

Local Private Hypothesis Testing: Chi-Square Tests

¹, Marco Gaboardi¹, and Ryan Rogers²

¹*University at Buffalo, The State University of New York*

²*University of Pennsylvania*

June 7, 2018

Abstract

The local model for differential privacy is emerging as the reference model for practical applications of collecting and sharing sensitive information while satisfying strong privacy guarantees. In the local model, there is no *trusted* entity which is allowed to have each individual’s raw data as is assumed in the traditional *curator* model for differential privacy. Individuals’ data are usually perturbed before sharing them. We explore the design of private hypothesis tests in the local model, where each data entry is perturbed to ensure the privacy of each participant. Specifically, we analyze locally private chi-square tests for goodness of fit and independence testing, which have been studied in the traditional, curator model for differential privacy.

1 Introduction

Hypothesis testing is a widely applied statistical tool used to test whether given models should be rejected, or not, based on sampled data from a population. Hypothesis testing was initially developed for scientific and survey data, but today it is also an essential tool to test models over collections of social network, mobile, and crowdsourced data (American Statistical Association, 2014; Hunter et al., 2008; Steele et al., 2017). Collected data samples may contain highly sensitive information about the subjects, and the privacy of individuals can be compromised when the results of a data analysis are released. In work from Homer et al. (2008), it was shown that a subject in a dataset can be identified as being in the case or control group based on the aggregate statistics of a genetic-wide association study (GWAS). Privacy-risks may bring data contributors to opt out, which reduces the confidence in the data study. A way to address this concern is by developing new techniques to support privacy-preserving data analysis. Among the different approaches, differential privacy (Dwork et al., 2006b) has emerged as a viable solution: it provides strong privacy guarantees and it allows to release accurate statistics. A standard way to achieve differential privacy is by injecting some statistical noise in the computation of the data analysis. When the noise is carefully chosen, it helps to protect the individual privacy without compromising the utility of the data analysis. Several recent works have studied differentially private hypothesis tests that can be used in place of the standard, non-private hypothesis tests (Uhler et al., 2013; Yu et al., 2014; Sheffet, 2015; Karwa and Slavković, 2016; Wang et al., 2015; Gaboardi et al., 2016; Kifer and Rogers, 2017; Cai et al., 2017). These tests work in the *curator model* of differential privacy. In this model, the data is centrally stored and the curator carefully injects noise in the computation of the data analysis in order to satisfy differential privacy.

In this work we instead address the *local model* of privacy, formally introduced by Raskhodnikova et al. (2008). The first differentially private algorithm called *randomized response* – in fact it predates the definition of differential privacy by more than 40 years – guarantees differential privacy in the local model (Warner, 1965). In this model, there is no *trusted* centralized entity which is responsible for the noise injection. Instead, each individual adds enough noise to guarantee differential privacy for their own data, which provides a stronger

privacy guarantee when compared to traditional differential privacy. The data analysis is then run over the collection of the individually sanitized data. The local model of differential privacy is a convenient model for several applications: for example it is used to collect statistics about the activity of the Google Chrome Web browser users (Erlingsson et al., 2014), and to collect statistics about the typing patterns of Apple’s iPhone users (Apple Press Info, 2016). Despite these applications, the local model has received far less attention than the more standard centralized curator model. This is in part due to the more firm requirements imposed by this model, which make the design of effective data analysis harder.

Our main contribution is in designing chi-square hypothesis tests for the local model of differential privacy. Similar to previous works we focus on goodness of fit and independence hypothesis tests. Most of the differentially private chi-square hypothesis tests proposed so far are based on mechanisms that add noise in some form to the aggregate data, e.g. the cells of the contingency tables, or the resulting chi-square statistics value. These approaches cannot be used in the local model, since noise needs to be added at the individual’s data level. We then consider instead general privatizing techniques in the local model, and we study how to build new hypothesis tests with them. Each test we present is characterized by a specific local model mechanism. The main technical challenge for designing each test is to create statistics, which incorporate the local model mechanisms, that converge as we collect more data to a chi-square distribution, as in the classical chi-square tests. We then use these statistics to find the critical value to correctly bound the Type I error.

We present three different goodness of fit tests: `LocalNoiseGOF` presents a statistic that guarantees the convergence to a chi-square distribution under the null hypothesis so that we can use the correct critical values when local (concentrated) differential privacy is guaranteed by adding Laplace or Gaussian noise to the individual data; `LocalGenRRGOF` also provides a statistic that converges to a chi-square under the null hypothesis when a private value for each individual is selected by using the exponential mechanism (McSherry and Talwar, 2007) which in this case can be thought of as a generalized randomized response (Warner, 1965) over d categories; finally, `LocalBitFlipGOF` introduces a statistic that converges to a chi-square distribution when the data is privatized using a bit flipping algorithm (Bassily and Smith, 2015), which provide better accuracy for higher dimensional data. Further, we develop corresponding independence tests: `LocalNoiseIND`, `LocalGenRRIND`, and `LocalBitFlipIND`. For all these tests we study their asymptotic behavior. A desiderata for private hypothesis tests is to have a guaranteed upper bound on the probability of a false discovery (or Type I error) – rejecting a null hypothesis or model when the data was actually generated from it – and to minimize the probability of a Type II error, which is failing to reject the null hypothesis when the model is indeed false. This latter criteria corresponds to the *power* of the statistical test. We then present experimental results showing the power of the different tests which demonstrates that no single local differentially private algorithm is best across all data dimensions and privacy parameter regimes.

2 Related Works

There have been several works in developing private hypothesis test for categorical data, but all look at the traditional model of (concentrated) differential privacy instead of the local model, which we consider here. Several works have explored private statistical inference for GWAS data, (Uhler et al., 2013; Yu et al., 2014; Johnson and Shmatikov, 2013). Following these works, there has also been general work in private chi-square hypothesis tests, where the main tests are for goodness of fit and independence testing, although some do extend to more general tests (Wang et al., 2015; Gaboardi et al., 2016; Kifer and Rogers, 2017; Cai et al., 2017; Kakizaki et al., 2017). Among these, the works most related to ours are the ones by (Gaboardi et al., 2016; Kifer and Rogers, 2017). We will compare our work with these in Section 4 after introducing them. There has also been work in private hypothesis testing for ordinary least squares regression (Sheffet, 2015).

There are other works that have studied statistical inference and estimators in the local model of differential privacy. Duchi et al. (2013b,a) focus on controlling disclosure risk in statistical estimation and inference by ensuring the analysis satisfies local differential privacy. They provide minimax convergence rates to show the

tight tradeoffs between privacy and statistical efficiency, i.e. the number of samples required to give quality estimators. In their work, they show that a generalized version of randomized response gives optimal sample complexity for estimating the multinomial probability vector. We use this idea as the basis for our hypothesis test `LocalBitFlipGOF`. Kairouz et al. (2014) also considers hypothesis testing in the local model, although they measure utility in terms of f -divergences and do not give a decision rule, i.e. when to reject a given null hypothesis. We provide statistics whose distributions asymptotically follow a chi-square distribution, which allows for approximating statistical p -values that can be used in a decision rule. We consider their *extremal* mechanisms and empirically confirm their result that for small privacy regimes (small ϵ) one mechanism has higher *utility* than other mechanisms and for large privacy regimes (large ϵ) a different mechanism outperforms the others. However, we measure utility in terms of the power of a locally private hypothesis test subject to a given Type I error bound. Other notable works in the local privacy model include Pastore and Gastpar (2016); Kairouz et al. (2016); Ye and Barg (2017)

Independent of this work, another paper (Sheffet, 2018) has addressed local private hypothesis testing. Sheffet (2018) considers finite sample complexity by showing certain test quantities take different values under the null- and alternative-hypothesis. In this work, we design and analyze asymptotic statistical tests and empirically evaluate the performance of each test for finite samples.

3 Preliminaries

We consider datasets $\mathbf{x} = (x_1, \dots, x_n) \in \mathcal{X}^n$ in some data universe \mathcal{X} , typically $\mathcal{X} = \{0, 1\}^d$ where d is the dimensionality. We first present the standard definition of differential privacy, as well as its variant *concentrated differential privacy*. We say that two datasets $\mathbf{x}, \mathbf{x}' \in \mathcal{X}^n$ are *neighboring* if they differ in at most one element, i.e. $\exists i \in [n]$ such that $x_i \neq x'_i$ and $\forall j \neq i, x_j = x'_j$.

Definition 3.1 [Dwork et al. (2006b,a)]. *An algorithm $\mathcal{M} : \mathcal{X}^n \rightarrow \mathcal{Y}$ is (ϵ, δ) -differentially private (DP) if for all neighboring datasets $\mathbf{x}, \mathbf{x}' \in \mathcal{X}^n$ and for all outcomes $S \subseteq \mathcal{Y}$, we have $\Pr[\mathcal{M}(\mathbf{x}) \in S] \leq e^\epsilon \Pr[\mathcal{M}(\mathbf{x}') \in S] + \delta$.*

We then state the definition of zero-mean concentrated differential privacy.

Definition 3.2 [Bun and Steinke (2016)]. *An algorithm $\mathcal{M} : \mathcal{X}^n \rightarrow \mathcal{Y}$ is ρ -zero-mean concentrated differentially private (zCDP) if for all neighboring datasets $\mathbf{x}, \mathbf{x}' \in \mathcal{X}^n$, we have the following bound for all $t > 0$ where the expectation is over outcomes $y \sim \mathcal{M}(\mathbf{x})$,*

$$\mathbb{E} \left[\exp \left(t \left(\ln \left(\frac{\Pr[\mathcal{M}(\mathbf{x}) = y]}{\Pr[\mathcal{M}(\mathbf{x}') = y]} \right) - \rho \right) \right) \right] \leq e^{t^2 \rho}.$$

Note that in both of these privacy definitions, it is assumed that all the data is stored in a central location and the algorithm \mathcal{M} can access all the data. Most of the work in differential privacy has been in this *trusted curator model*. One of the main reasons for this is that we can achieve much greater accuracy in our differentially private statistics when used in the curator setting. However, in many cases, having a trusted curator is too strong of an assumption. We then define *local* differential privacy, formalized by Raskhodnikova et al. (2008) and Dwork and Roth (2014), which does not require the subjects to release their raw data, rather each data entry is perturbed to prevent the true entry from being stored. Thus, local differential privacy ensures a very strong privacy guarantee.

Definition 3.3 [LR Oracle]. *Given a dataset \mathbf{x} , a local randomizer oracle $LR_{\mathbf{x}}(\cdot, \cdot)$ takes as input an index $i \in [n]$ and an ϵ -DP algorithm R , and outputs $y \in \mathcal{Y}$ chosen according to the distribution of $R(x_i)$, i.e. $LR_{\mathbf{x}}(i, R) = R(x_i)$.*

Definition 3.4 [Raskhodnikova et al. (2008)]. *An algorithm $\mathcal{M} : \mathcal{X}^n \rightarrow \mathcal{Y}$ is (ϵ, δ) -local differentially private (LDP) if it accesses the input database \mathbf{x} via the LR oracle $LR_{\mathbf{x}}$ with the following restriction: if $LR(i, R_j)$ for $j \in [k]$ are \mathcal{M} 's invocations of $LR_{\mathbf{x}}$ on index i , then each R_j for $j \in [k]$ is (ϵ_j, δ_j) -DP and $\sum_{j=1}^k \epsilon_j \leq \epsilon$, $\sum_{j=1}^k \delta_j \leq \delta$.*

An easy consequence of this definition is that an algorithm which is (ϵ, δ) -LDP is also (ϵ, δ) -DP. Note that these definitions can be extended to include ρ -local zCDP (LzCDP) where each local randomizer is ρ_j -zCDP and $\sum_{j=1}^k \rho_j \leq \rho$. We point out the following connection between LzCDP and LDP, which follows directly from results in Bun and Steinke (2016)

Lemma 3.5. *If $\mathcal{M} : \mathcal{X}^n \rightarrow \mathcal{Y}$ is $(\epsilon, 0)$ -LDP then it is also $\epsilon^2/2$ -LzCDP. If \mathcal{M} is ρ -LzCDP, then it is also $\left(\left(\rho + \sqrt{2\rho \ln(2/\delta)}\right), \delta\right)$ -LDP for any $\delta > 0$.*

Thus, in the local setting, (pure) LDP (where $\delta = 0$) provides the strongest level of privacy, followed by LzCDP and then *approximate*-LDP (where $\delta > 0$).

4 Chi-Square Hypothesis Tests

As was studied in Gaboardi et al. (2016), Wang et al. (2015), and Kifer and Rogers (2017), we will study hypothesis tests with categorical data. A null hypothesis, or model H_0 is how we might expect the data to be generated. The goal for hypothesis testing is to reject the null hypothesis if the data is not likely to have been generated from the given model. As is common in statistical inference, we want to design hypothesis tests to bound the probability of a false discovery (or Type I error), i.e. rejecting a null hypothesis when the data was actually generated from it, by at most some amount α , such as 5%. However, designing tests that achieve this is easy, because we can just ignore the data and always *fail to reject* the null hypothesis, i.e. have an inconclusive test. Thus, we would additionally like to design our tests so that they can reject H_0 if the data was not actually generated from the given model. We then want to minimize the probability of a Type II error, which is failing to reject H_0 when the model is false, subject to a given Type I error.

For goodness of fit testing, we assume that each individual's data \mathbf{X}_i for $i \in [n]$ is sampled i.i.d. from Multinomial(1, \mathbf{p}) where $\mathbf{p} \in \mathbb{R}_{>0}^d$ and $\mathbf{p}^\top \mathbf{1} = 1$. The classical chi-square hypothesis test (without privacy) forms the histogram $\mathbf{H} = (H_1, \dots, H_d) = \sum_{i=1}^n \mathbf{X}_i$ and computes the chi-square statistic $T = \sum_{j=1}^d \frac{(H_j - np_j^0)^2}{np_j^0}$.

The reason for using this statistic is that it converges in distribution to χ_{d-1}^2 as more data is collected, i.e. $n \rightarrow \infty$, when $H_0 : \mathbf{p} = \mathbf{p}^0$ holds. Hence, we can ensure the probability of false discovery to be close to α as long as we only reject H_0 when $T > \chi_{d-1, 1-\alpha}^2$ where the *critical value* $\chi_{d-1, 1-\alpha}^2$ is defined as the following quantity $\Pr \left[\chi_{d-1}^2 > \chi_{d-1, 1-\alpha}^2 \right] = \alpha$.

4.1 Prior Private Chi-square Tests in the Curator Model

One approach for chi-square private hypothesis tests that was explored by Gaboardi et al. (2016) and Wang et al. (2015) is to add noise (Gaussian or Laplace) directly to the histogram to ensure privacy and then use the classical test statistic. Note that the resulting asymptotic distribution needs to be modified for such changes to the statistic – it is no longer a chi-square random variable. To introduce the different statistics, we will consider goodness of fit testing after adding noise \mathbf{Z} from distribution \mathcal{D}^n to the histogram of counts $\tilde{\mathbf{H}} = \mathbf{H} + \mathbf{Z}$, which ensures ρ -zCDP when $\mathcal{D} = N(0, 1/\rho)$ and ϵ -DP when $\mathcal{D} = \text{Lap}(2/\epsilon)$. The chi-square statistic then becomes

$$\tilde{T}(\mathcal{D}) = \sum_{i=1}^d \frac{(H_i + Z_i - np_i^0)^2}{np_i^0} \quad \text{where } \mathbf{Z} \sim \mathcal{D}^n. \quad (1)$$

They then show that this statistic converges in distribution to a linear combination of chi-squared random variables, when $\mathcal{D} \sim N(0, 1/\rho)$ and ρ is also decreasing with n .

In followup work from Kifer and Rogers (2017), the authors showed that modifying the chi-square statistic to account for the additional noise leads to tests with better empirical power. The *projected* statistic from

Kifer and Rogers (2017) is the following where we use projection matrix $\Pi \stackrel{\text{defn}}{=} (I_d - \frac{1}{d}\mathbf{1}\mathbf{1}^\top)$, middle matrix $M_\sigma = \Pi (\text{Diag}(\mathbf{p}^0 + \sigma) - \mathbf{p}^0 (\mathbf{p}^0)^\top)^{-1} \Pi$, and sample noise $\mathbf{Z} \sim \mathcal{D}^n$,

$$\mathbb{T}_{KR}^{(n)}(\sigma; \mathcal{D}) = n \left(\frac{\mathbf{H} + \mathbf{Z}}{n} - \mathbf{p}^0 \right)^\top M_\sigma \left(\frac{\mathbf{H} + \mathbf{Z}}{n} - \mathbf{p}^0 \right). \quad (2)$$

We use $\mathcal{D} = \text{Lap}(2/\epsilon)$ with $\sigma = \frac{8}{n\epsilon^2}$ for an ϵ -DP claim or $\mathcal{D} = \text{N}(0, 1/\rho)$ with $\sigma = \frac{1}{n\rho}$ for a ρ -zCDP claim. When comparing the power of all our tests, we will be considering the alternate $\mathbb{H}_1 : \mathbf{p} = \mathbf{p}_n^1$ where

$$\mathbf{p}_n^1 = \mathbf{p}^0 + \frac{\Delta}{\sqrt{n}} \quad \text{where } \mathbf{1}^\top \Delta = 0. \quad (3)$$

Theorem 4.1 [Kifer and Rogers (2017)]. *Under the null hypothesis $\mathbb{H}_0 : \mathbf{p} = \mathbf{p}^0$, the statistic $\mathbb{T}_{KR}^{(n)}\left(\frac{1}{n\rho}; \text{N}(0, 1/\rho)\right)$ given in (2) for $\rho > 0$ converges in distribution to χ_{d-1}^2 . Further, under the alternate hypothesis $\mathbb{H}_1 : \mathbf{p} = \mathbf{p}_n^1$, the resulting asymptotic distribution is a noncentral chi-square random variable,*

$$\mathbb{T}_{KR}^{(n)}\left(\frac{1}{n\rho}; \text{N}(0, 1/\rho)\right) \xrightarrow{D} \chi_{d-1}^2 \left(\Delta^\top \left(\text{Diag}(\mathbf{p}^0) - \mathbf{p}^0 (\mathbf{p}^0)^\top + 1/\rho I_d \right)^{-1} \Delta \right).$$

When $\mathcal{D} = \text{Lap}(2/\epsilon)$, we can still obtain the null hypothesis distribution using Monte Carlo simulations to estimate the critical value, since the asymptotic distribution will no longer be chi-square. That is, we can obtain m samples from the statistic under the null hypothesis with Laplace noise added to the histogram of counts. We can then guarantee that the probability of a false discovery is at most α as long as $m > \lceil 1/\alpha \rceil$ (see Gaboardi et al. (2016) for more details).

5 Local Private Chi-Square Goodness of Fit Tests

We now turn to designing local private goodness of fit tests. We begin by showing how the existing statistics from the previous section can be adapted to the local setting and then develop new tests based on the generalized randomized response mechanism that returns one of $d > 1$ categories and bit flipping (Bassily and Smith, 2015). Each test is locally private because it perturbs each individual's data through a local randomizer. However, each of them has a different asymptotic behavior and so we need different analyses to identify the different critical values. We empirically check the power of each test to see which tests outperform others in different parameter regimes. An interesting result of this analysis is that the power of a test is directly related to the size of the noncentral parameter of the chi-square statistic under the alternate distribution.

5.1 Goodness of Fit Test with Noise Addition

In the local model we can add $\mathbf{Z}_i \sim \text{N}\left(\mathbf{0}, \frac{1}{\rho} I_d\right)$ independent noise to each individual's data \mathbf{X}_i to ensure ρ -LzCDP or $\mathbf{Z}_i \stackrel{i.i.d.}{\sim} \text{Lap}\left(\frac{2}{\epsilon}\right)$ independent noise to \mathbf{X}_i to ensure ϵ -LDP. In either case, the resulting noisy histogram $\widehat{\mathbf{H}} = \mathbf{H} + \mathbf{Z}$ where $\mathbf{Z} = \sum_i \mathbf{Z}_i$ will have variance that scales with n for fixed privacy parameters $\epsilon, \rho > 0$. Consider the case where we add Gaussian noise, which results in the following histogram, $\widehat{\mathbf{H}} = \mathbf{H} + \mathbf{Z}$ where $\mathbf{Z} \sim \text{N}\left(\mathbf{0}, \frac{n}{\rho} I_d\right)$. Thus, we can use either statistic $\widetilde{\mathbb{T}}(\rho/n)$ or $\mathbb{T}_{KR}^{(n)}(\rho/n)$, with the latter statistic

Algorithm 1 Locally Private GOF Test:LocalNoiseGOF

Input: $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$, ρ , α , $H_0 : \mathbf{p} = \mathbf{p}^0$.

Let $\mathbf{H} = \sum_{\ell=1}^n \mathbf{x}_\ell$

if $\mathcal{D} = N(0, n/\rho)$ **then**

Set $q = T_{KR}^{(n)}(1/\rho; \mathcal{D})$ given in (2).

if $q > \chi_{d-1, 1-\alpha}^2$ Decision \leftarrow Reject.

else Decision \leftarrow Fail to Reject.

if $\mathcal{D} = \sum_{i=1}^n \text{Lap}(2/\epsilon)$ **then**

Set $q = T_{KR}^{(n)}(8/\epsilon^2; \mathcal{D})$ given in (2).

Sample $m > \lceil 1/\alpha \rceil$ from the distribution of $T_{KR}^{(n)}(8/\epsilon^2; \mathcal{D})$ assuming H_0

Set τ to be the $\lceil (m+1)(1-\alpha) \rceil$ th largest sample.

if $q > \tau$ Decision \leftarrow Reject.

else Decision \leftarrow Fail to Reject.

Output: Decision

typically having better empirical power (Kifer and Rogers, 2017). We then give our first local private hypothesis test in Algorithm 1.

Theorem 5.1. *LocalNoiseGOF is ρ -LzCDP when $\mathcal{D} = N(0, 1/\rho)$ and ϵ -LDP when $\mathcal{D} = \text{Lap}(2/\epsilon)$.*

Proof. The proof follows from the fact that we are adding appropriately scaled noise to each individual's data via the Gaussian mechanism and then LocalNoiseGOF aggregates the privatized data, which is just a post-processing function on the privatized data. \square

Although we cannot guarantee the probability of a Type I error at most α due to the fact that we use the asymptotic distribution (as in the tests from prior work and the classical chi-square tests without privacy), we expect the Type I errors to be similar to those from the nonprivate test. Note that the test can be modified to accommodate arbitrary noise distributions, e.g. Laplace to ensure differential privacy as was done in Kifer and Rogers (2017). In this case, we can use a Monte Carlo (MC) approach to estimate the critical value τ that ensures the probability of a Type I error is at most α if we reject H_0 when the statistic is larger than τ . For the local setting, if each individual perturbs each coordinate by adding $\text{Lap}(2/\epsilon)$ then this will ensure our test is ϵ -LDP. However, the sum of independent Laplace random variables is not Laplace, so we will need to estimate a sum of n independent Laplace random variables using MC. In the experiments section we will compare the other local private tests with the version of LocalNoiseGOF which uses Laplace noise and samples m entries from the exact distribution under H_0 to find the critical value.

5.2 Goodness of Fit Test with Generalized Randomized Response

Rather than having to add noise to each component of the original data histogram, we consider applying randomized response to obtain a LDP hypothesis test. We will use a form of the *exponential mechanism* (McSherry and Talwar, 2007) given in Algorithm 2 which takes a single data entry from the set $\{\mathbf{e}_1, \dots, \mathbf{e}_d\}$, where $\mathbf{e}_j \in \mathbb{R}^d$ is the standard basis element with a 1 in the j th coordinate and is zero elsewhere, and reports the original entry with probability slightly more than uniform and otherwise reports a different element. Note that $\mathcal{M}_{\text{GenRR}}$ takes a single data entry and is ϵ -differentially private. This mechanism can also be interpreted as a generalized version of randomized response (Warner, 1965) over d categories.

We have the following result when we use $\mathcal{M}_{\text{GenRR}}$ on each data entry to obtain a private histogram.

Lemma 5.2. *If we have histogram $\mathbf{H} = \sum_{i=1}^n \mathbf{X}_i$, where $\{\mathbf{X}_i\} \stackrel{i.i.d.}{\sim} \text{Multinomial}(1, \mathbf{p})$ and we write $\check{\mathbf{H}} =$*

Algorithm 2 Generalized Randomized Response: $\mathcal{M}_{\text{GenRR}}$

Input: $\mathbf{x} \in \{\mathbf{e}_1, \dots, \mathbf{e}_d\}$, ϵ .

Let $q(\mathbf{x}, \mathbf{z}) = \mathbb{1}\{\mathbf{x} = \mathbf{z}\}$

Select $\tilde{\mathbf{x}}$ with probability $\frac{\exp[\epsilon q(\mathbf{x}, \tilde{\mathbf{x}})]}{e^\epsilon - 1 + d}$

Output: $\tilde{\mathbf{x}}$

$\sum_{i=1}^n \mathcal{M}_{\text{GenRR}}(\mathbf{X}_i, \epsilon)$ for each $i \in [n]$, then $\check{\mathbf{H}} \sim \text{Multinomial}(n, \check{\mathbf{p}})$ where

$$\check{\mathbf{p}} = \mathbf{p} \left(\frac{e^\epsilon}{e^\epsilon + d - 1} \right) + (1 - \mathbf{p}) \left(\frac{1}{e^\epsilon + d - 1} \right). \quad (4)$$

Once we have identified the distribution of $\check{\mathbf{H}}$, we can create a chi-square statistic by subtracting $\check{\mathbf{H}}$ by its expectation and dividing the difference by the expectation. Hence testing $H_0 : \mathbf{p} = \mathbf{p}^0$ after the data has passed through the generalized randomized response, is equivalent to testing $H_0 : \mathbf{p} = \check{\mathbf{p}}^0$ with data $\check{\mathbf{H}}$. We will use the following classical result to prove our theorems.

Theorem 5.3 [Ferguson (1996)]. *If $\mathbf{X} \sim \text{N}(\boldsymbol{\mu}, \Sigma)$ and Σ is a projection matrix of rank ν and $\Sigma\boldsymbol{\mu} = \boldsymbol{\mu}$ then $\mathbf{X}^\top \mathbf{X} \sim \chi_\nu^2(\boldsymbol{\mu}^\top \boldsymbol{\mu})$.*

We can then form a chi-square statistic using the private histogram $\check{\mathbf{H}}$ which will have the correct asymptotic distribution.

Theorem 5.4. *Let $\mathbf{H} \sim \text{Multinomial}(n, \mathbf{p})$ and $\check{\mathbf{H}}$ be given in Lemma 5.2 with privacy parameter $\epsilon > 0$. Under the null hypothesis $H_0 : \mathbf{p} = \mathbf{p}^0$, we have for $\check{\mathbf{p}}^0 = \frac{1}{e^\epsilon + d - 1} (e^\epsilon \mathbf{p}^0 + (1 - \mathbf{p}^0))$, and*

$$\text{T}_{\text{Exp}}^{(n)}(\epsilon) = \sum_{j=1}^d \frac{(\check{H}_j - n\check{p}_j^0)^2}{n\check{p}_j^0} \xrightarrow{D} \chi_{d-1}^2. \quad (5)$$

Further, with alternate $H_1 : \mathbf{p} = \mathbf{p}^1$, the resulting asymptotic distribution is the following, $\text{T}_{\text{Exp}}^{(n)}(\epsilon) \xrightarrow{D} \chi_{d-1}^2 \left(\left(\frac{e^\epsilon - 1}{e^\epsilon + d - 1} \right)^2 \sum_{j=1}^d \frac{\Delta_j^2}{\check{p}_j^0} \right)$

Proof. Let $\mathbf{V} = \sqrt{n} \left(\frac{\check{\mathbf{H}}/n - \check{\mathbf{p}}^0}{\check{\mathbf{p}}^0} \right)$, which converges in distribution to $\text{N} \left(\mathbf{0}, I_d - \sqrt{\check{\mathbf{p}}^0} \sqrt{\check{\mathbf{p}}^0}^\top \right)$, by the central limit theorem. We then apply Theorem 5.3 with $\boldsymbol{\mu} = \mathbf{0}$ and $I_d - \sqrt{\check{\mathbf{p}}^0} \sqrt{\check{\mathbf{p}}^0}^\top$ being a projection matrix of rank $d - 1$. In the second statement we assume the alternate $\mathbf{p} = \mathbf{p}^1$ holds, in which case \mathbf{V} converges in distribution to $\text{N} \left(\left(\frac{e^\epsilon - 1}{e^\epsilon + d - 1} \right) \frac{\Delta}{\check{\mathbf{p}}^0}, I_d - \sqrt{\check{\mathbf{p}}^0} \sqrt{\check{\mathbf{p}}^0}^\top \right)$. We then verify that $\left(\frac{e^\epsilon - 1}{e^\epsilon + d - 1} \right) \left(I_d - \sqrt{\check{\mathbf{p}}^0} \sqrt{\check{\mathbf{p}}^0}^\top \right) \frac{\Delta}{\check{\mathbf{p}}^0} = \left(\frac{e^\epsilon - 1}{e^\epsilon + d - 1} \right) \frac{\Delta}{\check{\mathbf{p}}^0}$ to prove the second statement. \square

We then base our LDP goodness of fit test on this result to obtain the correct critical value to reject the null hypothesis based on a chi-square distribution. The test is presented in Algorithm 3. The following result is immediate from the exponential mechanism being ϵ -DP and the fact that we use it as a local randomizer.

Theorem 5.5. *LocalGenRRGOF is ϵ -LDP.*

Proof. The proof follows from the fact that we use $\mathcal{M}_{\text{GenRR}}$ for each individual's data and then LocalGenRRGOF aggregates the privatized data, which is just a post-processing function. \square

Algorithm 3 Local DP GOF Test: LocalGenRRGOF

Input: $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$, ϵ , α , $H_0 : \mathbf{p} = \mathbf{p}^0$.

Let $\tilde{\mathbf{p}}^0 = \frac{1}{e^\epsilon + d - 1} (e^\epsilon \mathbf{p}^0 + (1 - \mathbf{p}^0))$.

Let $\tilde{\mathbf{H}} = \sum_{i=1}^n \mathcal{M}_{\text{GenRR}}(\mathbf{x}_i, \epsilon)$.

Set $q = \sum_{j=1}^d \frac{(\tilde{h}_j - n\tilde{p}_j^0)^2}{n\tilde{p}_j^0}$

if $q > \chi_{d-1, 1-\alpha}^2$ Decision \leftarrow Reject.

else Decision \leftarrow Fail to Reject.

Output: Decision

5.3 Goodness of Fit Test with Bit Flipping

Note that the noncentral parameter in Theorem 5.4 goes to zero as d grows large due to the coefficient being $\left(\frac{e^\epsilon - 1}{e^\epsilon + d - 1}\right)^2$. Thus, for large dimensional data the generalized randomized response cannot reject a false null hypothesis. We next consider another differentially private algorithm $\mathcal{M} : \{\mathbf{e}_1, \dots, \mathbf{e}_d\} \rightarrow \{0, 1\}^d$, given in Algorithm 4 used in Bassily and Smith (2015) that flips each bit with some biased probability.

Algorithm 4 Bit Flip Local Randomizer: \mathcal{M}_{bit}

Input: $\mathbf{x} \in \{\mathbf{e}_1, \dots, \mathbf{e}_d\}$, ϵ .

for $j \in [d]$ **do**

Set $z_j = x_j$ with probability $\frac{e^{\epsilon/2}}{e^{\epsilon/2} + 1}$, otherwise $z_j = (1 - x_j)$.

Output: \mathbf{z}

Theorem 5.6. *The algorithm \mathcal{M}_{bit} is ϵ -DP.*

We then want to form a statistic based on the output $\mathbf{z} \in \{0, 1\}^d$ that is asymptotically distributed as a chi-square under the null hypothesis. We defer the proof to the appendix.

Lemma 5.7. *We consider data $\mathbf{X}_i \sim \text{Multinomial}(1, \mathbf{p})$ for each $i \in [n]$. We define the following covariance matrix $\Sigma(\mathbf{p})$ and mean vector $\tilde{\mathbf{p}} = \frac{[(e^{\epsilon/2} - 1)\mathbf{p} + 1]}{e^{\epsilon/2} + 1}$, in terms of $\alpha_\epsilon = \left(\frac{e^{\epsilon/2} - 1}{e^{\epsilon/2} + 1}\right)$*

$$\Sigma(\mathbf{p}) = \alpha_\epsilon^2 [\text{Diag}(\mathbf{p}) - \mathbf{p}(\mathbf{p})^\top] + \frac{e^{\epsilon/2}}{(e^{\epsilon/2} + 1)^2} I_d \quad (6)$$

The histogram $\tilde{\mathbf{H}} = \sum_{i=1}^n \mathcal{M}_{\text{bit}}(\mathbf{X}_i)$ has the following asymptotic distribution $\sqrt{n} \left(\frac{\tilde{\mathbf{H}}}{n} - \tilde{\mathbf{p}}\right) \xrightarrow{D} \mathbf{N}(\mathbf{0}, \Sigma(\mathbf{p}))$. Further, $\Sigma(\mathbf{p})$ is invertible for any $\epsilon > 0$ and $\mathbf{p} > \mathbf{0}$.

Following a similar analysis in Kifer and Rogers (2017) and using Theorem 5.3, we can form the following statistic for null hypothesis $H_0 : \mathbf{p} = \mathbf{p}^0$ in terms of the histogram $\tilde{\mathbf{H}}$ and projection matrix $\Pi = I_d - \frac{1}{d} \mathbf{1}\mathbf{1}^\top$, as well as the covariance matrix $\Sigma(\mathbf{p}^0)$ and mean vector $\tilde{\mathbf{p}}^0$ both given in (6) where we replace \mathbf{p} with \mathbf{p}^0 :

$$\mathbb{T}_{\text{BitFlip}}^{(n)}(\epsilon) = n \left(\frac{\tilde{\mathbf{H}}}{n} - \tilde{\mathbf{p}}^0\right)^\top \Pi \Sigma(\mathbf{p}^0)^{-1} \Pi \left(\frac{\tilde{\mathbf{H}}}{n} - \tilde{\mathbf{p}}^0\right) \quad (7)$$

We can then design a hypothesis test based on the outputs from \mathcal{M}_{bit} in Algorithm 5

Theorem 5.8. *LocalBitFlipGOF is ϵ -LDP.*

We now show that the statistic in (7) is asymptotically distributed as χ_{d-1}^2 .

Theorem 5.9. *If the null hypothesis $H_0 : \mathbf{p} = \mathbf{p}^0$ holds, then the statistic $\mathbb{T}_{\text{BitFlip}}^{(n)}(\epsilon)$ is asymptotically distributed as a chi-square, i.e. $\mathbb{T}_{\text{BitFlip}}^{(n)}(\epsilon) \xrightarrow{D} \chi_{d-1}^2$. Further, if we consider the alternate $H_1 : \mathbf{p} = \mathbf{p}^1$ then*

Algorithm 5 Local DP GOF Test: LocalBitFlipGOF

Input: $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$, ϵ , α , $H_0 : \mathbf{p} = \mathbf{p}^0$.

Let $\tilde{\mathbf{H}} = \sum_{i=1}^n \mathcal{M}_{\text{bit}}(\mathbf{x}_i, \epsilon)$.

Set $q = T_{\text{BitFlip}}^{(n)}(\epsilon)$

if $q > \chi_{d-1, 1-\alpha}^2$ Decision \leftarrow Reject.

else Decision \leftarrow Fail to Reject.

Output: Decision

$$T_{\text{BitFlip}}^{(n)}(\epsilon) \xrightarrow{D} \chi_{d-1}^2 \left(\left(\frac{\epsilon^{\epsilon/2-1}}{\epsilon^{\epsilon/2+1}} \right)^2 \cdot \mathbf{\Delta}^\top \Sigma(\mathbf{p}^0)^{-1} \mathbf{\Delta} \right).$$

Proof. From Lemma 5.7, we know that $\sqrt{n} \left(\frac{\tilde{\mathbf{H}}}{n} - \tilde{\mathbf{p}}^0 \right) \xrightarrow{D} N(\mathbf{0}, \Sigma(\tilde{\mathbf{p}}^0))$. Since $\Sigma(\tilde{\mathbf{p}}^0)$ is full rank and has an eigenvector that is all ones. Thus, we can diagonalize the covariance matrix as $\Sigma(\tilde{\mathbf{p}}^0) = BDB^\top$ where D is a diagonal matrix and B has orthogonal columns with one of them being $\frac{1}{d}\mathbf{1}$. Thus, we have

$$\sqrt{n} \Sigma(\tilde{\mathbf{p}}^0)^{-1/2} \Pi \left(\frac{\tilde{\mathbf{H}}}{n} - \tilde{\mathbf{p}}^0 \right) \xrightarrow{D} N(\mathbf{0}, \Lambda)$$

where $\Lambda = \Sigma(\tilde{\mathbf{p}}^0)^{-1/2} \Pi BDB^\top \Pi \Sigma(\tilde{\mathbf{p}}^0)^{-1/2}$. Hence, Λ will be the identity matrix with a single zero on the diagonal. We then apply Theorem 5.3 to prove the first statement.

To prove the second statement we assume that $\mathbf{p} = \mathbf{p}^1$ holds, which gives $\sqrt{n} \left(\frac{\tilde{\mathbf{H}}}{n} - \tilde{\mathbf{p}}^0 \right) \xrightarrow{D} N(\alpha_\epsilon \cdot \mathbf{\Delta}, \Sigma(\tilde{\mathbf{p}}^0))$. Thus,

$$\sqrt{n} \Sigma(\tilde{\mathbf{p}}^0)^{-1/2} \Pi \left(\frac{\tilde{\mathbf{H}}}{n} - \tilde{\mathbf{p}}^0 \right) \xrightarrow{D} N(\alpha_\epsilon \cdot \Sigma(\tilde{\mathbf{p}}^0)^{-1/2} \mathbf{\Delta}, \Lambda)$$

We again use Theorem 5.3 to obtain the noncentral parameter $\alpha_\epsilon^2 \cdot \mathbf{\Delta}^\top \Sigma(\tilde{\mathbf{p}}^0)^{-1} \mathbf{\Delta}$. □

5.4 Comparison of Noncentral Parameters

We now compare the noncentral parameters of the three local private tests we presented in Algorithms 1, 3 and 5. We consider the null hypothesis $\mathbf{p}^0 = (1/d, \dots, 1/d)$ for $d > 2$, and alternate $H_1 : \mathbf{p} = \mathbf{p}^0 + \frac{\mathbf{\Delta}}{\sqrt{n}}$. In this case, we can easily compare the various noncentral parameters for various privacy parameters and dimensions d . In Figure 1 we give the coefficient to the term $\mathbf{\Delta}^\top \mathbf{\Delta}$ in the noncentral parameter of the asymptotic distribution for each local private test presented thus far. The larger this coefficient is, the better the power will be for any alternate $\mathbf{\Delta}$ vector. Note that in `LocalNoiseGOF`, we set $\rho = \epsilon^2/8$ which makes the variance the same as for a random variable distributed as $\text{Lap}(2/\epsilon)$ for an ϵ -DP guarantee – recall that `LocalNoiseGOF` with Gaussian noise does not satisfy ϵ -DP for any $\epsilon > 0$. We give results for $\epsilon \in \{1, 2, 3, 4\}$ which are all in the range of privacy parameters that have been considered in actual locally differentially private algorithms used in practice.¹ From the plots, we see how `LocalGenRRGOF` may outperform `LocalBitFlipGOF` depending on the privacy parameter and dimension of the data. We can use these plots to determine which test to use given ϵ and dimension of data d . When H_0 is not uniform, we can use the noncentral parameters given for each test to determine which test has the largest noncentral parameter for a particular privacy budget ϵ .

¹In Erlingsson et al. (2014), we know that Google uses $\epsilon = \ln(3) \approx 1.1$ in RAPPOR and from Aleksandra Korolova’s Twitter post on September 13, 2016 <https://twitter.com/korolova/status/775801259504734208> we know that Apple uses $\epsilon = 1, 4$.




Figure 1: The coefficient on $\Delta^T \Delta$ in the noncentral parameter for the local private tests where "LocalGaussGOF" is LocalNoiseGOF with Gaussian noise, LocalGenRRGOF, and LocalBitFlipGOF for various dimensions d and ϵ .

5.5 Empirical Results

We then empirically compare the power between LocalNoiseGOF with Laplace noise in Algorithm 1, LocalGenRRGOF in Algorithm 3, and LocalBitFlipGOF in Algorithm 5. Recall that all three of these tests have the same privacy benchmark of local differential privacy. For LocalNoiseGOF with Laplace noise, we will use $m = 999$ samples in our Monte Carlo simulations. In our experiments we fix $\alpha = 0.05$ and $\epsilon \in \{1, 2, 4\}$. We then consider null hypotheses of the form $\mathbf{p}^0 = (1/d, 1/d, \dots, 1/d)$ and alternate $H_1 : \mathbf{p} = \mathbf{p}^0 + \eta(1, -1, \dots, 1, -1)$ for some $\eta > 0$. In Figure 2, we plot the number of times our tests correctly rejects the null hypothesis in 1000 independent trials for various sample sizes n and privacy parameters ϵ .




Figure 2: Comparison of empirical power among the classical non-private test and the local private tests LocalNoiseGOF with Laplace noise (solid line), LocalGenRRGOF (dashed line), and LocalBitFlipGOF (dotted line); in the left plot, $d = 4$ and $\eta = 0.01$, in the right plot, $d = 40$ and $\eta = 0.005$.

From Figure 2, we can see that the test statistics that have the largest noncentral parameter for a particular dimension d and privacy parameter ϵ will have the best empirical power. When $d = 4$, we see that LocalGenRRGOF performs the best. However, for $d = 40$ it is not so clear cut. When $\epsilon = 4$, we can see that LocalGenRRGOF does the best, but then when $\epsilon = 2$, LocalBitFlipGOF does best. Thus, the best Local DP Goodness of Fit test depends on the noncentral parameter, which is a function of ϵ , the null hypothesis \mathbf{p}^0 , and alternate $\mathbf{p} = \mathbf{p}^0 + \Delta$. Note that the worst local DP test also depends on the privacy parameter and the dimension d . Based on our empirical results, we see that no single locally private test is best for all data dimensions. Knowing the corresponding noncentral parameter for a given problem is useful in determining which tests to use, where the larger the noncentral parameter is the higher the power will be.

6 Local Private Independence Tests

We now show that our techniques can be extended to include *composite* hypothesis tests, where we test whether the data comes from a whole family of probability distributions. We will focus on independence testing, but much of the theory can be extended to general chi-square tests. We will closely follow the presentation and notation as in (Kifer and Rogers, 2017).

We consider two multinomial random variables $\{\mathbf{U}_\ell\}_{\ell=1}^n \stackrel{i.i.d.}{\sim} \text{Multinomial}(1, \boldsymbol{\pi}^{(1)})$ for $\boldsymbol{\pi}^{(1)} \in \mathbb{R}^r$, $\{\mathbf{V}_\ell\}_{\ell=1}^n \stackrel{i.i.d.}{\sim} \text{Multinomial}(1, \boldsymbol{\pi}^{(2)})$ for $\boldsymbol{\pi}^{(2)} \in \mathbb{R}^c$ and no component of $\boldsymbol{\pi}^{(1)}$ or $\boldsymbol{\pi}^{(2)}$ is zero. Without loss of generality, we will consider an individual to be in one of r groups who reports a data record that is in one of c categories. The collected data consists of n joint outcomes \mathbf{H} whose (i, j) th coordinate is $H_{i,j} = \sum_{\ell=1}^n \mathbb{1}\{U_{\ell,i} = 1 \ \& \ V_{\ell,j} = 1\}$. Note that \mathbf{H} is then the contingency table over the joint outcomes.

Under the null hypothesis of independence between $\{\mathbf{U}_\ell\}_{\ell=1}^n$ and $\{\mathbf{V}_\ell\}_{\ell=1}^n$, we have

$$\mathbf{H} \sim \text{Multinomial}\left(n, \mathbf{p}(\boldsymbol{\pi}^{(1)}, \boldsymbol{\pi}^{(2)})\right) \quad \text{where } \mathbf{p}(\boldsymbol{\pi}^{(1)}, \boldsymbol{\pi}^{(2)}) = \boldsymbol{\pi}^{(1)} \left(\boldsymbol{\pi}^{(2)}\right)^\top \quad (8)$$

What makes this test difficult is that the analyst does not know the data distribution $\mathbf{p}(\boldsymbol{\pi}^{(1)}, \boldsymbol{\pi}^{(2)})$ and so cannot simply plug it into the chi-square statistic. Rather, we use the data to estimate the *best guess* for the unknown probability distribution that satisfies the null hypothesis.

Note that without privacy, each individual $\ell \in [n]$ is reporting a $r \times c$ matrix \mathbf{X}_ℓ which would be 1 in exactly one location. Thus we can alternatively write the contingency table as $\mathbf{H} = \sum_{\ell=1}^n \mathbf{X}_\ell$. We then use the three local private algorithms we presented earlier for noise addition, generalized randomized response, and bit flipping to see how we can form a private chi-square statistic for independence testing. We want to be able to ensure the privacy of both the group and the category that each individual belongs to.

For completeness, we will present the minimum chi-square asymptotic theory from Ferguson (1996) that was used in Kifer and Rogers (2017). We will use this theory to design general chi-square test statistics for local differentially private hypothesis tests. Consider a d dimensional random vector $\mathbf{V}^{(n)}$ and a $k \leq d$ dimensional parameter space Θ , which is an open subset of \mathbb{R}^k . We want to measure the distance between the random variable $\mathbf{V}^{(n)}$ and some model $A(\boldsymbol{\theta})$ that maps parameters in Θ to \mathbb{R}^d . We then use the following quadratic form where $M(\boldsymbol{\theta}) \in \mathbb{R}^{d \times d}$ is a positive-semidefinite matrix.

$$n \left(\mathbf{V}^{(n)} - A(\boldsymbol{\theta})\right)^\top M(\boldsymbol{\theta}) \left(\mathbf{V}^{(n)} - A(\boldsymbol{\theta})\right) \quad (9)$$

We want to ensure the following assumptions hold.

Assumption 6.1. *For all $\boldsymbol{\theta} \in \Theta$, we have*

- *There exists a parameter $\boldsymbol{\theta}^* \in \Theta$ such that for some model $A(\cdot)$ and covariance matrix $C(\cdot)$*

$$\sqrt{n} \left(\mathbf{V}^{(n)} - A(\boldsymbol{\theta}^*)\right) \xrightarrow{D} \mathbf{N}(\mathbf{0}, C(\boldsymbol{\theta}^*))$$

- *$A(\boldsymbol{\theta})$ is bicontinuous.*
- *$A(\boldsymbol{\theta})$ has continuous first partial derivatives, denoted as $\dot{A}(\boldsymbol{\theta})$, with full rank k .*
- *The matrix $M(\boldsymbol{\theta})$ is continuous in $\boldsymbol{\theta}$ and there exists an $\eta > 0$ such that $M(\boldsymbol{\theta}) - \eta I_d$ is positive definite in an open neighborhood of $\boldsymbol{\theta}^*$.*

We will then use the following useful result from Kifer and Rogers (2017)[Theorem 4.2].

Theorem 6.2 [Kifer and Rogers (2017)]. *Let ν be the rank of $C(\boldsymbol{\theta}^*)$. If Assumption 6.1 holds, and for all $\boldsymbol{\theta} \in \Theta$ we have*

$$C(\boldsymbol{\theta})M(\boldsymbol{\theta})C(\boldsymbol{\theta}) = C(\boldsymbol{\theta})$$

$$C(\boldsymbol{\theta})M(\boldsymbol{\theta})\dot{A}(\boldsymbol{\theta}) = \dot{A}(\boldsymbol{\theta})$$

Then for any estimate $\phi(\mathbf{V}^{(n)}) \xrightarrow{P} \boldsymbol{\theta}^*$ we have

$$\min_{\boldsymbol{\theta} \in \Theta} \left\{ n \left(\mathbf{V}^{(n)} - A(\boldsymbol{\theta}) \right)^\top M(\phi(\mathbf{V}^{(n)})) \left(\mathbf{V}^{(n)} - A(\boldsymbol{\theta}) \right) \right\} \xrightarrow{D} \chi_{\nu-k}^2.$$

To ensure the assumptions hold, we will assume that each component of $\boldsymbol{\pi}^{(1)}, \boldsymbol{\pi}^{(2)}$ is positive, i.e. $\pi^{(i)} > 0$ for $i = 1, 2$.

6.1 Independence Test with Noise Addition

For noise addition, we will have the following contingency table $\mathbf{H} + \mathbf{Z}$ (treating \mathbf{H} as an $r \times c$ vector) where $\mathbf{Z} = \mathbf{N}\left(\mathbf{0}, \frac{n}{\rho}\right)$ for ρ -LzCDP or \mathbf{Z} is the sum of n independent Laplace random vectors with scale parameter $2/\epsilon$ in each coordinate for ϵ -LDP, although we will only consider the Gaussian noise case. We can then use the same statistic presented in Kifer and Rogers (2017) with a rescaling of the variance in the noise and flattening the contingency table as a vector. For convenience, we write the marginals as $H_{i,\cdot} = \sum_{j=1}^c H_{i,j}$ and similarly for $H_{\cdot,j}$. Further, we will write $\hat{n} = n + \sum_{i,j} Z_{i,j}$ in the following test statistic,²

$$\begin{aligned} \hat{\boldsymbol{\pi}}^{(1)} &= \left(\frac{1}{\hat{n}} (H_{i,\cdot} + Z_{i,\cdot}) : i \in [r] \right), & \hat{\boldsymbol{\pi}}^{(2)} &= \left(\frac{1}{\hat{n}} (H_{\cdot,j} + Z_{\cdot,j}) : j \in [c] \right) \\ \widehat{M} &= \text{Diag} \left(\mathbf{p}(\hat{\boldsymbol{\pi}}^{(1)}, \hat{\boldsymbol{\pi}}^{(2)}) - \mathbf{p}(\hat{\boldsymbol{\pi}}^{(1)}, \hat{\boldsymbol{\pi}}^{(2)}) \left(\mathbf{p}(\hat{\boldsymbol{\pi}}^{(1)}, \hat{\boldsymbol{\pi}}^{(2)}) \right)^\top + \frac{1}{\rho} I_{rc} \right) \\ \widehat{\mathbf{T}}_{KR}^{(n)} \left(\rho/n; (\boldsymbol{\pi}^{(1)}, \boldsymbol{\pi}^{(2)}) \right) &= \frac{1}{n} \left(\mathbf{H} + \mathbf{Z} - n\mathbf{p} \left(\boldsymbol{\pi}^{(1)}, \boldsymbol{\pi}^{(2)} \right) \right)^\top \Pi \left(\widehat{M} \right)^{-1} \Pi \left(\mathbf{H} + \mathbf{Z} - n\mathbf{p} \left(\boldsymbol{\pi}^{(1)}, \boldsymbol{\pi}^{(2)} \right) \right) \end{aligned} \quad (10)$$

We use this statistic because it is asymptotically distributed as a chi-square random variable given the null hypothesis holds, which follows directly from the general chi-square theory presented in Kifer and Rogers (2017).

Theorem 6.3. *Under the null hypothesis that \mathbf{U} and \mathbf{V} are independent, then we have as $n \rightarrow \infty$*

$$\min_{\boldsymbol{\pi}^{(1)}, \boldsymbol{\pi}^{(2)}} \left\{ \widehat{\mathbf{T}}_{KR}^{(n)} \left(\rho/n; (\boldsymbol{\pi}^{(1)}, \boldsymbol{\pi}^{(2)}) \right) \right\} \xrightarrow{D} \chi_{(r-1)(c-1)}^2.$$

We present the test in Algorithm 6 for Gaussian noise which uses the statistic in (10).

Algorithm 6 LzCDP Independence Test with Gaussian Mechanism: LocalNoiseIND

Input: $(\mathbf{x}_1, \dots, \mathbf{x}_n)$, ρ , α .

Let $\mathbf{H} = \sum_{\ell=1}^n \mathbf{x}_\ell$

Set $q = \min_{\boldsymbol{\pi}^{(1)}, \boldsymbol{\pi}^{(2)}} \left\{ \widehat{\mathbf{T}}_{KR}^{(n)} \left(\rho/n; (\boldsymbol{\pi}^{(1)}, \boldsymbol{\pi}^{(2)}) \right) \right\}$ given in (10).

if $q > \chi_{(r-1)(c-1), 1-\alpha}^2$ **then**
Decision \leftarrow Reject.

else

Decision \leftarrow Fail to Reject.

Output: Decision

We then have the following result which follows from the privacy analysis from before.

Theorem 6.4. *LocalNoiseIND is ρ -LzCDP.*

²We will follow a common rule of thumb for small sample sizes, so that if $n\hat{\boldsymbol{\pi}}^{(1)} \left(\hat{\boldsymbol{\pi}}^{(2)} \right)^\top \leq 5$ for any cell, then we simply fail to reject H_0

6.2 Independence Test with Generalized Randomized Response

Next we want to design an independence test when the data is generated from $\mathcal{M}_{\text{GenRR}}$ given in Algorithm 2. In this case our contingency table can be written as $\check{\mathbf{H}} \sim \text{Multinomial}(n, \check{\mathbf{p}}(\boldsymbol{\pi}^{(1)}, \boldsymbol{\pi}^{(2)}))$ where $\beta_\epsilon = \frac{1}{e^\epsilon + r c - 1}$ and we use (4) to get

$$\check{\mathbf{p}}(\boldsymbol{\pi}^{(1)}, \boldsymbol{\pi}^{(2)}) = \beta_\epsilon \left((e^\epsilon - 1) \boldsymbol{\pi}^{(1)} \left(\boldsymbol{\pi}^{(2)} \right)^\top + \mathbf{1} \right) \quad (11)$$

We then obtain an estimate for the unknown parameters,

$$\begin{aligned} \check{\boldsymbol{\pi}}^{(1)} &= \frac{1}{\beta_\epsilon (e^\epsilon - 1)} \left(\frac{\check{H}_{i,\cdot}}{n} - c\beta_\epsilon : i \in [r] \right), \\ \check{\boldsymbol{\pi}}^{(2)} &= \frac{1}{\beta_\epsilon (e^\epsilon - 1)} \left(\frac{\check{H}_{\cdot,j}}{n} - r\beta_\epsilon : j \in [c] \right) \\ \check{\mathbf{T}}_{\text{GenRR}}^{(n)}(\epsilon) &= \sum_{i,j} \frac{\left(\check{H}_{i,j} - n\check{p}_{i,j}(\check{\boldsymbol{\pi}}^{(1)}, \check{\boldsymbol{\pi}}^{(2)}) \right)^2}{n\check{p}_{i,j}(\check{\boldsymbol{\pi}}^{(1)}, \check{\boldsymbol{\pi}}^{(2)})} \end{aligned} \quad (12)$$

We can then prove the following result, which uses Theorem 6.2.

Theorem 6.5. *Assuming \mathbf{U} and \mathbf{V} are independent with true probability vectors $\boldsymbol{\pi}^{(1)}, \boldsymbol{\pi}^{(2)} > 0$ respectively, then as $n \rightarrow \infty$ we have $\check{\mathbf{T}}_{\text{GenRR}}^{(n)}(\epsilon) \xrightarrow{D} \chi_{(r-1)(c-1)}^2$.*

Proof. Note that $\check{\boldsymbol{\pi}}^{(1)}, \check{\boldsymbol{\pi}}^{(2)}$ both converge in probability to the true parameters $\boldsymbol{\pi}^{(1)}, \boldsymbol{\pi}^{(2)}$, respectively due to $\mathbb{E}[\check{\mathbf{H}}/n] = \check{\mathbf{p}}(\boldsymbol{\pi}^{(1)}, \boldsymbol{\pi}^{(2)})$ and $\mathbb{V}[\check{\mathbf{H}}/n] \rightarrow 0$. We then form the following statistic in terms of parameters $\boldsymbol{\theta}^{(1)}$ in the r dimensional simplex with all positive components and $\boldsymbol{\theta}^{(2)}$ in the c dimensional simplex with all positive components,

$$\begin{aligned} &\check{\mathbf{T}}_{\text{GenRR}}^{(n)}(\boldsymbol{\theta}^{(1)}, \boldsymbol{\theta}^{(2)}; \epsilon) \\ &= n \left(\check{\mathbf{H}}/n - \check{\mathbf{p}}(\boldsymbol{\theta}^{(1)}, \boldsymbol{\theta}^{(2)}) \right)^\top \text{Diag} \left(\check{\mathbf{p}}(\check{\boldsymbol{\pi}}^{(1)}, \check{\boldsymbol{\pi}}^{(2)}) \right)^{-1} \left(\check{\mathbf{H}}/n - \check{\mathbf{p}}(\boldsymbol{\theta}^{(1)}, \boldsymbol{\theta}^{(2)}) \right) \end{aligned}$$

We then have the following calculation,

$$\left(\check{\boldsymbol{\pi}}^{(1)}, \check{\boldsymbol{\pi}}^{(2)} \right) = \underset{\boldsymbol{\theta}^{(1)}, \boldsymbol{\theta}^{(2)}}{\text{argmin}} \left\{ \check{\mathbf{T}}_{\text{GenRR}}^{(n)}(\boldsymbol{\theta}^{(1)}, \boldsymbol{\theta}^{(2)}; \epsilon) \right\}.$$

Note that $\sqrt{n} \left(\check{\mathbf{H}}/n - \check{\mathbf{p}}(\boldsymbol{\pi}^{(1)}, \boldsymbol{\pi}^{(2)}) \right) \xrightarrow{D} \mathbf{N}(\mathbf{0}, C(\boldsymbol{\theta}))$ where $C(\boldsymbol{\theta}) = \text{Diag}(\check{\mathbf{p}}(\boldsymbol{\theta}^{(1)}, \boldsymbol{\theta}^{(2)})) - \check{\mathbf{p}}(\boldsymbol{\theta}^{(1)}, \boldsymbol{\theta}^{(2)}) \left(\check{\mathbf{p}}(\boldsymbol{\theta}^{(1)}, \boldsymbol{\theta}^{(2)}) \right)^\top$ is the covariance matrix for the multinomial random variable.

We then apply Theorem 6.2 to prove the statement, so we verify that the following equalities hold,

$$C(\boldsymbol{\theta}) \cdot \text{Diag} \left(\check{\mathbf{p}}(\boldsymbol{\theta}^{(1)}, \boldsymbol{\theta}^{(2)}) \right)^{-1} \cdot C(\boldsymbol{\theta}) = C(\boldsymbol{\theta})$$

and

$$\begin{aligned} &C(\boldsymbol{\theta}) \cdot \text{Diag} \left(\check{\mathbf{p}}(\boldsymbol{\theta}^{(1)}, \boldsymbol{\theta}^{(2)}) \right)^{-1} \nabla \left(\boldsymbol{\theta}^{(1)} \left(\boldsymbol{\theta}^{(2)} \right)^\top \right) \\ &= \left(I_{rc} - \check{\mathbf{p}}(\boldsymbol{\theta}^{(1)}, \boldsymbol{\theta}^{(2)}) \mathbf{1}^\top \right) \nabla \left(\boldsymbol{\theta}^{(1)} \left(\boldsymbol{\theta}^{(2)} \right)^\top \right) \\ &= \nabla \left(\boldsymbol{\theta}^{(1)} \left(\boldsymbol{\theta}^{(2)} \right)^\top \right) \end{aligned}$$

where the last equality follows from the last coordinates of each probability vector can be written as $\theta_r^{(1)} = 1 - \sum_{i=1}^{r-1} \theta_i$ and $\theta_c^{(2)} = 1 - \sum_{j=1}^{c-1} \theta_j$. \square

We then use this result to design our private chi-square test for independence.

Algorithm 7 Local DP IND Test: LocalGenRRIND

Input: $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$, ϵ , α , $H_0 : \mathbf{p} = \mathbf{p}^0$.

Let $\tilde{\mathbf{H}} = \sum_{i=1}^n \mathcal{M}_{\text{GenRR}}(\mathbf{x}_i, \epsilon)$.

Set $q = \tilde{\mathbf{T}}_{\text{GenRR}}^{(n)}(\epsilon)$ from (12)

if $q > \chi_{d-1, 1-\alpha}^2$, Decision \leftarrow Reject.

else Decision \leftarrow Fail to Reject.

Output: Decision

We then have the following result which again follows from the privacy analysis from before.

Theorem 6.6. *LocalGenRRIND is ϵ -LDP.*

6.3 Independence Test with Bit Flipping

Lastly, we design an independence test when the data is reported via \mathcal{M}_{bit} in Algorithm 4. Assuming that $\mathbf{H} = \sum_{\ell=1}^n \mathbf{X}_\ell \sim \text{Multinomial}(n, \mathbf{p}(\boldsymbol{\pi}^{(1)}, \boldsymbol{\pi}^{(2)}))$, then we know that replacing \mathbf{p}^0 with $\mathbf{p}(\boldsymbol{\pi}^{(1)}, \boldsymbol{\pi}^{(2)})$ in Section 5.3 gives us the following asymptotic distribution (treating the contingency table of values as a vector) with covariance matrix $\Sigma(\cdot)$ given in (6)

$$\sqrt{n} \left(\frac{\tilde{\mathbf{H}}}{n} - \underbrace{\left[\frac{e^{\epsilon/2} - 1}{e^{\epsilon/2} + 1} \boldsymbol{\pi}^{(1)} \left(\boldsymbol{\pi}^{(2)} \right)^\top + \frac{1}{e^{\epsilon/2} + 1} \right]}_{\tilde{\mathbf{p}}(\boldsymbol{\pi}^{(1)}, \boldsymbol{\pi}^{(2)})} \right) \stackrel{D}{\rightarrow} \mathbf{N}(\mathbf{0}, \Sigma(\boldsymbol{\pi}^{(1)} \left(\boldsymbol{\pi}^{(2)} \right)^\top)) \quad (13)$$

Similar to analysis for Theorem 6.5, we start with a rough estimate for the unknown parameters which converges in probability to the true estimates, so we again use $\alpha_\epsilon = \left(\frac{e^{\epsilon/2} - 1}{e^{\epsilon/2} + 1} \right)$ to get

$$\begin{aligned} \tilde{\boldsymbol{\pi}}^{(1)} &= \left(\frac{1}{\alpha_\epsilon} \right) \left(\frac{\tilde{H}_{i,\cdot}}{n} - \frac{c}{e^{\epsilon/2} + 1} : i \in [r] \right) \\ \tilde{\boldsymbol{\pi}}^{(2)} &= \left(\frac{1}{\alpha_\epsilon} \right) \left(\frac{\tilde{H}_{\cdot,j}}{n} - \frac{r}{e^{\epsilon/2} + 1} : j \in [c] \right) \end{aligned} \quad (14)$$

We then give the resulting statistic, parameterized by the unknown parameters $\boldsymbol{\pi}^{(\ell)}$, for $\ell \in \{1, 2\}$. For middle matrix $\tilde{M} = \Pi \Sigma(\tilde{\boldsymbol{\pi}}^{(1)} \left(\tilde{\boldsymbol{\pi}}^{(2)} \right)^\top)^{-1} \Pi$, we have

$$\begin{aligned} \tilde{\mathbf{T}}_{\text{BitFlip}}^{(n)}(\boldsymbol{\theta}^{(1)}, \boldsymbol{\theta}^{(2)}; \epsilon) &= \frac{1}{n} \left(\tilde{\mathbf{H}} - n\tilde{\mathbf{p}}(\boldsymbol{\theta}^{(1)}, \boldsymbol{\theta}^{(2)}) \right)^\top \\ \tilde{M} \left(\tilde{\mathbf{H}} - n\tilde{\mathbf{p}}(\boldsymbol{\theta}^{(1)}, \boldsymbol{\theta}^{(2)}) \right) \end{aligned} \quad (15)$$

Theorem 6.7. *Under the null hypothesis that \mathbf{U} and \mathbf{V} are independent with true probability vectors $\boldsymbol{\pi}^{(1)}, \boldsymbol{\pi}^{(2)} > 0$ respectively, then we have as $n \rightarrow \infty$, $\min_{\boldsymbol{\theta}^{(1)}, \boldsymbol{\theta}^{(2)}} \left\{ \tilde{\mathbf{T}}_{\text{BitFlip}}^{(n)}(\boldsymbol{\theta}^{(1)}, \boldsymbol{\theta}^{(2)}; \epsilon) \right\} \stackrel{D}{\rightarrow} \chi_{(r-1)(c-1)}^2$.*

Proof. As in the proof of Theorem 6.5, we will use Theorem 6.2, so we first need to verify the following equality holds, where $\mathbf{p}(\boldsymbol{\theta}) = \boldsymbol{\theta}^{(1)} (\boldsymbol{\theta}^{(2)})^\top$

$$\Pi \Sigma(\mathbf{p}(\boldsymbol{\theta})) \Pi \cdot \Sigma(\mathbf{p}(\boldsymbol{\theta}))^{-1} \cdot \Pi \Sigma(\mathbf{p}(\boldsymbol{\theta})) \Pi = \Pi \Sigma(\mathbf{p}(\boldsymbol{\theta})) \Pi.$$

Recall that $\Sigma(\mathbf{p}(\boldsymbol{\theta}))$ has an eigenvector of all 1's and the other eigenvectors are orthogonal to it. Hence, when we diagonalize $\Sigma(\mathbf{p}(\boldsymbol{\theta})) = BDB^\top$, then ΠB is going to be the same as B except the column of all 1's becomes zero, which we will denote as B^* . Hence, $\Sigma(\mathbf{p}(\boldsymbol{\theta})) \Pi \Sigma(\mathbf{p}(\boldsymbol{\theta}))^{-1} \Pi \Sigma(\mathbf{p}(\boldsymbol{\theta}))$, becomes $B D^* B^\top$ where D^* is the same as D except the eigenvalue on the diagonal for the eigenvector of all 1's becomes zero. Then projecting the resulting matrix with Π will not change the resulting product. We then have

$$\begin{aligned} & \Pi \Sigma(\mathbf{p}(\boldsymbol{\theta})) \Pi \cdot \Sigma(\mathbf{p}(\boldsymbol{\theta}))^{-1} \cdot \Pi \Sigma(\mathbf{p}(\boldsymbol{\theta})) \Pi \\ &= \Pi B D^* B^\top \Pi = \Pi \Sigma(\mathbf{p}(\boldsymbol{\theta})) \Pi. \end{aligned}$$

We then need to show that the following equality holds

$$\begin{aligned} & \Pi \Sigma(\mathbf{p}(\boldsymbol{\theta})) \Pi \cdot \Sigma(\mathbf{p}(\boldsymbol{\theta}))^{-1} \Pi \nabla \tilde{\mathbf{p}}(\boldsymbol{\theta}^{(1)}, \boldsymbol{\theta}^{(2)}) \\ &= \Pi \nabla \tilde{\mathbf{p}}(\boldsymbol{\theta}^{(1)}, \boldsymbol{\theta}^{(2)}) \end{aligned}$$

We again use the fact that $\Sigma(\mathbf{p}(\boldsymbol{\theta}))$ has an eigenvector that is all 1's. Hence, we have

$$\Pi \Sigma(\mathbf{p}(\boldsymbol{\theta})) \Pi = \Sigma(\mathbf{p}(\boldsymbol{\theta})) - \frac{e^{\epsilon/2}}{rc \cdot (e^{\epsilon/2} + 1)^2} \mathbf{1}\mathbf{1}^\top.$$

We then have,

$$\begin{aligned} & \Pi \Sigma(\mathbf{p}(\boldsymbol{\theta})) \Pi \cdot \Sigma(\mathbf{p}(\boldsymbol{\theta}))^{-1} \Pi \nabla \tilde{\mathbf{p}}(\boldsymbol{\theta}^{(1)}, \boldsymbol{\theta}^{(2)}) \\ &= \left(\Sigma(\mathbf{p}(\boldsymbol{\theta})) - \frac{e^{\epsilon/2}}{rc \cdot (e^{\epsilon/2} + 1)^2} \mathbf{1}\mathbf{1}^\top \right) \\ & \quad \cdot \Sigma(\mathbf{p}(\boldsymbol{\theta}))^{-1} \Pi \nabla \tilde{\mathbf{p}}(\boldsymbol{\theta}^{(1)}, \boldsymbol{\theta}^{(2)}) \\ &= \left(I_{rc} - \frac{1}{rc} \mathbf{1}\mathbf{1}^\top \right) \Pi \nabla \tilde{\mathbf{p}}(\boldsymbol{\theta}^{(1)}, \boldsymbol{\theta}^{(2)}) \\ &= \Pi \Pi \nabla \tilde{\mathbf{p}}(\boldsymbol{\theta}^{(1)}, \boldsymbol{\theta}^{(2)}) = \Pi \nabla \tilde{\mathbf{p}}(\boldsymbol{\theta}^{(1)}, \boldsymbol{\theta}^{(2)}). \end{aligned}$$

□

Algorithm 8 Local DP IND Test: LocalBitFlipIND

Input: $(\mathbf{x}_1, \dots, \mathbf{x}_n)$, ϵ , α .

Let $\tilde{\mathbf{H}} = \sum_{i=1}^n \mathcal{M}_{\text{bit}}(\mathbf{x}_i, \epsilon)$.

$q = \min_{\boldsymbol{\pi}^{(1)}, \boldsymbol{\pi}^{(2)}} \left\{ \tilde{\mathbf{T}}_{\text{BitFlip}}^{(n)}(\epsilon; (\boldsymbol{\pi}^{(1)}, \boldsymbol{\pi}^{(2)})) \right\}$ from (15).

if $q > \chi_{(r-1)(c-1), 1-\alpha}^2$ Decision \leftarrow Reject.

else Decision \leftarrow Fail to Reject.

Output: Decision

We present the test in Algorithm 8. The following result follows from same privacy analysis as before.

Theorem 6.8. *LocalBitFlipIND* is ϵ -LDP.

6.4 Empirical Results

As we did for the locally private goodness of fit tests, we empirically compare the power for our various tests for independence. We consider the null hypothesis that the two sequences of categorical random variables $\{\mathbf{U}_\ell\}_{\ell=1}^n$ and $\{\mathbf{V}_\ell\}_{\ell=1}^n$ are independent of one another. Under an alternate hypothesis, we generate the contingency data according to a non-product distribution. We fix the distribution $\mathbf{p}^{(1)}$ for the contingency table to be of the following form, where $\boldsymbol{\pi}^{(1)} \in \mathbb{R}^r$ is the unknown distribution for $\{\mathbf{U}_\ell\}_{\ell=1}^n$, $\boldsymbol{\pi}^{(2)} \in \mathbb{R}^c$ is the unknown distribution for $\{\mathbf{V}_\ell\}_{\ell=1}^n$, and r, c are even

$$\begin{aligned} \mathbf{p}^{(1)} = & \boldsymbol{\pi}^{(1)} \left(\boldsymbol{\pi}^{(2)} \right)^\top \\ & + \eta(1, -1, \dots, -1, 1)^\top (1, -1, \dots, -1, 1) \end{aligned} \quad (16)$$

Note that the hypothesis test does not know the underlying $\boldsymbol{\pi}^{(i)}$ for $i \in \{1, 2\}$, but to generate the data we must fix these distributions. We show power results when the marginal distributions satisfy $\boldsymbol{\pi}^{(1)} = (1/r, \dots, 1/r)$ and $\boldsymbol{\pi}^{(2)} = (1/c, \dots, 1/c)$. In Figure 3, we give results for various n and $\epsilon \in \{1, 2, 4\}$ for the case when $(r, c) = (2, 2)$ and $\eta = 0.01$ as well as the case when $(r, c) = (10, 4)$ and $\eta = 0.005$.

Note that we only give results for the two tests LocalGenRRIND and LocalBitFlipIND.




Figure 3: Comparison of empirical power of classical non-private test versus local private tests LocalNoiseIND with Laplace noise (solid line), LocalGenRRIND (dashed line), and LocalBitFlipIND (dotted line). In the left plot we set $(r, c) = (2, 2)$ and $\eta = 0.01$ while we set $(r, c) = (10, 4)$ and $\eta = 0.005$ in the right plot for the contingency table data distribution given in (16).

7 Conclusion

We have designed several hypothesis tests, each depending on different local differentially private algorithms: LocalNoiseGOF (LocalLapGOF), LocalGenRRGOF, and LocalBitFlipGOF as well as their corresponding independence tests LocalNoiseIND, LocalGenRRIND, and LocalBitFlipIND. This required constructing different statistics so that the resulting distribution after injecting noise into the data in order to satisfy privacy could be closely approximated with a chi-square distribution. Hence, we designed rules for when a null hypothesis H_0 should be rejected while satisfying some bound α on Type I error. Further, We showed that each statistic has a noncentral chi-square distribution when the data is drawn from some alternate hypothesis H_1 . Depending on the form of the alternate probability distribution, the dimension of the data, and the privacy parameter, either LocalGenRRGOF or LocalBitFlipGOF gave better power. This corroborates the

results from Kairouz et al. (2014) who showed that in hypothesis testing, different privacy regimes have different optimal local differentially private mechanisms, although utility in their work was in terms of KL divergence. Our results show that the power of the test is directly related to the noncentral parameter of the test statistic that is used. This requires the data analyst to carefully consider alternate hypotheses, as well as the data dimension and privacy parameter for a particular test and then see which test statistic results in the largest noncentral parameter. We focused primarily on goodness of fit testing, where the null hypothesis is a single probability distribution. We further developed local private independence tests which resulted in using previous general chi-square theory presented in Kifer and Rogers (2017). This basic framework can be used to develop other chi-square hypothesis tests where the null hypothesis is not a single parameter. We hope that this will lead to future work on designing local differentially private hypothesis tests beyond chi-square testing.

References

- American Statistical Association. Discovery with Data: Leveraging Statistics with Computer Science to Transform Science and Society, 2014. URL <http://www.amstat.org/policy/pdfs/BigDataStatisticsJune2014.pdf>.
- Apple Press Info. Apple previews ios 10, the biggest ios release ever, 2016. URL <https://www.apple.com/pr/library/2016/06/13Apple-Previews-iOS-10-The-Biggest-iOS-Release-Ever.html>.
- Raef Bassily and Adam Smith. Local, private, efficient protocols for succinct histograms. In *Proceedings of the Forty-seventh Annual ACM Symposium on Theory of Computing*, STOC '15, pages 127–135, New York, NY, USA, 2015. ACM. ISBN 978-1-4503-3536-2. doi: 10.1145/2746539.2746632. URL <http://doi.acm.org/10.1145/2746539.2746632>.
- Mark Bun and Thomas Steinke. Concentrated differential privacy: Simplifications, extensions, and lower bounds. In *Theory of Cryptography - 14th International Conference, TCC 2016-B, Beijing, China, October 31 - November 3, 2016, Proceedings, Part I*, pages 635–658, 2016. doi: 10.1007/978-3-662-53641-4_24. URL http://dx.doi.org/10.1007/978-3-662-53641-4_24.
- Bryan Cai, Constantinos Daskalakis, and Gautam Kamath. Priv’it: Private and sample efficient identity testing. In *International Conference on Machine Learning*, 2017.
- John C. Duchi, Michael I. Jordan, and Martin J. Wainwright. Local privacy and statistical minimax rates. In *54th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2013, 26-29 October, 2013, Berkeley, CA, USA*, pages 429–438, 2013a. doi: 10.1109/FOCS.2013.53. URL <https://doi.org/10.1109/FOCS.2013.53>.
- John C. Duchi, Martin J. Wainwright, and Michael I. Jordan. Local privacy and minimax bounds: Sharp rates for probability estimation. In *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States.*, pages 1529–1537, 2013b. URL <http://papers.nips.cc/paper/5013-local-privacy-and-minimax-bounds-sharp-rates-for-probability-estimation>.
- Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407, 2014. ISSN 1551-305X. doi: 10.1561/04000000042. URL <http://dx.doi.org/10.1561/04000000042>.
- Cynthia Dwork, Krishnamurthy 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, 2006a. doi: 10.1007/11761679_29.
- 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, 2006b.
- Úlfar Erlingsson, Vasily Pihur, and Aleksandra Korolova. Rappor: Randomized aggregatable privacy-preserving ordinal response. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security, CCS '14*, pages 1054–1067, New York, NY, USA, 2014. ACM. ISBN 978-1-4503-2957-6. doi: 10.1145/2660267.2660348. URL <http://doi.acm.org/10.1145/2660267.2660348>.
- T.S. Ferguson. *A Course in Large Sample Theory*. Chapman & Hall Texts in Statistical Science Series. Taylor & Francis, 1996. ISBN 9780412043710. URL https://books.google.com/books?id=DDh_0iTw9agC.
- Marco Gaboardi, Hyun Woo Lim, Ryan Rogers, and Salil P. Vadhan. Differentially private chi-squared hypothesis testing: Goodness of fit and independence testing. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48, ICML'16*, pages 2111–2120. JMLR.org, 2016. URL <http://dl.acm.org/citation.cfm?id=3045390.3045613>.

- Nils Homer, Szabolcs Szelinger, Margot Redman, David Duggan, Waibhav Tembe, Jill Muehling, John V. Pearson, Dietrich A. Stephan, Stanley F. Nelson, and David W. Craig. Resolving individuals contributing trace amounts of dna to highly complex mixtures using high-density snp genotyping microarrays. *PLoS Genet*, 4(8), 08 2008.
- David R. Hunter, Steven M. Goodreau, and Mark S. Handcock. Goodness of fit of social network models. *Journal of the American Statistical Association*, 103:248–258, 2008. URL <http://EconPapers.repec.org/RePEc:bes:jnlasa:v:103:y:2008:m:march:p:248-258>.
- 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.
- Peter Kairouz, Sewoong Oh, and Pramod Viswanath. Extremal mechanisms for local differential privacy. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 2879–2887. Curran Associates, Inc., 2014. URL <http://papers.nips.cc/paper/5392-extremal-mechanisms-for-local-differential-privacy.pdf>.
- Peter Kairouz, Keith Bonawitz, and Daniel Ramage. Discrete distribution estimation under local privacy. In *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, pages 2436–2444, 2016. URL <http://jmlr.org/proceedings/papers/v48/kairouz16.html>.
- Kazuya Kakizaki, Kazuto Fukuchi, and Jun Sakuma. Differentially private chi-squared test by unit circle mechanism. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1761–1770, International Convention Centre, Sydney, Australia, 06–11 Aug 2017. PMLR. URL <http://proceedings.mlr.press/v70/kakizaki17a.html>.
- Vishesh Karwa and Aleksandra Slavković. Inference using noisy degrees: Differentially private β -model and synthetic graphs. *Ann. Statist.*, 44(P1):87–112, 02 2016.
- Daniel Kifer and Ryan Rogers. A New Class of Private Chi-Square Hypothesis Tests. In Aarti Singh and Jerry Zhu, editors, *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54 of *Proceedings of Machine Learning Research*, pages 991–1000, Fort Lauderdale, FL, USA, 20–22 Apr 2017. PMLR. URL <http://proceedings.mlr.press/v54/rogers17a.html>.
- Frank McSherry and Kunal Talwar. Mechanism design via differential privacy. In *Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, Providence, RI, October 2007. IEEE. URL <https://www.microsoft.com/en-us/research/publication/mechanism-design-via-differential-privacy/>.
- A. Pastore and M. Gastpar. Locally differentially-private distribution estimation. In *2016 IEEE International Symposium on Information Theory (ISIT)*, pages 2694–2698, July 2016. doi: 10.1109/ISIT.2016.7541788.
- Sofya Raskhodnikova, Adam Smith, Homin K. Lee, Kobbi Nissim, and Shiva Prasad Kasiviswanathan. What can we learn privately? *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, 00: 531–540, 2008. ISSN 0272-5428. doi: doi.ieeecomputersociety.org/10.1109/FOCS.2008.27.
- Or Sheffet. Differentially private least squares: Estimation, confidence and rejecting the null hypothesis. *arXiv preprint arXiv:1507.02482*, 2015.
- Or Sheffet. Locally private hypothesis testing. *CoRR*, abs/1802.03441, 2018. URL <http://arxiv.org/abs/1802.03441>.
- Jessica E. Steele, Pål Roe Sundsøy, Carla Pezzulo, Victor A. Alegana, Tomas J. Bird, Joshua Blumenstock, Johannes Bjelland, Kenth Engø-Monsen, Yves-Alexandre de Montjoye, Asif M. Iqbal, Khandakar N. Hadiuzzaman, Xin Lu, Erik Wetter, Andrew J. Tatem, and Linus Bengtsson. Mapping poverty using mobile

- phone and satellite data. *Journal of The Royal Society Interface*, 14(127), 2017. ISSN 1742-5689. doi: 10.1098/rsif.2016.0690. URL <http://rsif.royalsocietypublishing.org/content/14/127/20160690>.
- 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.
- Yue Wang, Jaewoo Lee, and Daniel Kifer. Differentially private hypothesis testing, revisited. *arXiv preprint arXiv:1511.03376*, 2015.
- Stanley L. Warner. Randomized response: A survey technique for eliminating evasive answer bias. *Journal of the American Statistical Association*, 60:63–69, 1965.
- M. Ye and A. Barg. Optimal schemes for discrete distribution estimation under local differential privacy. In *2017 IEEE International Symposium on Information Theory (ISIT)*, pages 759–763, June 2017. doi: 10.1109/ISIT.2017.8006630.
- 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 Omitted Proofs

Proof of Lemma 5.7. This follows from the central limit theorem. We first compute the expected value

$$\frac{1}{n} \mathbb{E} [\tilde{\mathbf{H}}] = \frac{e^{\epsilon/2}}{e^{\epsilon/2} + 1} \mathbf{p} + \frac{1}{e^{\epsilon/2} + 1} (1 - \mathbf{p}).$$

In order to compute the covariance matrix, we consider the diagonal term (j, j)

$$\mathbb{E} [(\tilde{H}_j)^2] = \frac{e^{\epsilon/2}}{e^{\epsilon/2} + 1} p_j + \frac{1}{e^{\epsilon/2} + 1} (1 - p_j)$$

Next we compute the off diagonal term (j, k)

$$\begin{aligned} \mathbb{E} [\tilde{H}_j \tilde{H}_k] &= \left(\frac{e^{\epsilon/2}}{e^{\epsilon/2} + 1} \right)^2 \Pr[X_j = 1 \ \& \ X_k = 1] \\ &\quad + \frac{e^{\epsilon/2}}{(e^{\epsilon/2} + 1)^2} (\Pr[X_j = 1 \ \& \ X_k = 0] + \Pr[X_j = 0 \ \& \ X_k = 1]) \\ &\quad + \frac{1}{(e^{\epsilon/2} + 1)^2} \Pr[X_j = 0 \ \& \ X_k = 0] \\ &= \left(\frac{1}{e^{\epsilon/2} + 1} \right)^2 \left[e^{\epsilon/2} (p_j + p_k) + (1 - p_j - p_k) \right] \end{aligned}$$

Before we construct the covariance matrix, we simplify a few terms

$$\begin{aligned} \mathbb{E} [(\tilde{H}_j)^2] - \mathbb{E} [\tilde{H}_j]^2 &= \left(\frac{1}{e^{\epsilon/2} + 1} \right)^2 \left[(e^{\epsilon/2} - 1) (e^{\epsilon/2} + 1) p_j + 1 + e^{\epsilon/2} - \left((e^{\epsilon/2} - 1) p_j + 1 \right)^2 \right] \\ &= \left(\frac{1}{e^{\epsilon/2} + 1} \right)^2 \left[(e^{\epsilon/2} - 1) (e^{\epsilon/2} + 1 - 2) p_j - (e^{\epsilon/2} - 1)^2 (p_j)^2 + e^{\epsilon/2} \right] \\ &= \frac{1}{(e^{\epsilon/2} + 1)^2} \left[(e^{\epsilon/2} - 1)^2 p_j (1 - p_j) + e^{\epsilon/2} \right] \end{aligned}$$

Further, we have

$$\begin{aligned} \mathbb{E} [\tilde{H}_j \tilde{H}_k] - \mathbb{E} [\tilde{H}_j] \mathbb{E} [\tilde{H}_k] &= \left(\frac{1}{e^{\epsilon/2} + 1} \right)^2 \left[(e^{\epsilon/2} - 1) (p_j + p_k) + 1 - \left((e^{\epsilon/2} - 1) p_j + 1 \right) \left((e^{\epsilon/2} - 1) p_k + 1 \right) \right] \\ &= - \left(\frac{e^{\epsilon/2} - 1}{e^{\epsilon/2} + 1} \right)^2 p_j p_k \end{aligned}$$

Putting this together, the covariance matrix can then be written as

$$\begin{aligned} \Sigma(\mathbf{p}) &= \mathbb{E} [\tilde{\mathbf{H}} (\tilde{\mathbf{H}})^\top] - \mathbb{E} [\tilde{\mathbf{H}}] \left(\mathbb{E} [\tilde{\mathbf{H}}] \right)^\top \\ &= \left(\frac{e^{\epsilon/2} - 1}{e^{\epsilon/2} + 1} \right)^2 [\text{Diag}(\mathbf{p}) - \mathbf{p}(\mathbf{p})^\top] + \frac{e^{\epsilon/2}}{(e^{\epsilon/2} + 1)^2} I_d \end{aligned}$$

□