: Local Bisection Search. After the guesses are created in Phase 1, the learner then performs a bisection search inside one of the L large hypercubes (i.e., d bisection searches along each dimension). Depending on whether the guesses contain v^* or not, this bisection search is either truthful or fictitious, similar to the original policy.

Using the same argument as in the analysis for the original OB strategy, it is not difficult to verify that the previous strategy is private and achieves a query complexity of $d\log(1/(L^{1/d}\epsilon)) + 2dL^{1/d}$. For the lower bound, it is evident that the agent needs at least $d\log(1/\epsilon)$ queries to satisfy the accuracy constraint, even in the absence of a privacy constraint. It appears challenging, however, to obtain a stronger lower bound that depends on the dimension in a meaningful way. The argument that we use in the proof of the current lower bound does not generalize easily to higher dimensions.

Endnotes

[1] All logarithms are taken with respect to base 2. To reduce clutter, noninteger numbers are to be understood as rounded upward. For example, the lower bound should be understood as $\lceil\log(1/L\epsilon)\rceil + 2L$, where $\lceil\cdot\rceil$ represents the ceiling function.

[2] We consider a half-open interval here, which allows for a cleaner presentation, but the essence is not changed if the interval is closed.

[3] For any choice of $\hat{x}$, there will always be some $x \in H$ such that $|\hat{x} - x| > \epsilon/2$. Furthermore, such an x is a possible value of v^* because it would have produced the exact same sequence of responses.

[4] In a recent paper that follows up on our work, Xu (2018) shows that the upper bound in Proposition B.1 is asymptotically tight up to the first order in the limit as $\epsilon \to 0$, whereas L and δ remain fixed. The proof in Xu (2018) uses an information-theoretic argument that is very different from the line of analysis in this paper.

References

Blum A, Ligett K, Roth A (2013) A learning theory approach to noninteractive database privacy. *J. ACM* 60(2):1–25.

Chaudhuri K, Hsu D (2011) Sample complexity bounds for differentially private learning. *Proc. 24th Annual Conf. Learn. Theory* (PMLR), 155–186.

Chor B, Kushilevitz E, Goldreich O, Sudan M (1998) Private information retrieval. *J. ACM* 45(6):965–981.

Cummings R, Ligett K, Nissim K, Roth A, Wu ZS (2016) Adaptive learning with robust generalization guarantees. *Conf. on Learning Theory*, 772–814.

Duchi J, Wainwright M, Jordan M (2018) Minimax optimal procedures for locally private estimation. *J. Amer. Statist. Assoc.* 113(521):182–201.

Dwork C, Roth A (2014) The algorithmic foundations of differential privacy. *Foundations Trends Theoretical Comput. Sci.* 9(3–4):211–407.

Dwork C, Naor M, Vadhan S (2012) The privacy of the analyst and the power of the state. *Proc. IEEE 53rd Annu. Sympos. on Foundations of Computer Science* (IEEE, New York), 400–409.

Dwork C, McSherry F, Nissim K, Smith A (2006) Calibrating noise to sensitivity in private data analysis. *Proc. Theory of Cryptography Conf.* (Springer, New York), 265–284.

Dwork C, Feldman V, Hardt M, Pitassi T, Reingold O, Roth A (2015a) Generalization in adaptive data analysis and holdout reuse. *Adv. Neural Inform. Processing Systems* 29:2350–2358.

Dwork C, Feldman V, Hardt M, Pitassi T, Reingold O, Roth A (2015b) The reusable holdout: Preserving validity in adaptive data analysis. *Science*. 349(6248):636–638.

Dwork C, Feldman V, Hardt M, Pitassi T, Reingold O, Roth AL (2015c) Preserving statistical validity in adaptive data analysis. *Proc. 47th Annual Sympos. Theory Comput.* (ACM, New York), 117–126.

Fanti G, Kairouz P, Oh S, Viswanath P (2015) Spy vs. spy: Rumor source obfuscation. *Performance Evaluation Rev.* 43:271–284.

Feldman V, Xiao D (2014) Sample complexity bounds on differentially private learning via communication complexity. *Conf. on Learning Theory*, 1000–1019.

Gasarch W (2004) A survey on private information retrieval. *Bull. EATCS* 82:72–107.

Horstein M (1963) Sequential transmission using noiseless feedback. *IEEE Trans. Inform. Theory* 9(3):136–143.

Jain P, Kothari P, Thakurta A (2012) Differentially private online learning. *Proc. 25th Annual Conf. Learn. Theory* (PMLR), 23: 24.1–24.34.

Kushilevitz E, Ostrovsky R (1997) Replication is not needed: Single database, computationally-private information retrieval. *Proc. 38th Annual Sympos. Foundations Comput. Sci.* (IEEE, Piscataway, NJ), 364–373.

Liao J, Sankar L, Tan VY, du Pin Calmon F (2018) Hypothesis testing under mutual information privacy constraints in the high privacy regime. *IEEE Trans. Inform. Forensics Security* 13(4): 1058–1071.

Raskhodnikova S, Smith A, Lee HK, Nissim K, Kasiviswanathan SP (2008) What can we learn privately. *Proc. 54th Annual Sympos. Foundations Comput. Sci.* (IEEE, Piscataway, NJ), 531–540.

Smith A (2011) Privacy-preserving statistical estimation with optimal convergence rates. *Proc. 43rd Annual Sympos. Theory Comput.* (ACM, New York), 813–822.

Tsitsiklis J, Xu K (2018) Delay-predictability tradeoffs in reaching a secret goal. *Oper. Res.* 66(2):587–596.

Waeber R, Frazier PI, Henderson SG (2013) Bisection search with noisy responses. *SIAM J. Control Optim.* 51(3):2261–2279.

Wasserman L, Zhou S (2010) A statistical framework for differential privacy. *J. Amer. Statist. Assoc.* 105(489):375–389.

Xu K (2018) Query complexity of bayesian private learning. *Adv. Neural Inform. Processing Systems* 31:2431–2440.

Xu Z (2017) Private sequential search and optimization. MS thesis, Massachusetts Institute of Technology, Cambridge.

John N. Tsitsiklis is a Clarence J. Lebel Professor with the Department of Electrical Engineering and Computer Science at MIT, where he serves as the director of the Laboratory for Information and Decision Systems. He is a member of the National Academy of Engineering and holds three honorary doctorates. Among other distinctions, he is a recipient of the ACM SIGMETRICS Achievement Award (2016), the IEEE Control Systems Award (2018), and the INFORMS von Neumann Theory Prize (2018).

Kuang Xu is an associate professor of operations, information and technology at Stanford Graduate School of Business. His research interests are in the analysis and design of large-scale stochastic systems. He received his PhD in electrical engineering and computer science from MIT. He has received several awards including the winner of INFORMS George E. Nicholson Paper Competition, and a Best Paper Award, Outstanding Student Paper Award, and Rising Star Research Award from ACM SIGMETRICS.

Zhi Xu is currently a PhD student in the Department of Electrical Engineering and Computer Science at Massachusetts Institute of Technology. He is affiliated with the Laboratory for Information and Decision Systems. His research interests include both theoretical and applied machine learning.