

Adaptive False Discovery Rate Control with Privacy Guarantee

Xintao Xia

*Department of Statistics
Iowa State University
Ames, IA 50011-1090, USA*

XINTAOX@IASTATE.EDU

Zhanrui Cai

*Faculty of Business and Economics
The University of Hong Kong
Hong Kong, China*

ZHANRUIC@HKU.HK

Editor: Po-Ling Loh

Abstract

Differentially private multiple testing procedures can protect the information of individuals used in hypothesis tests while guaranteeing a small fraction of false discoveries. In this paper, we propose a differentially private adaptive FDR control method that can control the classic FDR metric exactly at a user-specified level α with a privacy guarantee, which is a non-trivial improvement compared to the differentially private Benjamini-Hochberg method proposed in Dwork et al. (2021). Our analysis is based on two key insights: 1) a novel p -value transformation that preserves both privacy and the mirror conservative property, and 2) a mirror peeling algorithm that allows the construction of the filtration and application of the optimal stopping technique. Numerical studies demonstrate that the proposed DP-AdaPT performs better compared to the existing differentially private FDR control methods. Compared to the non-private AdaPT, it incurs a small accuracy loss but significantly reduces the computation cost.

Keywords: Selective inference, differential privacy, false discovery rate, model-free.

1. Introduction

1.1 Differential privacy

With the advancement of technology, researchers are able to collect and analyze data on a large scale and make decisions based on data-driven techniques. However, privacy issues could be encountered without a proper mechanism for data analysis and may lead to serious implications. For example, in bioinformatics, genomic data are usually very sensitive and irreplaceable, and it is of great importance to protect the individual's privacy in genome analysis, including GWAS (Kim et al., 2020). Many countries have classified genomic data as sensitive and must be handled according to certain regulations, such as the HIPAA in the USA and the Data Protection Directive in the European Union. The leak of individual genetic information can have severe consequences. It may underpin the trust of data-collecting agencies and discourage people or companies from sharing personal information.

The leak of individual genetic information may also lead to serious social problems, such as genome-based discrimination (Kamm et al., 2013).

In recent literature, a popular procedure to protect privacy is to apply *differentially private* algorithms in data analysis. First proposed by Dwork et al. (2006), the concept of differential privacy has seen successful applications in numerous fields, including but not limited to healthcare, information management, government agencies, etc. Privacy is achieved by adding proper noise to the algorithm (Fuller, 1993) and obscuring each individual’s characteristics. A differentially private procedure guarantees that an adversary can not judge whether a particular subject is included in the data set with high probability; thus is extremely useful in protecting personal information. During the past decades, considerable effort has been devoted to developing machine learning algorithms to guarantee differential privacy, such as differentially private deep learning (Abadi et al., 2016) or boosting (Dwork et al., 2010). In the statistics literature, Wasserman and Zhou (2010) estimated the convergence rate of density estimation in differential privacy. Other applications include but are not limited to differential privacy for synthetic data generation (Awan and Cai, 2020), functional data (Hall et al., 2013), network data (Karwa and Slavković, 2016), mean estimation, linear regression (Cai et al., 2021), etc. More recently, Dwork et al. (2021) proposed the private Benjamini-Hochberg procedure to control the false discovery rate in multiple hypothesis testing. We refer readers to the classic textbook by Dwork et al. (2014) for a comprehensive review of differential privacy.

1.2 False Discovery Rate Control

In modern statistical analysis, large-scale tests are often conducted to answer research questions from scientists or the technology sector’s management team. For example, in bioinformatics, researchers compare a phenotype to thousands of genetic variants and search for associations of potential biological interest. It is crucial to control the expected proportion of falsely rejected hypotheses, i.e., the *false discovery rate* (Benjamini and Hochberg, 1995). Controlling the false discovery rate (FDR) lets scientists increase power while maintaining a principled bound on the error. Let R be the number of total rejections and V be the number of false rejections; the FDR is defined as

$$\text{FDR} = \mathbb{E} \left[\frac{V}{\max\{R, 1\}} \right]. \quad (1)$$

The most famous multiple-testing procedure is the Benjamini–Hochberg (BH) procedure (Benjamini and Hochberg, 1995). Given n hypotheses and their ordered p -values $p_{(1)} < p_{(2)} < \dots < p_{(n)}$, the BH procedure rejects any null hypothesis whose p -value is non-greater than $\max\{p_{(i)} : p_{(i)} \leq \alpha i/n\}$, where α is a user-specified target level for FDR. Benjamini and Yekutieli (2001) extended the BH procedure to the setting where all the test statistics have positive regression dependency. Recent works focus on settings where prior information or extra data for hypotheses are available. The side information can be integrated to weighting the p -values (Genovese et al., 2006; Dobriban et al., 2015), exploring the group structures (Hu et al., 2010) or natural ordering (Barber and Candès, 2015; Li and Barber, 2017) among hypotheses, etc. Adaptively focusing on the more promising hypotheses can also lead to powerful multiple-testing procedures (Lei and Fithian, 2018; Tian and Ramdas,

2019). More recently, researchers developed sample-splitting-based procedures to control the false discovery rate (Dai et al., 2022; Tong et al., 2022).

Most multiple-testing procedures put assumptions on the p -values. A natural and mild assumption is that the p -values under the null hypothesis follow the uniform distribution on the interval $[0, 1]$. Because many statistical tests tend to be conservative under the null (Cai et al., 2022a; Cai, 2023), it is also common to assume that the p -values are stochastically larger than the uniform distribution, or *super-uniform*: $\mathbb{P}(p_i \leq t) \leq t, \forall t \in [0, 1]$ for $i \in \mathcal{H}_0$, where \mathcal{H}_0 denotes the true null hypotheses, see for example, Li and Barber (2017); Ramdas et al. (2019). Adaptive online FDR control procedure tends to require stronger assumptions. Tian and Ramdas (2019) assumes that all the null p -values are *uniformly conservative*, i.e., $\mathbb{P}(p_i/\tau \leq t \mid p_i \leq \tau) \leq t, \forall t, \tau \in (0, 1)$ for $i \in \mathcal{H}_0$. The AdaPT procedure proposed by Lei and Fithian (2018) assumes that the null p -values are *mirror conservative*:

$$\mathbb{P}(p_i \in [a_1, a_2]) \leq \mathbb{P}(p_i \in [1 - a_2, 1 - a_1]), \quad \forall 0 \leq a_1 \leq a_2 \leq 0.5. \quad (2)$$

Those assumptions on null p -values all cover the uniform distribution as a special case and hold under various scenarios, as discussed in the literature. Intuitively, mirror-conservatism allows us to control the quantity of small null p -values by referencing the number of large null p -values, thereby providing a way to control the FDR. It is important to note that mirror-conservatism doesn't automatically result in super-uniformity, and likewise, super-uniformity doesn't guarantee mirror-conservatism. Null p -values with a convex CDF or a monotonically increasing density are uniformly conservative, and such conservatism implies both super-uniformity and mirror-conservatism. This paper will build on the mirror conservative assumption to develop an adaptive differentially private FDR control procedure.

1.3 Related Work and Contributions

The most related work to our paper is the differentially private BH procedure proposed in Dwork et al. (2021). Dwork et al. (2021) provides conservative bounds for $\text{FDR}_k := \mathbb{E}[V/R \mid V \geq k], k \geq 2$, and $\text{FDR}^k := \mathbb{E}[V/R \mid R \geq k], k \geq 1$. However, conservatism is unavoidable in the approach proposed in Dwork et al. (2021) due to the additional noise required for privacy. Specifically, their procedure can only bound FDR_k (or FDR^k) to be non-greater than $C_k\alpha$, where C_k is a constant larger than 1.

In this paper, we propose an adaptive differentially private FDR control method that can control the FDR in the classical sense: $\text{FDR} \leq \alpha$, without any conditional component for the false discovery proportions or the constant term that inflates α . Our work is based on a novel private p -value transformation mechanism that can protect the privacy of individual p -values while maintaining the mirror conservative assumption on the null p -values. By further developing a mirror peeling algorithm, we can define a filtration and apply the optimal stopping technique to prove that the proposed DP-AdaPT method controls FDR at any user-specified level α with finite samples. Theoretically, the proposed method provides a stronger guarantee of false discovery rate control compared to the differentially private BH method. Numerically, the proposed method works as well as the differentially private BH method when only the p -values are available for each test and performs better when side information is available. The proposed method is also model-free when incorporating the side information for each hypothesis test. Lastly, the method is shown to only incur a

small accuracy loss compared to the non-private AdaPT (Lei and Fithian, 2018) but at the same time reduces huge computation costs.

This paper is organized as follows. Section 2 defines the basic concepts of differential privacy and briefly introduces the AdaPT procedure of FDR control. Section 3 provides the private p -value transformation mechanism, the definition of sensitivity for p -values, the DP-AdaPT algorithm, and the guaranteed FDR control. We demonstrate the numerical advantage of DP-AdaPT through extensive simulations in Section 4 and conclude the paper with some discussion on future work in Section 5.

2. Preliminaries

2.1 Differential Privacy

We first introduce the background for differential privacy. A dataset $\mathcal{S} = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ is a collection of n records, where $\mathbf{x}_i \in \mathcal{X}$ for $i = 1, \dots, n$ and \mathcal{X} is the domain of \mathbf{x} . The random variable $\{\mathbf{x}_i\}_{i=1}^n$ does not have to be independent. Researchers are usually concerned with certain statistics or summary information based on the dataset, denoted as $\mathcal{T}(\mathcal{S})$. For example, one might be interested in the sample mean, the regression coefficients, or specific test statistics. When the dataset \mathcal{S} is confidential and contains private individual information, researchers prefer to release a randomized version of $\mathcal{T}(\mathcal{S})$, which we denote as $\mathcal{M}(\mathcal{S})$. A *neighboring dataset* to \mathcal{S} is denoted by $\mathcal{S}' = \{\mathbf{x}'_1, \dots, \mathbf{x}'_n\}$, with the requirement that only one index $j \in \{1, \dots, n\}$ satisfies that $\mathbf{x}_j \neq \mathbf{x}'_j$. The classic (ϵ, δ) -differential privacy (Dwork et al., 2006) is defined as follows.

Definition 1 A randomized mechanism $\mathcal{M}(\cdot)$ is (ϵ, δ) -differentially private for $\epsilon > 0$ and $\delta > 0$, if for all neighboring datasets \mathcal{S} and \mathcal{S}' , and any measurable set E ,

$$\mathbb{P}(\mathcal{M}(\mathcal{S}) \in E) \leq e^\epsilon \mathbb{P}(\mathcal{M}(\mathcal{S}') \in E) + \delta. \quad (3)$$

When $\delta = 0$, Definition 1 is the *pure differential privacy* and denoted by ϵ -DP. When $\delta > 0$, it is called the *approximate differential privacy*. The two neighboring datasets are treated as fixed, and the mechanism $\mathcal{M}(\cdot)$ contains randomness that is independent of the dataset and protects privacy. The set E is measurable with respect to the random variable $\mathcal{M}(\cdot)$. In the definition, the two parameters ϵ and δ control the difference between the likelihood of $\mathcal{M}(\mathcal{S})$ and $\mathcal{M}(\mathcal{S}')$. A small value of ϵ and δ indicates that the difference between the distribution of $\mathcal{M}(\mathcal{S})$ and $\mathcal{M}(\mathcal{S}')$ is small and, as a result, using the outcome from the mechanism \mathcal{M} , one can hardly tell whether a single individual is included in the dataset \mathcal{S} . Thus, privacy is guaranteed with high probability for each individual in the dataset \mathcal{S} .

Dong et al. (2022) proposed to formulate privacy protection as a hypothesis-testing problem for two neighboring datasets \mathcal{S} and \mathcal{S}' :

$$H_0 : \text{the underlying dataset is } \mathcal{S}, \quad \text{versus} \quad H_1 : \text{the underlying dataset is } \mathcal{S}'. \quad (4)$$

Let \mathbf{x}_j denote the only individual in \mathcal{S} , but not in \mathcal{S}' . Accepting the null hypothesis implies identifying the presence of \mathbf{x}_j in the dataset \mathcal{S} , and rejecting the null hypothesis implies identifying the absence of \mathbf{x}_j . Thus privacy can be interpreted by the power function of testing (4). Specifically, the μ -Gaussian differential privacy is defined as a test that is at

least as hard as distinguishing between two normal distributions $N(0, 1)$ and $N(\mu, 1)$ based on one random draw of the data. For the readers' convenience, we rephrase the formal definition from Dong et al. (2022).

- Definition 2 (Gaussian Differential Privacy)** 1. A mechanism \mathcal{M} is f -differential private (f -DP) if any α -level test of (4) has power function $\beta(\alpha) \leq 1 - f(\alpha)$, where $f(\cdot)$ is a convex, continuous, non-increasing function satisfying $f(\alpha) \leq 1 - \alpha$ for all $\alpha \in [0, 1]$.
2. A mechanism \mathcal{M} is μ -Gaussian Differential Privacy (μ -GDP) if \mathcal{M} is G_μ -DP, where $G_\mu(\alpha) = \Phi(\Phi^{-1}(1 - \alpha) - \mu)$ and $\Phi(\cdot)$ is the cumulative distribution function of $N(0, 1)$.

The new definition has several advantages. For example, privacy can be fully described by a single mean parameter of a unit-variance Gaussian distribution, and this makes it easy to describe and interpret the privacy guarantees. The privacy definition maintains a tight privacy guarantee under multiple compositions of private mechanisms. Thus, it is particularly useful for statistical methods that require multiple or iterative operations of the data. We will use the definition of μ -GDP throughout the rest of this paper. The proposed method can be easily extended to the classic (ϵ, δ) -DP by Corollary 1 in Dong et al. (2022).

2.2 Adaptive False Discovery Rate Control

In this subsection, we introduce the AdaPT procedure proposed by Lei and Fithian (2018), which is described in Algorithm 1 for completeness. Assume we have p -values p_i and side information \mathbf{x}_i for each hypothesis H_i , $i = 1, \dots, n$. The procedure contains an iterative update of covariate-specific thresholds. At each step $t = 0, 1, \dots$, a rejection threshold $s_t(\mathbf{x})$ is decided based on the covariate \mathbf{x} . Let $R_t = |\{i : p_i \leq s_t(\mathbf{x}_i)\}|$, $A_t = |\{i : p_i \geq 1 - s_t(\mathbf{x}_i)\}|$ and the estimated false discovery rate $\widehat{\text{FDR}} = (1 + A_t)/(R_t \vee 1)$. If $\widehat{\text{FDR}} \leq \alpha$, then we stop and reject all the H_i with $p_i \leq s_t(\mathbf{x}_i)$. Otherwise, we update the thresholds $s_{t+1} \preceq s_t$, where $s_{t+1} \preceq s_t$ denotes $s_{t+1}(\mathbf{x}) \leq s_t(\mathbf{x})$ for every \mathbf{x} in the domain of $s(\cdot)$. The information that are used to update s_{t+1} contains A_t , R_t and $(\mathbf{x}_i, p_{pm,i})_{i=1}^n$, where

$$p_{pm,i} = \begin{cases} p_i & s_t(\mathbf{x}_i) < p_i < 1 - s_t(\mathbf{x}_i) \\ \{p_i, 1 - p_i\} & \text{otherwise,} \end{cases}$$

is the partially masked p -value and the subscript pm denotes ‘‘partially masked’’. The partially censored p -values restrict the analyst’s knowledge and enable the application of the optional stopping technique widely used in the FDR literature (Storey et al., 2004; Barber and Candès, 2015; Li and Barber, 2017). In this paper, we will develop a mirror-peeling algorithm that builds on this novel technique and prove the desired FDR guarantee under differential privacy.

3. Methodology

Consider n hypotheses H_1, H_2, \dots, H_n , and researchers can observe side information \mathbf{x}_i and estimate a p -value for each hypothesis H_i . In this section, we aim to develop a differentially private algorithm that protects the privacy of individual p -values and controls the FDR at

Algorithm 1 AdaPT (Lei and Fithian, 2018)

Require: $\{\mathbf{x}_i, p_i\}_{i=1}^n$, initialization s_0 , target FDR level α .
for $t = 0$ to \dots **do**
 $\widehat{\text{FDR}} \leftarrow \frac{1+A_t}{R_t \vee 1}$
 if $\widehat{\text{FDR}} \leq \alpha$
 Reject $\{H_i : p_i \leq s_t(\mathbf{x}_i)\}$
 Return s_t
 end if
 $s_{t+1} \leftarrow \text{Update}[(\mathbf{x}_i, p_{pm,i})_{i=1}^n, A_t, R_t, s_t]$
end for

the same time. The analysis does not rely on the threshold model $s(\mathbf{x})$ and is model-free. We assume that the auxiliary information \mathbf{x} is public and is not subject to privacy concerns. This assumption is reasonable because the auxiliary information is usually from scientific knowledge or previous experiments. With a specific model for the threshold function $s(\mathbf{x})$, the proposed method can also be easily extended to further protect the privacy of \mathbf{x} .

3.1 Private p-value

Following Lei and Fithian (2018), we assume that all the null p -values satisfy the mirror-conservative property as defined in (2). We first propose a novel differentially private mechanism on the individual p -values that protects privacy while still satisfying the mirror-conservative property. This is a crucial property because it helps us avoid the traditional technique in the differential privacy literature (e.g., Dwork et al. (2021)) that derives conservative error bounds on the noise added for privacy.

The proposed mechanism is based on the quantile function and cumulative distribution function of some symmetric distributions. Specifically, let $U \in \mathbb{R}^+$ be the boundary and can take the value of ∞ . Let $g(\cdot) : (-U, U) \rightarrow \mathbb{R}^+$ be an integrable function satisfying the following conditions:

1. Non-negative: $g(x) \geq 0$ for $x \in (-U, U)$, $g(x) = 0$ for $x \notin (-U, U)$ and the measure of the set $\{x \in (-U, U) : g(x) = 0\}$ is zero with respect to the measure μ on \mathbb{R} ;
2. Symmetric: $g(x) = g(-x)$ for $x \in (-\infty, \infty)$;
3. Unity: $\int_{-\infty}^{\infty} g(x)\mu(dx) = 1$.

The primitive function of $g(\cdot)$ is denoted by $G(x) = \int_{-\infty}^x g(x)\mu(dx)$. The function $g(\cdot)$ can be viewed as a symmetric probability density function, and the function $G(\cdot)$ can be viewed as a strictly increasing cumulative distribution function. The function $G(\cdot)$ is a one-to-one mapping from $(-U, U)$ to $[0, 1]$, which guarantees the existence of a quantile function $G^{-1}(\cdot)$. We will use $G^{-1}(\cdot)$ and $G(\cdot)$ to transform the p -values. When the distribution of p -value is continuous, the measure μ can be chosen as the Lebesgue measure.

Theorem 3 *Let the p -value p be mirror-conservative, and Z be an independent Gaussian random variable with mean zero and positive variance. Then the noisy p -value*

$$\tilde{p} := G\{G^{-1}(p) + Z\} \tag{5}$$

is also mirror-conservative.




Figure 1: Empirical density estimate of both the original and the transformed p -values. Case a: the null p -values follow the uniform distribution; case b: the null p -values have density $f(p) = 2p$ for $p \in [0, 1]$. In both cases, $G(\cdot) = \Phi(\cdot)$ and Z follows the standard normal distribution. The density curves are estimated by 100 000 samples.

The proof of Theorem 3 is provided in the appendix. Although Theorem 3 is based on the Gaussian noise, one can easily extend the theory to the case where Z follows Laplace distribution which is frequently used in the classic (ϵ, δ) -DP setting, see the discussions in Theorem 10. We provide two illustrative examples in Figure 1, where the empirical density estimates of both the original and the transformed p -values are plotted. On the top row, we show that when the original p -values follow the uniform distribution, the transformed noisy p -values are symmetric around 0.5. On the bottom row, we show that when the original p -values are strictly stochastically larger than the uniform distribution, the transformed noisy p -values also tend to concentrate on the right side of the curve. Figure 1 visually demonstrates that the two commonly encountered p -values can preserve the mirror conservative property. Note that in Figure 1, we implemented the standard normal distribution to transform the p -values: $G(\cdot) = \Phi(\cdot)$ and $g(\cdot) = \phi(\cdot)$. The normal density $\phi(\cdot)$ is monotone on either the positive or negative part of the horizontal axis. As a result, the transformed noisy p -values will concentrate on the two endpoints 0 and 1.

It is also interesting to note that Theorem 3 does not necessarily hold for the other conservative assumptions on p -values, such as the super uniform assumption or the uniformly conservative assumption as discussed in the introduction. One can easily construct a counterexample that violates the requirements. The mirror conservative condition is the

most appropriate in the sense of preserving differential privacy. Throughout the rest of the paper, we will use \tilde{p} to denote the noisy p -values defined in (5).

3.2 Sensitivity of p -values

The transformation in Theorem 3 is useful in preserving the mirror-conservative property, but we need to calibrate the variance of noise Z in order to provide a privacy guarantee with minimal losses on accuracy. In this section, we provide a definition for the *sensitivity* of p -values that directly fits into the framework of the transformation in Theorem 3. We begin with defining the *sensitivity* for any deterministic real-valued functions. The definition provides an upper bound of the difference in the outcome due to the change of one item in the dataset.

Definition 4 Let $g : \mathcal{X}^n \rightarrow \mathbb{R}$ be a deterministic function from the data set \mathcal{S} to \mathbb{R} . The *sensitivity* of $g(\cdot)$ is defined by

$$\Delta(g) := \sup_{\forall \mathcal{S}, \mathcal{S}'} \|g(\mathcal{S}) - g(\mathcal{S}')\|,$$

where \mathcal{S}' is a neighboring dataset of \mathcal{S} , and $\|\cdot\|$ is the Euclidean norm.

In this paper, we consider the case where the p -value is estimated from a non-randomized decision rule, and thus the p -value is a deterministic real-valued function of the data. However, due to its nature, the relative change of a p -value on two neighboring datasets is usually very small. Thus directly adding noise to p -values may easily overwhelm the signals and lead to unnecessary power losses. For example, Dwork et al. (2021) controls the sensitivity of p -value based on a truncated log transformation, and the truncation parameter has to be carefully tuned to make a tradeoff between privacy and accuracy.

In this paper, we define the sensitivity by considering a transformed p -value based on the function $G^{-1}(\cdot)$. The transformation is motivated by the fact that the p -values are usually obtained based on the limiting null distribution of the test statistics, which are, in most cases, asymptotically normal. For example, the p -value of a one-sided mean test is the quantile function of a normal distribution evaluated at the sample mean. We provide the formal definition as follows.

Definition 5 (Sensitivity) The sensitive of p -value function p is Δ_G if for all neighboring dataset \mathcal{S} and \mathcal{S}' ,

$$\sup_{\forall \mathcal{S}, \mathcal{S}'} \|G^{-1}\{p(\mathcal{S})\} - G^{-1}\{p(\mathcal{S}')\}\| \leq \Delta_G.$$

The choice of $G(\cdot)$ function in the Definition 5 is flexible. For example, if the density of the test statistics under the null hypothesis is symmetric, then one can choose $G(\cdot)$ as the CDF of the test statistic. We provide the following examples where the desired Δ_G is calculated under mild conditions.

Example 1 Assume X_1, X_2, \dots, X_n are i.i.d. random variables with mean β , variance 1, and are uniformly bounded by M . To test null $H_0 : \beta \geq 0$ against alternative $H_1 : \beta < 0$, we use the statistics $T = \sum_{i=1}^n X_i / \sqrt{n}$. With a large sample size, the p -value is estimated as $\hat{p} = \Phi(T)$. The $G(\cdot)$ function in Definition 5 can be chosen as $\Phi(\cdot)$, and $\Delta_G = 2M/\sqrt{n}$.

Example 2 Assume X_1, X_2, \dots, X_n are i.i.d. random variables with mean β , variance 1, and are uniformly bounded by M . To test null $H_0 : \beta = 0$ against alternative $H_1 : \beta \neq 0$, we also use the statistics $T = \sum_{i=1}^n X_i / \sqrt{n}$. With a large sample size, the p -value is estimated as $\hat{p} = 2\Phi(-|T|)$. Let $\text{supp}\{g(x)\} \subset [-M, M]$ with bounded density and $g(M) = g(-M) > 0$, $G(x) = \int_{-\infty}^x g(x)dx$. For example, $g(\cdot)$ can be the density of a truncated normal distribution supported on $[-M, M]$. Then $\Delta_G = 2MC/\sqrt{n}$, where C is a constant. Detailed proofs are provided in the appendix.

Example 3 Assume X_1, X_2, \dots, X_n are i.i.d. random variables with mean β , variance σ^2 , and are uniformly bounded by M . Let $h(X_1, X_2) = X_1X_2$ and

$$U_n = \binom{n}{2}^{-1} \sum_{i < j} h(X_i, X_j),$$

Then we have

$$\frac{nU_n}{\sigma^2} \xrightarrow{d} \chi_2^2 - 1,$$

as $n \rightarrow \infty$. In this case, the p -value for testing the null $H_0 : \beta = 0$ against the alternative $H_1 : \beta \neq 0$ is based on the statistics $T = nU_n/\sigma^2 + 1$. The $G(\cdot)$ function in Definition 5 can be chosen as $\Phi(\cdot)$ and the sensitivity is bounded by

$$\Delta_G = M^2/nC_1 + \frac{C_2}{1/2 - \delta} (M^2/n)^{1/2 - \delta},$$

where C_1 and C_2 are constants, and $0 < \delta < 1/2$. Detailed proofs are provided in the appendix.

The sensitivity in Example 1 is tight due to the nature of the one-sided test and normal transformation of the p -value. In Example 2, we used the truncated normal distribution to perform the transformation to simplify the technical calculation. In Example 3, the Δ_G is approximately the square root of the original sensitivity Δ because of the imperfect match between the tail of the transformation function (normal distribution) and the tail of the χ^2 distribution. In fact, the normal transformation, i.e., $G(\cdot) = \Phi(\cdot)$, works well for most cases, as we will show in the numerical studies.

3.3 The Differentially Private AdaPT Algorithm

In this section, we propose the DP-AdaPT algorithm. We begin the discussion by introducing the Gaussian mechanism. To report a statistic $\mathcal{T}(\mathcal{S})$ with the privacy guarantee, *Gaussian mechanism* adds noise to the target statistics $\mathcal{T}(\mathcal{S})$, with the scale of noise calibrated according to the sensitivity of $\mathcal{T}(\mathcal{S})$. We summarize some appealing properties of the Gaussian mechanism in Lemma 6.

Lemma 6 *The Gaussian Mechanism has the following properties (Dong et al., 2022):*

1. *GDP guarantee. The Gaussian mechanism that outputs $\mathcal{M}(\mathcal{T}) = \mathcal{T}(\mathcal{S}) + Z$ preserves μ -GDP, where Z is drawn independently from $\mathcal{N}(0, \Delta(\mathcal{T})^2/\mu^2)$ and $\Delta(\mathcal{T})$ is the sensitivity of \mathcal{T} defined in the Definition 4.*

2. *Composition.* Let \mathcal{M}_1 and \mathcal{M}_2 be two algorithms that are μ_1 -GDP and μ_2 -GDP, respectively. The composition algorithm $\mathcal{M}_1 \circ \mathcal{M}_2$ is $\sqrt{\mu_1^2 + \mu_2^2}$ -GDP.
3. *Post-processing.* Let $f(\cdot)$ be a deterministic function and \mathcal{M} be a μ -GDP algorithm. Then the post-processed algorithm $f \circ \mathcal{M}$ is μ -GDP.

In multiple testing, a common scenario is that the number of hypotheses is very large. If we report all the p -values under the private parameter μ , the standard deviation of the noise is proportional to the square root of the number of hypotheses by the composition lemma. Thus, in large-scale hypothesis testing, reporting all p -values adds very large noise to the signal and weakens the power of tests. To overcome the difficulty, the first step of our algorithm is to select a subset of p -values with more potential to be rejected, and the second step is to report the subset of p -values with a privacy guarantee. It is also common in real practice that the true signals are only a small subset of the total hypotheses. For example, only a few genes are truly related to the phenotype of interest. The following report noisy min algorithm (Dwork et al., 2021) builds the foundation of the selection algorithm.

Algorithm 2 The Report Noisy Min Algorithm

Require: data \mathcal{S} , p -values p_1, \dots, p_n each with sensitivity at most Δ , privacy parameter μ .

for $j = 1$ to n **do**

set $\tilde{f}_j = G[G^{-1}(p_j(\mathcal{S})) + Z_j]$, where Z_j is an independent sample from normal distribution with mean 0 and variance $8\Delta^2/\mu^2$;

end for

Output: $j^* = \arg \min_j \tilde{f}_j$ and $\tilde{p}_{j^*} := G[G^{-1}(p_{j^*}(\mathcal{S})) + Z]$, where Z is an independent, afresh drawn sample from the normal distribution with mean 0 and variance $8\Delta^2/\mu^2$.

Lemma 7 *The Report Noisy Min Algorithm, as detailed in Algorithm 2, is μ -GDP.*

A traditional way to select the most important signals is the peeling algorithm (Cai et al., 2021; Dwork et al., 2021), which repeats the report noisy min algorithm for a fixed number of times. However, the peeling algorithm creates complex dependent structures among the selected p -values and further complicates the analysis of FDR control. We believe it is the main issue in Dwork et al. (2021) that prevented the authors from bounding the classic FDR criterion and instead worked on the conditional quantity FDR_k , with $k \geq 2$.

In this paper, we propose a novel mirror peeling algorithm that perfectly suits the situation of adaptive FDR control. The selection procedure is based on the partially masked p -values and simultaneously selects both the largest and the smallest p -values. The largest p -values will be used to estimate the false discovery proportion as the control, which is a widely used technique in the multiple testing literature.

Lemma 8 *The Mirror Peeling Algorithm, as presented in Algorithm 3, is μ -GDP.*

The size m in the mirror peeling algorithm denotes the number of selected p -values. In practice, we suggest choosing a slightly large m to prevent potential power loss. Now we

Algorithm 3 The Mirror Peeling Algorithm

Require: p -values $p_1 \dots, p_n$ each with sensitivity at most Δ , privacy parameter μ , size m .

let $\mathcal{S} = \{1, \dots, n\}$ be the index set of p -values;

for $j = 1$ to m **do**

 let i_j be the returned index of the report noisy min algorithm applied to partially masked p -values

$$\{\min(p_i, 1 - p_i)\}_{i \in \mathcal{S}},$$

with privacy parameter μ/\sqrt{m} ;

 let \tilde{p}_{i_j} be the noisy p_{i_j} defined in (5) with Z following independent normal noise with mean 0 and variance $8m\Delta^2/\mu^2$;

 Update $\mathcal{S} = \mathcal{S} \setminus \{i_j\}$;

end for

Output: $\left\{ (i_1, \tilde{p}_{i_1}), \dots, (i_m, \tilde{p}_{i_m}) \right\}$.

are ready to state the DP-AdaPT procedure in Algorithm 4, which controls the FDR at a user-specified level α with guaranteed privacy. Theorem 9 follows directly by Lemma 7 and 8 and the post-processing property of GDP algorithms.

Algorithm 4 The DP-AdaPT Algorithm

Require: target FDR level α , $\{\mathbf{x}_i, p_i\}_{i=1}^n$, p -values with sensitivity at most Δ , privacy parameter μ and size m .

Apply the mirror peeling algorithm and obtain $\left\{ (i_1, \tilde{p}_{i_1}), \dots, (i_m, \tilde{p}_{i_m}) \right\}$.

for $t = 0, 1, \dots$ **do**

 Let $R_t = |\{i_j, j = 1, \dots, m : \tilde{p}_{i_j} \leq s_t(\mathbf{x}_{i_j})\}|$,

$A_t = |\{i_j, j = 1, \dots, m : \tilde{p}_{i_j} \geq 1 - s_t(\mathbf{x}_{i_j})\}|$

$$\widehat{\text{FDR}} \leftarrow \frac{1+A_t}{R_t \vee 1}$$

if $\widehat{\text{FDR}} \leq \alpha$

 Reject $\{H_{i_j} : \tilde{p}_{i_j} \leq s_t(\mathbf{x}_{i_j})\}$;

 Return s_t ;

end if

$s_{t+1} \leftarrow \text{Update}[(\mathbf{x}_{i_j}, \tilde{p}_{pm,i_j})_{j=1}^m, A_t, R_t, s_t]$, where

$$\tilde{p}_{pm,i_j} = \begin{cases} \tilde{p}_{i_j} & s_t(\mathbf{x}_i) < \tilde{p}_{i_j} < 1 - s_t(\mathbf{x}_i) \\ \{\tilde{p}_{i_j}, 1 - \tilde{p}_{i_j}\} & \text{otherwise.} \end{cases}$$

end for

Theorem 9 *The DP-AdaPT algorithm described in Algorithm 4 is μ -GDP.*

For completeness of the discussion, we provide a classic (ϵ, δ) -private version of the proposed DP-FDR control algorithm. By the relation between μ -GDP and (ϵ, δ) -DP, the Algorithm 4 is (ϵ, δ) -DP for $\forall \epsilon > 0$ and $\delta = \Phi(-\epsilon/\mu + \mu/2) - e^\epsilon \Phi(-\epsilon/\mu - \mu/2)$. With pre-specified private parameters ϵ and δ , we proposed a modification of Algorithm 4 which uses the Laplace mechanism in The Report Noisy Min Algorithm. Theorem 10 shows the proposed modified algorithm is (ϵ, δ) -DP.

Theorem 10 *Given private parameters $\epsilon \leq 0.5$, $\delta \leq 0.1$, the peeling size $m \geq 10$ and sensitivity at most Δ , the DP-AdaPT algorithm described in Algorithm 4 with Z_j and Z in Algorithm 2 following Laplace noise of scale $\lambda = \Delta\sqrt{10m \log(1/\delta)}/\epsilon$ is (ϵ, δ) -differentially private.*

Although the FDR control procedure of the proposed DP-AdaPT method is very different from the BH method used by Dwork et al. (2021), the privacy is protected by a similar procedure—the peeling mechanism. With the same sensitivity parameter Δ , the peeling size m , and privacy parameters (ϵ, δ) , our modified DP-AdaPT procedure uses the same level of noise as the DP-BH procedure proposed by Dwork et al. (2021). With the same noise level, our proposed method is superior to the DP-BH method in terms of the exact valid FDR control and the higher power of detecting the true hypothesis.

3.4 FDR Control

There are a few challenges in deriving the FDR bound for differentially private algorithms. Firstly, the privacy-preserving procedure is required to be randomized with noise independent of the data. However, most classic FDR procedures implement fixed thresholds to decide the rejection regions and are unsuitable for noisy or permuted private p -values. Secondly, the mirror peeling algorithm creates complicated dependence structures among the selected p -values. Classic tools in the literature that are used for proving FDR control crucially rely on the independence assumption or the positive dependence assumptions on the p -values, thus becoming inapplicable for differentially private algorithms. Thirdly, without the martingale technique (Storey et al., 2004), it is, in general, difficult to derive finite sample results with differential privacy. In fact, it is still unclear how to obtain valid finite sample FDR control for the differential private BH procedure. The DP-BH algorithm proposed in Dwork et al. (2021) addressed the challenges by conservatively bounding the noise and derived the upper bound for an unusual conditional version of FDR, i.e., $\text{FDR}_k := \mathbb{E}[V/R; V \geq k]$.

In this paper, we prove that the DP-AdaPT algorithm controls the FDR in finite samples. Our proof adopted the similar optional stopping argument in the multiple testing literature (Storey et al., 2004; Barber and Candès, 2015; Li and Barber, 2017; Lei and Fithian, 2018). We show that the adaptive procedure and the mirror conservative assumption work perfectly with the additional noise required to protect privacy. By only using partial information in the mirror peeling algorithm, we can construct a filtration and apply the martingale technique.

We first introduce the notations. For each hypothesis H_i , we observe p -value p_i and auxiliary information \mathbf{x}_i . Given a pre-specified sparsity level $m \leq n$, the DP-AdaPT algorithm first applies the mirror peeling algorithm, and we use \mathcal{P} to denote the index returned

by the mirror peeling algorithm. Let \mathcal{F}_t for $t = 0, \dots$, denote the filtration generated by all information available to the analyst at step t :

$$\mathcal{F}_t = \sigma(\{\mathbf{x}_i, \tilde{p}_{pm,t,i}\}_{i \in \mathcal{P}}, A_t, R_t)$$

where

$$\tilde{p}_{pm,t,i} = \begin{cases} \tilde{p}_i & s_t(\mathbf{x}_i) < \tilde{p}_i < 1 - s_t(\mathbf{x}_i) \\ \{\tilde{p}_i, 1 - \tilde{p}_i\} & \text{otherwise.} \end{cases}$$

The initial σ field is defined as $\mathcal{F}_{-1} = \sigma(\{\mathbf{x}_i, \{\tilde{p}_i, 1 - \tilde{p}_i\}\}_{i \in \mathcal{P}})$. The two updating thresholds principles: $s_{t+1} \preceq s_t$ and $s_{t+1} \in \mathcal{F}_t$, ensure that the $\{\mathcal{F}_t\}_{t=-1}^\infty$ is a filtration, i.e., $\mathcal{F}_t \subset \mathcal{F}_{t+1}$ for $t \geq -1$.

Theorem 11 *Assume that all the null p -values are independent of each other and of all the non-null p -values, and the null p -values are mirror-conservative. The DP-AdaPT procedure controls the FDR at level α .*

Theorem 11 has several important implications. First of all, it can control the FDR at a user-specified level α with a differential privacy guarantee while the existing DP-BH method fails. Secondly, due to the definition of filtration and the application of martingale tools, the DP-AdaPT can control the FDR for a finite number of tests. And lastly, when the side information is available to the hypothesis, the DP-AdaPT shares a similar property to the original AdaPT and can borrow the side information in a model-free sense to increase power. We demonstrate the numerical utilities in the next section.

3.5 The Two-groups Working Model and Selection Procedure

Our proposed DP-AdaPT procedure successfully controls FDR regardless of the strategy used in the threshold updating. In other words, it also enjoys the model-free property. But it is still important to provide a practical and powerful solution to update the threshold. Lei and Fithian (2018) proposed a two-group working model and used the local false discovery rate as the threshold. In this subsection, we illustrate a greedy procedure based on a two-group working model. Our procedure is similar to the method proposed by Lei and Fithian (2018) but has a simple illustration.

We begin with the working model specification. Assume that the distribution of hypothesis indicator H_i given side information \mathbf{x}_i follows Bernoulli distribution with probability $\pi(\mathbf{x}_i)$, $H_i \mid \mathbf{x}_i \sim \text{Bernoulli}(\pi(\mathbf{x}_i))$, where $H_i = 1$ if the i th hypothesis is true and $H_i = 0$ otherwise. The distribution of noisy p -value \tilde{p}_i given H_i and \mathbf{x}_i satisfies:

$$\tilde{p}_i \mid H_i, \mathbf{x}_i \sim \begin{cases} f_0(\tilde{p}_i \mid \mathbf{x}_i) & \text{if } H_i = 0 \\ f_1(\tilde{p}_i \mid \mathbf{x}_i) & \text{if } H_i = 1 \end{cases}.$$

In addition, we assume the data $\{(\mathbf{x}_i, \tilde{p}_i, H_i)\}_{i=1}^n$ is mutually independent with $\{H_i\}_{i=1}^n$ missing or being unobserved. The $\pi(\cdot)$, $f_0(\cdot \mid \mathbf{x})$ and $f_1(\cdot \mid \mathbf{x})$ are unknown functions and can be estimated by any user-specified methods. Lei and Fithian (2018) suggested using exponential families to model the $\pi(\cdot)$, $f_0(\cdot \mid \mathbf{x})$ and $f_1(\cdot \mid \mathbf{x})$. As mentioned by Lei and

Fithian (2018), the model $\{\pi(\cdot), f_0(\cdot | \mathbf{x}), f_1(\cdot | \mathbf{x})\}$ is not identifiable, and we use uniform distribution as the working model for the null hypothesis, $f_0(p | \mathbf{x}) = 1$ for $p \in [0, 1]$.

At the t -th iteration with available information \mathcal{F}_t , the first step is to fit the model using the data $\{\mathbf{x}_i, \tilde{p}_{pm,t,i}\}_{i \in \mathcal{P}}$. The complete log-likelihood at t -th iteration is

$$l_t(\pi(\cdot), f_1(\cdot)) = \sum_{i \in \mathcal{P}} [H_i \log(\pi(\mathbf{x}_i)) + (1 - H_i) \log(1 - \pi(\mathbf{x}_i))] + \sum_{i \in \mathcal{P}} [H_i \log(f_1(\tilde{p}_i | \mathbf{x}_i))], \quad (6)$$

where we use the fact $\log(f_0(\tilde{p}_i | \mathbf{x}_i)) = \log(1) = 0$. Because all H_i 's and parts of \tilde{p}_i 's are not observed, the Expectation-Maximization (EM) algorithm is an iterative algorithm to maximize the observed log-likelihood. For more information about missing data, see Chapter 3 in (Kim and Shao, 2021). We use \mathcal{T}_t to denote the index set, $\mathcal{T}_t := \{i \in \mathcal{P} : s_t(\mathbf{x}_i) < \tilde{p}_i < 1 - s_t(\mathbf{x}_i)\}$. The \tilde{p}_i is known for $i \in \mathcal{T}_t$ at the t -th iteration. The detailed procedure is shown in Algorithm 5.

Algorithm 5 The EM Algorithm

Require: data $\{\mathbf{x}_i, \tilde{p}_{pm,t,i}\}_{i \in \mathcal{P}}$, initial value $\{\hat{\pi}^{(0)}(\cdot), \hat{f}_1^{(0)}(\cdot | \mathbf{x})\}$, number of iteration k ;
for $r = 1, \dots, k - 1$ **do**

[E-step]: Calculate $\mathbb{E}[l_t(\hat{\pi}^{(r)}(\cdot), \hat{f}_1^{(r)}(\cdot | \mathbf{x})) | \{\mathbf{x}_i, \tilde{p}_{pm,t,i}\}_{i \in \mathcal{P}}]$;

For $i \in \mathcal{P}$ and $i \in \mathcal{T}_t$:

$$\hat{H}_i^{(r)} = \frac{\pi^{(r)}(\mathbf{x}_i) \hat{f}_1^{(r)}(\tilde{p}_i | \mathbf{x}_i)}{\pi^{(r)}(\mathbf{x}_i) \hat{f}_1^{(r)}(\tilde{p}_i | \mathbf{x}_i) + 1 - \pi^{(r)}(\mathbf{x}_i)};$$

$$\widehat{\log}(f_1(\tilde{p}_i | \mathbf{x}_i))^{(r)} = \log(f_1(\tilde{p}_i | \mathbf{x}_i));$$

For $i \in \mathcal{P}$ and $i \notin \mathcal{T}_t$:

$$\hat{H}_i^{(r)} = \frac{\pi^{(r)}(\mathbf{x}_i) [\hat{f}_1^{(r)}(\tilde{p}_i | \mathbf{x}_i) + \hat{f}_1^{(r)}(1 - \tilde{p}_i | \mathbf{x}_i)]}{\pi^{(r)}(\mathbf{x}_i) [\hat{f}_1^{(r)}(\tilde{p}_i | \mathbf{x}_i) + \hat{f}_1^{(r)}(1 - \tilde{p}_i | \mathbf{x}_i)] + 2[1 - \pi^{(r)}(\mathbf{x}_i)]};$$

$$\widehat{\log}(f_1(\tilde{p}_i | \mathbf{x}_i))^{(r)} = \frac{\hat{f}_1^{(r)}(\tilde{p}_i | \mathbf{x}_i) \log(f_1(\tilde{p}_i | \mathbf{x}_i)) + \hat{f}_1^{(r)}(1 - \tilde{p}_i | \mathbf{x}_i) \log(f_1(1 - \tilde{p}_i | \mathbf{x}_i))}{\hat{f}_1^{(r)}(\tilde{p}_i | \mathbf{x}_i) + \hat{f}_1^{(r)}(1 - \tilde{p}_i | \mathbf{x}_i)};$$

[M-step]: Solve

$$\begin{aligned} \{\hat{\pi}^{(r+1)}(\cdot), \hat{f}_1^{(r+1)}(\cdot | \mathbf{x})\} &= \arg \max \mathbb{E}[l_t(\hat{\pi}^{(r)}(\cdot), \hat{f}_1^{(r)}(\cdot | \mathbf{x})) | \{\mathbf{x}_i, \tilde{p}_{pm,t,i}\}_{i \in \mathcal{P}}] \\ &= \arg \max \sum_{i \in \mathcal{P}} [\hat{H}_i^{(r)} \log(\pi(\mathbf{x}_i)) + (1 - \hat{H}_i^{(r)}) \log(1 - \pi(\mathbf{x}_i))] \\ &\quad + \sum_{i \in \mathcal{P}} [\hat{H}_i^{(r)} \widehat{\log}(f_1(\tilde{p}_i | \mathbf{x}_i))^{(r)}]; \end{aligned}$$

end for

Output: $\{\hat{\pi}^{(k)}(\cdot), \hat{f}_1^{(k)}(\cdot | \mathbf{x})\}$.

At the t -th iteration with fitted model $\{\hat{\pi}^{(k)}(\cdot), \hat{f}_1^{(k)}(\cdot | \mathbf{x})\}$, the second step is to select one hypothesis from $\mathcal{T}_t^c := \{i \in \mathcal{P} : \tilde{p}_i \leq s_t(\mathbf{x}_i)\}$ and reject. The estimated probability of

$H = 0$ conditional on (\mathbf{x}, p) is

$$\mathbb{P}(H = 0 \mid \mathbf{x}, p, \{\hat{\pi}^{(k)}(\cdot), \hat{f}_1^{(k)}(\cdot \mid \mathbf{x})\}) = \frac{1 - \pi^{(k)}(\mathbf{x})}{\pi^{(k)}(\mathbf{x})\hat{f}_1^{(k)}(p \mid \mathbf{x}) + 1 - \pi^{(k)}(\mathbf{x})}. \quad (7)$$

We propose to select the hypothesis with the largest probability defined in equation (7) among the candidate set \mathcal{T}_t^c . Because all p -values in the candidate set are partially masked, we use the minimum elements in each pair and let $\tilde{p}'_{pm,t,i} = \min\{\tilde{p}_i, 1 - \tilde{p}_i\}$ for $i \in \mathcal{T}_t^c$. As a consequence, we reject the i -th hypothesis for $i \in \mathcal{T}_t^c$ satisfying

$$i = \arg \max_{j \in \mathcal{T}_t^c} \mathbb{P}(H = 0 \mid \mathbf{x}_i, \tilde{p}'_{pm,t,i}, \{\hat{\pi}^{(k)}(\cdot), \hat{f}_1^{(k)}(\cdot \mid \mathbf{x})\}). \quad (8)$$

We remark that the proposed selection criterion (8) is slightly different from the criterion in the AdaPT procedure by Lei and Fithian (2018). Under the conservative identifying assumption that

$$\begin{aligned} 1 - \pi(\mathbf{x}) &= \inf_{p \in [0,1]} [\pi(\mathbf{x})f_1(p \mid \mathbf{x}) + (1 - \pi(\mathbf{x}))f_0(p \mid \mathbf{x})] \\ &= \pi(\mathbf{x})f_1(1 \mid \mathbf{x}) + (1 - \pi(\mathbf{x}))f_0(1 \mid \mathbf{x}), \end{aligned}$$

the proposed selection criterion coincides with equation (23) in (Lei and Fithian, 2018).

4. Numerical Illustrations

In this section, we numerically evaluate the performance of the proposed DP-AdaPT in terms of false discovery rate and power. We compare with three other methods: the original AdaPT without privacy guarantee as proposed by Lei and Fithian (2018), the differentially private Benjamini–Hochberg procedure (“DP-BH”) proposed by Dwork et al. (2021), and the private Bonferroni’s method (“DP-Bonf”) as discussed in Dwork et al. (2021).

4.1 Without Side Information

We first consider the case where only the p -values are obtained for each hypothesis and side information is unavailable. To ensure a fair comparison, we adopt the same simulation settings as in Dwork et al. (2021) and apply the noises with the same variance for all the differentially private methods. Specifically, we set the total number of hypotheses to be $n = 100\,000$, with the number of true effects $t = 100$. We select $m = 500$ in the peeling step. Let $p_i = \Phi(\xi_i - \beta)$ for $i = 1, \dots, t$, where $\Phi(\cdot)$ is the CDF of standard normal distribution and ξ_1, \dots, ξ_m are i.i.d. standard normal distribution. We set the signal β to be 4 and the significance level $\alpha = 0.1$. Other parameters are required for the DP-BH algorithm, which is summarized in Algorithm 6. Two parameters are used to control the sensitivity of the p -values in the DP-BH algorithm: the multiplicative sensitivity η and the truncation threshold ν . We set η as 0,0001 and $\nu = 0.5\alpha/n$, which are the same as Dwork et al. (2021). The privacy parameters are also set to be the same as in Dwork et al. (2021): $\epsilon = 0.5$ and $\delta = 0,001$. For our proposed DP-AdaPT procedure, we set the privacy parameter $\mu = 4\epsilon/\sqrt{10\log(1/\delta)}$. We use Gaussian CDF as the sensitivity transformation

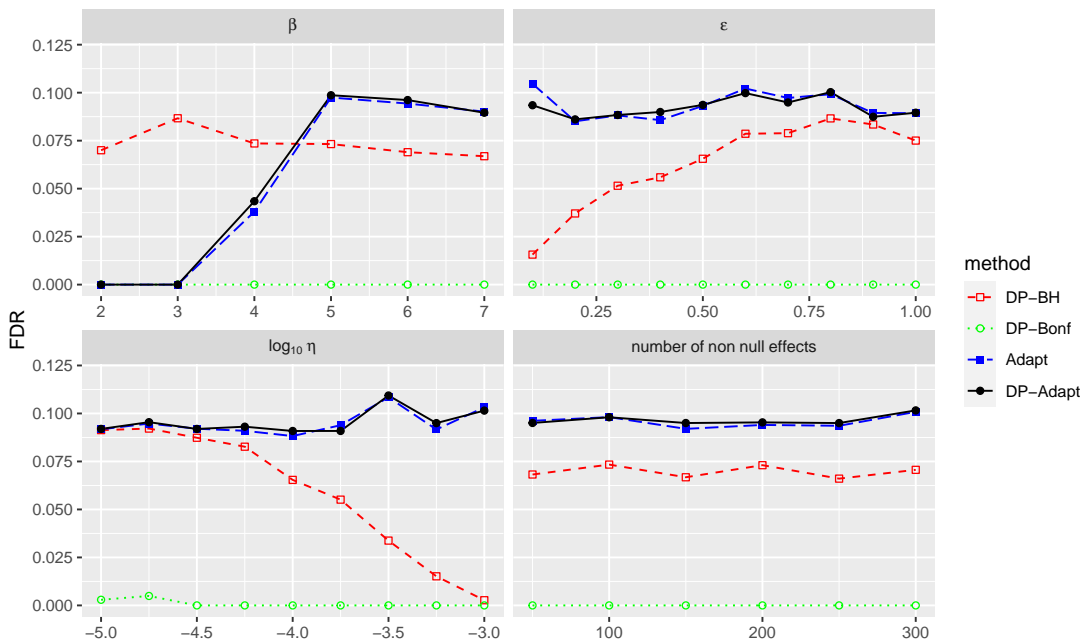


Figure 2: The FDR of DP-BH, DP-Bonf, and DP-AdaPT for varying parameters and averaged over 100 independent trials.

$G(\cdot) = \Phi(\cdot)$ and set the sensitivity parameter $\Delta_G = \eta$. The variance of the noise in our proposed DP-AdaPT procedure is the same as the variance of the noise in DP-BH.

We first consider the situation where the null p -values all follow an independent uniform distribution. Specifically, we generate p_i for $i = t + 1, \dots, n$ independently from $U(0, 1)$. In Figure 2, we report the empirical FDR control for all the methods by varying the signal parameter β , the privacy parameter ϵ , the sensitivity parameter η and the number of true effects t . Clearly, all methods successfully control the FDR below the specified level $\alpha = 0.1$. We report the power of all the methods in Figure 3. The naive DP-Bonf method is too conservative in detecting any positive signals. The power of DP-BH and DP-AdaPT performs similarly to each other in most cases. When the sensitivity parameter η is large, the proposed method has better power than the DP-BH procedure. The rationale is that when η is large, the variance of the noise is large, and the correction term in Algorithm 6, i.e., $\eta\sqrt{10m \log(1/\delta)} \log(6m/\alpha)/\epsilon$, is large. The correction term plays the role of ruling out the influence of adding noise and guarantees FDR control with high probability. However, the additional correction term is proportional to the standard deviation of noise multiplying $\log(m)$, where m is the number of p -values selected by peeling. This is the main weakness of the DP-BH procedure. On the other hand, the proposed DP-AdaPT is based on the symmetry of p -values and provides valid finite sample FDR control, which is more robust both theoretically and practically.

Next, we consider the situation where the null p -values follow conservative distributions compared to the uniform, which is a common phenomenon in practice. Specifically, we

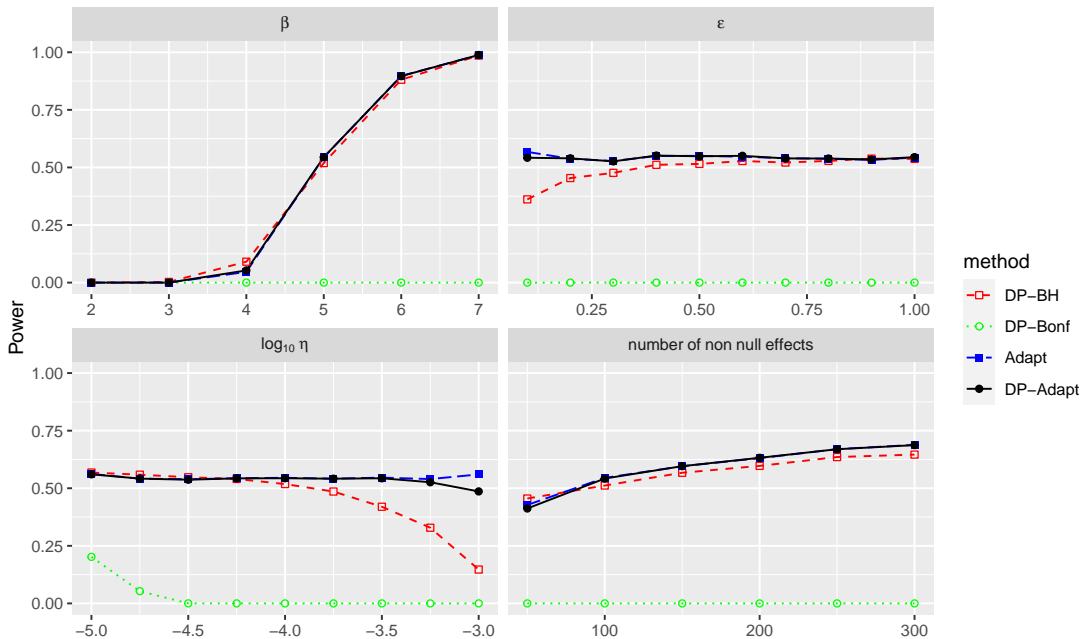


Figure 3: The power of DP-BH, DP-Bonf, and DP-AdaPT for varying parameters and averaged over 100 independent trials.

generate p_i for $i = t + 1, \dots, n$ independently from the Beta distribution with shape parameters $(2, 2)$. The empirical FDR and power of all the methods are summarized in Figure 4 and Figure 5. All methods successfully control the FDR below the $\alpha = 0.1$. DP-BH and DP-Bonf have nearly zero false discovery rates. Though Dwork et al. (2021) only have a theoretical proof for FDR control when the p -value of null hypotheses follows the uniform distribution. It is not surprising that the DP-BH procedure controls FDR at the target level because the false discovery rates of conservative null hypotheses are easier to control than non-conservative null hypotheses in principle. The power of our proposed method is uniformly better than the DP-BH procedure. In general, our proposed method has power close to 0.90 when the number of true effects is smaller than the number of invocations, and the signal size is reasonably strong.

4.2 With Side Information

In this subsection, we consider the case where the auxiliary side information is available for the hypothesis. We use similar simulation settings as in Lei and Fithian (2018). The auxiliary covariates (x_{1i}, x_{2i}) 's are generated from an equispaced 100×100 grid in the area $[-100, 100] \times [-100, 100]$. The p -values are generated i.i.d. from $p_i = 1 - \Phi(z_i)$, where $z_i \sim N(\mu_i, 1)$ and $\Phi(\cdot)$ is the CDF of $N(0, 1)$. For $i \in \mathcal{H}_0$, we set $\beta_i = 0$, and for $i \in \mathcal{H}_1$, we




Figure 4: The FDR of DP-BH, DP-Bonf, and DP-AdaPT for varying parameters and averaged over 100 independent trials under conservative p -values.

set $\beta_i = \beta$ for $\beta > 0$. Three different patterns of $i \in \mathcal{H}_1$ are considered.

$$i \in \mathcal{H}_1 \Leftrightarrow \begin{cases} x_{1i}^2 + x_{2i}^2 \leq 150 & \text{I} \\ (x_{1i} - 65)^2 + (x_{2i} - 65)^2 \leq 150 & \text{II} \\ 2(x_{1i} + x_{2i})^2/100^2 + (x_{2i} - x_{1i})^2/15^2 \leq 0.1 & \text{III} \end{cases}$$

The number of the true signals \mathcal{H}_1 are 120, 116 and 118 for case 1, case 2, and case 3, respectively. The number of selections in the peeling algorithm is set to $m = 500$. For our proposed DP-AdaPT procedure, we use Gaussian CDF as the sensitivity transformation $G(\cdot) = \Phi(\cdot)$ and set the sensitivity parameter $\Delta_G = 0.0001$. We set the privacy parameter $\mu = 0.24$, which matches the noise scale in Dwork et al. (2021). We compare the DP-AdaPT with the DP-BH method to evaluate the advantage of using side information. We also compare our proposed DP-AdaPT with the non-private AdaPT procedure to examine the privacy and accuracy tradeoff. For the AdaPT procedure, we follow a similar algorithm as in Lei and Fithian (2018) and fit two-dimensional Generalized Additive Models in M-step, using R package *mgcv* with the knots selected automatically in every step by GCV criterion. The procedure replicates 100 times, and the results are reported in terms of the average of 100 replications.

Figure 6 and Figure 7 present the empirical FDR and power of AdaPT, DP-BH, and DP-AdaPT, respectively. All methods can control the FDR at the desired level. The proposed DP-AdaPT procedure has larger power than the DP-BH procedure when the target FDR is greater than 0.03. Especially when the strength of signals is not very strong,




Figure 5: The power of DP-BH, DP-Bonf, and DP-AdaPT for varying parameters and averaged over 100 independent trials under conservative p -values.

our proposed method has more than 50% larger power than the DP-BH procedure. When the strength of signals is large enough, all procedures have similar powers, and AdaPT is slightly better than others. Compared to the original AdaPT, the proposed DP-AdaPT has uniformly smaller power, which is due to noise added for privacy guarantee. To be precise, the proposed procedure contains two pre-processing steps, peeling and adding noise. In the peeling algorithm, a subset of the original hypothesis is selected for further consideration. The selection procedure has two intrinsic weaknesses. First, some important variables can be potentially ignored due to random errors. Second, only 5% data are used in the model fitting procedure of the DP-AdaPT. In the simulation settings, the side information is perfect for separating null hypotheses and alternative hypotheses. Adding noise to the p -values after selection attenuates the strength of a valid signal, which bring difficulties in rejecting the null hypothesis. To sum up, it is not surprising that DP-AdaPT is less efficient than AdaPT. However, the loss is relatively mild and is the cost of privacy.

We also consider the case where the null p -values are conservative. We modify the previous settings and generate p_i for $i \in \mathcal{H}_0$ from density function $f(p) = 4p^3$, $p \in [0, 1]$. Figure 8 and Figure 9 present the empirical FDR and power of AdaPT, DP-BH, and DP-AdaPT, respectively. The distribution of null p -values is very conservative, and thus all methods successfully control the FDR. The proposed DP-AdaPT procedure has larger power than the DP-BH procedure when the target FDR is greater than 0.05. Especially when the strength of the signal is not very strong, our proposed method is around 70% better than the DP-BH procedure. Compared to the AdaPT procedure, the proposed DP-AdaPT has




Figure 6: The FDR of DP-BH, AdaPT, and DP-AdaPT for varying signal strength and averaged over 100 independent trials under uniform p -values.




Figure 7: The power of DP-BH, AdaPT and DP-AdaPT for varying signal strength and averaged over 100 independent trials under uniform p -values.

uniformly smaller power. When the strength of signals is reasonably strong, the difference between the power of our proposed DP-AdaPT and the power of AdaPT is smaller when the null is conservative. The power of the AdaPT procedure in the conservative case is smaller than that in the uniform case. Because the ratio of the density of p -values under

the null hypothesis near 1 and 0 is very large when the density function is $f(p) = 4p^3$, and thus the AdaPT is too conservative. However, the DP-AdaPT procedure uses noise to stable the ratio of the density of noisy p -values under the null hypothesis near 1 and 0, as illustrated in Figure 1. Thus, our proposed procedure is more robust to the distribution of p -values.

Although DP-AdaPT is not as accurate as the original AdaPT, the computation time of the proposed DP-AdaPT is much faster due to the mirror peeling algorithm. The model updating in AdaPT is computationally costly and should be performed every several steps. Thus when the dataset is large, the computation of AdaPT can be very slow. We provide the computation time in Table 1. The computation time is based on an HPC cluster with CPU Model Intel Xeon Gold 6152 and RAM 10 GB. From the table, the computation of DP-AdaPT is much faster than AdaPT because the data size of DP-AdaPT is only 5% of the data size of AdaPT. This provides a way to accelerate the computation of AdaPT in practice: when applying the AdaPT procedure to massive datasets, the proposed algorithm can be utilized to reduce the computation cost by applying noiseless screening of masked p -values.




Figure 8: The FDR of DP-BH, AdaPT, and DP-AdaPT for varying signal strength and averaged over 100 independent trials under conservative p -values.

4.3 Empirical Example

The Bottomly data set is an RNA-seq data set collected by Bottomly et al. (2011) to detect differential striatal gene expression between the C57BL/6J (B6) and DBA/2J (D2) inbred mouse strains. An average of 22 million short sequencing reads were generated per sample for 21 samples (10 B6 and 11 D2). The Bottomly data was analyzed by Ignatiadis et al. (2016) using *DESeq2* package (Love et al., 2014) and performing an independent hypothesis weighting (IHW) method. Approximate Wald test statistics were used to calculate p -values.




Figure 9: The power of DP-BH, AdaPT, and DP-AdaPT for varying signal strength and averaged over 100 independent trials under conservative p -values.

	$\beta = 2.5$		$\beta = 3.5$		$\beta = 4.5$	
Uniform Null	AdapT	DP-AdaPT	AdapT	DP-AdaPT	AdapT	DP-AdaPT
I	1254.47	18.22	1100.75	20.06	1043.71	23.50
II	1512.54	23.82	1314.59	23.01	1215.09	25.58
III	1515.57	25.42	1090.52	16.74	1165.79	22.62
Conservative Null	AdapT	DP-AdaPT	AdapT	DP-AdaPT	AdapT	DP-AdaPT
I	1810.04	30.34	1638.76	28.60	1662.77	26.49
II	1798.20	26.85	1508.15	32.04	1770.05	29.97
III	1921.82	31.59	1717.81	31.23	1675.85	25.46

Table 1: Computation Time in seconds averaged over 100 replications

The logarithm of each gene’s average normalized counts (across samples) is used as auxiliary information (Lei and Fithian, 2018). After removing missing records, the data set contains $n = 13932$ genes (p -values).

We pre-processed the data set by the *DESeq2* package and used our proposed DP-AdaPT to identify differentially expressed genes while controlling the FDR at a pre-set level $\alpha = 0.01, \dots, 0.1$. Because p -values are calculated based on Wald statistics, we used the standard normal CDF $G(\cdot) = \Phi(\cdot)$ as illustrated in Example 1. The privacy parameter was $\mu = 0.25$. The same privacy parameter was used by Avella-Medina et al. (2021). The sensitive parameter is $\Delta = 3 \times 10^{-5}$, which is roughly the inverse of the square root of the total sample size. The corresponding standard deviation of the added noise is $\sqrt{2} \sqrt{8m\Delta^2/\mu^2} \approx 0.024$. The number of pre-selected hypotheses in the peeling algorithm was $m = 2500$, which is

18% of the total number of hypotheses. We compare our proposed DP-AdaPT method with the non-private Adapt and private DP-BH methods. The results are shown in Figure 10.




Figure 10: Bottomly dataset: number of rejections of DP-BH, AdaPT, and DP-AdaPT for varying FDR

Our proposed DP-AdaPT and Adapt methods have significantly more discoveries than the DP-BH method because the auxiliary variable is highly correlated with hypotheses. The proposed DP-AdaPT procedure complies with the Adapt procedure, which matches the illustration in the simulation.

5. Discussion

In this paper, we propose a differentially private FDR control algorithm that accurately controls the FDR at a user-specified level with a privacy guarantee. The proposed algorithm is based on adding noise to a transformed p -value, and preserves the mirror conservative property of the p -values. By further integrating a mirror peeling algorithm, we can define a nice filtration and apply the classic optimal stopping argument. Our analysis provides a new perspective in analyzing differentially private statistical algorithms: instead of conservatively controlling the noise added for privacy, researchers can design transformations or modifications that can preserve the key desired property in the theoretical analysis.

There are several open problems left for future research. For example, one commonly assumed conditions on the null p -values in multiple testing is the *uniformly conservative* property: $\mathbb{P}(p_i/\tau \leq t \mid p_i \leq \tau) \leq t, \forall t, \tau \in (0, 1)$. However, it is still unknown whether a similar transformation can be applied. If given a positive answer, then many existing multiple testing procedures can be easily extended to the differentially private version. It is an open question whether *mirror conservative* or *uniformly conservative* assumption is necessary for differentially private FDR control. The other commonly assumed conditions is the *super uniform* property: $\mathbb{P}(p_i \leq t) \leq t, \forall t \in [0, 1]$. Developing a differentially private

algorithm with finite sample FDR control under a *super uniform* framework remains an open problem. Lastly, it is also interesting to further develop other differentially private modern statistical estimation and inference tools based on the similar idea of designing transformations. Examples include but are not limited to sliced inverse regression (Cai et al., 2020), covariance estimation (Cai et al., 2022b), test of independence (Cai et al., 2023), etc.

Acknowledgments and Disclosure of Funding

The authors thank Jing Lei, Lihua Lei, and Linjun Zhang for helpful discussions and suggestions. Zhanrui Cai’s research is supported by the Seed Fund for Basic Research for New Staff at the University of Hong Kong.

Appendix A.

A. Algorithm of the Private BHq procedure

Algorithm 6 The Private BHq procedure (Dwork et al., 2021)

Require: dataset \mathcal{S} , p -values p_1, \dots, p_n , threshold $\nu > 0$, multiplicative sensitivity η , significant level α , privacy level (ϵ, δ) , number of invocations m and Laplace noise scale $\lambda = \eta\sqrt{10m \log(1/\delta)}/\epsilon$

Output: a set of up to m rejected hypotheses

Apply the transformation and denote $f_i = \log \max\{\nu, p_i(\mathcal{S})\}$ for $i = 1, \dots, n$.

Perform the Peeling Mechanism and Report Noisy Min algorithm with Laplace noise with scale at λ . The results are denoted by $\left\{ (i_1, \tilde{f}_{i_1}), \dots, (i_{m'}, \tilde{f}_{i_{m'}}) \right\}$.

for $j = m$ to 1 **do**

if $\tilde{f}_{i_j} > \log(\alpha j/n) - \eta\sqrt{10m \log(1/\delta)} \log(6m/\alpha)/\epsilon$
 continue

else

 reject p_{i_1}, \dots, p_{i_j} and halt

end if

end for

B. Proof of Theorem 3

Proof Without loss of generality, we assume the bound $U = \infty$ and the measure μ is the Lebesgue measure. All other cases can be proved using the same technique by using a suitable measure.

We first observe that the random variable \tilde{p} is supported in $[0, 1]$. For any $0 \leq a_1 \leq a_2 \leq 0.5$, it remains to show the mirror-conservative condition,

$$\mathbb{P}(\tilde{p} \in [a_1, a_2]) \leq \mathbb{P}(\tilde{p} \in [1 - a_2, 1 - a_1]).$$

By the definition, $\tilde{p} = G\{G^{-1}(p) + Z\}$, the mirror-conservative condition is equivalent to $\mathbb{P}(G^{-1}(p) + Z \in [G^{-1}(a_1), G^{-1}(a_2)]) \leq \mathbb{P}(G^{-1}(p) + Z \in [G^{-1}(1 - a_2), G^{-1}(1 - a_1)])$, where we substitute \tilde{p} in the equation and apply the function $G^{-1}(\cdot)$ to both sides. By the symmetry of $g(\cdot)$, we have for $x < 0$,

$$G(x) = \int_{-\infty}^x g(t) dt = \int_{-\infty}^x g(-t) dt = \int_{-x}^{\infty} g(t) dt = 1 - \int_{-\infty}^{-x} g(t) dt = 1 - G(-x),$$

where we use the fact $g(x) = g(-x)$ and $\int_{-\infty}^{\infty} g(x) = 1$. After subtracting $1/2$ on both sides, we have $1/2 - G(x) = -1/2 + G(-x)$. Noticing that $G(x) < 1/2$ for $x < 0$, we have

$$G^{-1}(1/2 - t) = -G^{-1}(1/2 + t),$$

where $t = G(x) \in (0, 1/2)$ and we apply $G^{-1}(\cdot)$ to both sides. It follows that $G^{-1}(1 - a_1) = -G^{-1}(a_1)$ and $G^{-1}(1 - a_2) = -G^{-1}(a_2)$ by substituting t by $1/2 - a_1$ and $1/2 - a_2$, respectively. The mirror-conservative condition is equivalent to

$$\mathbb{P}(G^{-1}(p) + Z \in [G^{-1}(a_1), G^{-1}(a_2)]) \leq \mathbb{P}(G^{-1}(p) + Z \in [-G^{-1}(a_2), -G^{-1}(a_1)]),$$

where we use the relation $G^{-1}(1 - a_1) = -G^{-1}(a_1)$ and $G^{-1}(1 - a_2) = -G^{-1}(a_2)$.

To prove the mirror-conservative condition, we first perform a decomposition of the distribution of p and use the convolution formula. Let μ_p be the corresponding probability measure generated by the distribution of p . We define a symmetric measure μ_s on $[0, 1]$, which satisfies that $\mu_s(a, b) = \mu(a, b)$ for $0 \leq a \leq b \leq 1/2$ and $\mu_s(a, b) = \mu(1 - b, 1 - a)$ for $1/2 \leq a \leq b \leq 1$. Carathéodory's extension theorem guarantees the existence of μ_s . Intuitively, the μ_s is the mirror-symmetric part of μ_p . We have the decomposition $\mu_p = \mu_s + \mu_r$, where μ_r is defined by $\mu_r := \mu_p - \mu_s$. The μ_r is a positive measure by the fact that μ_p is mirror-conservative. Let μ_z be the probability measure generated by the standard normal distribution. It follows that

$$\begin{aligned} & \mathbb{P}(G^{-1}(p) + Z \in [G^{-1}(a_1), G^{-1}(a_2)]) - \mathbb{P}(G^{-1}(p) + Z \in [-G^{-1}(a_2), -G^{-1}(a_1)]) \\ &= (\mu_p \circ G) \otimes \mu_z [G^{-1}(a_1), G^{-1}(a_2)] - (\mu_p \circ G) \otimes \mu_z [-G^{-1}(a_2), -G^{-1}(a_1)] \\ &= (\mu_s \circ G) \otimes \mu_z [G^{-1}(a_1), G^{-1}(a_2)] - (\mu_s \circ G) \otimes \mu_z [-G^{-1}(a_2), -G^{-1}(a_1)] \\ &+ (\mu_r \circ G) \otimes \mu_z [G^{-1}(a_1), G^{-1}(a_2)] - (\mu_r \circ G) \otimes \mu_z [-G^{-1}(a_2), -G^{-1}(a_1)], \end{aligned}$$

where $\mu_p \circ G$ is the probability measure generated by the random variable $G^{-1}(p)$ and \otimes is the convolution operator.

We consider the first line in the above equation. For all $-\infty \leq a \leq b \leq 0$, using the fact that the two intervals $[G(a), G(b)]$ and $[G(-b), G(-a)] = [1 - G(b), 1 - G(a)]$ are symmetric around $1/2$ and the measure μ_s is symmetric around $1/2$, we know $\mu_s \circ G[a, b] = \mu_s \circ G[-b, -a]$. Thus, the measure $\mu_s \circ G$ is symmetric around 0. By the fact that μ_z is symmetric around 0, the convolution $(\mu_s \circ G) \otimes \mu_z$ is also symmetric around 0. Thus, we have

$$(\mu_s \circ G) \otimes \mu_z [G^{-1}(a_1), G^{-1}(a_2)] = (\mu_s \circ G) \otimes \mu_z [-G^{-1}(a_2), -G^{-1}(a_1)].$$

It remains to consider the second line of the equation,

$$(\mu_r \circ G) \otimes \mu_z [G^{-1}(a_1), G^{-1}(a_2)] - (\mu_r \circ G) \otimes \mu_z [-G^{-1}(a_2), -G^{-1}(a_1)].$$

By the decomposition of μ_p , we have $\mu_r(a, b) = 0$ for $0 \leq a \leq b \leq 1/2$. By the fact that $G(x) \leq 1/2$ for $x \leq 0$, we conclude that $(\mu_r \circ G)[a, b] = 0$ for all $a, b \leq 0$. Let $\phi(\cdot) = d\mu_z/d\mu$ be the probability density function of standard normal distribution, and we have

$$\begin{aligned} & (\mu_r \circ G) \otimes \mu_z [G^{-1}(a_1), G^{-1}(a_2)] - (\mu_r \circ G) \otimes \mu_z [-G^{-1}(a_2), -G^{-1}(a_1)] \\ &= \int_{-\infty}^{\infty} 1 \left\{ x \in [G^{-1}(a_1), G^{-1}(a_2)] \right\} \int_{-\infty}^{\infty} \phi(x-y) d(\mu_r \circ G)(y) d\mu(x) \\ & \quad - \int_{-\infty}^{\infty} 1 \left\{ x \in [-G^{-1}(a_2), -G^{-1}(a_1)] \right\} \int_{-\infty}^{\infty} \phi(x-y) d(\mu_r \circ G)(y) d\mu(x) \\ &= \int_{G^{-1}(a_1)}^{G^{-1}(a_2)} \int_0^{\infty} \phi(x-y) d(\mu_r \circ G)(y) dx - \int_{-G^{-1}(a_2)}^{-G^{-1}(a_1)} \int_0^{\infty} \phi(x-y) d(\mu_r \circ G)(y) d\mu(x) \\ &= \int_{G^{-1}(a_1)}^{G^{-1}(a_2)} \int_0^{\infty} \phi(x-y) d(\mu_r \circ G)(y) dx - \int_{G^{-1}(a_1)}^{G^{-1}(a_2)} \int_0^{\infty} \phi(-x-y) d(\mu_r \circ G)(y) d\mu(x) \\ &= \int_{G^{-1}(a_1)}^{G^{-1}(a_2)} \int_0^{\infty} [\phi(x-y) - \phi(-x-y)] d(\mu_r \circ G)(y) d\mu(x) \\ &= - \int_{G^{-1}(a_1)}^{G^{-1}(a_2)} \int_0^{-2x} \phi(-x-y) d(\mu_r \circ G)(y) d\mu(x). \end{aligned}$$

Remembering that $G^{-1}(a_1), G^{-1}(a_2) < 0$, we have $x < 0$. We conclude

$$- \int_{G^{-1}(a_1)}^{G^{-1}(a_2)} \int_0^{-2x} \phi(-x-y) d(\mu_r \circ G)(y) dx < 0.$$

So,

$$\mathbb{P}(\tilde{p} \in [a_1, a_2]) \leq \mathbb{P}(\tilde{p} \in [1 - a_2, 1 - a_1]),$$

and \tilde{p} is mirror conservative. ■

C. Proof of Example 2 and 3

C.1 Example 2

Proof For two-sided testing, we let the support $\text{supp}\{g(x)\} \subset [-M, M]$ and $g(M) = g(-M) > 0$, and $G(x) = \int_{-\infty}^x g(x) dx$. We have

$$\begin{aligned} & G^{-1}\{2\Phi[-|T(\mathcal{S}')|]\} - G^{-1}\{2\Phi[-|T(\mathcal{S})|]\} \\ &= \int_{-|T(\mathcal{S})|}^{-|T(\mathcal{S}')|} \frac{dG^{-1}[2\Phi(t)]}{dt} dt = \int_{-|T(\mathcal{S})|}^{-|T(\mathcal{S}')|} \frac{dG^{-1}[2\Phi(t)]}{d2\Phi(t)} \frac{d2\Phi(t)}{dt} dt \\ &= \int_{-|T(\mathcal{S})|}^{-|T(\mathcal{S}')|} \frac{2\phi(t)}{g(G^{-1}(2\Phi(t)))} dt. \end{aligned}$$

The function $2\phi(t)/g(G^{-1}(2\Phi(t)))$ is continuous on $(-\infty, 0)$. As $t \rightarrow -\infty$,

$$\lim_{t \rightarrow -\infty} \frac{2\phi(t)}{g(G^{-1}(2\Phi(t)))} = \sqrt{\frac{\pi}{2}} \lim_{t \rightarrow -\infty} \frac{\exp(-t^2/2)}{g(-M)} = 0,$$

where we used L'Hospital's rule. As $t \rightarrow 0^-$,

$$\lim_{t \rightarrow 0^-} \frac{2\phi(t)}{g(G^{-1}(2\Phi(t)))} = \lim_{t \rightarrow 0^-} \frac{2\phi(0)}{g(G^{-1}(2\Phi(t)))} = \frac{2\phi(0)}{g(M)},$$

where we used L'Hospital's rule again. Thus, the function is bounded by a constant C , and

$$\begin{aligned} & |G^{-1}\{2\Phi[-|T(\mathcal{S})|]\} - G^{-1}\{2\Phi[-|T(\mathcal{S}')|]\}| \leq \left| \int_{-|T(\mathcal{S}')|}^{-|T(\mathcal{S})|} C \right| \\ & \leq C \left| |T(\mathcal{S})| - |T(\mathcal{S}')| \right| \leq C \|T(\mathcal{S}) - T(\mathcal{S}')\| \leq 2CM/\sqrt{n}, \end{aligned}$$

where we use the fact $\sup \|T(\mathcal{S}') - T(\mathcal{S})\| \leq M/\sqrt{n}$ by the sensitivity of statistic T . \blacksquare

C.2 Example 3

Proof We let $G(\cdot) = \Phi(\cdot)$, $\xi(x) = (2\pi)^{-1/2} x^{-1/2} e^{-x/2}$ be the density function of χ_1^2 distribution and Ξ be the CDF of χ_1^2 distribution. We have

$$\begin{aligned} & \Phi^{-1}\{1 - \Xi[T(\mathcal{S})]\} - \Phi^{-1}\{1 - \Xi[T(\mathcal{S}')]\} \\ &= \int_{T(\mathcal{S})}^{T(\mathcal{S}')} \frac{d\Phi^{-1}[1 - \Xi(t)]}{dt} dt = \int_{T(\mathcal{S})}^{T(\mathcal{S}')} \frac{d\Phi^{-1}[1 - \Xi(t)]}{d[1 - \Xi(t)]} \frac{d[1 - \Xi(t)]}{dt} dt \\ &= \int_{T(\mathcal{S}')}^{T(\mathcal{S})} \frac{\xi(t)}{\phi(\Phi^{-1}(1 - \Xi(t)))} dt. \end{aligned}$$

The function $\xi(t)/\phi(\Phi^{-1}(1 - \Xi(t)))$ is continuous on $(0, \infty)$. We will use L'Hospital's rule to evaluate the limiting behavior of the function. The derivative of the denominator is,

$$\begin{aligned} & \frac{d\phi(\Phi^{-1}(1 - \Xi(t)))}{dt} = \frac{d\phi(\Phi^{-1}(1 - \Xi(t)))}{d\Phi^{-1}(1 - \Xi(t))} \frac{d\Phi^{-1}(1 - \Xi(t))}{d(1 - \Xi(t))} \frac{d(1 - \Xi(t))}{dt} \\ &= -\Phi^{-1}(1 - \Xi(t)) \phi(\Phi^{-1}(1 - \Xi(t))) \times \frac{1}{\phi(\Phi^{-1}(1 - \Xi(t)))} \times [-\xi(t)] \\ &= \Phi^{-1}(1 - \Xi(t)) \xi(t), \end{aligned}$$

where we use the fact

$$\frac{d\phi(t)}{dt} = -t \frac{\exp(-t^2/2)}{\sqrt{2\pi}} = -t\phi(t).$$

As $t \rightarrow \infty$,

$$\begin{aligned} & \lim_{t \rightarrow \infty} \frac{\xi(t)}{\phi(\Phi^{-1}(1 - \Xi(t)))} = \lim_{t \rightarrow \infty} \frac{1}{\sqrt{2\pi}} \frac{-\frac{1}{2}t^{-3/2}e^{-t/2} - \frac{1}{2}t^{-1/2}e^{-t/2}}{\Phi^{-1}(1 - \Xi(t)) \xi(t)} \\ &= \lim_{t \rightarrow \infty} -\frac{1}{2} \frac{\frac{1}{t} + 1}{\Phi^{-1}(1 - \Xi(t))} \stackrel{u:=1-\Xi(t)}{=} \lim_{u \rightarrow 0^+} -\frac{1}{2} \frac{1}{\Phi^{-1}(u)} = 0. \end{aligned}$$

As $t \rightarrow 0^+$, the function $\xi(t)/\phi(\Phi^{-1}(1-\Xi(t)))$ diverges, which means we can not bound the sensitivity using the method in Example 2. We investigate the limiting behavior of the function $\xi(t)/\phi(\Phi^{-1}(1-\Xi(t)))$ by multiplying a power of t . For $0 < \delta < 1/2$,

$$\begin{aligned} \lim_{t \rightarrow 0^+} t^{1/2+\delta} \frac{\xi(t)}{\phi(\Phi^{-1}(1-\Xi(t)))} &= \lim_{t \rightarrow 0^+} \frac{1}{\sqrt{2\pi}} \frac{t^\delta e^{-t/2}}{\phi(\Phi^{-1}(1-\Xi(t)))} \\ \stackrel{\text{L'Hospital's rule}}{=} \lim_{t \rightarrow 0^+} \frac{t^\delta e^{-t/2}}{\sqrt{2\pi}} \frac{\delta t^{-1} - 1/2}{\Phi^{-1}(1-\Xi(t)) \xi(t)} &= \lim_{t \rightarrow 0^+} \frac{(\delta t^{-1} - 1/2) t^{\delta+1/2}}{\Phi^{-1}(1-\Xi(t))} \\ &= \lim_{t \rightarrow 0^+} \frac{\delta t^{\delta-1/2}}{\Phi^{-1}(1-\Xi(t))}. \end{aligned}$$

To deal with the denominator, we use Chebyshev's approximation for the inverse of the Gaussian CDF (Blair et al., 1976). Specifically, $\Phi^{-1}(1-\Xi(t)) = \sqrt{-2 \log \Xi(t)}(1 + o(1))$ for $\Xi(t) \rightarrow 0$. Thus we have

$$\lim_{t \rightarrow 0^+} \frac{\delta t^{\delta-1/2}}{\Phi^{-1}(1-\Xi(t))} = \lim_{t \rightarrow 0^+} \frac{\delta t^{\delta-1/2}}{\sqrt{-2 \log \Xi(t)}}.$$

We approximate the CDF function $\Xi(t)$ by the limiting of $\Xi(t)$ as $t \rightarrow 0^+$. The CDF function $\Xi(t) = \gamma(1/2, t/2)/\Gamma(1/2)$, where $\gamma(\cdot, \cdot)$ is the *incomplete gamma function* and $\Gamma(\cdot)$ is the *gamma function*. Using the approximation of the incomplete gamma function, we have $\gamma(1/2, t/2) = \sqrt{2t}(1 + o(1))$ as $t \rightarrow 0^+$, and

$$\lim_{t \rightarrow 0^+} \frac{\delta t^{\delta-1/2}}{\sqrt{-2 \log \Xi(t)}} = \lim_{t \rightarrow 0^+} \frac{\delta t^{\delta-1/2}}{\sqrt{-\log t}} = 0.$$

Now we are ready to prove the statement on sensitivity. Because the function $\frac{\xi(t)}{\phi(\Phi^{-1}(1-\Xi(t)))}$ is continuous and converges to 0 as $t \rightarrow \infty$, there exists a threshold $C_3 > 0$ and the function $\frac{\xi(t)}{\phi(\Phi^{-1}(1-\Xi(t)))}$ has an upper bound $C_1 > 0$ on $[C_3, \infty)$. By the fact that the function $t^{1/2+\delta} \frac{\xi(t)}{\phi(\Phi^{-1}(1-\Xi(t)))}$ is continuous and converges to 0 as $t \rightarrow 0$. For the enough small threshold $C_3 > 0$, the function $t^{1/2+\delta} \frac{\xi(t)}{\phi(\Phi^{-1}(1-\Xi(t)))}$ has an upper bound $C_2 > 0$ on $(0, C_3]$. Thus, the function $\frac{\xi(t)}{\phi(\Phi^{-1}(1-\Xi(t)))}$ is bounded by $C_2 t^{-1/2-\delta}$ on $(0, C_3]$. Then, we have

inequality,

$$\begin{aligned}
 & \left| \Phi^{-1}\{1 - \Xi[T(\mathcal{S})]\} - \Phi^{-1}\{1 - \Xi[T(\mathcal{S}')] \} \right| \\
 & \leq \left| \int_{T(\mathcal{S}')}^{T(\mathcal{S})} \frac{\xi(t)}{\phi(\Phi^{-1}(1 - \Xi(t)))} dt \right| \\
 & \leq \left| \int_{\max\{T(\mathcal{S}'), C_3\}}^{\max\{T(\mathcal{S}), C_3\}} \frac{\xi(t)}{\phi(\Phi^{-1}(1 - \Xi(t)))} dt \right| + \left| \int_{\min\{T(\mathcal{S}'), C_3\}}^{\min\{T(\mathcal{S}), C_3\}} \frac{\xi(t)}{\phi(\Phi^{-1}(1 - \Xi(t)))} dt \right| \\
 & \leq \left| \int_{\max\{T(\mathcal{S}'), C_3\}}^{\max\{T(\mathcal{S}), C_3\}} C_1 dt \right| + \left| \int_{\min\{T(\mathcal{S}'), C_3\}}^{\min\{T(\mathcal{S}), C_3\}} C_2 t^{-1/2-\delta} dt \right| \\
 & \leq C_1 \|\max\{T(\mathcal{S}'), C_3\} - \max\{T(\mathcal{S}), C_3\}\| + C_2 \left| \int_{\min\{T(\mathcal{S}'), T(\mathcal{S}), C_3\}}^{\min\{T(\mathcal{S}'), T(\mathcal{S}), C_3\} + \|T(\mathcal{S}) - T(\mathcal{S}')\|} t^{-1/2-\delta} dt \right| \\
 & \leq M^2/nC_1 + C_2 \left| \int_0^{\|T(\mathcal{S}) - T(\mathcal{S}')\|} t^{-1/2-\delta} dt \right| \\
 & \leq M^2/nC_1 + \frac{C_2}{1/2-\delta} (M^2/n)^{1/2-\delta},
 \end{aligned}$$

where in the fourth inequality, we use the fact that $t^{-1/2-\delta}$ is decreasing, and in the fifth inequality, we use the fact that $\sup \|T(\mathcal{S}') - T(\mathcal{S})\| \leq M^2/n$. \blacksquare

D. Proof of Lemma 7

Proof We first show that *report noisy min algorithm* is $(\epsilon, \delta(\epsilon))$ -DP, where $\epsilon > 0$ and

$$\delta(\epsilon) = \Phi\left(-\frac{\epsilon}{\mu} + \frac{\mu}{2}\right) - e^\epsilon \Phi\left(-\frac{\epsilon}{\mu} - \frac{\mu}{2}\right).$$

Using the idea in Dwork et al. (2014), we first fix any $i \in \{1, \dots, n\}$. Let $f_1 = G^{-1}(p_1), \dots, f_n = G^{-1}(p_n)$ to denote the functions when the data is \mathcal{S} and $f'_1 = G^{-1}(p'_1), \dots, f'_n = G^{-1}(p'_n)$ to denote the functions when the data is \mathcal{S}' , where \mathcal{S}' is a neighborhood of \mathcal{S} . By the fact that the G function is monotone increasing, it is enough to consider reporting the noisy minimum of $\{f_1, \dots, f_n\}$.

Fix the random noise $Z_i, Z_1, \dots, Z_{i-1}, Z_{i+1}, \dots, Z_n$, which are draw independently from $N(0, 8\Delta^2/\mu^2)$. Define

$$Z^* = \max_{Z_i} : f_i + Z_i \leq f_j + Z_j \quad \forall j \neq i.$$

Then, i will be the algorithm's output when the data is \mathcal{S} if and only if $Z_i \leq Z^*$.

We have for $j \neq i$,

$$\begin{aligned}
 & f_i + Z^* \leq f_j + Z_j \\
 \Rightarrow & -\Delta + f'_i + Z^* \leq f_i + Z^* \leq f_j + Z_j \leq f'_j + Z_j + \Delta \\
 & \Rightarrow f'_i + Z^* - 2\Delta \leq f'_j + Z_j.
 \end{aligned}$$

Then, if $Z_i \leq Z^* - 2\Delta$, i will be the algorithm's output when the data is \mathcal{S}' . We propose the following claim: $1 - \Phi(x + \mu) \geq e^{-\epsilon} [1 - \Phi(x) - \delta(\epsilon)]$, where $\epsilon > 0$ and $\delta(\epsilon) = \Phi(-\epsilon/\mu + \mu/2) - e^\epsilon \Phi(-\epsilon/\mu - \mu/2)$. By the claim and letting $x = -Z^*\mu / (2^{3/2}\Delta)$, it follows that,

$$\begin{aligned} & \mathbf{P}(Z_i \geq Z^* + 2\Delta) \geq e^{-\epsilon} [\mathbf{P}(Z_i \geq Z^*) - \delta(\epsilon)] \\ \Rightarrow \mathbf{P}(i | \mathcal{S}', \mathbf{Z}_{-i}) & \geq \mathbf{P}(Z_i \geq Z^* + 2\Delta) \geq e^{-\epsilon} [\mathbf{P}(Z_i \geq Z^*) - \delta(\epsilon)] \geq e^{-\epsilon} [\mathbf{P}(i | \mathcal{S}, \mathbf{Z}_{-i}) - \delta(\epsilon)], \end{aligned}$$

where we use $\mathbf{P}(i | \mathcal{S}, \mathbf{Z}_{-i})$ and $\mathbf{P}(i | \mathcal{S}', \mathbf{Z}_{-i})$ to denote the probabilities of output i using data \mathcal{S} and \mathcal{S}' , respectively. We use \mathbf{Z}_{-i} to denote $\{Z_1, \dots, Z_{i-1}, Z_{i+1}, \dots, Z_m\}$. Then, after taking the expectation of \mathbf{Z}_{-i} , we conclude,

$$P(i | \mathcal{S}) \leq e^\epsilon P(i | \mathcal{S}') + \delta(\epsilon).$$

Then, we turn to prove the claim. Let $g(x) = 1 - \Phi(x + \mu) - e^{-\epsilon} [1 - \Phi(x) - \delta(\epsilon)]$. Take the derivative of $g(x)$, we have

$$\left. \frac{d}{dx} g(x) \right|_{x=\frac{\epsilon}{\mu} - \frac{\mu}{2}} = 0, \quad g(x) < 0 \text{ if } x < \frac{\epsilon}{\mu} - \frac{\mu}{2}, \quad g(x) > 0 \text{ if } x > \frac{\epsilon}{\mu} - \frac{\mu}{2}.$$

Then we conclude that $g(x)$ reaches global minimize at $x = \frac{\epsilon}{\mu} - \frac{\mu}{2}$, and $g\left(\frac{\epsilon}{\mu} - \frac{\mu}{2}\right) = 0$ by calculus. Then the claim holds. By Corollary 2.13 in Dong et al. (2022), report noisy min index is $2^{-1/2}\mu$ -GDP. Using the composition theorem, we conclude report noisy min algorithm is μ -GDP. \blacksquare

E. Proof of Theorem 10

Proof The Report Noisy Min algorithm with Laplace noise $\lambda = \Delta\sqrt{10m \log(1/\delta)}/\epsilon$ is $(2\epsilon/\sqrt{10m \log(1/\delta)}, 0)$ -differential private by Lemma 2.4 in Dwork et al. (2021).

Lemma 12 (Advanced Composition Theorem by Dwork et al. (2014)) *For all $\epsilon, \delta \geq 0$ and $\delta' > 0$, running l mechanisms sequentially that are each (ϵ, δ) -differentially private preserves $(\epsilon\sqrt{2l \log(1/\delta')} + l\epsilon(e^\epsilon - 1), l\delta + \delta')$ -differential privacy.*

By Lemma 12, the peeling algorithm is $(\tilde{\epsilon}, \delta)$ -differentially private, where

$$\begin{aligned} \tilde{\epsilon} &= \frac{2\epsilon}{\sqrt{10m \log(1/\delta)}} \sqrt{2m \log(1/\delta)} + m \frac{2\epsilon}{\sqrt{10m \log(1/\delta)}} \left(e^{\frac{2\epsilon}{\sqrt{10m \log(1/\delta)}}} - 1 \right) \\ &= \left[\frac{2}{\sqrt{5}} + \frac{2\sqrt{m}}{\sqrt{10 \log(1/\delta)}} \left(e^{\frac{2\epsilon}{\sqrt{10m \log(1/\delta)}}} - 1 \right) \right] \epsilon \\ &\leq \left(\frac{2}{\sqrt{5}} + 1.034 \frac{2\sqrt{m}}{\sqrt{10 \log(1/\delta)}} \frac{2\epsilon}{\sqrt{10m \log(1/\delta)}} \right) \epsilon \\ &= \left(\frac{2}{\sqrt{5}} + 1.034 \frac{4\epsilon}{10 \log(1/\delta)} \right) \epsilon, \end{aligned}$$

where we use the relation $e^x - 1 \leq 1.034x$ for $0 \leq x \leq 0.0660$ and $2\epsilon/\sqrt{10m \log(1/\delta)} \leq 0.0660$ for $\epsilon \leq 0.5, \delta \leq 0.1$ and $m \geq 10$. Furthermore,

$$\tilde{\epsilon} \leq \left(\frac{2}{\sqrt{5}} + 1.034 \frac{4 \times 0.5}{10 \log(10)} \right) \epsilon \leq \epsilon.$$

Then by post-processing property of (ϵ, δ) -differential privacy, the DP-AdaPT algorithm with Laplace noise is (ϵ, δ) -differential privacy. \blacksquare

F. Proof of Theorem 11

Proof We begin with a technique lemma. Let $[n]$ denote the set $\{1, \dots, n\}$.

Lemma 13 (Lei and Fithian (2018)) *Suppose that, conditionally on the σ -field \mathcal{G}_{-1} , b_1, \dots, b_n are independent Bernoulli random variables with $\mathbb{P}(b_i = 1 \mid \mathcal{G}_{-1}) = \rho_i \geq \rho > 0$, almost surely. Also, suppose that $[n] \supseteq \mathcal{C}_0 \supseteq \mathcal{C}_1 \supseteq \dots$, with each subset \mathcal{C}_t measurable with respect to*

$$\mathcal{G}_t = \sigma \left(\mathcal{G}_{-1}, \mathcal{C}_t, (b_i)_{i \notin \mathcal{C}_t}, \sum_{i \in \mathcal{C}_t} b_i \right).$$

If \hat{t} is an almost-surely finite stopping time with respect to the filtration $(\mathcal{G}_t)_{t \geq 0}$, then

$$\mathbb{E} \left[\frac{1 + |\mathcal{C}_{\hat{t}}|}{1 + \sum_{i \in \mathcal{C}_{\hat{t}}} b_i} \mid \mathcal{G}_{-1} \right] \leq \rho^{-1}.$$

Let \hat{t} denote the step at which we stop and reject. Then

$$\text{FDP}_{\hat{t}} = \frac{V_{\hat{t}}}{R_{\hat{t}} \vee 1} = \frac{1 + U_{\hat{t}}}{R_{\hat{t}} \vee 1} \frac{V_{\hat{t}}}{1 + U_{\hat{t}}} \leq \frac{1 + A_{\hat{t}}}{R_{\hat{t}} \vee 1} \frac{V_{\hat{t}}}{1 + U_{\hat{t}}} \leq \alpha \frac{V_{\hat{t}}}{1 + U_{\hat{t}}},$$

where we define

$$U_t = \left| \{i_j, j = 1, \dots, m : H_{i_j} \text{ is null and } \tilde{p}_{i_j} \geq 1 - s_t(\mathbf{x}_{i_j})\} \right|,$$

and

$$V_t = \left| \{i_j, j = 1, \dots, m : H_{i_j} \text{ is null and } \tilde{p}_{i_j} \leq s_t(\mathbf{x}_{i_j})\} \right|,$$

and we use the fact $U_t \leq A_t$ and the stopping criterion $(1 + A_t) / (R_t \vee 1) \leq \alpha$. It is enough to consider $V_{\hat{t}} / (1 + U_{\hat{t}})$.

We use \tilde{p}_{i_j} to denote the noisy p -values that are subject to release, for $j = 1, \dots, m$. Let $\tilde{m}_{i_j} = \min\{\tilde{p}_{i_j}, 1 - \tilde{p}_{i_j}\}$, $b_{i_j} = 1\{\tilde{p}_{i_j} \geq 0.5\}$ and $\tilde{p}_{i_j} = b_{i_j}(1 - \tilde{m}_{i_j}) + (1 - b_{i_j})\tilde{m}_{i_j}$. We use \mathcal{P} to denote the indexes returned by the mirror peeling algorithm. Let $\mathcal{C}_t = \{i \in \mathcal{H}_0, i \in \mathcal{P} : \tilde{p}_i \notin (s_t(\mathbf{x}_i), 1 - s_t(\mathbf{x}_i))\}$ be the index of the null p -values which are unavailable to the analyst. Then, we have the relation $U_t = \sum_{i \in \mathcal{C}_t} b_i$ and $V_t = \sum_{i \in \mathcal{C}_t} (1 - b_i) = |\mathcal{C}_t| - U_t$.

We then construct the auxiliary filtration $(\mathcal{G}_t)_{t \geq -2}$. First, define the σ -field,

$$\mathcal{G}_{-2} = \sigma(\{\mathbf{x}_i, m_i\}_{i=1}^n),$$

where $m_i = \min\{p_i, 1 - p_i\}$. In the mirror-peeling mechanism, we use $\tilde{p}_{i,j}$ to denote the noisy p -value of \tilde{p}_i in the j -th peeling mechanism. Define the peeling σ -field,

$$\mathcal{G}_{-2} \subseteq \mathcal{G}_{-2+\frac{1}{m+1}} \subseteq \cdots \subseteq \mathcal{G}_{-1-\frac{1}{m+1}},$$

where $\mathcal{G}_{-2+\frac{j}{m+1}} = \sigma\left(\mathcal{G}_{-2+\frac{j-1}{m+1}, \{\tilde{m}_{i,j}\}_{i=1}^n}\right)$ and $\tilde{m}_{i,j} = \min\{\tilde{p}_{i,j}, 1 - \tilde{p}_{i,j}\}$. It is not difficult to see the j -th peeling mechanism is measurable with respect to $\mathcal{G}_{-2+\frac{j}{m+1}}$. Then, we define

$$\mathcal{G}_{-1} = \sigma\left(\mathcal{G}_{-1-\frac{1}{m+1}, \{\tilde{m}_i\}_{i=1}^n, (b_i)_{i \notin \mathcal{H}_0}\right),$$

and

$$\mathcal{G}_t = \sigma\left(\mathcal{G}_{-1}, \mathcal{C}_t, (b_i)_{i \notin \mathcal{C}_t}, \sum_{i \in \mathcal{C}_t} b_i\right).$$

The assumptions of independence and mirror-conservatism guarantee $\mathbb{P}(b_i = 1 \mid \mathcal{G}_{-1}) \geq 0.5$ almost surely for each $i \in \mathcal{H}_0$. The reasons are

$$\begin{aligned} \mathbb{P}(b_i = 1 \mid \mathcal{G}_{-1}) &= \mathbb{P}(\tilde{p}_i \geq 0.5 \mid \mathcal{G}_{-1}) = \mathbb{P}(\tilde{p}_i \geq 0.5 \mid \tilde{m}_i, m_i, \{m_{i,j}\}_{j=1}^m) \\ &= \mathbb{P}(\tilde{p}_i \geq 0.5 \mid \tilde{m}_i), \end{aligned}$$

and that \tilde{p}_i is mirror-conservative by Theorem 3.

Notice that $\tilde{p}_i \in \mathcal{G}_t$ for $\tilde{p}_i \in (s_t(\mathbf{x}_i), 1 - s_t(\mathbf{x}_i))$, $A_t = U_t + |i \notin \mathcal{H}_0, i \in \mathcal{P} : \tilde{p}_i \geq 1 - s_t(\mathbf{x}_i)|$ and $R_t = V_t + |i \notin \mathcal{H}_0, i \in \mathcal{P} : \tilde{p}_i \leq s_t(\mathbf{x}_i)|$. Then, we conclude $\mathcal{F}_t \subseteq \mathcal{G}_t$. It follows that $\hat{t} = \min\{t : \widehat{\text{FDP}}_t \leq \alpha\}$ is a stopping time with respect to \mathcal{G}_t . Then we have

$$\mathbb{E}[\text{FDP} \mid \mathcal{G}_{-1}] \leq \alpha \mathbb{E}\left[\frac{V_{\hat{t}}}{1 + U_{\hat{t}}} \mid \mathcal{G}_{-1}\right] = \alpha \mathbb{E}\left[\frac{1 + |\mathcal{C}_{\hat{t}}|}{1 + U_{\hat{t}}} - 1 \mid \mathcal{G}_{-1}\right] \leq \alpha.$$

Finally, notice that $\mathcal{F}_{-1} \subset \mathcal{G}_{-1}$ and use the tower property of conditional expectation. ■

References

- Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, pages 308–318, 2016.
- Marco Avella-Medina, Casey Bradshaw, and Po-Ling Loh. Differentially private inference via noisy optimization. *arXiv preprint arXiv:2103.11003*, 2021.
- Jordan Awan and Zhanrui Cai. One step to efficient synthetic data. *arXiv preprint arXiv:2006.02397*, 2020.
- Rina Foygel Barber and Emmanuel J Candès. Controlling the false discovery rate via knockoffs. *The Annals of Statistics*, 43(5):2055–2085, 2015.

- Yoav Benjamini and Yosef Hochberg. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: series B (Methodological)*, 57(1):289–300, 1995.
- Yoav Benjamini and Daniel Yekutieli. The control of the false discovery rate in multiple testing under dependency. *Annals of Statistics*, pages 1165–1188, 2001.
- JM Blair, CA Edwards, and J Howard Johnson. Rational chebyshev approximations for the inverse of the error function. *Mathematics of Computation*, 30(136):827–830, 1976.
- Daniel Bottomly, Nicole AR Walter, Jessica Ezzell Hunter, Priscila Darakjian, Sunita Kawane, Kari J Buck, Robert P Searles, Michael Mooney, Shannon K McWeeney, and Robert Hitzemann. Evaluating gene expression in c57bl/6j and dba/2j mouse striatum using rna-seq and microarrays. *PloS one*, 6(3):e17820, 2011.
- T Tony Cai, Yichen Wang, and Linjun Zhang. The cost of privacy: Optimal rates of convergence for parameter estimation with differential privacy. *The Annals of Statistics*, 49(5):2825–2850, 2021.
- Zhanrui Cai. Robust model selection with application in single-cell multiomics data. *arXiv preprint arXiv:2305.05714*, 2023.
- Zhanrui Cai, Runze Li, and Liping Zhu. Online sufficient dimension reduction through sliced inverse regression. *The Journal of Machine Learning Research*, 21(1):321–345, 2020.
- Zhanrui Cai, Jing Lei, and Kathryn Roeder. Model-free prediction test with application to genomics data. *Proceedings of the National Academy of Sciences*, 119(34):e2205518119, 2022a. doi: 10.1073/pnas.2205518119.
- Zhanrui Cai, Changcheng Li, Jiawei Wen, and Songshan Yang. Asset splitting algorithm for ultrahigh dimensional portfolio selection and its theoretical property. *Journal of Econometrics*, 2022b.
- Zhanrui Cai, Jing Lei, and Kathryn Roeder. Asymptotic distribution-free independence test for high dimension data. *Journal of the American Statistical Association*, pages 1–20, 2023.
- Chenguang Dai, Buyu Lin, Xin Xing, and Jun S Liu. False discovery rate control via data splitting. *Journal of the American Statistical Association*, pages 1–18, 2022.
- Edgar Dobriban, Kristen Fortney, Stuart K Kim, and Art B Owen. Optimal multiple testing under a gaussian prior on the effect sizes. *Biometrika*, 102(4):753–766, 2015.
- Jinshuo Dong, Aaron Roth, and Weijie Su. Gaussian differential privacy. *Journal of the Royal Statistical Society, Series B*, 2022.
- Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*, pages 265–284. Springer, 2006.

- Cynthia Dwork, Guy N Rothblum, and Salil Vadhan. Boosting and differential privacy. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pages 51–60. IEEE, 2010.
- Cynthia Dwork, Aaron Roth, et al. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407, 2014.
- Cynthia Dwork, Weijie Su, and Li Zhang. Differentially private false discovery rate control. *Journal of Privacy and Confidentiality*, 11(2), 2021.
- W Fuller. Masking procedures for microdata disclosure. *Journal of Official Statistics*, 9(2):383–406, 1993.
- Christopher R Genovese, Kathryn Roeder, and Larry Wasserman. False discovery control with p-value weighting. *Biometrika*, 93(3):509–524, 2006.
- Rob Hall, Alessandro Rinaldo, and Larry Wasserman. Differential privacy for functions and functional data. *The Journal of Machine Learning Research*, 14(1):703–727, 2013.
- James X Hu, Hongyu Zhao, and Harrison H Zhou. False discovery rate control with groups. *Journal of the American Statistical Association*, 105(491):1215–1227, 2010.
- Nikolaos Ignatiadis, Bernd Klaus, Judith B Zaugg, and Wolfgang Huber. Data-driven hypothesis weighting increases detection power in genome-scale multiple testing. *Nature methods*, 13(7):577–580, 2016.
- Liina Kamm, Dan Bogdanov, Sven Laur, and Jaak Vilo. A new way to protect privacy in large-scale genome-wide association studies. *Bioinformatics*, 29(7):886–893, 2013.
- Vishesh Karwa and Aleksandra Slavković. Inference using noisy degrees: Differentially private β -model and synthetic graphs. *The Annals of Statistics*, 44(1):87–112, 2016.
- Duhyeong Kim, Yongha Son, Dongwoo Kim, Andrey Kim, Seungwan Hong, and Jung Hee Cheon. Privacy-preserving approximate gwas computation based on homomorphic encryption. *BMC Medical Genomics*, 13(7):1–12, 2020.
- Jae Kwang Kim and Jun Shao. *Statistical methods for handling incomplete data*. Chapman and Hall/CRC, 2021.
- Lihua Lei and William Fithian. Adapt: an interactive procedure for multiple testing with side information. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 80(4):649–679, 2018.
- Ang Li and Rina Foygel Barber. Accumulation tests for fdr control in ordered hypothesis testing. *Journal of the American Statistical Association*, 112(518):837–849, 2017.
- Michael I Love, Wolfgang Huber, and Simon Anders. Moderated estimation of fold change and dispersion for rna-seq data with deseq2. *Genome biology*, 15(12):1–21, 2014.

- Aaditya K Ramdas, Rina F Barber, Martin J Wainwright, and Michael I Jordan. A unified treatment of multiple testing with prior knowledge using the p-filter. *The Annals of Statistics*, 47(5):2790–2821, 2019.
- John D Storey, Jonathan E Taylor, and David Siegmund. Strong control, conservative point estimation and simultaneous conservative consistency of false discovery rates: a unified approach. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 66(1):187–205, 2004.
- Jinjin Tian and Aaditya Ramdas. Addis: an adaptive discarding algorithm for online fdr control with conservative nulls. *Advances in neural information processing systems*, 32, 2019.
- Zhaoxue Tong, Zhanrui Cai, Songsan Yang, and Runze Li. Model-free conditional feature screening with fdr control. *Journal of the American Statistical Association*, pages 1–13, 2022.
- Larry Wasserman and Shuheng Zhou. A statistical framework for differential privacy. *Journal of the American Statistical Association*, 105(489):375–389, 2010.