

# A New Class of Private Chi-Square Tests

Daniel Kifer\*      Ryan Rogers†

March 1, 2017

## Abstract

In this paper, we develop new test statistics for private hypothesis testing. These statistics are designed specifically so that their asymptotic distributions, after accounting for noise added for privacy concerns, match the asymptotics of the classical (non-private) chi-square tests for testing if the multinomial data parameters lie in lower dimensional manifolds (examples include goodness of fit and independence testing). Empirically, these new test statistics outperform prior work, which focused on noisy versions of existing statistics.

---

\*Department of Computer Science and Engineering, The Pennsylvania State University and the Census Bureau. Email: [dkifer@cse.psu.edu](mailto:dkifer@cse.psu.edu). Supported in part by NSF grant CNS-1228669.

†Department of Applied Mathematics and Computational Science, University of Pennsylvania. Email: [ryrogers@sas.upenn.edu](mailto:ryrogers@sas.upenn.edu).

# Contents

<b>1</b>	<b>Introduction</b>	<b>3</b>
<b>2</b>	<b>Related Work</b>	<b>4</b>
<b>3</b>	<b>Privacy Preliminaries</b>	<b>4</b>
<b>4</b>	<b>General Chi-Square Tests</b>	<b>6</b>
4.1	Private Asymptotics . . . . .	6
4.2	Minimum Chi-Square Theory . . . . .	7
<b>5</b>	<b>Private Goodness of Fit Tests</b>	<b>9</b>
5.1	Unprojected Private Test Statistic . . . . .	10
5.2	Projected Private Test Statistic . . . . .	12
5.3	Comparison of Statistics . . . . .	14
5.4	Power Analysis . . . . .	15
5.5	Experiments for Goodness of Fit Testing . . . . .	17
<b>6</b>	<b>General Chi-Square Private Tests</b>	<b>19</b>
6.1	Application - Independence Test . . . . .	22
6.2	Application - GWAS Testing . . . . .	26
<b>7</b>	<b>General Chi-Square Tests with Arbitrary Noise Distributions</b>	<b>26</b>
7.1	Application - Goodness of Fit Testing . . . . .	27
7.2	Application - Independence Testing . . . . .	28
<b>8</b>	<b>Conclusions</b>	<b>30</b>
<b>9</b>	<b>Acknowledgments</b>	<b>32</b>
	<b>References</b>	<b>33</b>

# 1 Introduction

In 2008, Homer et al. [13] published a proof-of-concept attack showing that participation of individuals in scientific studies can be inferred from aggregate data typically published in genome-wide association studies (GWAS). Since then, there has been renewed interest in protecting confidentiality of participants in scientific data [15, 22, 27, 20] using privacy definitions such as differential privacy and its variations [7, 6, 3, 5].

An important tool in statistical inference is *hypothesis testing*, a general framework for determining whether a given model – called the null hypothesis  $H_0$  – of a population should be rejected based on a sample from the population. One of the main benefits of hypothesis testing is that it gives a way to control the probability of false discovery or Type I error – falsely concluding that a model should be rejected when it is indeed true. Type II error is the probability of failing to reject  $H_0$  when it is false. Typically, scientists want a test that guarantees a pre-specified Type I error (say 0.05) and has high *power* – complement of Type II error.

The standard approach to hypothesis testing is to (1) estimate the model parameters from the data, (2) compute a *test statistic*  $T$  (a function of the data and the model parameters), (3) determine the (asymptotic) distribution of  $T$  under the assumption that the model generated the data, (4) compute the *p-value* (Type I error) as the probability of  $T$  being more extreme than the realized value from the data.<sup>1</sup>

Our main contribution is a general template for creating test statistics involving categorical data. Empirically, they improve on the power of previous work on differentially private hypothesis testing [12, 24], while maintaining at most some given Type I error. Our approach is to select certain properties of non-private hypothesis tests (e.g., their asymptotic distributions) and then build new test statistics that match these properties when Gaussian noise is added (e.g., to achieve *concentrated differential privacy* [5, 3] or (*approximate*) *differential privacy* [6]). Although the test statistics are designed with Gaussian noise in mind, other noise distributions can be applied, e.g. Laplace.<sup>2</sup>

We point out that implications of this work extend beyond simply alleviating privacy concerns. In *adaptive data analysis*, data may be reused for multiple analyses, each of which may depend on previous outcomes thus potentially overfitting. This problem was recently studied in the computer science literature by Dwork et al. [8], who show that differential privacy can help prevent overfitting despite reusing data. There have been several follow up works [9, 4, 1] that improve and extend the connection between differential privacy and generalization guarantees in adaptive data analysis. Specifically, [18] deals with *post-selection hypothesis testing* where they can ensure a bound on Type I error even for several adaptively chosen tests, as long as each test is differentially private.

We discuss related work in Section 2, provide background information about privacy in Section 3, present our extension of minimum chi-square theory in Section 4 and show how it can be applied to goodness of fit (Section 5) and independence testing (Section 6). Experiments appear in these latter two sections. We evaluate our test statistics with non-Gaussian noise in Section 7. We present conclusions in Section 8.

---

<sup>1</sup>For one-sided tests, the *p-value* is the probability of seeing the computed statistic or anything larger under  $H_0$ .

<sup>2</sup>If we use Laplace noise instead, we cannot match properties like the asymptotic distribution of the non-private statistics, but the new test statistics still empirically improve the power of the tests.

## 2 Related Work

One of the first works to study the asymptotic distributions of statistics that use differentially private data came from Wasserman and Zhou [25]. Smith [21] then showed that for a large family of statistics, there is a corresponding differentially private statistic that shares the same asymptotic distribution as the original statistic. However, these results do not ensure that statistically valid conclusions are made for finite samples. It is then the goal of a recent line of work to develop statistical inference tools that give valid conclusions for even reasonably sized datasets.

Prior work on private statistical inference for categorical data can be roughly grouped into two main approaches. The first group adds appropriately scaled noise to the sampled data (or histogram of data) to ensure differential privacy and uses existing classical hypothesis tests, disregarding the additional noise distribution [15]. This approach is based on the argument that the impact of the noise becomes small as the sample size grows large. Along these lines, [23] studies how many more samples would be needed before the test with additional noise recovers the same level of power as the original test on the actual data. However, as pointed out in [11, 16, 17, 12], even for moderately sized datasets, the impact of privacy noise is non-negligible and therefore such an approach can lead to misleading and statistically invalid results, specifically with much higher Type I error than the prescribed amount.

The second group of work consists of tests that focus on adjusting step (3) in the standard approach to hypothesis testing given in the introduction. That is, these tests use the same statistic in the classical hypothesis tests (without noise) and after making the statistic differentially private, they determine the resulting modified asymptotic distribution of the private statistic [22, 27, 24, 12]. Unfortunately, the resulting asymptotic distribution cannot be written analytically, and so Monte Carlo (MC) simulations or numerical approximations are commonly used to determine at what point to reject the null hypothesis.

We focus on a different technique from these two different approaches, namely modifying step (2) in our outline of hypothesis testing. Thus, we consider transforming the test statistic itself so that the resulting distribution is close to the original asymptotic distribution when additional Gaussian noise is used. If the noise is non-Gaussian, then this is followed by another step that appropriately adjusts the asymptotic distribution. The idea of modifying the test statistic for *regression coefficients* to obtain a *t*-statistic in ordinary least squares has also been considered in [19].

## 3 Privacy Preliminaries

Formal privacy definitions can be used to protect scientific data with the careful injection of noise. Hypothesis testing must then properly account for this noise to avoid generating false conclusions. We briefly discuss examples of privacy definitions that can be used and then elaborate on how to add noise to satisfy those definitions.

Let  $\mathcal{X}$  be an arbitrary domain for records. We define two datasets  $\mathbf{x} = (x_1, \dots, x_n), \mathbf{x}' = (x'_1, \dots, x'_n) \in \mathcal{X}^n$  to be *neighboring* if they differ in at most one entry, i.e. there is some  $i \in [n]$  where  $x_i \neq x'_i$ , but  $x_j = x'_j$  for all  $j \neq i$ . We now define differential privacy (DP)[7, 6].

**Definition 3.1** (Differential Privacy). A randomized algorithm  $\mathcal{M} : \mathcal{X}^n \rightarrow \mathcal{O}$  is  $(\epsilon, \delta)$ -DP if for all neighboring datasets  $\mathbf{x}, \mathbf{x}'$  and each subset of outcomes  $S \subseteq \mathcal{O}$ ,

$$\Pr [\mathcal{M}(\mathbf{x}) \in S] \leq e^\epsilon \Pr [\mathcal{M}(\mathbf{x}') \in S] + \delta.$$

If  $\delta = 0$ , we simply say  $\mathcal{M}$  is  $\epsilon$ -DP.

In this work, we focus on a recent variation of differential privacy, called *zero concentrated differential privacy* (zCDP) [3] (see also [5] for the definition of concentrated differential privacy which [3] modified). In order to define zCDP, we first define the Rényi-divergence between two probability distributions.

**Definition 3.2** (Rényi-Divergence). Let  $P_1$  and  $P_2$  be probability distributions on space  $\Omega$ . For  $\alpha \in (1, \infty)$ , we define the Rényi-Divergence of order  $\alpha$  of  $P_1$  from  $P_2$  as

$$D_\alpha(P_1||P_2) = \frac{1}{\alpha - 1} \log \left( \mathbb{E}_{x \sim P_1} \left[ \left( \frac{P_1(x)}{P_2(x)} \right)^{\alpha-1} \right] \right).$$

*Remark 3.3.* Note as  $\alpha \rightarrow 1$  we get KL-divergence and as  $\alpha \rightarrow \infty$  we get max-divergence.

We are now ready to define zCDP.

**Definition 3.4** (zCDP). A mechanism  $\mathcal{M} : \mathcal{X}^n \rightarrow \mathcal{O}$  is  $\rho$ -zero concentrated differentially private (zCDP), if for all neighboring datasets  $\mathbf{x}, \mathbf{x}' \in \mathcal{X}^n$  and all  $\alpha \in (1, \infty)$  we have

$$D_\alpha(\mathcal{M}(\mathbf{x})||\mathcal{M}(\mathbf{x}')) \leq \rho\alpha$$

The following result shows that zCDP lies between *pure*-DP where  $\delta = 0$  and *approximate*-DP where  $\delta$  may be positive.

**Theorem 3.5** ([3]). *If  $\mathcal{M}$  is  $\epsilon$ -DP, then  $\mathcal{M}$  is  $\frac{\epsilon^2}{2}$ -zCDP. Further, if  $\mathcal{M}$  is  $\rho$ -zCDP then  $\mathcal{M}$  is  $(\rho + 2\sqrt{\rho \ln(1/\delta)}, \delta)$ -DP for every  $\delta > 0$ .*

In order to compute some statistic  $f : \mathcal{X}^n \rightarrow \mathbb{R}^d$  on the data, a differentially private mechanism is to simply add symmetric noise to  $f(\mathbf{x})$  with standard deviation that depends on the *global sensitivity* of  $f$ , which we define as

$$\Delta_p(f) = \max_{\text{neighboring } \mathbf{x}, \mathbf{x}' \in \mathcal{X}^n} \{\|f(\mathbf{x}) - f(\mathbf{x}')\|_p\}.$$

In statistical hypothesis tests, it is typical to use the central limit theorem to form statistics of the data that are asymptotically normally distributed. Then we can determine whether to reject the given model in hypothesis testing by computing the corresponding  $p$ -values based on the asymptotic distribution of the statistic, which works well in practice. Because Gaussian random variables have nice composition guarantees, like the sum of two Gaussian random variables is again Gaussian (a property that is not shared with Laplace random variables), it is then desirable to use a privacy definition which is more accustomed to Gaussian perturbations. We then define the Gaussian mechanism  $\mathcal{M}_{Gauss} : \mathcal{X}^n \rightarrow \mathbb{R}^d$  for statistic  $f : \mathcal{X}^n \rightarrow \mathbb{R}^d$ , where  $\sigma = \frac{\Delta_2(f)}{\sqrt{2\rho}}$ , as

$$\mathcal{M}_{Gauss}(\mathbf{x}) \sim N(f(\mathbf{x}), \sigma^2 I_d). \tag{1}$$

**Theorem 3.6.** *For statistic  $f : \mathcal{X}^n \rightarrow \mathbb{R}^d$ , the Gaussian mechanism  $\mathcal{M}_{Gauss}$  is  $\rho$ -zCDP.*

We now state several of the nice properties that zCDP shares with DP.

**Theorem 3.7** (Post Processing [3]). *Let  $\mathcal{M} : \mathcal{X}^n \rightarrow \mathcal{O}$  and  $g : \mathcal{O} \rightarrow \mathcal{O}'$  be randomized algorithms. If  $\mathcal{M}$  is  $\rho$ -zCDP then  $\mathcal{M}' : \mathcal{X}^n \rightarrow \mathcal{O}'$  where  $\mathcal{M}'(\mathbf{x}) = g(\mathcal{M}(\mathbf{x}))$  is  $\rho$ -zCDP.*

**Theorem 3.8** (Composition [3]). *Let  $\mathcal{M}_1 : \mathcal{X}^n \rightarrow \mathcal{O}$  and  $\mathcal{M}_2 : \mathcal{X}^n \rightarrow \mathcal{O}'$  be randomized algorithms where  $\mathcal{M}_1$  is  $\rho_1$ -zCDP and  $\mathcal{M}_2$  is  $\rho_2$ -zCDP. Then the composition  $\mathcal{M} : \mathcal{X}^n \rightarrow \mathcal{O} \times \mathcal{O}'$  where  $\mathcal{M}(\mathbf{x}) = (\mathcal{M}_1(\mathbf{x}), \mathcal{M}_2(\mathbf{x}))$  is  $(\rho_1 + \rho_2)$ -zCDP.*

For this work we will be considering categorical data. That is, we assume the domain  $\mathcal{X}$  has been partitioned into  $d$  buckets or outcomes and the function  $f : \mathcal{X}^n \rightarrow \mathbb{R}^d$  returns a histogram counting how many records are in each bucket. Our test statistics will only depend on this histogram. Since neighboring datasets  $\mathbf{x}, \mathbf{x}'$  of size  $n$  differ on only one entry, their corresponding histograms differ by  $\pm 1$  in exactly two buckets. Hence, we will say that two histograms are neighboring if they differ in at most two entries by at most 1. In this case,  $\Delta_1(f) = 2$  and  $\Delta_2(f) = \sqrt{2}$ . To preserve privacy, we will add noise to the corresponding histogram  $X = (X_1, \dots, X_d)$  of our original dataset to get  $\tilde{X} = (\tilde{X}_1, \dots, \tilde{X}_d)$ . We perform hypothesis testing on this noisy histogram  $\tilde{X}$ . By Theorem 3.7, we know that each of our hypothesis tests will be  $\rho$ -zCDP as long as we add Gaussian noise with variance  $1/\rho$  to each count in  $X$ . Similarly, if we add Laplace noise with scale  $2/\epsilon$  to each count, we will achieve  $\epsilon$ -DP (this is just an instance of the Laplace Mechanism [7]).

## 4 General Chi-Square Tests

In the non-private setting, a chi-square test involves a histogram  $X$  and a model  $H_0$  that produces expected counts  $\bar{X}$  over the  $d$  buckets. In general,  $H_0$  will have  $k < d$  parameters and will estimate the parameters from  $X$ . The chi-square test statistic is defined as

$$T_{\text{chi}} = \sum_{i=1}^d (X_i - \bar{X}_i)^2 / \bar{X}_i.$$

If the data were generated from  $H_0$  and if  $k$  parameters had to be estimated, then the asymptotic distribution of  $T_{\text{chi}}$  is  $\chi_{d-k-1}^2$ , a chi-square random variable with  $d - k - 1$  degrees of freedom. This is the property we want our statistics to have when they are computed from the noisy histogram  $\tilde{X}$  instead of  $X$ . Note that in the classical chi-square tests (e.g. Pearson independence test), the statistic  $T_{\text{chi}}$  is computed and if it is larger than the  $1 - \alpha$  percentile of  $\chi_{d-k-1}^2$ , then the model is rejected.

The above facts are part of a more general *minimum chi-square asymptotic theory* [10], which we overview in Section 4.2. However, we first explain the differences between private and non-private asymptotics [24, 12].

### 4.1 Private Asymptotics

In non-private statistics, a function of  $n$  data records is considered a random variable, and non-private asymptotics considers this distribution as  $n \rightarrow \infty$ . In private asymptotics, there is another quantity  $\sigma_n^2$ , the variance of the added noise.

In the *classical private regime*, one studies what happens as  $n/\sigma_n^2 \rightarrow \infty$ ; i.e., when the variance due to privacy is insignificant compared to sampling variance in the data (i.e.  $O(n)$ ). In practice, asymptotic distributions derived under this regime result in unreliable hypothesis tests because privacy noise is significant [22].

In the *variance-aware private regime*, one studies what happens as  $n/\sigma_n^2 \rightarrow \text{constant}$  as  $n \rightarrow \infty$ ; that is, when the variance due to privacy is proportional to sampling variance. In practice, asymptotic distributions derived under this regime result in hypothesis tests with reliable Type I error (i.e. the  $p$ -values they generate are accurate) [12, 24]. From now on, we will be using the variance-aware privacy regime.<sup>3</sup>

## 4.2 Minimum Chi-Square Theory

In this section, we present important results about *minimum chi-square theory*. The discussion is based largely on [10] (Chapter 23). Our work relies on this theory to construct new private test statistics in Sections 5 and 6 whose asymptotic behavior matches the non-private asymptotic behavior of the classical chi-square test.

We consider a sequence of  $d$ -dimensional random vectors  $V^{(n)}$  for  $n \geq 1$  (e.g. the data histogram). The parameter space  $\Theta$  is a non-empty open subset of  $\mathbb{R}^k$ , where  $k \leq d$ . The model  $A$  maps a  $k$ -dimensional parameter  $\theta \in \Theta$  into a  $d$ -dimensional vector (e.g., the expected value of  $V^{(n)}$ ), hence it maps  $\Theta$  to a subset of a  $k$ -dimensional manifold in  $d$ -dimensional space.

In this abstract setting, the null hypothesis is that there exists a  $\theta^0 \in \Theta$  such that:<sup>4</sup>

$$\sqrt{n} \left( V^{(n)} - A(\theta^0) \right) \xrightarrow{D} N(0, C(\theta^0)) \quad (2)$$

where  $C(\theta) \in \mathbb{R}^{d \times d}$  is a covariance matrix. Intuitively, Equation 2 says that the Central Limit Theorem can be applied for  $\theta^0$ .

We measure the distance between  $V^{(n)}$  and  $A(\theta)$  with a test statistic given by the following quadratic form:

$$D^{(n)}(\theta) = n \left( V^{(n)} - A(\theta) \right)^\top M(\theta) \left( V^{(n)} - A(\theta) \right) \quad (3)$$

where  $M(\theta) \in \mathbb{R}^{d \times d}$  is a symmetric positive-semidefinite matrix; different choices of  $M$  will result in different test statistics. We make the following standard assumptions about  $A(\theta)$  and  $M(\theta)$ .

**Assumption 4.1.** *For all  $\theta \in \Theta$ , we have:*

- $A(\theta)$  is bicontinuous,<sup>5</sup>
- $A(\theta)$  has continuous first partial derivatives, which we denote as  $\dot{A}(\theta)$  with full rank  $k$ ,
- $M(\theta)$  is continuous in  $\theta$  and there exists an  $\eta > 0$  such that  $M(\theta) - \eta I_d$  is positive definite in an open neighborhood of  $\theta^0$ .

The following theorem will be useful in determining the distribution for the quadratic form  $D^{(n)}(\theta)$ .

**Theorem 4.2** ([10]). *Let  $W \sim N(0, \Lambda)$ .  $W^\top W \sim \chi_r^2$  (chi-square distribution with  $r$  degrees of freedom) if and only if  $\Lambda$  is a projection of rank  $r$ . If  $\Lambda$  is invertible,  $W^\top \Lambda^{-1} W \sim \chi_r^2$ .*

<sup>3</sup>Note that taking  $n$  and  $\sigma_n^2$  to infinity is just a mathematical tool for simplifying expressions while mathematically keeping privacy noise variance proportional to the data variance; it does not mean that the amount of actual noise added to the data depends on the data size.

<sup>4</sup>Here  $\xrightarrow{D}$  means convergence in distribution, as in the Central Limit Theorem [10].

<sup>5</sup>i.e.  $\theta_j \rightarrow \theta \Leftrightarrow A(\theta_j) \rightarrow A(\theta)$ .

If  $\theta^0$  is known, setting  $M(\theta) = C(\theta)^{-1}$  in (3) and applying Theorem 4.2 shows that then  $D^{(n)}(\theta^0)$  converges in distribution to  $\chi_d^2$ . However, as we show in Section 5, this can be a sub-optimal choice of  $M$ .

When  $\theta^0$  is not known, we need to estimate a good parameter  $\hat{\theta}^{(n)}$  to plug into (3). One approach is to set  $\hat{\theta}^{(n)} = \arg \min_{\theta \in \Theta} D^{(n)}(\theta)$ . However, this can be a difficult optimization. If there is a rough estimate of  $\theta^0$  based on the data, call it  $\phi(V^{(n)})$ , and if it converges in probability to  $\theta^0$  (i.e.  $\phi(V^{(n)}) \xrightarrow{P} \theta^0$  as  $n \rightarrow \infty$ ), then we can plug it into the middle matrix to get:

$$\widehat{D}^{(n)}(\theta) = n \left( V^{(n)} - A(\theta) \right)^\top M(\phi(V^{(n)})) \left( V^{(n)} - A(\theta) \right). \quad (4)$$

and then set our estimator  $\hat{\theta}^{(n)} = \arg \min_{\theta \in \Theta} \widehat{D}^{(n)}(\theta)$ . The test statistic becomes  $\widehat{D}^{(n)}(\hat{\theta}^{(n)})$  and the following theorems describe its asymptotic properties under the null hypothesis. We use the shorthand  $A = A(\theta^0)$ ,  $M = M(\theta^0)$ , and  $C = C(\theta^0)$ . See the appendix for the full proof, which follows a similar argument as in [10].

**Theorem 4.3.** *Let  $\hat{\theta}^{(n)} = \arg \min_{\theta \in \Theta} \widehat{D}^{(n)}(\theta)$ . Given Assumption 4.1 and (2), we have  $\sqrt{n}(\hat{\theta}^{(n)} - \theta^0) \xrightarrow{D} N(0, \Psi)$  where  $\theta^0$  is the true parameter and*

$$\Psi = \left( \dot{A}^\top M \dot{A} \right)^{-1} \dot{A}^\top M C M \dot{A} \left( \dot{A}^\top M \dot{A} \right)^{-1}.$$

*Proof.* Since  $\phi(V^{(n)})$  converges in probability to  $\theta^0$  and  $M(\cdot)$  is a continuous mapping, then for any  $b > 0, c > 0$  there exists an  $n_0$  such that when  $n \geq n_0$  then  $M(\phi(V^{(n)}))$  is within a distance  $b$  from  $M(\theta^0)$  with probability at least  $1 - c$ , which makes  $M(\phi(V^{(n)}))$  positive definite with high probability for sufficiently large  $n$ . Furthermore, for any  $d > 0$ , we can choose  $n$  large enough so that the smallest eigenvalue of  $M(\phi(V^{(n)}))$  is at least  $\gamma - d$ .

Since the parameter space is compact, we know a minimizer exists for  $\widehat{D}^{(n)}(\theta)$ . Together, this implies that for sufficiently large  $n$  and with high probability  $\widehat{D}^{(n)}(\hat{\theta}^{(n)}) \geq 0$ .

Also,  $\widehat{D}^{(n)}(\hat{\theta}^{(n)}) \leq \widehat{D}^{(n)}(\theta^0)$  but  $\widehat{D}^{(n)}(\theta^0)/n \xrightarrow{P} 0$  since  $M(\phi(V^{(n)})) \xrightarrow{P} M$  and  $V^{(n)} \xrightarrow{P} A$ . Thus  $\widehat{D}^{(n)}(\hat{\theta}^{(n)})/n \xrightarrow{P} 0$  which means  $V^{(n)} - A(\hat{\theta}^{(n)}) \xrightarrow{P} 0$  (since  $M(\phi(V^{(n)}))$  is positive definite with high probability and uniformly bounded away from 0 in a neighborhood of  $\theta^0$ ). This implies that  $A(\hat{\theta}^{(n)}) \xrightarrow{P} A$  and so  $\hat{\theta}^{(n)} \xrightarrow{P} \theta^0$  since  $A(\theta)$  is bicontinuous by assumption.

Thus, with high probability (e.g.,  $\geq 1 - c$  for large enough  $n$ ),  $\hat{\theta}^{(n)}$  satisfies the first order optimality condition  $\nabla \widehat{D}^{(n)}(\hat{\theta}^{(n)}) = 0$ . This is the same as

$$\dot{A}(\hat{\theta}^{(n)})^\top M(\phi(V^{(n)})) (V^{(n)} - A(\hat{\theta}^{(n)})) = 0 \quad (5)$$

Expanding  $A(\hat{\theta}^{(n)})$  around  $\theta^0$ .

$$A(\hat{\theta}^{(n)}) = A(\theta^0) + \underbrace{\int_0^1 \dot{A}(\theta^0 + t(\hat{\theta}^{(n)} - \theta^0)) dt}_{\equiv B(\hat{\theta}^{(n)})} (\hat{\theta}^{(n)} - \theta^0) \quad (6)$$

Substituting (6) into (5), we get:

$$\dot{A}(\hat{\theta}^{(n)})^\top M(\phi(V^{(n)})) \left( (V^{(n)} - A(\theta^0) - B(\hat{\theta}^{(n)})(\hat{\theta}^{(n)} - \theta^0)) \right) = 0 \quad (7)$$



$$\dot{A}(\widehat{\theta}^{(n)})^\top M(\phi(V^{(n)}))B(\widehat{\theta}^{(n)})\sqrt{n}(\widehat{\theta}^{(n)} - \theta^0) = \dot{A}(\widehat{\theta}^{(n)})^\top M(\phi(V^{(n)}))\sqrt{n}(V^{(n)} - p(\theta^0)) \quad (8)$$

Now, by the continuity of  $\dot{A}(\cdot)$  and the definition of  $B$  and the convergence in probability of  $\widehat{\theta}^{(n)}$  to  $\theta^0$ , we have  $B(\widehat{\theta}^{(n)}) \xrightarrow{P} \dot{A}(\theta^0)$ . Since  $\dot{A}(\theta)$  has full rank by assumption, then for sufficiently large  $n$ ,  $B(\widehat{\theta}^{(n)})$  has full rank with high probability. This leads to the following expression with high probability for sufficiently large  $n$ ,

$$\sqrt{n}(\widehat{\theta}^{(n)} - \theta^0) = \left( \dot{A}(\widehat{\theta}^{(n)})^\top M(\phi(V^{(n)}))B(\widehat{\theta}^{(n)}) \right)^{-1} \dot{A}(\widehat{\theta}^{(n)})^\top M(\phi(V^{(n)}))\sqrt{n}(V^{(n)} - A) \quad (9)$$

Since  $M(\phi(V^{(n)}))$  has smallest eigenvalue at least  $\gamma - d > 0$  with high probability for  $n$  large enough, and since  $\phi(V^{(n)}) \xrightarrow{P} \theta^0$ ,  $\widehat{\theta}^{(n)} \xrightarrow{P} \theta^0$ ,  $B(\widehat{\theta}^{(n)}) \rightarrow \dot{A}(\theta^0)$  in probability, using continuity in all of the above functions, and the assumption that  $\sqrt{n}(V^{(n)} - A) \rightarrow N(0, C)$  in distribution (and Slutsky's theorem) we get:

$$\sqrt{n}(\widehat{\theta}^{(n)} - \theta^0) \xrightarrow{D} N(0, \Psi) \quad \text{as} \quad n \rightarrow \infty. \quad (10)$$

□

We then state the following result using a slight modification of Theorem 24 in [10]. The proof can be in the appendix.

**Theorem 4.4.** *Let  $\nu$  be the rank of  $C(\theta_0)$ . If Assumption 4.1 and (2) hold, and, for all  $\theta \in \Theta$ ,*

$$C(\theta)M(\theta)C(\theta) = C(\theta)$$

and

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

then for  $\widehat{\theta}^{(n)}$  given in Theorem 4.3 and  $\widehat{D}^{(n)}(\theta)$  given in (4) we have:

$$\widehat{D}^{(n)}(\widehat{\theta}^{(n)}) \xrightarrow{D} \chi_{\nu-k}^2$$

*Proof.* Note that Theorem 24 in [10] shows that if the hypotheses hold then

$$n \left( V^{(n)} - A(\widehat{\theta}^{(n)}) \right)^\top M(\widehat{\theta}^{(n)}) \left( V^{(n)} - A(\widehat{\theta}^{(n)}) \right) \xrightarrow{D} \chi_{\nu-k}^2.$$

Note that we have  $\phi(V^{(n)}) \xrightarrow{P} \theta^0$  and  $\widehat{\theta}^{(n)} \xrightarrow{P} \theta^0$  for the true parameter  $\theta^0 \in \Theta$ . We can then apply Slutsky's Theorem due to  $M(\cdot)$  being continuous, to obtain the result for  $\widehat{D}^{(n)}(\widehat{\theta}^{(n)})$ . □

## 5 Private Goodness of Fit Tests

As a warmup, we will first cover goodness of fit testing where the null hypothesis is simply testing whether the underlying unknown parameter is equal to a particular value. We consider categorical data  $X^{(n)} = (X_1^{(n)}, \dots, X_d^{(n)}) \sim \text{Multinomial}(n, \mathbf{p})$  where  $\mathbf{p} = (p_1, \dots, p_d)$  is some probability vector over the  $d$  outcomes. We want to test the null hypothesis  $H_0 : \mathbf{p} = \mathbf{p}^0$ , where each component of  $\mathbf{p}^0$  is positive, but we want to do so in a private way. We then have the following classical result [2].

**Lemma 5.1.** *Under the null hypothesis  $H_0 : \mathbf{p} = \mathbf{p}^0$ ,  $X^{(n)}/n$  is asymptotically normal*

$$\sqrt{n} \left( \frac{X^{(n)}}{n} - \mathbf{p}^0 \right) \xrightarrow{D} N(0, \Sigma)$$

where  $\Sigma$  has rank  $d - 1$  and can be written as

$$\Sigma \stackrel{\text{defn}}{=} \text{Diag}(\mathbf{p}^0) - \mathbf{p}^0(\mathbf{p}^0)^\top. \quad (11)$$

## 5.1 Unprojected Private Test Statistic

To preserve  $\rho$ -zCDP, we will add appropriately scaled Gaussian noise to each component of the histogram  $X^{(n)}$ . We then define the zCDP statistic  $U_\rho^{(n)} = \left( U_{\rho,1}^{(n)}, \dots, U_{\rho,d}^{(n)} \right)$  where we write  $Z \sim N(0, 1/\rho \cdot I_d)$  and

$$U_\rho^{(n)} \stackrel{\text{defn}}{=} \sqrt{n} \left( \frac{X^{(n)} + Z}{n} - \mathbf{p}^0 \right). \quad (12)$$

We next derive the asymptotic distribution of  $U_\rho^{(n)}$  under both private asymptotic regimes in Section 4.1 (note that  $\sigma^2 = 1/\rho$ ).

**Lemma 5.2.** *The random vector  $U_{\rho_n}^{(n)}$  from (12) under the null hypothesis  $H_0 : \mathbf{p} = \mathbf{p}^0$  has the following asymptotic distribution. If  $n\rho_n \rightarrow \infty$  then  $U_{\rho_n}^{(n)} \xrightarrow{D} N(0, \Sigma)$ . Further, if  $n\rho_n \rightarrow \rho > 0$  then  $U_{\rho_n}^{(n)} \xrightarrow{D} N(0, \Sigma_\rho)$  where  $\Sigma_\rho$  has full rank and*

$$\Sigma_\rho \stackrel{\text{defn}}{=} \Sigma + 1/\rho \cdot I_d. \quad (13)$$

*Proof.* We know from the central limit theorem that  $U_\rho^{(n)}$  will converge in distribution to a multivariate normal with covariance matrix given in (13). We now show that  $\Sigma_\rho$  is full rank. From (11) we know that  $\Sigma$  is positive-semidefinite because it is a covariance matrix, hence it has all nonnegative eigenvalues. We then consider the eigenvalues of  $\Sigma_\rho$ . Let  $\mathbf{x} \in \mathbb{R}^d$  be an eigenvector of  $\Sigma_\rho$  with eigenvalue  $\lambda \in \mathbb{R}$ , i.e.

$$\Sigma_\rho \mathbf{x} = \lambda \mathbf{x} \implies \Sigma \mathbf{x} = (\lambda - 1/\rho) \mathbf{x}.$$

We then must have that  $\mathbf{x}$  is also an eigenvector of  $\Sigma$ . Because  $\Sigma$  is positive-semidefinite we have the following inequality

$$\lambda - 1/\rho \geq 0 \implies \lambda \geq 1/\rho > 0.$$

Thus, all the eigenvalues of  $\Sigma_\rho$  are positive, which results in  $\Sigma_\rho$  being nonsingular.  $\square$

Because  $\Sigma_\rho$  is invertible when the privacy parameter  $\rho > 0$ , we can create a new statistic based on  $U_\rho^{(n)}$  that has a chi-square asymptotic distribution under variance-aware privacy asymptotics.

**Theorem 5.3.** *Let  $U_{\rho_n}^{(n)}$  be given in (12) for  $n\rho_n \rightarrow \rho > 0$ . If the null hypothesis  $H_0 : \mathbf{p} = \mathbf{p}^0$  holds, then for  $\Sigma_{n\rho_n}$  given in (13), we have*

$$Q_{\rho_n}^{(n)} \stackrel{\text{defn}}{=} \left( U_{\rho_n}^{(n)} \right)^\top \Sigma_{n\rho_n}^{-1} U_{\rho_n}^{(n)} \xrightarrow{D} \chi_d^2. \quad (14)$$

*Proof.* We directly apply Theorem 4.2 with  $W^{(n)} = \Sigma_{n\rho_n}^{-1/2} U_{\rho_n}^{(n)}$  which is asymptotically multivariate normal with mean zero and covariance  $\Sigma_\rho^{-1/2} \Sigma_\rho \Sigma_\rho^{-1/2} = I_d$ .  $\square$

By computing the inverse of  $\Sigma_{n\rho_n}$  we can simplify the statistic  $Q_{\rho_n}^{(n)}$ .

**Lemma 5.4.** *We can rewrite the statistic in (14) as*

$$Q_\rho^{(n)} = \sum_{i=1}^d \frac{\left( U_{\rho,i}^{(n)} \right)^2}{p_i^0 + \frac{1}{n\rho}} + \frac{n\rho}{\sum_{\ell=1}^d \frac{p_\ell^0}{p_\ell^0 + \frac{1}{n\rho}}} \left( \sum_{j=1}^d \frac{p_j^0}{p_j^0 + \frac{1}{n\rho}} \cdot U_{\rho,j}^{(n)} \right)^2. \quad (15)$$

*Proof.* We begin by writing the inverse of the covariance matrix  $\Sigma_\rho$  from (13) by applying Woodbury's formula [26] which gives the inverse of a modified rank deficient matrix,

$$\Sigma_\rho^{-1} = \text{Diag}(\mathbf{p}^0 + 1/\rho \cdot \mathbf{1})^{-1} + \frac{1}{1 - \mathbf{p}^0 \cdot \omega(\rho)} \omega(\rho) \omega(\rho)^\top \quad (16)$$

where  $\omega(\rho) = \left( \frac{p_1^0}{p_1^0 + 1/\rho}, \dots, \frac{p_d^0}{p_d^0 + 1/\rho} \right)^\top = \frac{\mathbf{p}^0}{\mathbf{p}^0 + 1/\rho \cdot \mathbf{1}}$ .

We note that the vector  $\mathbf{1}$  is an eigenvector of  $\Sigma_\rho$  and  $\Sigma_\rho^{-1}$  with eigenvalue  $1/\rho$  and  $\rho$ , respectively. Letting  $\tilde{X}_i^{(n)} = X_i^{(n)} + Z_i$  be the perturbed version of  $X_i^{(n)}$  leads to the test statistic

$$\begin{aligned} \left( U_\rho^{(n)} \right)^\top \Sigma_{n\rho}^{-1} U_\rho^{(n)} &= \sum_{i=1}^d \frac{(\tilde{X}_i^{(n)} - np_i^0)^2}{np_i^0 + 1/\rho} + \frac{1}{1 - \sum_i \frac{(p_i^0)^2}{p_i^0 + \frac{1}{n\rho}}} \left( \sum_{i=1}^d \frac{(\tilde{X}_i^{(n)} - np_i^0)}{\sqrt{n}} \frac{p_i^0}{p_i^0 + \frac{1}{n\rho}} \right)^2 \\ &= \sum_{i=1}^d \frac{(\tilde{X}_i^{(n)} - np_i^0)^2}{np_i^0 + 1/\rho} + \frac{1}{1 - \sum_i \frac{(p_i^0)^2}{p_i^0 + \frac{1}{n\rho}}} \left( \sum_{i=1}^d \frac{(\tilde{X}_i^{(n)} - np_i^0)}{\sqrt{n}} \frac{p_i^0}{p_i^0 + \frac{1}{n\rho}} \right)^2 \end{aligned}$$

We can then rewrite the term in the denominator,

$$1 - \sum_{i=1}^d \frac{(p_i^0)^2}{p_i^0 + \frac{1}{n\rho}} = \sum_{i=1}^d \left( \frac{p_i^0(p_i^0 + \frac{1}{n\rho})}{p_i^0 + \frac{1}{n\rho}} - \frac{(p_i^0)^2}{p_i^0 + \frac{1}{n\rho}} \right) = \frac{1}{n\rho} \cdot \sum_{i=1}^d \frac{p_i^0}{p_i^0 + \frac{1}{n\rho}}.$$

Recalling the form of  $U_\rho^{(n)}$  from (12) concludes the proof.  $\square$

Note that the coefficient on the second term of (15) grows large as  $n\rho \rightarrow \infty$ , so this test statistic does not approach the nonprivate test for a fixed  $\rho$ . This is not surprising since  $\Sigma_{n\rho}$  must converge to a singular matrix as  $n\rho \rightarrow \infty$ . Further, the additional noise adds a degree of freedom to the asymptotic distribution of the original statistic. This additional degree of freedom results in increasing the point in which we reject the null hypothesis, i.e. the critical value. Thus, rejecting an incorrect model becomes harder as we increase the degrees of freedom, and hence decreases power.

## 5.2 Projected Private Test Statistic

Given that the test statistic in the previous section depends on a nearly singular matrix, we now derive a new test statistic for the private goodness of fit test. It has the remarkable property that its asymptotic distribution is  $\chi_{d-1}^2$  under both private asymptotics.

We start with the following observation. In the classical chi-square test, the random variables  $\left(\frac{X_i^{(n)} - np_i^0}{\sqrt{np_i^0}}\right)_{i=1}^d$  have covariance matrix  $I_d - \sqrt{\mathbf{p}^0}\sqrt{\mathbf{p}^0}^\top$  under the null hypothesis  $H_0 : \mathbf{p} = \mathbf{p}^0$ . The classical test essentially uncorrelates these random variables and projects them onto the subspace orthogonal to  $\sqrt{\mathbf{p}^0}$ . We will use a similar intuition for the privacy-preserving random vector  $U_\rho^{(n)}$ .

The matrix  $\Sigma_\rho$  in (13) has eigenvector  $\mathbf{1}$  with eigenvalue  $1/\rho$  – regardless of the true parameters of the data-generating distribution. Hence we think of this direction as pure noise. We therefore project  $U_\rho^{(n)}$  onto the space orthogonal to  $\mathbf{1}$  (i.e. enforce the constraint that the entries in  $U_\rho^{(n)}$  add up to 0, as they would in the noiseless case). We then define the *projected statistic*  $\mathcal{Q}_\rho^{(n)}$  as the following where we write the projection matrix  $\mathbf{P} \stackrel{\text{defn}}{=} I_d - \frac{1}{d}\mathbf{1}\mathbf{1}^\top$

$$\mathcal{Q}_\rho^{(n)} \stackrel{\text{defn}}{=} \left(U_\rho^{(n)}\right)^\top \mathbf{P} \Sigma_{n\rho}^{-1} \mathbf{P} U_\rho^{(n)}. \quad (17)$$

It will be useful to write out the middle matrix in  $\mathcal{Q}_{\rho_n}^{(n)}$  for analyzing its asymptotic distribution.

**Lemma 5.5.** *For the covariance matrix  $\Sigma_{n\rho_n}$  given in (13), we have the following identity when  $n\rho_n \rightarrow \rho > 0$*

$$\mathbf{P} \Sigma_{n\rho_n}^{-1} \mathbf{P} \rightarrow \Sigma_\rho^{-1} - \frac{\rho}{d} \cdot \mathbf{1}\mathbf{1}^\top \quad (18)$$

Further, when  $n\rho_n \rightarrow \infty$ , we have the following

$$\mathbf{P} \Sigma_{n\rho_n}^{-1} \mathbf{P} \rightarrow \mathbf{P} \text{Diag}(\mathbf{p}^0)^{-1} \mathbf{P} \quad (19)$$

*Proof.* To prove (18), we use the fact that  $\Sigma_\rho^{-1}$  has eigenvalue  $\rho$  with eigenvector  $\mathbf{1}$ . We then focus on proving (19). We use the identity for the inverse of  $\Sigma_{n\rho_n}^{-1}$  from (16).

$$\begin{aligned} & \mathbf{P} \Sigma_{n\rho_n}^{-1} \mathbf{P} \\ &= \mathbf{P} \text{Diag}\left(\mathbf{p}^0 + \frac{1}{n\rho_n} \cdot \mathbf{1}\right)^{-1} \mathbf{P} \\ & \quad + \frac{n\rho_n}{\sum_{i=1}^d \frac{p_i^0}{p_i^0 + \frac{1}{n\rho_n}}} \cdot \mathbf{P} \left( \frac{\mathbf{p}^0}{\mathbf{p}^0 + \frac{1}{n\rho_n} \mathbf{1}} \right) \left( \frac{\mathbf{p}^0}{\mathbf{p}^0 + \frac{1}{n\rho_n} \mathbf{1}} \right)^\top \mathbf{P} \end{aligned}$$

We then focus on the second term and write  $\lambda_n = \sum_{i=1}^d \frac{p_i^0}{p_i^0 + \frac{1}{n\rho_n}}$ .

$$\begin{aligned} & \frac{n\rho_n}{\lambda_n} \cdot \mathbf{P} \left( \frac{\mathbf{p}^0}{\mathbf{p}^0 + \frac{1}{n\rho_n} \mathbf{1}} \right) \left( \frac{\mathbf{p}^0}{\mathbf{p}^0 + \frac{1}{n\rho_n} \mathbf{1}} \right)^\top \mathbf{P} \\ &= \frac{n\rho_n}{\lambda_n} \cdot \left( \frac{\mathbf{p}^0}{\mathbf{p}^0 + \frac{1}{n\rho_n} \mathbf{1}} - \frac{\lambda_n}{d} \cdot \mathbf{1} \right) \left( \frac{\mathbf{p}^0}{\mathbf{p}^0 + \frac{1}{n\rho_n} \mathbf{1}} - \frac{\lambda_n}{d} \cdot \mathbf{1} \right)^\top \end{aligned}$$

$$\begin{aligned}
&= \frac{n\rho_n}{\lambda_n} \cdot \left( \frac{\mathbf{p}^0}{\mathbf{p}^0 + \frac{1}{n\rho_n} \cdot \mathbf{1}} \right) \left( \frac{\mathbf{p}^0}{\mathbf{p}^0 + \frac{1}{n\rho_n} \cdot \mathbf{1}} \right)^\top \\
&\quad - \frac{n\rho_n}{d} \cdot \left( \frac{\mathbf{p}^0}{\mathbf{p}^0 + \frac{1}{n\rho_n} \cdot \mathbf{1}} \right) \mathbf{1}^\top - \frac{n\rho_n}{d} \cdot \mathbf{1} \left( \frac{\mathbf{p}^0}{\mathbf{p}^0 + \frac{1}{n\rho_n} \cdot \mathbf{1}} \right)^\top + \frac{n\rho_n \lambda_n}{d^2} \mathbf{1} \mathbf{1}^\top
\end{aligned}$$

We consider entry  $(i, j)$  of the above matrix, which we can write as

$$\begin{aligned}
&\frac{n\rho_n}{\lambda_n} \cdot \left( \frac{p_i^0}{p_i^0 + \frac{1}{n\rho_n}} \right) \cdot \left( \frac{p_j^0}{p_j^0 + \frac{1}{n\rho_n}} \right) - \frac{n\rho_n}{d} \left( \frac{p_i^0}{p_i^0 + \frac{1}{n\rho_n}} + \frac{p_j^0}{p_j^0 + \frac{1}{n\rho_n}} \right) + \frac{n\rho_n \lambda_n}{d^2} \\
&= \frac{n\rho_n}{d\lambda_n} \left( \frac{\lambda_n^2}{d} - \frac{1}{(p_i^0 + \frac{1}{n\rho_n})(p_j^0 + \frac{1}{n\rho_n})} \left( \frac{\lambda_n}{n\rho_n} (p_i^0 + p_j^0) - p_i^0 p_j^0 (d - 2\lambda_n) \right) \right) \\
&= n\rho_n \left( \frac{\lambda_n}{d^2} - \frac{(2\lambda_n - d)p_i^0 p_j^0}{d(p_i^0 + \frac{1}{n\rho_n})(p_j^0 + \frac{1}{n\rho_n})} \right) - \frac{p_i^0 + p_j^0}{d\lambda_n(p_i^0 + \frac{1}{n\rho_n})(p_j^0 + \frac{1}{n\rho_n})}.
\end{aligned}$$

We then let  $n \rightarrow \infty$  to get

$$\frac{n\rho_n}{d} \left( \frac{\lambda_n}{d} - \frac{(2\lambda_n - d)p_i^0 p_j^0}{\lambda_n(p_i^0 + \frac{1}{n\rho_n})(p_j^0 + \frac{1}{n\rho_n})} \right) - \frac{p_i^0 + p_j^0}{d\lambda_n(p_i^0 + \frac{1}{n\rho_n})(p_j^0 + \frac{1}{n\rho_n})} \rightarrow \frac{1}{p_i^0} + \frac{1}{p_j^0} - \frac{1}{p_i^0} - \frac{1}{p_j^0} = 0.$$

Thus, we have shown that for  $n\rho_n \rightarrow \infty$ ,

$$\mathbf{P}\Sigma_{n\rho_n}^{-1}\mathbf{P} \rightarrow \mathbf{P}\text{Diag}(\mathbf{p}^0)^{-1}\mathbf{P}.$$

□

The projected statistic is asymptotically chi-square distributed in both private asymptotic regimes.

**Theorem 5.6.** Let  $U_\rho^{(n)}$  be given in (12). For null hypothesis  $H_0 : \mathbf{p} = \mathbf{p}^0$ , we can write the projected statistic  $\mathcal{Q}_\rho^{(n)}$  in the following way for  $\tilde{n} = \sum_{i=1}^d (X_i^{(n)} + Z_i)$

$$\begin{aligned}
\mathcal{Q}_\rho^{(n)} &= \sum_{i=1}^d \frac{\left( U_{\rho,i}^{(n)} \right)^2}{p_i^0 + \frac{1}{n\rho}} - \frac{\rho}{d} (\tilde{n} - n)^2 \\
&\quad + \frac{n\rho}{\sum_{\ell=1}^d \frac{p_\ell^0}{p_\ell^0 + \frac{1}{n\rho}}} \left( \sum_{j=1}^d \frac{p_j^0}{p_j^0 + \frac{1}{n\rho}} \cdot U_j^{(n)}(\rho) \right)^2.
\end{aligned} \tag{20}$$

Further for  $n\rho_n \rightarrow \rho > 0$ , if the null hypothesis holds then  $\mathcal{Q}_{\rho_n}^{(n)} \xrightarrow{D} \chi_{d-1}^2$ .

*Proof.* We first show that we can write the projected statistic in (17) in the proposed way. Using (18), we can write the projected statistic in terms of the unprojected statistic in (15), which will give the expression in (20)

$$\mathcal{Q}_{\rho_n}^{(n)} = \left( U_{\rho_n}^{(n)} \right)^\top \left( \Sigma_{n\rho_n}^{-1} - \frac{n\rho_n}{d} \cdot \mathbf{1} \mathbf{1}^\top \right) U_{\rho_n}^{(n)} = Q_{\rho_n}^{(n)} - \frac{n\rho_n}{d} \cdot \left( U_{\rho_n}^{(n)} \right)^\top \mathbf{1} \mathbf{1}^\top U_{\rho_n}^{(n)}.$$

We then turn to determining the asymptotic distribution of the projected statistics when  $n\rho_n \rightarrow \rho > 0$ . Recall that  $\mathbf{1}$  is an eigenvector of  $\Sigma_\rho$ . Note that  $\Sigma_\rho$  is diagonalizable, i.e.  $\Sigma_\rho = BDB^\top$  where  $D$  is a diagonal matrix and  $B$  is an orthogonal matrix with one column being  $1/d \cdot \mathbf{1}$ . For the following matrix  $\Lambda$ , we can write it as a  $d \times d$  identity matrix except one of the entries on the diagonal is zero.

$$\Lambda = \Sigma_\rho^{-1/2} \mathbf{P} B D B^\top \mathbf{P} \Sigma_\rho^{-1/2}.$$

Thus,  $\Lambda$  is idempotent and has rank  $d - 1$ . We define  $W \sim N(0, I_{d-1})$ . We then know that  $\mathbf{Q}_{\rho_n}^{(n)}$  has the same asymptotic distribution as  $W^\top W$  and so we can apply Theorem 4.2.  $\square$

**Theorem 5.7.** *For histogram data  $X^{(n)}$ , the projected statistic  $\mathbf{Q}_{\rho_n}^{(n)}$  in Theorem 5.6 converges in distribution to a  $\chi_{d-1}^2$  when  $H_0 : \mathbf{p} = \mathbf{p}^0$  holds and  $n\rho_n \rightarrow \infty$ . In fact, as  $n\rho_n \rightarrow \infty$ , the difference between  $\mathbf{Q}_{\rho_n}^{(n)}$  and the classical chi-square statistic  $\sum_{i=1}^d \frac{(X_i^{(n)} - np_i^0)^2}{np_i^0}$  converges in probability to 0.*

*Proof.* Although  $\Sigma_{n\rho_n}^{-1}$  does not exist as  $n\rho_n \rightarrow \infty$ , we can still write the asymptotic projected statistic. The middle matrix in the projected statistic when  $n\rho_n \rightarrow \infty$  is then  $\mathbf{P} \text{Diag}(\mathbf{p}^0)^{-1} \mathbf{P}$ .

When  $n\rho_n \rightarrow \infty$ , we also have that  $U_{\rho_n}^{(n)} \xrightarrow{D} N(0, \Sigma)$  from Lemma 5.2. We then analyze the asymptotic distribution of the projected statistic, where we write  $U \sim N(0, \Sigma)$  and study the distribution of  $U^\top \mathbf{P} \text{Diag}(\mathbf{p}^0)^{-1} \mathbf{P} U$ . We note that we have  $U^\top \mathbf{1} = 0$ , which simplifies the asymptotic distribution of the projected statistic.

$$U^\top \mathbf{P} \text{Diag}(\mathbf{p}^0)^{-1} \mathbf{P} U = \sum_{i=1}^d \frac{U_i^2}{p_i^0}$$

Note that this last final form is exactly the original chi-square statistic used in the classical test, which is known to converge to  $\chi_{d-1}^2$ .  $\square$

### 5.3 Comparison of Statistics

We now want to compare the two private chi-square statistics in (14) and (17) to see which may lead to a larger *power* (i.e. smaller Type II error). The following theorem shows that we can write the unprojected statistic (14) as a combination of both the projected statistic (17) and squared independent Gaussian noise.

**Theorem 5.8.** *Consider histogram data  $X^{(n)}$  that has Gaussian noise  $Z \sim N(0, 1/\rho \cdot I_d)$  added to it. For the statistics  $Q_\rho^{(n)}$  and  $\mathbf{Q}_\rho^{(n)}$  based on the noisy counts given in (14) and (17) respectively, we have*

$$Q_\rho^{(n)} = \mathbf{Q}_\rho^{(n)} + \frac{\rho}{d} \left( \sum_{i=1}^d Z_i \right)^2.$$

*Further, for any fixed data  $X^{(n)}$ ,  $\mathbf{Q}_\rho^{(n)}$  is independent of  $\left( \sum_{i=1}^d Z_i \right)^2$ .*

To prove this we will use the noncentral version of Craig's Theorem.

**Theorem 5.9** (Craig's Theorem [14]). *Let  $Y \sim N(\mu, V)$ . Then the quadratic forms  $Y^\top A Y$  and  $Y^\top B Y$  are independent if  $AVB = 0$ .*

We are now ready to prove our theorem.

*Proof of Theorem 5.8.* We first show that we can write  $Q_\rho^{(n)} - \mathcal{Q}_\rho^{(n)} = \frac{\rho}{d} \left( \sum_{i=1}^d Z_i \right)^2$ . Note that  $\left( U_\rho^{(n)} \right)^\top \mathbf{1} = \sum_{i=1}^d Z_i / \sqrt{n}$  and  $\Sigma_\rho^{-1}$  has eigenvalue  $\rho$  with eigenvector  $\mathbf{1}$ . We then have

$$\begin{aligned}
Q_\rho^{(n)} &= \left( U_\rho^{(n)} \right)^\top \Sigma_{n\rho}^{-1} U_\rho^{(n)} \\
&= \left( U_\rho^{(n)} \right)^\top \left( I_d - \frac{1}{d} \mathbf{1}\mathbf{1}^\top + \frac{1}{d} \mathbf{1}\mathbf{1}^\top \right)^\top \Sigma_{n\rho}^{-1} \left( I_d - \frac{1}{d} \mathbf{1}\mathbf{1}^\top + \frac{1}{d} \mathbf{1}\mathbf{1}^\top \right) U_\rho^{(n)} \\
&= \mathcal{Q}_\rho^{(n)} + \frac{2}{d} \left( U_\rho^{(n)} \right)^\top \mathbf{1}\mathbf{1}^\top \Sigma_{n\rho}^{-1} \left( I_d - \frac{1}{d} \mathbf{1}\mathbf{1}^\top \right) U_\rho^{(n)} + \frac{1}{d^2} \left( U_\rho^{(n)} \right)^\top \mathbf{1}\mathbf{1}^\top \Sigma_{n\rho}^{-1} \mathbf{1}\mathbf{1}^\top U_\rho^{(n)} \\
&= \mathcal{Q}_\rho^{(n)} + \frac{2}{d} \left( U_\rho^{(n)} \right)^\top \mathbf{1}\mathbf{1}^\top \Sigma_{n\rho}^{-1} \mathbf{P} U_\rho^{(n)} + \frac{\rho}{d} \left( \sum_{i=1}^d Z_i \right)^2 \\
&= \mathcal{Q}_\rho^{(n)} + \frac{2n\rho}{d} \left( \sum_{i=1}^d Z_i / \sqrt{n} \right) \mathbf{1}^\top U_\rho^{(n)} - \frac{2\rho}{d} \left( \sum_{i=1}^d Z_i \right)^2 + \frac{\rho}{d} \left( \sum_{i=1}^d Z_i \right)^2 \\
&= \mathcal{Q}_\rho^{(n)} + \frac{\rho}{d} \left( \sum_{i=1}^d Z_i \right)^2
\end{aligned}$$

We now apply Craig's Theorem to show that for a fixed histogram  $X^{(n)}$ , we have  $\mathcal{Q}_\rho^{(n)}$  is independent of  $\left( \sum_{i=1}^d Z_i \right)^2$ . When  $X^{(n)}$  is fixed, we can define the random variable  $Y \sim N(\mu, 1/\rho I_d)$  where  $\mu = (X^{(n)} - n\mathbf{p}^0) / \sqrt{n}$ . If we set  $A = \mathbf{P} \Sigma_{n\rho}^{-1} \mathbf{P}$ , then our projected statistic can be rewritten as  $Y^\top A Y$ . Further, if we define  $B = \mathbf{1}\mathbf{1}^\top$ , then  $\left( \sum_{i=1}^d Y_i \right)^2 = Y^\top B Y$ . We then have  $A(1/\rho \cdot I_d)B = 0$ , so that the projected statistic is independent of  $\left( \sum_{i=1}^d Y_i \right)^2$ . We next note that  $Y = Z + \mu$  and that  $\mathbf{1}^\top \mu = 0$ . Hence,

$$Y^\top B Y = (Z + \mu)^\top B (Z + \mu) = Z^\top B Z + 2\mu^\top B Z + \mu^\top B \mu = Z^\top B Z = \left( \sum_{i=1}^d Z_i \right)^2.$$

□

Algorithm 1 (zCDP-GOF) shows how to perform goodness of fit testing with either of these two test statistics, i.e. unprojected (14) or projected (17). We note that our test is zCDP for neighboring histogram datasets due to it being an application of the Gaussian mechanism and Theorem 3.7. Hence:

**Theorem 5.10.** *zCDP-GOF( $\cdot; \rho, \alpha, \mathbf{p}^0$ ) is  $\rho$ -zCDP.*

## 5.4 Power Analysis

From Theorem 5.8 we see that the difference between  $Q_\rho^{(n)}$  and  $\mathcal{Q}_\rho^{(n)}$  is the addition of squared independent noise. This additional noise can only hurt *power*, because for the same data the

---

**Algorithm 1** zCDP Chi-Square Goodness of Fit Test
 

---

**procedure** zCDP-GOF( $X^{(n)}$ ;  $\rho$ ,  $\alpha$ ,  $H_0 : \mathbf{p} = \mathbf{p}^0$ )

Set  $\tilde{X}^{(n)} \leftarrow X^{(n)} + Z$  where  $Z \sim N(0, 1/\rho \cdot I_d)$ .

**For the unprojected statistic:**

$$T \leftarrow \frac{1}{n} \left( \tilde{X}^{(n)} - n\mathbf{p}^0 \right)^\top \Sigma_{n\rho}^{-1} \left( \tilde{X}^{(n)} - n\mathbf{p}^0 \right)$$

$t \leftarrow (1 - \alpha)$  quantile of  $\chi_d^2$

**For the projected statistic:**

$$T \leftarrow \frac{1}{n} \left( \tilde{X}^{(n)} - n\mathbf{p}^0 \right)^\top \mathbf{P} \Sigma_{n\rho}^{-1} \mathbf{P} \left( \tilde{X}^{(n)} - n\mathbf{p}^0 \right)$$

$t \leftarrow (1 - \alpha)$  quantile of  $\chi_{d-1}^2$

**if**  $T > t$  **then** Reject

---

statistic  $Q_\rho^{(n)}$  has larger variance than  $\mathcal{Q}_\rho^{(n)}$  and does not depend on the underlying data. If we fix an alternate hypothesis, we can obtain asymptotic distributions for our two test statistics.

**Theorem 5.11.** Consider the null hypothesis  $H_0 : \mathbf{p} = \mathbf{p}^0$  and the alternate hypothesis  $H_1 : \mathbf{p} = \mathbf{p}^0 + \frac{\Delta}{\sqrt{n}}$  where  $\sum_{i=1}^d \Delta_i = 0$ . Assuming the data  $X^{(n)}$  comes from the alternate  $H_1$ , the two statistics  $\mathcal{Q}_{\rho_n}^{(n)}$ , and  $Q_{\rho_n}^{(n)}$  have noncentral chi-square distributions when  $n\rho_n \rightarrow \rho > 0$ , i.e.

$$Q_{\rho_n}^{(n)} \xrightarrow{D} \chi_d^2(\Delta^\top \Sigma_\rho^{-1} \Delta) \quad \& \quad \mathcal{Q}_{\rho_n}^{(n)} \xrightarrow{D} \chi_{d-1}^2(\Delta^\top \Sigma_\rho^{-1} \Delta).$$

Further, if  $n\rho_n \rightarrow \infty$  then

$$\mathcal{Q}_{\rho_n}^{(n)} \xrightarrow{D} \chi_{d-1}^2 \left( \sum_i \frac{\Delta_i^2}{p_i^0} \right)$$

We point out that in the case where  $n\rho_n \rightarrow \infty$ , the projected statistic has the same asymptotic distribution as the classical (nonprivate) chi-square test under the same alternate hypothesis.

We will use the following result to prove this theorem.

**Lemma 5.12** ([10]). Suppose  $Y \sim N(\boldsymbol{\delta}, V)$ . If  $V$  is a projection of rank  $\nu$  and  $V\boldsymbol{\delta} = \boldsymbol{\delta}$  then  $Y^\top Y \sim \chi_\nu^2(\boldsymbol{\delta}^\top \boldsymbol{\delta})$ .

*Proof of Theorem 5.11.* In this case we have the random vector  $U_{\rho_n}^{(n)}$  from (12) converging in distribution to  $N(\boldsymbol{\Delta}, \Sigma_\rho)$  if  $n\rho_n \rightarrow \rho > 0$  or  $N(\boldsymbol{\Delta}, \Sigma)$  if  $n\rho_n \rightarrow \infty$  by Lemma 5.2. We first consider the case when  $n\rho_n \rightarrow \rho > 0$ . Consider  $U \sim N(\boldsymbol{\Delta}, \Sigma_\rho)$  and  $Y = \Sigma_\rho^{-1/2} U \sim N((\Sigma_\rho^{-1/2} \boldsymbol{\Delta}, I_d)$ . We then know that  $Y^\top Y$  and the unprojected statistic  $Q_{\rho_n}^{(n)}$  have the same asymptotic distribution. In order to use Lemma 5.12, we need to verify that  $\Sigma_\rho^{-1/2} \Sigma_\rho \Sigma_\rho^{-1/2} (\Sigma_\rho^{-1/2} \boldsymbol{\Delta}) = \Sigma_\rho^{-1/2} \boldsymbol{\Delta}$ , which indeed holds.

We then consider the projected statistic  $\mathcal{Q}_{\rho_n}^{(n)}$  where  $n\rho_n \rightarrow \rho > 0$ . Similar to the proof of Theorem 5.6, we diagonalize  $\Sigma_\rho = BDB^\top$  where  $B$  is an orthogonal matrix with one column being  $1/d \cdot \mathbf{1}$  and  $D$  is a diagonal matrix. We then let  $U \sim N(\boldsymbol{\Delta}, \Sigma_\rho)$  and let

$$Y = \Sigma_\rho^{-1/2} \mathbf{P} U$$



We then have  $Y^\top Y$  and  $\mathcal{Q}_{\rho_n}^{(n)}$  will have the same asymptotic distribution. Recall that  $\Lambda = \Sigma_\rho^{-1/2} \mathbf{P} \Sigma_\rho \mathbf{P} \Sigma_\rho^{-1/2}$  is idempotent with rank  $d - 1$ . Lastly, to apply Lemma 5.12 we need to show the following

$$\Lambda \left( \Sigma_\rho^{-1/2} \mathbf{P} \Delta \right) = \Sigma_\rho^{-1/2} \mathbf{P} \Delta.$$

Let  $\widehat{B} \in \mathbb{R}^{d \times (d-1)}$  be the same as matrix  $B$  whose corresponding column for  $1/d \cdot \mathbf{1}$  is missing, which we assume to be the last column of  $B$ . Further, we define  $\widehat{D} \in \mathbb{R}^{(d-1) \times (d-1)}$  to be the same as  $D$  without the last row and column. We can then write  $\mathbf{P} \Sigma_\rho \mathbf{P} = \widehat{B} \widehat{D} \widehat{B}^\top$ . We can then simplify the left hand side to have

$$\begin{aligned} \Sigma_\rho^{-1/2} \mathbf{P} \Sigma_\rho \mathbf{P} \Sigma_\rho^{-1/2} \mathbf{P} &= BD^{-1/2} B^\top \mathbf{P} \Sigma_\rho \widehat{B} \widehat{D}^{-1} \widehat{B}^\top \\ &= BD^{-1/2} B^\top \widehat{B} \widehat{D} B^\top \widehat{B} \widehat{D}^{-1} \widehat{B}^\top \\ &= BD^{-1/2} B^\top \widehat{B} \widehat{D} \widehat{D}^{-1} \widehat{B}^\top \\ &= BD^{-1/2} \widehat{B}^\top \\ &= BD^{-1/2} B^\top \mathbf{P} \\ &= \Sigma_\rho^{-1/2} \mathbf{P} \end{aligned}$$

The noncentral parameter is then

$$\Delta^\top \mathbf{P} \Sigma_\rho^{-1} \mathbf{P} \Delta$$

We then note that  $\sum_i \Delta_i = 0$ .

For the case when  $n\rho_n \rightarrow \infty$ . From (19), we have  $\mathbf{P} \Sigma_{n\rho_n} \mathbf{P} \rightarrow M_\infty$ , which can be diagonalized. As we showed in Theorem 5.7, we have

$$\left( U_{\rho_n}^{(n)} \right)^\top M_\infty U_{\rho_n}^{(n)} = \left( U_{\rho_n}^{(n)} \right)^\top \text{Diag}(\mathbf{p}^0)^{-1} U_{\rho_n}^{(n)}$$

From Lemma 5.2, we know that  $\left( U_{\rho_n}^{(n)} \right) \xrightarrow{D} N(\Delta, \Sigma)$ . We then write  $U \sim N(\Delta, \Sigma)$  so that our projected chi-square statistic has the same asymptotic distribution as

$$U^\top \text{Diag}(\mathbf{p}^0)^{-1} U$$

which has a  $\chi_{d-1}^2(\Delta \text{Diag}(\mathbf{p}^0)^{-1} \Delta)$  distribution.  $\square$

Note that the noncentral parameters in the previous theorem are the same for both statistics and only the degrees of freedom are different.

## 5.5 Experiments for Goodness of Fit Testing

Throughout all of our experiments, we will fix  $\alpha = 0.05$  and privacy parameter  $\rho = 0.001$ . All of our tests are designed to achieve Type I error at most  $\alpha$ , as we empirically show for different null hypotheses  $\mathbf{p}^0$  and sample size  $n$  in Figure 1. We include 1.96 times the standard error of our 100,000 independent trials (giving a 95% confidence interval) for each sample size and each null hypothesis.

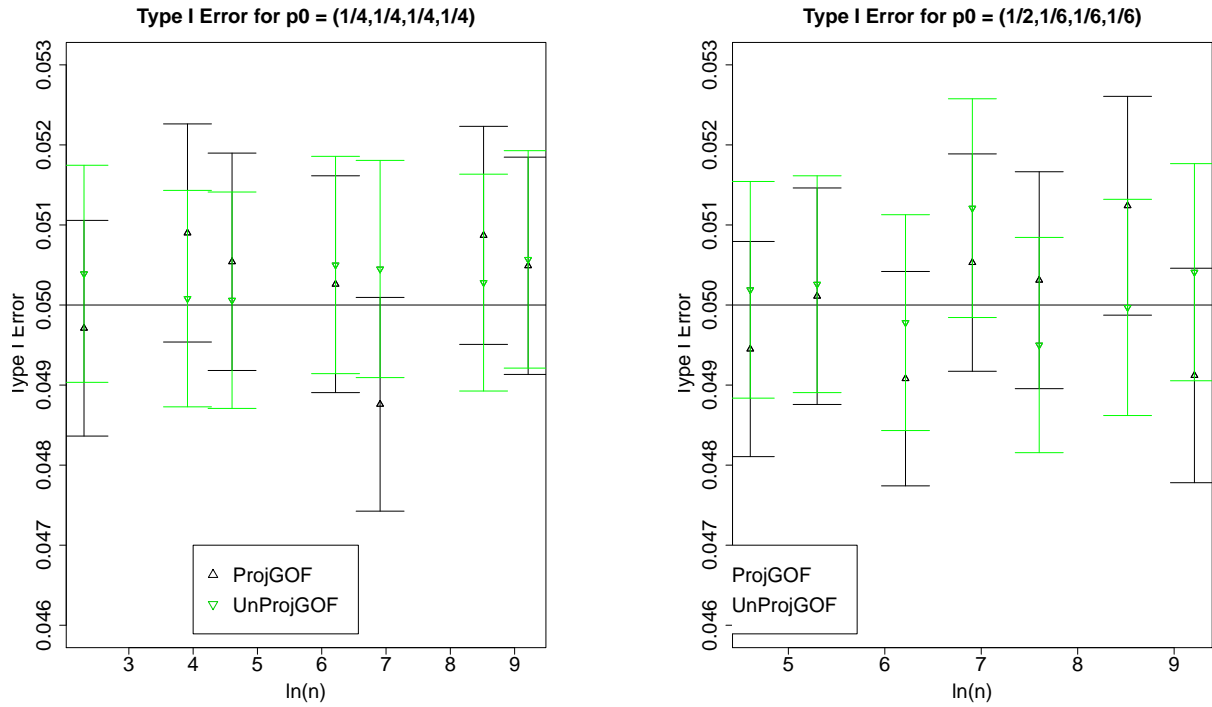


Figure 1: Empirical Type I Error for our new goodness of fit tests in  $\text{zCDP-GOF}$  with error bars corresponding to 1.96 times the standard error in 100,000 trials. We set  $\rho = 0.001$  which corresponds to a variance of 1,000 for the additional noise to the counts due to privacy. The horizontal line corresponds to the target  $\alpha = 0.05$  Type I error that we permit.

We then empirically check the power of our new tests in **zCDP-GOF** for both the projected and unprojected statistic. Subject to the constraint that our tests achieve Type I error at most  $\alpha$ , we seek to maximize *power*, or the probability of rejecting the null hypothesis when a distribution  $\mathbf{p}^1 \neq \mathbf{p}^0$ , called the *alternate hypothesis*, is true. We expect to see the projected statistic achieve higher power than the unprojected statistic due to Theorem 5.8. Further, the critical value we use for the projected statistic is smaller than the critical value for the unprojected statistic, which might improve the power of the projected statistic.

Here we present a typical experimental scenario. We set the null hypothesis  $\mathbf{p}^0 = (1/2, 1/6, 1/6, 1/6)$  and alternate hypothesis  $\mathbf{p}^1 = \mathbf{p}^0 + 0.01 \cdot (1, -1/3, -1/3, -1/3)$  for various sample sizes (we empirically found this to be a tough alternative hypothesis for our statistics). For each sample size  $n$ , we sample 5,000 independent datasets from the alternate hypothesis and test  $H_0 : \mathbf{p} = \mathbf{p}^0$  in **zCDP-GOF**. The resulting power plots are in Figure 2a for **zCDP-GOF** from Algorithm 1. We label “NonPrivate” as the classical chi-square goodness of fit test used on the actual data (and thus not private). Further, we write “ProjGOF” as the test from **zCDP-GOF** with the projected statistic whereas “UnProjGOF” uses the unprojected statistic. Clearly, the projected outperforms the unprojected statistic.

We then compare the projected and unprojected statistic in **zCDP-GOF** to prior work in Figure 2b. Since the projected statistic outperforms the other tests, we plot the difference in power between the projected statistic and the other tests. We label “GLRV\_MCGOF\_GAUSS” as the Monte-Carlo (MC) test with Gaussian noise from [12],<sup>6</sup> and “GLRV\_GOF\_Asympt” as the asymptotics-based test with Gaussian noise from [12, 24]. The error bars show 1.96 times the standard error in the difference of proportions from 100,000 trials, giving a 95% confidence interval.

## 6 General Chi-Square Private Tests

We now consider the case where the null hypothesis contains many distributions, so that the best fitting distribution must be estimated and used in the test statistics. The data is multinomial  $X^{(n)} \sim \text{Multinomial}(n, \mathbf{p}(\theta^0))$  and  $\mathbf{p}$  is a function that converts parameters into a  $d$ -dimensional multinomial probability vector. The null hypothesis is  $H_0 : \theta^0 \in \Theta$ ; i.e.  $\mathbf{p}(\theta^0)$  belongs to a subset of a lower-dimensional manifold. We again use Gaussian noise  $Z \sim N(0, 1/\rho \cdot I_d)$  to ensure  $\rho$ -zCDP, and we define

$$U_\rho^{(n)}(\theta) \stackrel{\text{defn}}{=} \sqrt{n} \left( \frac{X^{(n)} + Z}{n} - \mathbf{p}(\theta) \right). \quad (21)$$

With  $\theta^0$  being the unknown true parameter, we are now ready to define our two test statistics in terms of some function  $\phi : \mathbb{R}^d \rightarrow \mathbb{R}$ , such that  $\phi(X^{(n)} + Z) \xrightarrow{P} \theta^0$  (recall from Section 4.2 that  $\phi$  is a simple but possibly a suboptimal estimate of the true parameter  $\theta^0$  based on the noisy data) and the covariance matrix


$$\Sigma_\rho(\theta) \stackrel{\text{defn}}{=} \text{Diag}(\mathbf{p}(\theta)) - \mathbf{p}(\theta)\mathbf{p}(\theta)^\top + 1/\rho \cdot I_d.$$

We define the *unprojected* statistic  $R_\rho^{(n)}(\theta)$  as follows:


$$\begin{aligned} \widehat{M} &\stackrel{\text{defn}}{=} \left( \Sigma_{n\rho} \left( \phi(X^{(n)} + Z) \right) \right)^{-1} \\ R_\rho^{(n)}(\theta) &\stackrel{\text{defn}}{=} U_\rho^{(n)}(\theta)^\top \widehat{M} U_\rho^{(n)}(\theta). \end{aligned} \quad (22)$$

---

<sup>6</sup>We set the the number of MC trials  $m = 59$  in these experiments, which guarantees at most 5% Type I error.



(a) A comparison of power for goodness of fit testing between the projected and unprojected statistics in  $\mathbf{zCDP-GOF}$  with the classical nonprivate test for various  $n$  with 5,000 trials each.



(b) The empirical loss in power from using other private goodness of fit tests instead of using the projected statistic in  $\mathbf{zCDP-GOF}$  with error bars corresponding to 1.96 times the standard error of each difference for 100,000 trials.

Figure 2: Empirical power results for our new goodness of fit tests in  $\mathbf{zCDP-GOF}$  with  $\alpha = 0.05$  and a comparisons to previous private tests in [12]. We use  $\rho = 0.001$ , which corresponds to the variance of 1,000 for the additional noise to the counts.

This is a specialization of (4) in Section 4.2 with the following substitutions:  $V^{(n)} = \left(\frac{X^{(n)}+Z}{n}\right)$ ,  $A(\theta) = \mathbf{p}(\theta)$ , and  $M(\theta) = (\Sigma_{n\rho}(\theta))^{-1}$ .

For the *projected* statistic  $\mathcal{R}_\rho^{(n)}(\theta)$ , the corresponding substitutions are  $\mathbf{P} = I_d - \frac{1}{d}\mathbf{1}\mathbf{1}^\top$ ,  $V^{(n)} = \mathbf{P} \cdot \left(\frac{X^{(n)}+Z}{n}\right)$ ,  $A(\theta) = \mathbf{P} \cdot \mathbf{p}(\theta)$ , and again  $M(\theta) = (\Sigma_{n\rho}(\theta))^{-1}$  giving:

$$\mathcal{R}_\rho^{(n)}(\theta) \stackrel{\text{defn}}{=} U_\rho^{(n)}(\theta)^\top \cdot \mathbf{P}\widehat{M}\mathbf{P} \cdot U_\rho^{(n)}(\theta). \quad (23)$$

We then assume that for both the projected and unprojected statistic Assumption 4.1 holds using their relative vectors  $V^{(n)}$ ,  $A(\theta)$ , and matrix  $M(\theta)$ . We now present the asymptotic distribution of both statistics, which is proved using the result in Theorem 4.4 .

**Theorem 6.1.** *Under  $H_0 : \theta^0 \in \Theta$ , the following are true as  $n \rightarrow \infty$ . Setting  $\widehat{\theta}^{(n)} = \arg \min_{\theta \in \Theta} R_{\rho_n}^{(n)}(\theta)$  we have  $R_{\rho_n}^{(n)}(\widehat{\theta}^{(n)}) \xrightarrow{D} \chi_{d-k}^2$  if  $n\rho_n \rightarrow \rho > 0$ . Furthermore, setting  $\widehat{\theta}^{(n)} = \arg \min_{\theta \in \Theta} \mathcal{R}_{\rho_n}^{(n)}(\theta)$  we have  $\mathcal{R}_{\rho_n}^{(n)}(\widehat{\theta}^{(n)}) \xrightarrow{D} \chi_{d-k-1}^2$  if  $n\rho_n \rightarrow \rho$  or  $n\rho_n \rightarrow \infty$ .*

*Proof.* To prove this result, we appeal to Theorem 4.4. For the unprojected statistic  $R_\rho^{(n)}(\cdot)$  we have that  $C(\theta) = \Sigma_\rho(\theta)$  and the middle matrix  $M(\theta)$  is simply the inverse of it, which satisfies the hypotheses of Theorem 4.4.

For the projected statistic  $\mathcal{R}_{n\rho_n}^{(n)}(\cdot)$ , we will write  $C(\theta) = \mathbf{P} \cdot \Sigma_{n\rho_n}(\theta) \cdot \mathbf{P}$ ,  $M(\theta) = \Sigma_{n\rho_n}^{-1}(\theta)$ , and  $\dot{A}(\theta) = \mathbf{P} \cdot \nabla \mathbf{p}(\theta) \in \mathbb{R}^{d \times k}$ . Note that  $C(\theta)$  has rank  $d-1$  for all  $\theta \in \Theta$  in a neighborhood of  $\theta^0$  and all  $n$ . We will now show that we can satisfy the hypotheses in Theorem 4.4 with these matrices, i.e. we show the following two equalities hold for all  $\theta \in \Theta$

$$C(\theta) \cdot M(\theta) \cdot C(\theta) = C(\theta) \quad \& \quad C(\theta) \cdot M(\theta) \cdot \dot{A}(\theta) = \dot{A}(\theta).$$

We first focus on proving the first equality  $C(\theta) \cdot M(\theta) \cdot C(\theta) = C(\theta)$ . From (18), we can simplify the left hand side of the equality significantly by rewriting it as

$$\mathbf{P}\Sigma_{n\rho_n}(\theta)\mathbf{P} - \frac{n\rho_n}{d}\mathbf{P}\Sigma_{n\rho_n}(\theta)\mathbf{1}\mathbf{1}^\top\Sigma_{n\rho_n}(\theta)\mathbf{P}$$

We now show that  $\mathbf{P}\Sigma_{n\rho_n}(\theta)\mathbf{1}\mathbf{1}^\top\Sigma_{n\rho_n}(\theta) = 0$  for all  $n$ , which would prove this equality. Note that  $\Sigma_{n\rho_n}(\theta)$  is symmetric and has eigenvector  $\mathbf{1}$  with eigenvalue  $\frac{1}{n\rho_n}$ . Thus,

$$\mathbf{P}\Sigma_\rho(\theta)\mathbf{1}\mathbf{1}^\top\Sigma_{n\rho_n}(\theta) = \frac{1}{n^2\rho_n^2}\mathbf{P}\mathbf{1}\mathbf{1}^\top = 0 \quad \forall n.$$

We now prove the second equality  $C(\theta) \cdot M(\theta) \cdot \dot{A}(\theta) = \dot{A}(\theta)$ . We again use (18) to simplify the left hand side of the equality:

$$\begin{aligned} & \mathbf{P}\Sigma_{n\rho_n}(\theta) \left[ \Sigma_{n\rho_n}^{-1}(\theta) - \frac{n\rho_n}{d} \cdot \mathbf{1}\mathbf{1}^\top \right] \nabla \mathbf{p}(\theta) \\ &= \mathbf{P} \left[ I_d - \frac{n\rho_n}{d} \cdot \Sigma_{n\rho_n}(\theta)\mathbf{1}\mathbf{1}^\top \right] \nabla \mathbf{p}(\theta) \\ &= \mathbf{P}\mathbf{P}\nabla \mathbf{p}(\theta) \\ &= \mathbf{P}\nabla \mathbf{p}(\theta). \end{aligned}$$

This completes the proof for both cases  $n\rho_n \rightarrow \rho > 0$  and  $n\rho_n \rightarrow \infty$ . □

Again, the projected statistic has the same distribution under both private asymptotic regimes and matches the non-private chi-square test asymptotics. We present our more general test  $\mathbf{zCDP}\text{-Min-}\chi^2$  in Algorithm 2. The quick-and-dirty estimator  $\phi(\cdot)$  is application-specific (Section 6.1 gives independence testing as an example).<sup>7</sup> Further, for neighboring histogram data, we have the following privacy guarantee.

**Theorem 6.2.**  $\mathbf{zCDP}\text{-Min-}\chi^2(\cdot; \rho, \alpha, \phi, \Theta)$  is  $\rho$ -zCDP.

---

**Algorithm 2** zCDP General Chi-Square Test

---

**procedure**  $\mathbf{zCDP}\text{-Min-}\chi^2(X^{(n)}; \rho, \alpha, \phi, H_0 : \theta^0 \in \Theta)$

Set  $\tilde{X}^{(n)} \leftarrow X^{(n)} + Z$  where  $Z \sim N(0, 1/\rho \cdot I_d)$ .

Set  $\widehat{M} = \Sigma_{n\rho} \left( \phi(\tilde{X}^{(n)}) \right)^{-1}$

**For the *unprojected* statistic:**

$$T(\theta) = \frac{1}{n} \left( \tilde{X}^{(n)} - n\mathbf{p}(\theta) \right)^\top \widehat{M} \left( \tilde{X}^{(n)} - n\mathbf{p}(\theta) \right)$$

Set  $\widehat{\theta}^{(n)} = \operatorname{argmin}_{\theta \in \Theta} T(\theta)$

$t \leftarrow (1 - \alpha)$  quantile of  $\chi_{d-k}^2$

**For the *projected* statistic:**

$$T(\theta) = \frac{1}{n} \left( \tilde{X}^{(n)} - n\mathbf{p}(\theta) \right)^\top \mathbf{P} \widehat{M} \mathbf{P} \left( \tilde{X}^{(n)} - n\mathbf{p}(\theta) \right)$$

Set  $\widehat{\theta}^{(n)} = \operatorname{argmin}_{\theta \in \Theta} T(\theta)$

$t \leftarrow (1 - \alpha)$  quantile of  $\chi_{d-k-1}^2$

**if**  $T(\widehat{\theta}^{(n)}) > t$  **then** Reject

---

## 6.1 Application - Independence Test

We showcase our general chi-square test  $\mathbf{zCDP}\text{-Min-}\chi^2$  by giving results for independence testing. Conceptually, it is convenient to think of the data histogram as an  $r \times c$  table, with  $\mathbf{p}_{i,j}$  being the probability a person is in the bucket in row  $i$  and column  $j$ . We then consider two multinomial random variables  $Y \sim \text{Multinomial}(1, \pi^{(1)})$  for  $\pi^{(1)} \in \mathbb{R}^r$  (the marginal row probability vector) and  $Y' \sim \text{Multinomial}(1, \pi^{(2)})$  for  $\pi^{(2)} \in \mathbb{R}^c$  (the marginal column probability vector). Under the null hypothesis of independence between  $Y$  and  $Y'$ ,  $\mathbf{p}_{i,j} = \pi_i^{(1)} \pi_j^{(2)}$ . Generally, we write the probabilities as  $\mathbf{p}(\pi^{(1)}, \pi^{(2)}) = \pi^{(1)} (\pi^{(2)})^\top$  so that

$$X^{(n)} \sim \text{Multinomial} \left( n, \mathbf{p}(\pi^{(1)}, \pi^{(2)}) \right).$$

Thus we have the underlying parameter vector  $\theta^0 = \left( \pi_1^{(1)}, \dots, \pi_{r-1}^{(1)}, \pi_1^{(2)}, \dots, \pi_{c-1}^{(2)} \right)$  - we do not need the last component of  $\pi^{(1)}$  or  $\pi^{(2)}$  because we know that each must sum to 1. Also, we have

---

<sup>7</sup>For goodness-of-fit testing,  $\phi$  always returns  $\mathbf{p}^0$  and  $k = 0$  so  $\mathbf{zCDP}\text{-Min-}\chi^2$  is a generalization of  $\mathbf{zCDP}\text{-GOF}$ .

$d = rc$  and  $k = (r - 1) + (c - 1)$  in this case. We want to test whether  $Y$  is independent of  $Y'$ . For our data, we are given a collection of  $n$  independent trials of  $Y$  and  $Y'$ . We then count the number of joint outcomes in a contingency table given in Table 1. Each cell in the contingency table contains element  $X_{i,j}^{(n)}$  that gives the number of occurrences of  $Y_i = 1$  and  $Y'_j = 1$ . Since our test statistics notationally treat the data as a vector, when needed, we convert  $X^{(n)}$  to a vector that goes from left to right along each row of the table.

Table 1: Contingency Table.

$Y \setminus Y'$	1	2	$\dots$	$c$	Marginals
1	$X_{1,1}^{(n)}$	$X_{1,2}^{(n)}$	$\dots$	$X_{1,c}^{(n)}$	$X_{1,\cdot}^{(n)}$
2	$X_{2,1}^{(n)}$	$X_{2,2}^{(n)}$	$\dots$	$X_{2,c}^{(n)}$	$X_{2,\cdot}^{(n)}$
$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$	$\vdots$
$r$	$X_{r,1}^{(n)}$	$X_{r,2}^{(n)}$	$\dots$	$X_{r,c}^{(n)}$	$X_{r,\cdot}^{(n)}$
Marginals	$X_{\cdot,1}^{(n)}$	$X_{\cdot,2}^{(n)}$	$\dots$	$X_{\cdot,c}^{(n)}$	$n$

In order to compute the statistic  $R_\rho^{(n)}(\hat{\theta}^{(n)})$  or  $\mathcal{R}_\rho^{(n)}(\hat{\theta}^{(n)})$  in  $\mathbf{zCDP}\text{-Min-}\chi^2$ , we need to find a quick-and-dirty estimator  $\phi(X^{(n)} + Z)$  that converges in probability to  $\mathbf{p}(\pi^{(1)}, \pi^{(2)})$  as  $n \rightarrow \infty$ . We will use the estimator for the unknown probability vector based on the marginals of the table with noisy counts, so that for naïve estimates  $\tilde{\pi}_i^{(1)} = \frac{X_{i,\cdot}^{(n)} + Z_{i,\cdot}}{\tilde{n}}$ ,  $\tilde{\pi}_j^{(2)} = \frac{X_{\cdot,j}^{(n)} + Z_{\cdot,j}}{\tilde{n}}$  where  $\tilde{n} = n + \sum_{i,j} Z_{i,j}$  we have<sup>8</sup>

$$\phi(X^{(n)} + Z) = \left( \tilde{\pi}_1^{(1)}, \dots, \tilde{\pi}_{r-1}^{(1)}, \tilde{\pi}_1^{(2)}, \dots, \tilde{\pi}_{c-1}^{(2)} \right). \quad (24)$$

Note that as  $n \rightarrow \infty$ , the marginals converge in probability to the true probabilities even for  $Z \sim N(0, 1/\rho_n \cdot I_{rc})$  with  $\rho_n = \omega(1/n^2)$ , i.e. we have that  $\tilde{\pi}_i^{(1)} \xrightarrow{P} \pi_i^{(1)}$  and  $\tilde{\pi}_j^{(2)} \xrightarrow{P} \pi_j^{(2)}$  for all  $i \in [r]$  and  $j \in [c]$ . Recall that in Theorem 6.1, in order to guarantee the correct asymptotic distribution we require the  $n\rho_n \rightarrow \rho > 0$ , or in the case of the projected statistic, we need  $\rho_n = \Omega(1/n)$ . Thus, Theorem 6.1 imposes more restrictive settings of  $\rho_n$  for the unprojected statistic than what we need in order for the naïve estimate to converge to the true underlying probability. For the projected statistic, we only need  $\rho_n = \omega(1/n)$  to satisfy the conditions in Theorem 6.1 and for  $\phi(X^{(n)} + Z) \xrightarrow{P} \mathbf{p}(\pi^{(1)}, \pi^{(2)})$ .

We then use this statistic  $\phi(X^{(n)} + Z)$  in our unprojected and projected statistic in  $\mathbf{zCDP}\text{-Min-}\chi^2$  to have a  $\rho$ -zCDP hypothesis test for independence between two categorical variables. Note that in this setting, the projected statistic has a  $\chi_{(r-1)(c-1)}^2$  distribution, which is exactly the same asymptotic distribution used in the classical Pearson chi-square independence test.

For our results we will again fix  $\alpha = 0.05$  and  $\rho = 0.001$ . In Figure 3 we give the empirical Type I error for our independence tests given in  $\mathbf{zCDP}\text{-Min-}\chi^2$  for both the projected and unprojected statistic in 100,000 trials for various  $n$  and data distributions. We note that for small sample sizes we are achieving much smaller Type I Errors than the target  $\alpha$  due to the fact that sometimes the noise forces us to have small expected counts ( $< 5$  in any cell) in the contingency table based on the noisy counts, in which case our tests are inconclusive and fail to reject  $H_0$ .

<sup>8</sup>We note that in the case of small sample sizes, we follow a common rule of thumb where if any of the expected cell counts are less than 5, i.e. if  $n\tilde{\pi}_i^{(1)}\tilde{\pi}_j^{(2)} \leq 5$  for any  $(i, j) \in [r] \times [c]$ , then we do not make any conclusion.






Figure 3: Empirical Type I Error for our new independence tests in  $\mathbf{zCDP}\text{-Min-}\chi^2$  with 1.96 times the standard error in 100,000 trials. We set  $\rho = 0.001$  which corresponds to variance 1,000 due to noise in each cell count. It is desired to have Type I error at most  $\alpha = 0.05$ , which is given as the horizontal line.





(a) A comparison of power for independence testing between the projected and unprojected statistics in  $\text{zCDP-Min-}\chi^2$  with the classical non-private test. We set  $\alpha = 0.05$ ,  $\rho = 0.001$  and 5,000 to be the number of trials.



(b) The empirical power loss from using other private independence tests instead of using the projected statistic in  $\text{zCDP-Min-}\chi^2$  with error bars corresponding to 1.96 times the standard error of each difference in 50,000 trials. We set  $\alpha = 0.05$  and  $\rho = 0.001$ .

Figure 4: Empirical power results for our new independence tests in  $\text{zCDP-Min-}\chi^2$  and a comparisons to previous private tests in [12]

We then compare the power  $\text{zCDP-Min-}\chi^2$  achieves for either of our test statistics. As a sample of our experiments, we set  $r = c = 2$  and  $\pi^{(1)} = (2/3, 1/3), \pi^{(2)} = (1/2, 1/2)$ . We then sample our contingency table  $X^{(n)}$  from  $\text{Multinomial}(n, \mathbf{p}(\pi^{(1)}, \pi^{(2)}) + \mathbf{\Delta})$  where  $\mathbf{\Delta} = 0.01 \cdot (1, 0, -1, 0)$ , so that the null hypothesis is indeed false and should be rejected. We give the empirical power of  $\text{zCDP-Min-}\chi^2$  in Figure 4a using both the unprojected  $R_\rho^{(n)}(\hat{\theta}^{(n)})$  from (22) and projected statistic  $\mathcal{R}_\rho^{(n)}(\hat{\theta}^{(n)})$  from (23) for 5,000 independent trials and various sample sizes  $n$ . Note that again we pick  $\hat{\theta}^{(n)}$  from Theorem 4.3 relative to the statistic we use. We label “NonPrivate” as the classical Pearson chi-square test used on the actual data and “ProjIND” as the test from  $\text{zCDP-Min-}\chi^2$  with the projected statistic whereas “UnProjIND” uses the unprojected statistic.

The projected statistic again outperforms prior work, so in Figure 4b, we plot the difference in power between the projected statistic in  $\text{zCDP-Min-}\chi^2$  and the competitors (the unprojected statistic and independence tests from [12]) in 50,000 trials. Note that we label “GLRV\_MCIND\_GAUSS” (resp., “GLRV\_IND\_Asympt”) as the Monte Carlo (resp., asymptotics-based) test with Gaussian noise from [12].

## 6.2 Application - GWAS Testing

We next turn to demonstrating that our new general class of private hypothesis tests for categorical data significantly improves on existing private hypothesis tests even when extra structure is assumed about the dataset. Specifically, we will be interested in GWAS data, which was the primary reason for why hypothesis tests for independence should be made private [13]. We will then assume that  $r = 3$  and  $c = 2$  and the data is evenly split between the two columns - as is the case in a control and case group. For such tables, we can directly compute the sensitivity of the classical chi-square statistic  $q(\cdot)$ :

$$q(X^{(n)}) = \sum_{i=1}^3 \sum_{j=1}^2 \frac{n \cdot \left( X_{i,j}^{(n)} - \frac{X_{\cdot,j}^{(n)} \cdot X_{i,\cdot}^{(n)}}{n} \right)^2}{X_{\cdot,j}^{(n)} \cdot X_{i,\cdot}^{(n)}}$$

**Lemma 6.3** ([22, 27]). *The  $\ell_1$  and  $\ell_2$  global sensitivity of the chi-square statistic  $q(\cdot)$  based on a  $3 \times 2$  contingency table with positive margins and  $n/2$  cases and  $n/2$  controls is  $\Delta(q) = 4n/(n+2)$ .*

Hence, a different approach for a private independence test is to add Gaussian noise with variance  $\sigma^2 = \frac{\Delta_2(q)^2}{2\rho}$  to the statistic  $q(\cdot)$  itself, which we call *output perturbation*. Our statistic is then simply Gaussian mechanism  $\mathcal{M}_{\text{Gauss}}$  for statistic  $q$ . We then compare the private statistic value with the distribution of  $\mathcal{T}_{\text{Gauss}}(n, \rho) = \chi_2^2 + N(0, \sigma^2)$  where the degrees of freedom is 2 because we have  $(r-1) \cdot (c-1) = 2$ . Thus, given a Type I error of at most  $\alpha$ , we then set our critical value as  $\tau_{\text{Gauss}}(\alpha; n, \rho)$  where

$$\Pr[\mathcal{T}_{\text{Gauss}}(n, \rho) > \tau_{\text{Gauss}}(\alpha; n, \rho)] = \alpha$$

Hence, if  $\mathcal{M}_{\text{Gauss}}(X^{(n)})$  for the statistic  $q$  is larger than  $\tau_{\text{Gauss}}(\alpha; n, \rho)$  then we reject the null hypothesis.

For our experiments, we again set  $\rho = 0.001$  and  $\alpha = 0.05$ . We fix the probability vector  $(1/3, 1/3, 1/3)$  over the 3 rows in the first column whereas in the second column we set  $(1/2, 1/4, 1/4)$ , therefore the case and control groups do not produce the same outcomes. In Figure 5, we show a comparison in the power between our test with the projected statistic, which assumes no structure on the data, and the output perturbation test, which crucially relies on the fact that the data is evenly split between the case and control groups. We label “ProjIND” and “UnProjIND” as the tests from  $\text{zCDP-Min-}\chi^2$  with the projected statistic and unprojected statistic, respectively. Further, we label “YFSU\_Gauss” as the output perturbation tests for Gaussian noise proposed in [27]. Note that our new proposed test does significantly better than the output perturbation test, sometimes requiring 5 times more samples to achieve the same level of power than for our projected statistic test.

## 7 General Chi-Square Tests with Arbitrary Noise Distributions

We next show that we can apply our testing framework in Algorithm 2 for any type of noise distribution we want to include for privacy concerns. For example, we consider adding Laplace noise rather than Gaussian noise if our privacy benchmark were (pure) differential privacy (DP). In this case, we add Laplace noise with variance  $8/\epsilon^2$  when computing the two statistics  $R_{\epsilon^2/8}^{(n)}(\hat{\theta}^{(n)})$  from (22) and  $\mathcal{R}_{\epsilon^2/8}^{(n)}(\hat{\theta}^{(n)})$  from (23) so that the resulting tests will be  $\epsilon$ -DP and hence  $\frac{\epsilon^2}{2}$ -zCDP




Figure 5: A comparison of power between different hypothesis tests for independence testing for GWAS type datasets where the data is publicly known to be evenly split between the two columns and there are three rows in the contingency table.

from Theorem 3.5. Note that the resulting asymptotic distribution in this case will not be chi-square when we use noise other than Gaussian. We will then rely on Monte Carlo (MC) sampling to find the critical value in which to reject the null hypothesis. We give the MC based test which adds independent Laplace noise with variance  $8/\epsilon^2$  in Algorithm 3 and is thus  $\epsilon$ -DP, but any noise distribution can be used where we replace the parameter  $1/\rho$  in the two statistics to be the variance of the noise that is added to each count. In fact, Gaussian noise can be used in this framework although the asymptotic distribution seems to do well in practice even for small sample sizes.

## 7.1 Application - Goodness of Fit Testing

We first show that we can use the general chi-square test DP-MC-MIN with  $\epsilon$ -DP which uses Laplace noise in Algorithm 3 for goodness of fit testing  $H_0 : \mathbf{p} = \mathbf{p}^0$ . In this case we select  $\mathbf{p}(\hat{\theta}^{(n)}) = \mathbf{p}^0$  and  $\phi(X^{(n)} + Z) = \mathbf{p}^0$  in both the unprojected and projected statistics. From the way that we have selected the critical value  $\tau(\alpha, \epsilon)$  in Algorithm 3, we have the following result on Type I error, which follows directly from Theorem 5.3 in [12].

**Theorem 7.1.** *When the number of independent samples  $m$  we choose for our MC sampling is larger than  $1/\alpha$ , testing  $H_0 : \mathbf{p} = \mathbf{p}^0$  in Algorithm 3 guarantees Type I error at most  $\alpha$ .*

We then focus on empirically checking the power of DP-MC-MIN with  $\alpha = 0.05$  for the different statistics. As we did in the previous experiments, we will set the null hypothesis  $\mathbf{p}^0 = (1/2, 1/6, 1/6, 1/6)$  and alternate hypothesis  $\mathbf{p}^1 = \mathbf{p}^0 + 0.01 \cdot (1, -1/3, -1/3, -1/3)$  for various sample sizes. We set the privacy parameter  $\epsilon = \sqrt{2} \cdot .001 \approx 0.045$ , which implies  $(\rho = 0.001)$ -zCDP due

---

**Algorithm 3** DP Minimum Chi-Square Test using MC

---

**procedure** DP-MC-MIN(Histogram data  $X^{(n)} = (X_1^{(n)}, \dots, X_d^{(n)})$ ;  $\epsilon, \alpha, H_0 : \theta^0 \in \Theta, m$  trials)

Set  $\tilde{X}^{(n)} \leftarrow X^{(n)} + Z$  where  $Z = (Z_1, \dots, Z_d)$ , where  $Z_i \sim \text{Lap}(2/\epsilon)$ .

Set  $\widehat{M} = \Sigma_{n\rho} \left( \phi(\tilde{X}^{(n)}) \right)^{-1}$

**For the *unprojected* statistic:**

$$T(\theta) = \frac{1}{n} \left( \tilde{X}^{(n)} - n\mathbf{p}(\theta) \right)^\top \widehat{M} \left( \tilde{X}^{(n)} - n\mathbf{p}(\theta) \right)$$

Set  $\widehat{\theta}^{(n)} = \text{argmin}_{\theta \in \Theta} T(\theta)$

Sample  $\{r_1, \dots, r_m\}$  as  $m$  samples from the distribution of  $T(\widehat{\theta}^{(n)})$ .

Set  $\tau(\alpha, \epsilon)$  to be the  $\lceil (m+1)(1-\alpha) \rceil$ -largest value of  $\{r_1, \dots, r_m\}$ .

**For the *projected* statistic:**

$$T(\theta) = \frac{1}{n} \left( \tilde{X}^{(n)} - n\mathbf{p}(\theta) \right)^\top \mathbf{P} \widehat{M} \mathbf{P} \left( \tilde{X}^{(n)} - n\mathbf{p}(\theta) \right)$$

Set  $\widehat{\theta}^{(n)} = \text{argmin}_{\theta \in \Theta} T(\theta)$

Sample  $\{r_1, \dots, r_m\}$  as  $m$  samples from the distribution of  $T(\widehat{\theta}^{(n)})$ .

Set  $\tau(\alpha, \epsilon)$  to be the  $\lceil (m+1)(1-\alpha) \rceil$ -largest value of  $\{r_1, \dots, r_m\}$ .

**if**  $T(\widehat{\theta}^{(n)}) > \tau(\alpha, \epsilon)$  **then** Reject

---


to Theorem 3.5. We set the number of independent samples we draw from the distribution of the statistic under the null hypothesis as  $m = 59$ . In Figure 6a, we compare the power of the projected and unprojected statistic in DP-MC-GOF, labeled “ProjGOF\_LAP” and “UnProjGOF\_LAP” respectively, with the classical non-private chi-square test for various  $n$  each with 5,000 trials. Note that there is a drastically larger power when we use the projected statistic as opposed to the unprojected statistic.

We then show that the projected statistic using Laplace noise achieves significantly higher power than using the other unprojected test statistic as well as previous DP hypothesis tests with Laplace noise from [12]. We then label “GLRV\_MCGOF\_LAP” as the MC based test with Laplace noise from [12] and plot the power loss in Figure 6b that the other DP goodness of fit tests suffer when compared to the power that DP-MC-GOF achieves with the projected statistic. Note that the error bars in the figure show 1.96 times the standard error in the difference of proportions from 100,000 trials.

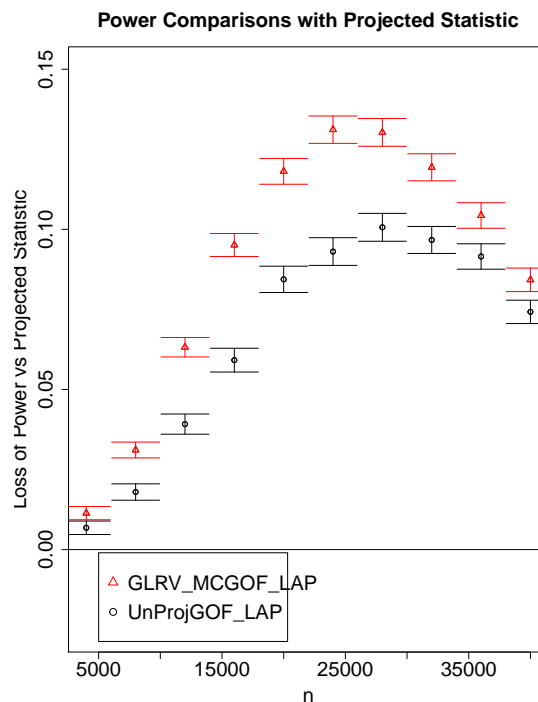
## 7.2 Application - Independence Testing

We then apply our general framework to independence testing as in Section 6.1. Unlike our goodness of fit testing, we are not guaranteed to have Type I error at most  $\alpha$  when we have composite tests, e.g. independence testing, in DP-MC-MIN because we are not sampling from the exact data distribution.

We then empirically show the Type I error is at most the desired level  $\alpha = 0.05$ . We again fix  $\epsilon = \sqrt{2 * 0.001} \approx 0.045$ , which ensures  $\epsilon$ -DP as well as  $(\rho = 0.001)$ -zCDP due to Theorem 3.5. We will use  $m = 59$  samples in all of our MC testing. We then give the empirical Type I error for



(a) A comparison of power for goodness of fit testing between the projected and unprojected statistics in DP-MC-MIN with the classical nonprivate test for various  $n$  with 5,000 trials each.



(b) The empirical loss in power from using other private goodness of fit tests instead of using the projected statistic in DP-MC-GOF. The error bars correspond to 1.96 times the standard error of each difference for 100,000 trials.

Figure 6: Empirical power results for our new DP goodness of fit tests in DP-MC-MIN and a comparisons to a previous private test in [12] that uses Laplace noise with variance 4,000 added to each cell count.




Figure 7: Empirical Type I Error for the new DP independence tests in DP-MC-MIN with 1.96 times the standard error in 10,000 trials. We set  $\epsilon = \sqrt{.001 * 2} \approx 0.045$  which corresponds to variance 1,000 due to noise for each cell count. It is desired to have Type I error at most  $\alpha = 0.05$ , which is given as the horizontal line.

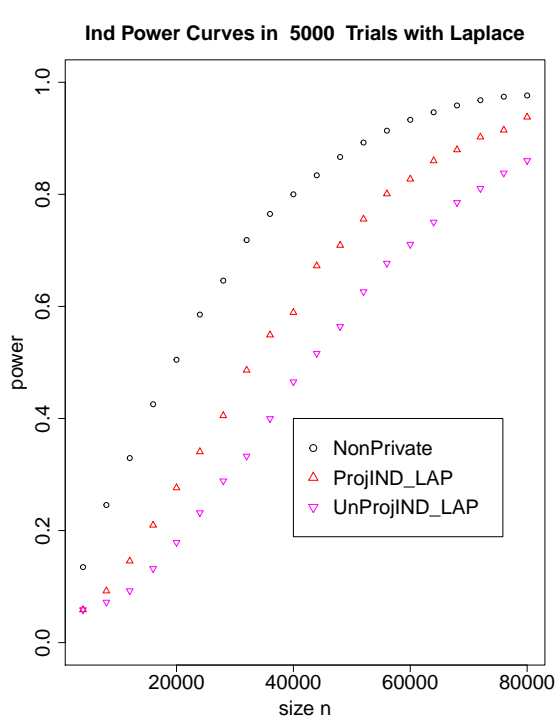
various  $n$  and data distributions in Figure 7. Note that we use the same rule of thumb as before where if our naïve estimate for the probability distribution produces expected cell counts smaller than 5, then our test is inconclusive and fails to reject. This is why in our experiments, the Type I error is close to zero for small sample sizes.

We also consider the power of our tests in DP-MC-MIN. As we did before, we will set the data distribution with  $\pi^{(1)} = (2/3, 1/3), \pi^{(2)} = (1/2, 1/2)$ . We then sample our contingency table  $X^{(n)}$  from  $\text{Multinomial}(n, \mathbf{p}(\pi^{(1)}, \pi^{(2)}) + \Delta)$  where  $\Delta = 0.01 \cdot (1, 0, -1, 0)$  for various sample sizes. In Figure 8a, we compare the power of the projected and unprojected statistic in DP-MC-MIN, labeled “ProjIND\_LAP” and “UnProjIND\_LAP” respectively, with the classical non-private chi-square test for 1,000 trials.


We then label “GLRV\_MCIND\_LAP” as the MC based test with Laplace noise from [12]. Note that the error bars show 1.96 times the standard error in the difference of proportions from 10,000 trials, giving a 95% confidence interval.

## 8 Conclusions

We have demonstrated a new broad class of private hypothesis tests  $\mathbf{zCDP}\text{-Min-}\chi^2$  for categorical data based on the minimum chi-square theory. We gave two statistics (*unprojected* and *projected*)



(a) A comparison of power for independence testing between the projected and unprojected statistics in DP-MC-MIN with the classical nonprivate test.



(b) The empirical loss in power from using other private goodness of fit tests instead of using the projected statistic in DP-MC-MIN with error bars corresponding to 1.96 times the standard error of each difference for 10,000 trials.

Figure 8: Empirical power results for our new DP independence tests in DP-MC-MIN and a comparisons to a previous private test in [12] that uses Laplace noise.

that converge to a chi-square distribution when we use Gaussian noise and thus lead to zCDP hypothesis tests. Unlike prior work, these statistics have the same asymptotic distributions in the private asymptotic regime as the classical chi-square tests have in the classical asymptotic regime.

Our simulations show that with both statistics our tests achieve at most  $\alpha$  Type I error . Empirically, the test using the projected statistic significantly improves the Type II error when compared to the unprojected statistic and prior work [12]. Further, our new tests give comparable power to the classical (nonprivate) chi-square tests. We then gave further applications to GWAS data and other privacy-preserving noise distributions (e.g. Laplace).

## 9 Acknowledgments

This work was partially supported by NSF grant 1228669.



## References

- [1] R. Bassily, K. Nissim, A. D. Smith, T. Steinke, U. Stemmer, and J. Ullman. Algorithmic stability for adaptive data analysis. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21*, pages 1046–1059, 2016.
- [2] Y. M. M. Bishop, S. E. Fienberg, and P. W. Holland. *Discrete multivariate analysis: Theory and practice*, 1975.
- [3] M. Bun and T. Steinke. Concentrated differential privacy: Simplifications, extensions, and lower bounds. *ArXiv e-prints*, May 2016.
- [4] R. Cummings, K. Ligett, K. Nissim, A. Roth, and Z. S. Wu. Adaptive learning with robust generalization guarantees. In *Proceedings of the 29th Conference on Learning Theory, COLT 2016, New York, USA, June 23-26, 2016*, pages 772–814, 2016.
- [5] C. Dwork and G. N. Rothblum. Concentrated differential privacy. *CoRR*, abs/1603.01887, 2016.
- [6] C. Dwork, K. Kenthapadi, F. McSherry, I. Mironov, and M. Naor. Our data, ourselves: Privacy via distributed noise generation. In *Proceedings of the 24th Annual International Conference on The Theory and Applications of Cryptographic Techniques, EUROCRYPT’06*, pages 486–503, Berlin, Heidelberg, 2006. Springer-Verlag.
- [7] C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In *TCC ’06*, pages 265–284, 2006.
- [8] C. Dwork, V. Feldman, M. Hardt, T. Pitassi, O. Reingold, and A. Roth. Preserving statistical validity in adaptive data analysis. In *STOC*, 2015.
- [9] C. Dwork, V. Feldman, M. Hardt, T. Pitassi, O. Reingold, and A. Roth. Generalization in adaptive data analysis and holdout reuse. In *Advances in Neural Information Processing Systems*, 2015.
- [10] T. Ferguson. *A Course in Large Sample Theory*. Chapman & Hall Texts in Statistical Science Series. Taylor & Francis, 1996. ISBN 9780412043710.
- [11] S. E. Fienberg, A. Rinaldo, and X. Yang. Differential privacy and the risk-utility tradeoff for multi-dimensional contingency tables. In *Proceedings of the 2010 International Conference on Privacy in Statistical Databases, PSD’10*, pages 187–199, Berlin, Heidelberg, 2010. Springer-Verlag.
- [12] M. Gaboardi, H. Lim, R. M. Rogers, and S. P. Vadhan. Differentially private chi-squared hypothesis testing: Goodness of fit and independence testing. In *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, pages 2111–2120, 2016.
- [13] N. Homer, S. Szlinger, M. Redman, D. Duggan, W. Tembe, J. Muehling, J. V. Pearson, D. A. Stephan, S. F. Nelson, and D. 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.
- [14] M. F. D. John G. Reid. An accessible proof of craig’s theorem in the noncentral case. *The American Statistician*, 42(2):139–142, 1988. ISSN 00031305. URL <http://www.jstor.org/stable/2684489>.
- [15] A. Johnson and V. 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.
- [16] V. Karwa and A. Slavković. Differentially private graphical degree sequences and synthetic graphs. In J. Domingo-Ferrer and I. Tinnirello, editors, *Privacy in Statistical Databases*, volume 7556 of *Lecture Notes in Computer Science*, pages 273–285. Springer Berlin Heidelberg, 2012.
- [17] V. Karwa and A. Slavković. Inference using noisy degrees: Differentially private  $\beta$ -model and synthetic graphs. *Ann. Statist.*, 44(P1):87–112, 02 2016.
- [18] R. Rogers, A. Roth, A. Smith, and O. Thakkar. Max-information, differential privacy, and post-selection hypothesis testing. In *Proceedings of the 57th Annual IEEE Symposium on Foundations of Computer Science, New Brunswick, NJ, USA, October 9 - 11*, pages 487–494, 2016.
- [19] O. Sheffet. Differentially private least squares: Estimation, confidence and rejecting the null hypothesis. *arXiv preprint arXiv:1507.02482*, 2015.
- [20] S. Simmons, C. Sahinalp, and B. Berger. Enabling privacy-preserving {GWASs} in heterogeneous human populations. *Cell Systems*, 3(1):54 – 61, 2016.
- [21] A. Smith. Privacy-preserving statistical estimation with optimal convergence rates. In *Proceedings of the Forty-third Annual ACM Symposium on Theory of Computing*, STOC ’11, pages 813–822, New York, NY, USA, 2011. ACM.
- [22] C. Uhler, A. Slavkovic, and S. E. Fienberg. Privacy-preserving data sharing for genome-wide association studies. *Journal of Privacy and Confidentiality*, 5(1), 2013.
- [23] D. Vu and A. Slavković. Differential privacy for clinical trial data: Preliminary evaluations. In *Proceedings of the 2009 IEEE International Conference on Data Mining Workshops*, ICDMW ’09, pages 138–143, Washington, DC, USA, 2009. IEEE Computer Society.
- [24] Y. Wang, J. Lee, and D. Kifer. Differentially private hypothesis testing, revisited. *CoRR*, abs/1511.03376, 2015.
- [25] L. Wasserman and S. Zhou. A statistical framework for differential privacy. *Journal of the American Statistical Association*, 105(489):375–389, 2010.
- [26] M. A. Woodbury. *Inverting Modified Matrices*. Number 42 in Statistical Research Group Memorandum Reports. Princeton University, Princeton, NJ, 1950.

- [27] F. Yu, S. E. Fienberg, A. B. Slavković, and C. Uhler. Scalable privacy-preserving data sharing methodology for genome-wide association studies. *Journal of Biomedical Informatics*, 50:133–141, 2014.