

# Preventing False Discovery in Interactive Data Analysis is Hard

Moritz Hardt

IBM Research Almaden

Jonathan Ullman

School of Engineering and Applied Sciences and  
Center for Research on Computation and Society  
Harvard University

**Abstract**—We show that, under a standard hardness assumption, there is no computationally efficient algorithm that given  $n$  samples from an unknown distribution can give valid answers to  $n^{3+o(1)}$  adaptively chosen statistical queries. A statistical query asks for the expectation of a predicate over the underlying distribution, and an answer to a statistical query is valid if it is “close” to the correct expectation over the distribution.

Our result stands in stark contrast to the well known fact that exponentially many statistical queries can be answered validly and efficiently if the queries are chosen non-adaptively (no query may depend on the answers to previous queries). Moreover, Dwork et al. [1], showed how to accurately answer exponentially many adaptively chosen statistical queries via a computationally inefficient algorithm. They also gave efficient algorithm that can answer nearly  $n^2$  adaptively chosen queries, which shows our result is almost quantitatively tight.

Conceptually, our result demonstrates that achieving statistical validity alone can be a source of computational intractability in adaptive settings. For example, in the modern large collaborative research environment, data analysts typically choose a particular approach based on previous findings. False discovery occurs if a research finding is supported by the data but not by the underlying distribution. While the study of preventing false discovery in Statistics is decades old, to the best of our knowledge our result is the first to demonstrate a computational barrier. In particular, our result suggests that the perceived difficulty of preventing false discovery in today’s collaborative research environment may be inherent.

## I. INTRODUCTION

Empirical research is commonly done by testing multiple hypotheses on a finite sample. A test outcome is deemed statistically significant if it is unlikely to have occurred by chance alone. False discovery arises if the analyst incorrectly declares an observation as statistically significant. For decades statisticians have been devising methods for preventing false discovery, such as the widely used and highly influential method for controlling the *false discovery rate* due to Benjamini and Hochberg [2].

Nevertheless the problem of false discovery persists across all empirical sciences today. Popular articles report on an increasing number of invalid research findings. Why is it seemingly so difficult to prevent false discovery? Today’s practice of data analysis diverges from classical statistics in its massive scale, heavy use of sophisticated algorithms,

and large number of participants in any given project. Importantly, the way modern data analysts interact with the data set is inherently *adaptive*—many design choices, including the choice and tuning of the algorithm itself, depend on previous interactions with the data set. An extreme example are data science competitions, in which hundreds of data scientists analyze the same data set and repeatedly evaluate their approach against the same data. This level of adaptivity makes it nearly impossible to give a precise a priori description of the experimental setup.

We formalize the way in which data analysts may interact with a data set using the *statistical-query model (SQ model)* of Kearns [3]. In the statistical query model, there is an algorithm called the *oracle* that gets access to  $n$  samples from an unknown distribution  $\mathcal{D}$  over some finite universe  $\mathcal{X}$ . We will assume throughout that  $\mathcal{X} = \{0, 1\}^d$  where we think of the parameter  $d$  as the dimensionality of the data. A statistical query  $q$  is specified by a predicate  $p: \mathcal{X} \rightarrow \{0, 1\}$  and the answer to a statistical query is defined as  $q(\mathcal{D}) = \mathbb{E}_{x \sim \mathcal{D}} p(x)$ . The goal of the oracle is to give an answer  $a$  to the query that is *accurate* (or *statistically valid*) in the sense that  $|a - q(\mathcal{D})| \leq \alpha$  with high probability. Throughout our work we only require  $\alpha$  to be a small constant. Put differently, the goal of the oracle is to provide answers that generalize to the underlying distribution rather than answers that are specific to the sample. The latter is always easy to achieve by outputting the empirical average of the query predicate on the sample.

The SQ model has a number of advantages for our purposes. First, almost all natural machine learning algorithms can be compiled into a sequence of statistical queries. Hence, the model does not give up much generality. Second, it makes it convenient to formalize adaptivity. In the adaptive/interactive setting, the analyst is modeled as an efficient algorithm that given a sequence of queries and answers  $q_1, a_1, q_2, a_2, \dots, q_i, a_i$  (previously exchanged with the oracle) produces a new query  $q_{i+1}$ . We say that an oracle is *accurate* given  $n$  samples for  $k$  adaptively chosen queries, if for every distribution  $\mathcal{D}$ , given  $n$  samples from  $\mathcal{D}$  the oracle accurately responds to any computationally efficient adaptive analyst that makes at most  $k$  queries. A computationally efficient oracle should run time polynomial in  $n$  and  $d$  on input of each query.

A recent work by Dwork, Feldman, Hardt, Pitassi, Rein-

Jonathan Ullman was supported by NSF grant CNS-1237235.

gold and Roth [1] addresses the problem of answering adaptive statistical queries. Their main result implies that there is a computationally inefficient oracle that accurately answers even an exponential number of adaptively chosen statistical queries. Moreover, they show that a quadratic number of queries can be answered accurately and efficiently. Our main theorem shows that these results are essentially as far as it goes. Under a standard cryptographic hardness assumption, we show that there is no efficient oracle that is accurate on more than a cubic number of adaptively chosen queries.

**Theorem I.1** (Informal). *Assuming one-way functions exist, there is no computationally efficient oracle that given  $n$  samples from an arbitrary distribution is accurate on  $n^{3+o(1)}$  arbitrary adaptively chosen queries.*

An intuitive interpretation of the theorem is that if an efficient oracle attempts to answer more than  $n^3$  statistical queries it cannot in general maintain that its answers are statistically valid with respect to the underlying distribution. Of course, the oracle can always report the exact answer of the query on its data set. However, one can show that this strategy does not maintain accuracy on adaptive queries in general and—as our theorem shows—neither does any other computationally efficient approach. From a technical perspective our result gives a strong computational lower bound in the statistical query model. Lower bounds in the statistical query model have been studied for more than two decades. But more broadly speaking, we interpret our result as pointing at an inherent computational obstruction to preventing false discovery in collaborative science.

Note that Theorem I.1 stands in sharp contrast to the non-adaptive setting. If we fix queries  $q_1, \dots, q_k$  and then sample  $n$  items from the distribution  $\mathcal{D}$ , the observed empirical answer to each query on the data set will be close to the correct answer with high probability so long as  $k = 2^{o(n)}$ . This guarantee follows from a Hoeffding bound together with the union bound.

Our hardness result applies when the dimensionality of the data grows with the sample size more than logarithmically so that  $2^d$  is no longer polynomial in  $n$ .<sup>1</sup> This requirement is rather mild, and is also necessary. If  $n \gg 2^d$  then the empirical distribution of the  $n$  samples will be close to the underlying distribution in statistical distance, and thus every statistical query can be answered accurately given the sample. More generally, as we discuss in Section I-B, there are algorithms that run in time polynomial in  $n$  and  $2^d$  and provide accuracy even on an exponential number of adaptively chosen queries [1]. Thus, our results show that the dimensionality of the data has a major effect on the hardness of the problem. In fact, we provide a second theorem that shows that if the dimensionality is polynomially large in  $n$ ,

<sup>1</sup>This is under the stronger but standard assumption that exponentially hard one-way-functions exist.

then we cannot even hope for a *computationally unbounded* oracle that provides accuracy on adaptive queries.

**Theorem I.2** (Informal). *There is no computationally unbounded oracle that given  $n$  samples of dimension  $d = n^{3+o(1)}$  from an arbitrary distribution is accurate on  $n^{3+o(1)}$  arbitrary adaptively chosen queries.*

While the dimension in the previous theorem has to be large, there are important data sets that exhibit this trade-off between sample size and dimension. A good example are genome wide association studies (GWAS). Here, the sample size corresponds to patients with a certain (possibly rare disease) and is often in the hundreds. The dimensionality corresponds to the number of relevant positions in the human genome and is often in the millions. Moreover, the genome resolution is increasing rapidly with new technology whereas the number of available patients is not.

To conclude this discussion of our results, we believe that adaptivity is an essential element of modern data analysis that ought to be taken into account by theoretical models. At the same time, our theorems demonstrate the intrinsic difficulty of coping with adaptivity.

#### A. Proof Overview

The intuition for our proof is rather simple. We will design a challenge distribution  $\mathcal{D}$  and a computationally efficient adaptive analyst  $\mathcal{A}$  so that the following is true. If any computationally efficient oracle  $\mathcal{O}$  is given  $n$  samples  $S = \{x_1, \dots, x_n\}$  drawn from  $\mathcal{D}$  then our adaptive analyst  $\mathcal{A}$  is able to reconstruct  $n' = n - O(1)$  samples  $\{y_1, \dots, y_{n'}\} \subseteq S$ . In other words, the analyst is able to find all but a constant number of samples that the oracle is using. While the analyst has a priori information about the distribution  $\mathcal{D}$  it has no information whatsoever about which sample  $\mathcal{O}$  received. Nevertheless, the analyst can reconstruct essentially all of the hidden sample. Quantitatively, the analyst proceeds in  $n - O(1)$  rounds and each round consists of roughly  $n^2$  queries. In each round the analyst successfully recovers one data item from the oracle provided that the oracle continues to give accurate answers. After the analyst has recovered almost all samples, the effective sample size of the oracle has shrunk down to a constant size. At this point it is easy for the analyst to find queries on which the oracle gives blatantly inaccurate answers.

The first problem is to recover even a single data point inside the oracle’s sample. To solve this problem we rely on a cryptographic primitive known as a fingerprinting code. Fingerprinting codes were introduced by Boneh and Shaw [4] for the problem of watermarking digital content. A fingerprinting code has two components. The first component generates a set of “challenge queries.” The second component is a “tracing algorithm” which takes answers to these queries and returns a data item. The fingerprinting code gives the guarantee that if the challenge queries are

answered accurately, and by looking only at how each challenge query is defined on  $S$ , then the tracing algorithm will successfully recover one element in  $S$ . Unfortunately, nothing prevents the oracle from evaluating the queries at points outside of  $S$ . In fact, information-theoretically the challenge queries used in our attack reveal information about the unknown distribution  $\mathcal{D}$  that the oracle didn't have previously. To circumvent this problem, we need to hide this additional information from the oracle. To do so, we use an encryption scheme to effectively hide the definition of the query on points outside of  $S$  from the oracle. The encryption is sufficient to show that, assuming that the oracle is computationally bounded, the tracing algorithm of the fingerprinting code will succeed. We note that encryption schemes suitable for our purpose exist under the standard assumption that one-way functions exist. With this one-round approach in mind, we can proceed iteratively. In the next round we exclude the previously learned data item from the definition of the challenge queries, which ensures that the analyst learns a new item in each round.

There is one important subtlety. The tracing algorithm of the fingerprinting code will only succeed if the oracle answers the challenge queries accurately with respect to its sample  $S$ . However, our assumption is that the oracle is accurate with respect to the underlying distribution  $\mathcal{D}$  rather than the sample  $S$ . We need to worry that eventually the sample and the distribution disagree on the challenge queries. In this case the oracle may be inaccurate on its sample (and hence tracing fails), yet still accurate on the distribution. To rule out this pathological situation we use a measure concentration property of our specific choice of fingerprinting code. Specifically, we use the fact that the challenge queries of the code are essentially random predicates with a certain bias. This property allows us to use the randomness of the challenge queries to argue that the sample  $S$  approximately agrees with the distribution  $\mathcal{D}$  on these queries with sufficiently high probability so long as there are at least  $O(1)$  elements in the sample that we haven't reconstructed yet. Due to the approximation error incurred here, we also need to use a somewhat stronger primitive called a robust fingerprinting code that was just recently provided in work by Bun, Ullman and Vadhan [5], which also satisfies the necessary measure concentration property.

### B. Connection to Privacy and Reconstruction Attacks

Our work builds on a close connection to the problem of designing privacy-preserving oracles. Here, the goal is to provide answers to statistical queries in such a way that the analyst does not learn the specifics of individual data records but rather global properties of the underlying distribution. A successful approach for formalizing this desideratum is the notion of differential privacy [6]. Differential privacy requires that the answers given by the oracle are randomized in such a way that the presence or absence of any single

data item in the sample cannot be detected. It is known that differential privacy prevents so-called *reconstruction attacks*. A reconstruction attack is an algorithm that is able to reconstruct most entries of a data set by interacting with the oracle. Such an attack demonstrates that the oracle is *blatantly non-private* (it fails to satisfy not only differential privacy, but any reasonable notion of privacy). Our work can be considered an efficient reconstruction attack as we give an efficient adaptive analyst that reconstructs almost all of the data points that the oracle uses if the oracle provides accuracy on  $n^{3+o(1)}$  queries. An immediate consequence of our work is therefore the following result.

**Theorem 1.3.** *Assuming one-way functions exist, any computationally efficient oracle that given  $n$  samples from an arbitrary distribution is accurate on  $n^{3+o(1)}$  arbitrary adaptively chosen queries must be blatantly non-private.*

This result should be compared with recent work of Ullman [7], which showed that oracles satisfying differential privacy cannot answer even  $n^{2+o(1)}$  non-adaptively chosen queries. Here we show that if the queries are chosen adaptively, then the same conclusion holds even for oracles that merely thwart blatant non-privacy, up to a factor of  $n$  loss in the number of queries.

An important difference to the privacy setting is how *accuracy* is defined. In the privacy setting, accuracy is defined with respect to the oracle's sample. It is trivial to maintain accuracy with respect to the sample by answering each query with the sample mean, which succeeds even when the oracle is blatantly non-private. In the setting of false discovery, we define accuracy with respect to the underlying distribution and show that achieving this notion of accuracy is hard for the oracle.

*Upper bounds for answering adaptive queries:* Differential privacy is also useful in establishing upper bounds in our setting. At a high-level, differential privacy is a *stability condition* on an algorithm requiring that the output varies only slightly with the addition or deletion of a sample point. On the other hand, it is known that stability implies generalization [8]. Hence, we can think of the interaction between an oracle and an analyst as a single algorithm that satisfies a stability guarantee strong enough to imply generalization bounds with respect to the underlying distribution. This approach was formalized by Dwork et al. [1] leading to upper bounds in the adaptive setting when combined with algorithms from the differential privacy literature. Specifically, work of Roth-Roughgarden [9] and Hardt-Rothblum [10] addresses differential privacy in the interactive setting. The latter work shows that  $2^{\tilde{\Omega}(n/\sqrt{d})}$  statistical queries can be answered with constant error under differential privacy. However, the running time is exponential in  $d$ . Using the results of [1] this leads to the same upper bound in the adaptive statistical query setting. Similarly, there is an efficient differentially private mechanism that

gives constant accuracy for  $\tilde{\Omega}(n^2)$  queries. This leads to a computationally efficient upper bound in our setting. To summarize we state the following theorem.

**Theorem I.4** ([1]). *There is an inefficient algorithm that accurately answers  $2^{\tilde{\Omega}(n/\sqrt{d})}$  adaptively chosen statistical queries. Moreover, there is an efficient algorithm that accurately answers  $\tilde{\Omega}(n^2)$  adaptively chosen queries.*

We emphasize that exponential running time was known to be inherent for differentially private algorithms that answer  $n^{2+o(1)}$  statistical queries [7], but prior to our results it was possible that there was an efficient oracle that accurately answered exponentially many adaptively chosen statistical queries via a different approach.

### C. Related Work

The combination of fingerprinting codes and encryption in our one-round approach is a common technique in the construction of “traitor-tracing schemes.” Traitor-tracing schemes were introduced by Chor, Fiat, and Naor [11], also for the problem of secure distribution of digital content. Dwork et al. [12] were the first to show that traitor-tracing schemes can be used to prove computational hardness results for differential privacy. Ullman [7] showed that traitor-tracing schemes with certain non-standard security properties can be used to prove strong computational hardness results for differential privacy, and showed how to construct such a scheme. In fact, the one-round approach described above closely mirrors the traitor-tracing scheme constructed in [7]. See [7] for a more detailed discussion of prior work on traitor-tracing and the issues that arise when using traitor-tracing schemes in the context of differential privacy.

Our work was also inspired by recent work of Hardt and Woodruff [13], which showed that no low-dimensional linear sketch can give valid answers to even a polynomial number of adaptively chosen queries. Technically our results are largely orthogonal to theirs, since we consider arbitrary computationally efficient statistical query oracles, rather than linear sketches. However, their work also noted the connection between differential privacy and validly answering adaptively chosen queries and inspired our iterative approach.

The work on false discovery and multiple hypothesis testing in Statistics is too vast to survey here. However, to our knowledge none of that work addresses adaptivity.

### Acknowledgments

We are extremely grateful to Aaron Roth for raising the issue of adaptivity in false discovery control at the Simons Workshop on Differential Privacy. We thank Cynthia Dwork and Omer Reingold for introducing us to the area of False Discovery Control. We benefited from discussions with Vitaly Feldman and Aaron Roth on the relation of Differential Privacy and adaptive statistical queries. We also thank Salil Vadhan for helpful discussions. We acknowledge

the Simons Institute for Theoretical Computer Science at Berkeley where this work started.

## II. PRELIMINARIES

### A. Oracles for Adaptive Queries

Let  $\mathcal{D}$  be a distribution over  $\{0, 1\}^d$ , for some parameter  $d \in \mathbb{N}$ . We are interested in answering *statistical queries* about the distribution  $\mathcal{D}$ . A statistical query on  $\{0, 1\}^d$  is specified by a predicate  $q : \{0, 1\}^d \rightarrow \{0, 1\}$  and (abusing notation) is defined to be  $q(\mathcal{D}) = \mathbb{E}_{x \leftarrow \mathcal{D}} [q(x)]$ . The goal is to design an *oracle*  $\mathcal{O}$  that answers statistical queries about the unknown distribution  $\mathcal{D}$ , given only iid samples  $x_1, \dots, x_n$  from  $\mathcal{D}$ . In this work, we are interested in the case where the queries may be adaptively and adversarially chosen.

Specifically,  $\mathcal{O}$  is a stateful algorithm that holds a tuple of samples  $x_1, x_2, \dots \in \{0, 1\}^*$ , takes a statistical query  $q$  as input, and returns a real-valued answer  $a \in [0, 1]$ . We require that when  $x$  consists of iid samples from  $\mathcal{D}$ , the answer  $a$  is close to  $q(\mathcal{D})$ , and moreover that this condition holds for every query in an adaptively chosen sequence  $q_1, q_2, \dots$ . Formally, we define the accuracy guarantee using the following game with a stateful adversary  $\mathcal{A}$ .

$\mathcal{A}$  chooses a distribution  $\mathcal{D}$  over  $\{0, 1\}^d$   
 Sample  $x_1, \dots, x_n \leftarrow_{\mathcal{R}} \mathcal{D}$ , let  $x = (x_1, \dots, x_n)$   
 For  $j = 1, \dots, k$   
 $\mathcal{A}(q_1, a_1, \dots, q_{j-1}, a_{j-1})$  outputs a query  $q_j$   
 $\mathcal{O}(x, q_j)$  outputs  $a_j$

Figure 1.  $\text{Acc}_{n,d,k}[\mathcal{O}, \mathcal{A}]$

**Definition II.1** (Accurate Oracle). An oracle  $\mathcal{O}$  is  $(\alpha, \beta)$ -accurate for  $k$  adaptively chosen queries given  $n$  samples in  $\{0, 1\}^d$  if for every adversary  $\mathcal{A}$ ,

$$\Pr_{\text{Acc}_{n,d,k}[\mathcal{O}, \mathcal{A}]} [\forall j \in [k] \quad |\mathcal{O}(x, q_j) - q_j(\mathcal{D})| \leq \alpha] \geq 1 - \beta.$$

As a shorthand, we will say that  $\mathcal{O}$  is  $\alpha$ -accurate for  $k$  queries if for every  $n, d \in \mathbb{N}$ ,  $\mathcal{O}$  is  $(\alpha, o_n(1))$ -accurate for  $k$  queries given  $n$  samples in  $\{0, 1\}^d$ . Here,  $k$  may depend on  $n$  and  $d$  and  $o_n(1)$  is a function of  $n$  that tends to 0.

We are interested in oracles that are both accurate and computationally efficient. We say that an oracle  $\mathcal{O}$  is *computationally efficient* if when given samples  $x_1, \dots, x_n \in \{0, 1\}^d$  and a query  $q : \{0, 1\}^d \rightarrow \{0, 1\}$  it runs in time  $\text{poly}(n, d, |q|)$ . Here  $q$  will be represented as a circuit that evaluates  $q(x)$  and  $|q|$  denotes the size of this circuit.

### B. Encryption Schemes

Our attack relies on the existence of a semantically secure private-key encryption scheme. An encryption scheme is

a triple of efficient algorithms  $(Gen, Enc, Dec)$  with the following syntax:

- $Gen$  is a randomized algorithm that takes as input a security parameter  $\lambda$  and outputs an  $\ell_{Enc}(\lambda)$ -bit secret key for some non-decreasing function  $\ell_{Enc} : \mathbb{N} \rightarrow \mathbb{N}$ . Formally,  $sk \leftarrow_r Gen(1^\lambda)$ .
- $Enc$  is a randomized algorithm that takes as input a secret key and a one-bit message  $m \in \{0, 1\}$  and outputs a ciphertext  $c$ . Formally,  $c \leftarrow_r Enc(sk, m)$ .
- $Dec$  is a deterministic algorithm that takes as input a secret key and a ciphertext  $c$  and outputs a decryption  $m'$ . If the ciphertext  $c$  was an encryption of  $m$  under the key  $sk$ , then  $m' = m$ . Formally, if  $c \leftarrow_r Enc(sk, m)$ , then  $Dec(sk, c) = m$  with probability 1.

Roughly, security of the encryption scheme asserts that no polynomial time adversary who does not know the secret key can distinguish encryptions of  $m = 0$  from encryptions of  $m = 1$ , even if the adversary has access to an oracle that returns the encryption of an arbitrary message under the unknown key. For convenience, we will require that this security property holds simultaneously for an arbitrary polynomial number of secret keys. The existence of an encryption scheme with this property follows immediately from the existence of an ordinary semantically secure encryption scheme, which in turn follows from the existence of one-way functions. We start with the stronger definition only to simplify our proofs. A secure encryption scheme exists under the minimal cryptographic assumption that one-way functions exist. Due to space requirements, we omit the formal definition of security.

### C. Fingerprinting Codes

Collusion-resilient fingerprinting codes were introduced by Boneh and Shaw [4] for the problem of watermarking digital content. A *fingerprinting code* is a pair of efficient algorithms  $(FPC.Gen, FPC.Trace)$ . The code generator  $FPC.Gen$  takes a number of users  $p$  as input and outputs a matrix  $F \in \{0, 1\}^{p \times \ell_{FPC}(p)}$ , for some function  $\ell_{FPC} : \mathbb{N} \rightarrow \mathbb{N}$ . We think of  $F$  as consisting of  $p$  codewords, one for each user  $i \in [p]$ , with each codeword being of length  $\ell_{FPC} = \ell_{FPC}(p)$ . For a subset of users  $S \subseteq [p]$ , we use  $F_S$  to denote the  $|S| \times \ell_{FPC}$  matrix consisting of the subset of codewords belonging to users in  $S$ .

The security property says that any codeword can be “traced” to its corresponding user. Moreover, the code is *fully collusion-resilient*—if any subset of users  $S \subseteq [p]$  “combines” their codewords in an arbitrary manner, then the combined codeword  $a \in \{0, 1\}^{\ell_{FPC}}$  can also be traced to one of the users in  $S$ , provided that the combined codeword is “consistent” with  $F_S$  in a very weak sense. For the standard definition of fingerprinting codes, the consistency condition would require that for every column  $j$  of  $F_S$ , if every entry of the  $j$ -th column shares the same bit  $b$ , then the  $j$ -th entry of  $a$  is also  $b$ . Formally, we will use the condition that

for every  $j$ ,  $|a_j - \mathbb{E}_{i \in S} [F_S(i, j)]| \leq 1/3$ . For our results we require a stronger, *error-robust fingerprinting code*, that can trace combined codewords that only respect a relaxed consistency condition, in which the above constraint on  $a$  is only required to hold for 99% of columns  $j$ .

Specifically, for any set of codewords  $F_S$ , we define  $\text{Con}(F_S)$  to be the set of  $a \in \{0, 1\}^{\ell_{FPC}}$  such that for .99 $\ell_{FPC}$  choices of  $j$ ,  $|a_j - \mathbb{E}_{i \in S} [F_S(i, j)]| \leq 1/3$ . We can now formally define error-robust fingerprinting codes

**Definition II.2.** For a function  $\ell_{FPC} : \mathbb{N} \rightarrow \mathbb{N}$ , a pair of efficient algorithms  $(FPC.Gen, FPC.Trace)$  is an error-robust fingerprinting code of length  $\ell_{FPC}$  if

- 1) for every  $p \in \mathbb{N}$ ,  $FPC.Gen(1^p)$  outputs a matrix  $F \in \{0, 1\}^{p \times \ell_{FPC}(p)}$  and
- 2) for every (possibly randomized) adversary  $\mathcal{A}_{FPC}$  and every  $S \subseteq [p]$ , if  $F \leftarrow_r FPC.Gen(1^p)$  and  $a \leftarrow_r \mathcal{A}_{FPC}(F_S)$ , then

$$\Pr \left[ \begin{array}{l} a \in \text{Con}(F_S) \\ \wedge FPC.Trace(F, a) \notin S \end{array} \right] \leq \text{negl}(p).$$

Bun, Ullman, and Vadhan [5] introduced error-robust fingerprinting codes. They gave a construction with nearly-optimal length, building on the nearly-optimal construction of standard (non-robust) fingerprinting codes by Tardos [14].

**Theorem II.3** ([5], building on [14]). *For every  $p \in \mathbb{N}$ , there exists an error-robust fingerprinting code of length  $\ell_{FPC}(p) = \tilde{\Theta}(p^2)$ .*

For our results, we will need an additional technical lemma about the fingerprinting code in [5] that we will use for our results. The lemma states that if  $|S|$  is at least a sufficiently large constant, then for most columns  $j$ , the mean of the  $j$ -th column of  $F$  and that of  $F_S$  are close. In order to prove the lemma, we need to partially describe the algorithm  $FPC.Gen$ .

Choose parameters  $p_1, \dots, p_{\ell_{FPC}} \in [0, 1]$ . (The parameters  $p_j$  must be chosen from a particular distribution, which is not relevant for our purposes)  
For every  $i \in [p], j \in [\ell_{FPC}]$ , let  $F(i, j) = 1$  with probability  $p_j$ , independently

Figure 2.  $FPC.Gen$

**Lemma II.4.** *For every  $p \geq 500$ , and every  $S \subseteq [p]$  such that  $|S| \geq 500$ , if  $F \leftarrow_r FPC.Gen(1^p)$  then with probability at least  $1 - \text{negl}(p)$ , it holds that for .99 $\ell_{FPC}$  choices of  $j$ ,  $|\mathbb{E}_{i \in [p]} [F(i, j)] - \mathbb{E}_{i \in S} [F_S(i, j)]| \leq 1/6$ .*

*Proof:* Fix any  $p_1, \dots, p_{\ell_{FPC}} \in [0, 1]$  and any  $j \in [\ell_{FPC}]$ . Then for every  $i \in [S]$ ,  $F(i, j)$  is an independent Bernoulli random variable with success probability  $p_j$ . Thus,

by a Hoeffding bound<sup>2</sup>

$$\Pr \left[ \left| \mathbb{E}_{i \in [p]} [F(i, j)] - p_j \right| > 1/12 \right] \leq 2e^{-p/72} \leq 1/400$$

and

$$\Pr \left[ \left| \mathbb{E}_{i \in S} [F(i, j)] - p_j \right| > 1/12 \right] \leq 2e^{-|S|/72} \leq 1/400.$$

Thus, by a triangle inequality, it holds that

$$\Pr \left[ \left| \mathbb{E}_{i \in S} [F(i, j)] - \mathbb{E}_{i \in [p]} [F(i, j)] \right| > 1/6 \right] \leq 1/200$$

If we let  $B_j$  be the indicator variable of the event  $\{|\mathbb{E}_{i \in S} [F(i, j)] - \mathbb{E}_{i \in [p]} [F(i, j)]| > 1/6\}$ , then  $\mathbb{E} \left[ \sum_{j=1}^{\ell_{FPC}} B_j \right] \leq \ell_{FPC}/200$ . Since the parameters  $p_j$  are fixed, the events  $B_j$  are independent. By a Hoeffding bound,  $\Pr \left[ \sum_{j=1}^{\ell_{FPC}} B_j > \ell_{FPC}/100 \right] \leq e^{-\ell_{FPC}/20000} \leq \text{negl}(p)$ . The final inequality holds because  $\ell_{FPC} = \ell_{FPC}(p) = \Theta(p^2)$ . Since the conclusion holds for every fixed choice of parameters  $p_j$ , it also holds when the parameters  $p_j$  are chosen randomly as in  $FPC.Gen$ . ■

### III. LOWER BOUNDS FOR ADAPTIVELY CHOSEN STATISTICS

In this section we will show that there is no computationally efficient oracle that is accurate for a sufficiently large number of adaptively chosen queries, and thereby formally establish Theorem I.1 in the introduction. To do, we will construct an adversary that chooses a distribution  $\mathcal{D}$ , and then issues queries to the oracle such that no computationally efficient oracle given samples from  $\mathcal{D}$  can answer all the queries correctly.

The adversary is specified in Figure 3. The adversary works in three phases. In the first phase the adversary chooses the distribution  $\mathcal{D}$  randomly. Then the oracle is given samples from  $\mathcal{D}$  and the adversary performs a recovery phase in order to identify (most of) the samples the oracle received. Finally, the adversary uses knowledge of (most of) the samples to find a query that the oracle cannot answer accurately. See Section I-A for more informal description of the adversary. In Figure 3,  $(Gen, Enc, Dec)$  is an encryption scheme with key length  $\ell_{Enc}(\lambda)$  and  $(FPC.Gen, FPC.Trace)$  is a fingerprinting code of length  $\ell_{FPC}(p)$ . Observe that  $\text{Attack}_{n,d}$  is only well defined for pairs  $n, d \in \mathbb{N}$  such that there exists  $\lambda \in \mathbb{N}$  for which  $\ell_{Enc}(1) + \lceil \log(2000n) \rceil \leq d$ . Through this section we will assume that  $n = n(d) = \text{poly}(d)$  and that  $d$  is a sufficiently large constant, which ensures that  $\text{Attack}_{n,d}$  is well defined.

<sup>2</sup>For independent random variables  $X_1, \dots, X_m \in [0, 1]$ , if  $X = \frac{1}{m} \sum_{i=1}^m X_i$ , then  $\Pr[X > \mathbb{E}[X] + \tau] \leq \exp(-2\tau^2 m)$  and  $\Pr[|X - \mathbb{E}[X]| > \tau] \leq 2 \exp(-2\tau^2 m)$ .

#### The distribution $\mathcal{D}$ :

Given parameters  $d, n$ , let  $p = 2000n$ , let  $R = n - 500$ . Let  $\lambda$  be the largest integer such that  $\ell_{Enc}(\lambda) + \lceil \log p \rceil \leq d$ . For  $i \in [p]$ , let  $sk_i \leftarrow_r Gen(1^\lambda)$  and let  $y_i = (i, sk_i)$ . Let  $\mathcal{D}$  be the uniform distribution over  $\{y_1, \dots, y_p\} \subseteq \{0, 1\}^d$ .

Choose samples  $x_1, \dots, x_n \leftarrow_r \mathcal{D}$ , let  $x = (x_1, \dots, x_n)$ . Let  $S \subseteq [p]$  be the set of unique indices  $i$  such that  $(i, sk_i)$  appears in  $x$ .

#### Recovery phase:

Setup fingerprinting codes and ciphertexts:

Let  $F^1, \dots, F^R \leftarrow_r FPC.Gen(1^p)$ , let  $\ell_{FPC} = \ell_{FPC}(p)$  be code length. For  $r = 1, \dots, R$ ,  $i = 1, \dots, p$ ,  $j = 1, \dots, \ell_{FPC}$ , let  $c^r(i, j) = Enc(sk_i, F^r(i, j))$ .

Iteratively reconstruct  $S$ :

Let  $T^0 = \emptyset$

For round  $r = 1, \dots, R$ :

For  $j = 1, \dots, \ell_{FPC}$ :

Define the query  $q_j^r(i', sk')$  to be

$Dec(sk', c^r(i', j))$  if  $i' \notin T^{r-1}$  and 0 otherwise

Let  $a_j^r = \mathcal{O}(x; q_j^r)$

Let  $a^r = (a_1^r, \dots, a_{\ell_{FPC}}^r)$

Let  $i^r = FPC.Trace(F^r, a^r)$ , and let  $T^r = T^{r-1} \cup \{i^r\}$

#### Attack phase:

Let  $\phi = 0$  with probability  $1/2$  and  $\phi = 1/500$  with probability  $1/2$ . Sample a random subset  $B \subseteq [p]$  of size  $\phi \cdot p$ . Let  $m_i = 1$  for all  $i \in B$  and  $m_i = 0$  for all  $i \in [p] \setminus B$ . Let  $c_i^* = Enc(sk_i, m_i)$  for all  $i \in [p]$ . Define the query  $q^*(i', sk')$  to be  $Dec(sk', c_{i'}^*)$  if  $i' \notin T^R$  and 0 otherwise. Let  $a^* = \mathcal{O}(x, q^*)$ .

Figure 3.  $\text{Attack}_{n,d}[\mathcal{O}]$

#### A. Analysis of the Recovery Phase

The goal of the recovery phase of the algorithm is to identify most of the samples  $x_1, \dots, x_n$  that are held by the oracle. Once the attacker has this information, he can use it to find queries that distinguish the oracle's keys from the population and force the oracle to be inaccurate.

In order to recover keys, the attacker will force the oracle to give answers that are consistent with the fingerprinting codes  $F^1, \dots, F^R$ , which are then given to  $FPC.Trace$  to recover an element of the sample. Our first claim establishes that an accurate oracle will indeed force the oracle to give answers consistent with the fingerprinting codes.

**Claim III.1.** *If  $\mathcal{O}$  is  $(1/12)$ -accurate for  $n \cdot \ell_{FPC}(2000n) + 1$  adaptively chosen queries then for every polynomial  $n = n(d)$  and every sufficiently large  $d \in \mathbb{N}$ , in the execution of  $\text{Attack}_{n,d}[\mathcal{O}]$ , with probability at least  $1 - o(1)$ , it holds that  $\forall r \in [R]$ , for  $.99\ell_{FPC}$  choices of  $j \in [\ell_{FPC}]$ ,  $|\mathcal{O}(x, q_j^r) - \mathbb{E}_{i \in S \setminus T^{r-1}} [F^r(i, j)]| \leq 1/3$ .*

*Proof:* First we show that,

$$\Pr_{\text{Attack}} \left[ \begin{array}{c} \forall r \in [R], j \in [\ell_{FPC}], \\ |\mathcal{O}(x, q_j^r) - \mathbb{E}_{i \in [p]} [F^r(i, j)]| \leq 1/6 \end{array} \right] \geq 1 - o(1)$$

Observe that by definition, for every  $r, j$ ,

$$\begin{aligned} & \mathbb{E}_{(i, sk_i) \leftarrow \mathcal{D}} [q_j^r(i, sk_i)] \\ &= \frac{1}{p} \sum_{i \in [p] \setminus T^{r-1}} \text{Dec}(sk_i, c^r(i, j)) + \frac{1}{p} \sum_{i \in T^{r-1}} 0 \\ &= \frac{1}{p} \sum_{i \in [p] \setminus T^{r-1}} F^r(i, j) \end{aligned}$$

where the last equality follows by construction, since  $c^r(i, j) \leftarrow_{\mathcal{R}} \text{Enc}(sk_i, F^r(i, j))$ . Since  $|T^{r-1}| \leq n$  and  $F^r(i, j) \in \{0, 1\}$ , for every  $r, j$ ,

$$\left| \mathbb{E}_{(i, sk_i) \leftarrow \mathcal{D}} [q_j^r(i, sk_i)] - \mathbb{E}_{i \in [p]} [F^r(i, j)] \right| \leq \frac{n}{p} \leq \frac{1}{12} \quad (1)$$

In  $\text{Attack}$ , the oracle's input  $x$  consists of  $n$  samples from  $\mathcal{D}$ . Moreover, the total number of queries issued to the oracle is at most  $k = n \cdot \ell_{FPC}(2000n) + 1$ . Since the oracle is assumed to be  $(1/12)$ -accurate for  $k$  queries given  $n$  samples in  $\{0, 1\}^d$ ,

$$\Pr_{\text{Attack}} \left[ \begin{array}{c} \forall r, j \\ |\mathcal{O}(x, q_j^r) - \mathbb{E}_{x \leftarrow \mathcal{D}} [q_j^r(i, sk_i)]| \leq \frac{1}{12} \end{array} \right] \geq 1 - o(1) \quad (2)$$

Applying the triangle inequality to (1) and (2), this shows

$$\Pr_{\text{Attack}} \left[ \begin{array}{c} \forall r, j \\ |\mathcal{O}(x, q_j^r) - \mathbb{E}_{i \in [p]} [F^r(i, j)]| \leq 1/6 \end{array} \right] \geq 1 - o(1). \quad (3)$$

By Lemma II.4, since  $|S \setminus T^{r-1}| \geq 500$ , for every  $r$ , if  $F^r \leftarrow_{\mathcal{R}} \text{FPC.Gen}(1^p)$ , then with probability at least  $1 - \text{negl}(n)$ , it holds that for  $.99\ell_{FPC}$  choices of  $j$ ,  $|\mathbb{E}_{i \in [p]} [F^r(i, j)] - \mathbb{E}_{i \in S} [F_S^r(i, j)]| \leq 1/6$ . where the probability is taken over the choice of  $F^r \leftarrow_{\mathcal{R}} \text{FPC.Gen}$ . By a union bound over  $r = 1, \dots, R$ , where  $R = n - 500$ , if  $F^1, \dots, F^R \leftarrow_{\mathcal{R}} \text{FPC.Gen}$ , then with probability at least  $1 - \text{negl}(n)$ , it holds that for every  $r$ , for  $.99\ell_{FPC}$  choices of  $j$ ,  $|\mathbb{E}_{i \in [p]} [F^r(i, j)] - \mathbb{E}_{i \in S} [F_S^r(i, j)]| \leq 1/6$ . The claim now follows by combining this last observation with (3). ■

Claim III.1, establishes that in every round  $r$ , the oracle  $\mathcal{O}$  holding  $x$  returns a set of answers that are consistent with the fingerprinting code  $F_{S \setminus T^{r-1}}^r$ . However, this fact alone is not enough to guarantee that  $\text{FPC.Trace}$  returns a user in  $S \setminus T^{r-1}$ , because the queries to the oracle depend on rows

of  $F^r$  for users outside of  $S \setminus T^{r-1}$ , whereas the security of the fingerprinting code applies only to algorithms that only have access to the rows of  $F^r$  for users in  $S \setminus T^{r-1}$ . To remedy this problem we rely on the fact that the rows of  $F^r$  outside of  $S \setminus T^{r-1}$  are encrypted under keys  $sk$  that are not known to the oracle. Thus, a computationally efficient oracle “does not know” those rows. We can formalize this argument by comparing  $\text{Attack}$  to an  $\text{IdealAttack}$  where rows of  $F^r$  for users outside of  $S \setminus T^{r-1}$  are replaced with zeros, and argue that the adversary cannot distinguish between these two attacks without breaking the security of the encryption scheme.

#### The distribution $\mathcal{D}$ :

Given parameters  $d, n$ , let  $p = 2000n$ , let  $R = n - 500$ . Let  $\lambda$  be the largest integer such that  $\ell_{\text{Enc}}(\lambda) + \lceil \log p \rceil \leq d$ . For  $i \in [p]$ , let  $sk_i \leftarrow_{\mathcal{R}} \text{Gen}(1^\lambda)$  and let  $y_i = (i, sk_i)$ . Let  $\mathcal{D}$  be the uniform distribution over  $\{y_1, \dots, y_p\} \subseteq \{0, 1\}^d$ .

Choose samples  $x_1, \dots, x_n \leftarrow_{\mathcal{R}} \mathcal{D}$ , let  $x = (x_1, \dots, x_n)$ . Let  $S \subseteq [p]$  be the set of unique indices  $i$  such that  $(i, sk_i)$  appears in  $x$ .

#### Recovery phase:

Setup fingerprinting codes and ciphertexts:

Let  $F^1, \dots, F^R \leftarrow_{\mathcal{R}} \text{FPC.Gen}(1^p)$ , let  $\ell_{FPC} = \ell_{FPC}(p)$  be code length. For  $r = 1, \dots, R$ ,  $j = 1, \dots, \ell_{FPC}$ , for every  $i \in S$  let  $c^r(i, j) = \text{Enc}(sk_i, F^r(i, j))$ , and for every  $i \in [p] \setminus S$  let  $c^r(i, j) = \text{Enc}(sk_i, 0)$ .

Iteratively reconstruct  $S$ :

Let  $T^0 = \emptyset$

For round  $r = 1, \dots, R$ :

For  $j = 1, \dots, \ell_{FPC}$ :

Define the query  $q_j^r(i', sk')$  to be

$\text{Dec}(sk', c^r(i', j))$  if  $i' \notin T^{r-1}$  and 0 otherwise

Let  $a_j^r = \mathcal{O}(x; q_j^r)$

Let  $a^r = (a_1^r, \dots, a_{\ell_{FPC}}^r)$

Let  $i^r = \text{FPC.Trace}(F^r, a^r)$ , and let  $T^r = T^{r-1} \cup \{i^r\}$

#### Attack phase:

Let  $\phi = 0$  with probability  $1/2$  and  $\phi = 1/500$  with probability  $1/2$ . Sample a random subset  $B \subseteq [p]$  of size  $\phi \cdot p$ . Let  $m_i = 1$  for all  $i \in B$  and  $m_i = 0$  for all  $i \in [p] \setminus B$ . For every  $i \in S$ , let  $c_i^* = \text{Enc}(sk_i, m_i)$  and for every  $i \in [p] \setminus S$  let  $c_i^* = \text{Enc}(sk_i, 0)$ . Define the query  $q^*(i', sk')$  to be  $\text{Dec}(sk', c_i^*)$  if  $i' \notin T^R$  and 0 otherwise. Let  $a^* = \mathcal{O}(x, q^*)$ .

Figure 4.  $\text{IdealAttack}_{n,d}[\mathcal{O}]$

**Claim III.2.** Let  $Z_1$  be the event that for every  $r \in [R]$ , for  $.99\ell_{FPC}$  choices of  $j \in [\ell_{FPC}]$ ,

$$\left| \mathcal{O}(x, q_j^r) - \mathbb{E}_{i \in S \setminus T^{r-1}} [F^r(i, j)] \right| \leq 1/3.$$

Assume  $(Gen, Enc, Dec)$  is a computationally secure encryption scheme with key length  $\ell_{Enc}(\lambda) = \lambda$  and let  $n = n(d)$  be any polynomial. Then if  $\mathcal{O}$  is computationally efficient, for every sufficiently large  $d \in \mathbb{N}$ ,

$$\left| \Pr_{\text{IdealAttack}_{n,d}[\mathcal{O}]} [Z_1] - \Pr_{\text{Attack}_{n,d}[\mathcal{O}]} [Z_1] \right| \leq \text{negl}(d)$$

Note that if  $n(d) = \text{poly}(d)$ , then for every sufficiently large  $d$ , the pair  $n(d), d$  is reasonable. The proof follows straightforwardly from the security of the encryption scheme. Due to space requirements we omit the details.

**Claim III.3.** If  $\mathcal{O}$  is computationally efficient and  $(1/12)$ -accurate for  $n \cdot \ell_{FPC}(2000n) + 1$  adaptively chosen queries, then for every polynomial  $n = n(d)$ , and every sufficiently large  $d \in \mathbb{N}$ , in the execution of  $\text{IdealAttack}_{n,d}[\mathcal{O}]$ , with probability at least  $1 - o(1)$  it holds that for every  $r \in [R]$ , for  $.99\ell_{FPC}$  choices of  $j \in [\ell_{FPC}]$ ,

$$\left| \mathcal{O}(x, q_j^r) - \mathbb{E}_{i \in S \setminus T^{r-1}} [F^r(i, j)] \right| \leq 1/3.$$

The proof is immediate by combining Claim III.1 and Claim III.2. Next, we argue that in  $\text{IdealAttack}$ , with high probability  $FPC.Trace$  only outputs users contained in the sample  $S$ .

**Claim III.4.** If  $\mathcal{O}$  is computationally efficient and  $(1/12)$ -accurate for  $n \cdot \ell_{FPC}(2000n) + 1$  adaptively chosen queries, then for any polynomial  $n = n(d)$ , and every sufficiently large  $d \in \mathbb{N}$ ,  $\Pr_{\text{IdealAttack}_{n,d}[\mathcal{O}]} [|S \setminus T^R| > 500] \leq o(1)$ .

*Proof:* Fix any round  $r \in \{1, \dots, |S| - 500\}$  and let  $U = S \setminus T^{r-1}$ . By the security of the fingerprinting code, for every algorithm  $\mathcal{A}$ , if  $F^r \leftarrow_r FPC.Gen(1^p)$ , then

$$\Pr \left[ \begin{array}{l} \mathcal{A}(F_U^r) \in \text{Con}(F_U^r) \\ \wedge FPC.Trace(F^r, \mathcal{A}(F_U^r)) \notin U \end{array} \right] \leq \text{negl}(n)$$

Observe that in  $\text{IdealAttack}$ , the oracle  $\mathcal{O}$  is never given any input that depends on rows of  $F^r$  that belong to users outside of  $U$ : The queries issued in rounds  $r' \neq r$  depend only on  $F^{r'}$ , which is independent from  $F^r$ . And in round  $r$  the query only depends on ciphertexts  $c^r(i, j)$  for  $i \notin T^{r-1}$ , which are all independent of  $F^r(i, j)$  whenever  $i \notin U$ . Therefore, if  $F^r \leftarrow_r FPC.Gen(1^p)$ , we have

$$\Pr \left[ \begin{array}{l} a^r \in \text{Con}(F_U^r) \\ \wedge FPC.Trace(F^r, a^r) \notin U \end{array} \right] \leq \text{negl}(n) \quad (4)$$

By a union bound over  $r = 1, \dots, |S| - 500$ , we also have that if  $F^1, \dots, F^{|S|-500} \leftarrow_r FPC.Gen(1^p)$ , then

$$\Pr \left[ \exists r \left( \begin{array}{l} a^r \in \text{Con}(F_{S \setminus T^{r-1}}^r) \\ \wedge FPC.Trace(F^r, a^r) \notin S \setminus T^{r-1} \end{array} \right) \right] \leq \text{negl}(n) \quad (5)$$

By Claim III.3, we have that with probability at least  $1 - o(1)$ , it holds that for all  $r \in \{1, \dots, |S| - 500\}$  and for  $.99\ell_{FPC}$  choices of  $j \in [\ell_{FPC}]$ ,  $|a_j^r - \mathbb{E}_{i \in U} [F^r(i, j)]| \leq 1/3$ . Note that in order to apply Claim III.3 we have used the fact that when  $r \in \{1, \dots, |S| - 500\}$ ,  $|S \setminus T^{r-1}| \geq 500$ . If this condition is satisfied, then indeed  $a^r \in \text{Con}(F_U^r)$ . Therefore, combining with (5), we have that if  $F^1, \dots, F^{|S|-500} \leftarrow_r FPC.Gen(1^p)$ , then

$$\Pr [\exists r, FPC.Trace(F^r, a^r) \notin S \setminus T^{r-1}] \leq o(1)$$

Now the claim follows by observing that

$$\Pr_{\text{IdealAttack}_{n,d}[\mathcal{O}]} [|S \setminus T^R| > 500] \leq \Pr_{\text{IdealAttack}_{n,d}[\mathcal{O}]} \left[ \begin{array}{l} \exists r \in \{1, \dots, |S| - 500\} \\ FPC.Trace(F^r, a^r) \notin S \setminus T^{r-1} \end{array} \right] \leq o(1)$$

Finally, we show that if  $|S \setminus T^R| \leq 500$  with high probability in  $\text{IdealAttack}$ , then  $|S \setminus T^R| \leq 500$  with high probability in  $\text{Attack}$ . Again, we do so by arguing that  $\text{Attack}$  and  $\text{IdealAttack}$  are computationally indistinguishable. ■

**Claim III.5.** Let  $Z_2$  be the event  $\{|S \setminus T^R| \leq 500\}$ . Assume  $(Gen, Enc, Dec)$  is a computationally secure encryption scheme with key length  $\ell_{Enc}(\lambda) = \lambda$  and let  $n = n(d)$  be any polynomial. Then if  $\mathcal{O}$  is computationally efficient, for every sufficient large  $d \in \mathbb{N}$ ,

$$\left| \Pr_{\text{IdealAttack}_{n,d}[\mathcal{O}]} [Z_2] - \Pr_{\text{Attack}_{n,d}[\mathcal{O}]} [Z_2] \right| \leq \text{negl}(d).$$

The proof follows straightforwardly from the security of the encryption scheme. We omit the details.

**Lemma III.6.** If  $\mathcal{O}$  is computationally efficient and  $(1/12)$ -accurate for  $n \cdot \ell_{FPC}(2000n) + 1$  adaptively chosen queries, then for any polynomial  $n = n(d)$ , and every sufficiently large  $d \in \mathbb{N}$ ,

$$\Pr_{\text{Attack}_{n,d}[\mathcal{O}]} [|S \setminus T^R| > 500] = o(1)$$

The proof is immediate from Claims III.4 and III.5.

## B. Analysis of the Attack Phase

By the arguments of Section III-A, we know that if the oracle is computationally efficient and accurately answers all the queries in the recovery phase, then with high probability  $|S \setminus T^R| \leq 500$ . In this section we will show that if these events indeed occur, then the probability that the oracle answers  $q^*$  accurately in the attack phase is bounded away from 1 by a constant. Since an accurate oracle is required to answer each query accurately with probability at least  $1 - o(1)$ , we will have obtained a contradiction.

To begin with we show that the population answer  $q^*(\mathcal{D})$  is close to the value  $\phi$  in the real attack.



**Claim III.7.** For every reasonable  $n, d \in \mathbb{N}$ , in the execution of  $\text{Attack}_{n,d}[\mathcal{O}]$ , it holds that  $|q^*(\mathcal{D}) - \phi| \leq 1/2000$ .

*Proof:* The case  $\phi = 0$  we have  $q^*(\mathcal{D}) = 0$ . If  $\phi = 1/500$ , then we have that  $\mathbb{E}_{i \in [p]} m_i = \phi$  since  $|B| = \phi \cdot p$ . Hence,

$$q^*(\mathcal{D}) = \mathbb{E}_{(i,sk) \sim \mathcal{D}} q^*(i, sk) = \phi - \Pr_{i \leftarrow_{\mathbb{R}} [p]} [i \in T^R \wedge m_i = 1]$$

where  $\Pr [i \in T^R \wedge m_i = 1] \leq \Pr [i \in T^R] = |T^R|/p = (n - 500)/p \leq 1/2000$ . ■

We will show that the oracle cannot guess the value of  $\phi$  with sufficiently high probability in  $\text{IdealAttack}$  provided that the recovery phase succeeded.

**Claim III.8.** For every reasonable  $n, d \in \mathbb{N}$ ,

$$\Pr_{\text{IdealAttack}_{n,d}[\mathcal{O}]} \left[ \begin{array}{l} |S \setminus T^R| \leq 500 \\ \wedge |\mathcal{O}(x, q^*) - \phi| \leq \frac{1}{2000} \end{array} \right] \leq 1 - \Omega(1)$$

*Proof:* Consider the case where  $\phi = 1/500$ . A straightforward calculation shows that

$$\Pr [\forall i \in S \setminus T^R, m_i = 0 \mid |S \setminus T^R| \leq 500] \geq \frac{1}{4e^2}$$

On the other hand, when  $\phi = 0$ , we have  $\Pr [\forall i \in S \setminus T^R, m_i = 0 \mid |S \setminus T^R| \leq 500] = 1$ . Note that in  $\text{IdealAttack}$  the oracle only sees  $m_i$  for  $i \in S \setminus T^R$ . When the oracle sees only that  $m_i = 0$  for every  $i \in S \setminus T^R$ , it cannot give an answer that is simultaneously accurate to within  $1/2000$  for both the case of  $\phi = 0$  and for the case of  $\phi = 1/500$ . The event  $m_i = 0$  for every  $i \in S \setminus T^R$  occurs with at least probability  $1/4e^2$  as shown above. Conditioned on this event, both cases  $\phi = 0$  and  $\phi = 1/500$  have constant probability. Hence, the answer of the oracle must be far from  $\phi$  with constant probability. Formally, in the execution of  $\text{IdealAttack}$ ,

$$\Pr \left[ |\mathcal{O}(x, q^*) - \phi| \leq \frac{1}{2000} \mid |S \setminus T^R| \leq 500 \right] \leq 1 - \Omega(1).$$

By Claim III.4, this implies

$$\Pr_{\text{IdealAttack}} \left[ \begin{array}{l} |\mathcal{O}(x, q^*) - \phi| \leq \frac{1}{2000} \\ \wedge |S \setminus T^R| \leq 500 \end{array} \right] \leq \frac{1 - \Omega(1)}{1 - o(1)} \leq 1 - \Omega(1). \quad \blacksquare$$

As in the analysis of the recovery phase, we will first argue that the probability the oracle is accurate for  $q^*$  in  $\text{Attack}$  is nearly the same as it is in an  $\text{IdealAttack}$  where the query  $q^*$  has been modified to contain no information about users outside of  $S$ .

**Claim III.9.** Let  $Z_3$  be the event

$$\left\{ |S \setminus T^R| \leq 500 \wedge |\mathcal{O}(x, q^*) - \phi| \leq \frac{1}{2000} \right\}.$$

Let  $(\text{Gen}, \text{Enc}, \text{Dec})$  is a computationally secure encryption scheme with key length  $\ell_{\text{Enc}}(\lambda) = \lambda$  and  $n = n(d)$  be any

polynomial. Then if  $\mathcal{O}$  is computationally efficient, for every sufficiently large  $d \in \mathbb{N}$ ,

$$\left| \Pr_{\text{IdealAttack}_{n,d}[\mathcal{O}]} [Z_3] - \Pr_{\text{Attack}_{n,d}[\mathcal{O}]} [Z_3] \right| \leq \text{negl}(d)$$

The proof follows straightforwardly from the security of the encryption scheme. We omit the details.

**Lemma III.10.** If  $\mathcal{O}$  is computationally efficient, then for any polynomial  $n = n(d)$ , and every sufficiently large  $d \in \mathbb{N}$ ,

$$\Pr_{\text{Attack}_{n,d}[\mathcal{O}]} \left[ \begin{array}{l} |S \setminus T^R| \leq 500 \\ \wedge |\mathcal{O}(x, q^*) - q^*(\mathcal{D})| \leq \frac{1}{2000} \end{array} \right] \leq 1 - \Omega(1)$$

*Proof:* By Claim III.7, we have that  $|q^*(\mathcal{D}) - \phi| \leq 1/2000$ . Combining Claim III.8 with Claim III.9, we further have

$$\Pr_{\text{Attack}} \left[ \begin{array}{l} |S \setminus T^R| \leq 500 \\ \wedge |\mathcal{O}(x, q^*) - \phi| \leq \frac{1}{2000} \end{array} \right] \leq 1 - \Omega(1).$$

The statement of the lemma now follows from a triangle inequality. ■

### C. Putting it Together

We can now prove Theorem I.1 from the introduction.

**Theorem III.11.** There is no computationally efficient oracle  $\mathcal{O}$  that is  $(1/2000)$ -accurate for  $n^{3+o(1)}$  adaptively chosen queries given  $n$  samples.

*Proof:* The entire attack will consist of  $k = R \cdot \ell_{\text{FPC}} + 1$  queries, where  $R = n - 500$  and  $\ell_{\text{FPC}} = \tilde{O}(p^2) = \tilde{O}(n^2)$ . Therefore the entire attack consists of  $k = n^{3+o(1)}$  queries.

Therefore, if  $\mathcal{O}$ , by Lemma III.6,

$$\Pr_{\text{Attack}_{n,d}} [|S \setminus T^R| > 500] = o_n(1)$$

By Lemma III.10,

$$\Pr_{\text{Attack}_{n,d}} \left[ \begin{array}{l} |S \setminus T^R| \leq 500 \\ \wedge |\mathcal{O}(x, q^*) - q^*(\mathcal{D})| \leq \frac{1}{2000} \end{array} \right] \leq 1 - \Omega(1)$$

Combining these two statements gives

$$\Pr_{\text{Attack}_{n,d}[\mathcal{O}]} \left[ |\mathcal{O}(x, q^*) - q^*(\mathcal{D})| \leq \frac{1}{2000} \right] \leq 1 - \Omega(1)$$

However, this contradicts the definition of an accurate oracle. ■

### D. An Information-Theoretic Lower Bound

We now state a precise version of our information theoretic lower bound (Theorem I.2). Due to space requirements we omit the proof of this Theorem. The key observation in the proof is that the encryption scheme used in our attack only needs to satisfy a certain weak security property, which can be achieved even against *computationally unbounded* adversaries, provided that the secret key is sufficiently long. This requirement that the keys are long is why the attack only applies for sufficiently large dimension  $d$  relative to  $n$ .

**Theorem III.12.** *There is no oracle  $\mathcal{O}$  that for every  $n \in \mathbb{N}$  is  $(1/2000, o_n(1))$ -accurate for  $n^{3+o(1)}$  adaptively chosen queries given  $n$  samples in  $\{0, 1\}^d$  for  $d = d(n) = n^{3+o(1)}$ .*

#### IV. LOWER BOUNDS FOR AVOIDING BLATANT NON-PRIVACY

We now state our results that computationally efficient oracles that guarantee accuracy for adaptively chosen queries must be blatantly non-private. Due to space requirements we defer all proofs to the full version.

First, we need to define a notion of accuracy that is more appropriate for the application to privacy. In contrast to Definition II.1 where accuracy is defined with respect to the distribution, here we define accurate with respect to the sample itself.

**Definition IV.1.** An oracle  $\mathcal{O}$  is  $(\alpha, \beta)$ -sample-accurate for  $k$  adaptively chosen queries given  $n$  samples in  $\{0, 1\}^d$  if for every adversary  $\mathcal{A}_{\text{priv}}$ ,

$$\Pr_{\text{NonPrivacy}_{n,d,k}[\mathcal{O}, \mathcal{A}_{\text{priv}}]} \left[ \forall j \in [k] \left[ |\mathcal{O}(x, q_j) - q_j(x)| \leq \alpha \right] \right] \geq 1 - \beta.$$

As a shorthand, we will say that  $\mathcal{O}$  is  $\alpha$ -sample-accurate for  $k$  queries if for every  $n, d \in \mathbb{N}$ ,  $\mathcal{O}$  is  $(\alpha, o_n(1))$ -accurate for  $k$  queries given  $n$  samples in  $\{0, 1\}^d$ . Here,  $k$  may depend on  $n$  and  $d$  and  $o_n(1)$  is a function of  $n$  that tends to 0.

With this change in mind, we model blatant non-privacy via the following game.

$\mathcal{A}_{\text{priv}}$  chooses a set  $y = \{y_1, \dots, y_{2n}\} \subseteq \{0, 1\}^d$   
 Sample a random subset  $x \subseteq y$  of size  $n$   
 For  $j = 1, \dots, k$   
 $\mathcal{A}(q_1, a_1, \dots, q_{j-1}, a_{j-1})$  outputs a query  $q_j$   
 $\mathcal{O}(x, q_j)$  outputs  $a_j$   
 $\mathcal{A}_{\text{priv}}$  outputs a set  $x' \subseteq y$

Figure 5.  $\text{NonPrivacy}_{n,d,k}[\mathcal{O}, \mathcal{A}_{\text{priv}}]$

**Definition IV.2.** *Giving  $\alpha$ -accurate answers to  $k$  adaptively chosen queries is blatantly non-private for efficient oracles* if there exists an adversary  $\mathcal{A}_{\text{priv}}$  such that for every oracle  $\mathcal{O}$  that is computationally efficient and  $\alpha$ -sample-accurate for  $k$  adaptively chosen queries,

$$\Pr_{\text{NonPrivacy}_{n,d,k}[\mathcal{O}, \mathcal{A}_{\text{priv}}]} [|x \Delta x'| > n/100] \leq o_n(1)$$

If the conclusion holds even for computationally inefficient oracles then we replace “for efficient oracles” with “for unbounded oracles” in the definition.

**Theorem IV.3.** *Giving accurate answers to  $n^{3+o(1)}$  adaptively chosen queries is blatantly non-private for computationally efficient oracles.*

We also obtain an information-theoretic analogue of our hardness result for avoiding blatant non-privacy.

**Theorem IV.4.** *Giving accurate answers to  $n^{3+o(1)}$  adaptively chosen queries on  $n$  samples of dimension  $d = n^{3+o(1)}$  is blatantly non-private for unbounded oracles.*

#### REFERENCES

- [1] C. Dwork, V. Feldman, M. Hardt, T. Pitassi, O. Reingold, and A. Roth, “Guilt-free data exploration (tentative title),” *Manuscript*, 2014.
- [2] Y. Benjamini and Y. Hochberg, “Controlling the false discovery rate: a practical and powerful approach to multiple testing,” *Journal of the Royal Statistical Society. Series B (Methodological)*, vol. 57, no. 1, pp. 289–300, 1995.
- [3] M. J. Kearns, “Efficient noise-tolerant learning from statistical queries,” in *STOC*, S. R. Kosaraju, D. S. Johnson, and A. Aggarwal, Eds. ACM, 1993, pp. 392–401.
- [4] D. Boneh and J. Shaw, “Collusion-secure fingerprinting for digital data,” *IEEE Transactions on Information Theory*, vol. 44, no. 5, pp. 1897–1905, 1998.
- [5] M. Bun, J. Ullman, and S. P. Vadhan, “Fingerprinting codes and the price of approximate differential privacy,” in *STOC*. ACM, May 31 – June 3 2014, pp. 1–10.
- [6] C. Dwork, F. McSherry, K. Nissim, and A. Smith, “Calibrating noise to sensitivity in private data analysis,” in *TCC*, ser. Lecture Notes in Computer Science, S. Halevi and T. Rabin, Eds., vol. 3876. Springer, 2006, pp. 265–284.
- [7] J. Ullman, “Answering  $n^{2+o(1)}$  counting queries with differential privacy is hard,” in *STOC*, D. Boneh, T. Roughgarden, and J. Feigenbaum, Eds. ACM, 2013, pp. 361–370.
- [8] O. Bousquet and A. Elisseeff, “Stability and generalization,” *Journal of Machine Learning Research*, vol. 2, pp. 499–526, 2002.
- [9] A. Roth and T. Roughgarden, “Interactive privacy via the median mechanism,” in *Proc. 42nd Symposium on Theory of Computing (STOC)*. ACM, 2010, pp. 765–774.
- [10] M. Hardt and G. Rothblum, “A multiplicative weights mechanism for privacy-preserving data analysis,” in *Proc. 51st Foundations of Computer Science (FOCS)*. IEEE, 2010.
- [11] B. Chor, A. Fiat, and M. Naor, “Tracing traitors,” in *CRYPTO*, ser. Lecture Notes in Computer Science, Y. Desmedt, Ed., vol. 839. Springer, 1994, pp. 257–270.
- [12] C. Dwork, M. Naor, O. Reingold, G. N. Rothblum, and S. P. Vadhan, “On the complexity of differentially private data release: efficient algorithms and hardness results,” in *STOC*. ACM, 2009, pp. 381–390.
- [13] M. Hardt and D. P. Woodruff, “How robust are linear sketches to adaptive inputs?” in *STOC*, D. Boneh, T. Roughgarden, and J. Feigenbaum, Eds. ACM, 2013, pp. 121–130.
- [14] G. Tardos, “Optimal probabilistic fingerprint codes,” *J. ACM*, vol. 55, no. 2, 2008.