Appendix A. Appendix

A.1. Detailed Discussion of Project Lighthouse

They consider a dataset of users with an anonymous ID and a feature vector, in their case the number of reservations made and the number of reservations rejected. An independent council labels each ID with a race/ethnicity using profile pictures and names. The features are made less granular to ensure that there are at least k users with the same features, and then joined with the IDs. If a particular feature has all the same race/ethnicity, then ℓ-diversity is used to change some of the groups. The resulting dataset can then be used to determine whether rejection rate is independent of race/ethnicity. It is not immediately clear how statistical tests might be modified to detect differences between groups, as sometimes k-anonymous features might result in very wide ranges of values. If it is later determined that the feature vector was not rich enough, then the whole procedure needs to be redone. Additionally, if some change had been implemented, we would like to know whether the rejection rate difference in groups had changed. Unfortunately, k-anonymous groups do not satisfy any composition property, so doing this procedure again with different features and hence different groups might allow someone to be identified.

A.2. Binary Outcomes

We include detailed derivations of the corrected Z-test and general χ^2 test, giving some additional results on performance for these tests.

There are traditionally two ways to test whether the success probabilities are different between two groups, either with a Z-test or with a χ^2-test. In fact, many built in proportion test packages in Python and R use χ^2, rather than Z-tests, as the default to determine p-values.[4] One useful property of the Z-test, compared to the χ^2 test is that it can readily be used to compute confidence intervals for the difference in proportions. Furthermore, the difference between the Z and χ^2 tests comes from the experimental design. For Z-test, we fix two sample sizes: n_1 for group 1 and n_2 for group 2. While for χ^2-tests, we sample n data points where there is some probability π of being in group 1, and otherwise probability $1-\pi$ of being in group 2. It is the latter experimental design that we will be primarily interested in, since we will assume that we do not know which group each person belongs to initially. However, we do analyze Z-tests to see if they can still be used to compute valid confidence intervals for the difference in proportions between two groups.

A.3. Z-test

We will assume that we have a binary variable and two groups. We will have data $\{(G_i, X_i)\}_{i=1}^n$, where we first sample which group each X_i belongs to, which we model with $G_i \sim \mathrm{Bern}(\pi)+1$ for an unknown $\pi \in [0,1]$, then we have $X_i | G_i \sim \mathrm{Bern}(p_{G_i})$, where $p_g \in [0,1]$ is the probability of success for group $g \in \{1,2\}$. One common test we may want to conduct is to test the null hypothesis $\mathtt{H}_0 : p_1 = p_2 + \Delta$. We will write $N[1] = \sum_{i=1}^n 1G_i = 1$, $N[2] = n - N[1]$, and $\bar{X}[g] = \frac{1}{N[g]} \sum_{i=1}^{N[g]} X_i \cdot 1G_i = g$ for $g \in \{1,2\}$. To carry out the

4. In fact, the `prop.test` method in R actually computes the χ^2 statistic, see RDocumentation.

difference in proportions test, we would form the following test statistic.[5]

$$T(\bar{X}[1], \bar{X}[2], N[1], N[2]; \Delta) = \frac{\bar{X}[1] - \bar{X}[2] - \Delta}{\sqrt{\bar{X}[1](1 - \bar{X}[2])/N[1] + \bar{X}[2](1 - \bar{X}[2])/N[2]}}$$

We then compare the test statistic with its asymptotic distribution under the null hypothesis, which is a standard normal. That is for significance level $1 - \alpha$, ($\alpha = 0.05$ or 0.01 typically, but we will use $\alpha = 0.05$ throughout), we reject $\mathtt{H}_0$ if

$$T(\bar{X}[1], \bar{X}[2], N[1], N[2]; \Delta) \notin [\Phi^{-1}(\alpha/2), \Phi^{-1}(1 - \alpha/2)].$$

We now turn to privatizing the group membership with a mechanism M so that our samples are $\{(M(G_i), X_i)\}_{i=1}^n$ and then analyze the resulting Z-test. We now change the set up where each individual belonging to group $g \in \{1, 2\}$ can be switched to the other group $3 - g$ with some probability $\leq \frac{1}{2}$. In particular, for ε-LGDP, we can use randomized response $M : \{1, 2\} \to \{1, 2\}$ where $\Pr[M(g) = g] = \frac{e^\varepsilon}{1+e^\varepsilon}$. Hence the privatized data for individual i in group g is a mixture of two Bernoulli's, which we write as

$$X_i^\varepsilon|(G_i = g) = \tfrac{e^\varepsilon}{1+e^\varepsilon}\mathrm{Bern}(p_g) + \tfrac{1}{1+e^\varepsilon}\mathrm{Bern}(p_{3-g}) = \mathrm{Bern}(\tfrac{e^\varepsilon}{1+e^\varepsilon}p_g + \tfrac{1}{1+e^\varepsilon}p_{3-g}). \tag{5}$$

Note that the sample sizes of each group changes due to some outcomes from group g switching to group $3 - g$. We will write the new (randomized) sample sizes for each group g as $N^\varepsilon[g]$. Note that $n = N^\varepsilon[1] + N^\varepsilon[2]$. Hence, the number of successful outcomes that we see in group g is then

$$\bar{X}^\varepsilon[g] = \sum_{i=1}^{N^\varepsilon[g]} X_i^\varepsilon \cdot 1M(G_i) = g.$$

Note that in the special case when $\Delta = 0$, we have $\mathtt{H}_0 : p_0 = p_1 = p$, which in this case we would still have $X_i^\varepsilon[g] \sim \mathrm{Bern}(p)$ where $p_1 = p_2 = p$. Carrying out the standard test statistic would give us $T(\bar{X}^\varepsilon[1], \bar{X}^\varepsilon[2], N^\varepsilon[1], N^\varepsilon[2]; \Delta = 0)$. The main difference now is that the number of samples $N^\varepsilon[g]$ in each group g is randomized. We then check to see if conducting the original Z-test as if there were no privacy, still provides valid results. When $\Delta \neq 0$, we compute the expected difference between the two proportions,

$$\mathbb{E}\left[\tfrac{1}{N^\varepsilon[1]} \sum_{j=1}^{N^\varepsilon[1]} X_i^\varepsilon \cdot 1M(G_i) = 1 - \tfrac{1}{N^\varepsilon[2]} \sum_{j=1}^{N^\varepsilon[2]} X_i^\varepsilon \cdot 1M(G_i) = 2\right]$$

However, the expectation becomes much more complicated, due to the random number of trials $N^\varepsilon[g]$. To approximate this expectation, we try treating $N^\varepsilon[g]$ as fixed for $g \in \{1, 2\}$, to get the following expression, which can then be used to *correct* confidence intervals.

$$\Delta^\varepsilon = \left(\tfrac{n\pi e^\varepsilon}{(1+e^\varepsilon)N^\varepsilon[1]} - \tfrac{n\pi}{(1+e^\varepsilon)N^\varepsilon[2]}\right)\Delta \tag{6}$$

5. If $\Delta = 0$, i.e. no difference and equal variances in both groups, the test statistic can use the *pooled* variance, which would result in the following term in the denominator of the test statistic: $\sqrt{\frac{N[1]\cdot\bar{X}[1]+N[2]\cdot\bar{X}[2]}{N[1]+N[2]}\left(1/N[1] + 1/N[2]\right)}$.

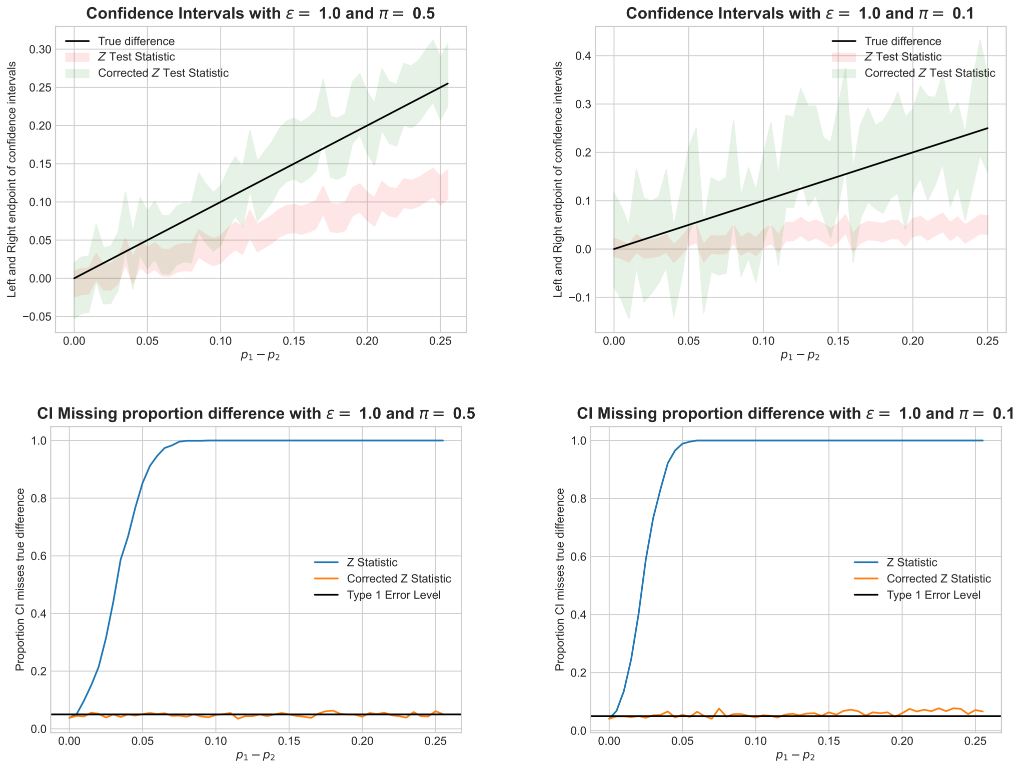

Figure 4: (Top Row) Confidence Intervals with $n = 10000$, $p_2 = 0.25$, and various $p_1 = p_2 + \Delta$ for $\varepsilon = 1.0$. We show that correcting Δ to be Δ^ε in (6) helps achieve valid confidence intervals for $\pi \in \{0.1, 0.5\}$. (Bottom Row) We use the same parameter settings as in the top row, but we compute the proportion of times the computed confidence interval actually overlaps with the true difference over 1000 independent trials.

Since π is not known, we can estimate it in the following way to get an unbiased estimator for it:

$$\hat{\pi} = \left(\frac{e^\varepsilon + 1}{e^\varepsilon - 1}\right)\left(\frac{N^\varepsilon[1] + N^\varepsilon[2]}{n} - \frac{1}{e^\varepsilon + 1}\right).$$

We then consider the following Z-statistic, $T(\bar{X}^\varepsilon[1], \bar{X}^\varepsilon[2], N^\varepsilon[1], N^\varepsilon[2]; \Delta = \Delta^\varepsilon)$, to account for privacy. We compute confidence intervals in Figure 4 with level of significance $1 - \alpha = 95\%$. Observe that this correction gives us valid confidence intervals (top panel), and appropriate type I error rates (bottom panel), compared to the uncorrected statistic.

A.4. Comparison with Fully Local DP

We plot power curves in Figure 5 (left, center) for our the private test, testing with $\Delta = 0$ and increasing $p_1 - p_2 > 0$ in the sample, showing the number of trials that rejected $\mathtt{H}_0 : p_1 = p_2$ over 1000 independent trials. We then compare with the local DP χ^2-test from Gaboardi

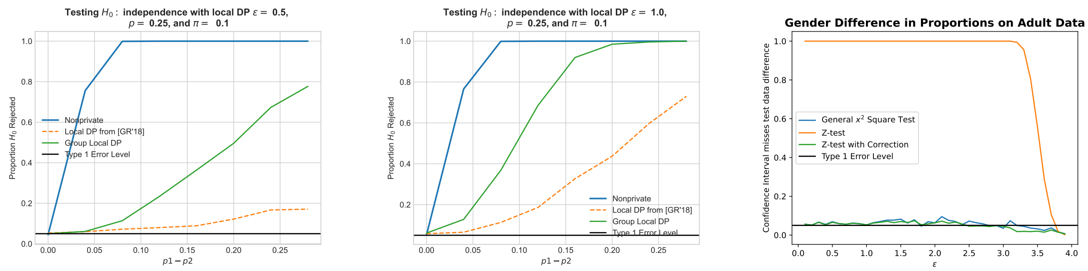

Figure 5: Comparing Group Local DP with (fully) LDP χ^2 tests from Gaboardi and Rogers (2018) with the probability of being in Group 1 being $\pi = 0.1$ and $\varepsilon \in \{0.5, 1\}$ and $n = 10000$ (left, center). Proportion of trials (right panel) on the adult dataset where the training data difference in proportion confidence interval misses the test data difference in proportion.

and Rogers (2018). Note that we can drastically improve the statistical power of our tests for the same level of ε in the less restrictive LGDP model compared with fully local DP tests.

We also present results on the UCI Adult dataset (Dua and Graff, 2017). Since we are working with two sensitive groups, we will use gender as the sensitive group category, where a sample is either *Male* or *Female*. We test whether there is a significant difference in whether males or females make more than 50k salary. We then compute the (non-private) sample difference in proportion on the adult test data. We then compare the traditional proportion test that ignores the additional noise due to privacy, the Z-test with the correction given in our paper, and our test in determining confidence intervals at the 95% significance level (Figure 5, right panel).

A.5. χ^2 test

Given that many statistical packages use χ^2 tests, rather than Z-tests, for testing the difference between two proportions, we then consider privatizing χ^2-tests and further using these tests to compute valid confidence intervals. There has been several works on privatizing χ^2 tests, even in the more restrictive (fully) local DP setting (Gaboardi and Rogers, 2018; Sheffet, 2018; Acharya et al., 2019a). We then compare the relevant tests in the local DP setting with the tests in this less restrictive LGDP setting. We will adopt the general minimum χ^2 theory outlined earlier from Kifer and Rogers (2017), which was also used to derive tests in the local DP setting by Gaboardi and Rogers (2018).

A.5.1. Non-private test

Recall Table 1, where Group 1 has probability of success p_1 and Group 2 has probability of success p_2. We then want to test $\mathtt{H}_0 : p_1 = p_2 + \Delta$. Recall also that when we sample n combined outcomes over groups 1 and 2, we can then consider a single multinomial random variable $Y = (Y[1,1], Y[1,2], Y[2,1], Y[2,2])^\intercal \sim \text{Multinom}(n, \boldsymbol{\theta}(\pi, p_1, p_2))$ where we flatten the outcome probabilities to an array

$$\boldsymbol{\theta}(\pi, p_1, p_2) = \left(\underbrace{\pi p_1}_{\text{Success in Group 1}}, \underbrace{(1-\pi)p_2}_{\text{Success in Group 2}}, \underbrace{\pi(1-p_1)}_{\text{Failure in Group 1}}, \underbrace{(1-\pi)(1-p_2)}_{\text{Failure in Group 2}} \right)^\intercal .$$

We use the χ^2 statistic by providing an estimate for the parameters p_1, p_2, and π, with $p_1 - p_2 = \Delta$ given. Note that Y can be written as the sum of i.i.d. random variables $\{Y_i\}$, i.e. $Y = \sum_{i=1} Y_i$. The covariance matrix $C(\theta)$ for Y_i is the following, where we write $\text{Diag}(\theta)$ to be the diagonal matrix with entries on the diagonal as θ.

$$C(\pi, p_1.p_2) = \text{Diag}\left(\boldsymbol{\theta}(\pi, p_1, p_2)\right) - \boldsymbol{\theta}(\pi, p_1, p_2)\boldsymbol{\theta}(\pi, p_1, p_2)^\intercal \tag{7}$$

Note that the covariance matrix is singular and has the all one's vector in its null space. It turns out that $\text{Diag}\left(\boldsymbol{\theta}(\pi, p_1, p_2)\right)^{-1}$ is the generalized inverse for $C(\pi, p_1, p_2)$ Ferguson (1996), which we will use in our χ^2 statistic. We then use the estimates $\hat{p}_1, \hat{p}_2$ and $\hat{\pi}$ for p_1, p_2, and π respectively where

$$\hat{p}_2 = \frac{Y[1,1] + Y[1,2]}{n} - \Delta\hat{\pi}, \quad \hat{p}_1 = \hat{p}_2 + \Delta, \qquad \hat{\pi} = \frac{Y[1,1] + Y[2,1]}{n} \tag{8}$$

The χ^2-statistic $\hat{D}$ then becomes the following,

$$\hat{D}(Y;\Delta) = n \cdot \left(Y/n - \boldsymbol{\theta}(\hat{\pi}, \hat{p}_1, \hat{p}_2)\right)^\intercal \text{Diag}\left(\boldsymbol{\theta}(\hat{\pi}, \hat{p}_1, \hat{p}_2)\right)^{-1} \left(Y/n - \boldsymbol{\theta}(\hat{\pi}, \hat{p}_1, \hat{p}_2)\right) \tag{9}$$

We then compare $\hat{D}(Y;\Delta)$ with a χ^2 with 1 degree of freedom, that is, if $\hat{D}(Y,\Delta) > \chi^2_{1,1-\alpha}$, then we reject $H_0 : p_1 = p_2 + \Delta$ with significance level $1-\alpha$. Note that this classical hypothesis test fits with the general χ^2 test outlined in Section 3 as Y actually has rank 3, due to the all ones vector being in its null space, and $\text{Diag}\left(\boldsymbol{\theta}(\hat{\pi}, \hat{p}_1, \hat{p}_2)\right)^{-1}$ is the generalized inverse of the covariance matrix evaluated at the estimates given.

One way to achieve valid confidence intervals for the difference $p_1 - p_2$ is to test for multiple values of Δ to see which intervals should be rejected under $\mathtt{H}_0$. That is, we search over the space $\Delta \in [-1, 1]$, with some tolerance level τ (say $\tau = 0.001$), and check whether $\hat{D}(Y;\Delta) \leq \chi^2_{1,1-\alpha}$. As we move from $\Delta = -1$, we will cross a point $\Delta = \Delta^L$ where $\hat{D}(Y;\Delta^L) > \chi^2_{1,1-\alpha}$, yet $\hat{D}(Y;\Delta^L + \tau) \leq \chi^2_{1,1-\alpha}$. This value Δ^L will be our left-end point of our confidence interval. We then continue searching until we reach a point $\Delta = \Delta^R$ where $\hat{D}(Y;\Delta^R - \tau) > \chi^2_{1,1-\alpha}$ yet $\hat{D}(Y;\Delta^R) > \chi^2_{1,1-\alpha}$. This value Δ^R will be our right-end point of our confidence interval. This simple grid search can also be replaced with a bisection root finding approach to the left and right of the Δ that minimizes the χ^2 statistic. We will use this method to compute confidence intervals with privatized groups in the following section and give very similar confidence intervals to the Z-test.

A.6. Confidence Intervals

We have shown that the classical Z-test may not need modifying if we test $\mathtt{H}_0 : p_1 = p_2$, i.e., it still provides valid results. We can also use the Z-test with a correction on the difference $\Delta \to \Delta^\varepsilon$ from (6) to compute valid confidence intervals, as in Figure 6. We also show that

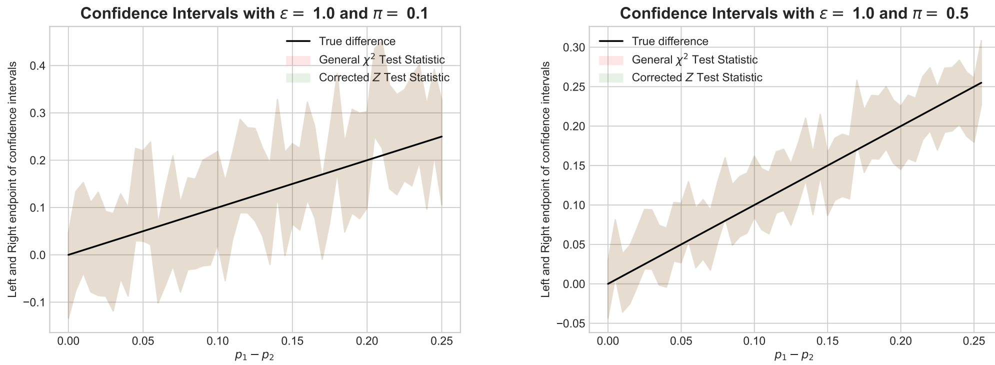

Figure 6: Comparing Confidence Intervals for Group Local DP between using the corrected Z-test and the χ^2 tests with $p_1 = 0.25$ and changing $p_1 - p_2$. The left plot uses group 1 probability $\pi = 0.10$ and the right plot uses $\pi = 0.5$. Note that we are showing 95% confidence intervals a single trial by generating data with different $p_1 - p_2$.

we can use the approach outlined above that uses the χ^2-test to compute the end points of a confidence interval. We then show the results in two cases, when $\pi = 1/2$ and when $\pi = 0.1$. Note that there is not much difference between the confidence intervals using the χ^2-test, and the Z-test with the correction factor.

Appendix B. Independence Testing with Categorical Data

We now consider testing whether the success probability is equal across several groups simultaneously. A common test is to use the classical Pearson χ^2 test for independence to see whether the outcome is independent of the group. Given that there are multiple groups, rather than just 2, we will compare the three private mechanisms presented in Section 2. To design private χ^2 tests for determining whether the success probability is the same across $g > 2$ groups simultaneously, we follow the general χ^2 test approach outlined in Section 3, which was also used to design (fully) local DP χ^2 tests in Gaboardi and Rogers (2018). We will compare these tests with those we develop in the less restrictive private setting of LGDP.

We first set up some notation. Let $Y_i \sim \text{Multinomial}(1, \theta(\pi, p))$ be the data entry for individual i, where $p \in [0, 1]$ is the success probability across all groups and $\pi \in [0, 1]^g$ is the probability vector over all g groups, so that $\sum_{i=1}^{g} \pi_i = 1$. Note that the covariance matrix of Y_i will still be of the form of the covariance matrix given in (7). Let $W_i \sim \text{Multinomial}(1, \pi)$, which will determine which group each sample belongs to, and let the outcome for a sample in group $j \in [g]$ be written as $X_i[j] \sim \text{Bern}(p_j)$. Hence, we have the following random vector that we use in the χ^2 statistic.

$$Y_i = \Bigg(\underbrace{W_i[1] \cdot X_i[1], W_i[2] \cdot X_i[2], \cdots, W_i[g] \cdot X_i[g]}_{\text{successes}}, \underbrace{W_i[1] \cdot (1 - X_i[1]), W_i[2] \cdot (1 - X_i[2]), \cdots, W_i[g] \cdot (1 - X_i[g])}_{\text{failures}} \Bigg)^{\intercal}.$$

To match the contingency table format we had in the previous section, we will write the coordinates of $Y_i = (Y_i[1,1], Y_i[1,2], \cdots, Y_i[1,g], Y_i[2,1], \cdots, Y_i[2,g])$, so that the entries whose first index is 1 are the successes and the entries whose first index is 2 are the failures. Once we use a privatization mechanism M, whether it be g-randomized response, bit flipping, or the subset mechanism, the resulting outcome will be Y_i^ε, which will consist of some successes/failures from group j being potentially replicated across various groups, and perhaps removed from group j. We then receive n i.i.d. samples from the distribution of Y_i^ε to obtain counts of successes and failures in each group $Y^\varepsilon = \sum_{i=1}^n Y_i^\varepsilon$. To make the randomness in the privacy mechanism explicit, we will write the mechanism as a matrix of noise terms multiplied by the original Y_i. With each mechanism $M(W_i) \in \{0,1\}^g$, we write the random matrix $Z_i^\varepsilon \in \{0,1\}^{g \times g}$ where column j is the corresponding random entries for $M(j)$. We can then succinctly write Y_i^ε in the following way where $\mathbf{0}$ is a g by g matrix of zeros:

$$Y_i^\varepsilon = \begin{bmatrix} Z_i^\varepsilon & \mathbf{0} \\ \mathbf{0} & Z_i^\varepsilon \end{bmatrix} Y_i \tag{10}$$

Here we can distinguish our work from the local DP setting considered in Gaboardi and Rogers (2018). In particular, the matrix multiplying Y_i in the (fully) local DP setting would include terms where there is a zero block submatrix. This would correspond to successes within group j being able to switch to a failure in another group j'. In our setting, we are not privatizing outcomes, i.e. successes or failures, hence the zero block matrices.

With our general set up, we can then compare the various privacy mechanisms for various levels of ε. That is, given privatized data Y_i^ε for $i \in [n]$ along with the privacy level ε and the mechanism M used, we then form the χ^2 statistic $\hat{D}$ from (**??**) and compare with the critical value $\chi^2_{g-1,1-\alpha}$ for $1 - \alpha$ level of significance.

B.1. Randomized Response

For g-randomized response, the matrix Z_i^ε will have column j be a multinomial of a single trial with probability vector that is $\frac{e^\varepsilon}{e^\varepsilon + g - 1}$ in position j and $\frac{1}{e^\varepsilon + g - 1}$ in every other coordinate. Note that the resulting privatized data Y^ε will still follow a multinomial distribution, where the probability vector is the following

$$\boldsymbol{\theta}^\varepsilon(\pi, p) = \begin{pmatrix} \frac{p}{e^\varepsilon + g - 1} \begin{bmatrix} e^\varepsilon & 1 & \cdots & 1 \\ 1 & e^\varepsilon & \cdots & 1 \\ & & \ddots & \\ 1 & 1 & \cdots & e^\varepsilon \end{bmatrix} \pi \\ \\ \frac{1-p}{e^\varepsilon + g - 1} \begin{bmatrix} e^\varepsilon & 1 & \cdots & 1 \\ 1 & e^\varepsilon & \cdots & 1 \\ & & \ddots & \\ 1 & 1 & \cdots & e^\varepsilon \end{bmatrix} \pi \end{pmatrix}$$

Note that because the resulting vector Y^ε is still a multinomial, we do not need to write out the covariance matrix in order to calculate the general χ^2 statistic, but rather the generalized inverse is the diagonal matrix whose entries are the inverse of the entries in $\boldsymbol{\theta}^\varepsilon(\pi, p)$. Note that we have $\boldsymbol{\theta}^\varepsilon(\pi, p) = \boldsymbol{\theta}(\pi^\varepsilon, p)$, where $\pi^\varepsilon = \left(\frac{e^\varepsilon}{e^\varepsilon + g - 1}\pi_j + \frac{1}{e^\varepsilon + g - 1}(1 - \pi_j) : j \in [g]\right)$, so we can use the same test statistic as in the non-private case, but with different group probabilities that can be estimated from the privatized samples. We then need to find estimates for p and π^ε based on the data Y^ε. We then use the generalized version of the estimate provided in (8) with $\Delta = 0$.

$$\hat{p} = \frac{\sum_{j \in [g]} Y^\varepsilon[1, j]}{n}, \qquad \hat{\pi}^\varepsilon = \left(\hat{\pi}^\varepsilon_j = \frac{Y^\varepsilon[1, j] + Y^\varepsilon[2, j]}{n} : j \in [g]\right).$$

We can then form the χ^2 statistic $\hat{D}$ as in (3), where $C(\hat{\theta}_n)^\dagger$ is the diagonal matrix whose entries are the inverse of $\boldsymbol{\theta}(\hat{\pi}^\varepsilon, \hat{p})$, which we write as $\mathrm{Diag}(\boldsymbol{\theta}(\hat{p}, \hat{\pi}^\varepsilon))^{-1}$

$$\hat{D} = \min_{\substack{p \in (0,1), \pi \in [0,1]^g \\ \text{s.t. } \sum_{j=1}^g \pi_j = 1}} \left\{ n \left(Y^\varepsilon/n - \boldsymbol{\theta}(\pi, p)\right)^\intercal \mathrm{Diag}(\boldsymbol{\theta}(\hat{\pi}^\varepsilon, \hat{p}))^{-1} \left(Y^\varepsilon/n - \boldsymbol{\theta}(\pi, p)\right) \right\}$$

We then compare $\hat{D}$ with a χ^2 distribution with $g - 1$ degrees of freedom, as in the non-private test, because the rank of the covariance matrix is at most $2g - 1$ and we are minimizing over g variables (p and $\pi_1, \cdots, \pi_{g-1}$ since $\pi_g = 1 - \sum_{j=1}^{g-1} \pi_j$).

B.2. Bit Flipping

The bit flipping mechanism will result in a random vector Y_i^ε that is not a multinomial, so a little more care will be needed in computing the general χ^2 statistic. The matrix of noise terms in (10) for the bit flipping mechanism will consist of the following entries,

$$Z_i^\varepsilon[j, j] \sim \mathrm{Bern}\left(\frac{e^{\varepsilon/2}}{e^{\varepsilon/2} + 1}\right), \quad j \in [g] \quad \text{and} \quad Z_i^\varepsilon[j, \ell] \sim \mathrm{Bern}\left(\frac{1}{e^{\varepsilon/2} + 1}\right), \quad j \neq \ell.$$

We will first compute $E[Y_i^\varepsilon] = \boldsymbol{\theta}^\varepsilon(\pi, p)$ when we privatize the group with the bit flipping mechanism

$$\boldsymbol{\theta}^\varepsilon(\pi,p) = \begin{pmatrix} \frac{p}{e^{\varepsilon/2}+1} \begin{bmatrix} e^{\varepsilon/2} & 1 & \cdots & 1 \\ 1 & e^{\varepsilon/2} & \cdots & 1 \\ \vdots & \vdots & \ddots & \vdots \\ 1 & 1 & \cdots & e^{\varepsilon/2} \end{bmatrix} \pi \\ \\ \frac{1-p}{e^{\varepsilon/2}+1} \begin{bmatrix} e^{\varepsilon/2} & 1 & \cdots & 1 \\ 1 & e^{\varepsilon/2} & \cdots & 1 \\ \vdots & \vdots & \ddots & \vdots \\ 1 & 1 & \cdots & e^{\varepsilon/2} \end{bmatrix} \pi \end{pmatrix}.$$

We now compute the covariance matrix, $C(\pi, p; \varepsilon) = \mathbb{E}[Y_i^\varepsilon \ (Y_i^\varepsilon)^\intercal] - \boldsymbol{\theta}^\varepsilon(\pi,p)\boldsymbol{\theta}^\varepsilon(\pi,p)^\intercal$

$$\mathbb{E}[Y_i^\varepsilon[1,j]^2] = \frac{p}{e^{\varepsilon/2}+1}(\pi_j e^{\varepsilon/2} + (1-\pi_j)), \quad \mathbb{E}[Y_i^\varepsilon[2,j]^2] = \frac{1-p}{e^{\varepsilon/2}+1}(\pi_j e^{\varepsilon/2} + (1-\pi_j)), \qquad \forall j \in [g].$$

$$\begin{aligned} \mathbb{E}[Y_i^\varepsilon[1,j] \cdot Y_i^\varepsilon[1,\ell]] &= \frac{p}{(e^{\varepsilon/2}+1)^2}\left(e^{\varepsilon/2}(\pi_j + \pi_\ell) + (1 - \pi_j - \pi_\ell)\right), && j,\ell \in [g], j \neq \ell \\ \mathbb{E}[Y_i^\varepsilon[2,j] \cdot Y_i^\varepsilon[2,\ell]] &= \frac{1-p}{(e^{\varepsilon/2}+1)^2}\left(e^{\varepsilon/2}(\pi_j + \pi_\ell) + (1 - \pi_j - \pi_\ell)\right), && j,\ell \in [g], j \neq \ell \\ \mathbb{E}[Y_i^\varepsilon[1,j] \cdot Y_i^\varepsilon[2,\ell]] &= \mathbb{E}[Y_i^\varepsilon[2,j] \cdot Y_i^\varepsilon[1,\ell]] = 0 && j,\ell \in [g], j \neq \ell \end{aligned}$$

Unlike the case for g-randomized response, the all ones vector is not in the null space. Gaboardi and Rogers (2018) showed that with the bit flipping algorithm in the local DP setting, the all ones vector is an eigenvector, whose eigenvalue depends solely on the privacy loss parameter ε. Using a technique from Kifer and Rogers (2017), they showed how one can project out this eigenvector and the resulting χ^2 statistic will have one fewer degree of freedom (asymptotically). We will not be able to do a similar technique here in the LGDP setting, since the all ones vector is not an eigenvector for general π, p. Hence, we will not be able to reduce the degrees of freedom in its asymptotic distribution, at least with similar techniques although it might be possible another way.

Next, we need to compute estimates for p and π.

$$\hat{p} = \frac{(e^{\varepsilon/2}+1)\sum_{j=1}^g Y^\varepsilon[j]}{n\left(e^{\varepsilon/2}+g-1\right)}, \qquad \hat{\pi} = \left(\frac{\frac{Y^\varepsilon[j]+Y^\varepsilon[g+j]}{n} - \frac{1}{e^{\varepsilon/2}+1}}{\frac{2e^{\varepsilon/2}}{e^{\varepsilon/2}+1} - 1} : j \in [g] \right).$$

We then plug in our estimates to the covariance matrix and form our χ^2 statistic, as in (3), and compare it to a χ^2 distribution with $2g - g = g$ degrees of freedom (one larger than the non-private version due to the covariance matrix being non-singular).

B.3. Subset Mechanism

The subset mechanism Ye and Barg (2017) from Definition 5 can also be used to privatize the group membership of each sample i, which takes an additional parameter $k < g$.

Column j of Z_i^ε will correspond to the outcome of the subset mechanism $M(j)$. That is. $Z_i^\varepsilon[j,j] \sim \text{Bern}(\frac{ke^\varepsilon}{ke^\varepsilon+g-k})$, and then the other coordinates $Z_i^\varepsilon[\ell,j]$ for $\ell \neq j$ will depend on the realization of $Z_i^\varepsilon[j,j]$. So if $Z_i^\varepsilon[j,j]=1$, then $(Z_i^\varepsilon[\ell,j] : \ell \neq j)$ will sample $k-1$ ones uniformly at random without replacement, while if $Z_i^\varepsilon[j,j]=0$, then $(Z_i^\varepsilon[\ell,j] : \ell \neq j)$ will sample k ones uniformly at random without replacement. Following our framework, we first compute $\boldsymbol{\theta}^\varepsilon(\pi,p;k) = \mathbb{E}[Y_i^\varepsilon]$ when we use the subset mechanism.

$$\boldsymbol{\theta}^\varepsilon(\pi,p;k) = \begin{pmatrix} \frac{p}{\binom{g-1}{k-1}e^\varepsilon + \binom{g-1}{k}} \begin{bmatrix} \binom{g-1}{k-1}e^\varepsilon & \left(\binom{g-2}{k-2}e^\varepsilon + \binom{g-2}{k-1}\right) & \cdots & \left(\binom{g-2}{k-2}e^\varepsilon + \binom{g-2}{k-1}\right) \\ \left(\binom{g-2}{k-2}e^\varepsilon + \binom{g-2}{k-1}\right) & \binom{g-1}{k-1}e^\varepsilon & \cdots & \left(\binom{g-2}{k-2}e^\varepsilon + \binom{g-2}{k-1}\right) \\ \vdots & \vdots & \ddots & \vdots \\ \left(\binom{g-2}{k-2}e^\varepsilon + \binom{g-2}{k-1}\right) & \left(\binom{g-2}{k-2}e^\varepsilon + \binom{g-2}{k-1}\right) & \cdots & \binom{g-1}{k-1}e^\varepsilon \end{bmatrix} \pi \\ \\ \frac{1-p}{\binom{g-1}{k-1}e^\varepsilon + \binom{g-1}{k}} \begin{bmatrix} \binom{g-1}{k-1}e^\varepsilon & \left(\binom{g-2}{k-2}e^\varepsilon + \binom{g-2}{k-1}\right) & \cdots & \left(\binom{g-2}{k-2}e^\varepsilon + \binom{g-2}{k-1}\right) \\ \left(\binom{g-2}{k-2}e^\varepsilon + \binom{g-2}{k-1}\right) & \binom{g-1}{k-1}e^\varepsilon & \cdots & \left(\binom{g-2}{k-2}e^\varepsilon + \binom{g-2}{k-1}\right) \\ \vdots & \vdots & \ddots & \vdots \\ \left(\binom{g-2}{k-2}e^\varepsilon + \binom{g-2}{k-1}\right) & \left(\binom{g-2}{k-2}e^\varepsilon + \binom{g-2}{k-1}\right) & \cdots & \binom{g-1}{k-1}e^\varepsilon \end{bmatrix} \pi \end{pmatrix}$$

We then compute the covariance matrix $C(\pi,p;\varepsilon,k) = \mathbb{E}[Y_i^\varepsilon (Y_i^\varepsilon)^\intercal] - \boldsymbol{\theta}^\varepsilon(\pi,p)\boldsymbol{\theta}^\varepsilon(\pi,p)^\intercal$

$$\begin{aligned} \mathbb{E}[Y_i^\varepsilon[1,j]^2] &= p\frac{\binom{g-1}{k-1}e^\varepsilon \pi_j + \left(\binom{g-2}{k-2}e^\varepsilon + \binom{g-2}{k-1}\right)(1-\pi_j)}{\binom{g-1}{k-1}e^\varepsilon + \binom{g-1}{k}}, && \forall j \in [g]. \\ \mathbb{E}[Y_i^\varepsilon[2,j]^2] &= (1-p)\frac{\binom{g-1}{k-1}e^\varepsilon \pi_j + \left(\binom{g-2}{k-2}e^\varepsilon + \binom{g-2}{k-1}\right)(1-\pi_j)}{\binom{g-1}{k-1}e^\varepsilon + \binom{g-1}{k}}, && \forall j \in [g]. \end{aligned}$$

Now for $\ell \neq j$ and $j,\ell \in [g]$,

$$\begin{aligned} \mathbb{E}[Y_i^\varepsilon[1,j]\cdot Y_i^\varepsilon[1,\ell]] &= \frac{p}{\binom{g-1}{k-1}e^\varepsilon + \binom{g-1}{k}}\left(e^\varepsilon\binom{g-2}{k-2}(\pi_j+\pi_\ell) + \left(e^\varepsilon\binom{g-3}{k-3} + \binom{g-3}{k-2}\right)(1-\pi_j-\pi_\ell)\right) \\ \mathbb{E}[Y_i^\varepsilon[2,j]\cdot Y_i^\varepsilon[2,\ell]] &= \frac{1-p}{p}\mathbb{E}[Y_i^\varepsilon[1,j]Y_i^\varepsilon[1,\ell]], \\ \mathbb{E}[Y_i^\varepsilon[1,j]\cdot Y_i^\varepsilon[2,\ell]] &= \mathbb{E}[Y_i^\varepsilon[2,j]Y_i^\varepsilon[1,\ell]] = 0 \end{aligned}$$

We then consider the rank of this covariance matrix. Note that the all ones vector is in the null space of the covariance matrix.

Lemma 6 *The covariance matrix $C(\pi,p;\varepsilon,k)$ corresponding to the subset mechanism has the all ones vector $\mathbf{1} \in \mathbb{R}^{2g}$ in its null space, i.e.*

$$C(\pi,p;\varepsilon,k)\mathbf{1} = \mathbf{0}$$

Proof Consider element $j \in [g]$ in $\mathbb{E}[Y_i^\varepsilon (Y_i^\varepsilon)^\intercal]\mathbf{1}$, where we ignore the coefficient $\frac{p}{\binom{g-1}{k-1}e^\varepsilon + \binom{g-1}{k}}$,

$$\begin{aligned}
&\pi_j \left(e^\varepsilon \binom{g-1}{k-1} + \binom{g-1}{k} \right) + (1-\pi_j)\left(e^\varepsilon \binom{g-2}{k-2} + \binom{g-2}{k-1} \right) \\
&\qquad + \sum_{j \neq \ell} \left((\pi_j + \pi_\ell) e^\varepsilon \binom{g-2}{k-2} + (1 - \pi_j - \pi_\ell) \left(e^\varepsilon \binom{g-3}{k-3} + \binom{g-3}{k-2} \right) \right) \\
&= \pi_j e^\varepsilon \left(\binom{g-1}{k-1} + (g-1)\binom{g-2}{k-2} \right) + (1-\pi_j) e^\varepsilon \left(2\binom{g-2}{k-2} + (g-2)\binom{g-3}{k-3} \right) \\
&\qquad + (1-\pi_j)\left(\binom{g-2}{k-1} + (g-2)\binom{g-3}{k-2} \right)
\end{aligned}$$

We then use the following properties

$$x\binom{x-1}{y-1} = y\binom{x}{y}, \qquad x\binom{x-1}{y} = (y+1)\binom{x}{y+1}$$

Hence, after simplifying, we get

$$\begin{aligned}
E[Y_i^\varepsilon (Y_i^\varepsilon)^\intercal]\mathbf{1} &= \frac{1}{\binom{g-1}{k-1}e^\varepsilon + \binom{g-1}{k}} \left[\pi_j e^\varepsilon k \binom{g-1}{k-1} + (1-\pi_j)\left(e^\varepsilon k \binom{g-2}{k-2} + k\binom{g-2}{k-1} \right) \right] \\
&= k\boldsymbol{\theta}^\varepsilon(\pi, p; k)
\end{aligned}$$

Furthermore, we have

$$\boldsymbol{\theta}^\varepsilon(\pi, p; k)^\intercal \mathbf{1} = k.$$

Putting everything together, we have

$$C(\pi, p; \varepsilon, k)\mathbf{1} = \left(\mathbb{E}[Y_i^\varepsilon (Y_i^\varepsilon)^\intercal] - \boldsymbol{\theta}^\varepsilon(\pi, p)\boldsymbol{\theta}^\varepsilon(\pi, p)^\intercal\right)\mathbf{1} = k\boldsymbol{\theta}^\varepsilon(\pi, p; k) - k\boldsymbol{\theta}^\varepsilon(\pi, p; k) = \mathbf{0}.$$

■

Next, we compute estimates for p and π based on the sample Y^ε as well as ε and k.

$$\begin{aligned}
\hat{p} &= \frac{\left(\binom{g-1}{k-1}e^\varepsilon + \binom{g-1}{k}\right) \cdot \sum_{j=1}^g Y[1,j]^\varepsilon}{n\left(\binom{g-1}{k-1}e^\varepsilon + (g-1)\left(\binom{g-2}{k-2}e^\varepsilon + \binom{g-2}{k-1}\right)\right)} = \frac{\sum_{j=1}^g Y[1,j]^\varepsilon}{nk}, \\
\hat{\pi} &= \left(\frac{\left(\binom{g-1}{k-1}e^\varepsilon + \binom{g-1}{k}\right) \cdot \left(\frac{Y[1,j]^\varepsilon + Y[2,j]^\varepsilon}{n}\right) - \left(e^\varepsilon\binom{g-2}{k-2} + \binom{g-2}{k-1}\right)}{e^\varepsilon\left(\binom{g-1}{k-1} - \binom{g-2}{k-2}\right) - \binom{g-2}{k-1}} : j \in [g] \right).
\end{aligned}$$

We can then use the χ^2 statistic in (3) with the covariance matrix $C(\pi, p; \varepsilon, k)$ that we computed along with the estimates $\hat{p}$ and $\hat{\pi}$. Note that we showed in Lemma 6 that the covariance is not full rank, so we will use $2g-1$ as the rank of the covariance matrix and $g-1$ degrees of freedom in the χ^2 distribution in our test.

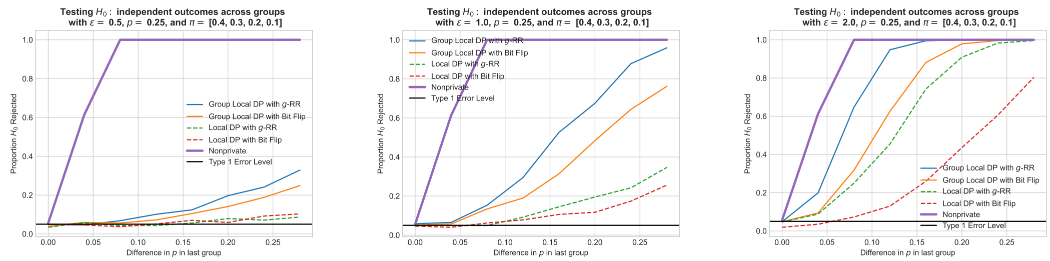

Figure 7: Comparing non-private (top bold line), Group Local DP (solid lines), and original LDP χ^2 tests (dashed lines) for independence testing with the probability vector over the $g = 4$ groups to be $\pi = (0.4, 0.3, 0.2, 0.1)$ with various ε and $n = 10000$.

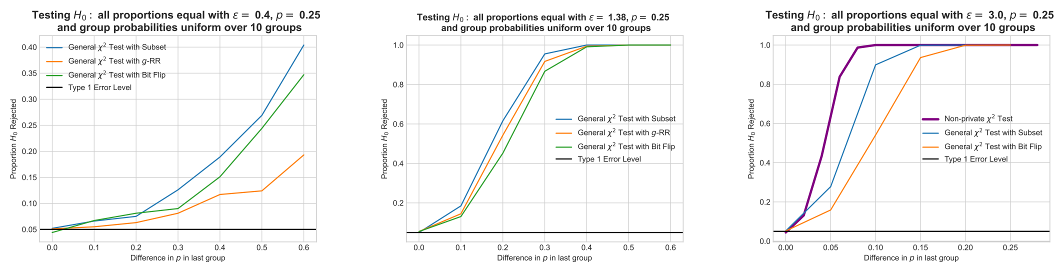

Figure 8: Comparing various Local Group DP mechanisms with corresponding χ^2 test for testing whether there is a difference in success probability across different sensitive groups with various ε and $n = 10000$. Note that with $\varepsilon = 3$, we get $k = \lceil g/(e^\varepsilon + 1) \rceil = 1$, which is equivalent to g-randomized response.

B.4. Results

In our results, we start by comparing with the existing (fully) local DP χ^2 tests for independence from Gaboardi and Rogers (2018) in Figure 7. As stated earlier, the difference between the local DP and group local DP setting is that in the latter the outcomes (successes or failures) cannot be changed but they can in the former. We can see that the power can be drastically improved in the less restrictive model with the same privacy loss parameter ε.

We include experiments for the subset mechanism Ye and Barg (2017), which were not considered before for local DP χ^2 independence tests. The main takeaway for the subset mechanism is that it seems to strictly dominate over g-randomized response and bit flipping for the same level of privacy, while in Gaboardi and Rogers (2018) there were privacy levels where the g-randomized response algorithm outperformed bit flipping for the high ε regime and vice versa in the low ε regime. See Figure 8 for experiments with various parameter settings where $g = 10$ and π is uniform over the g groups. We can see that the subset mechanism dominates the other two at various different privacy levels, which is to be expected as the subset mechanism is known to be optimal for certain tasks. Recall that as ε gets large, the Subset Mechanism and g-randomized response are the same.

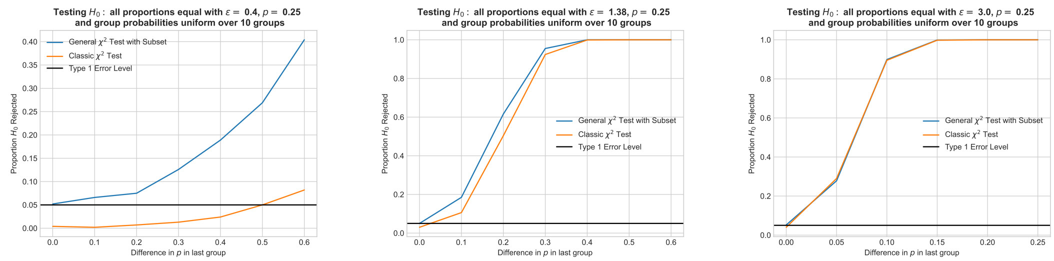

Figure 9: Comparing various Local Group DP mechanisms with corresponding χ^2 test for testing whether there is a difference in success probability across different sensitive groups with various ε and $n = 10000$. Note that with $\varepsilon = 3$, we get $k = \lceil g/(e^{\varepsilon} + 1) \rceil = 1$, which is equivalent to g-randomized response.

The tests we develop achieve higher empirical power than simply using the classical χ^2 tests after privatizing the groups. We present plots in Figure 9 that shows for data generated with the subset mechanism at various privacy levels, the general χ^2 test that accounts for the subset mechanism outperforms using the classic χ^2 test, which does not account for the privacy mechanism. We point out that when ε gets larger, the two tests seem to perform similarly. All plots consist of the proportion of times the null hypothesis was rejected over 1000 trials.

We also evaluate our method on the UCI Adult dataset Dua and Graff (2017) `adult.data`, where we will use Race as the sensitive group and the binary outcome as whether a sample makes more than \$50k salary. Note that race contains 5 groups, with labels White, Asian-Pac-Islander, Amer-Indian-Eskimo, Other, and Black. We now want to arrive at the same conclusion after privatizing the race of each sample as we would if we had not privatized it. Figure 10 gives our results, which considers various levels of privacy and for each privacy level we compute 1000 independent trials of the subset mechanism on each sample's race and use our general χ^2 test while comparing it to the traditional χ^2 test for independence, which ignores the privacy mechanism. We see that we can achieve more power for stronger levels of privacy, but since there is such a strong difference between proportions in the groups, i.e. we should reject the null hypothesis that there is no difference in proportions across all groups, for even moderate levels of privacy we arrive at the same conclusion as the non-private test almost all the time.

Appendix C. T-tests

Recall the classical t-test to test the difference between two means between samples $\{X_i[j]\}_{i=1}^{n_j} \overset{i.i.d.}{\sim} \mathtt{N}\left(\mu_j, \sigma_j^2\right)$ for $j \in \{1,2\}$. That is, we use the t-test statistic defined as follows where s_1^2, s_2^2 are the sample variances for groups 1 and 2, respectively.

$$T = \frac{\frac{1}{n_1}\sum_{i=1}^{n_1} X_i[1] - \frac{1}{n_2}\sum_{i=1}^{n_2} X_i[2]}{\sqrt{s_1^2/n_1 + s_2^2/n_2}}$$

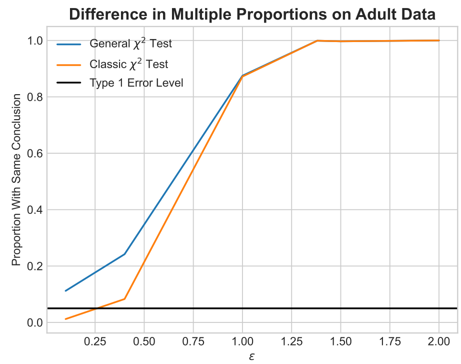

Figure 10: Proportion of times out of 1000 independent trials that we reject the null hypothesis of there being no difference in proportions across all groups on the UCI Adult dataset. We use race as the sensitive group and the outcome is whether a sample makes more thant \$50k salary or not.

We would then compare the test statistic to a t distribution, which converges to a standard Gaussian if n_1, n_2 are large. In this section, we show how to build the corresponding χ^2 test and how to modify the t-test, and give more detailed results.

C.1. A χ^2 Test for Difference in Means

For binary outcomes, we were able to move from the Z-test to a χ^2 test, and in fact χ^2 tests are typically used instead of Z-tests in many common statistical packages. However, for continuous outcomes, there is not a standard χ^2 version of the t-test. We then present a way to formulate the t-test as a χ^2 and show similar performance. Note that for binary outcomes, we could form a contingency table where the rows were outcomes (success/failure) and the columns were groups (group 1 or 2). To fit this framework, a first approach would be to discretize the outcomes into bins, and form the contingency table with r different rows, where r is a predetermined value for the number of bins the outcomes will be placed in. Unfortunately, this introduces additional complexity to the hypothesis test when adding privacy, and it is not clear how to discretize the outcome set and how this might impact statistical power.

Instead, we present a way to form a contingency table for continuous outcomes without discretizing, by considering the moments of the samples, which we presented in Table 2. It is easy to see that the contingency table for binary outcomes in Table 1, where instead of having failure outcomes, we could replace it with the 0th order moment, which would give the marginals. Recall that we use $W_i \sim \text{Bern}(\pi)$ to determine the group of sample i and $X_i[j] \sim \mathtt{N}\left(\mu_j, \sigma_j^2\right)$ for $j \in \{1, 2\}$. The entries in the table will suffice for estimating the population parameters $\mu_1, \mu_2 \in \mathbb{R}$, $\sigma_1, \sigma_2 > 0$, and $\pi \in (0, 1)$.

With this setup, we then want to test $H_0 : \mu_1 = \mu_2 + \Delta$, where $\Delta = 0$ is common. We will assume that the standard deviation of each group is different for each other as well. We then consider the random vector $Y = \sum_{i=1}^n Y_i$ which will consist of the entries from the contingency table above. Note that we do not require entries for both $\sum_{i=1}^n W_i$ as well as

$\sum_{i=1}^n (1 - W_i)$ as one can be written in terms of the other. Hence, we have

$$Y = \begin{pmatrix} Y[1] = & \sum_{i=1}^n W_i \\ Y[2] = & \sum_{i=1}^n W_i X_{i,0} \\ Y[3] = & \sum_{i=1}^n (1 - W_i) X_{i,1} \\ Y[4] = & \sum_{i=1}^n W_i X_{i,0}^2 \\ Y[5] = & \sum_{i=1}^n (1 - W_i) X_{i,1}^2 \end{pmatrix} \tag{11}$$

We then consider the individual i.i,d. samples Y_i where $Y = \sum_{i=1}^n Y_i$ so that we can compute the expectation of Y_i under the null hypothesis $\mu_1 = \mu_2 + \Delta$ for some $\Delta \in \mathbb{R}$

$$\boldsymbol{\theta}(\pi, \mu_1, \mu_2, \sigma_1, \sigma_2) = \mathbb{E}[Y_i] = \left(\pi, \pi\mu_1, (1-\pi)\mu_2, \pi\left(\mu_1^2 + \sigma_1^2\right), (1-\pi)\left(\mu_2^2 + \sigma_2^2\right)\right)^\intercal$$

Note that we can compute the covariance matrix and estimates for $\pi, \mu_1, \mu_2, \sigma_1, \sigma_2$, but it will help to simplify the χ^2 test first. In particular, when we write out the χ^2 statistic, the minimization will lead to the second moment terms (the last two entries in Y) to contribute nothing to the χ^2 value, since we can zero out those coordinates by setting $\sigma_1^2 = Y[4]/n - \mu_1^2$ as long as $\mu_1^2 \leq Y[4]/n$, and similarly for σ_2. Note that if it does turn out that for a particular μ_1 we have $Y[4]/n - \mu_1^2$, we should reject, i.e. return a large statistic of say $10g$ or so, since that would mean that zero variance would be our best estimate for that group. Hence, we will only consider $Y = (Y[1], Y[2], Y[3])^\intercal$ in our test. If it can be assumed that the variances are equal across groups, then we can keep the coordinates $Y[3], Y[4]$ in our test statistic. Hence, we will continue with Y_i denoting the first three coordinates of the random vector in (11).

We next calculate the covariance matrix $C(\pi, \mu_1, \mu_2, \sigma_1, \sigma_2)$ for the first 3 coordinates in Y_i

$$C(\pi, \mu_1, \mu_2, \sigma_1, \sigma_2) = \begin{bmatrix} \pi(1-\pi) & \pi\mu_1 - \pi^2\mu_1 & -\pi(1-\pi)\mu_2 \\ \pi\mu_1(1-\pi\mu_1) & \pi(\mu_1^2 + \sigma_1^2) - \pi\mu_1 & -\pi\mu_1(1-\pi)\mu_2 \\ \pi(1-\pi)\mu_2 & -\pi\mu_1(1-\pi)\mu_2 & (1-\pi)(\mu_2^2+\sigma_2^2)(1-(1-\pi)\mu_2) \end{bmatrix}$$

Under the null hypothesis $\mu_1 = \mu_2 + \Delta$, we will use the following estimates

$$\begin{aligned} &\hat{\pi} = Y[1]/n, \quad \hat{\mu} = \frac{Y[2] + Y[3]}{n}, \quad \hat{\mu}_1 = \hat{\mu} + (1 - \hat{\pi})\Delta, \quad \hat{\mu}_2 = \hat{\mu} - \hat{\pi}\Delta. \\ &\sigma_1^2 = \frac{Y[4]/n}{\hat{\pi}} - \hat{\mu_1}^2, \quad \sigma_2^2 = \frac{Y[5]/n}{1-\hat{\pi}} - \hat{\mu_2}^2 \end{aligned}$$

We are now ready to calculate the χ^2 statistic

$$D = \min_{\substack{\pi \in (0,1), \\ \mu_1, \mu_2 : \mu_1 = \mu_2 + \Delta}} \left\{ \begin{pmatrix} Y[1]/n - \pi \\ Y[2]/n - \pi\mu_1 \\ Y[3]/n - (1-\pi)\mu_2 \end{pmatrix}^\intercal C(\hat{\pi}, \hat{\mu_1}, \hat{\mu_2}, \hat{\sigma_1}, \hat{\sigma_2})^{-1} \begin{pmatrix} Y[1]/n - \pi \\ Y[2]/n - \pi\mu_1 \\ Y[3]/n - (1-\pi)\mu_2 \end{pmatrix} \right\} \tag{12}$$

We will evaluate the test statistic D against a χ^2 with $(3-2) = 1$ degree of freedom. We now compare this χ^2 based hypothesis test with the traditional t-test described in the previous subsection, at a significance level of 95%. We consider the null hypothesis $H_0 : \mu_1 = \mu_2 = \mu$ and generate data in two groups with equal variances. First, we modify the shared mean μ

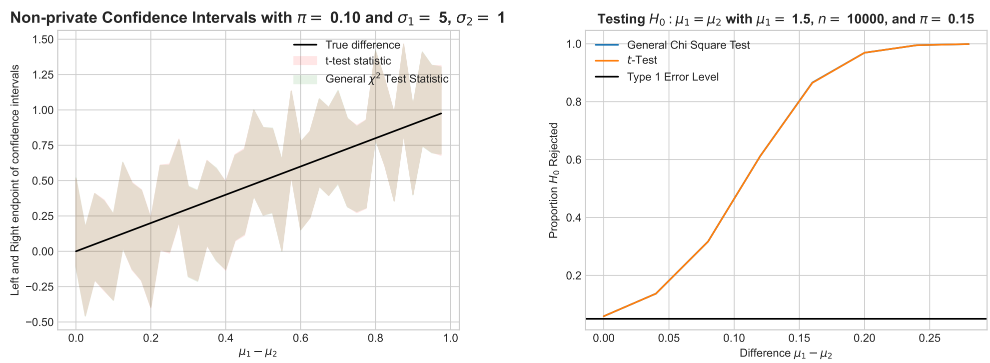

Figure 11: (Left) Comparing confidence intervals using the standard t-test statistic and the general χ^2 statistic in (12) with $\mu_2 = 0$. The bounds are completely overlapping, resulting in identical tests. (Right) Comparing power curves of the t-test and the general χ^2 test for $H_0 : \mu_1 = \mu_2$, which are on top of each other. In this case we use $\sigma_1 = 2, \sigma_2 = 1$ when we generate data.

as well as the common variance and the probability π of being in group 1. The left plot of Figure 11 shows confidence intervals using the standard t-test statistic (assuming unequal variances) and the general χ^2 statistic using an approach similar to when we had binary outcomes in Section A.6. When computing the left and right end points of the confidence interval for proportions, we could simply consider $\Delta \in [-1, 1]$. However, for differences in means, we will use the data to determine lower and upper bounds on the candidate confidence interval region. We can simply take the sample mean in both groups and add or subtract say 10 standard deviations within each group. The results show overlapping confidence intervals. The right plot of Figure 11 shows the average number of times over 1000 trials that each test rejected the null hypothesis $H_0 : \mu_1 = \mu_2$ as we change the difference $\mu_1 - \mu_2 > 0$, so that we should reject more frequently. The proportion of null hypotheses rejected is indistinguishable between the two tests.

C.2. Private Tests for Difference in Means

Recall that we use the traditional randomized response when privatizing group $j \in \{1, 2\}$. We will use the general χ^2 approach to design a private test for differences between two means. This test can then be used to calculate confidence intervals, in the same way that confidence intervals were derived from the χ^2 statistic for binary outcomes in Section A.6. We will then consider the privatized version of the random vector Y from (11), but we will only use the first 3 coordinates, as the last two coordinates will be eliminated in the minimization of the statistic, as we mentioned earlier.

For randomized response, we will write $Z_i^\varepsilon[j] \sim \mathrm{Bern}(\frac{e^\varepsilon}{e^\varepsilon+1})$ for $i \in [n]$ and $j \in \{1, 2\}$, and then write the privatized vector Y_i^ε in terms of $Z_i^\varepsilon[j]$,

$$Y_i^\varepsilon = \begin{pmatrix} Z_i^\varepsilon[1] \cdot W_i + (1 - Z_i^\varepsilon[2]) \cdot (1 - W_i) \\ Z_i^\varepsilon[1] \cdot W_i \cdot X_i[1] + (1 - Z_i^\varepsilon[2]) \cdot (1 - W_i) \cdot X_i[2] \\ (1 - Z_i^\varepsilon[1]) \cdot W_i \cdot X_i[1] + Z_i^\varepsilon[2] \cdot (1 - W_i) \cdot X_i[2] \end{pmatrix}$$

Next, we compute its expectation $\mathbb{E}[Y_i^\varepsilon] = \boldsymbol{\theta}^\varepsilon(\pi, \mu_1, \mu_2)$ where

$$\boldsymbol{\theta}^\varepsilon(\pi, \mu_1, \mu_2) = \begin{pmatrix} \frac{e^\varepsilon}{e^\varepsilon+1}\pi + \frac{1}{e^\varepsilon+1}(1-\pi) \\ \frac{e^\varepsilon}{e^\varepsilon+1}\pi\mu_1 + \frac{1}{e^\varepsilon+1}(1-\pi)\mu_2 \\ \frac{1}{e^\varepsilon+1}\pi\mu_1 + \frac{e^\varepsilon}{e^\varepsilon+1}(1-\pi)\mu_2 \end{pmatrix}$$

We can then compute Y_i^ε covariance matrix $C(\pi, \mu_1, \mu_2, \sigma_1, \sigma_2; \varepsilon)$.

$$\begin{aligned} C(\pi, \mu_1, \mu_2, \sigma_1, \sigma_2; \varepsilon) &= \mathbb{E}\left(Y_i^\varepsilon (Y_i^\varepsilon)^\intercal\right) - \mathbb{E}\left[Y_i^\varepsilon\right]\mathbb{E}\left[Y_i^\varepsilon\right]^\intercal \\ &= \frac{1}{e^\varepsilon+1}\begin{bmatrix} e^\varepsilon\pi + (1-\pi) & \pi e^\varepsilon\mu_1 + (1-\pi)\mu_2 & 0 \\ \pi e^\varepsilon\mu_1 + (1-\pi)\mu_2 & \pi e^\varepsilon(\mu_1^2+\sigma_1^2) + (1-\pi)(\mu_2^2+\sigma_2^2) & 0 \\ 0 & 0 & \pi(\mu_1^2+\sigma_1^2) + (1-\pi)e^\varepsilon(\mu_2^2+\sigma_2^2) \end{bmatrix} \\ &\quad - \boldsymbol{\theta}^\varepsilon(\pi, \mu_1, \mu_2)\boldsymbol{\theta}^\varepsilon(\pi, \mu_1, \mu_2)^\intercal \end{aligned}$$

Now, we need to assign estimates for the parameters $\pi, \mu_1, \mu_2, \sigma_1, \sigma_2$. Here, we have the null hypothesis $H_0 : \mu_1 = \mu_2 + \Delta$ and sample data Y^ε. Beginning with $\hat{\pi}, \hat{\mu}_1$, and $\hat{\mu}_2$, we use

$$\hat{\pi} = (e^\varepsilon + 1)\left(\frac{Y^\varepsilon[1]/n - \frac{1}{e^\varepsilon+1}}{e^\varepsilon - 1}\right)$$

$$\begin{bmatrix} \hat{\pi}\frac{e^\varepsilon}{e^\varepsilon+1} & (1-\hat{\pi})\frac{1}{e^\varepsilon+1} \\ \hat{\pi}\frac{1}{e^\varepsilon+1} & (1-\hat{\pi})\frac{e^\varepsilon}{e^\varepsilon+1} \end{bmatrix}\begin{pmatrix} \hat{\mu}_1 \\ \hat{\mu}_2 \end{pmatrix} = \begin{pmatrix} Y^\varepsilon[2]/n \\ Y^\varepsilon[3]/n \end{pmatrix}$$

We then use the null hypothesis to replace $\hat{\mu}_1 = \hat{\mu}_2 + \Delta$ and then solve the over constrained system of equations via least squares in which case we have

$$\hat{\mu}_2 = \frac{\hat{\pi}\frac{e^\varepsilon}{e^\varepsilon+1}\left(Y^\varepsilon[2]/n - \hat{\pi}\frac{e^\varepsilon}{e^\varepsilon+1}\Delta\right) + (1-\hat{\pi})\frac{1}{e^\varepsilon+1}\left(Y^\varepsilon[3]/n - \hat{\pi}\frac{1}{e^\varepsilon+1}\Delta\right)}{\left(\hat{\pi}\frac{e^\varepsilon}{e^\varepsilon+1}\right)^2 + \left((1-\hat{\pi})\frac{1}{e^\varepsilon+1}\right)^2}$$

For σ_1, σ_2, we will use different estimates in the two diagonal entries in which they appear, corresponding to the estimation sample. We replace the σ_1 in the $C[2,2]$ diagonal entry with the sample variance s_1^2 of $\{Y_i^\varepsilon[2]\}_{i=1}^n$. Analogously, for the $C[3,3]$ diagonal entry, we replace σ_2 with the sample variance s_2^2 of $\{Y_i^\varepsilon[3]\}_{i=1}^n$. We do point out that under the estimates $\hat{\mu}_1 = \hat{\mu}_2 + \Delta$, there might not be a possible $\sigma_1, \sigma_2 \geq 0$ that can achieve the sample variances s_1^2, s_2^2. We then compute the theoretical variance of our observations $Y^\varepsilon[2], Y^\varepsilon[3]$

$$\begin{aligned} \mathrm{Var}(Y^\varepsilon[2]) &= \mu_1^2\left(1 - \tfrac{\pi e^\varepsilon}{e^\varepsilon+1}\right)\tfrac{\pi e^\varepsilon}{e^\varepsilon+1} + \mu_2^2\left(1 - \tfrac{(1-\pi)}{e^\varepsilon+1}\right)\tfrac{(1-\pi)}{e^\varepsilon+1} - 2\tfrac{\pi(1-\pi)e^\varepsilon}{(e^\varepsilon+1)^2}\mu_1\mu_2 + \tfrac{\pi e^\varepsilon}{e^\varepsilon+1}\sigma_1^2 + \tfrac{(1-\pi)}{e^\varepsilon+1}\sigma_2^2 \\ \mathrm{Var}(Y^\varepsilon[3]) &= \mu_1^2\left(1 - \tfrac{\pi}{e^\varepsilon+1}\right)\tfrac{\pi}{e^\varepsilon+1} + \mu_2^2\left(1 - \tfrac{(1-\pi)e^\varepsilon}{e^\varepsilon+1}\right)\tfrac{(1-\pi)e^\varepsilon}{e^\varepsilon+1} - 2\tfrac{\pi(1-\pi)e^\varepsilon}{(e^\varepsilon+1)^2}\mu_1\mu_2 + \tfrac{\pi}{e^\varepsilon+1}\sigma_1^2 + \tfrac{(1-\pi)e^\varepsilon}{e^\varepsilon+1}\sigma_2^2 \end{aligned}$$

Because $\sigma_j^2 \geq 0$, we need to ensure our estimate s_j for $\mathrm{Var}(Y^\varepsilon[j])$ for $j \in \{1, 1\}$. Thus we check if the following inequalities are satisfied.

$$s_1^2 \geq \hat{\mu}_1^2 \left(1 - \frac{\hat{\pi}e^\varepsilon}{e^\varepsilon+1}\right) \frac{\hat{\pi}e^\varepsilon}{e^\varepsilon+1} + \hat{\mu}_2^2 \left(1 - \frac{(1-\hat{\pi})}{e^\varepsilon+1}\right) \frac{(1-\hat{\pi})}{e^\varepsilon+1} - 2\frac{\hat{\pi}(1-\hat{\pi})e^\varepsilon}{(e^\varepsilon+1)^2} \hat{\mu}_1 \hat{\mu}_2$$
$$s_2^2 \geq \hat{\mu}_1^2 \left(1 - \frac{\hat{\pi}}{e^\varepsilon+1}\right) \frac{\hat{\pi}}{e^\varepsilon+1} + \hat{\mu}_2^2 \left(1 - \frac{(1-\hat{\pi})e^\varepsilon}{e^\varepsilon+1}\right) \frac{(1-\hat{\pi})e^\varepsilon}{e^\varepsilon+1} - 2\frac{\hat{\pi}(1-\hat{\pi})e^\varepsilon}{(e^\varepsilon+1)^2} \hat{\mu}_1 \hat{\mu}_2$$

If they are not satisfied, we replace s_j^2 with the corresponding right hand side, essentially using $\sigma_1 = \sigma_2 = 0$ in our estimate.

Putting this all together, we have the following χ^2 statistic, D^ε, which will use Y^ε instead of Y in (11) which we compare to a χ^2 with 1 degree of freedom to base our hypothesis test.

$$D^\varepsilon = \min_{\substack{\pi \in (0,1), \\ \mu_1, \mu_2 : \mu_1 = \mu_2 s + \Delta}} \left\{ \left(Y^\varepsilon - \boldsymbol{\theta}^\varepsilon (\pi, \mu_1, \mu_2)\right)^\intercal C(\hat{\pi}, \hat{\mu_1}, \hat{\mu_2}, \hat{\sigma_1}, \hat{\sigma_2})^{-1} \left(Y^\varepsilon - \boldsymbol{\theta}^\varepsilon (\pi, \mu_1, \mu_2)\right) \right\} \tag{13}$$

where

$$Y^\varepsilon - \boldsymbol{\theta}^\varepsilon (\pi, \mu_1, \mu_2) = \begin{pmatrix} Y^\varepsilon[1]/n - \frac{e^\varepsilon}{e^\varepsilon+1}\pi + \frac{1}{e^\varepsilon+1}(1-\pi) \\ Y^\varepsilon[2]/n - \frac{e^\varepsilon}{e^\varepsilon+1}\pi\mu_1 + \frac{1}{e^\varepsilon+1}(1-\pi)\mu_2 \\ Y^\varepsilon[3]/n - \frac{1}{e^\varepsilon+1}\pi\mu_1 + \frac{e^\varepsilon}{e^\varepsilon+1}(1-\pi)\mu_2 \end{pmatrix}.$$

C.3. Results

We present results on the confidence intervals of the difference in means, as we did in Section A.6 for binary outcomes. As previously stated, we will use a similar approach to the binary outcome case and use the χ^2 statistic for various $\Delta = \mu_1 - \mu_2$.

For the t-test based confidence intervals, we correct the difference in means in the same way as in (6), and we compare the confidence intervals we get with the general χ^2 approach. We present our results in Figure 12, where we fix $\mu_1 = 0$, $\varepsilon = 1$, and $n = 1000$, while we vary $\mu_1 - \mu_2$ in each plot and change π, σ_1, σ_2 in the different plots. Note that the confidence intervals from the general χ^2 statistics sometimes produces wider confidence intervals, although they are very close to the confidence intervals from the t-test statistic with a correction.

We also present results in how often the various approaches provide confidence intervals that overlap the true difference in means in Figure 13. In each case we privatize the data using randomized response on the group and compute confidence intervals using the classical t-test with a correction from (6) and confidence intervals produced with the general χ^2 approach. We see that the general χ^2 approach produces slightly more conservative confidence intervals as the privacy parameter increases and we consistently achieve close to the target 5% proportion of missing the true difference in means.

Appendix D. ANOVA

We move to giving full derivations for testing whether there is a difference in means across $g > 2$ groups. It is straightforward to fit this hypothesis test to our generalized χ^2 test framework by considering the variable W_i, which will determine the group that sample i is in. That is, $W_i \sim \mathrm{Multinomial}(n, \pi)$ for $i \in [n]$ and $\pi \in [0,1]^g$ is a probability vector. We

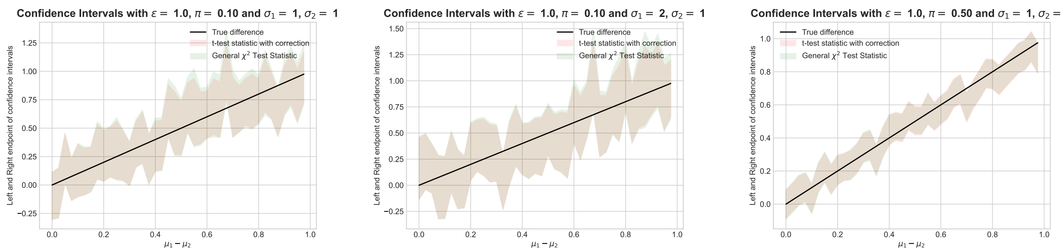

Figure 12: Comparing confidence intervals using either the t-test or with the approach we outline for the χ^2 statistic. We use $\mu_1 = 0$, $\epsilon = 1.0$, and $n = 10000$ for the plots and change σ_1, σ_2 as well as π.

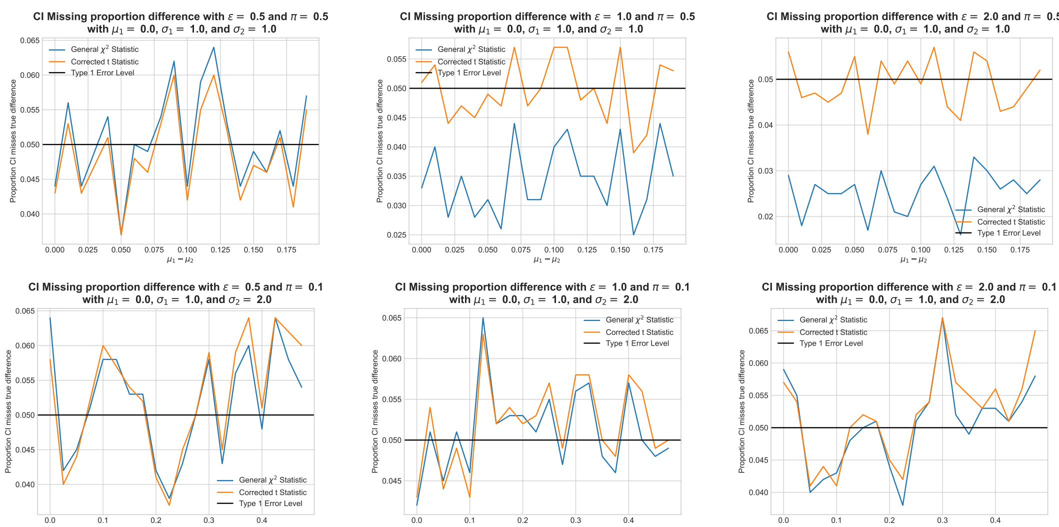

Figure 13: We plot the proportion of times different tests will produce confidence intervals that miss the true difference $\mu_1 - \mu_2$. The general χ^2 test that accounts for the privacy mechanism consistently achieves the target type I error of 5%. We compare our approach with using the classical t-test with correction in (6). All plots use data size $n = 10000$ over 1000 trials.

can then generalize Table 2 for multiple groups in Table 3, which we will use in our privacy model.

Sample Moments	Group 1 w.p. π_1	Group 2 w.p. π_2	$\cdots$	Group g w.p. π_g
0-th	$\sum_{i=1}^n W_i[1]$	$\sum_{i=1}^n W_i[2]$	$\cdots$	$\sum_{i=1}^n W_i[g]$
1-st	$\sum_{i=1}^n W_i[1] \cdot X_i[1]$	$\sum_{i=1}^n W_i[2] \cdot X_i[2]$	$\cdots$	$\sum_{i=1}^n W_i[g] \cdot X_i[g]$
2-nd	$\sum_{i=1}^n W_i[1] \cdot X_i^2[1]$	$\sum_{i=1}^n W_i[2] \cdot X_i^2[2]$	$\cdots$	$\sum_{i=1}^n W_i[g] \cdot X_i^2[g]$

Table 3: Contingency Table for continuous outcomes $\{X_i[j]\}_{i=1}^n \overset{i.i.d.}{\sim} \mathrm{N}(\mu_j, \sigma_j^2)$ with $j \in [g]$ and group variable $\{W_i\}_{i=1}^n \overset{i.i.d.}{\sim}$ Multinomial$(1, \pi)$.

We can then use our general χ^2 framework to design a χ^2 statistic based on the data in each cell of Table 3 in a similar way to testing the difference in means. Recall that for the difference in two means, we only considered the 0th and 1st sample orders in Y, which we will also consider. We use the random vector $Y = \sum_i^n Y_i$ in our test and then compute the expectation of $Y_i = (Y_i^\varepsilon[1], Y_i^\varepsilon[2], \cdots, Y_i^\varepsilon[2g])^\intercal$ and its covariance matrix. We then find estimates for the distribution parameters with common mean $\mu_j = \mu$ for all $j \in [g]$, group probability π, and standard deviations σ_j for $j \in [g]$, which will be used in the covariance matrix in place of the population parameters.

$$Y = \sum_{i=1}^n Y_i = \sum_{i=1}^n \begin{pmatrix} W_i[1] & W_i[2] & \ldots & W_i[g-1] & W_i[1] \cdot X_i[2] & W_i[2] \cdot X_i[2] & \ldots & W_i[g] \cdot X_i[g] \end{pmatrix}^T.$$

After minimizing the resulting χ^2 statistic over $\mu \in \mathbb{R}$ and $\pi \in [0,1]^g$ such that $\sum_{i=1}^g \pi_j = 1$, we then compare the statistic to a χ^2 distribution with $g-1$ degrees of freedom. Figure 14 shows that the χ^2 approach performs equivalently to the traditional one-way ANOVA test with equal variance $\sigma^2 = \sigma_j^2$ for $j \in [g]$.

To then introduce privacy of the group of each sample, but not the outcome, we will include the random matrix $Z_i^\varepsilon \in \{0,1\}^{g \times g}$, where column j corresponds to the outcome for $M(j)$ for various local DP mechanisms M. Hence, we can write out the privatized vector $Y^\varepsilon = \sum_{i=1}^n Y_i^\varepsilon$ in terms of Z_i^ε as we did for binary outcomes in (10), where $\mathbf{0}$ is the $g \times g$ zero matrix.

$$Y_i^\varepsilon = \begin{bmatrix} Z_i^\varepsilon & \mathbf{0} \\ \mathbf{0} & Z_i^\varepsilon \end{bmatrix} \begin{pmatrix} W_i[1] & W_i[2] & \ldots & W_i[g] & W_i[1] \cdot X_i[1] & W_i[2] \cdot X_i[2] & \ldots & W_i[g] \cdot X_i[g] \end{pmatrix}^T \tag{14}$$

Note that we include the coordinate for $W_i[g]$ because the privacy mechanism will modify the probability of being in the last group. For the g-randomized response mechanism, we can eliminate this entry as the first g coordinates of Y_i^ε will still form a multinomial distribution. However, for the bit flip mechanism, we will need to keep the full $2g$ dimensional vector, in which case we compare the resulting statistic to a χ^2 distribution with g degrees of freedom, one more than the non-private test. Lastly, for the subset mechanism, the resulting covariance

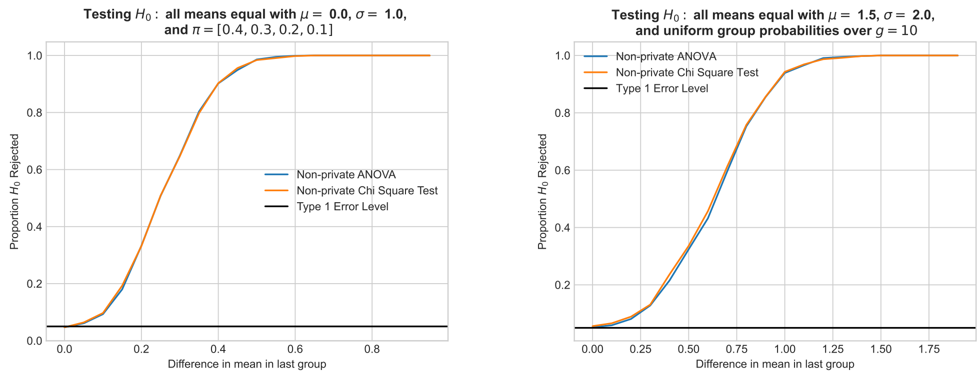

Figure 14: Comparing power curves for one-way ANOVA and the general χ^2 approach with $n = 1000$ samples.

matrix of Y_i^ε will have the vector comprising of ones in the first g coordinates and zeros in the second g coordinates be in its null space. We will cover each in turn in this section.

For each mechanism, we need to compute the expected vector $\boldsymbol{\theta}^\varepsilon(\pi, \mu)$ under the null hypothesis and the covariance matrix $C(\pi, \mu, \sigma_1^2, \cdots, \sigma_g^2; \varepsilon)$. Note that the mean vector $\boldsymbol{\theta}^\varepsilon(\pi, \mu) \in \mathbb{R}^{g \times g}$ will have the same form as the vectors in Section B, except the first g entries will have success probability $p = 1$ and the second g entries will have μ instead of $1 - p$ multiplied. We will be able to write the covariance matrix in the following block form with matrices $\Sigma, \Sigma' \in \mathbb{R}^{g \times g}$,

$$C(\pi, \mu, \sigma_1^2, \cdots, \sigma_g^2; \varepsilon) = \begin{bmatrix} \Sigma & \mu\Sigma \\ \mu\Sigma & \Sigma' \end{bmatrix} \tag{15}$$

Note that Σ will simply be the top left $g \times g$ submatrix of the covariance matrix for each mechanism in Section B, with probability of success $p = 1$. That is for $j, \ell \in [g]$,

$$\Sigma[j, \ell] = \sum_{m=1}^{g} \mathbb{E}[W_i[m] \cdot Z_i^\varepsilon[j, m] \cdot Z_i^\varepsilon[\ell, m]] - \mathbb{E}[Y_i^\varepsilon][j] \cdot \mathbb{E}[Y_i^\varepsilon][\ell]$$

Further, we can compute Σ' in terms of the expectation of Y_i^ε and the specific mechanism we use for $j, \ell \in [g]$

$$\Sigma'[j, \ell] = \sum_{m=1}^{g} (\mu^2 + \sigma_m^2)\mathbb{E}[W_i[m] \cdot Z_i^\varepsilon[j, m] \cdot Z_i^\varepsilon[\ell, m]] - \mu^2\mathbb{E}[Y_i^\varepsilon][j] \cdot \mathbb{E}[Y_i^\varepsilon][\ell]$$

We also need to form estimates for the population parameters in the covariance matrix, including $\mu, \pi, \sigma_1, \cdots, \sigma_g$. We will use our sample and the particular privacy mechanism to form these estimates. Note that if our estimates result in an expected group size to be less than 5, then we will simply fail to reject the null hypothesis, as we have stated earlier. We now cover each mechanism below.

D.1. Randomized Response

We first cover the g-randomized response mechanism we have $Y_i^\varepsilon = (Y_i^\varepsilon[1], \cdots, Y_i^\varepsilon[2g])^\intercal$, which will have the following mean vector

$$\boldsymbol{\theta}^\varepsilon(\pi,\mu) = \mathbb{E}[Y_i^\varepsilon] = \begin{pmatrix} \frac{1}{e^\varepsilon+g-1} \cdot \begin{bmatrix} e^\varepsilon & 1 & \cdots & 1 \\ 1 & e^\varepsilon & \cdots & 1 \\ & & \ddots & \\ 1 & 1 & \cdots & e^\varepsilon \end{bmatrix} \pi \\ \\ \frac{\mu}{e^\varepsilon+g-1} \cdot \begin{bmatrix} e^\varepsilon & 1 & \cdots & 1 \\ 1 & e^\varepsilon & \cdots & 1 \\ & & \ddots & \\ 1 & 1 & \cdots & e^\varepsilon \end{bmatrix} \pi \end{pmatrix}$$

Using our general form of the covariance matrix in (15) we can simplify the terms for $\Sigma[j,\ell]$ due to $Z_i^\varepsilon[j,m] \cdot Z_i[\ell,m] = 0$ for any $\ell \neq j$ and is 1 otherwise. We also need to form estimates to use in the χ^2 statistic, which gives us

$$\hat{\mu} = \frac{\sum_{j\in[g]} \sum_{i=}^n Y_i^\varepsilon[g+j]}{n}, \qquad \hat{\pi} = \left((e^\varepsilon+g-1)\left(\frac{\frac{\sum_{i=1}^n Y_i^\varepsilon[j]}{n} - \frac{1}{e^\varepsilon+g-1}}{e^\varepsilon - 1} \right) : j \in [g] \right).$$

We also need to estimate the variance for each group, so we will use the sample standard deviation for each group, i.e. $\{Y_i^\varepsilon[g+1:2g] : i \in [n]\}$ and use this for the main diagonal of Σ'. The terms of Σ' on the off diagonal will only consist of terms $-\mathbb{E}[Y_i^\varepsilon][j] \cdot \mathbb{E}[Y_i^\varepsilon][\ell]$ for $j \neq \ell$.

We also point out that for randomized response, the first g entries will still follow a multinomial distribution, where only the first $g-1$ entries are needed. Hence, we remove the g-th entry in Y_i^ε while also removing the g-th row and column of the covariance matrix. The result is then a covariance matrix of rank at most $2g-1$ and we optimize over g variables, with $\pi_1, \cdots, \pi_{g-1}$ and μ. We then form the χ^2-statistic and compare it to a χ^2 distribution with $g-1$ degrees of freedom.

D.2. Bit Flipping

We next turn to the bit flipping mechanism, where it is possible for a sample to be in multiple groups simultaneously. We first compute the expected vector of Y_i^ε, as we did for the binary outcome case in Section B.2

$$\boldsymbol{\theta}^{\varepsilon}(\pi,\mu) = \begin{pmatrix} \frac{1}{e^{\varepsilon/2}+1} \cdot \begin{bmatrix} e^{\varepsilon/2} & 1 & \cdots & 1 \\ 1 & e^{\varepsilon/2} & \cdots & 1 \\ & & \ddots & \\ 1 & 1 & \cdots & e^{\varepsilon/2} \end{bmatrix} \pi \\ \\ \frac{\mu}{e^{\varepsilon/2}+1} \cdot \begin{bmatrix} e^{\varepsilon/2} & 1 & \cdots & 1 \\ 1 & e^{\varepsilon/2} & \cdots & 1 \\ & & \ddots & \\ 1 & 1 & \cdots & e^{\varepsilon/2} \end{bmatrix} \pi \end{pmatrix}$$

We next compute the covariance matrix Σ', which consists of terms with $\mathbb{E}[Z_i^{\varepsilon}[j,m]Z_i^{\varepsilon}[\ell,m]]$ for $j,\ell,m \in [g]$. From the bit flip mechanism, we know each coordinate of Z_i^{ε} is independent of each other, so we have for $j \neq \ell$

$$\mathbb{E}[Z_i^{\varepsilon}[j,m]Z_i^{\varepsilon}[\ell,m]] = \mathbb{E}[Z_i^{\varepsilon}[j,m]] \cdot \mathbb{E}[Z_i^{\varepsilon}[\ell,m]] = \begin{cases} \frac{1}{(e^{\varepsilon/2}+1)^2} & j,\ell \neq m \\ \frac{e^{\varepsilon/2}}{(e^{\varepsilon/2}+1)^2} & j = m, \text{ or } \ell = m \end{cases}$$

and if $j = \ell$

$$\mathbb{E}[Z_i^{\varepsilon}[j,m]Z_i^{\varepsilon}[\ell,m]] = \mathbb{E}[Z_i^{\varepsilon}[j,m]] = \begin{cases} \frac{1}{e^{\varepsilon/2}+1} & j \neq m \\ \frac{e^{\varepsilon/2}}{e^{\varepsilon/2}+1} & j = m \end{cases}$$

We now form estimates for the population parameters in the covariance matrix.

$$\hat{\mu} = (e^{\varepsilon/2}+1) \cdot \frac{\sum_{i=1}^n \sum_{j=1}^g Y_i^{\varepsilon}[g+j]}{n(e^{\varepsilon/2}+(g-1))}, \quad \hat{\pi} = \left((e^{\varepsilon/2}+1) \cdot \left(\frac{\frac{\sum_{i=1}^n Y_i^{\varepsilon}[j]}{n} - \frac{1}{e^{\varepsilon/2}+1}}{e^{\varepsilon/2}-1} \right) : j \in [g] \right)$$

Next, we need to form an estimate for the variance. To help simplify things, we will assume equal variance across groups in our estimate, i.e $\sigma_j = \sigma$ for all $j \in [g]$. We point out that unequal variances can be used, but it just complicates the estimate we use. Further, we will write s_j^2 to denote the sample variance computed within each group, so that s_j^2 is an estimate for the variance of $Y_i^{\varepsilon}[g+j]$. This gives us the following estimate for σ,

$$\hat{\sigma}^2 = (e^{\varepsilon/2}+1) \cdot \frac{\sum_{j=1}^g s_j^2 - \hat{\mu}^2 \left(\frac{e^{\varepsilon/2}}{e^{\varepsilon/2}+1} + (g-1)\frac{1}{e^{\varepsilon/2}+1} - \sum_{j=1}^g (\pi_j \frac{e^{\varepsilon/2}}{e^{\varepsilon/2}+1} + (1-\pi_j)\frac{1}{e^{\varepsilon/2}+1})^2 \right)}{e^{\varepsilon/2}+(g-1)}.$$

We now use the covariance matrix with these estimates in the general χ^2 statistic and compare it with a χ^2 distribution with g degrees of freedom. Note that there is an extra degree of freedom compared to the previous test and the non-private version. This is because the covariance matrix may be full rank and some combination of elements in Y_i^{ε} cannot be used to determine other elements, as was possible when we had the g-randomized response mechanism. Hence, the covariance matrix is of rank at most $2g$ and we are minimizing over g variables $\pi_1, \cdots, \pi_g, \mu$.

Remark 7 *Recall that in the binary outcome case, we were using a technique from Kifer and Rogers (2017); Gaboardi and Rogers (2018) to* project out *the eigenvector associated with noise, which reduced the degrees of freedom by 1. A similar approach cannot be adopted here, as in the non-private case, the vector whose first g coordinates are 1 and the latter g coordinates are zero is in the null space of the covariance matrix, but the same vector is no longer an eigenvector of the covariance matrix after applying the bit flipping mechanism for general μ.*

D.3. Subset Mechanism

Lastly we cover the subset mechanism, which was shown to be the most powerful of the other privacy mechanisms for various privacy levels in the binary outcomes case. To ease notation, we will assume the variance is equal across all groups, i.e. $\sigma_j = \sigma$ for all $j \in [g]$. We follow the same procedure as for the other mechanisms, where we first give the expected value of Y_i^ε, denoted as $\boldsymbol{\theta}^\varepsilon(\pi, p; k)$ where the random entries in Z_i^ε from (14) come from the subset mechanism with parameter k.

$$\boldsymbol{\theta}^\varepsilon(\pi, p; k) = \begin{pmatrix} \frac{1}{\binom{g-1}{k-1}e^\varepsilon + \binom{g-1}{k}} \begin{bmatrix} \binom{g-1}{k-1}e^\varepsilon & \left(\binom{g-2}{k-2}e^\varepsilon + \binom{g-2}{k-1}\right) & \cdots & \left(\binom{g-2}{k-2}e^\varepsilon + \binom{g-2}{k-1}\right) \\ \left(\binom{g-2}{k-2}e^\varepsilon + \binom{g-2}{k-1}\right) & \binom{g-1}{k-1}e^\varepsilon & \cdots & \left(\binom{g-2}{k-2}e^\varepsilon + \binom{g-2}{k-1}\right) \\ & & \ddots & \\ \left(\binom{g-2}{k-2}e^\varepsilon + \binom{g-2}{k-1}\right) & \left(\binom{g-2}{k-2}e^\varepsilon + \binom{g-2}{k-1}\right) & \cdots & \binom{g-1}{k-1}e^\varepsilon \end{bmatrix} \pi \\ \\ \frac{\mu}{\binom{g-1}{k-1}e^\varepsilon + \binom{g-1}{k}} \begin{bmatrix} \binom{g-1}{k-1}e^\varepsilon & \left(\binom{g-2}{k-2}e^\varepsilon + \binom{g-2}{k-1}\right) & \cdots & \left(\binom{g-2}{k-2}e^\varepsilon + \binom{g-2}{k-1}\right) \\ \left(\binom{g-2}{k-2}e^\varepsilon + \binom{g-2}{k-1}\right) & \binom{g-1}{k-1}e^\varepsilon & \cdots & \left(\binom{g-2}{k-2}e^\varepsilon + \binom{g-2}{k-1}\right) \\ & & \ddots & \\ \left(\binom{g-2}{k-2}e^\varepsilon + \binom{g-2}{k-1}\right) & \left(\binom{g-2}{k-2}e^\varepsilon + \binom{g-2}{k-1}\right) & \cdots & \binom{g-1}{k-1}e^\varepsilon \end{bmatrix} \pi \end{pmatrix}$$

We write the j, ℓ entry of the submatrix $\Sigma' = \mathbb{E}[Y_i^\varepsilon[g+j] \cdot Y_i^\varepsilon[g+\ell]] - \mathbb{E}[Y_i^\varepsilon][g+j] \cdot \mathbb{E}[Y_i^\varepsilon][\ell]$, where the first term can be computed with $j = \ell$

$$\mathbb{E}[Y_i^\varepsilon[g+j]^2] = (\mu^2 + \sigma^2) \frac{\binom{g-1}{k-1}e^\varepsilon \pi_j + \left(\binom{g-2}{k-2}e^\varepsilon + \binom{g-2}{k-1}\right)(1 - \pi_j)}{\binom{g-1}{k-1}e^\varepsilon + \binom{g-1}{k}}, \qquad \forall j \in [g].$$

and for $\ell \neq j$,

$$\begin{aligned} &\mathbb{E}[Y_i^\varepsilon[g+j] Y_i^\varepsilon[g+\ell]] \\ &\quad = \frac{(\mu^2 + \sigma^2)}{\binom{g-1}{k-1}e^\varepsilon + \binom{g-1}{k}} \left(e^\varepsilon \binom{g-2}{k-2} (\pi_j + \pi_\ell) + \left(e^\varepsilon \binom{g-3}{k-3} + \binom{g-3}{k-2} \right) (1 - \pi_j - \pi_\ell) \right) \end{aligned}$$

We make the following observation about the covariance matrix of Y_i^ε that will show that it is not full rank, and hence we will lose a degree of freedom in the asymptotic χ^2 distribution of the test statistic.

Lemma 8 *The covariance matrix $C(\pi, \mu, \sigma_1^2, \cdots, \sigma_g^2; \varepsilon, k)$ of Y_i^ε in (14) for the subset mechanism has a nontrivial null space.*

Proof We can write the covariance matrix as a block matrix, where we will not care about the bottom right $g \times g$ block, since it will not be touched with the vector $(1, \cdots, 1, 0, \cdots, 0)$. Let Σ be the covariance matrix of $Z_i^\varepsilon W_i$, which we have actually computed in Lemma 6 with success probability $p = 1$. The top left $g \times g$ block matrix in $C(\cdot)$ will then be Σ, which has the all ones vector in its null space. Furthermore, the bottom left block matrix (as well as the top right, since the covariance matrix is symmetric) will be $\mu\Sigma$ where $\mu \in \mathbb{R}$ is the common mean across all groups under the null hypothesis. This gives us what we need for the lemma statement.

$$C(\pi, \mu, \sigma_1^2, \cdots, \sigma_g^2; \varepsilon) \begin{pmatrix} 1 \\ \vdots \\ 1 \\ 0 \\ \vdots \\ 0 \end{pmatrix} = \begin{bmatrix} \Sigma\mathbf{1} + \mu\Sigma\mathbf{0} \\ \mu\Sigma\mathbf{1} + 0 \end{bmatrix} = 0$$

■

We now turn to computing estimates for the population parameters to plug into our covariance matrix.

$$\hat{\mu} = \frac{\sum_{j=1}^{g}\sum_{i=1}^{n} Y_i^\varepsilon[g+j]}{nk},$$
$$\hat{\pi} = \left(\frac{\left(\binom{g-1}{k-1}e^\varepsilon + \binom{g-1}{k}\right) \cdot \left(\frac{\sum_{i=1}^{n} Y_i^\varepsilon[j]}{n}\right) - \left(e^\varepsilon \binom{g-2}{k-2} + \binom{g-2}{k-1}\right)}{e^\varepsilon \left(\binom{g-1}{k-1} - \binom{g-2}{k-2}\right) - \binom{g-2}{k-1}} : j \in [g] \right).$$

Lastly, we need to estimate the variance for each group, which we will again assume for ease of notation that $\sigma_j = \sigma$ for all groups $j \in [g]$. We will use the sample variance s_j^2 for $Y_i^\varepsilon[g+j]$ within each group $j \in [g]$. We then use the following estimate

$$\hat{\sigma}^2 = \frac{\sum_{j=1}^{g} s_j^2 - \mu^2 \left(k - \frac{1}{\left(\binom{g-1}{k-1}e^\varepsilon + \binom{g-1}{k}\right)^2} \sum_{j=1}^{g} \left(\binom{g-1}{k-1}e^\varepsilon \pi_j + (1-\pi_j)\left(\binom{g-2}{k-2}e^\varepsilon + \binom{g-2}{k-1}\right)\right)^2\right)}{k}$$

Hence, we will compare the resulting χ^2 statistic with the χ^2 distribution with $g-1$ degrees of freedom. As noted earlier, if any test computes a group probability $\hat{\pi}_j$ so that $n \cdot \hat{\pi} \leq 5$, we simply reject the null hypothesis.

D.4. Results

We show that the tests we develop achieve higher empirical power rather than simply using the classical one-way ANOVA tests after privatizing the groups. We present plots in Figure 15 that shows for data generated with the subset mechanism at various privacy levels, the

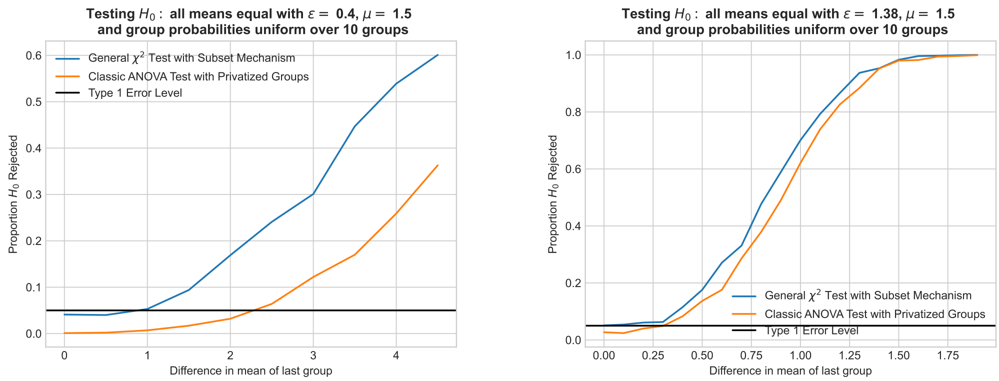

Figure 15: Comparing various Local Group DP mechanisms with corresponding one-way ANOVA test for testing whether there is a difference in means across different sensitive groups with various ε and $n = 10000$.

general χ^2 test that accounts for the subset mechanism outperforms using the classic ANOVA test, which does not account for the privacy mechanism. We point out that when ε gets larger, the two tests seem to perform similarly, similar to what we saw in Section B.4 with testing multiple proportions. All plots consist of the proportion of times the null hypothesis was rejected over 1000 trials.

D.5. Testing Difference in Two Groups

Once we have determined that there is indeed a difference across all group means, one typically wants to compute confidence intervals for the mean between two specific groups. In fact, one may want to directly compute a confidence interval between two groups, although the data has been privatized over several groups. In our privacy setting, we do not want to privatize the group membership each time we want to conduct a test, hence we will have samples with privatized groups which will mix with the two specific groups we want to compute a confidence interval for their difference. Ideally, we would privatize only the two groups we are interested in, but this would increase the privacy loss, something we want to avoid. Hence, we show how we can still obtain valid confidence intervals between two specific groups although the data has been privatized over $g > 2$ groups.

Since we are only interested in the difference in means between two groups, say $H_0 : \mu_j = \mu_\ell + \Delta$ for $j, \ell \in [g]$, we can change the optimization of the general χ^2 statistic to allow for any mean μ_m where $m \neq j, \ell$ based on the data samples. This will reduce our degrees of freedom of the asymptotic χ^2 by $g - 1$, thus if we privatize the groups with the Subset Mechanism, then the χ^2 statistic, after optimizing over all μ_m for $m \neq j, \ell$ and $\mu_j = \mu_\ell + \Delta$, should be compared to a χ^2 with 1 degree of freedom.

We show in Figure 16 that testing whether two means are equal can lead to invalid results if we were to use the classical t-test after the groups have been privatized. This is in contrast to when we would privatize only two groups, where the classical t-test empirically achieve the target level of Type 1 error, see Figure 5. The general χ^2 approach can then be used to

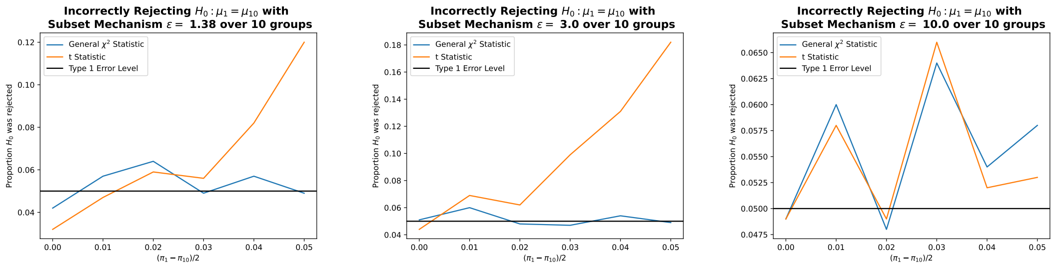

Figure 16: We give the proportion of times in 1000 trials that the classical t-test incorrectly rejects $H_0 : \mu_1 = \mu_{10}$ after we privatize the membership of $g = 10$ groups for each sample. We use the same standard deviation $\sigma = 2$ across all groups and have $\mu_j = 1.0$ for all $j \neq 1, 10$, $\mu_{10} = 1.5$ and $n = 10000$. We change the group probability π from uniform across all 10 groups and then change the first and last group probabilities.

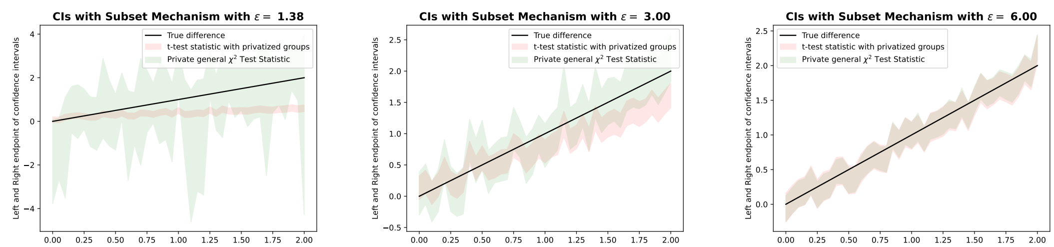

Figure 17: We give the confidence intervals between the first and 10th group mean when computed with the classical t-test and the general χ^2 test after we privatize the membership of $g = 10$ groups for each sample. We use the same standard deviation $\sigma = 2$ across all groups and have $\mu_j = 1.0$ for all $j \neq 1, 10$, $\mu_{10} = 1.5$ and $n = 10000$. The group probability π is uniform across all 10 groups except $\pi_1 = 0.15$ and $\pi_{10} = 0.05$.

achieve the target level of Type 1 error. We also give confidence intervals for the difference in mean in Figure 17 where the classical t-test misses the true difference while the general χ^2 test overlaps the true difference. The difference between the two tests becomes more pronounced when the group probabilities between the two groups of interest differ from each other.

Appendix E. A/B Testing

For this application, we assume that samples are randomly assigned to either a treatment or a control variant in an A/B test. We will denote the random variable $T_i \sim \text{Bern}(\lambda)$ to determine whether sample i is in the treatment $T_i = 1$ or in the control $T_i = 0$ group. Note that the parameter $\lambda \in [0, 1]$ is known and does not need estimating. In the treatment set of samples, data in group $j \in \{1,2\}$ will follow $X_i[j,t] \sim \mathtt{N}\left(\mu_{j,t}, \sigma^2_{j,t}\right)$, while in the control set of samples, data in group $j \in \{1,2\}$ will follow $X_i[j,c] \sim \mathtt{N}\left(\mu_{j,c}, \sigma^2_{j,c}\right)$. Again, we will let $W_i \sim \text{Bern}(\pi)$ determine the group that sample i belongs to, i.e. $W_i = 0$ for group 1 and $W_i = 1$ for group 2. Our goal here is to test whether the differences in means in the two groups has changed between the treatment and control. That is, we test $H_0 : \mu_{1,t} - \mu_{2,t} = \mu_{1,c} - \mu_{2,c}$. In order to compute confidence intervals, we will include a Δ term in the difference, so that we test $H_0 : \mu_{1,t} - \mu_{2,t} = \mu_{1,c} - \mu_{2,c}$, but $\Delta = 0$ is the typical hypothesis test.

We are not considering the membership of a sample to the control or treatment to be sensitive, and hence not privatizing it. Instead, we privatize the group membership $j \in \{1,2\}$ for each sample using randomized response. Recall that for randomized response, we will use random variables $Z_i^\varepsilon[j,j] \sim \text{Bern}(\frac{e^\varepsilon}{e^\varepsilon+1})$ for $j \in \{1,2\}$ with $Z_i^\varepsilon[1,2] = 1 - Z_i^\varepsilon[1,1]$ and $Z_i^\varepsilon[2,1] = 1 - Z_i^\varepsilon[2,2]$. We will consider the random vector $Y^\varepsilon = \sum_{i=1}^n Y_i^\varepsilon$ in our general χ^2 test framework, where

$$
Y_i^\varepsilon = \begin{pmatrix} T_i \cdot (Z_i^\varepsilon[1,1] \cdot W_i + Z_i^\varepsilon[1,2] \cdot (1 - W_i)) + (1 - T_i) \cdot (Z_i^\varepsilon[1,1] \cdot W_i + Z_i^\varepsilon[1,2] \cdot (1 - W_i)) \\ T_i \cdot (Z_i^\varepsilon[1,1] \cdot W_i \cdot X_i[1,t] + Z_i^\varepsilon[1,2] \cdot (1 - W_i) \cdot X_i[2,t]) \\ T_i \cdot (Z_i^\varepsilon[2,1] \cdot W_i \cdot X_i[1,t] + Z_i^\varepsilon[2,2] \cdot (1 - W_i) \cdot X_i[2,t]) \\ (1 - T_i) \cdot (Z_i^\varepsilon[1,1] \cdot W_i \cdot X_i[1,c] + Z_i^\varepsilon[1,2] \cdot (1 - W_i) X_i[2,c]) \\ (1 - T_i) \cdot (Z_i^\varepsilon[2,1] \cdot W_i \cdot X_i[1,c] + Z_i^\varepsilon[2,2] \cdot (1 - W_i) X_i[2,c]) \end{pmatrix}
$$

Since T_i is independent of the other variables, a lot of the calculations we have already done for t-tests in Section C.2 can be used. Next we compute its expectation in terms of the population parameters, where we will write $\boldsymbol{\mu} = (\mu_{1,t}, \mu_{2,t}, \mu_{1,c}, \mu_{2,c})$,

$$
\boldsymbol{\theta}^\varepsilon(\pi, \boldsymbol{\mu}; \varepsilon, \lambda) = \mathbb{E}[Y_i^\varepsilon] = \begin{pmatrix} \frac{e^\varepsilon}{e^\varepsilon+1}\pi + \frac{1}{e^\varepsilon+1}(1-\pi) \\ \lambda\left(\frac{e^\varepsilon}{e^\varepsilon+1}\pi\mu_{1,t} + \frac{1}{e^\varepsilon+1}(1-\pi)\mu_{2,t}\right) \\ \lambda\left(\frac{1}{e^\varepsilon+1}\pi\mu_{1,t} + \frac{e^\varepsilon}{e^\varepsilon+1}(1-\pi)\mu_{2,t}\right) \\ (1-\lambda)\left(\frac{e^\varepsilon}{e^\varepsilon+1}\pi\mu_{1,c} + \frac{1}{e^\varepsilon+1}(1-\pi)\mu_{2,c}\right) \\ (1-\lambda)\left(\frac{1}{e^\varepsilon+1}\pi\mu_{1,c} + \frac{e^\varepsilon}{e^\varepsilon+1}(1-\pi)\mu_{2,c}\right) \end{pmatrix}
$$

We next compute the covariance matrix $C(\pi, \boldsymbol{\mu}, \boldsymbol{\sigma}^2; \varepsilon, \lambda) = \mathbb{E}\left[Y_i^\varepsilon (Y_i^\varepsilon)^\intercal\right] - \mathbb{E}\left[Y_i^\varepsilon\right]\mathbb{E}\left[Y_i^\varepsilon\right]^\intercal$, where $\boldsymbol{\sigma}^2 = (\sigma_{1,t}^2, \sigma_{2,t}^2, \sigma_{1,c}^2, \sigma_{2,c}^2)$. We have

$$
\begin{aligned}
&\mathbb{E}\left[Y_i^\varepsilon (Y_i^\varepsilon)^\intercal\right][1,1] = \mathbb{E}\left[Y_i^\varepsilon\right][1] \\
&\mathbb{E}\left[Y_i^\varepsilon (Y_i^\varepsilon)^\intercal\right][1,2] = \lambda\left(\pi \tfrac{e^\varepsilon}{e^\varepsilon+1}\mu_{1,t} + (1-\pi)\tfrac{1}{e^\varepsilon+1}\mu_{2,t}\right) \\
&\mathbb{E}\left[Y_i^\varepsilon (Y_i^\varepsilon)^\intercal\right][1,3] = \mathbb{E}\left[Y_i^\varepsilon (Y_i^\varepsilon)^\intercal\right][1,5] = 0 \\
&\mathbb{E}\left[Y_i^\varepsilon (Y_i^\varepsilon)^\intercal\right][1,4] = (1-\lambda)\left(\pi \tfrac{1}{e^\varepsilon+1}\mu_{1,c} + (1-\pi)\tfrac{e^\varepsilon}{e^\varepsilon+1}\mu_{2,c}\right) \\
&\mathbb{E}\left[Y_i^\varepsilon (Y_i^\varepsilon)^\intercal\right][2,2] = \lambda\left(\pi \tfrac{e^\varepsilon}{e^\varepsilon+1}(\mu_{1,t}^2 + \sigma_{1,t}^2) + (1-\pi)\tfrac{1}{e^\varepsilon+1}(\mu_{2,t}^2 + \sigma_{2,t}^2)\right) \\
&\mathbb{E}\left[Y_i^\varepsilon (Y_i^\varepsilon)^\intercal\right][3,3] = \lambda\left(\pi \tfrac{1}{e^\varepsilon+1}(\mu_{1,t}^2 + \sigma_{1,t}^2) + (1-\pi)\tfrac{e^\varepsilon}{e^\varepsilon+1}(\mu_{2,t}^2 + \sigma_{2,t}^2)\right) \\
&\mathbb{E}\left[Y_i^\varepsilon (Y_i^\varepsilon)^\intercal\right][4,4] = (1-\lambda)\left(\pi \tfrac{e^\varepsilon}{e^\varepsilon+1}(\mu_{1,c}^2 + \sigma_{1,c}^2) + (1-\pi)\tfrac{1}{e^\varepsilon+1}(\mu_{2,c}^2 + \sigma_{2,c}^2)\right) \\
&\mathbb{E}\left[Y_i^\varepsilon (Y_i^\varepsilon)^\intercal\right][5,5] = (1-\lambda)\left(\pi \tfrac{1}{e^\varepsilon+1}(\mu_{1,c}^2 + \sigma_{1,c}^2) + (1-\pi)\tfrac{e^\varepsilon}{e^\varepsilon+1}(\mu_{2,c}^2 + \sigma_{2,c}^2)\right) \\
&\mathbb{E}\left[Y_i^\varepsilon (Y_i^\varepsilon)^\intercal\right][j,\ell] = 0, \quad j,\ell \in \{2,3,4,5\}, j \neq \ell.
\end{aligned}
$$

Under the null hypothesis $H_0 : \mu_{1,t} - \mu_{2,t} = \mu_{1,c} - \mu_{2,c} + \Delta$ we can solve for one of the means, so we will set $\mu_{1,t} = \mu_{1,c} - \mu_{2,c} + \mu_{2,t} + \Delta$ to reduce the number of parameters (note that we treat Δ as known). We now want to use our sample Y^ε to estimate the other means and the group probability $\hat{\pi}$. We start with the group probability, which we have estimated the same way in other sections:

$$
\hat{\pi} = (e^\varepsilon + 1)\left(\frac{Y^\varepsilon[0]/n - \frac{1}{e^\varepsilon+1}}{e^\varepsilon - 1}\right).
$$

We then solve for estimates of the means $\hat{\mu}_{1,t}, \hat{\mu}_{2,t}, \hat{\mu}_{1,c}, \hat{\mu}_{2,c}$ by setting the empirical averages $(\sum_{i=1}^n Y_i[j]/n : j \in \{2,3,4,5\})$ equal to the respective coordinates of $\boldsymbol{\theta}^\varepsilon(\hat{\pi}, \boldsymbol{\mu}; \varepsilon, \lambda)$ and solve for the means. Note that when we substitute in the null hypothesis $\mu_{1,t} = \mu_{2,t} + \mu_{1,c} - \mu_{2,c} + \Delta$, we get 4 equations and 3 unknowns. In this case, we choose the first three equations to solve for $\mu_{2,t}, \mu_{1,c}, \mu_{2,c}$ and then set $\hat{\mu}_{1,t} = \hat{\mu}_{2,t} + \hat{\mu}_{1,c} - \hat{\mu}_{2,c} + \Delta$. From these estimates, we can plug them into the covariance matrix. Note that we will not directly form estimates for $\boldsymbol{\sigma}^2$, instead we will compute the sample variance for $\{Y_i^\varepsilon[j]\}_{i=1}^n$ for $j \in \{2,3,4,5\}$ to use on the main diagonal of the covariance matrix and check to see that they indeed give valid sample variances (i.e. have positive variance of the true data), as we did for the test in Section C.2.

Our test statistic then becomes the following, where we will write $\hat{C}$ to denote the covariance matrix with the above parameter estimates and given $\varepsilon, \lambda, \Delta$,

$$
D^\varepsilon(\lambda, \Delta) = \min_{\substack{\pi \in (0,1), \\ \mu_{2,t}, \mu_{1,c}, \mu_{2,c} \in \mathbb{R} \\ \mu_{1,t} = \mu_{1,c} - \mu_{2,c} + \mu_{2,t} + \Delta}} \left\{ \left(Y^\varepsilon - \boldsymbol{\theta}^\varepsilon(\pi, \boldsymbol{\mu}; \varepsilon, \lambda)\right)^\intercal \hat{C}^{-1} \left(Y^\varepsilon - \boldsymbol{\theta}^\varepsilon(\pi, \boldsymbol{\mu}; \varepsilon, \lambda)\right) \right\}.
$$

We then compare the test statistic with a χ^2 random variable with 1 degrees of freedom. We give power results in Figure 18 for hypothesis tests comparing the general χ^2 test with the classical t-test on privatized groups, which seems to perform similarly. We also gave

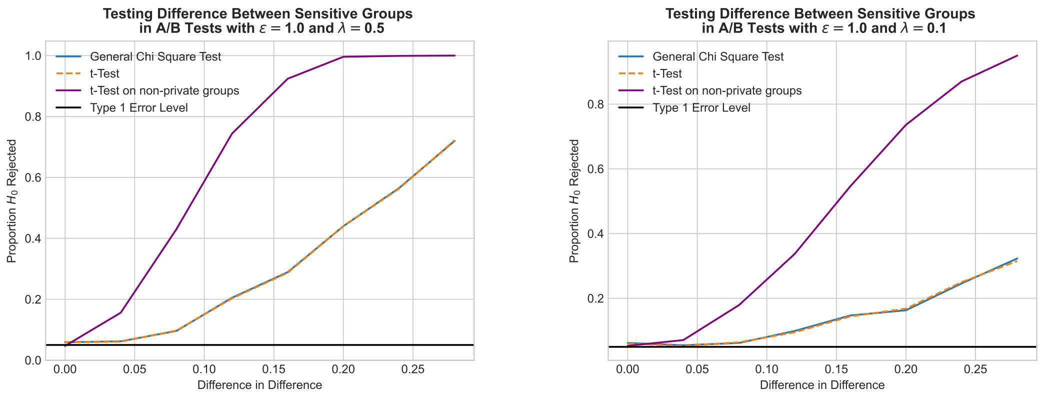

Figure 18: Power results for hypothesis testing for the difference in means across sensitive groups between treatment and control $H_0 : \mu_{1,t} - \mu_{2,t} = \mu_{1,c} - \mu_{2,c}$. We compare the (unmodified) t-test and the general χ^2 test on privatized groups with treatment probability $\lambda \in \{0.5, 0.1\}$ and $\varepsilon = 1$.

confidence intervals for the difference between means across treatment and control in Figure 3 using the general χ^2 test statistic and the unmodified t-test statistic. In our experimental setup we generate data with zero means across groups 1 and 2 in both treatment and control and keep the variance across all to be 1. We then will vary the mean in the control group of group 1, i.e. $\mu_{1,c}$ to change the difference between groups in the treatment and control.