

Max-Information, Differential Privacy, and Post-Selection Hypothesis Testing*

Ryan Rogers[†] Aaron Roth[‡] Adam Smith[§] Om Thakkar[§]

September 9, 2016

Abstract

In this paper, we initiate a principled study of how the generalization properties of approximate differential privacy can be used to perform adaptive hypothesis testing, while giving statistically valid p -value corrections. We do this by observing that the guarantees of algorithms with bounded approximate *max-information* are sufficient to correct the p -values of adaptively chosen hypotheses, and then by proving that algorithms that satisfy (ϵ, δ) -differential privacy have bounded approximate max-information when their inputs are drawn from a product distribution.

This substantially extends the existing connection between differential privacy and max-information, which previously was only known to hold for (pure) $(\epsilon, 0)$ -differential privacy. It also extends our understanding of max-information as a partially unifying measure controlling the generalization properties of adaptive data analyses. We also show a lower bound, proving that (despite the strong composition properties of max-information), when data is drawn from a product distribution, (ϵ, δ) -differentially private algorithms can come *first* in a composition with other algorithms satisfying max-information bounds, but not necessarily second if the composition is required to itself satisfy a nontrivial max-information bound. This, in particular, implies that the connection between (ϵ, δ) -differential privacy and max-information holds only for inputs drawn from particular distributions, unlike the connection between $(\epsilon, 0)$ -differential privacy and max-information.

*A preliminary version of this paper appeared in the proceedings of *FOCS 2016* [RRST16].

[†]Department of Applied Mathematics and Computational Science, University of Pennsylvania. Email: ryrogers@sas.upenn.edu. Supported in part by a grant from the Sloan foundation and NSF grant CNS-1253345

[‡]Department of Computer and Information Sciences, University of Pennsylvania. Email: aaroth@cis.upenn.edu. Supported in part by a grant from the Sloan foundation, a Google Faculty Research Award, and NSF grants CNS-1513694 and CNS-1253345.

[§]Computer Science and Engineering Department, The Pennsylvania State University. Email: asmith,omthkkr@cse.psu.edu. Supported in part by a grant from the Sloan foundation, a Google Faculty Research Award, and NSF grant IIS-1447700.

Contents

1	Introduction	1
1.1	Post-Selection Hypothesis Testing	1
1.2	Max-Information (and p -values)	3
1.3	Our Results	4
1.3.1	Further Interpretation	5
1.4	Other Related Work	6
2	Preliminaries	7
3	Max-Information for (ϵ, δ)-Differentially Private Algorithms	7
3.1	Preliminaries for Linear Codes	14
3.2	Proof of Theorem 3.9	16
A	Other Useful Probabilistic Tools	23
B	Sensitivity of p-values	25
C	General Conversion between Max-Information and Mutual Information, and its Consequences	25
D	Rederiving Generalization for Low Sensitivity Queries, and a Comparison with the Bounds of [BNS⁺16]	29
E	Omitted Proofs	30

1 Introduction

Adaptive Data Analysis refers to the reuse of data to perform analyses suggested by the outcomes of previously computed statistics on the same data. It is the common case when *exploratory data analysis* and *confirmatory data analysis* are mixed together, and both conducted on the same dataset. It models both well-defined, self-contained tasks, like selecting a subset of variables using the LASSO and then fitting a model to the selected variables, and also much harder-to-specify sequences of analyses, such as those that occur when the same dataset is shared and reused by multiple researchers.

Recently two lines of work have arisen, in statistics and computer science respectively, aimed at rigorous statistical understanding of adaptive data analysis. By and large, the goal in the statistical literature (often called “selective” or “post-selection” inference [BBB⁺13]) is to derive valid hypothesis tests and tight confidence intervals around parameter values that arise from very specific analyses, such as LASSO model selection followed by least squares regression (see e.g. [FST14, LSST13]). In contrast, the second line of work has aimed for generality (at the possible expense of giving tight application-specific bounds). This second literature imposes conditions on the algorithms performing each stage of the analysis, and makes no other assumptions on how, or in what sequence, the results are used by the data analyst. Two algorithmic constraints that have recently been shown to guarantee that future analyses will be statistically valid are differential privacy [DFH⁺15b, BNS⁺16] and bounded output description length, which are partially unified by a measure of information called *max-information* [DFH⁺15a]. This paper falls into the second line of research—specifically, we extend the connection made in [DFH⁺15b, BNS⁺16] between *differential privacy* and the adaptive estimation of *low-sensitivity queries* to a more general setting that includes adaptive hypothesis testing with statistically valid p -values.

Our main technical contribution is a quantitatively tight connection between differential privacy and a *max-information*. Max-information is a measure of correlation, similar to Shannon’s mutual information, which allows bounding the change in the conditional probability of events relative to their a priori probability. Specifically, we extend a bound on the max-information of $(\epsilon, 0)$ -differentially private algorithms, due to [DFH⁺15a], to the much larger class of (ϵ, δ) -differentially private algorithms.

1.1 Post-Selection Hypothesis Testing

To illustrate an application of our results, we consider a simple model of one-sided hypothesis tests on real valued test statistics. Let \mathcal{X} denote a data domain. A *dataset* \mathbf{x} consists of n elements in \mathcal{X} : $\mathbf{x} \in \mathcal{X}^n$. A hypothesis test is defined by a *test statistic* $\phi_i : \mathcal{X}^n \rightarrow \mathbb{R}$, where we use i to index different test statistics. Given an output $a = \phi_i(\mathbf{x})$, together with a distribution \mathcal{P} over the data domain, the p -value associated with a and \mathcal{P} is simply the probability of observing a value of the test statistic that is at least as extreme as a , assuming the data was drawn independently from \mathcal{P} : $p_i^{\mathcal{P}}(a) = \Pr_{\mathbf{X} \sim \mathcal{P}^n}[\phi_i(\mathbf{X}) \geq a]$. Note that there may be multiple distributions \mathcal{P} over the data that induce the same distribution over the test statistic. With each test statistic ϕ_i , we associate a *null hypothesis* $H_0^{(i)} \subseteq \Delta(\mathcal{X})$,¹ which is simply a collection of such distributions. The p -values are always computed with respect to a distribution $\mathcal{P} \in H_0^{(i)}$, and hence from now on, we elide the dependence on \mathcal{P} and simply write $p_i(a)$ to denote the p -value of a test statistic ϕ_i evaluated at a .

¹ $\Delta(\mathcal{X})$ denotes the set of probability distributions over \mathcal{X} .

The goal of a hypothesis test is to *reject the null hypothesis* if the data is not likely to have been generated from the proposed model, that is, if the underlying distribution from which the data were drawn was not in $H_0^{(i)}$. By definition, if \mathbf{X} truly is drawn from \mathcal{P}^n for some $\mathcal{P} \in H_0^{(i)}$, then $p_i(\phi_i(\mathbf{X}))$ is uniformly distributed over $[0, 1]$. A standard approach to hypothesis testing is to pick a *significance level* $\alpha \in [0, 1]$ (often $\alpha = 0.05$), compute the value of the test statistic $a = \phi_i(\mathbf{X})$, and then *reject* the null hypothesis if $p_i(a) \leq \alpha$. Under this procedure, the probability of incorrectly rejecting the null hypothesis—i.e., of rejecting the null hypothesis when $\mathbf{X} \sim \mathcal{P}^n$ for some $\mathcal{P} \in H_0^{(i)}$ —is at most α . An incorrect rejection of the null hypothesis is called a *false discovery*.

The discussion so far presupposes that ϕ_i , the test statistic in question, was chosen independently of the dataset \mathbf{X} . Let \mathcal{T} denote a collection of test statistics, and suppose that we select a test statistic using a data-dependent selection procedure $\mathcal{A} : \mathcal{X}^n \rightarrow \mathcal{T}$. If $\phi_i = \mathcal{A}(\mathbf{X})$, then rejecting the null hypothesis when $p_i(\phi_i(\mathbf{X})) \leq \alpha$ may result in a false discovery with probability much larger than α (indeed, this kind of naive approach to *post-selection* inference is suspected to be a primary culprit behind the prevalence of false discovery in empirical science [GL14, WL16, SNS11]). This is because even if the null hypothesis is true ($\mathbf{X} \sim \mathcal{P}^n$ for some $\mathcal{P} \in H_0^{(i)}$), the distribution on \mathbf{X} *conditioned on* $\phi_i = \mathcal{A}(\mathbf{X})$ *having been selected* need not be \mathcal{P}^n . Our goal in studying valid post-selection hypothesis testing is to find a *valid* p -value correction function $\gamma : [0, 1] \rightarrow [0, 1]$, which we define as follows:

Definition 1.1 (Valid p -value Correction Function). *A function $\gamma : [0, 1] \rightarrow [0, 1]$ is a valid p -value correction function for a selection procedure $\mathcal{A} : \mathcal{X}^n \rightarrow \mathcal{T}$ if for every significance level $\alpha \in [0, 1]$, the procedure:*

1. *Select a test statistic $\phi_i = \mathcal{A}(\mathbf{X})$ using selection procedure \mathcal{A} .*
 2. *Reject the null hypothesis $H_0^{(i)}$ if $p_i(\phi_i(\mathbf{X})) \leq \gamma(\alpha)$.*
- has probability at most α of resulting in a false discovery.*

Necessarily, to give a nontrivial correction function γ , we will need to assume that the selection procedure \mathcal{A} satisfies some useful property. In this paper, we focus on *differential privacy*, which is a measure of algorithmic stability, and more generally, max-information, which is defined in the next subsection. Differential privacy is of particular interest because it is closed under post-processing and satisfies strong composition properties. This means that, if the test statistics in \mathcal{T} are themselves differentially private, then the selection procedure \mathcal{A} can represent the decisions of a worst-case data analyst, who chooses which hypothesis tests to run in arbitrary ways as a function of the outcomes of previously selected tests.

Finally, we note that, despite the fact that previous works [DFH⁺15b, BNS⁺16] are explicitly motivated by the problem of false discovery in empirical science, most of the technical results to date have been about estimating the means of adaptively chosen predicates on the data (i.e., answering *statistical queries*) [DFH⁺15b], and more generally, estimating the values of low-sensitivity (i.e., Lipschitz continuous) functions on the dataset [BNS⁺16, RZ16, WLF16]. These kinds of results do not apply to the problem of adaptively performing hypothesis tests while generating statistically valid p -values, because p -values are by definition not low-sensitivity statistics. See Appendix B for a detailed discussion.

There is one constraint on the selection procedure \mathcal{A} that does allow us to give nontrivial p -value corrections—that \mathcal{A} should have bounded max-information. A condition of bounded mutual information has also been considered [RZ16] to give p -value corrections - but as we discuss in

Appendix C , it is possible to obtain a strictly stronger guarantee by instead reasoning via max-information. Max-information is a measure introduced by [DFH⁺15a], which we discuss next.

1.2 Max-Information (and p -values)

Given two (arbitrarily correlated) random variables X, Z , we let $X \otimes Z$ denote a random variable (in a different probability space) obtained by drawing independent copies of X and Z from their respective marginal distributions. We write \log to denote logarithms base 2.

Definition 1.2 (Max-Information [DFH⁺15a]). *Let X and Z be jointly distributed random variables over the domain $(\mathcal{X}, \mathcal{Z})$. The max-information between X and Z , denoted by $I_\infty(X; Z)$, is the minimal value of k such that for every x in the support of X and z in the support of Z , we have $\Pr[X = x | Z = z] \leq 2^k \Pr[X = x]$. Alternatively,*

$$I_\infty(X; Z) = \log \sup_{(x,z) \in (\mathcal{X}, \mathcal{Z})} \frac{\Pr[(X, Z) = (x, z)]}{\Pr[X \otimes Z = (x, z)]}.$$

The β -approximate max-information between X and Z is defined as

$$I_\infty^\beta(X; Z) = \log \sup_{\substack{\mathcal{O} \subseteq (\mathcal{X} \times \mathcal{Z}), \\ \Pr[(X, Z) \in \mathcal{O}] > \beta}} \frac{\Pr[(X, Z) \in \mathcal{O}] - \beta}{\Pr[X \otimes Z \in \mathcal{O}]}.$$

We say that an algorithm $\mathcal{A} : \mathcal{X}^n \rightarrow \mathcal{Y}$ has β -approximate max-information of k , denoted as $I_\infty^\beta(\mathcal{A}, n) \leq k$, if for every distribution \mathcal{S} over elements of \mathcal{X}^n , we have $I_\infty^\beta(\mathbf{X}; \mathcal{A}(\mathbf{X})) \leq k$ when $\mathbf{X} \sim \mathcal{S}$. We say that an algorithm $\mathcal{A} : \mathcal{X}^n \rightarrow \mathcal{Y}$ has β -approximate max-information of k over product distributions, written $I_{\infty, \mathcal{P}}^\beta(\mathcal{A}, n) \leq k$, if for every distribution \mathcal{P} over \mathcal{X} , we have $I_\infty^\beta(\mathbf{X}; \mathcal{A}(\mathbf{X})) \leq k$ when $\mathbf{X} \sim \mathcal{P}^n$.

It follows immediately from the definition that if an algorithm has bounded max-information, then we can control the probability of “bad events” that arise as a result of the dependence of $\mathcal{A}(\mathbf{X})$ on \mathbf{X} : for every event \mathcal{O} , we have $\Pr[(\mathbf{X}, \mathcal{A}(\mathbf{X})) \in \mathcal{O}] \leq 2^k \Pr[\mathbf{X} \otimes \mathcal{A}(\mathbf{X}) \in \mathcal{O}] + \beta$. For example, if \mathcal{A} is a data-dependent selection procedure for selecting a test statistic, we can derive a valid p -value correction function γ as a function of a max-information bound on \mathcal{A} :

Theorem 1.3. *Let $\mathcal{A} : \mathcal{X}^n \rightarrow \mathcal{T}$ be a data-dependent algorithm for selecting a test statistic such that $I_{\infty, \mathcal{P}}^\beta(\mathcal{A}, n) \leq k$. Then the following function γ is a valid p -value correction function for \mathcal{A} :*

$$\gamma(\alpha) = \max\left(\frac{\alpha - \beta}{2^k}, 0\right).$$

Proof. Fix a distribution \mathcal{P}^n from which the dataset \mathbf{X} is drawn. If $\frac{\alpha - \beta}{2^k} \leq 0$, then the theorem is trivial, so assume otherwise. Define $\mathcal{O} \subset \mathcal{X}^n \times \mathcal{T}$ to be the event that \mathcal{A} selects a test statistic for which the null hypothesis is true, but its p -value is at most $\gamma(\alpha)$:

$$\mathcal{O} = \{(\mathbf{x}, \phi_i) : \mathcal{P} \in H_0^{(i)} \text{ and } p_i(\phi_i(\mathbf{x})) \leq \gamma(\alpha)\}$$

Note that the event \mathcal{O} represents exactly those outcomes for which using γ as a p -value correction function results in a false discovery. Note also that, by definition of the null hypothesis, $\Pr[\mathbf{X} \otimes \mathcal{A}(\mathbf{X}) \in \mathcal{O}] \leq \gamma(\alpha) = \frac{\alpha - \beta}{2^k}$. Hence, by the guarantee that $I_{\infty, \mathcal{P}}^\beta(\mathcal{A}, n) \leq k$, we have that $\Pr[(\mathbf{X}, \mathcal{A}(\mathbf{X})) \in \mathcal{O}]$ is at most $2^k \cdot \left(\frac{\alpha - \beta}{2^k}\right) + \beta = \alpha$. \square

Because of Theorem 1.3, we are interested in methods for usefully selecting test statistics using data dependent algorithms \mathcal{A} for which we can bound their max-information. It was shown in [DFH⁺15a] that algorithms which satisfy *pure* differential privacy also have a guarantee of bounded max-information:

Theorem 1.4 (Pure Differential Privacy and Max-Information [DFH⁺15a]). *Let $\mathcal{A} : \mathcal{X}^n \rightarrow \mathcal{Y}$ be an $(\epsilon, 0)$ -differentially private algorithm. Then for every $\beta \geq 0$:*

$$I_\infty(\mathcal{A}, n) \leq \log(e) \cdot \epsilon n \quad \text{and} \quad I_{\infty, P}^\beta(\mathcal{A}, n) \leq \log(e) \cdot \left(\epsilon^2 n / 2 + \epsilon \sqrt{n \ln(2/\beta)/2} \right)$$

This connection is powerful, because there are a vast collection of data analyses for which we have differentially private algorithms— including a growing literature on differentially private hypothesis tests [JS13, USF13, YFSU14, KS16, DSZ15, She15, WLK15, GLRV16]. However, there is an important gap: Theorem 1.4 holds only for *pure* $(\epsilon, 0)$ -differential privacy, and not for approximate (ϵ, δ) -differential privacy. Many statistical analyses can be performed much more accurately subject to approximate differential privacy, and it can be easier to analyze private hypothesis tests that satisfy approximate differential privacy, because the approximate privacy constraint is amenable to perturbations using Gaussian noise (rather than Laplace noise) [GLRV16]. Most importantly, for pure differential privacy, the privacy parameter ϵ degrades *linearly* with the number of analyses performed, whereas for approximate differential privacy, ϵ need only degrade with the *square root* of the number of analyses performed [DRV10]. Hence, if the connection between max-information and differential privacy held also for approximate differential privacy, it would be possible to perform quadratically more adaptively chosen statistical tests without requiring a larger p -value correction factor.

1.3 Our Results

In addition to the framework just described for reasoning about adaptive hypothesis testing, our main technical contribution is to extend the connection between differential privacy and max-information to approximate differential privacy. We show the following (see Section 3 for a complete statement):

Theorem 3.1 (Informal). *Let $\mathcal{A} : \mathcal{X}^n \rightarrow \mathcal{Y}$ be an (ϵ, δ) -differentially private algorithm. Then,*

$$I_{\infty, P}^\beta(\mathcal{A}, n) = O\left(n\epsilon^2 + n\sqrt{\frac{\delta}{\epsilon}}\right) \quad \text{for } \beta = O\left(n\sqrt{\frac{\delta}{\epsilon}}\right).$$

It is worth noting several things. First, this bound nearly matches the bound for max-information over product distributions from Theorem 1.4, except Theorem 3.1 extends the connection to the substantially more powerful class of (ϵ, δ) -differentially private algorithms. The bound is qualitatively tight in the sense that despite its generality, it can be used to nearly recover the tight bound on the generalization properties of differentially private mechanisms for answering low-sensitivity queries that was proven using a specialized analysis in [BNS⁺16] (see Appendix D for a comparison).

We also only prove a bound on the max-information for product distributions on the input, and not for all distributions (that is, we bound $I_{\infty, P}^\beta(\mathcal{A}, n)$ and not $I_\infty^\beta(\mathcal{A}, n)$). A bound for general distributions would be desirable, since such bounds compose gracefully [DFH⁺15a]. Unfortunately,

a bound for general distributions based solely on (ϵ, δ) -differential privacy is impossible: a construction of De [De12] implies the existence of (ϵ, δ) -differentially private algorithms for which the max-information between input and output on arbitrary distributions is much larger than the bound in Theorem 3.1.

One might nevertheless hope that bounds on the max-information under product distributions can be meaningfully composed. Our second main contribution is a negative result, showing that such bounds do not compose when algorithms are selected adaptively. Specifically, we analyze the adaptive composition of two algorithms, the first of which has a small finite range (and hence, by [DFH⁺15a], small bounded max-information), and the second of which is (ϵ, δ) -differentially private. We show that the composition of the two algorithms can be used to exactly recover the input dataset, and hence, the composition does not satisfy any nontrivial max-information bound.

1.3.1 Further Interpretation

Although our presentation thus far has been motivated by p -values, an algorithm \mathcal{A} with bounded max-information allows a data analyst to treat *any event* $\mathcal{A}(\mathbf{x})$ that is a function of the output of the algorithm “as if” it is independent of the dataset \mathbf{x} , up to a correction factor determined by the max-information bound. Our results thus substantially broaden the class of analyses for which approximate differential privacy promises generalization guarantees—this class was previously limited to estimating the values of low-sensitivity numeric valued queries (and more generally, the outcomes of low-sensitivity optimization problems) [BNS⁺16].

Our result also further develops the extent to which max-information can be viewed as a unifying information theoretic measure controlling the generalization properties of adaptive data analysis. Dwork et al. [DFH⁺15a] previously showed that algorithms that satisfy bounded output description length, and algorithms that satisfy pure differential privacy (two constraints known individually to imply adaptive generalization guarantees), both have bounded max-information. Because bounded max-information satisfies strong composition properties, this connection implies that algorithms with bounded output description length and pure differentially private algorithms can be composed in arbitrary order and the resulting composition will still have strong generalization properties. Our result brings approximate differential privacy partially into this unifying framework. In particular, *when the data is drawn from a product distribution*, if an analysis that starts with an (arbitrary) approximate differentially private computation is followed by an arbitrary composition of algorithms with bounded max-information, then the resulting composition will satisfy a max-information bound. However, unlike with compositions consisting solely of bounded description length mechanisms and pure differentially private mechanisms, which can be composed in arbitrary order, in this case *it is important that the approximate differentially private computation come first*. This is because, even if the dataset \mathbf{x} is initially drawn from a product distribution, the conditional distribution on the data that results after observing the outcome of an initial computation need not be a product distribution any longer. In fact, the lower bound we prove in Section ?? is an explicit construction in which the composition of a bounded description length algorithm, followed by an approximate differentially private algorithm can be used to exactly reconstruct a dataset drawn from a product distribution (which can in turn be used to arbitrarily overfit that dataset).

Finally, we draw a connection between max-information and mutual information that allows us to improve on several prior results that dealt with mutual information [MMP⁺11, RZ16]. We present the proof in Appendix C.

Lemma 1.5. *Let \mathbf{X}, \mathbf{Y} be a pair of discrete random variables defined on the same probability space, where \mathbf{X} takes values in a finite set Σ .*

- *If $I(\mathbf{X}; \mathbf{Y}) \leq m$ then, for every $k > 0$, we have $I_\infty^{\beta(k)}(\mathbf{X}; \mathbf{Y}) \leq k$ where $\beta(k) = \frac{m+0.54}{k}$.*
- *If $I_\infty^\beta(\mathbf{X}; \mathbf{Y}) \leq k$ for $k > 0$ and $0 \leq \beta \leq \frac{3(1-2^{-k})}{20}$, then² $I(\mathbf{X}; \mathbf{Y}) \leq 2k \ln(2) + \frac{2\beta \log(|\Sigma|/2\beta)}{1-2^{-k}}$.*

We are able to improve on the p -value correction function implicitly given in [RZ16] given a mutual information bound, by first converting mutual information to a max-information bound and applying the p -value correction function from this paper. Our main theorem Theorem 3.1 combined with Lemma 1.5 also obtains an improved bound on the mutual information of approximately private mechanisms from Proposition 4.4 in [MMP⁺11]³. The following corollary improves the bound from [MMP⁺11] in its dependence on $|\mathcal{X}|$ from $|\mathcal{X}|^2 \cdot \log(1/|\mathcal{X}|)$ to $\log|\mathcal{X}|$.

Corollary 1.6. *Let $\mathcal{A} : \mathcal{X}^n \rightarrow \mathcal{T}$ be (ϵ, δ) -differentially private and $\mathbf{X} \sim \mathcal{P}^n$. If⁴ $\epsilon \in (0, 1/2)$, $\epsilon = \Omega\left(\frac{1}{\sqrt{n}}\right)$, and $\delta = O\left(\frac{\epsilon}{n^2}\right)$, we then have:*

$$I(\mathbf{X}; \mathcal{A}(\mathbf{X})) = O\left(n\epsilon^2 + n\sqrt{\frac{\delta}{\epsilon}}\left(1 + \ln\left(\frac{1}{n}\sqrt{\frac{\epsilon}{\delta}}\right) + n \log|\mathcal{X}|\right)\right).$$

1.4 Other Related Work

Differential Privacy is an algorithmic stability condition introduced by Dwork et al. [DMNS06]. Its connection to adaptive data analysis was made by Dwork et al. [DFH⁺15b] and both strengthened and generalized by Bassily et al. [BNS⁺16]. Dwork et al. [DFH⁺15a] showed that algorithms with bounded description length outputs have similar guarantees for adaptive data analysis, and introduced the notion of max-information. Cummings et al. [CLN⁺16] give a third method—compression schemes—which can also guarantee validity in adaptive data analysis in the context of learning. Computational and information theoretic lower bounds for adaptively estimating means in this framework were proven by Hardt and Ullman [HU14], and Steinke and Ullman [SU15].

Russo and Zou [RZ16] show how to bound the *bias* of sub-gaussian statistics selected in a data-dependent manner, in terms of the mutual information between the selection procedure and the value of the statistics. In particular (using our terminology), they show how to give a valid p -value correction function in terms of this mutual information. In Appendix C, we demonstrate that if a bound on the mutual information between the dataset and the output of the selection procedure is known, then it is possible to substantially improve on the p -value correction function given by [RZ16] by instead using the mutual information bound to prove a max-information bound on the selection procedure. [WLF16] study adaptive data analysis in a similar framework to [RZ16], and give a minimax analysis in a restricted setting.

McGregor et al. [MMP⁺11], and De [De12] also study (among other things) information theoretic bounds satisfied by differentially private algorithms. Together, they prove a result that is analogous to ours, for *mutual information*—that while pure differentially private algorithms have

²The bound stated here is looser than that claimed in the conference version [RRST16]; we do not know a proof of the original claim.

³Thanks to Salil Vadhan for pointing out to us this implication.

⁴The parameter ranges stated here are different than those claimed in the conference version [RRST16]; this modification is a result of the change in Lemma 1.5.

bounded mutual information between their inputs and their outputs, a similar bound holds for (approximate) (ϵ, δ) -differentially private algorithms only if the data is drawn from a product distribution. We improve quantitatively on this bound.

2 Preliminaries

We will use the following vector notation throughout: $\mathbf{x} = (x_1, \dots, x_n)$, $\mathbf{x}_a^b = (x_a, x_{a+1}, \dots, x_b)$, $(\mathbf{x}_{-i}, t) = (x_1, \dots, x_{i-1}, t, x_{i+1}, \dots, x_n)$. We denote the distribution of a random variable X as $p(X)$. In our analysis, jointly distributed random variables (X, Z) will typically be of the form $(\mathbf{X}, \mathcal{A}(\mathbf{X}))$ where $\mathbf{X} \sim \mathcal{P}^n$ is a dataset of n elements sampled from domain \mathcal{X} , and $\mathcal{A} : \mathcal{X}^n \rightarrow \mathcal{Y}$ is a (randomized) algorithm that maps a dataset to some range \mathcal{Y} . We denote by $\mathcal{A}(\mathbf{X})$ the random variable that results when \mathcal{A} is applied to a dataset $\mathbf{X} \sim \mathcal{P}^n$ (note that here, the randomness is both over the choice of dataset, and the internal coins of the algorithm). When the input variable is understood, we will sometimes simply write \mathcal{A} .

It will be useful in our analysis to compare the distributions of two random variables. In the introduction, we define (approximate-) max-information, and we now give some other measures between distributions. We first define indistinguishability, and then differential privacy.

Definition 2.1 (Indistinguishability [KS14]). *Two random variables X, Y taking values in a set \mathcal{D} are (ϵ, δ) -indistinguishable, denoted $X \approx_{\epsilon, \delta} Y$, if for all $\mathcal{O} \subseteq \mathcal{D}$,*

$$\Pr[X \in \mathcal{O}] \leq e^\epsilon \cdot \Pr[Y \in \mathcal{O}] + \delta \text{ and}$$

$$\Pr[Y \in \mathcal{O}] \leq e^\epsilon \cdot \Pr[X \in \mathcal{O}] + \delta.$$

Definition 2.2 (Point-wise indistinguishability [KS14]). *Two random variables X, Z taking values in a set \mathcal{D} are point-wise (ϵ, δ) -indistinguishable if with probability at least $1 - \delta$ over $a \sim p(X)$:*

$$e^{-\epsilon} \Pr[Z = a] \leq \Pr[X = a] \leq e^\epsilon \Pr[Z = a].$$

Before we define differential privacy, we say that two databases $\mathbf{x}, \mathbf{x}' \in \mathcal{X}^n$ are *neighboring* if they differ in at most one entry. We now define differential privacy in terms of indistinguishability:

Definition 2.3 (Differential Privacy [DMNS06, DKM⁺06]). *A randomized algorithm $\mathcal{A} : \mathcal{X}^n \rightarrow \mathcal{Y}$ is (ϵ, δ) -differentially private if for all neighboring datasets $\mathbf{x}, \mathbf{x}' \in \mathcal{X}^n$, we have $\mathcal{A}(\mathbf{x}) \approx_{\epsilon, \delta} \mathcal{A}(\mathbf{x}')$.*

In the appendix, we give several useful connections between these definitions along with other more widely known measures between distributions, e.g., KL-divergence, and total-variation distance.

3 Max-Information for (ϵ, δ) -Differentially Private Algorithms

In this section, we prove a bound on approximate max-information for (ϵ, δ) -differentially private algorithms over product distributions.

Theorem 3.1. Let $\mathcal{A} : \mathcal{X}^n \rightarrow \mathcal{Y}$ be an (ϵ, δ) -differentially private algorithm for $\epsilon \in (0, 1/2]$ and $\delta \in (0, \epsilon)$. For $\beta = e^{-\epsilon^2 n} + O\left(n\sqrt{\frac{\delta}{\epsilon}}\right)$, we have

$$I_{\infty, P}^{\beta}(\mathcal{A}, n) = O\left(\epsilon^2 n + n\sqrt{\frac{\delta}{\epsilon}}\right).$$

We will prove Theorem 3.1 over the course of this section, using a number of lemmas. We first set up some notation. We will sometimes abbreviate conditional probabilities of the form $\Pr[\mathbf{X} = \mathbf{x} | \mathcal{A} = a]$ as $\Pr[\mathbf{X} = \mathbf{x} | a]$ when the random variables are clear from context. Further, for any $\mathbf{x} \in \mathcal{X}^n$ and $a \in \mathcal{Y}$, we define

$$\begin{aligned} Z(a, \mathbf{x}) &\stackrel{\text{def}}{=} \log\left(\frac{\Pr[\mathcal{A} = a, \mathbf{X} = \mathbf{x}]}{\Pr[\mathcal{A} = a] \cdot \Pr[\mathbf{X} = \mathbf{x}]}\right) \\ &= \sum_{i=1}^n \log\left(\frac{\Pr[X_i = x_i | a, \mathbf{x}_1^{i-1}]}{\Pr[X_i = x_i]}\right) \end{aligned} \quad (1)$$

If we can bound $Z(a, \mathbf{x})$ with high probability over $(a, \mathbf{x}) \sim p(\mathcal{A}(\mathbf{X}), \mathbf{X})$, then we can bound the approximate max-information by using the following lemma:

Lemma 3.2 ([DFH⁺15a, Lemma 18]). *If $\Pr[Z(\mathcal{A}(\mathbf{X}), \mathbf{X}) \geq k] \leq \beta$, then $I_{\infty}^{\beta}(\mathcal{A}(\mathbf{X}); \mathbf{X}) \leq k$.*

We next define each term in the sum of $Z(a, \mathbf{x})$ as

$$Z_i(a, \mathbf{x}_1^i) \stackrel{\text{def}}{=} \log \frac{\Pr[X_i = x_i | a, \mathbf{x}_1^{i-1}]}{\Pr[X_i = x_i]}. \quad (2)$$

The plan of the proof is simple: our goal is to apply Azuma's inequality (Theorem A.8) to the sum of the Z_i 's to achieve a bound on Z with high probability. Applying Azuma's inequality requires both understanding the expectation of each term $Z_i(a, \mathbf{x}_1^i)$, and being able to argue that each term is bounded. Unfortunately, in our case, the terms are not always bounded – however, we will be able to show that they are bounded with high probability. This plan is somewhat complicated by the conditioning in the definition of $Z_i(a, \mathbf{x}_1^i)$.

First, we argue that we can bound each Z_i with high probability. This argument takes place over the course of Claims 3.3, 3.4, 3.5 and 3.6.

Claim 3.3. *If \mathcal{A} is (ϵ, δ) -differentially private and $\mathbf{X} \sim \mathcal{P}^n$, then for each $i \in [n]$ and each prefix $\mathbf{x}_1^{i-1} \in \mathcal{X}^{i-1}$, we have:*

$$(\mathcal{A}, X_i) |_{\mathbf{x}_1^{i-1}} \approx_{\epsilon, \delta} \mathcal{A} |_{\mathbf{x}_1^{i-1}} \otimes X_i.$$

Proof. Fix any set $\mathcal{O} \subseteq \mathcal{Y} \times \mathcal{X}$ and prefix $\mathbf{x}_1^{i-1} \in \mathcal{X}^{i-1}$. We then define the set $\mathcal{O}_{x_i} = \{a \in \mathcal{Y} : (a, x_i) \in \mathcal{O}\}$. Now, we have that:

$$\begin{aligned} \Pr[(\mathcal{A}(\mathbf{X}), X_i) \in \mathcal{O} | \mathbf{x}_1^{i-1}] &= \sum_{x_i \in \mathcal{X}} \Pr[X_i = x_i] \Pr[\mathcal{A}(\mathbf{X}) \in \mathcal{O}_{x_i} | \mathbf{x}_1^{i-1}, x_i] \\ &\leq \sum_{x_i \in \mathcal{X}} \Pr[X_i = x_i] (e^{\epsilon} \Pr[\mathcal{A}(\mathbf{X}) \in \mathcal{O}_{x_i} | \mathbf{x}_1^{i-1}, t_i] + \delta) \quad \forall t_i \in \mathcal{X} \end{aligned}$$

Thus, we can multiply both sides of the inequality by $\Pr[X_i = t_i]$ and sum over all $t_i \in \mathcal{X}$ to get:

$$\begin{aligned}
\Pr[(\mathcal{A}(\mathbf{X}), X_i) \in \mathcal{O} | \mathbf{x}_1^{i-1}] &= \sum_{t_i \in \mathcal{X}} \Pr[X_i = t_i] \Pr[(\mathcal{A}(\mathbf{X}), X_i) \in \mathcal{O} | \mathbf{x}_1^{i-1}] \\
&\leq \sum_{x_i \in \mathcal{X}} \sum_{t_i \in \mathcal{X}} \Pr[X_i = x_i] \Pr[X_i = t_i] (e^\epsilon \Pr[\mathcal{A}(\mathbf{X}) \in \mathcal{O}_{x_i} | \mathbf{x}_1^{i-1}, t_i] + \delta) \\
&\leq e^\epsilon \sum_{x_i \in \mathcal{X}} \Pr[X_i = x_i] \Pr[\mathcal{A}(\mathbf{X}) \in \mathcal{O}_{x_i} | \mathbf{x}_1^{i-1}] + \delta = e^\epsilon \Pr[\mathcal{A}(\mathbf{X}) \otimes X_i \in \mathcal{O} | \mathbf{x}_1^{i-1}] + \delta.
\end{aligned}$$

We follow a similar argument to prove:

$$\Pr[\mathcal{A}(\mathbf{X}) \otimes X_i \in \mathcal{O} | \mathbf{x}_1^{i-1}] \leq e^\epsilon \Pr[(\mathcal{A}(\mathbf{X}), X_i) \in \mathcal{O} | \mathbf{x}_1^{i-1}] + \delta.$$

□

We now define the following set of “good outcomes and prefixes” for any $\hat{\delta} > 0$:

$$\mathcal{E}_i(\hat{\delta}) = \left\{ (a, \mathbf{x}_1^{i-1}) : X_i \approx_{3\epsilon, \hat{\delta}} X_i |_{a, \mathbf{x}_1^{i-1}} \right\} \quad (3)$$

We use a technical lemma from [KS14] (stated in the appendix, Lemma A.7), and Claim 3.3 to derive the following result:

Claim 3.4. *If \mathcal{A} is (ϵ, δ) -differentially private and $\mathbf{X} \sim \mathcal{P}^n$, then for each $i \in [n]$ and each prefix $\mathbf{x}_1^{i-1} \in \mathcal{X}^{i-1}$ we have for $\hat{\delta} > 0$ and $\delta' \stackrel{\text{def}}{=} \frac{2\hat{\delta}}{\delta} + \frac{2\delta}{1-e^{-\epsilon}}$:*

$$\Pr[(\mathcal{A}, \mathbf{X}_1^{i-1}) \in \mathcal{E}_i(\hat{\delta}) | \mathbf{x}_1^{i-1}] \geq 1 - \delta'.$$

Proof. This follows directly from Lemma A.7:

$$\Pr[(\mathcal{A}, \mathbf{X}_1^{i-1}) \in \mathcal{E}_i(\hat{\delta}) | \mathbf{x}_1^{i-1}] = \Pr[X_i \approx_{3\epsilon, \hat{\delta}} X_i |_{\mathcal{A}, \mathbf{x}_1^{i-1}} | \mathbf{x}_1^{i-1}] = \Pr_{a \sim p(\mathcal{A} |_{\mathbf{x}_1^{i-1}})} [X_i \approx_{3\epsilon, \hat{\delta}} X_i |_{a, \mathbf{x}_1^{i-1}}] \geq 1 - \delta'$$

□

We now define the set of outcome/dataset prefix pairs for which the quantities Z_i are not large:

$$\mathcal{F}_i = \left\{ (a, \mathbf{x}_1^i) : |Z_i(a, \mathbf{x}_1^i)| \leq 6\epsilon \right\}. \quad (4)$$

Using another technical lemma from [KS14] (which we state in Lemma A.6 in the appendix), we prove:

Claim 3.5. *Given $(a, \mathbf{x}_1^{i-1}) \in \mathcal{E}_i(\hat{\delta})$ and $\delta'' \stackrel{\text{def}}{=} \frac{2\hat{\delta}}{1-e^{-3\epsilon}}$ we have:*

$$\Pr[(\mathcal{A}, \mathbf{X}_1^i) \in \mathcal{F}_i | a, \mathbf{x}_1^{i-1}] \geq 1 - \delta''.$$

Proof. Since $(a, \mathbf{x}_1^{i-1}) \in \mathcal{E}_i(\hat{\delta})$, we know that X_i is $(3\epsilon, \hat{\delta})$ -indistinguishable from $X_i|_{a, \mathbf{x}_1^{i-1}}$. Using Lemma A.6, we know that X_i and $X_i|_{a, \mathbf{x}_1^{i-1}}$ are point-wise $(6\epsilon, \delta'')$ -indistinguishable. Thus, by definition of \mathcal{F}_i and Z_i , we have:

$$\begin{aligned} \Pr [(\mathcal{A}, \mathbf{X}_1^i) \in \mathcal{F}_i | a, \mathbf{x}_1^{i-1}] &= \Pr [Z_i(\mathcal{A}, \mathbf{X}_1^i) \leq 6\epsilon | a, \mathbf{x}_1^{i-1}] \\ &= \Pr_{x_i \sim p(X_i|_{a, \mathbf{x}_1^{i-1}})} \left[\log \left(\frac{\Pr[X_i = x_i | a, \mathbf{x}_1^{i-1}]}{\Pr[X_i = x_i]} \right) \leq 6\epsilon \right] \geq 1 - \delta'' \end{aligned}$$

□

We now define the “good” tuples of outcomes and databases as

$$\mathcal{G}_i(\hat{\delta}) = \left\{ (a, \mathbf{x}_1^i) : (a, \mathbf{x}_1^{i-1}) \in \mathcal{E}_i(\hat{\delta}) \quad \& \quad (a, \mathbf{x}_1^i) \in \mathcal{F}_i \right\}, \quad (5)$$

$$\mathcal{G}_{\leq i}(\hat{\delta}) = \left\{ (a, \mathbf{x}_1^i) : (a, x_1) \in \mathcal{G}_1(\hat{\delta}), \dots, (a, \mathbf{x}_1^i) \in \mathcal{G}_i(\hat{\delta}) \right\} \quad (6)$$

Claim 3.6. *If \mathcal{A} is (ϵ, δ) -differentially private and $\mathbf{X} \sim \mathcal{P}^n$, then*

$$\Pr [(\mathcal{A}, \mathbf{X}_1^i) \in \mathcal{G}_i(\hat{\delta})] \geq 1 - \delta' - \delta''$$

for δ' and δ'' given in Claim 3.4 and Claim 3.5, respectively.

Proof. We have:

$$\begin{aligned} \Pr [(\mathcal{A}, \mathbf{X}_1^i) \notin \mathcal{G}_i(\hat{\delta})] &= \Pr [(\mathcal{A}, \mathbf{X}_1^{i-1}) \notin \mathcal{E}_i(\hat{\delta}) \quad \text{or} \quad (\mathcal{A}, \mathbf{X}_1^i) \notin \mathcal{F}_i] \\ &= 1 - \Pr [(\mathcal{A}, \mathbf{X}_1^{i-1}) \in \mathcal{E}_i(\hat{\delta}) \quad \text{and} \quad (\mathcal{A}, \mathbf{X}_1^i) \in \mathcal{F}_i] \\ &= 1 - \sum_{(a, \mathbf{x}_1^{i-1}) \in \mathcal{E}_i(\hat{\delta})} \Pr [(\mathcal{A}, \mathbf{X}_1^{i-1}) = (a, \mathbf{x}_1^{i-1})] \Pr [(\mathcal{A}, \mathbf{X}_1^i) \in \mathcal{F}_i | a, \mathbf{x}_1^{i-1}] \\ &\leq 1 - \sum_{(a, \mathbf{x}_1^{i-1}) \in \mathcal{E}_i(\hat{\delta})} \Pr [(\mathcal{A}, \mathbf{X}_1^{i-1}) = (a, \mathbf{x}_1^{i-1})] \cdot (1 - \delta'') \\ &= 1 - (1 - \delta'') \Pr [(\mathcal{A}, \mathbf{X}_1^{i-1}) \in \mathcal{E}_i(\hat{\delta})] \leq 1 - (1 - \delta'')(1 - \delta') = \delta' + \delta'' - \delta'\delta'' \end{aligned}$$

where the last two inequalities follow from Claim 3.5 and Claim 3.4, respectively. □

Having shown a high probability bound on the terms Z_i , our next step is to bound their expectation so that we can continue towards our goal of applying Azuma’s inequality. Note a complicating factor – throughout the argument, we need to condition on the event $(\mathcal{A}, \mathbf{X}_1^i) \in \mathcal{F}_i$ to ensure that Z_i has bounded expectation.

We will use the following shorthand notation for conditional expectation:

$$\begin{aligned} \mathbb{E} [Z_i(\mathcal{A}, \mathbf{X}_1^i) | a, \mathbf{x}_1^{i-1}, \mathcal{F}_i] \\ \stackrel{def}{=} \mathbb{E} [Z_i(\mathcal{A}, \mathbf{X}_1^i) | \mathcal{A} = a, \mathbf{X}_1^{i-1} = \mathbf{x}_1^{i-1}, (\mathcal{A}, \mathbf{X}_1^i) \in \mathcal{F}_i], \end{aligned}$$

with similar notation for sets $\mathcal{G}_i(\hat{\delta}), \mathcal{G}_{\leq i}(\hat{\delta})$.

Lemma 3.7. Let \mathcal{A} be (ϵ, δ) -differentially private and $\mathbf{X} \sim \mathcal{P}^n$. Given $(a, \mathbf{x}_1^{i-1}) \in \mathcal{E}_i(\hat{\delta})$, for all $\epsilon \in (0, 1/2]$ and $\hat{\delta} \in (0, \epsilon/15]$,

$$\mathbb{E} [Z_i(\mathcal{A}, \mathbf{X}_1^i) | a, \mathbf{x}_1^{i-1}, \mathcal{F}_i] = O(\epsilon^2 + \hat{\delta}).$$

More precisely, $\mathbb{E} [Z_i(\mathcal{A}, \mathbf{X}_1^i) | a, \mathbf{x}_1^{i-1}, \mathcal{F}_i] \leq \nu(\hat{\delta})$, where $\nu(\hat{\delta})$ is defined in (7).

Proof. Given an outcome and prefix $(a, \mathbf{x}_1^{i-1}) \in \mathcal{E}_i(\hat{\delta})$, we define the set of data entries

$$\mathcal{X}(a, \mathbf{x}_1^{i-1}) = \{x_i \in \mathcal{X} : (a, \mathbf{x}_1^i) \in \mathcal{F}_i\}.$$

We then have:

$$\mathbb{E} [Z_i(\mathcal{A}, \mathbf{X}_1^i) | a, \mathbf{x}_1^{i-1}, \mathcal{F}_i] = \sum_{x_i \in \mathcal{X}(a, \mathbf{x}_1^{i-1})} \Pr [X_i = x_i | a, \mathbf{x}_1^{i-1}, \mathcal{F}_i] \log \left(\frac{\Pr [X_i = x_i | a, \mathbf{x}_1^{i-1}]}{\Pr [X_i = x_i]} \right)$$

Here, our goal is to mimic the proof of the ‘‘advanced composition theorem’’ of [DRV10] by adding a term that looks like a KL divergence term (see Definition A.4). In our case, however, the sum is not over the entire set \mathcal{X} , and so it is not a KL-divergence, which leads to some additional complications. Consider the following term:

$$\begin{aligned} & \sum_{x_i \in \mathcal{X}(a, \mathbf{x}_1^{i-1})} \Pr [X_i = x_i] \log \left(\frac{\Pr [X_i = x_i | a, \mathbf{x}_1^{i-1}]}{\Pr [X_i = x_i]} \right) \\ &= \Pr [X_i \in \mathcal{X}(a, \mathbf{x}_1^{i-1})] \sum_{x_i \in \mathcal{X}(a, \mathbf{x}_1^{i-1})} \frac{\Pr [X_i = x_i]}{\Pr [X_i \in \mathcal{X}(a, \mathbf{x}_1^{i-1})]} \log \left(\frac{\Pr [X_i = x_i | a, \mathbf{x}_1^{i-1}]}{\Pr [X_i = x_i]} \right) \\ &\leq \log \left(\frac{\Pr [X_i \in \mathcal{X}(a, \mathbf{x}_1^{i-1}) | a, \mathbf{x}_1^{i-1}]}{\Pr [X_i \in \mathcal{X}(a, \mathbf{x}_1^{i-1})]} \right) = \log \left(\frac{1 - \Pr [X_i \notin \mathcal{X}(a, \mathbf{x}_1^{i-1}) | a, \mathbf{x}_1^{i-1}]}{1 - \Pr [X_i \notin \mathcal{X}(a, \mathbf{x}_1^{i-1})]} \right) \end{aligned}$$

where the inequality follows from Jensen’s inequality. Note that, because $(a, \mathbf{x}_1^{i-1}) \in \mathcal{E}_i(\hat{\delta})$, we have for $\hat{\delta} > 0$:

$$\Pr [X_i \notin \mathcal{X}(a, \mathbf{x}_1^{i-1})] \leq e^{3\epsilon} \Pr [X_i \notin \mathcal{X}(a, \mathbf{x}_1^{i-1}) | a, \mathbf{x}_1^{i-1}] + \hat{\delta}.$$

We now focus on the term $\Pr [X_i \notin \mathcal{X}(a, \mathbf{x}_1^{i-1}) | a, \mathbf{x}_1^{i-1}]$. Note that $x_i \notin \mathcal{X}(a, \mathbf{x}_1^{i-1}) \Leftrightarrow (a, \mathbf{x}_1^i) \notin \mathcal{F}_i$. Thus,

$$\Pr [X_i \notin \mathcal{X}(a, \mathbf{x}_1^{i-1}) | a, \mathbf{x}_1^{i-1}] = \Pr [(\mathcal{A}, \mathbf{X}_1^i) \notin \mathcal{F}_i | a, \mathbf{x}_1^{i-1}] \stackrel{\text{def}}{=} q$$

Note that $q \leq \delta''$ by Claim 3.5. Now, we can bound the following:

$$\begin{aligned} & \sum_{x_i \in \mathcal{X}(a, \mathbf{x}_1^{i-1})} \Pr [X_i = x_i] \log \left(\frac{\Pr [X_i = x_i | a, \mathbf{x}_1^{i-1}]}{\Pr [X_i = x_i]} \right) \leq \log(1 - q) - \log(1 - (e^{3\epsilon} q + \hat{\delta})) \\ &\leq \log(e) \cdot (-q + e^{3\epsilon} q + \hat{\delta} + 2(e^{3\epsilon} q + \hat{\delta})^2) = \log(e) \cdot ((e^{3\epsilon} - 1)q + \hat{\delta} + 2(e^{3\epsilon} q + \hat{\delta})^2) \\ &\leq \log(e) \cdot \left((e^{3\epsilon} - 1) \frac{2\hat{\delta}}{1 - e^{-3\epsilon}} + \hat{\delta} + 2\hat{\delta}^2 \cdot \left(\frac{2e^{3\epsilon}}{1 - e^{-3\epsilon}} + 1 \right)^2 \right) \\ &= \hat{\delta}(\log(e)(2e^{3\epsilon} + 1)) + \hat{\delta}^2 \left(2\log(e) \left(\frac{4e^{12\epsilon} + 4e^{9\epsilon} - 3e^{6\epsilon} - 2e^{3\epsilon} + 1}{e^{6\epsilon} - 2e^{3\epsilon} + 1} \right) \right) \stackrel{\text{def}}{=} \tau(\hat{\delta}) \end{aligned}$$

where the second inequality follows by using the inequality $(-x-2x^2)\log(e) \leq \log(1-x) \leq -x\log(e)$ for $0 < x \leq 1/2$, and as $(e^{3\epsilon}q + \hat{\delta}) \leq 1/2$ for ϵ and $\hat{\delta}$ bounded as in the lemma statement.

We then use this result to upper bound the expectation we wanted:

$$\begin{aligned}
& \mathbb{E} [Z_i(\mathcal{A}, \mathbf{X}_1^i) | a, \mathbf{x}_1^{i-1}, \mathcal{F}_i] \\
& \leq \sum_{x_i \in \mathcal{X}(a, \mathbf{x}_1^{i-1})} \Pr [X_i = x_i | a, \mathbf{x}_1^{i-1}, \mathcal{F}_i] \log \left(\frac{\Pr [X_i = x_i | a, \mathbf{x}_1^{i-1}]}{\Pr [X_i = x_i]} \right) \\
& \quad - \sum_{x_i \in \mathcal{X}(a, \mathbf{x}_1^{i-1})} \Pr [X_i = x_i] \log \left(\frac{\Pr [X_i = x_i | a, \mathbf{x}_1^{i-1}]}{\Pr [X_i = x_i]} \right) + \tau(\hat{\delta}) \\
& = \sum_{x_i \in \mathcal{X}(a, \mathbf{x}_1^{i-1})} (\Pr [X_i = x_i | a, \mathbf{x}_1^{i-1}, \mathcal{F}_i] - \Pr [X_i = x_i]) \log \left(\frac{\Pr [X_i = x_i | a, \mathbf{x}_1^{i-1}]}{\Pr [X_i = x_i]} \right) + \tau(\hat{\delta}) \\
& \leq 6\epsilon \sum_{x_i \in \mathcal{X}(a, \mathbf{x}_1^{i-1})} |\Pr [X_i = x_i | a, \mathbf{x}_1^{i-1}, \mathcal{F}_i] - \Pr [X_i = x_i]| + \tau(\hat{\delta}) \\
& \leq 6\epsilon \sum_{x_i \in \mathcal{X}(a, \mathbf{x}_1^{i-1})} \Pr [X_i = x_i] \max \left\{ \frac{e^{6\epsilon}}{\Pr [(\mathcal{A}, \mathbf{X}_1^i) \in \mathcal{F}_i | a, \mathbf{x}_1^{i-1}]} - 1, 1 - \frac{e^{-6\epsilon}}{\Pr [(\mathcal{A}, \mathbf{X}_1^i) \in \mathcal{F}_i | a, \mathbf{x}_1^{i-1}]} \right\} \\
& \quad + \tau(\hat{\delta}) \\
& \leq 6\epsilon \left(\frac{e^{6\epsilon}}{1 - \frac{2\hat{\delta}}{1 - e^{-3\epsilon}}} - 1 \right) + \tau(\hat{\delta}) \leq 6\epsilon \left(e^{6\epsilon} \left(1 + \frac{4\hat{\delta}}{1 - e^{-3\epsilon}} \right) - 1 \right) + \tau(\hat{\delta}) \\
& \leq 72\epsilon^2 + \hat{\delta} \left(\frac{24e^{6\epsilon}}{1 - e^{-3\epsilon}} + \log(e)(2e^{3\epsilon} + 1) \right) + \hat{\delta}^2 \left(2\log(e) \left(\frac{4e^{12\epsilon} + 4e^{9\epsilon} - 3e^{6\epsilon} - 2e^{3\epsilon} + 1}{e^{6\epsilon} - 2e^{3\epsilon} + 1} \right) \right) \\
& \stackrel{\text{def}}{=} \nu(\hat{\delta}) \tag{7}
\end{aligned}$$

where the third inequality follows from the definition of \mathcal{F}_i , the fourth inequality follows from Claim 3.5 and the last inequality follows by substituting the value of $\tau(\hat{\delta})$, and using the inequalities $1 + y \leq e^y$ and $e^{ky} \leq 1 + e^k y$ for $y \in (0, 0.5]$, $k > 1$. \square

Finally, we need to apply Azuma's inequality (stated in Theorem A.8) to a set of variables that are bounded with probability 1, not just with high probability. Towards this end, we define variables T_i that will match Z_i for "good events", and will be zero otherwise—and hence, are always bounded:

$$T_i(a, \mathbf{x}_1^i) = \begin{cases} Z_i(a, \mathbf{x}_1^i) & \text{if } (a, \mathbf{x}_1^i) \in \mathcal{G}_{\leq i}(\hat{\delta}) \\ 0 & \text{otherwise} \end{cases} \tag{8}$$

The next lemma verifies that the variables T_i indeed satisfy the requirements of Azuma's inequality:

Lemma 3.8. *Let \mathcal{A} be (ϵ, δ) -differentially private and $\mathbf{X} \sim \mathcal{P}^n$. The variables T_i defined in (8) are bounded by 6ϵ with probability 1, and for any $(a, \mathbf{x}_1^{i-1}) \in \mathcal{Y} \times \mathcal{X}^{i-1}$ and $\hat{\delta} \in [0, \epsilon/15]$,*

$$\mathbb{E} [T_i(\mathcal{A}, \mathbf{X}_1^i) | a, \mathbf{x}_1^{i-1}] = O(\epsilon^2 + \hat{\delta}/\epsilon), \tag{9}$$

where the bound does not depend on n or i . More precisely, $\mathbb{E} [T_i(\mathcal{A}, \mathbf{X}_1^i) | a, \mathbf{x}_1^{i-1}] \leq \nu(\hat{\delta})$, where $\nu(\hat{\delta})$ is defined in (7).

Proof. By definition, $T_i(\mathcal{A}, \mathbf{X}_1^i)$ takes values only in $[-6\epsilon, 6\epsilon]$. Thus,

$$\Pr [|T_i(\mathcal{A}, \mathbf{X}_1^i)| \leq 6\epsilon] = 1.$$

Now, given $(a, \mathbf{x}_1^{i-1}) \in \mathcal{E}_i(\hat{\delta}) \cap \mathcal{G}_{\leq i-1}(\hat{\delta})$, we can see that:

$$\mathbb{E} \left[T_i(\mathcal{A}, \mathbf{X}_1^i) \middle| a, \mathbf{x}_1^{i-1}, \mathcal{G}_{\leq i}^c(\hat{\delta}) \right] = 0.$$

Further, given $(a, \mathbf{x}_1^{i-1}) \in \mathcal{E}_i(\hat{\delta}) \cap \mathcal{G}_{\leq i-1}(\hat{\delta})$, we have:

$$\begin{aligned} \mathbb{E} \left[T_i(\mathcal{A}, \mathbf{X}_1^i) | a, \mathbf{x}_1^{i-1}, \mathcal{G}_{\leq i}(\hat{\delta}) \right] &= \sum_{x_i: (a, \mathbf{x}_1^i) \in \mathcal{F}_i} T_i(a, \mathbf{x}_1^i) \Pr \left[X_i = x_i | a, \mathbf{x}_1^{i-1}, \mathcal{G}_{\leq i}(\hat{\delta}) \right] \\ &= \sum_{x_i: (a, \mathbf{x}_1^i) \in \mathcal{F}_i} Z_i(a, \mathbf{x}_1^i) \Pr \left[X_i = x_i | a, \mathbf{x}_1^{i-1}, \mathcal{G}_{\leq i}(\hat{\delta}) \right] = \sum_{x_i: (a, \mathbf{x}_1^i) \in \mathcal{F}_i} Z_i(a, \mathbf{x}_1^i) \Pr \left[X_i = x_i | a, \mathbf{x}_1^{i-1}, \mathcal{F}_i \right] \\ &= \mathbb{E} \left[Z_i(\mathcal{A}, \mathbf{X}_1^i) | a, \mathbf{x}_1^{i-1}, \mathcal{F}_i \right] = O(\epsilon^2 + \hat{\delta}/\epsilon) \end{aligned}$$

where the second equality follows from (8), and the last equality follows from Lemma 3.7. For any $(a, \mathbf{x}_1^{i-1}) \notin \mathcal{E}_i(\hat{\delta}) \cap \mathcal{G}_{\leq i-1}(\hat{\delta})$, we have that the conditional expectation is zero. This proves the lemma. \square

We are now ready to prove our main theorem.

Proof of Theorem 3.1. For any constant ν , we have:

$$\begin{aligned} &\Pr \left[\sum_{i=1}^n Z_i(\mathcal{A}, \mathbf{X}_1^i) > n\nu + 6t\epsilon\sqrt{n} \right] \\ &\leq \Pr \left[\sum_{i=1}^n Z_i(\mathcal{A}, \mathbf{X}_1^i) > n\nu + 6t\epsilon\sqrt{n} \cap (\mathcal{A}, \mathbf{X}) \in \mathcal{G}_{\leq n}(\hat{\delta}) \right] + \Pr \left[(\mathcal{A}, \mathbf{X}) \notin \mathcal{G}_{\leq n}(\hat{\delta}) \right] \\ &= \Pr \left[\sum_{i=1}^n T_i(\mathcal{A}, \mathbf{X}_1^i) > n\nu + 6t\epsilon\sqrt{n} \cap (\mathcal{A}, \mathbf{X}) \in \mathcal{G}_{\leq n}(\hat{\delta}) \right] + \Pr \left[(\mathcal{A}, \mathbf{X}) \notin \mathcal{G}_{\leq n}(\hat{\delta}) \right] \end{aligned}$$

We then substitute ν by $\nu(\hat{\delta})$ as defined in Equation (7), and apply a union bound on $\Pr \left[(\mathcal{A}, \mathbf{X}) \notin \mathcal{G}_{\leq n}(\hat{\delta}) \right]$ using Claim 3.6 to get

$$\begin{aligned} \Pr \left[\sum_{i=1}^n Z_i(\mathcal{A}, \mathbf{X}_1^i) > n\nu(\hat{\delta}) + 6t\epsilon\sqrt{n} \right] &\leq \Pr \left[\sum_{i=1}^n T_i(\mathcal{A}, \mathbf{X}_1^i) > n\nu(\hat{\delta}) + 6t\epsilon\sqrt{n} \right] + n(\delta' + \delta'') \\ &\leq e^{-t^2/2} + n(\delta' + \delta'') \end{aligned}$$

where the two inequalities follow from Claim 3.6 and Theorem A.8, respectively. Therefore,

$$\Pr \left[Z(\mathcal{A}(\mathbf{X}), \mathbf{X}) > n\nu(\hat{\delta}) + 6t\epsilon\sqrt{n} \right] \leq e^{-t^2/2} + n(\delta' + \delta'') \stackrel{def}{=} \beta(t, \hat{\delta})$$

From Lemma 3.2, we have $I_\infty^{\beta(t,\delta)}(\mathbf{X}; \mathcal{A}(\mathbf{X})) \leq n\nu(\hat{\delta}) + 6t\epsilon\sqrt{n}$. We set the parameters $t = \epsilon\sqrt{2n}$ and $\hat{\delta} = \sqrt{\epsilon\delta}/15$ to obtain our result. Note that setting $\hat{\delta} = \sqrt{\epsilon\delta}/15$ does not violate the bounds on it stated in the statement of Lemma 3.7. □

Theorem 3.9. *Let $\mathcal{X} = \{0, 1\}$ and $\mathcal{Y} = \{\mathcal{X}^n \cup \{\perp\}\}$. Let \mathbf{X} be a uniformly distributed random variable over \mathcal{X}^n . For $n > 64e$, for every $\epsilon \in (0, \frac{1}{2}]$, $\delta \in (0, \frac{1}{4}]$, there exists an integer $r > 0$ and randomized algorithms $\mathcal{A} : \mathcal{X}^n \rightarrow \{0, 1\}^r$, and $\mathcal{B} : \mathcal{X}^n \times \{0, 1\}^r \rightarrow \mathcal{Y}$, such that:*

1. $r = O\left(\frac{\log(1/\delta)\log n}{\epsilon}\right)$ and $I_\infty^\beta(\mathbf{X}; \mathcal{A}(\mathbf{X})) \leq r + \log(\frac{1}{\beta})$ for all $\beta > 0$;
2. for every $\mathbf{a} \in \{0, 1\}^r$, $\mathcal{B}(\mathbf{X}, \mathbf{a})$ is (ϵ, δ) -differentially private and $I_\infty^\beta(\mathbf{X}; \mathcal{B}(\mathbf{X}, \mathbf{a})) \leq 1$ for all $\beta \geq 2\delta$;
3. for every $\mathbf{x} \in \mathcal{X}^n$, with probability at least $1 - \delta$, we have that $\mathcal{B}(\mathbf{x}; \mathcal{A}(\mathbf{x})) = \mathbf{x}$. In particular, $I_\infty^\beta(\mathbf{X}, \mathcal{B}(\mathbf{X}; \mathcal{A}(\mathbf{X}))) \geq n - 1$ for all $0 < \beta \leq \frac{1}{2} - \delta$.

De [De12] showed that the *mutual information* of (ϵ, δ) -differentially private protocols can be large: if $\frac{1}{\epsilon} \log(\frac{1}{\delta}) = O(n)$, then there exists an (ϵ, δ) -differentially private algorithm \mathcal{B} and a distribution \mathcal{S} such that for $\mathbf{X} \sim \mathcal{S}$, $I(\mathbf{X}; \mathcal{B}(\mathbf{X})) = \Omega(n)$, where I denotes mutual information. De's construction also has large approximate max-information.

By the composition theorem for approximate max-information (given in the appendix in Lemma A.3), our construction implies a similar bound:

Corollary 3.10. *There exists an (ϵ, δ) -differentially private mechanism $\mathcal{C} : \mathcal{X}^n \rightarrow \mathcal{Y}$ such that $I_\infty^{\beta_2}(\mathcal{C}, n) \geq n - 1 - r - \log(1/\beta_1)$ for all $\beta_1 \in (0, 1/2 - \delta)$ and $\beta_2 \in (0, 1/2 - \delta - \beta_1)$, where $r = O\left(\frac{\log(1/\delta)\log(n)}{\epsilon}\right)$.*

We adapt ideas from De's construction in order to prove Theorem 3.9. In De's construction, the input is not drawn from a product distribution—instead, the support of the input distribution is an error-correcting code, meaning that all points in the support are far from each other in Hamming distance. For such a distribution, De showed that adding the level of noise required for differential privacy does not add enough distortion to prevent decoding of the dataset.

Our construction adapts De's idea. Given as input a *uniformly random* dataset \mathbf{x} , we show a mechanism \mathcal{A} which outputs a short description of a code that contains \mathbf{x} . Because this description is short, \mathcal{A} has small max-information. The mechanism \mathcal{B} is then parameterized by this short description of a code. Given the description of a code and the dataset \mathbf{x} , \mathcal{B} approximates (privately) the distance from \mathbf{x} to the nearest *codeword*, and outputs that codeword when the distance is small. When \mathcal{B} is composed with \mathcal{A} , we show that it outputs the dataset \mathbf{x} with high probability.

Before getting into the details of our result, we present some preliminaries for linear codes.

3.1 Preliminaries for Linear Codes

For the current and the next subsections, we limit our scope to \mathbb{F}_2 , i.e., the finite field with 2 elements. First, we define linear codes:

Definition 3.11 (Linear Code). A code $C \subseteq \{0, 1\}^n$ of length n and rank k is called linear iff it is a k dimensional linear subspace of the vector space \mathbb{F}_2^n . The vectors in C are called codewords.

The minimum distance t of a linear code C is $t = \min_{c_1, c_2 \in C} \text{dist}_{\text{Ham}}(c_1, c_2)$, where, $\text{dist}_{\text{Ham}}(p, q)$ denotes the Hamming distance between binary vectors p and q .

We now define parity check matrices, which can be used to construct linear codes. Every linear code has a parity-check matrix corresponding to it. Thus, given a parity-check matrix, one can reconstruct the corresponding linear code.

Definition 3.12 (Parity-check matrix). For a linear code $C \subseteq \{0, 1\}^n$ of length n and rank k , $H \in \{0, 1\}^{(n-k) \times n}$ is a parity-check matrix of C iff H is a matrix whose null space is C , i.e., $c \in C$ iff $Hc = \mathbf{0}$, where $\mathbf{0}$ represents the zero vector.

Now, we state a theorem which shows the existence of high-rank linear codes when the minimum distance is less than half the code length:

Theorem 3.13 (From Theorem 5.1.8 in [Lin99]). For every $t \in (0, \frac{n}{2})$, there exists a linear code of rank k such that $k \geq n - 3t \log(n)$.

Next, we will define an affine code, which is a translation of a linear code by a fixed vector in the vector space of the linear code:

Definition 3.14 (Affine Code). Let $C \subseteq \{0, 1\}^n$ be a linear code of length n , rank k and minimum distance t . For any vector $b \in \{0, 1\}^n$, the code defined by $C_a = \{c + b : c \in C\}$, where $a = Hb$, is called an affine code.

Lemma 3.15. If C is a linear code with parity check matrix H and minimum distance t , then the affine code C_a also has minimum distance t . Further, for all $c' \in C_a$, we have $Hc' = a$.

Proof. Let $c' \in C_a$. We know that there exists a $c \in C$ such that

$$Hc' = H(c + b) = \mathbf{0} + Hb = a.$$

□

Lastly, we define the concept of a Hamming ball around a point, which is helpful in understanding the point's neighborhood – i.e., the points close to it with respect to Hamming distance.

Definition 3.16 (Hamming ball). A Hamming ball of radius r around a point $p \in \{0, 1\}^n$, denoted by $B_r(p)$, is the set of strings $x \in \{0, 1\}^n$ such that $\text{dist}_{\text{Ham}}(x, p) \leq r$.

The volume of a Hamming ball, denoted by $\text{Vol}(B_r)$, is independent of the point around which the ball is centered, i.e., for any point $p \in \{0, 1\}^n$:

$$\text{Vol}(B_r) = |B_r(p)| = \sum_{i=0}^r |\{x \in \{0, 1\}^n : \text{dist}_{\text{Ham}}(x, p) = i\}| = \sum_{i=0}^r \binom{n}{i}. \quad (10)$$

3.2 Proof of Theorem 3.9

In this section, we define the mechanisms \mathcal{A} and \mathcal{B} from the theorem statement, and then prove our result in three parts: First, we show that the first bullet in the theorem statement directly follows from setting the parameters appropriately and from [DFH⁺15a]. Next, we show the proof of the second bullet in two pieces. We start by showing that the algorithm \mathcal{B} that we define is differentially private, and then, we show that the approximate max-information of \mathcal{B} is small when its inputs are chosen independently. Lastly, we prove the third bullet by first showing that the adaptive composition of \mathcal{A} followed by \mathcal{B} results in the reconstruction of the input with high probability. Subsequently, we show that such a composition has large approximate max-information.

Before we define the mechanisms \mathcal{A} and \mathcal{B} , we must set up some notation. We fix t such that $t = \frac{8 \log(1/\delta)}{\epsilon} + 1$. We know that $t \geq 33$ because $\epsilon \in (0, 1/2]$ and $\delta \in (0, 1/4]$. Now, fix an $((n - k) \times n)$ parity-check matrix H for a linear code $C \subseteq \{0, 1\}^n$ of rank k over \mathbb{F}_2 where t is the minimum distance of C and $k = n - 3t \log n$, and let $r = n - k = 3t \log n$. We can ensure the existence of C from Theorem 3.13.

We define the mechanisms \mathcal{A} and \mathcal{B} from the theorem statement in Algorithm 1 and Algorithm 2, respectively.

Brief description of \mathcal{A} : For any input $\mathbf{x} \in \mathcal{X}^n$, mechanism \mathcal{A} returns a vector $\mathbf{a}_\mathbf{x} \in \{0, 1\}^r$ such that $\mathbf{x} \in C_{\mathbf{a}_\mathbf{x}}$, where $C_{\mathbf{a}_\mathbf{x}}$ is an affine code with minimum distance t . This follows as $\mathbf{a}_\mathbf{x} = \mathcal{A}(\mathbf{x}) = H\mathbf{x}$, and from Lemma 3.15, as $C_{\mathbf{a}_\mathbf{x}} = \{\mathbf{c} \in \mathcal{X}^n : H\mathbf{c} = \mathbf{a}_\mathbf{x}\}$.

Input: $\mathbf{x} \in \{0, 1\}^n$

Output: $\mathbf{a}_\mathbf{x} \in \{0, 1\}^r$

- 1 Return $H\mathbf{x}$ (multiplication in \mathbb{F}_2).

Algorithm 1: \mathcal{A}

Brief description of $\mathcal{B}_{\epsilon, \delta}$: For any input $\mathbf{x} \in \mathcal{X}^n$ and $\mathbf{a} \in \{0, 1\}^r$, mechanism $\mathcal{B}_{\epsilon, \delta}$ first computes $d_\mathbf{x}$, which is the distance of \mathbf{x} from $f(\mathbf{x})$, i.e., the nearest codeword to \mathbf{x} in code $C_\mathbf{a}$. Next, it sets $\hat{d}_\mathbf{x}$ to be $d_\mathbf{x}$ perturbed with Laplace noise $L \sim \text{Lap}(1/\epsilon)$. It returns $f(\mathbf{x})$ if $\hat{d}_\mathbf{x}$ is below a threshold $w \stackrel{\text{def}}{=} \left(\frac{t-1}{4} - \frac{\log(1/\delta)}{\epsilon} \right)$, and \perp otherwise.

Input: $\mathbf{x} \in \{0, 1\}^n$ (private) and $\mathbf{a} \in \{0, 1\}^r$ (public)

Output: $\mathbf{b} \in \mathcal{Y}$

- 1 Compute the distance of \mathbf{x} to the nearest codeword in code $C_\mathbf{a}$. Let $d_\mathbf{x} = \min_{\mathbf{c} \in C_\mathbf{a}} (\text{dist}_{\text{Hamm}}(\mathbf{x}, \mathbf{c}))$ and $f(\mathbf{x}) = \arg \min_{\mathbf{c} \in C_\mathbf{a}} (\text{dist}_{\text{Hamm}}(\mathbf{x}, \mathbf{c}))$ (breaking ties arbitrarily).
- 2 Let $\hat{d}_\mathbf{x} = d_\mathbf{x} + L$, where $L \sim \text{Lap}(1/\epsilon)$, and $\text{Lap}(c)$ denotes a random variable having Laplace(0, c) distribution.
- 3 **if** $\hat{d}_\mathbf{x} < \left(\frac{t-1}{4} - \frac{\log(1/\delta)}{\epsilon} \right)$ **then**
- 4 | Return $f(\mathbf{x})$.
- 5 **else**
- 6 | Return \perp .
- 7 **end**

Algorithm 2: $\mathcal{B}_{\epsilon, \delta}$

Now, we present the proof of our theorem.

Proof of Theorem 3.9, part 1. Observe that $r = O\left(\frac{\log(1/\delta) \log n}{\epsilon}\right)$ from the value assigned to t .

We know that the second statement holds by the max-information bound for mechanisms with bounded description length from [DFH⁺15a]. \square

Proof of Theorem 3.9, part 2. First, we show that $\mathcal{B}_{\epsilon, \delta}$ is indeed differentially private.

Lemma 3.17. $\mathcal{B}_{\epsilon, \delta}(\cdot, \mathbf{a})$ is (ϵ, δ) -differentially private for every $\mathbf{a} \in \{0, 1\}^r$.

Proof. We will prove this lemma by following the proof of Proposition 3 in [ST13]. Fix any $\mathbf{a} \in \{0, 1\}^r$. Firstly, observe that for every $\mathbf{x} \in \{0, 1\}^n$, there are only 2 possible outputs for $\mathcal{B}_{\epsilon, \delta}(\mathbf{x}, \mathbf{a})$: \perp or $f(\mathbf{x}) = \arg \min_{\mathbf{c} \in C_{\mathbf{a}}}(\text{dist}_{\text{Hammm}}(\mathbf{x}, \mathbf{c}))$. Also, $\mathcal{B}_{\epsilon, \delta}(\mathbf{x}, \mathbf{a}) = f(\mathbf{x})$ iff $\hat{d}_{\mathbf{x}} = d_{\mathbf{x}} + L < w$ in Algorithm 2, where $d_{\mathbf{x}} = \min_{\mathbf{c} \in C_{\mathbf{a}}}(\text{dist}_{\text{Hammm}}(\mathbf{x}, \mathbf{c}))$ and $L \sim \text{Lap}(1/\epsilon)$.

Now, for any pair of points \mathbf{x} and \mathbf{x}' such that $\text{dist}_{\text{Hammm}}(\mathbf{x}, \mathbf{x}') = 1$, there are two possible cases:

1. $f(\mathbf{x}) \neq f(\mathbf{x}')$:

In this case,

$$\begin{aligned} p &\stackrel{\text{def}}{=} \Pr[\mathcal{B}_{\epsilon, \delta}(\mathbf{x}, \mathbf{a}) = f(\mathbf{x})] = \Pr[\hat{d}_{\mathbf{x}} < w] = \Pr\left[d_{\mathbf{x}} + L < \frac{t-1}{4} - \frac{\log(1/\delta)}{\epsilon}\right] \\ &\leq \Pr\left[\frac{t-1}{2} + L < \frac{t-1}{4} - \frac{\log(1/\delta)}{\epsilon}\right] \leq \Pr\left[L < -\frac{\log(1/\delta)}{\epsilon}\right] \leq \delta \end{aligned}$$

where the first inequality follows as $f(\mathbf{x}) \neq f(\mathbf{x}')$ implies $d_{\mathbf{x}} > \frac{t-1}{2}$, and the last inequality follows from the tail property of the Laplace distribution. Therefore, $\Pr[\mathcal{B}_{\epsilon, \delta}(\mathbf{x}, \mathbf{a}) = \perp] = 1 - p$.

Similarly, $p' \stackrel{\text{def}}{=} \Pr[\mathcal{B}_{\epsilon, \delta}(\mathbf{x}', \mathbf{a}) = f(\mathbf{x}')] \leq \delta$, and consequently, $\Pr[\mathcal{B}_{\epsilon, \delta}(\mathbf{x}', \mathbf{a}) = \perp] = 1 - p'$.

Thus, for any set $\mathcal{O} \subseteq \mathcal{Y}$, we can bound the following difference in terms of the total variation distance $TV(\mathcal{B}_{\epsilon, \delta}(\mathbf{x}, \mathbf{a}), \mathcal{B}_{\epsilon, \delta}(\mathbf{x}', \mathbf{a}))$ (defined in the appendix)

$$\begin{aligned} |\Pr[\mathcal{B}_{\epsilon, \delta}(\mathbf{x}, \mathbf{a}) \in \mathcal{O}] - \Pr[\mathcal{B}_{\epsilon, \delta}(\mathbf{x}', \mathbf{a}) \in \mathcal{O}]| &\leq TV(\mathcal{B}_{\epsilon, \delta}(\mathbf{x}, \mathbf{a}), \mathcal{B}_{\epsilon, \delta}(\mathbf{x}', \mathbf{a})) \\ &= \frac{(p-0) + (p'-0) + |(1-p) - (1-p')|}{2} \\ &= \frac{p + p' + |p' - p|}{2} = \max\{p, p'\} \leq \delta \end{aligned}$$

2. $f(\mathbf{x}) = f(\mathbf{x}')$:

Observe that for every $\mathbf{x}'' \in \{0, 1\}^n$, the value of $d_{\mathbf{x}''}$ can change by at most 1 if exactly one coordinate is changed in \mathbf{x}'' . Computing $\hat{d}_{\mathbf{x}''}$ is then just an instantiation of the Laplace mechanism, given in the appendix (Theorem A.1). Therefore, $\hat{d}_{\mathbf{x}''}$ satisfies $(\epsilon, 0)$ -differential privacy. Notice that determining whether to output $f(\mathbf{x}) = f(\mathbf{x}')$ or \perp is a post-processing function of the $(\epsilon, 0)$ -differentially private $\hat{d}_{\mathbf{x}}$, and thus, by Theorem A.2, $\mathcal{B}_{\epsilon, \delta}(\cdot, \mathbf{a})$ is $(\epsilon, 0)$ -differentially private for such inputs.

Therefore, from the above two cases, for any set $\mathcal{O} \subseteq \mathcal{Y}$, we have that:

$$\Pr [\mathcal{B}_{\epsilon, \delta}(\mathbf{x}, \mathbf{a}) \in \mathcal{O}] \leq e^\epsilon \Pr [\mathcal{B}_{\epsilon, \delta}(\mathbf{x}', \mathbf{a}) \in \mathcal{O}] + \delta.$$

Thus, we can conclude that $\mathcal{B}_{\epsilon, \delta}(\cdot, \mathbf{a})$ is (ϵ, δ) -differentially private for every $\mathbf{a} \in \{0, 1\}^r$. \square

Next, we look at the outcome of $\mathcal{B}_{\epsilon, \delta}(\mathbf{X}, \mathbf{a})$ when \mathbf{X} is drawn uniformly over \mathcal{X}^n and \mathbf{a} is a fixed r -bit string. Note that $\mathcal{B}_{\epsilon, \delta}(\mathbf{X}, \mathbf{a})$ outputs either \perp or a codeword of $C_{\mathbf{a}}$. Thus,

$$\Pr [\mathcal{B}_{\epsilon, \delta}(\mathbf{X}, \mathbf{a}) \neq \perp] = \Pr [\hat{d}_{\mathbf{X}} < w] = \Pr \left[d_{\mathbf{X}} + L < \left(\frac{t-1}{4} - \frac{\log(1/\delta)}{\epsilon} \right) \right] \quad (11)$$

Now, let us define the set $\mathcal{D} = \left\{ \mathbf{x} \in \mathcal{X}^n : d_{\mathbf{x}} < \left(\frac{t-1}{4} \right) \right\}$. If $\left(d_{\mathbf{X}} + L < \left(\frac{t-1}{4} - \frac{\log(1/\delta)}{\epsilon} \right) \right)$, then either $\mathbf{X} \in \mathcal{D}$, or $L < -\frac{\log(1/\delta)}{\epsilon}$, or both. Thus,

$$\Pr [\mathcal{B}_{\epsilon, \delta}(\mathbf{X}, \mathbf{a}) \neq \perp] \leq \Pr [\mathbf{X} \in \mathcal{D}] + \Pr \left[L < -\frac{\log(1/\delta)}{\epsilon} \right] \quad (12)$$

From the tail bound of the Laplace distribution,

$$\Pr \left[L < -\frac{\log(1/\delta)}{\epsilon} \right] \leq \delta \quad (13)$$

Next, we will calculate the probability of $\mathbf{X} \in \mathcal{D}$. We then assign $s \stackrel{\text{def}}{=} \frac{t-1}{4}$. Notice that as the minimum distance of $C_{\mathbf{a}}$ is t , the Hamming balls B_{2s} of radius $2s$ around the codewords of $C_{\mathbf{a}}$ are disjoint and thus we can bound the volume (defined in (10)) of each,

$$|C_{\mathbf{a}}| \cdot \text{Vol}(B_{2s}) \leq 2^n \quad (14)$$

Therefore,

$$|C_{\mathbf{a}}| \cdot \text{Vol}(B_s) \leq \frac{2^n \cdot \text{Vol}(B_s)}{\text{Vol}(B_{2s})} = 2^n \cdot \frac{\sum_{i=0}^s \binom{n}{i}}{\sum_{j=0}^{2s} \binom{n}{j}} \leq 2^n \cdot \frac{s \cdot \binom{n}{s}}{\binom{n}{2s}} \leq 2^n \cdot \frac{s \cdot \left(\frac{ne}{s} \right)^s}{\left(\frac{n}{2s} \right)^{2s}} = 2^n s \left(\frac{4es}{n} \right)^s \quad (15)$$

where the first inequality follows from equation (14), and the last inequality follows as $\left(\frac{n}{k} \right)^k \leq \binom{n}{k} \leq \left(\frac{ne}{k} \right)^k$ for $k \geq 1$ (from Appendix C.1 in [CLRS09]).

Thus,

$$\Pr [\mathbf{X} \in \mathcal{D}] = \frac{|C_{\mathbf{a}}| \cdot \text{Vol}(B_s)}{2^n} \leq s \left(\frac{4es}{n} \right)^s \quad (16)$$

where the inequality follows from equation (15).

Hence,

$$\Pr [\mathcal{B}_{\epsilon,\delta}(\mathbf{X}, \mathbf{a}) \neq \perp] \leq s \left(\frac{4es}{n} \right)^s + \delta < s \cdot 2^{-s} + \delta \quad (17)$$

where the first inequality follows from equations (12),(13) and (16), and the last inequality follows from the fact that $n > 8es = 2e(t-1)$.

Bounding the term $s \cdot 2^{-s}$ from above, we have

$$\begin{aligned} s \cdot 2^{-s} &= \frac{t-1}{4} \cdot 2^{(1-t)/4} = \frac{2 \log(1/\delta)}{\epsilon} \cdot 2^{-2 \log(1/\delta)/\epsilon} = \frac{2 \log(1/\delta)}{\epsilon} \cdot \delta^{2/\epsilon} \\ &= (\delta \log(1/\delta)) \left(\frac{2}{\epsilon} \cdot \delta^{(2/\epsilon)-2} \right) \delta \leq \delta \end{aligned} \quad (18)$$

where the inequality follows as $\delta \log(1/\delta) \leq 1$ for $\delta \in (0, \frac{1}{4}]$, and $\frac{2}{\epsilon} \cdot \delta^{(2/\epsilon)-2} \leq 1$ for $\epsilon \in (0, \frac{1}{2}]$, $\delta \in (0, \frac{1}{4}]$. From equations (17) and (18),

$$\Pr [\mathcal{B}_{\epsilon,\delta}(\mathbf{X}, \mathbf{a}) = \perp] > 1 - 2\delta \quad (19)$$

Now, for any $\mathbf{x} \in \mathcal{X}^n$,

$$\begin{aligned} \log \left(\frac{\Pr [(\mathbf{X}, \mathcal{B}_{\epsilon,\delta}(\mathbf{X}, \mathbf{a})) = (\mathbf{x}, \perp)]}{\Pr [\mathbf{X} \otimes \mathcal{B}_{\epsilon,\delta}(\mathbf{X}, \mathbf{a}) = (\mathbf{x}, \perp)]} \right) &= \log \left(\frac{\Pr [\mathcal{B}_{\epsilon,\delta}(\mathbf{X}, \mathbf{a}) = \perp | \mathbf{X} = \mathbf{x}]}{\Pr [\mathcal{B}_{\epsilon,\delta}(\mathbf{X}, \mathbf{a}) = \perp]} \right) < \log \left(\frac{1}{1-2\delta} \right) \\ &\leq \log \left(\frac{1}{1-0.5} \right) = 1 \end{aligned} \quad (20)$$

where the first inequality follows from equation (19), and the second inequality follows from the fact that $\delta \leq \frac{1}{4}$.

We then apply Lemma 3.2 using (19) and (20) to get,

$$I_{\infty}^{\beta}(\mathbf{X}; \mathcal{B}_{\epsilon,\delta}(\mathbf{X}, \mathbf{a})) \leq 1, \text{ for } \beta \geq 2\delta.$$

□

Proof of Theorem 3.9, part 3. Let us look at the outcome of $\mathcal{B}_{\epsilon,\delta}(\mathbf{x}, \mathcal{A}(\mathbf{x}))$. First, as $\mathbf{x} \in C_{\mathcal{A}(\mathbf{x})}$, $f(\mathbf{x}) = \mathbf{x}$ and $d_{\mathbf{x}} = 0$. Thus, $\mathcal{B}_{\epsilon,\delta}(\mathbf{x}, \mathcal{A}(\mathbf{x}))$ will either return \mathbf{x} or \perp . Furthermore, we can show the probability of outputting \mathbf{x} is high:

$$\Pr_{\substack{\text{coins} \\ \text{of } \mathcal{B}_{\epsilon,\delta}}} [\mathcal{B}_{\epsilon,\delta}(\mathbf{x}, \mathcal{A}(\mathbf{x})) = \mathbf{x}] = \Pr [\hat{d}_{\mathbf{x}} < w] \geq \Pr \left[\hat{d}_{\mathbf{x}} < \frac{\log(1/\delta)}{\epsilon} \right] = \Pr \left[\text{Lap}(1/\epsilon) < \frac{\log(1/\delta)}{\epsilon} \right] \geq 1 - \delta$$

where the first inequality follows from the fact that $\left(\frac{t-1}{4} - \frac{\log(1/\delta)}{\epsilon} \right) \geq \frac{\log(1/\delta)}{\epsilon}$, the equality after it follows since $d_{\mathbf{x}} = 0$, and the last inequality follows from a tail bound of the Laplace distribution. Thus, for every $\mathbf{x} \in \mathcal{X}^n$,

$$\Pr [\mathcal{B}_{\epsilon,\delta}(\mathbf{x}, \mathcal{A}(\mathbf{x})) = \mathbf{x}] \geq 1 - \delta. \quad (21)$$

Consider the event $\mathcal{D}_{\mathcal{X}} = \{(\mathbf{x}, \mathbf{x}) : \mathbf{x} \in \mathcal{X}^n\}$. From equation (21),

$$\Pr [(\mathbf{X}, \mathcal{B}_{\epsilon, \delta}(\mathbf{X}, \mathcal{A}(\mathbf{X}))) \in \mathcal{D}_{\mathcal{X}}] \geq 1 - \delta. \quad (22)$$

Also, for $\mathbf{b} \in \mathcal{Y}$, if $\mathbf{b} = \perp$, then $\Pr [\mathbf{X} = \mathbf{b}] = 0$, and if $\mathbf{b} \in \mathcal{X}^n$, then $\Pr [\mathbf{X} = \mathbf{b}] = 2^{-n}$ as \mathbf{X} is drawn uniformly over \mathcal{X}^n . Thus, for all $\mathbf{b} \in \mathcal{Y}$,

$$\Pr [(\mathbf{X}, \mathbf{b}) \in \mathcal{D}_{\mathcal{X}}] \leq 2^{-n}.$$

Hence,

$$\Pr [(\mathbf{X} \otimes \mathcal{B}_{\epsilon, \delta}(\mathbf{X}, \mathcal{A}(\mathbf{X}))) \in \mathcal{D}_{\mathcal{X}}] = \sum_{\mathbf{b} \in \mathcal{Y}} \Pr [(\mathbf{X}, \mathbf{b}) \in \mathcal{D}_{\mathcal{X}}] \Pr [\mathcal{B}_{\epsilon, \delta}(\mathbf{X}, \mathcal{A}(\mathbf{X})) = \mathbf{b}] \leq 2^{-n} \quad (23)$$

Therefore, for $\beta \leq \frac{1}{2} - \delta$,

$$\begin{aligned} I_{\infty}^{\beta}(\mathbf{X}; \mathcal{B}_{\epsilon, \delta}(\mathbf{X}, \mathcal{A}(\mathbf{X}))) &= \log \left(\max_{\substack{\mathcal{O} \subseteq (\mathbf{X} \times \mathcal{Y}), \\ \Pr [(\mathbf{X}, \mathcal{B}_{\epsilon, \delta}(\mathbf{X}, \mathcal{A}(\mathbf{X}))) \in \mathcal{O}] > \beta}} \frac{\Pr [(\mathbf{X}, \mathcal{B}_{\epsilon, \delta}(\mathbf{X}, \mathcal{A}(\mathbf{X}))) \in \mathcal{O}] - \beta}{\Pr [(\mathbf{X} \otimes \mathcal{B}_{\epsilon, \delta}(\mathbf{X}, \mathcal{A}(\mathbf{X}))) \in \mathcal{O}]} \right) \\ &\geq \log \left(\frac{\Pr [(\mathbf{X}, \mathcal{B}_{\epsilon, \delta}(\mathbf{X}, \mathcal{A}(\mathbf{X}))) \in \mathcal{D}_{\mathcal{X}}] - \beta}{\Pr [(\mathbf{X} \otimes \mathcal{B}_{\epsilon, \delta}(\mathbf{X}, \mathcal{A}(\mathbf{X}))) \in \mathcal{D}_{\mathcal{X}}]} \right) \geq \log \left(\frac{1 - \delta - \beta}{2^{-n}} \right) \\ &= n + \log(1 - \delta - \beta) \geq n - 1 \end{aligned}$$

where the first inequality follows from equation (22) and as $(1 - \delta) > \beta$, the second inequality follows from equations (22) and (23), and the last inequality follows from the fact that $\beta \leq \frac{1}{2} - \delta$. \square

Acknowledgments

We thank Salil Vadhan for pointing out that our max-information bound can be used to bound mutual information, thus improving on a result in [MMP⁺11].

References

- [BBB⁺13] Richard Berk, Lawrence Brown, Andreas Buja, Kai Zhang, and Linda Zhao. Valid post-selection inference. *The Annals of Statistics*, 41(2):802–837, 2013.
- [BNS⁺16] Raef Bassily, Kobbi Nissim, Adam D. Smith, Thomas Steinke, Uri Stemmer, and Jonathan Ullman. Algorithmic stability for adaptive data analysis. In *Proceedings of the 48th Annual ACM on Symposium on Theory of Computing, STOC*, 2016.
- [CLN⁺16] Rachel Cummings, Katrina Ligett, Kobbi Nissim, Aaron Roth, and Zhiwei Steven Wu. Adaptive learning with robust generalization guarantees. *arXiv preprint arXiv:1602.07726*, 2016.
- [CLRS09] Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. *Introduction to Algorithms, Third Edition*. The MIT Press, 3rd edition, 2009.

- [De12] Anindya De. Lower bounds in differential privacy. In *Proceedings of the 9th International Conference on Theory of Cryptography*, TCC'12, pages 321–338, Berlin, Heidelberg, 2012. Springer-Verlag.
- [DFH⁺15a] Cynthia Dwork, Vitaly Feldman, Moritz Hardt, Toni Pitassi, Omer Reingold, and Aaron Roth. Generalization in adaptive data analysis and holdout reuse. In C. Cortes, N.D. Lawrence, D.D. Lee, M. Sugiyama, R. Garnett, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 2341–2349. Curran Associates, Inc., 2015.
- [DFH⁺15b] Cynthia Dwork, Vitaly Feldman, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Aaron Leon Roth. Preserving statistical validity in adaptive data analysis. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing*, STOC '15, pages 117–126, New York, NY, USA, 2015. ACM.
- [DKM⁺06] Cynthia Dwork, Krishnam Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our data, ourselves: Privacy via distributed noise generation. In *Advances in Cryptology - EUROCRYPT 2006, 25th Annual International Conference on the Theory and Applications of Cryptographic Techniques, St. Petersburg, Russia, May 28 - June 1, 2006, Proceedings*, pages 486–503, 2006.
- [DMNS06] Cynthia Dwork, Frank Mcsherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *In Proceedings of the 3rd Theory of Cryptography Conference*, pages 265–284. Springer, 2006.
- [DR14] Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3-4):211–407, 2014.
- [DRV10] Cynthia Dwork, Guy N. Rothblum, and Salil P. Vadhan. Boosting and differential privacy. In *51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010, October 23-26, 2010, Las Vegas, Nevada, USA*, pages 51–60, 2010.
- [DSZ15] Cynthia Dwork, Weijie Su, and Li Zhang. Private false discovery rate control. *arXiv preprint arXiv:1511.03803*, 2015.
- [FST14] William Fithian, Dennis Sun, and Jonathan Taylor. Optimal inference after model selection. *arXiv preprint arXiv:1410.2597*, 2014.
- [GL14] Andrew Gelman and Eric Loken. The statistical crisis in science. *American Scientist*, 102(6):460, 2014.
- [GLRV16] Marco Gaboardi, Hyun Lim, Ryan Rogers, and Salil Vadhan. Differentially private chi-squared hypothesis testing: Goodness of fit and independence testing. *arXiv preprint arXiv:1602.03090*, 2016.
- [HU14] Moritz Hardt and Jonathan Ullman. Preventing false discovery in interactive data analysis is hard. In *Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on*, pages 454–463. IEEE, 2014.
- [JS13] Aaron Johnson and Vitaly Shmatikov. Privacy-preserving data exploration in genome-wide association studies. In *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '13, pages 1079–1087, New York, NY, USA, 2013. ACM.

- [KS14] S.P. Kasiviswanathan and A. Smith. On the ‘Semantics’ of Differential Privacy: A Bayesian Formulation. *Journal of Privacy and Confidentiality*, Vol. 6: Iss. 1, Article 1, 2014. Available at <http://repository.cmu.edu/jpc/vol6/iss1/1>. The theorem numbers and exact statements refer to the Arxiv version (v3).
- [KS16] Vishesh Karwa and Aleksandra Slavković. Inference using noisy degrees: Differentially private beta-model and synthetic graphs. *The Annals of Statistics*, 44(1):87–112, 2016.
- [Lin99] Jacobus Hendricus van Lint. *Introduction to Coding Theory*. Springer, Berlin, 3rd edition, 1999.
- [LSST13] Jason D Lee, Dennis L Sun, Yuekai Sun, and Jonathan E Taylor. Exact post-selection inference, with application to the lasso. *arXiv preprint arXiv:1311.6238*, 2013.
- [MMP⁺11] Andrew McGregor, Ilya Mironov, Toniann Pitassi, Omer Reingold, Kunal Talwar, and Salil P. Vadhan. The limits of two-party differential privacy. *Electronic Colloquium on Computational Complexity (ECCC)*, 18:106, 2011.
- [RR10] Aaron Roth and Tim Roughgarden. Interactive privacy via the median mechanism. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 765–774. ACM, 2010.
- [RRST16] Ryan Rogers, Aaron Roth, Adam Smith, and Om Thakkar. Max-information, differential privacy, and post-selection hypothesis testing. In *Foundations of Computer Science (FOCS)*, 2016. arXiv:1604.03924 [cs.LG].
- [RZ16] Daniel Russo and James Zou. Controlling bias in adaptive data analysis using information theory. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics, AISTATS*, 2016.
- [She15] Or Sheffet. Differentially private least squares: Estimation, confidence and rejecting the null hypothesis. *arXiv preprint arXiv:1507.02482*, 2015.
- [SNS11] J. P. Simmons, L. D. Nelson, and U. Simonsohn. False-Positive Psychology: Undisclosed Flexibility in Data Collection and Analysis Allows Presenting Anything as Significant. *Psychological Science*, October 2011.
- [ST13] Adam Smith and Abhradeep Thakurta. Differentially private feature selection via stability arguments, and the robustness of the lasso. In *COLT 2013 - The 26th Annual Conference on Learning Theory, June 12-14, 2013, Princeton University, NJ, USA*, pages 819–850, 2013.
- [SU15] Thomas Steinke and Jonathan Ullman. Interactive fingerprinting codes and the hardness of preventing false discovery. In *Proceedings of The 28th Conference on Learning Theory*, pages 1588–1628, 2015.
- [USF13] Caroline Uhler, Aleksandra Slavkovic, and Stephen E. Fienberg. Privacy-preserving data sharing for genome-wide association studies. *Journal of Privacy and Confidentiality*, 5(1), 2013.
- [WL16] Ronald L. Wasserstein and Nicole A. Lazar. The asa’s statement on p-values: context, process, and purpose. *The American Statistician*, 0(ja):00–00, 2016.
- [WLF16] Yu-Xiang Wang, Jing Lei, and Stephen E. Fienberg. A minimax theory for adaptive data analysis. *CoRR*, abs/1602.04287, 2016.
- [WLK15] Yue Wang, Jaewoo Lee, and Daniel Kifer. Differentially private hypothesis testing, revisited. *arXiv preprint arXiv:1511.03376*, 2015.

[YFSU14] Fei Yu, Stephen E. Fienberg, Aleksandra B. Slavkovic, and Caroline Uhler. Scalable privacy-preserving data sharing methodology for genome-wide association studies. *Journal of Biomedical Informatics*, 50:133–141, 2014.

A Other Useful Probabilistic Tools

We use this section to give an overview of some useful probabilistic tools and their connections with one another. We start by giving a commonly used differentially private mechanism, called the *Laplace Mechanism*, which releases an answer to a query on the dataset with appropriately scaled Laplace noise. We use this mechanism in the proof of Theorem 3.9.

Theorem A.1 (Laplace Mechanism [DMNS06]). *Let $\phi : \mathcal{X}^n \rightarrow \mathbb{R}$ be a function such that for any pair of points \mathbf{x} and \mathbf{x}' where $\text{dist}_{\text{Ham}}(\mathbf{x}, \mathbf{x}') = 1$, then we have $|\phi(\mathbf{x}) - \phi(\mathbf{x}')| \leq 1$. The mechanism $\mathcal{M}(\mathbf{x}) = \phi(\mathbf{x}) + L$ where $L \sim \text{Lap}(1/\epsilon)$, is $(\epsilon, 0)$ -differentially private.*

A very useful fact about differentially private mechanisms is that one cannot take the output of a differentially private mechanism and perform any modification to it that does not depend on the input itself and make the output any less private.

Theorem A.2 (Post Processing [DMNS06]). *Let $\mathcal{M} : \mathcal{X}^n \rightarrow \mathcal{Y}$ be (ϵ, δ) -differentially private and $\psi : \mathcal{Y} \rightarrow \mathcal{Y}'$ be any function mapping to arbitrary domain \mathcal{Y}' . Then $\psi \circ \mathcal{M}$ is (ϵ, δ) -differentially private.*

We have been focusing on approximate max-information throughout this paper, which has the following strong composition guarantee:

Theorem A.3 (Composition [DFH⁺15a]). *Let $\mathcal{A}_1 : \mathcal{X}^n \rightarrow \mathcal{Y}$ and $\mathcal{A}_2 : \mathcal{X}^n \times \mathcal{Y} \rightarrow \mathcal{Z}$ be such that $I_\infty^{\beta_1}(\mathcal{A}_1, n) \leq k_1$ and $I_\infty^{\beta_2}(\mathcal{A}_2(\cdot, y), n) \leq k_2$ for every $y \in \mathcal{Y}$. Then the composition of \mathcal{A}_1 and \mathcal{A}_2 , defined to be $\mathcal{A}(\mathbf{X}) = \mathcal{A}_2(\mathbf{X}, \mathcal{A}_1(\mathbf{X}))$ satisfies:*

$$I_\infty^{\beta_1 + \beta_2}(\mathcal{A}, n) \leq k_1 + k_2.$$

In our analysis, in addition to max-information, we will use more familiar measures between two random variables, which we give here:

Definition A.4 (KL Divergence). *The KL Divergence between random variables X and Z , denoted as $D_{KL}(X||Z)$ over domain \mathcal{D} is defined as*

$$D_{KL}(X||Z) = \sum_{x \in \mathcal{D}} \Pr[X = x] \ln \left(\frac{\Pr[X = x]}{\Pr[Z = x]} \right)$$

Definition A.5 (Total Variation Distance). *The total variation distance between two random variables X and Z , denoted as $TV(X; Z)$, over domain \mathcal{D} is defined as*

$$TV(X, Z) = \frac{1}{2} \cdot \sum_{x \in \mathcal{D}} |\Pr[X = x] - \Pr[Z = x]|.$$

In the following lemma, we state some basic connections between max-information, total variation distance, differential privacy, and indistinguishability:

Lemma A.6. *Let X, Z be two random variables over the same domain. We then have:*

1. [DFH⁺15a] $I_\infty^\beta(X; Z) \leq k \Leftrightarrow (X, Z) \approx_{(k \ln 2), \beta} X \otimes Z$.
2. [KS14] If $X \approx_{\epsilon, \delta} Y$ then X and Y are pointwise $\left(2\epsilon, \frac{2\delta}{1-e^{-\epsilon}}\right)$ -indistinguishable.

Another useful result is from [KS14], which we use in the proof of our main result in Theorem 3.1:

Lemma A.7 (Conditioning Lemma). *Suppose that $(X, Z) \approx_{\epsilon, \delta} (X', Z')$. Then for every $\hat{\delta} > 0$, the following holds:*

$$\Pr_{t \sim p(Z)} \left[X|_{Z=t} \approx_{3\epsilon, \hat{\delta}} X'|_{Z'=t} \right] \geq 1 - \frac{2\delta}{\hat{\delta}} - \frac{2\delta}{1-e^{-\epsilon}}.$$

The proof of our main result in Theorem 3.1 also makes use of the following standard concentration inequality:

Theorem A.8 (Azuma's Inequality). *Let C_1, \dots, C_n be a sequence of random variables such that for every $i \in [n]$, we have*

$$\Pr [|C_i| \leq \alpha] = 1$$

and for every fixed prefix $\mathbf{C}_1^{i-1} = \mathbf{c}_1^{i-1}$, we have

$$\mathbb{E} [C_i | \mathbf{c}_1^{i-1}] \leq \gamma,$$

then for all $t \geq 0$, we have

$$\Pr \left[\sum_{i=1}^n C_i > n\gamma + t\sqrt{n\alpha} \right] \leq e^{-t^2/2}.$$

We say that a real-valued function $f : \mathcal{X}^n \rightarrow \mathbb{R}$ has sensitivity Δ if for all $i \in [n]$ and $\mathbf{x} \in \mathcal{X}^n$ and $x'_i \in \mathcal{X}$, $|f(\mathbf{x}_{-i}, x_i) - f(\mathbf{x}_{-i}, x'_i)| \leq \Delta$. One important use of max-information bounds is that they imply strong generalization bounds for low sensitivity functions when paired with McDiarmid's inequality.

Theorem A.9 (McDiarmid's Inequality). *Let X_1, \dots, X_n be independent random variables with domain \mathcal{X} . Further, let $f : \mathcal{X}^n \rightarrow \mathbb{R}$ be a function of sensitivity $\Delta > 0$. Then for every $\tau > 0$ and $\mu = \mathbb{E} [f(X_1, \dots, X_n)]$ we have*

$$\Pr [f(X_1, \dots, X_n) - \mu \geq \tau] \leq \exp \left(\frac{-2\tau^2}{n\Delta^2} \right).$$

B Sensitivity of p -values

In this section, we demonstrate that the theorem from [BNS⁺16], which shows that differentially private algorithms which select low-sensitivity queries cannot overfit, does not give nontrivial generalization bounds for p -values. We first show that p -values cannot have sensitivity smaller than $.37/\sqrt{n}$.

Lemma B.1. *Let $\phi : \mathcal{X}^n \rightarrow \mathbb{R}$ be a test statistic with null hypothesis H_0 , and $p : \mathbb{R} \rightarrow [0, 1]$, where $p(a) = \Pr_{\mathbf{x} \sim \mathcal{P}^n} [\phi(\mathbf{x}) \geq a]$, and $\mathcal{P} \in H_0$. The sensitivity of $p \circ \phi$ must be larger than $0.37/\sqrt{n}$.*

Proof. Note that if $\mathbf{X} \sim \mathcal{P}^n$, then $p \circ \phi(\mathbf{X})$ is uniform on $[0, 1]$, and thus, has mean $1/2$. From Theorem A.9, we know that if $p \circ \phi$ has sensitivity Δ , then for any $0 < \delta < 1/2$, we have:

$$\Pr \left[p \circ \phi(\mathbf{X}) \geq 1/2 + \Delta \sqrt{\frac{n}{2} \ln(1/\delta)} \right] \leq \delta.$$

However, we also know that $p \circ \phi(\mathbf{X})$ is uniform, so that

$$\Pr [p \circ \phi(\mathbf{X}) \geq 1 - \delta] = \delta.$$

Hence, if $\Delta < \frac{1/2 - \delta}{\sqrt{\frac{n}{2} \ln(1/\delta)}}$, we obtain a contradiction:

$$\delta \geq \Pr \left[p \circ \phi(\mathbf{X}) \geq 1/2 + \Delta \sqrt{\frac{n}{2} \ln(1/\delta)} \right] > \Pr [p \circ \phi(\mathbf{X}) \geq 1 - \delta] = \delta.$$

We then set $\delta = 0.08$ to get our stated bound on sensitivity. □

Thus, the sensitivity Δ for the p -value for any test statistic and any null hypothesis must be at least $0.37/\sqrt{n}$. This is too large for the following theorem, proven in [BNS⁺16], to give a nontrivial guarantee:

Theorem B.2 ([BNS⁺16]). *Let $\epsilon \in (0, 1/3)$, $\delta \in (0, \epsilon/4)$, and $n \geq \frac{\ln(4\epsilon/\delta)}{\epsilon^2}$. Let \mathcal{Y} denote the class of Δ -sensitive functions $f : \mathcal{X}^n \rightarrow \mathbb{R}$, and let $\mathcal{A} : \mathcal{X}^n \rightarrow \mathcal{Y}$ be an algorithm that is (ϵ, δ) -differentially private. Let $\mathbf{X} \sim \mathcal{P}^n$ for some distribution \mathcal{P} over \mathcal{X} , and let $q = \mathcal{A}(\mathbf{X})$. Then:*

$$\Pr_{\mathbf{X}, \mathcal{A}} [|q(\mathcal{P}^n) - q(\mathbf{X})| \geq 18\epsilon\Delta n] < \frac{\delta}{\epsilon}.$$

When we attempt to apply this theorem to a p -value, we see that the error it guarantees, by Lemma B.1, is at least: $18\epsilon\Delta n \geq 18\epsilon(.37)\sqrt{n}$. However, the theorem is only valid for $n \geq \frac{1}{\epsilon^2} \ln\left(\frac{4\epsilon}{\delta}\right)$. Plugging this in, we see that $18\epsilon(.37)\sqrt{n} \geq 1$, which is a trivial error guarantee for p -values (which take values in $[0, 1]$).

C General Conversion between Max-Information and Mutual Information, and its Consequences

We give here a proof of Lemma 1.5.

Proof of Lemma 1.5. We first prove the first direction of the lemma. We define a set $G(k) = \{(\mathbf{x}, \mathbf{y}) : Z(\mathbf{x}, \mathbf{y}) \leq k\}$ where $Z(\mathbf{x}, \mathbf{y}) = \log \left(\frac{\Pr[(\mathbf{X}, \mathbf{Y}) = (\mathbf{x}, \mathbf{y})]}{\Pr[\mathbf{X} \otimes \mathbf{Y} = (\mathbf{x}, \mathbf{y})]} \right)$, and $\beta(k)$ to be the quantity such that $\Pr[(\mathbf{X}, \mathbf{Y}) \in G(k)] = 1 - \beta(k)$. We then have:

$$\begin{aligned} m &\geq \sum_{(\mathbf{x}, \mathbf{y}) \in G(k)} \Pr[(\mathbf{X}, \mathbf{Y}) = (\mathbf{x}, \mathbf{y})] Z(\mathbf{x}, \mathbf{y}) + \sum_{(\mathbf{x}, \mathbf{y}) \notin G(k)} \Pr[(\mathbf{X}, \mathbf{Y}) = (\mathbf{x}, \mathbf{y})] Z(\mathbf{x}, \mathbf{y}) \\ &\geq \sum_{(\mathbf{x}, \mathbf{y}) \in G(k)} \Pr[(\mathbf{X}, \mathbf{Y}) = (\mathbf{x}, \mathbf{y})] Z(\mathbf{x}, \mathbf{y}) + k\beta(k) \end{aligned} \quad (24)$$

Also, we have:

$$- \sum_{(\mathbf{x}, \mathbf{y}) \in G(k)} \frac{\Pr[(\mathbf{X}, \mathbf{Y}) = (\mathbf{x}, \mathbf{y})]}{\Pr[(\mathbf{X}, \mathbf{Y}) \in G(k)]} Z(\mathbf{x}, \mathbf{y}) \leq \log \left(\frac{\Pr[\mathbf{X} \otimes \mathbf{Y} \in G(k)]}{\Pr[(\mathbf{X}, \mathbf{Y}) \in G(k)]} \right) \leq \log \left(\frac{1}{1 - \beta(k)} \right)$$

where the first inequality follows from applying Jensen's inequality.

By rearranging terms, we obtain:

$$\sum_{(\mathbf{x}, \mathbf{y}) \in G(k)} \Pr[(\mathbf{X}, \mathbf{Y}) = (\mathbf{x}, \mathbf{y})] \cdot Z(\mathbf{x}, \mathbf{y}) \geq -(1 - \beta(k)) \cdot \log \left(\frac{1}{1 - \beta(k)} \right).$$

Plugging the above in Equation (24), we obtain:

$$m \geq -(1 - \beta(k)) \log \left(\frac{1}{1 - \beta(k)} \right) + k\beta(k).$$

Thus,

$$k \leq \frac{m + (1 - \beta(k)) \log \left(\frac{1}{1 - \beta(k)} \right)}{\beta(k)} \leq \frac{m + 0.54}{\beta(k)}$$

where the last inequality follows from the fact that the function $(1 - w) \log(1/(1 - w))$ is maximized at $w = (e - 1)/e$ (and takes value < 0.54). Solving for $\beta(k)$ gives the claimed bound.

We now prove the second direction of the lemma. Note that we can use Lemma A.6 to say that (\mathbf{X}, \mathbf{Y}) and $(\mathbf{X} \otimes \mathbf{Y})$ are point-wise $(2k \ln(2), \beta')$ indistinguishable, for $\beta' = \frac{2\beta}{1 - 2^{-k}}$. Note that $\beta' \leq 0.3$ as $\beta \in \left[0, \frac{3(1 - 2^{-k})}{20}\right]$. We now define the following “bad” set \mathcal{B} :

$$\mathcal{B} = \left\{ (\mathbf{x}, \mathbf{y}) \in \Sigma \times \Sigma : \ln \left(\frac{(\mathbf{X}, \mathbf{Y}) = (\mathbf{x}, \mathbf{y})}{\mathbf{X} \otimes \mathbf{Y} = (\mathbf{x}, \mathbf{y})} \right) > 2k \ln(2) \right\}.$$

From the definition of point-wise indistinguishability, we have:

$$I(\mathbf{X}; \mathbf{Y}) \leq 2k \ln(2) + \sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{B}} \Pr[(\mathbf{X}, \mathbf{Y}) = (\mathbf{x}, \mathbf{y})] \ln \left(\frac{\Pr[(\mathbf{X}, \mathbf{Y}) = (\mathbf{x}, \mathbf{y})]}{\Pr[\mathbf{X} = \mathbf{x}] \Pr[\mathbf{Y} = \mathbf{y}]} \right). \quad (25)$$

Now, if set \mathcal{B} is empty, we get from (25) that:

$$I(\mathbf{X}; \mathbf{Y}) \leq 2k \ln(2)$$

which trivially gives us the claimed bound.

However, if set \mathcal{B} is non-empty, we then consider the mutual information conditioned on the event $(\mathbf{X}, \mathbf{Y}) \in \mathcal{B}$. We will write \mathbf{X}' for the random variable \mathbf{X} conditioned on $(\mathbf{X}, \mathbf{Y}) \in \mathcal{B}$, and \mathbf{Y}' for the random variable \mathbf{Y} conditioned on the same event. We can then obtain the following bound, where we write $H(W)$ to denote the entropy of random variable W :

$$I(\mathbf{X}'; \mathbf{Y}') \leq H(\mathbf{X}') \leq \log |\Sigma|.$$

We can then make a relation between the mutual information $I(\mathbf{X}', \mathbf{Y}')$ and the sum of terms over \mathcal{B} in (25). Note that, for $(\mathbf{x}, \mathbf{y}) \in \mathcal{B}$, we have from Bayes' rule:

$$\Pr [(\mathbf{X}, \mathbf{Y}) \in \mathcal{B}] \Pr [(\mathbf{X}, \mathbf{Y}) = (\mathbf{x}, \mathbf{y}) | \mathcal{B}] = \Pr [(\mathbf{X}, \mathbf{Y}) = (\mathbf{x}, \mathbf{y})].$$

This then gives us the following bound:

$$\sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{B}} \frac{\Pr [(\mathbf{X}, \mathbf{Y}) = (\mathbf{x}, \mathbf{y})]}{\Pr [(\mathbf{X}, \mathbf{Y}) \in \mathcal{B}]} \ln \left(\frac{\Pr [(\mathbf{X}, \mathbf{Y}) = (\mathbf{x}, \mathbf{y})]}{\Pr [(\mathbf{X}, \mathbf{Y}) \in \mathcal{B}] \Pr [\mathbf{X} = \mathbf{x} | \mathcal{B}] \Pr [\mathbf{Y} = \mathbf{y} | \mathcal{B}]} \right) \leq \log |\Sigma|. \quad (26)$$

Note that $\Pr [\mathbf{X} = \mathbf{x} | \mathcal{B}] \leq \frac{\Pr [\mathbf{X} = \mathbf{x}]}{\Pr [(\mathbf{X}, \mathbf{Y}) \in \mathcal{B}]}$, and similarly, $\Pr [\mathbf{Y} = \mathbf{y} | \mathcal{B}] \leq \frac{\Pr [\mathbf{Y} = \mathbf{y}]}{\Pr [(\mathbf{X}, \mathbf{Y}) \in \mathcal{B}]}$. From (26), we have:

$$\begin{aligned} & \sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{B}} \Pr [(\mathbf{X}, \mathbf{Y}) = (\mathbf{x}, \mathbf{y})] \ln \left(\frac{\Pr [(\mathbf{X}, \mathbf{Y}) = (\mathbf{x}, \mathbf{y})]}{\Pr [\mathbf{X} = \mathbf{x}] \Pr [\mathbf{Y} = \mathbf{y}]} \right) \\ & \leq \Pr [(\mathbf{X}, \mathbf{Y}) \in \mathcal{B}] \log |\Sigma| + \Pr [(\mathbf{X}, \mathbf{Y}) \in \mathcal{B}] \ln (1 / \Pr [(\mathbf{X}, \mathbf{Y}) \in \mathcal{B}]) \\ & \leq \beta' (\log |\Sigma| + \ln(1/\beta')) \\ & = \frac{2\beta}{1 - 2^{-k}} \left(\log |\Sigma| + \ln \left(\frac{1 - 2^{-k}}{2\beta} \right) \right) \\ & \leq \frac{2\beta \log(|\Sigma|/2\beta)}{1 - 2^{-k}} \end{aligned}$$

where the last inequality follows from the fact that $\Pr [(\mathbf{X}, \mathbf{Y}) \in \mathcal{B}] \leq \beta'$ and $w \ln(1/w) \leq \beta' \ln(1/\beta')$ when $0 \leq w \leq \beta' \leq 0.3$. Substituting the sum of terms over \mathcal{B} in (25) by the above gives us the claimed bound. \square

We have already shown how Lemma 1.5 along with Theorem 3.1 can be used to improve a bound from [MMP⁺11]. We now state another corollary which improves on a result from [RZ16].

In the introduction, we proved a simple theorem about how to correct p -values (using a *valid p -value correction function* – Definition 1.1) given a bound on the *max-information* between the input dataset and the test-statistic selection procedure \mathcal{A} (Theorem 1.3). We note that we can easily extend the definition of null hypotheses given in the introduction (and hence p -values and correction functions), to allow for distributions \mathcal{S} over \mathcal{X}^n that need not be product distributions. In fact, we can restate Theorem 1.3 in terms of non-product distributions:

Theorem C.1. *Let $\mathcal{A} : \mathcal{X}^n \rightarrow \mathcal{T}$ be a data-dependent algorithm for selecting a test statistic such that $I_\infty^\beta(\mathcal{A}, n) \leq k$. Then the following function γ is a valid p -value correction function for \mathcal{A} :*

$$\gamma(\alpha) = \max \left(\frac{\alpha - \beta}{2^k}, 0 \right).$$

Proof. The proof is exactly the same as the proof of Theorem 1.3, except we fix an arbitrary (perhaps non-product) distribution \mathcal{S} from which the dataset \mathbf{X} is drawn. \square

Previously, Russo and Zou [RZ16] have given a method to correct p -values given a bound on the *mutual information* between the input data and the test-statistic selection procedure.⁵ In Lemma 1.5, we observe that if we had a bound on the mutual information between the input data and the test-statistic procedure, this would imply a bound on the *max-information* that would be sufficiently strong so that our Theorem C.1 would give us the following corollary:

Corollary C.2. *Let $\mathcal{A} : \mathcal{X}^n \rightarrow \mathcal{T}$ be a test-statistic selection procedure such that $I(\mathbf{X}; \mathcal{A}(\mathbf{X})) \leq m$. Then $\gamma(\alpha)$ is a valid p -value correction function, where:*

$$\gamma(\alpha) = \frac{\alpha}{2} \cdot 2^{\frac{-2}{\alpha}(m+0.54)}.$$

Proof. From Lemma 1.5, we know that for any $k > 0$, $I_\infty^{\beta(k)}(\mathcal{A}, n) \leq k$, where $\beta(k) \leq \frac{m+0.54}{k}$. Hence, from Theorem C.1, we know that for any choice of $k > 0$, $\gamma(\alpha)$ is a valid p -value correction function where:

$$\gamma(\alpha) = \frac{\alpha - \frac{m+0.54}{k}}{2^k}.$$

Choosing $k = \frac{2(m+0.54)}{\alpha}$ gives our claimed bound. \square

We now show that this gives us a strictly improved p -value correction function than the bound given by [RZ16], which we state here using our terminology.

Theorem C.3 ([RZ16] Proposition 7). *Let $\mathcal{A} : \mathcal{X}^n \rightarrow \mathcal{T}$ be a test-statistic selection procedure such that $I(\mathbf{X}; \mathcal{A}(\mathbf{X})) \leq m$ (where I denotes mutual information). If we define $\phi_i = \mathcal{A}(\mathbf{X})$, then for every $\gamma \in [0, 1]$:*

$$\Pr [p_i(\phi_i(\mathbf{X})) \leq \gamma] \leq \gamma + \sqrt{\frac{m}{\ln(1/2\gamma)}}.$$

If we want to set parameters so that the probability of a false discovery is at most α , then in particular, we must pick γ such that $\sqrt{\frac{m}{\ln(1/2\gamma)}} \leq \alpha$. Equivalently, solving for α , the best valid p -value correction function implied by the bound of [RZ16] must satisfy:

$$\gamma_{RZ}(\alpha) \leq \min \left\{ \frac{\alpha}{2}, \frac{1}{2} \cdot 2^{-\log(e)m/\alpha^2} \right\}.$$

We can obtain a better bound by instead arguing via max-information combining Lemma 1.5 with Theorem C.1, we can derive a valid p -value correction function given a bound on the mutual information between the data and the selection procedure:

Comparing the p -value correction function $\gamma(\alpha)$ derived above, with the function $\gamma_{RZ}(\alpha)$ that arises from [RZ16], we see that the version above has an exponentially improved dependence on $1/\alpha$. Moreover, it almost always gives a better correction factor in practice: for any value of $\alpha \leq 0.05$, the function $\gamma(\alpha)$ derived above improves over $\gamma_{RZ}(\alpha)$ whenever the mutual information bound $m \geq 0.05$ (whereas, we would naturally expect the mutual information to be $m \gg 1$, and to scale with n).

⁵ Actually, [RZ16] do not explicitly model the dataset, and instead give a bound in terms of the mutual information between the test-statistics themselves and the test-statistic selection procedure. We could also prove bounds with this dependence, by viewing our input data to be the value of the given test-statistics, however for consistency, we will discuss all bounds in terms of the mutual information between the data and the selection procedure.

D Rederiving Generalization for Low Sensitivity Queries, and a Comparison with the Bounds of [BNS⁺16]

In this section, we use the bound from our main theorem (Theorem 3.1) to rederive known results for the generalization properties of differentially private algorithms which select *low sensitivity queries*. Our bounds do not exactly – but nearly – match the tight bounds for this problem, given in [BNS⁺16]. This implies a limit on the extent to which our main theorem can be quantitatively improved, despite its generality.

The goal of generalization bounds for low sensitivity queries is to argue that with high probability, if a low sensitivity function $f : \mathcal{X}^n \rightarrow \mathcal{R} = \mathcal{A}(\mathbf{x})$ is chosen in a data-dependent way when \mathbf{x} is sampled from a product distribution \mathcal{P}^n , then the value of the function on the realized data $f(\mathbf{x})$ is close to its expectation $f(\mathcal{P}^n) \stackrel{\text{def}}{=} \mathbb{E}_{\mathbf{X} \sim \mathcal{P}^n} [f(\mathbf{X})]$. If f were selected in a data-independent manner, this would follow from McDiarmid’s inequality. No bound like this is true for arbitrary selection procedures \mathcal{A} , but a tight bound by [BNS⁺16] is known when \mathcal{A} is (ϵ, δ) -differentially private – we have already quoted it, as Theorem B.2.

Using our main theorem (Theorem 3.1), together with McDiarmid’s inequality (Theorem A.9), we can derive a comparable statement to Theorem B.2:

Theorem D.1. *Let $\epsilon \in (0, 1)$, $\delta = O(\epsilon^5)$, and $n = \Omega\left(\frac{\log(\epsilon/\delta)}{\epsilon^2}\right)$. Let \mathcal{Y} denote the class of Δ -sensitive functions $f : \mathcal{X}^n \rightarrow \mathbb{R}$, and let $\mathcal{A} : \mathcal{X}^n \rightarrow \mathcal{Y}$ be an algorithm that is (ϵ, δ) -differentially private. Let $\mathbf{X} \sim \mathcal{P}^n$ for some distribution \mathcal{P} over \mathcal{X} , and let $q = \mathcal{A}(\mathbf{X})$. Then there exists a constant C such that:*

$$\Pr_{\mathbf{X}, \mathcal{A}} [|q(\mathcal{P}^n) - q(\mathbf{X})| \geq C\epsilon\Delta n] < n\sqrt{\frac{\delta}{\epsilon}}$$

Proof. If \mathcal{A} satisfied $I_\infty^\beta(\mathbf{X}; \mathcal{A}(\mathbf{X})) \leq k$, then McDiarmid’s inequality (Theorem A.9) paired with Definition 1.2 would imply:

$$\Pr_{\mathbf{X}, \mathcal{A}} [|q(\mathcal{P}^n) - q(\mathbf{X})| \geq C\epsilon\Delta n] < 2^k \exp(-2C^2\epsilon^2 n) + \beta$$

Because \mathcal{A} is (ϵ, δ) -differentially private, Theorem 3.1 implies that indeed $I_\infty^\beta(\mathbf{X}; \mathcal{A}(\mathbf{X})) \leq k$ for $k = O\left(n\epsilon^2 + n\sqrt{\frac{\delta}{\epsilon}}\right) = O(\epsilon^2 n)$ and $\beta = O\left(n\sqrt{\frac{\delta}{\epsilon}}\right) + e^{-\epsilon^2 n} = O\left(n\sqrt{\frac{\delta}{\epsilon}}\right)$, and the claimed bound follows. \square

Because Theorem B.2 is asymptotically tight, this implies a limit on the extent to which Theorem 3.1 can be quantitatively improved.

For comparison, note that the generalization bound for (ϵ, δ) -differentially private mechanisms given in Theorem B.2 differs by only constants from the generalization bound proven via the maximum information approach for (ϵ, δ') -differentially private mechanisms, where:

$$\delta' = \frac{\delta^2}{\epsilon n^2}$$

Note that in most applications (including the best known mechanism for answering large numbers of low sensitivity queries privately— the median mechanism [RR10] as analyzed in [DR14]

Theorem 5.10), the accuracy of a differentially private algorithm scales with $\sqrt{\frac{\log(1/\delta)}{n}}$ (ignoring other relevant parameters). In such cases, using the bound derived from the max-information approach yields an accuracy that is worse than the bound from Theorem B.2 by an additive term that is $O\left(\sqrt{\frac{\log(\epsilon n)}{n}}\right)$.

E Omitted Proofs

Proof of Corollary 3.10. We use the same algorithms \mathcal{A} and \mathcal{B} from Theorem 3.9. Suppose that for all $\mathbf{a} \in \{0, 1\}^r$ and $0 < \beta_2 < 1/2 - \delta - \beta_1$ for any $\beta_1 \in (0, 1/2 - \delta)$, we have:

$$I_{\infty}^{\beta_2}(\mathcal{B}(\cdot, \mathbf{a}), n) < n - 1 - r - \log(1/\beta_1).$$

Note that because \mathcal{A} has bounded description length r , we can bound $I_{\infty}^{\beta_1}(\mathcal{A}, n) \leq r + \log(1/\beta_1)$ for any $\beta_1 > 0$ [DFH⁺15a]. We then apply the composition theorem for max-information mechanisms, Theorem A.3, to obtain:

$$I_{\infty}^{\beta_2 + \beta_1}(\mathcal{B} \circ \mathcal{A}, n) < n - 1.$$

However, this contradicts Theorem 3.9, because for any $\beta < 1/2 - \delta$,

$$I_{\infty}^{\beta}(\mathcal{B} \circ \mathcal{A}, n) \geq I_{\infty, P}^{\beta}(\mathcal{B} \circ \mathcal{A}, n) \geq n - 1.$$

Thus, we know that there exists some $\mathbf{a}^* \in \{0, 1\}^r$ and (non-product) distribution $\mathbf{X} \sim \mathcal{S}$ such that:

$$I_{\infty}^{\beta_2}(\mathbf{X}; \mathcal{B}(\mathbf{X}, \mathbf{a}^*)) \geq n - 1 - r - \log(1/\beta_1).$$

We then define $\mathcal{C} : \mathcal{X}^n \rightarrow \mathcal{Y}$ to be $\mathcal{C}(\mathbf{x}) = \mathcal{B}(\mathbf{x}, \mathbf{a}^*)$. Hence,

$$I_{\infty}^{\beta_2}(\mathcal{C}, n) \geq I_{\infty}^{\beta_2}(\mathbf{X}; \mathcal{C}(\mathbf{X})) \geq n - 1 - r - \log(1/\beta_1)$$

which completes the proof. □