
Generalization in the Face of Adaptivity: A Bayesian Perspective

Moshe Shenfeld

Department of Computer Science
The Hebrew University of Jerusalem
moshe.shenfeld@mail.huji.ac.il

Katrina Ligett

Department of Computer Science
Federmann Center for the Study of Rationality
The Hebrew University of Jerusalem
katrina@cs.huji.ac.il

Abstract

Repeated use of a data sample via adaptively chosen queries can rapidly lead to overfitting, wherein the empirical evaluation of queries on the sample significantly deviates from their mean with respect to the underlying data distribution. It turns out that simple noise addition algorithms suffice to prevent this issue, and differential privacy-based analysis of these algorithms shows that they can handle an asymptotically optimal number of queries. However, differential privacy’s worst-case nature entails scaling such noise to the range of the queries even for highly-concentrated queries, or introducing more complex algorithms.

In this paper, we prove that straightforward noise-addition algorithms already provide variance-dependent guarantees that also extend to unbounded queries. This improvement stems from a novel characterization that illuminates the core problem of adaptive data analysis. We show that the harm of adaptivity results from the covariance between the new query and a Bayes factor-based measure of how much information about the data sample was encoded in the responses given to past queries. We then leverage this characterization to introduce a new data-dependent stability notion that can bound this covariance.

1 Introduction

Recent years have seen growing recognition of the role of *adaptivity* in causing overfitting and thereby reducing the accuracy of the conclusions drawn from data. Intuitively, allowing a data analyst to adaptively choose the queries that she issues potentially leads to misleading conclusions, because the results of prior queries might encode information that is specific to the available data sample and not relevant to the underlying distribution from which it was drawn. As a result, future queries might then be chosen to leverage these sample-specific properties, giving answers on the sample that differ wildly from the values of those queries on the data distribution. Such adaptivity has been blamed, in part, for the current reproducibility crisis in the data-driven sciences [Ioannidis, 2005, Gelman and Loken, 2014].

A series of works catalyzed by Dwork et al. [2015] recently established a formal framework for understanding and analyzing adaptivity in data analysis, and introduced a general toolkit for provably preventing the harms of choosing queries adaptively—that is, as a function of the results of previous queries. This line of work has established that enforcing that computations obey a constraint of *differential privacy* [Dwork et al., 2006]—which one can interpret as a robustness or stability requirement on computations—provably limits the extent to which adaptivity can cause overfitting. Practically, for simple numerical computations, these results translate into adding a level of noise to the query results that would be sufficient to ensure *worst-case* stability of the *worst-case* query on the *worst-case* dataset. In particular, when analyzed using differential privacy, simple noise-addition

mechanisms must add noise that scales with the *worst-case* change in query value that could be induced by changing a single input data element, in order to protect against adaptivity.

However, this can be overkill: for the purposes of statistical validity, we do not care about the worst case but the typical case. In the present work, we prove the following new and better generalization guarantees for simple Gaussian noise-addition algorithms:

Theorem 1.1 (Informal versions of main theorems). *With probability $> 1 - \delta$, the error of the responses produced by a mechanism which only adds Gaussian noise to the empirical values of the queries it receives is bounded by ϵ , even after responding to k adaptively chosen queries, if*

- *the range of the queries is bounded by Δ , their variance is bounded by σ^2 , and the size of the dataset $n = \Omega\left(\max\left\{\frac{\Delta}{\epsilon}, \frac{\sigma^2}{\epsilon^2}\right\} \sqrt{k} \cdot \ln\left(\frac{k\sigma}{\delta\epsilon}\right)\right)$ (Theorem 5.1), or*
- *the queries are σ^2 -sub-Gaussian and the size of the dataset $n = \Omega\left(\frac{\sigma^2\sqrt{k}}{\epsilon^2} \ln\left(\frac{k\sigma}{\delta\epsilon}\right)\right)$ (Theorem 5.3).*

In contrast, analyzing the same mechanism using differential privacy yields a sample size requirement that scales with $n = \Omega\left(\frac{\Delta^2\sqrt{k}}{\epsilon^2} \ln\left(\frac{k\sigma}{\delta\epsilon}\right)\right)$ in the first setting, and provides no guarantee in the latter.

We prove these theorems via a new notion, *pairwise concentration (PC)* (Definition 4.2), which captures the extent to which replacing one dataset by another would be “noticeable,” given a particular query-response sequence. This is thus a function of particular differing datasets (instead of worst-case over elements), and it also depends on the actual issued queries. We then build a composition toolkit (Theorem 4.4) that allows us to track PC losses over multiple computations. The PC notion allows for more careful analysis of the information encoded by the query-response sequence than differential privacy does.

In order to leverage this more careful analysis of the information encoded in query-response sequences, we rely on a simple new characterization (Lemma 3.5) that shows explicitly that the harms of adaptivity come from the covariance between the behavior of the future queries and a Bayes factor-based measure of how much information about the data sample was encoded in the responses to the issued queries. Our characterization gives insight into how differential privacy protects against adaptivity, and then allows us to step away from this worst-case approach. What differential privacy accomplishes is that it bounds a worst-case version of the Bayes factor; however, as this new characterization makes clear, it is sufficient to bound the “typical” Bayes factor in order to avoid overfitting.

We measure the harm that past adaptivity causes to a future query by considering the query as evaluated on a posterior data distribution and comparing this with its value on a prior. The prior is the true data distribution, and the posterior is induced by observing the responses to past queries and updating the prior. If the new query behaves similarly on the prior distribution as it does on this posterior (a guarantee we call *Bayes stability*; Definition 3.3), adaptivity has not led us too far astray.¹ If furthermore, the response given by the mechanism is close to the query result on the posterior, then by a triangle inequality argument, that mechanism is distribution accurate. This type of triangle inequality first appeared as an analysis technique in Jung et al. [2020].

The dependence of our PC notion on the actual adaptively chosen queries places it in the so-called *fully-adaptive* setting [Rogers et al., 2016, Whitehouse et al., 2023], which requires a fairly subtle analysis involving a set of tools and concepts that may be of independent interest. In particular, we establish a series of “dissimilarity” notions in Appendix B, which generalize the notion of divergence, replacing the scalar bound with a function. Our main stability notion (Definition 4.2) can be viewed as an instance-tailored variant of *zero-concentrated differential privacy* [Bun and Steinke, 2016], and we also make use of a similar extension of the classical max-divergence-based differential privacy definition (B.8).

Related work *Differential privacy* [Dwork et al., 2006] is a privacy notion based on a bound on the max divergence between the output distributions induced by any two neighboring input

¹This can be viewed as a generalization of the *Hypothesis Stability* notion of Bousquet and Elisseeff [2002]—which was proven to guarantee on-average generalization [Shalev-Shwartz et al., 2010]—where the hypothesis is a post-processing of the responses to past queries, and the future query is the loss function estimation.

datasets (datasets which differ in one element). One natural way to enforce differential privacy is by directly adding noise to the results of a numeric-valued query, where the noise is calibrated to the *global sensitivity* of the function to be computed—the maximal change in its value between any two neighboring datasets. Dwork et al. [2015] and Bassily et al. [2021] showed that differential privacy is also useful as tool for ensuring generalization in settings where the queries are chosen adaptively.

Differential privacy essentially provides the optimal asymptotic generalization guarantees given adaptive queries [Hardt and Ullman, 2014, Steinke and Ullman, 2015]. However, its optimality is for worst-case adaptive queries, and the guarantees that it offers only beat the naive intervention—of splitting a dataset so that each query gets fresh data—when the input dataset is quite huge [Jung et al., 2020]. A worst-case approach makes sense for privacy, but for statistical guarantees like generalization, we only need statements that hold with high probability with respect to the sampled dataset, and only on the actual queries issued.

One cluster of works that steps away from this worst-case perspective focuses on giving privacy guarantees that are tailored to the dataset at hand [Nissim et al., 2007, Ghosh and Roth, 2011, Ebadi et al., 2015, Wang, 2019]. In Feldman and Zrnic [2021] in particular, the authors elegantly manage to track the individual privacy loss of the elements in the dataset. However, their results do not enjoy a dependence on the standard deviation in place of the range of the queries. Several truncation-based specialized mechanisms have been proposed, both to provide differential privacy guarantees for Gaussian and sub-Gaussian queries even in the case of multivariate distribution with unknown covariance [Karwa and Vadhan, 2018, Ashtiani and Liaw, 2022, Duchi et al., 2023] and, remarkably, design specialized algorithms that achieve adaptive data analysis guarantees that scale like the standard deviation of the queries [Feldman and Steinke, 2017]. Recently, Blanc [2023] proved that randomized rounding followed by sub-sampling provides accuracy guarantees that scale with the queries’ variance. But none of these results apply to simple noise addition mechanisms.

Another line of work (e.g., Gehrke et al. [2012], Bassily et al. [2013], Bhaskar et al. [2011]) proposes relaxed privacy definitions that leverage the natural noise introduced by dataset sampling to achieve more average-case notions of privacy. This builds on intuition that average-case privacy can be viewed from a Bayesian perspective, by restricting some distance measure between some prior distribution and some posterior distribution induced by the mechanism’s behavior [Dwork et al., 2006, Kasiviswanathan and Smith, 2014]. This perspective was used Shenfeld and Ligett [2019] to propose a stability notion which is both necessary and sufficient for adaptive generalization under several assumptions. Unfortunately, these definitions have at best extremely limited adaptive composition guarantees. Bassily and Freund [2016] connect this Bayesian intuition to statistical validity via *typical stability*, an approach that discards “unlikely” databases that do not obey a differential privacy guarantee, but their results require a sample size that grows linearly with the number of queries even for iid distributions. Triastcyn and Faltings [2020] propose the notion of *Bayesian differential privacy* which leverages the underlying distribution to improve generalization guarantees, but their results still scale with the range in the general case.

An alternative route for avoiding the dependence on worst case queries and datasets was achieved using expectation based stability notions such as *mutual information* and *KL stability* Russo and Zou [2016], Bassily et al. [2021], Steinke and Zakyntinou [2020]. Using these methods Feldman and Steinke [2018] presented a natural noise addition mechanism, which adds noise that scales with the empirical variance when responding to queries with known range and unknown variance. Unfortunately, in the general case, the accuracy guarantees provided by these methods hold only for the expected error rather than with high probability.

A detailed comparison to other lines of work can be found in Appendix F.

2 Preliminaries

2.1 Setting

We study datasets, each of fixed size $n \in \mathbb{N}$, whose elements are drawn from some domain \mathcal{X} . We assume there exists some distribution $D_{\mathcal{X}^n}$ defined over the datasets $s \in \mathcal{X}^n$.² We consider a family

²Throughout the paper, the notation $D(\cdot)$ describing a distribution over a domain, will either represent the probability mass function defined by a discrete probability over a countable sample space or the probability density function (the Radon–Nikodym derivative) in case of a measure over measurable spaces.

of functions (*queries*) \mathcal{Q} of the form $q : \mathcal{X}^n \rightarrow \mathcal{R} \subseteq \mathbb{R}$. We refer to the functions’ outputs as *responses*.³

We denote by $q(D_{\mathcal{X}^n}) := \mathbb{E}_{s \sim D_{\mathcal{X}^n}} [q(S)]$ the mean of q with respect to the distribution $D_{\mathcal{X}^n}$, and think of it as the true value of the query q , which we wish to estimate. As is the case in many machine learning or data-driven science settings, we do not have direct access to the distribution over datasets, but instead receive a dataset s sampled from $D_{\mathcal{X}^n}$, which can be used to compute $q(s)$ —the empirical value of the query.

A *mechanism* M is a (possibly non-deterministic) function $M : \mathcal{X}^n \times \mathcal{Q} \times \Theta \rightarrow \mathcal{R}$ which, given a dataset $s \in \mathcal{X}^n$, a query $q \in \mathcal{Q}$, and some auxiliary parameters $\theta \in \Theta$, provides some response $r \in \mathcal{R}$ (we omit θ when not in use). The most trivial mechanism would simply return $r = q(s)$, but we will consider mechanisms that return a noisy version of $q(s)$. A mechanism induces a distribution over the responses $D_{\mathcal{R}|\mathcal{X}^n}^{q,\theta}(\cdot | s)$ —indicating the probability that each response r will be the output of the mechanism, given s , q , and θ as input. Combining this with $D_{\mathcal{X}^n}$ induces the marginal distribution $D_{\mathcal{R}}^{q,\theta}$ over the responses, and the conditional distribution $D_{\mathcal{X}^n|\mathcal{R}}^{q,\theta}(\cdot | r)$ over sample sets. Considering the uniform distribution over elements in the sample set induces similar distributions over the elements $D_{\mathcal{X}}$ and $D^r := D_{\mathcal{X}|\mathcal{R}}^{q,\theta}(\cdot | r)$. This last distribution, D^r , is central to our analysis, and it represents the *posterior distribution* over the sample elements given an observed response—a Bayesian update with respect to a prior of $D_{\mathcal{X}}$. All of the distributions discussed in this subsection are more formally defined in Appendix A.

An *analyst* A is a (possibly non-deterministic) function $A : \mathcal{R}^* \rightarrow (\mathcal{Q}, \Theta)$ which, given a sequence of responses, provides a query and parameters to be asked next. An interaction between a mechanism and an analyst—wherein repeatedly (and potentially adaptively) the analyst generates a query, then the mechanism generates a response for the analyst to observe—generates a sequence of k queries and responses for some $k \in \mathbb{N}$, which together we refer to as a *view* $v_k = (r_1, r_2, \dots, r_k) \in \mathcal{V}_k$ (we omit k when it is clear from context). In the case of a non-deterministic analyst, we add its coin tosses as the first entry in the view, in order to ensure each (q_i, θ_i) is a deterministic function of v_{i-1} , as detailed in Definition A.3; given a view $v \in \mathcal{V}$, we denote by \bar{q}_v the sequence of queries that it induces. Slightly abusing notation, we denote $v = M(s, A)$. The distributions $D_{\mathcal{R}|\mathcal{X}^n}^{q,\theta}$, $D_{\mathcal{R}}^{q,\theta}$, $D_{\mathcal{X}^n|\mathcal{R}}^{q,\theta}$, and D^r naturally extend to versions where a sequence of queries is generated by an analyst: $D_{\mathcal{V}|\mathcal{X}^n}^A$, $D_{\mathcal{V}}^A$, $D_{\mathcal{X}^n|\mathcal{V}}^A$, and $D^v := D_{\mathcal{X}|\mathcal{V}}^A(\cdot | v)$.

Notice that M holds s but has no access to D . On the other hand, A might have access to D , but her only information regarding s comes from M ’s responses as represented by D^v . This intuitively turns some metric of distance between D and D^v into a measure of potential overfitting, an intuition that we formalize in Definition 3.3 and Lemma 3.5.

2.2 Notation

Throughout the paper, calligraphic letters denote domains (e.g., \mathcal{X}), lower case letters denote elements of domains (e.g., $x \in \mathcal{X}$), and capital letters denote random variables (e.g., $X \sim D_{\mathcal{X}}$).

We omit most superscripts and subscripts when clear from context (e.g., $D(r | s) = D_{\mathcal{R}|\mathcal{X}^n}^{q,\theta}(r | s)$ is the probability to receive $r \in \mathcal{R}$ as a response, conditioned on a particular input dataset s , given a query q and parameters θ). Unless specified otherwise, we assume $D_{\mathcal{X}^n}$, n , k , and M are fixed, and omit them from notations and definitions.

We use $\|\cdot\|$ to denote the Euclidean norm, so $\|\bar{q}_v(x)\| = \sqrt{\sum_{i=1}^k (q_i(x))^2}$ will denote the norm of a concatenated sequence of queries.

2.3 Definitions

We introduce terminology to describe the accuracy of responses that a mechanism produces in response to queries.

³Most of the definitions and claims in this paper can be extended beyond \mathbb{R} and the Euclidean norm. To do so, the divergence in Lemma 3.5 and the φ function (Definition 4.2) must be chosen accordingly.

Definition 2.1 (Accuracy of a mechanism). Given a dataset s , an analyst A , and a view $v = M(s, A)$, we define three types of output error for the mechanism: *sample error* $\text{err}_S(s, v) := \max_{i \in [k]} |r_i - q_i(s)|$, *distribution error* $\text{err}_D(s, v) := \max_{i \in [k]} |r_i - q_i(D)|$, and *posterior error* $\text{err}_P(s, v) := \max_{i \in [k]} |r_i - q_i(D^v)|$, where q_i is the i th query and r_i is the i th response in v .

These errors can be viewed as random variables with a distribution that is induced by the underlying distribution D and the internal randomness of M and A . Given $\epsilon, \delta \geq 0$ we call a mechanism M (ϵ, δ) -*sample/distribution/posterior accurate* with respect to A if

$$\Pr_{S \sim D^{(n)}, V \sim M(S, A)} [(\text{err}(S, V)) > \epsilon] \leq \delta,$$

for $\text{err} = \text{err}_S/\text{err}_D/\text{err}_P$, respectively.

Notice that if each possible value of ϵ has a corresponding δ , then δ is essentially a function of ϵ . Given such a function $\delta : \mathbb{R}^+ \rightarrow [0, 1]$, we will say the mechanism is $(\epsilon, \delta(\epsilon))$ accurate, a perspective which will be used in Lemma 3.1.

We start by introducing a particular family of queries known as *linear queries*, which will be used to state the main results in this paper, but it should be noted that many of the claims extend to arbitrary queries as discussed in Section C.2

Definition 2.2 (Linear queries). A function $q : \mathcal{X}^n \rightarrow \mathbb{R}$ is a *linear query* if it is defined by a function $q_1 : \mathcal{X} \rightarrow \mathbb{R}$ such that $q(s) := \frac{1}{n} \sum_{i=1}^n q_1(s_i)$. For simplicity, we denote q_1 as q throughout.

The aforementioned error bounds implicitly assume some known scale of the queries, which is usually chosen to be their range, denoted by $\Delta_q := \sup_{x, y \in \mathcal{X}} |q(x) - q(y)|$ (we refer to such a query as Δ -bounded). This poses an issue for concentrated random variables, where the ‘‘typical’’ range can be arbitrarily smaller than the range, which might even be infinite. A natural alternative source of scale we consider in this paper is the query’s variance $\sigma_q^2 := \mathbb{E}_{X \sim D_{\mathcal{X}}} [(q(X) - q(D_{\mathcal{X}}))^2]$, or its variance proxy in the case of a sub-Gaussian query (Definition 5.2). The corresponding definitions for arbitrary queries are presented in Section C.2.

Our main tool for ensuring generalization is the Gaussian mechanism, which simply adds Gaussian noise to the empirical value of a query.

Definition 2.3 (Gaussian mechanism). Given $\eta > 0$ and a query q , the *Gaussian mechanism* with noise parameter η returns its empirical mean $q(s)$ after adding a random value, sampled from an unbiased Gaussian distribution with variance η^2 . Formally, $M(s, q) \sim \mathcal{N}(q(s), \eta^2)$.⁴

3 Analyzing adaptivity-driven overfitting

In this section, we give a clean, new characterization of the harms of adaptivity. Our goal is to bound the distribution error of a mechanism that responds to queries generated by an adaptive analyst. This bound will be achieved via a triangle inequality, by bounding both the posterior accuracy and the Bayes stability (Definition 3.3). Missing proofs from this section appear in Appendix C.

The simpler part of the argument is posterior accuracy, which we prove can be inherited directly from the sample accuracy of a mechanism. This lemma resembles Lemma 6 in Jung et al. [2020], but has the advantage of being independent of the range of the queries.

Lemma 3.1 (Sample accuracy implies posterior accuracy). *Given a function $\delta : \mathbb{R} \rightarrow [0, 1]$ and an analyst A , if a mechanism M is $(\epsilon, \delta(\epsilon))$ -sample accurate for all $\epsilon > 0$, then M is $(\epsilon, \delta'(\epsilon))$ -posterior accurate for $\delta'(\epsilon) := \inf_{\xi \in (0, \epsilon)} \left(\frac{1}{\xi} \int_{\epsilon - \xi}^{\infty} \delta(t) dt \right)$.*

We use this lemma to provide accuracy guarantees for the Gaussian mechanism.

⁴In the case of an adaptive process, one can also consider the case where η_i are adaptively chosen by the analyst and provided to the mechanism as the auxiliary parameter θ_i .

Lemma 3.2 (Accuracy of Gaussian mechanism). *Given $\eta > 0$, the Gaussian mechanism with noise parameter η that receives k queries is $(\epsilon, \delta(\epsilon))$ -sample accurate for $\delta(\epsilon) := \frac{2k}{\sqrt{\pi}} e^{-\frac{\epsilon^2}{2\eta^2}}$, and $(\epsilon, \delta(\epsilon))$ -posterior accurate for $\delta(\epsilon) := 4k \cdot e^{-\frac{\epsilon^2}{4\eta^2}}$.*

In order to complete the triangle inequality, we have to define the stability of the mechanism. Bayes stability captures the concept that the results returned by a mechanism and the queries selected by the adaptive adversary are such that the queries behave similarly on the true data distribution and on the posterior distribution induced by those results. This notion first appeared in Jung et al. [2020], under the name *Posterior Sensitivity*, as did the following theorem.

Definition 3.3 (Bayes stability). Given $\epsilon, \delta > 0$ and an analyst A , we say M is (ϵ, δ) -Bayes stable with respect to A , if

$$\Pr_{\substack{S \sim D^{(n)} \\ V \sim M(S, A), Q \sim A(V)}} [|Q(D^V) - Q(D)| > \epsilon] \leq \delta.$$

Bayes stability and sample accuracy (implying posterior accuracy) combine via a triangle inequality to give distribution accuracy.

Theorem 3.4 (Generalization). *Given two functions $\delta_1 : \mathbb{R} \rightarrow [0, 1]$, $\delta_2 : \mathbb{R} \rightarrow [0, 1]$, and an analyst A , if a mechanism M is $(\epsilon, \delta_1(\epsilon))$ -Bayes stable and $(\epsilon, \delta_2(\epsilon))$ -sample accurate with respect to A , then M is $(\epsilon, \delta'(\epsilon))$ -distribution accurate for $\delta'(\epsilon) := \inf_{\epsilon' \in (0, \epsilon), \xi \in (0, \epsilon - \epsilon')} \left(\delta_1(\epsilon') + \frac{1}{\xi} \int_{\epsilon - \epsilon' - \xi}^{\infty} \delta_2(t) dt \right)$.*

Since achieving posterior accuracy is relatively straightforward, guaranteeing Bayes stability is the main challenge in leveraging this theorem to achieve distribution accuracy with respect to adaptively chosen queries. The following lemma gives a useful and intuitive characterization of the quantity that the Bayes stability definition requires be bounded. Simply put, the Bayes factor $K(\cdot, \cdot)$ (defined in the lemma below) represents the amount of information leaked about the dataset during the interaction with an analyst, by moving from the prior distribution over data elements to the posterior induced by some view v . The degree to which a query q overfits to the dataset is expressed by the correlation between the query and that Bayes factor. This simple lemma is at the heart of the progress that we make in this paper, both in our intuitive understanding of adaptive data analysis, and in the concrete results we show in subsequent sections. Its corresponding version for arbitrary queries are presented in Section C.2.

Lemma 3.5 (Covariance stability). *Given a view $v \in \mathcal{V}$ and a linear query q ,*

$$q(D^v) - q(D_{\mathcal{X}}) = \underset{X \sim D_{\mathcal{X}}}{\text{Cov}}(q(X), K(X, v)).$$

Furthermore, given $\Delta, \sigma > 0$,

$$\sup_{q \in \mathcal{Q} \text{ s.t. } \Delta_q \leq \Delta} |q(D^v) - q(D_{\mathcal{X}})| = \Delta \cdot \mathbf{D}_{TV}(D^v \| D_{\mathcal{X}})$$

and

$$\sup_{q \in \mathcal{Q} \text{ s.t. } \sigma_q^2 \leq \sigma^2} |q(D^v) - q(D_{\mathcal{X}})| = \sigma \sqrt{\mathbf{D}_{\chi^2}(D^v \| D_{\mathcal{X}})},$$

where $K(x, v) := \frac{D(x|v)}{D(x)} = \frac{D(v|x)}{D(v)}$ is the Bayes factor of x given v (and vice-versa), \mathbf{D}_{TV} is the total variation distance (Definition B.1), and \mathbf{D}_{χ^2} is the chi-square divergence (Definition B.2).

Proof. By definition, $q(D^v) = \mathbb{E}_{X \sim D^v} [q(X)] = \mathbb{E}_{X \sim D} [K(X, v) q(X)]$, so

$$q(D^v) - q(D) = \mathbb{E}_{X \sim D} [K(X, v) q(X)] - \underbrace{\mathbb{E}_{X \sim D} [K(X, v)]}_{=1} \cdot \underbrace{\mathbb{E}_{X \sim D} [q(X)]}_{=q(D)} = \underset{X \sim D}{\text{Cov}}(q(X), K(X, v)).$$

The second part is a direct result of the known variational representation of total variation distance and χ^2 divergence, which are both f -divergences (see Equations 7.88 and 7.91 in Polyanskiy and Wu [2022] for more details). \square

Using the first part of the lemma, we guarantee Bayes stability by bounding the correlation between specific q and $K(\cdot, v)$ as discussed in Section 6. The second part of this Lemma implies that bounding the appropriate divergence is necessary and sufficient for bounding the Bayes stability of the worst query in the corresponding family, which is how the main theorems of this paper are all achieved, using the next corollary.

Corollary 3.6. *Given an analyst A , for any $\epsilon > 0$ we have*

$$\Pr_{\substack{S \sim D^{(n)} \\ V \sim M(S, A), Q \sim A(V)}} [|Q(D^V) - Q(D)| > \sigma_Q \cdot \epsilon] \leq \Pr_{\substack{S \sim D^{(n)} \\ V \sim M(S, A)}} [\mathbf{D}_{\mathcal{X}^2}(D_{\mathcal{X}}^v \| D_{\mathcal{X}}) > \epsilon^2].$$

The corresponding version of this corollary for bounded range queries provides an alternative proof to the generalization guarantees of the the LS stability notion [Shenfeld and Ligett, 2019].

4 Pairwise concentration

In order to leverage Lemma 3.5, we need a stability notion that implies Bayes stability of query responses in a manner that depends on the actual datasets and the actual queries (not just the worst case). In this section we propose such a notion and prove several key properties of it. Missing proofs from this section can be found in Appendix D.

We start by introducing a measure of the stability loss of the mechanism.

Definition 4.1 (Stability loss). Given two sample sets $s, s' \in \mathcal{X}^n$, a query q , and a response $r \in \mathcal{R}$, we denote by $\ell_q(s, s'; r) := \ln \left(\frac{D^q(r|s)}{D^q(r|s')} \right)$ the *stability loss* of r between s and s' for q (we omit q from notation for simplicity).⁵ Notice that if $R \sim M(s, q)$, the stability loss defines a random variable. Similarly, given two elements $x, y \in \mathcal{X}$, we denote $\ell(x, y; r) := \ln \left(\frac{D(r|x)}{D(r|y)} \right)$ and $\ell(x; r) := \ln \left(\frac{D(r|x)}{D(r)} \right)$. This definition extends to views as well.

Next we introduce a notion of a bound on the stability loss, where the bound is allowed to depend on the pair of swapped inputs and on the issued queries.

Definition 4.2 (Pairwise concentration). Given a non-negative function $\varphi : \mathcal{X}^n \times \mathcal{X}^n \times \mathcal{Q} \rightarrow \mathbb{R}^+$ which is symmetric in its first two arguments, a mechanism M will be called φ -Pairwise Concentrated (or PC, for short), if for any two datasets $s, s' \in \mathcal{X}^n$ and $q \in \mathcal{Q}$, for any $\alpha \geq 0$,⁶

$$\mathbb{E}_{R \sim M(s, q)} [\exp((\alpha - 1) \ell(s, s'; R))] \leq \exp(\alpha(\alpha - 1) \varphi(s, s'; q)).$$

Given a non-negative function $\varphi : \mathcal{X}^n \times \mathcal{X}^n \times \mathcal{V} \rightarrow \mathbb{R}^+$ which is symmetric in its first two arguments and an analyst A , a mechanism M will be called φ -Pairwise Concentrated with respect to A , if for any $s, s' \in \mathcal{X}^n$ and $\alpha \geq 0$,

$$\mathbb{E}_{V \sim M(s, A)} [\exp((\alpha - 1) (\ell(s, s'; V) - \alpha \varphi(s, s'; V)))] \leq 1.$$

We refer to such a function φ as a *similarity function* over responses or views, respectively.

The similarity function serves as a measure of the local sensitivity of the issued queries with respect to the replacement of the two datasets, by quantifying the extent to which they differ from each other with respect to the query q . The case of noise addition mechanisms provides a natural intuitive interpretation, where the φ scales with the difference between $q(s)$ and $q(s')$, which governs how much observing the response r distinguishes between the two datasets, as stated in the next lemma.

We note that the first part of this definition can be viewed as a refined version of zCDP (Definition B.18), where the bound on the Rényi divergence (Definition B.5) is a function of the sample sets

⁵This quantity is sometimes referred to as *privacy loss* in the context of differential privacy, or *information density* in the context of information theory.

⁶The requirement for the range $0 \leq \alpha \leq 1$ might appear somewhat counter-intuitive, and results from the fact that φ provides simultaneously a bound on the stability loss random variable's mean (the $-(\alpha - 1)$ coefficient) and variance proxy (the $(\alpha - 1)^2$ coefficient), as discusses in Definitions B.17, B.18.

and the query. As for the second part, since the bound depends on the queries, which themselves are random variables, it should be viewed as a bound on the Rényi dissimilarity notion that we introduce in the appendix (Definition B.9). This kind of extension is not limited to Rényi divergence, as discussed in Appendix B.

Lemma 4.3 (Gaussian mechanism is PC). *Given $\eta > 0$ and an analyst A , if M is a Gaussian mechanism with noise parameter η , then it is φ -PC for $\varphi(s, s'; q) := \frac{(q(s) - q(s'))^2}{2\eta^2}$ and φ -PC with respect to A for $\varphi(s, s'; v) := \frac{\|\bar{q}_v(s) - \bar{q}_v(s')\|^2}{2\eta^2}$.*

To leverage this stability notion in an adaptive setting, it must hold under adaptive composition, which we prove in the next theorem.

Theorem 4.4 (PC composition). *Given $k \in \mathbb{N}$, a similarity function over responses φ , and an analyst A issuing k queries, if a mechanism M is φ -PC, then it is $\tilde{\varphi}$ -PC with respect to A , where $\tilde{\varphi}(s, s'; v) := \sum_{i=1}^k \varphi(s, s'; q_i)$.*

Before we state the stability properties of PC mechanisms, we first transition to its element-wise version, which while less general than the original definition, is more suited for our use.

Lemma 4.5 (Element-wise PC). *Given a similarity function φ and an analyst A , if a mechanism M that is φ -PC with respect to A receives an iid sample from \mathcal{X}^n , then for any two elements $x, y \in \mathcal{X}$ and $\alpha \geq 1$ we have*

$$\mathbb{E}_{V \sim D(\cdot | x)} [\exp((\alpha - 1)(\ell(x, y; V) - \alpha\varphi(x, y; V)))] \leq 1$$

and

$$\mathbb{E}_{V \sim D(\cdot | x)} [\exp((\alpha - 1)(\ell(x; V) - \alpha\varphi(x; V)))] \leq 1,$$

where $\varphi(x, y; v) := \sup_{s \in \mathcal{X}^{n-1}} (\varphi((s, x), (s, y); v))$ and $\varphi(x; v) := \ln \left(\mathbb{E}_{Y \sim D} [e^{\varphi(x, Y; v)}] \right)$.

Finally, we bound the Bayes stability of PC mechanisms.

Theorem 4.6 (PC stability). *Given a similarity function over views φ and an analyst A , if a mechanism M that is φ -PC with respect to A receives an iid sample from \mathcal{X}^n , then for any $\epsilon, \delta > 0$, $\phi \geq \mathbb{E}_{\substack{X \sim D, S \sim D^{(n)} \\ V \sim M(S, A)}} [\varphi(X; V)]$, we have*

$$\begin{aligned} & \Pr_{\substack{S \sim D^{(n)} \\ V \sim M(S, A), Q \sim A(V)}} [|Q(D_{\mathcal{X}}^v) - Q(D_{\mathcal{X}})| > \sigma_Q \epsilon] \\ & \leq \Pr_{\substack{S \sim D^{(n)} \\ V \sim M(S, A)}} \left[\mathbb{E}_{X \sim D} \left[e^{27 \ln(\frac{1}{\delta})(\varphi(X; V) + \phi)} \right] > 1 + \frac{\epsilon^2}{6} \right] + O\left(\frac{e^{\phi} \delta}{\phi \epsilon^2}\right), \end{aligned}$$

where $\varphi(x; v)$ is as defined in Lemma 4.5.

An exact version of the bound can be found in Theorem D.10.

Proof outline. From Corollary 3.6,

$$\Pr_{\substack{S \sim D^{(n)} \\ V \sim M(S, A), Q \sim A(V)}} [|Q(D_{\mathcal{X}}^v) - Q(D_{\mathcal{X}})| > \sigma_Q \epsilon] \leq \Pr_{\substack{S \sim D^{(n)} \\ V \sim M(S, A)}} [\mathbf{D}_{\chi^2}(D_{\mathcal{X}}^v \| D_{\mathcal{X}}) > \epsilon^2].$$

Denoting $\varepsilon(x; v) := \varphi(x; v) + \phi + 4\sqrt{\ln(\frac{1}{\delta})(\varphi(x; v) + \phi)}$, this can be bounded by

$$\Pr_{\substack{S \sim D^{(n)} \\ V \sim M(S, A)}} [\mathbf{D}_{\chi^2}(D_{\mathcal{X}}^v \| D_{\mathcal{X}}) > \mathbb{E}_{X \sim D} \left[\left(e^{\varepsilon(X; V)} - 1 \right)^2 \right] + \frac{\epsilon^2}{2}] + \Pr_{\substack{S \sim D^{(n)} \\ V \sim M(S, A)}} \left[\mathbb{E}_{X \sim D} \left[\left(e^{\varepsilon(X; V)} - 1 \right)^2 \right] > \frac{\epsilon^2}{2} \right].$$

We then show via algebraic manipulation that this second term can be bounded using the inequality

$$\mathbb{E}_{X \sim D} \left[\left(e^{\varepsilon(X; v)} - 1 \right)^2 \right] \leq 3 \left(\mathbb{E}_{X \sim D} \left[e^{27 \ln(\frac{1}{\delta})(\varphi(X; v) + \phi)} \right] - 1 \right),$$

and, recalling the definition of the chi square divergence between P and Q (Definition B.2) $\mathbf{D}_{\chi^2}(P\|Q) := \mathbb{E}_{X \sim D_2} \left[\left(\frac{P(X)}{Q(X)} - 1 \right)^2 \right]$, the first term can be bounded using Markov's inequality followed by the Cauchy-Schwarz inequality by the term:

$$\begin{aligned} & \frac{2}{\epsilon^2} \mathbb{E}_{\substack{X \sim D, S \sim D^{(n)} \\ V \sim M(S, A)}} \left[\left(\frac{D(V|X)}{D(V)} - 1 \right)^2 \cdot \mathbb{1}_{B(\epsilon)}((X, V)) \right] \\ & \leq \frac{2}{\epsilon^2} \sqrt{\overbrace{\left(\mathbb{E}_{\substack{X \sim D, S \sim D^{(n)} \\ V \sim M(S, A)}} \left[\left(\frac{D(V|X)}{D(V)} \right)^4 \right] + 1 \right)^*} \cdot \overbrace{\Pr_{\substack{X \sim D, S \sim D^{(n)} \\ V \sim M(S, A)}} [(X, V) \in B(\epsilon)]}^{**}} \end{aligned}$$

where $B(\epsilon) := \left\{ (x, v) \in \mathcal{X} \times \mathcal{V} \mid \left| \ln \left(\frac{D(v|x)}{D(v)} \right) \right| > \epsilon(x; v) \right\}$.

We then bound the quantity $*$ by $O(e^\phi)$, and we bound $**$ by $O\left(\frac{\delta^2}{\phi}\right)$, which completes the proof. \square

Combining this theorem with Theorem 3.4 we get that PC and sample accuracy together imply distribution accuracy.

5 Variance-based generalization guarantees for the Gaussian mechanism

In this section we leverage the theorems of the previous sections to prove variance-based generalization guarantees for the Gaussian mechanism under adaptive data analysis. All proofs from this section appear in Appendix E.

We first provide generalization guarantees for bounded queries.

Theorem 5.1 (Generalization guarantees for bounded queries). *Given $k \in \mathbb{N}$; $\Delta, \sigma, \epsilon \geq 0$; $0 < \delta \leq \frac{1}{e}$; and an analyst A issuing k Δ -bounded linear queries with variance bounded by σ^2 , if M is a Gaussian mechanism with noise parameter $\eta = \Theta\left(\sigma\sqrt{\frac{\sqrt{k}}{n}}\right)$ that receives an iid dataset of size $n = \Omega\left(\max\left\{\frac{\Delta}{\epsilon}, \frac{\sigma^2}{\epsilon^2}\right\} \sqrt{k} \cdot \ln\left(\frac{k\sigma}{\delta\epsilon}\right)\right)$, then M is (ϵ, δ) -distribution accurate.*

An exact version of the bound can be found in Theorem E.4.

This theorem provides significant improvement over similar results that were achieved using differential privacy (see, e.g., Theorem 13 in Jung et al. [2020], which is defined for $\Delta = 1$), by managing to replace the Δ^2 term with σ^2 . The remaining dependence on Δ is inevitable, in the sense that such a dependence is needed even for non-adaptive analysis. This improvement resembles the improvement provided by Bernstein's inequality over Hoeffding's inequality.

This generalization guarantee, which nearly avoids dependence on the range of the queries, begs the question of whether it is possible to extend these results to handle unbounded queries. Clearly such a result would not be true without some bound on the tail distribution for a single query, so we focus in the next theorem on the case of sub-Gaussian queries. Formally, we will consider the case where $q(X) - q(D)$ is a sub-Gaussian random variable, for all queries.

Definition 5.2 (Sub-Gaussian random variable). Given $\sigma > 0$, a random variable $X \in \mathbb{R}$ will be called σ^2 -sub-Gaussian if for any $\lambda \in \mathbb{R}$ we have $\mathbb{E}[e^{\lambda X}] \leq e^{\frac{\lambda^2}{2}\sigma}$.

We can now state our result for unbounded queries.

Theorem 5.3 (Generalization guarantees for sub-Gaussian queries). *Given $k \in \mathbb{N}$; $\sigma, \epsilon \geq 0$; $0 < \delta \leq \frac{1}{e}$; and an analyst A issuing k σ^2 -sub-Gaussian linear queries, if M is a Gaussian mechanism with noise parameter $\eta = \Theta\left(\sigma\sqrt{\frac{\sqrt{k}}{n}}\right)$ that receives an iid dataset of size $n = \Omega\left(\frac{\sigma^2}{\epsilon^2} \sqrt{k} \ln\left(\frac{k\sigma}{\delta\epsilon}\right)\right)$, then M is (ϵ, δ) -distribution accurate.*

An exact version of the bound can be found in Theorem E.10.

These results extend to the case where the variance (or variance proxy) of each query q_i is bounded by a unique value σ_i^2 , by simply passing this value to the mechanism as auxiliary information and scaling the added noise η_i accordingly. Furthermore, using this approach we can quantify the extent to which incorrect bounds affect the accuracy guarantee. Overestimating the bound on a query’s variance would increase the error of the response to this query by a factor of square root of the ratio between the assumed and the correct variance, while the error of the other responses would only decrease. On the other hand, underestimating the bound on a query’s variance would only decrease the error of the response to this query, while increasing the error of each subsequent query by a factor of the square root of the ratio between the assumed and the correct variance, divided by the number of subsequent queries. A formal version of this claim can be found in Section E.3.

6 Discussion

The contribution of this paper is two-fold. In Section 3, we provide a tight measure of the level of overfitting of some query with respect to previous responses. In Sections 4 and 5, we demonstrate a toolkit to utilize this measure, and use it to prove new generalization properties of fundamental noise-addition mechanisms. The novelty of the PC definition stems replacing the fixed parameters that appear in the differential privacy definition with a function of the datasets and the query. The definition presented in this paper provides a generalization of zero-concentrated differential privacy, and future work could study similar generalizations of other privacy notions, as discussed in Section B.4.

One small extension of the present work would be to consider queries with range \mathbb{R}^d . It would also be interesting to extend our results to handle arbitrary normed spaces, using appropriate noise such as perhaps the Laplace mechanism. It might also be possible to relax our assumption that data elements are drawn iid to a weaker independence requirement. Furthermore, it would be interesting to explore an extension from linear queries to general low-sensitivity queries.

We hope that the mathematical toolkit that we establish in Appendix B to analyze our stability notion may find additional applications, perhaps also in context of privacy accounting. Furthermore, the max divergence can be generalized analogously to the “dynamic” generalization of Rényi divergence proposed in this paper (B.9), perhaps suggesting that this approach may be useful in analyzing other mechanisms as well.

Our Covariance Lemma (3.5) shows that there are two possible ways to avoid adaptivity-driven overfitting—by bounding the Bayes factor term, which induces a bound on $|q(D^v) - q(D)|$, as we do in this work, or by bounding the correlation between q and $K(\cdot, v)$. This second option suggests interesting directions for future work. For example, to capture an analyst that is non-worst-case in the sense that she “forgets” some of the information that she has learned about the dataset, both the posterior accuracy and the Bayes stability could be redefined with respect to the internal state of the analyst instead of with respect to the full view. This could allow for improved bounds in the style of Zrnic and Hardt [2019].

Acknowledgments and Disclosure of Funding

We gratefully acknowledge productive discussions with Etam Benger, Vitaly Feldman, Yosef Rinott, Aaron Roth, and Tomer Shoham. This work was supported in part by a gift to the McCourt School of Public Policy and Georgetown University, Simons Foundation Collaboration 733792, Israel Science Foundation (ISF) grant 2861/20, and a grant from the Israeli Council for Higher Education. Shenfeld’s work was also partly supported by the Apple Scholars in AI/ML PhD Fellowship. Part this work was completed while Ligett was visiting Princeton University’s Center for Information Technology Policy.

References

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

- Raef Bassily and Yoav Freund. Typical stability. *arXiv preprint arXiv:1604.03336*, 2016.
- Raef Bassily, Adam Groce, Jonathan Katz, and Adam Smith. Coupled-worlds privacy: Exploiting adversarial uncertainty in statistical data privacy. In *IEEE Symposium on Foundations of Computer Science*, pages 439–448, 2013.
- Raef Bassily, Kobbi Nissim, Adam Smith, Thomas Steinke, Uri Stemmer, and Jonathan Ullman. Algorithmic stability for adaptive data analysis. *SIAM Journal on Computing*, 2021.
- Raghav Bhaskar, Abhishek Bhowmick, Vipul Goyal, Srivatsan Laxman, and Abhradeep Thakurta. Noiseless database privacy. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 215–232. Springer, 2011.
- Guy Blanc. Subsampling suffices for adaptive data analysis. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, pages 999–1012, 2023.
- Olivier Bousquet and André Elisseeff. Stability and generalization. *The Journal of Machine Learning Research*, 2:499–526, 2002.
- Mark Bun and Thomas Steinke. Concentrated differential privacy: Simplifications, extensions, and lower bounds. In *Theory of Cryptography Conference*, pages 635–658. Springer, 2016.
- Konstantinos Chatzikokolakis, Miguel E Andrés, Nicolás Emilio Bordenabe, and Catuscia Palamidessi. Broadening the scope of differential privacy using metrics. In *Privacy Enhancing Technologies: 13th International Symposium, PETS 2013, Bloomington, IN, USA, July 10-12, 2013. Proceedings 13*, pages 82–102. Springer, 2013.
- Elbert Du and Cynthia Dwork. Improved generalization guarantees in restricted data models. In *3rd Symposium on Foundations of Responsible Computing*, 2022.
- John C. Duchi, Saminul Haque, and Rohith Kuditipudi. A fast algorithm for adaptive private mean estimation. *ArXiv*, abs/2301.07078, 2023.
- Cynthia Dwork and Guy N Rothblum. Concentrated differential privacy. *arXiv preprint arXiv:1603.01887*, 2016.
- Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography Conference*, pages 265–284. Springer, 2006.
- Cynthia Dwork, Vitaly Feldman, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Aaron Leon Roth. Preserving statistical validity in adaptive data analysis. In *ACM Symposium on Theory of Computing*, pages 117–126, 2015.
- Hamid Ebadi, David Sands, and Gerardo Schneider. Differential privacy: Now it’s getting personal. *Acm Sigplan Notices*, 50(1):69–81, 2015.
- Sam Elder. Challenges in bayesian adaptive data analysis. *arXiv preprint arXiv:1604.02492*, 2016.
- Vitaly Feldman and Thomas Steinke. Generalization for adaptively-chosen estimators via stable median. In *Conference on Learning Theory*, pages 728–757. PMLR, 2017.
- Vitaly Feldman and Thomas Steinke. Calibrating noise to variance in adaptive data analysis. In *Conference On Learning Theory*, pages 535–544. PMLR, 2018.
- Vitaly Feldman and Tijana Zrnic. Individual privacy accounting via a Renyi filter. *Advances in Neural Information Processing Systems*, 34:28080–28091, 2021.
- Johannes Gehrke, Michael Hay, Edward Lui, and Rafael Pass. Crowd-blending privacy. In *Annual Cryptology Conference*, pages 479–496. Springer, 2012.
- Andrew Gelman and Eric Loken. The statistical crisis in science. *American Scientist*, 102(6):460, 2014.
- Arpita Ghosh and Aaron Roth. Selling privacy at auction. In *ACM Conference on Electronic Commerce*, pages 199–208, 2011.

- Moritz Hardt and Jonathan Ullman. Preventing false discovery in interactive data analysis is hard. In *IEEE Symposium on Foundations of Computer Science*, pages 454–463, 2014.
- Jacob Imola, Shiva Kasiviswanathan, Stephen White, Abhinav Aggarwal, and Nathanael Teissier. Balancing utility and scalability in metric differential privacy. In *Uncertainty in Artificial Intelligence*, pages 885–894. PMLR, 2022.
- John PA Ioannidis. Why most published research findings are false. *PLoS medicine*, 2(8):e124, 2005.
- Christopher Jung, Katrina Ligett, Seth Neel, Aaron Roth, Saeed Sharifi-Malvajerdi, and Moshe Shenfeld. A new analysis of differential privacy’s generalization guarantees. In *Innovations in Theoretical Computer Science (ITCS)*, 2020.
- Vishesh Karwa and Salil Vadhan. Finite Sample Differentially Private Confidence Intervals. In *Innovations in Theoretical Computer Science Conference*, volume 94 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 44:1–44:9, 2018.
- Shiva P Kasiviswanathan and Adam Smith. On the ‘semantics’ of differential privacy: A Bayesian formulation. *Journal of Privacy and Confidentiality*, 6(1), 2014.
- Aryeh Kontorovich, Menachem Sadigurschi, and Uri Stemmer. Adaptive data analysis with correlated observations. In *International Conference on Machine Learning*, pages 11483–11498. PMLR, 2022.
- Ninghui Li, Wahbeh Qardaji, Dong Su, Yi Wu, and Weining Yang. Membership privacy: A unifying framework for privacy definitions. In *ACM SIGSAC Conference on Computer & Communications Security*, pages 889–900, 2013.
- Katrina Ligett, Seth Neel, Aaron Roth, Bo Waggoner, and Steven Z Wu. Accuracy first: Selecting a differential privacy level for accuracy constrained erm. *Advances in Neural Information Processing Systems*, 30, 2017.
- Ashwin Machanavajjhala, Johannes Gehrke, and Michaela Götz. Data publishing against realistic adversaries. *Proceedings of the VLDB Endowment*, 2(1):790–801, 2009.
- Ilya Mironov. Rényi differential privacy. In *IEEE Computer Security Foundations Symposium (CSF)*, pages 263–275, 2017.
- Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. Smooth sensitivity and sampling in private data analysis. In *ACM Symposium on Theory of Computing*, pages 75–84, 2007.
- Yury Polyanskiy and Yihong Wu. *Information Theory: From Coding to Learning (book draft)*. Cambridge university press, 2022.
- Ryan M Rogers, Aaron Roth, Jonathan Ullman, and Salil Vadhan. Privacy odometers and filters: Pay-as-you-go composition. *Advances in Neural Information Processing Systems*, 29:1921–1929, 2016.
- Daniel Russo and James Zou. Controlling bias in adaptive data analysis using information theory. In *Artificial Intelligence and Statistics*, pages 1232–1240, 2016.
- Shai Shalev-Shwartz, Ohad Shamir, Nathan Srebro, and Karthik Sridharan. Learnability, stability and uniform convergence. *The Journal of Machine Learning Research*, 11:2635–2670, 2010.
- Moshe Shenfeld and Katrina Ligett. A necessary and sufficient stability notion for adaptive generalization. *Advances in Neural Information Processing Systems*, 32:11485–11494, 2019.
- Thomas Steinke and Jonathan Ullman. Interactive fingerprinting codes and the hardness of preventing false discovery. In *Conference on Learning Theory*, pages 1588–1628. PMLR, 2015.
- Thomas Steinke and Lydia Zakyntinou. Reasoning about generalization via conditional mutual information. In *Conference on Learning Theory*, pages 3437–3452. PMLR, 2020.
- Aleksei Triastcyn and Boi Faltings. Bayesian differential privacy for machine learning. In *International Conference on Machine Learning*, pages 9583–9592. PMLR, 2020.

- Tim Van Erven and Peter Harremoës. Rényi divergence and kullback-leibler divergence. *IEEE Transactions on Information Theory*, 60(7):3797–3820, 2014.
- H Victor, Michael J Klass, and Tze Leung Lai. Pseudo-maximization and self-normalized processes. *Probability Surveys*, 4:172–192, 2007.
- Yu-Xiang Wang. Per-instance differential privacy. *Journal of Privacy and Confidentiality*, 9(1), 2019.
- Justin Whitehouse, Aaditya Ramdas, Ryan Rogers, and Steven Wu. Fully-adaptive composition in differential privacy. In *International Conference on Machine Learning*, pages 36990–37007. PMLR, 2023.
- Bin Yang, Issei Sato, and Hiroshi Nakagawa. Bayesian differential privacy on correlated data. In *Proceedings of the 2015 ACM SIGMOD international conference on Management of Data*, pages 747–762, 2015.
- Yuqing Zhu, Jinshuo Dong, and Yu-Xiang Wang. Optimal accounting of differential privacy via characteristic function. In *International Conference on Artificial Intelligence and Statistics*, pages 4782–4817. PMLR, 2022.
- Tijana Zrnica and Moritz Hardt. Natural analysts in adaptive data analysis. In *International Conference on Machine Learning*, pages 7703–7711. PMLR, 2019.

A Formal definitions

A.1 Distributions of interest

Definition A.1 (Distributions over \mathcal{X}^n and \mathcal{R}). A distribution $D_{\mathcal{X}^n}$, a query q , and a mechanism $M : \mathcal{X}^n \times \mathcal{Q} \rightarrow \mathcal{R}$, together induce a set of distributions over \mathcal{X}^n , \mathcal{R} , and $\mathcal{X}^n \times \mathcal{R}$.

The *conditional distribution* $D_{\mathcal{R}|\mathcal{X}^n}^q$ over \mathcal{R} represents the probability to get r as the output of $M(s, q)$. That is, $\forall s \in \mathcal{X}^n, r \in \mathcal{R}$, if \mathcal{R} is countable,

$$D_{\mathcal{R}|\mathcal{X}^n}^q(r | s) := \Pr_{R \sim M(s, q)} [R = r | s],$$

where the probability is taken over the internal randomness of M . In case \mathcal{R} is a measurable spaces, it represents the probability density function (the Radon–Nikodym derivative).

The *joint distribution* $D_{(\mathcal{X}^n, \mathcal{R})}^q$ over $\mathcal{X}^n \times \mathcal{R}$ represents the probability to sample a particular s and get r as the output of $M(s, q)$. That is, $\forall s \in \mathcal{X}^n, r \in \mathcal{R}$,

$$D_{(\mathcal{X}^n, \mathcal{R})}^q(s, r) := D_{\mathcal{X}^n}(s) \cdot D_{\mathcal{R}|\mathcal{X}^n}^q(r | s).$$

The *marginal distribution* $D_{\mathcal{R}}^q$ over \mathcal{R} represents the prior probability to get output r without any knowledge of s . That is, $\forall r \in \mathcal{R}$,

$$D_{\mathcal{R}}^q(r) := \int_{s \in \mathcal{X}^n} D_{(\mathcal{X}^n, \mathcal{R})}^q(s, r) dr.$$

The *conditional distribution* $D_{\mathcal{X}^n|\mathcal{R}}^q$ over \mathcal{X}^n represents the posterior probability that the input dataset to M was s , given that $M(\cdot, q)$ returns r . That is, $\forall s \in \mathcal{X}^n, r \in \mathcal{R}$,

$$D_{\mathcal{X}^n|\mathcal{R}}^q(s | r) := \frac{D_{(\mathcal{X}^n, \mathcal{R})}^q(s, r)}{D_{\mathcal{R}}^q(r)}.$$

Definition A.2 (Distributions over \mathcal{X} and \mathcal{R}). The *marginal distribution* $D_{\mathcal{X}}$ over \mathcal{X} represents the probability to get x by sampling a dataset and then sampling one element from that dataset uniformly at random. That is, $\forall x \in \mathcal{X}$,

$$D_{\mathcal{X}}(x) := \int_{s \in \mathcal{X}^n} D_{\mathcal{X}^n}(s) \cdot D_{\mathcal{X}|\mathcal{X}^n}(x | s) ds,$$

where $D_{\mathcal{X}|\mathcal{X}^n}(x | s)$ denotes the probability to get x by sampling s uniformly at random.

The *joint distribution* $D_{(\mathcal{X}, \mathcal{R})}^q$ over $\mathcal{X} \times \mathcal{R}$ represents the probability to get x by sampling a dataset uniformly at random and also get r as the output of $M(\cdot, q)$ from the same dataset. That is, $\forall x \in \mathcal{X}, r \in \mathcal{R}$,

$$D_{(\mathcal{X}, \mathcal{R})}^q(x, r) := \int_{s \in \mathcal{X}^n} D_{\mathcal{X}^n}(s) \cdot D_{\mathcal{X}|\mathcal{X}^n}(x | s) \cdot D_{\mathcal{R}|\mathcal{X}^n}^q(r | s) ds.$$

where $D_{\mathcal{X}|\mathcal{X}^n}(x | s)$ denotes the probability to get x by sampling s uniformly at random.

The *conditional distribution* $D_{\mathcal{X}|\mathcal{R}}^q$ over \mathcal{X} represents the probability to get x by sampling a dataset uniformly at random, given the fact that we got r as the output of $M(\cdot, q)$ from that dataset. That is, $\forall x \in \mathcal{X}, r \in \mathcal{R}$,

$$D^r(x) := D_{\mathcal{X}|\mathcal{R}}^q(x | r) := \int_{s \in \mathcal{X}^n} D_{\mathcal{X}^n|\mathcal{R}}^q(s | r) \cdot D_{\mathcal{X}|\mathcal{X}^n}(x | s) ds.$$

Although all of these definitions depend on $D_{\mathcal{X}^n}$ and M , we typically omit these from the notation for simplicity, and usually omit the superscripts and subscripts entirely. We include them only when necessary for clarity.

Though the conditional distributions $D_{\mathcal{R}|\mathcal{X}}^q$ and $D_{\mathcal{X}|\mathcal{R}}^q$ were not defined as the ratio between the joint and marginal distribution, the analogue of Bayes' rule still holds for these distributions. A formal proof of this claim can be found in Appendix A in Shenfeld and Ligett [2019].

A.2 A Formal treatment of adaptivity

Definition A.3 (Adaptive mechanism). Illustrated below, the *adaptive mechanism* $\text{Adp}_M : \mathcal{X}^n \times \mathcal{A} \rightarrow \mathcal{V}_k$ is a particular type of mechanism, which inputs an analyst as its query and which returns a view as its response and is parametrized by a mechanism M and number of iterations k . Given a dataset s and an analyst A as input, the adaptive mechanism iterates k times through the process where A sends a query and auxiliary information to M and receives its response to that query on the dataset. The adaptive mechanism returns the resulting sequence of k responses v_k . Naturally, this requires A to match M such that M 's output can be A 's input, and vice versa.

If A is randomized, we add one more step at the beginning where Adp_M randomly generates some bits c — A 's “coin tosses.” In this case, $v_k := (c, r_1, \dots, r_{ik})$ and A receives the coin tosses as an input as well. This addition turns q_{k+1} and θ_{k+1} into a deterministic function of v_i for any $i \in \mathbb{N}$, a fact that we later rely on. In this situation, the randomness of Adp_M results both from the randomness of the coin tosses and from that of M . We will denote $M(s, A) := \text{Adp}_M(s, A)$ for simplicity.

Adaptive mechanism Adp_M

Input: $s \in \mathcal{X}^n, A \in \mathcal{A}$

Output: $v_k \in \mathcal{V}_k$

$v_0 \leftarrow \emptyset$ or c

for $i \in [k]$:

$(q_i, \theta_i) \leftarrow A(v_{i-1})$

$r_i \leftarrow M(s, q_i, \theta_i)$

$v_i \leftarrow (v_{i-1}, r_i)$

return v_k

Using the adaptive mechanism we can extend the set of distributions presented in Definitions A.1 and A.2, to distributions over \mathcal{V} and \mathcal{X} , with \mathcal{V} taking the place of \mathcal{R} and A replacing q .

B Divergence, dissimilarity, and stability measures

Given two probability distributions P, Q over some domain Ω , there are many useful measures of their dissimilarity. In this section we develop what is as far as we know a new type of dissimilarity measure between distributions, which generalizes the notion of divergences. We start by recalling several common divergence measures in Section B.1. Next, Section B.2, we present in a generalization of the divergence notion, where a scalar is replaced by a function, yielding dynamic measures which we refer to as dissimilarities. We then prove (Section B.3) several convexity properties for max and Rényi dissimilarities. Finally (Section B.4) we propose several novel stability measures which build on the proposed dissimilarity measures and are used throughout the paper.

B.1 Divergence measures

Dissimilarity between distributions is typically measured using a divergence, that is, a function receiving two distributions as input, and outputting a non-negative value which is equal to 0 if and only if the two distributions are identical almost surely. One commonly used family of divergences is the f -divergences, defined as $\mathbf{D}_f(P\|Q) := \mathbb{E}_{\omega \sim Q} \left[f \left(\frac{P(\omega)}{Q(\omega)} \right) \right]$. Here we recall several instances which are used in this paper.

Definition B.1 (Total variation distance). The *total variation distance* (also known as *statistical distance*) between P and Q is defined as

$$\mathbf{D}_{\text{TV}}(P\|Q) := \frac{1}{2} \mathbb{E}_{\omega \sim Q} \left[\left| \frac{P(\omega)}{Q(\omega)} - 1 \right| \right].$$

Definition B.2 (Chi-square divergence). The *Chi-square divergence* (also known as Pearson's χ^2 divergence) between P and Q is defined as

$$\mathbf{D}_{\chi^2}(P\|Q) := \mathbb{E}_{\omega \sim Q} \left[\left(\frac{P(\omega)}{Q(\omega)} - 1 \right)^2 \right].$$

Definition B.3 (KL divergence). The *Kullback–Leibler divergence* (or *KL divergence*) between P and Q is defined as

$$\mathbf{D}_{\text{KL}}(P\|Q) := \mathbb{E}_{\omega \sim P} \left[\ln \left(\frac{P(\omega)}{Q(\omega)} \right) \right] = \mathbb{E}_{\omega \sim Q} \left[\frac{P(\omega)}{Q(\omega)} \cdot \ln \left(\frac{P(\omega)}{Q(\omega)} \right) \right].$$

Not all divergences are f -divergences. For example:

Definition B.4 ($(\delta$ -Approximate) max divergence). Given $\delta \geq 0$, the δ -*approximate max divergence* between P and Q is defined as

$$\mathbf{D}_{\infty}^{\delta}(P\|Q) := \sup_{\omega \subseteq \Omega \mid P(\omega) \geq \delta} \left(\ln \left(\frac{P(\omega) - \delta}{Q(\omega)} \right) \right).$$

The case of $\delta = 0$ is simply called the max divergence and is denoted by \mathbf{D}_{∞} .

These two last definitions are generalized by the following one:

Definition B.5 (Rényi divergence). Given $\alpha \in [0, \infty]$, the *Rényi divergence* between two probability distributions P and Q is defined as

$$\mathbf{D}_{\alpha}(P\|Q) := \frac{1}{\alpha - 1} \ln \left(\mathbb{E}_{\omega \sim P} \left[\left(\frac{P(\omega)}{Q(\omega)} \right)^{\alpha - 1} \right] \right).$$

The cases of $\alpha = 0, 1$, and ∞ are defined using the limit and converge to $-\ln \left(\Pr_{\omega \sim Q} [\omega > 0] \right)$, the KL divergence, and the max divergence, respectively.⁷

⁷This definition generalizes the original one introduced by Alfréd Rényi to include the range $[0, 1)$, which will become useful later in Theorem B.12. For more details, see Van Erven and Harremoës [2014]

B.2 Dissimilarity measures

In this section we present a generalization of the Rényi divergence and the max divergence. We start by recalling two useful identities that are a direct result of the definitions.

Fact B.6. Given $\epsilon, \delta \geq 0$ and two probability distributions P, Q , we have $\mathbf{D}_\infty^\delta(P||Q) \leq \epsilon$ if and only if for any $\omega \subseteq \Omega$, $\Pr_{\omega \sim P}[\omega \in \omega] \leq e^\epsilon \Pr_{\omega \sim Q}[\omega \in \omega] + \delta$.

Fact B.7. Given $\phi \geq 0$, $\alpha \geq 1$, and two probability distributions P, Q , we have $\mathbf{D}_\alpha(P||Q) \leq \phi$ if and only if

$$\mathbb{E}_{\omega \sim P} \left[\exp \left((\alpha - 1) \ln \left(\frac{P(\omega)}{Q(\omega)} \right) \right) \right] \leq e^{(\alpha-1)\phi}.$$

Motivated by these two identities, we generalize these divergences by replacing the scalar bounding the divergence by a function, leading to a more “dynamic” measure that can no longer be regarded a divergence, hence the names.

We first generalize the notion of max divergence.

Definition B.8 (Max dissimilarity). Given $\delta \geq 0$, and a non-negative function $\epsilon : \Omega \rightarrow \mathbb{R}^+$, we say the (δ -approximate) max dissimilarity between P and Q is bounded by ϵ if for any $\omega \subseteq \Omega$,

$$\int_{\omega \in \omega} P(\omega) d\omega \leq \int_{\omega \in \omega} e^{\epsilon(\omega)} Q(\omega) d\omega + \delta.$$

We denote this fact by $P \leq_{(\epsilon, \delta)}^+ Q$.

Similarly, we denote $P \leq_{(\epsilon, \delta)}^- Q$ if for any $\omega \subseteq \Omega$,

$$\int_{\omega \in \omega} e^{-\epsilon(\omega)} P(\omega) d\omega \leq \int_{\omega \in \omega} Q(\omega) d\omega + \delta.$$

Next we generalize the notion of Rényi divergence. This dissimilarity measure serves as the mathematical basis for the new stability measure that we introduce (Definition 4.2).

Definition B.9 (Rényi dissimilarity). Given $\alpha \geq 0$ and a non-negative function $\varphi : \Omega \rightarrow \mathbb{R}^+$, we say the α -Rényi dissimilarity between P and Q is bounded by φ if

$$\mathbb{E}_{\omega \sim P} \left[\exp \left((\alpha - 1) \left(\ln \left(\frac{P(\omega)}{Q(\omega)} \right) - \varphi(\omega) \right) \right) \right] \leq 1.$$

The following theorem states several implications between the various max dissimilarity definitions.

Theorem B.10. Given $\delta \geq 0$, a non-negative function $\epsilon : \Omega \rightarrow \mathbb{R}^+$, and two distributions P, Q :

1. $\Pr_{\omega \sim P} \left[\ln \left(\frac{P(\omega)}{Q(\omega)} \right) > \epsilon(\omega) \right] \leq \delta$ implies $P \leq_{(\epsilon, \delta)}^+ Q$,
2. $P \leq_{(\epsilon, \delta)}^+ Q$ implies $P \leq_{(\epsilon, \delta)}^- Q$,
3. $P \leq_{(\epsilon, \delta)}^+ Q$ implies $\Pr_{\omega \sim P} \left[\ln \left(\frac{P(\omega)}{Q(\omega)} \right) > \epsilon(\omega) + a \right] < \frac{e^a \delta}{e^a - 1}$ for any $a > 0$, and
4. $Q \leq_{(\epsilon, \delta)}^- P$ implies $\Pr_{\omega \sim P} \left[\ln \left(\frac{Q(\omega)}{P(\omega)} \right) > \epsilon(\omega) + a \right] < \frac{\delta}{e^a - 1}$ for any $a > 0$.

Proof. Part 1: Denoting $B_\epsilon(P, Q) := \left\{ \omega \in \Omega \mid \ln \left(\frac{P(\omega)}{Q(\omega)} \right) > \epsilon(\omega) \right\}$ we have

$$\begin{aligned}
\int_{\omega \in \omega} P(\omega) d\omega &= \int_{\omega \in \omega \setminus B_\varepsilon(P, Q)} P(\omega) d\omega + \int_{\omega \in \omega \cap B_\varepsilon(P, Q)} P(\omega) d\omega \\
&\stackrel{(a)}{\leq} \int_{\omega \in \omega \setminus B_\varepsilon(P, Q)} e^{\varepsilon(\omega)} Q(\omega) d\omega + \int_{\omega \in B_\varepsilon(P, Q)} P(\omega) d\omega \\
&= \int_{\omega \in \omega \setminus B_\varepsilon(P, Q)} e^{\varepsilon(\omega)} Q(\omega) d\omega + \Pr_{\omega \sim P} \left[\ln \left(\frac{P(\omega)}{Q(\omega)} \right) > \varepsilon(\omega) \right] \\
&\stackrel{(b)}{\leq} \int_{\omega \in \omega} e^{\varepsilon(\omega)} Q(\omega) d\omega + \delta,
\end{aligned}$$

where (a) results from the definition of $B_\varepsilon(P, Q)$ and (b) from the assumption that $\Pr_{\omega \sim P} \left[\ln \left(\frac{P(\omega)}{Q(\omega)} \right) > \varepsilon(\omega) \right] \leq \delta$.

Part 2: Denoting $\delta(\omega) := \max \{P(\omega) - e^{\varepsilon(\omega)} Q(\omega), 0\}$, which implies $P(\omega) \leq e^{\varepsilon(\omega)} Q(\omega) + \delta(\omega)$, we have

$$\begin{aligned}
\int_{\omega \in \omega} e^{-\varepsilon(\omega)} P(\omega) d\omega &\stackrel{(a)}{\leq} \int_{\omega \in \omega} e^{-\varepsilon(\omega)} \left(e^{\varepsilon(\omega)} Q(\omega) + \delta(\omega) \right) d\omega \\
&= \int_{\omega \in \omega} \left(Q(\omega) + e^{-\varepsilon(\omega)} \cdot \delta(\omega) \right) d\omega \\
&\stackrel{(b)}{\leq} \int_{\omega \in \omega} Q(\omega) d\omega + \int_{\omega \in \omega} \delta(\omega) d\omega \\
&\stackrel{(c)}{\leq} \int_{\omega \in \omega} Q(\omega) d\omega + \delta
\end{aligned}$$

where (a) results from the fact the inequality holds for any element $\omega \in \Omega$, (b) from the non-negativity of ε , and (c) from the assumption that $P \leq_{(\varepsilon, \delta)}^+ Q$.

Part 3: Denoting $\epsilon_Q^+ := \ln \left(\mathbb{E}_{\omega \sim Q} [e^{\varepsilon(X)}] \right)$ and $Q_P^+(\omega) := e^{\varepsilon(\omega) - \epsilon_Q^+} Q(\omega)$, we notice that from the definition Q_P^+ is a valid distribution, since

$$\int_{\omega \in \Omega} Q_P^+(\omega) d\omega = e^{-\epsilon_Q^+} \int_{\omega \in \Omega} e^{\varepsilon(\omega)} Q(\omega) d\omega = \frac{1}{\mathbb{E}_{\omega \sim Q} [e^{\varepsilon(X)}]} \mathbb{E}_{\omega \sim Q} [e^{\varepsilon(X)}] = 1.$$

We first transform the $P \leq_{(\varepsilon, \delta)}^+ Q$ bound into a bound on the max divergence between P and Q_P^+ .

$$\Pr_{\omega \sim P} [\omega \in \omega] \stackrel{(a)}{\leq} \int_{\omega \in \omega} e^{\varepsilon(\omega)} Q(\omega) d\omega + \delta \stackrel{(b)}{=} e^{\epsilon_Q^+} \int_{\omega \in \omega} Q_P^+(\omega) d\omega + \delta = e^{\epsilon_Q^+} \Pr_{\omega \sim Q_P^+} [\omega \in \omega] + \delta,$$

where (a) results from the assumption that $P \leq_{(\varepsilon, \delta)}^+ Q$ and (b) from the definition of $Q_P^+(\omega)$.

Next we denote $B_\varepsilon^a(P, Q) := \left\{ \omega \in \Omega \mid \ln \left(\frac{P(\omega)}{Q(\omega)} \right) > \varepsilon(\omega) + a \right\}$, and notice that for any $\omega \in B_\varepsilon^a(P, Q)$ we have $P(\omega) > e^{\varepsilon(\omega) + a} Q(\omega) = e^{\epsilon_Q^+ + a} Q_P^+(\omega)$, so

$$\begin{aligned}
\Pr_{\omega \sim P} [\omega \in B_\varepsilon^a(P, Q)] &\stackrel{(a)}{\leq} e^{\varepsilon_Q^+} \Pr_{\omega \sim Q_P^+} [\omega \in B_\varepsilon^a(P, Q)] + \delta \\
&= e^{\varepsilon_Q^+} \int_{\omega \in B_\varepsilon^a(P, Q)} Q_P^+(\omega) d\omega + \delta \\
&\stackrel{(b)}{<} e^{\varepsilon_Q^+} \int_{\omega \in B_\varepsilon^a(P, Q)} e^{-(\varepsilon_Q^+ + a)} P(\omega) d\omega + \delta \\
&= e^{-a} \Pr_{\omega \sim P} [\omega \in B_\varepsilon^a(P, Q)] + \delta
\end{aligned}$$

where (a) results from the previous bound and (b) from the fact the inequality holds for any element $\omega \in \Omega$.

Reordering the terms we get

$$\Pr_{\omega \sim Q} \left[\ln \left(\frac{P(\omega)}{Q(\omega)} \right) > \varepsilon(\omega) + a \right] = \Pr_{\omega \sim P} [\omega \in B_\varepsilon^a(P, Q)] < \frac{e^a \delta}{e^a - 1}.$$

Part 4: Denoting $\varepsilon_Q^- := -\ln \left(\mathbb{E}_{\omega \sim Q} [e^{-\varepsilon(\omega)}] \right)$ and $Q_P^-(\omega) := e^{-\varepsilon(\omega) + \varepsilon_Q^-} Q(\omega)$, we notice that from the definition Q_P^- is a valid distribution, since

$$\int_{\omega \in \Omega} Q_P^-(\omega) d\omega = e^{\varepsilon_Q^-} \int_{\omega \in \Omega} e^{-\varepsilon(\omega)} Q(\omega) d\omega = \frac{1}{\mathbb{E}_{\omega \sim Q} [e^{-\varepsilon(\omega)}]} \mathbb{E}_{\omega \sim Q} [e^{-\varepsilon(\omega)}] = 1.$$

We first transform the $Q \leq_{(\varepsilon, \delta)}^- P$ bound into a bound on the max divergence between Q_P^- and P .

$$\Pr_{\omega \sim Q_P^-} [\omega \in \omega] \stackrel{(a)}{=} e^{\varepsilon_Q^-} \int_{\omega \in \omega} e^{-\varepsilon(\omega)} Q(\omega) d\omega \stackrel{(b)}{\leq} e^{\varepsilon_Q^-} \left(\int_{\omega \in \omega} P(\omega) + \delta \right) = e^{\varepsilon_Q^-} \Pr_{\omega \sim P} [\omega \in \omega] + e^{\varepsilon_Q^-} \delta$$

where (a) results from the definition of Q_P^- and (b) from the assumption that $Q \leq_{(\varepsilon, \delta)}^- P$.

Next we denote $B_\varepsilon^a(Q, P) := \left\{ \omega \in \Omega \mid \ln \left(\frac{Q(\omega)}{P(\omega)} \right) > \varepsilon(\omega) + a \right\}$, and notice that for any $\omega \in B_\varepsilon^a(Q, P)$ we have $P(\omega) < e^{-(\varepsilon(\omega) + a)} Q(\omega) = e^{-(\varepsilon_Q^- + a)} Q_P^-(\omega)$, so

$$\begin{aligned}
\Pr_{\omega \sim P} [\omega \in B_\varepsilon^a(Q, P)] &= \int_{\omega \in B_\varepsilon^a(Q, P)} P(\omega) d\omega \\
&\stackrel{(a)}{<} \int_{\omega \in B_\varepsilon^a(Q, P)} e^{-(\varepsilon_Q^- + a)} Q_P^-(\omega) d\omega \\
&= e^{-(\varepsilon_Q^- + a)} \Pr_{\omega \sim Q_P^-} [\omega \in B_\varepsilon^a(Q, P)] \\
&\stackrel{(b)}{\leq} e^{-(\varepsilon_Q^- + a)} \left(e^{\varepsilon_Q^-} \Pr_{\omega \sim P} [\omega \in B_\varepsilon^a(Q, P)] + e^{\varepsilon_Q^-} \delta \right) \\
&= e^{-a} \left(\Pr_{\omega \sim P} [\omega \in B_\varepsilon^a(Q, P)] + \delta \right)
\end{aligned}$$

where (a) results from the fact the inequality holds for any element $\omega \in \Omega$ and (b) from the previous bound.

Reordering the terms we get

$$\Pr_{\omega \sim P} \left[\ln \left(\frac{Q(\omega)}{P(\omega)} \right) > \varepsilon(\omega) + a \right] = \Pr_{\omega \sim P} [\omega \in B_\varepsilon^a(Q, P)] < \frac{\delta}{e^a - 1}.$$

□

Next we prove an important connection between the two dissimilarity measures, starting with a supporting lemma which is a slight variation of Theorem 2.7 in Victor et al. [2007]. This technique is sometimes referred to as the *method of mixtures*.

Lemma B.11. *Given two jointly distributed random variables $(A, B) \sim D$, if $\mathbb{E}_{(A,B) \sim D} \left[\exp \left(\lambda A - \frac{\lambda^2 B^2}{2} \right) \right] \leq 1$ for all $\lambda \in \mathbb{R}$, then for any $\delta, b > 0$,*

$$\Pr_{(A,B) \sim D} \left[|A| > 2\sqrt{\ln \left(\frac{1}{\delta} \right) (B^2 + b^2)} \right] \leq \delta \sqrt{\mathbb{E}_{A,B \sim D} \left[\sqrt{\frac{B^2}{b^2} + 1} \right]}.$$

As Victor et al. [2007] note, at first glance this claim might seem trivial, and one can expect to prove it by simply optimizing over λ to show $\mathbb{E}_{A,B \sim D} \left[\exp \left(\frac{A}{2B^2} \right) \right] \leq 1$ and combining with Markov's inequality. The problem is, the optimal λ depends on the value of B which is a random variables, while λ must be set beforehand. To solve this, we define a distribution over λ and use Fubini's theorem.

Proof. Consider a Gaussian distribution over λ with parameters $\mu = 0, \sigma = \frac{1}{b}$.

$$\begin{aligned} 1 &\stackrel{(a)}{\geq} \mathbb{E}_{\lambda \sim \mathcal{N}(0, \frac{1}{b^2})} \left[\mathbb{E}_{A,B \sim D} \left[\exp \left(\lambda A - \frac{\lambda^2 B^2}{2} \right) \right] \right] \\ &\stackrel{(b)}{=} \mathbb{E}_{A,B \sim D} \left[\mathbb{E}_{\lambda \sim \mathcal{N}(0, \frac{1}{b^2})} \left[\exp \left(\lambda A - \frac{\lambda^2 B^2}{2} \right) \right] \right] \\ &= \mathbb{E}_{A,B \sim D} \left[\frac{b}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \exp \left(\lambda A - \frac{\lambda^2 B^2}{2} \right) \exp \left(-\frac{\lambda^2 b^2}{2} \right) d\lambda \right] \\ &= \mathbb{E}_{A,B \sim D} \left[\frac{b}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \exp \left(-\frac{B^2 + b^2}{2} \left(\lambda^2 - 2\lambda \frac{A}{B^2 + b^2} \right) \right) d\lambda \right] \\ &\stackrel{(c)}{=} \mathbb{E}_{A,B \sim D} \left[\frac{b}{\sqrt{B^2 + b^2}} \exp \left(\frac{A^2}{2(B^2 + b^2)} \right) \frac{\sqrt{B^2 + b^2}}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \exp \left(-\frac{B^2 + b^2}{2} \left(\lambda - \frac{A}{B^2 + b^2} \right)^2 \right) d\lambda \right] \\ &\stackrel{(d)}{=} \mathbb{E}_{A,B \sim D} \left[\frac{b}{\sqrt{B^2 + b^2}} \exp \left(\frac{A^2}{2(B^2 + b^2)} \right) \mathbb{E}_{\lambda \sim \mathcal{N}\left(\frac{A}{B^2 + b^2}, \frac{1}{B^2 + b^2}\right)} [1] \right] \\ &= \mathbb{E}_{A,B \sim D} \left[\frac{b}{\sqrt{B^2 + b^2}} \exp \left(\frac{A^2}{2(B^2 + b^2)} \right) \right]. \end{aligned}$$

where (a) results from the fact the inequality $\mathbb{E}_{A,B \sim D} \left[\exp \left(\lambda A - \frac{\lambda^2 B^2}{2} \right) \right] \leq 1$ holds for all λ , which implies it holds for the expectation taken with respect to any distribution over λ as well, (b) results from Fubini's theorem, (c) is simply completing the square, and (d) from the definition of the Gaussian distribution.

Using this bound we get

$$\begin{aligned}
& \Pr_{A,B,\sim D} \left[|A| > 2\sqrt{\ln\left(\frac{1}{\delta}\right)(B^2 + b^2)} \right] \\
&= \Pr_{A,B,\sim D} \left[\exp\left(\frac{A^2}{4(B^2 + b^2)}\right) > \frac{1}{\delta} \right] \\
&\stackrel{(a)}{\leq} \delta \mathbb{E}_{A,B,\sim D} \left[\exp\left(\frac{A^2}{4(B^2 + b^2)}\right) \right] \\
&= \delta \mathbb{E}_{A,B,\sim D} \left[\sqrt{\frac{\sqrt{B^2 + b^2}}{b}} \sqrt{\frac{b}{\sqrt{B^2 + b^2}}} \exp\left(\frac{A^2}{4(B^2 + b^2)}\right) \right] \\
&\stackrel{(b)}{\leq} \delta \sqrt{\mathbb{E}_{A,B,\sim D} \left[\frac{\sqrt{B^2 + b^2}}{b} \right] \mathbb{E}_{A,B,\sim D} \left[\frac{b}{\sqrt{B^2 + b^2}} \exp\left(\frac{A^2}{2(B^2 + b^2)}\right) \right]} \\
&\stackrel{(c)}{\leq} \delta \sqrt{\mathbb{E}_{A,B,\sim D} \left[\sqrt{\frac{B^2}{b^2} + 1} \right]},
\end{aligned}$$

where (a) results from Markov's inequality, (b) from the Cauchy-Schwarz inequality, and (c) from the previous inequality. \square

We can now prove the implication Theorem.

Theorem B.12. *Given a non-negative function $\varphi : \Omega \rightarrow \mathbb{R}^+$ and two distributions P, Q , if the α -Rényi dissimilarity between P and Q is bounded by $\alpha\varphi$ for any $\alpha \geq 0$, then*

$$\Pr_{\omega \sim P} \left[\left| \ln\left(\frac{P(\omega)}{Q(\omega)}\right) \right| > \varepsilon(\omega) \right] \leq \delta \sqrt{\mathbb{E}_{\omega \sim P} \left[\sqrt{\frac{\varphi(\omega)}{\phi} + 1} \right]} \text{ for any } \phi, \delta > 0, \text{ where } \varepsilon(\omega) := \varphi(\omega) + 2\sqrt{2\ln\left(\frac{1}{\delta}\right)(\varphi(\omega) + \phi)}.$$

Proof. Denoting $A := \ln\left(\frac{P(\omega)}{Q(\omega)}\right) - \varphi(\omega)$ and $B := \sqrt{2\varphi(\omega)}$, for any $\lambda \geq 0$ we denote $\alpha = \lambda + 1$ and from the assumption we get that

$$\mathbb{E}_{A,B \sim D} \left[\exp\left(\lambda A - \frac{\lambda^2 B^2}{2}\right) \right] \stackrel{(a)}{\leq} \mathbb{E}_{\omega \sim P} \left[\exp\left((\alpha - 1) \left(\ln\left(\frac{P(\omega)}{Q(\omega)}\right) - \alpha\varphi(\omega)\right)\right) \right] \stackrel{(b)}{\leq} 1$$

For any $\lambda < 0$ we denote $\alpha = -\lambda$ and from the assumption we get that

$$\begin{aligned}
\mathbb{E}_{A,B \sim D} \left[\exp\left(\lambda A - \frac{\lambda^2 B^2}{2}\right) \right] &\stackrel{(a)}{\leq} \mathbb{E}_{\omega \sim P} \left[\exp\left(\lambda \left(\ln\left(\frac{P(\omega)}{Q(\omega)}\right) - (\lambda + 1)\varphi(\omega)\right)\right) \right] \\
&\stackrel{(c)}{=} \mathbb{E}_{\omega \sim P} \left[\exp\left(\alpha \left(\ln\left(\frac{Q(\omega)}{P(\omega)}\right) - (\alpha - 1)\varphi(\omega)\right)\right) \right] \\
&= \mathbb{E}_{\omega \sim P} \left[\left(\frac{Q(\omega)}{P(\omega)}\right)^\alpha \exp(-\alpha(\alpha - 1)\varphi(\omega)) \right] \\
&\stackrel{(d)}{=} \mathbb{E}_{\omega \sim Q} \left[\left(\frac{Q(\omega)}{P(\omega)}\right)^{\alpha-1} \exp(-\alpha(\alpha - 1)\varphi(\omega)) \right] \\
&= \mathbb{E}_{\omega \sim Q} \left[\exp\left((\alpha - 1) \left(\ln\left(\frac{Q(\omega)}{P(\omega)}\right) - \alpha\varphi(\omega)\right)\right) \right] \\
&\stackrel{(b)}{\leq} 1
\end{aligned}$$

where (a) results from the definition of λ , (b) from the Rényi dissimilarity assumption, (c) from the fact $\ln\left(\frac{a}{b}\right) = -\ln\left(\frac{b}{a}\right)$, and (d) is a renormalization step.

Combining the two with Lemma B.11 and setting $\phi := 2b^2$ we get

$$\begin{aligned}
\Pr_{\omega \sim P} \left[\left| \ln \left(\frac{P(\omega)}{Q(\omega)} \right) \right| > \varepsilon(\omega) \right] &= \Pr_{\omega \sim P} \left[\left| \ln \left(\frac{P(\omega)}{Q(\omega)} \right) \right| > \varphi(\omega) + 2\sqrt{2\ln \left(\frac{1}{\delta} \right) (\varphi(\omega) + \phi)} \right] \\
&\stackrel{(a)}{\leq} \Pr_{\omega \sim P} \left[\left| \ln \left(\frac{P(\omega)}{Q(\omega)} \right) - \varphi(\omega) \right| > +2\sqrt{2\ln \left(\frac{1}{\delta} \right) (\varphi(\omega) + \phi)} \right] \\
&= \Pr_{(A,B) \sim D} \left[|A| > 2\sqrt{\ln \left(\frac{1}{\delta} \right) (B^2 + b^2)} \right] \\
&\stackrel{(b)}{\leq} \delta \sqrt{\mathbb{E}_{A,B \sim D} \left[\sqrt{\frac{B^2}{b^2} + 1} \right]} \\
&= \delta \sqrt{\mathbb{E}_{\omega \sim P} \left[\sqrt{\frac{\varphi(\omega)}{\phi} + 1} \right]},
\end{aligned}$$

where (a) results from the triangle inequality and (b) from Lemma B.11. \square

B.3 Convexity

Next we prove some convexity results for max and Rényi dissimilarities.

Theorem B.13 (Max dissimilarity convexity). *Given $\lambda \in (0, 1)$, $\delta_i \geq 0$, distributions P_i, Q_i , and non-negative functions $\varepsilon_i : \omega \rightarrow \mathbb{R}^+$ for $i \in \{0, 1\}$, denote $P_\lambda := \lambda P_0 + (1 - \lambda) P_1$, $Q_\lambda := \lambda Q_0 + (1 - \lambda) Q_1$, $\varepsilon_\lambda(\omega) := \ln(\lambda e^{\varepsilon_0(\omega)} + (1 - \lambda) e^{\varepsilon_1(\omega)})$, and $\delta_\lambda := \lambda \delta_0 + (1 - \lambda) \delta_1$.*

1. *Joint quasi-convexity: If $P_i \leq_{(\varepsilon_i, \delta_i)}^+ Q_i$ for $i \in \{0, 1\}$, then $P_\lambda \leq_{(\varepsilon_{\max}, \delta_\lambda)}^+ Q_\lambda$, where $\varepsilon_{\max}(\omega) := \max\{\varepsilon_0(\omega), \varepsilon_1(\omega)\}$*
2. *Left-hand log-convexity: If $P_i \leq_{(\varepsilon_i, \delta_i)}^+ Q$ for $i \in \{0, 1\}$, then $P_\lambda \leq_{(\varepsilon_\lambda, \delta_\lambda)}^+ Q$*
3. *Right-hand log-convexity: If $P \leq_{(\varepsilon_i, \delta_i)}^- Q_i$ for $i \in \{0, 1\}$, then $P \leq_{(\varepsilon_\lambda, \delta_\lambda)}^- Q_\lambda$*

Proof. Part 1: For any $\omega \subseteq \Omega$,

$$\begin{aligned}
\int_{\omega \in \omega} P_\lambda(\omega) d\omega &\stackrel{(a)}{=} \int_{\omega \in \omega} (\lambda P_0(\omega) + (1 - \lambda) P_1(\omega)) d\omega \\
&= \lambda \int_{\omega \in \omega} P_0(\omega) d\omega + (1 - \lambda) \int_{\omega \in \omega} P_1(\omega) d\omega \\
&\stackrel{(b)}{\leq} \lambda \left(\int_{\omega \in \omega} e^{\varepsilon_0(\omega)} Q_0(\omega) d\omega + \delta_0 \right) + (1 - \lambda) \left(\int_{\omega \in \omega} e^{\varepsilon_1(\omega)} Q_1(\omega) d\omega + \delta_1 \right) \\
&\stackrel{(c)}{\leq} \lambda \left(\int_{\omega \in \omega} e^{\varepsilon_{\max}(\omega)} Q_0(\omega) d\omega + \delta_0 \right) + (1 - \lambda) \left(\int_{\omega \in \omega} e^{\varepsilon_{\max}(\omega)} Q_1(\omega) d\omega + \delta_1 \right) \\
&\stackrel{(d)}{=} \int_{\omega \in \omega} e^{\varepsilon_{\max}(\omega)} Q_\lambda(\omega) d\omega + \delta_\lambda
\end{aligned}$$

where (a) results from the definition of P_λ , (b) from the assumption $P_i \leq_{(\varepsilon_i, \delta_i)}^+ Q_i$, (c) from the definitions of ε_{\max} and δ_λ , and (d) from the definition of Q_λ .

Part 2: For any $\omega \subseteq \Omega$,

$$\begin{aligned}
\int_{\omega \in \omega} P_\lambda(\omega) d\omega &\stackrel{(a)}{=} \int_{\omega \in \omega} (\lambda P_0(\omega) + (1 - \lambda) P_1(\omega)) d\omega \\
&= \lambda \int_{\omega \in \omega} P_0(\omega) d\omega + (1 - \lambda) \int_{\omega \in \omega} P_1(\omega) d\omega \\
&\stackrel{(b)}{\leq} \lambda \left(\int_{\omega \in \omega} e^{\varepsilon_0(\omega)} Q(\omega) d\omega + \delta_0 \right) + (1 - \lambda) \left(\int_{\omega \in \omega} e^{\varepsilon_1(\omega)} Q(\omega) d\omega + \delta_1 \right) \\
&\stackrel{(c)}{=} \int_{\omega \in \omega} \left(\lambda e^{\varepsilon_0(\omega)} + (1 - \lambda) e^{\varepsilon_1(\omega)} \right) Q(\omega) d\omega + \delta_\lambda \\
&\stackrel{(d)}{=} \int_{\omega \in \omega} e^{\varepsilon_\lambda(\omega)} Q(\omega) d\omega + \delta_\lambda
\end{aligned}$$

where (a) results from the definition of P_λ , (b) from the assumption $P_i \leq_{(\varepsilon_i, \delta_i)}^+ Q$, (c) from the definition of δ_λ , and (d) from the definition of ε_λ .

Part 3: For any $\omega \subseteq \Omega$,

$$\begin{aligned}
\int_{\omega \in \omega} e^{-\varepsilon_\lambda(\omega)} P(\omega) d\omega &\stackrel{(a)}{=} \int_{\omega \in \omega} \frac{1}{\lambda e^{\varepsilon_0(\omega)} + (1 - \lambda) e^{\varepsilon_1(\omega)}} P(\omega) d\omega \\
&\stackrel{(b)}{\leq} \int_{\omega \in \omega} \left(\lambda e^{-\varepsilon_0(\omega)} + (1 - \lambda) e^{-\varepsilon_1(\omega)} \right) P(\omega) d\omega \\
&= \lambda \left(\int_{\omega \in \omega} e^{-\varepsilon_0(\omega)} P(\omega) d\omega \right) + (1 - \lambda) \left(\int_{\omega \in \omega} e^{-\varepsilon_1(\omega)} P(\omega) d\omega \right) \\
&\stackrel{(c)}{\leq} \lambda \left(\int_{\omega \in \omega} Q_0(\omega) d\omega + \delta_0 \right) + (1 - \lambda) \left(\int_{\omega \in \omega} Q_1(\omega) d\omega + \delta_1 \right) \\
&\stackrel{(d)}{=} \int_{\omega \in \omega} Q_\lambda(\omega) d\omega + \delta_\lambda
\end{aligned}$$

where (a) results from the definition of ε_λ , (b) from Jensen's inequality for the convex function $\frac{1}{x}$, (c) from the assumption $P \leq_{(\varepsilon_i, \delta_i)}^- Q_i$, and (d) from the definitions of Q_λ and δ_λ . \square

Theorem B.14 (Rényi dissimilarity convexity). *Given $\alpha \geq 1$, $\lambda \in (0, 1)$, and non-negative functions $\varphi_i : \omega \rightarrow \mathbb{R}^+$ for $i \in \{0, 1\}$, denote $P_\lambda := \lambda P_0 + (1 - \lambda) P_1$, $Q_\lambda := \lambda Q_0 + (1 - \lambda) Q_1$, and $\varphi_\lambda(\omega) := \lambda \varphi_0(\omega) + (1 - \lambda) \varphi_1(\omega)$.*

1. *Joint quasi-convexity: If the α -Rényi dissimilarity between P_i and Q_i is bounded by φ_i for $i \in \{0, 1\}$, then the α -Rényi dissimilarity between P_λ and Q_λ is bounded by φ_{\max} , where $\varphi_{\max}(\omega) := \max\{\varphi_0(\omega), \varphi_1(\omega)\}$*
2. *Right-hand convexity: If the α -Rényi dissimilarity between P and Q_i is bounded by φ_i for $i \in \{0, 1\}$, then the α -Rényi dissimilarity between P and Q_λ is bounded by φ_λ*

Proof. The case of $\alpha = 1$ is trivial for both parts, so we focus on the $\alpha > 1$ case.

Part 1: We first notice that

$$\begin{aligned}
P_\lambda(\omega) &= \lambda P_0(\omega) + (1-\lambda) P_1(\omega) \\
&= \lambda (Q_0(\omega))^{\frac{\alpha-1}{\alpha}} \frac{P_0(\omega)}{(Q_0(\omega))^{\frac{\alpha-1}{\alpha}}} + (1-\lambda) (Q_1(\omega))^{\frac{\alpha-1}{\alpha}} \frac{P_1(\omega)}{(Q_1(\omega))^{\frac{\alpha-1}{\alpha}}} \\
&\stackrel{(a)}{\leq} \left(\lambda \left((Q_0(\omega))^{\frac{\alpha-1}{\alpha}} \right)^{\frac{\alpha}{\alpha-1}} + (1-\lambda) \left((Q_1(\omega))^{\frac{\alpha-1}{\alpha}} \right)^{\frac{\alpha}{\alpha-1}} \right)^{\frac{\alpha-1}{\alpha}} \\
&\quad \cdot \left(\lambda \left(\frac{P_0(\omega)}{(Q_0(\omega))^{\frac{\alpha-1}{\alpha}}} \right)^\alpha + (1-\lambda) \left(\frac{P_1(\omega)}{(Q_1(\omega))^{\frac{\alpha-1}{\alpha}}} \right)^\alpha \right)^{\frac{1}{\alpha}} \\
&= (\lambda Q_0(\omega) + (1-\lambda) Q_1(\omega))^{\frac{\alpha-1}{\alpha}} \cdot \left(\lambda \frac{(P_0(\omega))^\alpha}{(Q_0(\omega))^{\alpha-1}} + (1-\lambda) \frac{(P_1(\omega))^\alpha}{(Q_1(\omega))^{\alpha-1}} \right)^{\frac{1}{\alpha}} \\
&= Q_\lambda(\omega) \cdot \left(\lambda \frac{P_0(\omega)}{Q_\lambda(\omega)} \left(\frac{P_0(\omega)}{Q_0(\omega)} \right)^{\alpha-1} + (1-\lambda) \frac{P_1(\omega)}{Q_\lambda(\omega)} \left(\frac{P_1(\omega)}{Q_1(\omega)} \right)^{\alpha-1} \right)^{\frac{1}{\alpha}}
\end{aligned}$$

where (a) results from the generalized version of Hölder's inequality, $\mathbb{E}[X \cdot Y] \leq \left(\mathbb{E}[X^p] \right)^{\frac{1}{p}} \cdot \left(\mathbb{E}[Y^q] \right)^{\frac{1}{q}}$ for any $p, q > 0$ such that $\frac{1}{p} + \frac{1}{q} = 1$.

Using this bound we get

$$\begin{aligned}
&\mathbb{E}_{\omega \sim P_\lambda} \left[\exp \left((\alpha-1) \left(\ln \left(\frac{P_\lambda(\omega)}{Q_\lambda(\omega)} \right) - \varphi_{\max}(\omega) \right) \right) \right] \\
&= \mathbb{E}_{\omega \sim P_\lambda} \left[\left(\frac{P_\lambda(\omega)}{Q_\lambda(\omega)} \right)^{\alpha-1} \exp(-(\alpha-1) \varphi_{\max}(\omega)) \right] \\
&= \mathbb{E}_{\omega \sim Q_\lambda} \left[\left(\frac{P_\lambda(\omega)}{Q_\lambda(\omega)} \right)^\alpha \exp(-(\alpha-1) \varphi_{\max}(\omega)) \right] \\
&\stackrel{(a)}{\leq} \mathbb{E}_{\omega \sim Q_\lambda} \left[\left(\lambda \frac{P_0(\omega)}{Q_\lambda(\omega)} \left(\frac{P_0(\omega)}{Q_0(\omega)} \right)^{\alpha-1} + (1-\lambda) \frac{P_1(\omega)}{Q_\lambda(\omega)} \left(\frac{P_1(\omega)}{Q_1(\omega)} \right)^{\alpha-1} \right) \exp(-(\alpha-1) \varphi_{\max}(\omega)) \right] \\
&= \lambda \mathbb{E}_{\omega \sim Q_\lambda} \left[\frac{P_0(\omega)}{Q_\lambda(\omega)} \left(\frac{P_0(\omega)}{Q_0(\omega)} \right)^{\alpha-1} \exp(-(\alpha-1) \varphi_{\max}(\omega)) \right] \\
&\quad + (1-\lambda) \mathbb{E}_{\omega \sim Q_\lambda} \left[\frac{P_1(\omega)}{Q_\lambda(\omega)} \left(\frac{P_1(\omega)}{Q_1(\omega)} \right)^{\alpha-1} \exp(-(\alpha-1) \varphi_{\max}(\omega)) \right] \\
&\stackrel{(b)}{=} \lambda \mathbb{E}_{\omega \sim P_0} \left[\exp \left((\alpha-1) \left(\ln \left(\frac{P_0(\omega)}{Q_0(\omega)} \right) - \varphi_{\max}(\omega) \right) \right) \right] \\
&\quad + (1-\lambda) \mathbb{E}_{\omega \sim P_1} \left[\exp \left((\alpha-1) \left(\ln \left(\frac{P_1(\omega)}{Q_1(\omega)} \right) - \varphi_{\max}(\omega) \right) \right) \right] \\
&\stackrel{(c)}{\leq} \lambda \mathbb{E}_{\omega \sim P_0} \left[\exp \left((\alpha-1) \left(\ln \left(\frac{P_0(\omega)}{Q_0(\omega)} \right) - \varphi_0(\omega) \right) \right) \right] \\
&\quad + (1-\lambda) \mathbb{E}_{\omega \sim P_1} \left[\exp \left((\alpha-1) \left(\ln \left(\frac{P_1(\omega)}{Q_1(\omega)} \right) - \varphi_1(\omega) \right) \right) \right] \\
&\stackrel{(d)}{\leq} 1
\end{aligned}$$

where (a) results from the previous inequality, (b) is a renormalization step, (c) from the the definition of φ_{\max} , and (d) from the Rényi dissimilarity assumption.

Part 2:

$$\begin{aligned}
& \mathbb{E}_{\omega \sim P} \left[\exp \left((\alpha - 1) \left(\ln \left(\frac{P(\omega)}{Q_\lambda(\omega)} \right) - \varphi_\lambda(\omega) \right) \right) \right] \\
&= \mathbb{E}_{\omega \sim P} \left[\exp \left((\alpha - 1) \left(\ln \left(\frac{P(\omega)}{\lambda Q_0(\omega) + (1 - \lambda) Q_1(\omega)} \right) - \varphi_\lambda(\omega) \right) \right) \right] \\
&\stackrel{(a)}{\leq} \mathbb{E}_{\omega \sim P} \left[\exp \left((\alpha - 1) \left(\lambda \ln \left(\frac{P(\omega)}{Q_0(\omega)} \right) + (1 - \lambda) \ln \left(\frac{P(\omega)}{Q_1(\omega)} \right) - \varphi_\lambda(\omega) \right) \right) \right] \\
&\stackrel{(b)}{=} \mathbb{E}_{\omega \sim P} \left[\exp \left(\lambda (\alpha - 1) \left(\ln \left(\frac{P(\omega)}{Q_0(\omega)} \right) - \varphi_0(\omega) \right) + (1 - \lambda) (\alpha - 1) \left(\ln \left(\frac{P(\omega)}{Q_1(\omega)} \right) - \varphi_1(\omega) \right) \right) \right] \\
&\stackrel{(c)}{\leq} \lambda \mathbb{E}_{\omega \sim P} \left[\exp \left((\alpha - 1) \left(\ln \left(\frac{P(\omega)}{Q_0(\omega)} \right) - \varphi_0(\omega) \right) \right) \right] \\
&\quad + (1 - \lambda) \mathbb{E}_{\omega \sim P} \left[\exp \left((\alpha - 1) \left(\ln \left(\frac{P(\omega)}{Q_1(\omega)} \right) - \varphi_1(\omega) \right) \right) \right] \\
&\stackrel{(d)}{\leq} 1
\end{aligned}$$

where (a) results from Jensen's inequality for the convex function $\ln(\frac{1}{x})$, (b) from the definition of φ_λ , (c) from Jensen's inequality for the convex function e^x over the λ weighted combination, and (d) from the Rényi dissimilarity assumption. \square

B.4 Stability measures

The various divergence and dissimilarity measures considered in the previous sections were used to define several stability notions; the divergence (dissimilarity) between the distributions over responses induced by two differing input datasets serves as a measure of the stability of the mechanism producing those responses.

One major application of stability notions is in privacy-preserving mechanisms.

Definition B.15 (Differential privacy [Dwork et al., 2006]). Given $\epsilon, \delta \geq 0$, a mechanism M will be called (ϵ, δ) -differentially private (or DP, for short) if for any two datasets $s, s' \in \mathcal{X}^n$ that differ only in one element, and any query $q \in \mathcal{Q}$, the δ -approximate max divergence between the two distributions defined over \mathcal{R} by $M(s, q)$ and $M(s', q)$ is bounded by ϵ .

As in Definition 2.1, δ can be viewed as a function of ϵ . In this case, δ is essentially a tail bound on the distribution of the privacy loss random variable (Definition 4.1). An alternative stability notion can be based on a bound on the moments of the stability loss.

Definition B.16 (Rényi differential privacy [Mironov, 2017]). Given $\alpha, \phi \geq 0$, a mechanism M will be called (α, ϕ) -Rényi differentially private (or RDP, for short) if for any two datasets $s, s' \in \mathcal{X}^n$ that differ only in one element, and any query $q \in \mathcal{Q}$, the α -Rényi divergence between the two distributions defined over \mathcal{R} by $M(s, q)$ and $M(s', q)$ is bounded by ϕ .

The following two definitions bound the shape of the privacy loss curve.

Definition B.17 (Concentrated differential privacy [Dwork and Rothblum, 2016]). Given $\mu, \sigma \geq 0$, a mechanism M will be called (μ, σ) -concentrated differentially private (or CDP, for short) if for any two datasets $s, s' \in \mathcal{X}^n$ that differ only in one element, and any query $q \in \mathcal{Q}$, the stability loss random variable (Definition 4.1) is a σ^2 -sub-Gaussian random variable (Definition 5.2) with expectation bounded by μ . Formally, $\mathbb{E}_{R \sim M(s, q)} [\ell(s, s'; R)] \leq \mu$ and for any $\alpha \geq 1$ we have

$$\mathbb{E}_{R \sim M(s, q)} \left[\exp \left((\alpha - 1) \left(\ell(s, s'; R) - \mathbb{E}_{R' \sim M(s, q)} [\ell(s, s'; R')] \right) \right) \right] \leq e^{\frac{\alpha^2}{2\sigma^2}}.$$

Definition B.18 (Zero-concentrated differential privacy [Bun and Steinke, 2016]). Given $\phi \geq 0$, a mechanism M will be called ϕ -zero concentrated differentially private (or zCDP, for short) if for any two datasets $s, s' \in \mathcal{X}^n$ that differ only in one element, any query $q \in \mathcal{Q}$, and $\alpha \geq 1$, the α -Rényi

divergence between the two distributions defined over \mathcal{R} by $M(s, q)$ and $M(s', q)$ is bounded by $\alpha\phi$ in both directions.⁸

This definition can be viewed as a special case of the CDP definition, with $\mu = \phi$ and $\sigma^2 = 2\phi$. The comparison is formalized in Lemma 24 in [Bun and Steinke, 2016].

The above notions all hold over *all* pairs of neighboring datasets, which can be achieved by scaling the randomness added by the mechanism to the worst-case sensitivity—that is, to the worst-case change in the output that can be induced by replacing one dataset by a neighboring one. As discussed in detail in Appendix F, to avoid the dependence on the range of the queries, our approach instead focuses on a version of the “local sensitivity,” which, in turn, involves transitioning from divergences to dissimilarity notions.

Definition B.19 (Pairwise concentration; equivalent to Definition 4.2 but restated in terms of divergence and dissimilarity). Given a non-negative function $\varphi : \mathcal{X}^n \times \mathcal{X}^n \times \mathcal{Q} \rightarrow \mathbb{R}^+$ which is symmetric in its first two arguments, a mechanism M will be called φ -Pairwise Concentrated (or PC, for short), if for any $s, s' \in \mathcal{X}^n$, query $q \in \mathcal{Q}$, and $\alpha \geq 0$, the α -Rényi divergence between the two distributions defined over \mathcal{R} by $M(s, q)$ and $M(s', q)$ is bounded by $\alpha\varphi(s, s'; q)$ in both directions.

Given a non-negative function $\varphi : \mathcal{X}^n \times \mathcal{X}^n \times \mathcal{V} \rightarrow \mathbb{R}^+$ which is symmetric in its first two arguments and an analyst A , a mechanism M will be called φ -Pairwise Concentrated with respect to A , if for any $s, s' \in \mathcal{X}^n$ and $\alpha \geq 0$, the α -Rényi dissimilarity between the two distributions defined over \mathcal{V} by $M(s, A)$ and $M(s', A)$ is bounded by $\alpha\varphi(s, s'; \cdot)$ in both directions.

Analogously to how one can generalize the zCDP stability notion by replacing the divergence by a dissimilarity measure, one can also generalize other stability notions such as differential privacy.

Definition B.20 (Pairwise indistinguishability). Given $\delta \geq 0$ and a non-negative function $\varepsilon : \mathcal{X}^n \times \mathcal{X}^n \times \mathcal{Q} \rightarrow \mathbb{R}^+$ which is symmetric in its first two arguments, a mechanism M will be called ε -Pairwise Indistinguishable (or PI, for short), if for any $s, s' \in \mathcal{X}^n$, and any query $q \in \mathcal{Q}$, the max divergence between the two distributions defined over \mathcal{R} by $M(s, q)$ and $M(s', q)$ is bounded by $\varepsilon(s, s'; q)$ in both directions.

Given a non-negative function $\varepsilon : \mathcal{X}^n \times \mathcal{X}^n \times \mathcal{V} \rightarrow \mathbb{R}^+$ which is symmetric in its first two arguments and an analyst A , a mechanism M will be called ε -Pairwise Indistinguishable with respect to A , if for any $s, s' \in \mathcal{X}^n$, the max dissimilarity between the two distributions defined over \mathcal{V} by $M(s, A)$ and $M(s', A)$ is bounded by $\varepsilon(s, s'; \cdot)$ in both directions.

We note that Theorem 4.6 can be stated in terms of pairwise indistinguishability as well.

⁸The original definition includes an additional parameter that is equal to 0 in the case of the Gaussian mechanism, so we omit it here for simplicity.

C Missing parts from Section 3

C.1 Missing proofs

Lemma C.1. *Given a probability distribution P over a product domain $\mathcal{X} \times \mathcal{Y}$ and a function $f : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$, for any $\epsilon, \xi > 0$ we have,*

$$\Pr_{X \sim P_{\mathcal{X}}} \left[\left| \mathbb{E}_{Y \sim P_{\mathcal{Y}|X}} [f(X, Y) | X] \right| > \epsilon + \xi \right] \leq \frac{1}{\xi} \int_{\epsilon}^{\infty} \Pr_{(X, Y) \sim P} [|f(X, Y)| > t] dt.$$

Proof.

$$\begin{aligned} \Pr_{X \sim P_{\mathcal{X}}} \left[\left| \mathbb{E}_{Y \sim P_{\mathcal{Y}|X}} [f(X, Y) | X] \right| > \epsilon + \xi \right] &\stackrel{(a)}{\leq} \Pr_{X \sim P_{\mathcal{X}}} \left[\left[\left| \mathbb{E}_{Y \sim P_{\mathcal{Y}|X}} [f(X, Y) | X] \right| - \epsilon \right]^+ > \xi \right] \\ &\stackrel{(b)}{\leq} \frac{1}{\xi} \mathbb{E}_{X \sim P_{\mathcal{X}}} \left[\left[\left| \mathbb{E}_{Y \sim P_{\mathcal{Y}|X}} [f(X, Y) | X] \right| - \epsilon \right]^+ \right] \\ &\stackrel{(c)}{\leq} \frac{1}{\xi} \mathbb{E}_{X \sim P_{\mathcal{X}}} \left[\left[\mathbb{E}_{Y \sim P_{\mathcal{Y}|X}} [|f(X, Y)| | X] - \epsilon \right]^+ \right] \\ &\stackrel{(d)}{\leq} \frac{1}{\xi} \mathbb{E}_{X \sim P_{\mathcal{X}}} \left[\mathbb{E}_{Y \sim P_{\mathcal{Y}|X}} \left[[|f(X, Y)| - \epsilon]^+ | X \right] \right] \\ &= \frac{1}{\xi} \mathbb{E}_{(X, Y) \sim P} \left[[|f(X, Y)| - \epsilon]^+ \right] \\ &\stackrel{(e)}{=} \frac{1}{\xi} \int_0^{\infty} \Pr_{(X, Y) \sim P} \left[[|f(X, Y)| - \epsilon]^+ > t \right] dt \\ &= \frac{1}{\xi} \int_{\epsilon}^{\infty} \Pr_{(X, Y) \sim P} [|f(X, Y)| > t] dt \end{aligned}$$

where (a) $[x]^+ := \max\{x, 0\}$ is the ReLU function, (b) results from Markov's inequality, (c) from the triangle inequality which implies $|\mathbb{E}[X]| \leq \mathbb{E}[|X|]$ and $[\mathbb{E}[X]]^+ \leq \mathbb{E}[|X|^+]$, (d) from Jensen's inequality for the convex ReLU function, and (e) from the fact that the expectation is taken over a non-negative random variable. \square

Lemma C.2 (Lemma 3.1 restated). *Given a function $\delta : \mathbb{R} \rightarrow [0, 1]$ and an analyst A , if a mechanism M is $(\epsilon, \delta(\epsilon))$ -sample accurate for all $\epsilon > 0$, then M is $(\epsilon, \delta'(\epsilon))$ -posterior accurate for $\delta'(\epsilon) := \inf_{\xi \in (0, \epsilon)} \left(\frac{1}{\xi} \int_{\epsilon - \xi}^{\infty} \delta(t) dt \right)$.*

Proof. For any $v \in \mathcal{V}$, denote $I(v) := \arg \max_{i \in [k]} |r_i - q_i(D^v)|$, and $r_v := r_{I(v)}$, $q_v := q_{I(v)}$. When the view is a random variable V , the corresponding query Q_V and response R_V are also random

variables, and hence we denote them with capital letters.

$$\begin{aligned}
& \Pr_{V \sim D_V} \left[\max_{i \in [k]} |R_i - Q_i(D^V)| > \epsilon + \xi \right] \\
&= \Pr_{V \sim D} \left[|R_V - Q_V(D^V)| > \epsilon + \xi \right] \\
&= \Pr_{V \sim D} \left[\left| \mathbb{E}_{S \sim D(\cdot | V)} [R_V - Q_V(S)] \right| > \epsilon + \xi \right] \\
&\stackrel{(a)}{\leq} \frac{1}{\xi} \int_{\epsilon - \xi}^{\infty} \Pr_{S \sim D^{(n)}, V \sim M(S, A)} [|R_V - Q_V(S)| > u] du \\
&= \frac{1}{\xi} \int_{\epsilon - \xi}^{\infty} \Pr_{S \sim D^{(n)}, V \sim M(S, A)} \left[\max_{i \in [k]} |R_i - Q_i(S)| > u \right] du \\
&\stackrel{(b)}{\leq} \frac{1}{\xi} \int_{\epsilon - \xi}^{\infty} \delta(t) dt
\end{aligned}$$

where (a) results from part one of Lemma C.1 and (b) from the definition of sample accuracy (Definition 2.1). \square

Lemma C.3 (Lemma 3.2 restated). *Given $\eta > 0$, the Gaussian mechanism with noise parameter η that receives k queries is $(\epsilon, \delta(\epsilon))$ -sample accurate for $\delta(\epsilon) := \frac{2k}{\sqrt{\pi}} e^{-\frac{\epsilon^2}{2\eta^2}}$, and $(\epsilon, \delta(\epsilon))$ -posterior accurate for $\delta(\epsilon) := 4k \cdot e^{-\frac{\epsilon^2}{4\eta^2}}$.*

Proof. In the case of a Gaussian distribution, $\delta(\epsilon)$ is known as the Q-function and for any dataset s and query q ,

$$\Pr_{R \sim M(s, q)} [|R - q(s)| > \epsilon] = \operatorname{erfc} \left(\frac{\epsilon}{\sqrt{2}\eta} \right) \leq \frac{2}{\sqrt{\pi}} e^{-\frac{\epsilon^2}{2\eta^2}}.$$

Combining this with a union bound completes the proof of sample accuracy.

Invoking Lemma 3.1 with $\xi = \epsilon - \sqrt{\epsilon^2 - 2\eta^2}$ implies for all $\epsilon \geq \sqrt{2}\eta$ and any $i \in [k]$,

$$\begin{aligned}
\Pr_{S \sim D^{(n)}, V \sim M(S, A)} [\operatorname{err}_P(S, V, i) > \epsilon] &\leq \frac{1}{\xi} \int_{\epsilon - \xi}^{\infty} \operatorname{erfc} \left(\frac{t}{\sqrt{2}\eta} \right) dt \\
&\stackrel{(a)}{=} \frac{\sqrt{2}\eta}{\xi} \int_{\frac{\epsilon - \xi}{\sqrt{2}\eta}}^{\infty} \operatorname{erfc}(u) du \\
&\stackrel{(a)}{<} \frac{\sqrt{2}\eta \cdot e^{-\frac{\epsilon^2 - 2\eta^2}{2\eta^2}}}{\sqrt{\pi}\epsilon \left(1 - \sqrt{1 - \frac{2\eta^2}{\epsilon^2}} \right)} \\
&\stackrel{(b)}{\leq} \frac{\sqrt{2}\epsilon}{\sqrt{\pi}\eta} e^{-\frac{\epsilon^2}{2\eta^2}} \\
&\stackrel{(c)}{\leq} \frac{2e}{\sqrt{\pi}} e^{-\frac{\epsilon^2}{4\eta^2}} \\
&\leq 4e^{-\frac{\epsilon^2}{4\eta^2}}
\end{aligned}$$

where (a) results from the fact that $\int_a^{\infty} \operatorname{erfc}(x) dx = \frac{e^{-a^2}}{\sqrt{\pi}} - a \cdot \operatorname{erfc}(a) < \frac{e^{-a^2}}{\sqrt{\pi}}$ and the definition of ξ , (b) from the bound on ϵ and the inequality $\sqrt{1+x} \leq 1 + \frac{x}{2}$, and (c) from the inequality $x \cdot e^{-x^2} \leq e^{-\frac{x^2}{2}}$.

Combining this with a union bound completes the proof of posterior accuracy. \square

Theorem C.4 (Theorem 3.4 restated). *Given two functions $\delta_1 : \mathbb{R} \rightarrow [0, 1]$, $\delta_2 : \mathbb{R} \rightarrow [0, 1]$, and an analyst A , if a mechanism M is $(\epsilon, \delta_1(\epsilon))$ -Bayes stable and $(\epsilon, \delta_2(\epsilon))$ -sample accurate with respect to A , then M is $(\epsilon, \delta'(\epsilon))$ -distribution accurate for $\delta'(\epsilon) := \inf_{\epsilon' \in (0, \epsilon), \xi \in (0, \epsilon - \epsilon')} \left(\delta_1(\epsilon') + \frac{1}{\xi} \int_{\epsilon - \epsilon' - \xi}^{\infty} \delta_2(t) dt \right)$.*

Proof. Consider an analyst A , which once completed an interaction of length k resulting in a view v , chooses $q_{k+1} := \arg \max_{i \in [k]} |q_i(D^v) - q_i(D)|$. In this case,

$$\begin{aligned} & \Pr_{\substack{S \sim D^{(n)} \\ V \sim M(S, A)}} \left[\max_{i \in [k]} |R_i - Q_i(D)| > \epsilon \right] \\ & \leq \Pr_{\substack{S \sim D^{(n)} \\ V \sim M(S, A)}} \left[\max_{i \in [k]} |R_i - Q_i(D^V)| + \max_{i \in [k]} |Q_i(D^V) - Q_i(D)| > \epsilon \right] \\ & \leq \Pr_{\substack{S \sim D^{(n)} \\ V \sim M(S, A)}} \left[\max_{i \in [k]} |R_i - Q_i(D^V)| > \epsilon - \epsilon' \right] + \Pr_{\substack{S \sim D^{(n)} \\ V \sim M(S, A), Q \sim A(V)}} \left[|Q(D^V) - Q(D)| > \epsilon' \right] \\ & \stackrel{(a)}{\leq} \frac{1}{\xi} \int_{\epsilon - \epsilon' - \xi}^{\infty} \delta(t) dt + \delta(\epsilon'), \end{aligned}$$

where (a) results from Lemma 3.1 and the definition of Bayes stability (Definition 3.3). \square

C.2 Generalized results for arbitrary queries

While the main results in this paper were stated for linear queries, some of the intermediate results extend to arbitrary queries. Denoting by $\bar{\Delta}_q := \sup_{s, s' \in \mathcal{X}} |q(s) - q(s')|$ and $\bar{\sigma}_q^2 :=$

$\mathbb{E}_{X \sim D_{\mathcal{X}}} \left[(q(X) - q(D_{\mathcal{X}^n}))^2 \right]$ and extending the posterior distribution notations to $D_{\mathcal{X}^n}^v := D_{\mathcal{X}^n | \mathcal{Y}}^A(\cdot | v)$, Lemma 3.5 can be restated in terms of arbitrary queries.

Lemma C.5 (Covariance stability for arbitrary queries). *Given a view $v \in \mathcal{V}$ and a query q ,*

$$q(D_{\mathcal{X}^n}^v) - q(D_{\mathcal{X}^n}) = \mathop{\text{Cov}}_{S \sim D_{\mathcal{X}^n}}(q(S), K(S, v)).$$

Furthermore, given $\Delta, \sigma > 0$,

$$\sup_{q \in \mathcal{Q} \text{ s.t. } \bar{\Delta}_q \leq \Delta} |q(D_{\mathcal{X}^n}^v) - q(D_{\mathcal{X}^n})| = 2\Delta \cdot \mathbf{D}_{TV}(D_{\mathcal{X}^n}^v \| D_{\mathcal{X}^n})$$

and

$$\sup_{q \in \mathcal{Q} \text{ s.t. } \bar{\sigma}_q^2 \leq \sigma^2} |q(D_{\mathcal{X}^n}^v) - q(D_{\mathcal{X}^n})| = \sigma \sqrt{\mathbf{D}_{\chi^2}(D_{\mathcal{X}^n}^v \| D_{\mathcal{X}^n})},$$

where $K(s, v) := \frac{D(s|v)}{D(s)} = \frac{D(v|s)}{D(v)}$ is the Bayes factor of s given v (and vice-versa).

Combining this lemma with the PC definition for arbitrary sample sets, we can state a simple version of Theorem 4.6 for arbitrary queries, and achieve results similar to those guaranteed using typical stability [Bassily and Freund, 2016]. Like in the case of typical stability, this method does not achieve optimal accuracy rates for the iid case as discussed by Kontorovich et al. [2022].

D Missing parts from Section 4

D.1 PC mechanisms

Lemma D.1 (Lemma 4.3 restated). *Given $\eta > 0$ and an analyst A , if M is a Gaussian mechanism with noise parameter η , then it is φ -PC for $\varphi(s, s'; q) := \frac{(q(s) - q(s'))^2}{2\eta^2}$ and φ -PC with respect to A for $\varphi(s, s'; v) := \frac{\|\bar{q}_v(s) - \bar{q}_v(s')\|^2}{2\eta^2}$.*

Proof. From Lemma 16 in [Bun and Steinke, 2016], for any query q , we have $\varphi(s, s'; q) = \frac{(q(s) - q(s'))^2}{2\eta^2}$.⁹ Combining this result with the composition theorem (Theorem 4.4) completes the proof. \square

D.2 Composition

Theorem D.2 (Theorem 4.4 restated). *Given $k \in \mathbb{N}$, a similarity function over responses φ , and an analyst A issuing k queries, if a mechanism M is φ -PC, then it is $\tilde{\varphi}$ -PC with respect to A , where $\tilde{\varphi}(s, s'; v) := \sum_{i=1}^k \varphi(s, s'; q_i)$.*

Proof. Notice that

$$\ell(s, s'; v_k) = \ln \left(\frac{D(v_k | s)}{D(v_k | s')} \right) = \sum_{i=1}^k \ln \left(\frac{D(r_i | s, v_{i-1})}{D(r_i | s', v_{i-1})} \right) \stackrel{(a)}{=} \sum_{i=1}^k \ln \left(\frac{D^{q_i}(r_i | s)}{D^{q_i}(r_i | s')} \right) = \sum_{i=1}^k \ell(s, s'; r_i)$$

where (a) results from the fact that given s/s' the additional information of v_{i-1} does not change the distribution over r_i , since it can be encoded in the auxiliary parameter θ_i . In case of a non-deterministic analyst, the sum includes an additional term of the form $\ln \left(\frac{D(c | s)}{D(c | s')} \right)$ for the coin-tossing step, as discussed in Section A.2. But since c is independent of the dataset, the ratio is equal to 1 and this term has no effect.

Using this fact we can prove the theorem by induction over k . The base case $k = 0$ is trivial. For every $k > 0$,

$$\begin{aligned} & \mathbb{E}_{V_k \sim M(s, A)} \left[e^{(\alpha-1)(\ell(s, s'; V_k) - \alpha\varphi(s, s'; V_k))} \right] \\ &= \mathbb{E}_{V_{k-1} \sim M(s, A)} \left[\mathbb{E}_{V_k \sim M(s, A)} \left[e^{(\alpha-1)(\ell(s, s'; V_k) - \alpha\varphi(s, s'; V_k))} \mid V_{k-1} \right] \right] \\ & \stackrel{(a)}{=} \mathbb{E}_{V_{k-1} \sim M(s, A)} \left[e^{(\alpha-1)(\ell(s, s'; V_{k-1}) - \alpha\varphi(s, s'; V_{k-1}))} \cdot \overbrace{\mathbb{E}_{R_k \sim M(s, Q_k)} \left[e^{(\alpha-1)(\ell(s, s'; R_k) - \alpha\varphi(s, s'; Q_k))} \right]}^{\leq 1} \right] \\ & \stackrel{(b)}{\leq} \mathbb{E}_{V_{k-1} \sim M(s, A)} \left[e^{(\alpha-1)(\ell(s, s'; V_{k-1}) - \alpha\varphi(s, s'; V_{k-1}))} \right] \\ & \stackrel{(c)}{\leq} 1, \end{aligned}$$

where (a) results from the definition of $\varphi(\cdot, \cdot; \bar{q})$, the previous equation and the fact that both $\ell(s, s'; V_{k-1})$ and $\varphi(s, s'; V_{k-1})$ are independent of R_k , (b) follows from the PC definition, and (c) follows from the induction assumption. \square

In the next section we prove Lemma 4.5, as well as a bound on the Rényi divergence between the prior and posterior distributions over elements, which will later be used in the proof of Theorem 4.6.

⁹The lemma is stated for the range $\alpha \in [1, \infty]$, but holds for the range $[0, 1)$ as well, directly from the integral definition.

D.3 Expected stability bound

Lemma D.3 (Lemma 4.5 restated). *Given a similarity function φ and an analyst A , if a mechanism M that is φ -PC with respect to A receives an iid sample from \mathcal{X}^n , then for any two elements $x, y \in \mathcal{X}$ and $\alpha \geq 1$ we have*

$$\mathbb{E}_{V \sim D(\cdot | x)} [\exp((\alpha - 1)(\ell(x, y; V) - \alpha\varphi(x, y; V)))] \leq 1$$

and

$$\mathbb{E}_{V \sim D(\cdot | x)} [\exp((\alpha - 1)(\ell(x; V) - \alpha\varphi(x; V)))] \leq 1,$$

where $\varphi(x, y; v) := \sup_{s \in \mathcal{X}^{n-1}} (\varphi((s, x), (s, y); v))$ and $\varphi(x; v) := \ln \left(\mathbb{E}_{Y \sim D} [e^{\varphi(x, Y; v)}] \right)$.

Proof. The notation used in this proof is defined in Appendix B.

Notice that since the dataset was sampled iid, we have $D(v | x) = \mathbb{E}_{S \sim D^{(n-1)}} [D(v | (S, x))]$, so the first part is a direct result of Part 1 of Theorem B.14, where $P_\lambda := D(\cdot | x)$, $Q_\lambda := D(\cdot | y)$, and the convexity is taken over the convex combination of the sample sets of size $n - 1$.

For the second part, notice that $D(v) = \mathbb{E}_{X \sim D} [D(v | X)]$, so

$$\begin{aligned} & \mathbb{E}_{V \sim D(\cdot | x)} [\exp((\alpha - 1)(\ell(x; V) - \alpha\varphi(x; V)))] \\ & \stackrel{(a)}{\leq} \mathbb{E}_{V \sim D(\cdot | x)} \left[\exp \left((\alpha - 1) \left(\ell(x; V) - \alpha \mathbb{E}_{Y \sim D} [\varphi(x, Y; v)] \right) \right) \right] \\ & \stackrel{(b)}{\leq} 1, \end{aligned}$$

where (a) results from Jensen's inequality for the convex function e^x and (b) from combining the first part with Part 2 of Theorem B.14, where $P := D(\cdot | x)$, $Q_\lambda := D$, and the convexity is taken over the convex combination of the elements. \square

Next, we bound the Rényi divergence between the prior and posterior distributions for PC mechanisms.

Lemma D.4. *Given an element $x \in \mathcal{X}$, a similarity function over views φ and an analyst A , if a mechanism M that is φ -PC with respect to A receives an iid sample from \mathcal{X}^n , then for any $\alpha \geq 1$ we have*

$$\mathbf{D}_\alpha(D_V(\cdot | x) \| D_V) \leq \frac{1}{2(\alpha - 1)} \ln \left(\mathbb{E}_{V \sim D} [\exp(2\alpha(2\alpha - 1)\varphi(x; V))] \right),$$

where $\varphi(x; v)$ is as was defined in Lemma 4.5.

Proof.

$$\begin{aligned}
& \exp((\alpha - 1) \mathbf{D}_\alpha(D_{\mathcal{V}}(\cdot|X) \| D_{\mathcal{V}})) \\
&= \mathbb{E}_{V \sim D(\cdot|x)} \left[\left(\frac{D(V|x)}{D(V)} \right)^{\alpha-1} \right] \\
&= \mathbb{E}_{V \sim D} \left[\left(\frac{D(V|x)}{D(V)} \right)^\alpha \right] \\
&= \mathbb{E}_{V \sim D} [\exp(\alpha(\ell(x; V) - (2\alpha - 1)\varphi(x; V))) \cdot \exp(\alpha(2\alpha - 1)\varphi(x; V))] \\
&\stackrel{(a)}{\leq} \sqrt{\mathbb{E}_{V \sim D} [\exp(2\alpha(\ell(x; V) - (2\alpha - 1)\varphi(x; V)))] \cdot \mathbb{E}_{V \sim D} [\exp(2\alpha(2\alpha - 1)\varphi(x; V))]} \\
&= \sqrt{\mathbb{E}_{V \sim D} \left[\left(\frac{D(V|x)}{D(V)} \right)^{2\alpha} \exp(-2\alpha(2\alpha - 1)\varphi(x; V)) \right] \cdot \mathbb{E}_{V \sim D} [\exp(2\alpha(2\alpha - 1)\varphi(x; V))]} \\
&= \sqrt{\mathbb{E}_{V \sim D(\cdot|x)} \left[\left(\frac{D(V|x)}{D(V)} \right)^{2\alpha-1} \exp(-2\alpha(2\alpha - 1)\varphi(x; V)) \right] \cdot \mathbb{E}_{V \sim D} [\exp(2\alpha(2\alpha - 1)\varphi(x; V))]} \\
&= \sqrt{\mathbb{E}_{V \sim D(\cdot|x)} [\exp((2\alpha - 1)(\ell(x; V) - 2\alpha\varphi(x; V)))] \cdot \mathbb{E}_{V \sim D} [\exp(2\alpha(2\alpha - 1)\varphi(x; V))]} \\
&\stackrel{(b)}{\leq} \sqrt{\mathbb{E}_{V \sim D} [\exp(2\alpha(2\alpha - 1)\varphi(x; V))]}
\end{aligned}$$

where (a) results from the Cauchy-Schwarz inequality and (b) from the second part of Lemma 4.5.

Taking the log of both sides and dividing by $\alpha - 1$ completes the proof. \square

In the next section we prove a high probability bound on the stability loss in terms of the similarity function, as a second step toward proving Theorem 4.6.

D.4 High probability stability bound

Lemma D.5. *Given a similarity function over views φ and an analyst A , if a mechanism M that is φ -PC with respect to A receives an iid sample from \mathcal{X}^n , then*

$$\Pr_{\substack{X \sim D, S \sim D^{(n)} \\ V \sim M(S, A)}} [|\ell(X; V)| > \varepsilon(X; V)] \leq \frac{2^{\frac{1}{4}}(e^\phi + 1)\delta^2}{e^\phi - 1} \left(\mathbb{E}_{\substack{X \sim D \\ V \sim D}} [\exp(12\varphi(X; V))] \right)^{\frac{1}{8}}$$

for any $\delta > 0$, $\phi \geq \mathbb{E}_{\substack{X \sim D, S \sim D^{(n)} \\ V \sim M(S, A)}} [\varphi(X; V)]$, where $\varepsilon(x; v) := \varphi(x; v) + \phi + 2\sqrt{2\ln(\frac{1}{\delta})(\varphi(x; v) + \phi)}$ and $\varphi(x; v)$ is as was defined in Lemma 4.5.

Before we prove this lemma, we present several supporting claims. Throughout, we will fix a similarity function over views φ , an analyst A , and a mechanism M which is φ -PC with respect to A receiving an iid sample from \mathcal{X}^n . The notation used here is all defined in Appendix B.

Claim D.6. *Given $\phi, \delta > 0$ and two sample sets $s, s' \in \mathcal{X}^n$, we have $D(\cdot|s) \leq_{(\varepsilon_{s,s'}, \delta_{s,s'})}^+$*

$$\begin{aligned}
& D(\cdot|s'), \text{ where } \varepsilon_{s,s'}(v) := \varphi(s, s'; v) + 2\sqrt{2\ln(\frac{1}{\delta})(\varphi(s, s'; v) + \phi)} \text{ and } \delta_{s,s'} := \\
& \delta \sqrt{\mathbb{E}_{V \sim D(\cdot|s)} \left[\sqrt{\frac{\varphi(s, s'; V)}{\phi} + 1} \right]}.
\end{aligned}$$

Proof. Denote $\Omega := \mathcal{V}$, $P := D(\cdot|s)$, and $Q := D(\cdot|s')$. From the PC definition, for any $\alpha \geq 0$ the α -Rényi dissimilarity between P and Q is bounded by $\alpha \cdot \varphi(s, s'; \cdot)$.

Invoking Theorem B.12 we get $\Pr_{V \sim P} \left[\left| \ln \left(\frac{P(V)}{Q(V)} \right) \right| > \tilde{\varepsilon}_{s,s'}(V) \right] \leq \delta_{s,s'}$, where

$$\tilde{\varepsilon}_{s,s'}(v) := \varphi(s, s'; v) + 2\sqrt{2\ln\left(\frac{1}{\delta}\right) (\varphi(s, s'; v) + \phi)}.$$

Combining this with Part 1 of Theorem B.10 we get that $P \leq_{(\varepsilon_{s,s'}, \delta_{s,s'})}^+ Q$. \square

Claim D.7. Given $\phi, \delta > 0$ and two elements $x, y \in \mathcal{X}$, we have $D(\cdot | x) \leq_{(\varepsilon_{x,y}, \delta_{x,y})}^+ D(\cdot | y)$ and $D(\cdot | x) \leq_{(\varepsilon_{x,y}, \delta_{x,y})}^- D(\cdot | y)$, where $\varepsilon_{x,y}(v) := \sup_{s \in \mathcal{X}^{n-1}} (\varepsilon_{(s,x), (s,y)}(v))$, $\delta_{x,y} := \mathbb{E}_{S \sim D^{(n-1)}} [\delta_{(S,x), (S,y)}]$, and $\varepsilon_{s,s'}, \delta_{s,s'}$ are defined as in Claim D.6.

Proof. Notice that from the iid assumption, $D(v | x) = \mathbb{E}_{S \sim D^{(n-1)}} [D(v | (S, x))]$, so combining Claim D.6 and Part 1 of Theorem B.13 implies $D(\cdot | x) \leq_{(\varepsilon_{x,y}, \delta_{x,y})}^+ D(\cdot | y)$. Combining it with Part 2 of Theorem B.10 implies $D(\cdot | x) \leq_{(\varepsilon_{x,y}, \delta_{x,y})}^- D(\cdot | y)$. \square

Claim D.8. Given $\phi, \delta > 0$ and an element $x \in \mathcal{X}$, we have $D_{\mathcal{V}} \leq_{(\varepsilon_x, \delta_{x,*})}^+ D_{\mathcal{V}|\mathcal{X}}(\cdot | x)$ and $D_{\mathcal{V}|\mathcal{X}}(\cdot | x) \leq_{(\varepsilon_x, \delta_{x,*})}^- D_{\mathcal{V}}$, where $\varepsilon_x(v) := \ln \left(\mathbb{E}_{Y \sim D} [e^{\varepsilon(x, Y; v)}] \right)$, $\delta_{x,*} := \mathbb{E}_{Y \sim D} [\delta_{x, Y}]$, $\delta_{*,x} := \mathbb{E}_{Y \sim D} [\delta_{Y, x}]$, and $\varepsilon_{x,y}, \delta_{x,y}$ are defined as in Claim D.7.

Proof. Notice that $D(v) = \mathbb{E}_{Y \sim D} [D(v | Y)]$, so combining the first part of Claim D.7 and Part 2 of Theorem B.13 implies $D(\cdot) \leq_{(\varepsilon_x, \delta_{x,*})}^+ D(\cdot | x)$.

Similarly, combining the second part of Claim D.7 and Part 3 of Theorem B.13 implies $D(\cdot | x) \leq_{(\varepsilon_x, \delta_{*,x})}^- D(\cdot)$. Notice that the symmetry of φ implies ε is symmetric as well, so the same ε function bounds both terms. \square

Claim D.9. Given $\phi, \epsilon, \delta > 0$ and an element $x \in \mathcal{X}$, we have $\Pr_{V \sim D} \left[\ln \left(\frac{D(V)}{D(V|x)} \right) > \varepsilon_x(V) + \epsilon \right] < \frac{e^\epsilon \delta_{x,*}}{e^\epsilon - 1}$ and $\Pr_{V \sim D} \left[\ln \left(\frac{D(V|x)}{D(V)} \right) > \varepsilon_x(\omega) + \epsilon \right] < \frac{\delta_{*,x}}{e^\epsilon - 1}$, where $\varepsilon_x, \delta_{x,*}$, and $\delta_{*,x}$ are defined as in Claim D.8.

Proof. The first part results from combining the first part of Claim D.8 and Part 3 of Theorem B.10. The second part results from combining the second part of Claim D.8 and Part 4 of Theorem B.10. \square

Proof of Lemma D.5. Notice that for any $\phi, \tilde{\delta} > 0$ we have

$$\begin{aligned} \varepsilon_{x,y}(v) &= \sup_{s \in \mathcal{X}^{n-1}} (\varepsilon_{(s,x), (s,y)}(v)) \\ &= \sup_{s \in \mathcal{X}^{n-1}} \left(\varphi((s, x), (s, y); v) + 2\sqrt{2\ln\left(\frac{1}{\tilde{\delta}}\right) (\varphi((s, x), (s, y); v) + \phi)} \right) \\ &\leq \sup_{s \in \mathcal{X}^{n-1}} (\varphi((s, x), (s, y); v)) + 2\sqrt{2\ln\left(\frac{1}{\tilde{\delta}}\right) \left(\sup_{s \in \mathcal{X}^{n-1}} (\varphi((s, x), (s, y); v)) + \phi \right)} \\ &= \varphi(x, y; v) + 2\sqrt{2\ln\left(\frac{1}{\tilde{\delta}}\right) (\varphi(x, y; v) + \phi)} \end{aligned}$$

which implies

$$\begin{aligned}
\varepsilon_x(v) &= \mathbb{E}_{Y \sim D} [\varepsilon_{x,Y}(v)] \\
&\stackrel{(a)}{\leq} \mathbb{E}_{Y \sim D} \left[\varphi(x, Y; v) + 2\sqrt{2\ln\left(\frac{1}{\tilde{\delta}}\right)} (\varphi(x, Y; v) + \phi) \right] \\
&\stackrel{(b)}{\leq} \mathbb{E}_{Y \sim D} [\varphi(x, Y; v)] + 2\sqrt{2\ln\left(\frac{1}{\tilde{\delta}}\right)} \left(\mathbb{E}_{Y \sim D} [\varphi(x, Y; v)] + \phi \right) \\
&= \varphi(x; v) + 2\sqrt{2\ln\left(\frac{1}{\tilde{\delta}}\right)} (\varphi(x; v) + \phi),
\end{aligned}$$

where (a) results from the previous inequality and (b) from Jensen's inequality for the concave function \sqrt{x} .

Setting $\epsilon := \phi$ and $\tilde{\delta} := \delta^2$ we get

$$\begin{aligned}
&\Pr_{\substack{X \sim D, S \sim D^{(n)} \\ V \sim M(S, A)}} [|\ell(X; V)| > \varepsilon(X; V)] \\
&= \mathbb{E}_{X \sim D} \left[\Pr_{\substack{S \sim D^{(n)} \\ V \sim M(S, A)}} \left[\left| \ln \left(\frac{D(V|X)}{D(V)} \right) \right| > \varphi(X; V) + \phi + 2\sqrt{2\ln\left(\frac{1}{\tilde{\delta}}\right)} (\varphi(X; V) + \phi) \right] \right] \\
&\stackrel{(a)}{\leq} \mathbb{E}_{X \sim D} \left[\Pr_{\substack{S \sim D^{(n)} \\ V \sim M(S, A)}} \left[\left| \ln \left(\frac{D(V|X)}{D(V)} \right) \right| > \varepsilon_X(V) + \epsilon \right] \right] \\
&\stackrel{(b)}{\leq} \frac{e^\epsilon}{e^\epsilon - 1} \mathbb{E}_{X \sim D} [\delta_{*,X}] + \frac{1}{e^\epsilon - 1} \mathbb{E}_{X \sim D} [\delta_{X,*}] \\
&\stackrel{(c)}{=} \frac{e^\epsilon + 1}{e^\epsilon - 1} \mathbb{E}_{\substack{X \sim D, Y \sim D \\ S \sim D^{(n-1)}}} \left[\tilde{\delta} \sqrt{\mathbb{E}_{V \sim M((S, X), A)} \left[\sqrt{\frac{\varphi((S, X), (S, Y); V)}{\phi} + 1} \right]} \right] \\
&\stackrel{(d)}{\leq} \frac{(e^\phi + 1) \delta^2}{e^\phi - 1} \sqrt{\mathbb{E}_{\substack{X \sim D, Y \sim D, S \sim D^{(n-1)} \\ V \sim M((S, X), A)}} \left[\sqrt{\frac{\varphi((S, X), (S, Y); V)}{\phi} + 1} \right]},
\end{aligned}$$

where (a) results from the previous bound, (b) from the two parts of Claim D.9, (c) from the definitions of $\delta_{x,*}$ and $\delta_{*,x}$, and (d) from Jensen's inequality for the concave function \sqrt{x} .

Analyzing the expectation term we get

$$\begin{aligned}
& \mathbb{E}_{\substack{X \sim D, Y \sim D, S \sim D^{(n-1)} \\ V \sim M((S, X), A)}} \left[\sqrt{\frac{\varphi((S, X), (S, Y); V)}{\phi} + 1} \right] \\
& \stackrel{(a)}{\leq} \mathbb{E}_{\substack{X \sim D, Y \sim D \\ V \sim D(\cdot | X)}} \left[\sqrt{\frac{\varphi(X, Y; V)}{\phi} + 1} \right] \\
& = \mathbb{E}_{\substack{X \sim D, Y \sim D \\ V \sim D}} \left[\frac{D(V | X)}{D(V)} \sqrt{\frac{\varphi(X, Y; V)}{\phi} + 1} \right] \\
& \stackrel{(b)}{\leq} \sqrt{\mathbb{E}_{\substack{X \sim D \\ V \sim D}} \left[\left(\frac{D(V | X)}{D(V)} \right)^2 \right]} \cdot \mathbb{E}_{\substack{X \sim D, Y \sim D \\ V \sim D}} \left[\frac{\varphi(X, Y; V)}{\phi} + 1 \right] \\
& \stackrel{(c)}{\leq} \sqrt{2 \mathbb{E}_{\substack{X \sim D \\ V \sim D(\cdot | X)}} \left[\frac{D(X | V)}{D(X)} \right]} \\
& = \sqrt{2 \mathbb{E}_{X \sim D} [\exp(\mathbf{D}_2(D_V(\cdot | X) \| D_V))]} \\
& \stackrel{(d)}{\leq} \sqrt{2} \left(\mathbb{E}_{\substack{X \sim D \\ V \sim D}} [\exp(12\varphi(X; V))] \right)^{\frac{1}{4}},
\end{aligned}$$

where (a) results from the definition of $\varphi(x, y; v)$, (b) from the Cauchy-Schwarz inequality, (c) from the bound on ϕ , and (d) from Lemma D.4.

Plugging this term back in completes the proof. \square

We can now prove the main stability theorem for PC mechanisms.

D.5 Bayes stability of PC mechanisms

Theorem D.10 (Exact version of Theorem 4.6). *Given a similarity function over views φ and an analyst A , if a mechanism M that is φ -PC with respect to A receives an iid sample from \mathcal{X}^n , then*

$$\begin{aligned}
& \Pr_{\substack{S \sim D^{(n)} \\ V \sim M((S, A), Q \sim A(V))}} \left[|Q(D_{\mathcal{X}}^v) - Q(D_{\mathcal{X}})| > \sigma_Q \cdot \sqrt{\epsilon^2 + \xi} \right] \\
& \leq \Pr_{\substack{S \sim D^{(n)} \\ V \sim M((S, A), Q \sim A(V))}} \left[\mathbb{E}_{X \sim D} \left[e^{27 \ln(\frac{1}{8})(\varphi(X; v) + \phi)} \right] > 1 + \frac{\epsilon^2}{3} \right] \\
& \quad + \frac{\delta}{\xi} (4\psi(12))^{\frac{1}{16}} \sqrt{\left(\sqrt{\psi(56)} + 1 \right) \cdot \frac{e^\phi + 1}{e^\phi - 1}}
\end{aligned}$$

for any $\epsilon, \xi, \delta > 0$, $\phi \geq \mathbb{E}_{\substack{X \sim D, S \sim D^{(n)} \\ V \sim M((S, A)}} [\varphi(X; V)]$, where $\varphi(x; v)$ is as was defined in Lemma 4.5 and

$$\psi(\alpha) := \mathbb{E}_{\substack{X \sim D, S \sim D^{(n)} \\ V \sim M((S, A)}} [e^{\alpha \varphi(X; V)}].$$

We first prove several supporting claims.

Claim D.11. Given a function $f : \mathcal{X} \times \mathcal{V} \rightarrow \mathbb{R}$, an analyst A , and a mechanism M , then for any $\xi > 0$ we have

$$\begin{aligned} & \Pr_{\substack{S \sim D^{(n)} \\ V \sim M(S,A)}} \left[\mathbf{D}_{\mathcal{X}^2} (D_{\mathcal{X}}^v \| D_{\mathcal{X}}) > \mathbb{E}_{X \sim D} \left[\left(e^{f(X,V)} - 1 \right)^2 \right] + \xi \right] \\ & \leq \frac{1}{\xi} \sqrt{\left(\mathbb{E}_{\substack{X \sim D, S \sim D^{(n)} \\ V \sim M(S,A)}} \left[e^{4\ell(X;V)} \right] + 1 \right) \cdot \Pr_{\substack{X \sim D, S \sim D^{(n)} \\ V \sim M(S,A)}} [|\ell(X;V)| > f(X,V)]}. \end{aligned}$$

Proof. Denoting $B(f) := \{(x, v) \in \mathcal{X} \times \mathcal{V} \mid |\ell(x;v)| > \varepsilon(x,v)\}$ and $\bar{B}(f) := (\mathcal{X} \times \mathcal{V}) \setminus B(f)$ we get

$$\begin{aligned} & \mathbf{D}_{\mathcal{X}^2} (D_{\mathcal{X}}^v \| D_{\mathcal{X}}) - \mathbb{E}_{X \sim D} \left[\left(e^{f(X,V)} - 1 \right)^2 \right] \\ & = \mathbb{E}_{X \sim D} \left[\left(e^{\ell(X;V)} - 1 \right)^2 \right] - \mathbb{E}_{X \sim D} \left[\left(e^{f(X,V)} - 1 \right)^2 \right] \\ & = \mathbb{E}_{X \sim D} \left[\left(\left(e^{\ell(X;V)} - 1 \right)^2 - \left(e^{f(X,V)} - 1 \right)^2 \right) \cdot \mathbb{1}_{\bar{B}(f)}((X, V)) \right] \\ & \quad + \mathbb{E}_{X \sim D} \left[\left(\left(e^{\ell(X;V)} - 1 \right)^2 - \left(e^{f(X,V)} - 1 \right)^2 \right) \cdot \mathbb{1}_{B(f)}((X, V)) \right] \\ & \stackrel{(a)}{\leq} \mathbb{E}_{X \sim D} \left[\left(\left(e^{|\ell(X;V)|} - 1 \right)^2 - \left(e^{f(X,V)} - 1 \right)^2 \right) \cdot \mathbb{1}_{\bar{B}(f)}((X, V)) \right] \\ & \quad + \mathbb{E}_{X \sim D} \left[\left(\left(e^{\ell(X;V)} - 1 \right)^2 - \left(e^{f(X,V)} - 1 \right)^2 \right) \cdot \mathbb{1}_{B(f)}((X, V)) \right] \\ & \stackrel{(b)}{\leq} \mathbb{E}_{X \sim D} \left[\left(e^{\ell(X;V)} - 1 \right)^2 \cdot \mathbb{1}_{B(f)}((X, V)) \right] \end{aligned}$$

where (a) results from the fact that $(e^x - 1)^2 \leq (e^{|x|} - 1)^2$ and (b) from removing negative terms.

Using this bound we get

$$\begin{aligned} & \Pr_{\substack{S \sim D^{(n)} \\ V \sim M(S,A)}} \left[\mathbf{D}_{\mathcal{X}^2} (D_{\mathcal{X}}^v \| D_{\mathcal{X}}) > \mathbb{E}_{X \sim D} \left[\left(e^{f(X,V)} - 1 \right)^2 \right] + \xi \right] \\ & \stackrel{(a)}{\leq} \Pr_{\substack{S \sim D^{(n)} \\ V \sim M(S,A)}} \left[\mathbb{E}_{X \sim D} \left[\left(e^{\ell(X;V)} - 1 \right)^2 \cdot \mathbb{1}_{B(f)}((X, V)) \right] > \xi \right] \\ & \stackrel{(b)}{\leq} \frac{1}{\xi} \mathbb{E}_{\substack{X \sim D, S \sim D^{(n)} \\ V \sim M(S,A)}} \left[\left(e^{\ell(X;V)} - 1 \right)^2 \cdot \mathbb{1}_{B(f)}((X, V)) \right] \\ & \stackrel{(c)}{\leq} \frac{1}{\xi} \sqrt{\mathbb{E}_{\substack{X \sim D, S \sim D^{(n)} \\ V \sim M(S,A)}} \left[\left(e^{\ell(X;V)} - 1 \right)^4 \right] \cdot \mathbb{E}_{\substack{X \sim D, S \sim D^{(n)} \\ V \sim M(S,A)}} \left[\mathbb{1}_{B(f)}^2((X, V)) \right]} \\ & \stackrel{(d)}{\leq} \frac{1}{\xi} \sqrt{\left(\mathbb{E}_{\substack{X \sim D, S \sim D^{(n)} \\ V \sim M(S,A)}} \left[e^{4\ell(X;V)} \right] + 1 \right) \cdot \Pr_{\substack{X \sim D, S \sim D^{(n)} \\ V \sim M(S,A)}} [(X, V) \in B(f)]} \end{aligned}$$

where (a) results from the previous bound, (b) from Markov's inequality, (c) from the Cauchy-Schwarz inequality, and (d) from the inequality $(x - 1)^4 \leq x^4 + 1$ for any $x \geq 0$. \square

Claim D.12. Given $\phi > 0$, $\delta \in (0, \frac{1}{e})$, and a view $v \in \mathcal{V}$, we have

$$\mathbb{E}_{X \sim D} \left[\left(e^{\varepsilon(X;v)} - 1 \right)^2 \right] \leq 3 \left(\mathbb{E}_{X \sim D} \left[e^{27 \ln(\frac{1}{\delta}) (\varphi(X;v) + \phi)} \right] - 1 \right)$$

where $\varepsilon(x; v)$ was defined in Lemma D.5.

Proof. We first notice that for any $\alpha \geq 1$,

$$\mathbb{E}_{X \sim D} \left[e^{\alpha \varphi(X;v)} \right] \stackrel{(a)}{=} \mathbb{E}_{X \sim D} \left[\left(\mathbb{E}_{Y \sim D} \left[e^{\varphi(X,Y;v)} \right] \right)^\alpha \right] \stackrel{(b)}{\leq} \mathbb{E}_{X, Y \sim D} \left[e^{\alpha \varphi(X,Y;v)} \right] \quad (1)$$

where (a) follows from the definition of $\varphi(x; v)$ and (b) from Jensen's inequality for the convex function x^α .

Next we prove a useful identity. For any $a, b \geq 0$,

$$e^{a+b} - 1 = e^a (e^b - 1) + (e^a - 1) = e^a \sqrt{(e^b - 1)^2} + \sqrt{(e^a - 1)^2} \leq e^a \sqrt{(e^{\frac{3}{2}b^2} - 1)} + \sqrt{(e^{2a} - 1)}$$

where the inequality results from the inequalities $(e^x - 1)^2 \leq e^{\frac{3}{2}x^2} - 1$ for any $x \in \mathbb{R}$ and $(e^x - 1)^2 \leq e^{2x} - 1$ for any $x \geq 0$.

Using this bound we get

$$\begin{aligned} (e^{a+b} - 1)^2 &\stackrel{(a)}{\leq} \left(e^a \sqrt{(e^{\frac{3}{2}b^2} - 1)} + \sqrt{(e^{2a} - 1)} \right)^2 \\ &= e^{2a} (e^{\frac{3}{2}b^2} - 1) + 2e^a \sqrt{(e^{\frac{3}{2}b^2} - 1)} (e^{2a} - 1) + (e^{2a} - 1) \\ &= (e^{2a + \frac{3}{2}b^2} - 1) + 2e^a \sqrt{(e^{\frac{3}{2}b^2} - 1)} (e^{2a} - 1) \\ &\stackrel{(b)}{\leq} (e^{2a + \frac{3}{2}b^2} - 1) + 2(e^{3a + \frac{3}{2}b^2} - 1) \\ &\leq 3(e^{3a + \frac{3}{2}b^2} - 1) \end{aligned}$$

where (a) results from the previous bound and (b) from the inequalities $(e^x - 1)(e^y - 1) \leq (e^{x+y} - 1)^2$ and $e^x (e^y - 1) \leq (e^{x+y} - 1)$ for any $x, y \geq 0$.

Using this inequality we get

$$\begin{aligned} &\mathbb{E}_{X \sim D} \left[\left(e^{\varepsilon(X;v)} - 1 \right)^2 \right] \\ &= \mathbb{E}_{X \sim D} \left[\left(\exp \left(\varphi(X;v) + \phi + 4 \sqrt{\ln \left(\frac{1}{\delta} \right)} (\varphi(X;v) + \phi) \right) - 1 \right)^2 \right] \\ &\stackrel{(a)}{\leq} \mathbb{E}_{X \sim D} \left[3 \left(\exp \left(3 \left(1 + 8 \ln \left(\frac{1}{\delta} \right) \right) (\varphi(X;v) + \phi) \right) - 1 \right) \right] \\ &\stackrel{(b)}{\leq} 3 \left(\mathbb{E}_{X \sim D} \left[e^{27 \ln(\frac{1}{\delta}) (\varphi(X;v) + \phi)} \right] - 1 \right) \end{aligned}$$

where (a) results from the previous inequality and (b) from the bound on δ . □

Proof of Theorem D.10. First notice

$$\begin{aligned}
\mathbb{E}_{\substack{X \sim D, S \sim D^{(n)} \\ V \sim M(S, A)}} \left[e^{4\ell(X; V)} \right] &= \mathbb{E}_{X \sim D} \left[\mathbb{E}_{V \sim D} \left[\left(\frac{D(V|X)}{D(V)} \right)^4 \right] \right] \\
&= \mathbb{E}_{X \sim D} \left[\mathbb{E}_{V \sim D(\cdot|X)} \left[\left(\frac{D(V|X)}{D(V)} \right)^3 \right] \right] \\
&= \mathbb{E}_{X \sim D} \left[\exp(\mathbf{D}_4(D_V(\cdot|X) \| D_V)) \right] \\
&\stackrel{(a)}{\leq} \mathbb{E}_{X \sim D} \left[\sqrt{\mathbb{E}_{V \sim D} [e^{56\varphi(X; V)}]} \right] \\
&\stackrel{(b)}{\leq} \sqrt{\mathbb{E}_{X \sim D, V \sim D} [e^{56\varphi(X; V)}]} \\
&\stackrel{(c)}{=} \sqrt{\psi(56)}
\end{aligned}$$

where (a) results from Lemma D.4, (b) from Jensen's inequality for the concave function \sqrt{x} , and (c) from the definition of ψ .

Using this bound we get

$$\begin{aligned}
&\Pr_{\substack{S \sim D^{(n)} \\ V \sim M(S, A), Q \sim A(V)}} \left[|Q(D_{\mathcal{X}}^v) - Q(D_{\mathcal{X}})| > \sigma_Q \cdot \sqrt{\epsilon^2 + \xi} \right] \\
&\stackrel{(a)}{\leq} \Pr_{\substack{S \sim D^{(n)} \\ V \sim M(S, A), Q \sim A(V)}} \left[\mathbf{D}_{\mathcal{X}^2}(D_{\mathcal{X}}^v \| D_{\mathcal{X}}) > \epsilon^2 + \xi \right] \\
&\stackrel{(b)}{\leq} \Pr_{\substack{S \sim D^{(n)} \\ V \sim M(S, A), Q \sim A(V)}} \left[\mathbf{D}_{\mathcal{X}^2}(D_{\mathcal{X}}^v \| D_{\mathcal{X}}) > \mathbb{E}_{X \sim D} \left[\left(e^{\varepsilon(X; V)} - 1 \right)^2 \right] + \xi \right] \\
&\quad + \Pr_{\substack{S \sim D^{(n)} \\ V \sim M(S, A), Q \sim A(V)}} \left[\mathbb{E}_{X \sim D} \left[\left(e^{\varepsilon(X; V)} - 1 \right)^2 \right] > \epsilon^2 \right] \\
&\stackrel{(c)}{\leq} \frac{1}{\xi} \sqrt{\left(\mathbb{E}_{\substack{X \sim D, S \sim D^{(n)} \\ V \sim M(S, A)}} [e^{4\ell(X; V)}] + 1 \right) \cdot \Pr_{\substack{X \sim D, S \sim D^{(n)} \\ V \sim M(S, A)}} [|\ell(X; V)| > \varepsilon(X; V)]} \\
&\quad + \Pr_{\substack{S \sim D^{(n)} \\ V \sim M(S, A), Q \sim A(V)}} \left[3 \left(\mathbb{E}_{X \sim D} \left[e^{27\ln(\frac{1}{8})(\varphi(X; v) + \phi)} \right] - 1 \right) > \epsilon^2 \right] \\
&\stackrel{(d)}{\leq} \frac{1}{\xi} \sqrt{\left(\sqrt{\psi(56)} + 1 \right) \cdot \frac{2^{\frac{1}{4}} (e^\phi + 1) \delta^2}{e^\phi - 1} (\psi(12))^{\frac{1}{8}}} \\
&\quad + \Pr_{\substack{S \sim D^{(n)} \\ V \sim M(S, A), Q \sim A(V)}} \left[\mathbb{E}_{X \sim D} \left[e^{27\ln(\frac{1}{8})(\varphi(X; v) + \phi)} \right] > 1 + \frac{\epsilon^2}{3} \right]
\end{aligned}$$

where (a) results from Corollary 3.6, (b) from the inequality $\Pr[a > b] \leq \Pr[a > c \vee c > b] \leq \Pr[a > c] + \Pr[c > b]$, (c) from Claims D.11 and D.11, and (d) from the previous bound and Lemma D.5. \square

E Missing parts from Section 5

Using Theorem D.10, one only has to bound $\mathbb{E}_{X \sim D} \left[e^{27 \ln(\frac{1}{8}) (\varphi(X;v) + \phi)} \right]$ and $\psi(\alpha)$ to prove Bayes stability. To do so, we first prove a supporting claim and then provide separate bounds for the bounded case (Section E.1) and sub-Gaussian case (Section E.2).

Claim E.1. *Given $k \in \mathbb{N}$; $\eta > 0$ and a view $v \in \mathcal{V}$ consisting of responses to k linear queries, produced by a Gaussian mechanism M with noise parameter η which receives an iid sampled dataset of size n , then for any $\alpha \geq 1$ we have*

$$\mathbb{E}_{X \sim D} \left[e^{\alpha \varphi(X;v)} \right] \leq \mathbb{E}_{X, Y \sim D} \left[e^{\alpha \varphi(X, Y; v)} \right] = \mathbb{E}_{X, Y \sim D} \left[\exp \left(\frac{\alpha}{2n^2 \eta^2} \sum_{i=1}^k (q_i(X) - q_i(Y))^2 \right) \right],$$

where $\varphi(x; v)$ and ϕ are defined as in Theorem 4.6.

Proof. Notice that

$$\begin{aligned} \mathbb{E}_{X \sim D} \left[e^{\alpha \varphi(X;v)} \right] &\stackrel{(a)}{=} \mathbb{E}_{X \sim D} \left[\left(\mathbb{E}_{Y \sim D} \left[e^{\varphi(X, Y; v)} \right] \right)^\alpha \right] \\ &\stackrel{(b)}{\leq} \mathbb{E}_{X, Y \sim D} \left[e^{\alpha \varphi(X, Y; v)} \right] \\ &\stackrel{(c)}{=} \mathbb{E}_{X, Y \sim D} \left[\exp \left(\alpha \sup_{s \in \mathcal{X}^{n-1}} \left(\frac{\|\bar{q}_v((s, X)) - \bar{q}_v((s, Y))\|^2}{2\eta^2} \right) \right) \right] \\ &= \mathbb{E}_{X, Y \sim D} \left[\exp \left(\frac{\alpha}{2\eta^2} \sup_{s \in \mathcal{X}^{n-1}} \left(\sum_{i=1}^k (q_i((s, X)) - q_i((s, Y)))^2 \right) \right) \right] \\ &= \mathbb{E}_{X, Y \sim D} \left[\exp \left(\frac{\alpha}{2n^2 \eta^2} \sum_{i=1}^k (q_i(X) - q_i(Y))^2 \right) \right] \end{aligned}$$

where (a) results from the definition of $\varphi(x; v)$, (b) from Jensen's inequality for the convex function x^a for any $a \geq 1$, and (c) from Theorem 4.3. \square

E.1 Generalization guarantees, bounded case

Claim E.2. *Given $k \in \mathbb{N}$; $\eta, \Delta, \sigma > 0$; $\phi \geq 0$, and a view $v \in \mathcal{V}$ consisting of responses to k Δ -bounded linear queries with variance bounded by σ^2 , produced by a Gaussian mechanism M with noise parameter η which receives an iid sampled dataset of size n , then for any $1 \leq \alpha \leq \frac{1}{\frac{k\Delta^2}{2n^2\eta^2} + \phi}$ we have*

$$\mathbb{E}_{X \sim D} \left[e^{\alpha(\varphi(X;v) + \phi)} \right] \leq \mathbb{E}_{X, Y \sim D} \left[e^{\alpha(\varphi(X, Y; v) + \phi)} \right] \leq 1 + 2\alpha \left(\phi + \frac{k\sigma^2}{n^2\eta^2} \right),$$

where $\varphi(x; v)$ is as was defined in Theorem 4.6.

Proof. First notice that for any linear query q with variance σ^2 ,

$$\begin{aligned} \mathbb{E}_{X, Y \sim D} \left[(q(X) - q(Y))^2 \right] &= \mathbb{E}_{X, Y \sim D} \left[((q(X) - q(D)) + (q(D) - q(Y)))^2 \right] \\ &\stackrel{(a)}{=} \mathbb{E}_{X \sim D} \left[(q(X) - q(D))^2 \right] + \mathbb{E}_{Y \sim D} \left[(q(D) - q(Y))^2 \right] \\ &\quad + 2 \mathbb{E}_{X \sim D} [q(X) - q(D)] \cdot \mathbb{E}_{Y \sim D} [q(D) - q(Y)] \\ &\stackrel{(b)}{=} 2\sigma^2 \end{aligned} \tag{2}$$

where (a) results from the fact X, Y are independent and (b) from the fact that $q(X)$ is a random variable with expectation $q(D)$ and variance σ^2 .

For any two elements $x, y \in \mathcal{X}$ we have

$$\varphi(x, y; v) = \frac{1}{2n^2\eta^2} \sum_{i=1}^k (q_i(X) - q_i(Y))^2 \leq \frac{k\Delta^2}{2n^2\eta^2}. \quad (3)$$

Using this bound we get

$$\begin{aligned} \mathbb{E}_{X \sim D} \left[e^{\alpha(\varphi(X;v)+\phi)} \right] &\stackrel{(a)}{\leq} \mathbb{E}_{X, Y \sim D} \left[e^{\alpha(\varphi(X, Y;v)+\phi)} \right] \\ &= \mathbb{E}_{X, Y \sim D} \left[\exp \left(\alpha \left(\phi + \frac{1}{2n^2\eta^2} \sum_{i=1}^k (q_i(X) - q_i(Y))^2 \right) \right) \right] \\ &\stackrel{(b)}{\leq} \mathbb{E}_{X, Y \sim D} \left[1 + 2\alpha \left(\phi + \frac{1}{2n^2\eta^2} \sum_{i=1}^k (q_i(X) - q_i(Y))^2 \right) \right] \\ &= 1 + \alpha \left(2\phi + \frac{1}{n^2\eta^2} \sum_{i=1}^k \mathbb{E}_{X, Y \sim D} \left[(q_i(X) - q_i(Y))^2 \right] \right) \\ &\stackrel{(c)}{\leq} 1 + \alpha \left(2\phi + \frac{1}{n^2\eta^2} \sum_{i=1}^k 2\sigma^2 \right) \\ &= 1 + 2\alpha \left(\phi + \frac{k\sigma^2}{n^2\eta^2} \right) \end{aligned}$$

where (a) results from combining the assumption $\alpha \geq 1$ with Claim E.1, (b) from combining Equation 3, the assumption $\alpha \leq \frac{1}{\frac{k\Delta^2}{2n^2\eta^2} + \phi}$, and the inequality $e^x \leq 1 + 2x$ for any $0 \leq x \leq 1$, and (c) from Equation 2. \square

Lemma E.3. Given $k \in \mathbb{N}$; $\Delta, \sigma, \eta \geq 0$; $0 < \delta \leq \frac{1}{e}$; and an analyst A issuing k Δ -bounded linear queries with variance bounded by σ^2 , if M is a Gaussian mechanism with noise parameter η that receives an iid dataset of size $n \geq \frac{9\sqrt{k}\Delta}{\eta} \sqrt{\ln(\frac{1}{\delta})}$, then M is $\left(\frac{\sigma^2}{n\eta} \sqrt{488k \ln(\frac{1}{\delta})}, \frac{2n^3\eta^3\delta}{k^{\frac{3}{2}}\sigma^3} \right)$ -Bayes stable.

Proof. Setting $\phi := \frac{2k\sigma^2}{n^2\eta^2}$ we get

$$27 \ln \left(\frac{1}{\delta} \right) \stackrel{(a)}{\leq} \frac{2n^2\eta^2}{5k\Delta^2} = \frac{1}{\frac{5k\Delta^2}{2n^2\eta^2}} \stackrel{(b)}{\leq} \frac{1}{\frac{k\Delta^2}{2n^2\eta^2} + \phi}$$

where (a) results from the assumption $n \geq \frac{9\sqrt{k}\Delta}{\eta} \sqrt{\ln(\frac{1}{\delta})}$ and (b) from the definition of ϕ .

Using this bound we can invoke Claim E.2, and get that for any view $v \in \mathcal{V}$,

$$\mathbb{E}_{X \sim D} \left[e^{27 \ln(\frac{1}{\delta})(\varphi(X;v)+\phi)} \right] \leq 1 + 54 \left(\frac{k\sigma^2}{n^2\eta^2} + \phi \right) \ln \left(\frac{1}{\delta} \right). \quad (4)$$

Next, we notice that

$$\begin{aligned} \mathbb{E}_{X \sim D, V \sim D} [\varphi(X; V)] &= \mathbb{E}_{X \sim D, V \sim D} \left[\ln \left(\mathbb{E}_{Y \sim D} \left[e^{\varphi(X, Y; V)} \right] \right) \right] \\ &\stackrel{(a)}{\leq} \ln \left(\mathbb{E}_{X, Y \sim D, V \sim D} \left[e^{\varphi(X, Y; V)} \right] \right) \\ &\stackrel{(b)}{\leq} \ln \left(1 + \frac{2k\sigma^2}{n^2\eta^2} \right) \\ &\stackrel{(c)}{\leq} \frac{2k\sigma^2}{n^2\eta^2} \\ &\stackrel{(d)}{=} \phi \end{aligned} \quad (5)$$

where (a) results from Jensen's inequality for the concave function $\ln(x)$, (b) from Claim E.2 where the condition on α is fulfilled by the assumption that $n \geq \frac{9\sqrt{k}\Delta}{\eta} \sqrt{\ln\left(\frac{1}{\delta}\right)}$, (c) from the inequality $\ln(1+x) \leq x$ for any $x \geq 0$, and (d) from the definition of ϕ .

Similarly, for any $\alpha \leq \frac{2n^2\eta^2}{k\Delta^2}$ we have

$$\psi(\alpha) \stackrel{(a)}{\leq} 1 + \frac{2\alpha k\sigma^2}{n^2\eta^2} \stackrel{(b)}{\leq} 1 + \frac{\alpha k\Delta^2}{2n^2\eta^2} \stackrel{(c)}{\leq} 1 + \frac{\alpha k\Delta^2}{2n^2\eta^2} \stackrel{(d)}{\leq} 1 + \frac{\alpha}{162} \quad (6)$$

where (a) results from the assumption on α and Claim E.2, (b) from the fact that for any distribution, $\sigma^2 \leq \frac{\Delta^2}{4}$, (c) from the assumption that $\delta \leq \frac{1}{e}$, and (d) from the assumption that $n \geq \frac{9\sqrt{k}\Delta}{\eta} \sqrt{\ln\left(\frac{1}{\delta}\right)}$.

Setting $\xi := \frac{2k\sigma^2}{n^2\eta^2}$ we get

$$\begin{aligned} & \Pr_{\substack{S \sim D^{(n)} \\ V \sim M(S,A), Q \sim A(V)}} \left[|Q(D_{\mathcal{X}}^v) - Q(D_{\mathcal{X}})| > \frac{\sigma^2}{n\eta} \sqrt{488k \ln\left(\frac{1}{\delta}\right)} \right] \\ & \stackrel{(a)}{\leq} \Pr_{\substack{S \sim D^{(n)} \\ V \sim M(S,A), Q \sim A(V)}} \left[|Q(D_{\mathcal{X}}^v) - Q(D_{\mathcal{X}})| > \sigma_Q \sqrt{\frac{488k\sigma^2}{n^2\eta^2} \ln\left(\frac{1}{\delta}\right)} \right] \\ & \stackrel{(b)}{\leq} \Pr_{\substack{S \sim D^{(n)} \\ V \sim M(S,A), Q \sim A(V)}} \left[|Q(D_{\mathcal{X}}^v) - Q(D_{\mathcal{X}})| > \sigma_Q \sqrt{162 \left(\frac{k\sigma^2}{n^2\eta^2} + \frac{2k\sigma^2}{n^2\eta^2} \right) \ln\left(\frac{1}{\delta}\right) + \frac{2k\sigma^2}{n^2\eta^2}} \right] \\ & \stackrel{(c)}{=} \Pr_{\substack{S \sim D^{(n)} \\ V \sim M(S,A), Q \sim A(V)}} \left[|Q(D_{\mathcal{X}}^v) - Q(D_{\mathcal{X}})| > \sigma_Q \sqrt{162 \left(\frac{k\sigma^2}{n^2\eta^2} + \phi \right) \ln\left(\frac{1}{\delta}\right) + \xi} \right] \\ & \stackrel{(d)}{\leq} \Pr_{\substack{S \sim D^{(n)} \\ V \sim M(S,A), Q \sim A(V)}} \left[\mathbb{E}_{X \sim D} \left[e^{27 \ln\left(\frac{1}{\delta}\right) (\varphi(X;v) + \phi)} \right] > 1 + 54 \left(\frac{k\sigma^2}{n^2\eta^2} + \phi \right) \ln\left(\frac{1}{\delta}\right) \right] \\ & \quad + \frac{\delta}{\xi} (4\psi(12))^{1/6} \sqrt{\left(\sqrt{\psi(56)} + 1 \right) \frac{e^\phi + 1}{e^\phi - 1}} \\ & \stackrel{(e)}{\leq} \frac{\delta}{\xi} (4\psi(12))^{1/6} \sqrt{\left(\sqrt{\psi(56)} + 1 \right) \frac{e^\phi + 1}{e^\phi - 1}} \\ & \stackrel{(f)}{\leq} \frac{\delta}{\xi} \left(\frac{116}{27} \right)^{1/6} \sqrt{\left(\sqrt{\frac{109}{81}} + 1 \right) \frac{3}{\phi}} \\ & \leq \frac{3\delta}{\xi\sqrt{\phi}} \\ & \stackrel{(c)}{=} \frac{2n^3\eta^3\delta}{k^{3/2}\sigma^3} \end{aligned}$$

where (a) results from fact that the variance of all queries issued by A is bounded by σ^2 , (b) from the assumption that $\delta \leq \frac{1}{e}$, (c) from the definition of ϕ and ξ , (d) from Theorem D.10 where the condition was proven in Equation 5, (e) from Equation 4, and (f) from Equation 6, and the fact that $\frac{e^x+1}{e^x-1} \leq \frac{3}{x}$ for any $0 < x < 2$. \square

Theorem E.4 (Exact version of Theorem 5.1). *Given $k \in \mathbb{N}$; $\Delta, \sigma, \epsilon \geq 0$; $0 < \delta \leq \frac{1}{e}$; and an analyst A issuing k Δ -bounded linear queries with variance bounded by σ^2 , if M is a Gaussian mechanism with noise parameter $\eta = (122)^{\frac{1}{4}} \sigma \sqrt{\frac{\sqrt{k}}{n}}$ that receives an iid dataset of size $n \geq \max\left\{ \frac{9\Delta}{\epsilon}, \frac{89\sigma^2}{\epsilon^2} \right\} \sqrt{k} \ln\left(\frac{8k}{\delta}\right)$, and $k \geq 2n$, then M is (ϵ, δ) -distribution accurate.*

Proof. From Lemma E.3 setting $\delta' := \frac{\delta}{8k}$ (where δ' denotes the parameter δ used in the Lemma) we get that M is $\left(\frac{\sigma^2}{n\eta} \sqrt{488k \ln \left(\frac{4k}{\delta}\right)}, \frac{n^3 \eta^3 \delta}{k^{\frac{5}{2}} \sigma^3}\right)$ -Bayes stable. From Lemma 3.2, M is $\left(2\eta \sqrt{\ln \left(\frac{4k}{\delta}\right)}, \frac{\delta}{2}\right)$ -posterior accurate.

Optimizing over η we get that $\eta = (122)^{\frac{1}{4}} \sigma \sqrt{\frac{\sqrt{k}}{n}}$. Combining this with the bound on k implies $\frac{n^3 \eta^3}{k^{\frac{5}{2}} \sigma^3} = \frac{3(122)^{\frac{1}{4}} n^{\frac{3}{2}}}{\sqrt{2} k^{\frac{3}{4}}} \leq 2k^{\frac{3}{4}} \leq 2k$. Plugging the definition of η back in we get that M is $\left(\sigma \sqrt{\frac{\sqrt{488k}}{n} \ln \left(\frac{4k}{\delta}\right)}, \frac{\delta}{2}\right)$ -Bayes stable and posterior accurate.

Using the bound on n we get that M is $\left(\frac{\epsilon}{2}, \frac{\delta}{2}\right)$ -Bayes stable and posterior accurate. Combining this with Theorem 3.4 completes the proof. \square

E.2 Generalization guarantees, sub-Gaussian case

To prove the generalization guarantees in the sub-Gaussian case, we must first introduce the notion of sub-exponential random variables, and recall several folklore facts about sub-exponential random variables.

Definition E.5. Given $\sigma, \alpha \geq 0$, a random variable X is called (α, σ^2) -sub-exponential if for all $\lambda \in \left[-\frac{1}{\alpha}, \frac{1}{\alpha}\right]$ we have $\mathbb{E} [e^{\lambda X}] \leq e^{\frac{\lambda^2 \sigma^2}{2}}$.

Fact E.6. Given $\sigma \geq 0$ and a σ^2 -sub-Gaussian random variable X , the random variable $Y := X^2 - \mathbb{E} [X^2]$ is $(4\sigma^2, 64\sigma^4)$ -sub-exponential.

Fact E.7. Given $\sigma_i \geq 0$ and independent σ_i^2 -sub-Gaussian random variables X_i for $i \in \{0, 1\}$, the random variable $Y := X_0 + X_1$ is $(0, \sigma_0^2 + \sigma_1^2)$ -sub-Gaussian.

Claim E.8. Given $k \in \mathbb{N}$; $\eta, \sigma > 0$ and a view $v \in \mathcal{V}$ consisting of responses to k σ^2 -sub-Gaussian linear queries, produced by a Gaussian mechanism M with noise parameter η which receives an iid sampled dataset of size n , then for any $1 \leq \alpha \leq \frac{n^2 \eta^2}{4k\sigma^2}$ we have

$$\mathbb{E}_{X \sim D} [e^{\alpha \varphi(X; v)}] \leq \mathbb{E}_{X, Y \sim D} [e^{\alpha \varphi(X, Y; v)}] \leq \exp\left(\frac{9\alpha k \sigma^2}{n^2 \eta^2}\right),$$

where $\varphi(x; v)$ is as was defined in Theorem 4.6.

Proof. From Fact E.7 and the assumption, the random variable $q(X) - q(Y)$ is a $(0, 2\sigma^2)$ -sub-Gaussian random variable. Combining this with Fact E.6 we get that the random variable $(q(X) - q(Y))^2 - 2\sigma^2$ is a $(8\sigma^2, 256\sigma^4)$ -sub-exponential random variable, where we rely on the fact that

$$\begin{aligned} \mathbb{E}_{X, Y \sim D} [(q(X) - q(Y))^2] &= \mathbb{E}_{X, Y \sim D} [((q(X) - q(D)) + (q(D) - q(Y)))^2] \\ &\stackrel{(a)}{=} \mathbb{E}_{X \sim D} [(q(X) - q(D))^2] + \mathbb{E}_{Y \sim D} [(q(D) - q(Y))^2] \\ &\quad + 2 \mathbb{E}_{X \sim D} [q(X) - q(D)] \cdot \mathbb{E}_{Y \sim D} [q(D) - q(Y)] \\ &\stackrel{(b)}{\leq} 2\sigma^2 \end{aligned}$$

where (a) results from the fact X, Y are independent and (b) from the fact that $q(X)$ is a random variable with expectation $q(D)$ and variance proxy σ^2 which upper bounds its variance.

Using this we get

$$\begin{aligned}
\mathbb{E}_{X \sim D} \left[e^{\alpha \varphi(X;v)} \right] &\stackrel{(a)}{\leq} \mathbb{E}_{X, Y \sim D} \left[e^{\alpha \varphi(X, Y;v)} \right] \\
&= \mathbb{E}_{X, Y \sim D} \left[\exp \left(\frac{\alpha}{2n^2 \eta^2} \sum_{i=1}^k (q_i(X) - q_i(Y))^2 \right) \right] \\
&= e^{\frac{\alpha k \sigma^2}{n^2 \eta^2}} \mathbb{E}_{X, Y \sim D} \left[\exp \left(\frac{\alpha}{2n^2 \eta^2} \sum_{i=1}^k \left((q_i(X) - q_i(Y))^2 - 2\sigma^2 \right) \right) \right] \\
&= e^{\frac{\alpha k \sigma^2}{n^2 \eta^2}} \mathbb{E}_{X, Y \sim D} \left[\prod_{i=1}^k \exp \left(\frac{\alpha}{2n^2 \eta^2} \left((q_i(X) - q_i(Y))^2 - 2\sigma^2 \right) \right) \right] \\
&\stackrel{(b)}{\leq} e^{\frac{\alpha k \sigma^2}{n^2 \eta^2}} \prod_{i=1}^k \left(\mathbb{E}_{X, Y \sim D} \left[\exp \left(\frac{\alpha k}{2n^2 \eta^2} \left((q_i(X) - q_i(Y))^2 - 2\sigma^2 \right) \right) \right] \right)^{\frac{1}{k}} \\
&\stackrel{(c)}{\leq} e^{\frac{\alpha k \sigma^2}{n^2 \eta^2}} \prod_{i=1}^k \left(\exp \left(\frac{32\alpha^2 k^2 \sigma^4}{n^4 \eta^4} \right) \right)^{\frac{1}{k}} \\
&\stackrel{(d)}{\leq} e^{\frac{\alpha k \sigma^2}{n^2 \eta^2}} \prod_{i=1}^k \left(\exp \left(\frac{8\alpha k \sigma^2}{n^2 \eta^2} \right) \right)^{\frac{1}{k}} \\
&= \exp \left(\frac{9\alpha k \sigma^2}{n^2 \eta^2} \right)
\end{aligned}$$

where (a) results from combining the assumption $\alpha \geq 1$ with the Claim E.1, (b) from the generalized version of Hölder's inequality, (c) from the fact $(q(X) - q(Y))^2 - 2\sigma^2$ is a $(8\sigma^2, 256\sigma^4)$ -sub-exponential random variable and the assumption $\alpha \leq \frac{n^2 \eta^2}{4k\sigma^2}$, and (d) from the the assumption $\alpha \leq \frac{n^2 \eta^2}{4k\sigma^2}$. \square

Lemma E.9. *Given $k \in \mathbb{N}$; $\sigma, \eta > 0$; $0 < \delta \leq \frac{1}{e}$; and an analyst A issuing k σ^2 -sub-Gaussian linear queries, if M is a Gaussian mechanism with noise parameter η that receives an iid dataset of size $n \geq \frac{23\sqrt{k}\sigma}{\eta} \sqrt{\ln\left(\frac{1}{\delta}\right)}$, then M is $\left(\frac{\sigma^2}{n\eta} \sqrt{976k \ln\left(\frac{1}{\delta}\right)}, \frac{n^3 \eta^3 \delta}{k^{\frac{3}{2}} \sigma^3}\right)$ -Bayes stable.*

Proof. Setting $\phi := \frac{9k\sigma^2}{n^2 \eta^2}$ we get

$$\begin{aligned}
\mathbb{E}_{X \sim D} \left[e^{27 \ln\left(\frac{1}{\delta}\right) (\varphi(X;v) + \phi)} \right] &\stackrel{(a)}{=} \exp \left(\frac{243k\sigma^2}{n^2 \eta^2} \ln\left(\frac{1}{\delta}\right) \right) \mathbb{E}_{X \sim D} \left[e^{27 \ln\left(\frac{1}{\delta}\right) \varphi(X;v)} \right] \\
&\stackrel{(b)}{\leq} \exp \left(\frac{243k\sigma^2}{n^2 \eta^2} \ln\left(\frac{1}{\delta}\right) + \frac{243k\sigma^2}{n^2 \eta^2} \ln\left(\frac{1}{\delta}\right) \right) \\
&= \exp \left(\frac{486k\sigma^2}{n^2 \eta^2} \ln\left(\frac{1}{\delta}\right) \right) \\
&\stackrel{(c)}{\leq} 1 + \frac{972k\sigma^2}{n^2 \eta^2} \ln\left(\frac{1}{\delta}\right)
\end{aligned} \tag{7}$$

where (a) results from the definition of ϕ , (b) from Claim E.8 where the condition on α is fulfilled by the assumption that $n \geq \frac{23\sqrt{k}\sigma}{\eta} \sqrt{\ln\left(\frac{1}{\delta}\right)}$, and (c) from the inequality $e^x \leq 1 + 2x$ for any $x \leq 1$ where the condition on x is fulfilled by the assumption that $n \geq \frac{23\sqrt{k}\sigma}{\eta} \sqrt{\ln\left(\frac{1}{\delta}\right)}$.

Next, we notice that

$$\begin{aligned}
\mathbb{E}_{X \sim D, V \sim D} [\varphi(X; V)] &= \mathbb{E}_{X \sim D, V \sim D} \left[\ln \left(\mathbb{E}_{Y \sim D} \left[e^{\varphi(X, Y; V)} \right] \right) \right] \\
&\stackrel{(a)}{\leq} \ln \left(\mathbb{E}_{X, Y \sim D, V \sim D} \left[e^{\varphi(X, Y; V)} \right] \right) \\
&\stackrel{(b)}{\leq} \ln \left(\exp \left(\frac{9k\sigma^2}{n^2\eta^2} \right) \right) \\
&= \frac{9k\sigma^2}{n^2\eta^2} \\
&\stackrel{(d)}{=} \phi,
\end{aligned} \tag{8}$$

where (a) results from Jensen's inequality for the concave function $\ln(x)$, (b) from the assumption $n \geq \frac{23\sqrt{k}\sigma}{\eta} \sqrt{\ln\left(\frac{1}{\delta}\right)}$ and Claim E.8, and (c) from the definition of ϕ .

Similarly, for any $\alpha \leq \frac{n^2\eta^2}{4k\sigma^2}$ we have

$$\psi(\alpha) \stackrel{(a)}{\leq} \exp\left(\frac{9\alpha k\sigma^2}{n^2\eta^2}\right) \stackrel{(b)}{\leq} \exp\left(\frac{9\alpha}{529}\right), \tag{9}$$

where (a) results from the assumption on α and Claim E.8 and (b) from the fact that $n \geq \frac{23\sqrt{k}\sigma}{\eta} \sqrt{\ln\left(\frac{1}{\delta}\right)}$.

Setting $\xi := \frac{6k\sigma^2}{n^2\eta^2}$ we get

$$\begin{aligned}
&\Pr_{\substack{S \sim D^{(n)} \\ V \sim M(S, A), Q \sim A(V)}} \left[|Q(D_{\mathcal{X}}^v) - Q(D_{\mathcal{X}})| > \frac{\sigma^2}{n\eta} \sqrt{976k \ln\left(\frac{1}{\delta}\right)} \right] \\
&\stackrel{(a)}{\leq} \Pr_{\substack{S \sim D^{(n)} \\ V \sim M(S, A), Q \sim A(V)}} \left[|Q(D_{\mathcal{X}}^v) - Q(D_{\mathcal{X}})| > \sigma_Q \sqrt{\frac{976k\sigma^2}{n^2\eta^2} \ln\left(\frac{1}{\delta}\right)} \right] \\
&\stackrel{(b)}{\leq} \Pr_{\substack{S \sim D^{(n)} \\ V \sim M(S, A), Q \sim A(V)}} \left[|Q(D_{\mathcal{X}}^v) - Q(D_{\mathcal{X}})| > \sigma_Q \sqrt{162 \left(\frac{9k\sigma^2}{n^2\eta^2} + \frac{9k\sigma^2}{n^2\eta^2} \right) \ln\left(\frac{1}{\delta}\right) + \frac{6k\sigma^2}{n^2\eta^2}} \right] \\
&\stackrel{(c)}{=} \Pr_{\substack{S \sim D^{(n)} \\ V \sim M(S, A), Q \sim A(V)}} \left[|Q(D_{\mathcal{X}}^v) - Q(D_{\mathcal{X}})| > \sigma_Q \sqrt{162 \left(\frac{9k\sigma^2}{n^2\eta^2} + \phi \right) \ln\left(\frac{1}{\delta}\right) + \xi} \right] \\
&\stackrel{(d)}{\leq} \Pr_{\substack{S \sim D^{(n)} \\ V \sim M(S, A), Q \sim A(V)}} \left[\mathbb{E}_{X \sim D} \left[e^{27 \ln\left(\frac{1}{\delta}\right) (\varphi(X; v) + \phi)} \right] > 1 + 54 \left(\frac{9k\sigma^2}{n^2\eta^2} + \phi \right) \ln\left(\frac{1}{\delta}\right) \right] \\
&\quad + \frac{\delta}{\xi} (4\psi(12))^{\frac{1}{16}} \sqrt{\left(\sqrt{\psi(56)} + 1 \right) \frac{e^\phi + 1}{e^\phi - 1}} \\
&\stackrel{(e)}{\leq} \frac{\delta}{\xi} (4\psi(12))^{\frac{1}{16}} \sqrt{\left(\sqrt{\psi(56)} + 1 \right) \frac{e^\phi + 1}{e^\phi - 1}} \\
&\stackrel{(f)}{\leq} \frac{\delta}{\xi} e^{\frac{27}{2116}} \sqrt{\left(e^{\frac{252}{529}} + 1 \right) \frac{3}{\phi}} \\
&\leq \frac{\delta}{\xi \sqrt{\phi}} \\
&\stackrel{(c)}{=} \frac{n^3 \eta^3 \delta}{k^{\frac{3}{2}} \sigma^3}
\end{aligned}$$

where (a) results from fact that the variance of all queries issued by A is bounded by σ^2 , (b) from the assumption that $\delta \leq \frac{1}{e}$, (c) from the definition of ϕ and ξ , (d) from Theorem D.10 where the

condition was proven in Equation 8, (e) from Equation 7, and (f) from Equation 9, and the fact that $\frac{e^x+1}{e^x-1} \leq \frac{3}{x}$ for any $0 < x < 2$. \square

Theorem E.10 (Exact version of Theorem 5.3). *Given $k \in \mathbb{N}$; $\sigma, \epsilon \geq 0$; $0 < \delta \leq \frac{1}{e}$; and an analyst A issuing k σ^2 -sub-Gaussian linear queries, if M is a Gaussian mechanism with noise parameter $\eta = (244)^{\frac{1}{4}} \sigma \sqrt{\frac{\sqrt{k}}{n}}$ that receives an iid dataset of size $n \geq \frac{32\sqrt{k}\sigma^2}{\epsilon^2} \ln\left(\frac{2k}{\delta}\right)$, and $k \geq 2n$, then M is (ϵ, δ) -distribution accurate.*

Proof. From Lemma E.9 setting $\delta' := \frac{\delta}{4k}$ (where δ' denotes the parameter δ used in the Lemma) we get that M is $\left(\frac{\sigma^2}{n\eta} \sqrt{976k \ln\left(\frac{2k}{\delta}\right)}, \frac{n^3 \eta^3 \delta}{2k^{\frac{3}{2}} \sigma^3}\right)$ -Bayes stable. From Lemma 3.2, M is $\left(2\eta \sqrt{\ln\left(\frac{4k}{\delta}\right)}, \frac{\delta}{2}\right)$ Posterior accurate.

Optimizing over η we get that $\eta = (244)^{\frac{1}{4}} \sigma \sqrt{\frac{\sqrt{k}}{n}}$. Combining this with the bound on k implies $\frac{n^3 \eta^3}{k^{\frac{3}{2}} \sigma^3} = \frac{(244)^{\frac{3}{4}} n^{\frac{3}{2}}}{k^{\frac{3}{4}}} \leq 2k^{\frac{3}{4}} \leq 2k$. Plugging the definition of η back in we get that M is $\left(\sigma \sqrt{\frac{\sqrt{976k}}{n} \ln\left(\frac{4k}{\delta}\right)}, \frac{\delta}{2}\right)$ -Bayes stable and Posterior accurate.

Using the bound on n we get that M is $\left(\frac{\epsilon}{2}, \frac{\delta}{2}\right)$ -Bayes stable and Posterior accurate. Combining this with Theorem 3.4 completes the proof. \square

E.3 Different or incorrect variances

In this section we discuss in detail the case of a mechanism receiving queries with different variances, and then extend the discussion to the effects of analyst relying upon an incorrect bound on the variance of one of the queries.

In the case of a Gaussian mechanism with queries q_i , each with variance bounded by σ_i^2 , the results we achieved translate to a guarantee that with probability $\geq 1 - \delta$, for any $i \in k$, we have

$$|q_i(D) - r_i| \leq O\left(\eta_i \sqrt{\ln\left(\frac{1}{\delta}\right)} + \sigma_i \sqrt{\ln\left(\frac{1}{\delta}\right) \sum_{j=1}^k \frac{\sigma_j^2}{n^2 \eta_j^2}}\right),$$

where η_i is the noise parameter chosen at each iteration.

The $\eta_i \sqrt{\ln\left(\frac{1}{\delta}\right)}$ term represents the sample / posterior accuracy bound, and the $\sigma_i \sqrt{\ln\left(\frac{1}{\delta}\right) \sum_{j=1}^k \frac{\sigma_j^2}{n^2 \eta_j^2}}$ represents the Bayes stability. Optimizing over η_i , we get that this error is minimized by setting $\eta_i = O\left(\sigma_i \sqrt{\frac{\sqrt{k}}{n}}\right)$, which implies

$$|q_i(D) - r_i| \leq O\left(\sigma_i \sqrt{\frac{\sqrt{k}}{n} \ln\left(\frac{1}{\delta}\right)}\right).$$

Now consider a situation where the correct variances were each bounded by σ_i^2 , but the analyst mistakenly assumed a different bound τ_i^2 which might be higher or lower than σ_i^2 . In this case η_i would have been set to $O\left(\tau_i \sqrt{\frac{\sqrt{k}}{n}}\right)$, so

$$\begin{aligned} |q_i(D) - r_i| &\leq O\left(\eta_i \sqrt{\ln\left(\frac{1}{\delta}\right)} + \sigma_i \sqrt{\ln\left(\frac{1}{\delta}\right) \sum_{j=1}^k \frac{\sigma_j^2}{n^2 \eta_j^2}}\right) \\ &= O\left(\left(\frac{\tau_i}{\sigma_i} + \sqrt{\frac{1}{k} \sum_{j=1}^k \frac{\sigma_j^2}{\tau_j^2}}\right) \sigma_i \sqrt{\frac{\sqrt{k}}{n} \ln\left(\frac{1}{\delta}\right)}\right), \end{aligned}$$

which is the same term as in the case of correctly estimated variances, multiplied by the term

$$\frac{\tau_i}{\sigma_i} + \sqrt{\frac{1}{k} \sum_{j=1}^k \frac{\sigma_j^2}{\tau_j^2}}.$$

Analyzing this term allows us to precisely understand the implications of an analyst making a mistake in the assumed variance bound.

If $\tau_i \geq \sigma_i$ for a single query q_i (which means we added too much noise to the response to that query), the first term increased only for that query, and the second term decreased for all queries. On the other hand, if $\tau_i < \sigma_i$ (which means we did not add enough noise to the response to that query), the first term decreased for that query, but the second term increased for all queries.

The effect on the first term is proportional to the square root of the ratio of the wrong variance to the correct variance of that query, while the effect on the second term is proportional to the square root of average over all k queries of the ratio of the correct variance to the wrong variances (notice the switch between numerator and denominator).

F Relationships to other stability notions

Intuitively speaking, ensuring (ϵ, δ) -differential privacy means that for any query q and any two neighboring datasets s, s' , with probability $> 1 - \delta$ over the choice of the response r , the log conditional probability ratio $\ln \left(\frac{D(r|s)}{D(r|s')} \right)$ is bounded by ϵ (this quantity is sometimes referred to as the *privacy loss*).¹⁰ The most basic method for ensuring (ϵ, δ) -DP in the context of numerical queries is by randomizing the response through noise addition, calibrating the added noise to the *global sensitivity*, that is, the worst-case privacy loss over all queries and all pairs of neighboring datasets. A series of works initiated by Dwork et al. [2015] leverages bounds on the privacy loss to provide generalization guarantees for mean estimation under adaptive data analysis. Bassily et al. [2021] asymptotically improved this guarantee and extended it to low-sensitivity queries and optimization queries.

Differential privacy’s dependence on the global sensitivity naturally yields generalization guarantees that scale with the range of the query’s values. However, a dependence on the range of the queries is not required to protect against overfitting in the non-adaptive setting, where the generalization guarantees achievable via Bernstein’s inequality scale with the queries’ standard deviation. The worst-case nature of DP is not a bug but a feature, since privacy must be ensured for *all* participants under *any* circumstances against *every* potential attack. On the contrary, the task of statistical estimation focuses only on the *typical* case, even in the non-adaptive setting. To overcome the dependence on the queries’ range, we seek to decouple DP’s stability-preserving powers from its dependence on global sensitivity. This can be done by, in a sense, considering the distribution over the *local sensitivity* induced by the *underlying distribution* over the data, as we discuss in the next sections.

F.1 Scaling back the noise

One approach one might consider to attempt to circumvent DP’s worst-case nature in the context of adaptive data analysis is by adding noise that scales like the *local sensitivity* of each query, that is, the privacy loss given the specific query, the specific dataset under study, and the worst-case neighbor of that specific dataset. A variant of this notion was first introduced by Nissim et al. [2007], who illustrated that if the magnitude of noise addition depends naively on the local sensitivity, then the chosen magnitude “leaks information” about the dataset, resulting in a mechanism that is not differentially private. To address this, they introduce the *smoothed sensitivity* notion, which smooths out the addition of locally-tailored noise, achieving improved DP guarantees for the task of median estimation (or any other quantile estimation), but not for mean estimation.

Several works have leveraged this result to improve private mean estimation as well, using more involved mechanisms. For example, Feldman and Steinke [2017] propose an algorithm that splits a dataset into m subsets, estimates the empirical mean over each one in a non-private manner, and then privately reports the median of the means. They prove that this mechanism is still DP, and provide accuracy guarantees that scale with the queries’ standard deviation. Another line of work uses truncation-based specialized mechanisms to provide differential privacy guarantees for Gaussian and sub-Gaussian queries, even in the case of multivariate distribution unknown covariance [Karwa and Vadhan, 2018, Ashtiani and Liaw, 2022, Duchi et al., 2023]. Feldman and Steinke [2018] use a more natural mechanism for the purpose of generalization under adaptivity, which adds noise that scales like the empirical standard deviation of the query. The generalization guarantees they prove scale like the true standard deviation, but apply only to the expected error.

An alternative method for avoiding the privacy loss that results from scaling the added noise addition to parameters of the actual dataset is keeping the noise level constant, and instead using local sensitivity only when analyzing the mechanism’s accuracy and privacy/stability guarantees. This line of work was initiated by Ghosh and Roth [2011] and later extended by Ebadi et al. [2015], who use similar notions to provide privacy accounting for each individual in the dataset, but do not improve generalization (or privacy) guarantees in the general case. For a survey of this type of analysis, see Wang [2019]. One notable exception is a work by Feldman and Zrnic [2021], who provide improved guarantees under sparsity assumptions on the data—that is, assuming every sample element is relevant only to a small portion of the issued queries.

¹⁰This intuition is formalized with small degradation in the ϵ and δ terms in Lemma 3.3.2 of Kasiviswanathan and Smith [2014].

F.2 Accounting for the underlying distribution

The degradation in the guaranteed generalization capabilities caused by scaling the added noise to the global sensitivity of a query might be also mitigated by considering the “natural noise” induced by sampling the dataset from some underlying distribution. A variety of relaxations of the DP definition in this vein were introduced by Gehrke et al. [2012] and Bassily et al. [2013]. Bhaskar et al. [2011] even consider an alternative privacy notion that relies solely on noise from sampling without any noise added by the mechanism, but all of these approaches suffer from a similar problem: while the privacy guarantee of a single query might improve due to noise introduced by sampling, these definitions have little to no adaptive composition guarantees. A key step in DP’s composition proof is the identity $\ln \left(\frac{D(r_k | s, v_{k-1})}{D(r_k | s', v_{k-1})} \right) = \ln \left(\frac{D(r_k | s)}{D(r_k | s')} \right)$, which relies on the guarantee that the mechanism’s response is stable, given any auxiliary information (as discussed in the proof of Theorem 4.4). This guarantee can hold only if we condition on the full dataset, which one cannot do when using a definition that leverages the sampling of the data.

Failure to compose—or to compose well—is a common issue with alternative notions of stability that attempt to weaken DP. Shenefield and Ligett [2019] propose a stability notion that is both necessary and sufficient for adaptive generalization (under certain assumptions), but it suffers from limited composition guarantees. Triastcyn and Faltings [2020] propose the notion of *Bayesian differential privacy*, which does not enjoy advanced composition guarantees. Instead they prove that the generalization guarantees implied by this notion do compose, but their generalization guarantees still scale with the queries’ range in the general case. A notable exception to these composition limitations is the notion of *Typical stability* introduced by Bassily and Freund [2016]; their approach only requires that the stability notion hold with respect to a subset the of possible datasets, so long as the probability of drawing a dataset outside the set is negligible. Using this notion, they provide generalization guarantees that extend beyond datasets where each element is drawn independently and bounded linear queries, and they obtain guarantees that scale with the variance proxy in case of sub-Gaussian distributions. Unfortunately, this notion only provides accuracy guarantees for a number of queries that scales linearly with the sample size, rather than the quadratic guarantee achieved using DP in the iid case.

Considering the underlying data distribution also paves the way for a Bayesian interpretation of stability. This perspective actually dates back to the work that introduced differential privacy. Dwork et al. [2006] interpret the bound on the privacy loss as a bound on the change in the a posteriori probability belief of an analyst, over the sample set held by the mechanism. This interpretation was later formalized by Kasiviswanathan and Smith [2014], who prove that DP is essentially equivalent to a bound on the Total Variation distance (Definition B.1) between the two possible posteriors. Later, Jung et al. [2020] use the posterior distribution induced by the transcript of the interaction between the mechanism and the analyst as a key component in their analysis of DP’s generalization guarantees under adaptivity, which forms the basis for the ideas presented in this paper.

Some works consider a true Bayesian setting, where the analyst and/or the mechanism hold a prior over possible data distributions. Machanavajjhala et al. [2009] and Li et al. [2013] use such an assumption to provide stronger privacy guarantees under the assumption that the analyst’s knowledge about the distribution is limited. Yang et al. [2015] use this setting to account for non iid data distributions, where knowing of one element in the dataset changes the probability that another element is included in it, thus degrading the privacy of other elements. In the context of adaptive data analysis, Elder [2016] considered a more realistic setting, which removes the gap between the knowledge of the analyst and the mechanism regarding the underlying data distribution. Unfortunately, any assumption about the analyst’s knowledge is hard to justify in many realistic scenarios, where implicit priors might be embedded in her algorithm. The guarantees we provide hold even in the case of a fixed distribution which is known to the analyst.

F.3 Our approach

In this work we combine these two approaches, leveraging the potential improved stability resulting from low local sensitivity and the properties of the underlying distribution, to avoid the aforementioned potential pitfalls. We do so by introducing a localized stability notion that depends on the specific dataset(s) and on the queries issued during the interaction. Using this notion we can leverage the distribution over the local sensitivity induced by the underlying distribution over the domain. By

considering the local sensitivity only in the analysis stage, we avoid the induced stability loss that might be caused, had we scaled the noise with the local sensitivity. Similarly, by choosing a stability notion that depends on the full dataset(s), saving the underlying distribution for the analysis stage, we guarantee that this stability notion composes well.

By transitioning from a global stability notion to a local one, we provide high-probability generalization guarantees that scale with the queries’ standard deviation. Furthermore, by removing the dependency on the range, our guarantees extend to unbounded queries as well, so long as they are well-concentrated. Unlike previous works, both of these advantages are achieved even for simple noise-addition mechanisms.

Moving our focus from an a priori stability guarantee to an a posteriori guarantee provided after the fact is similar to the notion of *ex-post privacy* presented by Ligett et al. [2017]. This approach can also be viewed as a variant of the *metric differential privacy* (mDP) notion [Chatzikokolakis et al., 2013, Imola et al., 2022], for the purpose of adaptive data analysis. Unlike mDP, our notion does not rely on some external fixed metric over elements domain, but rather uses a separate metric at each iteration which is induced by the corresponding query.

The dependence of the stability function on the issued queries also introduces a new challenge. Since these queries are adaptively chosen, the sensitivity is not predefined even by fixing the two datasets, and, in fact, is a random variable. This challenge has been considered with different motivation under the name *full adaptivity*, first by Rogers et al. [2016] and later by Whitehouse et al. [2023]. In their framework, the sensitivity is fixed, but the privacy parameters are adaptively chosen. The dependence of our PC definition on the actual issued queries during the interaction is analogous to their *privacy odometer*.

Although differential privacy is defined with respect to given ϵ and δ parameters, most mechanisms in fact establish a relationship between the two, where each ϵ has a corresponding δ , and vice-versa (see Zhu et al. [2022] for a survey of possible representations of this relationship). [M: Be aware of the order things are presented] Mironov [2017] provided a simple representation of this relation, by introducing Rényi DP (Definition B.16), which replaces max divergence (Definition B.4) by Rényi divergence (Definition B.5) as the stability measure. Building on this notion, Bun and Steinke [2016] introduced a new privacy definition which they call *zero-concentration differential privacy* (zCDP) (Definition B.18), and prove that it implies a relationship between the ϵ and δ parameters for DP. This definition can also be viewed as a sub-Gaussian bound on the distribution of the privacy loss (Definition 4.1), in the style of *Concentrated differential privacy* introduced by Dwork and Rothblum [2016] (Definition B.17). Our PC stability notion (Definition 4.2) can be viewed as a “localized” version of zCDP, replacing the divergence by a dissimilarity measure as discussed in Appendix B, but the same technique could be used to creating a “localized” version of DP as well (Definition B.8). Triastcyn and Faltings [2020] also propose a notion similar to PC, but they do not leverage it to achieve variance-dependent accuracy guarantees.

Although we have managed to transition from worst-case to typical-case guarantees by considering the issued queries and chosen sample set, the guarantees we provide still protect against a worst-case analyst. Using the Cauchy-Schwarz inequality in the proof of the Covariance Lemma (Lemma 3.5), we bound the stability loss even for queries that are adversarially correlated with the previous view to increase the chance of overfitting. This might be viewed as overkill for a real-world setting that is typically non-adversarial. One possible extension for future work could be to consider a version of “natural analysts”, a notion introduced by Feldman and Zrnic [2021], who improve generalization guarantees for such analysts. Alternatively, Du and Dwork [2022] consider a setting where, in each iteration, the issued query is related to a different covariant of the data elements, which are weakly correlated. Although we do not consider this type of relaxation in the present work, the tools we develop here might be used to give improved guarantees under similar assumptions.