

DIFFERENTIALLY PRIVATE VARIABLE SELECTION VIA THE KNOCKOFF FILTER

Mehrdad Pournaderi and Yu Xiang

University of Utah
50 Central Campus Dr 2110, Salt Lake City, USA
Email: {mehrdad.pournaderi, yu.xiang}@utah.edu

ABSTRACT

The knockoff filter, recently developed by Barber and Candès, is an effective procedure to perform variable selection with a controlled false discovery rate (FDR). We propose a private version of the knockoff filter by incorporating Gaussian and Laplace mechanisms, and show that variable selection with controlled FDR can be achieved. Simulations demonstrate that our setting has reasonable statistical power.

Index Terms— False discovery rate control, differential privacy, knockoff filter, variable selection, Lasso.

1. INTRODUCTION

A common theme in statistical inference and machine learning is to select a few relevant features or variables (among many) that are associated with the responses of interest. Controlling the expected proportion of falsely selected features (or the FDR [1]) for the variable selection problem has attracted much attention since the development of the *knockoff filter* procedure due to its flexibility (in constructing test statistics) as well as good finite-sample performances [2, 3]. In this procedure, a test statistic is computed for every feature of the model and the variables are selected by (data-dependent) thresholding the statistics according to the target FDR. The knockoff filter has been generalized to various settings such as the high-dimensional case ($p > n$) [3], beyond linear models (the model-X framework) [4], and deep learning-based procedure [5], with applications in a variety of domains [6, 7].

In this work, we aim to develop (ϵ, δ) -differentially-private statistics that ensure (finite-sample) FDR control while being empirically powerful. Differential privacy is currently the standard framework concerning the privacy of individuals. The differential privacy of other variable selection methods like the Lasso [8] and BH procedure [1] are investigated in [9] and [10], respectively. However, the privacy of the knockoff procedure hasn't been discussed in the literature yet. Related works are [11, 12] where they consider the differential privacy of the sample covariance and the

least-squares estimator. Although our approach is similar, the work is different in the sense that we take care of the FDR control and the knockoff design which is a function of the data. We adopt the sensitivity analysis and output perturbation method [13] and randomize the statistics by adding noise according to Gaussian and Laplace mechanism [14].

2. PRELIMINARIES

2.1. Fixed-X Knockoff Filter

Let $\mathbf{y} \in \mathbb{R}^n$ denote a vector of responses, and $\mathbf{X} \in \mathbb{R}^{n \times p}$ a data matrix consisting of p (explanatory) variables and n samples. A canonical linear regression model is given by

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{w}, \quad (1)$$

where $\boldsymbol{\beta} = [\beta_1, \dots, \beta_p]^\top \in \mathbb{R}^p$ is the unknown regression vector and the model assumptions are as follows.

Assumption 1 \mathbf{w} is Gaussian with constant variance σ^2 i.e. $\mathbf{w} \sim \mathcal{N}(0, \sigma^2 \mathbf{I}_n)$.

Assumption 2 $\boldsymbol{\Sigma} = \mathbf{X}^\top \mathbf{X}$ is invertible.

We say the i -th variable is a null variable if $\beta_i = 0$. The goal is to select the non-null variables (i.e. $\beta_i \neq 0$) without selecting too many nulls. For the linear model (1), let $\hat{S} \subseteq \{1, 2, \dots, p\}$ denote the selected variables by some procedure. Then the FDR is defined as follows

$$\text{FDR} = \mathbb{E} \left[\frac{\#\{j : \beta_j = 0 \text{ and } j \in \hat{S}\}}{\#\{j : j \in \hat{S}\} \vee 1} \right], \quad (2)$$

where the expectation is taken over the randomness of the error terms, \mathbf{w} . A selection rule controls the FDR at level q if its corresponding FDR is guaranteed to be at most q . In the knockoff procedure, the FDR control is achieved by computing a vector of statistics $\mathbf{W} = (W_1, W_2, \dots, W_p)^\top$ with components corresponding to each variable and selecting $\{j : W_j \geq T(q, \Omega)\}$, where the threshold $T \geq 0$ depends

on the target FDR q and the set of computed statistics, i.e., $\Omega = \{W_1, \dots, W_p\}$. The main idea in computing the statistics is to use a fake copy of the design that preserve the correlation structure of \mathbf{X} but destroys the relationship between \mathbf{X} and \mathbf{y} . For example, if we shuffle (or permute) the samples in \mathbf{X} , we preserve the geometry of \mathbf{X} but the relation with \mathbf{y} would be diminished. The knockoff filter is based on a more elegant construction of such a fake design, namely, the knockoff design.

2.1.1. Knockoff Design

Knockoff copies are constructed using only the original features and the subtle geometrical relation between the data matrix \mathbf{X} and its knockoff $\tilde{\mathbf{X}} \in \mathbb{R}^{n \times p}$ makes them an appropriate tool for FDR control. Let $\Sigma = \mathbf{X}^\top \mathbf{X}$ denote the Gram matrix of the original features. The relationship between \mathbf{X} and $\tilde{\mathbf{X}}$ can be summarized as below

$$\begin{aligned} \mathbf{G} &:= [\mathbf{X} \ \tilde{\mathbf{X}}]^\top [\mathbf{X} \ \tilde{\mathbf{X}}] \\ &= \begin{bmatrix} \Sigma & \Sigma - \text{diag}\{\mathbf{s}\} \\ \Sigma - \text{diag}\{\mathbf{s}\} & \Sigma \end{bmatrix}, \end{aligned} \quad (3)$$

where $\mathbf{s} \in \mathbb{R}_+^p$ is chosen in a way that \mathbf{G} is positive semidefinite. If $n \geq 2p$, $\tilde{\mathbf{X}}$ can be constructed as follows

$$\tilde{\mathbf{X}} = \mathbf{X} (\mathbf{I}_p - \Sigma^{-1} \text{diag}\{\mathbf{s}\}) + \tilde{\mathbf{U}} \mathbf{C}, \quad (4)$$

where \mathbf{C} is obtained by Cholesky decomposition of the Schur complement of \mathbf{G} and $\tilde{\mathbf{U}}^{n \times p}$ is an orthonormal matrix that satisfies $\tilde{\mathbf{U}}^\top \mathbf{X} = 0$. The existence of $\tilde{\mathbf{U}}$ is guaranteed by $n \geq 2p$. By looking at (3), one can observe that the knockoff features matrix $\tilde{\mathbf{X}}$ preserves the correlation structure of \mathbf{X} and the vector \mathbf{s} gives us a degree of freedom to make knockoff features different than the original ones, e.g., if $\mathbf{s} = 0$ then $\tilde{\mathbf{X}} = \mathbf{X}$. For maximum detection power, the knockoff and original features should be as orthogonal to each other as possible. The restriction on this matter is imposed by $\mathbf{G} \succeq 0$ which is the necessary and sufficient condition for $\tilde{\mathbf{X}}$ to exist.

Remark 1 A suggestion in [2] regarding the choice of \mathbf{s} for a normalized design, is to set $\text{diag}\{\mathbf{s}\} = \min(2\lambda_{\min}(\Sigma), 1) \mathbf{I}_p$. If $2\lambda_{\min}(\Sigma) \leq 1$, this choice (as we shall see in Lemma 1) forces the minimum eigenvalue of \mathbf{G} to be zero. For the differential privacy purposes we recommend using $\text{diag}\{\mathbf{s}\} = \lambda_{\min}(\Sigma) \mathbf{I}_p$ to make the regression operators stable.

2.1.2. Statistics

Using the knockoff features, we are now able to construct a vector of statistics \mathbf{W} allowing FDR control. For instance, the Lasso penalized least-squares estimate can be used to construct the statistics.

$$\hat{\beta}^* := \arg \min_{\mathbf{b} \in \mathbb{R}^{2p}} \left\{ \left\| \mathbf{y} - [\mathbf{X} \ \tilde{\mathbf{X}}] \mathbf{b} \right\|_2^2 + \lambda \|\mathbf{b}\|_1 \right\}, \quad (5)$$

where $\lambda \geq 0$. Using this estimate, one way to define the statistics \mathbf{W} is as follows

$$W_i^{\text{LCD}} = |\hat{\beta}_i^*| - |\hat{\beta}_{i+p}^*|, \quad 1 \leq i \leq p. \quad (6)$$

We can also define the statistics differently,

$$W_i = \text{sgn}(|\hat{\beta}_i^*| - |\hat{\beta}_{i+p}^*|) \max(|\hat{\beta}_i^*|, |\hat{\beta}_{i+p}^*|). \quad (7)$$

The key feature of these statistics that make them suitable for the purpose of FDR control is the *i.i.d. sign property* of the nulls, i.e., the signs of null statistics (signs of $W = \{W_j : \beta_j = 0\}$) are independent of their magnitudes and have i.i.d. Bern(1/2) distribution. We will look at this property in more details in the subsequent subsection. Unlike the nulls, non-null statistics are expected to get large positive values. Thus, a selection rule of the form $\{j : W_j \geq T\}$ for some $T > 0$ seems reasonable. However, T should be characterized such that guarantees FDR control.

2.1.3. IID Sign Property of Nulls

In general, satisfying the following two properties is sufficient to guarantee the i.i.d. sign property of the null statistics. Let $\mathbf{X}^{(i)}$ denote the i -th column of \mathbf{X} .

Property I : Swapping $\mathbf{X}^{(i)}$ and $\tilde{\mathbf{X}}^{(i)}$ for all $i \in F$ and any $F \subseteq \{1, 2, \dots, p\}$ has the effect of switching the signs of $\{W_i : i \in F\}$. This property is called **antisymmetry** in [2].

In the least-squares problem (5), it can be observed that swapping $\mathbf{X}^{(i)}$ and $\tilde{\mathbf{X}}^{(i)}$ for any $i \in \{1, 2, \dots, p\}$ will swap $\hat{\beta}_i^*$ and $\hat{\beta}_{i+p}^*$ and this will result in switching the sign of W_i according to (6) and (7).

Property II : Let $\mathbf{W}_{\text{swap}(F)}$ denote the vector of statistics obtained when $\mathbf{X}^{(i)}$ and $\tilde{\mathbf{X}}^{(i)}$ are swapped in the regression problem for all $i \in F$ where $F \subseteq \{j : \beta_j = 0\}$. Then $\mathbf{W}_{\text{swap}(F)}$ and \mathbf{W} have the same distribution, i.e. $\mathbf{W}_{\text{swap}(F)} \stackrel{d}{=} \mathbf{W}$.

In the least-squares problem (5), it can be shown that the estimated coefficients for null variables and their corresponding knockoff variable are exchangeable which immediately implies **Property II**.

Using the i.i.d. sign property for the null statistics, it can be shown that the selection rule $\hat{S}(T) = \{j : W_j \geq T\}$ controls the FDR at level $q \in [0, 1]$, where

$$T = \min \left\{ t \in \Psi : \frac{1 + \#\{j : W_j \leq -t\}}{\#\{j : W_j \geq t\} \vee 1} \leq q \right\}, \quad (8)$$

if $\Psi = \{|W_i| : i = 1, 2, \dots, p\} \setminus \{0\}$ is not empty and $T = +\infty$ if $\Psi = \emptyset$.

2.2. Differential Privacy

Differential privacy provides a rigorous formulation of individual privacy that allows for the development of (randomizing) mechanisms that release some information about the data

without revealing too much about any individuals. In other words, the goal is to have a guarantee that the effect of each observation in the output is negligible.

Definition 1 *Two databases (matrices) are called neighbors or adjacent if they differ only in a single observation (row).*

Definition 2 (Differential Privacy) *A randomized mechanism \mathcal{M} with domain \mathcal{D} and range \mathcal{R} is (ϵ, δ) -differentially private if for all pairs of neighboring inputs $\mathbf{A}, \mathbf{A}' \in \mathcal{D}$ and all measurable subsets $\mathcal{S} \subset \mathcal{R}$, it holds that,*

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

The algorithms we are going to propose are based on the *output perturbation* mechanisms. In general, the class of additive noise mechanisms perturb the output by adding random noise proportional to the (differential) sensitivity of the output, that is, $\mathcal{M}(\mathbf{A}) = f(\mathbf{A}) + Z$, where f is some (vector-valued) mapping and Z denotes the noise. The classical Gaussian mechanism [14] and Laplace mechanism [13] adopt $Z \sim \mathcal{N}(0, \kappa^2 I)$ and $Z_i \stackrel{\text{i.i.d.}}{\sim} \text{Lap}(\eta)$, where κ^2 and η are determined by the ℓ_2 and ℓ_1 -sensitivity of f , respectively.

Definition 3 *The ℓ_2 -sensitivity of a function $f : \mathcal{D}(f) \rightarrow \mathbb{R}^p$ is defined as $\Delta_2 f = \max_{\mathbf{A} \& \mathbf{A}' \text{ adjacent}} \|f(\mathbf{A}) - f(\mathbf{A}')\|_2$. In case of a function with range in matrices we use the Frobenius norm $\|\cdot\|_F$ in the definition.*

Theorem 1 ([14]) *For any $\epsilon \in (0, 1)$ and $\delta \in (0, 1)$ the Gaussian mechanism with $\kappa^2 > 2 \ln(1.25/\delta) (\frac{\Delta_2 f}{\epsilon})^2$ is (ϵ, δ) -differentially private.*

Theorem 2 ([13]) *Defining ℓ_1 -sensitivity similar to Definition 3 $\Delta_1 f = \max_{\mathbf{A} \& \mathbf{A}' \text{ adjacent}} \|f(\mathbf{A}) - f(\mathbf{A}')\|_1$, adding i.i.d. $\text{Lap}(\frac{\Delta_1 f}{\epsilon})$ to the components of $f(\mathbf{A})$ is $(\epsilon, 0)$ -differentially private.*

3. MAIN RESULTS

In this section we provide two methods for computing differential private statistic under the assumptions 1,2, and the following assumption.

Assumption 3 *B is an upper bound for the ℓ_2 -norm of the rows in \mathbf{X} and a lower bound for the ℓ_2 -norm of the columns.*

For the purpose of variable selection with knockoffs, it is natural to normalize \mathbf{X} by columns. So in both methods, the normalization is considered as a part of the procedure. Also, we let $\text{diag}\{\mathbf{s}\} = \lambda_{\min}(\Sigma') \mathbf{I}_p$ in creating knockoffs (4), where $\Sigma' = \mathbf{X}'^\top \mathbf{X}'$ and \mathbf{X}' denotes the design matrix \mathbf{X} normalized by the ℓ_2 -norm of its columns.

3.1. Method I

In this method we consider statistics that depend on the data through the following form,

$$\widehat{\mathbf{W}}_1 = \mathbf{f}([\mathbf{X}' \widetilde{\mathbf{X}}']^\top [\mathbf{X}' \widetilde{\mathbf{X}}'] + E, [\mathbf{X}' \widetilde{\mathbf{X}}']^\top \mathbf{y} + \mathbf{e}), \quad (9)$$

where E and \mathbf{e} denote the additive noise terms ensuring the differential privacy. Define $C_{\min} = \min_{1 \leq i \leq p} \|\mathbf{X}^{(i)}\|_2$ and $\eta^2 = \frac{B^2}{C_{\min}^2 - B^2}$, where $\mathbf{X}^{(i)}$ denotes the i -th column of \mathbf{X} . In the first argument of (9), $E \in \mathbb{R}^{2p \times 2p}$ is a random term of the following structure,

$$E = \theta_1 \begin{bmatrix} \mathbf{O}_p & \mathbf{I}_p \\ \mathbf{I}_p & \mathbf{O}_p \end{bmatrix} + \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix} \otimes \theta_2,$$

where $\theta_1 \sim \text{Lap}(\eta^2(1 + \lambda_{\min}(\Sigma'))/\epsilon_1)$ and $\theta_2 \in \mathbb{R}^{p \times p}$ is a symmetric matrix with zero diagonal entries such that the elements of the upper triangle are drawn i.i.d. from $\mathcal{N}(0, \kappa_1^2)$, where $\kappa_1^2 \geq \frac{2 \ln(1.25/\delta)}{\epsilon_1^2} \eta^4 (\sqrt{2} + \|\Sigma'\|_F)^2$, and each lower triangle entry is copied from its upper triangle counterpart. Regarding the second argument of (9), \mathbf{e} is a vector with i.i.d. $\mathcal{N}(0, \kappa_2^2)$ components with $\kappa_2^2 > \frac{2 \ln(1.25/\delta_1)}{\epsilon_2^2} \Delta_2([\mathbf{X}' \widetilde{\mathbf{X}}']^\top \mathbf{y})^2$ computed according to the following ℓ_2 -sensitivity,

$$\begin{aligned} \Delta_2([\mathbf{X}' \widetilde{\mathbf{X}}']^\top \mathbf{y}) &= \zeta^{\frac{1}{2}} \left(2\sqrt{\gamma} + \eta (3 + 2\lambda_{\max}(\Sigma')) \right. \\ &+ \lambda_{\min}(\Sigma')^{\frac{1}{2}} \left. \right) + \|\beta\| \left(\sqrt{2} \left(\frac{\eta}{B} - \frac{1}{C_{\min}} \right) \|\Sigma'\|_F + 2\eta B \right. \\ &+ \left. \left(C_{\min} - \frac{B}{\eta} \right) \lambda_{\min}(\Sigma') + \eta^2 (\lambda_{\min}(\Sigma') + 1) \sqrt{C_{\min}^2 + B^2} \right), \end{aligned} \quad (10)$$

where, $\zeta = \frac{2p\sigma^2}{1 - \sqrt{\frac{2}{p} \ln(2/\delta_2)}}$, $\gamma = 2\lambda_{\max}(\Sigma') - \lambda_{\min}(\Sigma')$, and $\delta_2 > 2e^{-p/2}$.

Theorem 3 (Privacy) *Releasing $\widehat{\mathbf{W}}_1$ is $(\epsilon + \epsilon', \delta + \delta')$ -differentially private for arbitrary \mathbf{f} , where $\delta' = \delta_1 + \delta_2$ and $\epsilon' = \epsilon_1 + \epsilon_2$.*

Theorem 4 (FDR control) *Let $F \subseteq \{1, \dots, p\}$ and P_F denote a $2p \times 2p$ symmetric permutation matrix corresponding to swapping the i -th column (or row) with $(i + p)$ -th for all $i \in F$. Suppose the operator $\mathbf{f}(\cdot, \cdot) : \mathcal{D} \rightarrow \mathbb{R}^p$ (with $\mathcal{D} \subseteq \mathbb{R}^{2p \times 2p} \times \mathbb{R}^{2p}$) is antisymmetric in the sense that for all $(\mathbf{A}, \mathbf{b}) \in \mathcal{D}$, computing $\mathbf{f}(P_F^\top \mathbf{A} P_F, P_F \mathbf{b})$ has the effect of switching the signs of the components of $\mathbf{f}(\mathbf{A}, \mathbf{b})$ corresponding to F . Then, applying the knockoff selection rule on $\widehat{\mathbf{W}}_1$ controls the FDR.*

For instance, we observe that if $\mathbf{f}(\mathbf{A}, \mathbf{b})$ denotes the vector of statistics computed using the estimate $\mathbf{A}^{-1} \mathbf{b}$, then $\mathbf{f}(\cdot, \cdot)$ satisfies the antisymmetry condition in Theorem 4, because we have $(P_F \mathbf{A} P_F)^{-1} P_F \mathbf{b} = P_F \mathbf{A}^{-1} \mathbf{b}$.

3.2. Method II

In this method, we add noise directly to the least-squares estimate $\hat{\beta}'$, i.e. (5) with normalized design and $\lambda = 0$.

Theorem 5 (Privacy) Fix $\delta_2 > 2e^{-p/2}$. If $\lambda_{\min}(\Sigma') > \frac{\eta^2}{1-\eta^2}$, then by the Gaussian mechanism, releasing $\hat{\beta}' + \mathbf{e}$ with noise variance $\kappa^2 > \frac{2 \ln(1.25/\delta_1)}{\epsilon^2} \Delta_2(\hat{\beta}')^2$ and the following ℓ_2 -sensitivity,

$$\Delta_2(\hat{\beta}') = \frac{2\zeta^{\frac{1}{2}}}{\sqrt{(1-\eta^2)\lambda_{\min}(\Sigma') - \eta^2}} + \left(C_{\min} - \frac{B}{\eta}\right) \|\beta\|,$$

is $(\epsilon, \delta_1 + \delta_2)$ -differentially private, where $\zeta = \frac{2p\sigma^2}{1 - \sqrt{\frac{2}{p} \ln(2/\delta_2)}}$.

Theorem 6 (FDR control) Let $F \subseteq \{1, \dots, p\}$ and define P_F as in Theorem 4. Suppose the operator $\mathbf{g}(\cdot) : \mathcal{D} \rightarrow \mathbb{R}^p$ (with $\mathcal{D} \subseteq \mathbb{R}^{2p}$) is antisymmetric in the sense that for all $\mathbf{b} \in \mathcal{D}$, computing $\mathbf{g}(P_F \mathbf{b})$ has the effect of switching the signs of the components of $\mathbf{g}(\mathbf{b})$ corresponding to F . Then, for any vector \mathbf{e} with i.i.d. elements, applying the knockoff selection rule on $\widehat{\mathbf{W}}_2 = \mathbf{g}(\hat{\beta}' + \mathbf{e})$ will control the FDR.

Remark 2 The sensitivities we compute depend on the data, so the additive noise is calibrated according to the observed data set. As a result, the privacy holds in a local sense and not globally. In Method II, the term $\lambda_{\min}(\Sigma')$ in the denominator can make the ℓ_2 -sensitivity $\Delta_2(\hat{\beta}')$, sensitive to the data. This may be an issue in case the adversary has some side information about the data. However, we can use the Ridge regression to alleviate this issue and stabilize the regression operator, i.e., if we add $\omega^2 \mathbf{I}_{2p}$ to \mathbf{G}' we get $\lambda_{\min}(\mathbf{G}') = \lambda_{\min}(\Sigma') + \omega^2$, which is lower bounded by ω^2 . Also, the sensitivities depend on the unknown parameters σ^2 and β which should be estimated or bounded for the use in practice.

4. SIMULATION RESULTS

In this section we present simulation results on synthetic data sets for both Method I and II. The sample size n is swept between 10^3 to $10^{6.6}$. We have $p = 50$ measurements per sample and the entries of the design matrix \mathbf{X} are generated i.i.d. according to $\mathcal{N}(0, \mathbf{I}_p)$. The responses \mathbf{y} are generated according to the linear model (1) with noise variance $\sigma^2 = 1$ and $k = 15$ (non-zero) underlying coefficients of the same amplitude $A = 4.5$ and sign. We adopt the OLS estimator and the Coefficient Signed Max (CSM) statistic (7). The target FDR is $q = 0.2$ and the plots are based on averaging 250 trials. Each trial is $(0.2, \frac{2p}{n})$ -differentially private, meaning we let the overall δ parameter drop as the sample size grows, so the privacy hold in a more strict sense for larger sample sizes¹. In Method I we set $\epsilon = 0.1$, $\epsilon_1 = \epsilon_2 = 0.05$, and

¹In practice it is common to take δ to scale as $\frac{1}{n}$.

$\delta = \delta_1 = \delta_2$. In Method II we have $\epsilon = 0.2$ and $\delta_1 = \delta_2$. The power is defined as follows,

$$\text{Power} = \frac{1}{k} \mathbb{E}[\#\{j : \beta_j \neq 0 \text{ and } j \in \hat{S}\}]. \quad (11)$$




Fig. 1: Power and FDR for the two differentially private knockoff procedures. The FDR is controlled at the desired level $q = 0.2$ for both methods. However, Method II has the superior power in this experiment.

A. APPENDIX: PRIVACY PROOFS

In the differential privacy proofs we use the following lemma.

Lemma 1 If we set $\text{diag}\{\mathbf{s}\} = \lambda_{\min}(\Sigma) \mathbf{I}_p$ in creating knockoff variables (4), then

$$\lambda_{\max}(\mathbf{G}) = 2\lambda_{\max}(\Sigma) - \lambda_{\min}(\Sigma), \quad (12)$$

$$\lambda_{\min}(\mathbf{G}) = \lambda_{\min}(\Sigma), \quad (13)$$

where Σ denotes the Gram matrix $\mathbf{X}^\top \mathbf{X}$.

Proof: We note that $\Sigma - \lambda \mathbf{I}_p$ and $\Sigma - \lambda_{\min}(\Sigma) \mathbf{I}_p$ commute, therefore for the block matrix $\mathbf{G} - \lambda \mathbf{I}_{2p}$ we have [15],

$$\det(\mathbf{G} - \lambda \mathbf{I}_{2p}) = \det(2\Sigma - \lambda_{\min}(\Sigma) \mathbf{I}_p - \lambda \mathbf{I}_p) \det(\lambda_{\min}(\Sigma) \mathbf{I}_p - \lambda \mathbf{I}_p). \quad (14)$$

Hence, the set of the eigenvalues of \mathbf{G} is the union of eigenvalues of $2\Sigma - \lambda_{\min}(\Sigma) \mathbf{I}_p$ and $\lambda_{\min}(\Sigma) \mathbf{I}_p$. Therefore,

$$\begin{aligned} \lambda_{\max}(\mathbf{G}) &= \max(2\lambda_{\max}(\Sigma) - \lambda_{\min}(\Sigma), \lambda_{\min}(\Sigma)) \\ &= 2\lambda_{\max}(\Sigma) - \lambda_{\min}(\Sigma), \end{aligned} \quad (15)$$

$$\begin{aligned}\lambda_{\min}(\mathbf{G}) &= \min(2\lambda_{\min}(\boldsymbol{\Sigma}) - \lambda_{\min}(\boldsymbol{\Sigma}), \lambda_{\min}(\boldsymbol{\Sigma})) \\ &= \lambda_{\min}(\boldsymbol{\Sigma}).\end{aligned}\quad (16)$$

We observe that if we set $\text{diag}\{\mathbf{s}\} = 2\lambda_{\min}(\boldsymbol{\Sigma}) \mathbf{I}_p$, the same argument will result in $\lambda_{\min}(\mathbf{G}) = 0$.

A.1. Method I (Theorem 3)

Proof: To begin with, we compute the sensitivities for $\mathbf{G}' = [\mathbf{X}' \widetilde{\mathbf{X}}']^\top [\mathbf{X}' \widetilde{\mathbf{X}}']$. If $\text{diag}\{\mathbf{s}\} = \lambda_{\min}(\boldsymbol{\Sigma}') \mathbf{I}_p$ then,

$$\mathbf{G}' = \begin{bmatrix} \boldsymbol{\Sigma}' & \boldsymbol{\Sigma}' \\ \boldsymbol{\Sigma}' & \boldsymbol{\Sigma}' \end{bmatrix} - \begin{bmatrix} \mathbf{O}_p & \lambda_{\min}(\boldsymbol{\Sigma}') \mathbf{I}_p \\ \lambda_{\min}(\boldsymbol{\Sigma}') \mathbf{I}_p & \mathbf{O}_p \end{bmatrix}. \quad (17)$$

Suppose \mathbf{X}_1 and \mathbf{X}_2 differ only in one observation. We denote this data point by $\mathbf{x}_1^\top \in \mathbb{R}^p$ in \mathbf{X}_1 which is replaced by \mathbf{x}_2^\top in \mathbf{X}_2 . Let $\mathbf{X}'_1 = \mathbf{X}_1 D_1$ and $\mathbf{X}'_2 = \mathbf{X}_2 D_2$, where D_1 and D_2 are the (positive) diagonal matrices normalizing \mathbf{X}_1 and \mathbf{X}_2 by columns, respectively. In this case we have,

$$\begin{aligned}\mathbf{G}'_1 - \mathbf{G}'_2 &= \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix} \otimes (\boldsymbol{\Sigma}'_1 - \boldsymbol{\Sigma}'_2) \\ &\quad - (\lambda_{\min}(\boldsymbol{\Sigma}'_1) - \lambda_{\min}(\boldsymbol{\Sigma}'_2)) \begin{bmatrix} \mathbf{O}_p & \mathbf{I}_p \\ \mathbf{I}_p & \mathbf{O}_p \end{bmatrix}.\end{aligned}\quad (18)$$

By the triangle inequality we get,

$$\begin{aligned}\|\boldsymbol{\Sigma}'_1 - \boldsymbol{\Sigma}'_2\|_F &\leq \|D_2^\top (\boldsymbol{\Sigma}_1 - \boldsymbol{\Sigma}_2) D_2\| + \\ &\quad \|D^* (D_1^\top \boldsymbol{\Sigma}_1 D_1) D^* - D_1^\top \boldsymbol{\Sigma}_1 D_1\|,\end{aligned}\quad (19)$$

where $D^* = D_2 D_1^{-1}$. Let $C_i = \|\mathbf{X}_1^{(i)}\|_2$. By definition, $D_{1(ii)} = \frac{1}{C_i}$ for $1 \leq i \leq p$. Also, according to Assumption 3 we have $\frac{1}{\sqrt{C_i^2 + B^2}} \leq D_{2(ii)} \leq \frac{1}{\sqrt{C_i^2 - B^2}}$. Therefore, $\frac{C_i}{\sqrt{C_i^2 + B^2}} \leq D^*(ii) \leq \frac{C_i}{\sqrt{C_i^2 - B^2}}$. Let $C_{\min} = \min_{1 \leq i \leq p} \|\mathbf{X}_1^{(i)}\|_2$, then we have

$$\|\boldsymbol{\Sigma}'_1 - \boldsymbol{\Sigma}'_2\|_F \leq \frac{\|\boldsymbol{\Sigma}_1 - \boldsymbol{\Sigma}_2\| + B^2 \|\boldsymbol{\Sigma}'_1\|}{C_{\min}^2 - B^2}. \quad (20)$$

We note, $\|\boldsymbol{\Sigma}_1 - \boldsymbol{\Sigma}_2\|_F = \|\mathbf{x}_1^\top \mathbf{x}_1 - \mathbf{x}_2^\top \mathbf{x}_2\|_F \leq \sqrt{2} B^2$. Hence,

$$\|\boldsymbol{\Sigma}'_1 - \boldsymbol{\Sigma}'_2\|_F \leq \eta^2 (\sqrt{2} + \|\boldsymbol{\Sigma}'_1\|), \quad (21)$$

where $\eta^2 = \frac{B^2}{C_{\min}^2 - B^2}$. Regarding $\lambda_{\min}(\boldsymbol{\Sigma}'_1) - \lambda_{\min}(\boldsymbol{\Sigma}'_2)$, by the triangle inequality we have,

$$\begin{aligned}|\lambda_{\min}(\boldsymbol{\Sigma}'_1) - \lambda_{\min}(\boldsymbol{\Sigma}'_2)| &\leq \\ &|\lambda_{\min}(D_2^\top \boldsymbol{\Sigma}_2 D_2) - \lambda_{\min}(D_2^\top \boldsymbol{\Sigma}_1 D_2)| + \\ &|\lambda_{\min}(D_1^\top \boldsymbol{\Sigma}_1 D_1) - \lambda_{\min}(D^* D_1^\top \boldsymbol{\Sigma}_1 D_1 D^*)|.\end{aligned}\quad (22)$$

Weyl inequalities [16] imply,

$$|\lambda_{\min}(D_2^\top \boldsymbol{\Sigma}_2 D_2) - \lambda_{\min}(D_2^\top \boldsymbol{\Sigma}_1 D_2)| \leq \eta^2, \quad (23)$$

$$|\lambda_{\min}(\boldsymbol{\Sigma}'_1) - \lambda_{\min}(D^* \boldsymbol{\Sigma}'_1 D^*)| \leq \eta^2 \lambda_{\min}(\boldsymbol{\Sigma}'_1). \quad (24)$$

Thus, by summing up (23) and (24) we get

$$|\lambda_{\min}(\boldsymbol{\Sigma}'_1) - \lambda_{\min}(\boldsymbol{\Sigma}'_2)| \leq \eta^2 (1 + \lambda_{\min}(\boldsymbol{\Sigma}'_1)). \quad (25)$$

Therefore, \mathbf{G}' can be $(\epsilon_1 + \epsilon_2, \delta)$ -differentially protected by adding a noise term of the following structure,

$$E = \boldsymbol{\theta}_1 \begin{bmatrix} \mathbf{O}_p & \mathbf{I}_p \\ \mathbf{I}_p & \mathbf{O}_p \end{bmatrix} + \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix} \otimes \boldsymbol{\theta}_2,$$

where $\boldsymbol{\theta}_1$ and $\boldsymbol{\theta}_2$ are determined by Laplace and Gaussian mechanisms (Theorem 2 and 1), respectively.

Regarding the feature-response products $[\mathbf{X}' \widetilde{\mathbf{X}}']^\top \mathbf{y}$, suppose \mathbf{y}_1 is the vector of responses corresponding to \mathbf{X}_1 and \mathbf{y}_2 is the vector of responses corresponding to \mathbf{X}_2 , so they differ only in one element. Now we look at the ℓ_2 -sensitivity. By the triangle inequality we have,

$$\begin{aligned}\|[\mathbf{X}'_1 \widetilde{\mathbf{X}}'_1]^\top \mathbf{y}_1 - [\mathbf{X}'_2 \widetilde{\mathbf{X}}'_2]^\top \mathbf{y}_2\|_2 &\leq \\ &\|[\mathbf{X}'_1 \widetilde{\mathbf{X}}'_1]^\top \mathbf{y}_1 - [\mathbf{X}'_1 \widetilde{\mathbf{X}}'_1]^\top \mathbf{X}_1 \boldsymbol{\beta}\| \\ &+ \|[\mathbf{X}'_2 \widetilde{\mathbf{X}}'_2]^\top \mathbf{y}_2 - [\mathbf{X}'_2 \widetilde{\mathbf{X}}'_2]^\top \mathbf{X}_2 \boldsymbol{\beta}\| \\ &+ \|([\mathbf{X}'_1 \widetilde{\mathbf{X}}'_1]^\top \mathbf{X}_1 - [\mathbf{X}'_2 \widetilde{\mathbf{X}}'_2]^\top \mathbf{X}_2) \boldsymbol{\beta}\|.\end{aligned}\quad (26)$$

We note $\|[\mathbf{X}'_1 \widetilde{\mathbf{X}}'_1]^\top \mathbf{y}_1 - [\mathbf{X}'_1 \widetilde{\mathbf{X}}'_1]^\top \mathbf{X}_1 \boldsymbol{\beta}\| \sim \mathcal{N}(0, \sigma^2 \mathbf{G}'_1)$. Therefore, from the concentration results for the norm of the Gaussian measure [17, 18] we obtain

$$\|[\mathbf{X}'_1 \widetilde{\mathbf{X}}'_1]^\top \mathbf{y}_1 - [\mathbf{X}'_1 \widetilde{\mathbf{X}}'_1]^\top \mathbf{X}_1 \boldsymbol{\beta}\|_2 \leq \sqrt{\zeta \lambda_{\max}(\mathbf{G}'_1)}, \quad (27)$$

with probability at least $1 - \delta_2/2$, where $\zeta = \frac{2p\sigma^2}{1 - \sqrt{\frac{2}{p} \ln(2/\delta_2)}}$.

The same argument as (27) holds for the second term of (26). By adding up the first two terms of the bound in (26), we get

$$\begin{aligned}\zeta^{\frac{1}{2}} \left(\sqrt{\lambda_{\max}(\mathbf{G}'_1)} + \sqrt{\lambda_{\max}(\mathbf{G}'_2)} \right) &\leq \\ &\zeta^{\frac{1}{2}} \left(2\sqrt{\gamma} + \eta (3 + 2\lambda_{\max}(\boldsymbol{\Sigma}'_1) + \lambda_{\min}(\boldsymbol{\Sigma}'_1))^{\frac{1}{2}} \right),\end{aligned}\quad (28)$$

where $\gamma = 2\lambda_{\max}(\boldsymbol{\Sigma}'_1) - \lambda_{\min}(\boldsymbol{\Sigma}'_1)$ and the inequality follows from (15), (25), and the fact that the same argument as (22)-(25), holds for $\lambda_{\max}(\boldsymbol{\Sigma}')$ as well². Regarding the third term of (26) we have,

$$\|([\mathbf{X}'_1 \widetilde{\mathbf{X}}'_1]^\top \mathbf{X}_1 - [\mathbf{X}'_2 \widetilde{\mathbf{X}}'_2]^\top \mathbf{X}_2) \boldsymbol{\beta}\|_2 \leq \quad (29)$$

$$\left\| \begin{bmatrix} D_1 \boldsymbol{\Sigma}_1 - D_2 \boldsymbol{\Sigma}_2 \\ D_1 \boldsymbol{\Sigma}_1 - D_2 \boldsymbol{\Sigma}_2 \end{bmatrix} \right\|$$

$$\begin{aligned}&(\lambda_{\min}(\boldsymbol{\Sigma}'_1) D_1^{-1} - \lambda_{\min}(\boldsymbol{\Sigma}'_2) D_2^{-1}) \begin{bmatrix} \mathbf{O}_p \\ \mathbf{I}_p \end{bmatrix} \|\boldsymbol{\beta}\| \\ &\leq \left[\sqrt{2} \left(\frac{\eta}{B} - \frac{1}{C_{\min}} \right) \|\boldsymbol{\Sigma}_1\| + 2\eta B + (C_{\min} - \frac{B}{\eta}) \lambda_{\min}(\boldsymbol{\Sigma}'_1) \right. \\ &\quad \left. + \eta^2 (\lambda_{\min}(\boldsymbol{\Sigma}'_1) + 1) \sqrt{C_{\min}^2 + B^2} \right] \|\boldsymbol{\beta}\|,\end{aligned}$$

²We also use the inequality $(a+b)^{\frac{1}{2}} \leq a^{\frac{1}{2}} + b^{\frac{1}{2}}$ for $a, b \geq 0$.

where the last inequality is obtained by using the triangle inequality via $D_2 \Sigma_1$ and $\lambda_{\min}(\Sigma'_1) D_2^{-1}$. Thus, the overall ℓ_2 -sensitivity can be computed by adding up the bounds in (28) and (29), which results in the bound (10). The total bound holds with probability at least $1 - \delta_2$. Therefore, $[\mathbf{X} \tilde{\mathbf{X}}]^\top \mathbf{y}$ can be $(\epsilon, \delta_1 + \delta_2)$ -differentially protected by adding a noise vector $\mathbf{e} \sim \mathcal{N}(0, \kappa_2^2)$, where κ_2 is determined by Gaussian mechanism (Theorem 1).

A.2. Method II (Theorem 5)

Proof: Let $\bar{\beta}_1 = \begin{pmatrix} D_1^{-1} \beta \\ \mathbf{0}_{p \times 1} \end{pmatrix}$ and $\hat{\beta}'_1$ denote the OLS estimate based on the data set $[\mathbf{X}'_1, \mathbf{y}_1]$. By the triangle inequality,

$$\|\hat{\beta}'_1 - \hat{\beta}'_2\|_2 \leq \|\hat{\beta}'_1 - \bar{\beta}_1\| + \|\hat{\beta}'_2 - \bar{\beta}_2\| + \|\bar{\beta}_1 - \bar{\beta}_2\|. \quad (30)$$

We note $(\hat{\beta}'_1 - \bar{\beta}_1) \sim \mathcal{N}(0, \sigma^2 \mathbf{G}'_1^{-1})$. Therefore similar to (27) we have $\|\hat{\beta}'_1 - \bar{\beta}_1\| \leq \sqrt{\frac{\zeta}{\lambda_{\min}(\mathbf{G}'_1)}}$ with probability at least $1 - \delta_2/2$. Also,

$$\begin{aligned} \|\bar{\beta}_1 - \bar{\beta}_2\| &\leq \|D_2^{-1} - D_1^{-1}\|_2 \|\beta\|, \\ &\leq (C_{\min} - \sqrt{C_{\min}^2 - B^2}) \|\beta\|. \end{aligned} \quad (31)$$

Hence, if $\lambda_{\min}(\Sigma'_1) > \frac{\eta^2}{1-\eta^2}$, then using (16) and (25) we obtain the ℓ_2 -sensitivity of the OLS estimate,

$$\begin{aligned} \|\hat{\beta}'_1 - \hat{\beta}'_2\| &\leq \sqrt{\frac{\zeta}{\lambda_{\min}(\Sigma'_1)}} + \sqrt{\frac{\zeta}{(1-\eta^2)\lambda_{\min}(\Sigma'_1) - \eta^2}} \\ &\quad + (C_{\min} - \sqrt{C_{\min}^2 - B^2}) \|\beta\|. \end{aligned}$$

B. APPENDIX: FDR CONTROL PROOFS

B.1. Method I (Theorem 4)

Proof: According to the antisymmetry property of \mathbf{f} ,

$$\widehat{\mathbf{W}}_1^F = \mathbf{f}\left(P_F^\top (\mathbf{G} + E) P_F, P_F^\top ([\mathbf{X} \tilde{\mathbf{X}}]^\top \mathbf{y} + \mathbf{e})\right) \quad (32)$$

switches the signs of the components of $\widehat{\mathbf{W}}_1$ corresponding to $i \in F$. From [2] we know $P_F^\top \mathbf{G} P_F = \mathbf{G}$ and $P_F^\top [\mathbf{X} \tilde{\mathbf{X}}]^\top \mathbf{y} \stackrel{d}{=} [\mathbf{X} \tilde{\mathbf{X}}]^\top \mathbf{y}$ for all $F \subseteq \mathcal{S}_0$ where \mathcal{S}_0 denotes the set of null variables. To show the i.i.d. sign property of the nulls, we note that E , \mathbf{e} and \mathbf{y} are mutually independent. We note that $P_F^\top E P_F \stackrel{d}{=} E$ and $P_F^\top \mathbf{e} \stackrel{d}{=} \mathbf{e}$ hold for all $F \subseteq \{1, \dots, p\}$ since the entries that are swapped under P_F are generated i.i.d.. Therefore, $\widehat{\mathbf{W}}_1^F \stackrel{d}{=} \widehat{\mathbf{W}}_1$. This concludes the i.i.d. sign property and FDR control.

B.2. Method II (Theorem 6)

Proof: We note that from [2] we have $P_F^\top \hat{\beta}' \stackrel{d}{=} \hat{\beta}'$ for all $F \subseteq \mathcal{S}_0$. We observe $P_F^\top \mathbf{e} \stackrel{d}{=} \mathbf{e}$ for all $F \subseteq \{1, \dots, p\}$. Independence of $\hat{\beta}'$, and \mathbf{e} will immediately imply the i.i.d. sign property of nulls.

C. REFERENCES

- [1] 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)*, pp. 289–300, 1995.
- [2] Rina Foygel Barber and Emmanuel J Candès, “Controlling the false discovery rate via knockoffs,” *The Annals of Statistics*, vol. 43, no. 5, pp. 2055–2085, 2015.
- [3] Rina Foygel Barber and Emmanuel J Candès, “A knockoff filter for high-dimensional selective inference,” *The Annals of Statistics*, vol. 47, no. 5, pp. 2504–2537, 2019.
- [4] Emmanuel Candès, Yingying Fan, Lucas Janson, and Jinchi Lv, “Panning for gold: ‘model-X’ knockoffs for high dimensional controlled variable selection,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 80, no. 3, pp. 551–577, 2018.
- [5] Yaniv Romano, Matteo Sesia, and Emmanuel Candès, “Deep knockoffs,” *Journal of the American Statistical Association*, pp. 1–12, 2019.
- [6] Yang Lu, Yingying Fan, Jinchi Lv, and William Stafford Noble, “DeepPINK: reproducible feature selection in deep neural networks,” in *Advances in Neural Information Processing Systems*, 2018, pp. 8676–8686.
- [7] Yingying Fan, Jinchi Lv, Mahrad Sharifvaghefi, and Yoshimasa Uematsu, “IPAD: stable interpretable forecasting with knockoffs inference,” *Journal of the American Statistical Association*, pp. 1–13, 2019.
- [8] Robert Tibshirani, “Regression shrinkage and selection via the lasso,” *Journal of the Royal Statistical Society: Series B (Methodological)*, vol. 58, no. 1, pp. 267–288, 1996.
- [9] Abhradeep Guha Thakurta and Adam Smith, “Differentially private feature selection via stability arguments, and the robustness of the lasso,” in *Conference on Learning Theory*. PMLR, 2013, pp. 819–850.
- [10] Cynthia Dwork, Weijie J Su, and Li Zhang, “Differentially private false discovery rate control,” *arXiv:1807.04209*, 2018.
- [11] Cynthia Dwork, Kunal Talwar, Abhradeep Thakurta, and Li Zhang, “Analyze Gauss: optimal bounds for privacy-preserving principal component analysis,” in *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, 2014, pp. 11–20.
- [12] Or Sheffet, “Differentially private ordinary least squares,” in *International Conference on Machine Learning*. PMLR, 2017, pp. 3105–3114.
- [13] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith, “Calibrating noise to sensitivity in private data analysis,” in *Theory of cryptography conference*. Springer, 2006, pp. 265–284.

- [14] Cynthia Dwork and Aaron Roth, “The algorithmic foundations of differential privacy.,” *Foundations and Trends in Theoretical Computer Science*, vol. 9, no. 3-4, pp. 211–407, 2014.
- [15] John R Silvester, “Determinants of block matrices,” *The Mathematical Gazette*, vol. 84, no. 501, pp. 460–467, 2000.
- [16] Roger A. Horn and Charles R. Johnson, *Topics in Matrix Analysis*, Cambridge University Press, Cambridge; New York, 1994.
- [17] Roman Vershynin, *High-dimensional probability: An introduction with applications in data science*, vol. 47, Cambridge university press, 2018.
- [18] Avrim Blum, John Hopcroft, and Ravindran Kannan, “Foundations of data science,” *Vorabversion eines Lehrbuchs*, vol. 5, 2016.