

# On Differential Privacy and Adaptive Data Analysis with Bounded Space

Itai Dinur\*    Uri Stemmer†    David P. Woodruff‡    Samson Zhou§

February 11, 2023

## Abstract

We study the space complexity of the two related fields of *differential privacy* and *adaptive data analysis*. Specifically,

1. Under standard cryptographic assumptions, we show that there exists a problem  $P$  that requires exponentially more space to be solved efficiently with differential privacy, compared to the space needed without privacy. To the best of our knowledge, this is the first separation between the space complexity of private and non-private algorithms.
2. The line of work on adaptive data analysis focuses on understanding the number of *samples* needed for answering a sequence of adaptive queries. We revisit previous lower bounds at a foundational level, and show that they are a consequence of a space bottleneck rather than a sampling bottleneck.

To obtain our results, we define and construct an encryption scheme with multiple keys that is built to withstand a limited amount of key leakage in a very particular way.

## 1 Introduction

Query-to-communication lifting theorems allow translating lower bounds on the *query complexity* of a given function  $f$  to lower bounds on the *communication complexity* of a related function  $\hat{f}$ . Starting from the seminal work of Raz and McKenzie [74], several such lifting theorems were presented, and applied, to obtain new communication complexity lower bounds in various settings.

In the domain of cryptography, related results have been obtained, where the starting point is a lower bound on the *query complexity* of an adversary solving a cryptanalytic problem in an idealized model, such as the random oracle model [8]. The query complexity lower bound is then

---

\*Ben-Gurion University. [dinuri@bgu.ac.il](mailto:dinuri@bgu.ac.il). Partially supported by the Israel Science Foundation (grant 1903/20) and by the European Research Council under the ERC starting grant agreement no. 757731 (LightCrypt).

†Tel Aviv University and Google Research. [u@uri.co.il](mailto:u@uri.co.il). Partially supported by the Israel Science Foundation (grant 1871/19) and by Len Blavatnik and the Blavatnik Family foundation.

‡Carnegie Mellon University. [dwoodruf@andrew.cmu.edu](mailto:dwoodruf@andrew.cmu.edu). Work done in part while visiting Google Research. Also supported by a Simons Investigator Award and by the National Science Foundation under Grant No. CCF-1815840.

§UC Berkeley and Rice University. [samsonzhou@gmail.com](mailto:samsonzhou@gmail.com). Work done in part while at Carnegie Mellon University. Partially supported by a Simons Investigator Award and by the National Science Foundation under Grant No. CCF-1815840.

lifted to a *query-space* lower bound for a non-uniform (preprocessing) adversary solving the same problem [84, 25, 24].

Building on ideas developed in these lines of work, we present a new technique for translating sampling lower bounds to space lower bounds for problems in the context of *differential privacy* and *adaptive data analysis*. Before presenting our results, we motivate our settings.

## 1.1 Differential privacy

Differential privacy [35] is a mathematical definition for privacy that aims to enable statistical analyses of datasets while providing strong guarantees that individual-level information does not leak. Informally, an algorithm that analyzes data satisfies differential privacy if it is robust in the sense that its outcome distribution does not depend “too much” on any single data point. Formally,

**Definition 1.1** ([35]). *Let  $\mathcal{A} : X^* \rightarrow Y$  be a randomized algorithm whose input is a dataset  $D \in X^*$ . Algorithm  $\mathcal{A}$  is  $(\epsilon, \delta)$ -differentially private (DP) if for any two datasets  $D, D'$  that differ on one point (such datasets are called neighboring) and for any outcome set  $F \subseteq Y$  it holds that  $\Pr[\mathcal{A}(D) \in F] \leq e^\epsilon \cdot \Pr[\mathcal{A}(D') \in F] + \delta$ .*

To interpret the definition, let  $D$  be a dataset containing  $n$  data points, each of which represents the information of one individual. Suppose that Alice knows all but one of these data points (say Bob’s data point). Now suppose that we compute  $z \leftarrow \mathcal{A}(D)$ , and give  $z$  to Alice. If  $\mathcal{A}$  is differentially private, then Alice learns very little about Bob’s data point, because  $z$  would have been distributed roughly the same no matter what Bob’s data point is.

Over the last few years, we have witnessed an explosion of research on differential privacy in various settings. In particular, a fruitful line of work has focused on designing differentially private algorithms with small *space* complexity, mainly in streaming settings. Works in this vein include [36, 64, 9, 5, 14, 1, 59, 78, 69, 70, 13, 87, 10, 18, 85, 86, 90]. All of these works showed *positive* results and presented differentially private algorithms with small space complexity for various problems. In fact, some of these works showed that classical streaming algorithms are differentially private essentially *as is*. For example, Blocki et al. [9] showed that the Johnson-Lindenstrauss transform itself preserves differential privacy, and Smith et al. [78] showed this for the classical Flajolet-Martin Sketch.

In light of all these positive results, one might think that algorithms with small space are particularly suitable for differential privacy, because these algorithms are not keeping too much information about the input to begin with.

**Question 1.2.** *Does differential privacy require more space?*

### 1.1.1 Our results for differential privacy with bounded space

We answer Question 1.2 in the affirmative, i.e., we show that differential privacy *may require more space*. To this end, we come up with a problem that can be solved using a small amount of space without privacy, but requires a large amount of space to be solved with privacy. As a first step, let us examine the following toy problem, which provides *some* answer to the above question.

**A non-interesting toy problem.** Recall that  $F_2$  (the second frequency moment of a stream) estimation with multiplicative approximation error  $1 + \alpha$  has an  $\Omega(1/\alpha^2)$  space lower bound [88]. This immediately shows a separation for the problem of “output either the last element of the stream or a  $(1 + \alpha)$ -approximation to the  $F_2$  value of the stream”. In the non-private setting, the last element can be output using space independent of  $\alpha$ , but in the private setting the algorithm is forced to (privately) estimate  $F_2$  and thus use at least  $1/\alpha^2$  space. Of course, we could replace  $F_2$  with other tasks that have a large space lower bound in the standard non-private model.

We deem this toy problem non-interesting because, at a high level, our goal is to show that there are cases where computing something privately requires a lot more space than computing “the same thing” non-privately. In the toy problem, however, the private and non-private algorithms are arguably *not* computing “the same thing”. To reconcile this issue, we will focus on problems that are defined by a *function* (ranging over some metric space), and the desired task would be to approximate the value of this function. Note that this formulation disqualifies the toy problem from being a valid answer to Question 1.2, and that with this formulation there is a formal sense in which every algorithm for solving the task must compute (or approximate) “the same thing”.

Let us make our setting more precise. In order to simplify the presentation, instead of studying the streaming model, we focus on the following computation model.<sup>1</sup> Consider an algorithm that is instantiated on a dataset  $D$  and then aims to answer a query with respect to  $D$ . We say that such an algorithm has space  $s$  if, before it gets the query, it shrinks  $D$  to a *summary*  $z$  containing at most  $s$  bits. Then, when obtaining the query, the algorithm answers it using only the summary  $z$  (without additional access to the original dataset  $D$ ). Formally, we consider problems that are defined by a function  $P : X^* \times Q \rightarrow M$ , where  $X$  is the data domain,  $Q$  is a family of possible queries, and  $M$  is a metric space.

**Definition 1.3.** We say that  $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$  solves a problem  $P : X^* \times Q \rightarrow M$  with space complexity  $s$ , sample complexity  $n$ , error  $\alpha$ , and confidence  $\beta$  if

1.  $\mathcal{A}_1 : X^* \rightarrow \{0, 1\}^s$  is a preprocessing procedure that takes a dataset  $D$  and outputs an  $s$ -bit string.
2. For every input dataset  $D \in X^n$  and every query  $q \in Q$  it holds that

$$\Pr_{\substack{z \leftarrow \mathcal{A}_1(D) \\ a \leftarrow \mathcal{A}_2(z, q)}} [|a - P(D, q)| \leq \alpha] \geq 1 - \beta.$$

We show the following theorem.

**Theorem 1.4** (informal). Let  $d \in \mathbb{N}$  be a parameter controlling the size of the problem (specifically, data points from  $X$  can be represented using  $\text{polylog}(d)$  bits, and queries from  $Q$  can be represented using  $\text{poly}(d)$  bits). There exists a problem  $P : X^* \times Q \rightarrow M$  such that the following holds.

1.  $P$  can be solved non-privately using  $\text{polylog}(d)$  bits of space.  
% See Lemma 3.2 for the formal statement.
2.  $P$  can be solved privately using sample and space complexity  $\tilde{O}(\sqrt{d})$ .  
% See Lemma 3.4 for the formal statement.

---

<sup>1</sup>We remark, however, that all of our results extend to the streaming setting. See Remark 3.3.

3. Assuming the existence of a sub-exponentially secure symmetric-key encryption scheme, every computationally-efficient differentially-private algorithm  $\mathcal{A}$  for solving  $P$  must have space complexity  $\tilde{\Omega}(\sqrt{d})$ , even if its sample complexity is a large polynomial in  $d$ . Furthermore, this holds even if  $\mathcal{A}$  is only required to be computationally differentially private (namely, the adversary we build against  $\mathcal{A}$  is computationally efficient).

*% See Corollary 3.13 for the formal statement.*

Note that this is an exponential separation (in  $d$ ) between the non-private space complexity and the private space complexity. We emphasize that the hardness of privately solving  $P$  does not come from not having enough samples. Indeed, by Item 2,  $\tilde{O}(\sqrt{d})$  samples suffice for privately solving this problem. However, Item 3 states that unless the algorithm has large space, then it cannot privately solve this problem *even if it has many more samples than needed*.

To the best of our knowledge, this is the first result that separates the space complexity of private and non-private algorithms. Admittedly, the problem  $P$  we define to prove the above theorem is somewhat unnatural. In contrast, our negative results for adaptive data analysis (to be surveyed next) are for the canonical problem studied in the literature.

## 1.2 Adaptive data analysis

Consider a data analyst interested in testing a specific research hypothesis. The analyst acquires relevant data, evaluates the hypothesis, and (say) learns that it is false. Based on the findings, the analyst now decides on a second hypothesis to be tested, and evaluates it *on the same data* (acquiring fresh data might be too expensive or even impossible). That is, the analyst chooses the hypotheses *adaptively*, where this choice depends on previous interactions with the data. As a result, the findings are no longer supported by classical statistical theory, which assumes that the tested hypotheses are fixed before the data is gathered, and the analyst runs the risk of overfitting to the data.

Starting with [34], the line of work on *adaptive data analysis (ADA)* aims to design methods for provably guaranteeing statistical validity in such settings. Specifically, the goal is to design a mechanism  $\mathcal{A}$  that initially obtains a dataset  $D$  containing  $t$  i.i.d. samples from some unknown distribution  $\mathcal{P}$ , and then answers  $k$  *adaptively chosen queries* w.r.t.  $\mathcal{P}$ . Importantly,  $\mathcal{A}$ 's answers must be accurate w.r.t. the underlying distribution  $\mathcal{P}$ , and not just w.r.t. the empirical dataset  $D$ . The main question here is,

**Question 1.5.** *How many samples does  $\mathcal{A}$  need (i.e., what should  $t$  be) in order to support  $k$  such adaptive queries?*

As a way of dealing with worst-case analysts, the analyst is assumed to be adversarial in that it tries to cause the mechanism to fail. If a mechanism can maintain utility against such an adversarial analyst, then it maintains utility against any analyst. Formally, the canonical problem pursued by the line of work on ADA is defined as a two-player game between a mechanism  $\mathcal{A}$  and an adversary  $\mathcal{B}$ . See Figure 1.

**Definition 1.6** ([34]). *A mechanism  $\mathcal{A}$  is  $(\alpha, \beta)$ -accurate for  $k$  queries over a domain  $X$  using sample size  $t$  if for every adversary  $\mathcal{B}$  (interacting with  $\mathcal{A}$  in the game specified in Figure 1) it holds that*

$$\Pr [\exists j \in [k] \text{ s.t. } |z_j - h_j(\mathcal{P})| > \alpha] \leq \beta,$$

1. The adversary  $\mathcal{B}$  chooses a distribution  $\mathcal{P}$  over a data domain  $X$ .
2. The mechanism  $\mathcal{A}$  obtains a sample  $S \sim \mathcal{P}^t$  containing  $t$  i.i.d. samples from  $\mathcal{P}$ .
3. For  $k$  rounds  $j = 1, 2, \dots, k$ :
  - The adversary chooses a function  $h_j : X \rightarrow \{-1, 0, 1\}$ , possibly as a function of all previous answers given by the mechanism.
  - The mechanism obtains  $h_j$  and responds with an answer  $z_j$ , which is given to  $\mathcal{B}$ .

Figure 1: A two-player game between a mechanism  $\mathcal{A}$  and an adversary  $\mathcal{B}$ .

where  $h_j(\mathcal{P}) = \mathbb{E}_{x \sim \mathcal{P}}[h_j(x)]$  is the “true” value of  $h_j$  on the underlying distribution  $\mathcal{P}$ .

Following Dwork et al. [34], this problem has attracted a significant amount of work [53, 81, 4, 68, 61, 57, 77, 76, 44, 45, 33, 75, 26, 43, 42, 3], most of which is focused on understanding how many samples are needed for adaptive data analysis (i.e., focused on Question 1.5). In particular, the following almost matching bounds are known.

**Theorem 1.7** ([34, 4]). *There exists a computationally efficient mechanism that is (0.1, 0.1)-accurate for  $k$  queries using sample size  $t = \tilde{O}(\sqrt{k})$ .*

**Theorem 1.8** ([53, 81], informal). *Assuming the existence of one-way functions, every computationally efficient mechanism that is (0.1, 0.1)-accurate for  $k$  queries must have sample size at least  $t = \Omega(\sqrt{k})$ .*

### 1.2.1 Our results for adaptive data analysis with bounded space

All prior work on the ADA problem treated it as a *sampling* problem, conveying the message that “adaptive data analysis requires more samples than non adaptive data analysis”. In this work we revisit the ADA problem at a foundational level, and ask:

**Question 1.9.** *Is there a more fundamental bottleneck for the ADA problem than the number of samples?*

Consider a mechanism  $\mathcal{A}$  that initially gets the *full description of the underlying distribution*  $\mathcal{P}$ , but is required to shrink this description into a summary  $z$ , whose description length is identical to the description length of  $t$  samples from  $\mathcal{P}$ . (We identify the *space complexity* of  $\mathcal{A}$  with the size of  $z$  in bits.) Afterwards,  $\mathcal{A}$  needs to answer  $k$  adaptive queries using  $z$ , without additional access to  $\mathcal{P}$ . Does this give  $\mathcal{A}$  more power over a mechanism that only obtains  $t$  samples from  $\mathcal{P}$ ?

We show that, in general, the answer is no. Specifically, we show the following theorem.

**Theorem 1.10** (informal version of Theorem 4.1). *Assuming the existence of one-way functions, then every computationally efficient mechanism that is (0.1, 0.1)-accurate for  $k$  queries must have space complexity at least  $\Omega(\sqrt{k})$ .*

In fact, in the formal version of this theorem (see Theorem 4.1) we show that the space complexity must be at least  $\Omega(\sqrt{k})$  times the representation length of domain elements. We view this as a significant strengthening of the previous lower bounds for the ADA problem: it is not that the mechanism did not get enough information about  $\mathcal{P}$ ; it is just that it cannot shrink this information

in a way that allows for efficiently answering  $k$  adaptive queries. This generalizes the negative results of [53, 81], as sampling  $t = \sqrt{k}$  points from  $\mathcal{P}$  is just one particular way for storing information about  $\mathcal{P}$ .

### 1.3 Our techniques

We obtain our results through a combination of techniques across several research areas including cryptography, privacy, learning theory, communication complexity, and information theory.

#### 1.3.1 Our techniques: Multi-instance leakage-resilient scheme

The main cryptographic tool we define and construct is a suitable encryption scheme with multiple keys that is built to withstand a certain amount of key leakage in a very particular way. Specifically, the scheme consists of  $n$  instances of an underlying encryption scheme with independent keys (each of length  $\lambda$  bits). The keys are initially given to an adversary who shrinks them down to a summary  $z$  containing  $s \ll n \cdot \lambda$  bits. After this phase, each instance independently sets an additional parameter, which is public but unknown to the adversary in the initial phase. Then, the adversary obtains encryptions of plaintexts under the  $n$  keys. We require that given  $z$  and the public parameters, the plaintexts encrypted with each key remain computationally hidden, except for a small number of the keys (which depends on  $s$ , but not on  $n$ ).

We call the scheme a *multi-instance leakage-resilient scheme* (or MILR scheme) to emphasize the fact that although the leakage of the adversary is an arbitrary function of all the  $n$  keys, the scheme itself is composed of  $n$  instances that are completely independent.

The efficiency of the MILR scheme is measured by two parameters: (1) the number of keys under which encryptions are (potentially) insecure, and (2) the loss in the security parameter  $\lambda$ . The scheme we construct is optimal in both parameters up to a multiplicative constant factor. First, encryptions remain hidden for all but  $O\left(\frac{s}{\lambda}\right)$  of the keys. This is essentially optimal, as the adversary can define  $z$  to store the first  $\frac{s}{\lambda}$  keys. Second, we lose a constant factor in the security parameter  $\lambda$ . An additional advantage of our construction is that its internal parameters do not depend on  $s$ . If we did allow such a dependency, then in some settings (particularly when  $s \leq o(n \cdot \lambda)$ ) it would be possible to fine-tune the scheme to obtain a multiplicative  $1 + o(1)$  loss in the efficiency parameters, but this has little impact on our application.

Our construction is arguably the most natural one. To encrypt a plaintext with a  $\lambda$ -bit key after the initial phase, we first apply an extractor (with the public parameter as a seed) to hash it down to a smaller key, which is used to encrypt the plaintext with the underlying encryption scheme.

The MILR scheme is related to schemes developed in the area of leakage-resilient cryptography (cf., [17, 39, 40, 71, 79, 56, 58]) where the basic technique of randomness extraction is commonly used. However, leakage-resilient cryptography mainly deals with resilience of cryptosystems to side-channel attacks, whereas our model is not designed to formalize security against such attacks and has several properties that are uncommon in this domain (such as protecting independent multiple instances of an encryption scheme in a way that inherently makes some of them insecure). Consequently, the advanced cryptosystems developed in the area of leakage-resilient cryptography are either unsuitable, overly complex, or inefficient for our purposes.

Despite the simplicity of our construction, our proof that it achieves the claimed security property against leakage is somewhat technical. We stress that we do not rely on hardness assumptions for specific problems, nor assume that the underlying encryption scheme has special properties such as resilience to related-key attacks. Instead, our proof is based on the pre-sampling technique introduced by Unruh [84] to prove security of cryptosystems against non-uniform adversaries in the random oracle model. This technique has been recently refined and optimized in [25, 24, 32] based on tools developed in the area of communication complexity [49, 62]. The fact that we use the technique to prove security in a computational (rather than information-theoretic) setting seems to require the assumption that the underlying encryption scheme is secure against non-uniform adversaries (albeit this is considered a standard assumption).

### 1.3.2 Our techniques: Privacy requires more space

We design a problem that can be solved non-privately with very small space complexity, but requires a large space complexity with privacy. To achieve this, we lift a known negative result on the *sample* complexity of privately solving a specific problem to obtain a *space* lower bound for a related problem. The problem we start with, for which there exists a sampling lower bound, is the so-called *1-way marginals* problem with parameter  $d$ . In this problem, our input is a dataset  $D \in (\{0, 1\}^d)^*$  containing a collection of binary vectors, each of length  $d$ . Our goal is to output a vector  $\vec{a} \in [0, 1]^d$  that approximates the *average* of the input vectors to within small  $L_\infty$  error, say to within error  $1/10$ . That is, we want vector  $\vec{a} \in [0, 1]^d$  to satisfy:  $\left\| \vec{a} - \frac{1}{|D|} \sum_{\vec{x} \in D} \vec{x} \right\|_\infty \leq \frac{1}{10}$ .

We say that an algorithm for this problem has *sample complexity*  $n$  if, for every input dataset of size  $n$ , it outputs a good solution with probability at least 0.9. One of the most fundamental results in the literature of differential privacy shows that this problem requires a large dataset:

**Theorem 1.11** ([16], informal). *Every differentially private algorithm for solving the 1-way marginal problem with parameter  $d$  must have sample complexity  $n = \Omega(\sqrt{d})$ .*

To lift this sampling lower bound into a space lower bound, we consider a problem in which the input dataset contains  $n$  keys  $\vec{x} = (x_1, \dots, x_n)$  (sampled from our MILR scheme). The algorithm must then shrink this dataset into a summary  $z$  containing  $s$  bits. Afterwards, the algorithm gets a “query” that is specified by a collection of  $n$  ciphertexts, each of which is an encryption of a  $d$ -bit vector. The desired task is to approximate the average of the *plaintext* input vectors. Intuitively, if the algorithm has space  $s \ll \sqrt{d}$ , then by the properties of our MILR scheme, it can decrypt at most  $\approx s \ll \sqrt{d}$  of these  $d$ -bit vectors, and is hence trying to solve the 1-way marginal problem with fewer than  $\sqrt{d}$  samples. We show that this argument can be formalized to obtain a contradiction.

### 1.3.3 Our techniques: ADA is about space

As we mentioned, Hardt and Ullman [53] and Steinke and Ullman [81] showed that the ADA problem requires a large *sample* complexity (see Theorem 1.8). Specifically, they showed that there exists a computationally efficient adversary that causes every efficient mechanism to fail in answering adaptive queries. Recall that the ADA game (see Figure 1) begins with the adversary choosing the underlying distribution.

We lift the negative result of [53, 81] to a *space* lower bound. To achieve this, we design an alternative adversary that first samples a large collection of *keys* for our MILR scheme, and then

defines the target distribution to be uniform on these sampled keys. Recall that in our setting, the mechanism gets an exact description of this target distribution and afterwards it must shrink this description into  $s$  bits. However, by the properties of our MILR scheme, this means that the mechanism would only be able to decrypt ciphertexts that correspond to at most  $\approx s/\lambda$  of the keys. We then show that the adversary of [53, 81] can be simulated under these conditions, where the “input sample” from their setting corresponds to the collection of indices of keys for which the mechanism can decrypt ciphertexts.

## 1.4 Additional related work

Query-to-communication lifting theorems were applied by many works in the domain of communication complexity, e.g. [51, 49, 29, 22, 89, 52, 46, 21, 63, 55, 50, 20, 19]. As we mentioned, to obtain our negative result for DP algorithms, we lift the *sampling* lower bound of Bun et al. [16] for the 1-way marginal problem to a *space* lower bound. Many additional sampling lower bounds exist in the DP literature; here we only survey a few of them.

Dwork et al. [37] used *traitor-tracing schemes* to prove computational sampling lower bounds for differential privacy. Their results were extended by Ullman [83], who used *fingerprinting codes* to construct a novel traitor-tracing scheme and to obtain stronger computational sampling lower bounds for DP. Ullman’s construction can be viewed as an encrypted variant of the 1-way marginal problem. Bun et al. [15] and Alon et al. [2] showed that there are trivial learning tasks that require asymptotically more samples to solve with differential privacy (compared to the non-private sample complexity). Feldman [41] and Brown et al. [12] showed that there are learning problems for which every *near optimal* learning algorithm (obtaining near optimal error w.r.t. to the number of input *samples* it takes) must memorize a large portion of its input data. We stress that all of these results are fundamentally different than ours, as they are about sampling rather than space.

Steinhardt et al. [80] asked whether there exists a separation in learning a concept class given additional space constraints and conjectured that the problem of parity learning provided such a separation. Raz [73] proved the conjecture to be true by showing that although parity learning can be solved in  $O(n)$  samples and  $O(n^2)$  space, any algorithm for parity learning that uses  $o(n^2)$  samples requires an *exponential* number of samples. Subsequently, [60] showed that learning linear-size DNF formulas, linear-size decision trees and logarithmic-size juntas is infeasible under certain super-linear space constraints. An active line of work further characterized the limitations of learning concept classes [66, 72, 67, 6, 47, 27] and distributions subject to space constraints [23, 30, 31]. We remark that although the message of higher sample complexity due to space constraints is conceptually similar to our adaptive data analysis result, the problem of learning concept classes is not fundamentally inherent to accurately responding to adaptive queries.

## 1.5 Applications to Communication Complexity

Finally, our arguments also provide distributional one-way communication complexity lower bounds, which are useful when the goal is to compute a relation with a very low success probability. To the best of our knowledge, existing query-to-communication lifting theorems, e.g., [74, 51, 49] do not consider the problems and input distributions that we consider here. Roughly speaking, we show that if any sampling based protocol for computing a function  $f$  requires  $k$  samples  $(a_1, b_1), \dots, (a_k, b_k)$ , where  $a_i \in \{0, 1\}^t$  for each  $i \in [k]$ , then any one-way protocol that computes



$f(A, B)$  in this setting must use  $\Omega(kt)$  communication.

More precisely, in the two-player one-way communication game, inputs  $A$  and  $B$  are given to Alice and Bob, respectively, and the goal is for Alice to send a minimal amount of information to Bob, so that Bob can compute  $f(A, B)$  for some predetermined function  $f$ . The communication cost of a protocol  $\Pi$  is the size of the largest message in bits sent from Alice across all possible inputs  $A$  and the (randomized) communication complexity is the minimum communication cost of a protocol that succeeds with probability at least  $\frac{2}{3}$ . In the distributional setting,  $A$  and  $B$  are further drawn from a known underlying distribution.

In our distributional setting, suppose Alice has  $m$  independent and uniform numbers  $a_1, \dots, a_m$  so that either  $a_i \in GF(p)$  for all  $i \in [m]$  or  $a_i \in GF(2^t)$  for sufficiently large  $t$  for all  $i \in [m]$  and suppose Bob has  $m$  independent and uniform numbers  $b_1, \dots, b_m$  from the same field, either  $GF(p)$  or  $GF(2^t)$ . Then for any function  $f(\langle a_1, b_1 \rangle, \dots, \langle a_m, b_m \rangle)$ , where the dot products are taken over  $GF(2)$  or  $f(a_1 \cdot b_1, \dots, a_m \cdot b_m)$ , where the products are taken over  $GF(p)$ , has the property that the randomized one-way communication complexity of computing  $f$  with probability  $\sigma$  is the same as the number of samples from  $a_1, \dots, a_m$  that Alice needs to send Bob to compute  $f$  with probability  $\sigma - \varepsilon$ . It is easy to prove sampling lower bounds for many of these problems, sum as  $\sum_i a_i \cdot b_i \pmod{p}$  or MAJ( $\langle a_1, b_1 \rangle, \dots, \langle a_m, b_m \rangle$ ), and this immediately translates into communication complexity lower bounds. The main intuition for these lower bounds is that the numbers  $b_1, \dots, b_m$  can be viewed as the hash functions that Bob has and thus we can apply a variant of the leftover hash lemma if Bob only has a small subset of these numbers. We provide full details in Appendix E.

## 1.6 Paper structure

The rest of the paper is structured as follows. The MILR scheme is defined in Section 2. Our results for differential privacy and adaptive data analysis are described in Sections 3 and 4, respectively. We construct our MILR scheme in Section 5, and prove its security in Section 6. For readability, some of the technical details from these sections are deferred to Appendices B, C, and D. Appendix A gives standard preliminaries from cryptography. Finally, Appendix E details the relationship between sample complexity and communication complexity for one-way protocols.

## 2 Multi-instance Leakage-resilient Scheme

We define a *multi-instance leakage-resilient scheme* (or MILR scheme) to be a tuple of efficient algorithms (Gen, Param, Enc, Dec) with the following syntax:

- Gen is a randomized algorithm that takes as input a security parameter  $\lambda$  and outputs a  $\lambda$ -bit secret key. Formally,  $x \leftarrow \text{Gen}(1^\lambda)$ .
- Param is a randomized algorithm that takes as input a security parameter  $\lambda$  and outputs a  $\text{poly}(\lambda)$ -bit public parameter. Formally,  $p \leftarrow \text{Param}(1^\lambda)$ .
- Enc is a randomized algorithm that takes as input a secret key  $x$ , a public parameter  $p$ , and a message  $m \in \{0, 1\}$ , and outputs a ciphertext  $c \in \{0, 1\}^{\text{poly}(\lambda)}$ . Formally,  $c \leftarrow \text{Enc}(x, p, m)$ .
- Dec is a deterministic algorithm that takes as input a secret key  $x$ , a public parameter  $p$ , and a ciphertext  $c$ , and outputs a decrypted message  $m'$ . If the ciphertext  $c$  was an encryption of

$m$  under the key  $x$  with the parameter  $p$ , then  $m' = m$ . Formally, if  $c \leftarrow \text{Enc}(x, p, m)$ , then  $\text{Dec}(x, p, c) = m$  with probability 1.

To define the security of an MILR scheme, let  $n \in \mathbb{N}$ , let  $\vec{x} = (x_1, \dots, x_n)$  be a vector of keys, and let  $\vec{p} = (p_1, \dots, p_n)$  be a vector of public parameters (set once for each scheme by invoking Param). Let  $J \subseteq [n]$  be a subset, referred to as the ‘‘hidden coordinates’’. Now consider a pair of oracles  $\mathcal{E}_1(\vec{x}, \vec{p}, J, \cdot, \cdot)$  and  $\mathcal{E}_0(\vec{x}, \vec{p}, J, \cdot, \cdot)$  with the following properties.

1.  $\mathcal{E}_1(\vec{x}, \vec{p}, J, \cdot, \cdot)$  takes as input an index of a key  $j \in [n]$  and a message  $m$ , and returns  $\text{Enc}(x_j, p_j, m)$ .
2.  $\mathcal{E}_0(\vec{x}, \vec{p}, J, \cdot, \cdot)$  takes the same inputs. If  $j \in J$  then the outcome is  $\text{Enc}(x_j, p_j, 0)$ , and otherwise the outcome is  $\text{Enc}(x_j, p_j, m)$ .

**Definition 2.1.** *Let  $\lambda$  be a security parameter. Let  $\Gamma : \mathbb{R} \rightarrow \mathbb{R}$  and  $\bar{\tau} : \mathbb{R}^2 \rightarrow \mathbb{R}$  be real-valued functions. An MILR scheme (Gen, Param, Enc, Dec) is  $(\Gamma, \bar{\tau})$ -secure against space bounded preprocessing adversaries if the following holds.*

- (1) **Multi semantic security:** *For every  $n = \text{poly}(\Gamma(\lambda))$  and every  $\text{poly}(\Gamma(\lambda))$ -time adversary  $\mathcal{B}$  there exists a negligible function  $\text{negl}$  such that*

$$\left| \Pr_{\vec{x}, \vec{p}, \mathcal{B}, \text{Enc}} \left[ \mathcal{B}^{\mathcal{E}_0(\vec{x}, \vec{p}, [n], \cdot, \cdot)}(\vec{p}) = 1 \right] - \Pr_{\vec{x}, \vec{p}, \mathcal{B}, \text{Enc}} \left[ \mathcal{B}^{\mathcal{E}_1(\vec{x}, \vec{p}, [n], \cdot, \cdot)}(\vec{p}) = 1 \right] \right| \leq \text{negl}(\Gamma(\lambda)).$$

*That is, a computationally bounded adversary that gets the public parameters, but not the keys, cannot tell whether it is interacting with  $\mathcal{E}_0$  or with  $\mathcal{E}_1$ .*

- (2) **Multi-security against a bounded preprocessing adversary:** *For every  $n = \text{poly}(\Gamma(\lambda))$ , every  $s \leq n \cdot \lambda$ , and every preprocessing procedure  $F : (\{0, 1\}^\lambda)^n \rightarrow \{0, 1\}^s$  (possibly randomized), there exists a random function  $J = J(F, \vec{x}, z, \vec{p}) \subseteq [n]$  that given a collection of keys  $\vec{x}$  and public parameters  $\vec{p}$ , and an element  $z \leftarrow F(\vec{x})$ , returns a subset of size  $|J| \geq \tau := n - \bar{\tau}(\lambda, s)$  such that for every  $\text{poly}(\Gamma(\lambda))$ -time algorithm  $\mathcal{B}$  there exists a negligible function  $\text{negl}$  satisfying*

$$\left| \Pr_{\substack{\vec{x}, \vec{p}, \mathcal{B}, \text{Enc} \\ z \leftarrow F(\vec{x}) \\ J \leftarrow J(F, \vec{x}, z, \vec{p})}} \left[ \mathcal{B}^{\mathcal{E}_0(\vec{x}, \vec{p}, J, \cdot, \cdot)}(z, \vec{p}) = 1 \right] - \Pr_{\substack{\vec{x}, \vec{p}, \mathcal{B}, \text{Enc} \\ z \leftarrow F(\vec{x}) \\ J \leftarrow J(F, \vec{x}, z, \vec{p})}} \left[ \mathcal{B}^{\mathcal{E}_1(\vec{x}, \vec{p}, J, \cdot, \cdot)}(z, \vec{p}) = 1 \right] \right| \leq \text{negl}(\Gamma(\lambda)).$$

*That is, even if  $s$  bits of our  $n$  keys were leaked (computed by the preprocessing function  $F$  operating on the keys), then still encryptions w.r.t. the keys of  $J$  are computationally indistinguishable.*

**Remark 2.2.** *When  $\Gamma$  is the identity function, we simply say that the scheme is  $\bar{\tau}$ -secure. Note that in this case, security holds against all adversaries with runtime polynomial in the security parameter  $\lambda$ . We will further assume the existence of a sub-exponentially secure encryption scheme. By that we mean that there exists a constant  $\nu > 0$  such that the scheme is  $(\Gamma, \bar{\tau})$ -secure for  $\Gamma(\lambda) = 2^{\lambda^\nu}$ . That is, we assume the existence of a scheme in which security holds against all adversaries with runtime polynomial in  $2^{\lambda^\nu}$ .*

In Sections 5 and 6 we show the following theorem.

**Theorem 2.3.** *Let  $\Omega(\lambda) \leq \Gamma(\lambda) \leq 2^{o(\lambda)}$ . If there exists a  $\Gamma(\lambda)$ -secure encryption scheme against non-uniform adversaries then there exists an MILR scheme that is  $(\Gamma(\lambda), \bar{\tau})$ -secure against space bounded non-uniform preprocessing adversaries, where  $\bar{\tau}(\lambda, s) = \frac{2s}{\lambda} + 4$ .*

### 3 Space Hardness for Differential Privacy

Consider an algorithm that is instantiated on a dataset  $D$  and then aims to answer a query w.r.t.  $D$ . We say that such an algorithm has space  $s$  if, before it gets the query, it shrinks  $D$  to a *summary*  $z$  containing at most  $s$  bits. Then, when obtaining the query, the algorithm answers it using only the summary  $z$  (without additional access to the original dataset  $D$ ).

Let  $(\text{Gen}, \text{Param}, \text{Enc}, \text{Dec})$  be an MILR scheme with security parameter  $\lambda$ . In the  $(\lambda, d)$ -decoded average (DA) problem with sample complexity  $n$ , the input dataset contains  $n$  keys, that is  $D = (x_1, \dots, x_n) \in (\{0, 1\}^\lambda)^n$ . A query  $q = ((p_1, c_1), \dots, (p_n, c_n))$  is specified using  $n$  pairs  $(p_i, c_i)$  where  $p_i$  is a public parameter and  $c_i$  is a ciphertext, which is an encryption of a binary vector of length  $d$ . The goal is to release a vector  $\vec{a} = (a_1, \dots, a_d) \in [0, 1]^d$  that approximates the “decrypted average vector (dav)”, defined as  $\text{dav}_q(D) = \frac{1}{n} \sum_{i=1}^n \text{Dec}(x_i, p_i, c_i)$ .

**Definition 3.1.** *Let  $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$  be an algorithm where  $\mathcal{A}_1 : (\{0, 1\}^\lambda)^n \rightarrow \{0, 1\}^s$  is the preprocessing procedure that summarizes a dataset of  $n$  keys into  $s$  bits, and where  $\mathcal{A}_2$  is the “response algorithm” that gets the outcome of  $\mathcal{A}_1(D)$  and a query  $q$ . We say that  $\mathcal{A}$  solves the DA problem if with probability at least  $9/10$  the output is a vector  $\vec{a}$  satisfying  $\|\vec{a} - \text{dav}_q(D)\|_\infty \leq \frac{1}{10}$ .*

Without privacy considerations, the DA problem is almost trivial. Specifically,

**Lemma 3.2.** *The  $(\lambda, d)$ -DA problem can be solved efficiently using space  $s = O(\lambda \log(d))$ .*

*Proof.* The preprocessing algorithm  $\mathcal{A}_1$  samples  $O(\log d)$  of the input keys. Algorithm  $\mathcal{A}_2$  then gets the query  $q$  and estimates the dav vector using the sampled keys. The lemma then follows by the Chernoff bound.  $\square$

**Remark 3.3.** *As we mentioned, to simplify the presentation, in our computational model we identify the space complexity of algorithm  $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$  with the size of the output of algorithm  $\mathcal{A}_1$ . We remark, however, that our separation extends to a streaming model where both  $\mathcal{A}_1$  and  $\mathcal{A}_2$  are required to have small space. To see this, note that algorithm  $\mathcal{A}_1$  in the above proof already has small space (and not just small output length), as it merely samples  $O(\log d)$  keys from its input dataset. We now analyze the space complexity of  $\mathcal{A}_2$ , when it reads the query  $q$  in a streaming fashion. Recall that the query  $q$  contains  $n$  public parameters  $p_1, \dots, p_n$  and  $n$  ciphertexts  $c_1, \dots, c_n$ , where each  $c_i$  is an encryption of a  $d$ -bit vector, call it  $y_i \in \{0, 1\}^d$ . To allow  $\mathcal{A}_2$  to read  $q$  using small space, we order it as follows:  $q = (p_1, \dots, p_n), (c_{1,1}, \dots, c_{n,1}), \dots, (c_{1,d}, \dots, c_{n,d}) \triangleq q_0 \circ q_1 \circ \dots \circ q_d$ . Here  $c_{i,j} = \text{Enc}(x_i, p_i, y_i[j])$  is an encryption of the  $j$ th bit of  $y_i$  using key  $x_i$  and public parameter  $p_i$ . Note that the first part of the stream,  $q_0$ , contains the public parameters, and then every part  $q_j$  contains encryptions of the  $j$ th bit of each of the  $n$  input vectors. With this ordering of the query, algorithm  $\mathcal{A}_2$  begins by reading  $q_0$  and storing the  $O(\log d)$  public parameters corresponding to the keys that were stored by  $\mathcal{A}_1$ . Then, for every  $j \in [d]$ , when reading  $q_j$ , algorithm  $\mathcal{A}_2$  estimates the average of the  $j$ th coordinate using the sampled keys. Algorithm  $\mathcal{A}_2$  then outputs this estimated value, and proceeds to the next coordinate. This can be implemented using space complexity  $\text{poly}(\lambda \log(d))$ .*

So, without privacy constraints, the DA problem can be solved using small space. We now show that, assuming that the input dataset is large enough, the DA problem can easily be solved with differential privacy using *large* space. Specifically,

**Lemma 3.4.** *There is a computationally efficient  $(\varepsilon, \delta)$ -differentially private algorithm that solves the  $(\lambda, d)$ -DA problem using space  $s = O\left(\frac{1}{\varepsilon} \cdot \sqrt{d \cdot \log\left(\frac{1}{\delta}\right)} \cdot \lambda \cdot \log d\right)$ , provided that the size of the input dataset satisfies  $n = \Omega(s)$  (large enough).*

*Proof.* The preprocessing algorithm  $\mathcal{A}_1$  samples  $\approx \sqrt{d}$  of the keys. By standard composition theorems for differential privacy [38], this suffices for the response algorithm  $\mathcal{A}_2$  to privately approximate each of the  $d$  coordinates of the target vector.  $\square$

Thus the DA problem can be solved non-privately using small space, and it can be solved privately using large space. Our next goal is to show that large space is indeed necessary to solve this problem privately. Before showing that, we introduce several additional preliminaries on computational differential privacy and on fingerprinting codes.

### 3.1 Preliminaries on computational differential privacy and fingerprinting codes

Computational differential privacy was defined by Beimel et al [7] and Mironov et al. [65]. Let  $\mathcal{A}$  be a randomized algorithm (mechanism) that operates on datasets. Computational differential privacy is defined via a two player game between a *challenger* and an adversary, running a pair of algorithms  $(Q, T)$ . The game begins with the adversary  $Q$  choosing a pair of neighboring datasets  $(D_0, D_1)$  of size  $n$  each, as well as an arbitrary string  $r$  (which we think of as representing its internal state). Then the challenger samples a bit  $b$  and applies  $\mathcal{A}(D_b)$  to obtain an outcome  $a$ . Then  $T(r, \cdot)$  is applied on  $a$  and tries to guess  $b$ . Formally,

**Definition 3.5.** *Let  $\lambda$  be a security parameter, let  $\varepsilon$  be a constant, and let  $\delta : \mathbb{R} \rightarrow \mathbb{R}$  be a function. A randomized algorithm  $\mathcal{A} : X^* \rightarrow Y$  is  $(\varepsilon, \delta)$ -computationally differentially private (CDP) if for every  $n = \text{poly}(\lambda)$  and every non-uniform  $\text{poly}(\lambda)$ -time adversary  $(Q, T)$  there exists a negligible function  $\text{negl}$  such that*

$$\Pr_{\substack{(r, D_0, D_1) \leftarrow Q \\ \mathcal{A}, T}} [T(r, \mathcal{A}(D_0)) = 1] \leq e^\varepsilon \cdot \Pr_{\substack{(r, D_0, D_1) \leftarrow Q \\ \mathcal{A}, T}} [T(r, \mathcal{A}(D_1)) = 1] + \delta(n) + \text{negl}(\lambda).$$

**Definition 3.6.** *Let  $\varepsilon$  be a constant and let  $\delta = \delta(\lambda)$  be a function. Given two probability ensembles  $\mathcal{X} = \{X_\lambda\}_{\lambda \in \mathbb{N}}$  and  $\mathcal{Y} = \{Y_\lambda\}_{\lambda \in \mathbb{N}}$  we write  $\mathcal{X} \approx_{\varepsilon, \delta} \mathcal{Y}$  if for every non-uniform probabilistic polynomial-time distinguisher  $D$  there exists a negligible function  $\text{negl}$  such that  $\Pr_{x \leftarrow X_\lambda} [D(x) = 1] \leq e^\varepsilon \cdot \Pr_{y \leftarrow Y_\lambda} [D(y) = 1] + \delta(\lambda) + \text{negl}(\lambda)$ , and vice versa.*

We recall the concept of *fingerprinting codes*, which was introduced by Boneh and Shaw [11] as a cryptographic tool for watermarking digital content. Starting from the work of Bun, Ullman, and Vadhan [16], fingerprinting codes have played a key role in proving lower bounds for differential privacy in various settings.

A (collusion-resistant) fingerprinting code is a scheme for distributing codewords  $w_1, \dots, w_n$  to  $n$  users that can be uniquely traced back to each user. Moreover, if a group of (at most  $k$ ) users

combines its codewords into a pirate codeword  $\hat{w}$ , then the pirate codeword can still be traced back to one of the users who contributed to it. Of course, without any assumption on how the pirates can produce their combined codeword, no secure tracing is possible. To this end, the pirates are constrained according to a *marking assumption*, which asserts that the combined codeword must agree with at least one of the “real” codewords in each position. Namely, at an index  $j$  where  $w_i[j] = b$  for every  $i \in [n]$ , the pirates are constrained to output  $\hat{w}$  with  $\hat{w}[j] = b$  as well.<sup>2</sup>

**Definition 3.7** ([11, 82]). *A  $k$ -collusion resilient fingerprinting code of length  $d$  for  $n$  users with failure probability  $\gamma$ , or  $(n, d, k, \gamma)$ -FPC in short, is a pair of random variables  $C \in \{0, 1\}^{n \times d}$  and  $\text{Trace} : \{0, 1\}^d \rightarrow 2^{[n]}$  such that the following holds. For all adversaries  $P : \{0, 1\}^{k \times d} \rightarrow \{0, 1\}^d$  and  $S \subseteq [n]$  with  $|S| = k$ ,*

$$\Pr_{C, \text{Trace}, P} [(\forall 1 \leq j \leq d \exists i \in [n] \text{ s.t. } P(C_S)[j] = c_i[j]) \wedge (\text{Trace}(P(C_S)) = \emptyset)] \leq \gamma,$$

and

$$\Pr_{C, \text{Trace}, P} [\text{Trace}(P(C_S)) \cap ([n] \setminus S) \neq \emptyset] \leq \gamma,$$

where  $C_S$  contains the rows of  $C$  given by  $S$ .

**Remark 3.8.** *As mentioned, the condition  $\{\forall 1 \leq j \leq d \exists i \in [n] \text{ s.t. } P(C_S)[j] = c_i[j]\}$  is called the “marking assumption”. The second condition is called the “small probability of false accusation”. Hence, if the adversary  $P$  guarantees that its output satisfies the marking assumption, then with probability at least  $1 - 2\gamma$  it holds that algorithm  $\text{Trace}$  outputs an index  $i \in S$ .*

**Theorem 3.9** ([11, 82, 81]). *For every  $1 \leq k \leq n$  there is a  $k$ -collusion-resilient fingerprinting code of length  $d = O(k^2 \cdot \log n)$  for  $n$  users with failure probability  $\gamma = \frac{1}{n^2}$  and an efficiently computable  $\text{Trace}$  function.*

We remark that there exist both adaptive and non-adaptive constructions of fingerprinting codes with the guarantees of Theorem 3.9; we use the non-adaptive variant.

### 3.2 A negative result for the DA problem

Our main negative result for space bounded differentially private algorithms is the following.

**Theorem 3.10.** *Let  $\Pi = (\text{Gen}, \text{Param}, \text{Enc}, \text{Dec})$  be an MILR scheme with security parameter  $\lambda$  that is  $(\Gamma, \bar{\tau})$ -secure against space bounded preprocessing adversaries. Let  $d \leq \text{poly}(\Gamma(\lambda))$  and  $n \leq \text{poly}(\Gamma(\lambda))$  be functions of  $\lambda$ . Let  $\varepsilon$  be a constant and let  $\delta \leq \frac{1}{4n(e^\varepsilon + 1)} = \Theta(\frac{1}{n})$ . For every  $\text{poly}(\Gamma(\lambda))$ -time  $(\varepsilon, \delta)$ -CDP algorithm for the  $(\lambda, d)$ -DA problem with sample complexity  $n$  and space complexity  $s$  it holds that  $\bar{\tau}(\lambda, s) = \Omega\left(\sqrt{\frac{d}{\log n}}\right)$ .*

*Proof.* Let  $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$  be a  $\text{poly}(\Gamma(\lambda))$ -time CDP algorithm for the  $(\lambda, d)$ -DA problem using sample complexity  $n = \text{poly}(\Gamma(\lambda))$  and space complexity  $s$ . Denote  $\bar{\tau} = \bar{\tau}(\lambda, s)$ , and assume towards a contradiction that  $\bar{\tau} = O\left(\sqrt{\frac{d}{\log n}}\right)$  (small enough). We construct the following adversary  $\mathcal{B}$  to an  $(n+1, d, \bar{\tau}, \frac{1}{n^2})$ -FPC (such a code is guaranteed to exist by Theorem 3.9 and by the contradictory assumption).

<sup>2</sup>We follow the formulation of the marking assumption as given by [81], which is a bit different than the commonly considered one.

1. The input is  $n$  codewords  $w_1, \dots, w_n \in \{0, 1\}^d$ .
2. Sample  $n$  keys  $x_1, \dots, x_n \sim \text{Gen}(1^\lambda)$ .
3. Let  $z \leftarrow \mathcal{A}_1(x_1, \dots, x_n)$ .
4. Sample  $n$  public parameters  $p_1, \dots, p_n \sim \text{Param}(1^\lambda)$ .
5. For  $i \in [n]$  let  $c_i \leftarrow \text{Enc}(x_i, p_i, w_i)$ .
6. Let  $\vec{a} \leftarrow \mathcal{A}_2(z, (p_1, c_1), \dots, (p_n, c_n))$ .
7. Output  $\vec{a}$ , after rounding its coordinates to  $\{0, 1\}$ .

We think of  $\mathcal{B}$  as an adversary to an FPC, and indeed, its input is a collection of codewords and its output is a binary vector of length  $d$ . Observe that if  $\mathcal{A}$  solves the DA problem (i.e., approximates the decrypted average vector), then for every coordinate, the outcome of  $\mathcal{B}$  must agree with at least one of the input codewords, namely, it satisfies the marking assumption (see Remark 3.8).

**Remark 3.11.** *Before we proceed with the formal proof, we give an overview of its structure (this remark can be ignored, and is not needed for the formal proof). Informally, we will show that*

- (1) *Algorithm  $\mathcal{B}$  is computationally differentially private w.r.t. the collection of codewords (even though our assumption on  $\mathcal{A}$  is that it is private w.r.t. the keys).*
- (2) *Leveraging the properties of the MILR scheme, we will show that  $\mathcal{B}$  must effectively ignore most of its inputs, except for at most  $\bar{\tau}$  codewords. This means that  $\mathcal{B}$  is effectively an FPC adversary that operates on only  $\bar{\tau}$  codewords (rather than the  $n$  codewords it obtains as input).*
- (3) *A known result in the literature of differential privacy (starting from [16]) is that a successful FPC adversary cannot be differentially private, because this would contradict the fact that the tracing algorithm is able to recover one of its input points. Our gain here comes from the fact that  $\mathcal{B}$  only uses (effectively)  $\bar{\tau}$  codewords, and hence, in order to get a contradiction, it suffices to use an FPC with a much shorter codeword-length (only  $\approx \bar{\tau}^2$  instead of  $\approx n^2$ ). This will mean that the hardness of the DA problem depends on the space of  $\mathcal{A}$  (which controls  $\bar{\tau}$ ) rather than the size of the input (which is  $n$ ).*

Recall that  $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$  is computationally differentially private w.r.t. the *keys*. We first show that  $\mathcal{B}$  is computationally differentially private w.r.t. the *codewords*. To this end, let  $Q$  be an adversary (as in Definition 3.5) that outputs a pair of neighboring datasets  $(\vec{w}, \vec{w}')$ , each containing  $n$  codewords, together with a state  $r$ . Given  $(\vec{w}, \vec{w}')$ , we write  $\ell = \ell(\vec{w}, \vec{w}') \subseteq [n]$  to denote the index on which  $\vec{w}, \vec{w}'$  differ. We also write  $x_0$  to denote another key, independent of the keys  $x_1, \dots, x_n$  sampled by algorithm  $\mathcal{B}$ . By the privacy guarantees of algorithm  $\mathcal{A}$  and by the semantic security

of the encryption scheme (see Definition 2.1) we have that

$$\begin{aligned}
\langle r, \mathcal{B}(\vec{w}) \rangle &\equiv \\
&\equiv \langle r, \mathcal{A}_2(\mathcal{A}_1(x_1, \dots, x_\ell, \dots, x_n), \vec{p}, \text{Enc}(x_1, p_1, w_1), \dots, \text{Enc}(x_\ell, p_\ell, w_\ell), \dots, \text{Enc}(x_n, p_n, w_n)) \rangle \\
&\approx_{(\varepsilon, \delta)} \langle r, \mathcal{A}_2(\mathcal{A}_1(x_1, \dots, \mathbf{x}_0, \dots, x_n), \vec{p}, \text{Enc}(x_1, p_1, w_1), \dots, \text{Enc}(x_\ell, p_\ell, w_\ell), \dots, \text{Enc}(x_n, p_n, w_n)) \rangle \\
&\equiv_c \langle r, \mathcal{A}_2(\mathcal{A}_1(x_1, \dots, x_0, \dots, x_n), \vec{p}, \text{Enc}(x_1, p_1, w_1), \dots, \text{Enc}(x_\ell, p_\ell, \mathbf{w}'_\ell), \dots, \text{Enc}(x_n, p_n, w_n)) \rangle \\
&\approx_{(\varepsilon, \delta)} \langle r, \mathcal{A}_2(\mathcal{A}_1(x_1, \dots, \mathbf{x}_\ell, \dots, x_n), \vec{p}, \text{Enc}(x_1, p_1, w_1), \dots, \text{Enc}(x_\ell, p_\ell, \mathbf{w}'_\ell), \dots, \text{Enc}(x_n, p_n, w_n)) \rangle \\
&\equiv \langle r, \mathcal{B}(\vec{w}') \rangle.
\end{aligned}$$

So algorithm  $\mathcal{B}$  is  $(2\varepsilon, (e^\varepsilon + 1)\delta)$ -computationally differentially private. Now consider the following variant of algorithm  $\mathcal{B}$ , denoted as  $\hat{\mathcal{B}}$ . The modifications from  $\mathcal{B}$  are marked in red.

1. The input is  $n$  codewords  $w_1, \dots, w_n \in \{0, 1\}^d$ .
2. Sample  $n$  keys  $x_1, \dots, x_n \sim \text{Gen}(1^\lambda)$ .
3. Let  $z \leftarrow \mathcal{A}_1(x_1, \dots, x_n)$ .
4. Sample  $n$  public parameters  $p_1, \dots, p_n \sim \text{Param}(1^\lambda)$ .
5. Let  $J \leftarrow J(\mathcal{A}_1, \vec{x}, z, \vec{p}) \subseteq [n]$  be the subset of coordinates guaranteed to exist by Definition 2.1, of size  $|J| = n - \bar{\tau}$ .
6. For  $i \in J$  let  $c_i \leftarrow \text{Enc}(x_i, p_i, 0)$ .
7. For  $i \in [n] \setminus J$  let  $c_i \leftarrow \text{Enc}(x_i, p_i, w_i)$ .
8. Let  $\vec{a} \leftarrow \mathcal{A}_2(z, (p_1, c_1), \dots, (p_n, c_n))$ .
9. Output  $\vec{a}$ , after rounding its coordinates to  $\{0, 1\}$ .

**Remark 3.12.** Observe that algorithm  $\hat{\mathcal{B}}$  is not necessarily computationally efficient, since computing  $J(\mathcal{A}_1, \vec{x}, z, \vec{p})$  might not be efficient. Nevertheless, as we next show, this still suffices to obtain a contradiction and complete the proof of the lower bound. Specifically, we will show that  $\hat{\mathcal{B}}$  is computationally differentially private (w.r.t. the codewords) and that it is a successful adversary to the FPC. This will lead to a contradiction, even if  $\hat{\mathcal{B}}$  itself is a non-efficient mechanism.

We now show that, by the multi-security of the MILR scheme (see Definition 2.1), the outcome distributions of  $\mathcal{B}$  and  $\hat{\mathcal{B}}$  are computationally indistinguishable. Specifically, we want to show that for every efficient adversary  $(Q, T)$ , as in Definition 3.5, it holds that

$$|\Pr[T(r, \mathcal{B}(\vec{w})) = 1] - \Pr[T(r, \hat{\mathcal{B}}(\vec{w})) = 1]| \leq \text{negl}(\Gamma(\lambda)).$$

Note that here both expressions are with the same dataset  $\vec{w}$  (without the neighboring dataset  $\vec{w}'$ ). To show this, consider the following algorithm, denoted as  $\mathcal{W}$ , which we view as an adversary to the MILR scheme. This algorithm has only oracle access to encryptions, via an oracle  $\mathcal{E}$ .

1. The input of  $\mathcal{W}$  is  $n$  codewords  $w_1, \dots, w_n \in \{0, 1\}^d$ , an element  $z$  (supposedly computed by  $\mathcal{A}_1$ ), and a collection of  $n$  public parameters  $p_1, \dots, p_n$ .
2. For  $i \in [n]$  let  $c_i \leftarrow \mathcal{E}(i, w_i)$ .

3. Let  $\vec{a} \leftarrow \mathcal{A}_2(z, (p_1, c_1), \dots, (p_n, c_n))$ .
4. Output  $\vec{a}$ , after rounding its coordinates to  $\{0, 1\}$ .

Now by the multi-security of the MILR scheme we have

$$\begin{aligned}
& \left| \Pr_{Q,T,\mathcal{B}} [T(r, \mathcal{B}(\vec{w})) = 1] - \Pr_{Q,T,\hat{\mathcal{B}}} [T(r, \hat{\mathcal{B}}(\vec{w})) = 1] \right| \\
&= \left| \Pr_{\substack{\vec{x}, \vec{p}, \mathcal{W}, Q, T, \text{Enc} \\ z \leftarrow \mathcal{A}_1(\vec{x}) \\ J \leftarrow J(\mathcal{A}_1, \vec{x}, z, \vec{p})}} [Q(r, \mathcal{W}^{\mathcal{E}_1(\vec{x}, \vec{p}, J, \cdot)}(\vec{w}, z, \vec{p})) = 1] - \Pr_{\substack{\vec{x}, \vec{p}, \mathcal{W}, Q, T, \text{Enc} \\ z \leftarrow \mathcal{A}_1(\vec{x}) \\ J \leftarrow J(\mathcal{A}_1, \vec{x}, z, \vec{p})}} [Q(r, \mathcal{W}^{\mathcal{E}_0(\vec{x}, \vec{p}, J, \cdot)}(\vec{w}, z, \vec{p})) = 1] \right| \\
&\leq \text{negl}(\Gamma(\lambda)).
\end{aligned}$$

So we have that  $\hat{\mathcal{B}} \equiv_c \mathcal{B}$  and we have that  $\mathcal{B}$  is  $(2\varepsilon, (e^\varepsilon + 1)\delta)$ -computationally differentially private. Hence, algorithm  $\hat{\mathcal{B}}$  is also  $(2\varepsilon, (e^\varepsilon + 1)\delta)$ -computationally differentially private (w.r.t. the input codewords). Observe that algorithm  $\hat{\mathcal{B}}$  ignores all but  $N - |J| = \bar{\tau}$  of the codewords, and furthermore, the choice of which codewords to ignore is independent of the codewords themselves. Now consider the following thought experiment.

1. Sample a codebook  $w_0, w_1, \dots, w_n$  for the fingerprinting code.
2. Run  $\hat{\mathcal{B}}$  on  $(w_1, \dots, w_n)$ .
3. Run Trace on the outcome of  $\hat{\mathcal{B}}$  and return its output.

As  $\hat{\mathcal{B}}$  ignores all but  $\bar{\tau}$  codewords, by the properties of the FPC, with probability at least  $1 - \frac{1}{n^2} \geq \frac{1}{2}$  the outcome of Trace is a coordinate of a codeword that  $\hat{\mathcal{B}}$  did not ignore, and in particular, it is a coordinate between 1 and  $n$ . Therefore, there must exist a coordinate  $i^* \neq 0$  that is output by this thought experiment with probability at least  $\frac{1}{2n}$ . Now consider the following modified thought experiment.

1. Sample a codebook  $w_0, w_1, \dots, w_n$  for the fingerprinting code.
2. Run  $\hat{\mathcal{B}}$  on  $(w_1, \dots, w_{i^*-1}, w_0, w_{i^*+1}, \dots, w_n)$ .
3. Run Trace on the outcome of  $\hat{\mathcal{B}}$  and return its output.

As  $\hat{\mathcal{B}}$  is computationally differentially private and as Trace is an efficient algorithm, the probability of outputting  $i^*$  in this second thought experiment is roughly the same as in the previous thought experiment, specifically, at least

$$e^{-2\varepsilon} \left( \frac{1}{2n} - (e^\varepsilon + 1)\delta - \text{negl}(\Gamma(\lambda)) \right) \geq e^{-2\varepsilon} \left( \frac{1}{4n} - \text{negl}(\Gamma(\lambda)) \right) = \Omega\left(\frac{1}{n}\right).$$

However, by the guarantees of the FPC (small probability of false accusation), in the second experiment the probability of outputting  $i^*$  should be at most  $\frac{1}{n^2}$ . This is a contradiction to the existence of algorithm  $\mathcal{A}$ .  $\square$



The following corollary follows by instantiating Theorem 3.10 with our MILR scheme, as specified in Theorem 2.3.

**Corollary 3.13.** *Let  $\Omega(\lambda) \leq \Gamma(\lambda) \leq 2^{o(\lambda)}$ , and let  $d \leq \text{poly}(\Gamma(\lambda))$ . If there exists a  $\Gamma(\lambda)$ -secure encryption scheme against non-uniform adversaries then there exists an MILR scheme such that the corresponding  $(\lambda, d)$ -DA problem requires large space to be solved privately. Specifically, let  $n \leq \text{poly}(\Gamma(\lambda))$ , let  $\varepsilon$  be a constant, and let  $\delta \leq \frac{1}{4n(e^\varepsilon+1)} = \Theta(\frac{1}{n})$ . Every  $\text{poly}(\Gamma(\lambda))$ -time  $(\varepsilon, \delta)$ -CDP algorithm for the  $(\lambda, d)$ -DA problem with sample complexity  $n$  must have space complexity  $s = \Omega\left(\lambda \cdot \sqrt{\frac{d}{\log n}}\right)$ .*

## 4 Space Hardness for Adaptive Data Analysis (ADA)

Consider a mechanism that first gets as input a sample containing  $t$  i.i.d. samples from some underlying (unknown) distribution  $\mathcal{D}$ , and then answers  $k$  *adaptively chosen* statistical queries w.r.t.  $\mathcal{D}$ . Importantly, the answers must be accurate w.r.t. the underlying distribution and not just w.r.t. the empirical sample. The challenge here is that as the queries are being chosen adaptively, the interaction might quickly lead to *overfitting*, i.e., result in answers that are only accurate w.r.t. the empirical sample and not w.r.t. the underlying distribution. This fundamental problem, which we refer to as the ADA problem, was introduced by Dwork et al. [34] who connected it to differential privacy and showed that differential privacy can be used as a countermeasure against overfitting. Intuitively, overfitting happens when answers reveal properties that are specific to the input sample, rather than to the general population, and this is exactly what differential privacy aims to protect against.

Hardt, Steinke, and Ullman [53, 81] showed negative results for the ADA problem. Specifically, they showed that given  $t$  samples, it is computationally hard to answer more than  $k = O(t^2)$  adaptive queries. We show that the hardness of the ADA problem is actually more fundamental; it is, in fact, a result of a space bottleneck rather than a sampling bottleneck. Informally, we show that the same hardness result continues to hold even if in the preprocessing stage the mechanism is given the full description of the underlying distribution  $\mathcal{D}$ , and is then required to store only a limited amount of information about it (an amount that equals the representation length of  $t$  samples from  $\mathcal{D}$ ). So it is not that the mechanism did not get enough information about  $\mathcal{D}$ ; it is just that it cannot shrink this information in a way that supports  $t^2$  adaptive queries. This generalizes the negative results of [53, 81], as sampling  $t$  points from  $\mathcal{D}$  is just one particular way of trying to store information about  $\mathcal{D}$ .

Consider `AdaptiveGameSpace`, where the mechanism initially gets the full description of the underlying distribution, but it must shrink it into an  $s$ -bit summary  $z$ . To emphasize that the mechanism does not have additional access to the underlying distribution, we think about it as two mechanisms  $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$  where  $\mathcal{A}_1$  computes the summary  $z$  and where  $\mathcal{A}_2$  answers queries given  $z$ . We consider  $s = |z|$  as the space complexity of such a mechanism  $\mathcal{A}$ .

Our main theorem in the context of the ADA problem is the following.

**Theorem 4.1.** *Let  $\Omega(\lambda) \leq \Gamma(\lambda) \leq 2^{o(\lambda)}$ , and let  $k \leq \text{poly}(\Gamma(\lambda))$ . If there exists a  $\Gamma(\lambda)$ -secure encryption scheme against non-uniform adversaries then there exists a  $\text{poly}(\Gamma(\lambda))$ -time adversary  $\mathcal{B}$  such that the following holds. Let  $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$  be a  $\text{poly}(\Gamma(\lambda))$ -time mechanism with space com-*

---

**Algorithm 1** AdaptiveGameSpace( $\mathcal{A}=(\mathcal{A}_1, \mathcal{A}_2), \mathcal{B}, s, k$ )

---

1. The adversary  $\mathcal{B}$  chooses a distribution  $\mathcal{D}$  over a domain  $\mathcal{X}$ .
  2. The mechanism  $\mathcal{A}_1$  gets  $\mathcal{D}$  and summarizes it into  $s$  bits, denoted as  $z$ .
  3. The mechanism  $\mathcal{A}_2$  is instantiated with  $z$ .
  4. For round  $i = 1, 2, \dots, k$ :
    - (a) The adversary  $\mathcal{B}$  specifies a query  $q_i : \mathcal{X} \rightarrow \{-1, 0, 1\}$
    - (b) The mechanism  $\mathcal{A}_2$  obtains  $q_i$  and responds with an answer  $a_i \in [-1, 1]$
    - (c)  $a_i$  is given to  $\mathcal{A}$
  5. The outcome of the game is one if  $\exists i$  s.t.  $|a_i - \mathbb{E}_{y \sim \mathcal{D}}[q_i(y)]| > 1/10$ , and zero otherwise.
- 

plexity  $s \leq O(\lambda \cdot \sqrt{k})$  (small enough). Then,  $\Pr[\text{AdaptiveGameSpace}(\mathcal{A}, \mathcal{B}, s, k) = 1] > \frac{2}{3}$ .

Furthermore, the underlying distribution defined by the adversary  $\mathcal{B}$  can be fully described using  $O(\sqrt{k} \cdot \lambda)$  bits, it is sampleable in  $\text{poly}(\Gamma(\lambda))$ -time, and elements sampled from this distribution can be represented using  $O(\lambda + \log(k))$  bits.

In a sense, the “furthermore” part of the theorem shows that the distribution chosen by our adversary is not too complex. Specifically, our negative result continues to hold even if the space of the mechanism is *linear* in the full description length of the underlying distribution (in a way that allows for efficiently sampling it). If the space of the mechanism was just a constant times bigger, it could store the full description of the underlying distribution and answer an unbounded number of adaptive queries. The formal proof of Theorem 4.1 is deferred to Appendix B. Here we only provide an informal (and overly simplified) proof sketch.

#### 4.1 Informal proof sketch

Let  $k$  denote the number of queries that the adversary makes. Our task is to show that there is an adversary that fails every efficient mechanism  $\mathcal{A}_{\text{space}}$  that plays in AdaptiveGameSpace, provided that it uses space  $s \ll \sqrt{k}$ . What we know from [53, 81] is that there is an adversary  $\mathcal{B}_{\text{sample}}$  that fails every efficient mechanism that plays in the standard ADA game (the game specified in Figure 1), provided that its sample complexity is  $t \ll \sqrt{k}$ . We design an adversary  $\mathcal{B}_{\text{space}}$  that plays in AdaptiveGameSpace in a way that emulates  $\mathcal{B}_{\text{sample}}$ . We now elaborate on the key points in the construction of  $\mathcal{B}_{\text{space}}$ , and their connection to  $\mathcal{B}_{\text{sample}}$ .

Recall that both games begin with the adversary specifying the underlying distribution. A useful fact about the adversary  $\mathcal{B}_{\text{sample}}$  (from [53, 81]) is that the distribution it specifies is uniform on a small set of points of size  $n = \Theta(t)$  (these  $n$  points are unknown to the mechanism that  $\mathcal{B}_{\text{sample}}$  plays against). Our adversary,  $\mathcal{B}_{\text{space}}$ , first samples  $n$  independent keys  $(x_1, \dots, x_n)$  from our MILR scheme, and then defines the target distribution  $\mathcal{D}_{\text{space}}$  to be uniform over the set  $\{(j, x_j)\}_{j \in [n]}$ . Recall that in AdaptiveGameSpace this target distribution is given to the mechanism  $\mathcal{A}_{\text{space}}$ , who must shrink it into a summary  $z$  containing  $s$  bits. After this stage, by the security of our MILR scheme, there should exist a large set  $J \subseteq [n]$  of size  $|J| = n - \bar{\tau}$  corresponding to keys uncompromised by  $\mathcal{A}_{\text{space}}$ . Denote  $I = [n] \setminus J$ .

Our adversary  $\mathcal{B}_{\text{space}}$  now emulates  $\mathcal{B}_{\text{sample}}$  as follows. First, let  $\mathcal{D}_{\text{sample}}$  denote the target distribution chosen by  $\mathcal{B}_{\text{sample}}$ , and let  $m_1, \dots, m_n$  denote its support. Our adversary then samples  $n$  public parameters  $p_1, \dots, p_n$ , and encrypts every point in the support  $m_j$  using its corresponding

key and public parameter. Specifically,  $c_j \leftarrow \text{Enc}(x_j, p_j, m_j)$ . % This is an over-simplification. For technical reasons, the actual construction is somewhat different.

Now, for every query  $q$  specified by  $\mathcal{B}_{\text{sample}}$ , our adversary outputs the query  $f_q$  defined by  $f_q(j, x) = q(\text{Dec}(x, p_j, c_j))$ . Our adversary then obtains an answer  $a$  from the mechanism  $\mathcal{A}_{\text{space}}$ , and feeds  $a$  to  $\mathcal{B}_{\text{sample}}$ . Observe that the “true” value of  $f_q$  w.r.t.  $\mathcal{D}_{\text{space}}$  is the same as the “true” value of  $q$  w.r.t.  $\mathcal{D}_{\text{sample}}$ . Therefore, if  $\mathcal{A}_{\text{space}}$  maintains accuracy in this game against our adversary  $\mathcal{B}_{\text{space}}$ , then in the emulation that  $\mathcal{B}_{\text{space}}$  runs internally we have that  $\mathcal{B}_{\text{space}}$  maintains utility against  $\mathcal{B}_{\text{sample}}$ . Intuitively, we would like to say that this leads to a contradiction, since  $\mathcal{B}_{\text{sample}}$  fails every efficient mechanism it plays against. But this is not accurate, because  $\mathcal{B}_{\text{space}}$  saw the full description of the target distribution  $\mathcal{D}_{\text{sample}}$ , and  $\mathcal{B}_{\text{sample}}$  only fools mechanisms that get to see at most  $t$  samples from this target distribution.

To overcome this, we consider the following modified variant of our adversary, called  $\widehat{\mathcal{B}}_{\text{space}}$ . The modification is that  $\widehat{\mathcal{B}}_{\text{space}}$  does not get to see the full description of  $\mathcal{D}_{\text{sample}}$ . Instead it only gets to see points from the support of  $\mathcal{D}_{\text{sample}}$  that correspond to indices in the set  $I = [n] \setminus J$ . Then, when generating the ciphertexts  $c_j$ , the modified adversary  $\widehat{\mathcal{B}}_{\text{space}}$  encrypts zeroes instead of points  $m_j$  which it is missing. By the security of our MILR scheme, the mechanism  $\mathcal{A}_{\text{space}}$  cannot notice this modification, and hence, assuming that it maintains accuracy against our original adversary  $\mathcal{B}_{\text{space}}$  then it also maintains accuracy against our modified adversary  $\widehat{\mathcal{B}}_{\text{space}}$ . As before, this means that  $\widehat{\mathcal{B}}_{\text{space}}$  maintains accuracy against the emulated  $\mathcal{B}_{\text{sample}}$ . Intuitively, this leads to a contradiction, as  $\widehat{\mathcal{B}}_{\text{space}}$  is using only  $\bar{\tau} \leq t$  points from the target distribution  $\mathcal{D}_{\text{sample}}$ .

We stress that this proof sketch is over-simplified and inaccurate. In particular, the following two technical issues need to be addressed: (1) It is true that  $\widehat{\mathcal{B}}_{\text{space}}$  uses only  $\bar{\tau}$  points from the support of the target distribution  $\mathcal{D}_{\text{sample}}$ , but these points are not necessarily *sampled* from  $\mathcal{D}_{\text{sample}}$ ; and (2) The modified adversary  $\widehat{\mathcal{B}}_{\text{space}}$  is not computationally efficient because computing the set  $J$  is not efficient. We address these issues, and other informalities made herein, in Appendix B.

## 5 Construction of an MILR Scheme from a Semantically Secure Encryption Scheme

**Construction.** Let  $\lambda' \leq \lambda$  be such that  $\lambda = \text{poly}(\lambda')$ . Given an encryption scheme  $\Pi' = (\text{Gen}', \text{Enc}', \text{Dec}')$  such that  $\text{Gen}'(1^{\lambda'})$  outputs a key uniformly distributed on  $\{0, 1\}^{\lambda'}$  (i.e.,  $x' \leftarrow_R \{0, 1\}^{\lambda'}$ ), we construct an MILR scheme  $\Pi = (\text{Gen}, \text{Param}, \text{Enc}, \text{Dec})$  as follows:

- Gen: On input  $1^\lambda$ , return  $x \leftarrow_R \{0, 1\}^\lambda$ .
- Param: On input  $1^\lambda$ , let  $\mathcal{G}$  be a family of universal hash functions with domain  $\{0, 1\}^\lambda$  and range  $\{0, 1\}^{\lambda'}$ . Return (a description of)  $g \leftarrow_R \mathcal{G}$ .
- Enc: On input  $(x, p, m)$ , parse  $g := p$  (as a description of a hash function), let  $x' = g(x)$  and return  $\text{Enc}'(x', m)$ .
- Dec: On input  $(x, p, c)$ , parse  $g := p$ , let  $x' = g(x)$  and return  $\text{Dec}'(x', c)$ .

Using a standard construction of a universal hash function family, all the algorithms run in time polynomial in  $\lambda$ . Moreover, if  $c \leftarrow \text{Enc}(x, p, m)$ , then  $\text{Dec}(x, p, c) = m$  with probability 1 (as this

holds for  $\text{Enc}'$  and  $\text{Dec}'$ ).

The following two theorems (corresponding to the two security properties in Definition 2.1) establish the security of  $\Pi$  and prove Theorem 2.3.

**Theorem 5.1** (Multi semantic security). *Let  $\Omega(\lambda) \leq \Gamma(\lambda) \leq 2^{o(\lambda)}$  and  $\lambda' = 0.1\lambda$ . If  $\Pi'$  is  $\Gamma(\lambda')$ -secure against uniform (resp. non-uniform) adversaries, then  $\Pi$  is  $\Gamma(\lambda)$ -secure against uniform (resp. non-uniform) adversaries.*

**Theorem 5.2** (Multi-security against bounded preprocessing adversaries). *Let  $\Omega(\lambda) \leq \Gamma(\lambda) \leq 2^{o(\lambda)}$ ,  $\lambda' = 0.1\lambda$  (as in Theorem 5.1). If  $\Pi'$  is  $\Gamma(\lambda')$ -secure against non-uniform adversaries then  $\Pi$  is  $(\Gamma(\lambda), \bar{\tau})$ -secure against space bounded non-uniform preprocessing adversaries, where  $\bar{\tau}(\lambda, s) = \frac{2^s}{\lambda} + 4$ .*

**Remark 5.3.** *Since  $\Gamma(\lambda) \leq 2^{o(\lambda)}$  and  $\lambda' = 0.1\lambda$ , then  $\text{poly}(\Gamma(\lambda')) = \text{poly}(\Gamma(\lambda))$ . Therefore, for the sake of simplicity, we analyze the runtime and advantage of all adversaries (including those that run against  $\Pi'$ ) as functions of  $\lambda$ .*

We prove Theorem 5.1 in Appendix C and Theorem 5.2 in Section 6.

## 6 Multi-security Against a Bounded Preprocessing Adversary

In this section we prove Theorem 5.2. The proof requires specific definitions and notation, defined below.

### 6.1 Preliminaries

**Notation.** Given a sequence of elements  $X = (X_1, \dots, X_n)$  and a subset  $I \subseteq [n]$ , we denote by  $X_I$  the sequence composed of elements with coordinates in  $I$ .

For a random variable  $X$ , denote its min-entropy by  $H_\infty(X)$ . For random variables  $X, Y$  with the same range, denote by  $\Delta(X, Y)$  the statistical distance of their distributions. We say that  $X$  and  $Y$  are  $\gamma$ -close if  $\Delta(X, Y) \leq \gamma$ . We use the notation  $X \leftarrow_R \mathcal{X}$  to indicate that the random variable  $X$  is chosen uniformly at random from the set  $\mathcal{X}$ .

**Dense and bit-fixing sources.** We will use the following definition (see [25, Definition 1]).

**Definition 6.1.** *An  $(n, 2^\lambda)$ -source is a random variable  $X$  with range  $\{0, 1\}^\lambda$ . A source is called*

- $(1 - \delta)$ -dense if for every subset  $I \subseteq [n]$ ,  $H_\infty(X_I) \geq (1 - \delta) \cdot |I| \cdot \lambda$ ,
- $(k, 1 - \delta)$ -dense if it is fixed on at most  $k$  coordinates and is  $(1 - \delta)$ -dense on the rest,
- $k$ -bit-fixing if it is fixed on at most  $k$  coordinates and uniform on the rest.

Namely, the min-entropy of every subset of entries of a  $(1 - \delta)$ -dense source is at most a fraction of  $\delta$  less than what it would be for a uniformly random one.

## 6.2 Key leakage lemma

Let  $X = (X_1, \dots, X_n) \in (\{0, 1\}^\lambda)^n$  be a random variable for  $n$  keys of  $\Pi$  chosen independently and uniformly at random. Let  $Z := F(X)$  be a random variable for the leakage of the adversary. For  $z \in \{0, 1\}^s$ , let  $X_z$  be the random variable chosen from the distribution of  $X$  conditioned on  $F(X) = z$ .

We denote  $G := (G_1, \dots, G_n)$  and  $G(X) := (G_1(X_1), \dots, G_n(X_n))$  the random variable for the hash functions (public parameters) of  $\Pi$ . We will use similar notation for sequences of different lengths (which will be clear from the context).

The proof of Theorem 5.2 is based on the lemma below (proved in Section 6.4), which analyzes the joint distribution  $(G, Z, G(X))$ .

**Lemma 6.2.** *Let  $F : (\{0, 1\}^\lambda)^n \rightarrow \{0, 1\}^s$  be an arbitrary function,  $X = (X_1, \dots, X_n) \leftarrow_R (\{0, 1\}^\lambda)^n$  and denote  $Z := F(X)$ . Let  $\mathcal{G}$  be a family of universal hash functions with domain  $\{0, 1\}^\lambda$  and range  $\{0, 1\}^{\lambda'}$  and let  $G \leftarrow_R (\mathcal{G})^n$ . Let  $\delta > 0, \gamma > 0, s' > s$  be parameters such that  $(1 - \delta)\lambda > \lambda' + \log n + 1$ .*

*Then, there exists a family  $V_{G,Z} = \{V_{\vec{g},z}\}_{\vec{g} \in (\mathcal{G})^n, z \in \{0,1\}^s}$  of convex combinations  $V_{\vec{g},z}$  of  $k$ -bit-fixing  $(n, 2^{\lambda'})$ -sources for  $k = \frac{s' + \log 1/\gamma}{\delta \cdot \lambda}$  such that*

$$\Delta[(G, Z, G(X)), (G, Z, V_{G,Z})] \leq \sqrt{2^{-(1-\delta)\lambda + \lambda' + \log n}} + \gamma + 2^{s-s'}.$$

We obtain the following corollary (which implies the parameters of Theorem 5.2).

**Corollary 6.3.** *In the setting of Lemma 6.2, assuming  $n < 2^{0.15\lambda}$  and sufficiently large  $\lambda$ , the parameters  $s' = s + \lambda, \lambda' = 0.1\lambda, \delta = 0.5, \gamma = 2^{-\lambda}$  give*

$$\Delta[(G, Z, G(X)), (G, Z, V_{G,Z})] \leq 2^{-0.1\lambda}, \quad \text{and} \quad k = \frac{2s}{\lambda} + 4.$$

*Proof.* Set  $s' = s + \lambda, \lambda' = 0.1\lambda, \delta = 0.5, \gamma = 2^{-\lambda}$ . Then, for sufficiently large  $\lambda$ ,  $(1 - \delta)\lambda = 0.5\lambda > 0.1\lambda + 0.15\lambda + 1 > \lambda' + \log n + 1$  (and the condition of Lemma 6.2 holds). We therefore have (for sufficiently large  $\lambda$ ):  $\Delta[(G, Z, G(X)), (G, Z, V_{G,Z})] \leq \sqrt{2^{-(1-\delta)\lambda + \lambda' + \log n}} + \gamma + 2^{s-s'} = \sqrt{2^{-0.4\lambda + \log n}} + 2^{-\lambda} + 2^{-\lambda} \leq 2^{-0.1\lambda}$ , and  $k = \frac{s' + \log 1/\gamma}{\delta \cdot \lambda} = \frac{s + \lambda + \lambda}{0.5 \cdot \lambda} = \frac{2s}{\lambda} + 4$ .

□

## 6.3 The Proof of Theorem 5.2

**Using Lemma 6.2 to prove Theorem 5.2.** Before proving Theorem 5.2, we explain why Lemma 6.2 is needed and how it is used in the proof.

It is easy to prove some weaker statements than Lemma 6.2, but these do not seem to be sufficient for building the MILR scheme (i.e., proving Theorem 5.2). For example, one can easily prove that with high probability, given the leakage  $z$  and hash functions  $\vec{g}$ , there is a large subset of (hashed) keys such that each one of them is almost uniformly distributed. However, the adversary could have knowledge of various relations between the keys of this subset and it is not clear how to

prove security without making assumptions about the resistance of the encryption scheme against related-key attacks.

Moreover, consider a stronger statement, which asserts that with high probability, given the leakage  $z$  and hash functions  $\vec{g}$ , there is a large subset of (hashed) keys that are jointly uniformly distributed. We claim that even this stronger statement may not be sufficient to prove security, since it does not consider the remaining keys outside of the subset. In particular, consider a scenario in which the adversary is able to recover some weak keys outside of the subset. Given this extra knowledge and the leakage  $z$ , the original subset of keys may no longer be distributed uniformly (and may suffer from a significant entropy loss).

Lemma 6.2 essentially asserts that there is a subset of keys that is almost jointly uniformly distributed even if we give the adversary  $z, \vec{g}$  and all the remaining keys. More specifically, given the hash functions  $\vec{g}$  and the leakage  $z$ , according to the lemma, the distribution of the hashed keys  $\vec{g}(X)$  is (close to) a convex combinations  $V_{\vec{g},z}$  of  $k$ -bit-fixing sources. In the proof of Theorem 5.2 we will fix such a  $k$ -bit-fixing source by giving the adversary  $k$  hashed keys (we will do this carefully, making sure that the adversary's advantage does not change significantly). Since the remaining hashed keys are uniformly distributed from the adversary's view, security with respect to these keys follows from the semantic security of the underlying encryption scheme.

*Proof of Theorem 5.2.* Fix a preprocessing procedure  $F$  and let  $\lambda$  be sufficiently large,  $n < 2^{0.15\lambda}$ . By Lemma 6.2 (with parameters set in Corollary 6.3), there exists a family  $V_{G,Z}$  of convex combinations  $V_{\vec{g},z}$  of  $k$ -bit-fixing  $(n, 2^\lambda)$ -sources for  $k = \frac{2^s}{\lambda} + 4$  such that  $\Delta[(G, Z, G(X)), (G, Z, V_{G,Z})] \leq 2^{-0.1\lambda}$ .

**Sampling the index set  $J$  and simplifying the distribution.** We first define how the oracles  $\mathcal{E}_0$  and  $\mathcal{E}_1$  in Definition 2.1 sample  $J$ . The random variable  $J$  is naturally defined when sampling the random variables  $(G, Z, V_{G,Z})$ : given  $\vec{g} \in (\mathcal{G})^n, z \in \{0, 1\}^s$ , sample a  $k$ -bit-fixing source in the convex combination  $V_{\vec{g},z}$  (according to its weight) and let  $J$  be the set of (at least)  $n - k$  indices that are not fixed. This defines a joint distribution on the random variables  $(G, Z, V_{G,Z}, J)$ . Another way to sample from this distribution is to first sample the variables  $(G, Z, V_{G,Z})$  and then sample  $J$  according to its marginal distribution. This defines a randomized procedure for sampling  $J$ . Although the oracles do not sample  $(G, Z, V_{G,Z})$ , we reuse the same sampling procedure for sampling  $J$  given the sample  $(\vec{g}, z, \vec{g}(\vec{x}))$  (if the sample  $(\vec{g}, z, \vec{g}(\vec{x}))$  is not in the support of the distribution of  $(G, Z, V_{G,Z})$ , define  $J = [n]$ ).

Consider a poly( $\Gamma(\lambda)$ )-time algorithm  $\mathcal{B}$ . As encryption queries of  $\mathcal{B}$  to  $\mathcal{E}_0$  and  $\mathcal{E}_1$  are answered with the hash keys  $\vec{g}(\vec{x})$ , then given  $\vec{g}(\vec{x})$ , the interaction of  $\mathcal{B}$  with  $\mathcal{E}_0$  and  $\mathcal{E}_1$  no longer depends on  $\vec{x}$ . Therefore, for  $t = 0, 1$  we define  $\mathcal{E}_t^{(1)}(\vec{g}(\vec{x}), \vec{g}, J, \cdot, \cdot)$  that simulates the interaction of  $\mathcal{E}_t(\vec{x}, \vec{g}, J, \cdot, \cdot)$  with  $\mathcal{B}$ . Instead of sampling  $\vec{x}$ , the oracles directly sample  $(\vec{g}, z, \vec{g}(\vec{x}), J)$  according to their joint distribution before the interaction with  $\mathcal{B}$ . To simplify notation, we denote this joint distribution by  $\mathcal{D}_1$ . Denote

$$\text{Adv}_{\mathcal{B}}(\lambda) = \left| \Pr_{\substack{\mathcal{B}, \text{Enc} \\ (\vec{g}, z, \vec{g}(\vec{x}), J) \leftarrow \mathcal{D}_1}} \left[ \mathcal{B}^{\mathcal{E}_0^{(1)}(\vec{g}(\vec{x}), \vec{g}, J, \cdot, \cdot)}(z, \vec{g}) = 1 \right] - \Pr_{\substack{\mathcal{B}, \text{Enc} \\ (\vec{g}, z, \vec{g}(\vec{x}), J) \leftarrow \mathcal{D}_1}} \left[ \mathcal{B}^{\mathcal{E}_1^{(1)}(\vec{g}(\vec{x}), \vec{g}, J, \cdot, \cdot)}(z, \vec{g}) = 1 \right] \right|.$$

It remains to prove that  $\text{Adv}_{\mathcal{B}}(\lambda) \leq \text{negl}(\Gamma(\lambda))$ .

**Using Lemma 6.2 to switch to a family of convex combinations of bit-fixing sources.**  
We have

$$\Delta[(G, Z, G(X), J), (G, Z, V_{G,Z}, J)] \leq \Delta[(G, Z, G(X)), (G, Z, V_{G,Z})] \leq 2^{-0.1\lambda},$$

where the first inequality follows by the data processing inequality, since  $J$  is computed by applying the same function to the three variables of both distributions, and the second inequality is by Corollary 6.3. Hence, for  $t = 0, 1$  we replace  $\mathcal{E}_t^{(1)}$  that samples from  $\mathcal{D}_1$  with  $\mathcal{E}_t^{(2)}$  that samples from the joint distribution of  $(G, Z, V_{G,Z}, J)$ , which we denote by  $\mathcal{D}_2$ . Since  $\mathcal{B}$  and Enc use independent randomness, by the triangle inequality, the total penalty is at most  $2 \cdot 2^{-0.1\lambda}$ , namely

$$\left| \Pr_{\substack{\mathcal{B}, \text{Enc} \\ (\vec{g}, z, \vec{y}, J) \leftarrow \mathcal{D}_2}} \left[ \mathcal{B}^{\mathcal{E}_0^{(2)}(\vec{y}, \vec{g}, J, \cdot)}(z, \vec{g}) = 1 \right] - \Pr_{\substack{\mathcal{B}, \text{Enc} \\ (\vec{g}, z, \vec{y}, J) \leftarrow \mathcal{D}_2}} \left[ \mathcal{B}^{\mathcal{E}_1^{(2)}(\vec{y}, \vec{g}, J, \cdot)}(z, \vec{g}) = 1 \right] \right| \geq \text{Adv}_{\mathcal{B}}(\lambda) - 2^{-0.1\lambda+1},$$

where we denote a sample from  $\mathcal{D}_2$  by  $(\vec{g}, z, \vec{y}, J)$ .

**Giving the adversary additional input.** Consider a (potentially) more powerful  $\text{poly}(\Gamma(\lambda))$ -time algorithm  $\mathcal{B}_1$  against  $\Pi$  whose input consists of  $(z, \vec{g}, J, \vec{y}_{\bar{J}})$ , where  $\bar{J} = [n] \setminus J$ . Namely, in addition to  $(z, \vec{g})$  the input also consists of  $J$ , as well as the hashed keys  $\vec{y}_{\bar{J}} \in (\{0, 1\}^{\lambda'})^{n-|J|}$  (note that these parameters define a  $|\bar{J}|$ -bit-fixing source). We denote  $in = (z, \vec{g}, J, \vec{y}_{\bar{J}})$ , and

$$\text{Adv}_{\mathcal{B}_1}(\lambda) = \left| \Pr_{\substack{\mathcal{B}_1, \text{Enc} \\ (\vec{g}, z, \vec{y}, J) \leftarrow \mathcal{D}_2}} \left[ \mathcal{B}_1^{\mathcal{E}_0^{(2)}(\vec{y}, \vec{g}, J, \cdot)}(in) = 1 \right] - \Pr_{\substack{\mathcal{B}_1, \text{Enc} \\ (\vec{g}, z, \vec{y}, J) \leftarrow \mathcal{D}_2}} \left[ \mathcal{B}_1^{\mathcal{E}_1^{(2)}(\vec{y}, \vec{g}, J, \cdot)}(in) = 1 \right] \right|.$$

Next, we prove that for any such  $\text{poly}(\Gamma(\lambda))$ -time algorithm  $\mathcal{B}_1$ ,  $\text{Adv}_{\mathcal{B}_1}(\lambda) \leq \text{negl}(\Gamma(\lambda))$ . As any algorithm  $\mathcal{B}$  with input  $(z, \vec{g})$  can be simulated by an algorithm  $\mathcal{B}_1$  with input  $in$  and similar runtime, this implies that  $\text{Adv}_{\mathcal{B}}(\lambda) - 2^{-0.1\lambda+1} \leq \text{negl}(\Gamma(\lambda))$  and hence  $\text{Adv}_{\mathcal{B}}(\lambda) \leq \text{negl}(\Gamma(\lambda))$ , concluding the proof.

**Fixing the adversary's input.** Since both  $\mathcal{E}_0^{(2)}$  and  $\mathcal{E}_1^{(2)}$  sample the input of  $\mathcal{B}_1$  from the same distribution, by an averaging argument, there exists an input  $in^* = in_{\lambda}^* = (z^*, \vec{g}^*, J^*, \vec{y}_{\bar{J}^*}^*)$  such that the advantage of  $\mathcal{B}_1$  remains at least as large when fixing the input to  $in^*$  and sampling from  $\mathcal{D}_2$  conditioned on  $in^*$ . Note that given  $in^*$  sampling from  $\mathcal{D}_2$  reduces to sampling from the  $|\bar{J}^*|$ -bit-fixing source defined by  $(\bar{J}^*, \vec{y}_{\bar{J}^*}^*)$ , i.e., selecting  $\vec{w} \leftarrow_R (\{0, 1\}^{\lambda'})^{|\bar{J}^*|}$ . Therefore,

$$\left| \Pr_{\substack{\mathcal{B}_1, \text{Enc} \\ \vec{w} \leftarrow_R (\{0, 1\}^{\lambda'})^{|\bar{J}^*|}}} \left[ \mathcal{B}_1^{\mathcal{E}_0^{(2)}(in^*, \vec{w}, \cdot)}(in^*) = 1 \right] - \Pr_{\substack{\mathcal{B}_1, \text{Enc} \\ \vec{w} \leftarrow_R (\{0, 1\}^{\lambda'})^{|\bar{J}^*|}}} \left[ \mathcal{B}_1^{\mathcal{E}_1^{(2)}(in^*, \vec{w}, \cdot)}(in^*) = 1 \right] \right| \geq \text{Adv}_{\mathcal{B}_1}(\lambda).$$

**Reducing the security of  $\Pi$  with preprocessing from the (multi-instance) security of  $\Pi'$ .** We now use  $\mathcal{B}_1$  to define a non-uniform poly( $\Gamma(\lambda)$ )-time adversary  $\mathcal{B}_2$  (with no preprocessing) that runs against  $|J^*|$  instances of  $\Pi'$  and has advantage at least  $\text{Adv}_{\mathcal{B}_1}(\lambda)$ . By the semantic security of  $\Pi'$  and a hybrid argument (similarly to the proof of Theorem 5.1), this implies that  $\text{Adv}_{\mathcal{B}_1}(\lambda) \leq \text{negl}(\Gamma(\lambda))$ , concluding the proof.

---

**Algorithm 2**  $\mathcal{B}_2^{\mathcal{E}_{(\cdot)}^{(3)}(\vec{x}, [|J^*|], \cdot, \cdot)}$  ( $\cdot$ )

---

**Setting:**  $\mathcal{B}_2$  is a non-uniform adversary that runs against  $|J^*|$  instances of  $\Pi'$  (defined by  $\mathcal{E}_{(\cdot)}^{(3)}$ ). It gets  $in^* = in_\lambda^* = (z^*, \vec{g}^*, J^*, \vec{y}_{\overline{J^*}}^*)$  as advice.  $\mathcal{B}_2$  has access to  $\mathcal{B}_1^{\mathcal{E}_{(\cdot)}^{(2)}(in^*, \vec{w}, \cdot, \cdot)}$  ( $in^*$ ), which runs against  $\Pi$ .

1.  $\mathcal{B}_2$  gives  $in^*$  to  $\mathcal{B}_1$  as input.
  2.  $\mathcal{B}_2$  answers each query  $(j, m)$  of  $\mathcal{B}_1$  as follows:
    - If  $j \in \overline{J^*}$ ,  $\mathcal{B}_2$  uses the advice string  $in^*$  (which contains  $\vec{y}_j^*$ ) to compute the answer  $\text{Enc}'(\vec{y}_j^*, m)$  and gives it to  $\mathcal{B}_1$ .
    - If  $j \in J^*$ ,  $\mathcal{B}_2$  translates the query  $(j, m)$  to  $(j', m)$ , where  $j' \in [|J^*|]$  is obtained by mapping  $j$  to  $J^*$  (ignoring indices in  $\overline{J^*}$ ).  $\mathcal{B}_2$  then queries its oracle with  $(j', m)$  and forwards the answer to  $\mathcal{B}_1$ .
  3.  $\mathcal{B}_2$  outputs the same output as  $\mathcal{B}_1$ .
- 

The adversary  $\mathcal{B}_2$  is given in Algorithm 2. Note that  $\mathcal{B}_2$  perfectly simulates the oracles of  $\mathcal{B}_1$  given the input  $in^*$ , and hence its advantage is at least  $\text{Adv}_{\mathcal{B}_1}(\lambda)$  as claimed. Finally, it runs in time poly( $\Gamma(\lambda)$ ).  $\square$

## 6.4 Proof of Lemma 6.2

**Proof overview.** We first prove in Lemma 6.4 that  $G = (G_1, \dots, G_t)$  (for some  $t \in [n]$ ) is a good extractor, assuming its input  $Y = (Y_1, \dots, Y_t)$  is  $(1 - \delta)$ -dense, namely, it has sufficient min-entropy for each subset of coordinates (see Lemma 6.4 for the exact statement). Specifically, we prove that  $(G, G(Y))$  is statistically close to  $(G, U)$ , where  $U$  is uniformly distributed over  $(\{0, 1\}^{\lambda'})^t$ . The proof is by a variant of the leftover hash lemma [54] where a sequence of hash functions  $(G_1, \dots, G_t)$  are applied locally to each block of the input (instead of applying a single hash function to the entire input). We note that a related lemma was proved in [49, Lem. 13] in a different setting of communication complexity. Our variant is applicable to a different (mostly wider) range of parameters (such as various values of  $\delta$  and the number of bits extracted,  $t \cdot \lambda'$ ) that is relevant in our setting. Additional (somewhat less related) results were presented in [32, 28].

The remainder of the proof is given in Appendix D and is somewhat similar to [25, Lem. 1]. We fix  $z \in \{0, 1\}^s$  such that  $H_\infty(X_z)$  is sufficiently high. Lemma D.1 (proved in [25, 49, 62]) states that  $X_z$  (before applying  $G$ ) is close to a convex combination of  $(k, 1 - \delta)$ -dense sources (for a sufficiently small value of  $k$ ). We therefore invoke Lemma 6.4 on the high min-entropy coordinates of each such  $(1 - \delta)$ -dense source. This allows to derive Lemma D.2, which states that  $(G, G(X_z))$  is statistically close to  $(G, V_{G,z}^*)$ , where  $V_{G,z}^*$  is a family of convex combinations of  $k$ -bit-fixing  $(n, 2^{\lambda'})$ -sources, indexed by  $\vec{g} \in \mathcal{G}^n$ .

Proposition 1 (which is quite standard) shows that with very high probability (over  $X$  and  $F$ ,



setting  $z \leftarrow F(X)$ ,  $H_\infty(X_z)$  is indeed sufficiently high to apply Lemma D.2 with good parameters. For each such  $z$ , we define  $V_{G,z}$  to be  $V_{G,z}^*$  as above and this essentially results in the family  $V_{G,Z}$  (the remaining values of  $z$  with low  $H_\infty(X_z)$  are easy to handle with small loss in statistical distance). Thus, the statistical distance between  $(G, Z, G(X))$  and  $(G, Z, V_{G,Z})$  can be bounded by combining Proposition 1 and Lemma D.2, establishing Lemma 6.2.

### Block-wise extraction from dense sources.

**Lemma 6.4.** *Let  $Y = (Y_1, \dots, Y_t) \in (\{0, 1\}^\lambda)^t$  be a  $(t, 2^\lambda)$ -source that is  $(1-\delta)$ -dense for  $0 < \delta < 1$ . Let  $\mathcal{G}$  be a family of universal hash functions with domain  $\{0, 1\}^\lambda$  and range  $\{0, 1\}^{\lambda'}$ . Then, for  $G \leftarrow_R (\mathcal{G})^t$  and  $U \leftarrow_R (\{0, 1\}^{\lambda'})^t$ ,*

$$\Delta[(G, G(Y)), (G, U)] \leq \sqrt{2^{-(1-\delta)\lambda + \lambda' + \log t}},$$

assuming that  $(1-\delta)\lambda > \lambda' + \log t + 1$ .

*Proof.* Let  $d := \log |\mathcal{G}|$ . For a random variable  $Q$ , and  $Q'$  an independent copy of  $Q$ , we denote by  $\text{Col}[Q] = \Pr[Q = Q']$  the collision probability of  $Q$ . We have

$$\begin{aligned} \text{Col}[(G, G(Y))] &= \Pr_{G, Y, G', Y'}[(G, G(Y)) = (G', G'(Y'))] = \\ &= \Pr_{G, G'}[G = G'] \cdot \Pr_{G, Y, Y'}[G(Y) = G(Y')] = 2^{-t \cdot d} \cdot \Pr_{G, Y, Y'}[G(Y) = G(Y')]. \end{aligned} \quad (1)$$

For sequences  $Y_1, \dots, Y_t, Y'_1, \dots, Y'_t$ , define  $C = |\{i \mid Y_i = Y'_i\}|$ . We now upper bound the expression  $\Pr[C = c]$ .

Recall that  $Y$  is a  $(1-\delta)$ -dense source, i.e., for every subset  $I \subseteq [t]$ ,  $H_\infty(Y_I) \geq (1-\delta) \cdot |I| \cdot \lambda$ . Fix a subset  $I \subseteq [t]$  such that  $|I| = c$ . Then,

$$\begin{aligned} \Pr[Y_I = Y'_I] &= \sum_{y_I \in (\{0, 1\}^\lambda)^c} (\Pr[Y_I = y_I])^2 \leq \\ &= \max_{y_I} \{\Pr[Y_I = y_I]\} \cdot \sum_{y_I \in (\{0, 1\}^\lambda)^c} \Pr[Y_I = y_I] \leq 2^{-(1-\delta) \cdot c \cdot \lambda}. \end{aligned}$$

Therefore,

$$\begin{aligned} \Pr[C = c] &\leq \sum_{\{I \subseteq [t] \mid |I| = c\}} \Pr[Y_I = Y'_I] \leq \binom{t}{c} \cdot 2^{-(1-\delta) \cdot c \cdot \lambda} \leq \\ &= t^c \cdot 2^{-(1-\delta) \cdot c \cdot \lambda} = 2^{c \cdot (-(1-\delta)\lambda + \log t)}. \end{aligned} \quad (2)$$

We have

$$\Pr_{G, Y, Y'}[G(Y) = G(Y')] = \sum_{c=0}^t \Pr_{Y, Y'}[C = c] \cdot \Pr_{G, Y, Y'}[G(Y) = G(Y') \mid C = c].$$

For each coordinate  $i$  such that  $Y_i \neq Y'_i$ ,  $\Pr_{G_i}(G_i(Y_i) = G_i(Y'_i)) = 2^{-\lambda'}$  as  $G_i$  is selected uniformly from a family of universal hash functions. Since  $G = (G_1, \dots, G_t)$  contains  $t$  independent copies selected uniformly from  $\mathcal{G}$ ,

$$\Pr_G[G(Y) = G(Y') \mid C = c] = 2^{-\lambda' \cdot (t-c)}.$$

Hence, using (2) we obtain

$$\begin{aligned} \Pr_{G, Y, Y'}[G(Y) = G(Y')] &= \sum_{c=0}^t \Pr[C = c] \cdot 2^{-\lambda' \cdot (t-c)} \leq \\ &\sum_{c=0}^t 2^{c \cdot (-(1-\delta)\lambda + \log t)} \cdot 2^{-\lambda' \cdot (t-c)} = 2^{-\lambda' \cdot t} \cdot \sum_{c=0}^t 2^{-c \cdot ((1-\delta)\lambda - \lambda' - \log t)} = \\ &2^{-\lambda' \cdot t} \cdot \left(1 + \sum_{c=1}^t 2^{-c \cdot ((1-\delta)\lambda - \lambda' - \log t)}\right) \leq 2^{-\lambda' \cdot t} \cdot (1 + 2^{-(1-\delta)\lambda + \lambda' + \log t + 1}), \end{aligned}$$

where the last inequality uses the assumption that  $(1-\delta)\lambda > \lambda' + \log t + 1$ .

Treating distributions as vectors over  $\{0, 1\}^{t \cdot d + t \cdot \lambda'}$  (and abusing notation), we plug the above expression into (1) and deduce

$$\begin{aligned} \|(G, G(Y)) - (G, U)\|_2^2 &= \text{Col}[(G, G(Y))] - 2^{-t \cdot d - t \cdot \lambda'} \leq \\ &2^{-t \cdot d - t \cdot \lambda'} \cdot (1 + 2^{-(1-\delta)\lambda + \lambda' + \log t + 1}) - 2^{-t \cdot d - t \cdot \lambda'} = 2^{-t \cdot d - t \cdot \lambda' - (1-\delta)\lambda + \lambda' + \log t + 1}. \end{aligned}$$

Finally, using the Cauchy–Schwarz inequality, we conclude

$$\begin{aligned} \Delta[(G, G(Y)), (G, U)] &\leq 1/2 \cdot \sqrt{2^{t \cdot d + t \cdot \lambda'}} \cdot \|(G, G(Y)) - (G, U)\|_2 \leq \\ &1/2 \cdot \sqrt{2^{t \cdot d + t \cdot \lambda'}} \cdot \sqrt{2^{-t \cdot d - t \cdot \lambda' - (1-\delta)\lambda + \lambda' + \log t + 1}} < \sqrt{2^{-(1-\delta)\lambda + \lambda' + \log t}}. \end{aligned}$$

□

## A Additional preliminaries from cryptography

A *probability ensemble* is an infinite sequence of probability distributions. Computational indistinguishability is defined as follows.

**Definition A.1.** *Two probability ensembles  $\mathcal{X} = \{X_\lambda\}_{\lambda \in \mathbb{N}}$  and  $\mathcal{Y} = \{Y_\lambda\}_{\lambda \in \mathbb{N}}$  are computationally indistinguishable, denoted as  $\mathcal{X} \equiv_c \mathcal{Y}$ , if for every non-uniform probabilistic polynomial-time distinguisher  $D$  there exists a negligible function  $\text{negl}$  such that:*

$$\left| \Pr_{x \leftarrow X_\lambda} [D(x) = 1] - \Pr_{y \leftarrow Y_\lambda} [D(y) = 1] \right| \leq \text{negl}(\lambda).$$

**Definition A.2** ([48]). *A non-uniform polynomial-time algorithm is a pair  $(M, \bar{a})$ , where  $M$  is a two-input polynomial-time Turing machine and  $\bar{a} = a_1, a_2, \dots$  is an infinite sequence of strings*

such that  $|a_\lambda| = \text{poly}(\lambda)$ . For every  $x$ , we consider the computation of  $M$  on the input pair  $(x, a_{|x|})$ . Intuitively,  $a_\lambda$  can be thought of as extra “advice” supplied from the “outside” (together with the input  $x \in \{0, 1\}^\lambda$ ). We stress that  $M$  gets the same advice (i.e.,  $a_\lambda$ ) on all inputs of the same length (i.e.,  $\lambda$ ).

Equivalently, a non-uniform polynomial-time algorithm is an infinite sequence of Turing machines,  $M_1, M_2, \dots$ , such that both the length of the description of  $M_\lambda$  and its running time on inputs of length  $\lambda$  are bounded by polynomials in  $\lambda$  (fixed for the entire sequence). Machine  $M_\lambda$  is used only on inputs of length  $\lambda$ .

**Encryption schemes.** An encryption scheme is a tuple of efficient algorithms  $(\text{Gen}, \text{Enc}, \text{Dec})$  with the following syntax:

- $\text{Gen}$  is a randomized algorithm that takes as input a security parameter  $\lambda$  and outputs a  $\lambda$ -bit secret key. Formally,  $x \leftarrow \text{Gen}(1^\lambda)$ .
- $\text{Enc}$  is a randomized algorithm that takes as input a secret key  $x$ , a public parameter  $p$ , and a message  $m \in \{0, 1\}$ , and outputs a ciphertext  $c \in \{0, 1\}^{\text{poly}(\lambda)}$ . Formally,  $c \leftarrow \text{Enc}(x, m)$ .
- $\text{Dec}$  is a deterministic algorithm that takes as input a secret key  $x$ , a public parameter  $p$ , and a ciphertext  $c$ , and outputs a decrypted message  $m'$ . If the ciphertext  $c$  was an encryption of  $m$  under the key  $x$  with the parameter  $p$ , then  $m' = m$ . Formally, if  $c \leftarrow \text{Enc}(x, m)$ , then  $\text{Dec}(x, c) = m$  with probability 1.

Let us recall the formal definition of security of an encryption scheme. Consider a pair of oracles  $\mathcal{E}_0$  and  $\mathcal{E}_1$ , where  $\mathcal{E}_1(x, \cdot)$  takes as input a message  $m$  and returns  $\text{Enc}(m, x)$ , and where  $\mathcal{E}_0(x, \cdot)$  takes the same input but returns  $\text{Enc}(0, x)$ . An encryption scheme  $(\text{Gen}, \text{Enc}, \text{Dec})$  is *secure* if no computationally efficient adversary can tell whether it is interacting with  $\mathcal{E}_0$  or with  $\mathcal{E}_1$ . Formally,

**Definition A.3.** Let  $\lambda$  be a security parameter, and let  $\Gamma = \Gamma(\lambda)$  be a function of  $\lambda$ . An encryption scheme  $(\text{Gen}, \text{Enc}, \text{Dec})$  is  $\Gamma$ -secure if for every  $\text{poly}(\Gamma(\lambda))$ -time adversary  $\mathcal{B}$  there exists a negligible function  $\text{negl}$  such that:

$$\left| \Pr_{\substack{x \leftarrow \text{Gen}(1^\lambda) \\ \mathcal{B}, \text{Enc}}} \left[ \mathcal{B}^{\mathcal{E}_0(x, \cdot)} = 1 \right] - \Pr_{\substack{x \leftarrow \text{Gen}(1^\lambda) \\ \mathcal{B}, \text{Enc}}} \left[ \mathcal{B}^{\mathcal{E}_1(x, \cdot)} = 1 \right] \right| \leq \text{negl}(\Gamma(\lambda)),$$

where the probabilities are over sampling  $x \leftarrow \text{Gen}(1^\lambda)$  and over the randomness of  $\mathcal{B}$  and  $\text{Enc}$ .

## B Proof of Theorem 4.1

In this section we prove Theorem 4.1. We begin by recalling additional preliminaries from [53, 81]. As a stepping stone towards proving their negative results for the ADA problem, [53, 81] showed a negative result for a restricted family of mechanisms, called *natural* mechanisms. These are algorithms that can only evaluate the query on the sample points they are given. Formally,

**Definition B.1** ([53]). A mechanism  $\mathcal{A}$  is natural if, when given a sample  $Y = (y_1, \dots, y_t) \in X^t$  and a query  $q : X \rightarrow \{-1, 0, 1\}$ , the algorithm returns an answer that is a function only of  $(q(y_1), \dots, q(y_t))$ . In particular, it cannot evaluate  $q$  on other data points of its choice.

Consider the following variant of the adaptive game that appeared in Figure 1 (in the introduction), in which the underlying distribution is fixed to be uniform. See algorithm `AdaptiveGameOne`.

---

**Algorithm 3** `AdaptiveGameOne`( $\mathcal{A}, \mathcal{B}, t, k$ )

---

**Notation:**  $t$  is the sample size and  $n = 2000t$  is the domain size.

1. The mechanism  $\mathcal{A}$  gets a sample  $Y = (y_1, \dots, y_t)$  containing  $t$  i.i.d. samples from the uniform distribution on  $[n]$ . This sample is not given to  $\mathcal{B}$ .
  2. For round  $i = 1, 2, \dots, k$ :
    - (a) The adversary  $\mathcal{B}$  specifies a query  $q_i : [n] \rightarrow \{-1, 0, 1\}$
    - (b) The mechanism  $\mathcal{A}$  obtains  $q_i$  and responds with an answer  $a_i \in [-1, 1]$   
% If  $\mathcal{A}$  is *natural* then it only gets  $(q_i(y_1), \dots, q_i(y_t))$  instead of the full description of  $q_i$ .
    - (c)  $a_i$  is given to  $\mathcal{B}$
  3. The outcome of the game is one if  $\exists i$  s.t.  $\left| a_i - \frac{1}{n} \sum_{j \in [n]} q_i(j) \right| > 1/10$ , and zero otherwise.
- 

**Theorem B.2** ([53, 81]). *There exists a computationally efficient adversary  $\mathcal{B}$  such that for every natural mechanism  $\mathcal{A}$  and every large enough  $t$  and  $k = \Theta(t^2)$  it holds that*

$$\Pr[\text{AdaptiveGameOne}(\mathcal{A}, \mathcal{B}, t, k) = 1] > \frac{3}{4}.$$

We remark that this theorem holds even if  $M$  is computationally unbounded. That is, there exists a computationally efficient adversary  $\mathcal{A}$  that fails *every* natural mechanism  $M$  in the ADA problem, even when the underlying distribution is uniform.

Our first observation is that the proof of [53, 81] continues to hold even if in Step 1 of the game the mechanism *chooses* its  $t$  input points, rather than sampling them uniformly. We refer to such a mechanism as a *selecting* mechanism. See algorithm `AdaptiveGameTwo`.

---

**Algorithm 4** `AdaptiveGameTwo`( $\mathcal{A}, \mathcal{B}, t, k$ )

---

**Notations:**  $t$  is the sample size and  $n = 2000t$  is the domain size.

1. **The mechanism  $\mathcal{A}$  chooses a subset  $Y = (y_1, \dots, y_t) \subseteq [n]$  of size  $t$ . This subset is not given to  $\mathcal{B}$ .**
  2. For round  $i = 1, 2, \dots, k$ :
    - (a) The adversary  $\mathcal{B}$  specifies a query  $q_i : [n] \rightarrow \{-1, 0, 1\}$
    - (b) The mechanism  $\mathcal{A}$  obtains  $q_i$  and responds with an answer  $a_i \in [-1, 1]$   
% If  $\mathcal{A}$  is *natural* then it only gets  $(q_i(y_1), \dots, q_i(y_t))$  instead of the full description of  $q_i$ .
    - (c)  $a_i$  is given to  $\mathcal{B}$
  3. The outcome of the game is one if  $\exists i$  s.t.  $\left| a_i - \frac{1}{n} \sum_{j \in [n]} q_i(j) \right| > 1/10$ , and zero otherwise.
- 

**Theorem B.3** (extension of [53, 81]). *There exists a computationally efficient adversary  $\mathcal{B}$  such that for every natural mechanism  $\mathcal{A}$  and every large enough  $t$  and  $k = \Theta(t^2)$  it holds that*

$$\Pr[\text{AdaptiveGameTwo}(\mathcal{A}, \mathcal{B}, t, k) = 1] > \frac{3}{4}.$$

**Remark B.4.** *This extension is not directly necessary for our purposes, but it simplifies our analysis. Recall that we ultimately want to show a negative result for efficient mechanisms that can*

preprocess the underlying distribution and store a limited amount of space about it. Intuitively, using our notion of MILR, we will force such a mechanism to store useful information only about a bounded number of points from the underlying distribution. But these points are not necessarily sampled from the underlying distribution, and hence, it will be more convenient for us to rely on a hardness result for selecting natural mechanisms.

## B.1 From selecting and natural to space boundedness

We next “lift” the negative result for selecting and natural mechanisms (Theorem B.3) to obtain space a lower bound. To this end, recall `AdaptiveGameSpace`, presented in Section 4, where the mechanism initially gets the full description of the underlying distribution, but it must shrink it into an  $s$ -bit summary  $z$ .

We prove the following theorem.

**Theorem B.5.** *Assume the existence of an MILR scheme  $\Pi = (\text{Gen}, \text{Param}, \text{Enc}, \text{Dec})$  with security parameter  $\lambda$  that is  $(\Gamma, \bar{\tau})$ -secure against space bounded preprocessing adversaries. Let  $k \leq \text{poly}(\Gamma(\lambda))$  be a function of  $\lambda$ . There exists a  $\text{poly}(\Gamma(\lambda))$ -time adversary  $\mathcal{B}$  such that the following holds. Let  $\mathcal{A}=(\mathcal{A}_1, \mathcal{A}_2)$  be a  $\text{poly}(\Gamma(\lambda))$ -time mechanism with space complexity  $s$  such that  $\bar{\tau}(\lambda, s) \leq O(\sqrt{k})$  (small enough). Then,*

$$\Pr[\text{AdaptiveGameSpace}(\mathcal{A}, \mathcal{B}, s, k) = 1] > \frac{2}{3}.$$

Furthermore, the underlying distribution defined by the adversary  $\mathcal{B}$  can be fully described using  $O(\sqrt{k} \cdot \lambda)$  bits, it is sampleable in  $\text{poly}(\Gamma(\lambda))$ -time, and elements sampled from this distribution can be represented using  $O(\lambda + \log(k))$  bits.

*Proof.* Our task is to show that there is an adversary that fails every (efficient) mechanism that plays in `AdaptiveGameSpace`. What we know is that there is an adversary that fails every natural mechanism that plays in `AdaptiveGameTwo`. To bridge this gap we consider a wrapper algorithm, called `WrapA`, that can be used as a wrapper over a mechanism  $\mathcal{A}$  designed for `AdaptiveGameSpace`, in order to transform it into a *natural* mechanism that plays in `AdaptiveGameTwo`. Consider algorithm `WrapA`.

We stress that algorithm `WrapA` is not necessarily computationally efficient, since computing  $J(\mathcal{A}_1, \vec{x}, z, \vec{p})$  might not be efficient. This does not break our analysis because Theorem B.3 holds even for computationally unbounded mechanisms (that are natural). Indeed, for every mechanism  $\mathcal{A}$  (designed for `AdaptiveGameSpace`) we have that `WrapA` $\circ\mathcal{A}$  is a *natural* mechanism for playing in `AdaptiveGameTwo`. Therefore, the adversary  $\mathcal{B}$  guaranteed by Theorem B.3 causes `WrapA` $\circ\mathcal{A}$  to fail in `AdaptiveGameTwo`, provided that

$$\bar{\tau}(\lambda, s) = |I| \leq t = \Theta(\sqrt{k}), \tag{3}$$

where  $I$  is the set from Step 5 of algorithm `WrapA` and  $t$  is the sample size in `AdaptiveGameTwo`. We continue with the analysis assuming that Inequality (3) holds.

Recall that our task is to design an adversary that causes  $\mathcal{A}$  to fail in `AdaptiveGameSpace`. To this end, let us consider a modified version of algorithm `WrapA`, called `WrapB`.

---

**Algorithm 5 WrapA**

---

**Setting:** Algorithm `WrapA` has access to an algorithm  $\mathcal{A}=(\mathcal{A}_1, \mathcal{A}_2)$  with space complexity  $s$  (designed for `AdaptiveGameSpace`) and it interacts with an adversary  $\mathcal{B}$  (designed for `AdaptiveGameTwo`).

1. Let  $(\text{Gen}, \text{Param}, \text{Enc}, \text{Dec})$  be a  $(\Gamma, \bar{\tau})$ -secure MILR scheme with security parameter  $\lambda$ .
  2. Sample  $n$  keys  $x_1, \dots, x_n \in \{0, 1\}^\lambda$  from `Gen`, and define  $\mathcal{D}$  to be the uniform distribution over the set  $\{(j, x_j)\}_{j=1}^n$ .
  3. Let  $z \leftarrow \mathcal{A}_1(\mathcal{D})$ .
  4. Sample  $n$  public parameters  $p_1, \dots, p_n \in \{0, 1\}^{\text{poly}(\lambda)}$  from `Param`.
  5. Let  $J \leftarrow J(\mathcal{A}_1, \vec{x}, z, \vec{p}) \subseteq [n]$  be the set of “hidden coordinates”, as in Definition 2.1. Denote  $I = [n] \setminus J$ . The set  $I$  corresponds to the subset  $Y$  from Step 1 of `AdaptiveGameTwo`.
  6. Instantiate  $\mathcal{A}_2$  with  $z$ .
  7. For round  $i = 1, 2, \dots, k$ :
    - (a) Obtain a query  $q_i$ , restricted to coordinates in  $I$ . That is, for every  $j \in I$  we obtain  $q_i(j)$ .
    - (b) For  $j \in I$  let  $c_{i,j} = \text{Enc}(x_i, p_i, q_i(j))$ . For  $j \notin I$  let  $c_{i,j} = \text{Enc}(x_i, p_i, 0)$ .
    - (c) Define the query  $f_i : [n] \times \{0, 1\}^\lambda \rightarrow \{-1, 0, 1\}$  where  $f_i(j, x) = \text{Dec}(x, p_j, c_{i,j})$ . The description of  $f_i$  includes  $\{c_{i,j}\}_{j \in [n]}$  and  $\{p_j\}_{j \in [n]}$ .
    - (d) Give  $f_i$  to  $\mathcal{A}_2$  and obtain an answer  $a_i \in [-1, 1]$ .
    - (e) Output  $a_i$ .
- 

---

**Algorithm 6 WrapB**

---

**Setting:** Algorithm `WrapB` has access to an algorithm  $\mathcal{A}=(\mathcal{A}_1, \mathcal{A}_2)$  with space complexity  $s$  (designed for `AdaptiveGameSpace`) and it interacts with an adversary  $\mathcal{B}$  (designed for `AdaptiveGameTwo`).

1. Let  $(\text{Gen}, \text{Param}, \text{Enc}, \text{Dec})$  be a  $(\Gamma, \bar{\tau})$ -secure MILR scheme with security parameter  $\lambda$ .
  2. Sample  $n$  keys  $x_1, \dots, x_n \in \{0, 1\}^\lambda$  from `Gen`, and define  $\mathcal{D}$  to be the uniform distribution over the set  $\{(j, x_j)\}_{j=1}^n$ .
  3. Let  $z \leftarrow \mathcal{A}_1(\mathcal{D})$ .
  4. Sample  $n$  public parameters  $p_1, \dots, p_n \in \{0, 1\}^{\text{poly}(\lambda)}$  from `Param`.
  5. Let  $J \leftarrow J(\mathcal{A}_1, \vec{k}, z, \vec{p}) \subseteq [n]$  be the set of “hidden coordinates”, as in Definition 2.1. Denote  $I = [n] \setminus J$ . The set  $I$  corresponds to the subset  $Y$  from Step 1 of `AdaptiveGameTwo`.
  6. Instantiate  $\mathcal{A}_2$  with  $z$ .
  7. For round  $i = 1, 2, \dots, k$ :
    - (a) Obtain a query  $q_i$  (not restricted to coordinates in  $I$ )
    - (b) For  $j \in [n]$  let  $c_{i,j} = \text{Enc}(x_i, p_i, q_i(j))$ .
    - (c) Define the query  $f_i : [n] \times \{0, 1\}^\lambda \rightarrow \{-1, 0, 1\}$  where  $f_i(j, x) = \text{Dec}(x, p_j, c_{i,j})$ . The description of  $f_i$  includes  $\{c_{i,j}\}_{j \in [n]}$  and  $\{p_j\}_{j \in [n]}$ .
    - (d) Give  $f_i$  to  $\mathcal{A}_2$  and obtain an answer  $a_i \in [-1, 1]$ .
    - (e) Output  $a_i$ .
-

Unlike  $\text{WrapA}$ , the modified mechanism  $\text{WrapB}$  does not guarantee that the resulting  $\text{WrapB} \circ \mathcal{A}$  is natural. Nevertheless, as we next explain, for every computationally efficient mechanism  $\mathcal{A}$  it still holds that  $\text{WrapB} \circ \mathcal{A}$  fails in  $\text{AdaptiveGameTwo}$ . To see this, consider algorithm  $\text{WrapC}$ , which is another variant of  $\text{WrapA}$  and  $\text{WrapB}$  that has only oracle access to encryptions, via an oracle  $\mathcal{E}$ . We think of  $\text{WrapC}$  as an adversary to the MILLR scheme.

---

**Algorithm 7**  $\text{WrapC}$

---

**Setting:** Algorithm  $\text{WrapC}$  has access to an algorithm  $\mathcal{A}=(\mathcal{A}_1, \mathcal{A}_2)$  with space complexity  $s$  (designed for  $\text{AdaptiveGameSpace}$ ) and it interacts with an adversary  $\mathcal{B}$  (designed for  $\text{AdaptiveGameTwo}$ ). It also has access to an encryption oracle  $\mathcal{E}$ .

**Input:** An element  $z$  (supposedly computed by  $\mathcal{A}_1$ ), and a collection of  $n$  public parameters  $p_1, \dots, p_n \in \{0, 1\}^{\text{poly}(\lambda)}$ .

1. Instantiate  $\mathcal{A}_2$  with  $z$ .
  2. For round  $i = 1, 2, \dots, k$ :
    - (a) Obtain a query  $q_i$  (not restricted to a set of coordinates)
    - (b) **For**  $j \in [n]$  **let**  $c_{i,j} = \mathcal{E}(i, q_i(j))$ .
    - (c) Define the query  $f_i : [n] \times \{0, 1\}^\lambda \rightarrow \{-1, 0, 1\}$  where  $f_i(j, x) = \text{Dec}(x, p_j, c_{i,j})$ . The description of  $f_i$  includes  $\{c_{i,j}\}_{j \in [n]}$  and  $\{p_j\}_{j \in [n]}$ .
    - (d) Give  $f_i$  to  $\mathcal{A}_2$  and obtain an answer  $a_i \in [-1, 1]$ .
    - (e) Output  $a_i$ .
- 

Observe that if  $\text{WrapC}$  has access to the oracle  $\mathcal{E}_1$  then it behaves identically to  $\text{WrapA}$ , and if it has access to the oracle  $\mathcal{E}_0$  then it behaves identically to  $\text{WrapB}$ . Formally,

$$\begin{aligned}
& \Pr [\text{AdaptiveGameTwo} ((\text{WrapB} \circ \mathcal{A}), \mathcal{B}, t, k) = 1] \\
&= \Pr_{\substack{\vec{x}, \vec{p}, \mathcal{B}, \mathcal{A}_2, \text{Enc} \\ z \leftarrow \mathcal{A}_1(\vec{x}) \\ J \leftarrow J(\mathcal{A}_1, \vec{x}, z, \vec{p})}} \left[ \text{AdaptiveGameTwo} \left( (\text{WrapC}^{\mathcal{E}_0(\vec{x}, \vec{p}, J, \cdot)}(z, \vec{p}) \circ \mathcal{A}), \mathcal{B}, t, k \right) = 1 \right] \\
&\geq \Pr_{\substack{\vec{x}, \vec{p}, \mathcal{B}, \mathcal{A}_2, \text{Enc} \\ z \leftarrow \mathcal{A}_1(\vec{x}) \\ J \leftarrow J(\mathcal{A}_1, \vec{x}, z, \vec{p})}} \left[ \text{AdaptiveGameTwo} \left( (\text{WrapC}^{\mathcal{E}_1(\vec{x}, \vec{p}, J, \cdot)}(z, \vec{p}) \circ \mathcal{A}), \mathcal{B}, t, k \right) = 1 \right] - \text{negl}(\Gamma(\lambda)) \\
&= \Pr [\text{AdaptiveGameTwo} ((\text{WrapA} \circ \mathcal{A}), \mathcal{B}, t, k) = 1] - \text{negl}(\Gamma(\lambda)) \\
&\geq \frac{3}{4} - \text{negl}(\Gamma(\lambda)) \geq \frac{2}{3}.
\end{aligned}$$

Finally, observe that instead of composing  $\text{WrapB}$  with  $\mathcal{A}$  (and thinking about them as a single mechanism) we could compose it instead with the adversary  $\mathcal{B}$  to obtain a composed adversary that plays in  $\text{AdaptiveGameSpace}$ . That is,

$$\text{AdaptiveGameTwo} ((\text{WrapB} \circ \mathcal{A}), \mathcal{B}, t, k) \equiv \text{AdaptiveGameSpace} (\mathcal{A}, (\text{WrapB} \circ \mathcal{B}), s, k),$$

and hence,

$$\begin{aligned} & \Pr [\text{AdaptiveGameSpace}(\mathcal{A}, (\text{WrapB} \circ \mathcal{B}), s, k) = 1] \\ &= \Pr [\text{AdaptiveGameTwo}((\text{WrapB} \circ \mathcal{A}), \mathcal{B}, t, k) = 1] \\ &\geq \frac{2}{3}. \end{aligned}$$

□

Theorem 4.1 now follows by instantiating Theorem B.5 with our MILR scheme, as specified in Theorem 2.3.

## C Missing Proofs from Section 5

*Proof of Theorem 5.1.* We prove the result for a uniform adversary. The proof for a non-uniform adversary is similar. Let  $\mathcal{B}$  be a  $\text{poly}(\Gamma(\lambda))$  adversary for  $\Pi$ . Denote

$$\text{Adv}_{\mathcal{B}}(\lambda) = \left| \Pr_{\vec{x}, \vec{g}, \mathcal{B}, \text{Enc}} \left[ \mathcal{B}^{\mathcal{E}_0(\vec{x}, \vec{g}, [n], \cdot, \cdot)}(\vec{g}) = 1 \right] - \Pr_{\vec{x}, \vec{g}, \mathcal{B}, \text{Enc}} \left[ \mathcal{B}^{\mathcal{E}_1(\vec{x}, \vec{g}, [n], \cdot, \cdot)}(\vec{g}) = 1 \right] \right|.$$

Our goal is to prove that  $\text{Adv}_{\mathcal{B}}(\lambda) \leq \text{negl}(\Gamma(\lambda))$ .

**Reduction of the security of  $\Pi$  from (multi-instance) security of  $\Pi'$ .** We will use  $\mathcal{B}$  to construct an adversary  $\mathcal{B}_1$  for  $n$  instances of  $\Pi'$  with independent keys. We define the corresponding oracles for  $\Pi'$ .

1.  $\mathcal{E}_1^{(1)}(\vec{x}', [n], \cdot, \cdot)$  takes as input an index of a key  $j \in [n]$  and a message  $m$ , and returns  $\text{Enc}'(x'_j, m)$ .
2.  $\mathcal{E}_0^{(1)}(\vec{x}', [n], \cdot, \cdot)$  takes the same inputs and returns  $\text{Enc}'(x'_j, 0)$ .

The advantage of  $\mathcal{B}_1$  is defined as

$$\text{Adv}_{\mathcal{B}_1}(\lambda) = \left| \Pr_{\vec{x}', \mathcal{B}_1, \text{Enc}'} \left[ \mathcal{B}_1^{\mathcal{E}_0^{(1)}(\vec{x}', [n], \cdot, \cdot)}() = 1 \right] - \Pr_{\vec{x}', \mathcal{B}_1, \text{Enc}'} \left[ \mathcal{B}_1^{\mathcal{E}_1^{(1)}(\vec{x}', [n], \cdot, \cdot)}() = 1 \right] \right|.$$

---

**Algorithm 8**  $\mathcal{B}_1^{\mathcal{E}_{(\cdot)}^{(1)}(\vec{x}', [n], \cdot, \cdot)}()$

---

**Setting:**  $\mathcal{B}_1$  runs against  $n$  instances of  $\Pi'$  and has access to  $\mathcal{B}^{\mathcal{E}_{(\cdot)}(\vec{x}, \vec{g}, [n], \cdot, \cdot)}(\vec{g})$ , which runs against  $\Pi$ .

1.  $\mathcal{B}_1$  selects  $\vec{g} \leftarrow_R \mathcal{G}^n$  and gives it to  $\mathcal{B}$  as input.
  2.  $\mathcal{B}_1$  forwards each query  $(j, m)$  of  $\mathcal{B}$  to its oracle and forwards the answers back to  $\mathcal{B}$ .
  3.  $\mathcal{B}_1$  outputs the same output as  $\mathcal{B}$ .
- 

The adversary  $\mathcal{B}_1$  is given in Algorithm 8 and we analyze its runtime and advantage.

For this purpose, we first consider the distribution of the hashed keys of  $\Pi$ , given  $\vec{g}$ . Each key  $x \in \{0, 1\}^\lambda$  is selected uniformly and independently, and is hashed to  $\lambda' = \lambda/10$  bits using a



hash function selected uniformly from a universal hash function family. A standard application of the leftover hash lemma [54] shows that the distribution of each key is  $2^{-\Omega(\lambda)}$ -close to the uniform distribution (in terms of statistical distance). Since the keys are selected independently and  $n < 2^{-o(\lambda)}$ , applying the triangle inequality (summing the statistical distance over all  $n$  keys) shows that their joint distribution is  $2^{-\Omega(\lambda)}$ -close to uniform (given  $\vec{g}$ ).

Since the keys  $\vec{x}' \in (\{0, 1\}^{\lambda'})^n$  of  $\mathcal{B}_1$  are uniformly distributed independently of  $\vec{g}$ , the simulation of  $\mathcal{B}$  incurs an additive loss of  $2^{-\Omega(\lambda)}$ , namely

$$\text{Adv}_{\mathcal{B}_1}(\lambda) \geq \text{Adv}_{\mathcal{B}}(\lambda) - 2^{-\Omega(\lambda)}.$$

The runtime of  $\mathcal{B}_1$  is bounded by  $\text{poly}(\Gamma(\lambda))$  (as the runtime of  $\mathcal{B}$ ).

**Reduction of (multi-instance) security of  $\Pi'$  from single-instance security.** Next, we construct an adversary  $\mathcal{B}_2$  against a single instances of  $\Pi'$  (with  $n = 1$ ), using the multi-instance adversary  $\mathcal{B}_1$ . The advantage of  $\mathcal{B}_2$  is defined as

$$\text{Adv}_{\mathcal{B}_2}(\lambda) = \left| \Pr_{\vec{x}', \mathcal{B}_2, \text{Enc}'} \left[ \mathcal{B}_2^{\mathcal{E}_0^{(1)}(\vec{x}', [1], \cdot, \cdot)}() = 1 \right] - \Pr_{\vec{x}', \mathcal{B}_2, \text{Enc}'} \left[ \mathcal{B}_2^{\mathcal{E}_1^{(1)}(\vec{x}', [1], \cdot, \cdot)}() = 1 \right] \right|.$$

---

**Algorithm 9**  $\mathcal{B}_2^{\mathcal{E}_{(\cdot)}^{(1)}(\vec{x}', [1], \cdot, \cdot)}()$

---

**Setting:**  $\mathcal{B}_2$  runs against a single instance of  $\Pi'$  and has access to  $\mathcal{B}_1^{\mathcal{E}_{(\cdot)}^{(1)}(\vec{x}', [n], \cdot, \cdot)}()$ , which runs against  $n$  instances of  $\Pi'$ .

1.  $\mathcal{B}_2$  selects an index  $i \in [n]$  uniformly at random, and selects  $n - 1$  keys  $x'_1, \dots, x'_{i-1}, x'_{i+1}, \dots, x'_n$ , each chosen uniformly at random from  $\{0, 1\}^{\lambda'}$ .
  2.  $\mathcal{B}_2$  answers each query  $(j, m)$  of  $\mathcal{B}_1$  as follows:
    - If  $i = j$ ,  $\mathcal{B}_2$  forwards query  $(1, m)$  to its oracle and returns the answer to  $\mathcal{B}_1$ .
    - If  $j < i$ ,  $\mathcal{B}_2$  returns  $\text{Enc}'(x'_j, m)$  to  $\mathcal{B}_1$ .
    - If  $j > i$ ,  $\mathcal{B}_2$  returns  $\text{Enc}'(x'_j, 0)$  to  $\mathcal{B}_1$ .
  3.  $\mathcal{B}_2$  outputs the same output as  $\mathcal{B}_1$ .
- 

Adversary  $\mathcal{B}_2$  is given in Algorithm 9. In order to relate  $\text{Adv}_{\mathcal{B}_2}(\lambda)$  to  $\text{Adv}_{\mathcal{B}_1}(\lambda)$ , we use a hybrid argument. Specifically, for  $i \in \{0, \dots, n\}$  define  $\mathcal{E}^{(i)}$  as follows.

- $\mathcal{E}^{(i)}(\vec{x}', [n], \cdot, \cdot)$  takes as input an index of a key  $j \in [n]$  and a message  $m$ . If  $j \leq i$ , returns  $\text{Enc}'(x'_j, m)$ . Otherwise,  $j > i$  and it returns  $\text{Enc}'(x'_j, 0)$ .

Note that  $\mathcal{E}^{(0)} = \mathcal{E}_0^{(1)}$  and  $\mathcal{E}^{(n)} = \mathcal{E}_1^{(1)}$ . For  $i \in [n]$ , define

$$\text{Adv}_{\mathcal{B}_1}^{(i)}(\lambda) = \left| \Pr_{\vec{x}', \mathcal{B}_1, \text{Enc}'} \left[ \mathcal{B}_1^{\mathcal{E}^{(i-1)}(\vec{x}', [n], \cdot, \cdot)}() = 1 \right] - \Pr_{\vec{x}', \mathcal{B}_1, \text{Enc}'} \left[ \mathcal{B}_1^{\mathcal{E}^{(i)}(\vec{x}', [n], \cdot, \cdot)}() = 1 \right] \right|.$$

We have

$$\begin{aligned} \text{Adv}_{\mathcal{B}_1}(\lambda) &= \sum_{i=1}^n \left( \Pr_{\vec{x}', \mathcal{B}_1, \text{Enc}'} \left[ \mathcal{B}_1^{\mathcal{E}^{(i-1)}}(\vec{x}', [n], \cdot, \cdot)() = 1 \right] - \Pr_{\vec{x}', \mathcal{B}_1, \text{Enc}'} \left[ \mathcal{B}_1^{\mathcal{E}^{(i)}}(\vec{x}', [n], \cdot, \cdot)() = 1 \right] \right) \leq \\ & \sum_{i=1}^n \left| \Pr_{\vec{x}', \mathcal{B}_1, \text{Enc}'} \left[ \mathcal{B}_1^{\mathcal{E}^{(i-1)}}(\vec{x}', [n], \cdot, \cdot)() = 1 \right] - \Pr_{\vec{x}', \mathcal{B}_1, \text{Enc}'} \left[ \mathcal{B}_1^{\mathcal{E}^{(i)}}(\vec{x}', [n], \cdot, \cdot)() = 1 \right] \right| = \sum_{i=1}^n \text{Adv}_{\mathcal{B}_1}^{(i)}(\lambda). \end{aligned}$$

Depending on the oracle of  $\mathcal{B}_2$ ,  $\mathcal{B}_1$  interacts with either  $\mathcal{E}^{(i-1)}$  or  $\mathcal{E}^{(i)}$  and hence the advantage of  $\mathcal{B}_2$  given that  $i$  is selected is  $\text{Adv}_{\mathcal{B}_1}^{(i)}(\lambda)$ . Since  $i \in [n]$  is uniform,

$$\text{Adv}_{\mathcal{B}_2}(\lambda) = \frac{1}{n} \sum_{i=1}^n \text{Adv}_{\mathcal{B}_1}^{(i)}(\lambda) \geq \frac{1}{n} \text{Adv}_{\mathcal{B}_1}(\lambda).$$

The runtime of  $\mathcal{B}_2$  is bounded by  $n \cdot \text{poly}(\Gamma(\lambda)) = \text{poly}(\Gamma(\lambda))$ . Hence, the security assumption about  $\Pi'$  implies that  $\text{Adv}_{\mathcal{B}_2}(\lambda) \leq \text{negl}(\Gamma(\lambda))$ . We therefore conclude that  $\text{Adv}_{\mathcal{B}_1}(\lambda)$  and hence  $\text{Adv}_{\mathcal{B}}(\lambda)$  is negligible in  $\Gamma(\lambda)$ . □

## D Missing Proofs from Section 6

The following result is taken from [25, Claim 2] (its original variant is due to [49, 62]).

**Lemma D.1.** *Fix  $z \in \{0, 1\}^s$  and let  $s_z = \lambda \cdot n - H_\infty(X_z)$  be the min-entropy deficiency of  $X_z$ . Let  $\gamma > 0$  be arbitrary. Then, for every  $\delta > 0$ ,  $X_z$  is  $\gamma$ -close to a convex combination of finitely many  $(k, 1 - \delta)$ -dense sources for*

$$k = \frac{s_z + \log 1/\gamma}{\delta \cdot \lambda}.$$

We now prove the following lemma.

**Lemma D.2.** *Fix  $z \in \{0, 1\}^s$  and let  $s_z = \lambda \cdot n - H_\infty(X_z)$ .*

*Let  $\mathcal{G}$  be a family of universal hash functions with domain  $\{0, 1\}^\lambda$  and range  $\{0, 1\}^{\lambda'}$  and let  $G \leftarrow_R (\mathcal{G})^n$ . Let  $\delta > 0, \gamma > 0$  be parameters and assume that  $(1 - \delta)\lambda > \lambda' + \log n + 1$ .*

*Then, there exists a family of convex combinations  $V_{G,z}^* = \{V_{\vec{g},z}^*\}_{\vec{g} \in (\mathcal{G})^n}$  of  $k$ -bit-fixing  $(n, 2^{\lambda'})$ -sources where*

$$k = \frac{s_z + \log 1/\gamma}{\delta \cdot \lambda}$$

*such that*

$$\Delta[(G, G(X_z)), (G, V_{G,z}^*)] \leq \sqrt{2^{-(1-\delta)\lambda + \lambda' + \log n}} + \gamma.$$

Lemma D.2 is based on lemmas 6.4 and D.1.

*Proof.* Fix  $z \in \{0, 1\}^s$ . Let  $s_z, \gamma > 0, \delta > 0$  be as in the lemma.

Then, by Lemma D.1 (applied with parameters  $s_z, \gamma, \delta$ ),  $\Delta[X_z, X_z^*] \leq \gamma$ , where  $X_z^*$  is a convex combination of finitely many  $(k, 1 - \delta)$ -dense  $(n, 2^\lambda)$ -sources for  $k = \frac{s_z + \log 1/\gamma}{\delta \cdot \lambda}$ .

Let  $X^*$  be a  $(k, 1 - \delta)$ -dense source in the convex combination  $X_z^*$ . Let  $V^*$  be the corresponding  $k$ -bit-fixing source ( $X^*$  and  $V^*$  are fixed on the same indices to the same values).

Let  $J \subseteq [n]$  denote the set of indices on which  $X^*$  is not fixed and further denote  $t := |J| \geq n - k$ . Then,  $X_J^*$  is a  $(t, 2^\lambda)$ -source, which is  $(1 - \delta)$ -dense.

Since we assume that  $(1 - \delta)\lambda > \lambda' + \log n + 1 > \lambda' + \log t + 1$ , we can apply Lemma 6.4 to  $X_J^*$  (with  $\delta$  as above). We conclude that

$$\Delta[(G, G(X_J^*)), (G, U)] \leq \sqrt{2^{-(1-\delta)\lambda + \lambda' + \log t}} < \sqrt{2^{-(1-\delta)\lambda + \lambda' + \log n}},$$

where  $U \leftarrow_R (\{0, 1\}^{\lambda'})^t$ .

Since the remaining indices of  $\bar{J}$  in  $X^*$  and  $V^*$  are fixed to the same values,

$$\Delta[(G, G(X^*)), (G, V_G^*)] \leq \sqrt{2^{-(1-\delta)\lambda + \lambda' + \log n}},$$

where  $V_G^* \in (\{0, 1\}^{\lambda'})^n$  is a family of  $|\bar{J}|$ -bit-fixing sources, obtained from each  $\vec{g} \in \mathcal{G}^n$  and  $V^*$  by setting the indices of  $\bar{J}$  to their values in  $\vec{g}(V^*)$ , while the indices of  $J$  are uniform in  $(\{0, 1\}^{\lambda'})^t$ .

Let  $(G, V_{G,z}^*)$  be obtained by replacing every  $(G, G(X^*))$  in  $(G, G(X_z^*))$  with the corresponding  $(G, V_G^*)$ . Then, the above implies that

$$\Delta[(G, G(X_z^*)), (G, V_{G,z}^*)] \leq \sqrt{2^{-(1-\delta)\lambda + \lambda' + \log n}}.$$

Therefore,

$$\begin{aligned} \Delta[(G, G(X_z)), (G, V_{G,z}^*)] &\leq \\ \Delta[(G, G(X_z)), (G, G(X_z^*))] + \Delta[(G, G(X_z^*)), (G, V_{G,z}^*)] &\leq \gamma + \sqrt{2^{-(1-\delta)\lambda + \lambda' + \log n}}, \end{aligned}$$

where we have used the fact that

$$\Delta[(G, G(X_z)), (G, G(X_z^*))] \leq \Delta[X_z, X_z^*] \leq \gamma,$$

since  $G$  is independent of  $X_z$  and  $X_z^*$ .

□

**Proof of Lemma 6.2.** Lemma 6.2 is based on Lemma D.2 and the proposition below which upper bounds the probability over  $X$  and  $F$  that  $H_\infty(X_z)$  is low.

**Proposition 1.** *Let  $s' > s$ . Then,  $\Pr_{z \leftarrow F(X)}[H_\infty(X_z) < n \cdot \lambda - s'] < 2^{s-s'}$ .*

*Proof.* First, since a randomized function is a convex combination of deterministic ones, we assume without loss of generality that  $F$  is deterministic. Fix  $z \in \{0, 1\}^s$  such that  $H_\infty(X_z) < n \cdot \lambda - s'$ . Since  $X \in (\{0, 1\}^\lambda)^n$  is uniform, then  $\Pr_X[F(X) = z] < 2^{-s'}$ .

Let  $K = \{z \in \{0, 1\}^s \mid H_\infty(X_z) < n \cdot \lambda - s'\}$ . Then,

$$\Pr_{z \leftarrow F(X)}[H_\infty(X_z) < n \cdot \lambda - s'] = \sum_{z \in K} \Pr_X[F(X) = z] < \sum_{z \in K} 2^{-s'} \leq 2^s \cdot 2^{-s'} = 2^{s-s'}.$$

□

*Proof of Lemma 6.2.* Let  $\delta > 0, \gamma > 0, s' > s$  be parameters.

Consider  $z \in \{0, 1\}^s$  with  $s_z \leq s'$  (i.e.,  $H_\infty(X_z) \geq n \cdot \lambda - s'$ ). Since  $(1 - \delta)\lambda > \lambda' + \log n + 1$ , we invoke Lemma D.2 and deduce that there is a family of convex combinations  $V_{G,z}^*$  of finitely many  $k$ -bit-fixing such that

$$\Delta[(G, G(X_z)), (G, G(V_{G,z}^*))] \leq \sqrt{2^{-(1-\delta)\lambda + \lambda' + \log n}} + \gamma, \quad (4)$$

and

$$k = \frac{s' + \log 1/\gamma}{\delta \cdot \lambda}.$$

Define the family

$$V_{G,Z} = \{V_{G,z}\}_{z \in \{0,1\}^s} = \{V_{\bar{g},z}\}_{\bar{g} \in (\mathcal{G})^n, z \in \{0,1\}^s}$$

as follows: for each  $z$  such that  $s_z \leq s'$ , set  $V_{G,z} := V_{G,z}^*$ , while for each  $z$  such that  $s_z > s'$ , set  $V_{G,z}$  to be the uniform distribution over  $(\{0, 1\}^{\lambda'})^n$ . Then,  $V_{G,Z}$  is a family of convex combinations  $V_{\bar{g},z}$  of  $k$ -bit-fixing  $(n, 2^{\lambda'})$ -sources, for  $k = \frac{s' + \log 1/\gamma}{\delta \cdot \lambda}$ . To complete the proof, it remains to bound  $\Delta[(G, Z, G(X)), (G, Z, V_{G,Z})]$ .

Denote by  $\mathcal{E}$  the event that  $s_z < s'$ . Then,

$$\begin{aligned} \Delta[(G, Z, G(X)), (G, Z, V_{G,Z})] &= \mathbb{E}_{z \leftarrow F(X)} \Delta[(G, G(X_z)), (G, V_{G,z})] \leq \\ &\Pr[\mathcal{E}] \cdot 1 + \mathbb{E}_{z \leftarrow F(X) | \neg \mathcal{E}} \Delta[(G, G(X_z)), (G, V_{G,z})] \leq \\ &2^{s-s'} + \sqrt{2^{-(1-\delta)\lambda + \lambda' + \log n}} + \gamma, \end{aligned}$$

where the final inequality is based on Proposition 1 and (4). □

## E Application to Communication Complexity

In the two-player one-way communication game, inputs  $A$  and  $B$  are given to Alice and Bob, respectively, and the goal is for Alice to send a minimal amount of information to Bob, so that Bob can compute  $f(A, B)$  for some predetermined function  $B$ . The communication cost of a protocol  $\Pi$  is the size of the largest message in bits sent from Alice across all possible inputs  $X$  and the (randomized) communication complexity is the minimum communication cost of a protocol that succeeds with probability at least  $\frac{2}{3}$ . In the distributional setting,  $X$  and  $Y$  are further drawn from some known underlying distribution.

In our setting, suppose Alice has  $n$  independent and uniform numbers  $a_1, \dots, a_n$  so that either  $a_i \in GF(p)$  for all  $i \in [n]$  or  $a_i \in GF(2^t)$  for sufficiently large  $t$  for all  $i \in [n]$  and suppose Bob has  $n$  independent and uniform numbers  $b_1, \dots, b_n$  from the same field, either  $GF(p)$  or  $GF(2^t)$ . Then for any function  $f(\langle a_1, b_1 \rangle, \dots, \langle a_n, b_n \rangle)$ , where the dot products are taken over  $GF(2)$  or  $f(a_1 \cdot b_1, \dots, a_n \cdot b_n)$ , where the products are taken over  $GF(p)$ , has the property that the randomized one-way communication complexity of computing  $f$  with probability  $\sigma$  is the same as the number of samples from  $a_1, \dots, a_n$  that Alice needs to send Bob to compute  $f$  with probability

$\sigma - \varepsilon$ . It is easy to prove sampling lower bounds for many of these problems, sum as  $\sum_i a_i \cdot b_i \pmod{p}$  or MAJ( $\langle a_1, b_1 \rangle, \dots, \langle a_n, b_n \rangle$ ), and this immediately translates into communication complexity lower bounds by the above. The main intuition for these lower bounds is that the numbers  $b_1, \dots, b_n$  can be viewed as the hash functions that Bob has.

In the context of Corollary 6.3, the input  $A$  to Alice is the input  $X = (X_1, \dots, X_n)$  to the hash functions and the input  $B$  to Bob is the set of hash functions  $G = (G_1, \dots, G_n)$ . In particular, we have  $X_i \in \{0, 1\}^\lambda$  and  $G_i \in \{0, 1\}^t$  for each  $i \in [n]$ . Then informally speaking, Corollary 6.3 shows that if a sampling protocol requires  $\Omega(\eta^2 n)$  samples to achieve probability  $\frac{1}{2} + \frac{\eta}{2}$  of success, i.e.,  $\frac{\eta}{2}$  advantage, then any algorithm that achieves probability  $\frac{1}{2} + \eta$  of success, i.e.,  $\eta$  advantage, must use  $\Omega(\eta^2 n \lambda)$  communication.

**Theorem E.1.** *Let  $X = (X_1, \dots, X_n)$  with  $X_i \in \{0, 1\}^\lambda$  for all  $i \in [n]$  be given to Alice and let  $G = (G_1, \dots, G_n)$  with  $G_i \in \{0, 1\}^t$  for all  $i \in [n]$  be given to Bob. Suppose that for a function  $f : (\{0, 1\}^\lambda)^n \times (\{0, 1\}^\lambda)^n \rightarrow \{0, 1\}$ , any protocol  $\Pi$  that computes  $f(X, G)$  with probability at least  $\frac{1}{2} + \frac{\eta}{2}$  requires  $k$  samples  $(X_i, G_i)$  for some quantity  $k$  that depends on  $\eta$ . Then any one-way protocol that computes  $f(X, G)$  with probability at least  $\frac{1}{2} + \eta$  must use  $\Omega(s)$  communication, where  $s = \frac{k\lambda}{2} - 2 \log_2 \frac{6}{\eta}$ .*

*Proof.* We define  $h$  to be the function so that

$$f(X, G) = h(G_1(X_1), G_2(X_2), \dots, G_n(X_n)).$$

Setting  $t = \lambda' = 1$ ,  $\delta = 1/2$ ,  $n < 2^{\lambda/3}$ ,  $\gamma = \eta/6$ ,  $s = \Theta(\eta^2 n \lambda)$ , and  $s' = s + \log_2(6/\eta)$ , in the formulation of Corollary 6.3, then for sufficiently large  $\lambda$  so that  $2^{-c\lambda} < \frac{\eta}{6}$  for a fixed constant  $c > 0$ , we have

$$\Delta[P_{(G, Z, G(X))}, P_{(G, Z, V)}] \leq \sqrt{2^{-(1-\delta)\lambda + \lambda' + \log n} + \gamma} + 2^{s-s'} \leq \sqrt{2^{-\lambda/2 + 1 + \lambda/3} + \eta/6} + \eta/6 \leq \eta/2,$$

where  $V$  is a  $((G, Z), k)$ -convex-bit-fixing  $(n, 2^{\lambda'})$ -source, and  $k = \frac{2(s + 2 \log_2(6/\eta))}{\lambda}$ .

Suppose there exists a protocol  $\Pi$  that computes  $f(G, X)$  with probability more than  $\frac{1}{2} + \eta$  using  $o(s)$  space. By Corollary 6.3, the success probability is the convex combination of the success probabilities for  $((G, Z), k')$ -convex-bit-fixing  $(n, 2^{\lambda'})$ -sources with  $k' < k$ . Thus there exists a protocol  $\Pi'$  that produces a message  $z$  by Alice such that Bob's probability of success on  $g$  is at least  $\frac{1}{2} + \eta$ . By the law of total probability, the success probability of Bob is the weighted success probability of Bob conditioned on each  $k'$ -bit-fixing source in the convex combination and so there exists a  $k'$ -bit-fixing source with success probability at least  $\frac{1}{2} + \frac{\eta}{2}$ . Hence, Bob's posterior distribution on  $g_1, \dots, g_n$  is jointly uniform on all coordinates outside of a fixed set  $S$  of at most  $k'$  indices, up to a negligible statistical distance. Since each  $g_i$  is universal, Bob's posterior distribution is the same as if Alice had sampled the coordinates in  $S$  and sent the coordinates to Bob. However, since  $|S| \leq k' < k$ , this contradicts the assumption that any protocol  $\Pi$  that computes  $f(X, G)$  with probability at least  $\frac{1}{2} + \frac{\eta}{2}$  requires  $k$  samples  $(X_i, G_i)$  for some quantity  $k$  that depends on  $\eta$ .  $\square$

We can apply Corollary 6.3 in this setting to any relation  $f(X, G)$  which has the form

$$h(G_1(X_1), G_2(X_2), \dots, G_n(X_n))$$

for some relation  $h$ , provided the  $G_i$  are independent and drawn from a universal hash function family with range  $\{0, 1\}^{\lambda'}$ . Indeed, by viewing  $Z$  in that theorem as the message sent from Alice to Bob, Corollary 6.3 shows that the distributions  $P_{(G,Z,G(X))}$  and  $P_{(G,Z,V)}$  have small variation distance, where  $V$  is a  $((G, Z), k)$ -convex-bit-fixing source. Consequently, Bob, who is given  $G = g$  and  $Z = z$  for some particular realizations  $g$  and  $z$ , cannot distinguish  $g(X)$  from  $V$ , where  $V$  is a  $((g, z), k)$ -convex-bit-fixing source, which means that conditioned on  $G = g$  and  $Z = z$ ,  $V$  is a convex combination of distributions of  $k$ -bit-fixing sources, i.e., distributions on  $(\{0, 1\}^{\lambda'})^n$  that are fixed on at most  $k$  coordinates and uniform on the rest.

Since the average success probability over  $g$  is at least  $1/2 + \eta$ , this implies there exists a  $g$  for which the success probability is at least  $1/2 + \eta$ , and we can use this particular  $g = (g_1, \dots, g_n)$  to derive our lower bound. Similarly, we can fix a message  $z$  of Alice for which the success probability of Bob on  $g$  is at least  $1/2 + \eta$ . Now, by the law of total probability, the success probability of Bob is the weighted success probability of Bob conditioned on each  $k$ -bit-fixing source in the convex combination defining  $V$ , and so if Bob's success probability is at least  $1/2 + \eta$  on  $(g, z)$ , there exists a  $k$ -bit-fixing source with success probability at least  $1/2 + \eta$ . This in turn means that Bob's posterior distribution on the  $g_i(X_i)$  is jointly uniform on all coordinates outside of a fixed set  $S$  of indices, up to a negligible statistical distance.

The above means, in particular, that outside of  $S$ ,  $g_i$  cannot be the zero function, as otherwise  $g_i(X_i)$  could not be uniform. But for many universal families, if  $g_i \neq 0$  is fixed, then  $g_i(X_i)$  being uniform means  $X_i$  is uniform. This holds, in particular if  $g_i(X_i) = \langle g_i, X_i \rangle \bmod 2$ . In this case, Bob's posterior distribution is the same as that obtained from Alice just sampling the coordinates  $i \in S$  and sending the  $X_i$  values to Bob, and so if we prove a sampling lower bound for  $1/2 + \eta/2$  success probability, we obtain an arbitrary one-round communication lower bound for  $1/2 + \eta$  success probability, assuming the parameters are set so that the statistical distance is at most say,  $\eta/2$ .

As an example, suppose  $h$  is the majority function and for each  $i$ ,  $g_i(X_i) = \langle g_i, X_i \rangle \bmod 2$  for bitstrings  $g_i$  and  $X_i$ . It is well-known that the set of all such  $g_i$  form a universal family. To obtain  $1/2 + \eta$  success probability under the uniform distribution on  $(G, X)$ , assuming  $\lambda = \Omega(\log 1/\eta)$  so that  $g_i = 0^\lambda$  is unlikely (and for any fixed  $g_i \neq 0^\lambda$ , when  $X_i$  is uniform,  $\langle g_i, X_i \rangle$  is uniform), it is standard that it is necessary and sufficient for a sampling protocol to sample  $\Theta(\eta^2 n)$  coordinates  $i$  and send them to Bob, and then Bob outputs the empirical majority. In Corollary 6.3 we can set  $\lambda' = 1$ ,  $\delta = 1/2$ ,  $n < 2^{\lambda/3}$ ,  $\gamma = \eta/6$ ,  $s = \Theta(\eta^2 n \lambda)$ , and  $s' = s + \log_2(6/\eta)$  to conclude that for  $\eta/6 > 2^{-c\lambda}$  for a small constant  $c > 0$ ,

$$\Delta[P_{(G,Z,G(X))}, P_{(G,Z,V)}] \leq \sqrt{2^{-(1-\delta)\lambda + \lambda' + \log n} + \gamma} + 2^{s-s'} \leq \sqrt{2^{-\lambda/2 + 1 + \lambda/3} + \eta/6} + \eta/6 \leq \eta/2,$$

where  $V$  is a  $((G, Z), k)$ -convex-bit-fixing  $(n, 2^{\lambda'})$ -source, and  $k = \frac{2(s+2\log_2(6/\eta))}{\lambda}$ . One needs  $k = \Omega(\eta^2 n \lambda)$  by the sampling lower bound to achieve  $1/2 + \eta/2$  success probability, which implies  $s = \Omega(\eta^2 n \lambda - 2\log_2(6/\eta))$  as a 1-way arbitrary communication lower bound, since the success probability of the 1-way protocol is at most  $1/2 + \eta/2 + \eta/2 = 1/2 + \eta$ .

## References

- [1] Alabi, D., Ben-Eliezer, O., Chaturvedi, A.: Bounded space differentially private quantiles. CoRR abs/2201.03380 (2022) [2](#)
- [2] Alon, N., Livni, R., Malliaris, M., Moran, S.: Private PAC learning implies finite littlestone dimension. In: STOC. pp. 852–860. ACM (2019) [8](#)
- [3] Bassily, R., Freund, Y.: Typical stability. arXiv preprint arXiv:1604.03336 (2016) [5](#)
- [4] Bassily, R., Nissim, K., Smith, A.D., Steinke, T., Stemmer, U., Ullman, J.R.: Algorithmic stability for adaptive data analysis. SIAM J. Comput. 50(3) (2021) [5](#)
- [5] Bassily, R., Nissim, K., Stemmer, U., Thakurta, A.: Practical locally private heavy hitters. J. Mach. Learn. Res. 21, 16:1–16:42 (2020) [2](#)
- [6] Beame, P., Gharan, S.O., Yang, X.: Time-space tradeoffs for learning finite functions from random evaluations, with applications to polynomials. In: Conference On Learning Theory, COLT. pp. 843–856 (2018) [8](#)
- [7] Beimel, A., Nissim, K., Omri, E.: Distributed private data analysis: Simultaneously solving how and what. In: CRYPTO. Lecture Notes in Computer Science, vol. 5157, pp. 451–468. Springer (2008) [12](#)
- [8] Bellare, M., Rogaway, P.: Random oracles are practical: A paradigm for designing efficient protocols. In: Denning, D.E., Pyle, R., Ganesan, R., Sandhu, R.S., Ashby, V. (eds.) CCS 1993. pp. 62–73. ACM (1993) [1](#)
- [9] Blocki, J., Blum, A., Datta, A., Sheffet, O.: The johnson-lindenstrauss transform itself preserves differential privacy. In: FOCS. pp. 410–419. IEEE Computer Society (2012) [2](#)
- [10] Bolot, J., Fawaz, N., Muthukrishnan, S., Nikolov, A., Taft, N.: Private decayed predicate sums on streams. In: ICDT. pp. 284–295. ACM (2013) [2](#)
- [11] Boneh, D., Shaw, J.: Collusion-secure fingerprinting for digital data. IEEE Transactions on Information Theory 44(5), 1897–1905 (1998) [12](#), [13](#)
- [12] Brown, G., Bun, M., Feldman, V., Smith, A., Talwar, K.: When is memorization of irrelevant training data necessary for high-accuracy learning? In: Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing. pp. 123–132 (2021) [8](#)
- [13] Bu, Z., Gopi, S., Kulkarni, J., Lee, Y.T., Shen, H., Tantipongpipat, U.: Fast and memory efficient differentially private-sgd via jl projections. Advances in Neural Information Processing Systems 34, 19680–19691 (2021) [2](#)
- [14] Bun, M., Nelson, J., Stemmer, U.: Heavy hitters and the structure of local privacy. ACM Trans. Algorithms 15(4), 51:1–51:40 (2019) [2](#)
- [15] Bun, M., Nissim, K., Stemmer, U., Vadhan, S.P.: Differentially private release and learning of threshold functions. In: FOCS. pp. 634–649. IEEE Computer Society (2015) [8](#)
- [16] Bun, M., Ullman, J.R., Vadhan, S.P.: Fingerprinting codes and the price of approximate differential privacy. In: STOC. pp. 1–10. ACM (2014) [7](#), [8](#), [12](#), [14](#)

- [17] Canetti, R., Dodis, Y., Halevi, S., Kushilevitz, E., Sahai, A.: Exposure-resilient functions and all-or-nothing transforms. In: Preneel, B. (ed.) EUROCRYPT 2000. Lecture Notes in Computer Science, vol. 1807, pp. 453–469. Springer (2000) [6](#)
- [18] Chan, T.H.H., Li, M., Shi, E., Xu, W.: Differentially private continual monitoring of heavy hitters from distributed streams. In: International Symposium on Privacy Enhancing Technologies Symposium. pp. 140–159. Springer (2012) [2](#)
- [19] Chattopadhyay, A., Filmus, Y., Koroth, S., Meir, O., Pitassi, T.: Query-to-communication lifting for bpp using inner product. In: 46th International Colloquium on Automata, Languages, and Programming (ICALP 2019). Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik (2019) [8](#)
- [20] Chattopadhyay, A., Filmus, Y., Koroth, S., Meir, O., Pitassi, T.: Query-to-communication lifting using low-discrepancy gadgets. SIAM Journal on Computing 50(1), 171–210 (2021) [8](#)
- [21] Chattopadhyay, A., Koucký, M., Loff, B., Mukhopadhyay, S.: Simulation beats richness: New data-structure lower bounds. In: Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing. pp. 1013–1020 (2018) [8](#)
- [22] Chattopadhyay, A., Koucký, M., Loff, B., Mukhopadhyay, S.: Simulation theorems via pseudo-random properties. computational complexity 28(4), 617–659 (2019) [8](#)
- [23] Chien, S., Ligett, K., McGregor, A.: Space-efficient estimation of robust statistics and distribution testing. In: Innovations in Computer Science - ICS, Proceedings. pp. 251–265 (2010) [8](#)
- [24] Coretti, S., Dodis, Y., Guo, S.: Non-uniform bounds in the random-permutation, ideal-cipher, and generic-group models. In: Shacham, H., Boldyreva, A. (eds.) CRYPTO 2018. Lecture Notes in Computer Science, vol. 10991, pp. 693–721. Springer (2018) [2](#), [7](#)
- [25] Coretti, S., Dodis, Y., Guo, S., Steinberger, J.P.: Random Oracles and Non-uniformity. In: Nielsen, J.B., Rijmen, V. (eds.) EUROCRYPT 2018. Lecture Notes in Computer Science, vol. 10820, pp. 227–258. Springer (2018) [2](#), [7](#), [20](#), [24](#), [34](#)
- [26] Cummings, R., Ligett, K., Nissim, K., Roth, A., Wu, Z.S.: Adaptive learning with robust generalization guarantees. In: Conference on Learning Theory. pp. 772–814. PMLR (2016) [5](#)
- [27] Dagan, Y., Shamir, O.: Detecting correlations with little memory and communication. In: Conference On Learning Theory, COLT. pp. 1145–1198 (2018) [8](#)
- [28] Dai, W., Tessaro, S., Zhang, X.: Super-linear time-memory trade-offs for symmetric encryption. In: TCC 2020. pp. 335–365 (2020) [24](#)
- [29] De Rezende, S.F., Nordström, J., Vinyals, M.: How limited interaction hinders real communication (and what it means for proof and circuit complexity). In: 2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS). pp. 295–304. IEEE (2016) [8](#)
- [30] Diakonikolas, I., Gouleakis, T., Kane, D.M., Rao, S.: Communication and memory efficient testing of discrete distributions. In: Beygelzimer, A., Hsu, D. (eds.) Conference on Learning Theory, COLT. pp. 1070–1106 (2019) [8](#)



- [31] Diakonikolas, I., Kane, D.M., Pensia, A., Pittas, T.: Streaming algorithms for high-dimensional robust statistics. In: International Conference on Machine Learning, ICML. pp. 5061–5117 (2022) [8](#)
- [32] Dodis, Y., Farshim, P., Mazaheri, S., Tessaro, S.: Towards defeating backdoored random oracles: Indifferentiability with bounded adaptivity. In: Pass, R., Pietrzak, K. (eds.) TCC 2020. Lecture Notes in Computer Science, vol. 12552, pp. 241–273. Springer (2020) [7](#), [24](#)
- [33] Dwork, C., Feldman, V., Hardt, M., Pitassi, T., Reingold, O., Roth, A.: Generalization in adaptive data analysis and holdout reuse. In: NIPS. pp. 2350–2358 (2015) [5](#)
- [34] Dwork, C., Feldman, V., Hardt, M., Pitassi, T., Reingold, O., Roth, A.L.: Preserving statistical validity in adaptive data analysis. In: Proceedings of the forty-seventh annual ACM symposium on Theory of computing. pp. 117–126 (2015) [4](#), [5](#), [17](#)
- [35] Dwork, C., McSherry, F., Nissim, K., Smith, A.: Calibrating noise to sensitivity in private data analysis. *Journal of Privacy and Confidentiality* 7(3), 17–51 (2016) [2](#)
- [36] Dwork, C., Naor, M., Pitassi, T., Rothblum, G.N., Yekhanin, S.: Pan-private streaming algorithms. In: ICS. pp. 66–80. Tsinghua University Press (2010) [2](#)
- [37] Dwork, C., Naor, M., Reingold, O., Rothblum, G.N., Vadhan, S.: On the complexity of differentially private data release: efficient algorithms and hardness results. In: Proceedings of the forty-first annual ACM symposium on Theory of computing. pp. 381–390 (2009) [8](#)
- [38] Dwork, C., Rothblum, G.N., Vadhan, S.P.: Boosting and differential privacy. In: FOCS. pp. 51–60. IEEE Computer Society (2010) [12](#)
- [39] Dziembowski, S.: Intrusion-resilience via the bounded-storage model. In: Halevi, S., Rabin, T. (eds.) TCC 2006. vol. 3876, pp. 207–224. Springer (2006) [6](#)
- [40] Dziembowski, S., Pietrzak, K.: Leakage-resilient cryptography. In: FOCS 2008. pp. 293–302. IEEE Computer Society (2008) [6](#)
- [41] Feldman, V.: Does learning require memorization? a short tale about a long tail. In: Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing. pp. 954–959 (2020) [8](#)
- [42] Feldman, V., Steinke, T.: Generalization for adaptively-chosen estimators via stable median. In: Conference on Learning Theory. pp. 728–757. PMLR (2017) [5](#)
- [43] Feldman, V., Steinke, T.: Calibrating noise to variance in adaptive data analysis. In: Conference On Learning Theory. pp. 535–544. PMLR (2018) [5](#)
- [44] Fish, B., Reyzin, L., Rubinfeld, B.I.P.: Sublinear-time adaptive data analysis. In: ISAIM (2018) [5](#)
- [45] Fish, B., Reyzin, L., Rubinfeld, B.I.P.: Sampling without compromising accuracy in adaptive data analysis. In: ALT. Proceedings of Machine Learning Research, vol. 117, pp. 297–318. PMLR (2020) [5](#)

- [46] Garg, A., Göös, M., Kamath, P., Sokolov, D.: Monotone circuit lower bounds from resolution. In: Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing. pp. 902–911 (2018) [8](#)
- [47] Garg, S., Raz, R., Tal, A.: Extractor-based time-space lower bounds for learning. In: Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC. pp. 990–1002 (2018) [8](#)
- [48] Goldreich, O.: The Foundations of Cryptography - Volume 1: Basic Techniques. Cambridge University Press (2001) [26](#)
- [49] Göös, M., Lovett, S., Meka, R., Watson, T., Zuckerman, D.: Rectangles Are Nonnegative Juntas. In: Servedio, R.A., Rubinfeld, R. (eds.) STOC 2015. pp. 257–266. ACM (2015) [7](#), [8](#), [24](#), [34](#)
- [50] Göös, M., Pitassi, T.: Communication lower bounds via critical block sensitivity. In: Proceedings of the forty-sixth annual ACM symposium on Theory of computing. pp. 847–856 (2014) [8](#)
- [51] Göös, M., Pitassi, T., Watson, T.: Deterministic communication vs. partition number. In: 2015 IEEE 56th Annual Symposium on Foundations of Computer Science. pp. 1077–1088. IEEE (2015) [8](#)
- [52] Göös, M., Pitassi, T., Watson, T.: Query-to-communication lifting for bpp. SIAM Journal on Computing 49(4), FOCS17–441 (2020) [8](#)
- [53] Hardt, M., Ullman, J.R.: Preventing false discovery in interactive data analysis is hard. In: FOCS. pp. 454–463. IEEE Computer Society (2014) [5](#), [6](#), [7](#), [8](#), [17](#), [18](#), [27](#), [28](#)
- [54] Håstad, J., Impagliazzo, R., Levin, L.A., Luby, M.: A Pseudorandom Generator from any One-way Function. SIAM J. Comput. 28(4), 1364–1396 (1999) [24](#), [33](#)
- [55] Hatami, H., Hosseini, K., Lovett, S.: Structure of protocols for xor functions. SIAM Journal on Computing 47(1), 208–217 (2018) [8](#)
- [56] Hazay, C., López-Alt, A., Wee, H., Wichs, D.: Leakage-resilient cryptography from minimal assumptions. J. Cryptol. 29(3), 514–551 (2016) [6](#)
- [57] Jung, C., Ligett, K., Neel, S., Roth, A., Sharifi-Malvajerdi, S., Shenfeld, M.: A new analysis of differential privacy’s generalization guarantees (invited paper). In: STOC. p. 9. ACM (2021) [5](#)
- [58] Kalai, Y.T., Reyzin, L.: A survey of leakage-resilient cryptography. In: Goldreich, O. (ed.) Providing Sound Foundations for Cryptography: On the Work of Shafi Goldwasser and Silvio Micali, pp. 727–794. ACM (2019) [6](#)
- [59] Kaplan, H., Stemmer, U.: A note on sanitizing streams with differential privacy. CoRR abs/2111.13762 (2021) [2](#)
- [60] Kol, G., Raz, R., Tal, A.: Time-space hardness of learning sparse parities. In: Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC. pp. 1067–1080 (2017) [8](#)

- [61] Kontorovich, A., Sadigurschi, M., Stemmer, U.: Adaptive data analysis with correlated observations. In: ICML. Proceedings of Machine Learning Research, vol. 162, pp. 11483–11498. PMLR (2022) [5](#)
- [62] Kothari, P.K., Meka, R., Raghavendra, P.: Approximating rectangles by juntas and weakly-exponential lower bounds for LP relaxations of CSPs. In: Hatami, H., McKenzie, P., King, V. (eds.) STOC 2017. pp. 590–603. ACM (2017) [7](#), [24](#), [34](#)
- [63] Loff, B., Mukhopadhyay, S.: Lifting theorems for equality. In: 36th International Symposium on Theoretical Aspects of Computer Science (STACS 2019). Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik (2019) [8](#)
- [64] Mir, D., Muthukrishnan, S., Nikolov, A., Wright, R.N.: Pan-private algorithms via statistics on sketches. In: Proceedings of the thirtieth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems. pp. 37–48 (2011) [2](#)
- [65] Mironov, I., Pandey, O., Reingold, O., Vadhan, S.P.: Computational differential privacy. In: CRYPTO. Lecture Notes in Computer Science, vol. 5677, pp. 126–142. Springer (2009) [12](#)
- [66] Moshkovitz, D., Moshkovitz, M.: Mixing implies lower bounds for space bounded learning. In: Proceedings of the 30th Conference on Learning Theory, COLT 2017. pp. 1516–1566 (2017) [8](#)
- [67] Moshkovitz, D., Moshkovitz, M.: Entropy samplers and strong generic lower bounds for space bounded learning. In: 9th Innovations in Theoretical Computer Science Conference, ITCS 2018. pp. 28:1–28:20 (2018) [8](#)
- [68] Nissim, K., Smith, A., Steinke, T., Stemmer, U., Ullman, J.: The limits of post-selection generalization. Advances in Neural Information Processing Systems 31 (2018) [5](#)
- [69] Pagh, R., Stausholm, N.M.: Efficient differentially private  $f_0$  linear sketching. In: ICDT. LIPIcs, vol. 186, pp. 18:1–18:19. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2021) [2](#)
- [70] Pagh, R., Thorup, M.: Improved utility analysis of private countsketch. CoRR abs/2205.08397 (2022) [2](#)
- [71] Pietrzak, K.: A leakage-resilient mode of operation. In: Joux, A. (ed.) EUROCRYPT 2009. Lecture Notes in Computer Science, vol. 5479, pp. 462–482. Springer (2009) [6](#)
- [72] Raz, R.: A time-space lower bound for a large class of learning problems. In: 58th IEEE Annual Symposium on Foundations of Computer Science, FOCS. pp. 732–742 (2017) [8](#)
- [73] Raz, R.: Fast learning requires good memory: A time-space lower bound for parity learning. J. ACM 66(1), 3:1–3:18 (2019) [8](#)
- [74] Raz, R., McKenzie, P.: Separation of the monotone NC hierarchy. Comb. 19(3), 403–435 (1999) [1](#), [8](#)
- [75] Rogers, R., Roth, A., Smith, A., Thakkar, O.: Max-information, differential privacy, and post-selection hypothesis testing. In: 2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS). pp. 487–494. IEEE (2016) [5](#)
- [76] Shenfeld, M., Ligett, K.: A necessary and sufficient stability notion for adaptive generalization. In: NeurIPS. pp. 11481–11490 (2019) [5](#)

- [77] Shenfeld, M., Ligett, K.: Generalization in the face of adaptivity: A bayesian perspective. CoRR abs/2106.10761 (2021) [5](#)
- [78] Smith, A.D., Song, S., Thakurta, A.: The flajolet-martin sketch itself preserves differential privacy: Private counting with minimal space. In: NeurIPS (2020) [2](#)
- [79] Standaert, F., Pereira, O., Yu, Y., Quisquater, J., Yung, M., Oswald, E.: Leakage resilient cryptography in practice. In: Sadeghi, A., Naccache, D. (eds.) Towards Hardware-Intrinsic Security - Foundations and Practice, pp. 99–134. Information Security and Cryptography, Springer (2010) [6](#)
- [80] Steinhardt, J., Valiant, G., Wager, S.: Memory, communication, and statistical queries. In: Proceedings of the 29th Conference on Learning Theory, COLT. pp. 1490–1516 (2016) [8](#)
- [81] Steinke, T., Ullman, J.R.: Interactive fingerprinting codes and the hardness of preventing false discovery. In: COLT. JMLR Workshop and Conference Proceedings, vol. 40, pp. 1588–1628. JMLR.org (2015) [5](#), [6](#), [7](#), [8](#), [13](#), [17](#), [18](#), [27](#), [28](#)
- [82] Tardos, G.: Optimal probabilistic fingerprint codes. Journal of the ACM (JACM) 55(2), 1–24 (2008) [13](#)
- [83] Ullman, J.: 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. pp. 361–370 (2013) [8](#)
- [84] Unruh, D.: Random oracles and auxiliary input. In: Menezes, A. (ed.) CRYPTO 2007. Lecture Notes in Computer Science, vol. 4622, pp. 205–223. Springer (2007) [2](#), [7](#)
- [85] Upadhyay, J.: Sublinear space private algorithms under the sliding window model. In: ICML. Proceedings of Machine Learning Research, vol. 97, pp. 6363–6372. PMLR (2019) [2](#)
- [86] Upadhyay, J., Upadhyay, S.: A framework for private matrix analysis in sliding window model. In: ICML. Proceedings of Machine Learning Research, vol. 139, pp. 10465–10475. PMLR (2021) [2](#)
- [87] Wang, L., Pinelis, I., Song, D.: Differentially private fractional frequency moments estimation with polylogarithmic space. arXiv preprint arXiv:2105.12363 (2021) [2](#)
- [88] Woodruff, D.P.: Optimal space lower bounds for all frequency moments. In: Munro, J.I. (ed.) Proceedings of the Fifteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2004, New Orleans, Louisiana, USA, January 11-14, 2004. pp. 167–175. SIAM (2004) [3](#)
- [89] Wu, X., Yao, P., Yuen, H.S.: Raz-mckenzie simulation with the inner product gadget. In: Electron. Colloquium Comput. Complex. vol. 24, p. 10 (2017) [8](#)
- [90] Zhao, F., Qiao, D., Redberg, R., Agrawal, D., Abbadi, A.E., Wang, Y.X.: Differentially private linear sketches: Efficient implementations and applications. arXiv preprint arXiv:2205.09873 (2022) [2](#)