
Differentially Private Conditional Independence Testing

Iden Kalemaj
Boston University

Shiva Prasad Kasiviswanathan
Amazon

Aaditya Ramdas
Carnegie Mellon University

Abstract

Conditional independence (CI) tests are widely used in statistical data analysis, e.g., they are the building block of many algorithms for causal graph discovery. The goal of a CI test is to accept or reject the null hypothesis that $X \perp\!\!\!\perp Y \mid Z$, where $X \in \mathbb{R}, Y \in \mathbb{R}, Z \in \mathbb{R}^d$. In this work, we investigate conditional independence testing under the constraint of differential privacy. We design two private CI testing procedures: one based on the generalized covariance measure of Shah and Peters (2020) and another based on the conditional randomization test of Candès et al. (2016) (under the model-X assumption). We provide theoretical guarantees on the performance of our tests and validate them empirically. These are the first private CI tests with rigorous theoretical guarantees that work for the general case when Z is continuous.

1 INTRODUCTION

Conditional independence (CI) tests are a powerful tool in statistical data analysis, e.g., they are building blocks for graphical models, causal inference, and causal graph discovery [9, 20, 26]. These analyses are frequently performed on sensitive data, such as clinical datasets and demographic datasets, where concerns for privacy are foremost. For example, in clinical trials, CI tests are used to answer fundamental questions such as “After accounting for (conditioning on) a set of patient covariates Z (e.g., age or gender), does a treatment X lead to better patient outcomes Y ?”. Formally, given three random variables (X, Y, Z) where $X \in \mathbb{R}, Y \in \mathbb{R},$ and $Z \in \mathbb{R}^d$, denote the conditional independence of X and Y given Z by $X \perp\!\!\!\perp Y \mid Z$. Our problem is that of testing

$$H_0 \text{ (null)} : X \perp\!\!\!\perp Y \mid Z \text{ vs. } H_1 \text{ (alternate)} : X \not\perp\!\!\!\perp Y \mid Z$$

given data drawn i.i.d. from a joint distribution of (X, Y, Z) . CI testing is a much harder problem than (unconditional) independence testing, where the variable Z is omitted. Indeed, Shah and Peters [30] showed that CI testing is a statistically impossible task for continuous random variables.¹ Thus, techniques for independence testing *do not* extend to the CI testing problem.

When the underlying data is sensitive and confidential, publishing statistics (such as the value of a CI independence test statistic or the corresponding p-value) can leak private information about individuals in the data. For instance, Genome-Wide Association Studies (GWAS) involve finding (causal) relations between Single Nucleotide Polymorphisms (SNPs) and diseases. CI tests are building blocks for establishing these relations, and the existence of a link between a specific SNP and a rare disease may indicate the presence of a minority patient. *Differential privacy* [13] is a widely studied and deployed formal privacy guarantee for data analysis. The output distributions of a differentially private algorithm must look nearly indistinguishable for any two input datasets that differ only in the data of a single individual. In this work, we design the first differentially private (DP) CI tests that can handle continuous variables.

Our Contributions. We design two private CI tests, each based on a different set of assumptions about the data-generating distribution. We provide the *first* private CI test with rigorous type-I error and power guarantees. Given the aforementioned impossibility results for non-private CI testing, to obtain a CI test with meaningful theoretical guarantees, some assumptions are necessary; in particular we must restrict the space of possible null distributions. In designing our private tests, we start with non-private CI tests that provide rigorous guarantees on type-I error control.

Our first test (Section 2) is a private version of the generalized covariance measure (GCM) by Shah and Peters [30]. The type-I error guarantees of the GCM rely on the fairly weak assumption that the conditional means $\mathbb{E}[X \mid Z]$ and $\mathbb{E}[Y \mid Z]$ can be estimated sufficiently well given the dataset size. The test statistic of the GCM is a normalized sum of

¹Any test that uniformly controls the type-I error (false positive rate) for all absolutely continuous triplets (X, Y, Z) such that $X \perp\!\!\!\perp Y \mid Z$, even asymptotically, does not have nontrivial power against *any* alternative.

the product of the residuals of (nonlinearly) regressing X on Z and Y on Z . This test statistic has *unbounded* sensitivity, thus a more careful way of adding and analyzing the impact of the privacy noise is needed. Our private GCM test adds appropriately scaled, zero-mean noise to the residual products, and calculates the same statistic on the noisy residual products. We show that even with the added noise, the GCM score converges asymptotically to a standard Gaussian distribution under the null hypothesis. The magnitude of the noise added to the residuals is constant (it does not vanish with increasing sample size n), thus a priori it is surprising that convergence results can be obtained in the presence of such noise. We do so by bounding how the added noise interacts with the noise from the estimation of the residuals. While our asymptotics justify the threshold for rejecting the null, our private GCM test controls type-I error very well at small finite n , as we demonstrate empirically (because the central limit theorem kicks in rather quickly). Recall that finite sample guarantees on type-I error and power are impossible even for non-private CI testing [30].

The private GCM is the *first* private CI test with rigorous power guarantees. It achieves the same power as the non-private GCM test with a $O(1/\epsilon^2)$ -inflation of the dataset size. In addition, while private extensions of non-private hypothesis tests often suffer from introducing excessive type-I error (see the Related Work), the private GCM exhibits the opposite behavior: it can maintain type-I error control even when the non-private GCM test fails to do so. This occurs in scenarios where the regression methods used to estimate the conditional means either underfit or overfit.

Our second test (Section 3) relies on the *model-X assumption* that the conditional distribution of $X \mid Z$ is known or can be well-approximated. Recently introduced by Candès et al. [6], this assumption is useful in settings where one has access to abundant unlabeled data, such as in GWAS, but labeled data are scarce. The model-X assumption is also satisfied in experimental settings where a randomization mechanism is known or designed by the experimenter. CI tests utilizing this assumption provide exact, non-asymptotic, type-I error control [6, 3], thus bypassing the hardness result of Shah and Peters [30]. While this assumption has spurred a lot of recent research in (non-private) CI testing, there are no prior private tests in the literature that are designed to work under this assumption. In this work, we focus on the conditional randomization test (CRT) [6]. We design a private CRT and provide theoretical guarantees on the accuracy of its p-value. We adopt a popular framework for obtaining DP algorithms, known as Report Noisy Max (or the exponential mechanism), which requires defining a problem-specific score function of low sensitivity. To obtain good utility, our score function exploits the specific distribution of intermediate statistics calculated by the CRT. This score function is novel and can be used for solving a more general problem: given a set of queries on a dataset,

estimate privately the rank of a particular query amongst the rest of the queries.

We present a detailed empirical evaluation of the proposed tests, justifying their practicality across a wide range of settings. Our experiments confirm that our private CI tests provide the critical type-I error control, and can in fact do so more reliably than their non-private counterparts. As expected, our private tests achieve lower power due to the noise injected for privacy, which can be compensated for with a larger dataset size.

1.1 Related Work

Private Conditional Independence Testing. Wang et al. [37] is the only work, prior to ours, to explicitly study private CI testing, motivated by an application to causal discovery. Their tests (obtained from Kendall’s τ and Spearman’s ρ score) are designed for categorical Z . While these tests could be adapted to work for continuous Z via clustering, in practice this method does not seem to control type-I error, as we show in Fig. 1. The problem worsens with higher-dimensional Z . Our techniques also differ from those of Wang et al. [37], who obtain their tests by bounding the sensitivity of non-private CI scores and adding appropriately scaled noise to the true value of the score. They state two open problems: obtaining private CI tests for continuous Z and obtaining private tests from scores of unbounded sensitivity (as is the case with the GCM score). We solve both open problems, and manage to privatize the GCM score by instead adding noise to an intermediate statistic, the residuals of fitting X to Z and Y to Z .

Another line of work [31, 18, 27] has utilized the “subsample and aggregate” framework of differential privacy [25] to obtain private versions of existing hypothesis tests in a black-box fashion. In this approach, the dataset is partitioned into k smaller datasets; the non-private hypothesis test is evaluated on the smaller datasets; and finally, the results are privately aggregated. Based on this method, Kazan et al. [18] propose a test-of-tests (ToT) framework to construct a private version of any known (non-private) hypothesis test. However, they show guarantees on the power of their test based on finite-sample guarantees of the power of the non-private hypothesis test. Since finite-sample guarantees are impossible for CI testing, their method gives *no* power guarantees for CI testing, and thus cannot be reliably used in practice. In addition, in Fig. 1 we compare the type-I error control of our tests with the ToT framework and show that it can fail to control type-I error.

Smith [31] analyzed the asymptotic properties of subsample-and-aggregate and showed that for a large family of statistics, one can obtain a corresponding DP statistic with the same asymptotic distribution as the original statistic. In particular, the result of Smith [31] can be applied to obtain a DP version of the GCM statistic. However, compared to our results

on the private GCM, (a) only a weaker notion of privacy, known as *approximate* DP, would be guaranteed, and (b) an additional condition on the data-generating distribution would have to be introduced, to guarantee a bounded third moment of the GCM statistic. Finally, the test of Peña and Barrientos [27] only outputs a binary accept/reject decision and not a p-value as our tests provide, and was empirically outperformed by the test of Kazan et al. [18].

Private (non-conditional) Independence Testing. A line of work on private independence testing has focused on privatizing the chi-squared statistic [36, 17, 34, 41, 38, 16, 28]. These tests operate with categorical X and Y . Earlier works obtained private hypothesis tests by adding noise to the histogram of the data [17], but it was later pointed out that this approach does not provide reliable type-I error control at small sample sizes [14]. Consequent works used numerical approaches to obtain the distribution of the noisy statistic and calculate p-values with that distribution [34, 41, 38, 16], whereas Rogers and Kifer [28] obtain new statistics for chi-squared tests whose distribution after the privacy noise can be derived analytically. In this light, one important feature of our private GCM test is that its type-I error control can be more reliable than for the non-private GCM, even at small n , as our experiments demonstrate. For continuous X and Y , Kusner et al. [21] obtained DP versions of several dependence scores (Kendall’s τ , Spearman’s ρ , HSIC), however, they do not provide type-I error or power guarantees. In follow-up work, Kim and Schrab [19] obtained private versions of permutation tests, that were applied to kernel-based independence tests, such as HSIC. Note that CI testing is a much harder task than independence testing, and techniques for the latter do not necessarily translate to CI testing. Our work is part of the broader literature on private hypothesis testing [2, 5, 33, 8, 39, 1, 4, 10, 35].

1.2 Preliminaries

Notation. If $(V_{P,n})_{n \in \mathbb{N}, P \in \mathcal{P}}$ is a family of sequences of random variables whose distributions are determined by $P \in \mathcal{P}$, we say $V_{P,n} = o_{\mathcal{P}}(1)$ if for all $\delta > 0$, $\sup_{P \in \mathcal{P}} \Pr_P[|V_{P,n}| > \delta] \rightarrow 0$. Similarly, $V_{P,n} = O_{\mathcal{P}}(1)$ if for all $\delta > 0$, $\exists M > 0$ such that $\sup_{n \in \mathbb{N}} \sup_{P \in \mathcal{P}} \Pr_P[|V_{P,n}| > M] < \delta$.

The notion of neighboring datasets is central to differential privacy. In this work, we consider datasets $\mathbf{D} = (\mathbf{X}, \mathbf{Y}, \mathbf{Z})$ of n datapoints $\{(x_i, y_i, z_i)\}_{i=1}^n$, drawn i.i.d. from a joint distribution P on some domain $\mathcal{X} \times \mathcal{Y} \times \mathcal{Z}$. Let \mathcal{D} denote the universe of datasets. A dataset $\mathbf{D}' \in \mathcal{D}$ is a neighbor of \mathbf{D} if it can be obtained from \mathbf{D} by replacing at most one datapoint $(x_i, y_i, z_i) \in \mathbf{D}$ with an arbitrary entry $(x'_i, y'_i, z'_i) \in \mathcal{X} \times \mathcal{Y} \times \mathcal{Z}$, for some $i \in [n]$.

Definition 1.1 ([13]). A randomized algorithm Alg is ϵ -DP if for all neighboring datasets \mathbf{D}, \mathbf{D}' and all events \mathcal{R} in the output space of Alg, it holds $\Pr[\text{Alg}(\mathbf{D}) \in \mathcal{R}] \leq$

$\exp(\epsilon) \cdot \Pr[\text{Alg}(\mathbf{D}') \in \mathcal{R}]$, where the probability is over the randomness of the algorithm.

Additional background on hypothesis testing (including standard definitions of p-value, type-I error, uniform asymptotic level, power, etc.) and differential privacy is provided in Appendices A.2 and A.3. In our algorithms and experiments, we use kernel ridge regression (KRR) as a procedure for regressing \mathbf{X} and \mathbf{Y} on \mathbf{Z} , and rely on the following result by Kusner et al. [21] about the sensitivity of its residuals.²

Theorem 1.2 (Restated Theorem 5 of Kusner et al. [21]). *Let (\mathbf{U}, \mathbf{V}) be a dataset of n datapoints (u_i, v_i) , $i \in [n]$ from the domain $\mathcal{U} \times \mathcal{V} \subset \mathbb{R} \times \mathbb{R}^d$. Suppose that $|\mathcal{U}| \leq 1$. Given a Hilbert space \mathcal{H} , let \mathbf{w} be the vector that minimizes the kernel ridge regression objective $(\lambda/2)\|\mathbf{w}\|_{\mathcal{H}}^2 + (1/n)\sum_{i=1}^n (u_i - \mathbf{w}^\top \phi(v_i))^2$, for kernel $\phi: \mathbb{R}^d \rightarrow \mathcal{H}$ with $\|\phi(v)\|_{\mathcal{H}} \leq 1$ for all $v \in \mathcal{V}$. Define \mathbf{w}' analogously for a neighboring dataset $(\mathbf{U}', \mathbf{V}')$ that is obtained by replacing one datapoint in (\mathbf{U}, \mathbf{V}) . Then $\|\mathbf{w}\|_{\mathcal{H}} \leq \sqrt{2/\lambda}$ and for all $v \in \mathbf{V}$ it holds: $|\mathbf{w}^\top \phi(v) - \mathbf{w}'^\top \phi(v)| \leq 8\sqrt{2}/(\lambda^{3/2}n) + 8/(\lambda n)$.*

2 PRIVATE GCM TEST

Here, we present our private Generalized Covariance Measure (GCM) test. Missing proofs are in Appendix B.

GCM Test. We first describe the non-private GCM test of Shah and Peters [30]. Given a joint distribution P of the random variables (X, Y, Z) , we can always write:

$$X = f_P(Z) + \chi_P, \quad Y = g_P(Z) + \xi_P,$$

where $f_P(z) = \mathbb{E}_P[X|Z = z]$, $g_P(z) = \mathbb{E}_P[Y|Z = z]$, $\chi_P = X - f_P(z)$, and $\xi_P = Y - g_P(z)$.

Let $\mathbf{D} = (\mathbf{X}, \mathbf{Y}, \mathbf{Z})$ be a dataset of n i.i.d. samples from P . Let \hat{f} and \hat{g} be approximations to the conditional expectations f_P and g_P , obtained by fitting \mathbf{X} to \mathbf{Z} and \mathbf{Y} to \mathbf{Z} , respectively. We consider the products of the residuals from the fitting procedure:

$$R_i = ((x_i - \hat{f}(z_i))(y_i - \hat{g}(z_i))) \text{ for } i \in [n]. \quad (1)$$

The GCM test statistic T is defined as the normalized mean of the residual products, i.e.,

$$T(R_1, \dots, R_n) = \frac{\frac{1}{\sqrt{n}} \sum_{i=1}^n R_i}{\left(\frac{1}{n} \sum_{i=1}^n R_i^2 - \left(\frac{1}{n} \sum_{k=1}^n R_k\right)^2\right)^{1/2}}. \quad (2)$$

The normalization critically ensures that T follows a standard normal distribution asymptotically. However, it also leads to the unbounded sensitivity of the statistic T .

²One could also use other regression techniques within our private GCM and private CRT frameworks, and theoretical guarantees continue to hold if similar ($\approx O(1/n)$) bounds on the sensitivity of the residuals hold for those methods.

Private GCM Test. To construct a DP version of the GCM test, we focus on the vector of residual products, $\mathbf{R} = (R_1, \dots, R_n)$. Let Δ denote the ℓ_1 -sensitivity of \mathbf{R} . Given Δ , we use the Laplace mechanism (Lemma A.2) to add scaled Laplace noise to \mathbf{R} and then compute T on the noisy residual products. The private GCM test we present (in Algorithm 1) can be used with any fitting procedure, as long as a bound on the sensitivity of the residuals for that procedure is known.

Algorithm 1 Private Generalized Covariance Measure

- 1: **Input:** Dataset $(\mathbf{X}, \mathbf{Y}, \mathbf{Z}) = \{(x_i, y_i, z_i)\}_{i=1}^n$, privacy parameter $\varepsilon > 0$, fitting procedure \mathcal{F} , bound $\Delta > 0$ on the sensitivity of residual products of \mathcal{F} .
 - 2: Let $\hat{f} = \mathcal{F}(\mathbf{Z}, \mathbf{X})$ and $\hat{g} = \mathcal{F}(\mathbf{Z}, \mathbf{Y})$
 - 3: **for** $i = 1, \dots, n$ **do**
 - 4: $r_{X,i} \leftarrow x_i - \hat{f}(z_i)$, $r_{Y,i} \leftarrow y_i - \hat{g}(z_i)$
 - 5: $R_i \leftarrow r_{X,i} \cdot r_{Y,i}$
 - 6: $W_i \sim \text{Lap}(0, \Delta/\varepsilon)$
 - 7: Let $T^{(n)} \leftarrow T(R_1 + W_1, \dots, R_n + W_n)$ ▷ See (2)
 - 8: Output p-value = $2 \cdot (1 - \Phi(|T^{(n)}|))$
-

While the algorithm uses a simple noise addition strategy to privatize the GCM test, its asymptotic behavior is rather unexpected. Notice that the noise added to the residuals has constant variance that does not vanish with $n \rightarrow \infty$. Yet, we show (in Theorem 2.2) that $T^{(n)}$ converges to a standard Gaussian distribution similarly to the non-private GCM statistic. The key step in the analysis is to show that the error introduced by the noise random variables grows at a slower rate than the error introduced by the fitting procedure. Our algorithmic framework opens up a question of which fitting procedures have residuals with constant bounded sensitivity. We show such a result when Kernel Ridge Regression is used as the fitting method.

Next, we show theoretical guarantees on the type-I error control and power of Algorithm 1 under mild assumptions on the fitting procedures \hat{f} and \hat{g} , listed in Definition 2.1.

Definition 2.1 (Good fit). Consider $u_P(z) = \mathbb{E}_P[\chi_P^2 | Z = z]$, $v_P(z) = \mathbb{E}_P[\xi_P^2 | Z = z]$, and the following error terms:

$$\begin{aligned}
 A_f &= \frac{1}{n} \sum_{i=1}^n (f_P(z_i) - \hat{f}(z_i))^2, \\
 B_f &= \frac{1}{n} \sum_{i=1}^n (f_P(z_i) - \hat{f}(z_i))^2 v_P(z_i), \\
 A_g &= \frac{1}{n} \sum_{i=1}^n (g_P(z_i) - \hat{g}(z_i))^2, \\
 B_g &= \frac{1}{n} \sum_{i=1}^n (g_P(z_i) - \hat{g}(z_i))^2 u_P(z_i). \quad (3)
 \end{aligned}$$

Let \mathcal{P} be a class of distributions for (X, Y, Z) that are continuous with respect to the Lebesgue measure and such that

$\sup_{P \in \mathcal{P}} \mathbb{E}[|\chi_P \xi_P|^{2+\eta}] \leq c$ for some constants $c, \eta > 0$. We say the classifiers \hat{f} and \hat{g} are a good fit for \mathcal{P} if $A_f A_g = o_{\mathcal{P}}(n^{-1})$, $B_f = o_{\mathcal{P}}(1)$, and $B_g = o_{\mathcal{P}}(1)$.

The key part of Definition 2.1 is that the product of mean squared errors of the fitting method (for fitting X to Z and Y to Z) converge to zero at a sublinear rate in the sample size. This is known as the ‘‘doubly robust’’ assumption and is common in many analyses in theoretical ML [7]. It is a mild assumption because it only requires one of two regressions, either fitting X to Z or fitting Y to Z , to be sufficiently accurate. These assumptions are in fact slightly weaker than those of Shah and Peters [30] for guaranteeing uniformly asymptotic level³ and power of the GCM, as we do not require a lower bound on the variance $\mathbb{E}[\chi_P^2 \xi_P^2]$ of the true residuals. This requirement is no longer necessary as we add finite-variance noise to the residual products.

Type-I Error Control. We show that as with the GCM test of Shah and Peters [30], the private counterpart has uniformly asymptotic level. While the original GCM test of Shah and Peters [30] does not require the input variables to be bounded, we assume bounded random variables X and Y to obtain bounds on the sensitivity Δ of the residual products. For the rest of this section, we assume publicly known bounds a and b on the domain \mathcal{X} of X and \mathcal{Y} of Y , (i.e., $|x| \leq a, \forall x \in \mathcal{X}$ and $|y| \leq b, \forall y \in \mathcal{Y}$).⁴ Note that we do not assume such bounds on the domain of Z , which is important as Z could be high-dimensional.

Theorem 2.2. (Type-I Error Control of Private GCM) *Let a and b be known bounds on the domains of X and Y , respectively. Given a dataset $\mathbf{D} = (\mathbf{X}, \mathbf{Y}, \mathbf{Z})$, let $(\hat{\mathbf{X}}, \hat{\mathbf{Y}}, \mathbf{Z})$ be the rescaled dataset obtained by setting $\hat{\mathbf{X}} = \mathbf{X}/a$ and $\hat{\mathbf{Y}} = \mathbf{Y}/b$. Consider $R_i, i \in [n]$, as defined in (1), for the rescaled dataset $(\hat{\mathbf{X}}, \hat{\mathbf{Y}}, \mathbf{Z})$. Let $W_i \sim \text{Lap}(0, \Delta/\varepsilon)$ for $i \in [n]$, where $\Delta, \varepsilon > 0$ are constants. Let \mathcal{P} be the set of null distributions from Definition 2.1 for which \hat{f}, \hat{g} are a good fit. The statistic $T^{(n)} = T(R_1 + W_1, \dots, R_n + W_n)$, defined in Algorithm 1, converges uniformly over \mathcal{P} to the standard Gaussian distribution $\mathcal{N}(0, 1)$.*

Since $T^{(n)}$ converges uniformly to a standard Gaussian distribution, this implies that the CI test in Algorithm 1 has uniformly asymptotic level (see Appendix A.2). It also satisfies a weaker pointwise asymptotic level guarantee that holds under slightly weaker assumptions (Theorem B.1). Note that this result is independent of the bound on $|X|$ and $|Y|$. While the guarantee of Theorem 2.2 is asymptotic, type-I error control ‘‘kicks in’’ at very small sample sizes

³Given a level $\alpha \in (0, 1)$ and null hypothesis \mathcal{P}_0 , a test ψ_n has uniformly asymptotic level if its asymptotic type-I error is bounded by α over all distributions in \mathcal{P}_0 , i.e., $\lim_{n \rightarrow \infty} \sup_{P \in \mathcal{P}_0} \Pr_P[\psi_n \text{ rejects null}] \leq \alpha$. See Appendix A.2.

⁴These bounds can also be replaced with high probability bounds, but the privacy guarantees of our CI test would be replaced with what is known as *approximate differential privacy*.

(like $n = 100$) as confirmed by our experimental results. This behavior is typical for many statistics that converge to a standard Gaussian distribution. Recall that finite-sample guarantees are impossible even for non-private CI testing, even under the assumptions of Definition 2.1, since these assumptions are only asymptotic in nature.

Noise Addition Leads to Better Type-I Error Control. A beneficial consequence of the privacy noise is that there are scenarios, under the null hypothesis, where the non-private GCM fails to provide type-I error control, but our private GCM does. If the functions \hat{f} and \hat{g} fail to fit the data (i.e., the conditions on A_f, A_g, B_f, B_g in Theorem 2.2 are violated), private GCM can still provide type-I error control. We show in Section 4 one such scenario, when the learned model underfits the data. Consider on the other hand the case when the model overfits, and more extremely, when the model interpolates asymptotically, i.e. $\hat{f}(z_i) \rightarrow x_i$ and $\hat{g}(z_i) \rightarrow y_i$ as $n \rightarrow \infty$ for all $i \in [n]$ [22]. It is not too hard to show that convergence to the standard Gaussian still holds for the private GCM, and thus type-I error control is provided. Instead, the rejection rate of the non-private GCM converges to 1 when the model interpolates.

Power of the Private GCM. Next, we show a result on the power of our private GCM test.

Theorem 2.3. (Power of Private GCM). *Consider the setup of Theorem 2.2. Define the “signal” (ρ_P) and “noise” (σ_P) of the true residuals χ_P, ξ_P as:*

$$\rho_P = \mathbb{E}_P[\text{cov}(X, Y | Z)], \quad \sigma_P = \sqrt{\text{Var}_P(\chi_P \xi_P)}.$$

Let \mathcal{P} be the set of alternate distributions from Definition 2.1 for which \hat{f}, \hat{g} are a good fit. Then

$$T^{(n)} - \frac{\sqrt{n}\rho_P}{\sigma'_P}, \quad \text{where } \sigma'_P = \sqrt{\sigma_P^2 + \left(\frac{\sqrt{2}ab\Delta}{\varepsilon}\right)^2} \quad (4)$$

converges uniformly over \mathcal{P} to $\mathcal{N}(0, 1)$.

Discussion on Power. Theorem 2.3 implies that $T^{(n)}$ has uniform (asymptotic) power of 1 if $\rho_P \neq 0$. See Corollary B.5 for a short proof. In Theorem B.3, we also show a pointwise (asymptotic) power guarantee, under weaker assumptions. We remark that the bounds a and b on $|X|$ and $|Y|$ could depend on the dataset size n . Algorithm 1 has uniform asymptotic power of 1 as long as $a \cdot b = o(\sqrt{n})$.

Shah and Peters [30] show a similar result on the power of the (non-private) GCM, but with $\sigma'_P = \sigma_P$. Suppose $\sigma_P = 1$. Then, Theorem 2.3 states that a $(\frac{ab\Delta}{\varepsilon})^2$ -factor of the dataset size used in the non-private case is required to obtain the same power in the private case. A blow-up in the sample size is typical in DP analyses [11].

Private GCM with Kernel Ridge Regression (PrivGCM). To obtain a bound on the sensitivity of the vector of residual products, we use kernel ridge regression (KRR) as the

model for regressing X on Z and Y on Z , respectively. Let PrivGCM denote Algorithm 1 with KRR as the fitting procedure and the bound on Δ .

The vector of residual products has ℓ_1 -sensitivity $O_\lambda(1)$ (as formally shown in Lemma B.6 using Theorem 1.2). Along with Lemma A.2, this implies that PrivGCM is ε -DP. In addition, as shown by Shah and Peters [30], the requirements on A_f, A_g, B_f, B_g are satisfied when using KRR. If the additional conditions listed in Definition 2.1 are also satisfied, then PrivGCM has uniformly asymptotic level and uniform asymptotic power of 1 (see Corollary B.7).

3 PRIVATE CRT

Here, we propose a private version of the conditional randomization test (CRT), which uses access to the distribution of $X | Z$ as a key assumption. Recall that such an assumption is useful, for example, when one has access to abundant unlabeled data (X, Z) . Missing proofs are in Appendix C.

CRT. As before, consider a dataset $(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$ of n i.i.d. samples $(x_i, y_i, z_i), i \in [n]$ from the joint distribution P . For ease of notation, denote the original \mathbf{X} as $\mathbf{X}^{(0)}$. The key idea of CRT is to sample m copies of $\mathbf{X}^{(0)}$ from $X | Z$, where Z is fixed to the values in \mathbf{Z} . That is, for $j \in [m]$ and $i \in [n]$, a new datapoint $x_i^{(j)}$ is sampled from $X | Z = z_i$. Then $\mathbf{X}^{(j)} = (x_1^{(j)}, \dots, x_n^{(j)})$.

Under the null hypothesis, the triples $(\mathbf{X}^{(0)}, \mathbf{Y}, \mathbf{Z}), \dots, (\mathbf{X}^{(m)}, \mathbf{Y}, \mathbf{Z})$ are identically distributed. Thus, for every statistic T chosen independently of the data, the random variables $T(\mathbf{X}^{(0)}, \mathbf{Y}, \mathbf{Z}), \dots, T(\mathbf{X}^{(m)}, \mathbf{Y}, \mathbf{Z})$ are also identically distributed. Denote these random variables by T_0, \dots, T_m . The p-value is computed by ranking T_0 , obtained by using the original $\mathbf{X}^{(0)}$ vector, against T_1, \dots, T_m , obtained from the resamples:

$$\text{p-value} = \frac{1 + \sum_{j=1}^m \mathbf{1}(T_j \geq T_0)}{m + 1}.$$

For every choice of T , the p-value is uniformly distributed and finite-sample type-I error control is guaranteed.

Private CRT. Let $\mathbf{D} = (\mathbf{X}^{(0)}, \dots, \mathbf{X}^{(m)}, \mathbf{Y}, \mathbf{Z})$ denote the aggregated dataset. We say \mathbf{D}' is a neighbor of \mathbf{D} if they differ in at most one row. By defining \mathbf{D} to include the resamples $\mathbf{X}^{(1)}, \dots, \mathbf{X}^{(m)}$, we also protect the privacy of the data obtained in the resampling step.

Our private CRT test is shown in Algorithm 2: it obtains a private estimate of the rank of T_0 amongst the statistics T_1, \dots, T_m , sorted in decreasing order. Using the Laplace mechanism to privately estimate the rank is not a viable option, since the rank has high sensitivity: changing one point in \mathbf{D} could change all the values T_0, \dots, T_m and change the rank of T_0 by $O(m)$. Another straightforward approach is to employ the widely used Sparse Vector Technique [12, 11]

to privately answer questions “Is $T_i > T_0$?” for all $i \in [m]$. However, this algorithm pays a privacy price for each T_i that is above the “threshold” T_0 , which under the null is $\Omega(m)$, thus resulting in lower utility of the algorithm. Instead, we define a new score function and algorithm which circumvents this problem by intuitively only incurring a cost for the queries T_i that are very close to T_0 in value.

Our key algorithmic idea is to define an appropriate score function of bounded sensitivity. It assigns a score to each rank $c \in [0, m]$ that indicates how well c approximates the true rank of T_0 . The score of a rank c equals the negative absolute difference between T_0 and the statistic at rank c . The true rank of T_0 has the highest score (equal to 0), whereas all other ranks have negative scores. We show that this score function has bounded sensitivity for statistics T of bounded sensitivity. The rank with the highest score is privately selected using Report Noisy Max, a popular DP selection algorithm [11]. To obtain good utility, the design of the score function exploits the fact that for CRTs, the values T_i are distributed in a very controlled fashion, as explained in the remark following Theorem 3.6.

Algorithm 2 Private Conditional Randomization Test

- 1: **Input:** Dataset $(\mathbf{X}^{(0)}, \mathbf{Y}, \mathbf{Z})$, privacy parameter ε , bound Δ_T on the sensitivity of T , resamples m .
 - 2: $T_0 \leftarrow T(\mathbf{X}^{(0)}, \mathbf{Y}, \mathbf{Z})$, $s_0 \leftarrow 0$.
 - 3: **for** $i = 1, \dots, m$ **do**
 - 4: Sample $\mathbf{X}^{(i)} \mid \mathbf{Z}$ from $X \mid Z$.
 - 5: $T_i \leftarrow T(\mathbf{X}^{(i)}, \mathbf{Y}, \mathbf{Z})$.
 - 6: Let Q_0, \dots, Q_m denote the values $\{T_i\}_{i \in [0, m]}$ sorted in decreasing order.
 - 7: **for** $i = 0, 1, \dots, m$ **do**
 - 8: $s_i \leftarrow -\frac{|Q_i - T_0|}{2\Delta_T}$.
 - 9: $\hat{c} \leftarrow \text{ReportNoisyMax}(\{s_i\}_{i \in [0, m]}, \varepsilon)$. \triangleright Theorem 3.3
 - 10: Output p-value $\hat{p} = \frac{1 + \hat{c}}{m + 1}$.
-

Definition 3.1 (Score function for rank of query). Let $\{T_i\}_{i \in [0, m]}$ be $m + 1$ queries of sensitivity Δ_T on dataset \mathbf{D} . Let Q_0, \dots, Q_m denote the values $\{T_i\}_{i \in [0, m]}$ sorted in decreasing order. Let $k \in [0, m]$ be the index of the query whose rank we wish to know. For $c \in [0, m]$, define

$$s_k(c, \mathbf{D}) = -\frac{|Q_c - T_k|}{2\Delta_T}.$$

We bound the sensitivity of the score function in Lemma 3.2.

Lemma 3.2. (Sensitivity of the score function). *For all $c \in [0, m]$, the values Q_c (Definition 3.1) have sensitivity at most Δ_T , and the score $s_k(c, \mathbf{D})$ has sensitivity at most 1.*

Theorem 3.3 (Report Noisy Max [11, 24, 10]). *Let $\varepsilon > 0$. Given scores $s_i \in \mathbb{R}$, $i \in [B]$ evaluated from a score function of sensitivity at most 1, the algorithm ReportNoisyMax samples $Z_1, \dots, Z_B \sim \text{Exp}(2/\varepsilon)$ and returns $\hat{i} = \arg \max_{i \in [B]} (s_i + Z_i)$. This algorithm is ε -DP*

and for $\delta \in (0, 1)$, with probability at least $1 - \delta$, it holds $s_i \geq s^* - 2 \log(B/\delta)\varepsilon^{-1}$, where $s^* = \max_{i \in [B]} s_i$.

Statistic T and its Sensitivity. The statistic T that we use to obtain our private CRT test is defined as the numerator of the GCM statistic. The residuals of \mathbf{Y} with respect to \mathbf{Z} are calculated by fitting a KRR model of \mathbf{Y} to \mathbf{Z} . Denote such residuals $r_{Y,i}$, for $i \in [n]$. The residuals of \mathbf{X} with respect to \mathbf{Z} are exact, since we have access to the distribution $X \mid Z$. Denote such residuals $r_{X,i}$ for $i \in [n]$. The residual products are calculated as $R_i = r_{X,i}r_{Y,i}$ for $i \in [n]$.

Definition 3.4 (Statistic T for the private CRT). Given a dataset $(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$ of n points, let (R_1, \dots, R_n) be the vector of residual products of the exact residuals of \mathbf{X} with respect to \mathbf{Z} and the residuals of fitting a kernel ridge regression model of \mathbf{Y} to \mathbf{Z} . Define $T(\mathbf{X}, \mathbf{Y}, \mathbf{Z}) = \sum_{i=1}^n R_i$.

We obtain a bound of $O_\lambda(1)$ on the ℓ_1 -sensitivity of the statistic T by bounding the sensitivity of the R_i . To bound the sensitivity of the $r_{Y,i}$ we assume that the domain of the variable Y is bounded and use the result of Theorem 1.2. We assume a known bound on the magnitude of the residuals $r_{X,i}$, motivated by the fact that we have access to the distribution $X \mid Z$. This differs from the assumptions for our PrivGCM test, where we assumed bounds on both X and Y . Assuming a bound on the residuals $r_{X,i}$ gives a tighter sensitivity bound for T . Refer to Lemma C.1 for a formal statement.

Accuracy of the Private CRT. We define the accuracy of Algorithm 2 in terms of the difference between the private p-value it outputs and its non-private counterpart.

Definition 3.5 ((γ, δ) -accuracy). Let $G_\gamma = |\{i \in [m] \mid |T_i - T_0| \leq \gamma\}|$. Let c^* be the rank of T_0 given statistics $T_i, i \in [0, m]$, and $p^* = (1 + c^*)/(m + 1)$ be the non-private p-value. We say Algorithm 2 is (γ, δ) -accurate if with probability at least $1 - \delta$ it holds $|\hat{p} - p^*| \leq \frac{G_\gamma}{m+1}$.

Define PrivCRT as Algorithm 2 with T from Definition 3.4.

Theorem 3.6. *PrivCRT is ε -DP and (γ, δ) -accurate for $\gamma = \frac{4\Delta_T}{\varepsilon} \log\left(\frac{m}{\delta}\right)$.*

Remark on the Accuracy. While the guarantees from the above theorem does not translate into a meaningful type-I error control bound, empirically, we observe that, under the null, the p-values output by PrivCRT are uniformly distributed (Fig. 9, Appendix D), and thus the test provides type-I error control. Under the alternate, T_0 is much larger (or smaller) than all the other values $T_i, i \geq 1$ and thus G_γ is small. However, the power of PrivCRT can be affected if we increase m , as this can increase the value of G_γ (see Fig. 8, Appendix D). An interesting open question is whether the dependence on m in the accuracy of a private CRT is avoidable. For now, we recommend using $m = O(1/\alpha)$ for rejection level α .

4 EMPIRICAL EVALUATION

We evaluate our algorithms on a real-world dataset and synthetic data. We start with the latter as it has the advantage that we know the ground-truth of whether $X \perp\!\!\!\perp Y \mid Z$.

Setup. To generate synthetic data, we use a setup similar to that of Shah and Peters [30] proposed for evaluating the performance of GCM. Fix an RKHS \mathcal{H} that corresponds to a Gaussian kernel. The function $f_s(z) = \exp(-s^2/2) \sin(sz)$ satisfies $f_s \in \mathcal{H}$. $Z = (Z_1, \dots, Z_d)$ is a d -dimensional variable, where $d \in \{1, 5\}$. The distribution of (X, Y, Z) is as follows:

$$\begin{aligned} Z_1, \dots, Z_d &\sim N_Z, X = f_s(Z_1) + N_X, \\ Y &= -f_s(Z_1) + N_Y + \beta \cdot N_X, \end{aligned}$$

where $N_Z \sim \mathcal{N}(0, 4)$, $N_X \sim \mathcal{N}(0, 1)$, $N_Y \sim \mathcal{N}(0, 1)$, and $\beta \geq 0$ is a constant controlling the strength of dependence between X and Y . If $\beta = 0$, then $X \perp\!\!\!\perp Y \mid Z$, but not otherwise. For experiments with PrivGCM, the dataset $(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$ consists of n points sampled as above. For experiments with PrivCRT, we additionally sample m copies $\mathbf{X}^{(j)}$, $j \in [m]$, by fixing \mathbf{Z} . We study how varying β , s , and n affects the rejection rate of our tests (averaged over 500 resampled datasets). Shaded error bars represent 95% confidence intervals. We set type-I error level $\alpha = 0.05$.

We rescale X and Y so that all datapoints x_i and y_i satisfy $|x_i| \leq 1, |y_i| \leq 1$ (recall that we assume known bounds for the data; for this simulation, standard Gaussian concentration implies $\max_{i \leq n} X_i \leq \sqrt{2 \log n}$ with very high probability, so choosing $\sqrt{C \log n}$ suffices here for a sufficiently large constant C). We then fit a KRR model with a Gaussian kernel of \mathbf{X} to \mathbf{Z} and \mathbf{Y} to \mathbf{Z} . The best model is chosen via 5-fold cross-validation and grid search over the regularization parameter λ and the parameter of the Gaussian kernel.⁵ The choice of λ requires balancing the performance of the fitting step of the algorithm with the magnitude of noise added (see Lemma B.6), and thus some lower bound on λ is needed. We enforce $\lambda \geq 10$ and find that this does not hurt performance of the fitting step even for large n . See Fig. 6 (Appendix D) for an example.

Comparison to the Private Kendall CI test [37] and Test-of-Tests [18]. We start with a comparison to two other private CI tests in the literature. The first is the private Kendall’s CI test, proposed by Wang et al. [37] for categorical variables. The second, which we call PrivToT, is obtained from the Test-of-Tests framework of Kazan et al. [18] and uses the non-private GCM as a black-box. See Appendix D for details on the implementations of these two tests.

⁵Hyper-parameters are optimized non-privately, as is common in the literature on privately training ML models. Our algorithm can be combined in a black box fashion with methods for performing private hyper-parameter search such as [23].

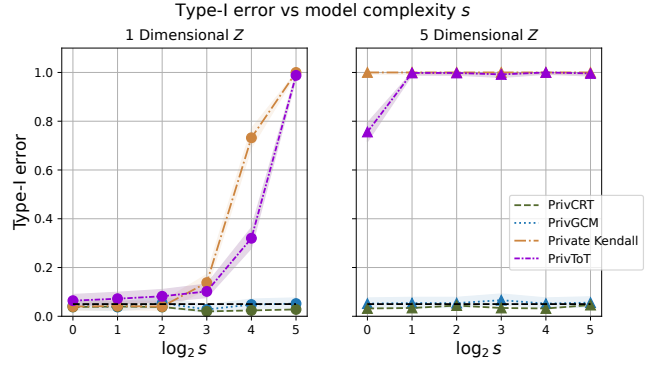


Figure 1: Type-I error control of PrivToT, private Kendall, PrivGCM, and PrivCRT (under the null): the first two fail to control Type-I error.

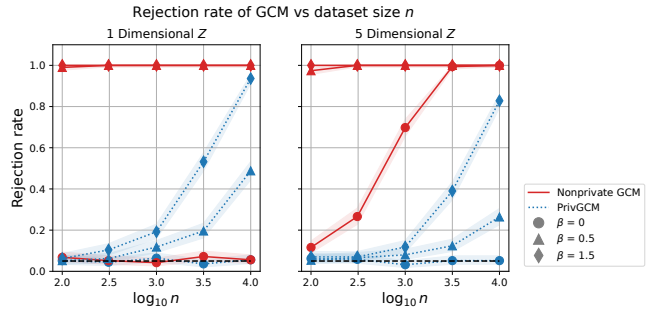


Figure 2: Comparison of the type-I error and power of private and nonprivate GCM tests as the dataset size n increases. Again, at $d = 5$ with $\beta = 0$, the (nonprivate) GCM fails to provide type-I error control even at large n (in fact, its type-I error gets worse with n).

In Fig. 1, we compare the performance of these two tests with our private tests under the null hypothesis. We vary s , the model complexity from 1 to 32, and use a sample size $n = 10^4$ and privacy parameter $\epsilon = 2$. The larger s , the harder it is to learn the function f_s . As the model complexity increases, the private Kendall test and PrivToT cannot control type-I error, even with a large sample size ($n = 10^4$). They perform even worse when Z is 5-dimensional. On the other hand, both PrivGCM and PrivCRT have consistent type-I error control across model complexity and dimensionality of Z . This experiment motivates the need for tests with rigorous theoretical type-I error guarantees, as we derive.⁶

Next, we compare our private CI tests with their non-private counterparts. We fix $s = 2$ for f_s .

Performance of PrivGCM. In Fig. 2, we examine the performance of PrivGCM on synthetic data. We vary the dataset size n from 10^2 to 10^4 , and plot the rejection rate of PrivGCM and GCM for $\beta = 0$ and $\beta > 0$. When $d = 5$, the privacy noise helps PrivGCM provide the critical type-I error control at $\beta = 0$, which non-private GCM fails at. At $d = 5$, the KRR model fails to fit the data (it returns

⁶Note that for tests without the desired type-I error control, statements about power are vacuous.

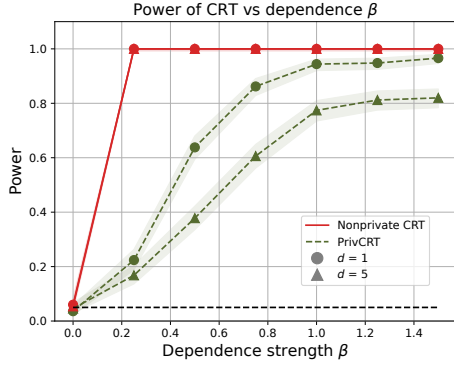


Figure 3: Comparing power of private and nonprivate CRT tests as we increase dependence β .

a predicted function that is nearly-zero). In this case, for $\beta = 0$, the GCM statistic converges to a Gaussian of standard deviation 1, but whose mean is removed from zero. The larger n , the further the mean of the Gaussian is from zero, and the worse the type-I error. The noise added for privacy brings the mean closer to zero since the standard deviation of the noisy residuals, σ'_P , is much larger than σ_P (see (4)). For $\beta > 0$, PrivGCM needs a higher dataset size to achieve the same power as GCM, concordant with our discussion following Theorem 2.3. As $n \rightarrow \infty$, the power of PrivGCM is expected to approach 1.

Performance of PrivCRT. We study the performance of PrivCRT in Figs. 3-4. PrivCRT achieves better power than PrivGCM for our setup, so we use a smaller privacy parameter of $\varepsilon = 2$ and set $m = 19$ (an extreme, but valid choice). In Fig. 3, we vary β , the dependence strength between X and Y , from 0 to 1.5, using $n = 10^3$. Both non-private CRT and PrivCRT provide type-I error control. Also, the power of both PrivCRT and (non-private) CRT converges to 1, with a faster convergence for the non-private test. In Fig. 4, we vary the dataset size n and $\beta \in \{0, 0.5, 1.5\}$.

In Fig. 7 (Appendix D), we compare the performance of PrivCRT with PrivGCM for different privacy parameters $\varepsilon \in [2^{-3}, 2^3]$. Both tests control type-I error, but PrivCRT achieves better power than PrivGCM for all privacy parameters ε . Therefore, PrivCRT appears preferable to PrivGCM when dataholders have access to the distribution $X | Z$. This result is consistent with the non-private case where the CRT has better power as it does not have to learn $\mathbb{E}[X | Z]$.

Real Data Experiments. We now demonstrate the performance of our private tests on real-world data. For these experiments, we use the ‘‘Concrete Compressive Strength’’ dataset that consists of 1030 datapoints and 9 continuous features [40]. We set Y to be the outcome feature: the concrete compressive strength. The choice of the X variable is varied, and Z consists of the 7 remaining features. Y is assumed to be a complex non-linear function of the other 8

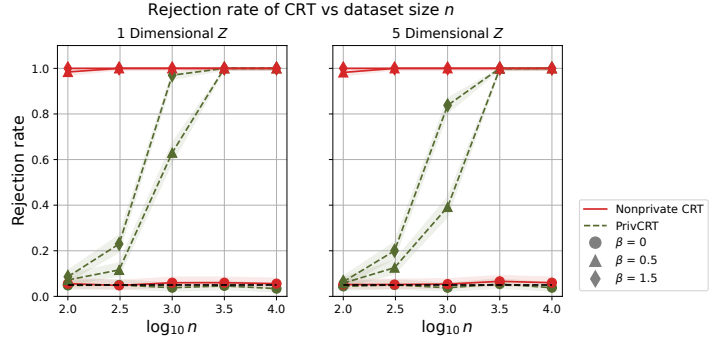


Figure 4: Comparison of the type-I error and power of private and nonprivate CRT tests as we increase the dataset size n .

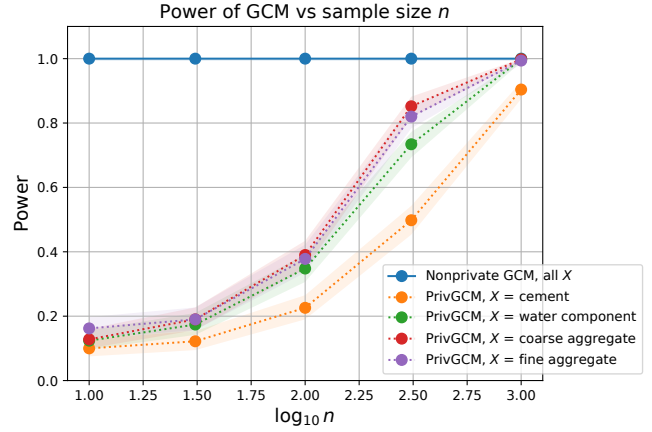


Figure 5: Power of the non-private GCM and PrivGCM on the ‘‘Concrete Compressive Strength’’ dataset. The power of PrivGCM tends to 1 with increasing sample size.

features, thus we expect the CI tests to reject. We focus on the private GCM because among our two algorithms, as it is the more generally applicable one, i.e., it does not require access to the conditional distribution. In Fig. 5, we evaluate the power of the PrivGCM. As desired, the power of PrivGCM tends to 1 as n increases approaching the power of the the non-private case.

5 CONCLUDING REMARKS

This work studies the fundamental statistical task of conditional independence testing under privacy constraints. We design the first DP conditional independence tests that support the general case of continuous variables and have strong theoretical guarantees. Our experiments support our theoretical results and additionally demonstrate that our private tests have more robust type-I error control than their non-private counterparts.

We envision two straightforward generalizations of our pri-

vate GCM test. First, our test can be generalized to handle multivariate X and Y , following Shah and Peters [30], who obtain the test statistic from the residual products of fitting each variable in X and each variable in Y to Z . A natural extension would be to compute the same statistic on our noisy residual products. Secondly, following Scheidegger et al. [29], a private version of the weighted GCM would allow the test to achieve power against a wider class of alternatives than the unweighted version. Finally, constructing private versions of other model- X tests, such as the Conditional Permutation Test [3], could be an interesting direction.

References

- [1] Jordan Awan and Aleksandra B. Slavkovic. Differentially private uniformly most powerful tests for binomial data. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 4212–4222, 2018.
- [2] Andrés F. Barrientos, Jerome P. Reiter, Ashwin Machanavajjhala, and Yan Chen. Differentially private significance tests for regression coefficients. *Journal of Computational and Graphical Statistics*, 28:440–453, 2017.
- [3] Thomas B. Berrett, Yi Wang, Rina Foygel Barber, and Richard J. Samworth. The conditional permutation test for independence while controlling for confounders. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 82, 2019.
- [4] Hai Brenner and Kobbi Nissim. Impossibility of differentially private universally optimal mechanisms. *SIAM Journal on Computing (SICOMP)*, 43(5):1513–1540, 2014.
- [5] Zachary Campbell, Andrew Bray, Anna M. Ritz, and Adam Groce. Differentially private ANOVA testing. In *International Conference on Data Intelligence and Security*, pages 281–285, 2018.
- [6] Emmanuel J. Candès, Yingying Fan, Lucas Janson, and Jinchi Lv. Panning for gold: ‘model- X ’ knock-offs for high dimensional controlled variable selection. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 80, 2016.
- [7] Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21:C1–C68, 2018.
- [8] Simon Couch, Zeki Kazan, Kaiyan Shi, Andrew Bray, and Adam Groce. Differentially private nonparametric hypothesis testing. In *Proceedings of the ACM Conference on Computer and Communications Security, CCS*, pages 737–751, 2019.
- [9] A Philip Dawid. Conditional independence in statistical theory. *Journal of the Royal Statistical Society: Series B (Methodological)*, 41(1):1–15, 1979.
- [10] Zeyu Ding, Daniel Kifer, Sayed M. Saghaian N. E., Thomas Steinke, Yuxin Wang, Yingtai Xiao, and Danfeng Zhang. The Permute-and-Flip mechanism is identical to Report-Noise-Max with exponential noise. *arXiv 2105.07260*, 2021.
- [11] Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Found. Trends Theor. Comput. Sci.*, 9(3-4):211–407, 2014.
- [12] Cynthia Dwork, Moni Naor, Omer Reingold, Guy N. Rothblum, and Salil P. Vadhan. On the complexity of differentially private data release: efficient algorithms and hardness results. In *Proceedings, ACM Symposium on Theory of Computing (STOC)*, pages 381–390, 2009.
- [13] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam D. Smith. Calibrating noise to sensitivity in private data analysis. *Journal of Privacy and Confidentiality*, 7(3):17–51, 2016.
- [14] Stephen E. Fienberg, Alessandro Rinaldo, and Xiaolin Yang. Differential privacy and the risk-utility tradeoff for multi-dimensional contingency tables. In *Privacy in Statistical Databases*, volume 6344 of *Lecture Notes in Computer Science*, pages 187–199. Springer, 2010.
- [15] Kenji Fukumizu, Arthur Gretton, Xiaohai Sun, and Bernhard Schölkopf. Kernel measures of conditional dependence. *Advances in Neural Information Processing Systems (NeurIPS)*, 20, 2007.
- [16] Marco Gaboardi, Hyun-Woo Lim, Ryan M. Rogers, and Salil P. Vadhan. Differentially private Chi-Squared hypothesis testing: Goodness of fit and independence testing. In *Proceedings, International Conference on Machine Learning (ICML)*, volume 48, 2016.
- [17] Aaron Johnson and Vitaly Shmatikov. Privacy-preserving data exploration in genome-wide association studies. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1079–1087, 2013.
- [18] Zeki Kazan, Kaiyan Shi, Adam Groce, and Andrew P. Bray. The test of tests: A framework for differentially private hypothesis testing. In *Proceedings, International Conference on Machine Learning (ICML)*, volume 202, pages 16131–16151, 2023.
- [19] Ilmun Kim and Antonin Schrab. Differentially private permutation tests: Applications to kernel methods. *CoRR*, abs/2310.19043, 2023. doi: 10.48550/ARXIV.2310.19043. URL <https://doi.org/10.48550/arXiv.2310.19043>.
- [20] Daphne Koller and Nir Friedman. *Probabilistic graphical models: principles and techniques*. MIT Press, 2009.
- [21] Matt J. Kusner, Yu Sun, Karthik Sridharan, and Kilian Q. Weinberger. Private causal inference. In *Pro-*

- ceedings, *International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 51, pages 1308–1317, 2016.
- [22] Tengyuan Liang and Alexander Rakhlin. Just interpolate: Kernel “ridgeless” regression can generalize. *The Annals of Statistics*, 48(3):1329 – 1347, 2020.
- [23] Jingcheng Liu and Kunal Talwar. Private selection from private candidates. In *Proceedings, ACM Symposium on Theory of Computing (STOC)*, pages 298–309. ACM, 2019.
- [24] Ryan McKenna and Daniel Sheldon. Permute-and-flip: A new mechanism for differentially private selection. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [25] Kobbi Nissim, Sofya Raskhodnikova, and Adam D. Smith. Smooth sensitivity and sampling in private data analysis. In *Proceedings, ACM Symposium on Theory of Computing (STOC)*, pages 75–84. ACM, 2007.
- [26] Judea Pearl. Models, reasoning and inference. *Cambridge, UK: Cambridge University Press*, 19(2), 2000.
- [27] Víctor Peña and Andrés F. Barrientos. Differentially private hypothesis testing with the subsampled and aggregated randomized response mechanism. *Statistica Sinica*, 2022.
- [28] Ryan Rogers and Daniel Kifer. A new class of private chi-square hypothesis tests. In *Proceedings, International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 54, pages 991–1000, 2017.
- [29] Cyrill Scheidegger, Julia Hörrmann, and Peter Bühlmann. The weighted generalised covariance measure. *Journal of Machine Learning Research*, 23(273): 1–68, 2022.
- [30] Rajen D. Shah and Jonas Peters. The hardness of conditional independence testing and the generalised covariance measure. *The Annals of Statistics*, 48(3), 2020.
- [31] Adam Smith. Privacy-preserving statistical estimation with optimal convergence rates. In *Proceedings, ACM Symposium on Theory of Computing (STOC)*, pages 813–822, 2011.
- [32] Eric V Strobl, Kun Zhang, and Shyam Visweswaran. Approximate kernel-based conditional independence tests for fast non-parametric causal discovery. *Journal of Causal Inference*, 7(1), 2019.
- [33] Marika Swanberg, Ira Globus-Harris, Iris Griffith, Anna M. Ritz, Adam Groce, and Andrew Bray. Improved differentially private analysis of variance. *Proc. Priv. Enhancing Technol.*, 2019(3):310–330, 2019.
- [34] Caroline Uhler, Aleksandra B. Slavkovic, and Stephen E. Fienberg. Privacy-preserving data sharing for genome-wide association studies. *J. Priv. Confidentiality*, 5(1), 2013.
- [35] Praneeth Vepakomma, Mohammad Mohammadi Amiri, Clément L. Canonne, Ramesh Raskar, and Alex Pentland. Private independence testing across two parties. *CoRR*, abs/2207.03652, 2022.
- [36] Duy Vu and Aleksandra B. Slavkovic. Differential privacy for clinical trial data: Preliminary evaluations. In *ICDM Workshops*, pages 138–143, 2009.
- [37] Lun Wang, Qi Pang, and Dawn Song. Towards practical differentially private causal graph discovery. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [38] Yue Wang, Jaewoo Lee, and Daniel Kifer. Revisiting differentially private hypothesis tests for categorical data. *arXiv 1511.03376*, 2015.
- [39] Yue Wang, Daniel Kifer, Jaewoo Lee, and Vishesh Karwa. Statistical approximating distributions under differential privacy. *J. Priv. Confidentiality*, 8(1), 2018.
- [40] I-Cheng Yeh. Concrete Compressive Strength. UCI Machine Learning Repository, 2007. DOI: <https://doi.org/10.24432/C5PK67>.
- [41] Fei Yu, Stephen E. Fienberg, Aleksandra B. Slavkovic, and Caroline Uhler. Scalable privacy-preserving data sharing methodology for genome-wide association studies. *J. Biomed. Informatics*, 50:133–141, 2014.
- [42] Kun Zhang, Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. Kernel-based conditional independence test and application in causal discovery. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 804–813, 2011.

Checklist

1. For all models and algorithms presented, check if you include:
 - (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model. [Yes]
 - (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm. [Yes]
 - (c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries. [Yes]
2. For any theoretical claim, check if you include:
 - (a) Statements of the full set of assumptions of all theoretical results. [Yes]
 - (b) Complete proofs of all theoretical results. [Yes]
 - (c) Clear explanations of any assumptions. [Yes]
3. For all figures and tables that present empirical results, check if you include:

- (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). [No]
 - (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). [Yes]
 - (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). [Yes]
 - (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
- (a) Citations of the creator If your work uses existing assets. [Yes]
 - (b) The license information of the assets, if applicable. [Not Applicable]
 - (c) New assets either in the supplemental material or as a URL, if applicable. [Not Applicable]
 - (d) Information about consent from data providers/curators. [Not Applicable]
 - (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. [Not Applicable]
5. If you used crowdsourcing or conducted research with human subjects, check if you include:
- (a) The full text of instructions given to participants and screenshots. [Not Applicable]
 - (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. [Not Applicable]
 - (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. [Not Applicable]

Supplementary Material: Differentially Private Conditional Independence Testing

A Missing Details from Section 1

A.1 Additional Related Work

Conditional Independence Testing. A popular category of CI tests are kernel-based tests, obtained by extending the Hilbert-Schmidt independence criterion to the conditional setting [15, 42, 32]. However, these tests only provide a weaker *pointwise* asymptotic validity guarantee. It is widely acknowledged that for a statistical test to be useful in practice, it needs to provide the stronger guarantees of either valid level at finite sample size or *uniformly* asymptotic level. Our private GCM test provides the latter guarantee.

One way of getting around the hardness result of Shah and Peters [30] is through the model-X assumption, where the conditional distribution of $X | Z$ is assumed to be accessible. Tests based on this assumption, such as CRT (conditional randomization test) [6] and CPT (conditional permutation test) [3], provide a general framework for conditional independence testing, where one can use their test statistic of choice and exactly (non-asymptotically) control the type-I error regardless of the data dimensionality.

A.2 Background on Hypothesis Testing

Let \mathcal{P} be the class of the joint distributions for the random variables (X, Y, Z) . We say $P \in \mathcal{P}$ is a null distribution if $X \perp\!\!\!\perp Y | Z$. The null-hypothesis, denoted \mathcal{P}_0 , is the class of null distributions,

Type-I error and validity. Consider a (potentially) randomized test φ_n that is run on n samples from a distribution $P \in \mathcal{P}$ and outputs a binary decision: 1 for rejecting the null hypothesis and 0 for accepting the null hypothesis. The quantity $\Pr_P[\varphi_n = 1]$, where P is a null distribution, refers to the type-I error of the test, i.e., the probability that it erroneously rejects the (true) null hypothesis. Given level $\alpha \in (0, 1)$ and the null hypothesis \mathcal{P}_0 , we say that the test has *valid level at sample size n* if the type-I error is bounded by α , i.e.,:

$$\sup_{P \in \mathcal{P}_0} \Pr_P[\varphi_n = 1] \leq \alpha.$$

The sequence $\{\varphi_n\}_{n=1}^\infty$ has

$$\begin{aligned} \text{uniformly asymptotic level} & \quad \text{if} \quad \lim_{n \rightarrow \infty} \sup_{P \in \mathcal{P}_0} \Pr_P[\varphi_n = 1] \leq \alpha, \\ \text{pointwise asymptotic level} & \quad \text{if} \quad \sup_{P \in \mathcal{P}_0} \lim_{n \rightarrow \infty} \Pr_P[\varphi_n = 1] \leq \alpha. \end{aligned}$$

Usually, we want at least uniformly asymptotic level to hold for a test. Otherwise, for any sample size n , there can be some null distribution that does not control type-I error at that sample size.

A hypothesis test is usually derived from a statistic T_n (such as the GCM statistic) calculated on n samples drawn i.i.d from the distribution P of (X, Y, Y) . Having obtained a value t for the statistic T_n , a two-sided p-value is calculated as:

$$\text{p-value} = \Pr_P[|T_n| \geq t].$$

The hypothesis test φ_n with desired validity level α is defined as:

$$\varphi_n = \mathbf{1}[\text{p-value} \leq \alpha]. \tag{5}$$

Therefore, to obtain a test with the desired validity we need to compute the p-value. The p-value is typically calculated using information about the distribution of T_n . We say that T_n converges uniformly over \mathcal{P}_0 to the standard Gaussian distribution if:

$$\lim_{n \rightarrow \infty} \sup_{P \in \mathcal{P}_0} \sup_{t \in \mathbb{R}} |\Pr_P[T_n \leq t] - \Phi(t)| = 0,$$

where Φ is the CDF of the standard Gaussian. For the GCM statistic, we are given that (under mild assumptions) T_n converges uniformly over the null hypothesis \mathcal{P}_0 to a standard Gaussian distribution [30]. Therefore, by setting p-value = $2(1 - \Phi(|t|))$ and defining the hypothesis test φ_n as in (5), we obtain that φ_n has uniformly asymptotic level α .

Power. Once we have a test with uniformly asymptotic level, we would also like the test to correctly accept the alternate hypothesis, when this hypothesis holds. Let \mathcal{Q} be the set of alternate distributions (for which $X \not\perp\!\!\!\perp Y \mid Z$). The power of a test is the probability that it correctly rejects the null hypothesis, given that the alternate hypothesis holds. A sequence $\{\varphi_n\}_{n=1}^\infty$ of tests has

$$\text{uniformly asymptotic power} \quad \text{if} \quad \lim_{n \rightarrow \infty} \inf_{P \in \mathcal{Q}} \Pr_P[\varphi_n = 1] = 1.$$

A.3 Background on Differential Privacy

Definition A.1 (ℓ_1 -sensitivity). For a function $f: \mathcal{D} \rightarrow \mathbb{R}^d$, its ℓ_1 -sensitivity Δ_f is defined as

$$\Delta_f = \max_{\mathbf{D}, \mathbf{D}' \text{ neighbors}} \|f(\mathbf{D}) - f(\mathbf{D}')\|_1.$$

Lemma A.2 (Laplace Mechanism [13]). Let $\varepsilon > 0$ and $f: \mathcal{D} \rightarrow \mathbb{R}^d$ be a function with ℓ_1 -sensitivity Δ_f . Let $W \sim \text{Lap}(0, \Delta_f/\varepsilon)$ be a noise vector from the Laplace distribution with scale parameter Δ_f/ε . The Laplace Mechanism that, on input \mathbf{D} and ε , outputs $f(\mathbf{D}) + W$ is ε -DP.

Lemma A.3 (Post-Processing [13]). If the algorithm \mathcal{A} is ε -differentially private, and g is any randomized function, then the algorithm $g(\mathcal{A}(x))$ is ε -differentially private.

B Proofs of Section 2

B.1 Proof of Theorem 2.2

In this section, we state and prove a longer version of Theorem 2.2. Item 1 of Theorem B.1 gives the pointwise asymptotic level guarantee of the private GCM, whereas Item 2 shows the more desirable uniform asymptotic level guarantee under a slightly stronger condition. Item 2 corresponds exactly to Theorem 2.2.⁷ See Appendix A.2 for definitions of pointwise asymptotic level and uniformly asymptotic level.

From a privacy perspective, the proof of Item 2 is more involved. While Shah and Peters [30] consider the asymptotic behavior of variables $\chi_i \xi_i$ (the product of the true residuals), we instead study the behavior of $\chi_i \xi_i W_i$ (the product of the true residuals with the noise random variables). Similarly, while they study the product of error terms of the fitting method, $(f(z_i) - \hat{f}(z_i))(g(z_i) - \hat{g}(z_i))$, we instead need to consider $(f(z_i) - \hat{f}(z_i))(g(z_i) - \hat{g}(z_i))W_i$, the product of error terms with the noise variables. A key step in the proof is to show that the noise variables grow at a slower rate than the rate of decay of the error terms, with increasing sample size n .

Let \mathcal{E}_0 be the set of distributions for (X, Y, Z) that are absolutely continuous with respect to the Lebesgue measure. The null hypothesis, $\mathcal{P}_0 \subset \mathcal{E}_0$, is the subset of distributions for which $X \perp\!\!\!\perp Y \mid Z$. Given $P \in \mathcal{P}$, let P' be the joint distribution of variables (X, Y, Z, W) where $W \sim \text{Lap}(\Delta/\varepsilon)$ is independent of (X, Y, Z) . For a set of distributions \mathcal{P} , let \mathcal{P}' denote the set of distributions P' for all $P \in \mathcal{P}$. Denote by Φ the CDF of the standard normal distribution.

Theorem B.1. (Type-I Error Control of Private GCM) Let a and b be known bounds on the domains of X and Y , respectively. Let $\mathcal{P}_0 \subset \mathcal{E}_0$ be the set of null distributions defined above. Given a dataset $\mathbf{D} = (\mathbf{X}, \mathbf{Y}, \mathbf{Z})$, let $(\hat{\mathbf{X}}, \hat{\mathbf{Y}}, \mathbf{Z})$ be the rescaled dataset obtained by setting $\hat{\mathbf{X}} = \mathbf{X}/a$ and $\hat{\mathbf{Y}} = \mathbf{Y}/b$. Consider $R_i, i \in [n]$, as defined in (1), for the rescaled dataset $(\hat{\mathbf{X}}, \hat{\mathbf{Y}}, \mathbf{Z})$. Let $W_i \sim \text{Lap}(0, \Delta/\varepsilon)$ for $i \in [n]$, where $\Delta, \varepsilon > 0$ are constants. Then $T^{(n)} = T(R_1 + W_1, \dots, R_n + W_n)$, defined in Algorithm 1, satisfies:

⁷The assumptions in Item 2, Theorem B.1 are identical to those of Definition 2.1. The assumptions in Item 1 of Theorem B.1 are slightly weaker than those of Definition 2.1.

1. For $P \in \mathcal{P}_0$ such that $A_f A_g = o_P(n^{-1})$, $B_f = o_P(1)$, $B_g = o_P(1)$, and $\mathbb{E}[\chi_P^2 \xi_P^2] < \infty$, then

$$\lim_{n \rightarrow \infty} \sup_{t \in \mathbb{R}} |\Pr_{P'}[T^{(n)} \leq t] - \Phi(t)| = 0.$$

2. Let $\mathcal{P} \subset \mathcal{P}_0$ be a set of distributions such that $A_f A_g = o_P(n^{-1})$, $B_f = o_P(1)$ and $B_g = o_P(1)$. If in addition $\sup_{P \in \mathcal{P}} \mathbb{E}[|\chi_P \xi_P|^{2+\eta}] \leq c$, for some constants $c, \eta > 0$, then

$$\lim_{n \rightarrow \infty} \sup_{P' \in \mathcal{P}'} \sup_{t \in \mathbb{R}} |\Pr_{P'}[T^{(n)} \leq t] - \Phi(t)| = 0. \quad (6)$$

Proof. Let $\sigma_P = \sqrt{\text{Var}(\chi_P \xi_P)}$, $\sigma_{\text{priv}} = \Delta/\varepsilon$, and $\sigma_{\text{joint}} = \sqrt{\frac{\sigma_P^2}{a^2 b^2} + 2\sigma_{\text{priv}}^2}$. Denote by τ_N the numerator of $T^{(n)}$ and by τ_D the denominator. We sometimes omit P from the notation for ease of presentation.

Item 1: We first show Item 1. Specifically, we show that $\tau_N \rightarrow \mathcal{N}(0, \sigma_{\text{joint}}^2)$ and $\tau_D \rightarrow \sigma_{\text{joint}}$, which, by Slutsky's lemma (Lemma B.8), would imply that $T^{(n)} \rightarrow \mathcal{N}(0, 1)$. For $i \in [n]$, let $\chi_i = x_i - f(z_i)$ and $\xi_i = y_i - g(z_i)$. Note that

$$\begin{aligned} \tau_N &= \frac{1}{\sqrt{n}} \sum_{i=1}^n (R_i + W_i) \\ &= \frac{1}{ab\sqrt{n}} \sum_{i=1}^n (f(z_i) - \hat{f}(z_i))(g(z_i) - \hat{g}(z_i)) + \frac{1}{ab\sqrt{n}} \sum_{i=1}^n (f(z_i) - \hat{f}(z_i))\xi_i \\ &\quad + \frac{1}{ab\sqrt{n}} \sum_{i=1}^n (g(z_i) - \hat{g}(z_i))\chi_i + \frac{1}{\sqrt{n}} \sum_{i=1}^n \left(\frac{\chi_i \xi_i}{ab} + W_i \right) \end{aligned} \quad (7)$$

We use the following claim obtained from the proof of Theorem 6 of Shah and Peters [30].

Claim B.2 ([30]). *Under the assumptions listed in Item 1 of Theorem B.1, the following hold*

1. $\frac{1}{\sqrt{n}} \sum_{i=1}^n (f(z_i) - \hat{f}(z_i))(g(z_i) - \hat{g}(z_i)) \xrightarrow{P} 0$.
2. $\frac{1}{\sqrt{n}} \sum_{i=1}^n (f(z_i) - \hat{f}(z_i))\xi_i \xrightarrow{P} 0$.
3. $\frac{1}{\sqrt{n}} \sum_{i=1}^n (g(z_i) - \hat{g}(z_i))\chi_i \xrightarrow{P} 0$.
4. $\frac{1}{n} \sum_{i=1}^n R_i \xrightarrow{P} 0$.
5. $\frac{1}{n} \sum_{i=1}^n R_i^2 \xrightarrow{P} \frac{\sigma_P^2}{a^2 b^2}$.

Additionally, under the assumptions of Item 2, all convergence statements above are uniform over \mathcal{P} .

By Claim B.2, Items 1-3, the first three terms of the sum in (7) are $o_P(1)$. We show that the last term of the sum converges to $\mathcal{N}(0, \sigma_{\text{joint}}^2)$. Note that $\mathbb{E}[\frac{\chi_i \xi_i}{ab} + W_i] = 0$ and $\text{Var}(\frac{\chi_i \xi_i}{ab} + W_i) = \text{Var}(\chi_i \xi_i) + \text{Var}(W_i) = \sigma_{\text{joint}}^2$. Since σ_{priv} is a constant and $\sigma_P < \infty$, we can apply the Central Limit Theorem to obtain the desired convergence to a Gaussian. By Slutsky's lemma, we obtain that $\tau_N \rightarrow \mathcal{N}(0, \sigma_{\text{joint}}^2)$.

We now consider τ_D . Since $\mathbb{E}[W_i] = 0$ for all $i \in [n]$, by the Weak Law of Large Numbers it holds that $\frac{1}{n} \sum_{i=1}^n W_i \rightarrow 0$. From Item 4 of Claim B.2, we obtain $\frac{1}{n} \sum_{i=1}^n (R_i + W_i) \xrightarrow{P} 0$. It remains to show that $\frac{1}{n} \sum_{i=1}^n (R_i + W_i)^2 \xrightarrow{P} \sigma_P$. We have

$$\frac{1}{n} \sum_{i=1}^n (R_i + W_i)^2 = \frac{1}{n} \sum_{i=1}^n R_i^2 + \frac{2}{n} \sum_{i=1}^n R_i W_i + \frac{1}{n} \sum_{i=1}^n W_i^2.$$

Since R_i and W_i are independent, we have $\mathbb{E}[R_i W_i] = \mathbb{E}[R_i] \mathbb{E}[W_i] = 0$. Another application of the Weak Law of Large Numbers yields $\frac{2}{n} \sum_{i=1}^n R_i W_i \xrightarrow{P} 0$. Finally, let $B^{(n)} = \frac{1}{n} \sum_{i=1}^n W_i^2$. We have $\mathbb{E}[W_i^2] = \text{Var}[W_i] = 2\sigma_{\text{priv}}^2$. The Weak

Law of Large numbers gives that $B^{(n)} \rightarrow 2\sigma_{\text{priv}}^2$. By Item 5 of Claim B.2 and Slutsky's lemma, we obtain that $\tau_D \rightarrow \sigma_{\text{joint}}$, as desired. This concludes the proof of Item 1.

Item 2: Turning to Item 2, more care is needed to replace statements of convergence with statements of uniform convergence due to the presence of privacy noise. We first show that τ_N converges to $\mathcal{N}(0, \sigma_{\text{joint}}^2)$ uniformly. Consider again the sum in (7). By Claim B.2, Items 1-3, the first three terms of the sum in (7) are $o_{\mathcal{P}}(1)$, and as a consequence, $o_{\mathcal{P}'}(1)$. We show that

$$\frac{1}{\sqrt{n}} \sum_{i=1}^n \left(\frac{\chi_i \xi_i}{ab} + W_i \right) \rightarrow \mathcal{N}\left(0, \sigma_{\text{joint}}^2\right) \text{ uniformly over } \mathcal{P}'. \quad (8)$$

Then, applying Item 1 of Lemma B.11 we would obtain the desired convergence of τ_N . We prove (8) by showing that the random variables $\frac{\chi_i \xi_i}{ab} + W_i, i \in [n]$ satisfy the conditions of Lemma B.9. Recall that $\mathbb{E}[\frac{\chi_i \xi_i}{ab} + W_i] = 0$ and $\mathbb{E}[(\frac{\chi_i \xi_i}{ab} + W_i)^2] = \text{Var}(\frac{\chi_i \xi_i}{ab} + W_i) = \sigma_{\text{joint}}^2$, which is a constant. It remains to bound the $(2 + \eta)$ -absolute moment of $\frac{\chi_i \xi_i}{ab} + W_i$. By Claim B.13 and our assumption on $\chi_i \xi_i$, we have

$$\mathbb{E}\left[\left|\frac{\chi_i \xi_i}{ab} + W_i\right|^{2+\eta}\right] \leq 2^{1+\eta} \left(\mathbb{E}\left[\left|\frac{\chi_i \xi_i}{ab}\right|^{2+\eta}\right] + \mathbb{E}[|W_i|^{2+\eta}] \right) \leq 2^{(1+\eta)} \cdot \left(\frac{c}{(ab)^{2+\eta}} + c' \right)$$

where c' is the (constant) bound on $\mathbb{E}[|W_i|^{2+\eta}]$ given by Claim B.14. Thus, all the conditions of Lemma B.9 are satisfied, and (8) holds. This concludes our argument on the convergence of τ_N .

We now show that τ_D converges uniformly over \mathcal{P}' to σ_{joint} . We first show that $n^{-1} \sum_{i=1}^n (R_i + W_i) = o_{\mathcal{P}'}(1)$. By Item 4 of Claim B.2, we have that $n^{-1} \sum_{i=1}^n R_i = o_{\mathcal{P}'}(1)$. We also showed that $\frac{1}{n} \sum_{i=1}^n W_i = o_{\mathcal{P}'}(1)$. Applying Lemma B.11, we obtain that $n^{-1} \sum_{i=1}^n (R_i + W_i) = o_{\mathcal{P}'}(1)$.

Next, we show that $n^{-1} \sum_{i=1}^n R_i W_i = o_{\mathcal{P}'}(1)$. Note that

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n R_i W_i &= \frac{1}{ab} \left(\frac{1}{n} \sum_{i=1}^n (f(z_i) - \hat{f}(z_i) + \chi_i)(g(z_i) - \hat{g}(z_i) + \xi_i) W_i \right) \\ &= \frac{1}{ab} \left(\frac{1}{n} \sum_{i=1}^n (f(z_i) - \hat{f}(z_i))(g(z_i) - \hat{g}(z_i)) W_i + \frac{1}{n} \sum_{i=1}^n (f(z_i) - \hat{f}(z_i)) \xi_i W_i \right. \\ &\quad \left. + \frac{1}{n} \sum_{i=1}^n (g(z_i) - \hat{g}(z_i)) \chi_{P,i} W_i + \frac{1}{n} \sum_{i=1}^n \chi_i \xi_i W_i \right) \end{aligned} \quad (9)$$

We show that each of the terms of the sum in (9) is $o_{\mathcal{P}'}(1)$, starting with $n^{-1} \sum_{i=1}^n \chi_i \xi_i W_i$. We do so by showing that the conditions of Lemma B.10 hold for the random variables $\chi_i \xi_i W_i$. More precisely, we need to show that for all $\mathcal{P}' \in \mathcal{P}'$ it holds that $\mathbb{E}_{\mathcal{P}'}[|\chi_i \xi_i W_i|^{1+\eta}] < c_2$ for some $c_2, \eta > 0$. By the assumptions in Item 2, we know $\mathbb{E}_{\mathcal{P}'}[|\chi_P \xi_P|^{1+\eta}] \leq \mathbb{E}_{\mathcal{P}'}[|\chi_P \xi_P|^{2+\eta}] \leq c$. By the independence of $\chi_i \xi_i$ and W_i we have

$$\mathbb{E}_{\mathcal{P}'}[|\chi_i \xi_i W_i|^{1+\eta}] = \mathbb{E}_{\mathcal{P}'}[|\chi_P \xi_P|^{1+\eta}] \mathbb{E}_{\mathcal{P}'}[|W_i|^{1+\eta}] \leq c \cdot c',$$

where c' is a bound on $\mathbb{E}_{\mathcal{P}'}[|W_i|^{1+\eta}]$ given by Claim B.14. Therefore, we have that the variables $\chi_i \xi_i W_i$ satisfy all properties listed in Lemma B.10, and thus

$$n^{-1} \sum_{i=1}^n \chi_i \xi_i W_i = o_{\mathcal{P}}(1). \quad (10)$$

Next, we show that

$$\frac{1}{n} \sum_{i=1}^n (f(z_i) - \hat{f}(z_i))(g(z_i) - \hat{g}(z_i)) W_i = o_{\mathcal{P}'}(1). \quad (11)$$

Denote the term in the LHS of (11) by b . We have

$$|b| \leq \max_{i \in [n]} |W_i| \cdot \left(\frac{1}{n} \sum_{i=1}^n (f(z_i) - \hat{f}(z_i))(g(z_i) - \hat{g}(z_i)) \right) \quad (12)$$

We show that $\max_{i \in [n]} |W_i| = O_{\mathcal{P}'}(\log n)$ in Claim B.12. By Item 1 of Claim B.2, the other term in the product in the RHS of (12) is $o_{\mathcal{P}}(n^{-1/2})$. As a result, $b = o_{\mathcal{P}'}(1)$. By a similar proof, and using Claim B.2, the other terms in the sum in (9) are equal to $o_{\mathcal{P}'}(1)$.

By Item 5 of Claim B.2, we have that $\frac{1}{n} \sum_{i=1}^n R_i^2$ converges uniformly over \mathcal{P}' to $\frac{\sigma_P^2}{a^2 b^2}$. In addition, we showed that $\frac{1}{n} \sum_{i \in [n]} W_i^2$ converges uniformly over \mathcal{P}' to $2\sigma_{\text{priv}}^2$. Applying Lemma B.11 we obtain that τ_D converges uniformly over \mathcal{P}' to σ_{joint} . One final application of Lemma B.11 yields that $T^{(n)}$ converges uniformly over \mathcal{P}' to $\mathcal{N}(0, 1)$. \square

B.2 Proof of Theorem 2.3

In this section, we state and prove a longer version of Theorem 2.3. Item 1 of Theorem B.3 gives the pointwise power guarantee, whereas Item 2 shows the uniform power guarantee under a slightly stronger condition. Item 2 corresponds exactly to Theorem 2.3.⁸

Following Shah and Peters [30], to facilitate the theoretical analysis of power, we separate the model fitting step from the calculation of the residuals. We calculate \hat{f} and \hat{g} on the first half of the dataset and calculate the residuals $R_i, i \in [n]$ on the second half. In practice, it is still advised to perform both steps on the full dataset.

Theorem B.3. (Power of Private GCM). *Consider the setup of Theorem 2.2. Let A_f, A_g, B_f, B_g be as defined in (3), with the difference that \hat{f} and \hat{g} are estimated on the first half of the dataset $(\hat{\mathbf{X}}, \hat{\mathbf{Y}}, \mathbf{Z})$, and $R_i, i \in [n/2, n]$ are calculated on the second half. Define the “signal” (ρ_P) and “noise” (σ_P) of the true residuals χ_P, ξ_P as:*

$$\rho_P = \mathbb{E}_P[\text{cov}(X, Y | Z)], \sigma_P = \sqrt{\text{Var}_P(\chi_P \xi_P)}.$$

1. If for $P \in \mathcal{E}_0$ we have $A_f A_g = o_P(n^{-1}), B_f = o_P(1), B_g = o_P(1)$ and $\sigma_P < \infty$, then

$$\lim_{n \rightarrow \infty} \sup_{t \in \mathbb{R}} \left| \Pr_{\mathcal{P}'} \left[T^{(n)} - \frac{\sqrt{n} \rho_P}{\sigma_P} \leq t \right] - \Phi(t) \right| = 0, \text{ where } \sigma'_P = \sqrt{\sigma_P^2 + \left(\frac{\sqrt{2} ab \Delta}{\varepsilon} \right)^2}. \quad (13)$$

2. Let $\mathcal{P} \subset \mathcal{E}_0$ such that $A_f A_g = o_{\mathcal{P}}(n^{-1}), B_f = o_{\mathcal{P}}(1)$ and $B_g = o_{\mathcal{P}}(1)$. If in addition $\sup_{P \in \mathcal{P}} \mathbb{E}[|\chi_P \xi_P|^{2+\eta}] \leq c$, for some constants $c, \eta > 0$, then (13) holds over \mathcal{P}' uniformly.

Proof. Note that $\mathbb{E}[\frac{\chi_P \xi_P}{ab}] = \frac{\rho}{ab}$ and $\sqrt{\text{Var}(\frac{\chi_P \xi_P}{ab})} = \frac{\sigma_P}{ab}$. The proof is similar to that of Theorem B.1, using Claim B.4 below from Shah and Peters [30]. \square

Claim B.4. *Under the assumptions listed in Theorem 2.3, as $n \rightarrow \infty$, the following hold.*

1. $\frac{1}{\sqrt{n}} \sum_{i=1}^n (f(z_i) - \hat{f}(z_i))(g(z_i) - \hat{g}(z_i)) \xrightarrow{P} 0$.
2. $\frac{1}{\sqrt{n}} \sum_{i=1}^n (f(z_i) - \hat{f}(z_i)) \xi_i \xrightarrow{P} 0$.
3. $\frac{1}{\sqrt{n}} \sum_{i=1}^n (g(z_i) - \hat{g}(z_i)) \chi_i \xrightarrow{P} 0$.
4. $\frac{1}{n} \sum_{i=1}^n R_i \xrightarrow{P} 0$.
5. $\frac{1}{n} \sum_{i=1}^n R_i^2 \xrightarrow{P} \frac{\sigma_P^2}{a^2 b^2}$.

Additionally, under the assumptions in Item 2 of Theorem 2.3, all convergence statements above are uniform over \mathcal{P} .

Corollary B.5. *Under the assumptions of Theorem B.3, Algorithm 1 has asymptotic power of 1 if $\rho_P \neq 0$.*

Next, we show that Theorem 2.3 implies that the private GCM has asymptotic power of 1. A similar claim and proof holds for uniformly asymptotic power.

⁸The assumptions in Item 2, Theorem B.3 are identical to those of Definition 2.1. The assumptions in Item 1 of Theorem B.3 are slightly weaker than those of Definition 2.1.

Proof. Note that Algorithm 1 has asymptotic power of 1 if, for all $M > 0$, it holds $\Pr_P[T^{(n)} > M] \rightarrow 1$ as $n \rightarrow \infty$. Given $M > 0$, note that

$$\Pr[T^{(n)} \leq M] = \Pr_{P'}\left[T^{(n)} - \frac{\sqrt{n}\rho_P}{\sigma'_P} \leq M - \frac{\sqrt{n}\rho_P}{\sigma'_P}\right] \rightarrow \Phi\left(M - \frac{\sqrt{n}\rho_P}{\sigma'_P}\right),$$

where the convergence statement follows from Theorem 2.3. Since $\rho_P \neq 0$ and σ'_P is a constant, then $M - \frac{\sqrt{n}\rho_P}{\sigma'_P} \rightarrow -\infty$ as $n \rightarrow \infty$. Therefore $\Phi\left(M - \frac{\sqrt{n}\rho_P}{\sigma'_P}\right) \rightarrow 0$, and as a result $\Pr[T^{(n)} \leq M] \rightarrow 0$, as desired. \square

B.3 Guarantees of PrivGCM

Lemma B.6. (Sensitivity of residual products). *Let \mathbf{R} be the vector of residual products, as defined in (1), of fitting a KRR model of \mathbf{X} to \mathbf{Z} and \mathbf{Y} to \mathbf{Z} with regularization parameter $\lambda > 0$. If $|x_i|, |y_i| \leq 1$ for all $i \in [n]$, then $\Delta_{\mathbf{R}} \leq C$ where $C = 4(1 + \frac{\sqrt{2}}{\sqrt{\lambda}})(1 + \frac{\sqrt{2}}{\sqrt{\lambda}} + \frac{4\sqrt{2}}{\lambda^{3/2}} + \frac{4}{\lambda})$.*

Proof. Consider two neighboring datasets $(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$ and $(\mathbf{X}', \mathbf{Y}', \mathbf{Z}')$. For $i \in [n]$, let $r_{X,i}, r_{Y,i}$ denote the residuals of fitting a kernel ridge regression model of \mathbf{X} to \mathbf{Z} and \mathbf{Y} to \mathbf{Z} , respectively. Suppose without loss of generality that $(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$ and $(\mathbf{X}', \mathbf{Y}', \mathbf{Z}')$ differ only in the last datapoint, i.e., $(x_i, y_i, z_i) = (x'_i, y'_i, z'_i)$ for $i \in [n-1]$. Then, by Theorem 1.2, for $i \in [n-1]$, we have

$$|r_{X,i} - r'_{X,i}| = |(x_i - \mathbf{w}^\top \phi(z_i)) - (x'_i - \mathbf{w}'^\top \phi(z'_i))| = |\mathbf{w}^\top \phi(z_i) - \mathbf{w}'^\top \phi(z_i)| \leq \frac{8\sqrt{2}}{\lambda^{3/2}n} + \frac{8}{\lambda n}.$$

For the last datapoint we have

$$\begin{aligned} |r_{X,n} - r'_{X,n}| &= |(x_n - \mathbf{w}^\top \phi(z_n)) - (x'_n - \mathbf{w}'^\top \phi(z'_n))| \\ &\leq |x_n - x'_n| + \|\mathbf{w}'\|_{\mathcal{H}} |\phi(z'_n) - \phi(z_n)| + |\mathbf{w}'^\top \phi(z_n) - \mathbf{w}^\top \phi(z_n)| \\ &\leq 2 + \frac{2\sqrt{2}}{\sqrt{\lambda}} + \frac{8\sqrt{2}}{\lambda^{3/2}n} + \frac{8}{\lambda n}. \end{aligned}$$

Lastly, note that for all $i \in [n]$, we have $|r_{X,i}| \leq |x_i| + \|\mathbf{w}\|_{\mathcal{H}} \|\phi(z_i)\|_{\mathcal{H}} \leq 1 + \frac{\sqrt{2}}{\sqrt{\lambda}}$. The same bound holds for $|r_{Y,i}|$. Let $c_1 = 2 + \frac{2\sqrt{2}}{\sqrt{\lambda}}$ and $c_2 = \frac{8\sqrt{2}}{\lambda^{3/2}} + \frac{8}{\lambda}$. This gives us that for all $i \in [n-1]$:

$$|R_i - R'_i| \leq |r_{X,i}| |r_{Y,i} - r'_{Y,i}| + |r'_{Y,i}| |r_{X,i} - r'_{X,i}| \leq 2\left(1 + \frac{\sqrt{2}}{\sqrt{\lambda}}\right) \left(\frac{8\sqrt{2}}{\lambda^{3/2}n} + \frac{8}{\lambda n}\right) = \frac{c_1 c_2}{n}.$$

For $i = n$ we have

$$|R_n - R'_n| \leq 2\left(1 + \frac{\sqrt{2}}{\sqrt{\lambda}}\right) \left(2 + \frac{2\sqrt{2}}{\sqrt{\lambda}} + \frac{8\sqrt{2}}{\lambda^{3/2}n} + \frac{8}{\lambda n}\right) = c_1(c_1 + \frac{c_2}{n}).$$

Finally,

$$\|\mathbf{R} - \mathbf{R}'\|_1 = \sum_{i=1}^n |R_i - R'_i| \leq (n-1) \cdot \frac{c_1 c_2}{n} + c_1(c_1 + \frac{c_2}{n}) = c_1^2 + c_1 c_2,$$

as desired. \square

Corollary B.7. *Let a and b be known bounds on the domains of X and Y , respectively. Given a dataset $\mathbf{D} = (\mathbf{X}, \mathbf{Y}, \mathbf{Z})$, let $(\hat{\mathbf{X}}, \hat{\mathbf{Y}}, \mathbf{Z})$ be the rescaled dataset obtained by setting $\hat{\mathbf{X}} = \mathbf{X}/a$ and $\hat{\mathbf{Y}} = \mathbf{Y}/b$. Let PrivGCM be the algorithm which runs Algorithm 1 with kernel ridge regression as the fitting procedure \mathcal{F} and sensitivity bound $\Delta = C$, where C is the constant from Lemma B.6. Algorithm PrivGCM is ε -differentially private.*

The statistic $T^{(n)} = T(R_1 + W_1, \dots, R_n + W_n)$, defined in Algorithm 1, satisfies the following.

1. Let $\mathcal{P} \subset \mathcal{P}_0$ be a family of distributions such that $A_f A_g = o_{\mathcal{P}}(n^{-1})$, $B_f = o_{\mathcal{P}}(1)$, $B_g = o_{\mathcal{P}}(1)$. If in addition $\sup_{P \in \mathcal{P}} \mathbb{E}[|\chi_P \xi_P|^{2+\eta}] \leq c$, for some constants $c, \eta > 0$, then

$$\lim_{n \rightarrow \infty} \sup_{P' \in \mathcal{P}'} \sup_{t \in \mathbb{R}} |\Pr[T^{(n)} \leq t] - \Phi(t)| = 0.$$

2. Define $\rho_P = \mathbb{E}_P[\text{cov}(X, Y \mid Z)]$ and $\sigma_P = \sqrt{\text{Var}_P(\chi_P \xi_P)}$. Let $\mathcal{P} \subset \mathcal{E}_0$ be a family of distributions such that $A_f A_g = o_{\mathcal{P}}(n^{-1})$, $B_f = o_{\mathcal{P}}(1)$ and $B_g = o_{\mathcal{P}}(1)$. If in addition we have $\sup_{P \in \mathcal{P}} \mathbb{E}[|\chi_P \xi_P|^{2+\eta}] \leq c$, for some constants $c, \eta > 0$, then

$$\lim_{n \rightarrow \infty} \sup_{P' \in \mathcal{P}'} \sup_{t \in \mathbb{R}} \left| \Pr \left[T^{(n)} - \frac{\sqrt{n} \rho_P}{\sigma_{P'}} \leq t \right] - \Phi(t) \right| = 0,$$

$$\text{where } \sigma_{P'} = \sqrt{\sigma_P^2 + \left(\frac{\sqrt{2ab} \cdot C}{\varepsilon}\right)^2}.$$

Proof. The fact that PrivGCM is ε -differentially private follows from Lemmas A.2 and B.6. Note that if the regularization parameter λ is chosen adaptively based on the data, then obtaining the value of the constant C by plugging in λ might not be ε -differentially private. This can be resolved by setting an a priori lower bound on λ , independent of the data, e.g., $\lambda \geq 1$, and plugging that lower bound to obtain C .

Item 1 follows from Theorem B.1 and Lemma B.6. Item 2 follows from Theorem B.3 and Lemma B.6 □

B.4 Auxiliary Lemmas

Lemma B.8 (Slutsky's Lemma). *Let A_n, B_n be sequences of random variables. If A_n converges in distribution to A and B_n converges in probability to a constant c , then (1) $A_n + B_n \rightarrow A + c$, (2) $A_n \cdot B_n \rightarrow c \cdot A$, and (3) $A_n/B_n \rightarrow A/c$.*

Lemma B.9 (Uniform version of the Central Limit Theorem (Lemma 18 of Shah and Peters [30])). *Let \mathcal{P} be a family of distributions for a random variable ζ such that for all $P \in \mathcal{P}$ it holds $\mathbb{E}_P[\zeta] = 0$, $\mathbb{E}_P[\zeta^2] = 1$, and $\mathbb{E}_P[|\zeta|^{2+\eta}] < c$ for some $\eta, c > 0$. Let $(\zeta_i)_{i \in \mathbb{N}}$ be i.i.d copies of ζ . For $n \in \mathbb{N}$, define $S_n = \frac{1}{\sqrt{n}} \sum_{i=1}^n \zeta_i$. Then*

$$\lim_{n \rightarrow \infty} \sup_{P \in \mathcal{P}} \sup_{t \in \mathbb{R}} |\Pr[S_n \leq t] - \Phi(t)| = 0.$$

Lemma B.10 (Uniform version of the Weak Law of Large Numbers (Lemma 19 of Shah and Peters [30])). *Let \mathcal{P} be a family of distributions for a random variable ζ such that for all $P \in \mathcal{P}$ it holds $\mathbb{E}_P[\zeta] = 0$ and $\mathbb{E}_P[|\zeta|^{1+\eta}] < c$ for some $\eta, c > 0$. Let $(\zeta_i)_{i \in \mathbb{N}}$ be i.i.d copies of ζ . For $n \in \mathbb{N}$, define $S_n = \frac{1}{n} \sum_{i=1}^n \zeta_i$. Then for all $\delta > 0$ it holds*

$$\lim_{n \rightarrow \infty} \sup_{P \in \mathcal{P}} \Pr[|S_n| > \delta] = 0.$$

Lemma B.11 (Uniform version of Slutsky's Lemma (Lemma 20 of Shah and Peters [30])). *Let \mathcal{P} be a family of distributions that determines the law of sequences $(A_n)_{n \in \mathbb{N}}$ and $(B_n)_{n \in \mathbb{N}}$ of random variables.*

1. *If A_n converges uniformly over \mathcal{P} to $\mathcal{N}(0, 1)$, and $B_n = o_{\mathcal{P}}(1)$, then $A_n + B_n$ converges uniformly over \mathcal{P} to $\mathcal{N}(0, 1)$.*
2. *If A_n converges uniformly over \mathcal{P} to $\mathcal{N}(0, 1)$, and $B_n = 1 + o_{\mathcal{P}}(1)$, then A_n/B_n converges uniformly over \mathcal{P} to $\mathcal{N}(0, 1)$.*
3. *If $A_n = M + o_{\mathcal{P}}(1)$ for some $M > 0$ and $B_n = o_{\mathcal{P}}(1)$, then $A_n + B_n = M + o_{\mathcal{P}}(1)$.*
4. *If $A_n = o_{\mathcal{P}}(1)$ and $B_n = o_{\mathcal{P}}(1)$, then $A_n + B_n = o_{\mathcal{P}}(1)$.*

Claim B.12. *Let $(W_i)_{i \in \mathbb{N}}$ be i.i.d copies of $W \sim \text{Lap}(0, \sigma)$. Let $S_n = \max_{i \in [n]} |W_i|$. Then $\mathbb{E}[S_n] \leq \sigma(1 + \ln n)$.*

Proof. If $W_i \sim \text{Lap}(0, \sigma)$, then $|W_i| \sim \text{Exponential}(\sigma^{-1})$. It is a known fact that $\mathbb{E}[\max\{|W_1|, \dots, |W_n|\}] = \sigma \cdot H(n)$, where $H(n)$ is the harmonic number $H(n) = \sum_{i=1}^n 1/i$. The claim follows from the fact that $H(n) \leq \ln(n) + 1$. □

Claim B.13. *Let A and B be random variables and $p > 1$. Then $\mathbb{E}[|A + B|^p] \leq 2^{p-1}(\mathbb{E}[|A|^p] + \mathbb{E}[|B|^p])$.*

Claim B.14. *Let $W \sim \text{Laplace}(0, \sigma)$. Then for $n \in \mathbb{N}$ we have $\mathbb{E}[|W|^n] = \frac{n!}{\sigma^n}$.*

Proof. This claim follows from the fact that $|W| \sim \text{Exponential}(\sigma^{-1})$. □

C Proofs of Section 3

Lemma 3.2. (Sensitivity of the score function). *For all $c \in [0, m]$, the values Q_c (Definition 3.1) have sensitivity at most Δ_T , and the score $s_k(c, \mathbf{D})$ has sensitivity at most 1.*

Proof. Let $\{T_i\}_{i \in [0, m]}$ be the values of $m + 1$ queries of sensitivity at most Δ_T on a dataset \mathbf{D} . Let $\{T'_i\}_{i \in [B]}$ be the values of the same queries on a neighboring dataset \mathbf{D}' . Let Q_0, \dots, Q_m (respectively Q'_0, \dots, Q'_m) denote the values $\{T_i\}_{i \in [0, m]}$ (respectively $\{T'_i\}_{i \in [0, m]}$) sorted in decreasing order.

We show that $|Q_c - Q'_c| \leq \Delta_T$ for all $c \in [0, m]$. Fix some $c \in [0, m]$. Suppose by contradiction that $|Q_c - Q'_c| > \Delta_T$. Consider the case when $Q_c > Q'_c$. Then $Q_c > Q'_c + \Delta_T$. Let $\text{Above}_c = \{i \in [0, m] \mid T_i \geq Q_c\}$. Define Above'_c similarly. Then, for all $i \in \text{Above}_c$, we have

$$T'_i \geq T_i - \Delta_T \geq Q_c - \Delta_T > Q'_c, \quad (14)$$

where the first inequality holds since the values T_i have sensitivity at most Δ_T , the second inequality holds from $i \in \text{Above}_c$, and the last inequality holds from our assumption by contradiction. Thus $i \in \text{Above}'_c$, and as a result $\text{Above}_c \subseteq \text{Above}'_c$. Moreover, for j such that $T'_j = Q'_c$ we have that $j \notin \text{Above}_c$ since it does not satisfy (14). Therefore $|\text{Above}'_c| > |\text{Above}_c|$, a contradiction.

For the case when $Q_c < Q'_c$ we obtain a contradiction by a symmetric argument. This concludes the proof on the sensitivity of Q_c . Next, we bound the sensitivity of the score function. We have,

$$|s_k(c, \mathbf{D}) - s_k(c, \mathbf{D}')| \leq \frac{|Q_c - T_k| - |Q'_c - T'_k|}{2\Delta_T} \leq \frac{|Q_c - Q'_c| + |T_k - T'_k|}{2\Delta_T}.$$

We just showed that $|Q_c - Q'_c| \leq \Delta_T$. Since the queries T_i have sensitivity at most Δ_T , we also have $|T_k - T'_k| \leq \Delta_T$. We obtain that $|s_k(c, \mathbf{D}) - s_k(c, \mathbf{D}')| \leq 1$ for all neighboring datasets \mathbf{D}, \mathbf{D}' . \square

Lemma C.1. (Sensitivity of T). *Consider two neighboring datasets $\mathbf{D} = (\mathbf{X}^{(0)}, \dots, \mathbf{X}^{(m)}, \mathbf{Y}, \mathbf{Z})$ and $\mathbf{D}' = (\mathbf{X}'^{(0)}, \dots, \mathbf{X}'^{(m)}, \mathbf{Y}', \mathbf{Z}')$. For $j \in [0, m]$, let $T_j = T(\mathbf{X}^{(j)}, \mathbf{Y}, \mathbf{Z})$. Define T'_j analogously. If $|y_i|, |y'_i| \leq 1$ for all $i \in [n]$ and $|r_{X,i}^{(j)}|, |r'_{X,i}{}^{(j)}| \leq 1^9$ for all $i \in [n], j \in [0, m]$, then $|T_j - T'_j| \leq C'$, where $C' = 4 \cdot (1 + \frac{\sqrt{2}}{\sqrt{\lambda}} + \frac{2\sqrt{2}}{\lambda^{3/2}} + \frac{2}{\lambda})$.*

Proof. Suppose that \mathbf{D} and \mathbf{D}' differ in the last row. In the following, we assume that $j \in [m]$ is fixed. To ease notation, we remove the superscript (j) from all $r_{X,i}^{(j)}, r_{Y,i}^{(j)}$ and $R_i^{(j)}$. Since we know $\mathbb{E}[X \mid Z]$ exactly, we have $r_{X,i} = x_i - \mathbb{E}[X \mid Z = z_i]$ for $i \in [n], j \in [m]$. Then for $i \in [n]$ we have

$$|r_{X,i} - r'_{X,i}| = |(x_i - \mathbb{E}[X \mid Z = z_i]) - (x'_i - \mathbb{E}[X \mid Z = z'_i])|.$$

If $i \in [n-1]$, then $x_i = x'_i$ and $z_i = z'_i$, so that $|r_{X,i} - r'_{X,i}| = 0$. For $i = n$, we have $|r_{X,n} - r'_{X,n}| \leq 2$ by the triangle inequality and since $|r_{X,n}| \leq 1$.

Let $c_1 = 2 + \frac{2\sqrt{2}}{\sqrt{\lambda}}$ and $c_2 = \frac{8\sqrt{2}}{\lambda^{3/2}} + \frac{8}{\lambda}$. Turning to the residuals of fitting \mathbf{Y} to \mathbf{Z} , by the same argument as in Lemma B.6 we have, for all $i \in [n-1]$,

$$|r_{Y,i} - r'_{Y,i}| \leq \frac{8\sqrt{2}}{\lambda^{3/2}n} + \frac{8}{\lambda n} = \frac{c_2}{n}.$$

For the last datapoint it holds

$$|r_{Y,n} - r'_{Y,n}| \leq 2 + \frac{2\sqrt{2}}{\sqrt{\lambda}} + \frac{8\sqrt{2}}{\lambda^{3/2}n} + \frac{8}{\lambda n} = c_1 + \frac{c_2}{n}.$$

Additionally, $|r_{Y,i}| \leq c_1/2$ for all $i \in [n]$.

We can now bound the sensitivity of the residual products $R_i = r_{X,i}r_{Y,i}$. For $i \in [n-1]$, we have

$$|R_i - R'_i| \leq |r_{X,i}| |r_{Y,i} - r'_{Y,i}| + |r'_{Y,i}| |r_{X,i} - r'_{X,i}| \leq \frac{c_2}{n} + 0.$$

⁹The bound of 1 can be replaced by any constant.

For $i = n$ we have

$$|R_n - R'_n| \leq \left(c_1 + \frac{c_2}{n}\right) + \frac{c_1}{2} \cdot 2 = 2c_1 + \frac{c_2}{n}.$$

Finally,

$$|T_j - T'_j| = \left| \sum_{i=1}^n R_i - \sum_{i=1}^n R'_i \right| \leq \sum_{i=1}^n |R_i - R'_i| \leq (n-1) \cdot \frac{c_2}{n} + 2c_1 + \frac{c_2}{n} = 2c_1 + c_2.$$

□

Theorem 3.6. PrivCRT is ε -DP and (γ, δ) -accurate for $\gamma = \frac{4\Delta_T}{\varepsilon} \log\left(\frac{m}{\delta}\right)$.

Proof. We first show that PrivCRT is ε -differentially private. The scores $T_i, i \in [0, m]$ have sensitivity at most $\Delta_T = O(1)$ by Lemma C.1. Therefore, the scores $s_i, i \in [0, m]$ have sensitivity at most 1 by Lemma 3.2. Finally, by Theorem 3.3 we have that outputting \hat{c} is ε -DP. Therefore, PrivCRT is ε -DP. Note that if the regularization parameter λ is chosen adaptively based on the data, then obtaining the value of the constant C' by plugging in λ might not be ε -differentially private. This can be resolved by setting an a priori lower bound on λ , independent of the data, e.g., $\lambda \geq 1$, and plugging that lower bound to obtain C' .

Next, we analyze the accuracy of PrivCRT. Let c^* be the true rank of T_0 amongst the statistics $\{T_i\}_{i \in [0, m]}$, sorted in decreasing order. Then $p^* = (1 + c^*)/(m + 1)$. Note that $s_{c^*} = 0$. By Theorem 3.3, with probability at least $1 - \delta$, it holds

$$s_{\hat{c}} \geq -\frac{2 \log(m/\delta)}{\varepsilon}.$$

As a result,

$$|Q_{\hat{c}} - T_0| = -2\Delta_T \cdot s_{\hat{c}} \leq \frac{4\Delta_T \log(m/\delta)}{\varepsilon}.$$

Let $\gamma = \frac{4\Delta_T \log(m/\delta)}{\varepsilon}$. The rank c^* of T_0 cannot differ from \hat{c} by more than G_γ , since $Q_{\hat{c}}$ is within distance γ of T_0 . Therefore, $|\hat{p} - p^*| \leq \frac{G_\gamma}{m+1}$. □

D Additional Experimental Details and Results

Example Dataset. In Fig. 6 we show one example dataset from our simulations. We fit a kernel ridge regression model of X to Z . As can be observed, the model we fit matches $\mathbb{E}[X|Z]$ closely.

Implementation of Private Kendall and PrivToT. The implementation of PrivToT requires setting a parameter k for the number of subsets into which the original dataset is divided. We use $k = 10$ for the experiments in Fig. 1. We additionally experimented with $k \in [20, 50, 100]$ and found that $k = 10$ performs the best in terms of type-I error. The implementation of PrivToT was adapted from the implementation of Kazan et al. [18]. The private Kendall test works for categorical Z . To adapt it to continuous Z we apply k -means clustering to Z to obtain 100 clusters. The implementation of private Kendall was adapted from Wang et al. [39].

PrivCRT vs. PrivGCM. In Fig. 7 we compare the performance of PrivCRT with PrivGCM for different privacy parameters $\varepsilon \in [2^{-3}, 2^3]$. We set $n = 10^4$ and, for PrivCRT, $m = 19$. Both tests control type-I error, but PrivCRT achieves better power than PrivGCM for all privacy parameters ε . Therefore, PrivCRT appears preferable to PrivGCM when dataholders have access to the distribution $X | Z$. This result is consistent with the non-private scenario where the CRT has higher power because it does not have to learn $\mathbb{E}[X | Z]$.

The Effect of Varying m on the Power of PrivCRT. In Fig. 8, we vary m , the number of resamples used in the CRT algorithm and run our experiments for $\beta \in \{0, 0.5, 1.5\}$. We set $n = 10^3$ and $\varepsilon = 2$. Increasing m does not affect the type-I error control of PrivCRT. However, we observe that the power of PrivCRT decreases as the number of resamples m increases. This is due to the increase of the variable G_γ with m (see Definition 3.5). More specifically, for a fixed γ , as m increases, T_0 stays fixed, while the number of other statistics T_i within distance γ of T_0 increases. Thus, the private test is more likely to select a rank that is further from the true rank. It is an interesting open question whether the dependence on m in the accuracy of a private CRT test is avoidable. For now, we recommend using $m = O(1/\alpha)$ when employing PrivCRT.

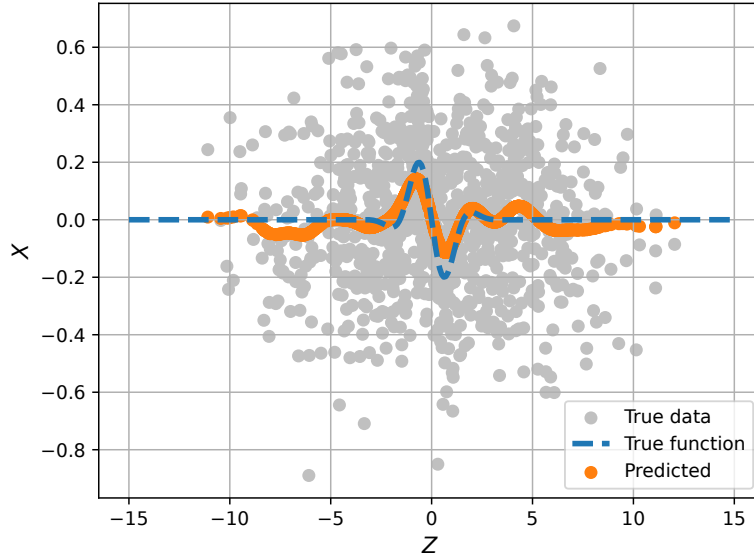


Figure 6: Values of X and Z (after rescaling) of one sampled dataset from our simulations, with $n = 1000$, $\beta = 0$, $s = 2$, $d = 1$. A kernel ridge regression model is fitted to the data. The model we fit closely matches $\mathbb{E}[X|Z]$.

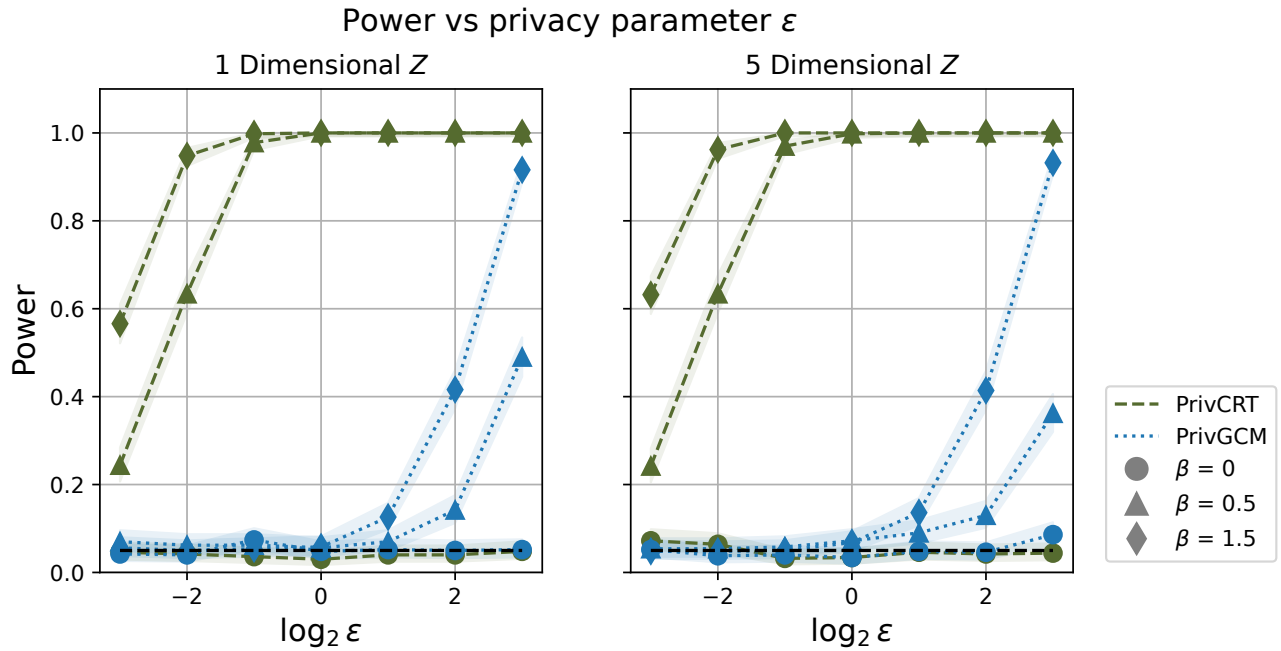


Figure 7: Power of PrivCRT and PrivGCM versus privacy ϵ .

Distribution of p-values Output by PrivCRT. In Fig. 9 we show the distribution of the p-values output by PrivCRT for different dependence strengths β . We set $n = 10^3$, $m = 19$, and $\epsilon = 2$. Under the null, i.e., when $\beta = 0$, the p-values output by PrivCRT are uniformly distributed in the interval $[\frac{1}{m+1}, 1]$. Thus, PrivCRT controls type-I error. When $\beta > 0$, most of the p-values are close to $\alpha = 0.05$, which is the desired outcome for PrivCRT to achieve power.

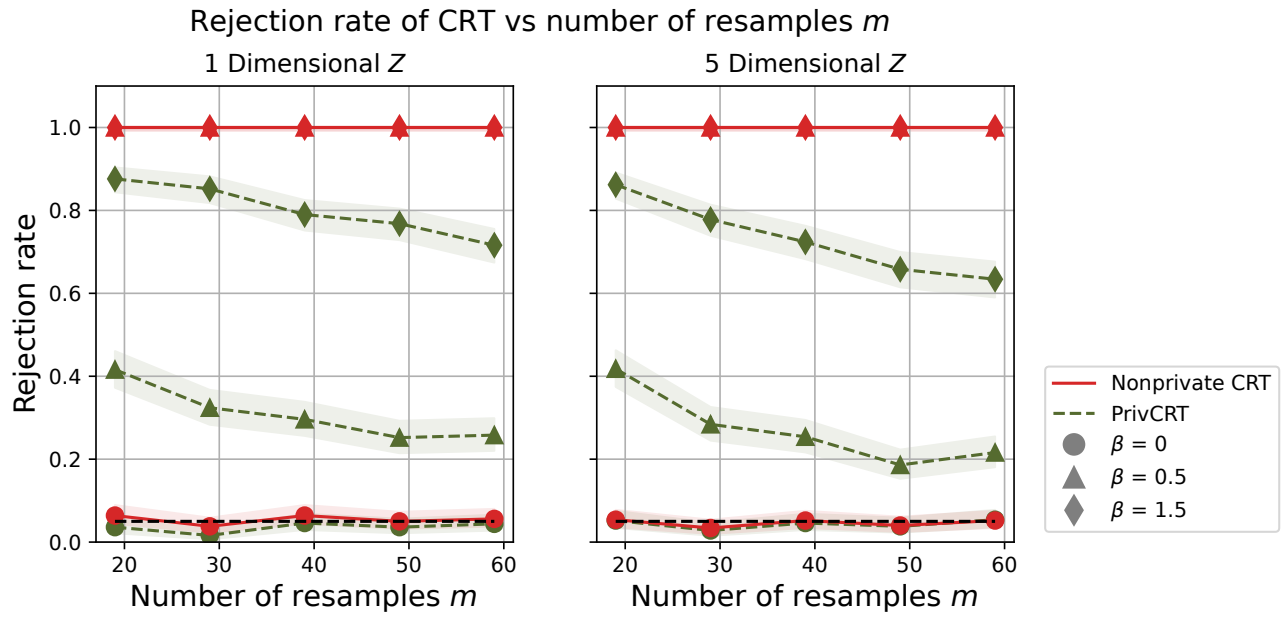


Figure 8: Effect on the power of PrivCRT with increasing m .

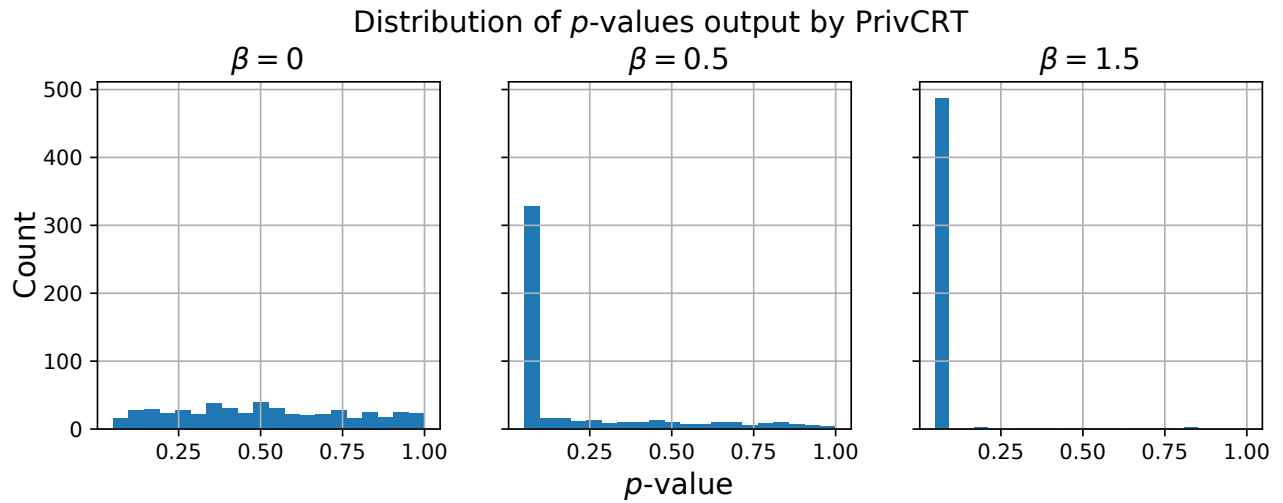


Figure 9: Distribution of p -values output by PrivCRT for different dependence strengths β under the setup in Section 4. Under the null, i.e., $\beta = 0$, the p -values are uniformly distributed as desired.