

---

# Adaptive Data Analysis with Correlated Observations

---

Aryeh Kontorovich<sup>1</sup> Menachem Sadigurschi<sup>1</sup> Uri Stemmer<sup>2,3</sup>

## Abstract

The vast majority of the work on adaptive data analysis focuses on the case where the samples in the dataset are independent. Several approaches and tools have been successfully applied in this context, such as *differential privacy*, *max-information*, *compression arguments*, and more. The situation is far less well-understood without the independence assumption.

We embark on a systematic study of the possibilities of adaptive data analysis with correlated observations. First, we show that, in some cases, differential privacy guarantees generalization even when there are dependencies within the sample, which we quantify using a notion we call *Gibbs-dependence*. We complement this result with a tight negative example. Second, we show that the connection between transcript-compression and adaptive data analysis can be extended to the non-iid setting.

## 1. Introduction

Statistical validity is a well known crucial aspect of modern science. In the past several years, the natural science and social science communities have come to realize that such validity was not in fact preserved in numerous peer-reviewed and widely cited studies, leading to many false discoveries. Known as the *replication crisis*, this phenomenon threatens to undermine the very basis for the public’s trust in science.

One of the main explanations for the prevalence of false discovery arises from the inherent *adaptivity* in the process of data analysis. To illustrate this issue, 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.

Before presenting our new results, we make the setting explicit. We give here the formulation presented by Dwork et al. (2015c). We consider a two-player game between a mechanism  $\mathbb{M}$  and an adversary  $\mathbb{A}$ , defined as follows (see Section 2 for precise definitions).

1. The adversary  $\mathbb{A}$  fixes a measure  $\mu$  over  $\mathcal{X}^n$  (satisfying some conditions).
2. The mechanism  $\mathbb{M}$  obtains a sample  $S \sim \mu$  containing  $n$  (possibly correlated) observations.
3. For  $k$  rounds  $j = 1, 2, \dots, k$ :
  - The adversary chooses a *query*  $h_j : \mathcal{X} \rightarrow \{0, 1\}$ , possibly as a function of all previous answers given by the mechanism.
  - The mechanism obtains  $h_j$  and responds with an answer  $z_j \in \mathbb{R}$ , which is given to  $\mathbb{A}$ .

We say that  $\mathbb{M}$  is  $(\alpha, \beta)$ -*empirically-accurate* if with probability at least  $1 - \beta$  for every  $j$  it holds that  $|z_j - h_j(S)| \leq \alpha$ , where  $h_j(S) = \frac{1}{n} \sum_{x \in S} h_j(x)$  is the empirical average of  $h_j$  on the sample  $S$ . We say that  $\mathbb{M}$  is  $(\alpha, \beta)$ -*statistically-accurate* if with probability at least  $1 - \beta$  for every  $j$  it holds that  $|z_j - h_j(\mu)| \leq \alpha$ , where  $h_j(\mu) = \mathbb{E}_{T \sim \mu} [h_j(T)] = \mathbb{E}_{T \sim \mu} [\frac{1}{n} \sum_{x \in T} h_j(x)]$  is the “true” value of the query  $h_j$  on the underlying distribution  $\mu$ . Our goal is to design mechanisms  $\mathbb{M}$  providing statistical-accuracy.

Starting from (Dwork et al., 2015b;a), it has been demonstrated that various notions of *algorithmic stability*, and in particular *differential privacy* (DP) (Dwork et al., 2006), allow for methods which maintain statistical validity under the adaptive setting. The vast majority of the works in this vein, however, strongly rely on the assumption that the data

---

<sup>\*</sup>Equal contribution <sup>1</sup>Department of Computer Science, Ben-Gurion University <sup>2</sup>Blavatnik School of Computer Science, Tel Aviv University <sup>3</sup>Google Research. Correspondence to: Menachem Sadigurschi <sadigurs@post.bgu.ac.il>.

is sampled in an i.i.d. fashion. This scenario excludes some natural and essential problems in learning theory such as Markov chains, active learning, and autoregressive models (Kontorovich and Ramanan, 2008; Kontorovich and Weiss, 2014; Kontorovich and Raginsky, 2017; Settles, 2009; Hanneke et al., 2014; Sacerdote, 2001).

A notable exception is a stability notion introduced by Bassily and Freund (2016), called *typical-stability*. This beautiful and natural notion has the advantage that, under some conditions on the underlying distribution, it can guarantee statistical validity even for non-i.i.d. settings. However, one downside of the results of Bassily and Freund (2016) is that they do not recover the i.i.d. generalization bounds in the limiting regime where the dependencies decay to zero. In particular, in the i.i.d. setting, it is possible to efficiently answer  $\tilde{O}(n^2)$  adaptive queries given a sample of size  $n$ . In contrast, the results of Bassily and Freund (2016) only allow to answer  $\tilde{O}(n)$  adaptive queries, *even if the dependencies in the data decay to zero*. Bridging this gap is one of the main motivations for our work.

## 1.1. Our Contributions

We reestablish the baseline for adaptive data analysis with correlated observations. Our first contribution is to extend existing generalization results for differential privacy from the i.i.d. setting to the correlated setting. To that end, we introduce a notion we call *Gibbs dependence* to quantify the dependencies between the covariates of a given joint distribution. We complement this result with a tight negative example. Our second contribution is to extend the connection between transcript-compression and adaptive data analysis also to the non-iid setting. Finally, we demonstrate an application of our results for when the underlying measure can be described as a Markov chain.

### 1.1.1. GIBBS DEPENDENCE

We extend the connection between differential privacy and generalization to the case where the observations are correlated. We quantify the correlations in the data using a new notion, called *Gibbs dependence*, which is closely related to the classical *Dobrushin* interdependence coefficient (Kontorovich and Raginsky, 2017; Levin and Peres, 2017). Intuitively, a measure which has  $\psi$ -Gibbs dependency is such that knowledge about almost the entire sample does not provide too much information about the remaining portion. Formally,

**Definition 1.1.** For a probability measure  $\mu$  over a product space  $\mathcal{X}^n$ , define

$$\psi(\mu) = \sup_{x \in \mathcal{X}^n} \mathbb{E}_{i \sim [n]} \left\| \mu_i(\cdot) - \mu_i(\cdot \mid x^{-i}) \right\|_{\text{TV}},$$

where  $\mu_i(\cdot)$  is the  $i^{\text{th}}$  marginal measure and  $\mu_i(\cdot \mid x^{-i})$

is the  $i^{\text{th}}$  marginal measure conditioned on all the coordinates other than  $i$  (given some  $n$ -tuple  $x$ ).

Given  $\psi$  we say the  $\mu$  has  $\psi$ -Gibbs dependence if  $\psi(\mu) \leq \psi$ .

**Example 1.2** (Product Measures). A probability measure  $\mu$  has  $\psi(\mu) = 0$  if and only if it is a product measure. This is since, for a product measure, for every  $i \in [n]$  we get that  $\mu_i(\cdot) = \mu_i(\cdot \mid x^{-i})$  and hence  $\sup_x \left\| \mu_i(\cdot) - \mu_i(\cdot \mid x^{-i}) \right\|_{\text{TV}} = 0$ , which means  $\psi(\mu) = 0$ . For the other direction, when  $\psi(\mu) = 0$ , as  $\left\| \mu_i(\cdot) - \mu_i(\cdot \mid x^{-i}) \right\|_{\text{TV}}$  is non-negative, we get that  $\left\| \mu_i(\cdot) - \mu_i(\cdot \mid x^{-i}) \right\|_{\text{TV}} = 0$ , meaning that  $\mu_i(\cdot) \equiv \mu_i(\cdot \mid x^{-i})$  for every possible  $x$ . This implies that  $\mu$  is a product measure (see Appendix C.1 for more details).

**Example 1.3** (Markov Random Fields). Let  $G = ([n], E)$  be a graph, and recall that  $\mu$  is a Markov Random Field w.r.t.  $G$  if for every  $i \in [n]$  we have  $\mu(x_i \mid x^{-i}) = \mu(x_i \mid x_{\Gamma(i)})$ , where  $\Gamma(i)$  is the set of neighbors of  $i$  on  $G$ . This enables modeling undirected dependencies. Such dependencies are common in computational biology (Banf and Rhee, 2017; Wei and Li, 2007), computer vision (Li, 2009; Blake et al., 2011) and physics (Rue and Held, 2005). Markov Random Fields are closely related to our notion of Gibbs dependence. We elaborate on this connection in Section B, and prove bounds on the Gibbs-dependency of undirected Markov chains, which are an important subclass of Markov random fields.

A naive way for leveraging our notion of Gibbs dependence would be to “union bound” the correlations across the  $n$  different coordinates. Specifically, one could show that if  $\mu$  has Gibbs-dependence  $\psi$  then  $\|\mu - \mu^*\|_{\text{TV}} \leq n\psi$ , where  $\mu^*$  is the product distribution in which every coordinate is sampled independently from the corresponding marginal distribution in  $\mu$ . Thus, if  $\psi \ll \frac{1}{n}$  then one could argue about generalization w.r.t.  $\mu$  by applying existing generalization bounds w.r.t.  $\mu^*$  in the independent case (since in this regime we have  $\|\mu - \mu^*\|_{\text{TV}} \ll 1$ ). This argument, however, only works when the dependencies in  $\mu$  are *very weak* (i.e., when  $\psi \ll \frac{1}{n}$ ). Our main result is to show that differential privacy still provides generalization even if  $\psi$  is much larger, e.g., a constant independent of the sample size  $n$ . Specifically,

**Theorem 1.4.** Let  $\mathbb{M}$  be an  $(\varepsilon, \delta)$ -differentially-private mechanism which is  $(\alpha, \beta)$ -empirically-accurate for  $k$  rounds given  $n$  samples. If  $n \geq \frac{\log(2k\varepsilon/\delta)}{\varepsilon^2}$ , then  $\mathbb{M}$  is also  $(\alpha + 10\varepsilon + 2\psi, \beta + \frac{\delta}{\varepsilon})$ -statistically-accurate.

**Remark 1.5.** For the case when  $\psi$  is zero, and hence  $\mu$  is a product measure (see Example 1.2), Theorem 1.4 recovers the results achieved by differential privacy for i.i.d. samples (Dwork et al., 2015c; Bassily et al., 2016). Thus, Theorem 1.4 generalizes the connection between differential privacy and generalization to the correlated setting.

Intuitively, the above theorem states that if the underlying distribution has Gibbs-dependence  $\psi$  then the additional generalization error incurred by DP algorithms (compared to the iid setting) is at most  $O(\psi)$ . We complement this result with a tight negative example showing that there exist a distribution  $\mu$  with Gibbs-dependence  $\psi$  and a DP algorithm  $\mathcal{A}$  that obtains generalization error  $\Omega(\psi)$ . This means that, in terms of the Gibbs-dependence, our result is tight.

By applying Theorem 1.4 with a known DP mechanism for answering queries while providing empirical accuracy, we get the following corollary.

**Corollary 1.6.** *There is a computationally efficient mechanism  $\mathbb{M}$  that is  $(\alpha + 2\psi, \beta)$ -statistically-accurate for  $k$  adaptively chosen queries given a sample (an  $n$ -tuple) from an underlying measure with Gibbs-dependency  $\psi$  satisfying  $n \geq \tilde{O}\left(\frac{\sqrt{k}}{\alpha^2} \log \frac{1}{\beta}\right)$ .*

This generalizes the state-of-the-art bounds for the i.i.d. setting, where  $\psi = 0$ . In particular, Corollary 1.6 shows that mild dependencies in the data, say  $\psi = \alpha$ , come *for free* in terms of the achievable bounds for adaptive data analysis. We emphasize that  $\psi = \alpha$  captures non-negligible dependencies. In particular,  $\alpha$  could be constant, independent of the sample size  $n$ .

### 1.1.2. TRANSCRIPT COMPRESSION

The second direction we examine is that of *transcript compression*. The concept of compression is a central idea in the learning literature. It is both an algorithmic tool and a statistical tool, used both for designing learning mechanisms and achieving a better understanding of the concept of generalization (see for example (Littlestone and Warmuth, 1986; Daniely and Granot, 2019; Moran and Yehudayoff, 2016; Ashtiani et al., 2020; Hanneke and Kontorovich, 2021)).

Compression has also been used in the context of adaptive data analysis. Dwork et al. (2015a) used the definition of *bounded description length* (referred to here as *transcript compression*) to present an algorithm that is able to adaptively answer queries when the data is i.i.d. sampled. Our contribution here is in generalizing this idea by showing that the same definition, when used in the right setting, allows maintaining adaptive accuracy even when the distribution includes dependencies.

Following the approach of Bassily and Freund (2016), we aim to provide the following guarantee: As long as the analyst chooses functions which, in the non-adaptive setting, are concentrated around their expected value, then the answers given by the mechanism should be accurate. Intuitively, the idea is that functions with large variance are hard to approximate even in the non-adaptive setting, and hence, we should not require our mechanism to approxi-

mate them well in the adaptive setting.

This is formalized as follows. For every query  $q$  and every distribution  $\mu$ , we write  $\gamma(q, \mu, \delta)$  to denote the length of a confidence interval around the expectation of  $q$  with confidence level  $(1 - \delta)$ . That is,  $\gamma(q, \mu, \delta)$  is such that when sampling  $T \sim \mu$ , with probability at least  $(1 - \delta)$  it holds that  $q(T)$  is within  $\gamma(q, \mu, \delta)$  from its expectation. We obtain the following theorem (see Section 4 for a precise statement).

**Theorem 1.7 (informal).** *Fix  $\alpha, \delta > 0$ . There exists a computationally efficient mechanism with the following properties. The mechanism obtains a sample (an  $n$ -tuple) from some unknown underlying distribution  $\mu$ . Then, for  $k$  rounds  $i = 1, 2, \dots, k$ , the mechanism obtains a query  $q_i$  and responds with an answer  $a_i$  such that*

$$\Pr[\exists i \text{ s.t. } |a_i - q_i(\mu)| > \alpha + \gamma(q_i, \mu, \delta)] \leq \delta \cdot k \cdot 2^{k \cdot \log \frac{1}{\alpha}}.$$

In particular, as long as the adversary poses queries  $q_i$  such that  $\gamma(q_i, \mu, \delta) \leq \alpha$ , the mechanism from Theorem 1.7 guarantees that all of its answers are  $2\alpha$ -accurate, with probability at least  $1 - \delta \cdot k \cdot 2^{k \cdot \log \frac{1}{\alpha}}$ . In order for such a statement to be meaningful, we want to assert that  $\delta \ll \frac{1}{k} \cdot 2^{-k \cdot \log \frac{1}{\alpha}}$ . This is easily obtained in many settings of interest by taking the sample size  $n$  to be big enough. For example, for sub-Gaussian or sub-exponential queries, we would get that  $\delta$  vanishes exponentially with  $n$ , and hence, for large enough  $n$  we would get that  $\delta \ll \frac{1}{k} \cdot 2^{-k \cdot \log \frac{1}{\alpha}}$ .

## 1.2. Comparison to Bassily and Freund (2016)

Bassily and Freund (2016) also studied the problem of adaptive data analysis with correlated observations; we now elaborate on the differences.

- 1. Results regarding transcript compression 4.** As we mentioned, Bassily and Freund (2016) introduced the beautiful framework where the mechanism is required to provide accurate answers only as long as the analyst poses “concentrated queries”. They obtained their results for this setting via a new notion they introduced, called *typical stability*. However, their analysis and definitions are quite complex. We show that essentially the same bounds can be obtained *in a significantly simpler way, using standard compression tools*. Specifically, our result (Theorem 1.7) recovers essentially the same bounds for all types of queries considered by Bassily and Freund (2016), including bounded-sensitivity queries, subgaussian queries, and subexponential queries. In addition to being significantly simpler, our result offers the following advantage: Using the results of Bassily and Freund (2016), we need to know *in advance* the parameter controlling the “concentration level” of the queries that will

be presented in runtime, and this parameter is used by their algorithm. In contrast, our algorithm is oblivious to this parameter, and the guarantee is that our accuracy depends on the “concentration level” of the given queries. Furthermore, with our algorithm, different queries throughout the execution can have different “concentration levels”, a feature which is not directly supported by Bassily and Freund (2016).

2. **Results regarding Gibbs-dependence (Section 3).** These results are in a different setting than that of Bassily and Freund, and the results are not directly comparable. In particular, Bassily and Freund do not study any specific dependence notion such as Gibbs dependence. Our results show that assuming low Gibbs dependence allows for improved bounds. Specifically, Bassily and Freund (2016) can answer at most  $\tilde{O}(n)$  adaptive queries efficiently, even if the dependencies within the sample are very weak. Using our notion of Gibbs-dependency, we can answer  $\tilde{O}(n^2)$  adaptive queries efficiently, while accommodating small (but non-negligible) dependencies.

See Appendix A for additional related works.

## 2. Preliminaries

Our main metric for similarity between probability measures will be the *total variation distance*.

**Definition 2.1** (Total Variation Distance). *Given two measures  $\nu$  and  $\mu$  on the same space  $\Omega$ , the total variation distance between them is defined as  $\|\nu - \mu\|_{\text{TV}} := \sup_{A \subseteq \Omega} |\nu(A) - \mu(A)|$ , where the supremum is over the Borel sets of  $\Omega$ . Equivalently,  $\|\nu - \mu\|_{\text{TV}} = \frac{1}{2} \sum_{a \in \Omega} |\nu(a) - \mu(a)| = \frac{1}{2} \|\mu - \nu\|_{\ell_1}$ .*

### 2.1. Preliminaries from differential privacy

Differential privacy (Dwork et al., 2006) 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 2.2** (Differential Privacy (Dwork et al., 2006)). *A randomized algorithm  $\mathcal{A} : \mathcal{X}^n \rightarrow \mathcal{Y}$  is  $(\epsilon, \delta)$ -differentially private if for every two datasets  $S, S'$  which differ on a single element and for any event  $F$  we have  $\Pr[\mathcal{A}(S) \in F] \leq e^\epsilon \cdot \Pr[\mathcal{A}(S') \in F] + \delta$ .*

### 2.2. Preliminaries on adaptive data analysis

The standard formulation of adaptive data analysis is defined as a game involving some (adversary) analyst and a

---

### Algorithm 1 $\text{Game}(\mathbb{M}, k, \mathbb{A}, S)$

---

**Inputs:** Mechanism  $\mathbb{M}$ , interaction length  $k$ , adversary  $\mathbb{A}$ , dataset  $S$ .  
 The dataset  $S$  is given to  $\mathbb{M}$ .  
**for**  $i \in [k]$  **do**  
      $\mathbb{A}$  picks a query  $q_i$ .  
     The query  $q_i$  is given to  $\mathbb{M}$ .  
      $\mathbb{M}$  outputs an answer  $a_i$ .  
     The answer  $a_i$  is given to  $\mathbb{A}$ .  
**end for**

---

query-answering mechanism. For the sake of this paper queries are *statistical queries*, meaning they are functions of the form  $q : \mathcal{X} \rightarrow [0, 1]$ . The goal of the mechanism is to make sure that the answers provided to the analyst are accurate w.r.t. the expected value of the corresponding queries over the underlying distribution. The idea is to formalize a utility notion that holds for *any* strategy of the data analyst. 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 any such and *adversarial* analyst, then it maintains utility against any analyst. This game is specified in Algorithm 1.

**Definition 2.3** (Adaptive Empirical Accuracy). *A mechanism  $\mathbb{M}$  is  $(\alpha, \beta)$ -empirically-accurate for  $k$  rounds given a dataset of size  $n$ , if for every dataset  $S$  of size  $n$  and every adversary  $\mathbb{A}$ , it holds that  $\Pr_{\text{Game}(\mathbb{M}, k, \mathbb{A}, S)} [\max_{i \in [k]} |q_i(S) - a_i| > \alpha] \leq \beta$ , where  $q_i(S) := \frac{1}{|S|} \sum_{x \in S} q_i(x)$ .*

**Definition 2.4** (Adaptive Statistical Accuracy). *A mechanism  $\mathbb{M}$  is  $(\alpha, \beta)$ -statistically-accurate for  $k$  rounds given  $n$  samples with Gibbs-dependence  $\psi$ , if for every distribution  $\mu$  over  $n$ -tuples with Gibbs dependency  $\psi$ , and every adversary  $\mathbb{A}$ , it holds that  $\Pr_{\substack{S \sim \mu \\ \text{Game}(\mathbb{M}, k, \mathbb{A}, S)}} [\max_{i \in [k]} |q_i(\mu) - a_i| > \alpha] \leq \beta$ , where  $q_i(\mu) := \mathbb{E}_{T \sim \mu} [q_i(T)] = \mathbb{E}_{T \sim \mu} \left[ \frac{1}{|T|} \sum_{x \in T} q_i(x) \right]$ .*

**Remark 2.5.** *The above definition is stated in general form, but in fact it is sufficient to show that a mechanism  $\mathbb{M}$  exhibits the above guarantee for every deterministic adversary  $\mathbb{A}$ . The reason is that for a randomized adversary one can fix the adversary’s random coins and use the total probability law in order to get the same result.*

## 3. Adaptive Generalization via Differential Privacy

We extend the connection between differential privacy and adaptive data analysis into settings where the data is not sampled in an i.i.d. fashion, but rather there are some small/bounded dependencies. We start by proving the following lemma, showing that differential privacy

guarantees generalization in expectation. The proof of this lemma mimics the analysis of Bassily et al. (2016) for the i.i.d. setting. We extend the proof to the case where there are dependencies in the data, and show that we can “pay” for these dependencies in a way that scales with  $\psi$ .

**Lemma 3.1** (Expectation bound). *Let  $\mathcal{A}' : (\mathcal{X}^n)^T \rightarrow 2^{\mathcal{X}} \times [T]$  be an  $(\varepsilon, \delta)$ -differentially private algorithm. Let  $\mu$  be a distribution over  $\mathcal{X}^n$  which has  $\psi$ -Gibbs-dependence let  $\vec{S} = (S_1, \dots, S_T)$  where for every  $i$   $S_i \sim \mu$ . Denote by  $(h, t)$  the output of  $\mathcal{A}'(\vec{S})$ . Then  $\left| \mathbb{E}_{\vec{S}, \mathcal{A}'} [h(\mu) - h(S_t)] \right| \leq e^\varepsilon + T\delta + \psi - 1$ .*

*Proof.* We consider a multi sample  $\vec{S} = (S_1, \dots, S_T)$ , where  $S_t = (x_{t,1}, \dots, x_{t,n}) \sim \mu$ . We calculate,

$$\begin{aligned}
 & \mathbb{E}_{\vec{S} \sim \mu^T} \left[ \mathbb{E}_{(h,t) \sim \mathcal{A}'(\vec{S})} [h(S_t)] \right] \\
 = & \mathbb{E}_{\vec{S} \sim \mu^T} \left[ \mathbb{E}_{(h,t) \sim \mathcal{A}'(\vec{S})} \left[ \frac{1}{n} \sum_{i=1}^n h(x_{t,i}) \right] \right] \\
 = & \frac{1}{n} \sum_{i=1}^n \left[ \mathbb{E}_{\vec{S} \sim \mu^T} \left[ \mathbb{E}_{(h,t) \sim \mathcal{A}'(\vec{S})} [h(x_{t,i})] \right] \right] \\
 = & \frac{1}{n} \sum_{i=1}^n \left[ \mathbb{E}_{\vec{S} \sim \mu^T} \left[ \Pr_{(h,t) \sim \mathcal{A}'(\vec{S})} [h(x_{t,i}) = 1] \right] \right] \\
 = & \frac{1}{n} \sum_{i=1}^n \left[ \mathbb{E}_{\vec{S} \sim \mu^T} \left[ \sum_{m=1}^T \Pr_{(h,t) \sim \mathcal{A}'(\vec{S})} [h(x_{m,i}) = 1 \wedge t = m] \right] \right] \\
 = & \frac{1}{n} \sum_{i=1}^n \left[ \mathbb{E}_{\vec{S} \sim \mu^T} \left[ \mathbb{E}_{\vec{z} \sim \vec{\mu}_i | \vec{S}} \left[ \Pr_{(h,t) \sim \mathcal{A}'(\vec{S})} [h(x_{m,i}) = 1 \wedge t = m] \right] \right] \right], \quad (1)
 \end{aligned}$$

where  $\vec{z} = (z_1, \dots, z_T)$  is a vector s.t.  $z_t \sim \mu_i(\cdot | S_t^{-i})$ . Given a multi-sample  $\vec{S}$  and an element  $z$ , we write  $\vec{S}^{(m,i) \leftarrow z}$  to denote the multi-sample  $\vec{S}$  after replacing the  $i^{\text{th}}$  element in the  $m^{\text{th}}$  sample  $S_m$  with  $z$ . Since  $\mathcal{A}'$  is

$(\varepsilon, \delta)$ -differentially private we get that the above is at most

$$\begin{aligned}
 (1) & \leq \frac{1}{n} \sum_{i=1}^n \left[ \mathbb{E}_{\vec{S} \sim \mu^T} \left[ \mathbb{E}_{\vec{z} \sim \vec{\mu}_i | \vec{S}} \left[ \sum_{m=1}^T e^\varepsilon \Pr_{(h,t) \sim \mathcal{A}'(\vec{S}^{(m,i) \leftarrow z_m})} \right. \right. \right. \\
 & \left. \left. \left. [h(x_{m,i}) = 1 \wedge t = m] + \delta \right] \right] \right] \\
 = & T\delta + e^\varepsilon \cdot \frac{1}{n} \sum_{i=1}^n \sum_{m=1}^T \left[ \mathbb{E}_{\vec{S} \sim \mu^T} \left[ \mathbb{E}_{\vec{z} \sim \vec{\mu}_i | \vec{S}} \left[ \Pr_{(h,t) \sim \mathcal{A}'(\vec{S}^{(m,i) \leftarrow z_m})} [h(x_{m,i}) = 1 \wedge t = m] \right] \right] \right] \\
 = & T\delta + e^\varepsilon \cdot \frac{1}{n} \sum_{i=1}^n \sum_{m=1}^T \left[ \mathbb{E}_{\vec{S} \sim \mu^T} \left[ \mathbb{E}_{\vec{z} \sim \vec{\mu}_i | \vec{S}} \left[ \Pr_{(h,t) \sim \mathcal{A}'(\vec{S})} [h(z_m) = 1 \wedge t = m] \right] \right] \right] \\
 = & T\delta + e^\varepsilon \cdot \frac{1}{n} \sum_{i=1}^n \sum_{m=1}^T \left[ \mathbb{E}_{\vec{S} \sim \mu^T} \left[ \mathbb{E}_{\vec{z} \sim \vec{\mu}_i | \vec{S}} \left[ \Pr_{(h,t) \sim \mathcal{A}'(\vec{S})} [h(z_t) = 1 \wedge t = m] \right] \right] \right] \\
 = & T\delta + e^\varepsilon \cdot \frac{1}{n} \sum_{i=1}^n \left[ \mathbb{E}_{\vec{S} \sim \mu^T} \left[ \mathbb{E}_{\vec{z} \sim \vec{\mu}_i | \vec{S}} \left[ \sum_{m=1}^T \Pr_{(h,t) \sim \mathcal{A}'(\vec{S})} [h(z_t) = 1 \wedge t = m] \right] \right] \right] \\
 = & T\delta + e^\varepsilon \cdot \frac{1}{n} \sum_{i=1}^n \left[ \mathbb{E}_{\vec{S} \sim \mu^T} \left[ \mathbb{E}_{\vec{z} \sim \vec{\mu}_i | \vec{S}} \left[ \Pr_{(h,t) \sim \mathcal{A}'(\vec{S})} [h(z_t) = 1] \right] \right] \right] \\
 = & T\delta + e^\varepsilon \frac{1}{n} \sum_{i=1}^n \left[ \mathbb{E}_{\vec{S} \sim \mu^T} \left[ \mathbb{E}_{\vec{z} \sim \vec{\mu}_i | \vec{S}} \left[ \mathbb{E}_{(h,t) \sim \mathcal{A}'(\vec{S})} [h(z_t)] \right] \right] \right] \\
 = & T\delta + e^\varepsilon \frac{1}{n} \sum_{i=1}^n \left[ \mathbb{E}_{\vec{S} \sim \mu^T} \left[ \mathbb{E}_{(h,t) \sim \mathcal{A}'(\vec{S})} \left[ \mathbb{E}_{\vec{z} \sim \vec{\mu}_i | \vec{S}} [h(z_t)] \right] \right] \right] \\
 = & T\delta + e^\varepsilon \cdot \frac{1}{n} \sum_{i=1}^n \left[ \mathbb{E}_{\vec{S} \sim \mu^T} \left[ \mathbb{E}_{(h,t) \sim \mathcal{A}'(\vec{S})} \left[ \mathbb{E}_{z \sim \mu_i(\cdot | S_t^{-i})} [h(z)] \right] \right] \right]. \quad (2)
 \end{aligned}$$

Since total variation is a special case of the Wasserstein metric  $\mathcal{W}_1$ , Kantorovich-Rubinstein duality implies that for two probability measures  $\mu, \nu$  on a space  $\mathcal{X}$  and any function  $h : \mathcal{X} \rightarrow [0, 1]$ , we have  $|\mathbb{E}_{z \sim \mu}[h(x)] - \mathbb{E}_{z \sim \nu}[h(z)]| \leq \|\mu - \nu\|_{\text{TV}}$ . Applying this

to  $\mu_i(\cdot \mid S_t^{-i})$  and  $\mu_i$  we get that the above is at most

$$\begin{aligned}
 (2) &\leq T\delta + e^\varepsilon \cdot \frac{1}{n} \sum_{i=1}^n \left[ \mathbb{E}_{\vec{s} \sim \mu^T} \left[ \mathbb{E}_{(h,t) \sim \mathcal{A}'(\vec{s})} \left[ \mathbb{E}_{z \sim \mu_i} [h(z)] + \|\mu_i(\cdot \mid S_t^{-i}) - \mu_i\|_{\text{TV}} \right] \right] \right] \\
 &\leq \psi + T\delta + e^\varepsilon \cdot \mathbb{E}_{\vec{s} \sim \mu^T} \left[ \mathbb{E}_{(h,t) \sim \mathcal{A}'(\vec{s})} \left[ \mathbb{E}_{z \sim \mu_i} [h(z)] \right] \right] \\
 &= \psi + T\delta + e^\varepsilon \cdot \mathbb{E}_{\vec{s}, \mathcal{A}'(\vec{s})} [h(\mu)] \\
 &\leq \psi + T\delta + e^\varepsilon - 1 + \mathbb{E}_{\vec{s}, \mathcal{A}'(\vec{s})} [h(\mu)],
 \end{aligned}$$

where the last inequality is due to the fact that  $ye^\varepsilon \leq e^\varepsilon - 1 + y$  for  $y \leq 1$  and  $\varepsilon \geq 0$ . In summary,

$$\begin{aligned}
 &\mathbb{E}_{\vec{s} \sim \mu^T} \left[ \mathbb{E}_{(h,t) \sim \mathcal{A}'(\vec{s})} [h(S_t)] \right] \\
 &\leq \psi + T\delta + e^\varepsilon - 1 + \mathbb{E}_{\vec{s}, \mathcal{A}'(\vec{s})} [h(\mu)].
 \end{aligned}$$

The other direction is symmetric.  $\square$

We use Lemma 3.1 to prove the following high-probability generalization bound for differentially private algorithms.

**Theorem 3.2** (High probability bound). *Let  $\varepsilon \in (0, 1/3)$ ,  $\delta \in (0, \varepsilon/4)$  and  $n \geq \frac{\log(2k\varepsilon/\delta)}{\varepsilon^2}$ . Let  $\mathcal{A} : \mathcal{X}^n \rightarrow (2^{\mathcal{X}})^k$  be an  $(\varepsilon, \delta)$ -differentially private algorithm. Let  $\mu$  be a distribution over  $\mathcal{X}^n$  and  $S$  be a sample of size  $n$  drawn from  $\mu$ , and let  $h_1, \dots, h_k$  be the output of  $\mathcal{A}(S)$ . Then*

$$\Pr_{S, \mathcal{A}(S)} \left[ \max_{i \in [k]} |h_i(\mu) - h_i(S)| \geq 10\varepsilon + 2\psi \right] \leq \frac{\delta}{\varepsilon}.$$

The proof of Theorem 3.2 is almost identical to the analysis of Bassily et al. (2016). It appears in the appendix for completeness. Intuitively, the proof is as follows. We assume, towards contradiction, that there may be a differentially private algorithm that does not enjoy strong generalization guarantees. We then use this mechanism to describe a different differentially private algorithm with a ‘‘boosted inability’’ to generalize. That is, the proof goes by saying that if there is a differentially private algorithm whose generalization properties are not ‘‘very good’’ then there must exist a differentially private algorithm whose generalization properties are ‘‘bad’’, to the extent that contradicts Lemma 3.1.

Our connection between Gibbs-dependence and differential privacy (Theorem 1.4) now follows as a corollary of Theorem 3.2.

*Proof of Theorem 1.4.*  $\mathbb{M}$  is  $(\varepsilon, \delta)$ -differentially private. Since  $\mathbb{A}$  can only access the data via  $\mathbb{M}$ , we can treat the pair  $\mathbb{A}, \mathbb{M}$  as a single algorithm  $\mathcal{A}$ , which gets a sample  $S \sim \mu$  as input and returns  $k$  predicates, as output. By closure to post-processing,  $\mathcal{A}$  is also  $(\varepsilon, \delta)$ -differentially private. Applying Theorem 3.2 on  $\mathcal{A}$  we get that

$$\Pr \left[ \max_{i \in [k]} |h_i(\mu) - h_i(S)| \geq 10\varepsilon + 2\psi \right] \leq \frac{\delta}{\varepsilon}.$$

Since  $\mathbb{M}$  is  $(\alpha, \beta)$ -empirically-accurate it holds that

$$\Pr \left[ \max_{i \in [k]} |q_i(S) - a_i| > \alpha \right] \leq \beta.$$

Combining these two bounds with the triangle inequality, we get

$$\Pr \left[ \max_{i \in [k]} |q_i(\mu) - a_i| > \alpha + 10\varepsilon + 2\psi \right] < \beta + \frac{\delta}{\varepsilon}.$$

$\square$

### 3.1. A Tight Negative Result for Differential Privacy and Gibbs-Dependence

In this section, we construct a distribution which is  $\psi$ -Gibbs-Dependant, and describe a differentially-private algorithm whose generalization gap w.r.t. this distribution is at least  $\psi$ . Hence, in general, the  $\psi$  factor attained on Theorem 1.4 is tight up to a constant. Let  $\mathcal{X} = [0, 1]$  and define a measure  $\mu$  over  $\mathcal{X}^n$  by the following random process:

1. Sample a point  $x^* \sim U([0, 1])$ .
2. For every  $i \in [n]$ :
  - (a) Sample  $\sigma \sim \text{Ber}(\psi)$ .
    - i. If  $\sigma = 1$  then  $x_i = x^*$ .
    - ii. Otherwise  $x_i \sim U([0, 1])$
3. Return  $S = (x_1, \dots, x_n)$

**Lemma 3.3.** *The measure defined by the above process has  $\psi$ -Gibbs-dependency.*

*Proof.* Initially, every marginal distribution is just uniform, i.e.  $\mu_i \sim U([0, 1])$  and hence, for every  $A \subseteq [0, 1]$  it holds that  $\mu_i(A) = |A|$ . After conditioning, for every possible  $x^{-i}$  and  $x^*$ , we get that

$$\begin{aligned}
 &\mu_i(A \mid x^{-i}, x^*) \\
 &= \mu_i(A \setminus \{x^*\} \mid x^{-i}, x^*) + \mu_i(A \cap \{x^*\} \mid x^{-i}, x^*) \\
 &\in \left( |A|(1 - \psi), |A|(1 - \psi) + \psi \right).
 \end{aligned}$$

Since the above holds for every choice of  $x^*$ , we also have that

$$\mu_i(A \mid x^{-i}) \in \left( |A|(1 - \psi), |A|(1 - \psi) + \psi \right).$$

Therefore, for every  $A \subseteq [0, 1]$  it holds that

$$\begin{aligned} & |\mu_i(A) - \mu_i(A \mid x^{-i})| \\ & \leq \max \{ |A| - |A|(1 - \psi), |A|(1 - \psi) + \psi - |A| \} \leq \psi. \end{aligned}$$

So  $\|\mu_i(\cdot) - \mu_i(\cdot \mid x^{-i})\|_{\text{TV}} \leq \psi$ . Plugging this bound to the Gibbs-dependency definition yields

$$\psi(\mu) = \sup_{x \in \mathcal{X}^n} \mathbb{E}_{i \sim [n]} \|\mu_i(\cdot) - \mu_i(\cdot \mid x^{-i})\|_{\text{TV}} \leq \psi.$$

□

We next describe an algorithm that, despite being differentially private, performs “badly” when executed on samples from the above measure  $\mu$ . Specifically, this algorithm is capable of identifying a predicate with generalization error  $\Omega(\psi)$ . This shows that our connection between differential privacy and generalization (in the correlated setting) is tight, in the sense that the generalization error of differentially private algorithms *can* grow with  $\psi$ . This matches our positive result (see Theorem 1.4).

Our algorithm is specified in Algorithm 2. As a subroutine, we use the following result of Bun et al. (2019) for privately computing histograms.

**Theorem 3.4** (Private histograms, (Bun et al., 2019)). *There exists an  $(\varepsilon, \delta)$ -differentially private algorithm that takes an input dataset  $S \in \mathcal{X}^n$  and returns a list  $L \subseteq \mathcal{X}$  such that with probability at least  $1 - \beta$ , every  $x \in \mathcal{X}$  that appears at least  $\mathcal{O}\left(\frac{1}{\varepsilon} \log \frac{1}{\beta\delta}\right)$  times in  $S$  is included in  $L$ , and furthermore, every  $x \in L$  appears at least twice in  $S$ .*

---

#### Algorithm 2 Deviating Private Algorithm

---

**Input:** A sample  $S$ , privacy parameters  $\varepsilon, \delta$ .  
**Tool used:** An  $(\varepsilon, \delta)$ -DP algorithm  $\mathcal{H}$  for histograms.  
 $L \leftarrow \mathcal{H}(S, \varepsilon, \delta)$   
**if**  $L$  is empty **then**  
     Return  $h \equiv 0$   
**else**  
     Let  $x$  be an arbitrary element in  $L$   
     Define  $h : \mathcal{X} \rightarrow [0, 1]$  as  $h = \mathbb{1}[x]$   
     Return  $h$   
**end if**

---

**Lemma 3.5.** *For every  $\beta > 0$ , every  $n \geq \mathcal{O}\left(\frac{1}{\psi\varepsilon} \log \frac{1}{\beta\delta}\right)$ , and for every  $\psi < 1$  Algorithm 2 is  $(\varepsilon, \delta)$ -differentially private and it outputs a predicate  $h : \mathcal{X} \rightarrow [0, 1]$  s.t.*

$$\Pr \left[ |h(S) - h(\mu)| \geq \frac{\psi}{2} \right] > 1 - \beta - \exp\left(-\frac{n}{8}\right).$$

*Proof.* First observe that Algorithm 2 is  $(\varepsilon, \delta)$ -differentially private, as it merely post-processes the outcome of the private histogram algorithm.

Next observe that, by the definition of the underlying measure  $\mu$ , and by our choice of  $n$ , w.h.p., there are many copies of  $x^*$  in the dataset  $S$ . Formally, by the Chernoff bound,

$$\begin{aligned} & \Pr \left[ \frac{1}{n} |\{x' \in S \mid x' = x^*\}| < \frac{1}{2}\psi \right] \\ & = \Pr \left[ \frac{1}{n} \sum_{i=1}^n \sigma_i < \frac{1}{2}\psi \right] \leq \exp\left(-\frac{n}{8}\right). \end{aligned}$$

In addition, the probability of any element  $x \neq x^*$  appearing more than once in  $S$  is simply zero. Thus, with probability at least  $1 - \exp\left(-\frac{n}{8}\right)$  we have that  $x^*$  appears in  $S$  at least  $n\psi/2 = \Omega\left(\frac{1}{\varepsilon} \log\left(\frac{1}{\beta\delta}\right)\right)$  times, and every other element appears in  $S$  at most once. By the properties of the private histogram algorithm (see Theorem 3.4), in such a case, with probability at least  $1 - \beta$  we have that  $L = \{x^*\}$ , and Algorithm 2 returns the hypothesis  $h = \mathbb{1}[x^*]$ . As  $x^*$  appears many times in  $S$ , this predicate has “large” empirical value. On the other hand, for such predicate it holds that

$$\begin{aligned} h(\mu) & = \mathbb{E}_{\bar{x}^*, \bar{x}_1, \dots, \bar{x}_n} \left[ \frac{1}{n} \left( \sum_{i=1}^n h(\bar{x}_i) \right) \right] \\ & = \frac{1}{n} \left( \sum_{i=1}^n \Pr_{\bar{x}^*, \bar{x}_1, \dots, \bar{x}_n} [\bar{x}_i = x^*] \right) = 0 \end{aligned}$$

as the probability that for a fresh new sampling we will get  $\bar{x}^* = x^*$  is zero, implying that the probability that any point in the sample to be  $x^*$  is also zero.

Overall, with probability at least  $1 - \beta - \exp\left(-\frac{n}{8}\right)$ , the algorithm returns a predicate  $h$  such that  $h(S) \geq \psi/2$  but  $h(\mu) = 0$ .

□

## 4. Adaptive Learning Via Transcript Compression

In this section we show how the notion of *transcript compressibility* can be used to derive generalization bounds even if the data is not i.i.d. distributed. We start by recalling the notion of *transcript compression* by Dwork et al. (2015a). We denote by  $AG_{n,k}(\mathcal{A}, \mathbb{M}, S)$  the *transcript* of the interaction between the mechanism  $\mathbb{M}$  and the analysis  $\mathcal{A}$  during the adaptive accuracy game defined in Algorithm 1 with sample of size  $n$  and  $k$  queries.

**Definition 4.1** (Transcript Compression (Dwork et al., 2015a)). *We say that a mechanism  $\mathbb{M}$  enables transcript*

compression to  $b(n, k)$ -bits, if for every deterministic analyst  $\mathcal{A}$  there exist a set of possible transcripts  $\mathcal{H}_{\mathcal{A}}$  of size  $|\mathcal{H}_{\mathcal{A}}| \leq 2^{b(n, k)}$ , s.t. for every sample  $S$  it holds that  $\Pr [AG_{n, k}(\mathcal{A}, \mathbb{M}, S) \in \mathcal{H}] = 1$ .

Following Bassily and Freund (2016), in this section we aim to design mechanisms that answer adaptively chosen queries while providing statistical accuracy, under the assumption that the given queries are *concentrated* around their expected value. Unlike Bassily and Freund (2016), we aim to achieve this goal using the notion of *transcript compression*, rather than *typical-stability*. As we show, this allows for a significantly simpler analysis (and definitions). Formally,

**Definition 4.2.** Given a measure  $\mu$  over  $\mathcal{X}$ , a query  $q : \mathcal{X}^n \rightarrow \mathbb{R}$ , and a parameter  $\delta \in [0, 1]$ , we write  $\gamma(q, \mu, \delta)$  to denote the minimal number  $\gamma \in [0, 1]$  such that

$$\Pr_{S \sim \mu} \left[ \left| q(S) - \mathbb{E}_{T \sim \mu} [q(T)] \right| > \gamma \right] < \delta.$$

That is,  $\gamma(q, \mu, \delta)$  denotes the minimal number such that, without adaptivity,  $q(S)$  deviates from its expectation by more than  $\gamma(q, \mu, \delta)$  with probability at most  $\delta$  when sampling  $S \sim \mu$ .

**Remark 4.3.** The results in this section are not restricted to statistical queries. The results in this section hold for arbitrary queries (mapping  $n$ -tuples to the reals).

Consider again Algorithm 1 and Definition 2.4 (the definition of statistical accuracy). We now use Definition 4.2 in order to introduce a relaxation for statistical accuracy, in which the mechanism is allowed to incur  $\gamma(q, \mu, \delta)$  as an additional error.

**Definition 4.4.** A mechanism  $\mathbb{M}$  is  $(\alpha, \beta, \delta)$ -statistically-query-accurate for  $k$  rounds given  $n$  samples, if for every distribution  $\mu$  over  $n$ -tuples, and every adversary  $\mathbb{A}$ , it holds that

$$\Pr_{\substack{S \sim \mu \\ \text{Game}(\mathbb{M}, k, \mathbb{A}, S)}} \left[ \max_{i \in [k]} |q_i(\mu) - a_i| > \alpha + \gamma(q_i, \mu, \delta) \right] \leq \beta.$$

**Remark 4.5.** For a statistical query  $q$  and a product measure  $\mu$ , by Hoeffding's inequality, we get that  $\gamma(q, \mu, \delta) = \sqrt{\frac{1}{2n} \ln \frac{2}{\delta}}$ . Hence, for the i.i.d. regime, for large enough samples, the definition of  $(\alpha, \beta, \delta)$ -statistical-query-accuracy is in fact equivalent (up to factor 2) to the original definition of  $(\alpha, \beta)$ -statistical-accuracy (Definition 2.4).

We observe that the analysis of Dwork et al. (2015a) for transcript compression easily extends to non-i.i.d. measures when given concentrated queries. Somewhat surprisingly,

this simple technique essentially matches the bounds obtained using typical stability (Bassily and Freund, 2016). In the next lemma we show that (w.h.p.) an analyst interacting with a transcript-compressing mechanism cannot identify a query that overfits to the data.

**Lemma 4.6.** Let  $\mathbb{M}$  be a mechanism which enables transcript compression to  $b(n, k)$ -bits. For every measure  $\mu$  and every analyst  $\mathcal{A}$ ,

$$\Pr_{S, AG_{n, k}} [\exists i : |q_i(S) - q_i(\mu)| \geq \gamma(q, \mu, \delta)] \leq \delta \cdot k \cdot 2^{b(n, k)}$$

*Proof.* Fix an analyst  $\mathcal{A}$ . By Definition 4.1, there exist a set of transcripts  $H_{\mathcal{A}}$  of size at most  $2^{b(n, k)}$ . As every transcript consists of at most  $k$  queries, there can be at most  $k \cdot 2^{b(n, k)}$  possible queries over all possible interactions between  $\mathcal{A}$  and  $\mathbb{M}$ . Denote this set of possible queries as  $Q_{\mathcal{A}}$ . By a union bound we get that

$$\Pr_{S \sim \mu} \left[ \bigvee_{q \in Q_{\mathcal{A}}} |q_i(S) - q_i(\mu)| \geq \gamma(q, \mu, \delta) \right] \leq \delta \cdot k \cdot 2^{b(n, k)},$$

and hence

$$\Pr_{S, AG_{n, k}} [\exists i : |q_i(S) - q_i(\mu)| \geq \gamma(q, \mu, \delta)] \leq k \cdot \delta \cdot 2^{b(n, k)}.$$

□

Using the above lemma, we prove our main theorem for this section.

**Theorem 4.7.** Let  $\mathbb{M}$  be a mechanism which enables transcript compression to  $b(n, k)$  bits and also exhibits  $(\alpha, \beta)$ -empirical-accuracy for  $k$  rounds given  $n$  samples. Then  $\mathbb{M}$  is also  $(\alpha, \beta + \delta k \cdot 2^{b(n, k)}, \delta)$ -statistically-query-accurate, for every choice of  $\delta$ .

*Proof.* As  $\mathbb{M}$  is  $(\alpha, \beta)$ -empirically-accurate and also enables transcript compression to  $b(n, k)$  bits, by Lemma 4.6 and the union bound

$$\begin{aligned} & \Pr_{S, AG_{n, k}} \left[ (\exists i : |q_i(S) - q_i(\mu)| > \gamma(q, \mu, \delta)) \right. \\ & \quad \left. \vee (\exists i : |q_i(S) - a_i| > \alpha) \right] \\ & \leq \beta + \delta \cdot k \cdot 2^{b(n, k)}. \end{aligned}$$

Hence by the triangle inequality

$$\begin{aligned} & \Pr_{S, AG_{n, k}} [\exists i : |a_i - q_i(\mu)| \geq \alpha + \gamma(q, \mu, \delta)] \\ & \leq \beta + \delta \cdot k \cdot 2^{b(n, k)}. \end{aligned}$$

□

Applying Theorem 4.7 together with the transcript-compressing mechanisms of Dwork et al. (2015a), we get the following two results.



**Theorem 4.8.** For every  $\alpha, \delta$ , there exists an  $(\alpha, \beta, \delta)$ -statistically-query-accurate mechanism for  $k$  rounds given  $n$  samples, where  $\beta = k \cdot \delta \cdot 2^{k \cdot \log \frac{1}{\alpha}}$ . The mechanism is computationally efficient.

**Theorem 4.9.** For every  $\delta$ , there exists an  $(\alpha, \beta, \delta)$ -statistically-query-accurate mechanism for  $k$  rounds given  $n$  samples, where  $\alpha = \mathcal{O}\left(\left(\frac{\ln k}{n}\right)^{1/4}\right)$  and  $\beta = k \cdot \delta \cdot 2^{\tilde{\mathcal{O}}(\sqrt{n} \cdot \log |\mathcal{X}| \cdot (\log k)^{3/2})}$ . The mechanism is computationally inefficient.

## Acknowledgements

M.S. was supported in part by the Lynn and William Frankel Center for Computer Science at Ben-Gurion University. A.K. and M.S. were supported in part by Ben-Gurion Data Science Research Center at Ben-Gurion University of the Negev. U.S. was Partially supported by Israel Science Foundation (grant 1871/19) and by Len Blavatnik and the Blavatnik Family foundation.

## References

- Ashtiani, H., Ben-David, S., Harvey, N. J., Liaw, C., Mehrabian, A., and Plan, Y. (2020). Near-optimal sample complexity bounds for robust learning of gaussian mixtures via compression schemes. *Journal of the ACM (JACM)*, 67(6):1–42.
- Attias, I., Cohen, E., Shechner, M., and Stemmer, U. (2021). A framework for adversarial streaming via differential privacy and difference estimators. *CoRR*, abs/2107.14527.
- Banf, M. and Rhee, S. Y. (2017). Enhancing gene regulatory network inference through data integration with markov random fields. *Scientific reports*, 7(1):1–13.
- Bassily, R. and Freund, Y. (2016). Typicality-based stability and privacy. *CoRR*, abs/1604.03336.
- Bassily, R., Nissim, K., Smith, A., Steinke, T., Stemmer, U., and Ullman, J. (2016). Algorithmic stability for adaptive data analysis. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 1046–1059.
- Beimel, A., Kaplan, H., Mansour, Y., Nissim, K., Saranurak, T., and Stemmer, U. (2021). Dynamic algorithms against an adaptive adversary: Generic constructions and lower bounds. *CoRR*, abs/2111.03980.
- Blake, A., Kohli, P., and Rother, C. (2011). *Markov random fields for vision and image processing*. MIT Press.
- Bousquet, O. and Elisseeff, A. (2002a). Stability and generalization. *Journal of machine learning research*, 2(Mar):499–526.
- Bousquet, O. and Elisseeff, A. (2002b). Stability and generalization. *J. Mach. Learn. Res.*, 2:499–526.
- Bun, M., Nissim, K., and Stemmer, U. (2019). Simultaneous private learning of multiple concepts. *J. Mach. Learn. Res.*, 20:94:1–94:34.
- Bun, M., Ullman, J., and Vadhan, S. (2018). Fingerprinting codes and the price of approximate differential privacy. *SIAM Journal on Computing*, 47(5):1888–1938.
- Clifford, P. (1990). Markov random fields in statistics. In Grimmett, G. and Welsh, D., editors, *Disorder in Physical Systems: A Volume in Honour of John M. Hammett*, pages 19–32. Oxford University Press, Oxford.
- Dagan, Y., Daskalakis, C., Dikkala, N., and Jayanti, S. (2019). Learning from weakly dependent data under dobrushin’s condition. In Beygelzimer, A. and Hsu, D., editors, *Conference on Learning Theory, COLT 2019, 25-28 June 2019, Phoenix, AZ, USA*, volume 99 of *Proceedings of Machine Learning Research*, pages 914–928. PMLR.
- Daniely, A. and Granot, E. (2019). Generalization bounds for neural networks via approximate description length. In Wallach, H., Larochelle, H., Beygelzimer, A., d’Alché-Buc, F., Fox, E., and Garnett, R., editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Daskalakis, C., Dikkala, N., and Panageas, I. (2019). Regression from dependent observations. In Charikar, M. and Cohen, E., editors, *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019, Phoenix, AZ, USA, June 23-26, 2019*, pages 881–889. ACM.
- Dwork, C., Feldman, V., Hardt, M., Pitassi, T., Reingold, O., and Roth, A. (2015a). Generalization in adaptive data analysis and holdout reuse. In *Advances in Neural Information Processing Systems*, pages 2350–2358.
- Dwork, C., Feldman, V., Hardt, M., Pitassi, T., Reingold, O., and Roth, A. (2015b). The reusable holdout: Preserving validity in adaptive data analysis. *Science*, 349(6248):636–638.
- Dwork, C., Feldman, V., Hardt, M., Pitassi, T., Reingold, O., and Roth, A. L. (2015c). Preserving statistical validity in adaptive data analysis. In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, pages 117–126.
- Dwork, C., McSherry, F., Nissim, K., and Smith, A. (2006). Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*, pages 265–284. Springer.

- Feldman, V. and Vondrák, J. (2018). Generalization bounds for uniformly stable algorithms. In Bengio, S., Wallach, H. M., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R., editors, *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, 3-8 December 2018, Montréal, Canada*, pages 9770–9780.
- Feldman, V. and Vondrák, J. (2019). High probability generalization bounds for uniformly stable algorithms with nearly optimal rate. In Beygelzimer, A. and Hsu, D., editors, *Conference on Learning Theory, COLT 2019, 25-28 June 2019, Phoenix, AZ, USA*, volume 99 of *Proceedings of Machine Learning Research*, pages 1270–1279. PMLR.
- Hammersley, J. M. and Clifford, P. E. (1971). Markov random fields on finite graphs and lattices. Unpublished manuscript.
- Hanneke, S. et al. (2014). Theory of disagreement-based active learning. *Foundations and Trends® in Machine Learning*, 7(2-3):131–309.
- Hanneke, S. and Kontorovich, A. (2021). Stable sample compression schemes: New applications and an optimal svm margin bound. In *Algorithmic Learning Theory*, pages 697–721. PMLR.
- Hardt, M., Recht, B., and Singer, Y. (2016). Train faster, generalize better: Stability of stochastic gradient descent. In Balcan, M. F. and Weinberger, K. Q., editors, *International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 1225–1234, New York, New York, USA. PMLR.
- Hardt, M. and Ullman, J. (2014). Preventing false discovery in interactive data analysis is hard. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, pages 454–463. IEEE.
- Hassidim, A., Kaplan, H., Mansour, Y., Matias, Y., and Stemmer, U. (2020). Adversarially robust streaming algorithms via differential privacy. In *NeurIPS*.
- Juang, B.-H. and Rabiner, L. R. (1991). Hidden markov models for speech recognition.
- Jung, C., Ligett, K., Neel, S., Roth, A., Sharifi-Malvajerdi, S., and Shenfeld, M. (2020). A new analysis of differential privacy’s generalization guarantees. In *ITCS*, volume 151 of *LIPICs*, pages 31:1–31:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.
- Kaplan, H., Mansour, Y., Nissim, K., and Stemmer, U. (2021). Separating adaptive streaming from oblivious streaming using the bounded storage model. In *CRYPTO* (3), volume 12827 of *Lecture Notes in Computer Science*, pages 94–121. Springer.
- Kontorovich, A. (2012). Obtaining measure concentration from Markov contraction. *Markov Processes and Related Fields*, 4:613–638.
- Kontorovich, A. and Raginsky, M. (2017). Concentration of measure without independence: a unified approach via the martingale method. In *Convexity and Concentration*, pages 183–210. Springer.
- Kontorovich, A. and Weiss, R. (2014). Uniform Chernoff and Dvoretzky-Kiefer-Wolfowitz-type inequalities for Markov chains and related processes. *Journal of Applied Probability*, 51(4):1100 – 1113.
- Kontorovich, L. A. and Ramanan, K. (2008). Concentration Inequalities for Dependent Random Variables via the Martingale Method. *Ann. Probab.*, 36(6):2126–2158.
- Lauritzen, S. L. (1996). *Graphical models*, volume 17. Clarendon Press.
- Levin, D. A. and Peres, Y. (2017). *Markov chains and mixing times*, volume 107. American Mathematical Soc.
- Li, S. Z. (2009). *Markov random field modeling in image analysis*. Springer Science & Business Media.
- Littlestone, N. and Warmuth, M. (1986). Relating data compression and learnability. Technical report.
- Marton, K. (1996). A measure concentration inequality for contracting markov chains. *Geometric & Functional Analysis GAFA*, 6(3):556–571.
- McSherry, F. and Talwar, K. (2007). Mechanism design via differential privacy. In *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS’07)*, pages 94–103. IEEE.
- Moran, S. and Yehudayoff, A. (2016). Sample compression schemes for vc classes. *Journal of the ACM (JACM)*, 63(3):1–10.
- Raginsky, M., Rakhlin, A., Tsao, M., Wu, Y., and Xu, A. (2016). Information-theoretic analysis of stability and bias of learning algorithms. In *2016 IEEE Information Theory Workshop, ITW 2016, Cambridge, United Kingdom, September 11-14, 2016*, pages 26–30. IEEE.
- Rogers, R. M., Roth, A., Smith, A. D., and Thakkar, O. (2016). Max-information, differential privacy, and post-selection hypothesis testing. In Dinur, I., editor, *IEEE 57th Annual Symposium on Foundations of Computer Science, FOCS 2016, 9-11 October 2016, Hyatt Regency, New Brunswick, New Jersey, USA*, pages 487–494. IEEE Computer Society.

- Rue, H. and Held, L. (2005). *Gaussian Markov random fields: theory and applications*. CRC press.
- Russo, D. and Zou, J. (2016). Controlling bias in adaptive data analysis using information theory. In Gretton, A. and Robert, C. C., editors, *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, volume 51 of *Proceedings of Machine Learning Research*, pages 1232–1240, Cadiz, Spain. PMLR.
- Russo, D. and Zou, J. (2019). How much does your data exploration overfit? controlling bias via information usage. *IEEE Transactions on Information Theory*, 66(1):302–323.
- Sacerdote, B. (2001). Peer Effects with Random Assignment: Results for Dartmouth Roommates\*. *The Quarterly Journal of Economics*, 116(2):681–704.
- Settles, B. (2009). Active learning literature survey. Technical report, University of Wisconsin-Madison Department of Computer Sciences.
- Shalev-Shwartz, S., Shamir, O., Srebro, N., and Sridharan, K. (2010). Learnability, stability and uniform convergence. *The Journal of Machine Learning Research*, 11:2635–2670.
- Shalev-Shwartz, S., Shamir, O., Srebro, N., and Sridharan, K. (2010). Learnability, stability and uniform convergence. *J. Mach. Learn. Res.*, 11:2635–2670.
- Shenfeld, M. and Ligett, K. (2019). A necessary and sufficient stability notion for adaptive generalization. In *NeurIPS*, pages 11481–11490.
- Shenfeld, M. and Ligett, K. (2021). Generalization in the face of adaptivity: A bayesian perspective. *CoRR*, abs/2106.10761.
- Steinke, T. and Zakyntinou, L. (2020). Reasoning about generalization via conditional mutual information. In Abernethy, J. D. and Agarwal, S., editors, *Conference on Learning Theory, COLT 2020, 9-12 July 2020, Virtual Event [Graz, Austria]*, volume 125 of *Proceedings of Machine Learning Research*, pages 3437–3452. PMLR.
- Ullman, J., Smith, A., Nissim, K., Stemmer, U., and Steinke, T. (2018). The limits of post-selection generalization. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R., editors, *Advances in Neural Information Processing Systems 31*, pages 6400–6409. Curran Associates, Inc.
- Wei, Z. and Li, H. (2007). A Markov random field model for network-based analysis of genomic data. *Bioinformatics*, 23(12):1537–1544.
- Wolfer, G. and Kontorovich, A. (2019). Minimax learning of ergodic markov chains. In Garivier, A. and Kale, S., editors, *Algorithmic Learning Theory, ALT 2019, 22-24 March 2019, Chicago, Illinois, USA*, volume 98 of *Proceedings of Machine Learning Research*, pages 903–929. PMLR.
- Xu, A. and Raginsky, M. (2017). Information-theoretic analysis of generalization capability of learning algorithms. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R., editors, *Advances in Neural Information Processing Systems 30*, pages 2524–2533. Curran Associates, Inc.

## A. Other related works

Algorithmic stability is known to be intimately connected (and, in some settings, equivalent) to learnability (Bousquet and Elisseeff, 2002a; Shalev-Shwartz et al., 2010). Most of the existing stability notions, however, are not sufficient for our goal of adaptive learnability. For example, *uniform stability*, which has recently been the subject of several interesting results, is not closed under post-processing and does not yield the same type of adaptive generalization bounds as we study in this paper. (Bousquet and Elisseeff, 2002b; Shalev-Shwartz et al., 2010; Hardt et al., 2016; Feldman and Vondrák, 2018; 2019) A notable exception is *local statistical stability*, which was shown to be both necessary and sufficient for adaptive generalization (Shenfeld and Ligett, 2019). However, so far, local statistical stability has not yielded new algorithmic insights.

A different line of research employs information-theoretic techniques, whereby overfitting is prevented by bounding the amount of mutual information between the input sample and the output hypothesis. However, these techniques generally only guarantee generalization in expectation, rather than high probability bounds. (Russo and Zou, 2016; Xu and Raginsky, 2017; Rogers et al., 2016; Raginsky et al., 2016; Russo and Zou, 2019; Steinke and Zakyntinou, 2020).

The formulation of the adaptive data analysis we consider was introduced by Dwork et al. (2015b) (in the context of i.i.d. sampling), and has since then been the subject of many interesting papers (Bassily et al., 2016; Bun et al., 2018; Hardt and Ullman, 2014; Ullman et al., 2018; Shenfeld and Ligett, 2019; Jung et al., 2020; Shenfeld and Ligett, 2021). The connection between differential privacy and adaptive generalization also originated from Dwork et al. (2015b). Interestingly, this connection has recently been repurposed for different settings, such as adversarial streaming and dynamic algorithms (Hassidim et al., 2020; Attias et al., 2021; Kaplan et al., 2021; Beimel et al., 2021).

We note that in the case of *non-adaptive* data analysis, learning from non-i.i.d samples is a well-known problem that has been heavily studied in various directions. This includes works on the Markovian criteria (Marton, 1996; Kontorovich and Raginsky, 2017; Wolfer and Kontorovich, 2019; Juang and Rabiner, 1991), as well as other criterias (Daskalakis et al., 2019; Dagan et al., 2019). These lines of work do not transfer, at least not in a way that we are aware of, to the adaptive setting.

## B. Application to Markov Chains

In this section, we demonstrate an application of our tools and results regarding Gibbs-dependency and differential privacy to the problem of learning *Markov chains* adaptively. For our notion of dependence, it will be more convenient to analyze the *Undirected Markov Chains*. By the Hammersley-Clifford theorem (Hammersley and Clifford, 1971; Clifford, 1990), every Markov measure on a chain graph with nonzero transition probabilities can be factorized according to pairwise potential functions (formalized below), which we refer to as the *undirected Markov chain* formalization (Kontorovich, 2012).

The formal definition of an undirected Markov chain measure is as follows.

**Definition B.1.** A measure  $\mu$  over  $\Omega^n$  is an undirected Markov chain if there are positive functions  $\{g_i\}_{i \in [n-1]}$ , called potential functions, such that for any  $x \in \Omega^n$

$$\mu(x) = \frac{\prod_{i=1}^{n-1} g_i(x_i, x_{i+1})}{\sum_{x' \in \Omega^n} \prod_{i=1}^{n-1} g_i(x'_i, x'_{i+1})}.$$

This is a special case of the more general *undirected graphical model* (see Lauritzen (1996)). For the sake of convenience, we will use the following notations.

**Definition B.2.** Let  $\mu$  be an undirected Markov chain with potential functions  $\{g_i\}_{i \in [n-1]}$ . We denote the maximal and minimal potentials as follows.

- $R_i(\mu) = \max_{a,b \in \Omega} g_i(a, b)$ ,
- $r_i(\mu) = \min_{a,b \in \Omega} g_i(a, b)$ ,
- $R(\mu) = \max_i \{R_i(\mu)\}$ ,
- $r(\mu) = \min_i \{r_i(\mu)\}$ ,
- $\bar{R}(\mu) := \frac{R(\mu)^2 - r(\mu)^2}{R(\mu)^2 + r(\mu)^2}$ .

When  $\mu$  is clear from the context, we simply write  $R_i, r_i, R, r, \bar{R}$  instead of  $R_i(\mu), r_i(\mu), R(\mu), r(\mu), \bar{R}(\mu)$ .

In order to apply our techniques to the case where the underlying distribution is an undirected Markov chain, we

need to bound the Gibbs-dependency of undirected Markov chains. We first show the following lemma. (The proof of this lemma is deferred to a later part of this section.)

**Lemma B.3.** For every undirected Markov chain  $\mu$  we have

$$\psi(\mu) \leq \bar{R} := \frac{R^2 - r^2}{R^2 + r^2}.$$

That is, the above lemma bounds the Gibbs-dependency of undirected Markov chains as a function of the potential functions. Combining this bound with Corollary 1.6, we obtain the following result.

**Corollary B.4.** There exists a computationally efficient mechanism for answering  $k$  adaptively chosen queries with the following properties. When given  $n \geq m = \tilde{O}\left(\frac{\sqrt{k}}{\alpha^2} \log \frac{1}{\beta}\right)$  samples (an  $n$ -tuple) from an (unknown) undirected Markov chain  $\mu$ , the mechanism guarantees  $(\alpha + 2\bar{R}(\mu), \beta)$ -statistical-accuracy (w.r.t. the underlying distribution  $\mu$ ).

In particular, Corollary B.4 shows that if the underlying chain  $\mu$  satisfies  $\bar{R}(\mu) \leq \alpha$ , then the dependencies in  $\mu$  can be “accommodated for free”, in the sense that we can efficiently answer the same amount of adaptive queries as if the underlying distribution is a product distribution. We are not aware of an alternative method for answering this amount of adaptive queries under these conditions. As we next explain, we can broaden the applicability of our techniques even further, by reducing dependencies in the data as follows. The idea is to access only a part of the chain, obtained by “skipping” a fixed number of elements between two random samples. Formally,

**Definition B.5** (Skipping Samples). Given a measure  $\mu$  over  $n$ -tuples, and an integer  $t$ , we define the measure  $\mu_{\times t}$  over  $\frac{n}{t}$ -tuples as follows.<sup>1</sup> To sample from  $\mu_{\times t}$ , let  $(x_0, x_1, x_2, x_3, \dots, x_{n-1}) \sim \mu$ , and return  $(x_0, x_t, x_{2t}, x_{3t}, \dots, x_{n-t})$ .

Intuitively, as Markov chains are “memoryless processes”, skipping points in our sample (as in Definition B.5), should significantly reduce dependencies within the remaining points. We formalize this intuition and prove the following theorem. (The proof of this theorem is deferred to a later part of this section.)

**Theorem B.6.** For every undirected Markov chain  $\mu$  and for every  $t$  we have

$$\psi(\mu_{\times t}) \leq \psi(\mu)^t.$$

That is, Theorem B.6 states that by reducing our sample size *linearly* with  $t$ , we could reduce dependencies within our sample *exponentially* in  $t$ . Combining this bound with Corollary 1.6, we obtain the following result.

<sup>1</sup>We assume here for simplicity that  $t$  divides  $n$ .

**Corollary B.7.** *There exists a computationally efficient mechanism that is  $(3\alpha, \beta)$ -statistically-accurate for  $k$  adaptively chosen queries, given a sample (an  $n$ -tuple) drawn from an underlying distribution  $\mu_{\times t}$ , where  $\mu$  is an undirected Markov-chain, and where*

$$n \geq \tilde{\mathcal{O}} \left( \frac{\log(1/\beta)\sqrt{k}}{\alpha^2} \right) \quad \text{and} \quad t \geq \frac{\log(1/\alpha)}{\log(1/\bar{R})}.$$

**Remark B.8.** *As a baseline, one can choose the “skipping parameter”  $t$  to be sufficiently big s.t. the Gibbs-dependency would drop below  $\beta/n$ . As we mentioned in Section 1.1.1, in that case the dependencies in the data would be small enough to the extent we could simply apply existing tools for answering queries w.r.t. product distributions, in order to answer adaptive queries w.r.t.  $\mu_{\times t}$ . However, this would require the skipping parameter  $t$  to be as big as  $\frac{\log(n/\beta)}{\log(1/\bar{R})}$ , i.e., to increase by (roughly) a  $\log(n)$  factor, which in turn, would result in a larger sample complexity.*

We next prove Lemma B.3 and Theorem B.6.

*Proof of Lemma B.3.* For any  $i \in [2, n-1]$ ,<sup>2</sup>  $a \in \Omega$  and  $u, v \in \Omega^n$

$$\mu_i(a \mid v^{-i}) = \frac{g_{i-1}(v_{i-1}, a)g_i(a, v_{i+1})}{\sum_{a'} g_{i-1}(v_{i-1}, a')g_i(a', v_{i+1})}$$

We will be using the following lemma of Kontorovich (2012):

**Lemma B.9.** *For  $n \in \mathbb{N}$  and  $0 \leq r \leq R$ , consider the vectors  $\alpha \in [0, \infty)^n$  and  $f, g \in [r, R]^n$ . Then,*

$$\frac{1}{2} \sum_{i=1}^n \left| \frac{\alpha_i f_i}{\sum_{j=1}^n \alpha_j f_j} - \frac{\alpha_i g_i}{\sum_{j=1}^n \alpha_j g_j} \right| \leq \frac{R-r}{R+r}.$$

We apply the lemma using

- $f_a = g_{i-1}(v_{i-1}, a)g_i(a, v_{i+1})$
- $h_a = g_{i-1}(u_{i-1}, a)g_i(a, u_{i+1})$
- $\alpha_a = 1$

and get that

$$\frac{1}{2} \sum_a |\mu_i(a \mid u^{-i}) - \mu_i(a \mid v^{-i})| \leq \frac{R_{i-1}R_i - r_{i-1}r_i}{R_{i-1}R_i + r_{i-1}r_i}.$$

<sup>2</sup>The case of  $i \in \{1, n\}$  has an almost identical argument; only the  $g_{i-1}(v_{i-1}, a)$  (respectively,  $g_n(a, v_{i+1})$ ) factor is omitted. This does not affect the rest of the argument for the upper bound.

It follows that

$$\begin{aligned} & \|\mu_i(\cdot) - \mu_i(\cdot \mid v^{-i})\|_{\text{TV}} \\ &= \frac{1}{2} \sum_a |\mu_i(a) - \mu_i(a \mid v^{-i})| \\ &= \frac{1}{2} \sum_a \left| \sum_{u^{-i}} \mu_i(a \mid u^{-i}) \mu^{-i}(u^{-i}) - \mu_i(a \mid v^{-i}) \right| \\ &= \frac{1}{2} \sum_a \left| \sum_{u^{-i}} \mu_i(a \mid u^{-i}) \mu^{-i}(u^{-i}) \right. \\ & \quad \left. - \sum_{u^{-i}} \mu^{-i}(u^{-i}) \mu_i(a \mid v^{-i}) \right| \\ &= \frac{1}{2} \sum_a \left| \sum_{u^{-i}} \mu^{-i}(u^{-i}) [\mu_i(a \mid u^{-i}) - \mu_i(a \mid v^{-i})] \right| \\ &\leq \frac{1}{2} \sum_a \sum_{u^{-i}} \mu^{-i}(u^{-i}) |\mu_i(a \mid u^{-i}) - \mu_i(a \mid v^{-i})| \\ &= \sum_{u^{-i}} \mu^{-i}(u^{-i}) \frac{1}{2} \sum_a |\mu_i(a \mid u^{-i}) - \mu_i(a \mid v^{-i})| \\ &\leq \sum_{u^{-i}} \mu^{-i}(u^{-i}) \frac{R_{i-1}R_i - r_{i-1}r_i}{R_{i-1}R_i + r_{i-1}r_i} = \frac{R_{i-1}R_i - r_{i-1}r_i}{R_{i-1}R_i + r_{i-1}r_i}. \end{aligned}$$

Finally,

$$\begin{aligned} \psi(\mu) &= \sup_v \mathbb{E}_i \|\mu_i(\cdot) - \mu_i(\cdot \mid v^{-i})\|_{\text{TV}} \\ &\leq \frac{R^2 - r^2}{R^2 + r^2}. \end{aligned}$$

□

In order to prove Theorem B.6, we first establish the following notations:

- $x_{i\pm t} = x_{i-1}, x_{i+t}$
- $c_i^t := \sup_{x_{i\pm t}, y_{i\pm t}} \|\mu_{i\pm(t-1)}(\cdot \mid x_{i\pm t}) - \mu_{i\pm(t-1)}(\cdot \mid y_{i\pm t})\|_{\text{TV}}$
- $\gamma_i^t := \sup_{x_{i\pm t}, y_{i\pm t}} \|\mu_i(\cdot \mid x_{i\pm t}) - \mu_i(\cdot \mid y_{i\pm t})\|_{\text{TV}}$

Note that

$$c_i^1 = \gamma_i^{i+1} = \sup_{x_{i\pm 1}, y_{i\pm 1}} \|\mu_i(\cdot \mid x_{i\pm 1}) - \mu_i(\cdot \mid y_{i\pm 1})\|_{\text{TV}}.$$

We will be using the following two lemmas (we prove these two lemmas after the proof of Theorem B.6).

**Lemma B.10.**  $\gamma_i^t \leq \prod_{j=1}^t c_i^j$

**Lemma B.11.** *For every  $t$  and every  $i$ , there exist some  $j$  s.t.  $c_i^t \leq c_j^1$ .*

We now prove Theorem B.6 using Lemmas B.10 and B.11.

*Proof of Theorem B.6.* Combining Lemma B.10 and Lemma B.11 yields that for every undirected Markov measure  $\mu$  and for every  $t$ ,

$$\begin{aligned} \max_i \gamma_i^t &\leq \max_i \prod_{j=1}^t c_i^j \leq \max_i \prod_{j=1}^t c_{l(j)}^1 \\ &\leq \max_i \max_l (c_l^1)^t = (\max_i c_l^1)^t = (\psi(\mu))^t, \end{aligned} \quad (3)$$

where the first inequality is due to Lemma B.10, the second is by Lemma B.11<sup>3</sup>. The Last equality holds by the definitions of  $\psi$  and  $c_l^1$ . Since

$$\mu_i(\cdot) = \sum_{x_{i\pm t} \in \Omega^2} \mu_i(\cdot | x_{i\pm t}) \mu_{i\pm t}(x_{i\pm t}),$$

we have, by the undirected Markov property,

$$\begin{aligned} \psi(\mu_{\times t}) &= \max_i \sup_{y_{i\pm t}} \|\mu_i(\cdot) - \mu_i(\cdot | y_{i\pm t})\|_{\text{TV}} \\ &= \max_i \sup_{y_{i\pm t}} \left\| \sum_{x_{i\pm t} \in \Omega^2} (\mu_i(\cdot | x_{i\pm t}) - \mu_i(\cdot | y_{i\pm t})) \mu(x_{i\pm t}) \right\|_{\text{TV}} \\ &\leq \max_i \sup_{y_{i\pm t}} \sum_{x_{i\pm t} \in \Omega^2} \|\mu_i(\cdot | x_{i\pm t}) - \mu_i(\cdot | y_{i\pm t})\|_{\text{TV}} \mu(x_{i\pm t}) \\ &\leq \max_i \gamma_i^t \leq (\psi(\mu))^t, \end{aligned}$$

where the last inequality is due to (3).  $\square$

*Proof of Lemma B.10.* Let  $x_{i\pm t}, y_{i\pm t}$  be some pairs of realization for the  $i-t, i+t$  variable in the chain. By the law of total probability,

$$\begin{aligned} &\|\mu_i(\cdot | x_{i\pm t}) - \mu_i(\cdot | y_{i\pm t})\|_{\text{TV}} \\ &= \left\| \sum_{x_{i\pm(t-1)}} \mu_i(\cdot | x_{i\pm(t-1)}) \mu_{i\pm(t-1)}(x_{i\pm(t-1)} | x_{i\pm t}) \right. \\ &\quad \left. - \sum_{y_{i\pm(t-1)}} \mu_i(\cdot | y_{i\pm(t-1)}) \mu_{i\pm(t-1)}(y_{i\pm(t-1)} | y_{i\pm t}) \right\|_{\text{TV}}. \end{aligned} \quad (4)$$

Define a coupling measure  $\Pi_{i\pm(t-1)}(\cdot, \cdot | x_{i\pm t}, y_{i\pm t})$  whose marginals are  $\mu_{i\pm(t-1)}(\cdot | x_{i\pm t})$  and  $\mu_{i\pm(t-1)}(\cdot | y_{i\pm t})$ .

<sup>3</sup>The function  $l : [n] \rightarrow [n]$  returns for every coordinate  $i$  the appropriate coordinate  $l(i)$  which is guaranteed by Lemma B.11 to bound it.

$y_{i\pm t}$ ). Then

$$\begin{aligned} (4) &= \left\| \sum_{x_{i\pm(t-1)}} \sum_{y_{i\pm(t-1)}} (\mu_i(\cdot | x_{i\pm(t-1)}) - \mu_i(\cdot | y_{i\pm(t-1)})) \right. \\ &\quad \left. \Pi_{i\pm(t-1)}(x_{i\pm(t-1)}, y_{i\pm(t-1)} | x_{i\pm t}, y_{i\pm t}) \right\|_{\text{TV}} \\ &\leq \sum_{x_{i\pm(t-1)}} \sum_{y_{i\pm(t-1)}} \|\mu_i(\cdot | x_{i\pm(t-1)}) - \mu_i(\cdot | y_{i\pm(t-1)})\|_{\text{TV}} \\ &\quad \Pi_{i\pm(t-1)}(x_{i\pm(t-1)}, y_{i\pm(t-1)} | x_{i\pm t}, y_{i\pm t}) \\ &\leq \gamma_i^{t-1} \sum_{x_{i\pm(t-1)}} \sum_{y_{i\pm(t-1)}} \mathbf{1}_{x_{i\pm(t-1)} \neq y_{i\pm(t-1)}} \\ &\quad \Pi_{i\pm(t-1)}(x_{i\pm(t-1)}, y_{i\pm(t-1)} | x_{i\pm t}, y_{i\pm t}). \end{aligned}$$

By the dual form of the total variation distance,<sup>4</sup> we can choose  $\Pi_{i\pm(t-1)}$  to be such that

$$\begin{aligned} &\left\| \mu_{i\pm(t-1)}(\cdot | x_{i\pm t}) - \mu_{i\pm(t-1)}(\cdot | y_{i\pm t}) \right\|_{\text{TV}} \\ &= \sum_{x_{i\pm(t-1)}} \sum_{y_{i\pm(t-1)}} \mathbf{1}_{x_{i\pm(t-1)} \neq y_{i\pm(t-1)}} \\ &\quad \Pi_{i\pm(t-1)}(x_{i\pm(t-1)}, y_{i\pm(t-1)} | x_{i\pm t}, y_{i\pm t}) \end{aligned}$$

and therefore

$$\begin{aligned} &\|\mu_i(\cdot | x_{i\pm t}) - \mu_i(\cdot | y_{i\pm t})\|_{\text{TV}} \\ &\leq \gamma_i^{t-1} \|\mu_{i\pm(t-1)}(\cdot | x_{i\pm t}) - \mu_{i\pm(t-1)}(\cdot | y_{i\pm t})\|_{\text{TV}} \\ &\leq \gamma_i^{t-1} c_i^t. \end{aligned}$$

Hence we get that

$$\gamma_i^t = \sup_{x_{i\pm t}, y_{i\pm t}} \|\mu_i(\cdot | x_{i\pm t}) - \mu_i(\cdot | y_{i\pm t})\|_{\text{TV}} \leq \gamma_i^{t-1} c_i^t$$

and by induction we get the lemma's result.  $\square$

*Proof of Lemma B.11.* First we will show that for any  $j, k$  the following holds

$$\begin{aligned} &\sup \|\mu_j(\cdot | x_{j-1}, x_{j+k}) - \mu_j(\cdot | y_{j-1}, y_{j+k})\|_{\text{TV}} \\ &\leq \sup \|\mu_j(\cdot | x_{j-1}, x_{j+k-1}) - \mu_j(\cdot | y_{j-1}, y_{j+k-1})\|_{\text{TV}}. \end{aligned} \quad (5)$$

Indeed,

$$\begin{aligned} &\sup \|\mu_j(\cdot | x_{j-1}, x_{j+k}) - \mu_j(\cdot | y_{j-1}, y_{j+k})\|_{\text{TV}} \\ &= \sup \left\| \sum_{x_{j+k-1}} \mu_j(\cdot | x_{j-1}, x_{j+k-1}) \mu_{j+k-1}(x_{j+k-1} | x_{j+k}) \right. \\ &\quad \left. - \sum_{y_{j+k-1}} \mu_j(\cdot | y_{j-1}, y_{j+k-1}) \mu_{j+k-1}(y_{j+k-1} | y_{j+k}) \right\|_{\text{TV}}. \end{aligned}$$

<sup>4</sup>By the Kantorovich-Rubinstein duality of the specific case of total-Variation distance  $\|P - Q\|_{\text{TV}} = \min_{\Pi \in \Delta(P, Q)} \int_{\Omega} \int_{\Omega} \mathbf{1}_{x \neq y} d\Pi(x, y)$  when  $\Delta(P, Q)$  is the set of all the possible coupling of  $P$  and  $Q$ .

Let  $\Pi_{j+k-1}(\cdot, \cdot \mid x_{j+k}, y_{j+k})$  be a coupling distribution whose marginal distributions are  $\mu_{j+k-1}(y_{j+k-1} \mid y_{j+k})$  and  $\mu_{j+k-1}(x_{j+k-1} \mid x_{j+k})$ , we get that the above is equal to

$$\begin{aligned} & \sup \left\| \sum_{x_{j+k-1}} \sum_{y_{j+k-1}} (\mu_j(\cdot \mid x_{j-1}, x_{j+k-1}) \right. \\ & \quad \left. - \mu_j(\cdot \mid y_{j-1}, y_{j+k-1})) \right. \\ & \quad \left. \Pi_{j+k-1}(x_{j+k-1}, y_{j+k-1} \mid x_{j+k}, y_{j+k}) \right\|_{\text{TV}} \\ & \leq \sup \sum_{x_{j+k-1}} \sum_{y_{j+k-1}} \left\| \mu_j(\cdot \mid x_{j-1}, x_{j+k-1}) \right. \\ & \quad \left. - \mu_j(\cdot \mid y_{j-1}, y_{j+k-1}) \right\|_{\text{TV}} \\ & \quad \Pi_{j+k-1}(x_{j+k-1}, y_{j+k-1} \mid x_{j+k}, y_{j+k}) \\ & \leq \sup \left\| \mu_j(\cdot \mid x_{j-1}, x_{j+k-1}) - \mu_j(\cdot \mid y_{j-1}, y_{j+k-1}) \right\|_{\text{TV}}. \end{aligned}$$

Now we turn to the quantity of interest:

$$\begin{aligned} & \sup_{x_{i\pm t}, y_{i\pm t}} \left\| \mu_{i\pm(t-1)}(\cdot \mid x_{i\pm t}) - \mu_{i\pm(t-1)}(\cdot \mid y_{i\pm t}) \right\|_{\text{TV}} \\ & = \sup_{x_{i\pm t}, y_{i\pm t}} \left\| \sum_{x_{i+t-1}} \mu_{i\pm(t-1)}(\cdot \mid x_{i\pm t}, x_{i+t-1}) \right. \\ & \quad \left. \mu_{i+t-1}(x_{i+t-1} \mid x_{i\pm t}) \right. \\ & \quad \left. - \sum_{y_{i+t-1}} \mu_{i\pm(t-1)}(\cdot \mid y_{i\pm t}, y_{i+t-1}) \right. \\ & \quad \left. \mu_{i+t-1}(y_{i+t-1} \mid y_{i\pm t}) \right\|_{\text{TV}} \\ & = \sup_{x_{i\pm t}, y_{i\pm t}} \left\| \sum_{x_{i+t-1}} \mu_{i-t+1}(\cdot \mid x_{i-t}, x_{i+t-1}) \right. \\ & \quad \left. \mu_{i+t-1}(x_{i+t-1} \mid x_{i\pm t}) \right. \\ & \quad \left. - \sum_{y_{i+t-1}} \mu_{i-t+1}(\cdot \mid y_{i-t}, y_{i+t-1}) \right. \\ & \quad \left. \mu_{i+t-1}(y_{i+t-1} \mid y_{i\pm t}) \right\|_{\text{TV}}. \quad (6) \end{aligned}$$

Let  $\Pi_{i+t-1}(\cdot, \cdot \mid x_{i\pm t}, y_{i\pm t})$  be a coupling distribution whose marginals are  $\mu_{i+t-1}(x_{i+t-1} \mid x_{i\pm t})$  and  $\mu_{i+t-1}(y_{i+t-1} \mid y_{i\pm t})$ . Then the above is then equal to

$$\begin{aligned} (6) & = \sup_{x_{i\pm t}, y_{i\pm t}} \left\| \sum_{x_{i+t-1}} \sum_{y_{i+t-1}} \mu_{i-t+1}(\cdot \mid x_{i-t}, x_{i+t-1}) \right. \\ & \quad \left. - \mu_{i-t+1}(\cdot \mid y_{i-t}, y_{i+t-1}) \right. \\ & \quad \left. \Pi_{i+t-1}(x_{i+t-1}, y_{i+t-1} \mid x_{i\pm t}, y_{i\pm t}) \right\|_{\text{TV}} \\ & \leq \sup_{x_{i\pm t}, y_{i\pm t}} \sum_{x_{i+t-1}} \sum_{y_{i+t-1}} \left\| \mu_{i-t+1}(\cdot \mid x_{i-t}, x_{i+t-1}) \right. \\ & \quad \left. - \mu_{i-t+1}(\cdot \mid y_{i-t}, y_{i+t-1}) \right\|_{\text{TV}} \\ & \quad \Pi_{i+t-1}(x_{i+t-1}, y_{i+t-1} \mid x_{i\pm t}, y_{i\pm t}). \end{aligned}$$

Plugging  $j = i - t + 1$  and  $k = t - 2$  into (5) yields

$$\begin{aligned} & \sup_{x_{i\pm t}, y_{i\pm t}} \left\| \mu_{i\pm(t-1)}(\cdot \mid x_{i\pm t}) - \mu_{i\pm(t-1)}(\cdot \mid y_{i\pm t}) \right\|_{\text{TV}} \\ & \leq \sup_{x_{i-t, i-t+2}, y_{i-t, i-t+2}} \left\| \mu_{i-t+1}(\cdot \mid x_{i-t, i-t+2}) \right. \\ & \quad \left. - \mu_{i-t+1}(\cdot \mid y_{i-t, i-t+2}) \right\|_{\text{TV}} \end{aligned}$$

---

### Algorithm 3 Auxiliary Algorithm $\mathcal{A}'$

---

**Input:**  $\vec{S} = (S_1, \dots, S_T)$ , where  $T = \frac{\varepsilon}{\delta}$ .  
 $F \leftarrow \emptyset$

**for**  $t \in [T]$  **do**

$(h_1^t, \dots, h_k^t) \leftarrow \mathcal{A}(S_t)$   
 $H_t \leftarrow \{(h_1^t, t), \dots, (h_k^t, t)\}$   
 $\bar{H}_t \leftarrow \{1 - h \mid h \in H_t\}$   
 $F \leftarrow F \cup H_t \cup \bar{H}_t$

**end for**

Sample  $(h^*, t^*)$  from  $F$  using the exponential mechanism. Specifically, sample  $(h^*, t^*) \in F$  with probability proportional to  $\exp\left(\frac{\varepsilon n}{2}(h^*(S_{t^*}) - h^*(\mu))\right)$ .  
 Return  $(h^*, t^*)$

---

which completes the proof.  $\square$

## C. Missing proofs

### C.1. Product measure

We show the following claim

**Claim C.1.** For a measure  $\mu \sim \mathcal{X}^n$ , if for every  $i \in [n]$  and for every possible  $x \in \mathcal{X}^n$  it holds that  $\mu_i = \mu_i(\cdot \mid x^{-i})$ , then  $\mu$  is a product measure.

*Proof.* For convenience, we denote for every  $i \leq j$  the following notation  $a_{i:j} = a_i, \dots, a_j$ . Now, for every  $a \in \mathcal{X}^n$ ,  $\mu(a) = \prod_{i \in [n]} \mu(a_i \mid a_{1:i-1})$ . For start, we show that  $\mu(a_2 \mid a_1) = \mu(a_2)$ . Indeed,

$$\begin{aligned} \mu(a_2 \mid a_1) & = \sum_{a_3, \dots, a_n} \mu(a_2 \mid a_1, a_{3:n}) \cdot \mu(a_{3:n} \mid a_1) \\ & = \sum_{a_3, \dots, a_n} \mu(a_2) \cdot \mu(a_{3:n} \mid a_1) = \mu(a_2) \end{aligned}$$

In the same way it can be shown that  $\mu(a_3) = \mu(a_3 \mid a_{1:2})$  and so on.  $\square$

### C.2. Proofs from Section 3

*Proof of Theorem 3.2.* Fix a measure  $\mu$  on  $\mathcal{X}^n$  with Gibbs-dependence  $\psi_n$ , and fix an  $(\varepsilon, \delta)$ -differentially private algorithm that takes a sample  $S \in \mathcal{X}^n$  and returns  $k$  predicates  $h_1, \dots, h_k : \mathcal{X} \rightarrow \{0, 1\}$ . Assume towards contradiction that

$$\Pr_{S, \mathcal{A}(S)} \left[ \max_{i \in [k]} |h_i(\mu) - h_i(S)| \geq 10\varepsilon + 2\psi \right] \geq \frac{\delta}{\varepsilon}. \quad (7)$$

Consider the procedure described in Algorithm 3. As differential private algorithms are immune to post-processing

and by the composition theorem,  $\mathcal{A}'$  is by itself  $(2\varepsilon, \delta)$ -differentially private. Given a multi-set  $\vec{S}$  sampled from  $\mu^T$ , by (7) we get that

$$\forall t : \Pr_{S_t, \mathcal{A}(S_t)} \left[ \max_{i \in [k]} |h_i^t(\mu) - h_i^t(S_t)| \geq 10\varepsilon + 2\psi \right] \geq \frac{\delta}{\varepsilon},$$

and hence, by setting  $T = \frac{\varepsilon}{\delta}$ , we have that

$$\begin{aligned} \Pr_{\vec{S}, \mathcal{A}'(\vec{S})} \left[ \max_{t \in [T], i \in [k]} |h_i^t(\mu) - h_i^t(S_t)| \geq 10\varepsilon + 2\psi \right] \\ \geq 1 - \left(1 - \frac{\delta}{\varepsilon}\right)^T \geq \frac{1}{2}. \end{aligned}$$

By Markov's inequality,

$$\mathbb{E}_{\vec{S}, \mathcal{A}'(\vec{S})} \left[ \max_{t \in [T], i \in [k]} |h_i^t(\mu) - h_i^t(S_t)| \right] \geq 5\varepsilon + \psi.$$

Now the set constructed in the algorithm's run,  $F$ , contains also the negation of each predicate, and hence

$$\begin{aligned} \mathbb{E}_{\vec{S}, \mathcal{A}'(\vec{S})} \left[ \max_{(h,t) \in F} \{h(S_t) - h(\mu)\} \right] \\ = \mathbb{E}_{\vec{S}, \mathcal{A}'(\vec{S})} \left[ \max_{t \in [T], i \in [k]} |h_i^t(\mu) - h_i^t(S_t)| \right] \geq 5\varepsilon + \psi. \end{aligned}$$

By the properties of the exponential mechanism (see [McSherry and Talwar \(2007\)](#) or [Bassily et al. \(2016\)](#)), denoting the output of the algorithm by  $(h^*, t^*)$  we get that

$$\begin{aligned} \mathbb{E}_{(h^*, t^*)} [h^*(S_{t^*}) - h^*(\mu)] \\ \geq \max_{(h,t) \in F} \{h^*(S_{t^*}) - h^*(\mu)\} - \frac{2}{\varepsilon n} \log(2Tk). \end{aligned}$$

Taking expectation on both sides yields

$$\begin{aligned} \mathbb{E}_{\vec{S}, \mathcal{A}'(\vec{S})} [h^*(S_{t^*}) - h^*(\mu)] \\ \geq \mathbb{E}_{\vec{S}, \mathcal{A}'(\vec{S})} \left[ \max_{(h,t) \in F} \{h^*(S_{t^*}) - h^*(\mu)\} \right] - \frac{2}{\varepsilon n} \log(2Tk) \\ \geq 5\varepsilon + \psi - \frac{2}{\varepsilon n} \log(2k\varepsilon/\delta). \end{aligned}$$

For  $n \geq \frac{\log(2k\varepsilon/\delta)}{\varepsilon^2}$ , this is at least  $2\varepsilon + \psi$  which contradicts [Lemma 3.1](#).  $\square$