Private Sequential Learning

John N. Tsitsiklis JNT@MIT.EDU
LIDS, Massachusetts Institute of Technology, Cambridge, MA 02139

Kuang Xu KUANGXU@STANFORD.EDU
Graduate School of Business, Stanford University, Stanford, CA 94305

Zhi Xu ZHIXU@MIT.EDU
LIDS, Massachusetts Institute of Technology, Cambridge, MA 02139

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 with binary responses. In the meantime, an adversary observes the learner's queries, though 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 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.[1]

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.

The present 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 her to tailor future queries based on past responses and thus reduce the number of queries needed, while, at the same time, exposes the learner to substantial privacy risk: the more her queries depend on past responses, the easier it might be for an adversary to use the observed queries to infer past responses. Our main objective is to articulate and understand an intrinsic *privacy versus query complexity tradeoff* in the context of 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*, which lies in a bounded subset of $\mathbb{R}$. To search for v^*, she 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 $v^* \geq q_k$,

1. This version: February 2018.

and $r_k = 0$, otherwise. The interaction is sequential in the sense that the learner may choose a query depending on the responses of all previous queries. Meanwhile, there is an *adversary* who eavesdrops on the learner's actions: she observes all of 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 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 for her 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.

Two illustrative motivating applications of our model, on sequential price learning and online optimization, are discussed in Appendix A.

1.1. 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^*. We denote by ϵ and δ the absolute estimation error the learner and the adversary is willing to tolerate, respectively. We will employ 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 δ with probability at most $1/L$. A private query strategy for the learner must be able to produce an estimate of the true value within an error of at most ϵ, while simultaneously guaranteeing that the desired privacy level L holds against the adversary.

Our main objective is to quantify the *query complexity* of a private learner strategy, $N^*(\epsilon, \delta, L)$, 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 $2\epsilon < \delta \leq 1/L$. 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 the query complexity:

1. We establish an upper bound[2] 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 show 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.

2. All logarithms are taken with respect to base 2. To reduce clutter, non-integer numbers are to be understood as rounded upwards.

1.2. 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 $\log(1/\epsilon)$ (cf. 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 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 contrast, our definition of privacy measures the adversary's ability to perform a *specific* inference task.

In a different model, Tsitsiklis and Xu (2017) 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 her intended destination against an adversary who observes her past trajectories. 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 her destination. Similar to the setting of Tsitsiklis and Xu (2017), 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 she herself has no prior knowledge, in contrast to the agent in Tsitsiklis and Xu (2017) who tries to conceal private information already in her possession. In a way, the central conflict of trying to learn something while preventing others from learning the same information sets our work apart from the extant literature.

Our model is close in spirit to the private information retrieval problem 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 to 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.

2. The Private Sequential Learning Model

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[3] $[0, 1)$ and that the learner knows that this is the case. The true value is stored in an external database, and 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, i.e.,

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

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

Denote by N the total number of learner queries, and by $\epsilon > 0$ the learner's desired accuracy. After having received 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 all the queries but not the responses, and knows the learner's query strategy. Based on this information, and after observing all of 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 until 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 will 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, as well as 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^*.

Observe that knowing the value of the random seed Y and the responses to the queries is sufficient for reconstructing the values of these queries. Through induction, it then suffices to let ϕ_k be a function of just the responses, $r_1, \ldots, r_{k-1}$, and Y. This leads to the following formal definition of learner strategies:

1. If $k = 1$, then $\phi_1 : \{1, 2, \ldots, \mathcal{Y}\} \to [0, 1)$, and $q_1 = \phi_1(Y)$;

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

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

We will consider learner strategies that submit distinct queries, as 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 as above.

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 subsection 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 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, she would know exactly what q_k is for step k. While it may seem that an adversary who sees both the learner strategy and her actions is too powerful to defend against, we will see in the sequel that even in such a scenario the learner will still be able to implement effective and efficient obfuscation by exploiting the randomness of Y.

3. Private Learner Strategies

We now introduce and formally define private learning strategies, the central concept of this paper. 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 the adversary's probability of correct detection of v^* sufficiently small. Our goal in this section is to formalize these ideas by introducing the notions of information set and cover number.

3.1. Information Set

Recall from Section 2.2 that the adversary knows the values of the queries and the learner strategy. We will 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 family of sequences of queries that have a positive probability of appearing under ϕ, when the true value v^* is equal to x:

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

where Q is a vector-valued random variable representing the sequence of learner queries, whereas $\overline{q}$ stands for a typical realization, and that 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:

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

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. $(\epsilon, \delta, L)-$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, on the other hand, is more subtle, as it depends on what the adversary is able to infer. To this end, we develop in this subsection a privacy metric by measuring the "effective size" of the information set $\mathcal{I}(\overline{q})$ described in Definition 1. Intuitively, since 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. Since the adversary has an accuracy parameter of δ, we could require that the diameter of $\mathcal{I}(\overline{q})$ be greater than δ. The diameter, however, is not a good metric, as it paints an overly optimistic picture for the learner. Consider the example where the information set is the union of two intervals of length δ each, placed far apart from each other. By setting her estimate to be the center of one of the two intervals, chosen at random with equal probabilities, the adversary always has a probability of $1/2$ of correctly predicting the true value, even though the diameter of the information set could be large. The Lebesgue measure of the information set appears to be another plausible candidate. However, it also fails to accurately describe the learner's privacy. Consider again the example 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 though the Lebesgue measure of the set could be made arbitrarily small by reducing the lengths of the intervals.

The shortcomings of the above 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 will 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 $[a_1, b_1]$, $[a_2, b_2]$, $\ldots$, $[a_L, b_L]$, 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 $j = 1, 2, \ldots, L$.*

*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*

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

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

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

1. *Accuracy constraint: the learner estimate accurately recovers the true value, with probability one:*
$$\mathbb{P}\Big(\big|\hat{x}(x, Y) - x\big| \le \epsilon/2\Big) = 1, \quad \forall\, x \in [0, 1),$$
where the probability is measured with respect to the randomness in Y.

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

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

The accuracy constraint requires that a private learner strategy always produces 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: since 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 $1/L$. We make the link between $1/L$ and the adversary's probability of correct estimation precise in Appendix C.

4. Main Result

The learner's overall objective is to employ the minimum 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. Define $N^*(\epsilon, \delta, L)$ to be the minimum number of queries needed across all (ϵ, δ, L)-private learner strategies,

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

Our result will focus on the regime of parameters where

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

Having $2\epsilon < \delta$ corresponds to a scenario where the learner would like to identify the true value with high accuracy, while the adversary is aiming for a coarse estimate. Note that the regime where $\delta < \epsilon$ is arguably much less interesting, because it is not natural to expect the adversary, who is not engaged in the querying process, to have a higher accuracy requirement than the learner. The requirement that $\delta \leq 1/L$ stems from the following argument. If $\delta \geq 1/(L-1)$, then the entire interval $[0,1)$ is trivially $(\delta, L-1)$-coverable, and $C_\delta\big(\mathcal{I}(\overline{q})\big) \leq C_\delta\big([0,1)\big) \leq L - 1 < L$. 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 sightly broader than the regime $\delta \leq 1/L$ that we consider. The following theorem is the main result of this paper.

Theorem 4 (Query Complexity of Private Sequential Learning) *Fix $\epsilon > 0$, $\delta > 0$, and a positive integer $L \geq 2$, such that $2\epsilon < \delta \leq 1/L$. Then,*

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

The proof of the upper bound in Theorem 4 is constructive, providing a specific learner strategy that satisfies the bound. If we set $\delta = 1/L$, which corresponds to the worst case where the adversary's accuracy requirement is as loose as possible, Theorem 4 leads to the following corollary. It yields upper and lower bounds that are tight up to an additive constant of 4. In other words, the private learner strategy that we construct achieves essentially the optimal query-complexity in this scenario.

Corollary 5 *Fix $\epsilon > 0$ and a positive integer $L \geq 2$ such that $2\epsilon < 1/L$. The following holds.*

1. *If $L = 2$, we have that*
$$\log \frac{1}{\epsilon} \leq N^*\Big(\epsilon, \frac{1}{L}, L\Big) \leq \log \frac{1}{\epsilon} + 4. \tag{8}$$
2. *If $L \geq 3$, we have that*
$$\log \frac{1}{L\epsilon} + 2L - 4 \leq N^*\Big(\epsilon, \frac{1}{L}, L\Big) \leq \log \frac{1}{L\epsilon} + 2L. \tag{9}$$

A main take-away from the above 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 $2L$.

5. Proof of the Upper Bound: Opportunistic Bisection Strategy

We prove in this section the upper bound on the query complexity in Theorem 4. To obtain some insight on the difficulties in designing an efficient private learner strategy, before delving into the proofs in this section, we encourage the reader to examine three learner strategies in Appendix B, which are situated at different locations along the complexity-privacy tradeoff curve. Notably, the most sophisticated strategy among them (Replicated Bisection) achieves a query-complexity of $L\log(\delta/\epsilon)$, where the level of privacy, L, incurs a *multiplicative* overhead on query-complexity. We will show in this section that this query-complexity can be significantly improved so that the overhead becomes only *additive* in L. This will be achieved by constructing a specific learner strategy, Opportunistic Bisection (OB). We start with some terminology, to facilitate the exposition.

Definition 6 *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 i.i.d. Bernoulli random variables, with $\mathbb{P}(Z_1 = 0) = 1/2$. Let $(q_1, q_2, \ldots, q_M)$ be a sequence of M queries, where q_1 is equal to the mid-point of $\mathcal{J}$, and let $(r_1, r_2, \ldots, r_M)$ their corresponding responses.*

1. *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) *q_i is set to the mid-point of interval $\mathcal{J}_i$.*
   (b) *$\mathcal{J}_{i+1}$ is set to*
$$\mathcal{J}_{i+1} = \begin{cases} [\inf \mathcal{J}_i\,,\, q_i)\,, & \text{if } r_i = 0, \\ [q_i\,,\, \sup \mathcal{J}_i)\,, & \text{if } r_i = 1. \end{cases} \tag{10}$$
2. *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) q_i *is set to the mid-point of interval* $\mathcal{J}_i$.

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

$$\mathcal{J}_{i+1} = \begin{cases} [\inf \mathcal{J}_i \,,\, q_i)\,, & \text{if } Z_i = 0, \\ [q_i \,,\, \sup \mathcal{J}_i)\,, & \text{if } Z_i = 1. \end{cases} \tag{11}$$

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 Opportunistic Bisection 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

$$q_i = (i-1)\frac{1}{L}, \quad i = 1, \ldots, L, \tag{12}$$

and

$$(q_{L+1}, q_{L+2}, \ldots, q_{2L}) = (q_1 + \epsilon, q_2 + \epsilon, \ldots, q_L + \epsilon). \tag{13}$$

Notice that 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 will 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 hence 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:

$$\mathcal{J}^{(i)} = [q_{L+i},\, q_{i+1}) = \left[(i-1)\frac{1}{L} + \epsilon,\, \frac{i}{L}\right), \quad i = 1, 2, \ldots, L. \tag{14}$$

We will refer to $\mathcal{J}^{(i)}$ as the *ith sub-interval*. Importantly, by the end of Phase 1, if none of the guesses contains the true value, then the learner knows which sub-interval contains the true value, which we will 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 $M = \log\left(\frac{1}{\epsilon L}\right)$.

2. If one of the guesses in Phase 1 contains v^*, then, let $\tilde{\mathcal{J}}$ be a sub-interval chosen uniformly at random among all L sub-intervals, and let $(q_{2L+1}, q_{2L+2}, \ldots, q_{2L+M})$ be a **fictitious bisection search** of $\tilde{\mathcal{J}}$ with $M = \log\left(\frac{1}{\epsilon L}\right)$, using the randomization provided by Y (i.e., using Y to generate the sequence of i.i.d. Bernoulli random variables, Z_i).

Remark. It is interesting to contrast Opportunistic Bisection with the Replicated Bisection strategy in Appendix B. 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*. Crucially, the

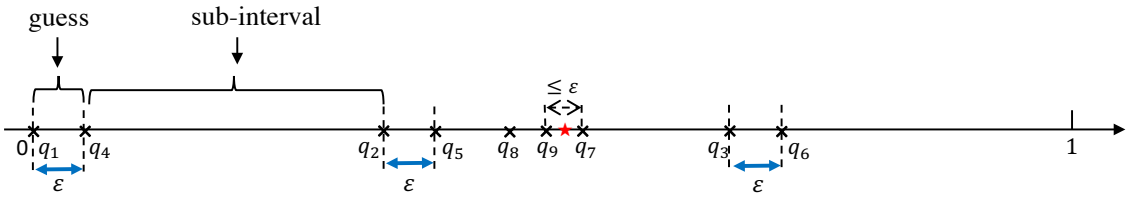

Figure 1: An example of the Opportunistic Bisection Strategy, with $L = 3$.

guesses make it possible to immediately discover the location of the true value in the first phase, albeit such discoveries might be unlikely. In the second stage, while the Replicated Bisection strategy conducts a bisection search in *each* of the L sub-intervals, the Opportunistic Bisection strategy does so in only *one* of the sub-intervals, hence drastically reducing the number of queries.

It follows directly from the definition that the number of queries submitted under the Opportunistic Bisection strategy is:

$$N = 2L + \log\left(\frac{1}{\epsilon L}\right). \tag{15}$$

To complete the proof of the upper bound in Theorem 4, 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. The proof is given in Appendix E.1.

Proposition 7 *Fix $\epsilon > 0$, $\delta > 0$, and a positive integer $L \geq 2$, such that $2\epsilon < \delta \leq 1/L$. Then, the Opportunistic Bisection (OB) strategy is (ϵ, δ, L)-private.*

6. Proof of the Lower Bound: Elementary Intervals

We now derive the two lower bounds on query complexity in Theorem 4. Note that the query-complexity of the Opportunistic Bisection strategy carries a $2L$ overhead compared to the (non-private) 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 2 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 this is not possible: in order to successfully obfuscate the true value, one needs $2L$ that are *distinct* from those that participate in the bisection algorithm.

We first introduce some notation to facilitate our discussion. Recall that $(q_1, q_2, \ldots, q_N)$ is the sequence of learner queries. In the remainder of this section, we will 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 it is trivially true that $0 \leq v^* < 1$, and hence adding q_0 and q_{N+1} does not provide additional information to either the learner or the adversary. For a sequence of queries, $\overline{q} = (q_0, q_1, \ldots, q_{N+1})$ and $x \in [0, 1)$, we will denote by $H(\overline{q}, x)$ the smallest interval formed by points in $\overline{q}$ that contains x:

$$H(\overline{q}, x) = \Big[\max\{z \in \overline{q} : z \leq x\},\ \min\{z \in \overline{q} : z > x\}\Big). \tag{16}$$

We first examine some basic properties of a learner strategy ϕ that satisfies the accuracy constraint, but without necessarily being private. In such a setting, the goal of the learner is solely to learn the true value without ever worrying about the information revealed. Intuition suggests that in order to produce an accurate estimate, the learner must find an interval that contains the true value v^* and whose length is at most ϵ. We confirm this is indeed necessary in the next lemma. Fix a learner strategy ϕ. Define the event $\mathcal{E}$

$$\mathcal{E} = \Big\{ \exists\, q_i, q_j \in Q, \text{ such that } 0 < q_j - q_i \leq \epsilon \text{ and } v^* \in [q_i, q_j) \Big\} \tag{17}$$

In words, $\mathcal{E}$ is the event that there exist two queries separated by a distance less than or equal to ϵ, such that the interval whose end points correspond to the two queries contains v^*. The proof of the lemma is given in Appendix E.2.

Lemma 8 *Fix any learner strategy ϕ that satisfies the accuracy constraint in Definition 3. Then, for every $v^* \in [0,1)$, we have $\mathbb{P}(\mathcal{E}) = 1$, where the probability is measured with respect to the randomness in Y.*

Lemma 8 states that under any accurate learner strategy, the true value will always be situated between two queries spaced at most ϵ apart. Building on this insight, we are now ready to introduce the main building block for proving the lower bound: elementary intervals, which are intervals formed by the queries that are at most ϵ-long.

Definition 9 (Elementary Interval and Query) *Fix $\epsilon > 0$, and a sequence of queries $\overline{q} = (q_0, q_1, \ldots, q_{N+1})$. For two queries q_i, q_j that are part of $\overline{q}$, with $q_j > q_i$, we say that the interval $[q_i, q_j)$ is an* elementary interval *if $q_j - q_i \leq \epsilon$. Similarly, we say that a query q_i is* elementary*, if there exists some q_j, such that $|q_i - q_j| \leq \epsilon$.*

Define $\overline{\mathcal{I}}(\epsilon, \overline{q})$ to be the union of all elementary intervals formed by the queries in $\overline{q} = (q_0, q_1, \ldots, q_{N+1})$:

$$\overline{\mathcal{I}}(\epsilon, \overline{q}) := \bigcup_{0 < q_j - q_i \leq \epsilon} [q_i, q_j), \tag{18}$$

The next lemma states an important structural property of the adversary's information set: the information set only contains points on $[0,1)$ that fall within some elementary interval. The proof is a straightforward application of Lemma 8, and is given in Appendix E.3.

Lemma 10 *Fix $\epsilon > 0$ and a learner strategy ϕ that satisfies the accuracy constraint in Definition 3. Then, for any $v^* \in [0,1)$, the information set for the adversary must contain only points that lie in some elementary interval, i.e., $\mathcal{I}\Big(\overline{q}(v^*, Y)\Big) \subset \overline{\mathcal{I}}\Big(\epsilon, \overline{q}(v^*, Y)\Big)$, with probability one.*

6.1. Completing the Proof of the Lower Bound

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

Lemma 11 *Fix a learner strategy ϕ that satisfies the accuracy constraint, a constant $b \in [0,1)$, and an interval $J \subset (0,1)$ of length b. Then, for every $y \in \{1, 2, \ldots, \mathcal{Y}\}$, there exists $x \in J$ such that*

$$|\overline{q}(x, y) \cap J| \geq \log(b/\epsilon), \tag{19}$$

where $|\cdot|$ stands for the cardinality of a set.

A corollary of the lemma is that there exists $x \in J$ such that when $v^* = x$, then at least $\log(b/\epsilon)$ queries must fall within J. By setting b to 1 (i.e., J being the interval $(0,1)$), we recover the (classical) lower bound that at least $\log(1/\epsilon)$ queries are necessary to achieve accuracy; This proves the first term of the lower bound in Theorem 4. Lemma 11 generalizes the classical result that $\log(1/\epsilon)$ query complexity of the bisection strategy is optimal (cf. Waeber et al. (2013)), and can be similarly proved using a recursive argument. We omit the proof, which is fairly standard, but provide an intuitive argument on why the lemma holds. Fix $y \in \{1, 2, \ldots, \mathcal{Y}\}$. Note that the interval J consists of b/ϵ disjoint sub-intervals of length ϵ each. An accurate learner strategy, therefore, must be able to distinguish in which one of these sub-intervals the true value resides, and distinguishing among b/ϵ possibilities using binary feedback therefore implies that there will be some $v^* \in J$ whose accurate identification requires $\log(b/\epsilon)$ queries in J.

To strengthen the lower bound provided by Lemma 11, we will incorporate the privacy constraint by exploiting the structural property of the information set obtained in Lemma 10. Fix an (ϵ, δ, L)-private learner strategy $\phi \in \Phi_N$. Applying Lemma 11 with $J = (0, \delta)$ and $b = \delta$, we conclude that there exists $x_0 \in (0, \delta)$ and $y_0 \in \{1, 2, \ldots, \mathcal{Y}\}$, such that

$$|\overline{q}(x_0, y_0) \cap (0, \delta)| \geq \log(\delta/\epsilon), \tag{20}$$

that is, if $v^* = x_0$ and $Y = y_0$, then there will be at least $\log(\delta/\epsilon)$ queries in the interval $[0, \delta)$. Fix $v^* = x_0$ and $Y = y_0$.

The next lemma is the main technical result of this subsection, and shows that there must be at least $2L - 4$ queries in the interval $[\delta, 1)$; the proof is given in Appendix E.4.

Lemma 12 *If $\mathcal{I}(\overline{q}(x_0, y_0))$ is not $(\delta, L-1)$-coverable, i.e., $C_\delta\Big(\mathcal{I}(\overline{q}(x_0, y_0))\Big) \geq L$, then there must be at least $2L - 4$ queries in $\overline{q}(x_0, y_0)$ that belong to the interval $[\delta, 1)$.*

Combining Lemma 12 with Eq. (20), we conclude that there must be at least $\log(\delta/\epsilon) + 2L - 4$ queries (excluding $q_0 = 0$ and $q_{N+1} = 1$). We have thus proven the lower bound, and together with the upper bound in Section 5, completed the proof of Theorem 4.

7. Conclusions and Future Work

This paper studies an intrinsic privacy-complexity tradeoff faced by a learner in a sequential learning problem while trying to conceal her findings from an observant adversary. We use the notion of information set, the set of possible true values, to capture the amount of 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, i.e., that the resulting information set requires at least L intervals of size δ to be fully covered, it is necessary for the learner to employ 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.

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 $v^* \in \mathbb{R}^n$. A query in this setting will be a hyperplane in $\mathbb{R}^n$, where the response indicates whether the true value is to the right or to the left of the queried hyperplane. Another interesting direction is to consider adversaries with

varying levels of risk aversion. The present model considers a risk-averse adversary who considers all points in $[0, 1)$ as equally "plausible" as long as they have a positive probability of being close to the true value. In contrast, a less risk-averse adversary may entirely ignore points that are less likely to be close to the true value, thus reducing the size of the information set. One such variation, the Bayesian Private Learning model, is explored in Appendix D.

References

Avrim Blum, Katrina Ligett, and Aaron Roth. A learning theory approach to noninteractive database privacy. *Journal of the ACM (JACM)*, 60(2):12, 2013.

Kamalika Chaudhuri and Daniel Hsu. Sample complexity bounds for differentially private learning. In *Proceedings of the Conference on Learning Theory*, pages 155–186, 2011.

Benny Chor, Eyal Kushilevitz, Oded Goldreich, and Madhu Sudan. Private information retrieval. *Journal of the ACM*, 45(6):965–981, 1998.

John Duchi, Martin Wainwright, and Michael Jordan. Minimax optimal procedures for locally private estimation. In *arXiv preprint arXiv:1604.02390*, 2016.

Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3–4):211–407, 2014.

Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography Conference*, pages 265–284. Springer, 2006.

Vitaly Feldman and David Xiao. Sample complexity bounds on differentially private learning via communication complexity. In *Conference on Learning Theory*, pages 1000–1019, 2014.

William Gasarch. A survey on private information retrieval. In *Bulletin of the EATCS*. Citeseer, 2004.

Michael Horstein. Sequential transmission using noiseless feedback. *IEEE Transactions on Information Theory*, 9(3):136–143, 1963.

Prateek Jain, Pravesh Kothari, and Abhradeep Thakurta. Differentially private online learning. In *Conference on Learning Theory*, pages 24–1, 2012.

Eyal Kushilevitz and Rafail Ostrovsky. Replication is not needed: Single database, computationally-private information retrieval. In *Foundations of Computer Science (FOCS)*, volume 97, pages 364–373, 1997.

Sofya Raskhodnikova, Adam Smith, Homin K Lee, Kobbi Nissim, and Shiva Prasad Kasiviswanathan. What can we learn privately. In *Proceedings of the 54th Annual Symposium on Foundations of Computer Science*, pages 531–540, 2008.

Adam Smith. Privacy-preserving statistical estimation with optimal convergence rates. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 813–822. ACM, 2011.

John Tsitsiklis and Kuang Xu. Delay-predictability tradeoffs in reaching a secret goal. *To appear in Operations Research*, 2017. URL http://dx.doi.org/10.2139/ssrn.2784502.

Rolf Waeber, Peter I Frazier, and Shane G Henderson. Bisection search with noisy responses. *SIAM Journal on Control and Optimization*, 51(3):2261–2279, 2013.

Larry Wasserman and Shuheng Zhou. A statistical framework for differential privacy. *Journal of the American Statistical Association*, 105(489):375–389, 2010.

Appendix A. Motivating Examples

We examine two illustrative example applications of our model.

Example 1 - learning an optimal price. A firm is to release a new product and would like to identify a revenue maximizing price, p^*, prior to the product launch. The firm believes that the revenue function, $f(p)$, is strictly concave and differentiable as a function of the price, p, but has otherwise little additional information. A sequential learning process is employed 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 firm would like to estimate p^* with reasonable accuracy over a small number of epochs, but is wary that a competitor might be able to observe the surveys and deduce from them 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 - online optimization with private weights. In the previous example, the adversary is a third-party entity who does not observe the responses to the queries. We now illustrate in this example that the Private Sequential Learning model can also describe a situation where 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. The query strategies we develop for Private Sequential Learning can also be applied in this setting.

Appendix B. Examples of Learner Strategies

To provide some intuition and motivation for the design of the Opportunistic Bisection strategy in Section 5, we examine three representative learner strategies.

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 $[0, 1)$, i.e., $q_1 = 0.5$. Then, based on the response, the learner identifies the half interval that contains the true value, and subsequently submits its midpoint as the next query, q_2. The process continues recursively until the learner finds an interval of length at most ϵ that contains the true value v^*.

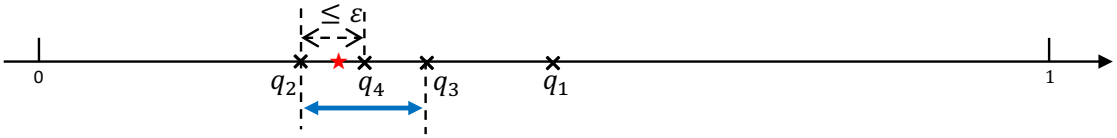

Figure 2: An example of the bisection strategy where the red star represents the true value v^*. The dashed line with arrows represents the learner's error tolerance, and the solid line with arrows represents the information set for the adversary, $\mathcal{I}(\overline{q})$.

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 to the last query the learner submitted, and hence the δ-cover number of the information set is always 1 whenever $\delta > 2\epsilon$. As such, the bisection strategy lies at one extreme end of the complexity-privacy tradeoff, with a minimum query complexity but no privacy.

Strategy 2: ϵ-Dense. Sitting on the opposite end of the spectrum is the ϵ-Dense strategy, where the learner submits a sequence of $N = 1/\epsilon - 1$ pre-determined queries, with $q_1 = \epsilon$, $q_2 = 2\epsilon, \ldots, q_N = N\epsilon$. The strategy is accurate because the distance between two adjacent queries is equal to the error tolerance, ϵ. Moreover, because the sequence of queries is pre-determined and does not depend on the location of the true value, the adversary obtains no information from the learner's query patterns, and the information set remains the interval $[0, 1)$ throughout. Thus, if $\delta \leq 1/L$, the strategy is (ϵ, δ, L)-private. Compared to the Bisection strategy, the perfect privacy of 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.

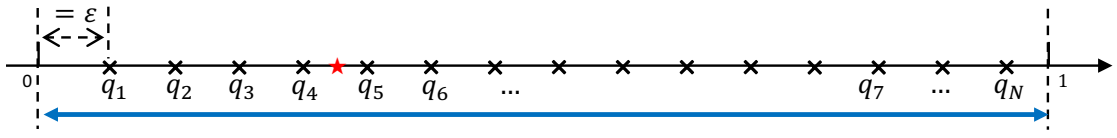

Figure 3: An example of the ϵ-Dense strategy. The dashed line with arrows represents the learner's error tolerance, and the solid line with arrows represents the information set for the adversary, $\mathcal{I}(\overline{q})$.

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. This strategy has two phases:

1. *Phase 1 - Deterministic Queries*. The learner submits $L-1$ queries, chosen deterministically:
$$q_1 = \frac{1}{L},\ q_2 = \frac{2}{L}, \ldots,\ q_{L-1} = \frac{L-1}{L}, \tag{21}$$
which partition the unit interval into L disjoint sub-intervals of length $1/L$ each: $[0, 1/L)$, $[1/L, 2/L), \ldots, [1-1/L, 1)$. At this point, the learner has learned which one of the L sub-intervals contains the true value, while the adversary has gained no additional information about the true value. We will refer to the sub-interval that contains the true value as the *true sub-interval*, and all other sub-intervals as *false sub-intervals*. This phase uses $L-1$ queries.

2. *Phase 2 - Replicated Bisection*. In the second phase, the learner conducts a bisection strategy within the true sub-interval until the true value has been located, while in the meantime submitting translated replicas of these queries in each false sub-intervals, 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 sub-interval. 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 sub-interval.

When the process is completed, the learner will have identified the true value via the bisection strategy within the true sub-interval, while 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 to that of the Bisection strategy ($N = \log(1/\epsilon)$).

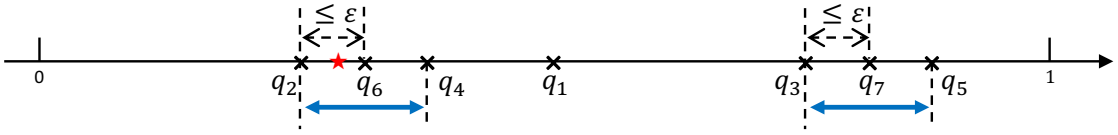

Figure 4: An example of the Replicated Bisection strategy, with $L = 2$. The dashed line with arrows represents the learner's error tolerance, and the solid line with arrows represents the information set for the adversary, $\mathcal{I}(\overline{q})$.

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 queries than ϵ-Dense ($N = 1/\epsilon$). Is it an optimal strategy, achieving the minimal query complexity for a given privacy level, L? Perhaps surprisingly, the answer is negative. Our proof for the upper bound of Theorem 4 will show that the query complexity of the Replicated Bisection strategy can be much improved, so that the query complexity overhead, compared to the Bisection strategy, is only an *additive* factor of L.

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

In this section we argue that the quantity $1/L$ can be interpreted as a probability of correct detection for the adversary. Recall the definition of set $\mathcal{Q}(x)$ of possible sequences (cf. Eq. (1)), and let $\mathcal{Q} = \cup_{x\in[0,1)}\mathcal{Q}(x)$. We consider here adversary estimates $\hat{x}^a$ that are random variables, determined by the observed query sequence $\overline{q}$, together with an independent randomization seed.

Definition 13 *Fix $\delta > 0$, $L \geq 2$, a learner strategy $\phi \in \Phi_N$, and a sequence of queries $\overline{q} \in \mathcal{Q}$. We say that an adversary estimate, $\hat{x}^a$, is* (δ, L)-correct *given $\overline{q}$, if*

$$\mathbb{P}\Big(|\hat{x}^a(\overline{q}) - x| \leq \delta/2\Big) > \frac{1}{L}, \quad \forall\, x \in \mathcal{I}(\overline{q}). \tag{22}$$

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

In words, an adversary estimate is (δ, L)-correct given $\overline{q}$ if, as soon as the learner deploys the queries $\overline{q}$, the adversary will *know* that the estimate will incur an error of at most $\delta/2$ with probability at least $1/L$. In a sense, this means that $\overline{q}$ conceals the true value "poorly." The following proposition, the main result of this section, shows that the $(\delta, L-1)$ coverability of the information set is effectively equivalent to the existence of a (δ, L)-correct estimate.

Proposition 14 *Fix $\delta > 0$, $L \geq 2$, a learner strategy $\phi \in \Phi_N$, and a sequence of queries $\overline{q} \in \mathcal{Q}$. The following hold.*

1. *If the δ-cover number of the information set $\mathcal{I}(\overline{q})$, $C_\delta\big(\mathcal{I}(\overline{q})\big)$, is at most $L-1$, then there exists an adversary estimate that is (δ, L)-correct given $\overline{q}$.*
2. *Conversely, if $C_\delta\left(\mathcal{I}(\overline{q})\right)$ is at least L, then, for any $\delta' < \delta$, there does not exist an adversary estimate that is (δ', L)-correct given $\overline{q}$.*

Proof To prove the first statement, fix $\overline{q} \in \mathcal{Q}$ such that $\mathcal{I}(\overline{q})$ is $(\delta, L-1)$-coverable. Then, there exist $L-1$ intervals, $[a_1, b_1]$, $[a_2, b_2]$, $\ldots$, $[a_{L-1}, b_{L-1}]$, each of length δ, that cover $\mathcal{I}(\overline{q})$. Consider a randomized adversary estimate $\hat{x}^a$ that is distributed uniformly at random among the $L-1$ midpoints of the intervals. It is not difficult to show that with probability $1/(L-1)$, the estimate $\hat{x}^a$ will lie in the same interval as the true value; since all intervals have length at most δ, such a $\hat{x}^a$ will be at a distance of at most $\delta/2$ from the true value, i.e.,

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

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

We now prove the second statement. We will make use of the following lemma.

Lemma 15 *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*

$$|x_i - x_j| \geq \delta, \quad \forall i \neq j. \tag{23}$$

Proof We will prove the lemma by constructing the x_j's explicitly. Consider the following procedure:

1. $x_1 = \inf J$.
2. For $i = 2, 3, \ldots$, define x_i recursively as:

$$x_i = \inf\{x \in J : x \geq x_{i-1} + \delta\}. \tag{24}$$

The procedure terminates at some step T when $x_T + \delta \geq 1$ or $[x_T + \delta, 1) \cap J = \emptyset$. Note that by construction, all x_i's belong to the closure of J. Furthermore, the intervals

$$W_i := [x_i, x_i + \delta], \quad i = 1, 2, \ldots, T, \tag{25}$$

form a cover of J. Since $C_\delta(J) \geq L$ by assumption, it follows that we must have $T \geq L$. It is easy to verify that the points $\{x_1, x_2, \ldots, x_L\}$ satisfy the conditions outlined in the lemma, and which completes the proof. ■

Fix any adversary estimate $\hat{x}^a$, $\gamma \in (\delta', \delta)$, and $\overline{q} \in \mathcal{Q}$ such that $C_\delta\big(\mathcal{I}(\overline{q})\big) \geq L$. Apply Lemma 15 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's belong to the closure of $\mathcal{I}(\overline{q})$, by 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

$$|\tilde{x}_i - \tilde{x}_j| \geq \gamma > \delta', \quad \forall i \neq j. \tag{26}$$

Define intervals

$$U_i := [\tilde{x}_i - \delta'/2,\ \tilde{x}_i + \delta'/2], \quad i = 1, 2, \ldots, L. \tag{27}$$

Since the distance between any two distinct $\tilde{x}_i$'s is greater than δ', we know that the intervals U_i are disjoint, which implies that there exists $i^* \in \{1, 2, \ldots, L\}$ such that

$$\mathbb{P}\Big(|\widehat{x}^a(\overline{q}) - \tilde{x}_{i^*}| \leq \delta'/2\Big) \leq \mathbb{P}(\widehat{x}^a(\overline{q}) \in U_{i^*}) \leq 1/L. \tag{28}$$

Since $\tilde{x}_{i^*} \in \mathcal{I}(\overline{q})$ by construction, we conclude that the adversary estimate $\hat{x}^a$ is not (δ', L)-correct given $\overline{q}$. This completes the proof of Proposition 14. ■

Appendix D. 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 regarding the true value v^*, and they can obtain such information only gradually, through queries. In this section, we discuss a Bayesian variation of the model, where the true value v^* admits a prior distribution, known to both parties. In particular, suppose that v^* is distributed according to a distribution P_{v^*}, where, for simplicity, we assume that the support of P_{v^*} is equal to $[0, 1)$.

We modify the definition of a learner strategy ϕ in Section 2.1 by adding an extra input, P_{v^*}, to each query function ϕ_i and the estimation function ϕ^E. In addition, since the adversary aims to produce an estimate $\hat{x}^a$ that is close to the true value, we define an adversary strategy ψ to be a function that maps the adversary's available information to a probability distribution over $[0, 1)$. Formally, ψ is a function that takes as input the prior distribution P_{v^*}, the learner strategy ϕ, and the sequence of observed queries $\overline{q}$, and outputs a random variable, $\hat{x}^a$, defined over $[0, 1)$. Denote by Ψ the set of all such functions, i,e., the set of all adversary strategies.

Note that, since 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, which leads to the following definition of private learner strategies.

Definition 16 *Fix $\epsilon > 0, \delta > 0$, $L \in \mathbb{N}$, and a prior distribution P_{v^*}. A learner strategy $\phi \in \Phi_N$ is (ϵ, δ, L)-private if it satisfies the following:*

1. *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.

2. *Privacy constraint: for every adversary strategy $\psi \in \Psi$, we have*

$$\mathbb{P}\big(|\hat{x}^a - v^*| \leq \delta/2\big) \leq \frac{1}{L}, \tag{29}$$

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

Notice the resemblance of the above definition with Definition 3. The parameters ϵ and δ have the same meaning as in the original model, while 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 $N^*(\epsilon, \delta, L)$ the minimum number of queries needed for there to exist an (ϵ, δ, r)-private learner strategy:

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

Similar to the original model, we would like to obtain lower and upper bounds on $N^*(\epsilon, \delta, L)$. 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 17 *Fix $\epsilon > 0$, $\delta > 0$, and a positive integer $L \geq 2$, such that $2\epsilon < \delta \leq 1/L$. Suppose that the prior distribution P_{v^*} is uniform over $[0, 1)$. Then,*

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

Note that compared to the bounds in Theorem 4, the leading terms in the upper and lower bounds in Proposition 17 differ by a multiplicative factor of L. The lower bound, $\log(1/\epsilon)$ follows directly from the accuracy constraint (see Section 6). However, the proof for the stronger lower bound in Theorem 4, $\log(\delta/\epsilon) + 2L - 4$, does not apply to this Bayesian setting. This is due to the fact that a large information set, measured by its coverability, does not necessarily imply a high probability of error for the adversary in the Bayesian setting: this is because parts of the information set could have very small posterior probability of containing the true value, and thus could effectively be ignored by the adversary.

For the same reason, the Opportunistic Bisection strategy in Section 2.1 is no longer private in the Bayesian setting, since the guesses now have negligible posterior probability of containing v^*, and hence do not provide effective obfuscation against the adversary. The upper bound in Proposition 17 can be achieved by the Replicated Bisection strategy in Appendix B, with L replications.

Because the Replicated Bisection strategy creates L identical copies of query patterns across L sub-intervals, the symmetry ensures that the posterior distribution of v^* will be evenly spread across all sub-intervals, forcing the adversary's probability of correct detection to be at most $1/L$.

How to close the gap between the upper and lower bounds in Proposition 17 is an interesting open question. Intuition seems to suggest that the lower bound can be substantially improved, and that the upper bound could be (close to) optimal. Unfortunately, our current proof technique for the lower bound is highly dependent on the structure of the information set and is unlikely to be able to close this gap.

Appendix E. Proofs

E.1. Proof of Proposition 7

Proof We first show that the OB strategy is accurate, and specifically, that it will allow 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 mid-point of the guess, i.e., $\hat{x} = (q_{i'} + q_{i'+L})/2$. Since the length of each guess is exactly ϵ, we must 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 sub-interval 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 sub-interval 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 mid-point 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, $C_\delta\big(\mathcal{I}(\overline{q})\big)$, is at least L. Denote by $\mathcal{G}$ the union of the guesses, i.e.,

$$\mathcal{G} = \bigcup_{i=1}^{L} [q_i, q_{i+L}). \tag{30}$$

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

Claim 18 *The δ-cover number of $\mathcal{G}$, $C_\delta(\mathcal{G})$, is at least L.*

Claim 19 *The information set, $\mathcal{I}(\overline{q})$, contains $\mathcal{G}$.*

We first show Claim 18. Fix an interval $J \subset [0, 1)$ with length δ. Note that by construction, each guess has length ϵ, and two adjacent guesses are separated by a distance of $1/L - \epsilon$. Since $\delta \leq 1/L$, this implies that the Lebesgue measure of $J \cap \mathcal{G}$ is at most ϵ. Since the Lebesgue measure of $\mathcal{G}$ is ϵL, we conclude that it will require at least L intervals of size δ to cover $\mathcal{G}$. Therefore, $C_\delta(\mathcal{G}) \geq L$. This proves Claim 18.

We next show Claim 19. Recall that, given the sequence of queries $\overline{q}$, the information set contains a point $x \in [0, 1)$ if, in the case of $v^* = x$, the sequence $\overline{q}$ has a strictly positive probability of occurring under the learner strategy. Therefore, showing that $\mathcal{G} \subset \mathcal{I}(\overline{q})$ amounts to demonstrating

that, regardless of the realization of the queries, the adversary will never be able to assert, with absolute certainty, that a point in $\mathcal{G}$ is not the true value.

We consider two possibilities. First, suppose that the true value is inside one of the guesses, i.e., $v^* = x$ for some $x \in \mathcal{G}$. Because the guesses are deterministic, and the fictitious bisection search that follows in Phase 2 is independent from the location of v^*, we conclude that all $x \in \mathcal{G}$ would have produced the same distribution for the queries. This implies that all points in $\mathcal{G}$ must belong to the information set, with probability one.

Secondly, suppose that $v^* \notin \mathcal{G}$. Let Q^{true} be the set of all queries that could be produced by a truthful bisection search in Phase 2. Under our construction, we have that,

$$\mathbb{P}_{v^*=x}\Big((q_{2L+1}, q_{2L+2}, \ldots, q_N) = \overline{q}\Big) = L^{-1} \cdot \big(2^{-1}\big)^{\log\left(\frac{1}{\epsilon L}\right)-1} = 2\epsilon > 0, \quad \text{for all } x \in \mathcal{G} \text{ and } \overline{q} \in Q^{\text{true}}, \tag{31}$$

where the notation $\mathbb{P}_{v^*=x}(\cdot)$ denotes the probability law of the queries when $v^* = x$, and where the quantity $L^{-1} \cdot \big(2^{-1}\big)^{\log\left(\frac{1}{\epsilon L}\right)-1}$ follows from the definition of a fictitious bisection search. Eq. (31) demonstrates that, when $v^* \notin \mathcal{G}$, the resulting queries submitted during Phase 2 (a truthful bisection search) could have equally been produced under *any* alternative value of v^* that belongs to $\mathcal{G}$, with strictly positive probability. This implies that $\mathcal{G} \subset \mathcal{I}(\overline{q})$, and thus proves Claim 19.

Claims 18 and 19 together show that the OB strategy is private, and this completes the proof of Proposition 7. ■

E.2. Proof of Lemma 8

Proof In the remainder of the proof, we will use the notation $\mathbb{P}_{v^*=x}(\cdot)$ to signify the probability law when the true value is equal to x. Suppose, for the sake of contradiction, that there exists an accurate learner strategy, ϕ_0, and $x_0 \in [0, 1)$, such that $\mathbb{P}_{v^*=x_0}(\mathcal{E})$ is strictly less than 1 under ϕ. We will show that there exist x_1, x_2 that are sufficiently far from each other, so that the learner estimates under ϕ, $\hat{x}$, satisfy

$$\mathbb{P}_{v^*=x_1}(\hat{x} = x_3) > 0, \text{ and } \mathbb{P}_{v^*=x_2}(\hat{x} = x_3) > 0, \quad \text{for some } x_3 \in [0, 1). \tag{32}$$

Since x_1 and x_2 are separated with a sufficient distance, x_3 cannot be simultaneously $\epsilon/2$-close to both x_1 and x_2. Eq. (32) would therefore contradict the assumption of ϕ being accurate.

Fix the learner strategy to ϕ_0. Recall the definition of the set $\mathcal{Q}(x)$ in Eq. (1), and of $H(\overline{q}, x)$, the smallest interval formed by queries in $\overline{q}$ that contains x, in Eq. (16). Let $\mathcal{Q}_\epsilon(x)$ be the subset of $\mathcal{Q}(x)$ comprised of those $\overline{q}$ such that the length of $H(\overline{q}, x)$ is greater than ϵ:

$$\mathcal{Q}_\epsilon(x) = \{\overline{q} \in \mathcal{Q}(x) \ : \ |H(\overline{q}, x)| > \epsilon\}.$$

We are now ready to prove our claim. The assumption that $\mathbb{P}_{v^*=x_0}(\mathcal{E}) < 1$ implies

$$\mathbb{P}_{v^*=x_0}\left(\overline{q}(x_0, Y) \in \mathcal{Q}_\epsilon(x_0)\right) > 0. \tag{33}$$

In other words, there exists $y_0 \in \{1, 2, \ldots, \mathcal{Y}\}$ such that $\overline{q}(x_0, y_0) \in \mathcal{Q}_\epsilon(x_0)$. We make the following observation: by the definition of y_0, we have that

$$\overline{q}(x, y_0) = \overline{q}(x_0, y_0), \quad \text{for all } x \in H(\overline{q}(x_0, y_0), x_0). \tag{34}$$

To see why the above equality is true, note that, for a fixed learner strategy, the query submitted at any point in time only depends on the value of Y and the responses from earlier queries. When $v^* = x_0$ and $Y = y_0$, no queries are submitted in the interval $H(\overline{q}(x_0, y_0), x_0)$, and therefore, when $Y = y_0$ the learner will see the exact same responses, and thus same queries, whenever $v^* \in H(\overline{q}(x_0, y_0), x_0)$.

Crucially, Eq. (34) implies that the learner will produce the same estimate $\hat{x}$ whenever v^* lies within the interval $H(\overline{q}(x_0, y_0), x_0)$ and $Y = y_0$. Note that by definition, $\overline{q}(x_0, y_0) \in \mathcal{Q}_\epsilon(x_0)$, and hence

$$|H(\overline{q}(x_0, y_0), x_0)| > \epsilon. \tag{35}$$

We thus conclude that there exist $x_1, x_2 \in H(\overline{q}(x_0, y_0), x_0)$, with $|x_1 - x_2| > \epsilon$, such that the learner estimates satisfy:

$$\hat{x}(x_1, y_0) = \hat{x}(x_2, y_0) \stackrel{\triangle}{=} x_3. \tag{36}$$

This implies that x_3 cannot be simultaneously close to x_1 and x_2 by a distance of $\epsilon/2$, and hence either $\hat{x}(x_1, y_0)$ or $\hat{x}(x_2, y_0)$ must be inaccurate. We thus reach a contradiction with the accuracy of ϕ_0, which completes the proof of Lemma 8. ■

E.3. Proof of Lemma 10

Proof For the sake of contradiction, suppose that there exist $x_0, x' \in [0, 1)$ and $y' \in \{1, 2, \ldots, \mathcal{Y}\}$ such that

$$x' \in \mathcal{I}\Big(\overline{q}(x_0, y')\Big), \text{ and } x' \notin \overline{\mathcal{I}}\Big(\epsilon, \overline{q}(x_0, y')\Big). \tag{37}$$

Since x' belongs to the set $\mathcal{I}\left(\overline{q}(x_0, y')\right)$, the definition of the information set implies that

$$\mathbb{P}_{v^*=x'}\Big(Q = \overline{q}(x_0, y')\Big) > 0. \tag{38}$$

The assumption that $x' \notin \overline{\mathcal{I}}\left(\epsilon, \overline{q}(x_0, y')\right)$ implies that $H\left(\overline{q}(x_0, y'), x'\right)$, the smallest interval formed by queries in $\overline{q}(x_0, y')$ that contains x', must have length greater than ϵ. By Lemma 8, this implies that $\mathbb{P}_{v^*=x'}\left(Q = \overline{q}(x_0, y')\right) = 0$, which leads to a contradiction with Eq. (38). This proves the lemma. ■

E.4. Proof of Lemma 12

Proof Suppose, for the sake of contradiction, that there are at most $2L - 5$ queries in the interval $[\delta, 1)$, or, equivalently, at most $2L - 4$ queries in the interval $[\delta, 1]$ (including the trivial query $q_{N+1} = 1$). We will show that under this assumption the information set, $\mathcal{I}\left(\overline{q}(x_0, y_0)\right)$, must be $(\delta, L-1)$-coverable, i.e., $C_\delta\left(\mathcal{I}\left(\overline{q}(x_0, y_0)\right)\right) \leq L - 1$, which will lead to a contradiction.

Recall from Lemma 10 that $\mathcal{I}\left(\overline{q}(x_0, y_0)\right)$ is a subset of the union of all elementary intervals in $\overline{q}(x_0, y_0)$, $\overline{\mathcal{I}}(\epsilon, \overline{q}(x_0, y_0))$. In what follows, we will construct a $(\delta, L-1)$ cover for $\overline{\mathcal{I}}(\epsilon, \overline{q}(x_0, y_0))$. Let $q' = (q'_1, q'_2, \ldots, q'_K)$ be the set of all elementary queries (Definition 9) in $\overline{q}(x_0, y_0) = (q_0, q_1, \ldots, q_N, q_{N+1})$ that lie in the interval $[\delta, 1]$ (Note that this interval is closed at 1.) Without loss of generality, assume that $q'_1 < q'_2 < \cdots < q'_K$. Let q^* be the maximum value of the queries in $[0, \delta)$:

$$q^* = \max(\overline{q}(x_0, y_0) \cap [0, \delta)). \tag{39}$$

We now construct an explicit collection of intervals of length at most δ, $V_0, V_1, \ldots$, that form a cover for the union of all elementary intervals. Let $V_0 = [0, \delta]$. The remaining intervals $V_1, V_2, \ldots$ are constructed recursively, as follows:

1. V_1: if $q'_1 - q^* \leq \epsilon$, then let $v_1 = 0$ and $V_1 = [\delta, 2\delta]$;

   if $q'_1 - q^* > \epsilon$, then let $v_1 = \min\{l \in \{1, 2, \ldots, K-1\} : q'_{l+1} - q'_l \leq \epsilon\}$ and $V_1 = [q'_{v_1}, q'_{v_1} + \delta]$.

2. For $k = 2, 3, \ldots$, let $v_k = \min\{l \in \{1, 2, \ldots, K-1\} : l > v_{k-1} + 1 \text{ and } q'_{l+1} - q'_l \leq \epsilon\}$ and $V_k = [q'_{v_k}, q'_{v_k} + \delta]$.

The above procedure continues until no more v_k can be found. Since there is a finite number of queries, the procedure terminates after a finite number of steps, and we will denote by T the number of intervals generated. Figure 5 provides an example of this procedure.

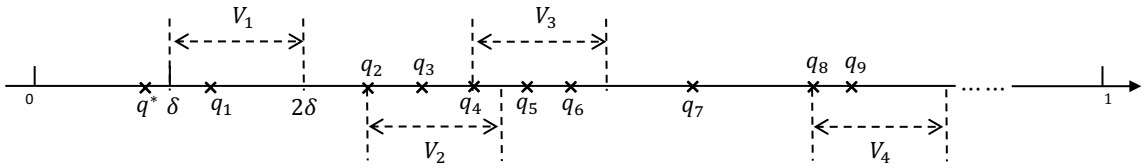

Figure 5: An example of the $(\delta, L-1)$ cover when $q_1 - q^* \leq \epsilon$. Note that q_7 in this example is far away from the rest of the queries, and hence not an elementary query.

We claim that the union of the intervals $\{V_0, V_1, V_2, \ldots, V_T\}$ contains all elementary intervals. Recall that $\delta > 2\epsilon$. It is clear from the first step that V_0 and V_1 together contain any elementary interval that intersects with $[0, \delta]$. Fix an elementary interval $[q'_a, q'_b)$ in $(\delta, 1)$, i.e., $q'_a > \delta$. By construction, there are two possibilities:

1. Suppose $a = v_k$ for some k. Then q'_a is equal to the left endpoint of V_k. Since $\delta > \epsilon$, we obtain $[q'_a, q'_b) \subset V_k$.

2. Suppose $a \neq v_k$ for all k. We know that $a \geq 2$, and we further consider two sub-cases:

   (a) Suppose that $a - 1 = v_{k'}$ for some k'. Then, we must have $q'_a \leq q'_{a-1} + \epsilon$ and $q'_b \leq q'_{a-1} + 2\epsilon$. Because $\delta > 2\epsilon$, this further implies that $[q'_a, q'_b) \subset V_{k'}$.

   (b) Suppose that $a - 1 \neq v_k$ for all k. By construction of the V_k's, this implies that $q'_{a+1} - q'_a > \epsilon$, which contradicts the assumption that $[q'_a, q'_b)$ is an elementary interval. Therefore, this possibility cannot materialize.

The above arguments demonstrate that $[q'_a, q'_b) \subset V_k$ for some $k \in \{1, 2, \ldots, T\}$. Recall also that $V_0 = [0, \delta]$. We have thus demonstrated that

$$\overline{\mathcal{I}}(\epsilon, \overline{q}(x_0, y_0)) \subset \bigcup_{0 \leq k \leq T} V_k. \tag{40}$$

To complete the proof, it remains to show that the collection $\{V_0, V_1, V_2, \ldots, V_T\}$ consists of at most $L-1$ intervals, i.e., that T is at most $L-2$. We will do so by counting the number of elementary queries contained in each V_k, as follows:

1. V_1 always contains at least one query in $[\delta, 1]$: if $q'_1 - q^* \leq \epsilon$ then $q'_1 \in V_1$; otherwise, $q'_{v_1} \in V_1$ and $q'_{v_1+1} \in V_1$.

2. For every $k \geq 2$, V_k contains at least two queries: q'_{v_k} and q'_{v_k+1}. Furthermore, since $v_k > v_{k-1}+1$, we know that q'_{v_k} and q'_{v_k+1} are not contained in any other $V_{k'}$, for $k' > k$. Therefore, there exist $f(k) = \{\tilde{q}_a, \tilde{q}_b\} \subset q'$, such that

$$f(k) \subset V_k, \, \forall k \geq 2, \quad \text{and } f(k) \cap f(k') = \emptyset, \, \forall k \neq k'. \tag{41}$$

   In other words, the different V_k can be associated with disjoint 2-element sets of queries.

We conclude that the set $\bigcup_{1 \leq k \leq T} V_k$ contains at least $2(T-1)+1 = 2T-1$ distinct queries in the interval $[\delta, 1]$ (possibly including the trivial query $q_{N+1} = 1$). Note that by assumption, there are at most $2L-4$ queries in the interval $[\delta, 1]$. It thus follows that $T \leq L - 1.5$. Since T and L are integers, this further implies that $T \leq L - 2$. Together with Eq. (40) and the fact that the V_k's have length δ, we conclude that $\overline{\mathcal{I}}(\epsilon, \overline{q}(x_0, y_0))$ is $(\delta, L-1)$-coverable. This completes the proof of Lemma 12. ■