Proceedings of Machine Learning Research vol 247:1–36, 2024 37th Annual Conference on Learning Theory

Sample-Optimal Locally Private Hypothesis Selection and the Provable Benefits of Interactivity

Alireza F. Pour ALIREZA.FATHOLLAHPOUR@UWATERLOO.CA
Cheriton School of Computer Science, University of Waterloo

Hassan Ashtiani ZOKAEIAM@MCMASTER.CA
Department of Computing and Software, McMaster University

Shahab Asoodeh ASOODEH@MCMASTER.CA
Department of Computing and Software, McMaster University

Editors: Shipra Agrawal and Aaron Roth

Abstract

We study the problem of hypothesis selection under the constraint of local differential privacy. Given a class $\mathcal{F}$ of k distributions and a set of i.i.d. samples from an unknown distribution h, the goal of hypothesis selection is to pick a distribution $\hat{f}$ whose total variation distance to h is comparable with the best distribution in $\mathcal{F}$ (with high probability). We devise an ε-locally-differentially-private (ε-LDP) algorithm that uses $\Theta\left(\frac{k}{\alpha^2 \min\{\varepsilon^2,1\}}\right)$ samples to guarantee that $d_{TV}(h,\hat{f}) \leq \alpha + 9 \min_{f\in\mathcal{F}} d_{TV}(h,f)$ with high probability. This sample complexity is optimal for $\varepsilon < 1$, matching the lower bound of Gopi et al. (2020). All previously known algorithms for this problem required $\Omega\left(\frac{k \log k}{\alpha^2 \min\{\varepsilon^2,1\}}\right)$ samples to work.

Moreover, our result demonstrates the power of interaction for ε-LDP hypothesis selection. Namely, it breaks the known lower bound of $\Omega\left(\frac{k\log k}{\alpha^2\varepsilon^2}\right)$ for the sample complexity of non-interactive hypothesis selection. Our algorithm achieves this using only $\Theta(\log\log k)$ rounds of interaction.

To prove our results, we define the notion of *critical queries* for a Statistical Query Algorithm (SQA) which may be of independent interest. Informally, an SQA is said to use a small number of critical queries if its success relies on the accuracy of only a small number of queries it asks. We then design an LDP algorithm that uses a smaller number of critical queries.

Keywords: hypothesis selection, local differential privacy, statistical query model

1. Introduction

One of the basic problems in statistical learning is hypothesis selection where we are given a set of i.i.d. samples from an unknown distribution h and a set $\mathcal{F}$ of k distributions. The goal is to select a distribution in $\mathcal{F}$ whose total variation distance to h is close to the minimum total variation distance between h and any distribution in $\mathcal{F}$. This problem has been studied extensively in the literature (Yatracos, 1985; Mahalanabis and Stefankovic, 2007; Bousquet et al., 2019, 2022; Aliakbarpour et al., 2024) and it is known that the number of samples required to solve this problem has a tight logarithmic dependency on k for a general class of distributions.

In many practical statistical estimation scenarios, data points contain sensitive information such as medical and financial records. This urges the study of hypothesis selection under differential privacy (DP) (Dwork et al., 2006b) —the de-facto privacy standard in machine learning. Central and local models are two common settings of DP. In the central model, a learning algorithm is differentially private if its output, given full access to a database, does not significantly change with a small perturbation in the input database (e.g., by changing one entry in the database). In the local differential privacy (LDP) model (Warner, 1965; Kasiviswanathan et al., 2011; Duchi et al., 2013), however, the algorithm cannot access the database directly. Instead, it receives a privatized (e.g., randomized) version of each data point through privacy-preserving mechanisms. In this paper, we work with the notion of ε-LDP; see Section 5 for a more formal definition. In fact, the local model of privacy ensures a *user-level* privacy and it has been the preferred model for implementation in industry, e.g., by Google, Apple, and Microsoft (Erlingsson et al., 2014; Differential Privacy Team, December 2017; Ding et al., 2017; Davidson et al., May 9 2017).

It is important to distinguish between non-adaptive and adaptive locally private algorithms. A non-adaptive algorithm chooses a privacy mechanism for each data point independently. On the other hand, an adaptive algorithm runs in multiple rounds and chooses each mechanism at each round based on the outcomes of mechanisms in the previous rounds (Duchi et al., 2018; Duchi and Rogers, 2019; Joseph et al., 2019b). This *(sequential) interactivity* allows the algorithm to be more flexible in hiding private information and has led to a smaller sample complexity in several problems (Han et al., 2018; Joseph et al., 2022; Acharya et al., 2020, 2022). Nevertheless, the cost of interactions can become a bottleneck in designing adaptive algorithms (Kairouz et al., 2021) and therefore designing algorithms that use a small number of adaptive rounds is crucial.

More recently, hypothesis selection has been studied under the constraint of differential privacy. The sample complexity of hypothesis selection is well understood in the central DP model (Bun et al., 2019a; Aden-Ali et al., 2021a), and it has a tight logarithmic dependence on k. In the local model, however, there is still a gap between the best known upper and lower bounds for hypothesis selection as well as several other statistical learning tasks (Duchi et al., 2018; Duchi and Rogers, 2019; Joseph et al., 2019a; Bun et al., 2019b; Acharya et al., 2020; Chen et al., 2020; Edmonds et al., 2020; Gopi et al., 2020; Acharya et al., 2022; Asi et al., 2022, 2023). In this work, we fill this gap for the problem of hypothesis selection and characterize the optimal sample complexity in the local model. In particular, we propose a new iterative algorithm and a novel analysis technique, which together establish a *linear* sample complexity in k for LDP hypothesis selection. Before stating our results in detail, it is useful to define the problem of hypothesis selection formally.

1.1. Hypothesis selection

A hypothesis selector $\mathcal{A}$ is a randomized function that receives a set of samples S, the description of a class of distribution $\mathcal{F}$, accuracy and failure parameters α and β, and returns a distribution in $\mathcal{F}$. Next, we give the formal definition of the hypothesis selection problem.

Definition 1 (Hypothesis Selection) *We say a hypothesis selector $\mathcal{A}$ can do hypothesis selection with $m : \mathbb{N} \times (0,1)^2 \to \mathbb{N}$ samples if the following holds: for any unknown distribution h, any class $\mathcal{F}$ of $k \in \mathbb{N}$ distributions, and $\alpha, \beta > 0$, if S is a set of at least $m(k,\alpha,\beta)$ i.i.d. samples from h, then with probability at least $1-\beta$ (over the randomness of S and $\mathcal{A}$) $\mathcal{A}(\mathcal{F}, S, \alpha, \beta)$ returns $\hat{f} \in \mathcal{F}$ such that $d_{TV}(h, \hat{f}) \leq C \min_{f\in\mathcal{F}} d_{TV}(h, f) + \alpha$, where $C \geq 1$ is a universal constant and is called the approximation factor of $\mathcal{A}$. The sample complexity of hypothesis selection is the minimum $m(k,\alpha,\beta)$ among all hypothesis selectors.*

Note that C is a multiplicative *approximation factor* and is typically a small constant. A smaller value for C signifies a more accurate estimate $\hat{f}$. The sample complexity of hypothesis selection is proportional to $\Theta\left(\frac{\log k}{\alpha^2}\right)$ as a function of k and α (see Devroye and Lugosi (2001) for a detailed discussion). Other work has studied hypothesis selection with additional considerations, including computational efficiency, robustness, and more (Yatracos, 1985; Mahalanabis and Stefankovic, 2007; Acharya et al., 2014, 2018; Bousquet et al., 2019, 2022). Recently, Bun et al. (2019a) showed that, similar to the non-private setting, one can perform hypothesis selection under the constraint of central differential privacy with a logarithmic number of samples in k. However, the same statement does not hold in the local differential privacy setting. In fact, Gopi et al. (2020), building on a result of Duchi and Rogers (2019) and Braverman et al. (2016) showed that the sample complexity of locally private hypothesis selection problem scales *at least linearly* in k.

Theorem 2 (Informal, Theorem 1.2 of Gopi et al. (2020), Corollary 6 of Duchi and Rogers (2019)) *There exists a family of k distributions for which any (interactive) ε-LDP selection method requires at least $\Omega\left(\frac{k}{\alpha^2 \min\{\varepsilon,\varepsilon^2\}}\right)$ samples to learn it.*

A basic idea to perform LDP hypothesis selection is using the classical tournament-based (round-robin) hypothesis selection method of Devroye and Lugosi (2001). While it is straightforward to come up with a locally private version of this approach, it would require $\Omega(k^2)$ samples as it compares all pairs of hypotheses. Gopi et al. (2020) improved over this baseline and proposed a multi-round LDP hypothesis selection algorithm whose sample complexity scales as $O(k \log k \log\log k)$.

Theorem 3 (Informal, Corollary 5.10 of Gopi et al. (2020)) *There exists an algorithm with failure probability of $1/10$, that solves the problem of hypothesis selection under the constraint of ε-local differential privacy (for $\varepsilon \in (0,1)$) with $O(\log\log k)$ rounds, with approximation factor of 27, and using $O\left(\frac{k \log k \log\log k}{\alpha^2 \min\{\varepsilon^2,1\}}\right)$ samples.*

The algorithm of Gopi et al. (2020) is also a tournament-based method. However, unlike the round-robin approach, it only makes $O(k \log\log k)$ comparisons between the hypotheses. This is achieved by comparing the distributions *adaptively* in $\log\log k$ LDP rounds. A union bound argument is then required to make sure that every comparison is accurate, resulting in the $O(k \log k \log\log k)$ term in the sample complexity. Given the gap between the upper bound of Theorem 3 and the lower bound of Theorem 2 in Gopi et al. (2020), the following question remains open:

> Is it possible to perform hypothesis selection in the LDP model using $O(k)$ samples?
> If so, can it still be done in $O(\log\log k)$ adaptive LDP rounds?

One barrier in answering the above question is the sample complexity lower bound of Gopi et al. (2020) for *non-interactive* hypothesis selection which is based on the work of Ullman (2018).

Theorem 4 (Informal, Theorem 3.3 of Gopi et al. (2020)) *There is a family of k distributions for which any non-interactive ε-LDP hypothesis selection method requires at least $\Omega(\frac{k \log k}{\alpha^2 \varepsilon^2})$ samples.*

In other words, we have to use interactivity to break the $O(k \log k)$ barrier on the sample complexity and achieve an $O(k)$ upper bound. This raises the following question:

> Does interactivity offer a provable sample-complexity benefit for locally private hypothesis selection?

In this work, we show that the answers to the above questions are affirmative and, indeed, a small number of interactions helps us to achieve a linear sample complexity dependence on k.

1.2. Results and discussion

Our main result shows that locally private hypothesis selection can be solved with a sample complexity that is linear in k. To achieve this, we both design a *new algorithm* and propose *a new analysis technique*.

Theorem 5 (Informal Version of Theorem 23) *There exists an ε-LDP algorithm that solves the problem of hypothesis selection in $O(\log\log k)$ rounds, with an approximation factor of 9, and uses $O\left(\frac{k(\log 1/\beta)^2}{\alpha^2 \min\{\varepsilon^2,1\}}\right)$ samples.*

Note that, unlike the existing upper bound given in Theorem 3, this result works for all values of β; a detailed discussion on this will follow. The following corollary, which is a direct consequence of the above theorem and Theorem 2, yields the optimal sample complexity.

Corollary 6 (Sample Complexity of LDP Hypothesis Selection) *Let $\varepsilon \in (0,1)$. For any constant $\beta \in (0,1)$, the sample complexity of (interactive) ε-LDP hypothesis selection is $\Theta\left(\frac{k}{\alpha^2\varepsilon^2}\right)$.*

We now highlight some important aspects of our results and contributions.

Optimal dependence on k. The private hypothesis selector proposed in Gopi et al. (2020) relies on $\Theta(\log\log k)$ adaptive rounds and requires $O(k \log\log k)$ pairwise comparisons. Applying the union bound then yields the sub-optimal sample complexity that scales with $k \log k \log\log k$. On the other hand, we develop a new analysis technique that avoids the simple union bound over all comparisons. Instead, it exploits the fact that not all the comparisons are *critical* for establishing the guarantees of an algorithm (See Section 3). We also propose a new algorithm that runs in the same number of rounds as the algorithm in Gopi et al. (2020) but uses only $O(k)$ queries, from which only $O(k/\log k)$ are critical queries (See Section 4).

Approximation factor. Theorem 5 solves the hypothesis selection with an approximation factor of 9, compared to the factor of 27 in Theorem 3. We achieve this by using a variant of Minimum Distance Estimate (Mahalanabis and Stefankovic, 2007) as the final sub-routine in our algorithm (see Section 4).

High probability bound. A drawback of the upper bound in Gopi et al. (2020) is its loose dependence on the failure parameter β. In fact, Theorem 3 is stated only for $\beta = 1/10$. In contrast, our result holds for any value of $\beta > 0$ with a mild cost of $(\log 1/\beta)^2$ in the sample complexity. We achieve this by boosting the success probability of our algorithm in various steps[1].

Computational complexity. Similar to Gopi et al. (2020), our algorithm runs in linear time (as a function of the number of samples) assuming an oracle access to the Scheffe sets of distributions.

Number of rounds. A main feature of the algorithm of Gopi et al. (2020) is that it runs in only $\Theta(\log\log k)$ adaptive rounds. This sets it apart from similar approaches (Acharya et al., 2014, 2018) that require $\Theta(\log k)$ adaptive rounds. Our method also runs in $\Theta(\log\log k)$ rounds. Moreover, our result demonstrates the power of interaction: any non-adaptive LDP method for hypothesis selection requires at least $\Omega(k\log k)$ samples (Ullman, 2018; Gopi et al., 2020), while our algorithm works with $O(k)$ samples. It remains open whether a sample-optimal LDP algorithm can run in $o(\log\log k)$ rounds.

1. It is possible extend the analysis of Gopi et al. (2020) to make it work for any $\beta \in (0,1)$ but with with a substantial $(1/\beta)^2$ cost in the sample complexity (as opposed to the poly-logarithmic dependence in our bound).

Hypothesis selection method	App. factor	#Queries	#Rounds	#Samples for LDP
Round-Robin (Devroye and Lugosi, 2001)	9	$O(k^2)$	1	$O(k^2 \log k)$
Comb (Acharya et al., 2014)	9	$O(k)$	$O(\log k)$	$O(k \log k)$
Gopi et al. (2020)	27	$O(k \log \log k)$	$O(\log \log k)$	$O(k \log k \log \log k)$
Gopi et al. (2020), $\forall t \in [k]$	9^t	$O(k^{1+\frac{1}{2^t-1}} t)$	t	$O(k^{1+\frac{1}{2^t-1}} t \log k)$
MDE-Variant (Mahalanabis and Stefankovic, 2007)	3	$O(k^2)$	1	$O(k^2 \log k)$
BOKSERR [THIS WORK]	9	$O(k)$	$O(\log \log k)$	$O(k)$

Table 1: Hypothesis selection methods, their approximation factors, number of statistical queries they ask, number of rounds and samples required to implement them in the LDP model.

The statistical query viewpoint. Our analysis is presented in the more general Statistical Query (SQ) model (Kearns, 1998). SQ algorithms can be readily implemented in the LDP model (Kasiviswanathan et al., 2011); see Section 2.2. Moreover, most (if not all) existing hypothesis selection methods are SQ algorithms (and therefore can be implemented in the LDP model). We provide a summary of existing hypothesis selection methods and the cost of implementing them in the LDP model in Table 1.

2. Preliminaries

Notation. We denote by $[N]$ the set of numbers $\{1, 2, \ldots, N\}$. The total variation distance between two probability densities f and g over domain $\mathcal{X}$ is defined by $d_{TV}(f, g) = \frac{1}{2} \int_{\mathcal{X}} |f(x) - g(x)|\, dx = \frac{1}{2}\|f - g\|_1$. We define the total variation distance between a class of distributions $\mathcal{F}$ and a distribution h as $d_{TV}(h, \mathcal{F}) = \inf_{f \in \mathcal{F}} d_{TV}(h, f)$. We abuse the notation and use f both as a probability measure (e.g., in $f[B]$ where $B \subset \mathcal{X}$ is a measurable set) and as its corresponding probability density function (e.g., in $f(x)$ where $x \in \mathcal{X}$).

2.1. The statistical query viewpoint of hypothesis selection

Assume h is an (unknown) probability distribution over domain $\mathcal{X}$. In a typical scenario, the learning algorithm is assumed to have direct access to random samples from h. Many learning algorithms, however, can be implemented in the more limited *Statistical Query (SQ)* model (Kearns, 1998). Namely, instead of accessing random samples, the learning algorithm chooses a (measurable) function $q : \mathcal{X} \to \{0, 1\}$ and receives an estimate of $\mathbb{E}_{x \sim h}[q(x)]$.

The hypothesis selection algorithms that we study in this paper are based on pair-wise comparisons between distributions. In Section 2.3, we show that such a comparison can be executed by a statistical query. Therefore, it is helpful to view these hypothesis selection algorithms as Statistical Query Algorithms (SQAs). We define a Statistical Query Oracle (SQO) and a SQA in below.

Definition 7 (Statistical Query Oracle) *Let $(\mathcal{X}, \mathcal{B}, h)$ be a probability space. A Statistical Query Oracle (SQO) for h is a random function $\mathcal{O}_h$ with the following property: for any $\alpha, \beta \in (0, 1)$, and* *a finite workload* $W = (W_i)_{i=1}^n$ *where* $W_i \in \mathcal{B}$*, the oracle outputs* n *real values such that*

$$\mathbb{P}\left[\sup_{i\in[n]} \left|\mathcal{O}_h\left(W,\alpha,\beta\right)_i - \mathbb{E}_{x\sim h}\left[1\left\{x \in W_i\right\}\right]\right| \geq \alpha\right] \leq \beta,$$

where $\mathcal{O}_h\left(W,\alpha,\beta\right)_i$ *is the* i*-th output of* $\mathcal{O}_h$ *and the probability is over the randomness of oracle.*

Based on the above definition, we define the statistical query algorithm that runs in t rounds and calls a SQO in each round adaptively (i.e., based on the outcomes of previous rounds).

Definition 8 (Statistical Query Algorithm) *We say* $\mathcal{A}$ *is a Statistical Query Algorithm (SQA) with* t *rounds if it returns its output by making* t *(adaptive) calls to a statistical query oracle* $\mathcal{O}_h$*, where in each call* $1 < i \leq t$*, the workload* $W^{(i)}$ *of queries may depend on the output of* $\mathcal{O}_h$ *in the previous round* $i-1$ *rounds.*

2.2. Locally private hypothesis selection

In the local model of DP, each data point goes through a private mechanism (called a local randomizer) and only the privatized outcomes are used by the algorithm. A canonical example of a local randomizer is the composition of the randomized response mechanism with a binary function. The randomized response is a mechanism that gets as input a bit and outputs it with probability p or flips it with probability $1-p$. An LDP algorithm operates in multiple rounds. In each round, the algorithm chooses a series of local randomizers based on the information that it has received in the previous rounds. Each selected randomizer then generates a noisy version of a new data point for the use of algorithm. The local privacy model is defined more formally in Section 5.

The hypothesis selection algorithms discussed in this paper all fit into the SQA framework described in Definition 8. We show in Section 5 that any SQA can be implemented within the local privacy model using the same number of rounds. We also discuss how many samples are required to simulate a SQO with a single-round LDP protocol. It turns out that in general $\Theta(\frac{n\log n}{\alpha^2\min\{\varepsilon^2,1\}})$ samples are sufficient to simulate a SQO that answers a workload of n queries with accuracy α and failure $\beta = 2/3$ using a single-round LDP method. This allows us to simplify the presentation of the results by focusing on the query complexity of the hypothesis selection approaches.

We find it useful to give an informal discussion on how a SQO can be implemented with a single round LDP algorithm. A natural way of answering a SQ is to return the empirical average of the query values on a set of i.i.d. samples. However, to satisfy the privacy constraint, the algorithm cannot use the query values on the actual samples. Instead, the algorithm can observe the private version of these values through a (de-biased) randomized response mechanism. An application of Hoeffding's inequality implies that $\Theta(\frac{\log 1/\beta}{\alpha^2\min\{\varepsilon^2,1\}})$ samples are sufficient to make sure that the empirical average of the randommized response values is α-close to the true value of query with probability at least $1-\beta$. To ensure α accuracy for all queries with probability at least $1-\beta$, one needs to resort to a union-bound argument over all n queries. Therefore, each query needs to be estimated with a higher confidence of $1-\beta/n$ (i.e., failure of β/n), which results in the factor of $O(\log n)$ in sample complexity. It is important to note that, due to the privacy constraint, we need to use fresh samples for answering each query (otherwise, we incur a larger privacy budget). Hence, the sample complexity of this approach is $\Theta(\frac{n\log n/\beta}{\alpha^2\min\{\varepsilon^2,1\}})$. In Section 3, we introduce the notion of *critical* queries and show how one can avoid the typical union bound analysis and, thus, the sub-optimal logarithmic term in the sample complexity. In Section 5, we expand on this and show how to privately implement any SQA.

2.3. Existing tournament-based approaches for hypothesis selection

In this section, we review some of the classical algorithms for hypothesis selection, and how they can be implemented as SQAs. As discussed in the previous section, any SQA can be implemented in the LDP setting too. Some of these algorithms will be used later as building blocks of our method. We relegate the pseudo-codes for these algorithms to Appendix G.

2.3.1. CLASSES OF TWO DISTRIBUTIONS

We begin with the simplest case where $\mathcal{F} = \{f_1, f_2\}$ and the classical Shceffé test (Scheffé, 1947; Devroye and Lugosi, 2001). This algorithm is a building block of other algorithms in the sequel. Let us first define the Scheffé set between two distributions.

Definition 9 (Scheffé Set) *The Scheffé set between an ordered pair of distributions (f, g) is defined as $Sch(f, g) = \{x \in \mathcal{X} : f(x) > g(x)\}$.*

Let h be the (unknown) probability measure that generates the samples. The Shceffé test, delineated in Algorithm 4, chooses between f_1 and f_2 by estimating $h[Sch(f_1, f_2)]$. Note that $h[Sch(f_1, f_2)]$ can be estimated using a single statistical query. The following folklore theorem states the formal guarantee of the Shceffé test.

Theorem 10 (Analysis of Shceffé Test) *Let $\mathcal{F} = \{f_1, f_2\}$, $\alpha, \beta \in (0, 1)$ and $\mathcal{O}_h$ be a SQO for the (unknown) distribution h. Let $y = \mathcal{O}_h(Sch(f_1, f_2), \alpha, \beta)$. If $\hat{f} = \text{SCHEFFÉ}(f_1, f_2, y)$, then we have $d_{TV}(h, \hat{f}) \leq 3d_{TV}(h, \mathcal{F}) + \alpha$ with probability at least $1 - \beta$.*

2.3.2. CLASSES OF k DISTRIBUTIONS

In the more general setting where $|\mathcal{F}| = k$, one possible approach is to simply run a Scheffé test for all $\Theta(k^2)$ pairs of distributions and output the distribution that is returned the most. The round-robin tournament (Devroye and Lugosi, 2001) implements this idea (see see Algorithm 5), and achieves an approximation factor of 9. The following establishes the guarantee of this method; see Theorem 6.2 in Devroye and Lugosi (2001) for a proof.

Theorem 11 (Analysis of Round-Robin) *Let $\mathcal{F} = \{f_1, \ldots, f_k\}$ be a set of k distributions, $\alpha, \beta \in (0, 1)$, and $\mathcal{O}_h$ be a SQO for the (unknown) distribution h. Algorithm* ROUND-ROBIN$(\mathcal{F}, \mathcal{O}_h, \alpha, \beta)$ *is a SQA with a single round that makes at most $\frac{k(k-1)}{2}$ queries with accuracy α to $\mathcal{O}_h$, and returns a distribution $\hat{f}$ such that $d_{TV}(h, \hat{f}) \leq 9d_{TV}(h, \mathcal{F}) + \alpha$ with probability at least $1 - \beta$.*

Improving the approximation factor. It is possible to improve over the approximation factor of round-robin by using the MDE-variant algorithm (Mahalanabis and Stefankovic, 2007) described in Algorithm 6. This algorithm is an “advanced” version of the well-known Minimum Distance Estimate (MDE) algorithm (Yatracos, 1985) that solves the hypothesis selection problem with the approximation factor of 3 and $\Theta(k^2)$ queries. This is significantly better than the classic MDE that requires $\Theta(k^3)$ queries. Similar to round-robin, MDE-variant also relies on pair-wise comparisons.

[2] The precise guarantee of MDE-variant is given in the following theorem. The proof can be found in Mahalanabis and Stefankovic (2007).

Theorem 12 (Analysis of MDE-Variant) *Let $\mathcal{F} = \{f_1, \dots, f_k\}$ be a set of k distributions and let $\alpha, \beta \in (0,1)$. Let $\mathcal{O}_h$ be SQO for the (unknown) distribution h. Algorithm* $\text{MDE-VARIANT}(\mathcal{F}, \mathcal{O}_h, \alpha, \beta)$ *is a SQA that runs in a single round and returns a distribution $\hat{f} \in \mathcal{F}$ such that $d_{TV}(h, \hat{f}) \leq 3d_{TV}(h, \mathcal{F}) + \alpha$ with probability at least $1 - \beta$. Moreover, it makes $\frac{k(k-1)}{2}$ queries of accuracy α to the oracle.*

Both the round-robin and MDE-variant are non-adaptive (i.e., run in a single round) and make $\Theta(k^2)$ queries. Essentially, these algorithms query about every possible Scheffé set between pairs of distributions. This may not be a significant issue in settings where answering many queries does not incur huge cost, e.g., in the non-private setting where we can reuse the samples. However, as mentioned in Section 2.2, the number of samples under LDP constraint scales (roughly) linearly with the number of queries. Hence, simply simulating round-robin or MDE-variant in the local model results in a sub-optimal (with respect to k) sample complexity bound of $O(k^2)$. In the next section, we review adaptive methods and how they help to reduce the number of required queries for LDP hypothesis selection.

2.3.3. Hypothesis selection with adaptive statistical queries

Gopi et al. (2020) developed an adaptive algorithm that only makes $O(k \log\log k)$ queries to solve hypothesis selection. The algorithm partitions the candidate distributions in several groups with small size and runs a round-robin in each group. It then moves the winners to the next round and eliminates the rest of the candidates. The algorithm runs in $O(\log\log k)$ rounds and the size of groups increases in each round. Running the round-robin tournaments on groups with small size makes it possible to reduce the number of queries from $O(k^2)$ to $O(k \log\log k)$. However, the sequential nature of this method causes the resulting approximation factor to grow exponentially with the number of rounds. To address this issue, they proposed a modification of this algorithm using a general tool that we will discuss in detail in Section 4. They then showed that the resulting approximation factor becomes 27 for the modified algorithm (see Theorem 3).

We should mention that the hypothesis selection method of Acharya et al. (2014) can also be implemented using $\Theta(k)$ queries. However, compared to Gopi et al. (2020), it uses a significantly larger number of interactions (i.e., it runs in $\Theta(\log k)$ round rather than $\Theta(\log\log k)$).

As mentioned in Section 2.2, it is possible to privately implement a SQA that makes $O(k)$ queries using $O(\frac{k \log k/\beta}{\alpha^2 \varepsilon^2})$ samples. Such implementation was a main technical tool in the sample complexity analysis of Theorem 3. Note that the $O(\log k)$ term in the sample complexity comes from taking a union-bound argument over the accuracy of all the $O(k)$ queries. In the next section, we present the framework of *critical queries* which can help to get a better sample complexity.

3. SQ model with critical queries

In this section, we present a new framework for the sample complexity analysis of answering multiple statistical queries. As stated earlier, it is possible to implement an oracle that α-accurately

2. MDE and MDE-variant can also be used in the setting where $|\mathcal{F}|$ is unbounded. In this case, the sample complexity will depend on the VC dimension of the Yatracose class. See Devroye and Lugosi (2001) for details.

answers n queries with probability at least $1-\beta$ using $O(\frac{n\log n/\beta}{\alpha^2\varepsilon^2})$ samples. The fact that the oracle is required to be accurate in *all* the queries results in the $O(\log n)$ factor in the sample complexity (see Section 2.2). But what if the algorithm can establish its guarantees without the need to be *confident* about the accuracy of all queries? In other words, can we improve the sample complexity if we know that the success of the algorithm hinges on the correctness of only a small subset of the queries (i.e., *critical* queries)? To formulate this intuition, we define the following.

Definition 13 (Statistical Query Oracle with Critical Queries) *Let $(\mathcal{X}, \mathcal{B}, h)$ be a probability space. A Statistical Query Oracle with Critical queries (SQOC) for h is a random function that receives $\alpha, \beta \in (0,1)$, a finite workload $W = (W_i)_{i=1}^n$ where $W_i \in \mathcal{B}$, and $m \in [n]$, and outputs n real values. Formally, we say $\mathcal{O}_h$ is an SQOC for h if for all $\alpha, \beta \in (0,1), n \in \mathbb{N}, i \in [n], W = (W_1, \dots, W_n) \in \mathcal{B}^n$ we have*

$$\forall U \subset [n], |U| = m,\ \mathbb{P}\left[\sup_{i\in U} \left|\mathcal{O}_h\left(W, \alpha, \beta, m\right)_i - \mathbb{E}_{x\sim h}\left[1\left\{x \in W_i\right\}\right]\right| \geq \alpha\right] \leq \beta,$$

where $\mathcal{O}_h\left(W, \alpha, \beta, m\right)_i$ is the i-th output of $\mathcal{O}_h$ and the probability is over the randomness of oracle (for the fixed choice of input and U). We refer to m as the number of critical queries and to α, β as the approximation accuracy and the failure probability, respectively.

Compared to an SQO, SQOC receives an additional input of m that indicates the number of critical queries. An SQOC offers a weaker guarantee than a standard SQO (Definition 7) in that it ensures α-accuracy only for subsets of size $\leq m$ of (out of n) queries. That is, Definition 13 reduces to Definition 7 only for $m = n$. Note that, crucially, an algorithm that uses an SQOC does not need to know which specific queries are critical for its success (otherwise, it would have asked just those queries). Instead, the algorithm only needs to know (a bound on) the number of such critical queries. The potential benefit of using an SQOC over an SQO is that it can be implemented in the LDP setting using $O(\frac{n\log m}{\alpha^2\min\{\varepsilon^2,1\}})$ samples (See Observation 22 in Section 5).

3.1. Critical queries for maximal selection with adversarial comparators

In this section, we motivate the concept of critical queries further by providing a simple example in which critical queries can help improving the analysis. We first describe (and slightly modify) the setting of maximal selection with adversarial comparators (Acharya et al., 2014). Assume we have a set of k items $X = \{x_1, \dots, x_k\}$ and a value function $V : X \to \mathbb{R}$ that maps each item x_i to a value $V(x_i)$. Our goal is to find the item with the largest value by making pairwise comparisons. As an intermediate step, we want to pair the items randomly, compare them, and eliminate half of them. We, however, do not know the value function and only have access to an adversarial comparator $\mathcal{C}_V : X^2 \to X$. The comparator gets as input two items x_i and x_j, consumes $m(\alpha, \beta)$ fresh samples and *(i)* if $|V(x_i) - V(x_j)| > \alpha$ then with probability at least $1-\beta$ outputs $\arg\max_{l\in\{i,j\}}\{V(x_l)\}$ and *(ii)* if $|V(x_i) - V(x_j)| \leq \alpha$ then returns either item randomly. Here, $m : (0,1)^2 \to \mathbb{N}$ is decreasing in α and β, i.e., higher accuracy and smaller failure probability requires more samples. We now define the problem formally.

Example 1 *Let X be a set of k items, $V : X \to \mathbb{R}$ be a value function, and $\mathcal{C}_V : X^2 \to X$ be an adversarial comparator for the value function, and $\alpha, \beta \in (0,1)$. Assume $X^* = \{x \in X : |V(x) - \max_{x^*\in X} V(x^*)| \leq \alpha\}$ is the set of items with values comparable to the maximum value up* *to error α. Assume there is an algorithm that pairs the items randomly and invokes $\mathcal{C}_V$ to compare them. The algorithm returns the set Y of the $k/2$ winners. How many samples do we need to make sure that with probability at least $1-\beta$ we have a "good" item in Y, i.e., $X^* \cap Y \neq \emptyset$?*

Approach 1. Recall that if we have $m(\alpha, \beta)$ samples, then any pair of items (x, x') can be compared with accuracy α and with probability at least $1-\beta$. Applying union bound, we deduce that the algorithm can correctly compare all $k/2$ pairs with accuracy α if it has access to $\frac{k}{2}m(\alpha, \frac{2}{k}\beta)$ samples. Since the comparator is α-accurate, it follows that either a member $x^* \in X^*$ is included in Y or it has lost the comparison to another item $\hat{x}$ with $|V(\hat{x}) - V(x^*)| \leq \alpha$, which implies that $x' \in X^*$. Either way, $X^* \cap Y \neq \emptyset$ with probability at least $1-\beta$.
Approach 2 (Critical query). Let $x^* \in X^*$ and x' be the item paired with x^*. Clearly, if x^* is compared accurately, then $X^* \cap Y \neq \emptyset$. Thus, unlike the previous approach, we only need to inspect a single comparison. In other words, whether two other items x_i and x_j are compared accurately can be ignored in analyzing $X^* \cap Y \neq \emptyset$. This shows that we need only $\frac{k}{2}m(\alpha, \beta)$ samples, which is smaller than what we established with the previous approach. Notice that although the algorithm is unaware of the comparison involving x^*, there is only one critical query for the desired task.

The above example demonstrates that the concept of critical queries enables us to derive the same guarantees but with a smaller sample complexity, compared to what would be obtained simply by union bound.

4. A Sample-optimal algorithm for hypothesis selection

The algorithm of Gopi et al. (2020), basically, has as many critical queries as the total number of queries. Thus, using the concept of critical queries does not lead to a better sample complexity for their approach. In this section, we propose a new locally private hypothesis selector (which we term BOKSERR) that uses a smaller number of critical queries and achieves the optimal (i.e., linear) sample complexity dependence on k. BOKSERR consists of three sub-routines, namely boosted knockout, boosted sequential round-robin, and MDE-variant. All the algorithms can be found in Appendices C, D, and G, respectively. Before discussing each sub-routine, it is useful to introduce the following design technique that is borrowed from Acharya et al. (2014).

A design technique and some notations. Let $f^* \in \mathcal{F}$ be such that $d_{TV}(h, f^*) = d_{TV}(h, \mathcal{F})$ and $f \in \mathcal{F}$ be another distribution in $\mathcal{F}$. Theorem 10 suggests that if $d_{TV}(h, f) > 3d_{TV}(h, \mathcal{F}) + \alpha$, then with high probability, SCHEFFÉ(f, f^*) returns f^*. Assume $S = \{f \in \mathcal{F} : d_{TV}(h, f) \leq 3d_{TV}(h, \mathcal{F}) + \alpha\}$ and $\zeta = |S|/|\mathcal{F}|$. Intuitively, if ζ is small, then f^* is likely to win when compared to a randomly chosen distribution from $\mathcal{F}$. Otherwise, if ζ is large, then any appropriately-sized random sub-sample of $\mathcal{F}$ will contain a distribution from S with high probability. This intuition has been incorporated in our algorithm design: The first two sub-routines (namely, boosted knockout and boosted sequential round-robin) return two lists of distributions $\mathcal{L}_1$ and $\mathcal{L}_2$ such that with high probability either $f^* \in \mathcal{L}_1$ or $S \cap \mathcal{L}_2 \neq \emptyset$. Put differently, there exists a "good" distribution in $\mathcal{L}_1 \cup \mathcal{L}_2$ with high probability. The proof appears in Appendix C.

4.1. Boosted knockout

Given $\mathcal{F}$, oracle $\mathcal{O}_h$, and $t \geq 1$, boosted knockout runs in t rounds and returns two lists of distributions. The first one is constructed as follows. In each round, it randomly pairs distributions in $\mathcal{F}$ for $r = O(\log 1/\beta)$ times and makes pairwise comparisons, where each comparison is carried out by a

Scheffé test. It then selects the distributions that win at least $\frac{3}{4}$ of the comparisons in which they are involved and moves them to the next round. This ensures that the number of candidate distributions is reduced by at least a factor of $\frac{3}{2}$ in each round (see Appendix C for a proof). The distributions that have progressed to the last round constitute the first list. The second list is a random sub-sample of size $O(2^t \log(1/\beta))$ from the initial set of distributions $\mathcal{F}$. As mentioned before, the union of the above two lists includes at least one distribution from S, as formalized in the following theorem.

Theorem 14 (Analysis of Algorithm 1) *Let $\mathcal{F} = \{f_1, \dots, f_k\}$ be a set of k distributions, and let $\alpha, \beta \in (0,1)$ and $t \geq 1$. Let $\mathcal{O}_h$ be a SQOC for the (unknown) distribution h. Algorithm* BOOSTED-KNOCKOUT$(\mathcal{F}, \mathcal{O}_h, \alpha, \beta, t)$ *is a SQA that runs in t rounds. The number of statistical queries (i.e., size of the workload) to $\mathcal{O}_h$ in round $i \in [t]$ is at most $16k(\frac{4}{3})^i \log \frac{1}{\beta}$ with $32(\frac{4}{3})^i \log \frac{1}{\beta}$ many critical queries. Moreover, it returns two lists of distributions $\mathcal{K}_1$ and $\mathcal{K}_2$ with $|\mathcal{K}_1| \leq \frac{k}{2^{t.\log\frac{3}{2}}}$ and $|\mathcal{K}_2| \leq 8 \log \frac{1}{\beta} 2^{t.\log\frac{3}{2}}$ such that with probability at least $1-\beta$ either $d_{TV}(h, \mathcal{K}_1) = d_{TV}(h, \mathcal{F})$ or $d_{TV}(h, \mathcal{K}_2) \leq 3d_{TV}(h, \mathcal{F}) + \alpha$.*

We highlight two key properties of boosted knockout. First, it makes a small number of critical queries in each round, which are only those that correspond to comparing f^* with other distributions. To see this, note that we bound the probability that f^* is not in the first list by the probability that it will be paired with a distributions in S for at least $1/4$ of its comparisons. The analysis does not depend on the result of other comparisons. This implies that in each round only the r comparisons that concern f^* are critical. Second, it eliminates a subset of candidate distributions and prepares a smaller list of distributions for the next sub-routine sequential-round-robin. A more careful analysis of boosted knockout is given in the next theorem.

4.2. Boosted sequential round-robin

We now discuss the second sub-routine, namely, boosted sequential-round-robin (BSRR for short) described in Algorithm 2. Similar to boosted knockout, this algorithm relies on sequentially reducing the list of potential candidates starting from the $\mathcal{F}$, and eventually returning two lists of distributions. The first list is determined by sequential (i.e., adaptive) comparisons between distributions, while the second list is an appropriately-sized sub-sample from $\mathcal{F}$. However, instead of pairing the distributions, BSRR partitions the candidate distributions into a number of groups in each round, runs a round-robin tournament in each group, and keeps only the winners for the next round. The size of the groups are squarred in each round of BSRR and, therefore, the size of the candidate distributions decreases very quickly. If $\zeta = |S|/|\mathcal{F}|$ is small then with high probability f^* is not grouped with any distribution from S and will not be eliminated until the very last round. If ζ is large, then a distribution in S will be included in a random sub-sample of $\mathcal{F}$ with high probability. Hence, the union of the two lists will include a distribution from S with high probability.

BSRR is adapted from Gopi et al. (2020) with one main difference. Similar to boosted knockout, instead of only partitioning the distributions once in every round, we repeat this process for $O(\log 1/\beta)$ times and keep the winners of all of the round-robin tournaments. This results in boosting the probability that, in each round, there exists at least one group that includes f^* but not any other distributions from S. Thus, the probability that f^* is included in the first list also increases. The guarantee of BSRR is characterized in the following. In Appendix D, We show why these repetitions are necessary to achieve the logarithmic dependency on $1/\beta$ in the sample complexity and provide a proof for the theorem.

Theorem 15 (Analysis of Algorithm 2) *Let $\mathcal{F} = \{f_1, \ldots, f_k\}$ be a set of k distributions, and let $\alpha, \beta \in (0,1)$, $\eta > 0$, and $t \geq 1$. Let $\mathcal{O}_h$ be a SQOC for the (unknown) distribution h. Algorithm* BOOSTED-SEQUENTIAL-ROUND-ROBIN$(\mathcal{F}, \mathcal{O}_h, \alpha, \beta, \eta, t)$ *is a SQA that runs in t rounds. The number of statistical queries (i.e., size of the workload) to $\mathcal{O}_h$ in round $i \in [t]$ is at most $k\eta \left(\log \frac{1}{\beta}\right)^i$, all of which are critical. The algorithm returns two lists of distributions $\mathcal{R}_1$ and $\mathcal{R}_2$ with $|\mathcal{R}_1| \leq \frac{k(\log \frac{1}{\beta})^t}{\eta^{2^t-1}}$ and $|\mathcal{R}_2| \leq 2\eta^{2^t} \log \frac{1}{\beta}$ such that with probability at least $1-\beta$ either $d_{TV}(h, \mathcal{R}_1) = d_{TV}(h, \mathcal{F})$ or $d_{TV}(h, \mathcal{R}_2) \leq 3d_{TV}(h, \mathcal{F}) + \alpha$.*

4.3. Boosted-sequential-round-robin-MDE-variant (BOKSERR)

We are now ready to delineate the overall algorithm, that we call boosted-sequential-round-robin-MDE-variant (BOKSERR for short). It starts by calling the boosted knockout with $\Theta(\log\log k)$ rounds to construct two lists of distributions $\mathcal{K}_1$ and $\mathcal{K}_2$. The set $\mathcal{K}_1$ is then fed to BSRR with the same $\Theta(\log\log k)$ number of rounds to further generate two lists of distributions $\mathcal{R}_1$ and $\mathcal{R}_2$. Recall that all of the queries made by BSRR are critical. However, knockout ensures that the size of the input list to BSRR, namely $\mathcal{K}_1$, is $O(\frac{k}{(\log k)^{\log\log 1/\beta}})$. This makes sure that that the number of queries that BSRR makes is small and Observation 22 suggests that the queries can be answered with a sample complexity that is sub-linear in k.

Finally, BOKSERR uses MDE-variant (that was discussed in Section 2.3) to select its output distribution from $\mathcal{R}_1 \cup \mathcal{R}_2 \cup \mathcal{K}_2$. The union of the three lists is of size $O\left(\frac{\sqrt{k}(\log \frac{1}{\beta})}{(\log k)^{\log\log \frac{1}{\beta}}}\right)$, and therefore MDE-variant sub-routine (similar to round-robin) needs $O(k \log^2 \frac{1}{\beta})$ samples to choose the output distribution from this lists. The guarantee of BOKSERR is given as follows.

Theorem 16 (Analysis of Algorithm 3) *Let $\mathcal{F} = \{f_1, \ldots, f_k\}$ be a set of k distributions and let $\alpha, \beta \in (0,1)$. Let $\mathcal{O}_h$ be SQOC for the (unknown) distribution h. Algorithm* BOKSERR$(\mathcal{F}, \mathcal{O}_h, \alpha, \beta)$ *is a SQA that runs in $\left(6 + 4\log\log \frac{3}{\beta}\right) \log\log k$ rounds and returns a distribution $\hat{f} \in \mathcal{F}$ such that with probability at least $1-\beta$ we have $d_{TV}(h, \hat{f}) \leq 9d_{TV}(h, \mathcal{F}) + \alpha$. The total number of statistical queries is $\Theta(k \log^2 \frac{1}{\beta})$, and the number of critical queries is $O\left(\frac{k \log^2 \frac{1}{\beta}}{(\log k)^{\log\log \frac{1}{\beta}}}\right)$.*

5. Locally private selection

We start this section by defining the local privacy model. We then demonstrate that every SQO(C) and SQA can be implemented under the local privacy constraint, enabling us to construct a locally private version of Algorithm 3.

Definition 17 (ε-Local Randomizer, (Kasiviswanathan et al., 2011)) *We say a randomized function $\mathcal{M} : \mathcal{X} \to \mathcal{Y}$ is an ε-Local Randomizer (ε-LR) if for all $x, x' \in \mathcal{X}$ and any measurable subset $Y \subseteq \mathcal{Y}$ it satisfies*

$$\mathbb{P}\left[\mathcal{M}(x) \in Y\right] \leq e^{\varepsilon} \mathbb{P}\left[\mathcal{M}(x') \in Y\right].$$

The following definition enables us to abstract the "single round access" to a database through an LR oracle.

Definition 18 (ε-LR Oracle) *Let $D = \{x_1, \dots, x_m\} \in \mathcal{X}^m$ be a database of size m. An ε-LR oracle for D, denoted by Ψ_D, operates as follows: initially, it sets the total number of "questions" to 0, i.e., $s = 0$. Then in each call to the oracle:*

- *Ψ_D receives a series of ε-LRs $(\mathcal{M}_i)_{i=1}^n$ (for some $n > 0$).*
- *If $n + s > m$, then the oracle outputs NULL.*
- *Otherwise, it outputs $\Psi_D\left((\mathcal{M}_i)_{i=1}^n\right) = (\mathcal{M}_i(x_{i+s})_{i=1}^n)$ and updates $s = s + n$.*

We stress that the above oracle Ψ_D uses each data point only *once*. As a result, each $\mathcal{M}_i$ being an ε-LR implies that Ψ_D is also ε-LR. The construction laid out in Definition 18 provides a machinery for implementing oracles under the local privacy constraint. More specifically, one can construct a locally private version of an SQOC by composing a given local randomizer with the set indicator functions. To this goal, we choose the most well-known local randomizer, namely the randomized response (Warner, 1965), defined below.

Definition 19 (Randomized Response (RR)) *The randomized response mechanism R_ε is a randomized function that receives $x \in \{0, 1\}$ and outputs x with probability $\frac{e^\varepsilon}{e^\varepsilon+1}$ and $1 - x$ with probability $\frac{1}{e^\varepsilon+1}$.*

It can be easily verified that R_ε is an ε-LR. Moreover, according to the post-processing property of differential privacy (see, e.g., (Dwork and Roth, 2014, Proposition 2.1)), the composition of R_ε with any (measurable) function is still an ε-LR. The following lemma expounds on how to privately implement an SQOC by post-processing of randomized response mechanisms.

Lemma 20 (Simulating an SQOC with an Unbiased RR) *For every probability space $(\mathcal{X}, \mathcal{B}, h)$ and every $\varepsilon > 0$, $\mathcal{O}_h^{RR}$ defined as*

$$\mathcal{O}_h^{RR}((W_i)_{i=1}^k, \alpha, \beta, m) = \left(\frac{e^\varepsilon+1}{e^\varepsilon-1}\left(\frac{1}{p}\sum_{j=p(i-1)+1}^{p.i} R_\varepsilon\left(1\left\{x_j \in W_i\right\}\right) - \frac{1}{e^\varepsilon+1}\right)\right)_{i=1}^k,$$

is a valid SQOC for h, where $W_i \in \mathcal{B}$, $p = O\left(\frac{\log m/\beta}{\alpha^2 \min\{\varepsilon^2, 1\}}\right)$ and $D = \{x_1, \dots, x_{pk}\}$ is fresh i.i.d. samples generated from h.

Note that by setting the number of critical queries to the total number of queries the above lemma can be also used to simulate any SQO.

Next, we seek to show how to privately implement our hypothesis selection algorithm. To do so, we first need to formally define locally private algorithms. Informally speaking, an ε-LDP algorithm is a multi-round algorithm with access to an ε-LR oracle. In each round, the algorithm observes the outputs of the previous round and defines a series of ε-LRs. It then calls the oracle and computes a function of the returned values of the oracle as its output for that round.

Definition 21 (ε-LDP Algorithm) *An ε-LDP algorithm with t rounds is a randomized function $\mathcal{A}_\varepsilon$ that has access to an ε-LR oracle Ψ_D (for some database D), and outputs y^t using the following recursive procedure: at each round $j \in [t]$, the algorithm defines a function f^j depending on the outcomes of the previous round (i.e., y^{j-1}), picks a series of n_j ε-LRs $(\mathcal{M}_i^j)_{i=1}^{n_j}$, and computes $y^j = f^j\left(\Psi_D\left((\mathcal{M}_i^j)_{i=1}^{n_j}\right)\right)$.*

Intuitively, one can employ a strategy akin to that in Lemma 20 to simulate SQA by an ε-LDP algorithm using $\mathcal{O}_h^{RR}$, thus restricting the access to the database only through an ε-LR oracle.

Observation 22 *Any SQA with t rounds can be simulated with an ε-LDP algorithm $\mathcal{A}_\varepsilon$ with t rounds. Moreover, the ε-LR oracle associated with $\mathcal{A}_\varepsilon$ requires $O\left(\frac{n_i \log m_i/\beta_i}{\alpha_i^2 \min\{\varepsilon^2,1\}}\right)$ data points in round $i \in [t]$, where n_i is the size of the workload of queries, α_i, β_i are the accuracy and failure parameters, respectively, and $m_i \in [n_i]$ is the number of critical queries.*

This observation can be proved by a direct application of Lemma 20 and provides a systemic framework for constructing an ε-LDP version of Algorithm 3. The following theorem, whose proof is given in Appendix E, presents the guarantees of such a locally private hypothesis selector.

Theorem 23 (ε-LDP Implementation of Algorithm 3) *Let $\mathcal{F} = \{f_1, \dots, f_k\}$ be a set of k distributions, $\varepsilon > 0$ and $\alpha, \beta \in (0,1)$. Let $\mathcal{O}_h^{RR}$, as defined in Lemma 20, be the oracle for the (unknown) distribution h. Algorithm* $\text{BOKSERR}(\mathcal{F}, \mathcal{O}_h^{RR}, \alpha, \beta)$ *is an ε-LDP algorithm that requires a database of $O\left(\frac{k(\log\frac{1}{\beta})^2}{\alpha^2\min\{\varepsilon^2,1\}}\right)$ i.i.d. samples from h, runs in $\left(6 + 4\log\log\frac{3}{\beta}\right)\log\log k$ rounds, and returns a distribution $\hat{f} \in \mathcal{F}$ such that with probability at least $1-\beta$ we have $d_{TV}(h,\hat{f}) \le 9 d_{TV}(h,\mathcal{F}) + \alpha$.*

Acknowledgments

Hassan Ashtiani is affiliated with the Vector Institute and was supported by an NSERC Discovery grant. Shahab Asoodeh is also affiliated with the Vector Institute and was supported by an NSERC Discovery grant. Alireza F. Pour is supported by Cheriton Graduate scholarship.

References

Jayadev Acharya, Ashkan Jafarpour, Alon Orlitsky, and Ananda Theertha Suresh. Sorting with adversarial comparators and application to density estimation. In *2014 IEEE International Symposium on Information Theory*, pages 1682–1686. IEEE, 2014.

Jayadev Acharya, Ilias Diakonikolas, Chinmay Hegde, Jerry Zheng Li, and Ludwig Schmidt. Fast and near-optimal algorithms for approximating distributions by histograms. In *Proceedings of the 34th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*, pages 249–263, 2015.

Jayadev Acharya, Ilias Diakonikolas, Jerry Li, and Ludwig Schmidt. Sample-optimal density estimation in nearly-linear time. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1278–1289. SIAM, 2017.

Jayadev Acharya, Moein Falahatgar, Ashkan Jafarpour, Alon Orlitsky, and Ananda Theertha Suresh. Maximum selection and sorting with adversarial comparators. *The Journal of Machine Learning Research*, 19(1):2427–2457, 2018.

Jayadev Acharya, Clément Canonne, Cody Freitag, and Himanshu Tyagi. Test without trust: Optimal locally private distribution testing. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 2067–2076. PMLR, 2019.

Jayadev Acharya, Clément L Canonne, Ziteng Sun, and Himanshu Tyagi. Unified lower bounds for interactive high-dimensional estimation under information constraints. *arXiv preprint arXiv:2010.06562*, 2020.

Jayadev Acharya, Clément L Canonne, Cody Freitag, Ziteng Sun, and Himanshu Tyagi. Inference under information constraints iii: Local privacy constraints. *IEEE Journal on Selected Areas in Information Theory*, 2(1):253–267, 2021.

Jayadev Acharya, Clément L Canonne, Ziteng Sun, and Himanshu Tyagi. The role of interactivity in structured estimation. In *Conference on Learning Theory*, pages 1328–1355. PMLR, 2022.

Ishaq Aden-Ali, Hassan Ashtiani, and Gautam Kamath. On the sample complexity of privately learning unbounded high-dimensional gaussians. In *Algorithmic Learning Theory*, pages 185–216. PMLR, 2021a.

Ishaq Aden-Ali, Hassan Ashtiani, and Christopher Liaw. Privately learning mixtures of axis-aligned gaussians. *Advances in Neural Information Processing Systems*, 34:3925–3938, 2021b.

Mohammad Afzali, Hassan Ashtiani, and Christopher Liaw. Mixtures of gaussians are privately learnable with a polynomial number of samples. *arXiv preprint arXiv:2309.03847*, 2023.

Daniel Alabi, Pravesh K Kothari, Pranay Tankala, Prayaag Venkat, and Fred Zhang. Privately estimating a gaussian: Efficient, robust, and optimal. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, pages 483–496, 2023.

Maryam Aliakbarpour, Mark Bun, and Adam Smith. Hypothesis selection with memory constraints. *Advances in Neural Information Processing Systems*, 36, 2024.

Jamil Arbas, Hassan Ashtiani, and Christopher Liaw. Polynomial time and private learning of unbounded gaussian mixture models. In *International Conference on Machine Learning*. PMLR, 2023.

Hassan Ashtiani and Christopher Liaw. Private and polynomial time algorithms for learning gaussians and beyond. In *Conference on Learning Theory*, pages 1075–1076. PMLR, 2022.

Hassan Ashtiani and Abbas Mehrabian. Some techniques in density estimation. *arXiv preprint arXiv:1801.04003*, 2018.

Hassan Ashtiani, Shai Ben-David, Nicholas JA Harvey, Christopher Liaw, Abbas Mehrabian, and Yaniv Plan. Near-optimal sample complexity bounds for robust learning of gaussian mixtures via compression schemes. *Journal of the ACM (JACM)*, 67(6):1–42, 2020.

Hilal Asi, Vitaly Feldman, and Kunal Talwar. Optimal algorithms for mean estimation under local differential privacy. In *International Conference on Machine Learning*, pages 1046–1056. PMLR, 2022.

Hilal Asi, Vitaly Feldman, Jelani Nelson, Huy L Nguyen, and Kunal Talwar. Fast optimal locally private mean estimation via random projections. *arXiv preprint arXiv:2306.04444*, 2023.

Shahab Asoodeh and Huanyu Zhang. Contraction of locally differentially private mechanisms. *arXiv preprint arXiv:2210.13386*, 2022.

Tugkan Batu, Lance Fortnow, Ronitt Rubinfeld, Warren D Smith, and Patrick White. Testing that distributions are close. In *Proceedings 41st Annual Symposium on Foundations of Computer Science*, pages 259–269. IEEE, 2000.

Shai Ben-David, Alex Bie, Clément L Canonne, Gautam Kamath, and Vikrant Singhal. Private distribution learning with public data: The view from sample compression. *arXiv preprint arXiv:2308.06239*, 2023.

Alex Bie, Gautam Kamath, and Vikrant Singhal. Private estimation with public data. *Advances in Neural Information Processing Systems*, 35:18653–18666, 2022.

Lucien Birge. The grenader estimator: A nonasymptotic approach. *The Annals of Statistics*, pages 1532–1549, 1989.

Sourav Biswas, Yihe Dong, Gautam Kamath, and Jonathan Ullman. Coinpress: Practical private mean and covariance estimation. *Advances in Neural Information Processing Systems*, 33:14475–14485, 2020.

Olivier Bousquet, Daniel Kane, and Shay Moran. The optimal approximation factor in density estimation. In *Conference on Learning Theory*, pages 318–341. PMLR, 2019.

Olivier Bousquet, Mark Braverman, Gillat Kol, Klim Efremenko, and Shay Moran. Statistically near-optimal hypothesis selection. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 909–919. IEEE, 2022.

Mark Braverman, Ankit Garg, Tengyu Ma, Huy L Nguyen, and David P Woodruff. Communication lower bounds for statistical estimation problems via a distributed data processing inequality. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 1011–1020, 2016.

Mark Bun, Gautam Kamath, Thomas Steinke, and Steven Z Wu. Private hypothesis selection. *Advances in Neural Information Processing Systems*, 32, 2019a.

Mark Bun, Jelani Nelson, and Uri Stemmer. Heavy hitters and the structure of local privacy. *ACM Transactions on Algorithms (TALG)*, 15(4):1–40, 2019b.

Clément L. Canonne. *A Survey on Distribution Testing: Your Data is Big. But is it Blue?* Number 9 in Graduate Surveys. Theory of Computing Library, 2020. doi: 10.4086/toc.gs.2020.009. URL http://www.theoryofcomputing.org/library.html.

Clément L Canonne, Gautam Kamath, Audra McMillan, Adam Smith, and Jonathan Ullman. The structure of optimal private tests for simple hypotheses. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 310–321, 2019.

Clément L Canonne, Gautam Kamath, Audra McMillan, Jonathan Ullman, and Lydia Zakynthinou. Private identity testing for high-dimensional distributions. *Advances in Neural Information Processing Systems*, 33:10099–10111, 2020.

Siu-On Chan, Ilias Diakonikolas, Rocco A Servedio, and Xiaorui Sun. Efficient density estimation via piecewise polynomial approximation. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 604–613, 2014.

Wei-Ning Chen, Peter Kairouz, and Ayfer Ozgur. Breaking the communication-privacy-accuracy trilemma. *Advances in Neural Information Processing Systems*, 33:3312–3324, 2020.

Edith Cohen, Haim Kaplan, Yishay Mansour, Uri Stemmer, and Eliad Tsfadia. Differentially-private clustering of easy instances. In *International Conference on Machine Learning*, pages 2049–2059. PMLR, 2021.

Yuval Dagan and Gil Kur. A bounded-noise mechanism for differential privacy. In *Conference on Learning Theory*, pages 625–661. PMLR, 2022.

Amit Daniely and Vitaly Feldman. Locally private learning without interaction requires separation. *Advances in neural information processing systems*, 32, 2019.

Constantinos Daskalakis and Gautam Kamath. Faster and sample near-optimal algorithms for proper learning mixtures of gaussians. In *Conference on Learning Theory*, pages 1183–1213. PMLR, 2014.

Constantinos Daskalakis, Ilias Diakonikolas, and Rocco A Servedio. Learning poisson binomial distributions. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 709–728, 2012.

Doug Davidson, Abhradeep Guha Thakurta, Julien Freudiger, Gaurav Kapoor, Vivek Rangarajan Sridhar, Umesh S. Vaishampayan, and Andrew H. Vyrros. Learning new words, May 9 2017.

Luc Devroye and Gábor Lugosi. A universally acceptable smoothing factor for kernel density estimates. *The Annals of Statistics*, pages 2499–2512, 1996.

Luc Devroye and Gábor Lugosi. Nonasymptotic universal smoothing factors, kernel complexity and yatracos classes. *The Annals of Statistics*, pages 2626–2637, 1997.

Luc Devroye and Gábor Lugosi. *Combinatorial methods in density estimation*. Springer Science & Business Media, 2001.

Ilias Diakonikolas. Learning structured distributions. *Handbook of Big Data*, 267:10–1201, 2016.

Ilias Diakonikolas and Daniel M Kane. A new approach for testing properties of discrete distributions. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 685–694. IEEE, 2016.

Ilias Diakonikolas and Daniel M Kane. *Algorithmic high-dimensional robust statistics*. Cambridge University Press, 2023.

Ilias Diakonikolas, Daniel M Kane, and Alistair Stewart. Statistical query lower bounds for robust estimation of high-dimensional gaussians and gaussian mixtures. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 73–84. IEEE, 2017.

Ilias Diakonikolas, Gautam Kamath, Daniel Kane, Jerry Li, Ankur Moitra, and Alistair Stewart. Robust estimators in high-dimensions without the computational intractability. *SIAM Journal on Computing*, 48(2):742–864, 2019.

Apple Differential Privacy Team. Learning with privacy at scale. https://machinelearning.apple.com/research/learning-with-privacy-at-scale, December 2017.

Bolin Ding, Janardhan Kulkarni, and Sergey Yekhanin. Collecting telemetry data privately. *Advances in Neural Information Processing Systems*, 30, 2017.

John Duchi and Ryan Rogers. Lower bounds for locally private estimation via communication complexity. In *Conference on Learning Theory*, pages 1161–1191. PMLR, 2019.

John C Duchi, Michael I Jordan, and Martin J Wainwright. Local privacy and statistical minimax rates. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, pages 429–438. IEEE, 2013.

John C Duchi, Michael I Jordan, and Martin J Wainwright. Minimax optimal procedures for locally private estimation. *Journal of the American Statistical Association*, 113(521):182–201, 2018.

Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. 2014. URL http://www.cis.upenn.edu/~aaroth/Papers/privacybook.pdf.

Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our data, ourselves: Privacy via distributed noise generation. In *Advances in Cryptology-EUROCRYPT 2006: 24th Annual International Conference on the Theory and Applications of Cryptographic Techniques, St. Petersburg, Russia, May 28-June 1, 2006. Proceedings 25*, pages 486–503. Springer, 2006a.

Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam 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 3*, pages 265–284. Springer, 2006b.

Alexander Edmonds, Aleksandar Nikolov, and Jonathan Ullman. The power of factorization mechanisms in local and central differential privacy. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pages 425–438, 2020.

Úlfar Erlingsson, Vasyl Pihur, and Aleksandra Korolova. Rappor: Randomized aggregatable privacy-preserving ordinal response. In *Proceedings of the 2014 ACM SIGSAC conference on computer and communications security*, pages 1054–1067, 2014.

Alexandre Evfimievski, Johannes Gehrke, and Ramakrishnan Srikant. Limiting privacy breaches in privacy preserving data mining. In *Proceedings of the twenty-second ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, pages 211–222, 2003.

Vitaly Feldman. A general characterization of the statistical query complexity. In *Conference on Learning Theory*, pages 785–830. PMLR, 2017.

Badih Ghazi, Ravi Kumar, and Pasin Manurangsi. On avoiding the union bound when answering multiple differentially private queries. In *Conference on Learning Theory*, pages 2133–2146. PMLR, 2021.

Oded Goldreich and Dana Ron. On testing expansion in bounded-degree graphs. *Studies in Complexity and Cryptography. Miscellanea on the Interplay between Randomness and Computation: In Collaboration with Lidor Avigad, Mihir Bellare, Zvika Brakerski, Shafi Goldwasser, Shai Halevi, Tali Kaufman, Leonid Levin, Noam Nisan, Dana Ron, Madhu Sudan, Luca Trevisan, Salil Vadhan, Avi Wigderson, David Zuckerman*, pages 68–75, 2011.

Sivakanth Gopi, Gautam Kamath, Janardhan Kulkarni, Aleksandar Nikolov, Zhiwei Steven Wu, and Huanyu Zhang. Locally private hypothesis selection. In *Conference on Learning Theory*, pages 1785–1816. PMLR, 2020.

Yanjun Han, Ayfer Özgür, and Tsachy Weissman. Geometric lower bounds for distributed parameter estimation under communication constraints. In *Conference On Learning Theory*, pages 3163–3188. PMLR, 2018.

Rafael Hasminskii and Ildar Ibragimov. On density estimation in the view of kolmogorov's ideas in approximation theory. *The Annals of Statistics*, pages 999–1010, 1990.

Samuel B Hopkins, Gautam Kamath, Mahbod Majid, and Shyam Narayanan. Robustness implies privacy in statistical estimation. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, pages 497–506, 2023.

Matthew Joseph, Janardhan Kulkarni, Jieming Mao, and Steven Z Wu. Locally private gaussian estimation. *Advances in Neural Information Processing Systems*, 32, 2019a.

Matthew Joseph, Jieming Mao, Seth Neel, and Aaron Roth. The role of interactivity in local differential privacy. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 94–105. IEEE, 2019b.

Matthew Joseph, Jieming Mao, and Aaron Roth. Exponential separations in local differential privacy. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 515–527. SIAM, 2020.

Matthew Joseph, Jieming Mao, and Aaron Roth. Exponential separations in local privacy. *ACM Trans. Algorithms*, 18(4), oct 2022.

Peter Kairouz, H Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Kallista Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, et al. Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning*, 14(1–2):1–210, 2021.

Gautam Kamath, Jerry Li, Vikrant Singhal, and Jonathan Ullman. Privately learning high-dimensional distributions. In *Conference on Learning Theory*, pages 1853–1902. PMLR, 2019a.

Gautam Kamath, Or Sheffet, Vikrant Singhal, and Jonathan Ullman. Differentially private algorithms for learning mixtures of separated gaussians. *Advances in Neural Information Processing Systems*, 32, 2019b.

Gautam Kamath, Vikrant Singhal, and Jonathan Ullman. Private mean estimation of heavy-tailed distributions. In *Conference on Learning Theory*, pages 2204–2235. PMLR, 2020.

Gautam Kamath, Argyris Mouzakis, Vikrant Singhal, Thomas Steinke, and Jonathan Ullman. A private and computationally-efficient estimator for unbounded gaussians. In *Conference on Learning Theory*, pages 544–572. PMLR, 2022.

Vishesh Karwa and Salil Vadhan. Finite sample differentially private confidence intervals. In *9th Innovations in Theoretical Computer Science Conference (ITCS 2018)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.

Shiva Prasad Kasiviswanathan, Homin K Lee, Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. What can we learn privately? *SIAM Journal on Computing*, 40(3):793–826, 2011.

Michael Kearns. Efficient noise-tolerant learning from statistical queries. *Journal of the ACM (JACM)*, 45(6):983–1006, 1998.

Pravesh Kothari, Pasin Manurangsi, and Ameya Velingker. Private robust estimation by stabilizing convex relaxations. In *Conference on Learning Theory*, pages 723–777. PMLR, 2022.

Satyaki Mahalanabis and Daniel Stefankovic. Density estimation in linear time. *arXiv preprint arXiv:0712.2869*, 2007.

Ankur Moitra and Gregory Valiant. Settling the polynomial learnability of mixtures of gaussians. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pages 93–102. IEEE, 2010.

Shyam Narayanan. Private high-dimensional hypothesis testing. In *Conference on Learning Theory*, pages 3979–4027. PMLR, 2022.

Jerzy Neyman and Egon Sharpe Pearson. Ix. on the problem of the most efficient tests of statistical hypotheses. *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, 231(694-706):289–337, 1933.

Liam Paninski. A coincidence-based test for uniformity given very sparsely sampled discrete data. *IEEE Transactions on Information Theory*, 54(10):4750–4755, 2008.

Ankit Pensia, Po-Ling Loh, and Varun Jog. Simple binary hypothesis testing under communication constraints. In *2022 IEEE International Symposium on Information Theory (ISIT)*, pages 3297–3302. IEEE, 2022.

Ankit Pensia, Amir R Asadi, Varun Jog, and Po-Ling Loh. Simple binary hypothesis testing under local differential privacy and communication constraints. *arXiv preprint arXiv:2301.03566*, 2023.

Henry Scheffé. A useful convergence theorem for probability distributions. *The Annals of Mathematical Statistics*, 18(3):434–438, 1947.

Vikrant Singhal. A polynomial time, pure differentially private estimator for binary product distributions. *arXiv preprint arXiv:2304.06787*, 2023.

Thomas Steinke and Jonathan Ullman. Between pure and approximate differential privacy. *Journal of Privacy and Confidentiality*, 7(2), 2016.

Ananda Theertha Suresh, Alon Orlitsky, Jayadev Acharya, and Ashkan Jafarpour. Near-optimal-sample estimators for spherical gaussian mixtures. *Advances in Neural Information Processing Systems*, 27, 2014.

Balázs Szörényi. Characterizing statistical query learning: simplified notions and proofs. In *International Conference on Algorithmic Learning Theory*, pages 186–200. Springer, 2009.

Eliad Tsfadia, Edith Cohen, Haim Kaplan, Yishay Mansour, and Uri Stemmer. Friendlycore: Practical differentially private aggregation. In *International Conference on Machine Learning*, pages 21828–21863. PMLR, 2022.

Jonathan Ullman. Answering n {2+ o (1)} counting queries with differential privacy is hard. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 361–370, 2013.

Jonathan Ullman. Tight lower bounds for locally differentially private selection. *arXiv preprint arXiv:1802.02638*, 2018.

Stanley L Warner. Randomized response: A survey technique for eliminating evasive answer bias. *Journal of the American Statistical Association*, 60(309):63–69, 1965.

Yuhong Yang and Andrew Barron. Information-theoretic determination of minimax rates of convergence. *Annals of Statistics*, pages 1564–1599, 1999.

Yannis G Yatracos. Rates of convergence of minimum distance estimators and kolmogorov's entropy. *The Annals of Statistics*, 13(2):768–774, 1985.

Appendix A. Related work

Hypothesis selection is a classical problem in statistics (Scheffé, 1947; Yatracos, 1985; Devroye and Lugosi, 1996, 1997, 2001); see Devroye and Lugosi (2001) for an overview. More recent papers have studied hypothesis selection under various considerations such as computational efficiency, robustness and more (Mahalanabis and Stefankovic, 2007; Daskalakis et al., 2012; Daskalakis and Kamath, 2014; Suresh et al., 2014; Acharya et al., 2014, 2018; Diakonikolas et al., 2019; Bousquet et al., 2019, 2022). Other related problems are hypothesis testing (Neyman and Pearson, 1933; Paninski, 2008; Batu et al., 2000; Goldreich and Ron, 2011; Diakonikolas and Kane, 2016; Canonne, 2020) and distribution learning (Yatracos, 1985; Devroye and Lugosi, 2001; Birge, 1989; Hasminskii and Ibragimov, 1990; Yang and Barron, 1999; Devroye and Lugosi, 2001; Moitra and Valiant, 2010; Chan et al., 2014; Acharya et al., 2015; Diakonikolas, 2016; Acharya et al., 2017; Ashtiani and Mehrabian, 2018; Ashtiani et al., 2020; Diakonikolas and Kane, 2023).

Differential privacy was introduced in the seminal work of Dwork et al. (2006b). The central models of DP (including the pure (Dwork et al., 2006b) and approximate (Dwork et al., 2006a) models) have been studied extensively. In the central model, the first private hypothesis selection method was proposed by Bun et al. (2019a). This result was later improved by Aden-Ali et al. (2021a). Other problems that have been investigated in the central model are distribution testing Canonne et al. (2019, 2020); Narayanan (2022) and distribution learning Karwa and Vadhan (2018); Kamath et al. (2019b,a); Bun et al. (2019a); Biswas et al. (2020); Kamath et al. (2020); Aden-Ali et al. (2021a,b); Kamath et al. (2022); Ashtiani and Liaw (2022); Kothari et al. (2022); Alabi et al. (2023); Hopkins et al. (2023); Cohen et al. (2021); Tsfadia et al. (2022); Bie et al. (2022); Singhal (2023); Arbas et al. (2023); Ben-David et al. (2023); Afzali et al. (2023).

The local model of DP Warner (1965); Evfimievski et al. (2003); Kasiviswanathan et al. (2011); Duchi et al. (2013) is more stringent than the central one, and has been the model of choice for various real-world applications (Erlingsson et al., 2014; Differential Privacy Team, December 2017; Ding et al., 2017; Davidson et al., May 9 2017). Several lower and upper bounds have been established for various statistical estimation tasks under local DP (Duchi et al., 2018; Joseph et al., 2019a; Duchi and Rogers, 2019; Bun et al., 2019b; Chen et al., 2020; Asi et al., 2022, 2023; Acharya et al., 2021, 2020; Asoodeh and Zhang, 2022), including hypothesis testing (Pensia et al., 2022, 2023; Acharya et al., 2019) and selection (Gopi et al., 2020).

Studying statistical estimation tasks in the statistical query model is another related topic (Kearns, 1998; Szörényi, 2009; Feldman, 2017; Diakonikolas et al., 2017). More related to our work are those that study the sample complexity of answering multiple statistical queries under differential privacy (Ullman, 2013; Steinke and Ullman, 2016; Edmonds et al., 2020; Ghazi et al., 2021; Dagan and Kur, 2022).

The role of interaction in local models has been investigated in several works (Kasiviswanathan et al., 2011; Duchi et al., 2018; Duchi and Rogers, 2019; Joseph et al., 2019b; Daniely and Feldman, 2019; Joseph et al., 2020) and it has been shown that it is a powerful tool in designing algorithms that are sample efficient (Han et al., 2018; Acharya et al., 2020; Gopi et al., 2020; Acharya et al., 2022). We also prove in this work that using adaptive algorithms can solve the task of hypothesis selection more efficiently.

Appendix B. Miscellaneous facts

Lemma 24 (Hoeffding's Inequality) *Let $X_1, \ldots, X_n$ be i.i.d. random variables with $\mu = \mathbb{E}[X_i]$ and $\mathbb{P}[a \leq X_i \leq b] = 1$ for all $i \in [n]$. Then for every $\alpha > 0$ we have*

$$\mathbb{P}\left[\left|\frac{1}{n}\sum_{i=1}^{n} X_i - \mu\right| > \alpha\right] \leq \exp\left(-\frac{2n\alpha^2}{(b-a)^2}\right).$$

Appendix C. Analysis of Theorem 14 (Boosted knockout)

We prove each part of Theorem 14 in one of the following sections.

C.1. Size of the returned lists

We know that at the end of each round the set of distributions $\mathcal{F}$ is updated and only the distributions that are in $\mathcal{H}$ in that round will proceed to the next round. In the following lemma, we show that the number of remaining distributions decays exponentially with the number of rounds.

Lemma 25 *Let $k_0 = k$ and denote by k_i be the number of distributions that are remaining from the initial set of distributions $\mathcal{F} = \{f_1, \ldots, f_k\}$ at the end of round $i \in [t]$ of Algorithm 1, i.e., the distributions in $\mathcal{H}$ at the end of round i. Then we have $|k_i| \leq \frac{k}{(\frac{3}{2})^i}$.*

Proof We know that at round i we pair the remaining distributions randomly for r times and each time we run Scheffé tests between the pairs of distributions. Therefore, the total number of Scheffé

Algorithm 1 Boosted Knockout

Input: A set $\mathcal{F} = \{f_1, \dots, f_k\}$ of k distributions, Oracle $\mathcal{O}_h$, parameters $\alpha, \beta > 0, t \geq 1$
Output: Two lists of candidate distributions from $\mathcal{F}$
procedure BOOSTED-KNOCKOUT($\mathcal{F}, \mathcal{O}_h, \alpha, \beta, t$)**:**
    Sample $n = 8 \log \frac{1}{\beta} . 2^{t. \log \frac{3}{2}}$ distributions randomly from $\mathcal{F}$ and copy them to $\mathcal{K}_2$
    $\mathcal{F}_1 \leftarrow \mathcal{F}$
    for $i \in [t]$ **do**
        $\forall f \in \mathcal{F}_i$, let $w[f] \leftarrow 0$ ▷ To record the number of wins of f
        $\mathcal{H} \leftarrow \emptyset, \mathcal{G} \leftarrow \emptyset, r \leftarrow \lceil 32(\frac{4}{3})^i \log \frac{1}{\beta} \rceil$
        for $j \in [r]$ **do**
            Randomly pair the distribution in $\mathcal{F}_i$, copy the pairs to $\mathcal{G}$ ▷ $\frac{|\mathcal{F}_i|}{2} \leq |\mathcal{G}| \leq \frac{|\mathcal{F}_i|}{2} r$
        end
        for *every pair* $p = (f, f') \in \mathcal{G}$ **do**
            $B_p \leftarrow Sch(f, f')$
        end
        $W = (B_p)_{p \in \mathcal{G}}, (y_p)_{p \in \mathcal{G}} = \mathcal{O}_h(W, \alpha, \beta, r)$
        for *every pair* $p = (f, f') \in \mathcal{G}$ **do**
            if SCHEFFÉ(f, f', y_p) *returns* f **then**
                $w[f] \leftarrow w[f] + 1$
            else
                $w[f'] \leftarrow w[f'] + 1$
            end
        end
        for *every* $f \in \mathcal{F}_i$ **do**
            if $w[f] \geq \frac{3}{4} r$ **then**
                $\mathcal{H} \leftarrow \mathcal{H} \cup f$
        end
        $\mathcal{F}_{i+1} \leftarrow \mathcal{H}$
    end
    $\mathcal{K}_1 \leftarrow \mathcal{F}_{t+1}$
    return $\mathcal{K}_1, \mathcal{K}_2$
end

tests run in round i is equal to $\frac{k_i}{2}r$. We also know that only the distributions that are returned by at least $\frac{3}{4}r$ of the Scheffé tests will be included in $\mathcal{H}$ and remain for the next round. Therefore, each one of the k_{i+1} distributions that will remain for the next round (i.e., round $i+1$) must have been involved in at least $\frac{3}{4}r$ tests. We can then write

$$k_{i+1} \cdot \frac{3r}{4} \leq \frac{k_i}{2}r.$$

Thus, we can conclude that $k_{i+1} \leq \frac{k_i}{(\frac{3}{2})}$. The lemma can then be proved by using induction on $i \in [t]$. ■

Proof The fact that $|\mathcal{K}_1| \leq \frac{k}{2^{t.\log\frac{3}{2}}}$ can be verified easily from Lemma 25 and $|\mathcal{K}_2| \leq 8\log\frac{1}{\beta}2^{t.\log\frac{3}{2}}$ comes from the properties of the algorithm. ■

C.2. TV guarantees of the returned lists

Proof Let f^* be a distribution in $\mathcal{F}$ such that $d_{TV}(f^*, h) = \min_{f\in\mathcal{F}} d_{TV}(h, f) = d_{TV}(h, \mathcal{F})$ and let $\gamma = \frac{\alpha}{d_{TV}(h,\mathcal{F})}$. We only use γ for the analysis and we do not need to know its value– in fact the values of γ and $d_{TV}(h, \mathcal{F})$ are unknown to us. Denote by $S = \{f : d_{TV}(h, f) \leq (3+\gamma)d_{TV}(h, \mathcal{F}), f \in \mathcal{F}\}$ the set of distributions in $\mathcal{F}$ that their TV distance to h is within a $3+\gamma$ factor of the minimum TV distance $d_{TV}(h, \mathcal{F})$. We define $\zeta = \frac{|S|}{|\mathcal{F}|}$ as the ratio of these distributions and base our analysis on two different ranges for the value of ζ. We show that when ζ is smaller than $\frac{1}{8.\left(\frac{3}{2}\right)^t}$ then with high probability $f^* \in \mathcal{K}_1$ and if it is larger than this value then, with high probability, $S \cap \mathcal{K}_2 \neq \emptyset$.

Case 1: $\zeta \leq \frac{1}{8.\left(\frac{3}{2}\right)^t}$. We want to prove that in this case f^* will not be eliminated in any round $i \in [t]$ and makes it to the last round, i.e., $f^* \in \mathcal{K}_1$. In each round of boosted knockout, the distributions are randomly paired for r times and each time a Scheffé test is done between every pair, meaning that each distribution is involved in r tests. We then select the distributions that are returned by Scheffé tests for at least $\frac{3r}{4}$ times and move them to the next round, while eliminating the rest. Since $\alpha = \gamma.d_{TV}(h, \mathcal{F})$ we know from the guarantees of Scheffé test that for any $f' \in \mathcal{F}$, with probability at least $1-\beta$, the SCHEFFÉ$(f^*, f', \mathcal{O}_h, \alpha, \beta)$ test may not return f^* only if $d_{TV}(h, f') \leq 3.d_{TV}(h, f^*) + \alpha = (3+\gamma)d_{TV}(h, f^*) = (3+\gamma)d_{TV}(h, \mathcal{F})$, i.e., it may not return f^* only if f' is in S. For all $j \in [r]$, let Z_{ij} be a Bernoulli random variable that is equal to 1 if f^* is paired with a distribution in S at the jth time and is equal to 0 otherwise. Consequently, the number of times that f^* is tested against a distribution in S at round i is equal to $\sum_{j=1}^{r} Z_{ij}$. Moreover, let $\mathcal{F}_i$ refer to the set of distributions that are remained at the beginning of round $i \in [t]$ and denote by $\zeta_i = \frac{|S|}{\mathcal{F}_i}$ the ratio of the distributions in S relative to the distributions that are remained at the beginning of round $i \in [t]$, e.g., $\mathcal{F}_1 = \mathcal{F}$ and $\zeta_1 = \zeta$. Since $\mathbb{E}[Z_{ij}] = \zeta_i$ for all $i \in [t]$ and $j \in [r]$, we can write that

$$\begin{aligned}
&\mathbb{P}\left[f^* \text{ is eliminated at round } i\right] \\
&= \mathbb{P}\left[f^* \text{ is not returned by at least } \frac{r}{4} \text{ of the Scheffé tests in which it is tested at round } i\right] \\
&\leq \mathbb{P}\left[\sum_{j=1}^{r} Z_{ij} > \frac{r}{4}\right] = \mathbb{P}\left[\frac{1}{r}\sum_{j=1}^{r} Z_{ij} > \frac{1}{4}\right] = \mathbb{P}\left[\frac{1}{r}\sum_{j=1}^{r} Z_{ij} - \zeta_i > \frac{1}{4} - \zeta_i\right] \qquad (1) \\
&\leq \exp\left(-2r(\frac{1}{4} - \zeta_i)^2\right) = \exp\left(-2r(\frac{1}{4} - \zeta_i)^2\right) = \exp\left(-64(\frac{4}{3})^i \log\frac{1}{\beta}(\frac{1}{4} - \zeta_i)^2\right) \\
&= \beta^{64(\frac{4}{3})^i(\frac{1}{4}-\zeta_i)^2},
\end{aligned}$$

where the last inequality follows from Hoeffding's inequality (Lemma 24). From Lemma 25 we know that at the beginning of round i we have $\mathcal{F}_i = k_{i-1} \leq \frac{k}{(\frac{3}{2})^{i-1}}$. Therefore, we can conclude that $\zeta_i \leq (\frac{3}{2})^{i-1}\zeta$, i.e., the ratio of the distributions in S relative to the distributions that are remained grows at most by a factor $3/2$ at each round. We also know that $\zeta \leq \frac{1}{8.\left(\frac{3}{2}\right)^t}$. Hence $\zeta_i \leq \frac{1}{8}(\frac{2}{3})^{t-i+1} \leq \frac{1}{8}$ for all $i \in [t]$ and

$$64(\frac{1}{4} - \zeta_i)^2 \geq 64.\left(\frac{1}{4} - \frac{1}{8}\right)^2 \geq 1.$$

From the above equation and Equation 1 we can write that

$$\mathbb{P}\left[f^* \notin \mathcal{K}_1\right] = \sum_{i=1}^{t} \mathbb{P}\left[f^* \text{ is eliminated at round } i\right] \leq \sum_{i=1}^{t} \beta^{(\frac{4}{3})^i} \leq \sum_{i=1}^{\infty} \beta^{(\frac{4}{3})^i},$$

which is smaller than β for $\beta < 1/2$. Therefore, we have proved that if $\zeta \leq \frac{1}{8.\left(\frac{3}{2}\right)^t}$, then with probability at least $1-\beta$ (for $\beta < 0.5$) we have $f^* \in \mathcal{K}_1$ and thus $d_{TV}(h, \mathcal{K}_1) = d_{TV}(h, \mathcal{F})$.

Case 2: $\zeta > \frac{1}{8.\left(\frac{3}{2}\right)^t}$. In this case we prove that with probability at least $1-\beta$ a distribution ins S will be returned in $\mathcal{K}_2$. We can write

$$\mathbb{P}\left[S \cap \mathcal{K}_2 = \emptyset\right] \leq (1-\zeta)^n \leq e^{-\zeta n} = e^{-8\zeta \log\frac{1}{\beta}.2^{t.\log\frac{3}{2}}} = \left(e^{\log\beta}\right)^{8\zeta 2^{t.\log\frac{3}{2}}} = \beta^{8\zeta 2^{t.\log\frac{3}{2}}}.$$

We know that $\zeta > \frac{1}{8.\left(\frac{3}{2}\right)^t}$ and , therefore,

$$8\zeta 2^{t.\log\frac{3}{2}} > \frac{1}{8.\left(\frac{3}{2}\right)^t}.8.2^{t.\log\frac{3}{2}} = 1.$$

Combining the above two equations we can conclude that $\mathbb{P}\left[S \cap \mathcal{K}_2 = \emptyset\right] \leq \beta$. ■

C.3. Number of total and critical queries

Proof We know that a single Scheffé test can be done by one statistical query to the oracle. It is easy to verify that at round $i \in [t]$ the total number of statistical queries is equal to $r.\frac{k_{i-1}}{2}$, which is exactly equal to the number of Scheffé tests run in that round. From Lemma 25 we know that $k_{i-1} \leq \frac{k}{(\frac{3}{2})^{i-1}}$. Therefore the number of queries at round i is less than or equal to $r.\frac{k}{2(\frac{3}{2})^{i-1}} = \frac{32(\frac{4}{3})^i \log\frac{1}{\beta}.k}{2(\frac{3}{2})^{i-1}}$. Therefore, the total number of queries is less than or equal to

$$\sum_{i=1}^{t}\frac{32(\frac{4}{3})^i k\log\frac{1}{\beta}}{2(\frac{3}{2})^{i-1}} = 16\frac{4}{3}\sum_{i=1}^{t}\frac{(\frac{4}{3})^{i-1}k\log\frac{1}{\beta}}{(\frac{3}{2})^{i-1}} = \frac{64k\log\frac{1}{\beta}}{3}\sum_{i=1}^{\infty}(\frac{8}{9})^{i-1} \leq 192k\log\frac{1}{\beta} = O\left(k\log\frac{1}{\beta}\right).$$

Now we derive an upper bound on the number of critical queries. Note that if $\zeta \leq \frac{1}{8.(\frac{3}{2})^t}$ the guarantees of boosted knockout requires that with probability at least $1-\beta$ we have $f^* \in \mathcal{K}_1$ and the way that we proved this fact relied on finding the probability that f^* is returned by at least $\frac{3}{4}$ fraction of the rScheffé tests in which it is involved at each round. Also note that the output of all the other Scheffé tests that do not have f^* as their input is not important since we do not base our analysis on the exact value of ζ_i. We rather simply bound it by its worst-case– when none of the distributions in S are eliminated even until the last round t and ζ_i is multiplied by a factor of $3/2$ at each round. Therefore, the only queries that are critical and can change the output of the algorithm are those that are related to the Scheffé tests in which f^* is included. On the other hand if $\zeta \leq \frac{1}{8.(\frac{3}{2})^t}$, we proved that with probability at least $1-\beta$ a distribution in S is sampled and copied to $\mathcal{K}_2$ and this does not require any queries at all. Consequently, we can say that the number of critical queries in round i is less than or equal $r = 32(\frac{4}{3})^i\log\frac{1}{\beta}$, and thus the total number of critical queries is less than or equal to $96(\frac{4}{3})^{t_1+1}\log\frac{1}{\beta}$. ■

Appendix D. Analysis of Theorem 15 (Boosted sequential round-robin)

Before proving Theorem 15 we provide an intuition as to why the repetition of the round-robin tournament in each round of boosted SRR is necessary to obtain the logarithmic dependency of sample complexity on $1/\beta$ in Theorem 3.

Analysis of β in SRR. As already discussed, we wish to make sure that if ζ is smaller than some threshold then f^* will not be grouped with any distribution in S with high probability. It is easy to verify that in the "original" SRR (Gopi et al., 2020), where the distribution are grouped only once in each round, the failure can happen with probability $O(\zeta\sqrt{k})$. In order to make sure that this probability is smaller than β, we have to set the threshold to $\beta/\sqrt{k}$. On the other hand, we have to make sure that when ζ is larger than the threshold, then a sub-sample of the input distributions will include a distribution from S with probability at least $1-\beta$. The probability of failure for this event will depend on the size of the sub-sample $\mathcal{R}_2$ and is bounded by $O(\exp(-\zeta n))$. To make sure that this is smaller than β, we would have to set $n = O(\frac{\sqrt{k}\log(1/\beta)}{\beta})$. But notice that the sample complexity of MDE-variant is quadratic in the size of its input and, therefore, we would get a quadratic dependence on $1/\beta$ in the sample complexity. However, repeating the round-robins for several times in each round of boosted SRR makes sure that the probability that f^* is not eliminated increases sufficiently such that we can set a smaller threshold for ζ and, thus, sub-sample smaller number of distributions.

Similar to Theorem 14, we prove each part of Theorem 15 in one of the following sections.

Algorithm 2 Boosted Sequential-Round-Robin

```
Input: A set F = {f_1, ..., f_k} of k distributions, Oracle O_h, parameters α, β, η > 0, t ≥ 1
Output: Two lists of candidate distributions from F
procedure MULTI-ROUND-ROBIN(F, O_h, α, β, η):
    H ← ∅, G ← ∅
    Randomly partition the distributions in F into |F|/η sets of size η, copy the partitions to G
    for every partition G_i ∈ G do
        g_i ← ROUND-ROBIN(G_i, O_h, α, β)
        H ← H ∪ g_i
    end
    return H
end
procedure BOOSTED-SEQUENTIAL-ROUND-ROBIN(F, O_h, α, β, η, t):
    Sample 2η^{2^t} log(1/β) distributions randomly from F and copy them to R_2
    F_1 ← F
    for i ∈ [t] do
        H ← ∅
        for j ∈ ⌈log(1/β)⌉ do
            H_j ← MULTI-ROUND-ROBIN(F_i, O_h, α, β, η)
            H ← H ∪ H_j
        end
        F_{i+1} ← H
        η ← η^2
    end
    R_1 ← F_{t+1}
    return R_1, R_2
end
```

D.1. Size of the returned lists

We first state the following fact regarding the size of the output class of distributions of a single call to multi-round-robin.

Fact 26 *Let $\mathcal{F} = \{f_1, \ldots, f_k\}$ be a set of k distributions and let $\alpha, \beta > 0$. Let $\mathcal{O}_h$ be a SQOC of accuracy α for the (unknown) distribution h. Algorithm* MULTI-ROUND-ROBIN$(\mathcal{F}, \mathcal{O}_h, \alpha, \beta, \eta)$ *returns a set of distributions $\mathcal{H}$ with $|\mathcal{H}| = k/\eta$.*

Proof The fact is an immediate consequence of the facts that the number of partitions is $|\mathcal{G}| = k/\eta$ and that round-robin outputs a single distribution. ∎

We know that at each round of the boosted sequential-round-robin the set of distributions $\mathcal{F}$ is updated and only the distributions in $\mathcal{H}$ will remain for the next round. We first state the following fact and in Lemma 28 we state the total number of remaining distributions at each round of the boosted sequential-round-robin.

Fact 27 *Let $\eta_1 = \eta$ and η_i denote the updated value of η at the beginning of round $i \in [t]$ of Algorithm 2. We have $\eta_i = \eta^{2^{i-1}}$.*

Lemma 28 *Let $k_0 = k$ and k_i denote the number of distributions that are remaining from the initial set of distribtuions $\mathcal{F} = \{f_1, \ldots, f_k\}$ at the end of round $i \in [t]$ of Algorithm 2, i.e., the distributions in $\mathcal{H}$ at the end of round i. Then we have $k_i \leq \frac{k(\log \frac{1}{\beta})^i}{\eta^{2^i - 1}}$.*

Proof In round $i \in [t]$ of the boosted sequential-round-robin, the remaining set of distributions are used as an input to multi-round-robin for $\log \frac{1}{\beta}$ times. From Facts 26 and 27 we know that for each one of the $\log \frac{1}{\beta}$ runs of multi-round-robin at round i, the size of the output set of distributions is $\frac{k_{i-1}}{\eta_i} = \frac{k_{i-1}}{\eta^{2^{i-1}}}$. Therefore, the total number of distributions that are remained at the end of round i is at most $\frac{k_{i-1}. \log \frac{1}{\beta}}{\eta^{2^{i-1}}}$. We can then conclude the lemma using an induction on $i \in [t]$. It is easy to verify that at the end of round 1 we have $k_1 = \frac{k \log \frac{1}{\beta}}{\eta}$. Assume we have $k_i \leq \frac{k(\log \frac{1}{\beta})^i}{\eta^{2^i-1}}$. Then from the above argument we can conclude that $k_{i+1} \leq \frac{k_i \log \frac{1}{\beta}}{\eta^{2^i}} = \frac{k(\log \frac{1}{\beta})^{i+1}}{\eta^{2^i}\eta^{2^i-1}} = \frac{k(\log \frac{1}{\beta})^{i+1}}{\eta^{2^{i+1}-1}}$. ■

Proof From Lemma 28 we know that size of the returned list $\mathcal{R}_1$ is equal to $|\mathcal{R}_1| \leq \frac{k(\log \frac{1}{\beta})^t}{\eta^{2^t-1}}$. The fact that $|\mathcal{R}_2| = 2\eta^{2^t} \log \frac{1}{\beta}$ simply follows from the definition of algorithm. ■

D.2. TV guarantees of the returned lists

Proof Let f^* be a distribution in $\mathcal{F}$ such that $d_{TV}(f^*, h) = \min_{f \in \mathcal{F}} d_{TV}(h, f) = d_{TV}(h, \mathcal{F})$ and let $\gamma = \frac{\alpha}{d_{TV}(h,\mathcal{F})}$. We only use γ for the analysis and we do not need to know its value– in fact the values of γ and $d_{TV}(h, \mathcal{F})$ are unknown to us. Denote by $S = \{f : d_{TV}(h, f) \leq (3+\gamma)d_{TV}(h, \mathcal{F}), f \in \mathcal{F}\}$ the set of distributions in $\mathcal{F}$ that their TV distance to h is within a $3+\gamma$ factor of the minimum TV distance $d_{TV}(h, \mathcal{F})$. We define $\zeta = \frac{|S|}{|\mathcal{F}|}$ as the ratio of these distributions and base our analysis on two different ranges for the value of ζ. We show that when ζ is smaller than $\frac{1}{2\eta^{2^t}}$ then with high probability $f^* \in \mathcal{R}_1$ and if it is larger than this value then, with high probability, $S \cap \mathcal{R}_2 \neq \emptyset$.

Case 1: $\zeta \leq \frac{1}{2\eta^{2^t}}$. We want to prove that in this case f^* will not be eliminated in any round $i \in [t]$ and makes it to the last round, i.e., $f^* \in \mathcal{R}_1$. In each round of the boosted sequential-round-robin, we run the multi-round-robin sub-routine for a total of $\log \frac{1}{\beta}$. For each run of multi-round-robin at round $i \in [t]$ the distributions are partitioned into k_{i-1}/η_i groups and a round-robin is run on the distributions in each group. The outputs of all round-robins will then proceed to the next round. Since $\alpha = \gamma.d_{TV}(h, \mathcal{F})$ we know from the guarantees of Scheffé test that for any $f' \in \mathcal{F}$, with probability at least $1 - \beta$, the SCHEFFÉ$(f^*, f', \mathcal{O}_h, \alpha, \beta)$ test may not return f^* only if $d_{TV}(h, f') \leq 3.d_{TV}(h, f^*) + \alpha = (3 + \gamma)d_{TV}(h, f^*) = (3 + \gamma)d_{TV}(h, \mathcal{F})$, i.e., it may not return f^* only if f' is in S. In the following, we will bound the probability that f^* is eliminated by the probability, for all $\log \frac{1}{\beta}$ repetitions of multi-round-robin, the partition that contains f^* will also contain at least one distribution from S . Otherwise, f^* will be returned by every Scheffé test in that partition and thus by the multi-round-robin and proceeds to the next round.

Let $\mathcal{F}_i$ refer to the set of distributions that are remained at the beginning of round $i \in [t]$ and denote by $\zeta_i = \frac{|S|}{\mathcal{F}_i}$ the ratio of the distributions in S relative to the distributions that are remained at the beginning of round $i \in [t]$, e.g., $\mathcal{F}_1 = \mathcal{F}$ and $\zeta_1 = \zeta$. We can write that

$$\begin{aligned} &\mathbb{P}\left[f^* \text{ is eliminated at round } i\right] \\ &\mathbb{P}\left[f^* \text{ is grouped with at least one distribution in } S \text{ for all } \log\frac{1}{\beta} \text{ repetitions at round } i\right] \\ &\leq \left(\zeta_i \eta_i\right)^{\log\frac{1}{\beta}}, \end{aligned} \tag{2}$$

where the last inequality follows from the fact that the ratio of distributions in S and the size of the groups at round $i \in [t]$ are ζ_i and η_i, respectively. We also know that $\zeta_i \leq \eta^{2^{i-1}-1}\zeta$. To see this, note that in each of the $\log \frac{1}{\beta}$ runs of multi-round-robin the same set $\mathcal{F}_i$ of distributions are used as input and in the worst-case every distribution in S will be among the $\mathcal{F}_i/\eta_i$ distributions that are returned by each multi-round-robin. Similar to Lemma 28, we can prove the above bound on ζ_i using an induction on i.

Taking Fact 27 and $\zeta \leq \frac{1}{2\eta^{2^t}}$ into account, we can write that for all $i \in [t]$,

$$(\zeta_i\eta_i)^{\log\frac{1}{\beta}} \leq \left(\frac{\eta^{2^{i-1}-1}}{2\eta^{2^t}}\eta^{2^{i-1}}\right)^{\log\frac{1}{\beta}} = \left(\frac{\eta^{2^i-1}}{2\eta^{2^t}}\right)^{\log\frac{1}{\beta}}$$

From the above equation and Equation 2 we can write that

$$\begin{aligned} &\mathbb{P}\left[f^* \notin \mathcal{R}_1\right] = \sum_{i=1}^{t} \mathbb{P}\left[f^* \text{ is eliminated at round } i\right] \\ &\leq \sum_{i=1}^{t}\left(\frac{\eta^{2^i-1}}{2\eta^{2^t}}\right)^{\log\frac{1}{\beta}} = \left(\frac{1}{2\eta^{2^t}}\right)^{\log\frac{1}{\beta}} \sum_{i=1}^{t}\left(\eta^{2^i-1}\right)^{\log\frac{1}{\beta}} = \beta^{\log\left(2\eta^{2^t}\right)} \sum_{i=1}^{t} \beta^{\log\frac{1}{\eta^{2^i-1}}} \\ &= \beta^{\log\left(2\eta^{2^t}\right)} \sum_{i=1}^{t} \beta^{(2^i-1)\log\frac{1}{\eta}} = \beta^{\log\left(2\eta^{2^t}\right)} \frac{\left(\beta^{\log\frac{1}{\eta}}\right)^{2^t}-1}{\left(\beta^{\log\frac{1}{\eta}}\right)-1} \\ &\leq \beta^{\log\left(2\eta^{2^t}\right)}.\left(\beta^{\log\frac{1}{\eta}}\right)^{2^t} = \beta^{\log 2\eta^{2^t}\left(\frac{1}{\eta}\right)^{2^t}} = \beta. \end{aligned}$$

Therefore, we have proved that if $\zeta \leq \frac{1}{2\eta^{2^t}}$, then with probability at least $1-\beta$ we have $f^* \in \mathcal{R}_1$ and thus $d_{TV}(h, \mathcal{R}_1) = d_{TV}(h, \mathcal{F})$.

Case 2: $\zeta > \frac{1}{2\eta^{2^t}}$. In this case we prove that with probability at least $1-\beta$ a distribution in S will be returned in $\mathcal{R}_2$. We can write

$$\mathbb{P}\left[S \cap \mathcal{R}_2 = \emptyset\right] \leq (1-\zeta)^n \leq e^{-\zeta n} = e^{-\zeta 2\eta^{2^t}\log\frac{1}{\beta}} = \left(e^{\log\beta}\right)^{\zeta 2\eta^{2^t}} \leq \beta.$$

■

D.3. Number of total and critical queries

Proof We know from Theorem 11 that a round-robin run on a set of k distributions requires answer to $\frac{k(k-1)}{2}$ queries, all of which are critical. Thus, the number of total and critical queries are equal in boosted sequential-round-robin. Particularly, in round i, we run multi-round-robin for $\log \frac{1}{\beta}$ times. Each multi-round-robin partitions the set of distributions in round i into k/η_i groups of size η_i and runs a round-robin on each group. Therefore, the total number of queries is less than or equal to

$$\sum_{i=1}^{t} \frac{k_{i-1}}{\eta_i} \eta_i^2 \log \frac{1}{\beta} = \sum_{i=1}^{t} \frac{k\left(\log \frac{1}{\beta}\right)^{i-1}}{\eta^{2^{i-1}-1}} \eta_i \log \frac{1}{\beta} = \sum_{i=1}^{t} \frac{k\left(\log \frac{1}{\beta}\right)^{i-1}}{\eta^{2^{i-1}-1}} \eta^{2^{i-1}} \log \frac{1}{\beta} = \sum_{i=1}^{t} k\eta \left(\log \frac{1}{\beta}\right)^{i}$$
$$\leq k\eta \left(\log \frac{1}{\beta}\right)^{t+1},$$

where the last inequality holds if $\beta \leq 1/4$. ■

Appendix E. Analysis of Theorem 16 (BOKSERR)

We first prove the TV guarantees of the returned distribution and then find the number of total and critical queries.

E.1. TV guarantees of the returned distribution

Proof Let f^* be a distribution in $\mathcal{F}$ such that $d_{TV}(f^*, h) = \min_{f \in \mathcal{F}} d_{TV}(h, f) = d_{TV}(h, \mathcal{F})$ and let $\gamma = \frac{\alpha}{6 d_{TV}(h,\mathcal{F})}$. We only use γ for the analysis and we do not need to know its value– in fact the values of γ and $d_{TV}(h, \mathcal{F})$ are unknown to us. Denote by $S = \{f : d_{TV}(h, f) \leq (3 + \gamma) d_{TV}(h, \mathcal{F}), f \in \mathcal{F}\}$ the set of distributions in $\mathcal{F}$ that their TV distance to h is within a $3+\gamma$ factor of the minimum TV distance $d_{TV}(h, \mathcal{F})$. From Theorem 14 we know that with probability at least $1-\frac{\beta}{3}$ either $f^* \in \mathcal{K}_1$ or a distribution from S is in $\mathcal{K}_2$. Assume that $f^* \in \mathcal{K}_1$. Then it is easy to verify that for $S' = \{f : d_{TV}(h, f) \leq (3 + \gamma) d_{TV}(h, \mathcal{K}_1), f \in \mathcal{K}_1\}$ we have $S' \subseteq S$. Theorem 15 suggests that after running BOOSTED-SEQUENTIAL-ROUND-ROBIN$(\mathcal{K}_1, \mathcal{O}_h, \alpha/3, \beta/3, \eta, t')$, given that $f^* \in \mathcal{K}_1$, with probability at least $1 - \frac{\beta}{3}$ either $f^* \in \mathcal{R}_1$ or a distribution in S' (and thus in S) is in $\mathcal{R}_2$. Let $\tilde{\mathcal{F}} = \mathcal{R}_1 \cup \mathcal{R}_2 \cup \mathcal{K}_2$. By a union bound argument, we can conclude that with probability at least $1 - \frac{2\beta}{3}$ we have $\tilde{\mathcal{F}} \cap S \neq \emptyset$, i.e., $d_{TV}(h, \tilde{\mathcal{F}}) \leq (3+\gamma) d_{TV}(h, \mathcal{F})$. Theorem 12 together with a union bound argument suggests that for $\hat{f} \leftarrow$ MDE-VARIANT$(\tilde{\mathcal{F}}, \mathcal{O}_h, \alpha/2, \beta/3)$ with probability at least $1 - \beta$ we have

$$d_{TV}(h, \hat{f}) \leq 3 d_{TV}(h, \tilde{\mathcal{F}}) + \frac{\alpha}{2} \leq 3(3 + \gamma) d_{TV}(h, \mathcal{F}) + \frac{\alpha}{2} = 9 d_{TV}(h, \mathcal{F}) + 3\gamma . d_{TV}(h, \mathcal{F}) + \frac{\alpha}{2}$$
$$= 9 d_{TV}(h, \mathcal{F}) + \frac{\alpha}{2} + \frac{\alpha}{2} = 9 d_{TV}(h, \mathcal{F}) + \alpha.$$

■

E.2. Number of rounds

Proof It is easy to prove that the total number of rounds of Algorithm 3 is equal to the sum of the number of rounds required for Algorithms 1 and 2 plus one additional round for MDE-variant. Therefore, we can conclude that the total number of rounds is equal to $t + t' + 1 = \left(6 + 4\log\log\frac{3}{\beta}\right)\log\log k$. ■

E.3. Number of total and critical queries

Proof The number of total and critical queries are equal to the sum of the queries in boosted knockout, boosted sequential-round-robin, and MDE-variant sub-routines. Let $\tilde{\mathcal{F}} = \mathcal{R}_1 \cup \mathcal{R}_2 \cup \mathcal{K}_2$. From Theorems 14, 15, and 12, we can conclude that the total number of queries is less than or equal to

$$
\begin{aligned}
&192|\mathcal{F}|\log\frac{3}{\beta} + |\mathcal{K}_1|\eta\left(\log\frac{3}{\beta}\right)^{t'+1} + \frac{|\tilde{\mathcal{F}}|(|\tilde{\mathcal{F}}|-1)}{2} \\
&\leq 192k\log\frac{3}{\beta} + \frac{k\cdot k'^{\frac{1}{2(t'+1)}}\cdot(\log\frac{3}{\beta})^{t'+1}}{2^{t\log\frac{3}{2}}} + \left(\frac{k(\log\frac{3}{\beta})^{t_2}}{2^{t\log\frac{3}{2}}k'^{\frac{2t'-1}{2(t'+1)}}} + 8\log\frac{3}{\beta}2^{t\cdot\log\frac{3}{2}} + 2k'^{\frac{2t'}{2(t'+1)}}\log\frac{3}{\beta}\right)^2 \\
&\leq 192k\log\frac{3}{\beta} + O\left(\frac{k(\log\frac{1}{\beta})^2}{(\log k)^{\log\log\frac{1}{\beta}}}\right) = O\left(\frac{k(\log\frac{1}{\beta})^2}{(\log k)^{\log\log\frac{1}{\beta}}}\right),
\end{aligned}
$$

where the last line follows by plugging the values of k' and t' based on Algorithm 3. The number of critical queries is less than equal to

$$
96(\frac{4}{3})^{t+1}\log\frac{1}{\beta} + |\mathcal{K}_1|\eta\left(\log\frac{3}{\beta}\right)^{t'+1} + \frac{|\tilde{\mathcal{F}}|(|\tilde{\mathcal{F}}|-1)}{2} \leq O\left((\log k)^{\log\log\frac{1}{\beta}}\right) + O\left(\frac{k(\log\frac{1}{\beta})^2}{(\log k)^{\log\log\frac{1}{\beta}}}\right).
$$

■

Algorithm 3 BOKSERR

Input: A set $\mathcal{F} = \{f_1, \ldots, f_k\}$ of k distributions, Oracle $\mathcal{O}_h$, parameters $\alpha, \beta > 0$
Output: $\hat{f} \in \mathcal{F}$ such that $d_{TV}(h, \hat{f}) \leq 9d_{TV}(h, \mathcal{F}) + \alpha$ with probability at least $1-\beta$
$t \leftarrow \left(5 + 4\log\log\frac{3}{\beta}\right)\log\log k$, $t' \leftarrow \log\log k - 1$
$k' \leftarrow \frac{k}{2^{t\log\frac{3}{2}}}$, $\eta \leftarrow k'^{\frac{1}{2(t'+1)}}$
procedure BOKSERR($\mathcal{F}, \mathcal{O}_h, \alpha, \beta$):
  $\mathcal{K}_1, \mathcal{K}_2, \mathcal{R}_1, \mathcal{R}_2 \leftarrow \emptyset$
  $\mathcal{K}_1, \mathcal{K}_2 \leftarrow$ BOOSTED-KNOCKOUT($\mathcal{F}, \mathcal{O}_h, \alpha/6, \beta/3, t$)
  $\mathcal{R}_1, \mathcal{R}_2 \leftarrow$ BOOSTED-SEQUENTIAL-ROUND-ROBIN($\mathcal{K}_1, \mathcal{O}_h, \alpha/6, \beta/3, \eta, t'$)
  return MDE-VARIANT($\mathcal{R}_1 \cup \mathcal{R}_2 \cup \mathcal{K}_2, \mathcal{O}_h, \alpha/2, \beta/3$)
end

Appendix F. Missing proofs from Section 5

F.1. Proof of Theorem 23

Proof We find the sample complexity of the algorithm by finding the samples needed to run the three sub-routines MDE-variant, which we denote by m_1, m_2, and m_3, respectively. From Theorem 14, we know in round $i \in [t]$ of boosted knockout the algorithm asks a workload of $\frac{32(\frac{4}{3})^i k \log\frac{1}{\beta}}{2(\frac{3}{2})^{i-1}}$ queries and at most $32(\frac{4}{3})^i \log\frac{1}{\beta}$ of them are critical. Therefore, from Lemma 20, the total number of samples that the oracle requires for the run of boosted knockout is equal to

$$\begin{aligned} m_1 &= \sum_{i=1}^{t} \frac{1}{\alpha^2 \min\{\varepsilon^2, 1\}} \frac{32(\frac{4}{3})^i k \log\frac{3}{\beta}}{2(\frac{3}{2})^{i-1}} \cdot \log\left(\frac{32(\frac{4}{3})^i \log\frac{3}{\beta}}{\beta}\right) \\ &= \frac{64k}{3\alpha^2 \min\{\varepsilon^2, 1\}} \left(\log\frac{3}{\beta}\right) \sum_{i=1}^{t} \frac{(\frac{4}{3})^{i-1} \log\left(\frac{32(\frac{4}{3})^i \log\frac{3}{\beta}}{2^i \beta}\right)}{(\frac{3}{2})^{i-1}} \\ &= \frac{64k\log\frac{3}{\beta}}{3\alpha^2 \min\{\varepsilon^2, 1\}} \left(\log(\frac{3}{\beta}) \sum_{i=1}^{\infty} (\frac{8}{9})^{i-1} + \sum_{i=1}^{\infty} (\frac{8}{9})^{i-1} \log\left(\frac{128}{3} (\frac{4}{3})^{i-1} \log\frac{3}{\beta}\right)\right) \\ &= O\left(\frac{k(\log\frac{1}{\beta})^2}{\alpha^2 \min\{\varepsilon^2, 1\}}\right). \end{aligned} \tag{3}$$

Theorem 15 suggests that in round $i \in [t']$ of boosted sequential-round-robin there are at most $|\mathcal{K}_1| \eta \left(\log\frac{3}{\beta}\right)^i$ queries and all of them are critical. Therefore, we can find the samples required to run this sub-routine by writing

$$\begin{aligned} m_2 &= \sum_{i=1}^{t'} \frac{1}{\alpha^2 \min\{\varepsilon^2, 1\}} \cdot \frac{k}{2^{t\log\frac{3}{2}}} . k'^{\frac{1}{2(t'+1)}} \left(\log\frac{3}{\beta}\right)^i . \log\left(\frac{\frac{k}{2^{t\log\frac{3}{2}}} . k'^{\frac{1}{2(t'+1)}} \left(\log\frac{3}{\beta}\right)^i}{\beta}\right) \\ &= \frac{1}{\alpha^2 \min\{\varepsilon^2, 1\}} \frac{2k}{(\log k)^{\log\frac{3}{2}(5+4\log\log\frac{3}{\beta})(1+1/\log k)}} \cdot \\ & \sum_{i=1}^{t'} \left(\log\frac{3}{\beta}\right)^i \cdot \log\left(\frac{2k}{(\log k)^{\log\frac{3}{2}(5+4\log\log\frac{3}{\beta})(1+1/\log k)} \cdot \beta} \cdot \left(\log\frac{3}{\beta}\right)^i\right) \end{aligned}$$

We first upper bound the term with the sum.

$$\begin{aligned}
&\sum_{i=1}^{t'} \left(\log\frac{3}{\beta}\right)^i \cdot \log\left(\frac{2k}{(\log k)^{\log\frac{3}{2}(5+4\log\log\frac{3}{\beta})(1+1/\log k)}\cdot\beta}\cdot\left(\log\frac{3}{\beta}\right)^i\right)\\
&=\log\left(\frac{2k}{(\log k)^{\log\frac{3}{2}(5+4\log\log\frac{3}{\beta})(1+1/\log k)}\cdot\beta}\right)\sum_{i=1}^{t'}\left(\log\frac{3}{\beta}\right)^i+\sum_{i=1}^{t'}\log\left(\frac{3}{\beta}\right)^i\log\left(\left(\log\frac{3}{\beta}\right)^i\right)\right)\\
&\le\log\left(\frac{2k}{(\log k)^{\log\frac{3}{2}(5+4\log\log\frac{3}{\beta})(1+1/\log k)}\cdot\beta}\right)\frac{\left(\log\frac{3}{\beta}\right)^{t'+1}-1}{\left(\log\frac{3}{\beta}\right)-1}+\log\log k\cdot\log\log\frac{3}{\beta}\cdot\sum_{i=1}^{t'}\left(\log\frac{3}{\beta}\right)^i\\
&\le\log\left(\frac{2k}{(\log k)^{\log\frac{3}{2}(5+4\log\log\frac{3}{\beta})(1+1/\log k)}\cdot\beta}\right)\left(\log\frac{3}{\beta}\right)^{t'+1}+\log\log k.\log\log\frac{3}{\beta}\left(\log\frac{3}{\beta}\right)^{t'+1}\\
&=\log\left(\frac{2k}{(\log k)^{\log\frac{3}{2}(5+4\log\log\frac{3}{\beta})(1+1/\log k)}.\beta}\right)(\log k)^{\log\log\frac{3}{\beta}}+\log\log k.\log\log\frac{3}{\beta}\,(\log k)^{\log\log\frac{3}{\beta}},
\end{aligned}$$

where in the last line we used the fact that $t' = \log\log k - 1$. We can therefore upper bound m_2 as follows.

$$\begin{aligned}
&m_2=\frac{1}{\alpha^2\min\{\varepsilon^2,1\}}\frac{2k}{(\log k)^{\log\frac{3}{2}(5+4\log\log\frac{3}{\beta})(1+1/\log k)}}\cdot\\
&\sum_{i=1}^{t'}\left(\log\frac{3}{\beta}\right)^i\cdot\log\left(\frac{2k}{(\log k)^{\log\frac{3}{2}(5+4\log\log\frac{3}{\beta})(1+1/\log k)}\cdot\beta}\cdot\left(\log\frac{3}{\beta}\right)^i\right)\\
&\le\frac{1}{\alpha^2\min\{\varepsilon^2,1\}}\cdot\frac{2k\,(\log k)^{\log\log\frac{3}{\beta}}}{(\log k)^{\log\frac{3}{2}(5+4\log\log\frac{3}{\beta})(1+1/\log k)}}\log\left(\frac{2k}{(\log k)^{\log\frac{3}{2}(5+4\log\log\frac{3}{\beta})(1+1/\log k)}.\beta}\right)\\
&+\frac{1}{\alpha^2\min\{\varepsilon^2,1\}}\cdot\frac{2k\,(\log k)^{\log\log\frac{3}{\beta}}}{(\log k)^{\log\frac{3}{2}(5+4\log\log\frac{3}{\beta})(1+1/\log k)}}\log\log k.\log\log\frac{3}{\beta}\\
&=O\left(\frac{k\log\frac{1}{\beta}}{\alpha^2\min\{\varepsilon^2,1\}}\right).
\end{aligned}\tag{4}$$

Finally we find the number of samples needed to run MDE-VARIANT. Let $\tilde{\mathcal{F}} = \mathcal{R}_1\cup\mathcal{R}_2\cup\mathcal{K}_1$. We know that MDE-variant makes $\frac{|\mathcal{F}|(|\mathcal{F}|-1)}{2}$ queries and all of them are critical. Therefore, we can write

$$\begin{aligned}
m_3&\le\frac{1}{\alpha^2\min\{\varepsilon^2,1\}}\left(\frac{k(\log\frac{1}{\beta})^{t'}}{2^{t\log\frac{3}{2}}k'^{\frac{2^{t'}-1}{2^{(t'+1)}}}}+8\log\frac{3}{\beta}2^{t\cdot\log\frac{3}{2}}+2k'^{\frac{2^{t'}}{2^{(t'+1)}}}\right)^2\cdot\\
&\quad\log\left(\frac{k(\log\frac{1}{\beta})^{t'}}{2^{t\cdot\log\frac{3}{2}}k'^{\frac{2^{t'}-1}{2^{(t'+1)}}}}+8\log\frac{3}{\beta}2^{t\cdot\log\frac{3}{2}}+2k'^{\frac{2^{t'}}{2^{(t'+1)}}}\right)\\
&=O\left(\frac{k(\log\frac{1}{\beta})^2}{\alpha^2\min\{\varepsilon^2,1\}}\right).
\end{aligned}\tag{5}$$

Combining Equations 3, 4, and 5 concludes that $m_1 + m_2 + m_3 = O\left(\frac{k(\log\frac{1}{\beta})^2}{\alpha^2 \min\{\varepsilon^2,1\}}\right)$ as desired. ■

F.2. Proof of Lemma 20

Proof We first prove that for any $i \in [n]$ with probability at least $\beta' = \beta/m$ we have that the output of oracle for the query W_i is close to $\mathbb{E}_{x\sim h}\left[1\left\{x \in W_i\right\}\right]$. For any $p(i-1)+1 \leq j \leq p.i$, let $y_j = R_\varepsilon(1\left\{x_j \in W_i\right\})$ and z_j be the random variable defined below

$$z_j = \frac{e^\varepsilon+1}{e^\varepsilon-1}\left(R_\varepsilon\left(1\left\{x_j \in W_i\right\}\right) - \frac{1}{e^\varepsilon+1}\right) = \frac{e^\varepsilon+1}{e^\varepsilon-1}\left(y_j - \frac{1}{e^\varepsilon+1}\right).$$

We now want to apply the Hoeffding's inequality on random variables z_j. We know that $\mathbb{E}\left[y_j\right] = \frac{e^\varepsilon}{e^\varepsilon+1}1\left\{x_j \in W_i\right\} + \frac{1}{e^\varepsilon+1}(1 - 1\left\{x_j \in W_i\right\}) = \frac{e^\varepsilon-1}{e^\varepsilon+1}1\left\{x_j \in W_i\right\} + \frac{1}{e^\varepsilon+1}$ where the expectation is taken over the randomness of R_ε. We can then compute the expected value of z_j as follows.

$$\mathbb{E}_{x_j,R_\varepsilon}\left[z_j\right] = \frac{e^\varepsilon+1}{e^\varepsilon-1}\left(\mathbb{E}\left[y_j\right] - \frac{1}{e^\varepsilon+1}\right) = \frac{e^\varepsilon+1}{e^\varepsilon-1}\left(\mathbb{E}\left[\frac{e^\varepsilon-1}{e^\varepsilon+1}1\left\{x_j \in W_i\right\} + \frac{1}{e^\varepsilon+1}\right] - \frac{1}{e^\varepsilon+1}\right)$$
$$= \frac{e^\varepsilon+1}{e^\varepsilon-1}\left(\frac{e^\varepsilon-1}{e^\varepsilon+1}\mathbb{E}\left[1\left\{x_j \in W_i\right\}\right] + \frac{1}{e^\varepsilon+1} - \frac{1}{e^\varepsilon+1}\right) = \mathbb{E}\left[1\left\{x_j \in W_i\right\}\right].$$

We can therefore apply the Hoeffding's inequality (Lemma 24) on z_j and write that

$$\mathbb{P}\left[\left|\frac{e^\varepsilon+1}{e^\varepsilon-1}\left(\frac{1}{p}\sum_{j=p(i-1)+1}^{p.i} R_\varepsilon\left(1\left\{x_j \in W_i\right\}\right) - \frac{1}{e^\varepsilon+1}\right) - \mathbb{E}\left[1\left\{x_j \in W_i\right\}\right]\right| > \alpha\right] \leq \exp\left(-\frac{2p\alpha^2}{(\frac{e^\varepsilon+1}{e^\varepsilon-1})^2}\right).$$

If $\varepsilon \leq 1$ then $\left(\frac{e^\varepsilon+1}{e^\varepsilon-1}\right)^2 = O\left((\frac{1}{\varepsilon})^2\right)$. If $\varepsilon > 1$ then for any constant $c > 3$ we know that $\ln\frac{c+1}{c-1} < 1 < \varepsilon$ and, thus, $\left(\frac{e^\varepsilon+1}{e^\varepsilon-1}\right)^2 < c^2 = O(1)$. Therefore, we can conclude that $\left(\frac{e^\varepsilon+1}{e^\varepsilon-1}\right)^2 = O\left(\frac{1}{\min\{\varepsilon^2,1\}}\right)$. Setting $p = \frac{\log 1/\beta'}{\min \alpha^2\{\varepsilon^2,1\}} = \frac{\log m/\beta}{\alpha^2 \min\{\varepsilon^2,1\}}$ implies that

$$\mathbb{P}\left[\left|\frac{e^\varepsilon+1}{e^\varepsilon-1}\left(\frac{1}{p}\sum_{j=p(i-1)+1}^{p.i} R_\varepsilon\left(1\left\{x_j \in W_i\right\}\right) - \frac{1}{e^\varepsilon+1}\right) - \mathbb{E}\left[1\left\{x_j \in W_i\right\}\right]\right| > \alpha\right] \leq \beta' = \frac{\beta}{m}.$$

A simple union bound argument concludes concludes that

$$\forall U \subset [k], |U| = m, \mathbb{P}\left[\sup_{i\in U}\left|\mathcal{O}_h^{RR}\left(W,\alpha,\beta,pk\right)_i - \mathbb{E}_{x\sim h}\left[1\left\{x \in W_i\right\}\right]\right| \geq \alpha\right] \leq \beta,$$

which proves that $\mathcal{O}_h^{RR}$ is a valid SQO for h. ■

Appendix G. Pseudocode of Classical Algorithms

Algorithm 4 Scheffé test

Input: A pair of distributions f_1 and f_2 on $\mathcal{X}$, $y \in \mathbb{R}$

Output: f_1 or f_2

procedure SCHEFFÉ(f_1, f_2, y)**:**

  $B_s \leftarrow Sch(f_1, f_2)$

  if $|f_1[B_s] - y| \leq |f_1[B_s] - y|$ **then**

    return f_1

  else

    return f_2

  end

end

Algorithm 5 Round-Robin

Input: A set $\mathcal{F} = \{f_1, \ldots, f_k\}$ of k distributions, Oracle $\mathcal{O}_h$, parameters $\alpha, \beta > 0$

Output: A distribution $\hat{f} \in \mathcal{F}$ such that $d_{TV}(h, \hat{f}) \leq 9d_{TV}(h, \mathcal{F}) + \alpha$ with probability at least $1 - \beta$.

procedure ROUND-ROBIN$(\mathcal{F}, \mathcal{O}_h, \alpha, \beta)$**:**

  $\forall f \in \mathcal{F}, w[f] \leftarrow 0$ ▷ To record the number of wins of f

  for *all* $i, j \in [k]$ *where* $j > i$ **do**

    $B_{ij} \leftarrow Sch(f_i, f_j)$

  end

  $W = (B_{ij})_{i=1,j>i}^{k}, (y_{ij})_{i=0,j>i}^{k} \leftarrow \mathcal{O}_h(W, \alpha, \beta)$

  for *all* $i, j \in [k]$ *where* $j > i$ **do**

    if SCHEFFÉ(f_i, f_j, y_{ij}) *returns* f_i **then**

      $w[f_i] \leftarrow w[f_i] + 1$

    else

      $w[f_j] \leftarrow w[f_j] + 1$

    end

  end

  return $\hat{f}$ with maximum number of wins $w[\hat{f}]$

end

Algorithm 6 MDE-Variant

Input: A set $\mathcal{F} = \{f_1, \ldots, f_k\}$ of k distributions, Oracle $\mathcal{O}_h$, parameters $\alpha, \beta > 0$

Output: A distribution $\hat{f} \in \mathcal{F}$ such that $d_{TV}(h, \hat{f}) \leq 3d_{TV}(h, \mathcal{F}) + \alpha$ with probability at least $1 - \beta$.

procedure MDE-VARIANT($\mathcal{F}, \mathcal{O}_h, \alpha, \beta$)**:**

$\forall f \in \mathcal{F}, w[f] \leftarrow 0$

for *every* $f \in \mathcal{F}$ **do**

for *every* $f' \in \mathcal{F} \setminus f$ **do**

$B_i \leftarrow Sch(f, f'), i \leftarrow i + 1$

end

$W = (B_i)_{i=1}^{k-1}$, $(y_i)_{i=1}^{k-1} = \mathcal{O}_h(W, \alpha, \beta)$

$w[f] \leftarrow \sup_{B_i} |f[B_i] - y_i|$

end

return $\hat{f} = \arg\min_{f \in \mathcal{F}} w[f]$

end