arXiv:2011.03900v2 [stat.ML] 6 Dec 2020

The Cost of Privacy in Generalized Linear Models: Algorithms and Minimax Lower Bounds

T. Tony Cai[1], Yichen Wang[1], and Linjun Zhang[2]

Department of Statistics, University of Pennsylvania[1]

Department of Statistics, Rutgers University[2]

Abstract

We propose differentially private algorithms for parameter estimation in both low-dimensional and high-dimensional sparse generalized linear models (GLMs) by constructing private versions of projected gradient descent. We show that the proposed algorithms are nearly rate-optimal by characterizing their statistical performance and establishing privacy-constrained minimax lower bounds for GLMs. The lower bounds are obtained via a novel technique, which is based on Stein's Lemma and generalizes the tracing attack technique for privacy-constrained lower bounds. This lower bound argument can be of independent interest as it is applicable to general parametric models. Simulated and real data experiments are conducted to demonstrate the numerical performance of our algorithms.

KEYWORDS: Differential privacy; High-dimensional data; Iterative hard thresholding algorithm; Generalized linear models; Optimal rate of convergence; Stein's Lemma

1 Introduction

Statistical and machine learning algorithms are gaining prominence in our daily lives, and so are demands for data privacy guarantees by these algorithms. The need for data privacy protection has in turn inspired the development of formal criteria and frameworks for data privacy, with differential privacy [20, 21] (and its variants [33, 23, 41, 16]) being the most widely studied in theory [24, 22, 26, 1], and adopted in practice [14, 2, 15, 28]. Much of its popularity can be attributed to the ease of building privacy-preserving algorithms that the

differential privacy framework affords [21, 40, 24, 22], but privacy can also come at a cost: it has been observed that requiring algorithms to be differentially private may sacrifice their statistical accuracy [6, 29, 37].

The quest for privacy-preserving yet statistically accurate algorithms has since become a vibrant field of research. On the methodological front, a variety of popular computational and statistical methods has seen their differentially private counterparts, for some examples: causal inference [36, 35], deep learning [1, 46], and multiple testing [27]. On the theoretical side, however, the study of statistical optimality of differentially private algorithms focuses more heavily on the simpler and more stylized problems, such as mean estimation [6, 31, 32], top-k selection [5, 51], and linear regression [11].

In this paper, we study a broadly applicable model, the generalized linear model (GLM) [44, 39], by proposing differentially private algorithms for parameter estimation with theoretical guarantees. We characterize their statistical performance, and prove their near-optimality by establishing minimax lower bounds. In this work, we consider both the classical low-dimensional setting and the contemporary high-dimensional setting.

1.1 Our Contribution and Related Literature

Our main contribution is two-fold: constructing differentially private algorithms for GLM parameter estimation (Section 3), and establishing the near-optimality for the algorithms via privacy-constrained minimax lower bounds for GLM parameter estimation (Section 4).

Private algorithms for GLMs. We construct algorithms, based on noisy gradient descent [8, 7] and noisy iterative hard thresholding [11, 30, 9], for privately estimating the vector of GLM parameters.

There has been an extensive literature on private logistic regression [12, 13, 56], and more broadly, private empirical risk minimization [8, 34, 7]. Our work is inspired by but distinct from previous works in its focus on parameter estimation accuracy, as opposed to excess risk of the solution; for statisticians, parameter estimation accuracy is also arguably a more informative measure of performance than excess risks. Since the log-likelihood function of GLM in general lacks strong convexity [52], bounding the distance between an estimator and the true parameter requires more refined analysis of the algorithms. Theorem 1 shows that (ε, δ)-differentially private estimation of the GLM parameters can be achieved by the noisy gradient descent algorithm (Algorithm 1 based on [8]) with an extra privacy cost of $\tilde{O}\left(\frac{d^2 \log(1/\delta)}{n^2\varepsilon^2}\right)$ in terms of the squared ℓ_2 risk, where n, d respectively denote the sample size and the dimension of the parameter vector.

The difficulty posed by a "flat" log-likelihood landscape [43, 3] is even more salient in the high-dimensional setting: when the number of parameters exceeds the sample size,

strong convexity is categorically impossible for any objective function. We instead leverage the sparsity of the parameter vector to design a noisy iterative hard thresholding algorithm (Algorithm 4), which attains convergence in $O(\log n)$ iterations and incurs an extra privacy cost of $\tilde{O}\left(\frac{(s^* \log d)^2 \log(1/\delta)}{n^2\varepsilon^2}\right)$ in terms of the squared ℓ_2 risk (Theorem 3), where s^* denotes the sparsity of the parameter vector. In particular, the linear dependence on sparsity and logarithmic dependence on the ambient dimension suggest that differentially private estimation remains feasible in high dimensions, which contrasts with the impossibility of high-dimensional estimation [18, 19] under the more restrictive local differential privacy framework [33]. The technical analysis of our algorithm can also be extended to private estimation in other sparse M-estimation problems that enjoy restricted strong convexity and restricted smoothness properties.

Minimax lower bounds. We develop a novel lower bound technique based on Stein's Lemma and show that the statistical accuracy of our algorithms are optimal up to logarithmic factors in the sample size, by establishing privacy-constrained minimax lower bounds for GLM parameter estimation (Theorems 5 and 6).

Our strategy for establishing these lower bounds entails a broad generalization of the "tracing attack" techniques, first developed by [10, 26] and further applied to various statistical problems, including sharp lower bounds for classical Gaussian mean estimation and linear regression [31, 11], as well as lower bounds for sparse mean estimation and linear regression in the high-dimensional setting [51, 11]. In these previous works, the design of tracing attacks seems to be largely ad hoc and catered to specific distribution families such as Gaussian or Beta-Binomial; a general principle for designing attacks has not been observed. Although some promising proposals have been made along this direction [47, 42], it remains unclear whether the suggested attacks in these works actually imply any lower bound results.

In the present paper, we address this problem by proposing a "score attack" based on and named after the score statistic, that is the gradient of the log-likelihood function with respect to the parameter vector. Not only does the score attack imply lower bounds for the GLM problems in the present paper, it also opens paths for lower bound analysis in a much greater range of statistical problems with differential privacy constraints, as the form of our score attack and its theoretical properties (Theorem 4) are applicable to general parametric families of distributions .

1.2 Structure of the Paper

The paper is organized as follows. Section 2 formulates the problem and provides necessary background information. Section 3 describes the algorithms and analyzes in detail their

privacy guarantees as well as statistical accuracy. Section 4 introduces the score attack framework for minimax lower bounds, and applies the framework to establish minimax lower bounds for the GLM problems. Section 5 provides simulated and real data examples that illustrate the numerical performance of our algorithms. Section 6 summarizes our work and discusses its implication for future research. Main technical results are proved in Section 7, and other auxiliary results in the Appendix.

Notation. For real-valued sequences $\{a_n\}, \{b_n\}$, we write $a_n \lesssim b_n$ if $a_n \leq cb_n$ for some universal constant $c \in (0, \infty)$, and $a_n \gtrsim b_n$ if $a_n \geq c'b_n$ for some universal constant $c' \in (0, \infty)$. We say $a_n \asymp b_n$ if $a_n \lesssim b_n$ and $a_n \gtrsim b_n$. $c, C, c_0, c_1, c_2, \cdots$, and so on refer to universal constants in the paper, with their specific values possibly varying from place to place.

For a vector $\boldsymbol{v} \in \mathbb{R}^d$ and a subset $S \subseteq [d]$, we use $\boldsymbol{v}_S$ to denote the restriction of vector $\boldsymbol{v}$ to the index set S. We write $\mathrm{supp}(\boldsymbol{v}) := \{j \in [d] : v_j \neq 0\}$. $\|\boldsymbol{v}\|_p$ denotes the vector ℓ_p norm for $1 \leq p \leq \infty$, with an additional convention that $\|\boldsymbol{v}\|_0$ denotes the number of non-zero coordinates of $\boldsymbol{v}$. For a function $f : \mathbb{R} \to \mathbb{R}$, $\|f\|_\infty$ denotes the the essential supremum of $|f|$. For $t \in \mathbb{R}$ and $R > 0$, let $\Pi_R(t)$ denote the projection of t onto the closed interval $[-R, R]$. For a random variable X, we use $\operatorname{ess\,sup}(X) = \inf\{c : \mathbb{P}(X < c) = 1\}$ to denote the essential supremum of X.

2 Problem Formulation

In this section, we present a detailed description of the scope of statistical models (generalized linear models) and algorithms (differentially private algorithms) to be studied in this paper, and formally define the "cost of privacy" in terms of minimax risks.

2.1 Generalized Linear Models

We study parameter estimation in generalized linear models. In a generalized linear model, the response variable $y \in \mathbb{R}$, conditional on the design vector $\boldsymbol{x} \in \mathbb{R}^d$, follows a distribution of the natural exponential family form,

$$f_{\boldsymbol{\beta}^*}(y|\boldsymbol{x}) = h(y, \sigma) \exp\left(\frac{(\boldsymbol{x}^\top \boldsymbol{\beta}^*) g(y) - \psi(\boldsymbol{x}^\top \boldsymbol{\beta}^*)}{c(\sigma)}\right), \tag{2.1}$$

where $c(\sigma)$ is a nuisance scale parameter and $\psi(\cdot)$ is the cumulant generating function of y given $\boldsymbol{x}$. The generalized linear model is, first of all, a generalization of the linear model: setting $g(y) = y$, $\psi(u) = u^2/2$ and $c(\sigma) = \sigma^2$ in (2.1) recovers the (Gaussian) linear model.

Model (2.1) also subsumes other special cases such as logistic and multinomial regression.

Throughout the paper, our goal is estimating $\boldsymbol{\beta}^* \in \mathbb{R}^d$ using an i.i.d. sample $\{(y_i, \boldsymbol{x}_i)\}_{i\in[n]}$ drawn from the model (2.1). We shall consider both the classical setting, where the dimension d is dominated by the sample size n, and the high-dimensional sparse setting where d potentially dominates n but only a small proportion of $\boldsymbol{\beta}^*$'s coordinates are non-zero.

In either case, the issue of data privacy is relevant, as any nontrivial estimator of $\boldsymbol{\beta}^*$ must take the data $\{(y_i, \boldsymbol{x}_i)\}_{i\in[n]}$ as input. Before considering concrete estimators and their performance, let us first define the desired criteria of privacy protection.

2.2 Differential Privacy

Intuitively speaking, an algorithm M applied over a data set $\boldsymbol{X}$ compromises data privacy if an adversary is able to correctly infer from the algorithm's output $M(\boldsymbol{X})$ whether an individual datum $\boldsymbol{x}$ belongs to $\boldsymbol{X}$ or not.

The notion of differential privacy formalizes this idea by requiring that, for every pair of data sets $\boldsymbol{X}$ and $\boldsymbol{X}'$ that differ by a single datum, hereafter called "adjacent data sets", the algorithm M is randomized so that the distributions of $M(\boldsymbol{X})$ and of $M(\boldsymbol{X}')$ are close to each other.

Definition 1 (Differential Privacy, [21]). *A randomized algorithm $M : \mathcal{X}^n \to \mathcal{R}$ is (ε, δ)-differentially private if for every pair of adjacent data sets $\boldsymbol{X}, \boldsymbol{X}' \in \mathcal{X}^n$ that differ by one individual datum and every (measurable) $S \subseteq \mathcal{R}$,*

$$\mathbb{P}\left(M(\boldsymbol{X}) \in S\right) \leq e^{\varepsilon} \cdot \mathbb{P}\left(M(\boldsymbol{X}') \in S\right) + \delta,$$

where the probability measure $\mathbb{P}$ is induced by the randomness of M only.

The definition guarantees that, for small values of $\varepsilon, \delta \geq 0$, the distributions of $M(\boldsymbol{X})$ and $M(\boldsymbol{X}')$ are almost indistinguishable. But beyond its strong privacy guarantees, the notion of differential privacy is desirable also for the ease and flexibility of constructing differentially private algorithms. We summarize here some useful facts for our construction of algorithms in this paper.

First, a large class of non-private algorithms can be made differentially private via random perturbations.

Fact 1 (The Laplace and Gaussian mechanisms, [21, 22]). *Let $M : \mathcal{X}^n \to \mathbb{R}^d$ be an algorithm that is not necessarily differentially private.*

- *Suppose $\sup_{\boldsymbol{X},\boldsymbol{X}' adjacent} \|M(\boldsymbol{X}) - M(\boldsymbol{X}')\|_1 < B < \infty$. For $\boldsymbol{w} \in \mathbb{R}^d$ with its coordinates $w_1, w_2, \cdots, w_d \overset{i.i.d.}{\sim}$ Laplace(B/ε), $M(\boldsymbol{X}) + \boldsymbol{w}$ is $(\varepsilon, 0)$-differentially private.*

- *If instead we have* $\sup_{\boldsymbol{X},\boldsymbol{X}' adjacent} \|M(\boldsymbol{X}) - M(\boldsymbol{X}')\|_2 < B < \infty$*, for* $\boldsymbol{w} \sim N_d(\boldsymbol{0}, \sigma^2 \boldsymbol{I})$ *with* $\sigma^2 = 2B^2 \log(2/\delta)/\varepsilon$*,* $M(\boldsymbol{X}) + \boldsymbol{w}$ *is* (ε, δ)*-differentially private.*

That is, if a non-private algorithm's output is not too sensitive to changing any single datum in the input data set, perturbing the algorithm with Laplace or Gaussian noises produces a differentially private algorithm.

Second, differential privacy is preserved under compositions, albeit with weaker privacy parameters. The composition theorems, stated below, explicitly quantify how privacy parameters degrade as private algorithms are composited.

Fact 2 (Composition Theorems)**.** *Consider an* (ε, δ)*-differentially private algorithm* $M_0 : \mathcal{X}^n \to \mathcal{R}_0$ *and* (ε, δ)*-differentially private algorithms* $M_i : \mathcal{X}^n \times \mathcal{R}_{i-1} \to \mathcal{R}_i$ *for* $i = 1, 2, \cdots, k-1$*. Consider the composite algorithm* $M = M_{k-1} \circ M_{k-2} \circ \cdots \circ M_0$*.*

- *Composition theorem [21].* M *is* $(k\varepsilon, k\delta)$*-differentially private.*
- *Advanced composition [24]. For every* $\delta' > 0$*,* M *is* $(\sqrt{2k \log(1/\delta')}\varepsilon + k(e^\varepsilon - 1)\varepsilon, k\delta + \delta')$*-differentially private.*

It is worth noting that the notion of "composition" considered here, termed "k-fold adaptive composition" in the literature, is more general than the composition of functions in the usual sense: each part of the composite algorithm may access the same data set, or a different data set, after receiving the output from its previous part. In fact, if a differentially private algorithm is simply post-processed independent of data, there is no deterioration of privacy whatsoever.

Fact 3 (Post-processing [21, 55])**.** *Consider an* (ε, δ)*-differentially private algorithm* $M : \mathcal{X}^n \to \mathcal{R}$*. If* g *is a measurable function, then* $g(M)$ *is also* (ε, δ)*-differentially private.*

The composition and post-processing properties will be particularly useful for analyzing the privacy guarantees of the iterative algorithms considered in Section 3.

2.3 The Cost of Privacy in Generalized Linear Models

Once an algorithm is known to be differentially private, it is natural to ask whether the privacy guarantees come at the expense of accuracy: as seen in Fact 1, often random perturbations are introduced to achieve differential privacy. In this paper, we assess the accuracy of algorithms via the lens of *minimax risk*, defined as follows.

Let $\{f_{\boldsymbol{\theta}} : \boldsymbol{\theta} \in \Theta\}$ be a family of statistical models supported over $\mathcal{X}$ and indexed by the parameter $\boldsymbol{\theta}$. Let $\boldsymbol{X} = \{\boldsymbol{x}_1, \boldsymbol{x}_2, \cdots, \boldsymbol{x}_n\}$ be an i.i.d. sample drawn from $f_{\boldsymbol{\theta}^*}$ for some

unknown $\boldsymbol{\theta}^* \in \Theta$, $M : \mathcal{X}^n \to \Theta$ be an estimator. Let $\ell : \Theta \times \Theta \to \mathbb{R}_+$ be a metric on Θ and $\rho : \mathbb{R}_+ \to \mathbb{R}_+$ be an increasing function. The (statistical) risk of M is given by $\mathbb{E}\rho(\ell(M(\boldsymbol{X}), \boldsymbol{\theta}^*))$, where the expectation is taken over the data distribution $f_{\boldsymbol{\theta}^*}$ and the randomness of estimator M.

Because the risk $\mathbb{E}\rho(\ell(M(\boldsymbol{X}), \boldsymbol{\theta}^*))$ depends on the unknown $\boldsymbol{\theta}^*$ and can be trivially minimized by setting $M(\boldsymbol{X}) \equiv \boldsymbol{\theta}^*$, a more sensible measure of performance is the maximum risk over the entire class of distributions $\{f_{\boldsymbol{\theta}} : \boldsymbol{\theta} \in \Theta\}$, $\sup_{\boldsymbol{\theta} \in \Theta} \mathbb{E}\rho(\ell(M(\boldsymbol{X}), \boldsymbol{\theta}))$.

The minimax risk of estimating $\boldsymbol{\theta} \in \Theta$ is then given by

$$\inf_{M} \sup_{\boldsymbol{\theta} \in \Theta} \mathbb{E}\rho(\ell(M(\boldsymbol{X}), \boldsymbol{\theta})). \tag{2.2}$$

By definition, this quantity characterizes the best possible worst-case performance that an estimator can hope to achieve over the class of models $\{f_{\boldsymbol{\theta}} : \boldsymbol{\theta} \in \Theta\}$.

In this paper, we study a *privacy-constrained* version of the minimax risk: let $\mathcal{M}_{\varepsilon,\delta}$ be the collection of all (ε, δ)-differentially private algorithms mapping from $\mathcal{X}^n$ to Θ, we consider

$$\inf_{M \in \mathcal{M}_{\varepsilon,\delta}} \sup_{\boldsymbol{\theta} \in \Theta} \mathbb{E}\rho(\ell(M(\boldsymbol{X}), \boldsymbol{\theta})). \tag{2.3}$$

As $\mathcal{M}_{\varepsilon,\delta}$ is a proper subset of all possible estimators, the privacy-constrained minimax risk as defined above will be at least as large as the unconstrained minimax risk, with the difference between these two minimax risks, (2.2) and (2.3) being the "cost of privacy".

In our case, the statistical models of interest are the generalized linear models (2.1) indexed by the parameter vector $\boldsymbol{\beta}^*$, and we would like to precisely characterize the cost of privacy in GLM parameter estimation problems. This goal will be achieved in two steps: Section 3 provides upper bounds of the privacy-constrained minimax risk (2.3) via analysis of differentially private algorithms; Section 4 establishes corresponding lower bounds of the privacy-constrained minimax risk.

3 Differentially Private Algorithms for GLMs

In this section, we develop differentially private algorithms for estimating parameters $\boldsymbol{\beta}^* \in \mathbb{R}^d$ in the generalized linear model

$$f_{\boldsymbol{\beta}^*}(y|\boldsymbol{x}) = h(y, \sigma) \exp\left(\frac{\boldsymbol{x}^\top \boldsymbol{\beta}^* g(y) - \psi(\boldsymbol{x}^\top \boldsymbol{\beta}^*)}{c(\sigma)}\right); \boldsymbol{x} \sim f_{\boldsymbol{x}}. \tag{3.1}$$

With an i.i.d. sample $\boldsymbol{Z} = \{\boldsymbol{z}_i\}_{i\in[n]} = \{(y_i, \boldsymbol{x}_i)\}_{i\in[n]}$ drawn from the model (3.1), the general approach is to minimize the following (scaled) negative log-likelihood function in a differentially private fashion:

$$\mathcal{L}_n(\boldsymbol{\beta}; \boldsymbol{Z}) = \frac{1}{n}\sum_{i=1}^{n} \left(\psi(\boldsymbol{x}_i^\top \boldsymbol{\beta}) - g(y_i)\boldsymbol{x}_i^\top \boldsymbol{\beta}\right). \tag{3.2}$$

We may write $\mathcal{L}_n(\boldsymbol{\beta})$ as a shorthand when the relevant data set is unambiguous.

Since $\mathcal{L}_n$ is convex in $\boldsymbol{\beta}$, the problem is an instance of differentially private convex optimization, for which there are many well-studied methods. Roughly speaking, these methods can be organized into two categories depending on the form of random perturbations involved: "one-shot" methods [12, 13, 34] in which random noises are added only once to the objective function or before reporting the final solution, or iterative, gradient-descent type methods [8, 7] in which random noises are added to each iteration of the algorithm.

As discussed in Section 1.1, existing convergence results for these methods are focused on the excess risk of a differentially private minimizer $\boldsymbol{\beta}^{\text{priv}}$ of (3.2), compared to the non-private solution $\hat{\boldsymbol{\beta}}$, that is, $\mathbb{E}\mathcal{L}_n(\boldsymbol{\beta}^{\text{priv}}) - \mathbb{E}\mathcal{L}_n(\hat{\boldsymbol{\beta}})$. The lack of strong convexity in the generalized linear model (3.1) precludes the possibility of obtaining bounds for parameter estimation from excess risk bounds. For example, for the logistic regression model, which is obtained from (3.1) by setting $g(y) = y$ and $\psi(u) = \log(1 + e^u)$. The Hessian of $\mathcal{L}_n$,

$$\nabla^2\mathcal{L}_n(\boldsymbol{\beta}) = \frac{1}{n}\sum_{i=1}^{n} \psi''(\boldsymbol{x}_i^\top \boldsymbol{\beta})\boldsymbol{x}_i\boldsymbol{x}_i^\top = \frac{1}{n}\sum_{i=1}^{n} \frac{e^{\boldsymbol{x}_i^\top \boldsymbol{\beta}}}{(1 + e^{\boldsymbol{x}_i^\top \boldsymbol{\beta}})^2}\boldsymbol{x}_i\boldsymbol{x}_i^\top,$$

has its smallest eigenvalue approaching 0 as $\|\boldsymbol{\beta}\|_2 \to \infty$ even in the favorable setting where n is much greater than the dimension of $\boldsymbol{\beta}$. When n is dominated by the dimension, $\nabla^2\mathcal{L}_n(\boldsymbol{\beta})$ is simply rank-deficient and therefore cannot be positive-definite. The absence of strong convexity in GLMs also implies that the "one-shot" algorithms are not guaranteed to be differentially private (see [13, 34] and the references therein), unless a quadratic penalty term is added to $\mathcal{L}_n$.

Our approach, then, is to consider gradient descent type algorithms. Although strong convexity fails to hold for $\mathcal{L}_n$ globally, it turns out that $\mathcal{L}_n$ satisfies a "restricted" and "local" sense of strong convexity [43], to be made precise in Section 3.1, is sufficient for the noisy gradient descent algorithm to enjoy fast convergence and optimal statistical accuracy. In Section 3.1, we analyze in detail the privacy guarantees and convergence rates of the noisy gradient descent algorithm, which works well for the classical setting of $d = o(n)$.

The high-dimensional, $d \gtrsim n$ setting is considered in Section 3.2. In this case, consis-

tent estimation of $\boldsymbol{\beta}^*$ is not possible even without privacy constraints, unless additional assumptions such as sparsity of $\boldsymbol{\beta}^*$ are imposed. When $\boldsymbol{\beta}^*$ is indeed sparse, we introduce a noisy iterative hard thresholding algorithm that allows the random perturbations to scale with the sparsity ("intrinsic dimension") of $\boldsymbol{\beta}^*$ rather than the ambient dimension d, thereby achieving the optimal statistical accuracy with privacy constraints.

3.1 The Classical Low-dimensional Setting

We first consider the classical low-dimensional setting where $d = o(n)$. For minimizing the negative GLM log-likelihood

$$\mathcal{L}_n(\boldsymbol{\beta}; \boldsymbol{Z}) = \frac{1}{n} \sum_{i=1}^{n} \left(\psi(\boldsymbol{x}_i^\top \boldsymbol{\beta}) - g(y_i) \boldsymbol{x}_i^\top \boldsymbol{\beta} \right)$$

in a differentially private fashion, we consider the noisy gradient descent algorithm, first proposed by [8] in its generic form for arbitrary convex functions. The following algorithm is a specialization the generic algorithm to GLMs.

Algorithm 1: Differentially Private Generalized Linear Regression

Input : $\mathcal{L}_n(\boldsymbol{\beta}, \boldsymbol{Z})$, data set $\boldsymbol{Z}$, step size η^0, privacy parameters ε, δ, noise scale B, number of iterations T, truncation parameter R, initial value $\boldsymbol{\beta}^0 \in \mathbb{R}^d$.

for t *in* 0 **to** $T-1$ **do**

  Generate $\boldsymbol{w}_t \in \mathbb{R}^d$ with $w_{t1}, w_{t2}, \cdots, w_{td} \overset{\text{i.i.d.}}{\sim} N\left(0, (\eta^0)^2 2B^2 \frac{d \log(2T/\delta)}{n^2(\varepsilon/T)^2}\right)$;

  Compute $\boldsymbol{\beta}^{t+1} = \boldsymbol{\beta}^t - (\eta_0/n) \sum_{i=1}^{n} (\psi'(\boldsymbol{x}_i^\top \boldsymbol{\beta}^t) - \Pi_R(y_i)) \boldsymbol{x}_i + \boldsymbol{w}_t$;

end

Output: $\boldsymbol{\beta}^{(T)}$.

Before delving into the analysis of its privacy guarantees and convergence rates, we collect some necessary assumptions here for the clarity of ensuing technical results.

(D1) Bounded design: there is a constant $\sigma_{\boldsymbol{x}} < \infty$ such that $\|\boldsymbol{x}\|_2 < \sigma_{\boldsymbol{x}} \sqrt{d}$ almost surely.

(D2) Bounded moments of design: $\mathbb{E}\boldsymbol{x} = \boldsymbol{0}$, and the covariance matrix $\Sigma_{\boldsymbol{x}} = \mathbb{E}\boldsymbol{x}\boldsymbol{x}^\top$ satisfies $0 < 1/C < \lambda_{\min}(\Sigma_{\boldsymbol{x}}) \leq \lambda_{\max}(\Sigma_{\boldsymbol{x}}) < C$ for some constant $0 < C < \infty$.

(G1) The function ψ in the GLM (3.1) satisfies $\|\psi'\|_\infty < c_1$ for some constant $c_1 < \infty$.

(G2) The function ψ satisfies $\|\psi''\|_\infty < c_2$ for some constant $c_2 < \infty$.

These assumptions are comparable to those required for the theoretical analysis of GLMs in the non-private setting; for examples, see [43, 38, 54] and the references therein.

Let us first consider the privacy guarantees of Algorithm 1. Because the algorithm is a composition of T individual steps, if each step is $(\varepsilon/T, \delta/T)$-differentially private, the overall algorithm would be (ε, δ)-differentially private in view of Fact 2. This is indeed the case under appropriate assumptions.

Lemma 1. *If assumptions (D1) and (G1) hold, then choosing* $B = 4(R + c_1)\sigma_{\boldsymbol{x}}$ *guarantees that Algorithm 1 is* (ε, δ)*-differentially private.*

Lemma 1 is proved in Section A.1. Although the privacy guarantee holds for any number of iterations T, choosing T properly has significant implications for the accuracy of Algorithm 1, as a larger value of T introduces greater noises into the algorithm in order to achieve privacy.

Existing results on noisy gradient descent typically call for $O(n)$ [7] or $O(n^2)$ [8] iterations for minimizing generic convex functions. For the GLM problem, we shall show that $O(\log n)$ iterations suffice, thanks to the restricted strong convexity and restricted smoothness of generalized linear models.

Fact 4 ([38], Proposition 1 paraphrased)**.** *If assumptions (D1) and (D2) hold, there is a constant* $\alpha > 0$ *that depends on* $\sigma_{\boldsymbol{x}}, C, \psi$ *and satisfies*

$$\langle \nabla \mathcal{L}_n(\boldsymbol{\beta}_1) - \nabla \mathcal{L}_n(\boldsymbol{\beta}_2), \boldsymbol{\beta}_1 - \boldsymbol{\beta}_2 \rangle \geq \begin{cases} \alpha \|\boldsymbol{\beta}_1 - \boldsymbol{\beta}_2\|_2^2 - \frac{c^2 \sigma_{\boldsymbol{x}}^2}{2\alpha} \frac{\log d}{n} \|\boldsymbol{\beta}_1 - \boldsymbol{\beta}_2\|_1^2 & \text{if } \|\boldsymbol{\beta}_1 - \boldsymbol{\beta}_2\|_2 \leq 3, \\ 3\alpha \|\boldsymbol{\beta}_1 - \boldsymbol{\beta}_2\|_2 - \sqrt{2} c \sigma_{\boldsymbol{x}} \sqrt{\frac{\log d}{n}} \|\boldsymbol{\beta}_1 - \boldsymbol{\beta}_2\|_1 & \text{if } \|\boldsymbol{\beta}_1 - \boldsymbol{\beta}_2\|_2 > 3, \end{cases} \tag{3.3}$$

with probability at least $1 - c_3 \exp(-c_4 n)$*. If we further assume (G2), there is a constant* $\gamma \geq \alpha > 0$ *that depends on* $\sigma_{\boldsymbol{x}}, M, c_2$ *and satisfies*

$$\langle \nabla \mathcal{L}_n(\boldsymbol{\beta}_1) - \nabla \mathcal{L}_n(\boldsymbol{\beta}_2), \boldsymbol{\beta}_1 - \boldsymbol{\beta}_2 \rangle \leq \gamma \|\boldsymbol{\beta}_1 - \boldsymbol{\beta}_2\|_2^2 + \frac{4\gamma}{3} \frac{\log d}{n} \|\boldsymbol{\beta}_1 - \boldsymbol{\beta}_2\|_1^2. \tag{3.4}$$

with probability at least $1 - c_3 \exp(-c_4 n)$*.*

These weaker versions of strong convexity and smoothness, as it turns out, are sufficient for Algoirthm 1 to attain linear convergence, which is the same rate for minimizing strongly convex and smooth functions, and cannot be further improved in general [45]. Therefore, $O(\log n)$ iterations would allow the algorithm to converge to an accuracy of $O(n^{-1})$ within $\hat{\boldsymbol{\beta}}$, the true minimizer of $\mathcal{L}_n$, in terms of squared ℓ_2 norm; as $\mathbb{E}\|\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*\|_2^2$ is of order d/n, from a statistical perspective, there is little reason to run the algorithm further than $O(\log n)$ iterations.

Theorem 1. *Let $\{(y_i, \boldsymbol{x}_i)\}_{i\in[n]}$ be an i.i.d. sample from the GLM (3.1). Suppose assumptions (D1), (D2), (G1) and (G2) are true. Let the parameters of Algorithm 1 be chosen as follows.*

- *Set step size $\eta^0 = 3/4\gamma$, where γ is the smoothness constant defined in Fact 4.*
- *Set $R = \min\left(\operatorname{ess\,sup}|y_1|, c_1 + \sqrt{2c_2 c(\sigma)\log n}\right) \lesssim \sqrt{c(\sigma)\log n}$.*
- *Noise scale B. Set $B = 4(R + c_1)\sigma_{\boldsymbol{x}}$.*
- *Number of iterations T. Let $T = (2\gamma/\alpha)\log(9n)$, where α, γ are the strong convexity and smoothness constants defined in in Fact 4.*
- *Initialization $\boldsymbol{\beta}^0$. Choose $\boldsymbol{\beta}^0$ so that $\|\boldsymbol{\beta}^0 - \hat{\boldsymbol{\beta}}\|_2 \le 3$, where $\hat{\boldsymbol{\beta}} = \arg\min \mathcal{L}_n(\boldsymbol{\beta}; Z)$.*

If $n \ge K \cdot \left(Rd\sqrt{\log(1/\delta)}\log n \log\log n/\varepsilon\right)$ for a sufficiently large constant K, the output of Algorithm 1 satisfies

$$\|\boldsymbol{\beta}^{(T)} - \boldsymbol{\beta}^*\|_2^2 \lesssim c(\sigma)\left(\frac{d}{n} + \frac{d^2\log(1/\delta)\log^3 n}{n^2\varepsilon^2}\right), \tag{3.5}$$

with probability at least $1 - c_3\exp(-c_4 n) - c_3\exp(-c_4 d) - c_3\exp(-c_4\log n)$.

Theorem 1 is proved in Section A.2. Some further comments may help clarify the theorem. For the choice of various algorithm tuning parameters, we note that the step size, number of iterations and initialization are chosen to assure convergence; in particular the initialization condition, as in [38], is standard in the literature and can be extended to $\|\boldsymbol{\beta}^0 - \hat{\boldsymbol{\beta}}\|_2 \le 3\max(1, \|\boldsymbol{\beta}^*\|_2)$.

The choice of truncation level R is to ensure privacy while keeping as many data intact as possible; when the distribution of y has bounded support, for example in the logistic model, it can be chosen to be an $O(1)$ constant and therby saving an extra factor of $O(\log n)$ in the second term of (3.5). The choice of B which depends on R then ensures the privacy of Algorithm 1 as seen in Lemma 1.

Finally, the scaling of n versus d, ε and δ in Theorem 1 is nearly optimal, as our lower bound result, Theorem 5, shall imply that no estimator can achieve low ℓ_2 error unless the assumed scaling holds, and that the statistical accuracy of Algorithm 1 cannot be further improved except possibly for factors of $\log n$.

3.2 The High-dimensional Sparse Setting

In this section, we construct differentially private algorithms for estimating GLM parameters when the dimension d dominates the sample size n. In this setting, even without

privacy requirements, directly minimizing the negative log-likelihood function $\mathcal{L}_n(\boldsymbol{\beta})$ no longer achieves any meaningful statistical accuracy, because the objective function $\mathcal{L}_n$ can have infinitely many minimizers due to a rank-deficient Hessian matrix $\nabla^2 \mathcal{L}_n(\boldsymbol{\beta}) = \frac{1}{n}\sum_{i=1}^n \psi''(\boldsymbol{x}_i^\top \boldsymbol{\beta}) \boldsymbol{x}_i \boldsymbol{x}_i^\top$.

The problem is nevertheless solvable when the true parameter vector $\boldsymbol{\beta}^*$ is s^*-sparse with $s^* = o(d)$, that is when at most s^* out of d coordinates of $\boldsymbol{\beta}^*$ are non-zero. For estimating a sparse $\boldsymbol{\beta}^*$, the primary challenge lies in (approximately) solving the non-convex optimization problem $\hat{\boldsymbol{\beta}} = \arg\min_{\boldsymbol{\beta}:\|\boldsymbol{\beta}\|_0 \le s^*} \mathcal{L}_n(\boldsymbol{\beta}; \boldsymbol{Z})$. Some popular non-private approaches include convex relaxation via ℓ_1 regularization of $\mathcal{L}_n$ [43, 3], or projected gradient descent onto the non-convex feasible set $\{\boldsymbol{\beta} : \|\boldsymbol{\beta}\|_0 \le s^*\}$, also known as iterative hard thresholding [9, 30]:

Algorithm 2: Iterative Hard Thresholding (IHT)

Input : Objective function $f(\boldsymbol{\theta})$, sparsity s, step size η, number of iterations T.

Initialize $\boldsymbol{\theta}^0$ with $\|\boldsymbol{\theta}^0\|_0 \le s$, set $t = 0$;

for t *in* 0 **to** $T-1$ **do**

  $\boldsymbol{\theta}^{t+1} = P_s\left(\boldsymbol{\theta}^t - \eta \nabla f(\boldsymbol{\theta}^t)\right)$, where $P_s(\boldsymbol{v}) = \arg\min_{\boldsymbol{z}:\|\boldsymbol{z}\|_0 = s} \|\boldsymbol{v} - \boldsymbol{z}\|_2^2$;

end

Output: $\boldsymbol{\theta}^{(T)}$.

In each iteration, the algorithm updates the solution via gradient descent, keeps its largest s coordinates in magnitude, and sets the other coordinates to 0.

For privately fitting high-dimensional sparse GLMs, we shall construct a noisy version of Algorithm 2, and show in Section 3.2.2 that it again enjoys a linear rate of convergence, as a consequence of Fact 4 and sparsity of $\boldsymbol{\beta}^*$. As a first step towards this goal, we consider in Section 3.2.1 a noisy, differentially private version of the projection operator P_s, as well as a noisy iterative hard thresholding algorithm applicable to any objective function that satisifies restricted strong convexity and restricted smoothness.

3.2.1 The Noisy Iterative Hard Thresholding Algorithm

At the core of our algoirthm is a noisy, differentially private algorithm that identifies the top-s largest coordinates of a given vector with good accuracy. The following "Peeling" algorithm [27] serves this purpose, with fresh Laplace noises added to the underlying vector and one coordinate "peeled" from the vector in each iteration.

The algorithm is guaranteed to be (ε, δ)-differentially private when the vector-valued function $\boldsymbol{v}(\boldsymbol{Z})$ is not sensitive to replacing any single datum.

Algorithm 3: Noisy Hard Thresholding (NoisyHT)

Input : vector-valued function $\boldsymbol{v} = \boldsymbol{v}(\boldsymbol{Z}) \in \mathbb{R}^d$, data $\boldsymbol{Z}$, sparsity s, privacy parameters ε, δ, noise scale λ.

Initialize $S = \emptyset$;

for i *in* 1 **to** s **do**

  Generate $\boldsymbol{w}_i \in \mathbb{R}^d$ with $w_{i1}, w_{i2}, \cdots, w_{id} \overset{\text{i.i.d.}}{\sim} \text{Laplace}\left(\lambda \cdot \frac{2\sqrt{3s\log(1/\delta)}}{\varepsilon}\right)$;

  Append $j^* = \arg\max_{j \in [d] \setminus S} |v_j| + w_{ij}$ to S;

end

Set $\tilde{P}_s(\boldsymbol{v}) = \boldsymbol{v}_S$;

Generate $\tilde{\boldsymbol{w}}$ with $\tilde{w}_1, \cdots, \tilde{w}_d \overset{\text{i.i.d.}}{\sim} \text{Laplace}\left(\lambda \cdot \frac{2\sqrt{3s\log(1/\delta)}}{\varepsilon}\right)$;

Output: $\tilde{P}_s(\boldsymbol{v}) + \tilde{\boldsymbol{w}}_S$.

Lemma 2 ([27, 11])**.** *If for every pair of adjacent data sets $\boldsymbol{Z}, \boldsymbol{Z}'$ we have $\|\boldsymbol{v}(\boldsymbol{Z}) - \boldsymbol{v}(\boldsymbol{Z}')\|_\infty < \lambda$, then NoisyHT is an (ε, δ)-differentially private algorithm.*

The accuracy of Algorithm 3 is quantified by the next lemma.

Lemma 3. *Let $\tilde{P}_s$ be defined as in Algorithm* 3*. For any index set I, any $\boldsymbol{v} \in \mathbb{R}^I$ and $\hat{\boldsymbol{v}}$ such that $\|\hat{\boldsymbol{v}}\|_0 \le \hat{s} \le s$, we have that for every $c > 0$,*

$$\|\tilde{P}_s(\boldsymbol{v}) - \boldsymbol{v}\|_2^2 \le (1 + 1/c)\frac{|I| - s}{|I| - \hat{s}}\|\hat{\boldsymbol{v}} - \boldsymbol{v}\|_2^2 + 4(1 + c)\sum_{i \in [s]} \|\boldsymbol{w}_i\|_\infty^2.$$

Lemma 3 is proved in Section A.3. In comparison, the exact, non-private projection operator P_s satisfies ([30], Lemma 1)

$$\|P_s(\boldsymbol{v}) - \boldsymbol{v}\|_2^2 \le \frac{|I| - s}{|I| - \hat{s}}\|\hat{\boldsymbol{v}} - \boldsymbol{v}\|_2^2.$$

Algorithm 3, therefore, is as accurate as its non-private counterpart up to a constant multiplicative factor and some additive noise. Taking the private top-s projection algorithm, we have the following noisy iterative hard thresholding algorithm.

Compared to the non-private Algorithm 2, we simply replaced the exact projection P_s with the noisy projection given by Algorithm 3. The privacy guarantee of Algorithm 4 is then inherited from that of Algorithm 3.

Lemma 4. *If for every pair of adjacent data $\boldsymbol{z}, \boldsymbol{z}'$ and every $\boldsymbol{\theta} \in \Theta$ we have $\|\nabla l(\boldsymbol{\theta}; \boldsymbol{z}) - \nabla l(\boldsymbol{\theta}; \boldsymbol{z}')\|_\infty < B$, then NoisyIHT is an (ε, δ)-differentially private algorithm.*

Algorithm 4: Noisy Iterative Hard Thresholding (NoisyIHT)

Input : Objective function $\mathcal{L}_n(\boldsymbol{\theta}, \boldsymbol{Z}) = n^{-1}\sum_{i=1}^n l(\boldsymbol{\theta}, \boldsymbol{z}_i)$, data set $\boldsymbol{Z}$, sparsity level s, step size η^0, privacy parameters ε, δ, noise scale B, number of iterations T.

Initialize $\boldsymbol{\theta}^0$ with $\|\boldsymbol{\theta}^0\|_0 \leq s$, set $t = 0$;

for t *in* 0 **to** $T-1$ **do**

  $\boldsymbol{\theta}^{t+1} = \texttt{NoisyHT}\left(\boldsymbol{\theta}^t - \eta^0 \nabla \mathcal{L}_n(\boldsymbol{\theta}^t; \boldsymbol{Z}), \boldsymbol{Z}, s, \varepsilon/T, \delta/T, (\eta^0/n)B\right)$;

end

Output: $\boldsymbol{\theta}^{(T)}$.

The lemma is proved in Section A.4. Similar to the noisy gradient descent (Algorithm 1), the privacy guarantee of Algorithm 4 is valid for any choice of T, however a fast rate of convergence would allow us to select a small T and thereby introducing less noise into the algorithm. To our delight, restricted strong convexity and restricted smoothness again lead to a linear rate of convergence even in the high-dimensional sparse setting.

Theorem 2. *Let* $\hat{\boldsymbol{\theta}} = \arg\min_{\|\boldsymbol{\theta}\|_0 \leq s^*} \mathcal{L}_n(\boldsymbol{\theta}; \boldsymbol{Z})$. *For iteration number* $t \geq 0$, *suppose*

$$\langle \nabla \mathcal{L}_n(\boldsymbol{\theta}^t) - \nabla \mathcal{L}_n(\hat{\boldsymbol{\theta}}), \boldsymbol{\theta}^t - \hat{\boldsymbol{\theta}} \rangle \geq \alpha \|\boldsymbol{\theta}^t - \hat{\boldsymbol{\theta}}\|_2^2 \tag{3.6}$$

$$\langle \nabla \mathcal{L}_n(\boldsymbol{\theta}^{t+1}) - \nabla \mathcal{L}_n(\hat{\boldsymbol{\theta}}), \boldsymbol{\theta}^{t+1} - \hat{\boldsymbol{\theta}} \rangle \leq \gamma \|\boldsymbol{\theta}^{t+1} - \hat{\boldsymbol{\theta}}\|_2^2. \tag{3.7}$$

for constants $0 < \alpha < \gamma$. *Let* $\boldsymbol{w}_1, \boldsymbol{w}_2, \cdots, \boldsymbol{w}_s$ *be the noise vectors added to* $\boldsymbol{\theta}^t - \eta^0 \nabla \mathcal{L}_n(\boldsymbol{\theta}^t; \boldsymbol{Z})$ *when the support of* $\boldsymbol{\theta}^{t+1}$ *is iteratively selected,* S^{t+1} *be the support of* $\boldsymbol{\theta}^{t+1}$, *and* $\tilde{\boldsymbol{w}}$ *be the noise vector added to the selected* s*-sparse vector. Then, for* $\eta_0 = 2/3\gamma$, *there exists an absolute constant* c_0 *so that, choosing* $s \geq c_0(\gamma/\alpha)^2 s^*$ *guarantees*

$$\mathcal{L}_n(\boldsymbol{\theta}^{t+1}) - \mathcal{L}_n(\hat{\boldsymbol{\theta}}) \leq \left(1 - \rho \cdot \frac{\alpha}{\gamma} - \frac{2s^*}{s + s^*}\right)\left(\mathcal{L}_n(\boldsymbol{\theta}^t) - \mathcal{L}_n(\hat{\boldsymbol{\theta}})\right) + C_\gamma \left(\sum_{i \in [s]} \|\boldsymbol{w}_i\|_\infty^2 + \|\tilde{\boldsymbol{w}}_{S^{t+1}}\|_2^2\right),$$

where $0 < \rho < 1$ *is an absolute constant, and* $C_\gamma > 0$ *is a constant depending on* γ.

Theorem 2 is proved in Section A.5. While conditions (3.6) and (3.7) are similar to the ordinary strong convexity and smoothness conditions in appearance, they are in fact much weaker because $\hat{\boldsymbol{\theta}}$, $\boldsymbol{\theta}^t$ are both s-sparse. It is unclear yet, however, whether these weaker conditions are satisfied by the GLM log-likelihood function, and whether the linear convergence in terms of $\mathcal{L}_n$ implies any positive result for parameter estimation accuracy $\|\boldsymbol{\theta}^{(T)} - \hat{\boldsymbol{\theta}}\|_2^2$. In the next section, we resolve these issues for high-dimensional sparse GLMs and obtain a parameter estimation accuracy result.

3.2.2 Noisy Iterative Hard Thresholding for High-Dimensional Sparse GLMs

Assuming that the true GLM parameter vector $\boldsymbol{\beta}^*$ satisfies $\|\boldsymbol{\beta}^*\|_0 \leq s^*$, we now specialize the results of Section 3.2.1 to the GLM negative log-likelihood function

$$\mathcal{L}_n(\boldsymbol{\beta}; \boldsymbol{Z}) = \frac{1}{n} \sum_{i=1}^{n} \left(\psi(\boldsymbol{x}_i^\top \boldsymbol{\beta}) - g(y_i) \boldsymbol{x}_i^\top \boldsymbol{\beta} \right).$$

Algorithm 5: Differentially Private Sparse Generalized Linear Regression

Input : $\mathcal{L}_n(\boldsymbol{\beta}, \boldsymbol{Z})$, data set $\boldsymbol{Z}$, sparsity level s, step size η^0, privacy parameters ε, δ, noise scale B, number of iterations T, truncation parameter R.

Initialize $\boldsymbol{\beta}^0$ with $\|\boldsymbol{\beta}^0\|_0 \leq s$, set $t = 0$;

for t *in* 0 **to** $T-1$ **do**

  Compute $\boldsymbol{\beta}^{t+0.5} = \boldsymbol{\beta}^{(T)} - (\eta_0/n) \sum_{i=1}^{n} (\psi'(\boldsymbol{x}_i^\top \boldsymbol{\beta}^t) - \Pi_R(y_i)) \boldsymbol{x}_i$;

  $\boldsymbol{\beta}^{t+1} = \texttt{NoisyHT}\left(\boldsymbol{\beta}^{t+0.5}, \boldsymbol{Z}, s, \varepsilon/T, \delta/T, \eta^0 B/n\right)$;

end

Output: $\boldsymbol{\beta}^{(T)}$.

Some assumptions about the data set $\{(y_i, \boldsymbol{x}_i)\}_{i \in [n]}$ and its distribution will be helpful for analyzing the accuracy and privacy guarantees of Algorithm 5. The necessary assumptions for the high-dimensional sparse case are identical to those for the low-dimensional case, except with (D1) replaced by (D1'), as follows.

(D1') Bounded design: there is a constant $\sigma_{\boldsymbol{x}} < \infty$ such that $\|\boldsymbol{x}\|_\infty < \sigma_{\boldsymbol{x}}$ almost surely.

Because Algorithm 5 is a special case of the general Algorithm 4, the privacy guarantee of Algorithm 5 reduces to specializing Lemma 4 to GLMs, as follows.

Lemma 5. *If assumptions (D1') and (G1) are true, then choosing* $B = 4(R + c_1)\sigma_{\boldsymbol{x}}$ *guarantees that Algorithm 5 is* (ε, δ)*-differentially private.*

The lemma is proved in Section A.6.

For the parameter estimation accuracy of Algorithm 5, Fact 4 combined with the sparsity of $\hat{\boldsymbol{\beta}}$, $\boldsymbol{\beta}^*$ and $\boldsymbol{\beta}^t$ for every t are sufficient for conditions (3.6) and (3.7) in Theorem 2 to hold. Invoking Theorem 2 in a proof by induction then leads to an upper bound for $\|\boldsymbol{\beta}^{(T)} - \boldsymbol{\beta}^*\|_2^2$. Below we state the main result; the detailed proof is in Section 7.1.

Theorem 3. *Let* $\{(y_i, \boldsymbol{x}_i)\}_{i \in [n]}$ *be an i.i.d. sample from the GLM* (3.1) *with the true parameter vector* $\|\boldsymbol{\beta}^*\|_0 \leq s^*$*. Suppose assumptions (D1'), (D2), (G1) and (G2) are true. Let the parameters of Algorithm* 5 *be chosen as follows.*

- *Set sparsity level* $s = 4c_0(\gamma/\alpha)^2 s^*$ *and step size* $\eta^0 = 1/(2\gamma)$*, where the constant* c_0 *is defined in Theorem 2 and constants* α*,* γ *are defined in Proposition 4.*
- *Set* $R = \min\left(\operatorname{ess\,sup}|y_1|, c_1 + \sqrt{2c_2 c(\sigma)\log n}\right) \lesssim \sqrt{c(\sigma)\log n}$.
- *Noise scale* B*. Set* $B = 4(R + c_1)\sigma_{\boldsymbol{x}}$.
- *Number of iterations* T*. Let* $T = (2\gamma/\rho\alpha)\log(6\gamma n)$*, where* ρ *is an absolute constant defined in Theorem 1.1.*
- *Initialization* $\boldsymbol{\beta}^0$*. Choose* $\boldsymbol{\beta}^0$ *so that* $\|\boldsymbol{\beta}^0\|_0 \le s$ *and* $\|\boldsymbol{\beta}^0 - \hat{\boldsymbol{\beta}}\|_2 \le 3$*, where* $\hat{\boldsymbol{\beta}} = \arg\min_{\|\boldsymbol{\beta}\|_0 \le s^*} \mathcal{L}_n(\boldsymbol{\beta}; Z)$.

If $n \ge K \cdot \left(Rs^* \log d\sqrt{\log(1/\delta)}\log n/\varepsilon\right)$ *for a sufficiently large constant* K*, it holds with probability at least* $1 - c_3\exp(-c_4\log(d/s^*\log n)) - c_3\exp(-c_4 n) - c_3\exp(-c_4\log n)$ *that* $\boldsymbol{\beta}^{(T)}$*, the output of Algorithm 5 satisfies*

$$\|\boldsymbol{\beta}^{(T)} - \boldsymbol{\beta}^*\|_2^2 \lesssim c(\sigma)\left(\frac{s^*\log d}{n} + \frac{(s^*\log d)^2\log(1/\delta)\log^3 n}{n^2\varepsilon^2}\right). \tag{3.8}$$

Theorem 3 is proved in Section 7.1. Similar to the low-dimensional GLM algorithm, the step size, number of iterations and initialization are chosen to ensure convergence; the initialization condition, as in [38], is standard in the literature and can be extended to $\|\boldsymbol{\beta}^0 - \hat{\boldsymbol{\beta}}\|_2 \le 3\max(1, \|\boldsymbol{\beta}^*\|_2)$.

The choice of truncation level R is to ensure privacy while keeping as many data intact as possible; when the distribution of y has bounded support, for example in the logistic model, it can be chosen to be an $O(1)$ constant and therby saving an extra factor of $O(\log n)$ in the second term of (3.5). The scaling of n versus d, s^*, ε and δ in Theorem 3 is nearly optimal, as the corresponding lower bound, Theorem 6, shall show that no estimator can achieve low ℓ_2 error unless the assumed scaling holds, and that the statistical accuracy of Algorithm 5 cannot be further improved except possibly for factors of $\log n$.

4 Privacy-constrained Minimax Lower Bounds

Section 3 proposed differentially private algorithms for estimating GLM parameters and obtained convergence rates for these algorithms. We shall show in this section that the convergence rates cannot be improved by any other (ε, δ)-differentially private estimator beyond possibly factors of $\log n$, via privacy-constrained lower bounds of the form

$$\inf_{M \in \mathcal{M}_{\varepsilon,\delta}} \sup_{\boldsymbol{\beta}\in\Theta} \mathbb{E}\|M(\boldsymbol{y}, \boldsymbol{X}) - \boldsymbol{\beta}\|_2^2 \gtrsim r(n, d, \Theta, \sigma, \varepsilon, \delta), \tag{4.1}$$

where $\mathcal{M}_{\varepsilon,\delta}$ is the collection of all (ε, δ)-differentially private estimators, $\Theta \subseteq \mathbb{R}^d$ is a parameter space to which the true value of $\boldsymbol{\beta}$ is assumed to belong, and the expectation is taken over $\boldsymbol{y}, \boldsymbol{X}$ and the randomness of M.

We shall provide precise forms of the lower bound $r(n, d, \Theta, \sigma, \varepsilon, \delta)$ for both the low-dimensional and high-dimensional sparse GLMs, via a broad generalization of the "tracing attack" argument [10, 26, 25, 50] for privacy-constrained minimax lower bounds.

A tracing attack is an algorithm that takes a single candidate datum as input and attempts to infer whether this candidate belongs to a given data set or not, by comparing the candidate with some summary statistics computed from the data set. Statisticians can think of a tracing attack as a hypothesis test which rejects the null hypothesis that the candidate is out of the data set for large values of some test statistic. The hypothesis testing formulation naturally motivates some desiderata for a tracing attack:

- Soundness (type I error control): if the candidate does not belong to the data set, the tracing attack is likely to takes small values.
- Completeness (type II error control): if the candidate does belong, the tracing attack is likely to take large values.

For example, [26, 31, 11] showed that, if the random sample $\boldsymbol{X}$ and the candidate $\boldsymbol{z}$ are drawn from a Gaussian distribution with mean $\boldsymbol{\mu}$, tracing attacks of the form $\langle M(\boldsymbol{X}) - \boldsymbol{\mu}, \boldsymbol{z} - \boldsymbol{\mu} \rangle$ is sound and complete provided that $M(\boldsymbol{X})$ is an accurate estimator of $\boldsymbol{\mu}$. This accuracy requirement in turn connects tracing attacks with risk lower bounds for differentially private algorithms: if an estimator $M(\boldsymbol{X})$ is differentially private, it cannot possibly be too close to the estimand, or the existence of tracing attacks leads to a contradiction with the guarantees of differential privacy.

Designing sound and complete tracing attacks, therefore, is crucial to the sharpness of privacy-constrained minimax lower bounds. Besides the Gaussian mean tracing attack mentioned above, there are some successful tracing attacks proposed for specific problems, such as top-k selection [51] or linear regression [11], but a general recipe for the design and analysis of tracing attacks has not been available. In Section 4.1, we construct a tracing attack applicable to general parametric families of distributions, and describe its utility for privacy-constrained minimax lower bounds. This general approach is then specialized to low-dimensional and high-dimensional sparse GLMs, in Sections 4.2 and 4.3 respectively, to establish lower bound results that match the upper bound results in Section 3 up to factors of $\log n$.

4.1 The Score Attack

Given a parametric family of distributions $\{f_{\boldsymbol{\theta}}(\boldsymbol{x}) : \boldsymbol{\theta} \in \Theta\}$ with $\Theta \subseteq \mathbb{R}^d$, the score statistics, or simply the score, is given by $S_{\boldsymbol{\theta}}(\boldsymbol{x}) := \nabla_{\boldsymbol{\theta}} \log f_{\boldsymbol{\theta}}(\boldsymbol{x})$. If $\boldsymbol{x} \sim f_{\boldsymbol{\theta}}$, we have $\mathbb{E}S_{\boldsymbol{\theta}}(\boldsymbol{x}) = \mathbf{0}$ and $\mathrm{Var}S_{\boldsymbol{\theta}}(\boldsymbol{x}) = \mathcal{I}(\boldsymbol{\theta})$, where $\mathcal{I}(\boldsymbol{\theta})$ is the Fisher information matrix of $f_{\boldsymbol{\theta}}$.

Using the score statistic, we define the score attack as

$$\mathbf{A}_{\boldsymbol{\theta}}(\boldsymbol{z}, M(\boldsymbol{X})) := \langle M(\boldsymbol{X}) - \boldsymbol{\theta}, S_{\boldsymbol{\theta}}(\boldsymbol{z}) \rangle. \tag{4.2}$$

The score attack conjectures that $\boldsymbol{z}$ belongs to $\boldsymbol{X}$ for large values of $\mathbf{A}_{\boldsymbol{\theta}}(\boldsymbol{z}, M(\boldsymbol{X}))$. In particular, if $f_{\boldsymbol{\theta}}(\boldsymbol{x})$ is the density of $N(\boldsymbol{\theta}, \boldsymbol{I})$, the score attack coincides with the tracing attacks for Gaussian means studied in [26, 31, 11].

As argued earlier, an effective tracing attack should ideally be "sound" (low type I error) and "complete" (low Type II error). This is indeed the case for our score attack.

Theorem 4. *Let $\boldsymbol{X} = \{\boldsymbol{x}_1, \boldsymbol{x}_2, \cdots, \boldsymbol{x}_n\}$ be an i.i.d. sample drawn from $f_{\boldsymbol{\theta}}$. For each $i \in [n]$, let $\boldsymbol{X}_i'$ denote the data set obtained from $\boldsymbol{X}$ by replacing $\boldsymbol{x}_i$ with an independent copy $\boldsymbol{x}_i' \sim f_{\boldsymbol{\theta}}$.*

1. *Soundness: for each $i \in [n]$,*

$$\mathbb{E}\mathbf{A}_{\boldsymbol{\theta}}(\boldsymbol{x}_i, M(\boldsymbol{X}_i')) = 0;\ \mathbb{E}|\mathbf{A}_{\boldsymbol{\theta}}(\boldsymbol{x}_i, M(\boldsymbol{X}_i'))| \le \sqrt{\mathbb{E}\|M(\boldsymbol{X}) - \boldsymbol{\theta}\|_2^2}\sqrt{\lambda_{\max}(\mathcal{I}(\boldsymbol{\theta}))}. \tag{4.3}$$

2. *Completeness: if for every $j \in [d]$, $\log f_{\boldsymbol{\theta}}(\boldsymbol{X})$ is continuously differentiable with respect to θ_j and $|\frac{\partial}{\partial \theta_j} \log f_{\boldsymbol{\theta}}(\boldsymbol{X})| < g_j(X)$ such that $\mathbb{E}|g_j(\boldsymbol{X})M(\boldsymbol{X})_j| < \infty$, we have*

$$\sum_{i \in [n]} \mathbb{E}\mathbf{A}_{\boldsymbol{\theta}}(\boldsymbol{x}_i, M(\boldsymbol{X})) = \sum_{j \in [d]} \frac{\partial}{\partial \theta_j} \mathbb{E}M(\boldsymbol{X})_j. \tag{4.4}$$

Theorem 4 is proved in Section 7.2. The special form of "completeness" for Gaussian and Beta-Binomial families have been discovered as "fingerprinting lemma" in the literature [53, 10, 51, 31].

It may not be clear yet how the soundness and completeness properties would imply lower bounds for $\mathbb{E}\|M(\boldsymbol{X}) - \boldsymbol{\theta}\|_2^2$. For the specific attacks designed for Gaussian mean estimation [31] and top-k selection [51], it has been observed that, if M is an (ε, δ)-differentially private algorithm, one can prove inequalities of the form $\mathbb{E}\mathbf{A}_{\boldsymbol{\theta}}(\boldsymbol{x}_i, M(\boldsymbol{X})) \le \mathbb{E}\mathbf{A}_{\boldsymbol{\theta}}(\boldsymbol{x}_i, M(\boldsymbol{X}_i')) + O(\varepsilon)\mathbb{E}|\mathbf{A}_{\boldsymbol{\theta}}(\boldsymbol{x}_i, M(\boldsymbol{X}_i'))|$. Suppose such relations hold for the score attack

as well, the soundness property (4.3) would then imply

$$\sum_{i\in[n]} \mathbb{E}\mathbf{A}_{\boldsymbol{\theta}}(\boldsymbol{x}_i, M(\boldsymbol{X})) \le \sqrt{\mathbb{E}\|M(\boldsymbol{X}) - \boldsymbol{\theta}\|_2^2} \cdot n\sqrt{\lambda_{\max}(\mathcal{I}(\boldsymbol{\theta}))}O(\varepsilon).$$

We give precise statement of such an inequality in Section 4.1.1.

On the other hand, if we can also bound $\sum_{i\in[n]} \mathbb{E}\mathbf{A}_{\boldsymbol{\theta}}(\boldsymbol{x}_i, M(\boldsymbol{X}))$ from below by some positive quantity, a lower bound for $\mathbb{E}\|M(\boldsymbol{X}) - \boldsymbol{\theta}\|_2^2$ is immediately implied. Completeness may help us in this regard: when $\mathbb{E}M(\boldsymbol{X})_j$ is close to θ_j, it is reasonable to expect that $\frac{\partial}{\partial\theta_j}\mathbb{E}M(\boldsymbol{X})_j$ is bounded away from zero. Indeed several versions of this argument, often termed "strong distribution", exist in the literature [26, 50] and have led to lower bounds for Gaussian mean estimation and top-k selection. In Section 4.1.2, we consider a systematic approach to lower bounding $\frac{\partial}{\partial\theta_j}\mathbb{E}M(\boldsymbol{X})_j$ via Stein's Lemma [48, 49]. The technical results in Sections 4.1.1 and 4.1.2 combined with Theorem 4 would enable us to later prove concrete minimax lower bounds for GLMs.

4.1.1 Score Attacks and Differential Privacy

In Theorem 4, we have found that, when the data set $\boldsymbol{X}_i'$ does not include $\boldsymbol{x}_i$, the score attack is unlikely to take large values:

$$\mathbb{E}\mathbf{A}_{\boldsymbol{\theta}}(\boldsymbol{x}_i, M(\boldsymbol{X}_i')) = 0;\ \mathbb{E}|\mathbf{A}_{\boldsymbol{\theta}}(\boldsymbol{x}_i, M(\boldsymbol{X}_i'))| \le \sqrt{\mathbb{E}\|M(\boldsymbol{X}) - \boldsymbol{\theta}\|_2^2}\sqrt{\lambda_{\max}(\mathcal{I}(\boldsymbol{\theta}))}.$$

If M is differentially private, the distribution of $M(\boldsymbol{X}_i')$ is close to that of $M(\boldsymbol{X})$; as a result, the inequalities above can be related to the case where the data set $\boldsymbol{X}$ does include the candidate $\boldsymbol{x}_i$.

Lemma 6. *If M is an (ε, δ)-differentially private algorithm with $0 < \varepsilon < 1$ and $\delta \ge 0$, then for every $T > 0$,*

$$\mathbb{E}\mathbf{A}_{\boldsymbol{\theta}}(\boldsymbol{x}_i, M(\boldsymbol{X})) \le 2\varepsilon\sqrt{\mathbb{E}\|M(\boldsymbol{X}) - \boldsymbol{\theta}\|_2^2}\sqrt{\lambda_{\max}(\mathcal{I}(\boldsymbol{\theta}))} + 2\delta T + \int_T^\infty \mathbb{P}\left(|\mathbf{A}_{\boldsymbol{\theta}}(\boldsymbol{x}_i, M(\boldsymbol{X}))| > t\right). \tag{4.5}$$

Lemma 6 is proved in Section 7.2.1. The quantity on the right side of (4.5) is determined by the statistical model $f_{\boldsymbol{\theta}}(\boldsymbol{x})$ and the choice of T. In Sections 4.2 and 4.3, we work out its specific forms for low-dimensional and high-dimensional sparse GLMs.

4.1.2 Score Attacks and Stein's Lemma

Let us denote $\mathbb{E}_{X|\boldsymbol{\theta}} M(\boldsymbol{X})$ by $g(\boldsymbol{\theta})$, then g is a map from Θ to Θ, and we are interested in bounding $\frac{\partial}{\partial \theta_j} g_j(\boldsymbol{\theta})$ from below. Stein's Lemma [48, 49], as stated below, suggests some promising directions.

Lemma 7 (Stein's Lemma)**.** *Let Z be distributed according to some density $p(z)$ that is continuously differentiable with respect to z and let $h : \mathbb{R} \to \mathbb{R}$ be a differentiable function such that $\mathbb{E}|h'(Z)| < \infty$. We have*

$$\mathbb{E}h'(Z) = \mathbb{E}\left[\frac{-h(Z)p'(Z)}{p(Z)}\right].$$

In particular, if $p(z) = (2\pi)^{-1/2} e^{-z^2/2}$, we have $\mathbb{E}h'(Z) = \mathbb{E}Zh(Z)$.

Stein's Lemma implies that, by imposing appropriate prior distributions on $\boldsymbol{\theta}$, one can obtain a lower bound for $\frac{\partial}{\partial \theta_j} g_j(\boldsymbol{\theta})$ on average over the prior distribution of $\boldsymbol{\theta}$, as follows.

Lemma 8. *Let $\boldsymbol{\theta}$ be distributed according to a density $\boldsymbol{\pi}$ with marginal densities $\{\pi_j\}_{j\in[d]}$. If for every $j \in [d]$, π_j, g_j satisfy the regularity conditions in Lemma 7, we have*

$$\mathbb{E}_{\boldsymbol{\pi}}\left(\sum_{j\in[d]} \frac{\partial}{\partial \theta_j} g_j(\boldsymbol{\theta})\right) \geq \mathbb{E}_{\boldsymbol{\pi}}\left(\sum_{j\in[d]} \frac{-\theta_j \pi_j'(\theta_j)}{\pi_j(\theta_j)}\right) - \sqrt{\mathbb{E}_{\boldsymbol{\pi}}\|g(\boldsymbol{\theta}) - \boldsymbol{\theta}\|_2^2 \cdot \mathbb{E}_{\boldsymbol{\pi}}\left[\sum_{j\in[d]} \left(\frac{\pi_j'(\theta_j)}{\pi_j(\theta_j)}\right)^2\right]} \tag{4.6}$$

Lemma 8 is proved in Section 7.2.2. Often we may assume without the loss of generality that $\mathbb{E}_{\boldsymbol{\pi}}\|g(\boldsymbol{\theta}) - \boldsymbol{\theta}\|_2^2 \leq \mathbb{E}_{\boldsymbol{\pi}}\mathbb{E}_{\boldsymbol{X}|\boldsymbol{\theta}}\|M(\boldsymbol{X}) - \boldsymbol{\theta}\|_2^2 < C$ for some constant C when the sample size n is sufficiently large, the right side is completely determined by the choice of π, as the following example illustrates:

Example 1. *Let $\boldsymbol{\pi}$ be the density of $N(\mathbf{0}, \boldsymbol{I})$, then (4.6) reduces to*

$$\mathbb{E}_{\boldsymbol{\pi}}\left(\sum_{j\in[d]} \frac{\partial}{\partial \theta_j} g_j(\boldsymbol{\theta})\right) \geq \sum_{j\in[d]} \mathbb{E}_{\pi_j}\theta_j^2 - \sqrt{C}\sqrt{\sum_{j\in[d]} \mathbb{E}_{\pi_j}\theta_j^2} = d - \sqrt{Cd} \gtrsim d.$$

In view of the completeness property (4.4), Lemma 8 suggests an *average* lower bound for $\sum_{i\in[n]} \mathbb{E}\mathbf{A}_{\boldsymbol{\theta}}(\boldsymbol{x}_i, M(\boldsymbol{X}))$ over some prior distribution $\boldsymbol{\pi}(\boldsymbol{\theta})$, with the specific form of this average lower bound entirely determined by the choice of $\boldsymbol{\pi}$.

4.1.3 From Score Attacks to Lower Bounds

Let us combine Theorem 4 with Lemmas 6 and 8 to understand how the score attack leads to privacy-constrained minimax lower bounds.

Let $\boldsymbol{\pi}$ be a prior distribution supported over the parameter space Θ with marginal densities $\{\pi_j\}_{j\in[d]}$, and assume without the loss of generality that $\mathbb{E}_{\boldsymbol{X}|\boldsymbol{\theta}}\|M(\boldsymbol{X})-\boldsymbol{\theta}\|_2^2 < C$ for every $\boldsymbol{\theta}\in\Theta$. The completeness part of Theorem 4 and Lemma 8 imply that

$$\sum_{i\in[n]}\mathbb{E}_{\boldsymbol{\pi}}\mathbb{E}_{\boldsymbol{X}|\boldsymbol{\theta}}\mathbf{A}_{\boldsymbol{\theta}}(\boldsymbol{x}_i, M(\boldsymbol{X})) \geq \mathbb{E}_{\boldsymbol{\pi}}\left(\sum_{j\in[d]}\frac{-\theta_j\pi_j'(\theta_j)}{\pi_j(\theta_j)}\right) - \sqrt{C}\sqrt{\mathbb{E}_{\boldsymbol{\pi}}\left[\sum_{j\in[d]}\left(\frac{\pi_j'(\theta_j)}{\pi_j(\theta_j)}\right)^2\right]}$$

Since Lemma 6 holds for every $\boldsymbol{\theta}$, it follows from the Lemma that

$$\begin{aligned}&\sum_{i\in[n]}\mathbb{E}_{\boldsymbol{\pi}}\mathbb{E}_{\boldsymbol{X}|\boldsymbol{\theta}}\mathbf{A}_{\boldsymbol{\theta}}(\boldsymbol{x}_i, M(\boldsymbol{X}))\\ &\leq 2n\varepsilon\sqrt{\mathbb{E}_{\boldsymbol{\pi}}\mathbb{E}_{\boldsymbol{X}|\boldsymbol{\theta}}\|M(\boldsymbol{X})-\boldsymbol{\theta}\|_2^2}\sqrt{\lambda_{\max}(\mathcal{I}(\boldsymbol{\theta}))} + 2n\delta T + \sum_{i\in[n]}\int_T^\infty \mathbb{P}\left(|\mathbf{A}_{\boldsymbol{\theta}}(\boldsymbol{x}_i, M(\boldsymbol{X}))| > t\right).\end{aligned}$$

These two inequalities are true for every (ε,δ)-differentially private M, and they therefore suggest a lower bound for $\inf_{M\in\mathcal{M}_{\varepsilon,\delta}}\mathbb{E}_{\boldsymbol{\pi}}\mathbb{E}_{\boldsymbol{X}|\boldsymbol{\theta}}\|M(\boldsymbol{X})-\boldsymbol{\theta}\|_2^2$, which in turn lower bounds $\inf_{M\in\mathcal{M}_{\varepsilon,\delta}}\sup_{\boldsymbol{\theta}\in\Theta}\mathbb{E}_{\boldsymbol{X}|\boldsymbol{\theta}}\|M(\boldsymbol{X})-\boldsymbol{\theta}\|_2^2$ since the maximum risk is greater than the average risk regardless of the prior distribution.

Following this strategy, we shall obtain the privacy-constrained minimax lower bounds for GLM problems, by choosing an appropriate prior distribution $\boldsymbol{\pi}$ and working out the specific forms of the two inequalities (4.5) and (4.6) in the context of GLMs.

4.2 The Classical Low-dimensional Setting

We first consider the low-dimensional $d = o(n)$ setting. For the generalized linear model

$$f_{\boldsymbol{\beta}}(y|\boldsymbol{x}) = h(y,\sigma)\exp\left(\frac{\boldsymbol{x}^\top\boldsymbol{\beta}y - \psi(\boldsymbol{x}^\top\boldsymbol{\beta})}{c(\sigma)}\right); \boldsymbol{x}\sim f_{\boldsymbol{x}},$$

and a candidate datum $(\tilde{y},\tilde{\boldsymbol{x}})$, the score attack, as defined by (4.2), takes the form

$$\mathbf{A}_{\boldsymbol{\beta}}((\tilde{y},\tilde{\boldsymbol{x}}), M(\boldsymbol{y},\boldsymbol{X})) = \frac{1}{c(\sigma)}\big\langle M(\boldsymbol{y},\boldsymbol{X})-\boldsymbol{\beta}, [\tilde{y}-\psi'(\tilde{\boldsymbol{x}}^\top\boldsymbol{\beta})]\tilde{\boldsymbol{x}}\big\rangle. \tag{4.7}$$

For the prior distribution of $\boldsymbol{\beta}$, we choose $\boldsymbol{\pi}(\boldsymbol{\beta})$ to be the density of $N(\mathbf{0}, \boldsymbol{I})$. The strategy outlined in Section 4.1 implies the following lower bound result.

Theorem 5. *Consider i.i.d. observations* $(y_1, \boldsymbol{x}_1), \cdots, (y_n, \boldsymbol{x}_n) \in \mathbb{R} \times \mathbb{R}^d$*, where* $\boldsymbol{x} \sim f_{\boldsymbol{x}}$ *such that* $\mathbb{E}(\boldsymbol{x}\boldsymbol{x}^\top)$ *is diagonal with* $0 < \lambda_{\max}(\mathbb{E}(\boldsymbol{x}\boldsymbol{x}^\top)) < C < \infty$*,* $\|\boldsymbol{x}\|_2 \lesssim \sqrt{d}$ *almost surely, and* y *given* $\boldsymbol{x}$ *follows the conditional distribution*

$$f_{\boldsymbol{\beta}}(y|\boldsymbol{x}) = h(y, \sigma) \exp\left(\frac{\boldsymbol{x}^\top \boldsymbol{\beta} y - \psi(x^\top \boldsymbol{\beta})}{c(\sigma)}\right).$$

If $0 < \|\psi''\|_\infty < c_2 < \infty$*,* $0 < \varepsilon < 1$*,* $0 < \delta < n^{-(1+\gamma)}$ *for some* $\gamma > 0$*, then for sufficiently large* n *and every* (ε, δ)*-differentially private* M *such that* $\|M(\boldsymbol{y}, \boldsymbol{X}) - \boldsymbol{\beta}\|_2^2 \lesssim d$ *and* $\mathbb{E}\|M(\boldsymbol{y}, \boldsymbol{X}) - \boldsymbol{\beta}\|_2^2 = o(1)$*,*

$$\sup_{\boldsymbol{\beta} \in \mathbb{R}^d} \mathbb{E}\|M(\boldsymbol{y}, \boldsymbol{X}) - \boldsymbol{\beta}\|_2^2 \gtrsim c(\sigma) \frac{d^2}{n^2 \varepsilon^2}. \tag{4.8}$$

Theorem 5 is proved in Section B.1. The (ε, δ)-differentially private estimators $\mathcal{M}_{\varepsilon,\delta}$ are also subject to the non-private minimax risks lower bound for GLMs, $\inf_M \sup_{\boldsymbol{\beta}} \mathbb{E}\|M(\boldsymbol{y}, \boldsymbol{X}) - \boldsymbol{\beta}\|_2^2 \gtrsim c(\sigma) d/n$. It then follows from (4.8) that

$$\inf_{M \in \mathcal{M}_{\varepsilon,\delta}} \sup_{\boldsymbol{\beta}} \mathbb{E}\|M(\boldsymbol{y}, \boldsymbol{X}) - \boldsymbol{\beta}\|_2^2 \gtrsim c(\sigma) \left(\frac{d}{n} + \frac{d^2}{n^2 \varepsilon^2}\right).$$

The lower bound matches the statistical accuracy of noisy gradient descent, Theorem 1, up to factors of $\log n$ under the usual setting of $\delta = n^{-\alpha}$ for some constant $\alpha > 1$. Besides showing the optimality of noisy gradient descent, this comparison also suggests that the cost of privacy, as measured by the squared ℓ_2-norm, in GLM parameter estimation is of the order $d^2/n^2\varepsilon^2$.

4.3 The High-Dimensional Sparse Setting

We now consider the setting where d, the dimension of Θ, dominates the sample size n, but each $\boldsymbol{\beta} \in \Theta$ is assumed to be s^*-sparse, that is $\|\boldsymbol{\beta}\|_0 \leq s^*$. As seen in the following theorem, the sparsity assumption leads to a lower bound that depends primarily on the sparsity, or the "intrinsic dimension" of $\boldsymbol{\beta}$, and only logarithmically on the ambient dimension d.

For high-dimensional sparse GLMs, we consider a modification of the classical GLM score attack (4.7), the sparse GLM score attack:

$$\mathbf{A}_{\boldsymbol{\beta},s^*}((\tilde{y}, \tilde{\boldsymbol{x}}), M(\boldsymbol{y}, \boldsymbol{X})) = \frac{1}{c(\sigma)} \big\langle (M(\boldsymbol{y}, \boldsymbol{X}) - \boldsymbol{\beta})_{\mathrm{supp}(\boldsymbol{\beta})}, [\tilde{y} - \psi'(\tilde{\boldsymbol{x}}^\top \boldsymbol{\beta})] \tilde{\boldsymbol{x}} \big\rangle. \tag{4.9}$$

For the prior $\boldsymbol{\pi}$, we have to choose some distribution supported over the set $\{\boldsymbol{\beta} : \boldsymbol{\beta} \in \mathbb{R}^d, \|\boldsymbol{\beta}\|_0 \leq s^*\}$. Specifically, we consider $\boldsymbol{\beta}$ generated as follows: let $\tilde{\beta}_1, \tilde{\beta}_2, \cdots, \tilde{\beta}_d$ be drawn i.i.d. from $N(0,1)$, let I_{s^*} be be the index set of $\tilde{\boldsymbol{\beta}}$ with top s^* greatest absolute values so that $|I_{s^*}| = s^*$ by definition, and define $\beta_j = \tilde{\beta}_j \mathbb{1}(j \in I_{s^*})$.

The score attack strategy then leads to the following lower bound result.

Theorem 6. *Consider n i.i.d. observations $(y_1, \boldsymbol{x}_1), \cdots, (y_n, \boldsymbol{x}_n)$, where $\boldsymbol{x} \sim f_{\boldsymbol{x}}$ such that $\mathbb{E}(\boldsymbol{x}\boldsymbol{x}^\top)$ is diagonal with $0 < \lambda_{\max}(\mathbb{E}(\boldsymbol{x}\boldsymbol{x}^\top)) < C < \infty$, $\|\boldsymbol{x}\|_\infty < c < \infty$ almost surely, and y given $\boldsymbol{x}$ follows the conditional distribution*

$$f_{\boldsymbol{\beta}}(y|\boldsymbol{x}) = h(y,\sigma) \exp\left(\frac{\boldsymbol{x}^\top\boldsymbol{\beta} y - \psi(\boldsymbol{x}^\top\boldsymbol{\beta})}{c(\sigma)}\right).$$

If $0 < \|\psi''\|_\infty = c_2 < \infty$, $0 < \varepsilon < 1$, $0 < \delta < n^{-(1+\gamma)}$ for some $\gamma > 0$, $s = o\left(d^{1-\gamma}\right)$ for some $\gamma > 0$, then for sufficiently large n and every (ε, δ)-differentially private M such that $\|M(\boldsymbol{y}, \boldsymbol{X}) - \boldsymbol{\beta}\|_2^2 \lesssim s^*$ and $\mathbb{E}\|M(\boldsymbol{y}, \boldsymbol{X}) - \boldsymbol{\beta}\|_2^2 = o(1)$,

$$\sup_{\boldsymbol{\beta} \in \mathbb{R}^d, \|\boldsymbol{\beta}\|_0 \leq s^*} \mathbb{E}\|M(\boldsymbol{y}, \boldsymbol{X}) - \boldsymbol{\beta}\|_2^2 \gtrsim c(\sigma) \frac{(s^* \log d)^2}{n^2 \varepsilon^2}. \tag{4.10}$$

Theorem 6 is proved in Section B.2. In conjunction with the non-private minimax lower bound $\inf_M \sup_{\boldsymbol{\beta} \in \mathbb{R}^d, \|\boldsymbol{\beta}\|_0 \leq s} \mathbb{E}\|M(\boldsymbol{y}, \boldsymbol{X}) - \boldsymbol{\beta}\|_2^2 \gtrsim c(\sigma) s^* \log d / n$, (4.10) implies

$$\inf_{M \in \mathcal{M}_{\varepsilon,\delta}} \sup_{\boldsymbol{\beta} \in \mathbb{R}^d, \|\boldsymbol{\beta}\|_0 \leq s} \mathbb{E}\|M(\boldsymbol{y}, \boldsymbol{X}) - \boldsymbol{\beta}\|_2^2 \gtrsim c(\sigma) \left(\frac{s^* \log d}{n} + \frac{(s^* \log d)^2}{n^2 \varepsilon^2}\right).$$

By comparing the privacy-constrained minimax lower bound with Theorem 3, we can see that the noisy iterative hard thresholding algorithm for sparse GLMs is optimal up to factors of $\log n$ under the usual setting of $\delta = n^{-\alpha}$, and that the cost of privacy, as measured by squared ℓ_2 norm, in sparse GLM parameter estimation is of the order $(s^* \log d)^2 / n^2 \varepsilon^2$.

5 Numerical Results

In this section, we investigate the numerical performance of the proposed privacy-preserving algorithms by conducting experiments with both simulated and real data sets. The numerical results also illustrate our theoretical findings on differentially private GLM parameter estimation.

5.1 Simulated Data

For the low-dimensional GLM, our simulated data set is constructed as follows. For our desired choice of d and n, we sample $\boldsymbol{\beta}$ uniformly at random from the unit sphere in $\mathbb{R}^d$, draw coordinates of the design vector $\boldsymbol{x}_i$ independently from the uniform distribution over $(-1, 1)$ for each $i \in [n]$, and sample y_i from the logistic regression model, that is y_i following the Bernoulli distribution with success probability $\frac{1}{1+\exp(-\boldsymbol{x}_i^\top \boldsymbol{\beta})}$. Using the simulated data, we study the numerical performance of Algorithm 1 via three sets of experiments. In each experiment, the algorithm is initialized with $\boldsymbol{\beta} = \mathbf{0} \in \mathbb{R}^d$, with step size $\eta^0 = 1$ for each iteration.

(a). Fix $n = 40000, \varepsilon = 0.5$ and $\delta = (2n)^{-1}$, and compare the iterates of Algorithm 1 with the true $\boldsymbol{\beta}$ for $d = 10, 20$, or 40. As displayed in Figure 1(a), the log error $\log(\|\boldsymbol{\beta}^t - \boldsymbol{\beta}\|_2^2)$ is linear in t when $d = 10$ but deteriorates as d increases, confirming the theoretical result in Theorem 1.

(b). Fix $d = 20, \varepsilon = 0.5$ and $\delta = (2n)^{-1}$, and compare the iterates of Algorithm 1 with the true $\boldsymbol{\beta}$ for $n = 20000, 40000$, or 80000. As predicted by Theorem 1, $\log(\|\boldsymbol{\beta}^t - \boldsymbol{\beta}\|_2^2)$ is linear in t when $n = 80000$ but deteriorates as n decreases.

(c). Fix $d = 20, n = 40000$ and $\delta = (2n)^{-1}$, and compare the iterates of Algorithm 1 with the true $\boldsymbol{\beta}$ for $\varepsilon = 0.2, 0.5, 0.8$, or ∞ (non-private). The decrease in $\log(\|\boldsymbol{\beta}^t - \boldsymbol{\beta}\|_2^2)$ as ε increases is consistent with Theorem 1.

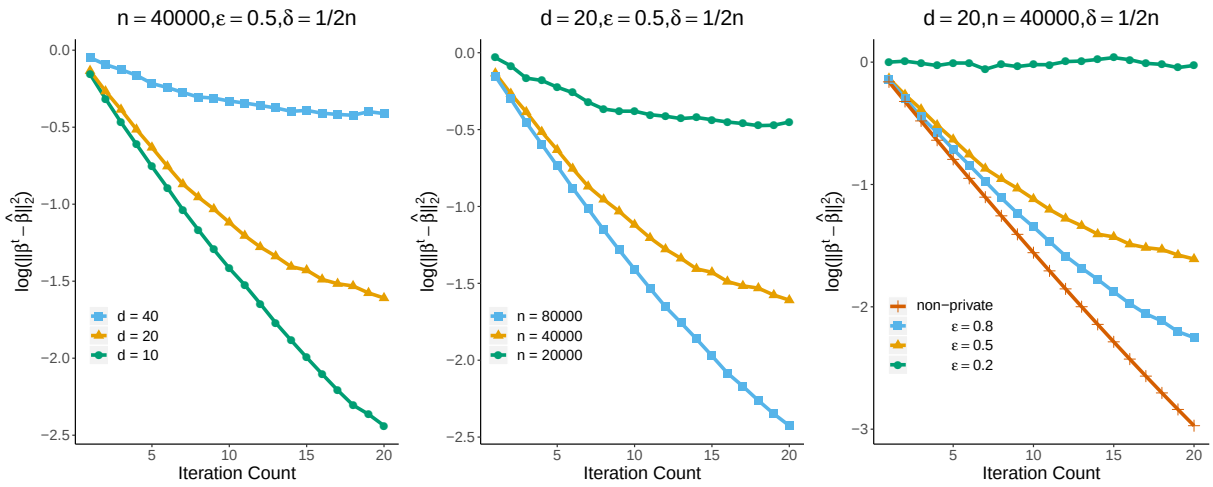

Figure 1: Log-distance between the iterates of Algorithm 1 and the true parameter $\boldsymbol{\beta}$ under various settings of n, d, ε and δ.

For the high-dimensional sparse GLM, the simulated data set is constructed in the identical way as the low-dimensional case, except that the s-sparse true parameter $\boldsymbol{\beta}$ is obtained by concatenating a random draw from the unit sphere in $\mathbb{R}^s$ with $\mathbf{0} \in \mathbb{R}^{d-s}$. We have three sets of experiments to study the numerical performance of Algorithm 5. In each experiment, the algorithm is initialized with $\boldsymbol{\beta} = \mathbf{0} \in \mathbb{R}^d$, with step size $\eta^0 = 1$ for each iteration and the sparsity level set at twice of the true sparsity.

(a). Fix $d = 10000, n = 40000, \varepsilon = 0.5$ and $\delta = (2n)^{-1}$, and compare the iterates of Algorithm 5 with the true $\boldsymbol{\beta}$ for $s = 10, 20$, or 40. As suggested by Theorem 3, the log error $\log(\|\boldsymbol{\beta}^t - \boldsymbol{\beta}\|_2^2)$ is linear in t when $s = 10$ but deteriorates as s increases.

(b). Fix $d = 10000, s = 10, \varepsilon = 0.5$ and $\delta = (2n)^{-1}$, and compare the iterates of Algorithm 5 with the true $\boldsymbol{\beta}$ for $n = 20000, 40000$, or 80000. $\log(\|\boldsymbol{\beta}^t - \boldsymbol{\beta}\|_2^2)$ is linear in t when $n = 80000$ or $n = 40000$, but deteriorates as n decreases.

(c). Fix $d = 10000, n = 40000, s = 10$ and $\delta = (2n)^{-1}$, and compare the iterates of Algorithm 5 with the true $\boldsymbol{\beta}$ for $\varepsilon = 0.2, 0.5, 0.8$, or ∞ (non-private). The decrease in $\log(\|\boldsymbol{\beta}^t - \boldsymbol{\beta}\|_2^2)$ as ε increases confirms Theorem 3.

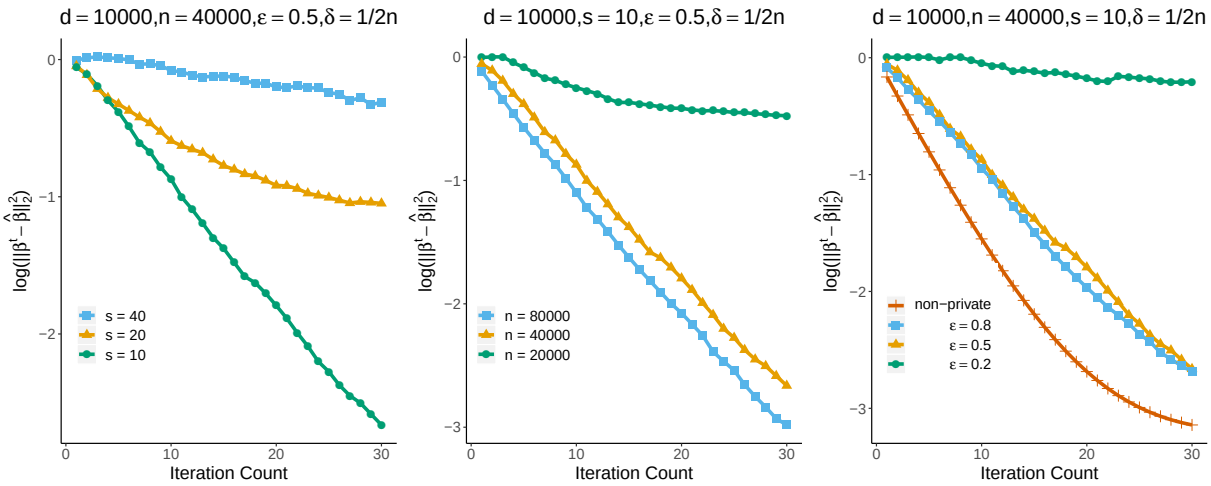

Figure 2: Log-distance between the iterates of Algorithm 5 and the true parameter $\boldsymbol{\beta}$ under various settings of n, d, s, ε and δ.

5.2 Real Data

For the real data experiment, we consider the Swarm Behavior Data Set, collected by the Human Perception of Swarming project at the University of New South Wales (`https:`

//unsw-swarm-survey.netlify.app/) and made publicly available at the UCI Machine Learning Repository [17]. In this data set, each of $n = 24016$ instances contains $d = 2400$ attributes describing the behavior (velocity, direction, location, etc.) of 200 individuals in the system, and with each instance assigned a binary class label, "flocking" or "not flocking". A system of individual birds, insects, or people are said to be "flocking" if they are perceived moving as a group with the same velocity without colliding each other.

In our experiment, we attempt to classify these instances into "flocking" or "not flocking" by our Algorithm 5 for high-dimensional sparse GLMs. We randomly split the data set into two halves, train a sparse logistic regression model using one half, and predict the labels of the other half by this logistic model. For fitting the sparse logistic model on the training set, we run Algorithm 5 for 50 iterations with step size $\eta^0 = 0.5$ and initial value $\boldsymbol{\beta}^0 = \mathbf{0} \in \mathbb{R}^{2401}$ (including the intercept). For various settings of s, ε and δ, the average misclassification rate (and its standard error) over repetitions of the experiment are displayed in the tables below. The results suggest that the classification accuracy indeed worsens as the privacy requirement becomes more stringent, but the loss of accuracy is mild compared to the non-private $\varepsilon = \infty$ case.

	$s = 25$	$s = 50$	$s = 100$
$\varepsilon = 0.2$	0.33(.05)	0.21(.05)	0.13(.05)
$\varepsilon = 0.5$	0.28(.05)	0.20(.05)	0.10(.02)
$\varepsilon = \infty$	0.30(.05)	0.21(.05)	0.09(.03)

(a) $\delta = 1/2n$

	$s = 25$	$s = 50$	$s = 100$
$\varepsilon = 0.2$	0.32(.05)	0.22(.05)	0.13(.06)
$\varepsilon = 0.5$	0.33(.04)	0.19(.05)	0.08(.02)
$\varepsilon = \infty$	0.30(.05)	0.21(.05)	0.09(.02)

(b) $\delta = 1/n^2$

Figure 3: Mean and standard error of misclassification rates of Algorithm 5 in the randomly drawn test subset of the Swarm Behavior Data Set.

6 Discussion

In this paper, we studied the cost of differential privacy in estimating the parameters of the GLMs. We designed differentially private algorithms, based on projected gradient descent, that achieve fast, linear convergence to the optimal non-private solution, and analyzed their statistical accuracy with respect to the true parameters. The theoretical properties of our algorithms are demonstrated in numerical experiments with real and simulated data sets.

The accuracy of these algorithms are shown to be optimal up to logarithmic factors, via

lower bounds of the privacy-constrained minimax risk. These lower bounds are established by the score attack framework, which generalizes prior works on tracing attacks for privacy-constrained minimax lower bounds. The upper bounds and lower bounds together have led to a clear characterization of the cost of privacy in estimating GLM parameters.

This paper suggests several promising directions of further research. On the algorithmic side, since our convergence analysis of differentially private algorithms can be applied to other M-estimation problems satisfying restricted strong convexity and restricted smoothness, it is of interest to study their performance in problems such as low-rank matrix recovery and regression. Our results on high-dimensional sparse GLMs also raise questions on the interplay between privacy and other structural assumptions, for example, group-structured sparsity, approximate sparsity, or low-rankness as mentioned above.

On the statistical optimality side, our score attack framework may lead to lower bounds for a much larger variety of statistical models than generalized linear models. It is also of significant value to prove sharper lower bounds that can potentially capture the remaining logarithmic gap between upper and lower bounds, and develop sharp lower bounds for differentially private confidence intervals or hypothesis testing.

7 Proofs

In this section, we prove the main technical results of this paper, Theorems 3 and 4.

7.1 Proof of Theorem 3

Proof of Theorem 3. We shall first define several favorable events under which the desired convergence does occur, and then show that the probability that any of the favorable events fails to happen is negligible. These events are,

$$\mathcal{E}_1 = \{(3.3) \text{ and } (3.4) \text{ hold}\}, \mathcal{E}_2 = \{\Pi_R(y_i) = y_i, \forall i \in [n]\}, \mathcal{E}_3 = \{\|\boldsymbol{\beta}^t - \hat{\boldsymbol{\beta}}\|_2 \leq 3, 0 \leq t \leq T\}.$$

We first analyze the behavior of Algorithm 5 under these events. The assumed scaling of $n \geq K \cdot \left(Rs^* \log d \sqrt{\log(1/\delta)} \log n/\varepsilon\right)$ implies that $n \geq K' s^* \log d/n$ for a sufficiently large K'. Since $\|\boldsymbol{\beta}^t\|_0 \leq s \asymp s^*$ for every t and $\|\hat{\boldsymbol{\beta}}\|_0 \leq s^*$ by definition, the RSM condition (3.4) implies that for every t,

$$\langle \nabla \mathcal{L}_n(\boldsymbol{\beta}^t) - \nabla \mathcal{L}_n(\hat{\boldsymbol{\beta}}), \boldsymbol{\beta}^t - \hat{\boldsymbol{\beta}} \rangle \leq \frac{4\gamma}{3} \|\boldsymbol{\beta}^t - \hat{\boldsymbol{\beta}}\|_2^2. \tag{7.1}$$

Similarly, under event $\mathcal{E}_3$, the RSC condition (3.3) implies that

$$\langle \nabla \mathcal{L}_n(\boldsymbol{\beta}^t) - \nabla \mathcal{L}_n(\hat{\boldsymbol{\beta}}), \boldsymbol{\beta}^t - \hat{\boldsymbol{\beta}} \rangle \geq \frac{2\alpha}{3} \|\boldsymbol{\beta}^t - \hat{\boldsymbol{\beta}}\|_2^2. \tag{7.2}$$

These two inequalities and our choice of parameters s, η now allow Theorem 2 to apply. Let $\boldsymbol{w}_1^t, \boldsymbol{w}_2^t, \cdots, \boldsymbol{w}_s^t$ be the noise vectors added to $\boldsymbol{\beta}^t - \eta^0 \nabla \mathcal{L}_n(\boldsymbol{\beta}^t; Z)$ when the support of $\boldsymbol{\beta}^{t+1}$ is iteratively selected, S^{t+1} be the support of $\boldsymbol{\beta}^{t+1}$, and $\tilde{\boldsymbol{w}}^t$ be the noise vector added to the selected s-sparse vector. Define $\boldsymbol{W}_t = C_\gamma \left(\sum_{i \in [s]} \|\boldsymbol{w}_i^t\|_\infty^2 + \|\tilde{\boldsymbol{w}}_{S^{t+1}}^t\|_2^2 \right)$, then Theorem 2 leads to

$$\begin{aligned}
\mathcal{L}_n(\boldsymbol{\beta}^{(T)}) - \mathcal{L}_n(\hat{\boldsymbol{\beta}}) &\leq \left(1 - \rho \frac{\alpha}{2\gamma}\right)^T \left(\mathcal{L}_n(\boldsymbol{\beta}^0) - \mathcal{L}_n(\hat{\boldsymbol{\beta}}) \right) + \sum_{k=0}^{T-1} \left(1 - \rho \frac{\alpha}{2\gamma}\right)^{T-k-1} \boldsymbol{W}_k \\
&\leq \left(1 - \rho \frac{\alpha}{2\gamma}\right)^T \frac{2\gamma}{3} \|\boldsymbol{\beta}_0 - \hat{\boldsymbol{\beta}}\|_2^2 + \sum_{k=0}^{T-1} \left(1 - \rho \frac{\alpha}{2\gamma}\right)^{T-k-1} \boldsymbol{W}_k \\
&\leq \left(1 - \rho \frac{\alpha}{2\gamma}\right)^T 6\gamma + \sum_{k=0}^{T-1} \left(1 - \rho \frac{\alpha}{2\gamma}\right)^{T-k-1} \boldsymbol{W}_k.
\end{aligned} \tag{7.3}$$

The second inequality is a consequence of (7.1), and the third inequality follows from the assumption that $\|\boldsymbol{\beta}_0 - \hat{\boldsymbol{\beta}}\|_2 \leq 3$. On the other hand, we can lower bound $\mathcal{L}_n(\boldsymbol{\beta}^{(T)}) - \mathcal{L}_n(\hat{\boldsymbol{\beta}})$ as follows: by (7.2),

$$\mathcal{L}_n(\boldsymbol{\beta}^{(T)}) - \mathcal{L}_n(\hat{\boldsymbol{\beta}}) \geq \mathcal{L}_n(\boldsymbol{\beta}^{(T)}) - \mathcal{L}_n(\boldsymbol{\beta}^*) \geq \frac{\alpha}{3} \|\boldsymbol{\beta}^{(T)} - \boldsymbol{\beta}^*\|_2^2 - \langle \nabla \mathcal{L}_n(\boldsymbol{\beta}^*), \boldsymbol{\beta}^* - \boldsymbol{\beta}^{(T)} \rangle. \tag{7.4}$$

Combining (7.3) and (7.4) yields

$$\begin{aligned}
\frac{\alpha}{3} \|\boldsymbol{\beta}^{(T)} - \boldsymbol{\beta}^*\|_2^2 &\leq \langle \nabla \mathcal{L}_n(\boldsymbol{\beta}^*), \boldsymbol{\beta}^* - \boldsymbol{\beta}^{(T)} \rangle + \left(1 - \rho \frac{\alpha}{2\gamma}\right)^T 6\gamma + \sum_{k=0}^{T-1} \left(1 - \rho \frac{\alpha}{2\gamma}\right)^{T-k-1} \boldsymbol{W}_k \\
&\leq \|\nabla \mathcal{L}_n(\boldsymbol{\beta}^*)\|_\infty \sqrt{s + s^*} \|\boldsymbol{\beta}^* - \boldsymbol{\beta}^{(T)}\|_2 + \left(1 - \rho \frac{\alpha}{2\gamma}\right)^T 6\gamma + \sum_{k=0}^{T-1} \left(1 - \rho \frac{\alpha}{2\gamma}\right)^{T-k-1} \boldsymbol{W}_k \\
&= \|\nabla \mathcal{L}_n(\boldsymbol{\beta}^*)\|_\infty \sqrt{s + s^*} \|\boldsymbol{\beta}^* - \boldsymbol{\beta}^{(T)}\|_2 + \frac{1}{n} + \sum_{k=0}^{T-1} \left(1 - \rho \frac{\alpha}{2\gamma}\right)^{T-k-1} \boldsymbol{W}_k.
\end{aligned} \tag{7.5}$$

The last step follows from our choice of $T = (2\gamma/\rho\alpha) \log(6\gamma n)$. Now let us define two events

that allow for high-probability bounds of the right side.

$$\mathcal{E}_4 = \left\{ \max_t \boldsymbol{W}_t \le K \left(\frac{Rs^* \log d \sqrt{\log(1/\delta)} \log n}{n\varepsilon} \right)^2 \right\}, \quad \mathcal{E}_5 = \left\{ \|\nabla \mathcal{L}_n(\boldsymbol{\beta}^*)\|_\infty \le 4\sigma_x \sqrt{c_2} \sqrt{\frac{\log d}{n}} \right\}.$$

Under $\mathcal{E}_4, \mathcal{E}_5$, we can conclude from (7.5) that

$$\|\boldsymbol{\beta}^{(T)} - \boldsymbol{\beta}^*\|_2 \lesssim \sqrt{c(\sigma)} \left(\sqrt{\frac{s^* \log d}{n}} + \frac{s^* \log d \sqrt{\log(1/\delta)} \log^{3/2} n}{n\varepsilon}. \right)$$

We have shown so far that the desired rate of convergence (3.8) holds when $\mathcal{E}_i$ occurs for $1 \le i \le 5$; we now turn to controlling the probability that any of the five events fails to happen, $\sum_{i=1}^5 \mathbb{P}(\mathcal{E}_i^c)$.

- By Proposition 4, $\mathbb{P}(\mathcal{E}_1^c) \le c_3 \exp(-c_4 n)$ under the assumptions of Theorem 3.
- We have $\mathbb{P}(\mathcal{E}_2^c) \le c_3 \exp(-c_4 \log n)$ by the choice of R, and assumptions (G1), (G2) which imply the following bound of moment generating function of y_i: we have

$$\begin{aligned} \log \mathbb{E} \exp\left(\lambda \cdot \frac{y_i - \psi'(\boldsymbol{x}_i^\top \boldsymbol{\beta})}{c(\sigma)} \Big| \boldsymbol{x}_i \right) &= \frac{1}{c(\sigma)} \left(\psi(\boldsymbol{x}_i^\top \boldsymbol{\beta} + \lambda) - \psi(\boldsymbol{x}_i^\top \boldsymbol{\beta}) - \lambda \psi'(\boldsymbol{x}_i^\top \boldsymbol{\beta}) \right) \\ &\le \frac{1}{c(\sigma)} \cdot \frac{\lambda^2 \psi''(\boldsymbol{x}_i^\top \boldsymbol{\beta} + \tilde{\lambda})}{2} \end{aligned}$$

  for some $\tilde{\lambda} \in (0, \lambda)$. It follows that $\mathbb{E} \exp\left(\lambda \cdot \frac{y_i - \psi'(\boldsymbol{x}_i^\top \boldsymbol{\beta})}{c(\sigma)} \Big| \boldsymbol{x}_i \right) \le \exp\left(\frac{c_2 \lambda^2}{2c(\sigma)} \right)$ because $\|\psi''\|_\infty < c_2$.
- For $\mathcal{E}_3$, we have $\mathbb{P}(\mathcal{E}_3^c) \le T \cdot c_3 \exp(-c_4 \log(d/s^*)) = c_3 \exp(-c_4 \log(d/s^* \log n))$ by the initial condition $\|\boldsymbol{\beta}^0 - \hat{\boldsymbol{\beta}}\|_2^3$ and proof by induction via the following lemma, to be proved in Section A.7.1.

  Lemma 9. *Under the assumptions of Theorem 3 Let $\boldsymbol{\beta}^k, \boldsymbol{\beta}^{k+1}$ be the kth and $(k+1)$th iterates of Algorithm 5. If $\|\boldsymbol{\beta}^k - \hat{\boldsymbol{\beta}}\|_2 \le 3$, we have $\|\boldsymbol{\beta}^{k+1} - \hat{\boldsymbol{\beta}}\|_2 \le 3$ with probability at least $1 - c_3 \exp(-c_4 \log(d/s^*))$.*
- For $\mathcal{E}_4$, we invoke an auxiliary lemma to be proved in Section A.7.2.

  Lemma 10. *Consider $\boldsymbol{w} \in \mathbb{R}^k$ with $w_1, w_2, \cdots, w_k \overset{i.i.d.}{\sim}$ Laplace(λ). For every $C > 1$,*

$$\begin{aligned} &\mathbb{P}\left(\|\boldsymbol{w}\|_2^2 > kC^2\lambda^2 \right) \le ke^{-C} \\ &\mathbb{P}\left(\|\boldsymbol{w}\|_\infty^2 > C^2\lambda^2 \log^2 k \right) \le e^{-(C-1)\log k}. \end{aligned}$$

For each iterate t, the individual coordinates of $\tilde{\boldsymbol{w}}^t$, $\boldsymbol{w}_i^t$ are sampled i.i.d. from the Laplace distribution with scale $(2\gamma)^{-1} \cdot \frac{2B\sqrt{3s\log(T/\delta)}}{n\varepsilon/T}$, where the noise scale $B \lesssim R$ and $T \asymp \log n$ by our choice. If $n \geq K \cdot \left(Rs^* \log d\sqrt{\log(1/\delta)}\log n/\varepsilon\right)$ for a sufficiently large constant K, Lemma 10 implies that, with probability at least $1 - c_3\exp(-c_4\log(d/(s^*\log n))$, $\max_t \boldsymbol{W}_t$ is bounded by $K\left(\frac{Rs^*\log d\sqrt{\log(1/\delta)}\log n}{n\varepsilon}\right)^2$ for some appropriate constant K.

- Under assumptions of Theorem 3, it is a standard probabilistic result (see, for example, [54] pp. 288) that $\mathbb{P}(\mathcal{E}_5^c) \leq 2e^{-2\log d}$.

We have $\sum_{i=1}^5 \mathbb{P}(\mathcal{E}_i^c) \leq c_3\exp(-c_4\log(d/s^*\log n)) + c_3\exp(-c_4 n) + c_3\exp(-c_4\log n)$. The proof is complete. □

7.2 Proof of Theorem 4

Proof. For soundness, we note that $\boldsymbol{x}_i$ and $M(\boldsymbol{X}_i')$ are independent, and therefore

$$\mathbb{E}\mathcal{A}_{\boldsymbol{\theta}}(\boldsymbol{x}_i, M(\boldsymbol{X}_i')) = \mathbb{E}\langle M(\boldsymbol{X}_i') - \boldsymbol{\theta}, S_{\boldsymbol{\theta}}(\boldsymbol{x}_i)\rangle = \langle \mathbb{E}M(\boldsymbol{X}_i') - \boldsymbol{\theta}, \mathbb{E}S_{\boldsymbol{\theta}}(\boldsymbol{x}_i)\rangle = \mathbf{0}.$$

The last equality is true by the property of the score that $\mathbb{E}S_{\boldsymbol{\theta}}(\boldsymbol{z}) = \mathbf{0}$ for any $\boldsymbol{z} \sim f_{\boldsymbol{\theta}}$. As to the first absolute moment, we apply Jensen's inequality,

$$\begin{aligned}\mathbb{E}|\mathcal{A}_{\boldsymbol{\theta}}(\boldsymbol{x}_i, M(\boldsymbol{X}_i'))| &\leq \sqrt{\mathbb{E}\langle M(\boldsymbol{X}_i') - \boldsymbol{\theta}, S_{\boldsymbol{\theta}}(\boldsymbol{x}_i)\rangle^2} \\ &\leq \sqrt{\mathbb{E}(M(\boldsymbol{X}_i') - \boldsymbol{\theta})^\top(\mathrm{Var}S_{\boldsymbol{\theta}}(\boldsymbol{x}_i))(M(\boldsymbol{X}_i') - \boldsymbol{\theta})} \leq \sqrt{\mathbb{E}\|M(\boldsymbol{X}) - \boldsymbol{\theta}\|_2^2}\sqrt{\lambda_{\max}(\mathcal{I}(\boldsymbol{\theta}))}.\end{aligned}$$

For completeness, we first simplify

$$\sum_{i\in[n]} \mathbb{E}\mathbf{A}_{\boldsymbol{\theta}}(\boldsymbol{x}_i, M(\boldsymbol{X})) = \mathbb{E}\Big\langle M(\boldsymbol{X}) - \boldsymbol{\theta}, \sum_{i\in[n]} S_{\boldsymbol{\theta}}(\boldsymbol{x}_i)\Big\rangle = \mathbb{E}\Big\langle M(\boldsymbol{X}), \sum_{i\in[n]} S_{\boldsymbol{\theta}}(\boldsymbol{x}_i)\Big\rangle.$$

By the definition of score and that $\boldsymbol{x}_1, \cdots, \boldsymbol{x}_n$ are i.i.d., $\sum_{i\in[n]} S_{\boldsymbol{\theta}}(\boldsymbol{x}_i) = S_{\boldsymbol{\theta}}(\boldsymbol{x}_1, \cdots, \boldsymbol{x}_n) = S_{\boldsymbol{\theta}}(\boldsymbol{X})$. It follows that

$$\mathbb{E}\Big\langle M(\boldsymbol{X}), \sum_{i\in[n]} S_{\boldsymbol{\theta}}(\boldsymbol{x}_i)\Big\rangle = \mathbb{E}\Big\langle M(\boldsymbol{X}), S_{\boldsymbol{\theta}}(\boldsymbol{X})\Big\rangle = \sum_{j\in[d]} \mathbb{E}\left[M(\boldsymbol{X})_j \frac{\partial}{\partial\theta_j}\log f_{\boldsymbol{\theta}}(\boldsymbol{X})\right].$$

For each term in the right-side summation, one may exchange differentiation and inte-

gration thanks to the regularity conditions on $f_{\boldsymbol{\theta}}$, and therefore

$$\begin{aligned}\mathbb{E}\left[M(\boldsymbol{X})_j \frac{\partial}{\partial\theta_j}\log f_{\boldsymbol{\theta}}(\boldsymbol{X})\right] &= \mathbb{E}\left[M(\boldsymbol{X})_j (f_{\boldsymbol{\theta}}(\boldsymbol{X}))^{-1}\frac{\partial}{\partial\theta_j} f_{\boldsymbol{\theta}}(\boldsymbol{X})\right] \\ &= \frac{\partial}{\partial\theta_j}\mathbb{E}\left[M(\boldsymbol{X})_j (f_{\boldsymbol{\theta}}(\boldsymbol{X}))^{-1} f_{\boldsymbol{\theta}}(\boldsymbol{X})\right] = \frac{\partial}{\partial\theta_j}\mathbb{E}M(\boldsymbol{X})_j.\end{aligned}$$

□

7.2.1 Proof of Lemma 6

Proof. Let $A_i := \mathbf{A}_{\boldsymbol{\theta}}(\boldsymbol{x}_i, M(\boldsymbol{X}))$, $A_i' := \mathbf{A}_{\boldsymbol{\theta}}(\boldsymbol{x}_i, M(\boldsymbol{X}_i'))$, and let $Z^+ = \max(Z, 0)$ and $Z^- = -\min(Z, 0)$ denote the positive and negative parts of a random variables Z respectively. We have

$$\mathbb{E}A_i = \mathbb{E}A_i^+ - \mathbb{E}A_i^- = \int_0^\infty \mathbb{P}(A_i^+ > t)\mathrm{d}t - \int_0^\infty \mathbb{P}(A_i^- > t)\mathrm{d}t.$$

For the positive part, if $0 < T < \infty$ and $0 < \varepsilon < 1$, we have

$$\begin{aligned}\int_0^\infty \mathbb{P}(A_i^+ > t)\mathrm{d}t &= \int_0^T \mathbb{P}(A_i^+ > t)\mathrm{d}t + \int_T^\infty \mathbb{P}(A_i^+ > t)\mathrm{d}t \\ &\leq \int_0^T \left(e^\varepsilon \mathbb{P}(A_i^+ > t) + \delta\right)\mathrm{d}t + \int_T^\infty \mathbb{P}(A_i^+ > t)\mathrm{d}t \\ &\leq \int_0^\infty \mathbb{P}(A_i'^+ > t)\mathrm{d}t + 2\varepsilon\int_0^\infty \mathbb{P}(A_i'^+ > t)\mathrm{d}t + \delta T + \int_T^\infty \mathbb{P}(|A_i| > t)\mathrm{d}t.\end{aligned}$$

Similarly for the negative part,

$$\begin{aligned}\int_0^\infty \mathbb{P}(A_i^- > t)\mathrm{d}t &= \int_0^T \mathbb{P}(A_i^- > t)\mathrm{d}t + \int_T^\infty \mathbb{P}(A_i^- > t)\mathrm{d}t \\ &\geq \int_0^T \left(e^{-\varepsilon}\mathbb{P}(A_i'^- > t) - \delta\right)\mathrm{d}t + \int_T^\infty \mathbb{P}(A_i^- > t)\mathrm{d}t \\ &\geq \int_0^T \mathbb{P}(A_i'^- > t)\mathrm{d}t - 2\varepsilon\int_0^T \mathbb{P}(A_i'^- > t) - \delta T + \int_T^\infty \mathbb{P}(A_i^- > t)\mathrm{d}t \\ &\geq \int_0^\infty \mathbb{P}(A_i'^- > t)\mathrm{d}t - 2\varepsilon\int_0^\infty \mathbb{P}(A_i'^- > t) - \delta T.\end{aligned}$$

It then follows that

$$\mathbb{E}A_i \leq \int_0^\infty \mathbb{P}(A_i'^+ > t)\mathrm{d}t - \int_0^\infty \mathbb{P}(A_i'^- > t)\mathrm{d}t + 2\varepsilon\int_0^\infty \mathbb{P}(|A_i'| > t)\mathrm{d}t + 2\delta T + \int_T^\infty \mathbb{P}(|A_i| > t)\mathrm{d}t$$

$$= \mathbb{E}A_i' + 2\varepsilon\mathbb{E}|A_i| + 2\delta T + \int_T^\infty \mathbb{P}(|A_i| > t)\mathrm{d}t.$$

The proof is now complete by soundness (4.3). □

7.2.2 Proof of Lemma 8

Proof. For each $j \in [d]$, by Lemma 7, we have

$$\mathbb{E}_{\pi_j}\left(\frac{\partial}{\partial\theta_j}g_j(\boldsymbol{\theta})\right) = \mathbb{E}_{\pi_j}\left(\frac{\partial}{\partial\theta_j}\mathbb{E}[g_j(\boldsymbol{\theta})|\theta_j]\right) = \mathbb{E}_{\pi_j}\left[\frac{-\mathbb{E}[g_j(\boldsymbol{\theta})|\theta_j]\pi_j'(\theta_j)}{\pi_j(\theta_j)}\right]$$

Because $|g_j(\boldsymbol{\theta}) - \theta_j| \le \|g(\boldsymbol{\theta}) - \boldsymbol{\theta}\|_2 \le \mathbb{E}_{\boldsymbol{X}|\boldsymbol{\theta}}\|M(\boldsymbol{X}) - \boldsymbol{\theta}\|_2$ for every $\boldsymbol{\theta} \in \Theta$, we have

$$\begin{aligned}\mathbb{E}_{\pi_j}\left[\frac{-\mathbb{E}[g(\boldsymbol{\theta})|\theta_j]\pi_j'(\theta_j)}{\pi_j(\theta_j)}\right] &\ge \mathbb{E}_{\pi_j}\left[\frac{-\theta_j\pi_j'(\theta_j)}{\pi_j(\theta_j)}\right] - \mathbb{E}_{\pi_j}\left[\mathbb{E}_{\boldsymbol{X}|\boldsymbol{\theta}}\|M(\boldsymbol{X}) - \boldsymbol{\theta}\|_2 \cdot \left|\frac{\pi_j'(\theta_j)}{\pi_j(\theta_j)}\right|\right] \\ &\ge \mathbb{E}_{\pi_j}\left[\frac{-\theta_j\pi_j'(\theta_j)}{\pi_j(\theta_j)}\right] - \sqrt{\mathbb{E}_{\pi_j}\mathbb{E}_{\boldsymbol{X}|\boldsymbol{\theta}}\|M(\boldsymbol{X}) - \boldsymbol{\theta}\|_2^2}\sqrt{\mathbb{E}_{\pi_j}\left[\left(\frac{\pi_j'(\theta_j)}{\pi_j(\theta_j)}\right)^2\right]}.\end{aligned}$$

So we have obtained

$$\mathbb{E}_{\pi_j}\left(\frac{\partial}{\partial\theta_j}g_j(\boldsymbol{\theta})\right) \ge \mathbb{E}_{\pi_j}\left[\frac{-\theta_j\pi_j'(\theta_j)}{\pi_j(\theta_j)}\right] - \sqrt{\mathbb{E}_{\pi_j}\mathbb{E}_{\boldsymbol{X}|\boldsymbol{\theta}}\|M(\boldsymbol{X}) - \boldsymbol{\theta}\|_2^2}\sqrt{\mathbb{E}_{\pi_j}\left[\left(\frac{\pi_j'(\theta_j)}{\pi_j(\theta_j)}\right)^2\right]}.$$

Now we take expectation over $\boldsymbol{\pi}(\boldsymbol{\theta})/\pi_j(\theta_j)$ and sum over $j \in [d]$ to complete the proof. □

Acknowledgments

The research of T. Cai was supported in part by the National Science Foundation grants DMS-1712735 and DMS-2015259 and the National Institutes of Health grants R01-GM129781 and R01-GM123056. The research of L. Zhang was supported in part by the National Science Foundation grant DMS-2015378.

References

[1] Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In *ACM CCS 2016*, pages 308–318. ACM, 2016.

[2] John M. Abowd. The challenge of scientific reproducibility and privacy protection for statistical agencies. In *Census Scientific Advisory Committee*, 2016.

[3] Alekh Agarwal, Sahand Negahban, and Martin J Wainwright. Fast global convergence rates of gradient methods for high-dimensional statistical recovery. In *Advances in Neural Information Processing Systems*, pages 37–45, 2010.

[4] Richard Arratia and Louis Gordon. Tutorial on large deviations for the binomial distribution. *Bulletin of mathematical biology*, 51(1):125–131, 1989.

[5] Mitali Bafna and Jonathan Ullman. The price of selection in differential privacy. *arXiv preprint arXiv:1702.02970*, 2017.

[6] Rina Foygel Barber and John C Duchi. Privacy and statistical risk: Formalisms and minimax bounds. *arXiv preprint arXiv:1412.4451*, 2014.

[7] Raef Bassily, Vitaly Feldman, Kunal Talwar, and Abhradeep Guha Thakurta. Private stochastic convex optimization with optimal rates. In *Advances in Neural Information Processing Systems*, pages 11282–11291, 2019.

[8] Raef Bassily, Adam Smith, and Abhradeep Thakurta. Private empirical risk minimization: Efficient algorithms and tight error bounds. In *FOCS 2014*, pages 464–473. IEEE, 2014.

[9] Thomas Blumensath and Mike E Davies. Iterative hard thresholding for compressed sensing. *Applied and computational harmonic analysis*, 27(3):265–274, 2009.

[10] Mark Bun, Jonathan Ullman, and Salil Vadhan. Fingerprinting codes and the price of approximate differential privacy. In *STOC 2014*, pages 1–10. ACM, 2014.

[11] T Tony Cai, Yichen Wang, and Linjun Zhang. The cost of privacy: optimal rates of convergence for parameter estimation with differential privacy. *arXiv preprint arXiv:1902.04495*, 2019.

[12] Kamalika Chaudhuri and Claire Monteleoni. Privacy-preserving logistic regression. In *Advances in neural information processing systems*, pages 289–296, 2009.

[13] Kamalika Chaudhuri, Claire Monteleoni, and Anand D Sarwate. Differentially private empirical risk minimization. *Journal of Machine Learning Research*, 12(Mar):1069–1109, 2011.

[14] Apple Differential Privacy Team. Privacy at scale. Apple, 2017.

[15] Bolin Ding, Janardhan Kulkarni, and Sergey Yekhanin. Collecting telemetry data privately. In *NeurIPS 2017*, pages 3571–3580, 2017.

[16] Jinshuo Dong, Aaron Roth, and Weijie J Su. Gaussian differential privacy. *arXiv preprint arXiv:1905.02383*, 2019.

[17] Dheeru Dua and Casey Graff. UCI machine learning repository, 2017.

[18] John C Duchi, Michael I Jordan, and Martin J Wainwright. Minimax optimal procedures for locally private estimation. *J. Am. Stat. Assoc.*, 113(521):182–201, 2018.

[19] John C Duchi and Feng Ruan. The right complexity measure in locally private estimation: It is not the fisher information. *arXiv preprint arXiv:1806.05756*, 2018.

[20] Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our data, ourselves: Privacy via distributed noise generation. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 486–503. Springer, 2006.

[21] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *TCC 2006*, pages 265–284. Springer, 2006.

[22] Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407, 2014.

[23] Cynthia Dwork and Guy N Rothblum. Concentrated differential privacy. *arXiv preprint arXiv:1603.01887*, 2016.

[24] 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.

[25] Cynthia Dwork, Adam Smith, Thomas Steinke, and Jonathan Ullman. Exposed! a survey of attacks on private data. *Annu. Rev. Stat. Appl.*, 4:61–84, 2017.

[26] Cynthia Dwork, Adam Smith, Thomas Steinke, Jonathan Ullman, and Salil Vadhan. Robust traceability from trace amounts. In *FOCS 2015*, pages 650–669. IEEE, 2015.

[27] Cynthia Dwork, Weijie J Su, and Li Zhang. Differentially private false discovery rate control. *arXiv preprint arXiv:1807.04209*, 2018.

[28] Úlfar Erlingsson, Vasyl Pihur, and Aleksandra Korolova. Rappor: Randomized aggregatable privacy-preserving ordinal response. In *ACM CCS 2014*, pages 1054–1067. ACM, 2014.

[29] Stephen E Fienberg, Alessandro Rinaldo, and Xiaolin Yang. Differential privacy and the risk-utility tradeoff for multi-dimensional contingency tables. In *International Conference on Privacy in Statistical Databases*, pages 187–199. Springer, 2010.

[30] Prateek Jain, Ambuj Tewari, and Purushottam Kar. On iterative hard thresholding methods for high-dimensional m-estimation. In *NeurIPS 2014*, pages 685–693, 2014.

[31] Gautam Kamath, Jerry Li, Vikrant Singhal, and Jonathan Ullman. Privately learning high-dimensional distributions. *arXiv preprint arXiv:1805.00216*, 2018.

[32] Gautam Kamath, Vikrant Singhal, and Jonathan Ullman. Private mean estimation of heavy-tailed distributions. *arXiv preprint arXiv:2002.09464*, 2020.

[33] Shiva Prasad Kasiviswanathan, Homin K Lee, Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. What can we learn privately? *SIAM Journal on Computing*, 40(3):793–826, 2011.

[34] Daniel Kifer, Adam Smith, and Abhradeep Thakurta. Private convex empirical risk minimization and high-dimensional regression. In *COLT 2012*, pages 25.1–25.40, 2012.

[35] Si Kai Lee, Luigi Gresele, Mijung Park, and Krikamol Muandet. Privacy-preserving causal inference via inverse probability weighting. *arXiv preprint arXiv:1905.12592*, 2019.

[36] Si Kai Lee, Luigi Gresele, Mijung Park, and Krikamol Muandet. Private causal inference using propensity scores. *arXiv preprint arXiv:1905.12592*, 2019.

[37] Jing Lei. Differentially private m-estimators. In *NeurIPS 2011*, pages 361–369, 2011.

[38] Po-Ling Loh and Martin J Wainwright. Regularized m-estimators with nonconvexity: Statistical and algorithmic theory for local optima. *The Journal of Machine Learning Research*, 16(1):559–616, 2015.

[39] P McCullagh and John A Nelder. *Generalized Linear Models*, volume 37. CRC Press, 1989.

[40] Frank McSherry and Kunal Talwar. Mechanism design via differential privacy. In *FOCS 2007*, volume 7, pages 94–103, 2007.

[41] Ilya Mironov. Renyi differential privacy. In *CSF 2017*, pages 263–275. IEEE, 2017.

[42] Sasi Kumar Murakonda, Reza Shokri, and George Theodorakopoulos. Ultimate power of inference attacks: Privacy risks of high-dimensional models. *arXiv preprint arXiv:1905.12774*, 2019.

[43] Sahand Negahban, Bin Yu, Martin J Wainwright, and Pradeep K Ravikumar. A unified framework for high-dimensional analysis of m-estimators with decomposable regularizers. In *Advances in neural information processing systems*, pages 1348–1356, 2009.

[44] John Ashworth Nelder and Robert WM Wedderburn. Generalized linear models. *Journal of the Royal Statistical Society: Series A (General)*, 135(3):370–384, 1972.

[45] Yurii Nesterov. *Introductory Lectures on Convex Optimization: A Basic Course*, volume 87. Springer Science & Business Media, 2003.

[46] NhatHai Phan, Yue Wang, Xintao Wu, and Dejing Dou. Differential privacy preservation for deep auto-encoders: an application of human behavior prediction. In *Aaai*, volume 16, pages 1309–1316, 2016.

[47] Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. Membership inference attacks against machine learning models. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 3–18. IEEE, 2017.

[48] Charles Stein. A bound for the error in the normal approximation to the distribution of a sum of dependent random variables. In *Proceedings of the Sixth Berkeley Symposium on Mathematical Statistics and Probability, Volume 2: Probability Theory*. The Regents of the University of California, 1972.

[49] Charles Stein, Persi Diaconis, Susan Holmes, and Gesine Reinert. Use of exchangeable pairs in the analysis of simulations. In *Stein's Method*, pages 1–25. Institute of Mathematical Statistics, 2004.

[50] Thomas Steinke and Jonathan Ullman. Between pure and approximate differential privacy. *Journal of Privacy and Confidentiality*, 7(2), 2017.

[51] Thomas Steinke and Jonathan Ullman. Tight lower bounds for differentially private selection. In *FOCS 2017*, pages 552–563. IEEE, 2017.

[52] Pragya Sur, Yuxin Chen, and Emmanuel J Candès. The likelihood ratio test in high-dimensional logistic regression is asymptotically a rescaled chi-square. *arXiv preprint arXiv:1706.01191*, 2017.

[53] Gábor Tardos. Optimal probabilistic fingerprint codes. *Journal of the ACM (JACM)*, 55(2):10, 2008.

[54] Martin J Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge University Press, 2019.

[55] Larry Wasserman and Shuheng Zhou. A statistical framework for differential privacy. *J. Am. Stat. Assoc.*, 105(489):375–389, 2010.

[56] Huanyu Zhang, Gautam Kamath, Janardhan Kulkarni, and Zhiwei Steven Wu. Privately learning markov random fields. *arXiv preprint arXiv:2002.09463*, 2020.

A Omitted Proofs in Section 3

A.1 Proof of Lemma 1

Proof of Lemma 1. Consider two data sets $\boldsymbol{Z}$ and $\boldsymbol{Z}'$ that differ only by one datum, $(y, \boldsymbol{x}) \in \boldsymbol{Z}$ versus $(y', \boldsymbol{x}') \in \boldsymbol{Z}'$. For any t, we have

$$\begin{aligned}\|\boldsymbol{\beta}^{t+1}(\boldsymbol{Z}) - \boldsymbol{\beta}^{t+1}(\boldsymbol{Z}')\|_2 &\le \frac{\eta^0}{n}\left(|\psi'(\boldsymbol{x}^\top\boldsymbol{\beta}^t) - \Pi_R(y)|\|\boldsymbol{x}\|_2 + |\psi'((\boldsymbol{x}')^\top\boldsymbol{\beta}^t) - \Pi_R(y')|\|\boldsymbol{x}'\|_2\right)\\ &\le \frac{\eta^0}{n} 4(R + c_1)\sigma_{\boldsymbol{x}}\sqrt{d},\end{aligned}$$

where the last step follows from (D1) and (G1). By the Gaussian mechanism, Fact 1, $\boldsymbol{\beta}^{t+1}(\boldsymbol{Z})$ is $(\varepsilon/T, \delta/T)$-differentially private, implying that Algorithm 1 is (ε, δ)-differentially private. □

A.2 Proof of Theorem 1

Proof of Theorem 1. We shall first define several favorable events under which the desired convergence does occur, and then show that the probability that any of the favorable events fails to happen is negligible. The events are,

$$\mathcal{E}_1 = \{(3.3) \text{ and } (3.4) \text{ hold}\}, \mathcal{E}_2 = \{\Pi_R(y_i) = y_i, \forall i \in [n]\}, \mathcal{E}_3 = \{\|\boldsymbol{\beta}^t - \hat{\boldsymbol{\beta}}\|_2 \le 3, 0 \le t \le T\}.$$

Let us first analyze the behavior of Algorithm 1 under these events. The scaling of $n \ge K \cdot \left(Rd\sqrt{\log(1/\delta)}\log n \log\log n/\varepsilon\right)$ for a sufficiently large K implies that $n \ge K' d \log d$ for a sufficiently large K'. Since $\|\boldsymbol{\beta}_1 - \boldsymbol{\beta}\|_1 \le \sqrt{d}\|\boldsymbol{\beta}_1 - \boldsymbol{\beta}_2\|_2$ for all $\boldsymbol{\beta}_1, \boldsymbol{\beta}_2 \in \mathbb{R}^d$, the RSM condition (3.4) implies that for every t,

$$\langle \nabla\mathcal{L}_n(\boldsymbol{\beta}^t) - \nabla\mathcal{L}_n(\hat{\boldsymbol{\beta}}), \boldsymbol{\beta}^t - \hat{\boldsymbol{\beta}}\rangle \le \frac{4\gamma}{3}\|\boldsymbol{\beta}^t - \hat{\boldsymbol{\beta}}\|_2^2. \tag{A.1}$$

Similarly, under event $\mathcal{E}_3$, the RSC condition (3.3) implies that

$$\langle \nabla\mathcal{L}_n(\boldsymbol{\beta}^t) - \nabla\mathcal{L}_n(\hat{\boldsymbol{\beta}}), \boldsymbol{\beta}^t - \hat{\boldsymbol{\beta}}\rangle \ge \frac{2\alpha}{3}\|\boldsymbol{\beta}^t - \hat{\boldsymbol{\beta}}\|_2^2. \tag{A.2}$$

To analyze the convergence of Algorithm 1, define $\tilde{\boldsymbol{\beta}}^{t+1} = \boldsymbol{\beta}^t - \eta^0\nabla\mathcal{L}_n(\boldsymbol{\beta}^t)$, so that $\boldsymbol{\beta}^{t+1} = \tilde{\boldsymbol{\beta}}^{t+1} + \boldsymbol{w}_t$. Let $\hat{\boldsymbol{\beta}} = \arg\min_{\boldsymbol{\beta}} \mathcal{L}_n(\boldsymbol{\beta})$. It follows that

$$\|\boldsymbol{\beta}^{t+1} - \hat{\boldsymbol{\beta}}\|_2^2 \le \left(1 + \frac{\alpha}{4\gamma}\right)\|\tilde{\boldsymbol{\beta}}^{t+1} - \hat{\boldsymbol{\beta}}\|_2^2 + \left(1 + \frac{4\gamma}{\alpha}\right)\|\boldsymbol{w}_t\|_2^2. \tag{A.3}$$

Now for $\|\tilde{\boldsymbol{\beta}}^{t+1} - \hat{\boldsymbol{\beta}}\|_2^2$,

$$\|\tilde{\boldsymbol{\beta}}^{t+1} - \hat{\boldsymbol{\beta}}\|_2 = \|\boldsymbol{\beta}^t - \hat{\boldsymbol{\beta}}\|_2^2 - 2\eta^0 \langle \nabla \mathcal{L}_n(\boldsymbol{\beta}^t), \boldsymbol{\beta}^t - \hat{\boldsymbol{\beta}} \rangle + \left(\eta^0\right)^2 \|\nabla \mathcal{L}_n(\boldsymbol{\beta}^t)\|_2^2. \tag{A.4}$$

We would like to bound the last two terms via the strong convexity (A.2) and smoothness (A.1), as follows

$$\begin{aligned}
&\mathcal{L}_n(\tilde{\boldsymbol{\beta}}^{t+1}) - \mathcal{L}_n(\hat{\boldsymbol{\beta}}) = \mathcal{L}_n(\tilde{\boldsymbol{\beta}}^{t+1}) - \mathcal{L}_n(\boldsymbol{\beta}^t) + \mathcal{L}_n(\boldsymbol{\beta}^t) - \mathcal{L}_n(\hat{\boldsymbol{\beta}}) \\
&\leq \langle \nabla \mathcal{L}_n(\boldsymbol{\beta}^t), \tilde{\boldsymbol{\beta}}^{t+1} - \boldsymbol{\beta}^t \rangle + \frac{2\gamma}{3} \|\tilde{\boldsymbol{\beta}}^{t+1} - \boldsymbol{\beta}^t\|_2^2 + \langle \nabla \mathcal{L}_n(\boldsymbol{\beta}^t), \boldsymbol{\beta}^t - \hat{\boldsymbol{\beta}} \rangle - \frac{\alpha}{3} \|\boldsymbol{\beta}^t - \hat{\boldsymbol{\beta}}\|_2^2 \\
&= \langle \nabla \mathcal{L}_n(\boldsymbol{\beta}^t), \tilde{\boldsymbol{\beta}}^{t+1} - \hat{\boldsymbol{\beta}} \rangle + \frac{3}{8\gamma} \|\nabla \mathcal{L}_n(\boldsymbol{\beta}^t)\|_2^2 - \frac{\alpha}{3} \|\boldsymbol{\beta}^t - \hat{\boldsymbol{\beta}}\|_2^2 \\
&= \langle \nabla \mathcal{L}_n(\boldsymbol{\beta}^t), \tilde{\boldsymbol{\beta}}^{t} - \hat{\boldsymbol{\beta}} \rangle - \frac{3}{8\gamma} \|\nabla \mathcal{L}_n(\boldsymbol{\beta}^t)\|_2^2 - \frac{\alpha}{3} \|\boldsymbol{\beta}^t - \hat{\boldsymbol{\beta}}\|_2^2 \\
&= \langle \nabla \mathcal{L}_n(\boldsymbol{\beta}^t), \tilde{\boldsymbol{\beta}}^{t} - \hat{\boldsymbol{\beta}} \rangle - \frac{\eta^0}{2} \|\nabla \mathcal{L}_n(\boldsymbol{\beta}^t)\|_2^2 - \frac{\alpha}{3} \|\boldsymbol{\beta}^t - \hat{\boldsymbol{\beta}}\|_2^2.
\end{aligned}$$

Since $\mathcal{L}_n(\tilde{\boldsymbol{\beta}}^{t+1}) - \mathcal{L}_n(\hat{\boldsymbol{\beta}}) \geq 0$, the calculations above imply that

$$-2\eta^0 \langle \nabla \mathcal{L}_n(\boldsymbol{\beta}^t), \boldsymbol{\beta}^t - \hat{\boldsymbol{\beta}} \rangle + \left(\eta^0\right)^2 \|\nabla \mathcal{L}_n(\boldsymbol{\beta}^t)\|_2^2 \leq -\frac{\alpha}{2\gamma} \|\boldsymbol{\beta}^t - \hat{\boldsymbol{\beta}}\|_2^2.$$

Substituting back into (A.4) and (A.3) yields

$$\|\boldsymbol{\beta}^{t+1} - \hat{\boldsymbol{\beta}}\|_2^2 \leq \left(1 - \frac{\alpha}{4\gamma}\right) \|\boldsymbol{\beta}^t - \boldsymbol{\beta}^0\|_2^2 + \left(1 + \frac{4\gamma}{\alpha}\right) \|\boldsymbol{w}_t\|_2^2.$$

It follows by induction over t, the choice of $T = \frac{4\gamma}{\alpha} \log(9n)$ and $\|\boldsymbol{\beta}^0 - \hat{\boldsymbol{\beta}}\|_2 \leq 3$ that

$$\|\boldsymbol{\beta}^T - \hat{\boldsymbol{\beta}}\|_2^2 \leq \frac{1}{n} + \left(1 + \frac{4\gamma}{\alpha}\right) \sum_{k=0}^{T-1} \left(1 - \frac{\alpha}{4\gamma}\right)^{T-k-1} \|\boldsymbol{w}_k\|_2^2. \tag{A.5}$$

The noise term can be controlled by the following lemma:

Lemma 11. *For $X_1, X_2, \cdots, X_k \overset{i.i.d.}{\sim} \chi_d^2$, $\lambda > 0$ and $0 < \rho < 1$,*

$$\mathbb{P}\left(\sum_{j=1}^{k} \lambda \rho^j X_j > \frac{\rho \lambda d}{1-\rho} + t\right) \leq \exp\left(-\min\left(\frac{(1-\rho^2)t^2}{8\rho^2\lambda^2 d}, \frac{t}{8\rho\lambda}\right)\right).$$

To apply the tail bound, we let $\lambda = (\eta^0)^2 2B^2 \frac{d \log(2T/\delta)}{n^2(\varepsilon/T)^2}$. It follows that, with $t = K\lambda d$ for a sufficiently large constant K, the noise term in (A.5) is bounded by $K\lambda d \asymp$

$\left(\frac{Rd\sqrt{\log(1/\delta)}\log n}{n\varepsilon}\right)^2$ with probability at least $1 - c_3\exp(-c_4 d)$.

Therefore, we have shown so far that, under events $\mathcal{E}_1, \mathcal{E}_2, \mathcal{E}_3$, it holds with probability at least $1 - c_3\exp(-c_4 d)$ that

$$\|\boldsymbol{\beta}^T - \hat{\boldsymbol{\beta}}\|_2 \lesssim \sqrt{\frac{1}{n}} + \frac{Rd\sqrt{\log(1/\delta)}\log n}{n\varepsilon}. \tag{A.6}$$

Combining with the statistical rate of convergence of $\|\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*\|$ yields the desired rate of

$$\|\boldsymbol{\beta}^T - \boldsymbol{\beta}^*\|_2 \lesssim \sqrt{c(\sigma)}\left(\sqrt{\frac{d}{n}} + \frac{d\sqrt{\log(1/\delta)}\log^{3/2} n}{n\varepsilon}\right).$$

It remains to show that the events $\mathcal{E}_1, \mathcal{E}_2, \mathcal{E}_3$ occur with overwhelming probability.

- By Proposition 4, $\mathbb{P}(\mathcal{E}_1^c) \le c_3\exp(-c_4 n)$ under the assumptions of Theorem 1.
- We have $\mathbb{P}(\mathcal{E}_2^c) \le c_3\exp(-c_4\log n)$ by the choice of R, and assumptions (G1), (G2) which imply the following bound of moment generating function of y_i: we have

$$\begin{aligned}\log\mathbb{E}\exp\left(\lambda\cdot\frac{y_i - \psi'(\boldsymbol{x}_i^\top\boldsymbol{\beta})}{c(\sigma)}\Big|\boldsymbol{x}_i\right) &= \frac{1}{c(\sigma)}\left(\psi(\boldsymbol{x}_i^\top\boldsymbol{\beta}+\lambda) - \psi(\boldsymbol{x}_i^\top\boldsymbol{\beta}) - \lambda\psi'(\boldsymbol{x}_i^\top\boldsymbol{\beta})\right)\\ &\le \frac{1}{c(\sigma)}\cdot\frac{\lambda^2\psi''(\boldsymbol{x}_i^\top\boldsymbol{\beta}+\tilde{\lambda})}{2}\end{aligned}$$

  for some $\tilde{\lambda}\in(0,\lambda)$. It follows that $\mathbb{E}\exp\left(\lambda\cdot\frac{y_i-\psi'(\boldsymbol{x}_i^\top\boldsymbol{\beta})}{c(\sigma)}\Big|\boldsymbol{x}_i\right) \le \exp\left(\frac{c_2\lambda^2}{2c(\sigma)}\right)$ because $\|\psi''\|_\infty < c_2$.

- For $\mathcal{E}_3$, we have the following lemma to be proved in A.2.2

  Lemma 12. *Under the assumptions of Theorem 1, if $\|\boldsymbol{\beta}^0 - \hat{\boldsymbol{\beta}}\|_2 \le 3$, then $\|\boldsymbol{\beta}^t - \hat{\boldsymbol{\beta}}\|_2 \le 3$ for all $0\le t\le T$ with probability at least $1 - c_3\exp(-c_4 d)$.*

We have shown that $\sum_{i=1}^3\mathbb{P}(\mathcal{E}_i^c) \le c_3\exp(-c_4 d) + c_3\exp(-c_4 n) + c_3\exp(-c_4\log n)$. The proof is complete. □

A.2.1 Proof of Lemma 11

Proof of Lemma 11. Since $\mathbb{E}\sum_{j=1}^k\lambda\rho^j X_j \le \lambda d\sum_{j=1}^k\rho^j < \frac{\rho\lambda d}{1-\rho}$, we have

$$\mathbb{P}\left(\sum_{j=1}^k\lambda\rho^j X_j > \frac{\rho\lambda d}{1-\rho} + t\right) \le \mathbb{P}\left(\sum_{j=1}^k\lambda\rho^j(X_j - \mathbb{E}X_j) > t\right).$$

The (centered) χ^2_d random variable is sub-exponential with parameters $(2\sqrt{d}, 4)$, the weighted sum is also sub-exponential, with parameters at most $\left(2\lambda\sqrt{d}\sqrt{\sum_{j=1}^{k}\rho^{2j}}, 4\lambda\rho\right)$. The desired tail bound now follows directly from standard sub-exponential tail bounds. □

A.2.2 Proof of Lemma 12

Proof of Lemma 12. We prove the lemma by induction. Suppose $\|\boldsymbol{\beta}^t - \hat{\boldsymbol{\beta}}\|_2 \leq 3$, by (A.1) we have

$$\begin{aligned}
&\mathcal{L}_n(\boldsymbol{\beta}^{t+1}) - \mathcal{L}_n(\hat{\boldsymbol{\beta}}) = \mathcal{L}_n(\boldsymbol{\beta}^{t+1}) - \mathcal{L}_n(\boldsymbol{\beta}^t) + \mathcal{L}_n(\boldsymbol{\beta}^t) - \mathcal{L}_n(\hat{\boldsymbol{\beta}}) \\
&\leq \langle \nabla\mathcal{L}_n(\boldsymbol{\beta}^t), \boldsymbol{\beta}^{t+1} - \boldsymbol{\beta}^t\rangle + \frac{2\gamma}{3}\|\tilde{\boldsymbol{\beta}}^{t+1} - \boldsymbol{\beta}^t\|_2^2 + \langle \nabla\mathcal{L}_n(\boldsymbol{\beta}^t), \boldsymbol{\beta}^t - \hat{\boldsymbol{\beta}}\rangle \\
&= \frac{4\gamma}{3}\langle \boldsymbol{\beta}^t - \boldsymbol{\beta}^{t+1}, \boldsymbol{\beta}^{t+1} - \hat{\boldsymbol{\beta}}\rangle + \frac{2\gamma}{3}\|\tilde{\boldsymbol{\beta}}^{t+1} - \boldsymbol{\beta}^t\|_2^2 + \frac{4\gamma}{3}\langle \boldsymbol{w}_t, \boldsymbol{\beta}^{t+1} - \hat{\boldsymbol{\beta}}\rangle \\
&\leq \frac{2\gamma}{3}\left(\|\boldsymbol{\beta}^t - \hat{\boldsymbol{\beta}}\|_2^2 - \|\boldsymbol{\beta}^{t+1} - \hat{\boldsymbol{\beta}}\|_2^2\right) + \frac{16\gamma^2}{\alpha}\|\boldsymbol{w}_t\|_2^2 + \frac{\alpha}{9}\|\boldsymbol{\beta}^{t+1} - \hat{\boldsymbol{\beta}}\|_2^2.
\end{aligned}$$

Assume by contradiction that $\|\boldsymbol{\beta}^{t+1} - \hat{\boldsymbol{\beta}}\|_2 > 3$. By (3.3) and (A.2), we have $\mathcal{L}_n(\boldsymbol{\beta}^{t+1}) - \mathcal{L}_n(\hat{\boldsymbol{\beta}}) \geq \alpha\|\boldsymbol{\beta}^{t+1} - \hat{\boldsymbol{\beta}}\|_2$ and therefore

$$\left(2\gamma + \frac{2\alpha}{3}\right)\|\boldsymbol{\beta}^{t+1} - \hat{\boldsymbol{\beta}}\|_2 \leq 6\gamma + \frac{16\gamma^2}{\alpha}\|\boldsymbol{w}_t\|_2^2.$$

Recall that the coordinates of $\boldsymbol{w}_t$ are i.i.d. Gaussian with variance of the order $\frac{d\log(1/\delta)\log^2 n}{n^2\varepsilon^2}$. By the scaling of $n \geq K \cdot \left(Rd\sqrt{\log(1/\delta)}\log n \log\log n/\varepsilon\right)$ and the choice of $T \asymp \log n$, it holds with probability at least $1 - c_3\exp(-c_4 d)$ that $\frac{16\gamma^2}{\alpha}\|\boldsymbol{w}_t\|_2^2 < 2\alpha$ for every $0 \leq t \leq T$. We then have $\left(2\gamma + \frac{2\alpha}{3}\right)\|\boldsymbol{\beta}^{t+1} - \hat{\boldsymbol{\beta}}\|_2 \leq 6\gamma + 2\alpha$, which is a contradiction with the original assumption. □

A.3 Proof of Lemma 3

Proof of Lemma 3. Let T be the index set of the top s coordinates of $\boldsymbol{v}$ in terms of absolute values. We have

$$\begin{aligned}
\|\tilde{P}_s(\boldsymbol{v}) - \boldsymbol{v}\|_2^2 &= \sum_{j\in S^c} v_j^2 = \sum_{j\in S^c\cap T^c} v_j^2 + \sum_{j\in S^c\cap T} v_j^2 \\
&\leq \sum_{j\in S^c\cap T^c} v_j^2 + (1 + 1/c)\sum_{j\in S\cap T^c} v_j^2 + 4(1+c)\sum_{i\in[s]}\|\boldsymbol{w}_i\|_\infty^2.
\end{aligned}$$

The last step is true by observing that $|S \cap T^c| = |S^c \cap T|$ and applying the following lemma.

Lemma 13. *Let S and $\{\boldsymbol{w}\}_{i\in[s]}$ be defined as in Algorithm 3. For every $R_1 \subseteq S$ and $R_2 \in S^c$ such that $|R_1| = |R_2|$ and every $c > 0$, we have*

$$\|\boldsymbol{v}_{R_2}\|_2^2 \le (1+c)\|\boldsymbol{v}_{R_1}\|_2^2 + 4(1+1/c)\sum_{i\in[s]}\|\boldsymbol{w}_i\|_\infty^2.$$

Now, for an arbitrary $\hat{\boldsymbol{v}}$ with $\|\hat{\boldsymbol{v}}\|_0 = \hat{s} \le s$, let $\hat{S} = \mathrm{supp}(\hat{\boldsymbol{v}})$. We have

$$\frac{1}{|I|-s}\sum_{j\in T^c} v_j^2 = \frac{1}{|T^c|}\sum_{j\in T^c} v_j^2 \overset{(*)}{\le} \frac{1}{|(\hat{S})^c|}\sum_{j\in(\hat{S})^c} v_j^2 = \frac{1}{|I|-\hat{s}}\sum_{j\in(\hat{S})^c} v_j^2 \le \frac{1}{|I|-\hat{s}}\sum_{j\in(\hat{S})^c}\|\hat{\boldsymbol{v}}-\boldsymbol{v}\|_2^2$$

The (*) step is true because T^c is the collection of indices with the smallest absolute values, and $|T^c| \le |\hat{S}^c|$. We then combine the two displays above to conclude that

$$\begin{aligned}
\|\tilde{P}_s(\boldsymbol{v}) - \boldsymbol{v}\|_2^2 &\le \sum_{j\in S^c\cap T^c} v_j^2 + (1+1/c)\sum_{j\in S\cap T^c} v_j^2 + 4(1+c)\sum_{i\in[s]}\|\boldsymbol{w}_i\|_\infty^2 \\
&\le (1+1/c)\sum_{j\in T^c} v_j^2 + 4(1+c)\sum_{i\in[s]}\|\boldsymbol{w}_i\|_\infty^2 \\
&\le (1+1/c)\frac{|I|-s}{|I|-\hat{s}}\|\hat{\boldsymbol{v}}-\boldsymbol{v}\|_2^2 + 4(1+c)\sum_{i\in[s]}\|\boldsymbol{w}_i\|_\infty^2.
\end{aligned}$$

□

A.3.1 Proof of Lemma 13

Proof of Lemma 13. Let $\psi : R_2 \to R_1$ be a bijection. By the selection criterion of Algorithm 3, for each $j \in R_2$ we have $|v_j| + w_{ij} \le |v_{\psi(j)}| + w_{i\psi(j)}$, where i is the index of the iteration in which $\psi(j)$ is appended to S. It follows that, for every $c > 0$,

$$\begin{aligned}
v_j^2 &\le \left(|v_{\psi(j)}| + w_{i\psi(j)} - w_{ij}\right)^2 \\
&\le (1+1/c)v_{\psi(j)}^2 + (1+c)(w_{i\psi(j)} - w_{ij})^2 \le (1+1/c)v_{\psi(j)}^2 + 4(1+c)\|\boldsymbol{w}_i\|_\infty^2
\end{aligned}$$

Summing over j then leads to

$$\|\boldsymbol{v}_{R_2}\|_2^2 \le (1+1/c)\|\boldsymbol{v}_{R_1}\|_2^2 + 4(1+c)\sum_{i\in[s]}\|\boldsymbol{w}_i\|_\infty^2.$$

□

A.4 Proof of Lemma 4

Proof of Lemma 4. In view of Lemma 2, it suffices to control

$$\|\eta^0\nabla\mathcal{L}_n(\boldsymbol{\theta}^t;\boldsymbol{Z})-\eta^0\nabla\mathcal{L}_n(\boldsymbol{\theta}^t;\boldsymbol{Z}')\|_\infty\le(\eta^0/n)\|\nabla l(\boldsymbol{\theta};\boldsymbol{z})-\nabla l(\boldsymbol{\theta};\boldsymbol{z}')\|_\infty<(\eta^0/n)B.$$

It follows that each iteration of Algorithm 4 is $(\varepsilon/T,\delta/T)$ differentially private. The overall privacy of Algorithm 4 is then a consequence of composition theorem, Fact 2. □

A.5 Proof of Theorem 2

Proof of Theorem 2. We first introduce some notation useful throughout the proof.

- Let $S^t=\mathrm{supp}(\boldsymbol{\theta}^t)$, $S^{t+1}=\mathrm{supp}(\boldsymbol{\theta}^{t+1})$ and $S^*=\mathrm{supp}(\hat{\boldsymbol{\theta}})$, and define $I^t=S^{t+1}\cup S^t\cup S^*$.
- Let $\boldsymbol{g}^t=\nabla\mathcal{L}_n(\boldsymbol{\theta}^t)$ and $\eta_0=\eta/\gamma$, where γ is the constant in (3.7).
- Let $\boldsymbol{w}_1,\boldsymbol{w}_2,\cdots,\boldsymbol{w}_s$ be the noise vectors dded to $\boldsymbol{\theta}^t-\eta^0\nabla\mathcal{L}_n(\boldsymbol{\theta}^t;Z)$ when the support of $\boldsymbol{\theta}^{t+1}$ is iteratively selected. We define $\boldsymbol{W}=4\sum_{i\in[s]}\|\boldsymbol{w}_i\|_\infty^2$.

We start by analyzing $\mathcal{L}_n(\boldsymbol{\theta}^{t+1})-\mathcal{L}_n(\boldsymbol{\theta}^t)$. By the restricted smoothness property (3.7),

$$\begin{aligned}\mathcal{L}_n(\boldsymbol{\theta}^{t+1})-\mathcal{L}_n(\boldsymbol{\theta}^t)&\le\langle\boldsymbol{\theta}^{t+1}-\boldsymbol{\theta}^t,\boldsymbol{g}^t\rangle+\frac{\gamma}{2}\|\boldsymbol{\theta}^{t+1}-\boldsymbol{\theta}^t\|_2^2\\&=\frac{\gamma}{2}\left\|\boldsymbol{\theta}^{t+1}_{I^t}-\boldsymbol{\theta}^t_{I^t}+\frac{\eta}{\gamma}\boldsymbol{g}^t_{I^t}\right\|_2^2-\frac{\eta^2}{2\gamma}\left\|\boldsymbol{g}^t_{I^t}\right\|_2^2+(1-\eta)\langle\boldsymbol{\theta}^{t+1}-\boldsymbol{\theta}^t,\boldsymbol{g}^t\rangle.\quad\text{(A.7)}\end{aligned}$$

We make use of this expansion to analyze each term separately. We first branch out to the third term and obtain the following expression after some calculations.

Lemma 14. *For every $c>0$, we have*

$$\langle\boldsymbol{\theta}^{t+1}-\boldsymbol{\theta}^t,\boldsymbol{g}^t\rangle\le-\frac{\eta}{2\gamma}\left\|\boldsymbol{g}^t_{S^t\cup S^{t+1}}\right\|_2^2+(1/c)\left(4+\frac{\eta}{2\gamma}\right)\|\boldsymbol{g}^t_{S^{t+1}}\|_2^2+c\|\tilde{\boldsymbol{w}}_{S^{t+1}}\|_2^2+(1+c)\frac{\gamma}{2\eta}\boldsymbol{W}.$$

The lemma is proved in Section A.5.1. Combining Lemma 14 with (A.7) yields

$$\begin{aligned}&\mathcal{L}_n(\boldsymbol{\theta}^{t+1})-\mathcal{L}_n(\boldsymbol{\theta}^t)\\\le\;&\frac{\gamma}{2}\left\|\boldsymbol{\theta}^{t+1}_{I^t}-\boldsymbol{\theta}^t_{I^t}+\frac{\eta}{\gamma}\boldsymbol{g}^t_{I^t}\right\|_2^2-\frac{\eta^2}{2\gamma}\left\|\boldsymbol{g}^t_{I^t}\right\|_2^2-\frac{\eta(1-\eta)}{2\gamma}\left\|\boldsymbol{g}^t_{S^t\cup S^{t+1}}\right\|_2^2\end{aligned}$$

$$
\begin{aligned}
&+\frac{1-\eta}{c}\left(4+\frac{\eta}{2\gamma}\right)\|\boldsymbol{g}^t_{S^{t+1}}\|_2^2+(1-\eta)c\|\tilde{\boldsymbol{w}}_{S^{t+1}}\|_2^2+(1-\eta)(1+c)\frac{\gamma}{2\eta}\boldsymbol{W}\\
\leq\ &\frac{\gamma}{2}\left\|\boldsymbol{\theta}^{t+1}_{I^t}-\boldsymbol{\theta}^t_{I^t}+\frac{\eta}{\gamma}\boldsymbol{g}^t_{I^t}\right\|_2^2-\frac{\eta^2}{2\gamma}\left\|\boldsymbol{g}^t_{I^t\setminus(S^t\cup S^*)}\right\|_2^2-\frac{\eta^2}{2\gamma}\left\|\boldsymbol{g}^t_{S^t\cup S^*}\right\|_2^2-\frac{\eta(1-\eta)}{2\gamma}\left\|\boldsymbol{g}^t_{S^t\cup S^{t+1}}\right\|_2^2\\
&+\frac{1-\eta}{c}\left(4+\frac{\eta}{2\gamma}\right)\|\boldsymbol{g}^t_{S^{t+1}}\|_2^2+(1-\eta)c\|\tilde{\boldsymbol{w}}_{S^{t+1}}\|_2^2+(1-\eta)(1+c)\frac{\gamma}{2\eta}\boldsymbol{W}\\
\leq\ &\frac{\gamma}{2}\left\|\boldsymbol{\theta}^{t+1}_{I^t}-\boldsymbol{\theta}^t_{I^t}+\frac{\eta}{\gamma}\boldsymbol{g}^t_{I^t}\right\|_2^2-\frac{\eta^2}{2\gamma}\left\|\boldsymbol{g}^t_{I^t\setminus(S^t\cup S^*)}\right\|_2^2-\frac{\eta^2}{2\gamma}\left\|\boldsymbol{g}^t_{S^t\cup S^*}\right\|_2^2-\frac{\eta(1-\eta)}{2\gamma}\left\|\boldsymbol{g}^t_{S^{t+1}\setminus(S^t\cup S^*)}\right\|_2^2\\
&+\frac{1-\eta}{c}\left(4+\frac{\eta}{2\gamma}\right)\|\boldsymbol{g}^t_{S^{t+1}}\|_2^2+(1-\eta)c\|\tilde{\boldsymbol{w}}_{S^{t+1}}\|_2^2+(1-\eta)(1+c)\frac{\gamma}{2\eta}\boldsymbol{W}. \qquad \text{(A.8)}
\end{aligned}
$$

The last step is true because $S^{t+1}\setminus(S^t\cup S^*)$ is a subset of $S^t\cup S^{t+1}$. Now we analyze the first two terms $\frac{\gamma}{2}\left\|\boldsymbol{\theta}^{t+1}_{I^t}-\boldsymbol{\theta}^t_{I^t}+\frac{\eta}{\gamma}\boldsymbol{g}^t_{I^t}\right\|_2^2-\frac{\eta^2}{2\gamma}\left\|\boldsymbol{g}^t_{I^t\setminus(S^t\cup S^*)}\right\|_2^2$

Lemma 15. *Let α be the restricted strong convexity constant as stated in condition* (3.6)*. For every $c>1$, we have*

$$
\begin{aligned}
&\frac{\gamma}{2}\left\|\boldsymbol{\theta}^{t+1}_{I^t}-\boldsymbol{\theta}^t_{I^t}+\frac{\eta}{\gamma}\boldsymbol{g}^t_{I^t}\right\|_2^2-\frac{\eta^2}{2\gamma}\left\|\boldsymbol{g}^t_{I^t\setminus(S^t\cup S^*)}\right\|_2^2\\
&\leq\frac{3s^*}{s+s^*}\left(\eta\mathcal{L}_n(\hat{\boldsymbol{\theta}})-\eta\mathcal{L}_n(\boldsymbol{\theta}^t)+\frac{\gamma-\eta\alpha}{2}\|\hat{\boldsymbol{\theta}}-\boldsymbol{\theta}^t\|_2^2+\frac{\eta^2}{2\gamma}\|\boldsymbol{g}^t_{I^t}\|_2^2\right)\\
&\quad+\frac{\eta^2}{2c\gamma}(1+1/c)\|\boldsymbol{g}^t_{S^{t+1}}\|_2^2+\frac{(c+3)\gamma}{2}\boldsymbol{W}+\frac{\gamma}{2}\|\tilde{\boldsymbol{w}}_{S^{t+1}}\|_2^2.
\end{aligned}
$$

The lemma is proved in Section A.5.1. Substitution into (A.8) leads to

$$
\begin{aligned}
&\mathcal{L}_n(\boldsymbol{\theta}^{t+1})-\mathcal{L}_n(\boldsymbol{\theta}^t)\\
\leq\ &\frac{3s^*}{s+s^*}\left(\eta\mathcal{L}_n(\hat{\boldsymbol{\theta}})-\eta\mathcal{L}_n(\boldsymbol{\theta}^t)+\frac{\gamma-\eta\alpha}{2}\|\hat{\boldsymbol{\theta}}-\boldsymbol{\theta}^t\|_2^2+\frac{\eta^2}{2\gamma}\|\boldsymbol{g}^t_{I^t}\|_2^2\right)\\
&-\frac{\eta^2}{2\gamma}\left\|\boldsymbol{g}^t_{S^t\cup S^*}\right\|_2^2-\frac{\eta(1-\eta)}{2\gamma}\left\|\boldsymbol{g}^t_{S^{t+1}\setminus(S^t\cup S^*)}\right\|_2^2\\
&+(1/c)\left(4(1-\eta)+\frac{\eta}{2\gamma}+\frac{(1+1/c)\eta^2}{2\gamma}\right)\|\boldsymbol{g}^t_{S^{t+1}}\|_2^2+\frac{\gamma}{2}\left(c+3+\frac{(1+c)(1-\eta)}{\eta}\right)\boldsymbol{W}\\
&+\left((1-\eta)c+\frac{\gamma}{2}\right)\|\tilde{\boldsymbol{w}}_{S^{t+1}}\|_2^2.
\end{aligned}
$$

Up to this point, the inequality holds for every $0<\eta<1$ and $c>1$. We now specify the choice of these parameters: let $\eta=2/3$ and set c large enough so that

$$
\mathcal{L}_n(\boldsymbol{\theta}^{t+1})-\mathcal{L}_n(\boldsymbol{\theta}^t)\leq\ \frac{3s^*}{s+s^*}\left(\eta\mathcal{L}_n(\hat{\boldsymbol{\theta}})-\eta\mathcal{L}_n(\boldsymbol{\theta}^t)+\frac{\gamma-\eta\alpha}{2}\|\hat{\boldsymbol{\theta}}-\boldsymbol{\theta}^t\|_2^2+\frac{\eta^2}{2\gamma}\|\boldsymbol{g}^t_{I^t}\|_2^2\right)
$$

$$
\begin{aligned}
&- \frac{\eta^2}{4\gamma} \left\| \boldsymbol{g}^t_{S^t \cup S^*} \right\|_2^2 - \frac{\eta(1-\eta)}{4\gamma} \left\| \boldsymbol{g}^t_{S^{t+1} \setminus (S^t \cup S^*)} \right\|_2^2 \\
&+ \frac{\gamma}{2} \left(\frac{3c+7}{2} \right) \boldsymbol{W} + \left(\frac{c}{3} + \frac{\gamma}{2} \right) \| \tilde{\boldsymbol{w}}_{S^{t+1}} \|_2^2.
\end{aligned}
$$

Such a choice of c is available because γ is an absolute constant determined by the RSM condition. Now we set $s = 72(\gamma/\alpha)^2 s^*$, so that $\frac{3s^*}{s+s^*} \leq \frac{\alpha^2}{24\gamma(\gamma - \eta\alpha)}$, and $\frac{\alpha^2}{24\gamma(\gamma-\eta\alpha)} \leq 1/8$ because $\alpha < \gamma$. It follows that

$$
\begin{aligned}
\mathcal{L}_n(\boldsymbol{\theta}^{t+1}) - \mathcal{L}_n(\boldsymbol{\theta}^t) \leq \; & \frac{3s^*}{s+s^*} \left(\eta \mathcal{L}_n(\hat{\boldsymbol{\theta}}) - \eta \mathcal{L}_n(\boldsymbol{\theta}^t) \right) + \frac{\alpha^2}{48\gamma} \| \hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^t \|_2^2 + \frac{1}{36\gamma} \| \boldsymbol{g}^t_{I^t} \|_2^2 \\
& - \frac{1}{9\gamma} \left\| \boldsymbol{g}^t_{S^t \cup S^*} \right\|_2^2 - \frac{1}{18\gamma} \left\| \boldsymbol{g}^t_{S^{t+1} \setminus (S^t \cup S^*)} \right\|_2^2 \\
& + \frac{\gamma}{2} \left(\frac{3c+7}{2} \right) \boldsymbol{W} + \left(\frac{c}{3} + \frac{\gamma}{2} \right) \| \tilde{\boldsymbol{w}}_{S^{t+1}} \|_2^2.
\end{aligned}
$$

Because $\| \boldsymbol{g}^t_{I^t} \|_2^2 = \| \boldsymbol{g}^t_{S^t \cup S^*} \|_2^2 + \left\| \boldsymbol{g}^t_{S^{t+1} \setminus (S^t \cup S^*)} \right\|_2^2$, we have

$$
\begin{aligned}
\mathcal{L}_n(\boldsymbol{\theta}^{t+1}) - \mathcal{L}_n(\boldsymbol{\theta}^t) \leq \; & \frac{3s^*}{s+s^*} \left(\eta \mathcal{L}_n(\hat{\boldsymbol{\theta}}) - \eta \mathcal{L}_n(\boldsymbol{\theta}^t) \right) + \frac{\alpha^2}{48\gamma} \| \hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^t \|_2^2 - \frac{3}{36\gamma} \left\| \boldsymbol{g}^t_{S^t \cup S^*} \right\|_2^2 \\
& + \frac{\gamma}{2} \left(\frac{3c+7}{2} \right) \boldsymbol{W} + \left(\frac{c}{3} + \frac{\gamma}{2} \right) \| \tilde{\boldsymbol{w}}_{S^{t+1}} \|_2^2 \\
\leq \; & \frac{3s^*}{s+s^*} \left(\eta \mathcal{L}_n(\hat{\boldsymbol{\theta}}) - \eta \mathcal{L}_n(\boldsymbol{\theta}^t) \right) - \frac{3}{36\gamma} \left(\left\| \boldsymbol{g}^t_{S^t \cup S^*} \right\|_2^2 - \frac{\alpha^2}{4} \| \hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^t \|_2^2 \right) \\
& + \frac{\gamma}{2} \left(\frac{3c+7}{2} \right) \boldsymbol{W} + \left(\frac{c}{3} + \frac{\gamma}{2} \right) \| \tilde{\boldsymbol{w}}_{S^{t+1}} \|_2^2. \qquad \text{(A.9)}
\end{aligned}
$$

To continue the calculations, we invoke a lemma from [30]:

Lemma 16 ([30], Lemma 6)**.**

$$
\left\| \boldsymbol{g}^t_{S^t \cup S^*} \right\|_2^2 - \frac{\alpha^2}{4} \| \hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^t \|_2^2 \geq \frac{\alpha}{2} \left(\mathcal{L}_n(\boldsymbol{\theta}^t) - \mathcal{L}_n(\hat{\boldsymbol{\theta}}) \right).
$$

It then follows from (A.9) and the lemma that, for an appropriate constant C_γ,

$$
\mathcal{L}_n(\boldsymbol{\theta}^{t+1}) - \mathcal{L}_n(\boldsymbol{\theta}^t) \leq - \left(\frac{3\alpha}{72\gamma} + \frac{2s^*}{s+s^*} \right) \left(\mathcal{L}_n(\boldsymbol{\theta}^t) - \mathcal{L}_n(\hat{\boldsymbol{\theta}}) \right) + C_\gamma (\boldsymbol{W} + \| \tilde{\boldsymbol{w}}_{S^{t+1}} \|_2^2).
$$

The proof is now complete by adding $\mathcal{L}_n(\boldsymbol{\theta}^t) - \mathcal{L}_n(\hat{\boldsymbol{\theta}})$ to both sides of the inequality. $\square$

A.5.1 Proofs of Lemma 14 and Lemma 15

Proof of Lemma 14. Since $\boldsymbol{\theta}^{t+1}$ is an output from Noisy Hard Thresholding, we may write $\boldsymbol{\theta}^{t+1} = \tilde{\boldsymbol{\theta}}^{t+1} + \tilde{\boldsymbol{w}}_{S^{t+1}}$, so that $\tilde{\boldsymbol{\theta}}^{t+1} = \tilde{P}_s(\boldsymbol{\theta}^t - \eta^0 \nabla \mathcal{L}(\boldsymbol{\theta}^t; Z))$ and $\tilde{\boldsymbol{w}}$ is a vector consisting of d i.i.d. draws from Laplace $\left(\eta_0 B \cdot \frac{2\sqrt{3s\log(T/\delta)}}{n\varepsilon/T}\right)$.

$$\begin{aligned}\langle \boldsymbol{\theta}^{t+1} - \boldsymbol{\theta}^t, \boldsymbol{g}^t \rangle &= \langle \boldsymbol{\theta}^{t+1}_{S^{t+1}} - \boldsymbol{\theta}^t_{S^{t+1}}, \boldsymbol{g}^t_{S^{t+1}} \rangle - \langle \boldsymbol{\theta}^t_{S^t \setminus S^{t+1}}, \boldsymbol{g}^t_{S^t \setminus S^{t+1}} \rangle \\ &= \langle \tilde{\boldsymbol{\theta}}^{t+1}_{S^{t+1}} - \boldsymbol{\theta}^t_{S^{t+1}}, \boldsymbol{g}^t_{S^{t+1}} \rangle + \langle \tilde{\boldsymbol{w}}_{S^{t+1}}, \boldsymbol{g}^t_{S^{t+1}} \rangle - \langle \boldsymbol{\theta}^t_{S^t \setminus S^{t+1}}, \boldsymbol{g}^t_{S^t \setminus S^{t+1}} \rangle.\end{aligned}$$

It follows that, for every $c > 0$,

$$\langle \boldsymbol{\theta}^{t+1} - \boldsymbol{\theta}^t, \boldsymbol{g}^t \rangle \le -\frac{\eta}{\gamma} \|\boldsymbol{g}^t_{S^{t+1}}\|_2^2 + c\|\tilde{\boldsymbol{w}}_{S^{t+1}}\|_2^2 + (1/4c)\|\boldsymbol{g}^t_{S^{t+1}}\|_2^2 - \langle \boldsymbol{\theta}^t_{S^t \setminus S^{t+1}}, \boldsymbol{g}^t_{S^t \setminus S^{t+1}} \rangle. \quad \text{(A.10)}$$

Now for the last term in the display above, we have

$$\begin{aligned}-\langle \boldsymbol{\theta}^t_{S^t \setminus S^{t+1}}, \boldsymbol{g}^t_{S^t \setminus S^{t+1}} \rangle &\le \frac{\gamma}{2\eta} \left(\left\| \boldsymbol{\theta}^t_{S^t \setminus S^{t+1}} - \frac{\eta}{\gamma} \boldsymbol{g}^t_{S^t \setminus S^{t+1}} \right\|_2^2 - \left(\frac{\eta}{\gamma}\right)^2 \|\boldsymbol{g}^t_{S^t \setminus S^{t+1}}\|_2^2 \right) \\ &\le \frac{\gamma}{2\eta} \left\| \boldsymbol{\theta}^t_{S^t \setminus S^{t+1}} - \frac{\eta}{\gamma} \boldsymbol{g}^t_{S^t \setminus S^{t+1}} \right\|_2^2 - \frac{\eta}{2\gamma} \|\boldsymbol{g}^t_{S^t \setminus S^{t+1}}\|_2^2.\end{aligned}$$

We apply Lemma 13 to $\left\| \boldsymbol{\theta}^t_{S^t \setminus S^{t+1}} - \frac{\eta}{\gamma} \boldsymbol{g}^t_{S^t \setminus S^{t+1}} \right\|_2^2$ to obtain that, for every $c > 0$,

$$\begin{aligned}-\langle \boldsymbol{\theta}^t_{S^t \setminus S^{t+1}}, \boldsymbol{g}^t_{S^t \setminus S^{t+1}} \rangle &\le \frac{\gamma}{2\eta} \left[(1 + 1/c) \left\| \tilde{\boldsymbol{\theta}}^{t+1}_{S^{t+1} \setminus S^t} \right\|_2^2 + (1 + c)\boldsymbol{W} \right] - \frac{\eta}{2\gamma} \|\boldsymbol{g}^t_{S^t \setminus S^{t+1}}\|_2^2 \\ &= \frac{\eta}{2\gamma} \left[(1 + 1/c) \left\| \boldsymbol{g}^t_{S^{t+1} \setminus S^t} \right\|_2^2 + (1 + c)\frac{\gamma}{2\eta}\boldsymbol{W} \right] - \frac{\eta}{2\gamma} \|\boldsymbol{g}^t_{S^t \setminus S^{t+1}}\|_2^2.\end{aligned}$$

Plugging the inequality above back into (A.10) yields

$$\begin{aligned}\langle \boldsymbol{\theta}^{t+1} - \boldsymbol{\theta}^t, \boldsymbol{g}^t \rangle \le \; & -\frac{\eta}{\gamma} \|\boldsymbol{g}^t_{S^{t+1}}\|_2^2 + c\|\tilde{\boldsymbol{w}}_{S^{t+1}}\|_2^2 + (1/4c)\|\boldsymbol{g}^t_{S^{t+1}}\|_2^2 \\ & + \frac{\eta}{2\gamma} \left[(1 + 1/c) \left\| \boldsymbol{g}^t_{S^{t+1} \setminus S^t} \right\|_2^2 + (1 + c)\frac{\gamma}{2\eta}\boldsymbol{W} \right] - \frac{\eta}{2\gamma} \|\boldsymbol{g}^t_{S^t \setminus S^{t+1}}\|_2^2 \\ \le \; & \frac{\eta}{2\gamma} \left\| \boldsymbol{g}^t_{S^{t+1} \setminus S^t} \right\|_2^2 - \frac{\eta}{2\gamma} \|\boldsymbol{g}^t_{S^t \setminus S^{t+1}}\|_2^2 - \frac{\eta}{\gamma} \|\boldsymbol{g}^t_{S^{t+1}}\|_2^2 \\ & + (1/c)\left(4 + \frac{\eta}{2\gamma}\right) \|\boldsymbol{g}^t_{S^{t+1}}\|_2^2 + c\|\tilde{\boldsymbol{w}}_{S^{t+1}}\|_2^2 + (1 + c)\frac{\gamma}{2\eta}\boldsymbol{W}.\end{aligned}$$

Finally, we have

$$\langle \boldsymbol{\theta}^{t+1} - \boldsymbol{\theta}^t, \boldsymbol{g}^t \rangle \leq -\frac{\eta}{2\gamma} \left\| \boldsymbol{g}^t_{S^t \cup S^{t+1}} \right\|_2^2 + (1/c) \left(4 + \frac{\eta}{2\gamma} \right) \| \boldsymbol{g}^t_{S^{t+1}} \|_2^2 + c \| \tilde{\boldsymbol{w}}_{S^{t+1}} \|_2^2 + (1+c) \frac{\gamma}{2\eta} \boldsymbol{W}.$$

□

Proof of Lemma 15. Let R be a subset of $S^t \setminus S^{t+1}$ such that $|R| = |I^t \setminus (S^t \cup S^*)| = |S^{t+1} \setminus (S^t \cup S^*)|$. By the definition of $\tilde{\boldsymbol{\theta}}^{t+1}$ and Lemma 13, we have, for every $c > 1$,

$$\frac{\eta^2}{\gamma^2} \left\| \boldsymbol{g}^t_{I^t \setminus (S^t \cup S^*)} \right\|_2^2 = \| \tilde{\boldsymbol{\theta}}^{t+1}_{I^t \setminus (S^t \cup S^*)} \|_2^2 \geq (1 - 1/c) \left\| \boldsymbol{\theta}^t_R - \frac{\eta}{\gamma} \boldsymbol{g}^t_R \right\|_2^2 - c\boldsymbol{W}.$$

It follows that

$$\begin{aligned}
&\frac{\gamma}{2} \left\| \boldsymbol{\theta}^{t+1}_{I^t} - \boldsymbol{\theta}^t_{I^t} + \frac{\eta}{\gamma} \boldsymbol{g}^t_{I^t} \right\|_2^2 - \frac{\eta^2}{2\gamma} \left\| \boldsymbol{g}^t_{I^t \setminus (S^t \cup S^*)} \right\|_2^2 \\
&\leq \frac{\gamma}{2} \| \tilde{\boldsymbol{w}}_{S^{t+1}} \|_2^2 + \frac{\gamma}{2} \left\| \tilde{\boldsymbol{\theta}}^{t+1}_{I^t} - \boldsymbol{\theta}^t_{I^t} + \frac{\eta}{\gamma} \boldsymbol{g}^t_{I^t} \right\|_2^2 - \frac{\gamma}{2}(1 - 1/c) \left\| \boldsymbol{\theta}^t_R - \frac{\eta}{\gamma} \boldsymbol{g}^t_R \right\|_2^2 + \frac{c\gamma}{2} \boldsymbol{W} \\
&= \frac{\gamma}{2} \left\| \tilde{\boldsymbol{\theta}}^{t+1}_{I^t} - \boldsymbol{\theta}^t_{I^t} + \frac{\eta}{\gamma} \boldsymbol{g}^t_{I^t} \right\|_2^2 - \frac{\gamma}{2} \left\| | \tilde{\boldsymbol{\theta}}^{t+1}_R - \boldsymbol{\theta}^t_R + \frac{\eta}{\gamma} \boldsymbol{g}^t_R \right\|_2^2 + \frac{\gamma}{2}(1/c) \left\| \boldsymbol{\theta}^t_R - \frac{\eta}{\gamma} \boldsymbol{g}^t_R \right\|_2^2 + \frac{c\gamma}{2} \boldsymbol{W} + \frac{\gamma}{2} \| \tilde{\boldsymbol{w}}_{S^{t+1}} \|_2^2 \\
&\leq \frac{\gamma}{2} \left\| \tilde{\boldsymbol{\theta}}^{t+1}_{I^t \setminus R} - \boldsymbol{\theta}^t_{I^t \setminus R} + \frac{\eta}{\gamma} \boldsymbol{g}^t_{I^t \setminus R} \right\|_2^2 + \frac{\eta^2}{2c\gamma}(1 + 1/c) \left\| \boldsymbol{g}^t_{I^t \setminus (S^t \cup S^*)} \right\|_2^2 + \frac{c\gamma}{2} \boldsymbol{W} + \frac{\gamma}{2} \| \tilde{\boldsymbol{w}}_{S^{t+1}} \|_2^2.
\end{aligned}$$

The last inequality is obtained by applying Lemma 13 to $\left\| \boldsymbol{\theta}^t_R - \frac{\eta}{\gamma} \boldsymbol{g}^t_R \right\|_2^2$. Now we apply Lemma 3 to obtain

$$\begin{aligned}
&\frac{\gamma}{2} \left\| \boldsymbol{\theta}^{t+1}_{I^t} - \boldsymbol{\theta}^t_{I^t} + \frac{\eta}{\gamma} \boldsymbol{g}^t_{I^t} \right\|_2^2 - \frac{\eta^2}{2\gamma} \left\| \boldsymbol{g}^t_{I^t \setminus (S^t \cup S^*)} \right\|_2^2 \\
&\leq \frac{3\gamma}{4} \frac{|I^t \setminus R| - s}{|I^t \setminus R| - s^*} \left\| \tilde{\boldsymbol{\theta}}_{I^t \setminus R} - \boldsymbol{\theta}^t_{I^t \setminus R} + \frac{\eta}{\gamma} \boldsymbol{g}^t_{I^t \setminus R} \right\|_2^2 + \frac{3\gamma}{2} \boldsymbol{W} + \frac{\eta^2(1 + c^{-1})}{2c\gamma} \left\| \boldsymbol{g}^t_{I^t \setminus (S^t \cup S^*)} \right\|_2^2 + \frac{c\gamma}{2} \boldsymbol{W} + \frac{\gamma}{2} \| \tilde{\boldsymbol{w}}_{S^{t+1}} \|_2^2 \\
&\leq \frac{3\gamma}{4} \frac{2s^*}{s + s^*} \left\| \tilde{\boldsymbol{\theta}}_{I^t \setminus R} - \boldsymbol{\theta}^t_{I^t \setminus R} + \frac{\eta}{\gamma} \boldsymbol{g}^t_{I^t \setminus R} \right\|_2^2 + \frac{3\gamma}{2} \boldsymbol{W} + \frac{\eta^2}{2c\gamma}(1 + 1/c) \left\| \boldsymbol{g}^t_{S^{t+1}} \right\|_2^2 + \frac{c\gamma}{2} \boldsymbol{W} + \frac{\gamma}{2} \| \tilde{\boldsymbol{w}}_{S^{t+1}} \|_2^2.
\end{aligned}$$

The last step is true by observing that $|I^t \setminus R| \leq 2s^* + s$, and the inclusion $I^t \setminus (S^t \cup S^*) \subseteq S^{t+1}$. We continue to simplify,

$$\frac{\gamma}{2} \left\| \boldsymbol{\theta}^{t+1}_{I^t} - \boldsymbol{\theta}^t_{I^t} + \frac{\eta}{\gamma} \boldsymbol{g}^t_{I^t} \right\|_2^2 - \frac{\eta^2}{2\gamma} \left\| \boldsymbol{g}^t_{I^t \setminus (S^t \cup S^*)} \right\|_2^2$$

$$
\begin{aligned}
&\leq \frac{\gamma}{2}\frac{3s^*}{s+s^*}\left\|\tilde{\boldsymbol{\theta}}_{I^t} - \boldsymbol{\theta}^t_{I^t} + \frac{\eta}{\gamma}\boldsymbol{g}^t_{I^t}\right\|_2^2 + \frac{3\gamma}{2}\boldsymbol{W} + \frac{\eta^2}{2c\gamma}(1+1/c)\left\|\boldsymbol{g}^t_{S^{t+1}}\right\|_2^2 + \frac{c\gamma}{2}\boldsymbol{W} + \frac{\gamma}{2}\|\tilde{\boldsymbol{w}}_{S^{t+1}}\|_2^2 \\
&\leq \frac{3s^*}{s+s^*}\left(\eta\langle\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^t, \boldsymbol{g}^t\rangle + \frac{\gamma}{2}\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^t\|_2^2 + \frac{\eta^2}{2c\gamma}\|\boldsymbol{g}^t_{I^t}\|_2^2\right) \\
&+ \frac{\eta^2}{2c\gamma}(1+1/c)\left\|\boldsymbol{g}^t_{S^{t+1}}\right\|_2^2 + \frac{(c+3)\gamma}{2}\boldsymbol{W} + \frac{\gamma}{2}\|\tilde{\boldsymbol{w}}_{S^{t+1}}\|_2^2 \\
&\leq \frac{3s^*}{s+s^*}\left(\eta\mathcal{L}_n(\hat{\boldsymbol{\theta}}) - \eta\mathcal{L}_n(\boldsymbol{\theta}^t) + \frac{\gamma-\eta\alpha}{2}\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^t\|_2^2 + \frac{\eta^2}{2c\gamma}\|\boldsymbol{g}^t_{I^t}\|_2^2\right) \\
&\quad + \frac{\eta^2}{2c\gamma}(1+1/c)\left\|\boldsymbol{g}^t_{S^{t+1}}\right\|_2^2 + \frac{(c+3)\gamma}{2}\boldsymbol{W} + \frac{\gamma}{2}\|\tilde{\boldsymbol{w}}_{S^{t+1}}\|_2^2.
\end{aligned}
$$

□

A.6 Proof of Lemma 5

Proof of Lemma 5. For every pair of adjacent data sets $\boldsymbol{Z}, \boldsymbol{Z}'$ we have

$$
\begin{aligned}
\|\boldsymbol{\beta}^{t+0.5}(\boldsymbol{Z}) - \boldsymbol{\beta}^{t+0.5}(\boldsymbol{Z}')\|_\infty &\leq \frac{\eta^0}{n}\left(|\psi'(\boldsymbol{x}^\top\boldsymbol{\beta}^t) - \Pi_R(y)|\|\boldsymbol{x}\|_\infty + |\psi'((\boldsymbol{x}')^\top\boldsymbol{\beta}^t) - \Pi_R(y')|\|\boldsymbol{x}'\|_\infty\right) \\
&\leq \frac{\eta^0}{n}4(R+c_1)\sigma_{\boldsymbol{x}},
\end{aligned}
$$

where the last step follows from (D1') and (G1). Algorithm 5 is (ε, δ)-differentially private by Lemma 4. □

A.7 Omitted Steps in Section 7.1, Proof of Theorem 3

A.7.1 Proof of Lemma 9

Proof of Lemma 9. By Algorithm 5, $\boldsymbol{\beta}^k, \boldsymbol{\beta}^{k+1}$ are both s-sparse with $s = 4c_0(\gamma/\alpha)^2 s^*$. The scaling assumed in Theorem 3 guarantees that $n \geq Ks^* \log d\sqrt{\log(T/\delta)}/(\varepsilon/T)$ for a sufficiently large constant K, (3.4) implies

$$
\langle\nabla\mathcal{L}_n(\boldsymbol{\beta}^{k+1}) - \nabla\mathcal{L}_n(\boldsymbol{\beta}^k), \boldsymbol{\beta}^{k+1} - \boldsymbol{\beta}^k\rangle \leq \frac{4\gamma}{3}\|\boldsymbol{\beta}^{k+1} - \boldsymbol{\beta}^k\|_2^2. \tag{A.11}
$$

Similarly, because $\|\boldsymbol{\beta}^k - \hat{\boldsymbol{\beta}}\|_2 \leq 3$ by assumption, the RSC condition (3.3) implies that

$$
\langle\nabla\mathcal{L}_n(\boldsymbol{\beta}^k) - \nabla\mathcal{L}_n(\hat{\boldsymbol{\beta}}), \boldsymbol{\beta}^k - \hat{\boldsymbol{\beta}}\rangle \geq \frac{2\alpha}{3}\|\boldsymbol{\beta}^k - \hat{\boldsymbol{\beta}}\|_2^2. \tag{A.12}
$$

Let $\boldsymbol{g}^k = \nabla \mathcal{L}_n(\boldsymbol{\beta}^k; Z)$. It follows from (A.11) and (A.12) that,

$$
\begin{aligned}
&\mathcal{L}_n(\boldsymbol{\beta}^{k+1}) - \mathcal{L}_n(\hat{\boldsymbol{\beta}}) \\
&= \mathcal{L}_n(\boldsymbol{\beta}^{k+1}) - \mathcal{L}_n(\boldsymbol{\beta}^k) + \mathcal{L}_n(\boldsymbol{\beta}^k) - \mathcal{L}_n(\hat{\boldsymbol{\beta}}) \\
&\leq \langle \boldsymbol{g}^k, \boldsymbol{\beta}^{k+1} - \boldsymbol{\beta}^k \rangle + \frac{2\gamma}{3} \|\boldsymbol{\beta}^{k+1} - \boldsymbol{\beta}^k\|_2^2 + \langle \boldsymbol{g}^k, \boldsymbol{\beta}^k - \hat{\boldsymbol{\beta}} \rangle - \frac{\alpha}{3} \|\boldsymbol{\beta}^k - \hat{\boldsymbol{\beta}}\|_2^2 \\
&\leq \langle \boldsymbol{g}^k, \boldsymbol{\beta}^{k+1} - \hat{\boldsymbol{\beta}} \rangle + \gamma \|\boldsymbol{\beta}^{k+1} - \boldsymbol{\beta}^k\|_2^2 - \frac{\alpha}{3} \|\boldsymbol{\beta}^k - \hat{\boldsymbol{\beta}}\|_2^2 \\
&= \langle 2\gamma(\boldsymbol{\beta}^k - \boldsymbol{\beta}^{k+1}), \boldsymbol{\beta}^{k+1} - \hat{\boldsymbol{\beta}} \rangle + \gamma \|\boldsymbol{\beta}^{k+1} - \boldsymbol{\beta}^k\|_2^2 - \frac{\alpha}{3} \|\boldsymbol{\beta}^k - \hat{\boldsymbol{\beta}}\|_2^2 + \langle \boldsymbol{g}^k - 2\gamma(\boldsymbol{\beta}^k - \boldsymbol{\beta}^{k+1}), \boldsymbol{\beta}^{k+1} - \hat{\boldsymbol{\beta}} \rangle \\
&= \left(\gamma - \frac{\alpha}{3}\right) \|\boldsymbol{\beta}^k - \hat{\boldsymbol{\beta}}\|_2^2 - \gamma \|\boldsymbol{\beta}^{k+1} - \hat{\boldsymbol{\beta}}\|_2^2 + \langle \boldsymbol{g}^k - 2\gamma(\boldsymbol{\beta}^k - \boldsymbol{\beta}^{k+1}), \boldsymbol{\beta}^{k+1} - \hat{\boldsymbol{\beta}} \rangle. \qquad \text{(A.13)}
\end{aligned}
$$

Let S^{k+1}, $\hat{S}$ denote the supports of $\boldsymbol{\beta}^{k+1}$, $\hat{\boldsymbol{\beta}}$ respectively. Since $\boldsymbol{\beta}^{k+1}$ is an output from Noisy Hard Thresholding, we may write $\boldsymbol{\beta}^{k+1} = \tilde{\boldsymbol{\beta}}^{k+1} + \tilde{\boldsymbol{w}}_{S^{k+1}}$, so that $\tilde{\boldsymbol{\theta}}^{k+1} = \tilde{P}_s(\boldsymbol{\beta}^k - (1/2\gamma)\nabla \mathcal{L}_n(\boldsymbol{\beta}^k; Z))$ and $\tilde{\boldsymbol{w}}$ is the Laplace noise vector.

Now we continue the calculation. For the last term of (A.13),

$$
\begin{aligned}
&\langle \boldsymbol{g}^k - 2\gamma(\boldsymbol{\beta}^k - \boldsymbol{\beta}^{k+1}), \boldsymbol{\beta}^{k+1} - \hat{\boldsymbol{\beta}} \rangle \\
&= 2\gamma \langle \tilde{\boldsymbol{w}}_{S^{k+1}}, \boldsymbol{\beta}^{k+1} - \hat{\boldsymbol{\beta}} \rangle + 2\gamma \langle \tilde{\boldsymbol{\beta}}^{k+1} - \boldsymbol{\beta}^k + (1/2\gamma)\boldsymbol{g}^k, \boldsymbol{\beta}^{k+1} - \hat{\boldsymbol{\beta}} \rangle \\
&\leq \frac{36\gamma^2}{\alpha} \|\tilde{\boldsymbol{w}}_{S^{k+1}}\|_2^2 + \frac{36\gamma^2}{\alpha} \|(\tilde{\boldsymbol{\beta}}^{k+1} - \boldsymbol{\beta}^k + (1/2\gamma)\boldsymbol{g}^k)_{S^{k+1} \cup \hat{S}}\|_2^2 + \frac{2\alpha}{9} \|\boldsymbol{\beta}^{k+1} - \hat{\boldsymbol{\beta}}\|_2^2 \qquad \text{(A.14)}
\end{aligned}
$$

For the middle term of (A.14), since $S^{k+1} \subseteq S^{k+1} \cup \hat{S}$, we have $\tilde{P}_s((\boldsymbol{\beta}^k + (1/2\gamma)\boldsymbol{g}^k)_{S^{k+1} \cup \hat{S}}) = \tilde{\boldsymbol{\beta}}^{k+1}_{S^{k+1} \cup \hat{S}}$, and therefore Lemma 3 applies. Because $|S^{k+1} \cup \hat{S}| \leq s + s^*$, we have

$$
\begin{aligned}
&\|(\tilde{\boldsymbol{\beta}}^{k+1} - \boldsymbol{\beta}^k + (1/2\gamma)\boldsymbol{g}^k)_{S^{k+1} \cup \hat{S}}\|_2^2 \\
&\leq \frac{5}{4} \frac{s^*}{s} \|(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}^k + (1/2\gamma)\boldsymbol{g}^k)_{S^{k+1} \cup \hat{S}}\|_2^2 + 20 \sum_{i \in [s]} \|\boldsymbol{w}_i\|_\infty^2 \\
&\leq \frac{5\alpha^2}{16 c_0 \gamma^2} \left(\frac{5}{3} \|\boldsymbol{\beta}^k - \hat{\boldsymbol{\beta}}\|_2^2 + \frac{5/2}{4\gamma^2} \|\boldsymbol{g}^k\|_2^2 \right) + 20 \sum_{i \in [s]} \|\boldsymbol{w}_i\|_\infty^2 \leq \frac{125\alpha^2}{16 c_0 \gamma^2} + 20 \sum_{i \in [s]} \|\boldsymbol{w}_i\|_\infty^2.
\end{aligned}
$$

For the last step to go through, we invoke the assumption that $\|\boldsymbol{\beta}^k - \hat{\boldsymbol{\beta}}\|_2 < 3$ and we have $\|\boldsymbol{g}^k\|_2^2 = \|\nabla \mathcal{L}_n(\boldsymbol{\beta}^k) - \nabla \mathcal{L}_n(\hat{\boldsymbol{\beta}})\|_2^2 \leq (4\gamma/3)^2 \|\boldsymbol{\beta}^k - \hat{\boldsymbol{\beta}}\|_2^2 \leq 16\gamma^2$ by (A.11). We recall from the proof of Theorem 2 that $c_0 = 72$; substituting the inequality above into (A.14) yields

$$
\langle \boldsymbol{g}^k - 2\gamma(\boldsymbol{\beta}^k - \boldsymbol{\beta}^{k+1}), \boldsymbol{\beta}^{k+1} - \hat{\boldsymbol{\beta}} \rangle
$$

$$\leq \frac{125\alpha}{32} + \frac{36\gamma^2}{\alpha}\left(\|\tilde{\boldsymbol{w}}_{S^{k+1}}\|_2^2 + 20\sum_{i\in[s]}\|\boldsymbol{w}_i\|_\infty^2\right) + \frac{2\alpha}{9}\|\boldsymbol{\beta}^{k+1} - \hat{\boldsymbol{\beta}}\|_2^2. \tag{A.15}$$

To analyze the noise term in the middle, we apply Lemma 10. Because the individual coordinates of $\tilde{\boldsymbol{w}}$, $\boldsymbol{w}_i$ are sampled i.i.d. from the Laplace distribution with scale $(2\gamma)^{-1} \cdot \frac{2\sqrt{3s\log(T/\delta)}}{n\varepsilon/T}$, if $n \geq Ks^* \log d\sqrt{\log(T/\delta)}/(\varepsilon/T)$ for a sufficiently large constant K, Lemma 10 implies that, with probability at least $1 - c_3\exp(-c_4\log(d/s^*))$ for some appropriate constants c_3, c_4, the noise term $(36\gamma^2/\alpha)\left(\|\tilde{\boldsymbol{w}}_{S^{k+1}}\|_2^2 + 20\sum_{i\in[s]}\|\boldsymbol{w}_i\|_\infty^2\right) < 3\alpha/32$. We substitute this upper bound back into (A.15), and then combine (A.15) with (A.13) to obtain

$$\mathcal{L}_n(\boldsymbol{\beta}^{k+1}) - \mathcal{L}_n(\hat{\boldsymbol{\beta}}) \leq \left(\gamma - \frac{\alpha}{3}\right)\|\boldsymbol{\beta}^k - \hat{\boldsymbol{\beta}}\|_2^2 - \left(\gamma - \frac{2\alpha}{9}\right)\|\boldsymbol{\beta}^{k+1} - \hat{\boldsymbol{\beta}}\|_2^2 + 4\alpha. \tag{A.16}$$

Let us now assume by contradiction that $\|\boldsymbol{\beta}^{k+1} - \hat{\boldsymbol{\beta}}\|_2 > 3$. From (3.3) and (A.12) we know that $\mathcal{L}_n(\boldsymbol{\beta}^{k+1}) - \mathcal{L}_n(\hat{\boldsymbol{\beta}}) \geq \alpha\|\boldsymbol{\beta}^{k+1} - \hat{\boldsymbol{\beta}}\|_2$. We combine this observation, the assumptions that $\|\boldsymbol{\beta}^{k+1} - \hat{\boldsymbol{\beta}}\|_2 > 3, \|\boldsymbol{\beta}^k - \hat{\boldsymbol{\beta}}\|_2 < 3$ and (A.16) to obtain

$$\left(3\gamma + \frac{\alpha}{3}\right)\|\boldsymbol{\beta}^{k+1} - \hat{\boldsymbol{\beta}}\|_2 \leq 9\gamma + \alpha,$$

which contradicts the original assumption that $\|\boldsymbol{\beta}^{k+1} - \hat{\boldsymbol{\beta}}\|_2 > 3$. □

A.7.2 Proof of Lemma 10

Proof of Lemma 10. By union bound and the i.i.d. assumption,

$$\mathbb{P}\left(\|\boldsymbol{w}\|_2^2 > kC^2\lambda^2\right) \leq k\mathbb{P}(w_1^2 > C^2\lambda^2) \leq ke^{-C}.$$

It follows that

$$\mathbb{P}\left(\|\boldsymbol{w}\|_\infty^2 > C^2\lambda^2\log^2 k\right) \leq k\mathbb{P}(w_1^2 > C^2\lambda^2\log^2 k) \leq ke^{-C\log k} = e^{-(C-1)\log k}.$$

□

B Omitted Proofs in Section 4

B.1 Proof of Theorem 5

Proof of Theorem 5. It can be shown via Theorem 4 that the score attack (4.7) is indeed sound and complete:

Lemma 17. *Under the assumptions of Theorem 5, the score attack* (4.7) *satisfies the following properties.*

1. *Soundness: For each $i \in [n]$ let $(\boldsymbol{y}_i', \boldsymbol{X}_i')$ denote the data set obtained by replacing $(y_i, \boldsymbol{x}_i)$ in $(\boldsymbol{y}, \boldsymbol{X})$ with an independent copy, then $\mathbb{E}\mathbf{A}_{\boldsymbol{\beta}}((y_i, \boldsymbol{x}_i), M(\boldsymbol{y}_i', \boldsymbol{X}_i')) = 0$ and $\mathbb{E}|\mathbf{A}_{\boldsymbol{\beta}}((y_i, \boldsymbol{x}_i), M(\boldsymbol{y}_i', \boldsymbol{X}_i'))| \le \sqrt{\mathbb{E}\|M(\boldsymbol{y}, \boldsymbol{X}) - \boldsymbol{\beta}\|_2^2}\sqrt{Cc_2/c(\sigma)}$.*
2. *Completeness: $\sum_{i\in[n]} \mathbb{E}\mathbf{A}_{\boldsymbol{\beta}}((y_i, \boldsymbol{x}_i), M(\boldsymbol{y}, \boldsymbol{X})) = \sum_{j\in[d]} \frac{\partial}{\partial \beta_j}\mathbb{E}M(\boldsymbol{y}, \boldsymbol{X})_j$.*

We follow the strategy outlined in Section 4.1 to establish appropriate upper and lower bounds for $\sum_{i\in[n]} \mathbb{E}\mathbf{A}_{\boldsymbol{\beta}}((y_i, \boldsymbol{x}_i), M(\boldsymbol{y}, \boldsymbol{X}))$, using Lemma 17.

Step 1. upper bounding the score attacks. We first work on the upper bound. Define $A_i = \mathbf{A}_{\boldsymbol{\beta}}((y_i, \boldsymbol{x}_i), M(\boldsymbol{y}, \boldsymbol{X}))$; the soundness part of Lemma 17 and Lemma 6 together imply that

$$\mathbb{E}A_i \le 2\varepsilon\sqrt{\mathbb{E}\|M(\boldsymbol{y}, \boldsymbol{X}) - \boldsymbol{\beta}\|_2^2}\sqrt{Cc_2/c(\sigma)} + 2\delta T + \int_T^\infty \mathbb{P}(|A_i| > t)\mathrm{d}t.$$

We need to choose T so that the remainder terms are controlled. We have

$$\begin{aligned}\mathbb{P}(|A_i| > t) &= \mathbb{P}\left(\left|\frac{y_i - \psi'(\boldsymbol{x}_i^\top\boldsymbol{\beta})}{c(\sigma)}\right| |\langle \boldsymbol{x}_i, M(\boldsymbol{y}, \boldsymbol{X}) - \boldsymbol{\beta}\rangle| > t\right)\\ &\le \mathbb{P}\left(\left|\frac{y_i - \psi'(\boldsymbol{x}_i^\top\boldsymbol{\beta})}{c(\sigma)}\right| d > t\right).\end{aligned}$$

For the first term, consider $f_\theta(y) = h(y, \sigma)\exp\left(\frac{y\theta - \psi(\theta)}{c(\sigma)}\right)$ and we have

$$\mathbb{E}\exp(\lambda y) = \int e^{\lambda y}h(y, \sigma)\exp\left(\frac{y\theta - \psi(\theta)}{c(\sigma)}\right)\mathrm{d}y = \exp\left(\frac{\psi(\theta + c(\sigma)\lambda) - \psi(\theta)}{c(\sigma)}\right).$$

We may then compute the moment generating function of $\frac{y_i - \psi'(\langle \boldsymbol{x}_i, \boldsymbol{\beta}\rangle)}{c(\sigma)}$, conditional on $\boldsymbol{x}_i$:

$$\log\mathbb{E}\exp\left(\lambda\cdot\frac{y_i - \psi'(\boldsymbol{x}_i^\top\boldsymbol{\beta})}{c(\sigma)}\Big|\boldsymbol{x}_i\right) = \frac{1}{c(\sigma)}\left(\psi(\boldsymbol{x}_i^\top\boldsymbol{\beta} + \lambda) - \psi(\boldsymbol{x}_i^\top\boldsymbol{\beta}) - \lambda\psi'(\boldsymbol{x}_i^\top\boldsymbol{\beta})\right)$$

$$\leq \frac{1}{c(\sigma)} \cdot \frac{\lambda^2 \psi''(\boldsymbol{x}_i^\top \boldsymbol{\beta} + \tilde{\lambda})}{2}$$

for some $\tilde{\lambda} \in (0, \lambda)$. It follows that $\mathbb{E} \exp\left(\lambda \cdot \frac{y_i - \psi'(\boldsymbol{x}_i^\top \boldsymbol{\beta})}{c(\sigma)} \middle| \boldsymbol{x}_i\right) \leq \exp\left(\frac{c_2 \lambda^2}{2c(\sigma)}\right)$ because $\|\psi''\|_\infty < c_2$. The bound for moment generating function implies that

$$\mathbb{P}(|A_i| > t) \leq \mathbb{P}\left(\left|\frac{y_i - \psi'(\boldsymbol{x}_i^\top \boldsymbol{\beta})}{c(\sigma)}\right| d > t\right) \leq \exp\left(-\frac{c(\sigma) t^2}{2 c_2 d^2}\right).$$

It follows that

$$\begin{aligned}
\mathbb{E} A_i &\leq 2\varepsilon \sqrt{\mathbb{E}\|M(\boldsymbol{y}, \boldsymbol{X}) - \boldsymbol{\beta}\|_2^2} \sqrt{C c_2 / c(\sigma)} + 2\delta T + \int_T^\infty \mathbb{P}(|A_i| > t) \mathrm{d}t \\
&\leq 2\varepsilon \sqrt{\mathbb{E}\|M(\boldsymbol{y}, \boldsymbol{X}) - \boldsymbol{\beta}\|_2^2} \sqrt{C c_2 / c(\sigma)} + 2\delta T + 2\sqrt{c_2 / c(\sigma)} d \exp\left(-\frac{c(\sigma) T^2}{2 c_2 d^2}\right).
\end{aligned}$$

We set $T = \sqrt{2c_2 / c(\sigma)} d \sqrt{\log(1/\delta)}$ to obtain

$$\sum_{i \in [n]} \mathbb{E} A_i \leq 2n\varepsilon \sqrt{\mathbb{E}\|M(\boldsymbol{y}, \boldsymbol{X}) - \boldsymbol{\beta}\|_2^2} \sqrt{C c_2 / c(\sigma)} + 4\sqrt{2} \delta d \sqrt{c_2 \log(1/\delta) / c(\sigma)}. \tag{B.1}$$

Step 2. lower bounding the score attacks. Next we prove a lower bound for $\sum_{i \in [n]} \mathbb{E} \mathbf{A}_{\boldsymbol{\beta}}((y_i, \boldsymbol{x}_i), M(\boldsymbol{y}, \boldsymbol{X}))$, or more precisely, an *average* lower bound with respect to an appropriately chosen prior distribution of $\boldsymbol{\beta}$. By completeness in Lemma 17,

$$\sum_{i \in [n]} \mathbb{E} \mathbf{A}_{\boldsymbol{\beta}}((y_i, \boldsymbol{x}_i), M(\boldsymbol{y}, \boldsymbol{X})) = \sum_{j \in [d]} \frac{\partial}{\partial \beta_j} \mathbb{E} M(\boldsymbol{y}, \boldsymbol{X})_j.$$

By Lemma 8 and the choice of $\boldsymbol{\pi}(\boldsymbol{\beta})$ as the density of $N(\mathbf{0}, \boldsymbol{I})$, Example 1 implies

$$\sum_{i \in [n]} \mathbb{E}_{\boldsymbol{\pi}} \mathbb{E}_{\boldsymbol{y}, \boldsymbol{X} | \boldsymbol{\beta}} \mathbf{A}_{\boldsymbol{\beta}}((y_i, \boldsymbol{x}_i), M(\boldsymbol{y}, \boldsymbol{X})) \gtrsim d. \tag{B.2}$$

Step 3. establishing the minimax risk lower bound. We combine (B.1) and (B.2) to prove the minimax risk lower bound (4.8). Since (B.1) holds for every fixed $\boldsymbol{\beta}$, for any choice of prior $\boldsymbol{\pi}$, we have

$$\begin{aligned}
d &\lesssim \mathbb{E}_{\boldsymbol{\pi}} \left[\sum_{i \in [n]} \mathbb{E} \mathbf{A}_{\boldsymbol{\beta}}((y_i, \boldsymbol{x}_i), M(\boldsymbol{y}, \boldsymbol{X}))\right] \\
&\leq 2n\varepsilon \mathbb{E}_{\boldsymbol{\pi}} \sqrt{\mathbb{E}_{\boldsymbol{y}, \boldsymbol{X} | \boldsymbol{\beta}} \|M(\boldsymbol{y}, \boldsymbol{X}) - \boldsymbol{\beta}\|_2^2} \sqrt{C c_2 / c(\sigma)} + 4\sqrt{2} n \delta d \sqrt{c_2 \log(1/\delta) / c(\sigma)}.
\end{aligned}$$

It follows that

$$2n\varepsilon\mathbb{E}_{\boldsymbol{\pi}}\sqrt{\mathbb{E}_{\boldsymbol{y},\boldsymbol{X}|\boldsymbol{\beta}}\|M(\boldsymbol{y},\boldsymbol{X})-\boldsymbol{\beta}\|_2^2}\sqrt{Cc_2/c(\sigma)} \gtrsim d - 4\sqrt{2}n\delta d\sqrt{c_2\log(1/\delta)/c(\sigma)}.$$

The assumption of $\delta < n^{-(1+\gamma)}$ implies $d - 4\sqrt{2}n\delta d\sqrt{C_1\log(1/\delta)/c(\sigma)} \gtrsim d$. We can then conclude that

$$\mathbb{E}_{\boldsymbol{\pi}}\mathbb{E}_{\boldsymbol{y},\boldsymbol{X}|\boldsymbol{\beta}}\|M(\boldsymbol{y},\boldsymbol{X})-\boldsymbol{\beta}\|_2^2 \gtrsim \frac{c(\sigma)d^2}{n^2\varepsilon^2}$$

The proof is complete because the minimax risk is always greater than the Bayes risk. □

B.1.1 Proof of Lemma 17

Proof. On the basis of Theorem 4, it suffices to calculate the score statistic of $f(y,\boldsymbol{x})$ with respect to $\boldsymbol{\beta}$ and the Fisher information matrix. In particular, all regularity conditions required for exchanging integration and differentiation are satisfied since $f_{\boldsymbol{\beta}}(y|x)$ is an exponential family. We have

$$\begin{aligned}\frac{\partial}{\partial\boldsymbol{\beta}}\log f(y,\boldsymbol{x}) &= \frac{\partial}{\partial\boldsymbol{\beta}}\log\left(f_{\boldsymbol{\beta}}(y|\boldsymbol{x})f(\boldsymbol{x})\right) = \frac{\partial}{\partial\boldsymbol{\beta}}\log f_{\boldsymbol{\beta}}(y|\boldsymbol{x}) \\ &= \frac{\partial}{\partial\boldsymbol{\beta}}\left(\frac{\boldsymbol{x}^\top\boldsymbol{\beta}y-\psi(\boldsymbol{x}^\top\boldsymbol{\beta})}{c(\sigma)}\right) = \frac{[y-\psi'(\boldsymbol{x}^\top\boldsymbol{\beta})]\boldsymbol{x}}{c(\sigma)}.\end{aligned}$$

For the Fisher information, we have

$$\mathcal{I}(\boldsymbol{\beta}) = -\mathbb{E}\left(\frac{\partial^2}{\partial\boldsymbol{\beta}^2}\log f(y,\boldsymbol{x})\right) = \mathbb{E}\left(\frac{\psi''(\boldsymbol{x}^\top\boldsymbol{\beta})}{c(\sigma)}\boldsymbol{x}\boldsymbol{x}^\top\right) \preceq \frac{c_2}{c(\sigma)}\mathbb{E}[\boldsymbol{x}\boldsymbol{x}^\top],$$

where the last inequality holds by $\|\psi''\|_\infty \le c_2$. We then have $\lambda_{\max}(\mathcal{I}(\boldsymbol{\beta})) \le Cc_2/c(\sigma)$ by $\lambda_{\max}(\mathbb{E}[\boldsymbol{x}\boldsymbol{x}^\top]) \le C$. □

B.2 Proof of Theorem 6

Proof of Theorem 6. By Theorem 4, the sparse GLM score attack is sound and complete. Similar to the classical GLM score attack, the sparse attack (4.9) also satisfies soundness and completeness.

Lemma 18. *Under the assumptions of Theorem 6, the score attack* (4.9) *satisfies the following properties.*

1. *Soundness: For each $i \in [n]$ let $(\boldsymbol{y}_i', \boldsymbol{X}_i')$ denote the data set obtained by replacing*

$(y_i, \boldsymbol{x}_i)$ *in* $(\boldsymbol{y}, \boldsymbol{X})$ *with an independent copy, then* $\mathbb{E}\mathbf{A}_{\boldsymbol{\beta},s^*}((y_i, \boldsymbol{x}_i), M(\boldsymbol{y}_i', \boldsymbol{X}_i')) = 0$ *and* $\mathbb{E}|\mathbf{A}_{\boldsymbol{\beta},s^*}((y_i, \boldsymbol{x}_i), M(\boldsymbol{y}_i', \boldsymbol{X}_i'))| \le \sqrt{\mathbb{E}\|(M(\boldsymbol{y}, \boldsymbol{X}) - \boldsymbol{\beta})_{\mathrm{supp}(\boldsymbol{\beta})}\|_2^2}\sqrt{Cc_2/c(\sigma)}$.

2. *Completeness:* $\sum_{i\in[n]} \mathbb{E}\mathbf{A}_{\boldsymbol{\beta},s^*}((y_i, \boldsymbol{x}_i), M(\boldsymbol{y}, \boldsymbol{X})) = \sum_{j\in[d]} \frac{\partial}{\partial \beta_j}\mathbb{E}(M(\boldsymbol{y}, \boldsymbol{X})_j \mathbb{1}(\beta_j \neq 0))$.

From the soundness and completeness properties, we can follow the strategy in Section 4.1 to derive the minimax risk lower bound.

Step 1. upper bounding the score attacks. Define $A_i = \mathbf{A}_{\boldsymbol{\beta},s^*}((y_i, \boldsymbol{x}_i), M(\boldsymbol{y}, \boldsymbol{X}))$; the soundness part of Lemma 18 and Lemma 6 together imply that

$$\mathbb{E}A_i \le 2\varepsilon\sqrt{\mathbb{E}\|(M(\boldsymbol{y}, \boldsymbol{X}) - \boldsymbol{\beta})_{\mathrm{supp}(\boldsymbol{\beta})}\|_2^2}\sqrt{Cc_2/c(\sigma)} + 2\delta T + \int_T^\infty \mathbb{P}(|A_i| > t)\mathrm{d}t.$$

We look for T such that the remainder terms are controlled. We have

$$\begin{aligned}\mathbb{P}(|A_i| > t) &= \mathbb{P}\left(\left|\frac{y_i - \psi'(\boldsymbol{x}_i^\top\boldsymbol{\beta})}{c(\sigma)}\right| \left|\langle \boldsymbol{x}_i, (M(\boldsymbol{y}, \boldsymbol{X}) - \boldsymbol{\beta})_{\mathrm{supp}(\boldsymbol{\beta})}\rangle\right| > t\right) \\ &\le \mathbb{P}\left(\left|\frac{y_i - \psi'(\boldsymbol{x}_i^\top\boldsymbol{\beta})}{c(\sigma)}\right| s^* > t\right).\end{aligned}$$

In the proof of Theorem 5, we have found $\mathbb{E}\exp\left(\lambda \cdot \frac{y_i - \psi'(\boldsymbol{x}_i^\top\boldsymbol{\beta})}{c(\sigma)}\middle|\boldsymbol{x}_i\right) \le \exp\left(\frac{c_2\lambda^2}{2c(\sigma)}\right)$. The bound for moment generating function then yields

$$\mathbb{P}(|A_i| > t) \le \mathbb{P}\left(\left|\frac{y_i - \psi'(\boldsymbol{x}_i^\top\boldsymbol{\beta})}{c(\sigma)}\right| s^* > t\right) \le \exp\left(-\frac{c(\sigma)t^2}{2c_2(s^*)^2}\right).$$

It follows that

$$\begin{aligned}\mathbb{E}A_i &\le 2\varepsilon\sqrt{\mathbb{E}\|(M(\boldsymbol{y}, \boldsymbol{X}) - \boldsymbol{\beta})_{\mathrm{supp}(\boldsymbol{\beta})}\|_2^2}\sqrt{Cc_2/c(\sigma)} + 2\delta T + \int_T^\infty \mathbb{P}(|A_i| > t)\mathrm{d}t \\ &\le 2\varepsilon\sqrt{\mathbb{E}\|M(\boldsymbol{y}, \boldsymbol{X}) - \boldsymbol{\beta}\|_2^2}\sqrt{Cc_2/c(\sigma)} + 2\delta T + 2s\sqrt{c_2/c(\sigma)}\exp\left(-\frac{c(\sigma)T^2}{2c_2(s^*)^2}\right).\end{aligned}$$

We choose $T = \sqrt{2c_2/c(\sigma)}s^*\sqrt{\log(1/\delta)}$ to obtain

$$\sum_{i\in[n]} \mathbb{E}A_i \le 2n\varepsilon\sqrt{\mathbb{E}\|M(\boldsymbol{y}, \boldsymbol{X}) - \boldsymbol{\beta}\|_2^2}\sqrt{Cc_2/c(\sigma)} + 4\sqrt{2}\delta s^*\sqrt{c_2\log(1/\delta)/c(\sigma)}. \tag{B.3}$$

Step 2. lower bounding the score attacks. By completeness in Lemma 18,

$$\sum_{i\in[n]} \mathbb{E}\mathbf{A}_{\boldsymbol{\beta},s^*}((y_i, \boldsymbol{x}_i), M(\boldsymbol{y}, \boldsymbol{X})) = \sum_{j\in[d]} \frac{\partial}{\partial \beta_j}\mathbb{E}(M(\boldsymbol{y}, \boldsymbol{X})_j \mathbb{1}(\beta_j \neq 0)).$$

For Lemma 8 to apply, we will choose some appropriate $\boldsymbol{\pi}(\boldsymbol{\beta})$; unlike the proof of Theorem 5, we have to find some prior distribution for $\boldsymbol{\beta}$ that obeys the sparsity condition $\|\boldsymbol{\beta}\|_0 \le s^*$. To this end, consider $\boldsymbol{\beta}$ generated as follows: let $\tilde{\beta}_1, \tilde{\beta}_2, \cdots, \tilde{\beta}_d$ be drawn i.i.d. from $N(0,1)$, let I_{s^*} be a subset of $[d]$ with $|I_{s^*}| = s^*$, and define $\beta_j = \tilde{\beta}_j \mathbb{1}(j \in I_{s^*})$, so that $\|\boldsymbol{\beta}\|_0 \le s^*$ by construction, and $\mathrm{supp}(\boldsymbol{\beta}) = I_{s^*}$.

It now follows from Lemma 8 that, if $\mathbb{E}_{\boldsymbol{y},\boldsymbol{X}|\boldsymbol{\beta}}\|(M(\boldsymbol{y},\boldsymbol{X}) - \boldsymbol{\beta})_{\mathrm{supp}(\boldsymbol{\beta})}\|_2^2 < 1$,

$$\mathbb{E}_{\boldsymbol{\pi}}\left(\sum_{j\in[d]} \frac{\partial}{\partial \beta_j}\mathbb{E}(M(\boldsymbol{y},\boldsymbol{X})_j \mathbb{1}(\beta_j \neq 0))\right) \ge \mathbb{E}_{\boldsymbol{\pi}}\left(\sum_{j\in[d]} \beta_j^2 \mathbb{1}(j \in I_{s^*})\right) - \sqrt{\mathbb{E}_{\boldsymbol{\pi}}\left(\sum_{j\in[d]} \beta_j^2 \mathbb{1}(j \in I_{s^*})\right)}.$$

We choose I_{s^*} to be the index set of $\tilde{\boldsymbol{\beta}}$ with top s^* greatest absolute values, and invoke the following lemma:

Lemma 19. *Consider $Z_1, Z_2, \cdots, Z_d$ drawn i.i.d. from $N(0,1)$, and let $|Z|_{(1)}, |Z|_{(2)}, \cdots, |Z|_{(d)}$ be the order statistics of $\{|Z_j|\}_{j\in[d]}$. If $s = o(d)$, for sufficiently large d we have*

$$\mathbb{E}|Z|_{(d-s+1)}^2 > c\log(d/s).$$

for some constant $c > 0$ and consequently $\sum_{k=0}^{s-1} \mathbb{E}|Z|_{(d-k)}^2 \gtrsim s\log(d/s)$.

Therefore, by our choice of prior for $\boldsymbol{\beta}$, we have

$$\sum_{i\in[n]} \mathbb{E}_{\boldsymbol{\pi}}\mathbb{E}\mathbf{A}_{\boldsymbol{\beta},s^*}((y_i,\boldsymbol{x}_i), M(\boldsymbol{y},\boldsymbol{X})) \gtrsim s^*\log(d/s^*) - \sqrt{s^*\log(d/s^*)} \gtrsim s^*\log(d/s^*). \quad \text{(B.4)}$$

Step 3. establishing the minimax risk lower bound. We combine (B.3) and (B.4) to prove the minimax risk lower bound (4.10). Since (B.3) holds for every fixed $\boldsymbol{\beta}$, for our choice of prior $\boldsymbol{\pi}$, we have

$$\begin{aligned} s^*\log(d/s^*) &\lesssim \mathbb{E}_{\boldsymbol{\pi}}\left[\sum_{i\in[n]} \mathbb{E}\mathbf{A}_{\boldsymbol{\beta},s^*}((y_i,\boldsymbol{x}_i), M(\boldsymbol{y},\boldsymbol{X}))\right] \\ &\le 2n\varepsilon\mathbb{E}_{\boldsymbol{\pi}}\sqrt{\mathbb{E}_{\boldsymbol{y},\boldsymbol{X}|\boldsymbol{\beta}}\|M(\boldsymbol{y},\boldsymbol{X}) - \boldsymbol{\beta}\|_2^2}\sqrt{Cc_2/c(\sigma)} + 4\sqrt{2}n\delta s^*\sqrt{c_2\log(1/\delta)/c(\sigma)}. \end{aligned}$$

It follows that

$$2n\varepsilon\mathbb{E}_{\boldsymbol{\pi}}\sqrt{\mathbb{E}_{\boldsymbol{y},\boldsymbol{X}|\boldsymbol{\beta}}\|M(\boldsymbol{y},\boldsymbol{X}) - \boldsymbol{\beta}\|_2^2}\sqrt{Cc_2/c(\sigma)} \gtrsim s^*\log(d/s^*) - 4\sqrt{2}n\delta s^*\sqrt{c_2\log(1/\delta)/c(\sigma)}.$$

The assumption that $\delta < n^{-(1+\gamma)}$ for some $\gamma > 0$ implies that for n sufficiently large,

$s^* \log(d/s^*) - 4\sqrt{2} n\delta s^* \sqrt{c_2 \log(1/\delta)/c(\sigma)} \gtrsim s^* \log(d/s^*)$. We then conclude that

$$\mathbb{E}_{\boldsymbol{\pi}} \mathbb{E}_{\boldsymbol{y},\boldsymbol{X}|\boldsymbol{\beta}} \|M(\boldsymbol{y}, \boldsymbol{X}) - \boldsymbol{\beta}\|_2^2 \gtrsim \frac{c(\sigma)(s^* \log(d/s^*))^2}{n^2 \varepsilon^2}$$

The proof is complete because the minimax risk is always greater than the Bayes risk. □

B.2.1 Proof of Lemma 18

Proof. We observe that $(M(\boldsymbol{y}, \boldsymbol{X}) - \boldsymbol{\beta})_{\mathrm{supp}(\boldsymbol{\beta})} = M(\boldsymbol{y}, \boldsymbol{X})_{\mathrm{supp}(\boldsymbol{\beta})} - \boldsymbol{\beta}$. The lemma is then a consequence of Theorem 4 and the score and Fisher information calculations in the proof of Lemma 17. □

B.2.2 Proof of Lemma 19

Proof. We denote $Y = |Z|_{(d-s+1)}$ and observe that

$$\mathbb{P}(Y > t) = 1 - \mathbb{P}(Y \leq t) = 1 - \mathbb{P}\left(\sum_{j \in [d]} \mathbb{1}(|Z_j| > t) \leq s \right)$$

Since $\mathbb{P}(|Z_i| > t) \geq t^{-1} \exp(-t^2/2)$ for $t \geq \sqrt{2}$ by Mills ratio, we choose $t = \sqrt{\log(d/2s)}$, so that $t^{-1} \exp(-t^2/2) > 2s/d$ as long as $d > 2s$. Now consider $N \sim \mathrm{Binomial}(d, 2s/d)$; we have $\mathbb{P}\left(\sum_{j \in [d]} \mathbb{1}(|Z_j| > t) \leq s \right) \leq \mathbb{P}(N \leq s)$. By standard Binomial tail bounds [4],

$$\begin{aligned} \mathbb{P}(N \leq s) &\leq \exp\left[-d \left((s/d) \log(1/2) + (1 - s/d) \log\left(\frac{1 - s/d}{1 - 2s/d} \right) \right) \right] \\ &\leq 2^s \left(1 - \frac{s}{d-s} \right)^{d-s} < (2/e)^s \end{aligned}$$

It follows that $\mathbb{P}\left(Y > \sqrt{\log(d/2s)} \right) > 1 - (2/e)^s > 0$. Because Y is non-negative, the proof is complete. □