

---

## MAKE UP YOUR MIND: THE PRICE OF ONLINE QUERIES IN DIFFERENTIAL PRIVACY

MARK BUN, THOMAS STEINKE, AND JONATHAN ULLMAN

Department of Computer Science, Princeton University  
*e-mail address:* markmbun@gmail.com

IBM Research - Almaden  
*e-mail address:* jpc@thomas-steinke.net

College of Computer and Information Science, Northeastern University  
*e-mail address:* jullman@ccs.neu.edu

---

**ABSTRACT.** We consider the problem of answering queries about a sensitive dataset subject to differential privacy. The queries may be chosen adversarially from a larger set  $Q$  of allowable queries in one of three ways, which we list in order from easiest to hardest to answer:

- *Offline:* The queries are chosen all at once and the differentially private mechanism answers the queries in a single batch.
- *Online:* The queries are chosen all at once, but the mechanism only receives the queries in a streaming fashion and must answer each query before seeing the next query.
- *Adaptive:* The queries are chosen one at a time and the mechanism must answer each query before the next query is chosen. In particular, each query may depend on the answers given to previous queries.

Many differentially private mechanisms are just as efficient in the adaptive model as they are in the offline model. Meanwhile, most lower bounds for differential privacy hold in the offline setting. This suggests that the three models may be equivalent.

We prove that these models are all, in fact, distinct. Specifically, we show that there is a family of statistical queries such that exponentially more queries from this family can be answered in the offline model than in the online model. We also exhibit a family of search queries such that exponentially more queries from this family can be answered in the online model than in the adaptive model. We also investigate whether such separations might hold for simple queries like threshold queries over the real line.

---

*Key words and phrases:* differential privacy, online query, offline query, pricing.

Earlier versions of this work appeared at the 3rd Workshop on Theory and Practice of Differential Privacy (TPDP) 2016, and the Symposium on Discrete Algorithms (SODA) 2017.

This work was done while Mark Bun and Thomas Steinke were at Harvard University John A. Paulson School of Engineering and Applied Sciences. Supported by NSF grants CCF-1116616, CCF-1420938, and CNS-1237235.

## 1. INTRODUCTION

*Differential privacy* [Dwork et al., 2006] is a formal guarantee that an algorithm run on a sensitive dataset does not reveal too much about any individual in that dataset. Since its introduction, a rich literature has developed to determine what statistics can be computed accurately subject to differential privacy. For example, suppose we wish to approximate a real-valued *query*  $q(x)$  on some dataset  $x$  that consists of the private data of many individuals. Then, this question has a clean answer—we can compute a differentially private estimate of  $q(x)$  with error proportional to the *global sensitivity* of  $q$ , and we cannot have smaller error in the worst case.

But how much error do we need to answer a large set of queries  $q_1, \dots, q_k$ ? Before we can answer this question, we have to define a model of how the queries are asked and answered. The literature on differential privacy has considered three different interactive models<sup>1</sup> for specifying the queries:

- *The Offline Model:* The sequence of queries  $q_1, \dots, q_k$  are given to the algorithm together in a batch and the mechanism answers them together.
- *The Online Model:* The sequence of queries  $q_1, \dots, q_k$  is chosen in advance and then the mechanism must answer each query  $q_j$  before seeing  $q_{j+1}$ .
- *The Adaptive Model:* The queries are not fixed in advance, each query  $q_{j+1}$  may depend on the answers to queries  $q_1, \dots, q_j$ .

In all three cases, we assume that  $q_1, \dots, q_k$  are chosen from some family of allowable queries  $Q$ , but may be chosen adversarially from this family.

Differential privacy seems well-suited to the adaptive model. Arguably its signature property is that any adaptively-chosen sequence of differentially private algorithms remains collectively differentially private, with a graceful degradation of the privacy parameters [Dwork et al., 2006, 2010b]. As a consequence, there is a simple differentially private algorithm that takes a dataset of  $n$  individuals and answers  $\tilde{\Omega}(n)$  statistical queries in the adaptive model with error  $o(1/\sqrt{n})$ , simply by perturbing each answer independently with carefully calibrated noise. In contrast, the seminal lower bound of Dinur and Nissim and its later refinements [Dinur and Nissim, 2003, Dwork and Yekhanin, 2008] shows that there exists a fixed set of  $O(n)$  queries that cannot be answered by any differentially private algorithm with such little error, even in the easiest offline model. For an even more surprising example, the private multiplicative weights algorithm of Hardt and Rothblum [Hardt and Rothblum, 2010] can in many cases answer an exponential number of arbitrary, adaptively-chosen statistical queries with a strong accuracy guarantee, whereas Bun et al. [2014] show that the accuracy guarantee of private multiplicative weights is nearly optimal even for a simple, fixed family of queries.

These examples might give the impression that answering adaptively-chosen queries comes “for free” in differential privacy—that everything that can be achieved in the offline model can be matched in the adaptive model. Beyond just the lack of any separation between the models, many of the most powerful differentially private algorithms in all of these models use techniques from no-regret learning, which are explicitly designed for adaptive models.

<sup>1</sup>Usually, the “interactive model” refers only to what we call the “adaptive model.” We prefer to call all of these models interactive, since they each require an interaction with a data analyst who issues the queries. We use the term “interactive” to distinguish these models from one where the algorithm only answers a fixed set of queries.

Another motivation for studying the relationship between these models is the recent line of work connecting differential privacy to statistical validity for *adaptive data analysis* [Hardt and Ullman, 2014, Dwork et al., 2015a, Steinke and Ullman, 2015a, Bassily et al., 2016], which shows that differentially private algorithms for adaptively-chosen queries in fact yield state-of-the-art algorithms for statistical problems unrelated to privacy. This connection further motivates studying the adaptive model and its relationship to the other models in differential privacy.

In this work, we show for the first time that these three models are actually distinct. In fact, we show exponential separations between each of the three models. These are the first separations between these models in differential privacy.

**1.1. Our Results.** Given a dataset  $x$  whose elements come from a data universe  $X$ , a *statistical query on  $X$*  is defined by a predicate  $\phi$  on  $X$  and asks “what fraction of elements in the dataset satisfy  $\phi$ ?” The answer to a statistical query lies in  $[0, 1]$  and our goal is to answer these queries up to some small additive error  $\pm\alpha$ , for a suitable choice of  $0 < \alpha < 1$ . If the mechanism is required to answer *arbitrary* statistical queries, then the offline, online, and adaptive models are essentially equivalent — the upper bounds in the adaptive model match the lower bounds in the offline model [Dwork et al., 2010b, Hardt and Rothblum, 2010, Bun et al., 2014, Steinke and Ullman, 2015b]. However, we show that when the predicate  $\phi$  is required to take a specific form, then it becomes strictly easier to answer a set of these queries in the offline model than it is to answer a sequence of queries presented online.

**Theorem 1.1 Informal.** *There exists a data universe  $X$  and a family of statistical queries  $Q$  on  $X$  such that for every  $n \in \mathbb{N}$ ,*

- (1) *there is a differentially private algorithm that takes a dataset  $x \in X^n$  and answers any set of  $k = 2^{\Omega(\sqrt{n})}$  offline queries from  $Q$  up to error  $\pm 1/100$  from  $Q$ , but*
- (2) *no differentially private algorithm can take a dataset  $x \in X^n$  and answer an arbitrary sequence of  $k = O(n^2)$  online (but not adaptively-chosen) queries from  $Q$  up to error  $\pm 1/100$ .*

This result establishes that the online model is strictly harder than the offline model. We also demonstrate that the adaptive model is strictly harder than the online model. Here, the family of queries we use in our separation is not a family of statistical queries, but is rather a family of *search queries* with a specific definition of accuracy that we will define later.

**Theorem 1.2 Informal.** *For every  $n \in \mathbb{N}$ , there is a family of “search queries”  $Q$  on datasets in  $X^n$  such that*

- (1) *there is a differentially private algorithm that takes a dataset  $x \in \{\pm 1\}^n$  and accurately answers any online (but not adaptively-chosen) sequence of  $k = 2^{\Omega(n)}$  queries from  $Q$ , but*
- (2) *no differentially private algorithm can take a dataset  $x \in \{\pm 1\}^n$  and accurately answer an adaptively-chosen sequence of  $k = O(1)$  queries from  $Q$ .*

We leave it as an interesting open question to separate the online and adaptive models for statistical queries, or to show that the models are equivalent for statistical queries.

Although Theorems 1.1 and 1.2 separate the three models, these results use somewhat contrived families of queries. Thus, we also investigate whether the models are distinct for

*natural* families of queries that are of use in practical applications. One very well studied class of queries is *threshold queries*. These are a family of statistical queries  $Q_{\text{thresh}}$  defined on the universe  $[0, 1]$  and each query is specified by a point  $\tau \in [0, 1]$  and asks “what fraction of the elements of the dataset are at most  $\tau$ ?” If we restrict our attention to so-called pure differential privacy (i.e.  $(\epsilon, \delta)$ -differential privacy with  $\delta = 0$ ), then we obtain an exponential separation between the offline and online models for answering threshold queries.

**Theorem 1.3 Informal.** *For every  $n \in \mathbb{N}$ ,*

- (1) *there is a pure differentially private algorithm that takes a dataset  $x \in [0, 1]^n$  and answers any set of  $k = 2^{\Omega(n)}$  offline queries from  $Q_{\text{thresh}}$  up to error  $\pm 1/100$ , but*
- (2) *no pure differentially private algorithm takes a dataset  $x \in [0, 1]^n$  and answers an arbitrary sequence of  $k = O(n)$  online (but not adaptively-chosen) queries from  $Q_{\text{thresh}}$  up to error  $\pm 1/100$ .*

We also ask whether or not such a separation exists for arbitrary differentially private algorithms (i.e.  $(\epsilon, \delta)$ -differential privacy with  $\delta > 0$ ). Theorem 1.3 shows that, for pure differential privacy, threshold queries have near-maximal sample complexity. That is, up to constants, the lower bound for online threshold queries matches what is achieved by the Laplace mechanism, which is applicable to arbitrary statistical queries. This may lead one to conjecture that adaptive threshold queries also require near-maximal sample complexity subject to approximate differential privacy. However, we show that this is not the case:

**Theorem 1.4 .** *For every  $n \in \mathbb{N}$ , there is a differentially private algorithm that takes a dataset  $x \in [0, 1]^n$  and answers any set of  $k = 2^{\Omega(n)}$  adaptively-chosen queries from  $Q_{\text{thresh}}$  up to error  $\pm 1/100$ .*

In contrast, for any offline set of  $k$  thresholds  $\tau_1, \dots, \tau_k$ , we can round each element of the dataset up to an element in the finite universe  $X = \{\tau_1, \dots, \tau_k, 1\}$  without changing the answers to any of the queries. Then we can use known algorithms for answering all threshold queries over any finite, totally ordered domain [Beimel et al., 2013, Bun et al., 2015] to answer the queries using a very small dataset of size  $n = 2^{O(\log^* k)}$ . We leave it as an interesting open question to settle the complexity of answering threshold queries in the adaptive model.

## 1.2. Techniques.

*Separating Offline and Online Queries.* To prove Theorem 1.1, we construct a sequence of queries  $q_1, \dots, q_k$  such that, for all  $j \in [k]$ ,

- $q_j$  “reveals” the answers to  $q_1, \dots, q_{j-1}$ , but
- $q_1, \dots, q_{j-1}$  do not reveal the answer to  $q_j$ .

Thus, given the sequence  $q_1, \dots, q_k$  in the offline setting, the answers to  $q_1, \dots, q_{k-1}$  are revealed by  $q_k$ . So only  $q_k$  needs to be answered and the remaining query answers can be inferred. However, in the online setting, each query  $q_{j-1}$  must be answered before  $q_j$  is presented and this approach does not work. This is the intuition for our separation.

To prove the online lower bound, we build on a lower bound for marginal queries [Bun et al., 2014], which is based on the existence of short secure fingerprinting codes [Boneh and Shaw, 1998, Tardos, 2008]. Consider the data universe  $\{\pm 1\}^k$ . Given a dataset  $x \in \{\pm 1\}^{n \times k}$ , a *marginal query* is a specific type of statistical query that asks for the mean of a given

column of  $x$ . Bun et al. [2014] showed that unless  $k \ll n^2$ , there is no differentially private algorithm that answers all  $k$  marginal queries with non-trivial accuracy. This was done by showing that such an algorithm would violate the security of a short fingerprinting code due to Tardos [2008]. We are able to “embed”  $k$  marginal queries into the sequence of online queries  $q_1, \dots, q_k$ . Thus a modification of the lower bound for marginal queries applies in the online setting.

To prove the offline upper bound, we use the fact that every query reveals information about other queries. However, we must handle arbitrary sequences of queries, not just the specially-constructed sequences used for the lower bound. The key property of our family of queries is the following. Each element  $x$  of the data universe  $X$  requires  $k$  bits to specify. On the other hand, for any set of queries  $q_1, \dots, q_k$ , we can specify  $q_1(x), \dots, q_k(x)$  using only  $O(\log(nk))$  bits. Thus the effective size of the data universe given the queries is  $\text{poly}(nk)$ , rather than  $2^k$ . Then we can apply a differentially private algorithm that gives good accuracy as long as the data universe has subexponential size [Blum et al., 2013]. Reducing the size of the data universe is only possible once the queries have been specified; hence this approach only works in the offline setting.

*Separating Online and Adaptive Queries.* To prove Theorem 1.2, we start with the classical randomized response algorithm [Warner, 1965]. Specifically, given a dataset  $x \in \{\pm 1\}^n$ , randomized response produces a new dataset  $y \in \{\pm 1\}^n$  where each coordinate  $y_i$  is independently set to  $+x_i$  with probability  $(1 + \alpha)/2$  and is set to  $-x_i$  with probability  $(1 - \alpha)/2$ . It is easy to prove that this algorithm is  $(O(\alpha), 0)$ -differentially private. What accuracy guarantee does this algorithm satisfy? By design, it outputs a vector  $y$  that has correlation approximately  $\alpha$  with the dataset  $x$ , i.e.  $\langle y, x \rangle \approx \alpha n$ . On the other hand, it is also easy to prove that there is no differentially private algorithm (for any reasonable privacy parameters) that can output a vector that has correlation at least  $1/2$  with the sensitive dataset.

Our separation between the online and adaptive models is based on the observation that, if we can obtain  $O(1/\alpha^2)$  “independent” vectors  $y_1, \dots, y_k$  that are each roughly  $\alpha$ -correlated with  $x$ , then we can obtain a vector  $z$  that is  $(1/2)$ -correlated with  $x$ , simply by letting  $z$  be the coordinate-wise majority of the vectors  $y_1, \dots, y_k$ . Thus, no differentially private algorithm can output such a set of vectors. More precisely, we require that  $\langle y_i, y_j \rangle \approx \alpha^2 n$  for  $i \neq j$ , which is achieved if each  $y_j$  is an independent sample from randomized response.

Based on this observation, we devise a class of queries such that, if we are allowed to choose  $k$  of these queries adaptively, then we obtain a set of vectors  $y_1, \dots, y_k$  satisfying the conditions above. This rules out differential privacy for  $k = O(1/\alpha^2)$  adaptive queries. The key is that we can use adaptivity to ensure that each query asks for an “independent”  $y_j$  by adding the previous answers  $y_1, \dots, y_{j-1}$  as constraints in the search query.

On the other hand, randomized response can answer each such query with high probability. If a subexponential number of these queries are fixed in advance, then, by a union bound, the vector  $y$  output by randomized response is simultaneously an accurate answer to every query with high probability. Since randomized response is oblivious to the queries, we can also answer the queries in the online model, as long as they are not chosen adaptively.

At a high level, the queries that achieve this property are of the form “output a vector  $y \in \{\pm 1\}^n$  that is approximately  $\alpha$ -correlated with  $x$  and is approximately as uncorrelated as possible with the vectors  $v_1, \dots, v_m$ .” A standard concentration argument shows that randomized response gives an accurate answer to all the queries simultaneously with high

probability. On the other hand, if we are allowed to choose the queries adaptively, then for each query  $q_i$ , we can ask for a vector  $y_i$  that is correlated with  $x$  but is as uncorrelated as possible with the previous answers  $y_1, \dots, y_{i-1}$ .

*Threshold Queries.* For pure differential privacy, our separation between offline and online threshold queries uses a simple argument based on binary search. Our starting point is a lower bound showing that any purely differentially private algorithm that takes a dataset of  $n$  points  $x_1, \dots, x_n \in \{1, \dots, T\}$  and outputs an *approximate median* of these points requires  $n = \Omega(\log T)$ . This lower bound follows from a standard application of the “packing” technique of [Hardt and Talwar \[2010\]](#). On the other hand, by using binary search, any algorithm that can answer  $k = O(\log T)$  adaptively-chosen threshold queries can be used to find an approximate median. Thus, any purely differentially private algorithm for answering such queries requires a dataset of size  $n = \Omega(k)$ . We can also show that the same lower bound holds for online non-adaptive queries as well. In contrast, using the algorithms of [Dwork et al. \[2010a\]](#), [Chan et al. \[2011\]](#), [Dwork et al. \[2015b\]](#), we can answer  $k$  offline threshold queries on a dataset with only  $n = O(\log k)$  elements, giving an exponential separation.

The aforementioned separation depends on strong lower bounds for pure differential privacy that follow from its *group privacy* guarantees. Approximate differential privacy, on the other hand, does not give such strong group privacy, and is hence not susceptible to these lower bounds. In fact, circumventing these lower bounds allows us to obtain an improved algorithm for answering adaptively chosen threshold queries with approximate differential privacy.

The technical basis for our improved algorithm is a generalization of the *sparse vector* technique [[Dwork et al., 2010a](#), [Roth and Roughgarden, 2010](#), [Hardt and Rothblum, 2010](#)] (see [Dwork and Roth \[2014, §3.6\]](#) for a textbook treatment). Our algorithm makes crucial use of a *stability argument* similar to the propose-test-release techniques of [Dwork and Lei \[2009\]](#). To our knowledge, this is the first use of a stability argument for any online or adaptive problem in differential privacy and may be of independent interest. In particular, our algorithm is given an input  $x \in X^n$ , a threshold  $t \in (0, 1)$ , and an adaptive sequence of statistical (or low-sensitivity) queries  $q_1, \dots, q_k : X^n \rightarrow [0, 1]$  and, for each query  $q_j$ , it reports (i)  $q_j(x) \geq t$ , (ii)  $q_j(x) \leq t$ , or (iii)  $t - \alpha \leq q_j(x) \leq t + \alpha$ . The sample complexity of this algorithm is  $n = O(\sqrt{c} \log(k/\varepsilon\delta)/\varepsilon\alpha)$ , where  $k$  is the total number of queries,  $c$  is an upper bound on the number of times (iii) may be reported, and  $(\varepsilon, \delta)$ -differential privacy is provided. We call this the *Between Thresholds algorithm*.

Once we have this algorithm, we can use it to answer adaptively-chosen thresholds using an approach inspired by [Bun et al. \[2015\]](#). The high-level ideal is to sort the dataset  $x_{(1)} < x_{(2)} < \dots < x_{(n)}$  and then partition it into chunks of consecutive sorted elements. For any chunk, and a threshold  $\tau$ , we can use the between thresholds algorithm to determine (approximately) whether  $\tau$  lies below all elements in the chunk, above all elements in the chunk, or inside the chunk. Obtaining this information for every chunk is enough to accurately estimate the answer to the threshold query  $\tau$  up to an error proportional to the size of the chunks. The sample complexity is dominated by the  $O(\log k)$  sample complexity of our *Between Thresholds* algorithm multiplied by the number of chunks needed, namely  $O(1/\alpha)$ .

2. PRELIMINARIES

**2.1. Datasets and Differential Privacy.** A *dataset*  $x \in (x_1, \dots, x_n) \in X^n$  is an ordered tuple of  $n$  elements from some *data universe*  $X$ . We say that two datasets  $x, x'$  are *adjacent* if they differ on only a single element and denote this relation by  $x \sim x'$ .

**Definition 2.1** (Differential Privacy [Dwork et al., 2006]). A randomized algorithm  $M : X^n \rightarrow \mathcal{R}$  is  $(\epsilon, \delta)$ -differentially private if for every two adjacent datasets  $x \sim x'$ , and every  $R \subseteq \mathcal{R}$ ,

$$\mathbb{P}[M(x) \in R] \leq e^\epsilon \mathbb{P}[M(x') \in R] + \delta.$$

We also use the following well known *group privacy* property of  $(\epsilon, 0)$ -differential privacy. We say that two datasets  $x, x'$  are  $c$ -adjacent if they differ on at most  $c$ -elements, and denote this relation by  $x \sim_c x'$ .

**Lemma 2.2** [Dwork et al., 2006]. If  $M : X^n \rightarrow \mathcal{R}$  is  $(\epsilon, 0)$ -differentially private, then for every  $c \in \mathbb{N}$  and every two  $c$ -adjacent datasets  $x \sim_c x'$ , and every  $R \subseteq \mathcal{R}$ ,

$$\mathbb{P}[M(x) \in R] \leq e^{c\epsilon} \mathbb{P}[M(x') \in R].$$

**2.2. Queries.** In this work we consider two general classes of queries on the dataset: *statistical queries*, and *search queries*. Although statistical queries are a very special case of search queries, we will present each of them independently to avoid having to use overly abstract notation to describe statistical queries.

**Statistical Queries.** A statistical query on a data universe  $X$  is defined by a Boolean predicate  $q : X \rightarrow \{0, 1\}$ . Abusing notation, we define the evaluation of a statistical query  $q$  on a dataset  $x = (x_1, \dots, x_n)$  to be the average of the predicate over the rows

$$q(x) = \frac{1}{n} \sum_{i=1}^n q(x_i) \in [0, 1].$$

For a dataset  $x$ , a statistical query  $q$ , and an answer  $a \in [0, 1]$ , the answer is  $\alpha$ -accurate for  $q$  on  $x$  if

$$|q(x) - a| \leq \alpha.$$

**Search Queries.** A *search query*  $q$  on  $X^n$  is defined by a *loss function*  $L_q : X^n \times \mathcal{R} \rightarrow [0, \infty)$ , where  $\mathcal{R}$  is an arbitrary set representing the range of possible outputs. For a dataset  $x \in X^n$  and an output  $y \in \mathcal{R}$ , we will say that  $y$  is  $\alpha$ -accurate for  $q$  on  $x$  if  $L_q(x, y) \leq \alpha$ . In some cases the value of  $L_q$  will always be either 0 or 1. Thus we simply say that  $y$  is *accurate for  $q$  on  $x$*  if  $L_q(x, y) = 0$ . For example, if  $X^n = \{\pm 1\}^n$ , we can define a search query by  $\mathcal{R} = \{\pm 1\}^n$ , and  $L_q(x, y) = 0$  if  $\langle x, y \rangle \geq \alpha n$  and  $L_q(x, y) = 1$  otherwise. In this case, the search query would ask for any vector  $y$  that has correlation  $\alpha$  with the dataset.

To see that statistical queries are a special case of search queries, given a statistical query  $q$  on  $X^n$ , we can define a search query  $L_q$  with  $\mathcal{R} = [0, 1]$  and  $L_q(x, a) = |q(x) - a|$ . Then both definitions of  $\alpha$ -accurate align.

**2.3. Models of Interactive Queries.** The goal of this work is to understand the implications of different ways to allow an adversary to query a sensitive dataset. In each of these models there is an algorithm  $M$  that holds a dataset  $x \in X^n$ , and a fixed family of (statistical or search) queries  $Q$  on  $X^n$ , and a bound  $k$  on the number of queries that  $M$  has to answer. There is also an adversary  $A$  that chooses the queries. The models differ in how the queries chosen by  $A$  are given to  $M$ .

*Offline.* In the *offline* model, the queries  $q_1, \dots, q_k \in Q$  are specified by the adversary  $A$  in advance and the algorithm  $M$  is given all the queries at once and must provide answers. Formally, we define the following function  $\text{Offline}_{A \Leftarrow M} : X^n \rightarrow Q^k \times \mathcal{R}^k$  depending  $A$  and  $M$ .

Input:  $x \in X^n$ .  
 $A$  chooses  $q_1, \dots, q_k \in Q$ .  
 $M$  is given  $x$  and  $q_1, \dots, q_k$  and outputs  $a_1, \dots, a_k \in \mathcal{R}$ .  
Output:  $(q_1, \dots, q_k, a_1, \dots, a_k) \in Q^k \times \mathcal{R}^k$ .

Figure 1:  $\text{Offline}_{A \Leftarrow M} : X^n \rightarrow Q^k \times \mathcal{R}^k$

*Online Non-Adaptive.* In the *online non-adaptive* model, the queries  $q_1, \dots, q_k \in Q$  are again fixed in advance by the adversary, but are then given to the algorithm one at a time, and the algorithm must give an answer to query  $q_j$  before it is shown  $q_{j+1}$ . We define a function  $\text{Online}_{A \Leftarrow M} : X^n \rightarrow Q^k \times \mathcal{R}^k$  depending on the adversary  $A$  and the algorithm  $M$  as follows.


Input:  $x \in X^n$ .  
 $A$  chooses  $q_1, \dots, q_k \in Q$ .  
 $M$  is given  $x$ .  
For  $j = 1, \dots, k$ :  
 $M$  is given  $q_j$  and outputs  $a_j \in \mathcal{R}$ .<sup>2</sup>  
Output:  $(q_1, \dots, q_k, a_1, \dots, a_k) \in Q^k \times \mathcal{R}^k$ .

Figure 2:  $\text{Online}_{A \Leftarrow M} : X^n \rightarrow Q^k \times \mathcal{R}^k$

*Online Adaptive.* In the *online adaptive* model, the queries  $q_1, \dots, q_k \in Q$  are not fixed, and the adversary may choose each  $q_j$  based on the answers that the algorithm gave to the previous queries. We define a function  $\text{Adaptive}_{A \Leftarrow M} : X^n \rightarrow Q^k \times \mathcal{R}^k$  depending on the adversary  $A$  and the algorithm  $M$  as follows.

**Definition 2.3** (Differential Privacy for Interactive Mechanisms). In each of the three cases — Offline, Online Non-Adaptive, or Online Adaptive — we say that  $M$  is  $(\epsilon, \delta)$ -differentially private if, for all adversaries  $A$ , respectively  $\text{Offline}_{A \Leftarrow M}$ ,  $\text{Online}_{A \Leftarrow M}$ , or  $\text{Adaptive}_{A \Leftarrow M}$  is  $(\epsilon, \delta)$ -differentially private.




 Figure 3:  $\text{Adaptive}_{A \rightleftharpoons M} : X^n \rightarrow Q^k \times \mathcal{R}^k$ 

**Definition 2.4** (Accuracy for Interactive Mechanisms). In each case — Offline, Online Non-Adaptive, or Online Adaptive queries — we say that  $M$  is  $(\alpha, \beta)$ -accurate if, for all adversaries  $A$  and all inputs  $x \in X^n$ ,

$$\mathbb{P}_{q_1, \dots, q_k, a_1, \dots, a_k} \left[ \max_{j \in [k]} L_{q_j}(x, a_j) \leq \alpha \right] \geq 1 - \beta, \quad (2.1)$$

where  $(q_1, \dots, q_k, a_1, \dots, a_k)$  is respectively drawn from one of  $\text{Offline}_{A \rightleftharpoons M}(x)$ ,  $\text{Online}_{A \rightleftharpoons M}(x)$ , or  $\text{Adaptive}_{A \rightleftharpoons M}(x)$ . We also say that  $M$  is  $\alpha$ -accurate if the above holds with (2.1) replaced by

$$\mathbb{E}_{q_1, \dots, q_k, a_1, \dots, a_k} \left[ \max_{j \in [k]} L_{q_j}(x, a_j) \right] \leq \alpha.$$

### 3. A SEPARATION BETWEEN OFFLINE AND ONLINE QUERIES

In this section we prove that online accuracy is strictly harder to achieve than offline accuracy, even for statistical queries. We prove our results by constructing a set of statistical queries that we call *prefix queries* for which it is possible to take a dataset of size  $n$  and accurately answer superpolynomially many offline prefix queries in a differentially private manner, but it is impossible to answer more than  $O(n^2)$  online prefix queries while satisfying differential privacy.

We now define the family of prefix queries. These queries are defined on the universe  $X = \{\pm 1\}^* = \bigcup_{j=0}^{\infty} \{\pm 1\}^j$  consisting of all finite length binary strings.<sup>3</sup> For  $x, y \in \{\pm 1\}^*$ , we use  $y \preceq x$  to denote that  $y$  is a *prefix* of  $x$ . Formally

$$y \preceq x \iff |y| \leq |x| \text{ and } \forall i = 1, \dots, |y| \ x_i = y_i.$$

**Definition 3.1.** For any finite set  $S \subseteq \{\pm 1\}^*$  of finite-length binary strings, we define the *prefix query*  $q_S : \{\pm 1\}^* \rightarrow \{\pm 1\}$  by

$$q_S(x) = 1 \iff \exists y \in S \ y \preceq x.$$

We also define

$$\begin{aligned} Q_{\text{prefix}} &= \{q_S \mid S \subseteq \{\pm 1\}^*\} \\ Q_{\text{prefix}}^B &= \{q_S \mid S \subseteq \{\pm 1\}^*, |S| \leq B\} \end{aligned}$$

to be the set of all prefix queries and the set of prefix queries with sizes bounded by  $B$ , respectively.

<sup>3</sup>All of the arguments in this section hold if we restrict to strings of length at most  $k + \log n$ . However, we allow strings of arbitrary length to reduce notational clutter.

**3.1. Answering Offline Prefix Queries.** We now prove that there is a differentially private algorithm that answers superpolynomially many prefix queries, provided that the queries are specified offline.

**Theorem 3.2** Answering Offline Prefix Queries. *For every  $\alpha, \varepsilon \in (0, 1/10)$ , every  $B \in \mathbb{N}$ , and every  $n \in \mathbb{N}$ , there exists a*

$$k = \min \left\{ 2^{\Omega(\sqrt{\alpha^3 \varepsilon n})}, 2^{\Omega(\alpha^3 \varepsilon n / \log B)} \right\}$$

and an  $(\varepsilon, 0)$ -differentially private algorithm  $M_{\text{prefix}} : X^n \times (Q_{\text{prefix}}^B)^k \rightarrow \mathbb{R}^k$  that is  $(\alpha, 1/100)$ -accurate for  $k$  offline queries from  $Q_{\text{prefix}}^B$ .

We remark that it is possible to answer even more offline prefix queries by relaxing to  $(\varepsilon, \delta)$ -differential privacy for some negligibly small  $\delta > 0$ . However, we chose to state the results for  $(\varepsilon, 0)$ -differential privacy to emphasize the contrast with the lower bound, which applies even when  $\delta > 0$ , and to simplify the statement.

Our algorithm for answering offline queries relies on the existence of a good differentially private algorithm for answering *arbitrary* offline statistical queries. For concreteness, the so-called ‘‘BLR mechanism’’ of Blum et al. [2013] suffices, although different parameter tradeoffs can be obtained using different mechanisms. Differentially private algorithms with this type of guarantee exist only when the data universe is bounded, which is not the case for prefix queries. However, as we show, when the queries are specified offline, we can replace the infinite universe  $X = \{\pm 1\}^*$  with a finite, restricted universe  $X'$  and run the BLR mechanism. Looking ahead, the key to our separation will be the fact that this universe restriction is only possible in the offline setting. Before we proceed with the proof of Theorem 3.2, we will state the guarantees of the BLR mechanism.

**Theorem 3.3** [Blum et al., 2013]. *For every  $0 < \alpha, \varepsilon \leq 1/10$  and every finite data universe  $X$ , if  $\mathcal{Q}_{\text{SQ}}$  is the set of all statistical queries on  $X$ , then for every  $n \in \mathbb{N}$ , there is a*

$$k = 2^{\Omega(\alpha^3 \varepsilon n / \log |X|)}$$

and an  $(\varepsilon, 0)$ -differentially private algorithm  $M_{\text{BLR}} : X^n \times \mathcal{Q}_{\text{SQ}}^k \rightarrow \mathbb{R}^k$  that is  $(\alpha, 1/100)$ -accurate for  $k$  offline queries from  $\mathcal{Q}_{\text{SQ}}$ .

We are now ready to prove Theorem 3.2.

*Proof of Theorem 3.2.* Suppose we are given a set of queries  $q_{S_1}, \dots, q_{S_k} \in Q_{\text{prefix}}^B$  and a dataset  $x \in X^n$  where  $X = \{\pm 1\}^*$ . Let  $S = \bigcup_{j=1}^k S_j$ . We define the universe  $X_S = S \cup \{\emptyset\}$  where  $\emptyset$  denotes the empty string of length 0. Note that this universe depends on the choice of queries, and that  $|X_S| \leq kB + 1$ . Since  $X_S \subset X$ , it will be well defined to restrict the domain of each query  $q_{S_j}$  to elements of  $X_S$ .

Next, given a dataset  $x = (x_1, \dots, x_n) \in X^n$ , and a collection of sets  $S_1, \dots, S_k \subset X$ , we give a procedure for mapping each element of  $x$  to an element of  $X_S$  to obtain a new dataset  $x^S = (x_1^S, \dots, x_n^S) \in X_S^n$  that is equivalent to  $x$  with respect to the queries  $q_{S_1}, \dots, q_{S_k}$ . Specifically, define  $r_S : X \rightarrow X_S$  by

$$r_S(x) = \arg \max_{y \in X_S, y \preceq x} |y|.$$

That is,  $r_S(x)$  is the longest string in  $X_S$  that is a prefix of  $x$ . We summarize the key property of  $r_S$  in the following claim

**Claim 3.4 .** For every  $x \in X$ , and  $j = 1, \dots, k$ ,  $q_{S_j}(r_S(x)) = q_{S_j}(x)$ .

*Proof of Claim 3.4.* First, we state a simple but important fact about prefixes: If  $y, y'$  are both prefixes of a string  $x$  with  $|y| \leq |y'|$ , then  $y$  is a prefix of  $y'$ . Formally,

$$\forall x, y, y' \in \{0, 1\}^* \quad (y \preceq x \wedge y' \preceq x \wedge |y| \leq |y'|) \implies y \preceq y'. \quad (3.1)$$

Now, fix any  $x \in X$  and any query  $q_{S_j}$  and suppose that  $q_{S_j}(x) = 1$ . Then there exists a string  $y \in S_j$  such that  $y \preceq x$ . By construction, we have that  $r_S(x) \preceq x$  and that  $|r_S(x)| \geq |y|$ . Thus, by (3.1), we have that  $y \preceq r_S(x)$ . Thus, there exists  $y \in S_j$  such that  $y \preceq r_S(x)$ , which means  $q_{S_j}(r_S(x)) = 1$ , as required.

Next, suppose that  $q_{S_j}(r_S(x)) = 1$ . Then, there exists  $y \in S_j$  such that  $y \preceq r_S(x)$ . By construction,  $r_S(x) \preceq x$ , so by transitivity we have that  $y \preceq x$ . Therefore,  $q_{S_j}(x) = 1$ , as required.  $\square$

Given this lemma, we can replace every row  $x_i$  of  $x$  with  $x_i^S = r_S(x_i)$  to obtain a new dataset  $x^S$  such that for every  $j = 1, \dots, k$ ,

$$q_{S_j}(x^S) = \frac{1}{n} \sum_{i=1}^n q_{S_j}(x_i^S) = \frac{1}{n} \sum_{i=1}^n q_{S_j}(x_i) = q_{S_j}(x).$$

Thus, we can answer  $q_{S_1}, \dots, q_{S_k}$  on  $x^S \in X_S^n$ , rather than on  $x \in X^n$ . Note that each row of  $x^S$  depends only on the corresponding row of  $x$ . Hence, for every set of queries  $q_{S_1}, \dots, q_{S_k}$ , if  $x \sim x'$  are adjacent datasets, then  $x^S \sim x'^S$  are also adjacent datasets. Consequently, applying a  $(\epsilon, \delta)$ -differentially private algorithm to  $x^S$  yields a  $(\epsilon, \delta)$ -differentially private algorithm as a function of  $x$ .

In particular, we can give  $\alpha$ -accurate answers to these queries using the algorithm  $M_{\text{BLR}}$  as long as

$$k \leq 2^{\Omega(\alpha^3 \epsilon n / \log |X_S|)} = 2^{\Omega(\alpha^3 \epsilon n / \log(kB+1))}.$$

Rearranging terms gives the bound in Theorem 3.2. We specify the complete algorithm  $M_{\text{prefix}}$  in Figure 4.




Figure 4:  $M_{\text{prefix}}$

$\square$

**3.2. A Lower Bound for Online Prefix Queries.** Next, we prove a lower bound for online queries. Our lower bound shows that the simple approach of perturbing the answer to each query with independent noise is essentially optimal for prefix queries. Since this approach is only able to answer  $k = O(n^2)$  queries, we obtain an exponential separation between online and offline statistical queries for a broad range of parameters.

**Theorem 3.5** Lower Bound for Online Prefix Queries. *There exists a function  $k = O(n^2)$  such that for every sufficiently large  $n \in \mathbb{N}$ , there is no  $(1, 1/30n)$ -differentially private algorithm  $M$  that takes a dataset  $x \in X^n$  and is  $(1/100, 1/100)$ -accurate for  $k$  online queries from  $Q_{\text{prefix}}^n$ .*

In this parameter regime, our algorithm from Section 3.1 answers  $k = \exp(\tilde{\Omega}(\sqrt{n}))$  offline prefix queries, so we obtain an exponential separation.

Our lower bound relies on a connection between *fingerprinting codes* and differential privacy [Ullman, 2013, Bun et al., 2014, Steinke and Ullman, 2015b, Dwork et al., 2015c]. However, instead of using fingerprinting codes in a black-box way, we will make a direct use of the main techniques. Specifically, we will rely heavily on the following key lemma. The proof appears in Appendix A.

**Lemma 3.1** (Fingerprinting Lemma). *Let  $f : \{\pm 1\}^n \rightarrow [-1, 1]$  be any function. Suppose  $p$  is sampled from the uniform distribution over  $[-1, 1]$  and  $c \in \{\pm 1\}^n$  is a vector of  $n$  independent bits, where each bit has expectation  $p$ . Letting  $\bar{c}$  denote the coordinate-wise mean of  $c$ , we have*

$$\mathbb{E}_{p,c} \left[ f(c) \cdot \sum_{i \in [n]} (c_i - p) + 2|f(c) - \bar{c}| \right] \geq \frac{1}{3}.$$

Roughly the fingerprinting lemma says that if we sample a vector  $c \in \{\pm 1\}^n$  in a specific fashion, then for any bounded function  $f(c)$ , we either have that  $f(c)$  has “significant” correlation with  $c_i$  for some coordinate  $i$ , or that  $f(c)$  is “far” from  $\bar{c}$  on average. In our lower bound, the vector  $c$  will represent a *column* of the dataset, so each coordinate  $c_i$  will correspond to the value of some *row* of the dataset. The function  $f(c)$  will represent the answer to some prefix query. We will use the accuracy of a mechanism for answering prefix queries to argue that  $f(c)$  is not far from  $\bar{c}$ , and therefore conclude that  $f(c)$  must be significantly correlated with some coordinate  $c_i$ . On the other hand, if  $c_i$  were excluded from the dataset, then  $c_i$  is sufficiently random that the mechanism’s answers cannot be significantly correlated with  $c_i$ . We will use this to derive a contradiction to differential privacy.

*Proof of Theorem 3.5.* First we define the distribution on the input dataset  $x = (x_1, \dots, x_n)$  and the queries  $q_{S_1}, \dots, q_{S_k}$ .

Input dataset  $x$ :

- Sample  $p^1, \dots, p^k \in [-1, 1]$  independently and uniformly at random.
- Sample  $c^1, \dots, c^k \in \{\pm 1\}^n$  independently, where each  $c^j$  is a vector of  $n$  independent bits, each with expectation  $p^j$ .
- For  $i \in [n]$ , define

$$x_i = (\text{binary}(i), c_i^1, \dots, c_i^k) \in \{\pm 1\}^{\lceil \log_2 n \rceil + k},$$

where  $\text{binary}(i) \in \{\pm 1\}^{\lceil \log_2 n \rceil}$  is the binary representation of  $i$  where 1 is mapped to +1 and 0 is mapped to -1.<sup>4</sup> Let  $x = (x_1, \dots, x_n) \in (\{\pm 1\}^{\lceil \log_2 n \rceil + k})^n$ .

Queries  $q_{S_1}, \dots, q_{S_k}$ :

- For  $i \in [n]$  and  $j \in [k]$ , define

$$z_{i,j} = (\text{binary}(i), c_i^1, \dots, c_i^{j-1}, 1) \in \{\pm 1\}^{\lceil \log_2 n \rceil + j}.$$

- For  $j \in [k]$ , define  $q_{S_j} \in Q_{\text{prefix}}^n$  by  $S_j = \{z_{i,j} \mid i \in [n]\}$ .

These queries are designed so that the correct answer to each query  $j \in [k]$  is given by  $q_{S_j}(x) = \bar{c}^j$ :

**Claim 3.6 .** *For every  $j \in [k]$ , if the dataset  $x$  and the queries  $q_{S_1}, \dots, q_{S_k}$  are constructed as above, then with probability 1,*

$$q_{S_j}(x) = \frac{1}{n} \sum_{i=1}^n q_{S_j}(x_i) = \frac{1}{n} \sum_{i=1}^n c_i^j = \bar{c}^j$$

*Proof of Claim 3.6.* We have

$$q_{S_j}(x_i) = 1 \iff \exists w \in S_j (w \preceq x_i) \iff \exists \ell \in [n] (z_{\ell,j} \preceq x_i).$$

By construction, we have  $z_{\ell,j} \preceq x_i$  if and only if  $\ell = i$  and  $x_i^j = c_i^j = 1$ , as required. Here, we have used the fact that the strings  $\text{binary}(i)$  are unique to ensure that  $z_{\ell,j} \preceq x_i$  if and only if  $\ell = i$ .  $\square$

We now show no differentially private algorithm  $M$  is capable of giving accurate answers to these queries. Let  $M$  be an algorithm that answers  $k$  online queries from  $Q_{\text{prefix}}^n$ . Suppose we generate an input dataset  $x$  and queries  $q_{S_1}, \dots, q_{S_k}$  as above, and run  $M(x)$  on this sequence of queries. Let  $a^1, \dots, a^k \in [-1, 1]$  denote the answers given by  $M$ .

First, we claim that, if  $M(x)$  is accurate for the given queries, then each answer  $a^j$  is close to the corresponding value  $\bar{c}^j = \frac{1}{n} \sum_{i=1}^n c_i^j$ .

**Claim 3.2.** If  $M$  is  $(1/100, 1/100)$ -accurate for  $k$  online queries from  $Q_{\text{prefix}}^n$ , then with probability 1 over the choice of  $x$  and  $q_{S_1}, \dots, q_{S_k}$  above,

$$\mathbb{E}_M \left[ \sum_{j \in [k]} |a^j - \bar{c}^j| \right] \leq \frac{k}{10}.$$

*Proof of Claim 3.2.* By Claim 3.6, for every  $j \in [k]$ ,  $q_{S_j}(x) = \bar{c}^j$ . Since, by assumption,  $M$  is  $(1/100, 1/100)$ -accurate for  $k$  online queries from  $Q_{\text{prefix}}^n$ , we have that with probability at least  $99/100$ ,

$$\forall j \in [k] \quad |a^j - q_{S_j}(x)| \leq \frac{1}{100} \implies \forall j \in [k] \quad |a^j - \bar{c}^j| \leq \frac{1}{100}$$

By linearity of expectation, this case contributes at most  $k/100$  to the expectation. On the other hand,  $|a^j - q_{S_j}(x)| \leq 2$ , so by linearity of expectation the case where  $M$  is inaccurate contributes at most  $2k/100$  to the expectation. This suffices to prove the claim.  $\square$

<sup>4</sup>This choice is arbitrary, and is immaterial to our lower bound. The only property we need is that  $\text{binary}(i)$  uniquely identifies  $i$  and, for notational consistency, we require  $\text{binary}(i)$  to be a string over the alphabet  $\{\pm 1\}$ .

The next claim shows how the fingerprinting lemma (Lemma 3.1) can be applied to  $M$ .

**Claim 3.3.**

$$\mathbb{E}_{p,x,q,M} \left[ \sum_{j \in [k]} \left( a^j \sum_{i \in [n]} (c_i^j - p^j) + 2 |a^j - \bar{c}^j| \right) \right] \geq \frac{k}{3}.$$

*Proof.* By linearity of expectation, it suffices to show that, for every  $j \in [k]$ ,

$$\mathbb{E}_{p,x,q,M} \left[ a^j \sum_{i \in [n]} (c_i^j - p^j) + 2 |a^j - \bar{c}^j| \right] \geq \frac{1}{3}.$$

Since each column  $c^j$  is generated independently from the columns  $c^1, \dots, c^{j-1}$ ,  $c^j$  and  $p^j$  are independent from  $q_{S_1}, \dots, q_{S_j}$ . Thus, at the time  $M$  produces the output  $a^j$ , it does not have any information about  $c^j$  or  $p^j$  apart from its private input. (Although  $M$  later learns  $c^j$  when it is asked  $q_{S_{j+1}}$ .) For any fixed values of  $c^1, \dots, c^{j-1}$  and the internal randomness of  $M$ , the answer  $a^j$  is a deterministic function of  $c^j$ . Thus we can apply Lemma 3.1 to this function to establish the claim.  $\square$

Combining Claims 3.2 and 3.3 gives

$$\mathbb{E}_{p,x,q,M} \left[ \sum_{j \in [k]} a^j \sum_{i \in [n]} (c_i^j - p^j) \right] \geq \frac{2k}{15}.$$

In particular, there exists some  $i^* \in [n]$  such that

$$\mathbb{E}_{p,x,q,M} \left[ \sum_{j \in [k]} a^j (c_{i^*}^j - p^j) \right] \geq \frac{2k}{15n}. \quad (3.2)$$

To complete the proof, we show that (3.2) violates the differential privacy guarantee unless  $n \geq \Omega(\sqrt{k})$ .

To this end, fix any  $p^1, \dots, p^k \in [-1, 1]$ , whence  $c_{i^*}^1, \dots, c_{i^*}^k \in \{\pm 1\}$  are independent bits with  $\mathbb{E}[c^j] = p^j$ . Let  $\tilde{c}^1, \dots, \tilde{c}^k \in \{\pm 1\}$  be independent bits with  $\mathbb{E}[\tilde{c}^j] = p^j$ . The random variables  $c_{i^*}^1, \dots, c_{i^*}^k$  have the same marginal distribution as  $\tilde{c}^1, \dots, \tilde{c}^k$ . However,  $\tilde{c}^1, \dots, \tilde{c}^k$  are independent from  $a^1, \dots, a^k$ , whereas  $a^1, \dots, a^k$  depend on  $c_{i^*}^1, \dots, c_{i^*}^k$ . Consider the quantities

$$Z = \sum_{j \in [k]} a^j (c_{i^*}^j - p^j) \quad \text{and} \quad \tilde{Z} = \sum_{j \in [k]} a^j (\tilde{c}^j - p^j).$$

Differential privacy implies that  $Z$  and  $\tilde{Z}$  have similar distributions. Specifically, if  $M$  is  $(1, 1/30n)$ -differentially private, then

$$\mathbb{E}[|Z|] = \int_0^{2k} \mathbb{P}[|Z| > z] dz \leq \int_0^{2k} \left( e\mathbb{P}[|\tilde{Z}| > z] + \frac{1}{30n} \right) dz = e\mathbb{E}[|\tilde{Z}|] + \frac{k}{15n},$$

as  $|Z|, |\tilde{Z}| \leq 2k$  with probability 1.

Now  $\mathbb{E}[|Z|] \geq \mathbb{E}[Z] \geq 2k/15n$ , by (3.2). On the other hand,  $a^j$  is independent from  $\tilde{c}^j$  and  $\mathbb{E}[\tilde{c}^j - p^j] = 0$ , so  $\mathbb{E}[\tilde{Z}] = 0$ . We now observe that

$$\mathbb{E}[|\tilde{Z}|^2] \leq \mathbb{E}[\tilde{Z}^2] = \text{Var}[\tilde{Z}] = \sum_{j \in [k]} \text{Var}[a^j(\tilde{c}^j - p^j)] \leq \sum_{j \in [k]} \mathbb{E}[(\tilde{c}^j - p^j)^2] \leq k.$$

Thus, we have

$$\frac{2k}{15n} \leq \mathbb{E}[|Z|] \leq e\mathbb{E}[|\tilde{Z}|] + \frac{k}{15n} \leq e\sqrt{k} + \frac{k}{15n}.$$

The condition  $2k/15n \leq e\sqrt{k} + k/15n$  is a contradiction unless  $k \leq 225e^2n^2$ . Thus, we can conclude that there exists a  $k = O(n^2)$  such that no  $(1, 1/30n)$ -differentially private algorithm is accurate for more than  $k$  online queries from  $Q_{\text{prefix}}^n$ , as desired. This completes the proof.  $\square$

#### 4. A SEPARATION BETWEEN ADAPTIVE AND NON-ADAPTIVE ONLINE QUERIES

In this section we prove that even among online queries, answering adaptively-chosen queries can be strictly harder than answering non-adaptively-chosen queries. Our separation applies to a family of search queries that we call *correlated vector queries*. We show that for a certain regime of parameters, it is possible to take a dataset of size  $n$  and privately answer an exponential number of fixed correlated vector queries, even if the queries are presented online, but it is impossible to answer more than a constant number of adaptively-chosen correlated vector queries under differential privacy.

The queries are defined on datasets  $x \in \{\pm 1\}^n$  (hence the data universe is  $X = \{\pm 1\}$ ). For every query, the range  $\mathcal{R} = \{\pm 1\}^n$  is the set of  $n$ -bit vectors. We fix some parameters  $0 < \alpha < 1$  and  $m \in \mathbb{N}$ . A query  $q$  is specified by a set  $V$  where  $V = \{v^1, \dots, v^m\} \subseteq \{\pm 1\}^n$  is a set of  $n$ -bit vectors. Roughly, an accurate answer to a given search query is any vector  $y \in \{\pm 1\}^n$  that is approximately  $\alpha$ -correlated with the input dataset  $x \in \{\pm 1\}^n$  and has nearly as little correlation as possible with every  $v^j$ . By “as little correlation as possible with  $v^j$ ” we mean that  $v^j$  may itself be correlated with  $x$ , in which case  $y$  should be correlated with  $v^j$  only insofar as this correlation comes through the correlation between  $y$  and  $x$ . Formally, for a query  $q_V$ , we define the loss function  $L_{q_V} : X^n \times X^n \rightarrow \{0, 1\}$  by

$$L_{q_V}(x, y) = 0 \iff |\langle y - \alpha x, x \rangle| \leq \frac{\alpha^2 n}{100} \wedge \forall v^j \in V \quad |\langle y - \alpha x, v^j \rangle| \leq \frac{\alpha^2 n}{100}.$$

We remark that the choice of  $\alpha^2 n/100$  is somewhat arbitrary, and we can replace this choice with  $C$  for any  $\sqrt{n} \ll C \ll n$  and obtain quantitatively different results. We chose to fix this particular choice in order to reduce notational clutter. We let

$$Q_{\text{corr}}^{n, \alpha, m} = \{q_V \mid V \subseteq \{\pm 1\}^n, |V| \leq m\}$$

be the set of all correlated vector queries on  $\{\pm 1\}^n$  for parameters  $\alpha, m$ .

**4.1. Answering Online Correlated Vector Queries.** Provided that all the queries are fixed in advance, we can privately answer correlated vector queries using the randomized response algorithm. This algorithm simply takes the input vector  $x \in \{\pm 1\}^n$  and outputs a new vector  $y \in \{\pm 1\}^n$  where each bit  $y_i$  is independent and is set to  $x_i$  with probability  $1/2 + \rho$  for a suitable choice of  $\rho > 0$ . The algorithm will then answer every correlated vector query with this same vector  $y$ . The following theorem captures the parameters that this mechanism achieves.

**Theorem 4.1** Answering Online Correlated Vector Queries. *For every  $0 < \alpha < 1/2$ , there exists  $k = 2^{\Omega(\alpha^4 n)}$  such that, for every sufficiently large  $n \in \mathbb{N}$ , there is a  $(3\alpha, 0)$ -differentially private algorithm  $M_{\text{corr}}$  that takes a dataset  $x \in \{\pm 1\}^n$  and is  $(1/k)$ -accurate for  $k$  online queries from  $Q_{\text{corr}}^{n, \alpha, k}$ .*

*Proof Theorem 4.1.* Our algorithm based on randomized response is presented in Figure 5 below.

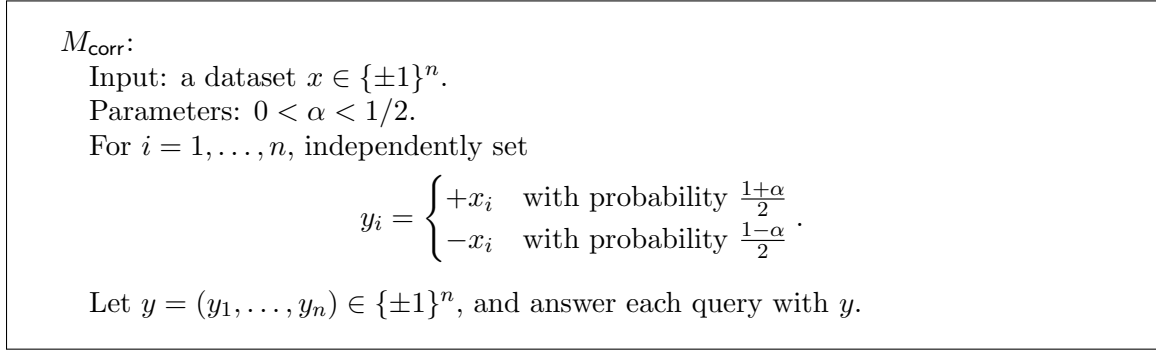


Figure 5:  $M_{\text{corr}}$

To establish privacy, observe that by construction each output bit  $y_i$  depends only on  $x_i$  and is independent of all  $x_j, y_j$  for  $j \neq i$ . Therefore, it suffices to observe that if  $0 < \alpha < 1/2$ ,

$$1 \leq \frac{\mathbb{P}[y_i = +1 \mid x_i = +1]}{\mathbb{P}[y_i = +1 \mid x_i = -1]} = \frac{1 + \alpha}{1 - \alpha} \leq e^{3\alpha}$$

and similarly

$$1 \geq \frac{\mathbb{P}[y_i = -1 \mid x_i = +1]}{\mathbb{P}[y_i = -1 \mid x_i = -1]} = \frac{1 - \alpha}{1 + \alpha} \geq e^{-3\alpha}.$$

To prove accuracy, observe that since the output  $y$  does not depend on the sequence of queries, we can analyze the mechanism as if the queries  $q_{V_1}, \dots, q_{V_k} \in Q_{\text{corr}}^{n, \alpha, k}$  were fixed and given all at once. Let  $V = \bigcup_{j=1}^k V_j$ , and note that  $|V| \leq k^2$ . First, observe that  $\mathbb{E}[y] = \alpha x$ . Thus we have

$$\mathbb{E}_y [\langle y - \alpha x, x \rangle] = 0 \text{ and } \forall v \in V \mathbb{E}_y [\langle y - \alpha x, v \rangle] = 0$$

Since  $x$  and every vector in  $V$  is fixed independently of  $y$ , and the coordinates of  $y$  are independent by construction, the quantities  $\langle y, x \rangle$  and  $\langle y, v \rangle$  are each the sum of  $n$  independent



$\{\pm 1\}$ -valued random variables. Thus, we can apply Hoeffding's inequality<sup>5</sup> and a union bound to conclude

$$\begin{aligned} \mathbb{P}_y \left[ |\langle y - \alpha x, x \rangle| > \frac{\alpha^2 n}{100} \right] &\leq 2 \exp \left( \frac{-\alpha^4 n}{20000} \right) \\ \mathbb{P}_y \left[ \exists v \in V \text{ s.t. } |\langle y - \alpha x, v \rangle| > \frac{\alpha^2 n}{100} \right] &\leq 2k^2 \exp \left( \frac{-\alpha^4 n}{20000} \right) \end{aligned}$$

The theorem now follows by setting an appropriate choice of  $k = 2^{\Omega(\alpha^4 n)}$  such that  $2(k^2 + 1) \cdot \exp \left( \frac{-\alpha^4 n}{20000} \right) \leq 1/k$ .  $\square$

**4.2. A Lower Bound for Adaptive Correlated Vector Queries.** We now prove a contrasting lower bound showing that if the queries may be chosen adaptively, then no differentially private algorithm can answer more than a constant number of correlated vector queries. The key to our lower bound is that fact that adaptively-chosen correlated vector queries allow an adversary to obtain many vectors  $y^1, \dots, y^k$  that are correlated with  $x$  but pairwise nearly orthogonal with each other. As we prove, if  $k$  is sufficiently large, this information is enough to recover a vector  $\tilde{x}$  that has much larger correlation with  $x$  than any of the vectors  $y^1, \dots, y^k$  have with  $x$ . By setting the parameters appropriately, we will obtain a contradiction to differential privacy.

**Theorem 4.2** Lower Bound for Correlated Vector Queries. *For every  $0 < \alpha < 1/2$ , there is a  $k = O(1/\alpha^2)$  such that for every sufficiently large  $n \in \mathbb{N}$ , there is no  $(1, 1/20)$ -differentially private algorithm that takes a dataset  $x \in \{\pm 1\}^n$  and is  $1/100$ -accurate for  $k$  adaptive queries from  $Q_{\text{corr}}^{n, \alpha, k}$*

We remark that the value of  $k$  in our lower bound is optimal up to constants, as there is a  $(1, 1/20)$ -differentially private algorithm that can answer  $k = \Omega(1/\alpha^2)$  adaptively-chosen queries of this sort. The algorithm simply answers each query with an independent invocation of randomized response. Randomized response is  $O(\alpha)$ -differentially private for each query, and we can invoke the adaptive composition theorem [Dwork et al., 2006, 2010b] to argue differential privacy for  $k = \Omega(1/\alpha^2)$ -queries.

Before proving Theorem 4.2, we state and prove the combinatorial lemma that forms the foundation of our lower bound.

**Lemma 4.1** (Reconstruction Lemma). *Fix parameters  $0 \leq a, b \leq 1$ . Let  $x \in \{\pm 1\}^n$  and  $y^1, \dots, y^k \in \{\pm 1\}^n$  be vectors such that*

$$\begin{aligned} \forall 1 \leq j \leq k \quad \langle y^j, x \rangle &\geq an \\ \forall 1 \leq j < j' \leq k \quad |\langle y^j, y^{j'} \rangle| &\leq bn. \end{aligned}$$

<sup>5</sup>We use the following statement of Hoeffding's Inequality: if  $Z_1, \dots, Z_n$  are independent  $\{\pm 1\}$ -valued random variables, and  $Z = \sum_{i=1}^n Z_i$ , then

$$\mathbb{P} \left[ \left| Z - \mathbb{E}[Z] \right| > C\sqrt{n} \right] \leq 2e^{-C^2/2}$$

Then, if we let  $\tilde{x} = \text{sign}(\sum_{j=1}^k y^j) \in \{\pm 1\}^n$  be the coordinate-wise majority of  $y^1, \dots, y^k$ , we have

$$\langle \tilde{x}, x \rangle \geq \left(1 - \frac{2}{a^2 k} - \frac{2(b - a^2)}{a^2}\right) n.$$

*Proof of Lemma 4.1.* Let

$$\bar{y} = \frac{1}{k} \sum_{j=1}^k y^j \in [-1, 1]^n.$$

By linearity,  $\langle \bar{y}, x \rangle \geq an$  and

$$\|\bar{y}\|_2^2 = \frac{1}{k^2} \sum_{j,j'=1}^k \langle y^j, y^{j'} \rangle \leq \frac{1}{k^2} (kn + (k^2 - k)bn) \leq \left(\frac{1}{k} + b\right) n.$$

Define a random variable  $W \in [-1, 1]$  to be  $x_i \bar{y}_i$  for a uniformly random  $i \in [n]$ . Then

$$\mathbb{E}[W] = \frac{1}{n} \langle x, \bar{y} \rangle \geq a \text{ and } \mathbb{E}[W^2] = \frac{1}{n} \sum_{i=1}^n x_i^2 \bar{y}_i^2 = \frac{1}{n} \|\bar{y}\|_2^2 \leq \frac{1}{k} + b$$

By Chebyshev's inequality,

$$\mathbb{P}[W \leq 0] \leq \mathbb{P}[|W - \mathbb{E}[W]| \geq a] \leq \frac{\text{Var}[W]}{a^2} = \frac{\mathbb{E}[W^2] - \mathbb{E}[W]^2}{a^2} \leq \frac{\frac{1}{k} + b - a^2}{a^2}.$$

Meanwhile,

$$\mathbb{P}[W \leq 0] = \frac{1}{n} \sum_{i=1}^n \mathbb{I}[x_i \bar{y}_i \leq 0] \geq \frac{1}{n} \sum_{i=1}^n \mathbb{I}[\text{sign}(\bar{y}_i) \neq x_i] = \frac{1}{2} - \frac{1}{2n} \langle \text{sign}(\bar{y}), x \rangle.$$

Thus we conclude

$$\langle \text{sign}(\bar{y}), x \rangle \geq n - 2n\mathbb{P}[W \leq 0] \geq n - 2n \left(\frac{\frac{1}{k} + b - a^2}{a^2}\right)$$

To complete the proof, we rearrange terms and note that  $\text{sign}(\bar{y}) = \text{sign}(\sum_{j=1}^k y^j)$ .  $\square$

Now we are ready to prove our lower bound for algorithms that answer adaptively-chosen correlated vector queries.

*Proof of Theorem 4.2.* We will show that the output  $y^1, \dots, y^k$  of any algorithm  $M$  that takes a dataset  $x \in \{\pm 1\}^n$  and answers  $k = 100/\alpha^2$  adaptively-chosen correlated vector queries can be used to find a vector  $\tilde{x} \in \{\pm 1\}^n$  such that  $\langle \tilde{x}, x \rangle > n/2$ . In light of Lemma 4.1, this vector will simply be  $\tilde{x} = \text{sign}(\sum_{j=1}^k y^j)$ . We will then invoke the following elementary fact that differentially private algorithms do not admit this sort of reconstruction of their input dataset.

**Fact 4.2.** For every sufficiently large  $n \in \mathbb{N}$ , there is no  $(1, 1/20)$ -differentially private algorithm  $M : \{\pm 1\}^n \rightarrow \{\pm 1\}^n$  such that for every  $x \in \{\pm 1\}^n$ , with probability at least  $99/100$ ,  $\langle M(x), x \rangle > n/2$ .

The attack works as follows. For  $j = 1, \dots, k$ , define the set  $V_j = \{y^1, \dots, y^{j-1}\}$  and ask the query  $q_{V_j}(x) \in Q_{\text{corr}}^{n, \alpha, k}$  to obtain some vector  $y^j$ . Since  $M$  is assumed to be accurate for

$k$  adaptively-chosen queries, with probability  $99/100$ , we obtain vectors  $y^1, \dots, y^k \in \{\pm 1\}^n$  such that

$$\begin{aligned} \forall 1 \leq j \leq k \quad \langle y^j, x \rangle &\geq \langle \alpha x, x \rangle - |\langle y - \alpha x, x \rangle| \\ &\geq \alpha n - \frac{\alpha^2 n}{100} \\ &\geq an, \end{aligned}$$

$$\begin{aligned} \forall 1 \leq j < j' \leq k \quad |\langle y^j, y^{j'} \rangle| &\leq |\langle \alpha x, y^j \rangle| + |\langle y^{j'} - \alpha x, y^j \rangle| \\ &\leq \alpha |\langle y^j, x \rangle| + \frac{\alpha^2 n}{100} \\ &\leq \alpha (|\langle \alpha x, x \rangle| + |\langle y^j - \alpha x, x \rangle|) + \frac{\alpha^2 n}{100} \\ &\leq \alpha^2 n + \frac{\alpha^3 n}{100} + \frac{\alpha^2 n}{100} \\ &\leq \frac{51}{50} \alpha^2 n \\ &= bn, \end{aligned}$$

where  $a = 99\alpha/100$  and  $b = 51\alpha^2/50$ . Thus, by Lemma 4.1, if  $\tilde{x} = \text{sign}(\sum_{j=1}^k y^j)$ , and  $k = 100/\alpha^2$ , we have

$$\begin{aligned} \langle \tilde{x}, x \rangle &\geq \left( 1 - \frac{2}{a^2 k} - \frac{2(b-a^2)}{a^2} \right) n \\ &= \left( 1 - \frac{2}{(99\alpha/100)^2 k} - \frac{2(51\alpha^2/50 - (99\alpha/100)^2)}{(99\alpha/100)^2} \right) n \\ &= \left( 1 - \frac{2(100/99)^2}{100} - 2 \left( \frac{(51/50) - (99/100)^2}{(99/100)^2} \right) \right) n \\ &\geq 0.89n \geq n/2. \end{aligned}$$

By Fact 4.2, this proves that  $M$  cannot be  $(1, 1/20)$ -differentially private.  $\square$

## 5. THRESHOLD QUERIES

First we define threshold queries, which are a family of statistical queries.

**Definition 5.1.** Let  $\text{Thresh}_X$  denote the class of threshold queries over a totally ordered domain  $X$ . That is,  $\text{Thresh}_X = \{c_x : x \in X\}$  where  $c_x : X \rightarrow \{0, 1\}$  is defined by  $c_x(y) = 1$  iff  $y \leq x$ .

**5.1. Separation for Pure Differential Privacy.** In this section, we show that the sample complexity of answering adaptively-chosen thresholds can be exponentially larger than that of answering thresholds offline.

**Proposition 5.1** ([Dwork et al., 2010a, Chan et al., 2011, Dwork et al., 2015b]). *Let  $X$  be any totally ordered domain. Then there exists a  $(\varepsilon, 0)$ -differentially private mechanism  $M$  that, given  $x \in X^n$ , gives  $\alpha$ -accurate answers to  $k$  offline queries from  $\text{Thresh}_X$  for*

$$n = O\left(\min\left\{\frac{\log k + \log^2(1/\alpha)}{\alpha\varepsilon}, \frac{\log^2 k}{\alpha\varepsilon}\right\}\right)$$

On the other hand, we show that answering  $k$  adaptively-chosen threshold queries can require sample complexity as large as  $\Omega(k)$  – an exponential gap. Note that this matches the upper bound given by the Laplace mechanism [Dwork et al., 2006].

**Proposition 5.2.** *Answering  $k$  adaptively-chosen threshold queries on  $X = [2^{k-1}]$  to accuracy  $\alpha$  subject to  $\varepsilon$ -differential privacy requires sample complexity  $n = \Omega(k/\alpha\varepsilon)$ .*

The idea for the lower bound is that an analyst may adaptively choose  $k$  threshold queries to binary search for an “approximate median” of the dataset. However, a packing argument shows that locating an approximate median requires sample complexity  $\Omega(k)$ .

**Definition 5.2** (Approximate Median). Let  $X$  be a totally ordered domain,  $\alpha > 0$ , and  $x \in X^n$ . We call  $y \in X$  an  $\alpha$ -approximate median of  $x$  if

$$\frac{1}{n} |\{i \in [n] : x_i \leq y\}| \geq \frac{1}{2} - \alpha \quad \text{and} \quad \frac{1}{n} |\{i \in [n] : x_i \geq y\}| \geq \frac{1}{2} - \alpha.$$

Proposition 5.2 is obtained by combining Lemmas 5.3 and 5.4 below.

**Lemma 5.3.** *Suppose  $M$  answers  $k = \lceil 1 + \log_2 T \rceil$  adaptively-chosen queries from  $\text{Thresh}_{[T]}$  with  $\varepsilon$ -differential privacy and  $(\alpha, \beta)$ -accuracy. Then there exists an  $\varepsilon$ -differentially private  $M' : [T]^n \rightarrow [T]$  that computes an  $\alpha$ -approximate median with probability at least  $1 - \beta$ .*

*Proof.* The algorithm  $M'$ , formalized in Figure 6, uses  $M$  to perform a binary search.

Input:  $x \in X^n$ .  
 $M$  is given  $x$ .  
Initialize  $\ell_1 = 0$ ,  $u_1 = T$ , and  $j = 1$ .  
While  $u_j - \ell_j > 1$  repeat:  
  Let  $m_j = \lceil (u_j + \ell_j)/2 \rceil$ .  
  Give  $M$  the query  $c_{m_j} \in \text{Thresh}_{[T]}$  and obtain the answer  $a_j \in [0, 1]$ .  
  If  $a_j \geq \frac{1}{2}$ , set  $(\ell_{j+1}, u_{j+1}) = (\ell_j, m_j)$ ; otherwise set  $(\ell_{j+1}, u_{j+1}) = (m_j, u_j)$ .  
  Increment  $j$ .  
Output  $u_j$ .

Figure 6:  $M' : X^n \rightarrow X$

We have  $u_1 - \ell_1 = T$  and, after every query  $j$ ,  $u_{j+1} - \ell_{j+1} \leq \lceil (u_j - \ell_j)/2 \rceil$ . Since the process stops when  $u_j - \ell_j = 1$ , it is easy to verify that  $M'$  makes at most  $\lceil 1 + \log_2(T - 1) \rceil$  queries to  $M$ .

Suppose all of the answers given by  $M$  are  $\alpha$ -accurate. This happens with probability at least  $1 - \beta$ . We will show that, given this,  $M'$  outputs an  $\alpha$ -approximate median, which completes the proof.

We claim that  $c_{u_j}(x) \geq \frac{1}{2} - \alpha$  for all  $j$ . This is easily shown by induction. The base case is  $c_T(x) = 1 \geq \frac{1}{2} - \alpha$ . At each step either  $u_{j+1} = u_j$  (in which case the induction hypothesis can be applied) or  $u_{j+1} = m_j$ ; in the latter case our accuracy assumption gives

$$c_{u_{j+1}}(x) = c_{m_j}(x) \geq a_j - \alpha \geq \frac{1}{2} - \alpha.$$

We also claim that  $c_{\ell_j}(x) < \frac{1}{2} + \alpha$  for all  $j$ . This follows from a similar induction and completes the proof.  $\square$

**Lemma 5.4.** *Let  $M : [T]^n \rightarrow [T]$  be an  $\varepsilon$ -differentially private algorithm that computes an  $\alpha$ -approximate median with confidence  $1 - \beta$ . Then*

$$n \geq \Omega\left(\frac{\log T + \log(1/\beta)}{\alpha\varepsilon}\right).$$

*Proof.* Let  $m = \lceil(\frac{1}{2} - \alpha)n\rceil - 1$ . For each  $t \in [T]$ , let  $x^t \in [T]^n$  denote the dataset containing  $m$  copies of 1,  $m$  copies of  $T$ , and  $n - 2m$  copies of  $t$ . Then for each  $t \in [T]$ ,

$$\mathbb{P}[M(x^t) = t] \geq 1 - \beta.$$

On the other hand, by the pigeonhole principle, there must exist  $t_* \in [T - 1]$  such that

$$\mathbb{P}[M(x^T) = t_*] \leq \frac{\mathbb{P}[M(x^T) \in [T - 1]]}{T - 1} \leq \frac{\beta}{T - 1}.$$

The inputs  $x^T$  and  $x^{t_*}$  differ in at most  $n - 2m \leq 2\alpha n + 2$  entries. By group privacy,

$$1 - \beta \leq \mathbb{P}[M(x^{t_*}) = t_*] \leq e^{\varepsilon(2\alpha n + 2)} \mathbb{P}[M(x^T) = t_*] \leq e^{\varepsilon(2\alpha n + 2)} \frac{\beta}{T - 1}.$$

Rearranging these inequalities gives

$$O(\varepsilon\alpha n) \geq \varepsilon(2\alpha n + 2) \geq \log\left(\frac{(1 - \beta)(T - 1)}{\beta}\right) \geq \Omega(\log(T/\beta)),$$

which yields the result.  $\square$

**Remark 5.3 .** *Proposition 5.2 can be extended to online non-adaptive queries, which yields a separation between the online non-adaptive and offline models for pure differential privacy and threshold queries.*

The key observation behind Remark 5.3 is that, while Lemma 5.3 in general requires making adaptive queries, for the inputs  $x^t \in [T]^n$  ( $t \in [T]$ ) used in Lemma 5.4 the queries are “predictable.” In particular, on input  $x^t$ , the algorithm  $M'$  from the proof of Lemma 5.3 will (with probability at least  $1 - \beta$ ) always make the same sequence queries. This allows the queries to be specified in advance in a non-adaptive manner. More precisely, we can produce an algorithm  $M'_t$  that produces non-adaptive online queries by simulating  $M'$  on input  $x^t$  and using those queries. Given the answers to these online non-adaptive queries,  $M'_t$  can either accept or reject its input depending on whether the answers are consistent with the input  $x^t$ ;  $M'_t$  will accept  $x^t$  with high probability and reject  $x^{t'}$  for  $t' \neq t$  with high probability. The proof of Lemma 5.4 can be carried out using  $M'_{t_*}$  instead of  $M'$  at the end.

**5.2. The BetweenThresholds Algorithm.** The key technical novelty behind our algorithm for answering adaptively-chosen threshold queries is a refinement of the “Above Threshold” algorithm [Dwork and Roth, 2014, §3.6], which underlies the ubiquitous “sparse vector” technique [Dwork et al., 2009, Roth and Roughgarden, 2010, Dwork et al., 2010a, Hardt and Rothblum, 2010].

The sparse vector technique addresses a setting where we have a stream of  $k$  (adaptively-chosen) low-sensitivity queries and a threshold parameter  $t$ . Instead of answering all  $k$  queries accurately, we are interested in answering only the ones that are above the threshold  $t$  – for the remaining queries, we only require a signal that they are below the threshold. Intuitively, one would expect to only pay in privacy for the queries that are actually above the threshold. And indeed, one can get away with sample complexity proportional to the number of queries that are above the threshold, and to the *logarithm* of the total number of queries.

We extend the sparse vector technique to settings where we demand slightly more information about each query beyond whether it is below a single threshold. In particular, we set two thresholds  $t_\ell < t_u$ , and for each query, release a signal as to whether the query is below the lower threshold, above the upper threshold, or between the two thresholds.

As long as the thresholds are sufficiently far apart, whether (the noisy answer to) a query is below the lower threshold or above the upper threshold is *stable*, in that it is extremely unlikely to change on neighboring datasets. As a result, we obtain an  $(\varepsilon, \delta)$ -differentially private algorithm that achieves the same accuracy guarantees as the traditional sparse vector technique, i.e. sample complexity proportional to  $\log k$ .

Our algorithm is summarised by the following theorem.<sup>6</sup>

**Theorem 5.5.** *Let  $\alpha, \beta, \varepsilon, \delta, t \in (0, 1)$  and  $n, k \in \mathbb{N}$  satisfy*

$$n \geq \frac{1}{\alpha\varepsilon} \max \{12 \log(30/\varepsilon\delta), 16 \log((k+1)/\beta)\}.$$

*Then there exists a  $(\varepsilon, \delta)$ -differentially private algorithm that takes as input  $x \in X^n$  and answers a sequence of adaptively-chosen queries  $q_1, \dots, q_k : X^n \rightarrow [0, 1]$  of sensitivity  $1/n$  with  $a_1, \dots, a_{\leq k} \in \{\text{L}, \text{R}, \text{T}\}$  such that, with probability at least  $1 - \beta$ ,*

- $a_j = \text{L} \implies q_j(x) \leq t$ ,
- $a_j = \text{R} \implies q_j(x) \geq t$ , and
- $a_j = \text{T} \implies t - \alpha \leq q_j(x) \leq t + \alpha$ .

*The algorithm may halt before answering all  $k$  queries; however, it only halts after outputting  $\text{T}$ .*

Our algorithm is given in Figure 7. The analysis is split into Lemmas 5.6 and 5.7.

**Lemma 5.6** (Privacy for BetweenThresholds). *Let  $\varepsilon, \delta \in (0, 1)$  and  $n \in \mathbb{N}$ . Then BetweenThresholds (Figure 7) is  $(\varepsilon, \delta)$ -differentially private for any adaptively-chosen sequence of queries as long as the gap between the thresholds  $t_\ell, t_u$  satisfies*

$$t_u - t_\ell \geq \frac{12}{\varepsilon n} (\log(10/\varepsilon) + \log(1/\delta) + 1).$$

---

<sup>6</sup>In Theorem 5.5, only one threshold is allowed. However, our algorithm is more general and permits the setting of two thresholds. We have chosen this statement for simplicity.

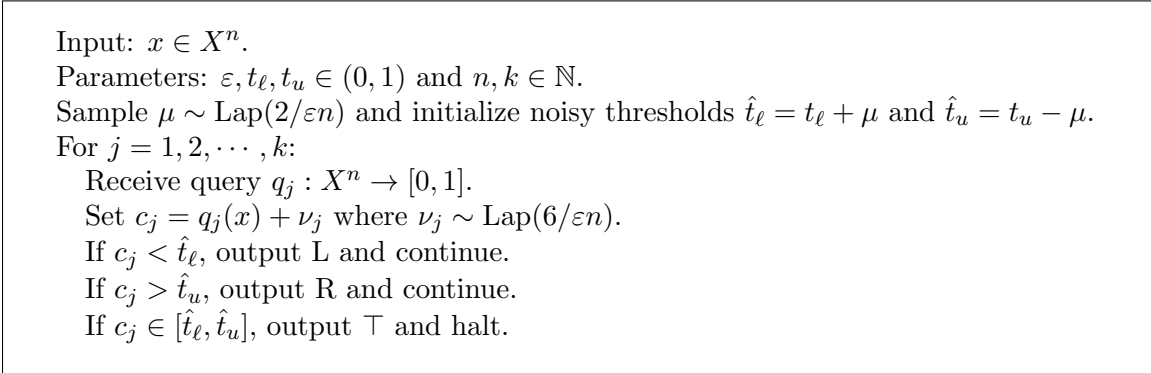


Figure 7: BetweenThresholds

**Lemma 5.7** (Accuracy for BetweenThresholds). *Let  $\alpha, \beta, \varepsilon, t_\ell, t_u \in (0, 1)$  and  $n, k \in \mathbb{N}$  satisfy*

$$n \geq \frac{8}{\alpha\varepsilon} (\log(k+1) + \log(1/\beta)).$$

*Then, for any input  $x \in X^n$  and any adaptively-chosen sequence of queries  $q_1, q_2, \dots, q_k$ , the answers  $a_1, a_2, \dots, a_{\leq k}$  produced by BetweenThresholds (Figure 7) on input  $x$  satisfy the following with probability at least  $1 - \beta$ . For any  $j \in [k]$  such that  $a_j$  is returned before BetweenThresholds halts,*

- $a_j = \text{L} \implies q_j(x) \leq t_\ell + \alpha,$
- $a_j = \text{R} \implies q_j(x) \geq t_u - \alpha,$  and
- $a_j = \top \implies t_\ell - \alpha \leq q_j(x) \leq t_u + \alpha.$

Combining Lemmas 5.6 and 5.7 and setting  $t_\ell = t - \alpha/2$  and  $t_u = t + \alpha/2$  yields Theorem 5.5.

*Proof of Lemma 5.6.* Our analysis is an adaptation of Dwork and Roth [2014, §3.6]’s analysis of the AboveThreshold algorithm. Recall that a transcript of the execution of BetweenThresholds is given by  $a \in \{\text{L}, \text{R}, \top\}^*$ . Let  $\mathcal{M} : X^n \rightarrow \{\text{L}, \text{R}, \top\}^*$  denote the function that simulates BetweenThresholds interacting with a given adaptive adversary (cf. Figure 3) and returns the transcript.

Let  $S \subset \{\text{L}, \text{R}, \top\}^*$  be a set of transcripts. Our goal is to show that for adjacent datasets  $x \sim x'$ ,

$$\mathbb{P}[\mathcal{M}(x) \in S] \leq e^\varepsilon \mathbb{P}[\mathcal{M}(x') \in S] + \delta.$$

Let

$$z^* = \frac{1}{2}(t_u - t_\ell) - \frac{6}{\varepsilon n} \log(10/\varepsilon) - 1/n \geq \frac{2}{\varepsilon n} \log(1/\delta).$$

Our strategy will be to show that as long as the noise value  $\mu$  is under control, in particular if  $\mu \leq z^*$ , then the algorithm behaves in essentially the same way as the standard AboveThreshold algorithm. Meanwhile, the event  $\mu > z^*$  which corresponds to the (catastrophic) event where the upper and lower thresholds are too close or overlap, happens with probability at most  $\delta$ .

The following claim reduces the privacy analysis to examining the probability of obtaining any single transcript  $a$ :

**Claim 5.8.** Suppose that for any transcript  $a \in \{\text{L}, \text{R}, \top\}^*$ , and any  $z \leq z^*$ , that

$$\mathbb{P}[\mathcal{M}(x) = a | \mu = z] \leq e^{\varepsilon/2} \mathbb{P}[\mathcal{M}(x') = a | \mu = z + 1/n].$$

Then  $\mathcal{M}$  is  $(\varepsilon, \delta)$ -differentially private.

*Proof.* By properties of the Laplace distribution, since  $\mu \sim \text{Lap}(2/\varepsilon n)$ , for any  $z \in \mathbb{R}$ , we have

$$\mathbb{P}[\mu = z] \leq e^{\varepsilon/2} \mathbb{P}[\mu = z + 1/n],$$

and

$$\mathbb{P}[\mu > z^*] = \frac{1}{2} e^{-\varepsilon n z^*/2} \leq \delta.$$

Fix a set of transcripts  $S$ . Combining these properties allows us to write

$$\begin{aligned} \mathbb{P}[\mathcal{M}(x) \in S] &= \int_{\mathbb{R}} \mathbb{P}[\mathcal{M}(x) \in S | \mu = z] \mathbb{P}[\mu = z] dz \\ &\leq \left( \int_{-\infty}^{z^*} \mathbb{P}[\mathcal{M}(x) \in S | \mu = z] \mathbb{P}[\mu = z] dz \right) + \mathbb{P}[\mu > z^*] \\ &\leq \left( e^{\varepsilon/2} \int_{-\infty}^{z^*} \mathbb{P}[\mathcal{M}(x') \in S | \mu = z + 1/n] \mathbb{P}[\mu = z] dz \right) + \delta \\ &\leq \left( e^{\varepsilon} \int_{-\infty}^{z^*} \mathbb{P}[\mathcal{M}(x') \in S | \mu = z + 1/n] \mathbb{P}[\mu = z + 1/n] dz \right) + \delta \\ &\leq e^{\varepsilon} \mathbb{P}[\mathcal{M}(x') \in S] + \delta \end{aligned}$$

□

Returning to the proof of Lemma 5.6, fix a transcript  $a \in \{\text{L}, \text{R}, \top\}^*$ . Our goal is now to show that  $\mathcal{M}$  satisfies the hypotheses of Claim 5.8, namely that for any  $z \leq z^*$ ,

$$\mathbb{P}[\mathcal{M}(x) = a | \mu = z] \leq e^{\varepsilon/2} \mathbb{P}[\mathcal{M}(x') = a | \mu = z + 1/n]. \quad (5.1)$$

For some  $k \geq 1$ , we can write the transcript  $a$  as  $(a_1, a_2, \dots, a_k)$ , where  $a_j \in \{\text{L}, \text{R}\}$  for each  $j < k$ , and  $a_k = \top$ .

For convenience, let  $A = \mathcal{M}(x)$  and  $A' = \mathcal{M}(x')$ . We may decompose

$$\begin{aligned} \mathbb{P}[\mathcal{M}(x) = a | \mu = z] &= \mathbb{P}[(\forall j < k, A_j = a_j) \wedge q_k(x) + \nu_k \in [\hat{t}_\ell, \hat{t}_u] | \mu = z] \\ &= \mathbb{P}[(\forall j < k, A_j = a_j) | \mu = z] \\ &\quad \cdot \mathbb{P}[q_k(x) + \nu_k \in [\hat{t}_\ell, \hat{t}_u] | \mu = z \wedge (\forall j < k, A_j = a_j)]. \end{aligned} \quad (5.2)$$

We upper bound each factor on the right-hand side separately.

**Claim 5.9.**

$$\mathbb{P}[(\forall i < k, A_i = a_i) | \mu = z] \leq \mathbb{P}[(\forall i < k, A'_i = a_i) | \mu = z + 1/n]$$

*Proof.* For fixed  $z$ , let  $A_z(x)$  denote the set of noise vectors  $(\nu_1, \dots, \nu_{k-1})$  for which  $(A_1, \dots, A_{k-1}) = (a_1, \dots, a_{k-1})$  when  $\nu = z$ . We claim that as long as  $z \leq z^*$ , then  $A_z(x) \subseteq A_{z+1/n}(x')$ . To argue this, let  $(\nu_1, \dots, \nu_{k-1}) \in A_z(x)$ . Fix an index  $j \in \{1, \dots, k-1\}$



and suppose  $a_j = \text{L}$ . Then  $q_j(x) + \nu_j < t_\ell + z$ , but since  $q_j$  has sensitivity  $1/n$ , we also have  $q_j(x') + \nu_j < t_\ell + (z + 1/n)$ . Likewise, if  $a_j = \text{R}$ , then  $q_j(x) + \nu_j > t_u - z$ , so

$$q_j(x') + \nu_j > t_u - z - 1/n \geq t_\ell + (z + 1/n)$$

as long as  $z \leq z^* \leq \frac{1}{2}(t_u - t_\ell) - 1/n$ . (This ensures that  $\mathcal{M}(x')$  does not output L on the first branch of the “if” statement, and proceeds to output R.)

Since  $A_z(x) \subseteq A_{z+1/n}(x')$ , this proves that

$$\begin{aligned} \mathbb{P}[(\forall i < k, A_i = a_i) | \mu = z] &= \mathbb{P}[(\nu_1, \dots, \nu_{k-1}) \in A_z(x)] \\ &\leq \mathbb{P}[(\nu_1, \dots, \nu_{k-1}) \in A_{z+1/n}(x')] \\ &= \mathbb{P}[(\forall i < k, A'_i = a_i) | \mu = z + 1/n]. \end{aligned}$$

□

Given Claim 5.9, all that is needed to prove (5.1) and, thereby, prove Lemma 5.6 is to bound the second factor in (5.2) — that is, we must only show that

$$\begin{aligned} &\mathbb{P}[q_k(x) + \nu_k \in [\hat{t}_\ell, \hat{t}_u] | \mu = z \wedge (\forall j < k, A_j = a_j)] \\ &\leq e^{\varepsilon/2} \mathbb{P}[q_k(x') + \nu_k \in [\hat{t}_\ell, \hat{t}_u] | \mu = z + 1/n \wedge (\forall j < k, A'_j = a_j)]. \end{aligned}$$

Let  $\Delta = (q_k(x') - q_k(x)) \in [-1/n, 1/n]$ . Then

$$\begin{aligned} &\mathbb{P}[q_k(x) + \nu_k \in [\hat{t}_\ell, \hat{t}_u] | \mu = z \wedge (\forall j < k, A_j = a_j)] \\ &= \mathbb{P}[t_\ell + z \leq q_k(x) + \nu_k \leq t_u - z] \\ &= \mathbb{P}[t_\ell + z + \Delta \leq q_k(x') + \nu_k \leq t_u - z + \Delta] \\ &= \mathbb{P}[t_\ell + (z + 1/n) + (\Delta - 1/n) \leq q_k(x') + \nu_k \leq t_u - (z + 1/n) + (\Delta + 1/n)] \\ &= \mathbb{P}[q_k(x') + \nu_k \in [\hat{t}_\ell + \Delta - 1/n, \hat{t}_u + \Delta + 1/n] | \mu = z + 1/n] \\ &\leq e^{\varepsilon/2} \mathbb{P}[q_k(x') + \nu_k \in [\hat{t}_\ell, \hat{t}_u] | \mu = z + 1/n] \\ &= e^{\varepsilon/2} \mathbb{P}[q_k(x') + \nu_k \in [\hat{t}_\ell, \hat{t}_u] | \mu = z + 1/n \wedge (\forall j < k, A'_j = a_j)] \end{aligned}$$

where the last inequality follows from Claim 5.10 below (setting  $\eta = 2/n$ ,  $\lambda = 6/\varepsilon n$ ,  $[a, b] = [\hat{t}_\ell, \hat{t}_u]$ , and  $[a', b'] = [\hat{t}_\ell + \Delta - 1/n, \hat{t}_u + \Delta + 1/n]$ ) and the fact that  $z \leq z^* = \frac{1}{2}(t_u - t_\ell) - \frac{6}{\varepsilon n} \log(10/\varepsilon) - 1/n$  implies

$$b - a = \hat{t}_u - \hat{t}_\ell = t_u - t_\ell - 2\mu \geq \frac{12}{\varepsilon n} \log\left(\frac{10}{\varepsilon}\right) \geq 2\lambda \log\left(\frac{1}{1 - e^{-\varepsilon/6}}\right)$$

whenever  $0 \leq \varepsilon \leq 1$ .

**Claim 5.10.** Let  $\nu \sim \text{Lap}(\lambda)$  and let  $[a, b], [a', b'] \subset \mathbb{R}$  be intervals satisfying  $[a, b] \subset [a', b']$ . If  $\eta \geq (b' - a') - (b - a)$ , then

$$\mathbb{P}[\nu \in [a', b']] \leq \frac{e^{\eta/\lambda}}{1 - e^{-(b-a)/2\lambda}} \cdot \mathbb{P}[\nu \in [a, b]].$$

*Proof.* Recall that the probability density function of the Laplace distribution is given by  $f_\lambda(x) = \frac{1}{2\lambda}e^{-|x|/\lambda}$ . There are four cases to consider: In the first case,  $a < b \leq 0$ . In the second case,  $a < 0 < b$  with  $|a| \leq |b|$ . In the third case,  $0 \leq a < b$ . Finally, in the fourth case,  $a < 0 < b$  with  $|a| \geq |b|$ . Since the Laplace distribution is symmetric, it suffices to analyze the first two cases.

Case 1: Suppose  $a < b \leq 0$ . Then

$$\begin{aligned} \mathbb{P}[\nu \in [a', b']] &\leq \mathbb{P}[\nu \in [a, b]] + \int_b^{b+\eta} \frac{1}{2\lambda} e^{x/\lambda} dx \\ &= \frac{1}{2}(e^{(b+\eta)/\lambda} - e^{a/\lambda}) \\ &= \frac{1}{2} \cdot \left( \frac{e^{\eta/\lambda} - e^{(a-b)/\lambda}}{1 - e^{(a-b)/\lambda}} \right) \cdot (e^{b/\lambda} - e^{a/\lambda}) \\ &= \left( \frac{e^{\eta/\lambda} - e^{-(b-a)/\lambda}}{1 - e^{-(b-a)/\lambda}} \right) \cdot \mathbb{P}[\nu \in [a, b]]. \end{aligned}$$

Case 2: Suppose  $a < 0 < b$  and  $|a| \leq |b|$ . Note that this implies  $b \geq (b-a)/2$ . Then

$$\begin{aligned} \mathbb{P}[\nu \in [a', b']] &\leq \mathbb{P}[\nu \in [a, b]] + \eta \cdot \frac{1}{2\lambda} e^{a/\lambda} \\ &\leq \mathbb{P}[\nu \in [a, b]] \left( 1 + \frac{\eta}{2\lambda} \frac{e^{a/\lambda}}{\mathbb{P}[\nu \in [0, b]]} \right) \\ &= \mathbb{P}[\nu \in [a, b]] \frac{1 - e^{-b/\lambda} + \frac{\eta}{\lambda} e^{a/\lambda}}{1 - e^{-b/\lambda}} \\ &\leq \mathbb{P}[\nu \in [a, b]] \frac{1 + \eta/\lambda}{1 - e^{-b/\lambda}} \\ &\leq \mathbb{P}[\nu \in [a, b]] \frac{e^{\eta/\lambda}}{1 - e^{-(b-a)/2\lambda}}. \end{aligned}$$

□

□

*Proof of Lemma 5.7.* We claim that it suffices to show that with probability at least  $1 - \beta$  we have

$$\forall 1 \leq j \leq k \quad |\nu_j| + |\mu| \leq \alpha.$$

To see this, suppose  $|\nu_j| + |\mu| \leq \alpha$  for every  $j$ . Then, if  $a_j = \text{L}$ , we have

$$c_j = q_j(x) + \nu_j < \hat{t}_\ell = t_\ell + \mu, \quad \text{whence} \quad q_j(x) < t_\ell + |\mu| + |\nu_j| \leq t_\ell + \alpha.$$

Similarly, if  $a_j = \text{R}$ , then

$$c_j = q_j(x) + \nu_j > \hat{t}_u = t_u - \mu, \quad \text{whence} \quad q_j(x) > t_u - (|\mu| + |\nu_j|) \geq t_u - \alpha.$$

Finally, if  $a_j = \text{T}$ , then

$$c_j = q_j(x) + \nu_j \in [\hat{t}_\ell, \hat{t}_u] = [t_\ell + \mu, t_u - \mu], \quad \text{whence} \quad t_\ell - \alpha \leq q_j(x) \leq t_u + \alpha.$$

We now show that indeed  $|\nu_j| + |\mu| \leq \alpha$  for every  $j$  with high probability. By tail bounds for the Laplace distribution,

$$\mathbb{P}[|\mu| > \alpha/4] = \exp\left(-\frac{\varepsilon\alpha n}{8}\right) \quad \text{and} \quad \mathbb{P}[|\nu_j| > 3\alpha/4] = \exp\left(-\frac{\varepsilon\alpha n}{8}\right)$$

for all  $j$ . By a union bound,

$$\mathbb{P}[|\mu| > \alpha/4 \vee \exists j \in [k] |\nu_j| > 3\alpha/4] \leq (k+1) \cdot \exp\left(-\frac{\varepsilon\alpha n}{8}\right) \leq \beta,$$

as required.  $\square$

**5.3. The Adaptive Interior Point Problem.** Our algorithm extends a result of Bun et al. [2015] showing how to reduce the problem of privately releasing thresholds to the much simpler *interior point problem*. By analogy, our algorithm for answering adaptively-chosen thresholds relies on solving multiple instances of an online variant of the interior point problem in parallel. In this section, we present the AIP problem and give an  $(\varepsilon, \delta)$ -differentially private solution that can handle  $k$  adaptively-chosen queries with sample complexity  $O(\log k)$ . Our AIP algorithm is a direct application of the `BetweenThresholds` algorithm from Section 5.2.

**Definition 5.4** (Adaptive Interior Point Problem). An algorithm  $M$  solves the *Adaptive Interior Point (AIP) Problem* for  $k$  queries with confidence  $\beta$  if, when given as input any private dataset  $x \in [0, 1]^n$  and any adaptively-chosen sequence of real numbers  $y_1, \dots, y_k \in [0, 1]$ , with probability at least  $1 - \beta$  it produces a sequence of answers  $a_1, \dots, a_k \in \{\text{L}, \text{R}\}$  such that

$$\forall j \in \{1, 2, \dots, k\} \quad y_j < \min_{i \in [n]} x_i \implies a_j = \text{L}, \quad y_j \geq \max_{i \in [n]} x_i \implies a_j = \text{R}.$$

(If  $\min_{i \in [n]} x_i \leq y_j < \max_{i \in [n]} x_i$ , then  $M$  may output either symbol L or R.)

Input: Dataset  $x \in [0, 1]^n$ .  
 Initialize a `BetweenThresholds` instance (Figure 7)  $\mathcal{B}$  on dataset  $x$  with thresholds  $t_\ell = \frac{1}{3}$ ,  $t_u = \frac{2}{3}$ .  
 For  $j = 1, 2, \dots, k$ :  
   Receive query  $y_j \in [0, 1]$ .  
   If  $\mathcal{B}$  already halted on some query  $q_{y^*}$ , output L if  $y_j < y^*$  and output R if  $y_j \geq y^*$ .  
  
   Otherwise, give  $\mathcal{B}$  the query  $c_{y_j} \in \text{Thresh}_{[0,1]}$ .  
   If  $\mathcal{B}$  returns  $\top$ , output R. Otherwise, output the answer produced by  $\mathcal{B}$ .

Figure 8: Adaptive Interior Point Algorithm

**Proposition 5.11.** *The algorithm in Figure 8 is  $(\varepsilon, \delta)$ -differentially private and solves the AIP Problem with confidence  $\beta$  as long as*

$$n \geq \frac{36}{\varepsilon} (\log(k+1) + \log(1/\beta) + \log(10/\varepsilon) + \log(1/\delta) + 1).$$

*Proof.* Privacy follows immediately from Lemma 5.6, since Algorithm 8 is obtained by post-processing Algorithm 7, run using thresholds with a gap of size  $1/3$ .

To argue utility, let  $\alpha = 1/3$  so that

$$n \geq \frac{8}{\varepsilon\alpha}(\log(k+1) + \log(1/\beta)).$$

By Lemma 5.7, with probability at least  $1 - \beta$ , the following events occur:

- If the `BetweenThresholds` instance  $\mathcal{B}$  halts when it is queried on  $c_{y^*}$ , then  $\min_{i \in [n]} x_i \leq y^* < \max_{i \in [n]} x_i$ .
- If  $\mathcal{B}$  has not yet halted and  $y_j < \min_{i \in [n]} x_i$ , its answer to  $c_{y_j}$  is L.
- If  $\mathcal{B}$  has not yet halted and  $y_j \geq \max_{i \in [n]} x_i$ , its answer to  $c_{y_j}$  is R.

Thus, if  $\mathcal{B}$  has not yet halted, the answers provided are accurate answers for the AIP Problem. On the other hand, when  $\mathcal{B}$  halts, it has successfully identified an “interior point” of the dataset  $x$ , i.e. a  $y^*$  such that  $\min_{i \in [n]} x_i \leq y^* < \max_{i \in [n]} x_i$ . Thus, for any subsequent query  $y$ , we have that

$$y < \min_{i \in [n]} x_i \implies y < y^*,$$

so Algorithm 8 correctly outputs L. Similarly,

$$y \geq \max_{i \in [n]} x_i \implies y \geq y^*,$$

so Algorithm 8 correctly outputs R on such a query.  $\square$

**5.4. Releasing Adaptive Thresholds with Approximate Differential Privacy.** We are now ready to state our reduction from releasing thresholds to solving the AIP Problem.

**Theorem 5.5 .** *If there exists an  $(\varepsilon, \delta)$ -differentially private algorithm solving the AIP problem for  $k$  queries with confidence  $\alpha\beta/8$  and sample complexity  $n'$ , then there is a  $(4\varepsilon, (1 + e^\varepsilon)\delta)$ -differentially private algorithm for releasing  $k$  threshold queries with  $(\alpha, \beta)$ -accuracy and sample complexity*

$$n = \max \left\{ \frac{6n'}{\alpha}, \frac{24 \log^{2.5}(4/\alpha) \cdot \log(2/\beta)}{\alpha\varepsilon} \right\}.$$

Combining this reduction with our algorithm for the AIP Problem (Proposition 5.11) yields:

**Corollary 5.6 .** *There is an  $(\varepsilon, \delta)$ -differentially private algorithm for releasing  $k$  adaptively-chosen threshold queries with  $(\alpha, \beta)$ -accuracy for*

$$n = O \left( \frac{\log k + \log^{2.5}(1/\alpha) + \log(1/\beta\varepsilon\delta)}{\alpha\varepsilon} \right).$$

*Proof of Theorem 5.5.* Our algorithm and its analysis follow the reduction of Bun et al. [2015] for reducing the (offline) query release problem for thresholds to the offline interior point problem.

Let  $T$  be an  $(\varepsilon, \delta)$ -differentially private algorithm solving the AIP Problem with confidence  $\alpha\beta/8$  and sample complexity  $n'$ . Without loss of generality, we may assume that  $T$  is differentially private in “add-or-remove-an-item sense”—i.e. if  $x \in [0, 1]^*$  and  $x'$  differs from  $x$  up to the addition or removal of a row, then for every adversary  $A$  and set  $S$  of outcomes of the interaction between  $A$  and  $T$ , we have  $\mathbb{P}[\text{Adaptive}_{A=T}(x) \in S] \leq e^\varepsilon \mathbb{P}[\text{Adaptive}_{A=T}(x') \in S] + \delta$ .

Moreover,  $T$  provides accurate answers to the AIP Problem with probability at least  $1 - \alpha\beta/8$  whenever its input is of size at least  $n'$ . To force an algorithm  $T$  to have these properties, we may pad any dataset of size less than  $n'$  with an arbitrary fixed element. On the other hand, we may subsample the first  $n'$  elements from any dataset with more than this many elements.

Consider the algorithm `AdaptiveThresholds $_T$`  in Figures 9 and 10.

Input: Dataset  $x \in [0, 1]^n$ .  
 Parameter:  $\alpha \in (0, 1)$ .  
 Let  $(x^{(1)}, \dots, x^{(M)}) \leftarrow_{\mathbf{R}} \text{Partition}(x_1, \dots, x_n, \alpha)$ .  
 Initialize an instance of the AIP algorithm  $T^{(m)}$  on each chunk  $x^{(m)} \in [0, 1]^*$ , for  $m \in [M]$ .  
 For each  $j = 1, \dots, k$ :  
     Receive query  $c_{y_j} \in \text{Thresh}_{[0,1]}$ .  
     Give query  $y_j \in [0, 1]$  to every AIP instance  $T^{(m)}$ , receiving answers  $a_j^{(1)}, \dots, a_j^{(M)} \in \{\mathbf{L}, \mathbf{R}\}$ .  
     Return  $a_j = \frac{1}{M} \cdot \left| \left\{ m \in [M] : a_j^{(m)} = \mathbf{R} \right\} \right|$ .

Figure 9: `AdaptiveThresholds $_T$`

Input: Dataset  $x \in [0, 1]^n$ .  
 Parameter:  $\alpha \in (0, 1)$ .  
 Output: (Random) partition  $(x^{(1)}, \dots, x^{(M)}) \in ([0, 1]^*)^M$  of  $x$ , where  $2/\alpha \leq M < 4/\alpha$ .  
  
 Let  $M = 2^{\lceil \log_2(2/\alpha) \rceil}$ .  
 Sort  $x$  in nondecreasing order  $x_1 \leq x_2 \leq \dots \leq x_n$ .  
 For each  $0 \leq \ell \leq \log_2 M$  and  $s \in \{0, 1\}^\ell$ , sample  $\nu_s \sim \text{Lap}((\log_2 M)/\varepsilon)$  independently.  
  
 For each  $1 \leq m \leq M - 1$ , let  $\eta_m = \sum_{s \in P(m)} \nu_s$ , where  $P(m)$  is the set of all prefixes of the binary representation of  $m$ .  
 Let  $t_0 = 1, t_1 = \lfloor \frac{n}{M} + \eta_1 \rfloor, \dots, t_m = \lfloor \frac{m \cdot n}{M} + \eta_m \rfloor, \dots, t_M = n + 1$ .  
 Let  $x^{(m)} = (x_{t_{m-1}}, \dots, x_{t_m - 1})$  for all  $m \in [M]$ .

Figure 10: `Partition`

The proof of Theorem 5.5 relies on the following two claims about the `Partition` subroutine, both of which are implicit in the work of Bun et al. [2015, Appendix C] and are based on ideas of Dwork et al. [2010a]. Claim 5.12 shows that for neighboring databases  $x \sim x'$ , the behaviors of the `Partition` subroutine on  $x$  and  $x'$  are “similar” the following sense: for any fixed partition of  $x$ , one is roughly as likely (over the randomness of the partition algorithm) to obtain a partition of  $x'$  that differs on at most two chunks, where the different chunks themselves differ only up to the addition or removal of a single item. This will allow us to

show that running  $M$  parallel copies of the AIP algorithm on the chunks remains roughly  $(\varepsilon, \delta)$ -differentially private. Claim 5.13 shows that, with high probability, each chunk is simultaneously large enough for the corresponding AIP algorithm to succeed, but also small enough so that treating all of the elements in a chunk as if they were the same element still permits us to get  $\alpha$ -accurate answers to arbitrary threshold queries.

**Claim 5.12.** Fix neighboring datasets  $x, x' \in [0, 1]^n$ . Then there exists a (measurable) bijection  $\varphi : \mathbb{R}^{2M} \rightarrow \mathbb{R}^{2M}$  with the following properties:

- (1) Let  $z \in \mathbb{R}^{2M}$  be any noise vector. Let  $x^{(1)}, \dots, x^{(M)}$  denote the partition of  $x$  obtained with random noise set to  $\nu = z$ . Similarly, let  $x'^{(1)}, \dots, x'^{(M)}$  denote the partition of  $x'$  obtained under noise  $\nu = \varphi(z)$ . Then there exist indices  $i_1, i_2$  such that: 1) For  $i \in \{i_1, i_2\}$ , the chunks  $x^{(i)}$  and  $x'^{(i)}$  differ up to the addition or removal of at most one item and 2) For every index  $i \notin \{i_1, i_2\}$ , we have  $x^{(i)} = x'^{(i)}$ .
- (2) For every noise vector  $z \in \mathbb{R}^{2M}$ , we have  $\mathbb{P}[\nu = \varphi(z)] \leq e^{2\varepsilon} \mathbb{P}[\nu = z]$ .

**Claim 5.13.** With probability at least  $1 - \beta/2$ , we have that  $|t_m - m \cdot n/M| \leq \alpha n/24$  for all  $m \in [M]$ .

Privacy of Algorithm 9. We first show how to use Claim 5.12 to show that Algorithm 9 is differentially private. Fix an adversary  $A$ , and let  $B = \text{Adaptive}_{A \equiv \text{AdaptiveThresholds}_T}$  simulate the interaction between  $A$  and Algorithm 9. Let  $S$  be a subset of the range of  $B$ . Then, by Property (1) of Claim 5.12 and group privacy, we have that for any  $z \in \mathbb{R}^{2M}$ :

$$\mathbb{P}[B(x) \in S | \nu = z] \leq e^{2\varepsilon} \mathbb{P}[B(x') \in S | \nu = \varphi(z)] + (1 + e^\varepsilon)\delta.$$

By Property (2) of Claim 5.12, we also have  $\Pr[\nu = z] \leq e^{2\varepsilon} \Pr[\nu = \varphi(z)]$  for every  $z \in \mathbb{R}^{2M}$ . Therefore,

$$\begin{aligned} \mathbb{P}[B(x) \in S] &= \int_{\mathbb{R}^{2M}} \mathbb{P}[B(x) \in S | \nu = z] \cdot \mathbb{P}[\nu = z] dz \\ &\leq \int_{\mathbb{R}^{2M}} \left( e^{2\varepsilon} \mathbb{P}[B(x') \in S | \nu = \varphi(z)] + (1 + e^\varepsilon)\delta \right) \cdot \mathbb{P}[\nu = z] dz \\ &\leq (1 + e^\varepsilon)\delta + \int_{\mathbb{R}^{2M}} e^{2\varepsilon} \mathbb{P}[B(x') \in S | \nu = \varphi(z)] \cdot e^{2\varepsilon} \mathbb{P}[\nu = \varphi(z)] dz \\ &\leq (1 + e^\varepsilon)\delta + e^{4\varepsilon} \mathbb{P}[B(x') \in S]. \end{aligned}$$

Hence,  $B$  is  $(e^{4\varepsilon}, (1 + e^\varepsilon)\delta)$ -differentially private, as claimed.

Accuracy of Algorithm 9. We now show how to use Claim 5.13 to show that Algorithm 9 produces  $(\alpha, \beta)$ -accurate answers. By a union bound, the following three events occur with probability at least  $1 - \beta$ :

- (1) For all  $m \in [M]$ ,  $|\frac{m}{M} - \frac{t_m}{n}| \leq \frac{\alpha}{6}$ .
- (2) Every chunk  $x^{(m)}$  has size  $|x^{(m)}| = t_m - t_{m-1} \in [\alpha n/6, 2\alpha n/3]$ .
- (3) Every instance of  $T$  succeeds.

Now we need to show that if these three events occur, we can produce  $\alpha$ -accurate answers to every threshold query  $c_{y_1}, \dots, c_{y_k}$ . Write the sorted input database as  $x_1 \leq x_2 \leq \dots \leq x_n$ . We consider two cases for the  $j^{\text{th}}$  query: As our first case, suppose  $x_n \leq y_j$ . Then for every

chunk  $x^{(m)}$ , we have  $\max\{x^{(m)}\} \leq y_j$ . Then the success condition of  $T^{(m)}$  guarantees that  $a_j^{(m)} = \mathbf{R}$ . Thus, the answer  $a_j = 1$  is (exactly) accurate for the query  $c_j$ .

As our second case, let  $i$  be the smallest index for which  $x_i > y_j$ , and suppose the item  $x_i$  is in some chunk  $x^{(m_i)}$ . Note that this means that the true answer to the query  $c_{y_j}$  is  $(i - 1)/n$  and that  $t_{m_i-1} \leq i \leq t_{m_i} - 1$ . Then again, for every  $m < m_i$  we have  $\max\{x^{(m)}\} \leq y_j$ , so every such  $T^{(m)}$  instance yields  $a_j^{(m)} = \mathbf{R}$ . Thus,

$$a_j = \frac{1}{M} \cdot \left| \left\{ m \in [M] : a_j^{(m)} = \mathbf{R} \right\} \right| \geq \frac{m_i - 1}{M} \geq \frac{t_{m_i} - \alpha}{n} - \frac{\alpha}{2} \geq \frac{(i - 1)}{n} - \alpha,$$

since  $M \geq 2/\alpha$ .

On the other hand, for every  $m > m_i$ , we have  $\min\{x^{(m)}\} > y_j$ , so every such  $T^{(m)}$  instance instead yields  $a_j^{(m)} = \mathbf{L}$ .

$$a_j \leq \frac{m_i}{M} \leq \frac{t_{m_i}}{n} + \frac{\alpha}{6} \leq \frac{t_{m_i-1} + 2\alpha n/3}{n} + \frac{\alpha}{6} \leq \frac{i}{n} + \frac{2\alpha}{3} + \frac{\alpha}{6} \leq \frac{i - 1}{n} + \alpha,$$

since  $n \geq 6/\alpha$ .

□

**Acknowledgements.** We thank Salil Vadhan for many helpful discussions. We also thank the anonymous reviewers for suggestions on the presentation of this work.

## REFERENCES

- R. Bassily, K. Nissim, A. D. Smith, T. Steinke, U. Stemmer, and J. Ullman. Algorithmic stability for adaptive data analysis. In *Proceedings of the Forty-Eighth Annual ACM on Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, 2016*.
- A. Beimel, K. Nissim, and U. Stemmer. Private learning and sanitization: Pure vs. approximate differential privacy. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques - 16th International Workshop, APPROX 2013, and 17th International Workshop, RANDOM 2013, Berkeley, CA, USA, August 21-23, 2013. Proceedings*, pages 363–378, 2013.
- A. Blum, K. Ligett, and A. Roth. A learning theory approach to noninteractive database privacy. *J. ACM*, 60(2):12, 2013. doi: 10.1145/2450142.2450148. URL <http://doi.acm.org/10.1145/2450142.2450148>.
- D. Boneh and J. Shaw. Collusion-secure fingerprinting for digital data. *IEEE Trans. Information Theory*, 44(5):1897–1905, 1998.
- M. Bun, J. Ullman, and S. P. Vadhan. Fingerprinting codes and the price of approximate differential privacy. In *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, pages 1–10, 2014.
- M. Bun, K. Nissim, U. Stemmer, and S. P. Vadhan. Differentially private release and learning of threshold functions. In *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015*, pages 634–649, 2015.
- T. H. Chan, E. Shi, and D. Song. Private and continual release of statistics. *ACM Trans. Inf. Syst. Secur.*, 14(3):26, 2011.

- I. Dinur and K. Nissim. Revealing information while preserving privacy. In *Proceedings of the Twenty-Second ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, June 9-12, 2003, San Diego, CA, USA*, pages 202–210, 2003.
- C. Dwork and J. Lei. Differential privacy and robust statistics. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31 - June 2, 2009*, pages 371–380, 2009.
- C. Dwork and A. Roth. The algorithmic foundations of differential privacy. *Found. Trends Theor. Comput. Sci.*, 9(3–4):211–407, Aug. 2014. ISSN 1551-305X. doi: 10.1561/04000000042. URL <http://dx.doi.org/10.1561/04000000042>.
- C. Dwork and S. Yekhanin. New efficient attacks on statistical disclosure control mechanisms. In *Advances in Cryptology - CRYPTO 2008, 28th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 17-21, 2008. Proceedings*, pages 469–480, 2008.
- C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography, Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006, Proceedings*, pages 265–284, 2006.
- C. Dwork, M. Naor, O. Reingold, G. N. Rothblum, and S. P. Vadhan. On the complexity of differentially private data release: efficient algorithms and hardness results. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31 - June 2, 2009*, pages 381–390, 2009.
- C. Dwork, M. Naor, T. Pitassi, and G. N. Rothblum. Differential privacy under continual observation. In *Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010, Cambridge, Massachusetts, USA, 5-8 June 2010*, pages 715–724, 2010a.
- C. Dwork, G. N. Rothblum, and S. P. Vadhan. Boosting and differential privacy. In *IEEE Symposium on Foundations of Computer Science (FOCS '10)*, pages 51–60. IEEE, 23–26 October 2010b.
- C. Dwork, V. Feldman, M. Hardt, T. Pitassi, O. Reingold, and A. L. Roth. Preserving statistical validity in adaptive data analysis. In *Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing, STOC 2015, Portland, OR, USA, June 14-17, 2015*, pages 117–126, 2015a.
- C. Dwork, M. Naor, O. Reingold, and G. N. Rothblum. Pure differential privacy for rectangle queries via private partitions. In *Advances in Cryptology - ASIACRYPT 2015 - 21st International Conference on the Theory and Application of Cryptology and Information Security, Auckland, New Zealand, November 29 - December 3, 2015, Proceedings, Part II*, pages 735–751, 2015b.
- C. Dwork, A. D. Smith, T. Steinke, J. Ullman, and S. P. Vadhan. Robust traceability from trace amounts. In *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015*, pages 650–669, 2015c. doi: 10.1109/FOCS.2015.46. URL <http://dx.doi.org/10.1109/FOCS.2015.46>.
- M. Hardt and G. N. Rothblum. A multiplicative weights mechanism for privacy-preserving data analysis. In *51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010, October 23-26, 2010, Las Vegas, Nevada, USA*, pages 61–70, 2010.
- M. Hardt and K. Talwar. On the geometry of differential privacy. In *Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010, Cambridge, Massachusetts, USA, 5-8 June 2010*, pages 705–714, 2010.
- M. Hardt and J. Ullman. Preventing false discovery in interactive data analysis is hard. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014*,



- Philadelphia, PA, USA, October 18-21, 2014, pages 454–463, 2014.
- A. Roth and T. Roughgarden. Interactive privacy via the median mechanism. In *STOC*, pages 765–774. ACM, June 5–8 2010.
- T. Steinke and J. Ullman. Interactive fingerprinting codes and the hardness of preventing false discovery. In *Proceedings of The 28th Conference on Learning Theory, COLT 2015, Paris, France, July 3-6, 2015*, pages 1588–1628, 2015a.
- T. Steinke and J. Ullman. Between pure and approximate differential privacy. *CoRR*, abs/1501.06095, 2015b. URL <http://arxiv.org/abs/1501.06095>.
- G. Tardos. Optimal probabilistic fingerprint codes. *J. ACM*, 55(2), 2008.
- J. Ullman. Answering  $n^{2+o(1)}$  counting queries with differential privacy is hard. In *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013*, pages 361–370, 2013.
- S. L. Warner. Randomized response: A survey technique for eliminating evasive answer bias. *Journal of the American Statistical Association*, 60(309):63–69, 1965.

#### APPENDIX A. THE FINGERPRINTING LEMMA

In this section we prove the fingerprinting lemma (Lemma 3.1). The proof is broken into several lemmata.

**Lemma A.1.** *Let  $f : \{\pm 1\}^n \rightarrow \mathbb{R}$ . Define  $g : [\pm 1] \rightarrow \mathbb{R}$  by*

$$g(p) = \mathbb{E}_{x_1 \dots x_n \sim p} [f(x)].$$

*Then*

$$\mathbb{E}_{x_1 \dots x_n \sim p} \left[ f(x) \cdot \sum_{i \in [n]} (x_i - p) \right] = g'(p) \cdot (1 - p^2).$$

A rescaling of this lemma appears in Steinke and Ullman [2015a]. The following proof is taken from Dwork et al. [2015c].

*Proof.* We begin by establishing several identities.

Since  $x^2 = 1$  for  $x \in \{\pm 1\}$ , we have the identity

$$\frac{d}{dp} \frac{1 + xp}{2} = \frac{x}{2} = \frac{1 + xp}{2} \frac{x - p}{1 - p^2}$$

for all  $x \in \{\pm 1\}$  and  $p \in (-1, 1)$ . By the product rule, we have

$$\frac{d}{dp} \prod_{i \in [n]} \frac{1 + x_i p}{2} = \sum_{i \in [n]} \left( \frac{d}{dp} \frac{1 + x_i p}{2} \right) \prod_{k \in [n] \setminus \{i\}} \frac{1 + x_k p}{2} = \sum_{i \in [n]} \frac{x_i - p}{1 - p^2} \prod_{k \in [n]} \frac{1 + x_k p}{2}$$

for all  $x \in \{\pm 1\}^n$  and  $p \in (-1, 1)$ .

Sampling  $x \sim p$  samples each  $x \in \{\pm 1\}$  with probability  $\frac{1+xp}{2}$ . Thus sampling  $x_1 \dots x_n \sim p$ , samples each  $x \in \{\pm 1\}^n$  with probability  $\prod_{i \in [n]} \frac{1+x_i p}{2}$ .

Now we can write

$$g(p) = \mathbb{E}_{x_1 \dots x_n \sim p} [f(x)] = \sum_{x \in \{\pm 1\}^n} f(x) \prod_{i \in [n]} \frac{1 + x_i p}{2}.$$

Using the above identities gives

$$\begin{aligned}
g'(p) &= \sum_{x \in \{\pm 1\}^n} f(x) \frac{d}{dp} \prod_{i \in [n]} \frac{1 + x_i p}{2} \\
&= \sum_{x \in \{\pm 1\}^n} f(x) \sum_{i \in [n]} \frac{x_i - p}{1 - p^2} \prod_{k \in [n]} \frac{1 + x_k p}{2} \\
&= \mathbb{E}_{x_1 \dots x_n \sim p} \left[ f(x) \sum_{i \in [n]} \frac{x_i - p}{1 - p^2} \right]
\end{aligned}$$

□

**Lemma A.2.** *Let  $g : [\pm 1] \rightarrow \mathbb{R}$  be a polynomial. Then*

$$\mathbb{E}_{p \in [\pm 1]} [g'(p) \cdot (1 - p^2)] = 2 \mathbb{E}_{p \in [\pm 1]} [g(p) \cdot p].$$

*Proof.* Let  $u(p) = 1 - p^2$ . By integration by parts and the fundamental theorem of calculus,

$$\begin{aligned}
\mathbb{E}_{p \in [\pm 1]} [g'(p) \cdot (1 - p^2)] &= \frac{1}{2} \int_{-1}^1 g'(p)(1 - p^2) dp \\
&= \frac{1}{2} \int_{-1}^1 g'(p)u(p) dp \\
&= \frac{1}{2} \int_{-1}^1 \left( \frac{d}{dp} g(p)u(p) \right) - g(p)u'(p) dp \\
&= \frac{1}{2} (g(1)u(1) - g(-1)u(-1)) - \frac{1}{2} \int_{-1}^1 g(p)(-2p) dp \\
&= 0 + \int_{-1}^1 g(p)p dp \\
&= 2 \mathbb{E}_{p \in [\pm 1]} [g(p) \cdot p].
\end{aligned}$$

□

**Proposition A.3.** *Let  $f : \{\pm 1\}^n \rightarrow \mathbb{R}$ . Then*

$$\mathbb{E}_{p \in [\pm 1], x_1 \dots x_n \sim p} \left[ f(x) \cdot \sum_{i \in [n]} (x_i - p) + (f(x) - \bar{x})^2 \right] \geq \frac{1}{3}.$$

*Proof.* Define  $g : [\pm 1] \rightarrow \mathbb{R}$  by

$$g(p) = \mathbb{E}_{x_1 \dots x_n \sim p} [f(x)].$$

By Lemmas A.1 and A.2,

$$\mathbb{E}_{p \in [\pm 1], x_1 \dots x_n \sim p} \left[ f(x) \cdot \sum_{i \in [n]} (x_i - p) \right] = \mathbb{E}_{p \in [\pm 1]} [g'(p)(1 - p^2)] = \mathbb{E}_{p \in [\pm 1]} [2g(p)p].$$

Moreover, by Jensen's inequality,

$$\begin{aligned}
 \mathbb{E}_{p \in [\pm 1], x_1 \dots x_n \sim p} [(f(x) - \bar{x})^2] &\geq \mathbb{E}_{p \in [\pm 1]} \left[ \left( \mathbb{E}_{x_1 \dots x_n \sim p} [f(x) - \bar{x}] \right)^2 \right] \\
 &= \mathbb{E}_{p \in [\pm 1]} [(g(p) - p)^2] \\
 &= \mathbb{E}_{p \in [\pm 1]} [g(p)^2 - 2g(p)p + p^2] \\
 &= \mathbb{E}_{p \in [\pm 1]} [g(p)^2] - \mathbb{E}_{p \in [\pm 1], x_1 \dots x_n \sim p} \left[ f(x) \cdot \sum_{i \in [n]} (x_i - p) \right] + \frac{1}{3}.
 \end{aligned}$$

Rearranging yields the result:

$$\mathbb{E}_{p \in [\pm 1], x_1 \dots x_n \sim p} \left[ f(x) \cdot \sum_{i \in [n]} (x_i - p) + (f(x) - \bar{x})^2 \right] \geq \mathbb{E}_{p \in [\pm 1]} [g(p)^2] + \frac{1}{3} \geq \frac{1}{3}.$$

□

We also have an alternative version of Proposition A.3:

**Proposition A.4.** *Let  $f : \{\pm 1\}^n \rightarrow \mathbb{R}$ . Then*

$$\mathbb{E}_{p \in [\pm 1], x_1 \dots x_n \sim p} \left[ f(x) \cdot \sum_{i \in [n]} (x_i - p) + (f(x) - p)^2 \right] \geq \frac{1}{3}.$$

*Proof.* Define  $g : [\pm 1] \rightarrow \mathbb{R}$  by

$$g(p) = \mathbb{E}_{x_1 \dots x_n \sim p} [f(x)].$$

By Lemmas A.1 and A.2,

$$\mathbb{E}_{p \in [\pm 1], x_1 \dots x_n \sim p} \left[ f(x) \cdot \sum_{i \in [n]} (x_i - p) \right] = \mathbb{E}_{p \in [\pm 1]} [g'(p)(1 - p^2)] = \mathbb{E}_{p \in [\pm 1]} [2g(p)p].$$

Moreover,

$$\mathbb{E}_{p \in [\pm 1], x_1 \dots x_n \sim p} [(f(x) - p)^2] = \mathbb{E}_{p \in [\pm 1], x_1 \dots x_n \sim p} [f(x)^2 - 2g(p)p + p^2] \geq 0 - \mathbb{E}_{p \in [\pm 1]} [2g(p)p] + \frac{1}{3}.$$

The result follows by combining the above equality and inequality. □

Finally we restate and prove Lemma 3.1

**Lemma A.5** (Fingerprinting Lemma). *Let  $f : \{\pm 1\}^n \rightarrow [\pm 1]$ . Then*

$$\mathbb{E}_{p \in [\pm 1], x_1 \dots x_n \sim p} \left[ f(x) \cdot \sum_{i \in [n]} (x_i - p) + 2|f(x) - \bar{x}| \right] \geq \frac{1}{3}.$$

*Proof.* Since  $|f(x) - \bar{x}| \leq 2$ , we have  $|f(x) - \bar{x}|^2 \leq 2|f(x) - \bar{x}|$ . The result thus follows from Proposition A.3. □