High-dimensional variable selection with heterogeneous signals: A precise asymptotic perspective

SAPTARSHI ROY^{1,a}, AMBUJ TEWARI^{1,b} and ZIWEI ZHU^{2,c}

¹Department of Statistics, University of Michigan, Ann Arbor, USA, Email: ^aroysapta@umich.edu, ^btewaria@umich.edu ²Radix Trading, Chicago, USA Email: ^czzw9348ustc@gmail.com

We study the problem of exact support recovery for high-dimensional sparse linear regression under independent Gaussian design when the signals are weak, rare, and possibly heterogeneous. Under a suitable scaling of the sample size and signal sparsity, we fix the minimum signal magnitude at the information-theoretic optimal rate and investigate the asymptotic selection accuracy of best subset selection (BSS) and marginal screening (MS) procedures. We show that despite the ideal setup, somewhat surprisingly, marginal screening can fail to achieve exact recovery with probability converging to one in the presence of heterogeneous signals, whereas BSS enjoys model consistency whenever the minimum signal strength is above the information-theoretic threshold. To mitigate the computational intractability of BSS, we also propose an efficient two-stage algorithmic framework called ETS (Estimate Then Screen) comprised of an estimation step and gradient coordinate screening step, and under the same scaling assumption on sample size and sparsity, we show that ETS achieves model consistency under the same information-theoretic optimal requirement on the minimum signal strength as BSS. Finally, we present a simulation study comparing ETS with LASSO and marginal screening. The numerical results agree with our asymptotic theory even for realistic values of the sample size, dimension and sparsity.

Keywords: Heterogeneous signals; high-dimensional statistics; iterative hard thresholding; marginal screening; model consistency; variable selection

1. Introduction

Consider *n* independent observations $(x_i, y_i)_{i \in [n]}$ of a random pair (x, y) drawn from the following linear regression model:

$$(x,\varepsilon) \sim \mathcal{P}_x \times \mathcal{P}_{\varepsilon},$$

$$y = x^{\mathsf{T}} \beta + \varepsilon,$$
 (1)

where \mathcal{P}_x is the *p*-dimensional isotropic Gaussian distribution N_p(0, I_p), and $\mathcal{P}_{\varepsilon}$ is the standard Gaussian distribution on \mathbb{R} . In matrix notation, the observations can be represented as

$$Y = X\beta + E,$$

where $Y = (y_1, y_2, ..., y_n)^\top$, $X = (x_1, x_2, ..., x_n)^\top$ and $E = (\varepsilon_1, \varepsilon_2, ..., \varepsilon_n)^\top$. The vector β is unknown but sparse in the sense that $\|\beta\|_0 := \sum_{j=1}^p \mathbb{1}(\beta_j \neq 0) = s$, which is much smaller than p. Denote by $\mathcal{S}(v)$ the set of non-zero coordinates of a vector $v \in \mathbb{R}^p$. Lastly, we denote by $\mathbb{P}_{\beta_0}(\cdot)$ and $\mathbb{E}_{\beta_0}(\cdot)$ the probability measure and the expectation with $\beta = \beta_0$ respectively. In this paper, we focus on the variable selection problem, i.e., identifying $S(\beta)$. We primarily use the 0-1 loss, i.e., $\mathbb{P}_{\beta}(S(\hat{\beta}) \neq S(\beta))$, to assess the quality of the selected model $S(\hat{\beta})$. The isotropic Gaussian design has been widely used to conduct precise analysis of variable selection procedures (Fletcher, Rangan and Goyal, 2009, Genovese et al., 2012, Kowshik and Polyanskiy, 2021, Ndaoud and Tsybakov, 2020, Su, Bogdan and Candes, 2017). Specifically, these works either derive the necessary and sufficient condition for exact model recovery (Aeron, Saligrama and Zhao, 2010, Akçakaya and Tarokh, 2009, Fletcher, Rangan and Goyal, 2009, Jin, Kim and Rao, 2011, Rad, 2011), or establish tight asymptotic bounds of model selection error (Genovese et al., 2012, Ji et al., 2012, Su, Bogdan and Candes, 2017). The isotropic Gaussian design is also used in compressed sensing to generate a measurement matrix (Candès, Romberg and Tao, 2006, Candes and Tao, 2006, Donoho, 2006) so that one can sense the sparse signals with few measurements of the high-dimensional signal. Scarlett and Cevher (2016), Wang, Wainwright and Ramchandran (2010) considered the variable selection problem and studied the information-theoretic limit of support recovery under non-Gaussian setup. It is worth emphasizing that these works also assumed that the entries of X are independent and identically distributed, which is also the setup of this paper.

Recently there has been growing interest in the variable selection problem in the presence of weak and rare signal regimes (Genovese et al., 2012, Ji et al., 2012) where the active signals are highly sparse with very low magnitude of the order $O(\sqrt{(\log p)/n})$, which is known to be the informationtheoretic optimal rate necessary to achieve model consistency. This regime is ubiquitous in modern data analytics such as those in Genome-Wide Association Study (GWAS). There the genes that exhibit detectable association with the trait of interest can be extremely few with weak effects (Consortium et al., 2007, Marttinen et al., 2013). Moreover, the number of subjects n typically ranges in thousands, while the number of features p can range from tens of thousands to hundreds of thousands. Such a high dimension further adds to the difficulty of identifying the weak signals. Weak and rare signals also arise in multi-user detection problems (Arias-Castro, Candès and Plan, 2011) where one typically uses linear model of the form (1). There the *j*th column of X, denoted by X_j , is the channel impulse response for user j. The signal received from user j is $\beta_i X_j$. Thus $\beta_i = 0$ means that j th user is not sending any signal. It is a common practice to model the mixing matrix X as random with i.i.d. entries. Under the presence of strong noise, one might be interested in knowing whether information is being transmitted or not. Typically, in some applications, it is reasonable to assume that very few users are sending signals. Also due to strong noise environment the signals become quite weak, making them harder to detect. Therefore, from an application point of view, understanding variable selection in weak and rare signal regimes is crucial. Despite its importance, typically most of the popular methods such as LASSO (Tibshirani, 1996), SCAD (Fan and Li, 2001), adaptive LASSO (Huang, Ma and Zhang, 2008) have been extensively analyzed in terms of 0-1 loss when the signals are uniformly strong (Guo et al., 2015, Huang, Ma and Zhang, 2008, Zhang and Huang, 2008, Zhao and Yu, 2006, Zheng, Fan and Lv, 2014) in the sense that

$$a := \min_{j \in \mathcal{S}(\beta)} \left| \beta_j \right| \gg \left(\frac{\log p}{n} \right)^{1/2}.$$

However, Wainwright (2009a) established a sharp phase transition for LASSO in terms of exact recovery under a general combination of (n, p, s). Under some regularity conditions on the design matrix, the author shows that $a \gtrsim \{(\log p)/n\}^{1/2}$ is necessary and sufficient for the model consistency of LASSO in terms of 0-1 loss. Zhang and Huang (2008) proposed MC+ method based on *minimax concave penalty*, which also achieves model consistency under the optimal rate for *a*. Therefore, these works can accommodate weak and rare signal regimes, and it is important to emphasize that the analyses presented in them are *non-asymptotic*. Other works on weak and rare signal regimes include Genovese et al. (2012), Ji et al. (2012), and Jin, Zhang and Zhang (2014) that establish sharp *asymptotic* requirement on the signal strength.

Besides the weakness and rarity of signals, heterogeneity in the signal *strength* is another important feature of modern data applications that has not yet received sufficient attention. Roughly, *heterogeneity* in the signal allows the magnitude of the active β_j 's to differ in an arbitrary fashion, whereas *homogeneity* restricts the magnitude of the active signals to be in the same order. One limitation of the existing literature on variable selection in the *weak* and *rare* signals regime is that it typically assumes that the true signals are homogeneous (Genovese et al., 2012, Ji et al., 2012, Jin, Zhang and Zhang, 2014). Ji et al. (2012) refer to this setup as the *Asymptotically Rare and Weak* (ARW) signal regime. Many popular approaches have been shown to enjoy satisfactory variable selection properties under the ARW regime. For instance, Genovese et al. (2012) showed that both LASSO and marginal screening enjoy model consistency in terms of Hamming loss under independent random design. Ji et al. (2012) and Jin, Zhang and Zhang (2014) investigated the same problem under sparsely correlated design. They proposed two-stage screen and clean algorithms that also exhibit model consistency in terms of Hamming loss. However, their theory heavily relies on homogeneous signals and does not extend to the heterogeneous case that is of interest to us. In reality, the ARW setup seldom occurs: the signals almost always have different strengths (Li et al., 2019).

To underscore the contrasting effects of homogeneous and heterogeneous signals in terms of exact model recovery, we study the variable selection property of marginal screening (see Section 3). We show that under the presence of strong heterogeneity in the signal, marginal screening fails to recover the exact model with probability converging to 1, whereas under homogeneous signal it can recover the exact model asymptotically (Genovese et al., 2012). It turns out that due to heterogeneity, the spurious correlations become large and create impediment to selecting the exact model. In correlated design, a different problem known as unfaithfulness (Robins et al., 2003, Wasserman and Roeder, 2009) prevents marginal screening from achieving model consistency. Specifically, due to "correlation cancellation", the marginal correlation between Y and X_i becomes negligible even when β_i is large and this ultimately leads to false negatives. In this paper, we study independent random design model in which correlation cancellation does not occur. Instead, we identify a different source of problem under the presence of signal heterogeneity that affects the exact variable selection performance of marginal screening. Varying effect of signal heterogeneity in variable selection was also identified in the case of LASSO by Su, Bogdan and Candes (2017) and Wang, Yang and Su (2022) for i.i.d. Gaussian design under a special asymptotic setting. In particular, Su, Bogdan and Candes (2017) studied the tradeoff between power and type-I error of LASSO and showed that strong heterogeneity in the signal helps to reduce the false discovery in the LASSO path. The same effect was also analyzed in more detail in Wang, Yang and Su (2022). These works use approximate message passing (AMP) theory to obtain the exact asymptotic behavior of LASSO estimator in terms of variable selection and show that it is unable to achieve model consistency under linear sparsity regime.

On the computational side, modern methods like LASSO, SCAD, MC+ were initially motivated as alternatives to Best Subset Selection (BSS). BSS is in general an NP-hard optimization problem and was believed to be practically intractable even for p as small as 30. Thanks to recent advancements in algorithms and hardware, the optimal solution to the BSS problem can now be computed, sometimes with approximations, for some practical settings. Jain, Tewari and Kar (2014) showed that a wide family of iterative hard thresholding (IHT) algorithms can approximately solve the BSS problem, in the sense that they can achieve similar goodness of fit with the best subset with slight violation of the sparsity constraint. Liu and Foygel Barber (2020) studied the optimal thresholding operator for such iterative thresholding algorithms, which manages to exploit fewer variables than IHT to achieve the same goodness fit as BSS. Bertsimas, King and Mazumder (2016) viewed the BSS problem through the lens of mixed integer optimization (MIO) and showed that for n in 100s and p in 1000s, the MIO algorithm can obtain a near optimal solution reasonably fast. Bertsimas and Parys (2020) developed a new cutting plane method that solves to provable optimality the Tikhonov-regularized (Tikhonov, 1943) BSS problem for n and p in the 100,000s. Xie and Deng (2020) considered solving the Tikhonov-regularized

BSS via mixed integer second order cone formulation and the largest problem instance they considered has $p \sim 10^3$. Most recently, Hazimeh, Mazumder and Saab (2022) developed a Branch-and-Bound method that solves the ℓ_0/ℓ_2 -regularized BSS problem for $p \sim 10^7$. A recent work (Zhu et al., 2020) proposed an iterative splicing method called Adaptive Best Subset Selection (ABESS) to solve the BSS problem. They also showed that ABESS enjoys both statistical accuracy and polynomial computational complexity when the design matrix satisfies sparse Reisz condition and minimum signal strength is of order $\Omega\{(s \log p \log \log n/n)^{1/2}\}$.

Given these recent advances in solving BSS, there has been growing acknowledgment that BSS enjoys significant statistical superiority over the aforementioned alternative methods. Bertsimas, King and Mazumder (2016) and Bertsimas and Parys (2020) numerically demonstrated higher predictive power and lower false discovery rate (FDR) respectively of the BSS solution compared to LASSO. Guo, Zhu and Fan (2020) and Zhu and Wu (2021) reported that the approximate BSS solutions provided by IHT have much fewer false discoveries than LASSO, SCAD and SIS, especially in the presence of highly correlated design. They also theoretically showed that the model selection behavior of BSS does not explicitly depend on the restricted eigenvalue condition for the design (Bickel, Ritov and Tsybakov, 2009, Van De Geer and Bühlmann, 2009), a condition which appears unavoidable (assuming a standard computational complexity conjecture) for any polynomial-time method (Zhang, Wainwright and Jordan, 2014). This suggests that BSS is robust against design collinearity in terms of model selection.

In this paper, we mainly focus on the *precise* asymptotic bound, i.e., the bound with the optimal constant for the minimum signal strength that allows BSS to achieve model consistency. Under a specific asymptotic setup, we show that BSS achieves asymptotic exact recovery of the true model once the minimum signal strength parameter is above the information-theoretic lower bound, meaning that BSS is optimal in terms of the requirement on the signal strength. In contrast, previous works such as Aeron, Saligrama and Zhao (2010), Rad (2011), Wainwright (2009b) analyze BSS from a sample complexity point of view: they show that BSS can achieve model consistency under the optimal rate of the sample complexity and different asymptotic regimes. Later Ndaoud and Tsybakov (2020) showed the existence of a polynomial-time method that achieves model consistency under the same sufficient condition on n as BSS for i.i.d. Gaussian design. For general Gaussian design, Wainwright (2009b) showed a similar result for BSS. However, the analyses of all these works are non-asymptotic and thus may yield a suboptimal constant in the minimum signal strength. Next, we summarize the main contributions of our paper below:

- 1. Under suitable asymptotic scaling of the sample size and signal sparsity, we present a sharp and novel asymptotic analysis of BSS that leverages results from Fan, Shao and Zhou (2018) and achieves the optimal constant in minimum signal strength required for model consistency of BSS (Theorem 4.1).
- 2. To establish the sharpness of Theorem 4.1, we first present a novel impossibility result of BSS (Theorem 4.3) at the information-theoretic boundary which also relies on asymptotic analysis of BSS. This result together with the impossibility result in Wang, Wainwright and Ramchandran (2010), establishes the sharp optimality of the constant obtained in Theorem 4.1. However, it is important to point out that the result in the Theorem 4.1 is only orderwise optimal with respect to the sparsity.
- 3. To achieve computational expediency, we propose a computationally tractable two-stage algorithmic framework that also enjoys model consistency under the same condition on the minimum signal strength as BSS (Theorem 5.2).

The rest of the paper is organized as follows. Section 2 introduces the Asymptotically Ultra-Rare and Weak Minimum signal (AURWM) regime that accommodates heterogeneous signal strengths. In Section 3, we point out a potential adverse effect of signal heterogeneity on model selection under

AURWM regime. For demonstration, we present an illustrative example of marginal screening to substantiate our claim. In Section 4, we derive an impossibility result for BSS under the AURWM regime and show that BSS is optimal in terms of the requirement on the minimum signal strength. In Section 5, we propose the aforementioned two-stage algorithmic framework and establish its optimality property in terms of the constant in minimum signal strength. Finally, in Section 6, we carry out simulation studies and numerically demonstrate the superiority of our method over other competing methods.

Notation. Let \mathbb{R} and \mathbb{R}_+ denote the set of real numbers and the set of non-negative real numbers respectively. Denote by \mathbb{R}^p the *p*-dimensional Euclidean space and by $\mathbb{R}^{p\times q}$ the space of real matrices of order $p \times q$. For a positive integer *K*, denote by [K] the set $\{1, 2, ..., K\}$.

Regarding vectors and matrices, for a vector $v \in \mathbb{R}^p$, we denote by $||v||_2$ the ℓ_2 -norm of v. We use $I_p \in \mathbb{R}^{p \times p}$ to denote the *p*-dimensional identity matrix. For a matrix $A \in \mathbb{R}^{p \times p}$, we denote by A_j and a_j the *j*th column and the transposed *j*th row of A respectively.

Throughout the paper, let $O(\cdot)$ (respectively $\Omega(\cdot)$) denote the standard big-O (respectively big-Omega) notation, i.e., we say $a_n = O(b_n)$ (respectively $a_n = \Omega(b_n)$) if there exists a universal constant C > 0, such that $a_n \leq Cb_n$ (respectively $a_n \geq Cb_n$) for all $n \in \mathbb{N}$. Sometimes for notational convenience, we write $a_n \leq b_n$ in place of $a_n = O(b_n)$ and $a_n \geq b_n$ in place of $a_n = \Omega(b_n)$. We write $a_n \approx b_n$ if $a_n = O(b_n)$ and $a_n = \Omega(b_n)$. We denote by $\Omega_{\mathbb{P}}$ the big-Omega in probability: for a sequence of random variables $\{Z_n\}_{n\geq 1}$ and a sequence of constants $\{a_n\}_{n\geq 1}$, $X_n = \Omega_{\mathbb{P}}(a_n)$ means that for any $\varepsilon_0 > 0$, there exist $C_{\varepsilon_0} > 0$ and $n_{\varepsilon_0} \in \mathbb{N}$, both of which depend on ε_0 , such that

$$\mathbb{P}(|X_n/a_n| < C_{\varepsilon_0}) \le \varepsilon_0, \quad \forall n \ge n_{\varepsilon_0}.$$

We use \xrightarrow{p} and \xrightarrow{d} to denote convergence in probability and distribution respectively. Also we say $X \stackrel{d}{=} Y$ for two random variables X, Y if their distributions are equal. We denote by $\mathbb{1}(\cdot)$ the indicator function.

Finally, regarding probabilistic distributions, we use N(0, 1) to denote the standard Gaussian distribution. We use N_p(0, Σ) to denote the *p*-dimensional Gaussian distribution with mean zero and variance-covariance matrix $\Sigma \in \mathbb{R}^{p \times p}$. We denote by Ber(π) the Bernoulli distribution with success probability $\pi \in [0, 1]$.

2. Ultra rare and weak minimum signal regime

In this section, we focus on a specific asymptotic setup that allows *heterogeneity* among the sparse signals in high dimension. Throughout our paper, we consider the following signal class:

$$\mathcal{M}_s^a := \{ \beta \in \mathbb{R}^p : ||\beta||_0 = s, \min_{j \in \mathcal{S}(\beta)} ||\beta_j| \ge a \}.$$

Here *a* denotes the minimum signal strength of β . Note that the signal class \mathcal{M}_s^a only imposes a lower bound for the minimum signal strength and thus allows arbitrarily large magnitudes across the true signals. This implicitly accommodates heterogeneity in the signal, which is in sharp contrast with the homogeneous signal setup considered by Genovese et al. (2012).

Now we are in a position to introduce the *Asymptotically Ultra Rare and Weak Minimum signal* regime (AURWM), in which we mainly consider the signal class above with

$$a = \left(\frac{2r\log p}{n}\right)^{1/2} \quad \text{and} \quad s = O(\log p), \tag{2}$$

where the parameter r controls the magnitude of the minimum signal strength. As we will see in Section 4, the model consistency of BSS will depend on the value of r. Besides, we set the sample size n as

$$n = \lfloor p^k \rfloor, \quad 0 < k < 1.$$

The condition $s \leq \log p$ characterizes the ultra-rarity of the signals, which is common in genetic studies such as GWAS (Yang et al., 2020). Unless stated otherwise, from now on our statistical analysis follows the scalings of n, p, s, a in this AURWM regime. We say a support estimator \hat{S} achieves *asymptotic consistent recovery* in the AURWM regime if

$$\lim_{p \to \infty} \sup_{\beta \in \mathcal{M}_{\delta}^{\alpha}} \mathbb{P}_{\beta}(\widehat{S} \neq \mathcal{S}(\beta)) = 0.$$
(3)

This paper mainly focuses on the criterion (3) to measure the quality of exact recovery performance for an estimator \hat{S} .

It is also worth mentioning that a relevant but different asymptotic setup is studied by Genovese et al. (2012) and Ji et al. (2012). There the authors assumed a Bayesian model such that all the signals are independent and identically distributed and that the sparsity $s \sim p^{1-\vartheta}$ for some $\vartheta \in (0, 1)$. Under such a setup they obtained asymptotically tight phase transition boundaries with respect to Bayesian Hamming risk, which partitions the $r-\vartheta$ plane into three regions: (a) Region of exact recovery, (b) Region of almost recovery, (c) Region of no recovery. We skip the details of these results for brevity. The major differences between their setup and ours are twofold: (1) They essentially assume homogeneous signals; (2) They assume s to grow in a polynomial fashion with respect to p.

3. Marginal screening under heterogeneous signal

Marginal screening (MS) is one of the most widely used variable selection methods in practice. It selects the variables with top absolute marginal correlation with the response. Formally, for any $j \in [p]$, write $\mu_j := X_j^{\top} Y/n$. Given any possibly data-driven threshold $\tau(X, Y)$, define the marginal screening estimator as follows:

$$\widehat{\mathcal{S}}_{\tau} := \{ j \in [p] : |\mu_j| \ge \tau(X, Y) \}.$$

$$\tag{4}$$

Note that μ_j is essentially equivalent to the marginal correlation between X_j and Y because of isotropy of X. Marginal screening has been applied in various fields for feature selection and dimension reduction, including biomedicine (Huang, McKeague and Qian, 2019, Leisenring, Pepe and Longton, 1997, Lu, 2005), survival data analysis (Hong, Kang and Li, 2018, Li et al., 2016), economics and econometrics (Huang, Li and Wang, 2014, Wang et al., 2022).

Besides the broad applications, marginal screening has been shown to enjoy some desirable statistical properties. Fan and Lv (2008) established the sure screening property of marginal screening under an ultra-high dimensional setup, which serves as theoretical justification for MS to be used for dimension reduction in many applications. Later, Genovese et al. (2012) showed that MS enjoys the minimax optimal rate under Hamming loss with homogeneous signals. Nevertheless, as mentioned in Section 1, precise asymptotic characterization of the 0-1 loss of MS remains fairly underexplored under high dimension, especially in the presence of heterogeneity in signal strength.

3.1. Failure of MS in the AURWM regime

In this section, we study the 0-1 risk of the MS estimator. Define $\mathcal{T} := \{\widehat{S}_{\tau} \mid \tau : \mathbb{R}^{n \times p} \times \mathbb{R}^n \to \mathbb{R}_+\}$, which is the class of all possible marginal screening estimators. Under the AURWM regime, we show that MS fails to achieve exact model recovery in the minimax sense.

Theorem 3.1. Under the AURWM regime with $n = \lfloor p^k \rfloor$ for some $k \in (0, 1)$, none of the MS estimators of the form (4) can achieve asymptotic exact recovery, i.e.,

$$\lim_{p \to \infty} \inf_{\widehat{S}_{\tau} \in \mathcal{T} \beta \in \mathcal{M}_{s}^{a}} \mathbb{P}_{\beta}(\widehat{S}_{\tau} \neq \mathcal{S}(\beta)) = 1.$$

To understand the main message of this theorem, it is instructive to compare it with the parallel result in Genovese et al. (2012) with homogeneous signal. Specifically, Genovese et al. (2012) consider a Bayesian setup where all the signal coefficients are independent and identically distributed Bernoulli random variables (up to a universal constant). Under the AURWM regime, $s = O(\log p)$, which implies that $\vartheta = 1$ in Theorem 10 of Genovese et al. (2012). Then Theorem 10 in Genovese et al. (2012) says that when r > 1, MS enjoys consistency in terms of Hamming risk and thus 0-1 risk too. In contrast, when we broaden the signal class to \mathcal{M}_{s}^{a} that embraces possibly heterogeneous signals, the same model consistency fails to hold anymore for MS as shown in Theorem 3.1. This comparison clearly reveals the curse of signal heterogeneity on MS. However, the above impossibility result does not contradict Theorem 2 in Fletcher, Rangan and Goyal (2009). The result therein states that asymptotically r > 1 $(1 + \|\beta\|_2^2)$ is sufficient for the model consistency of MS. However, in the AURWM regime, r can be smaller than $1 + \|\beta\|_2^2$, which would violate the previous condition. In fact, the proof of the above theorem essentially relies on constructing a sequence of signal patterns that violates the condition r > $(1+||\beta||_2^2)$ asymptotically. Thus, in a way, the proof techniques of Theorem 3.1 show that $r > (1+||\beta||_2^2)$ is also necessary for MS to achieve model consistency in the AURWM regime. Hence, this establishes the sharpness of Theorem 2 of Fletcher, Rangan and Goyal (2009), at least in AURWM regime.

To see how signal heterogeneity hurts MS, for any $j \in [p]$, write μ_j as

$$\mu_j = (\beta_j/n) \left\| X_j \right\|_2^2 + X_j^\top (\sum_{\ell \neq j} X_\ell \beta_\ell + E)/n =: \mu_j^{(1)} + \mu_j^{(2)}.$$
(5)

Here $\mu_j^{(1)} = n^{-1}\beta_j ||X_j||_2^2$ represents the marginal contribution from β_j to μ_j , and $\mu_j^{(2)}$ represents the random error of μ_j due to the cross covariance between X_j and the other signals and noise. Suppose there are spiky signals among $\{\beta_\ell\}_{\ell \neq j}$. Though $\mathbb{E}(\mu_j^{(2)}) = 0$ regardless of the magnitude of β_j , the spiky signals may incur large variance of $\mu_j^{(2)}$ and overwhelm the magnitude of $\mu_j^{(1)}$, which is the essential indicator of the significance of β_j . Consequently, for weak signals, one cannot tell if β_j is a true variable based on only μ_j in the presence of spiky signals. For further understanding of this phenomenon, it is instructive to note that

$$\mu_{j} \mid X_{j} \sim \mathsf{N}\left(\frac{1}{n} \|X_{j}\|_{2}^{2} \beta_{j}, \frac{1}{n^{2}} \|X_{j}\|_{2}^{2} \left(1 + \sum_{\ell \neq j} \beta_{\ell}^{2}\right)\right).$$

If the non-zero components of β are arbitrarily large, then the variance of the above Gaussian distribution becomes extremely large. Therefore, even if $\beta_j = 0$, there is a high probability that the correlation μ_j will take larger values. Hence, there is a chance that the true variables associated with weak signals would lose to a noise variable, thereby ultimately leading to false discovery. To rigorously show these claims, we construct a specific example as mentioned before and we study the asymptotic limits of $\max_{j\notin S(\beta)} \mu_j^{(2)}$ and μ_{j_0} , where j_0 denotes the index of a weak signal. While the asymptotic analysis of μ_{j_0} is rather straightforward, we borrow some non-trivial results from Fan, Shao and Zhou (2018) to obtain the asymptotic properties of $\max_{j\notin S(\beta)} \mu_j^{(2)}$. Details of the proof can be found in Section A of the supplementary material.

In contrast, the AMP line of works on LASSO in Su, Bogdan and Candes (2017) and Wang, Yang and Su (2022) show that under a certain asymptotic regime, signal heterogeneity helps LASSO in terms of variable selection. Specifically, under i.i.d. Gaussian design and *linear sparsity* regime (i.e. $s/p \rightarrow \alpha$ for some constant $\alpha \in (0, 1)$, Wang, Yang and Su (2022) show that higher signal heterogeneity delays the inclusion of false variables in the LASSO solution path whereas, under signal homogeneity, false discovery occurs in a much earlier stage in the solution path. The effect is somewhat opposite to what we discussed for MS. The reason perhaps lies in the fact that LASSO tries to select the features that are highly correlated with the shrinkage noise (see Su, Bogdan and Candes (2017)), whereas, MS tries to select the features that have a higher correlation with the response. In the case of LASSO, higher signal heterogeneity makes the magnitudes of the correlations between features and shrinkage noise more distinguishable compared to a homogeneous signal pattern, thus preventing early false discovery in the first case. In the case of MS, higher signal heterogeneity increases the variance of $\mu_i^{(2)}$, which essentially dwarfs the influence of weak signals and leads to false discovery. However, these two phenomena are not directly comparable as the asymptotic settings are different for the two cases. In fact, when $s = O(\log p)$, the effect of shrinkage noise is much smaller (see Section 3.2 in Su, Bogdan and Candes (2017)) and such phenomenon does not occur for LASSO.

4. Best subset selection

Now we shift our focus to BSS, one of the most classical variable selection approaches. With the oracle knowledge of true sparsity *s*, BSS solves for

$$\hat{\beta}_{\text{best}} \in \arg\min_{\beta \in \mathbb{R}^{P}, \|\beta\|_{0} = s} n^{-1} \|Y - X\beta\|_{2}^{2}$$

Define $P_{\mathcal{D}} := X_{\mathcal{D}} (X_{\mathcal{D}}^{\top} X_{\mathcal{D}})^{-1} X_{\mathcal{D}}^{\top}$, which is the orthogonal projection operator onto the column space of $X_{\mathcal{D}}$. The BSS above can be alternatively viewed as solving for

$$\widehat{\mathcal{S}}_{\text{best}} := \arg\min_{\mathcal{D}\subseteq [p]:|\mathcal{D}|=s} n^{-1} Y^{\top} (\mathsf{I}_n - P_{\mathcal{D}}) Y = \arg\max_{\mathcal{D}\subseteq [p]:|\mathcal{D}|=s} n^{-1} Y^{\top} P_{\mathcal{D}} Y.$$
(6)

Using a union bound as in Wainwright (2009b) or Guo, Zhu and Fan (2020), one can show that there exists a universal positive constant φ (approximately equal to 0.618) such that whenever $r > 4/(1 - \varphi)$, BSS achieves model consistency, i.e.,

$$\lim_{p \to \infty} \sup_{\beta \in \mathcal{M}_{c}^{a}} \mathbb{P}_{\beta}(\mathcal{S}_{\text{best}} \neq \mathcal{S}(\beta)) = 0.$$

We emphasize that the requirement on r here is more stringent than needed: we will show that BSS achieves model consistency whenever r > 1. Moreover, a direct application of Theorem 1 of Wang, Wainwright and Ramchandran (2010) tells us that $r \ge 1$ is necessary for model consistency in AURWM regime. Therefore, our result (Theorem 4.1) essentially resolves the gap between the existing necessary $(r \ge 1)$ and sufficient $(r > 4/(1 - \varphi))$ conditions on the parameter r by showing that r > 1 is sufficient.

4.1. Exact support recovery of BSS

In the following theorem, we show that r > 1 is sufficient for BSS to achieve asymptotic exact recovery. Recall that $n = \lfloor p^k \rfloor$ with 0 < k < 1.

Theorem 4.1. Let r > 1 and write $\delta = r - 1$. Then there exists a universal positive constant C_0 such that whenever

$$s < C_0 \min\left\{2k, \frac{\delta^2}{\{(1+0.75\delta)^{1/2} + (1+0.5\delta)^{1/2}\}^2}\right\} \log p,$$

we have

$$\lim_{p \to \infty} \sup_{\beta \in \mathcal{M}_s^{\alpha}} \mathbb{P}_{\beta}(\widehat{\mathcal{S}}_{\text{best}} \neq \mathcal{S}(\beta)) = 0.$$

In order for BSS to achieve model consistency, we need to ensure that the maximum spurious correlation, i.e., the correlation between the spurious variables and the response, is well controlled so that the best subset does not involve any false discovery. One important ingredient of our analysis is the asymptotic distribution of the maximum spurious correlation due to Fan, Shao and Zhou (2018), based on which we can derive the sharp constant in the minimum signal strength for BSS to be modelconsistent. It is worth emphasizing that pursuing the exact asymptotic distributions is crucial to obtain constant-sharp results; typically, standard non-asymptotic analysis can only yield optimal rates rather than optimal constants. Detailed proof can be found in Section B.1 of the supplementary material.

Note also that Theorem 4.1 requires *s* to grow slowly, i.e., we need $s = O(\log p)$. Given that we have at least $(p - s)!/\{s!(p - 2s)!\}$ spurious models and that this number increases with respect to *s* when *s* is small, a larger *s* implies a higher maximum spurious correlation due to randomness and thus a thinner chance for the best subset to remain the true model.

Remark 4.2. The result of Theorem 4.1 can be extended to the sub-Gaussian case. In particular, if the coordinates of x follow i.i.d. distribution with *mean-zero* and *unite variance*, and ϵ is also distributed as a *mean-zero* sub-Gaussian distribution with *unit variance* and independently from x, then BSS is model consistent under a similar condition on sparsity s. Details can be found in Section B.2 of the supplementary material.

In the next section, we show that r > 1 is the weakest possible requirement on the minimum signal strength for BSS to achieve asymptotic model consistency.

4.2. Necessary condition for model consistency of BSS

In this section, we establish that under the AURWM regime, it is impossible for BSS to achieve model consistency if $r \le 1$, i.e., r > 1 is necessary for BSS to exactly recover the true support of β . To begin with, Theorem 1 of Fletcher, Rangan and Goyal (2009) in our setting yields that whenever r < 1, i.e., $r < 1 - \delta_0$ for some $\delta_0 \in (0, 1)$, the 0-1 loss of BSS approaches to 1 as p grows to infinity. This immediately shows that if r < 1, BSS is not model consistent. However, to the best of our knowledge, there is no existing research that has explored the compatibility of the condition r = 1 with the model consistency of BSS. In the following theorem, we answer this question negatively.

Theorem 4.3. Under AURWM regime (2) with r = 1 and $n = \lfloor p^k \rfloor$ for some $k \in (0, 1)$, BSS is unable to achieve model consistency, i.e.,

$$\lim_{p \to \infty} \sup_{\beta \in \mathcal{M}_{s}^{n}} \mathbb{P}(\widehat{\mathcal{S}}_{\text{best}} \neq \mathcal{S}(\beta)) > \frac{1}{10}.$$

The above theorem together with Theorem 1 of Fletcher, Rangan and Goyal (2009) shows that r > 1 is necessary for model consistency of BSS which is an improvement over the prior existing necessary condition $r \ge 1$. Therefore, These results along with Theorem 4.1, provide a *complete characterization* of model consistency of BSS in terms of the magnitude of r and resolve the existing theoretical gap at r = 1.

It is worth mentioning that a more general information-theoretic impossibility result is true for the regime r < 1. In other words, if r < 1, all methods (including BSS) fail to achieve model consistency. This claim is a direct consequence of Theorem 1 of Wang, Wainwright and Ramchandran (2010) that states that the minimax 0-1 loss is bounded away from 0 as p diverges to infinity, i.e.,

$$\lim_{p \to \infty} \inf_{\widehat{S}} \sup_{\beta \in \mathcal{M}_{s}^{a}} \mathbb{P}_{\beta}(\widehat{S} \neq \mathcal{S}(\beta)) \geq c,$$

where c > 0 is a universal constant and the infimum is taken over the class of all possible measurable functions $\hat{S} : (X, Y) \to \{\mathcal{D} \subseteq [p] : |\mathcal{D}| = s\}$. The above inequality suggests that $r \ge 1$ is a necessary condition for exact support recovery. Combining this with Theorem 4.1 and Theorem 4.3, we can see that BSS is *almost* optimal in terms of the requirement on the constant r in minimum signal strength to achieve model consistency. More discussion on this can be found in Section B.4 of the supplementary material.

Next, we briefly discuss the intuition behind the Theorem 4.3. In the regime r = 1, the main difficulty for BSS arises from the fact that it gets confused between $S(\beta)$ and its closest competitors. To be precise, let j_0 denote the index of a weak signal, i.e., $\beta_{j_0} = \{(2r \log p)/n\}^{1/2}$ with r = 1. Due to the weak magnitude of β_{j_0} , it becomes indistinguishable from 0 and as a result, BSS confuses $S(\beta)$ with other candidate models $\{\mathcal{D} \subset [p] : S(\beta) \setminus \mathcal{D} = \{j_0\}, |\mathcal{D}| = s\}$ of size *s* that differ only at j_0 with non-negligible probability. To prove Theorem 4.3, we also analyze the asymptotic distribution of an appropriate maximum spurious correlation statistics using results from Fan, Shao and Zhou (2018). We point the readers to Section B.3 of the supplementary material for further details of the proof.

Comparison with previous literature: As pointed out before in Section 1, there is a sharp contrast between the above results and the results in the previous works like Aeron, Saligrama and Zhao (2010), Rad (2011), Wainwright (2009b), where the authors study the necessary and sufficient conditions for model consistency of BSS in terms of sample complexity under different asymptotic regimes. For example, under *strong-noise* regime, Aeron, Saligrama and Zhao (2010) showed that the necessary and sufficient conditions for model consistency in terms of 0-1 loss are given by $n = \Omega(s \log(p/s))$ and $a^2 = \Omega(\log(p - s))$, and BSS is optimal in the sense that it achieves exact recovery under these conditions. For the *fixed noise-variance* regime, the results are different. Firstly, Wang, Wainwright and Ramchandran (2010) showed that the following condition is necessary for any method to achieve exact recovery:

$$n = \Omega\left(\frac{s\log(p/s)}{\log(1+sa^2)} \lor \frac{\log(p-s)}{\log(1+a^2)}\right),\tag{7}$$

where $u \lor v := \max\{u, v\}$. Under the restriction that a = O(1) and $a = \Omega(1/\sqrt{s})$, which represents *strong-signal* regime, Rad (2011) showed that BSS achieves model consistency under the necessary

condition (7). In the general case, that is with no assumption on the joint behavior of (n, p, s, a), Wainwright (2009b) established that $n = \Omega(\max\{s \log(p/s), a^{-2}\log(p-s)\})$ is a sufficient condition for model consistency of BSS. One can check that the previous condition matches with the condition in (7) under the weak signal regime $a = O(1/\sqrt{s})$. This indicates that BSS is also optimal in this regime in terms of sample complexity. It is interesting to note that the AURWM regime (2) also falls under this regime as $a = O(\sqrt{(\log p)/n}) \ll 1/\sqrt{s}$, and $n \asymp p^k$. However, all of these results fail to capture the precise dependence on *a* in terms of sharp requirement on the constant *r*.

5. Achieving information-theoretic optimality with computational efficiency

Despite the optimality of BSS in terms of model selection, its NP-hardness seriously restricts its practical applicability. To address the computational issue, we propose a two-stage algorithm framework called ETS (Estimate then Screen) that combines an estimation step with a follow-up coordinate screening step. Under this framework, one has the flexibility to use any sensible algorithm in the first stage that outputs an estimate with a good estimation guarantee for β . For example, one choice could be the well-known *iterative hard thresholding* (IHT) algorithm (Blumensath and Davies, 2009) which is a computational surrogate for BSS and enjoys a desirable estimation guarantee (Jain, Tewari and Kar, 2014). Other choices may include algorithms like *pathwise calibrated sparse shooting algorithm* (PICASSO) or *prox-gradient homotopy* (PGH) method that are known to produce good approximate solutions for LASSO (see Xiao and Zhang (2013), Zhao, Liu and Zhang (2018)) in high-dimensional setup. We show that in the AURWM regime, ETS enjoys the same selection optimality as BSS in terms of the requirement on the minimum signal strength, i.e., ETS asymptotically achieves model consistency whenever *r* is greater than the information-theoretic threshold 1, which is also the the optimal requirement for BSS to achieve exact recovery.

The above framework is similar to the methodology introduced in Ndaoud and Tsybakov (2020). In that paper, the authors used the square-root SLOPE estimator (Bogdan et al., 2015) for the estimation step, and under i.i.d. Gaussian design they showed that their algorithm achieves model consistency under the same *sample complexity* as BSS. However, they do not study optimal dependence on r, which is the main focus of our paper.

5.1. The ETS algorithm

In this section, we introduce our ETS algorithm (Algorithm 1) in detail. Given a partition parameter $0 < \gamma < 1$, ETS first splits the full sample $(x_i, Y_i)_{i \in [n]}$ into two subsamples $\mathcal{D}_1, \mathcal{D}_2$ of respective sizes $n_1 = \lfloor \gamma n \rfloor$ and $n_2 = n - n_1$. Then ETS performs two main steps on these two sub-samples respectively:

1. Given an objective function $f_{n_1}(\cdot; \mathcal{D}_1)$ and a constraint set $C \subseteq \mathbb{R}^p$, in the estimation step, ETS procures a close approximation to β by solving for an approximate solution to the optimization problem

$$\operatorname{minimize}_{\theta \in C} f_{n_1}(\theta; \mathcal{D}_1) \tag{8}$$

via a suitable iterative algorithm $\mathcal{A}(\cdot, \cdot)$ that takes the objective function $f_{n_1}(\cdot; \mathcal{D}_1)$ and a set of tuning parameters $\mathcal{T}_{\mathcal{A}}$ as inputs. In particular, in this step, ETS outputs an estimator $\hat{\beta} := \mathcal{A}(f_{n_1}(\cdot; \mathcal{D}_1), \mathcal{T}_{\mathcal{A}})$ of the true signal vector β .

2. In the second step, ETS performs a coordinatewise screening based on \mathcal{D}_2 and $\hat{\beta}$ to select the true variables.

Algorithm 1: ETS

Input: Data $\mathcal{D} = \{(x_i, Y_i)\}_{i=1}^n$, objective function $f_{n_1}(\cdot; \mathcal{D}_1)$, partition parameter γ , threshold parameter ς ; 1. Randomly partition the whole dataset \mathcal{D} into two disjoint subsets $\mathcal{D}_1 = (X^{(1)}, Y^{(1)})$ and $\mathcal{D}_2 = (X^{(2)}, Y^{(2)})$; 2.Apply the algorithm \mathcal{A} to compute an approximate solution $\hat{\beta}$ of the optimization problem (8);

3. Construct the statistics $\{\Delta_i\}_{i=1}^p$ and thresholds $\{\kappa_{\varsigma}(X_i^{(2)})\}_{i=1}^p$ using (9)-(10);

4. Finally compute the selector $\hat{\eta}(X, Y)$;

Output: The selector $\hat{\eta}(X.Y)$.

To elaborate more on the method, for $\ell \in \{1,2\}$, let $X^{(\ell)} \in \mathbb{R}^{n_\ell \times p}$ and $Y^{(\ell)} \in \mathbb{R}^{n_\ell}$ denote the design matrix and the response vector of the ℓ th sub-sample respectively. ETS computes $\hat{\beta}$ based on the first sub-sample $\mathcal{D}_1 := (X^{(1)}, Y^{(1)})$ by finding an approximate solution the optimization problem (8) via algorithm \mathcal{A} . In practice, there could be several choices for both the objective function $f_{n_1}(\cdot; \mathcal{D}_1)$ and the algorithm \mathcal{A} . For example, one of the most common choices is to consider the ℓ_0 -constrained squared-error loss, i.e., $f_{n_1}(\theta; \mathcal{D}_1) = n_1^{-1} ||Y^{(1)} - X^{(1)}\theta||_2^2$ with $C = \{\theta \in \mathbb{R}^p : ||\theta||_0 \leq s\}$. In this case, a natural choice for \mathcal{A} is the IHT algorithm which is basically a projected gradient descent method. Another popular choice for the objective function is the well-known ℓ_1 -regularized LASSO objective function $f_{n_1}(\theta; \lambda, \mathcal{D}_1) = n_1^{-1} ||Y^{(1)} - X^{(1)}\theta||_2^2 + \lambda ||\theta||_1$ with $C = \mathbb{R}^p$ and one can choose either PICASSO, PGH or the composite gradient method proposed in Agarwal, Negahban and Wainwright (2012) as the algorithm \mathcal{A} . Besides these, another choice could be to solve the square-root LASSO problem (Bogdan et al., 2015) via proximal-gradient descent algorithm proposed in Li et al. (2020).

Next comes the screening step of ETS. For each $i \in [p]$, define

$$\Delta_{i} := \frac{X_{i}^{(2)\top} \left(Y^{(2)} - \sum_{j \neq i} X_{j}^{(2)} \hat{\beta}_{j}\right)}{\|X_{i}^{(2)}\|_{2}}$$
(9)

and

$$\kappa_{\varsigma}(u) := \frac{a \, \|u\|_2}{2} + \frac{\varsigma^2 \log p}{a \, \|u\|_2}, \quad \forall u \in \mathbb{R}^{n_2}, \tag{10}$$

where $\varsigma > 0$ is specified later. ETS selects the *i*th variable if and only if $|\Delta_i| > \kappa_{\varsigma}(X_i^{(2)})$. To see why we can screen variables based on $\{\Delta_i\}_{i \in [p]}$, note that

$$\Delta_{i} = \beta_{i} \|X_{i}^{(2)}\|_{2} + \frac{X_{i}^{(2)\top} \left(\sum_{j \neq i} X_{j}^{(2)} (\beta_{j} - \hat{\beta}_{j}) + E\right)}{\|X_{i}^{(2)}\|_{2}}.$$
(11)

A straightforward argument shows that conditioned on \mathcal{D}_1 and $X_i^{(2)}$, Δ_i is distributed as:

$$\Delta_i \left| \left(\mathcal{D}_1, X_i^{(2)} \right) \stackrel{\mathrm{d}}{=} \beta_i \left\| X_i^{(2)} \right\|_2 + \left\{ 1 + \sum_{j \neq i} (\beta_j - \hat{\beta}_j)^2 \right\}^{1/2} g_i,$$

where $g_i \sim N(0, 1)$ and is independent of $X_i^{(2)}$. If estimation method performs well in the sense that $\|\hat{\beta} - \beta\|_2$ is small, then for all $i \in S(\beta)$, $\beta_i \|X_i^{(2)}\|_2$ becomes the dominant term in Δ_i . In contrast, for all $i \notin S(\beta)$, $\beta_i \|X_i^{(2)}\|_2 = 0$ and we thus expect Δ_i to be small. This suggests the existence of a threshold $t(\cdot)$ on $(\Delta_i)_{i \in [p]}$ that distinguishes the true support $S(\beta)$ from the irrelevant variables. We

follow Ndaoud and Tsybakov (2020) to choose the threshold function in (10), which is shown to be a reasonable choice to identify the true variables.

For each $i \in [p]$, define $\hat{\eta}_i(X, Y) := \mathbb{1}\{|\Delta_i| > \kappa_{\varsigma}(X_i^{(2)})\}$ and write

$$\hat{\eta}(X,Y) := (\hat{\eta}_1(X,Y),\ldots,\hat{\eta}_p(X,Y))^\top$$

The selector $\hat{\eta}(X, Y)$ is the final estimate of the support $S(\beta)$ produced by the ETS algorithm. Algorithm 1 shows the detailed steps of the ETS algorithm.

5.1.1. Model consistency of ETS

In this section, we establish theoretical guarantees for ETS-IHT. First, we introduce a technical assumption that concerns how fast algorithm \mathcal{A} can generate a good approximation of the true signal β .

Assumption 5.1. The following holds with a probability converging to 1 as p diverges to infinity: For any given tolerance level $\epsilon > 0$, there exists a suitable set of deterministic tuning parameters $\mathcal{T}_{\mathcal{R}}$ such that the algorithm \mathcal{R} requires no more than $T(\epsilon, p, \beta)$ iterations to produce a solution $\hat{\beta} := \mathcal{R}(f_{n_1}(\cdot; \mathcal{D}_1), \mathcal{T}_{\mathcal{R}})$ such that $\|\hat{\beta} - \beta\|_2^2 \le \epsilon$.

The above assumption essentially tells that the $T(\epsilon, p, \beta)$ th iterate of algorithm \mathcal{A} is already ϵ -close to β in squared ℓ_2 -distance with high probability for large enough p. If ϵ is small, then $\hat{\beta}$ is a good estimate of β and we can use it in the screening step to select the variables. Typically, as ϵ decreases towards 0, the iteration counts $T(\epsilon, p, \beta)$ increases to infinity as higher accuracy generally demands more computation. However, in many examples, as we will see in Section 5.1.2, $T(\epsilon, p, \beta)$ depends only poly-logarithmically on ϵ^{-1} , which alleviates the computational cost.

Next, we define the binary decoder of the true support $S(\beta)$ as $\eta_{\beta} := (\mathbb{1}\{\beta_1 \neq 0\}, \dots, \mathbb{1}\{\beta_p \neq 0\})^{\top}$. The following theorem shows that ETS can achieve exact recovery under suitable choices of tuning parameters.

Theorem 5.2. Assume the condition in Assumption 5.1 hold and the sample size $n = \lfloor p^k \rfloor$ for some $k \in (0, 1)$. Let r > 1 and write $\delta = r - 1$. Then, under AURWM regime (2), there exist universal positive constants A_1, A_2 such that with overall iteration count no more than $T(A_1\delta, p, \beta)$ for algorithm \mathcal{A} , $\gamma \in (0, \delta/(8+8\delta))$ and $\varsigma = (1+A_2\delta)^{1/2}$, we have that $\lim_{p\to\infty} \sup_{\beta \in \mathcal{M}_{\infty}^{s}} \mathbb{P}_{\beta}(\hat{\eta} \neq \eta_{\beta}) = 0$.

Note that as the signal strength parameter r approaches the information-theoretic boundary, i.e., as δ approaches 0, ETS may require more iterations to achieve model consistency as $T(A_1\delta, p, \beta)$ generally increases as δ decreases to 0. This is not surprising: intuitively, weaker signals are harder to identify than strong ones.

Besides, ETS does not require the knowledge of the true sparsity *s*, but requires the knowledge of *a* in the second stage for accurate screening. If the true sparsity *s* is known, then we can enforce ETS to select exactly *s* features as follows: Let $|\Delta|_{(m)}$ denote the *m*th largest value of $\{|\Delta_i|\}_{i \in [p]}$. For each $i \in [p]$, define

$$\hat{\eta}_i(X,Y;s) = \mathbb{1}\{|\Delta_i| \ge |\Delta|_{(s)}\}.$$
(12)

Hence $\hat{\eta}(s) := (\hat{\eta}_1(s), \dots, \hat{\eta}_p(s))^\top$ selects exactly *s* features and the knowledge of *a* is not required in this case. The following corollary shows that under the same conditions of Theorem 5.2, $\hat{\eta}(s)$ achieves model consistency.

Corollary 5.3. Assume the condition in Assumption 5.1 holds and the sample size $n = \lfloor p^k \rfloor$ for some $k \in (0, 1)$. Let A_1 be the same universal constant as in Theorem 5.2, r > 1 and write $\delta = r - 1$. Then, under AURWM regime (2), with overall iteration count no more than $T(A_1\delta, p, \beta)$ for algorithm \mathcal{A} and $\gamma \in (0, \delta/(8 + 8\delta))$, we have that $\lim_{p \to \infty} \sup_{\beta \in \mathcal{M}_{\mathcal{B}}} \mathbb{P}_{\beta}(\hat{\eta}(s) \neq \eta_{\beta}) = 0$.

Remark 5.4. The algorithm can be made adaptive to *a* in some certain regime of *r*. In particular, if there exists a known positive constant δ_* such that $r > 1 + \delta_*$, then a threshold as in (10) can be constructed without the knowledge of *a* or *r* so that ETS still enjoys model consistency. In this case, δ_* can be arbitrarily small, and as long as δ_* is known, an adaptive choice of threshold exists.

Detailed proofs of Theorem 5.2, Corollary 5.3 and Remark 5.4 can be found in Section C of the supplementary materials. Next, we will discuss some concrete examples of ETS methods that enjoy model consistency.

5.1.2. Examples of ETS methods

In this section, we will present a few examples of ETS methods. In particular, we will consider the ETS methods with different choices for the base algorithm \mathcal{A} : (1) the IHT algorithm which solves the ℓ_0 -constrained optimization problem, (2) the PICASSO and PGH algorithm which solves the ℓ_1 -regularized optimization problem. We will show that Assumption 5.1 is met in all these cases and we will explicitly derive the dependence of $T(\epsilon, p, \beta)$ on (ϵ, p, β) . Hence, this will automatically establish the model consistency of these three variants of the ETS method due to the result in Theorem 5.2. For clarity, depending on the algorithm used in the estimation step, we will refer to these methods as ETS-IHT, ETS-PICASSO, and ETS-PGH. To be self-contained, we describe the steps of IHT in Algorithm 2. However, we do not add the description of PICASSO and PGH as those are too involved to add in this paper. Detailed descriptions of PICASSO and PGH can be found in Zhao, Liu and Zhang (2018) and Xiao and Zhang (2013) respectively. We remind the readers that throughout the discussion in this section, we will consider the AURWM regime defined in (2) with sample size $n = \lfloor p^k \rfloor$ for some $k \in (0, 1)$. More details and proofs related to the examples can be found in Section C.4 of the supplementary material.

Solving ℓ_0 -constrained problem: As discussed in Section 5.1, In this case, the optimization problem (8) takes the form

We consider the ETS-IHT in this case which uses IHT (Algorithm 2) to obtain an approximate solution to the above optimization problem. In this case $\mathcal{T}_{\mathcal{R}} = \{\hat{s}, h\}$, where \hat{s} is the sparsity level and h is the gradient step-size. Following the discussion of Section 4 in Jain, Tewari and Kar (2014), in particular, using Theorem 3 of that paper we have that the final output $\hat{\beta}$ of IHT satisfies $\|\hat{\beta} - \beta\|_2^2 \le \epsilon$ with probability converging to 1, when $\hat{s} = 2592s$, $h \le 8/27$ and $T(\epsilon, p, \beta) = O(\log p + \log((1 + \|\beta\|_{\infty})/\epsilon))$. Hence, the conditions in Assumption 5.1 hold. Moreover, Theorem 3 of Jain, Tewari and Kar (2014) suggests that if $f_{n_1}(\hat{\beta}; \mathcal{D}_1) - \min_{\theta:\|\theta\|_0 \le s} f_{n_1}(\theta; \mathcal{D}_1) \le (\epsilon/16)$, then for large values of p, we have $\|\hat{\beta} - \beta\|_2^2 \le \epsilon$. Hence, it is enough to output an estimator $\hat{\beta}$ which incurs a sub-optimality gap of the order $O(\epsilon)$.

Solving ℓ_1 -regularized problem: In this case, the objective function is

$$f_{n_1}(\theta; \lambda, \mathcal{D}_1) = n_1^{-1} \left\| Y^{(1)} - X^{(1)} \theta \right\|_2^2 + \lambda \|\theta\|_1,$$

14

AI2011000 4. 1111

Input: Objective function f, sparsity level \hat{s} , step size h; $\beta^{(0)} = 0$; t = 0; while not converged **do** $\begin{vmatrix} \beta^{(t+1)} = P_{\hat{s}}^{0}(\beta^{(t)} - h\nabla_{\theta}f(\beta^{t})), \text{ where } P_{\hat{s}}^{0}(v) = \arg\min_{z: ||z||_{0} = \hat{s}} ||v - z||_{2};$ $t \leftarrow t + 1$ end Output: $\hat{\beta} = \beta^{(t)}$.

and $C = \mathbb{R}^p$. For brevity of discussion, we only consider PICASSO and PGH as the candidate methods for solving the above optimization problem. We omit the details of the tuning parameters for these algorithms in this paper, but details of those can be found in Zhao, Liu and Zhang (2018) and Xiao and Zhang (2013) respectively.

• (ETS-PICASSO): For ETS-PICASSO, Theorem 3.12 of Zhao, Liu and Zhang (2018) yields that with $T(\epsilon, p, \beta) \leq (\log p)^3 (\log p + \log \|\beta\|_{\infty}) (\log \log p + \log(\epsilon^{-1} \vee C_1))^2$ and the regularization parameter $\lambda = C_2 \{(\log p)/n_1\}^{1/2}$ for appropriate absolute constants $C_1, C_2 > 0$, the approximate solution $\hat{\beta}$ has the property $f_{n_1}(\hat{\beta}; \lambda, \mathcal{D}_1) - f_{n_1}(\hat{\beta}_L; \lambda, \mathcal{D}_1) = O(\epsilon)$ with probability converging to 1, where

$$\hat{\beta}_L := \arg\min_{\theta} f_{n_1}(\theta; \lambda, \mathcal{D}_1).$$

Then, the strong convexity property of the Gram matrix $X^{(1)\top}X^{(1)}/n_1$, and the good estimation property of $\hat{\beta}_L$ yields that $\|\hat{\beta} - \beta\|_2^2 \le \epsilon$.

• (ETS-PGH): For ETS-PGH, Theorem 3.2 of Xiao and Zhang (2013) yields that with $T(\epsilon, p, \beta) = O((\log p + \log \|\beta\|_{\infty}) \log \log p + \log(\epsilon^{-1} \vee \tilde{C}_1))$ and $\lambda = \tilde{C}_2\{(\log p)/n_1\}^{1/2}$ for appropriate absolute constants $\tilde{C}_1, \tilde{C}_2 > 0$, the approximate solution $\hat{\beta}$ satisfies $f_{n_1}(\hat{\beta}; \lambda, \mathcal{D}_1) - f_{n_1}(\hat{\beta}_L; \lambda, \mathcal{D}_1) = O(\epsilon)$ with probability converging to 1. Then, again by the strong convexity property of the Gram matrix $X^{(1)\top}X^{(1)}/n_1$, and the good estimation property of $\hat{\beta}_L$, it follows that $\|\hat{\beta} - \beta\|_2^2 \le \epsilon$.

It is worth mentioning that the choice of PICASSO or PGH is not special for solving the ℓ_1 -regularized problem. As long as the base algorithm \mathcal{A} outputs $\hat{\beta}$ which enjoys a sub-optimality gap of the order $O(\epsilon)$ in the functional value, it follows that $\|\hat{\beta} - \beta\|_2^2 \le \epsilon$. We formalize this result in the next proposition.

Proposition 5.5. Consider the AURWM regime in (2) and let the sample size $n = \lfloor p^k \rfloor$ for some $k \in (0, 1)$. Then, there exists a positive universal constant C_3 such that for all $\epsilon \in (0, C_3)$, the following holds with probability at least $1 - 3p^{-0.5}$ for large enough p and $\lambda = 8\{(\log p)/n_1\}^{1/2}$:

$$f_{n_1}(\hat{\beta};\lambda,\mathcal{D}_1) - f_{n_1}(\hat{\beta}_L;\lambda,\mathcal{D}_1) \le \epsilon/C_3 \quad implies \quad \|\hat{\beta} - \beta\|_2^2 \le \epsilon.$$

The proof of the above result is deferred to Section C.4 of the supplementary material. The above proposition basically shows that in the case of solving the LASSO problem, if $\hat{\beta}$ can be produced efficiently, then an optimality-gap of $C_3^{-1}\epsilon$ in the functional value is enough to guarantee that $\hat{\beta}$ falls inside the $\epsilon^{1/2}$ neighborhood of the true parameter β , i.e., the conditions in Assumption 5.1 hold with probability at least $1 - O(p^{-0.5})$. Hence, this provides us the flexibility to use any sensible algorithm for the ℓ_1 -regularized problem such as the composite gradient method proposed in Agarwal, Negahban and Wainwright (2012).

5.2. Discussion on information-theoretic optimality, statistical accuracy and computational efficiency

In the previous section, we have shown that ETS achieves the model consistency under the same information-theoretic optimal requirement on r with computational expediency. However, the computational efficiency of ETS heavily depends on the magnitude of r. Theorem 5.2 suggests that as r approaches the information-theoretic boundary 1, the demand on the number of iterations in the estimation step increases. It could be possible that the computational load of ETS surpasses the computational load of BSS when r is extremely close to 1 for a fixed ambient dimension. Hence, even though ETS is able to recover the weak signals *asymptotically* (as n approaches infinity) under the optimal requirement on r, it may suffer from high computational costs. Moreover, as r approaches 1, it turns out that the decaying rate of the error probability worsens. This fact can be verified from the rates obtained in the proof of Theorem 5.2 in the supplementary material and we do not include those in the main theorem for conciseness. This suggests that weak signals also hurt the statistical power or accuracy of the ETS methods. Hence, both statistical accuracy and computational efficiency suffer as the signals get weaker.

6. Numerical experiments

In this section, we first numerically investigate¹ the probability for MS to achieve exact recovery of the true model with growing ambient dimension p under both homogeneous and heterogeneous signal setups. Our results show that while MS exhibits model consistency under the homogeneous signal regime, it completely fails to do so under the heterogeneous signal regime, which is consistent with Theorem 3.1. We then conduct simulation experiments to demonstrate the superiority of ETS methods over competing methods including LASSO and MS as signal strength grows or signal heterogeneity grows. For ETS methods, we only include ETS-IHT and ETS-PICASSO methods. For ETS-PICASSO we used picasso package in R which uses the PICASSO method for solving the LASSO problem. To this end, we mention that we do not numerically compare *exact* BSS in this section mainly due to computational issues. In most of our simulation setups, we consider p in thousands and exact BSS suffers from high computational costs in such regimes, which is also a limitation of the commercial solver Gurobi (Hastie, Tibshirani and Tibshirani, 2020). Bertsimas and Parys (2020), Hazimeh, Mazumder and Saab (2022), Xie and Deng (2020) have made efforts to overcome this computational bottleneck by considering different methods for solving approximate versions of BSS. In particular, they all consider different regularized versions of BSS which is beyond the scope of this paper, and hence we do not include those in the numerical experiments. Instead, we focus on LASSO and ETS, both of which are two different computational surrogates of the BSS problem.

6.1. Exact recovery performance of MS

In Figure 1, we demonstrate the asymptotics of MS under both homogeneous and heterogeneous signal patterns. We consider $p \in \{1000, 2000, \dots, 8000\}$ and signal strength parameter $r \in \{2, 3, 4, 5, 6\}$. We set $s = \lfloor 2 \log p \rfloor$ and $n = \lfloor p^{0.9} \rfloor$. We let $\tau(X, Y)$ in (4) be equal to the *s*th largest value of $\{|\mu_1|, \dots, |\mu_p|\}$, so that MS always chooses a model of size *s*. For the homogeneous signal setup, we consider β with $\|\beta\|_0 = s$ and $\beta_j = a$ for all $j \in S(\beta)$, where *a* is defined in (2). This implies that

¹Codes are available online: Github link

the SNR varies between 0.19 and 2.15 across different choices of (r, p). For the heterogeneous signal setup, we consider β with (s - 1) active coordinates equal to a and one "spiky" coordinate equal to $\{10 - (s - 1)a^2\}^{1/2}$. This ensures that the SNR is fixed at 10 for all choices of r, p.

Figure 1(a) shows that under homogeneous signal MS is able to recover the exact model with probability converging to 1 as p grows. In contrast, Figure 1(b) shows that under heterogeneous signal MS never achieves exact model recovery: plots for all values of r are at level 0. Such a contrast corroborates Theorem 3.1: signal spikes can give rise to substantial spurious correlation and jeopardize the accuracy of MS.



(a) Proportion of exact recovery under weak homogeneous signal. $(0.19 \le \text{SNR} \le 2.15)$

(b) Proportion of exact recovery under heterogeneous signal. (SNR = 10)

Figure 1: Asymptotics of MS with growing dimension *p*.

6.2. Effect of growing signal strength

Here we numerically compare the probability of exact support recovery of ETS with those of LASSO and MS as signal strength parameter *r* grows. We investigate both homogeneous and heterogeneous signal patterns. We set p = 2000, $s \in \{13, 52\}$ and $n = \lfloor p^{0.9} \rfloor = 935$. We set signal strength parameter $r \in \{1.5, 2, 2.5, \ldots, 9\}$ in (2). The support *S* is chosen uniformly over all the size-*s* subsets of [p], and each support coordinate of β is chosen as follows:

$$\beta_j = (1 - b_j)(1 + Z_j^2/n)^{1/2}a + b_j r^{1/2} \quad \forall j \in \mathcal{S},$$

where $(Z_j)_{j \in S} \stackrel{i.i.d.}{\sim} N(0,1)$, and where $(b_j)_{j \in S} \stackrel{i.i.d.}{\sim} Ber(\pi)$ with $\pi \in \{0, 0.2\}$. $\pi = 0$ corresponds to the homogeneous signal pattern, and $\pi = 0.2$ corresponds to the heterogeneous signal pattern, where spiky signals are present with probability 0.2. Each entry x_{ij} of the design matrix X is generated independently from N(0, 1).

In this experiment, we grant all the approaches with the knowledge of *s*, so that the comparison is fair. Using this oracle knowledge, we only look at the solutions of the aforementioned three methods with sparsity exactly equal to *s*. Specifically, for LASSO, we look at the solution path and select the model of size exactly equal to *s*. For MS, we just select the top *s* variables corresponding to the largest absolute values of μ 's. For ETS, we do not split data for estimation and screening separately; instead we use the full data in both steps. Specifically, we replace $X^{(1)}$ and $Y^{(1)}$ with X and Y respectively in



(a) Homogeneous signal pattern ($\pi = 0, s = 13$)



(b) Heterogeneous signal pattern ($\pi = 0.2, s = 13$)



(c) Homogeneous signal pattern ($\pi = 0, s = 52$)



Figure 2: Plot of the proportion of exact recovery for varying r. (Gaussian case)

(13) and replace $X^{(2)}$ and $Y^{(2)}$ with X and Y respectively in (9). We set gradient step size h = 0.5 in IHT. We choose projection size \hat{s} by cross-validation in terms of mean squared prediction error. Lastly, for selecting exactly *s* features we use (12) in the screening stage of ETS. It is worthwhile to mention that from an application point of view, incorporating data splitting in ETS is not necessary as we are only interested in identifying the active signals, which is akin to point estimation. Also, given the fact that $n \ll p$ in a high dimensional regime, using the full sample in both the estimation and screening step delivers greater sample efficiency and provides better inference.

Next, for each choice of r, we run LASSO, MS, and ETS over 200 independent Monte Carlo experiments to compute the empirical probability of exact recovery. Figure 2 presents the results. We make the following important observations:

- 1. All three methods enjoy a higher chance of exact support recovery as the signal strength grows.
- 2. MS completely fails to achieve exact support recovery when *s* becomes large (compare panels (a) and (c)) or the signal becomes heterogeneous (compare panels (a) and (b)).
- 3. LASSO and ETS algorithms are insensitive to the heterogeneity of the signal. However, LASSO suffers from larger sparsity, while ETS algorithms are much more robust against it.
- 4. Overall, ETS is the best among all the three methods in terms of exact support recovery. However, ETS-IHT is somewhat better than ETS-PICASSO, and the difference between their performance is more prominent when *s* is large.

Remark 6.1. Genovese et al. (2012) established model consistency of LASSO under rare and weak signal regime for i.i.d. Gaussian design under a different asymptotic setting and the scaling $s = O(\log p)$ is a special case of their setting. However, they assume homogeneous signal, and our simulation results in Figure 2(a) concur with their theoretical findings. However, the performance of LASSO degrades significantly for larger sparsity, even under homogeneous signal (see Figure 2(b)), which perhaps shows the limitations of the results in Genovese et al. (2012) for realistic values of p and s. In contrast, Wainwright (2009a) obtains the sharpest possible results for model consistency of LASSO under a very general setting. To be precise, under i.i.d. Gaussian design and for general combination of (n, p, s), the paper shows that LASSO achieves model consistency for $a = \Omega(\lambda)$, where λ is the regularization parameter and $\lambda \gtrsim {(\log p)/n}^{1/2}$. Hence, it is also valid for rare and weak signal regimes and also accommodates heterogeneity in the signal when $\lambda \approx {(\log p)/n}^{1/2}$. However, those are tight only up to multiplicative constants. Moreover, as pointed out in Section B of Wainwright (2009a), the model consistency of LASSO depends on whether or not the following is achieved:

$$n > s \log(p - s) + \lambda^{-2} \log(p - s).$$

The above condition is harder to satisfy if *s* becomes large keeping other parameters fixed, which could be a possible explanation for the phenomenon observed in the third point of the prior observations.



Figure 3: Plot of the proportion of exact recovery for varying r. (non-Gaussian case)

To demonstrate the effectiveness of the ETS framework beyond Gaussian design, we also evaluate the performance of ETS-IHT and ETS-PICASSO under the case where the entries of X are generated from Unif $(-\sqrt{3}, \sqrt{3})$ in an i.i.d. fashion with the same choices of (n, p, s) as in the Gaussian example. Based on Figure 3, in this case, ETS also enjoys superior performance than LASSO and MS. Moreover, we observe a similar effect of signal heterogeneity and sparsity on the model recovery rate of the methods. This provides empirical evidence in favor of the good model selection performance of ETS under sub-Gaussian design.

6.3. Effect of growing heterogeneity

In this numerical experiment, we study the effect of growing heterogeneity on ETS-IHT, LASSO and MS. We set $p = 2000, s = 13, n = \lfloor p^{0.9} \rfloor = 935$ and $r \in \{2, 6\}$ in (2). Next, we introduce n_{spike} , the number of "spiky" signals in S. We vary n_{spike} in $\{0\} \cup [6]$. The case $n_{spike} = 0$ corresponds to the homogeneous signal setup where the true signals are set as *a* uniformly. For $n_{spike} > 0$, we randomly set $(s - n_{spike})$ signals in S to be equal to *a* and the remaining signals to be equal to a_{spike} , which is defined as

$$a_{\rm spike} := \left\{ \frac{(2 - sa^2)}{\mathsf{n}_{\rm spike}} + a^2 \right\}^{1/2}.$$

Such a choice of a_{spike} ensures that the SNR always equals 2 whenever $n_{spike} > 0$. We perform ETS-IHT, LASSO, and MS over 200 Monte Carlo simulations for each choice of r and n_{spike} to obtain the empirical probability of exact support recovery. Similarly to the previous sections, we assume that the true sparsity s is known and we apply the three methods in the same fashion as before.

Figure 4 shows again the detrimental effect of heterogeneity on MS in terms of exact recovery. In both panels, we see a significant drop in the proportion of exact recovery for MS when n_{spike} changes from 0 to 1. This is consistent with the theory in Section 3. However, in Figure 4(b) we see that the proportion of exact recovery is slowly increasing as n_{spike} grows from 1 to 6. This is because as n_{spike} increases, a_{spike} monotonically decreases, so that the signals become more homogeneous. MS is then able to recover the exact model more frequently. We do not see a similar phenomenon in Figure 4(a) because a_{spike} is too large. Another important observation is that while ETS-IHT and LASSO are both performing nearly perfectly when r = 6, ETS-IHT significantly outperforms both LASSO and MS when r = 2. Therefore, ETS-IHT is again the overall winner.

7. Conclusion

In this paper, we study exact support recovery in high-dimensional sparse linear regression with independent Gaussian design. We focus on the AURWM regime that not only accommodates *rare* and *weak* signals as the ARW regime but also allows *heterogeneity* in the signal strength. Our first theoretical result (Theorem 3.1) shows that marginal screening fails to achieve exact support recovery under the AURWM regime. The main reason is that the presence of "spiky" signals increases the maximum spurious marginal correlation, thereby blinding the marginal screening for variable selection in practice.

In contrast, we show that BSS is robust to signal spikes and is able to achieve model consistency under the AURWM regime with the optimal requirement on signal strength (Theorem 4.1, 4.3). The primary reason behind this is that, unlike MS, BSS takes into account multiple features simultaneously



Figure 4: Plot of proportion of exact recovery with varying n_{spike}.

and thus selects variables based on their capability of fitting the residualized responses given the other variables rather than the responses themselves. Therefore, spiky signals do not affect BSS: They are very likely to be in plausible candidate models in the first place and their effect on the response has been removed in the residualization procedure. Given the recent computational advancements in solving BSS, our positive result on BSS makes it more appealing from an application point of view.

However, it is worth mentioning that even with modern advances in optimization, BSS suffers from high computational costs when the ambient dimension is extremely high. To address this issue, we propose a computationally tractable two-stage method ETS that delivers essentially the same optimal exact recovery performance as BSS (Theorem 5.2). Similar to BSS, ETS seeks for the features that exhibit high explanation power for the residuals from the model that excludes these features themselves (see (9)). Therefore, ETS is robust to spiky signals. This fact together with the slowly growing sparsity condition in (2) yields the optimal exact recovery accuracy of ETS.

Our work naturally raises several important questions for future research. One question is whether similar optimality results hold for BSS and ETS when the sparsity *s* grows faster than log *p*. The same question can also be asked for *correlated* random design. Another direction of our interest is studying the problem of exact recovery in a *distributed* setting where data are stored at different places and communication between them is restricted.

Acknowledgements

The authors would like to thank the anonymous referees and the Associate Editor for their constructive comments that improved the quality of this paper.

Supplementary Material

Supplement to "High-dimensional variable selection with heterogeneous signals: A precise asymptotic perspective"

The supplementary material (Roy, Tewari and Zhu, 2024) contains detailed proofs of all the main theorems, corollaries, and Remark 5.4.

References

- AERON, S., SALIGRAMA, V. and ZHAO, M. (2010). Information theoretic bounds for compressed sensing. *IEEE Trans. Inform. Theory* 56 5111–5130. https://doi.org/10.1109/TIT.2010.2059891
- AGARWAL, A., NEGAHBAN, S. and WAINWRIGHT, M. J. (2012). Fast global convergence of gradient methods for high-dimensional statistical recovery. *Ann. Statist.* **40** 2452–2482. https://doi.org/10.1214/12-AOS1032
- AKÇAKAYA, M. and TAROKH, V. (2009). Shannon-theoretic limits on noisy compressive sampling. *IEEE Trans. Inform. Theory* 56 492–504. https://doi.org/10.1109/TIT.2009.2034796
- ARIAS-CASTRO, E., CANDÈS, E. J. and PLAN, Y. (2011). Global testing under sparse alternatives: ANOVA, multiple comparisons and the higher criticism. Ann. Statist. 39 2533–2556. https://doi.org/10.1214/11-AOS910
- BERTSIMAS, D., KING, A. and MAZUMDER, R. (2016). Best subset selection via a modern optimization lens. Ann. Statist. 44 813–852. https://doi.org/10.1214/15-AOS1388
- BERTSIMAS, D. and PARYS, B. V. (2020). Sparse high-dimensional regression: exact scalable algorithms and phase transitions. Ann. Statist. 48 300–323. https://doi.org/10.1214/18-AOS1804
- BICKEL, P. J., RITOV, Y. and TSYBAKOV, A. B. (2009). Simultaneous analysis of lasso and Dantzig selector. Ann. Statist. 37 1705–1732. https://doi.org/10.1214/08-AOS620
- BLUMENSATH, T. and DAVIES, M. E. (2009). Iterative hard thresholding for compressed sensing. Appl. Comput. Harmon. Anal. 27 265–274. https://doi.org/10.1137/100806278
- BOGDAN, M., VAN DEN BERG, E., SABATTI, C., SU, W. and CANDÈS, E. J. (2015). SLOPE—adaptive variable selection via convex optimization. Ann. Appl. Stat. 9 1103–1140. https://doi.org/10.1214/15-AOAS842
- CANDÈS, E. J., ROMBERG, J. and TAO, T. (2006). Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *IEEE Trans. Inform. Theory* 52 489–509. https://doi.org/10.1109/TIT.2005.862083
- CANDES, E. J. and TAO, T. (2006). Near-optimal signal recovery from random projections: universal encoding strategies? *IEEE Trans. Inform. Theory* 52 5406–5425. https://doi.org/10.1109/TIT.2006.885507
- CONSORTIUM, W. T. C. C. et al. (2007). Genome-wide association study of 14,000 cases of seven common diseases and 3,000 shared controls. *Nature* 447 661–678. https://doi.org/10.1038/nature05911
- DONOHO, D. L. (2006). Compressed sensing. *IEEE Trans. Inform. Theory* **52** 1289–1306. https://doi.org/10.1109/TIT.2006.871582
- FAN, J. and LI, R. (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. J. Amer. Statist. Assoc. 96 1348–1360. https://doi.org/10.1198/016214501753382273
- FAN, J. and LV, J. (2008). Sure independence screening for ultrahigh dimensional feature space. J. R. Statist. Soc. Ser. B 70 849–911. https://doi.org/10.1111/j.1467-9868.2008.00674.x
- FAN, J., SHAO, Q.-M. and ZHOU, W.-X. (2018). Are discoveries spurious? Distributions of maximum spurious correlations and their applications. Ann. Statist. 46 989–1017. https://doi.org/10.1214/17-AOS1575
- FLETCHER, A. K., RANGAN, S. and GOYAL, V. K. (2009). Necessary and sufficient conditions for sparsity pattern recovery. *IEEE Trans. Inform. Theory* 55 5758–5772. https://doi.org/10.1109/TIT.2009.2032726
- GENOVESE, C. R., JIN, J., WASSERMAN, L. and YAO, Z. (2012). A comparison of the lasso and marginal regression. J. Mach. Learn. Res. 13 2107–2143. http://jmlr.org/papers/v13/genovese12b.html.
- GUO, Y., ZHU, Z. and FAN, J. (2020). Best subset selection is robust against design dependence. *arXiv preprint arXiv:2007.01478*. https://doi.org/10.48550/arXiv.2007.01478
- GUO, X., ZHANG, H., WANG, Y. and WU, J.-L. (2015). Model selection and estimation in high dimensional regression models with group SCAD. *Stat. Probab. Lett.* **103** 86–92. https://doi.org/10.1016/j.spl.2015.04.017
- HASTIE, T., TIBSHIRANI, R. and TIBSHIRANI, R. (2020). Best subset, forward stepwise or lasso? Analysis and recommendations based on extensive comparisons. *Statist. Sci.* 35 579–592. https://doi.org/10.1214/19-STS733
- HAZIMEH, H., MAZUMDER, R. and SAAB, A. (2022). Sparse regression at scale: branch-and-bound rooted in first-order optimization. *Math. Program.* 196 347–388. https://doi.org/10.1007/s10107-021-01712-4
- HONG, H. G., KANG, J. and LI, Y. (2018). Conditional screening for ultra-high dimensional covariates with survival outcomes. *Lifetime Data Anal.* 24 45–71. https://doi.org/10.1007/s10985-016-9387-7
- HUANG, D., LI, R. and WANG, H. (2014). Feature screening for ultrahigh dimensional categorical data with applications. J. Bus. Econ. Stat. 32 237–244. https://doi.org/10.1080/07350015.2013.863158
- HUANG, J., MA, S. and ZHANG, C.-H. (2008). Adaptive Lasso for sparse high-dimensional regression models. *Stat. Sin.* 18 1603–1618. https://www.jstor.org/stable/24308572.

22

- HUANG, T.-J., MCKEAGUE, I. W. and QIAN, M. (2019). Marginal screening for high-dimensional predictors of survival outcomes. *Stat. Sin.* 29 2105–2139. https://doi.org/10.5705/ss.202017.0298
- JAIN, P., TEWARI, A. and KAR, P. (2014). On Iterative Hard Thresholding Methods for High-dimensional M-Estimation. In Advances in Neural Information Processing Systems 27. https://proceedings.neurips.cc/paper/ 2014/hash/218a0aefd1d1a4be65601cc6ddc1520e-Abstract.html.
- JI, P., JIN, J. et al. (2012). UPS delivers optimal phase diagram in high-dimensional variable selection. Ann. Statist. 40 73–103. https://doi.org/10.1214/11-AOS947
- JIN, Y., KIM, Y.-H. and RAO, B. D. (2011). Limits on support recovery of sparse signals via multiple-access communication techniques. *IEEE Trans. Inform. Theory* 57 7877–7892. https://doi.org/10.1109/TIT.2011.2170116
- JIN, J., ZHANG, C.-H. and ZHANG, Q. (2014). Optimality of graphlet screening in high dimensional variable selection. J. Mach. Learn. Res. 15 2723–2772. https://www.jmlr.org/papers/v15/jin14a.html.
- KOWSHIK, S. S. and POLYANSKIY, Y. (2021). Fundamental limits of many-user MAC with finite payloads and fading. *IEEE Trans. Inform. Theory* 67 5853–5884. https://doi.org/10.1109/TIT.2021.3091423
- LEISENRING, W., PEPE, M. S. and LONGTON, G. (1997). A marginal regression modelling framework for evaluating medical diagnostic tests. *Stat. Med.* 16 1263–1281. https://doi.org/10.1002/(SICI)1097-0258(19970615)16:11<1263::AID-SIM550>3.0.CO;2-M
- LI, J., ZHENG, Q., PENG, L. and HUANG, Z. (2016). Survival impact index and ultrahigh-dimensional model-free screening with survival outcomes. *Biometrics* 72 1145–1154. https://doi.org/10.1111/biom.12499
- LI, Y., HONG, H. G., AHMED, S. E. and LI, Y. (2019). Weak signals in high-dimensional regression: Detection, estimation and prediction. Appl. Stoch. Model. Bus. Ind. 35 283–298. https://doi.org/10.1002/asmb.2340
- LI, X., JIANG, H., HAUPT, J., ARORA, R., LIU, H., HONG, M. and ZHAO, T. (2020). On fast convergence of proximal algorithms for SQRT-Lasso optimization: Don't worry about its nonsmooth loss function. In *Proceedings of The 35th Uncertainty in Artificial Intelligence Conference. Proceedings of Machine Learning Research* 115 49–59. PMLR https://proceedings.mlr.press/v115/li20a.html.
- LIU, H. and FOYGEL BARBER, R. (2020). Between hard and soft thresholding: optimal iterative thresholding algorithms. *Inform. Inference: J. IMA* **9** 899–933. https://doi.org/10.1093/imaiai/iaz027
- LU, W. (2005). Marginal regression of multivariate event times based on linear transformation models. *Lifetime Data Anal.* 11 389–404. https://doi.org/10.1007/s10985-005-2969-4
- MARTTINEN, P., GILLBERG, J., HAVULINNA, A., CORANDER, J. and KASKI, S. (2013). Genomewide association studies with high-dimensional phenotypes. *Statist. Appl. Genet. Mol. Biol.* 12 413–431. https://doi.org/10.1515/sagmb-2012-0032
- NDAOUD, M. and TSYBAKOV, A. B. (2020). Optimal variable selection and adaptive noisy compressed sensing. IEEE Trans. Inform. Theory 66 2517–2532. https://doi.org/10.1109/TIT.2020.2965738
- RAD, K. R. (2011). Nearly sharp sufficient conditions on exact sparsity pattern recovery. *IEEE Trans. Inform. Theory* 57 4672–4679. https://doi.org/10.1109/TIT.2011.2145670
- ROBINS, J. M., SCHEINES, R., SPIRTES, P. and WASSERMAN, L. (2003). Uniform consistency in causal inference. *Biometrika* 90 491–515. https://doi.org/10.1093/biomet/90.3.491
- ROY, S., TEWARI, A. and ZHU, Z. (2024). Supplement to "High-dimensional variable selection with heterogeneous signals: A precise asymptotic perspective".
- SCARLETT, J. and CEVHER, V. (2016). Limits on support recovery with probabilistic models: an informationtheoretic framework. *IEEE Trans. Inform. Theory* 63 593–620. https://doi.org/10.1109/TIT.2016.2606605
- SU, W., BOGDAN, M. and CANDES, E. (2017). False discoveries occur early on the lasso path. Ann. Statist. 2133–2150. https://www.jstor.org/stable/26362897.
- TIBSHIRANI, R. (1996). Regression shrinkage and selection via the lasso. J. R. Statist. Soc. Ser. B 58 267–288. https://doi.org/10.1111/j.2517-6161.1996.tb02080.x
- TIKHONOV, A. (1943). On the stability of inverse problems. Proc. USSR Acad. Sci. 39 195–198. https://api. semanticscholar.org/CorpusID:202866372.
- VAN DE GEER, S. A. and BÜHLMANN, P. (2009). On the conditions used to prove oracle results for the lasso. *Electron. J. Statist.* **3** 1360–1392. https://doi.org/10.1214/09-EJS506
- WAINWRIGHT, M. J. (2009a). Sharp thresholds for high-dimensional and noisy sparsity recovery using l₁-constrained quadratic programming (Lasso). *IEEE Trans. Inform. Theory* 55 2183–2202. https://doi.org/10.1109/TIT.2009.2016018

- WAINWRIGHT, M. J. (2009b). Information-theoretic limits on sparsity recovery in the high-dimensional and noisy setting. *IEEE Trans. Inform. Theory* 55 5728–5741. https://doi.org/10.1109/TIT.2009.2032816
- WANG, W., WAINWRIGHT, M. J. and RAMCHANDRAN, K. (2010). Information-theoretic limits on sparse signal recovery: dense versus sparse measurement matrices. *IEEE Trans. Inform. Theory* 56 2967–2979. https://doi.org/10.1109/TIT.2010.2046199
- WANG, H., YANG, Y. and SU, W. J. (2022). The price of competition: effect size heterogeneity matters in high dimensions. *IEEE Trans. Inform. Theory* 68 5268–5294. https://doi.org/10.1109/TIT.2022.3166720
- WANG, C. D., CHEN, Z., LIAN, Y. and CHEN, M. (2022). Asset selection based on high frequency Sharpe ratio. J. Econom. 227 168–188. Annals Issue: Time Series Analysis of Higher Moments and Distributions of Financial Data. https://doi.org/10.1016/j.jeconom.2020.05.007
- WASSERMAN, L. and ROEDER, K. (2009). High dimensional variable selection. Ann. Statist. 37 2178–2201. https://doi.org/10.1109/TIT.2022.3166720
- XIAO, L. and ZHANG, T. (2013). A proximal-gradient homotopy method for the sparse least-squares problem. SIAM J. Optim. 23 1062–1091. https://doi.org/10.1137/12086999
- XIE, W. and DENG, X. (2020). Scalable algorithms for the sparse ridge regression. SIAM J. Optim. 30 3359–3386. https://doi.org/10.1137/19M1245414
- YANG, S., WEN, J., ECKERT, S. T., WANG, Y., LIU, D. J., WU, R., LI, R. and ZHAN, X. (2020). Prioritizing genetic variants in GWAS with lasso using permutation-assisted tuning. *Bioinformatics* 36 3811–3817. https://doi.org/10.1093/bioinformatics/btaa229
- ZHANG, C.-H. and HUANG, J. (2008). The sparsity and bias of the lasso selection in high-dimensional linear regression. *Ann. Statist.* **36** 1567–1594. https://doi.org/10.1214/07-AOS520
- ZHANG, Y., WAINWRIGHT, M. J. and JORDAN, M. I. (2014). Lower bounds on the performance of polynomialtime algorithms for sparse linear regression. In *Conference on Learning Theory* 921–948. PMLR https: //proceedings.mlr.press/v35/zhang14.html.
- ZHAO, T., LIU, H. and ZHANG, T. (2018). Pathwise coordinate optimization for sparse learning: algorithm and theory. Ann. Statist. 46 180–218. https://doi.org/10.1214/17-AOS1547
- ZHAO, P. and YU, B. (2006). On model selection consistency of lasso. The J. Mach. Learn. Res. 7 2541–2563.
- ZHENG, Z., FAN, Y. and LV, J. (2014). High dimensional thresholded regression and shrinkage effect. J. R. Statist. Soc. Ser. B 627–649. https://doi.org/10.1111/rssb.12037
- ZHU, Z. and WU, S. (2021). On the early solution path of best subset selection. arXiv preprint arXiv:2107.06939.
- ZHU, J., WEN, C., ZHU, J., ZHANG, H. and WANG, X. (2020). A polynomial algorithm for best-subset selection problem. Proc. Natl. Acad. Sci. U.S.A. 117 33117–33123. https://doi.org/10.1073/pnas.2014241117

Supplement to "High-dimensional variable selection with heterogeneous signals: A precise asymptotic perspective"

SAPTARSHI ROY^{1,a}, AMBUJ TEWARI^{1,b} and ZIWEI ZHU^{2,c}

¹Department of Statistics, University of Michigan, Ann Arbor, USA, Email: ^aroysapta@umich.edu, ^btewaria@umich.edu ²Radix Trading, Chicago, USA Email: ^czzw9348ustc@gmail.com

In this document, we collect the proofs of all the main theorems, propositions, and corollaries. We also discuss about Remark 5.4 and prove some important results related to the examples of ETS.

Appendix A: Proof of Theorem 3.1

Consider a MS procedure $\hat{S}_{\tau} \in \mathcal{T}$. Now, there are mainly two steps of the proof:

- 1. Upper bound the probability of recovery in terms of the probability of an event depending only on $\max_{i \in S^c} |\mu_i|$ and $\min_{i \in S} |\mu_i|$.
- 2. Find the asymptotic limits of the above random variables and find the limiting probability of the aforementioned event.

To start with note that

$$\mathbb{P}_{\beta}(\widehat{\mathcal{S}}_{\tau} = \mathcal{S}(\beta)) = \mathbb{P}_{\beta}\left(\max_{j \in \mathcal{S}^{c}} |\mu_{j}| < \tau(X, Y) \le \min_{j \in \mathcal{S}} |\mu_{j}|\right) \le \mathbb{P}\left(\max_{j \in \mathcal{S}^{c}} |\mu_{j}| < \min_{j \in \mathcal{S}} |\mu_{j}|\right).$$

Recall that for for $j \in [p]$ we have

$$\mu_j = \begin{cases} \beta_j \|X_j\|_2^2 / n + \omega_j \|X_j\|_2 g_j / n & \text{if } j \in \mathcal{S}(\beta) \\ \omega_j \|X_j\|_2 g_j / n & \text{if } j \notin \mathcal{S}(\beta), \end{cases}$$

where $\omega_j^2 = 1 + \sum_{k \neq j} \beta_k^2$ and $g_j = X_j^\top (\sum_{k \neq j} X_k \beta_k + E) / (\omega_j ||X_j||_2) \sim \mathsf{N}(0, 1)$. Thus we have

$$\mathbb{P}_{\beta}(\widehat{S}_{\tau} = S) \le \mathbb{P}_{\beta}\left(\max_{j \in S^{c}} \omega_{j} \left\|X_{j}\right\|_{2} |g_{j}| / (2n\log p)^{1/2} < \min_{j \in S} |\beta_{j}| \left\|X_{j}\right\|_{2}^{2} + \omega_{j} \left\|X_{j}\right\|_{2} |g_{j}| / (2n\log p)^{1/2}\right).$$
(1)

The right-hand side of Equation (1) does not depend on \hat{S}_{τ} hence the above inequality is valid uniformly over the class \mathcal{T} . Now choose a sequence $\{c_p\}_{p=1}^{\infty}$ such that $\lim_{p\to\infty} c_p^2/r \ge 1$. Next construct a sequence of $\beta^{(p)}$ in the following manner:

- Consider the set $S_0 = \{1, \dots, s\} \subseteq [p]$ with $s = O(\log p)$.
- Set $\beta_1^{(p)} = c_p$. For all other $i \in \mathcal{S}_0 \setminus \{1\}$ set $\beta_i^{(p)} = a = (2r(\log p)/n)^{1/2}$.

• Set $\beta_i^{(p)} = 0$ if $i \notin S_0$.

In this setup we have $\omega_j \sim (1 + c_p^2)^{1/2}$ for all $j \neq 1$. Now fix $k_0 \in S_0 \setminus \{1\}$ (say $k_0 = 2$). From Equation (1) it can be concluded that

$$\sup_{\widehat{S}_{\tau} \in \mathcal{T}} \mathbb{P}_{\beta^{(p)}}\left(\widehat{S}_{\tau} = S_{0}\right)$$

$$\leq \mathbb{P}_{\beta^{(p)}}\left(\max_{j \in S^{c}} \omega_{j} \left\|X_{j}\right\|_{2} |g_{j}| / (2n \log p)^{1/2} < |\beta_{k_{0}}| \left\|X_{k_{0}}\right\|_{2}^{2} + \omega_{k_{0}} \left\|X_{k_{0}}\right\|_{2} g_{k_{0}}| / (2n \log p)^{1/2}\right).$$

Also note that $\omega_j > (1 + c_p^2)^{1/2}$ for all $j \in S_0^c$. Using these facts, we get

$$\frac{1}{(1+c_p^2)^{1/2}} \max_{j \in \mathcal{S}_0^c} \omega_j \frac{\|X_j\|_2 |g_j|}{(2n\log p)^{1/2}} \ge \min_{j \in \mathcal{S}_0^c} \frac{\|X_j\|_2}{n^{1/2}} \max_{j \in \mathcal{S}_0^c} \frac{|g_j|}{(2\log p)^{1/2}}$$

Now, note that for $j \in S_0^c$, we have all $\omega_j^2 = 1 + \|\beta\|_2^2$ and

$$g_j = \frac{X_j^\top z}{\left\|X_j\right\|_2}$$

where $z = (\sum_{k \in S_0} X_k \beta_k + E)/\omega_j \sim N_n(0, I_n)$ and it is independent of $\{X_j\}_{j \in S_0^c}$. Thus, $g_j^2 = ||P_j z||_2^2$, where P_j is the orthogonal projection operator onto the subspace span $\{X_j\}$. This shows that the random quantity $\max_{j \in S_0^c} g_j^2$ is the scaled version of maximum spurious correlation (defined in Section 7.2 of Fan, Shao and Zhou (2018)) between $\{X_j\}_{j \in S_0^c}$ and the noise z with sparsity level 1, i.e., $\max_{j \in S_0^c} g_j^2 =$ $n \widehat{R}_n^2(1, p - s)$, where

$$\widehat{R}_{n}(1, p-s) := \sup_{\alpha: \|\alpha\|_{2}=1, \|\alpha\|_{0}=1} \frac{1}{n} \sum_{i=1}^{n} \frac{\alpha^{\top}(z_{i}x_{i}, S_{0}^{c})}{(\alpha^{\top}\widehat{\Sigma}_{n, S_{0}^{c}}\alpha)^{1/2}},$$

with $\widehat{\Sigma}_{n,S_0^c} = n^{-1} \sum_{i=1}^n x_{i,S_0^c} x_{i,S_0^c}^{\top}$. Thus, following the arguments of Fan, Shao and Zhou (2018), in particular, using Theorem 3.1 and Remark 3.3 of Fan, Shao and Zhou (2018) we have,

$$\max_{j \in \mathcal{S}_0^c} \frac{|g_j|}{(2\log p)^{1/2}} \xrightarrow{p} 1.$$

Using the above fact and Lemma 3 from Fletcher, Rangan and Goyal (2009), we get

$$\frac{1}{(1+c_p^2)^{1/2}} \max_{j \in \mathcal{S}_0^c} \omega_j \frac{\|X_j\|_2 |g_j|}{(2n\log p)^{1/2}} \ge \min_{j \in \mathcal{S}_0^c} \frac{\|X_j\|_2}{n^{1/2}} \max_{j \in \mathcal{S}_0^c} \frac{|g_j|}{(2\log p)^{1/2}} \xrightarrow{p} 1$$

Now, recall that $\beta_{k_0} = \{2r(\log p)/n\}^{1/2}$ and and define $u_r := \frac{1}{2}(1 + \frac{r^{1/2}}{\sqrt{1+r}}) < 1$. Thus we have the following:

$$\mathbb{P}(|\beta_{k_0} \| X_{k_0} \|_2^2 + \omega_{k_0} \| X_{k_0} \|_2 g_{k_0} | / \{ (1 + c_p^2) (2n \log p) \}^{1/2} < u_r) \to 1.$$

2

Variable selection with heterogeneous signals

This tells that,

$$\lim_{p \to \infty} \sup_{\widehat{S}_{\tau} \in \mathcal{T}} \inf_{\beta \in \mathcal{M}_{s}^{a}} \mathbb{P}_{\beta}(\widehat{S}_{\tau} = \mathcal{S}(\beta)) \leq \lim_{p \to \infty} \sup_{\widehat{S}_{\tau} \in \mathcal{T}} \mathbb{P}_{\beta^{(p)}}(\widehat{S}_{\tau} = \mathcal{S}_{0}) = 0.$$

This finishes the proof.

Appendix B: Best subset selection

In this section, we prove the main results related to BSS.

B.1. Proof of Theorem 4.1

In this proof, we reparametrize δ by $8\delta_0$ for algebraic convenience. The main result can be salvaged by back substituting $8\delta_0$ by δ in all the main equations in this section. Also, for brevity of notation, in this proof we use \hat{S} and S to denote that oracle BSS estimator and $S(\beta)$ respectively. We highlight the three main steps of the proof:

- 1. Convert the BSS problem in the problem of selecting the model with maximum spurious correlation.
- 2. Use results from Fan, Shao and Zhou (2018) to find the asymptotic distribution of the maximum spurious correlation statistics.
- 3. Use the asymptotic distribution along with non-asymptotic concentration inequalities to upper bound the error probability.

Recall that BSS is defined as

$$\widehat{\mathcal{S}} = \arg \max_{\mathcal{D}, |\mathcal{D}|=s} \|P_{\mathcal{D}}Y\|_2^2 = \arg \min_{\mathcal{D}: |\mathcal{D}|=s} \|(I_n - P_{\mathcal{D}})Y\|_2^2.$$

Thus from the above definition we have the following equality:

$$\mathbb{P}(\widehat{\mathcal{S}} \neq \mathcal{S}) = \mathbb{P}\left(\|P_{\mathcal{S}}Y\|_{2}^{2} < \max_{\mathcal{D} \neq \mathcal{S}} \|P_{\mathcal{D}}Y\|_{2}^{2} \right).$$

Now we will try to understand how the quantity $||P_SY||_2^2$ behaves asymptotically. First it is easy to see that $P_SY = \sum_{j \in S} X_j\beta_j + P_SE$. Note that $\sum_{j \in S} X_j\beta_j = ||\beta||_2 \tilde{\varepsilon}$ where $\tilde{\varepsilon} \sim N_n(0, I_n)$ and independent of the noise *z*. Hence we up with the following:

$$\begin{split} \|P_{\mathcal{S}}Y\|_{2}^{2} &= \left\|\sum_{j \in S} X_{j}\beta_{j}\right\|_{2}^{2} + 2E^{\top}P_{\mathcal{S}}\left(\sum_{j \in S} X_{j}\beta_{j}\right) + E^{\top}P_{\mathcal{S}}E \\ &= \left\|\sum_{j \in S} X_{j}\beta_{j}\right\|_{2}^{2} + 2E^{\top}\left(\sum_{j \in S} X_{j}\beta_{j}\right) + E^{\top}P_{\mathcal{S}}E \\ &= \|\beta\|_{2}^{2} \|\tilde{\varepsilon}\|_{2}^{2} + 2\|\beta\|_{2} E^{\top}\tilde{\varepsilon} + E^{\top}P_{\mathcal{S}}E. \end{split}$$

Recall that $|\beta_j| \ge \{2r(\log p)/n\}^{1/2}$ for all $j \in S$. This is presumably the hardest setup as increasing signal strength can only decrease the error probability. Then $\|\beta\|_2^2 \ge (2rs \log p)/n$. also note that

 $E^{\top} P_{S} E \sim \chi_{s}^{2}$. Hence we have ,

$$\frac{\|P_{\mathcal{S}}Y\|_2^2}{s\log p} \ge 2r\frac{\|\tilde{\varepsilon}\|_2^2}{n} + 2\left(\frac{2r}{s\log p}\right)^{1/2}\frac{\tilde{\varepsilon}^\top E}{n^{1/2}} + \frac{E^\top P_{\mathcal{S}}E}{s\log p} \xrightarrow{\mathbf{p}} 2r.$$

Thus $\lim_{p\to\infty} \mathbb{P}(\widehat{S} \neq S) \leq \mathbb{P}(2r \leq \limsup_{p\to\infty} \max_{\mathcal{D}\neq S} ||P_{\mathcal{D}}Y||_2^2/(s \log p))$. The limiting behaviour of the obtained maximal process turns out to be very challenging to analyze and hence we do not directly study this maximal process. Instead we focus on a related maximal process (will be defined shortly) derived from the earlier one and we use the results from Fan, Shao and Zhou (2018) to study its asymptotic behaviour. Now let us denote the set $S \cap \mathcal{D}$ by I_0 and $\mathcal{D} \setminus S$ by I_1 , i.e., $\mathcal{D} = I_0 \cup I_1$. Next define the class $\mathcal{J}_{I_0} = \{I_1 \subseteq [p] : I_1 \cap S = \emptyset, |I_1 \cup I_0| = s\}$ for each $I_0 \subset S$. Note that $0 \leq |I_0| \leq s - 1$ from the construction (if $|I_0| = s$ then $\mathcal{D} = S$). The random variable of interest can be rewritten as follows:

$$\max_{\mathcal{D}\neq \mathcal{S}} \frac{\|P_{\mathcal{D}}Y\|_2^2}{s\log p} = \max_{I_0:I_0\subset \mathcal{S}} \max_{I_1:I_1\in \mathcal{J}_{I_0}} \frac{\|P_{I_0\cup I_1}Y\|_2^2}{s\log p}.$$

Using union bound we get,

$$\mathbb{P}(\widehat{\mathcal{S}} \neq \mathcal{S}) \leq \sum_{I_0 \subset \mathcal{S}} \mathbb{P}\left(\max_{I_1: I_1 \in \mathcal{J}_{I_0}} \frac{\left\|P_{I_0 \cup I_1}Y\right\|_2^2}{s \log p} > \frac{\left\|P_{\mathcal{S}}Y\right\|_2^2}{s \log p}\right).$$
(2)

Now fix a subset I_0 of the true support S. Similar to previous section define $\tilde{Y} = Y - X_{I_0}\beta_{I_0}$ and this independent of the features in $I_0 \cup I_1$. Also we have

$$\begin{aligned} \left\| P_{I_0 \cup I_1} Y \right\|_2^2 &= \left\| P_{I_0 \cup I_1} \tilde{Y} \right\|_2^2 + \left\| X_{I_0} \beta_{I_0} \right\|_2^2 + 2\beta_{I_0}^\top X_{I_0}^\top \tilde{Y} \\ \\ \left\| P_{\mathcal{S}} Y \right\|_2^2 &= \left\| P_{\mathcal{S}} \tilde{Y} \right\|_2^2 + \left\| X_{I_0} \beta_{I_0} \right\|_2^2 + 2\beta_{I_0}^\top X_{I_0}^\top \tilde{Y}. \end{aligned}$$

Thus the summands in the right hand side of (2) can be written as the probability of the event $\{\max_{I_1:I_1\in \mathcal{J}_{I_0}} \|P_{I_0\cup I_1}g\|_2^2/(s\log p) > \|P_{\mathcal{S}\mathcal{S}}\|_2^2/(s\log p)\}$, where $g := (1 + \|\beta_{\mathcal{S}\setminus I_0}\|_2^2)^{-1/2}\tilde{Y}$. Note that $g \sim N_n(0, I_n)$ and is independent of the features in \mathcal{D} . Now fix a specific I_0 . In the analysis we encounter the maximal process

$$\max_{I_1:I_1 \in \mathcal{J}_{I_0}} \frac{\left\| P_{I_0 \cup I_1} g \right\|_2^2}{s \log p}$$

Now consider the set of indices $F_{I_0} = (\{1, \dots, p\} \setminus S) \cup I_0$. Hence it is easy to see that $\tilde{p} := |F_{I_0}| = p - s + |I_0|$. Without loss of generality, let $F_{I_0} = \{1, \dots, \tilde{p}\}$. Also define $\tilde{s} := s - |I_0|$. Let the set $\mathcal{V}_{I_0} = \{\alpha \in \mathbb{R}^p : \|\alpha\|_0 = s, \|\alpha\|_2 = 1, I_0 \subseteq S(\alpha), \alpha_{F_{I_0}}^c = 0\}$. Here α_J denotes the sub-vector of α corresponding to the indices in $J \subseteq [p]$. Next we will focus on the random variable,

$$\widehat{L}_n := \widehat{L}_n(\widetilde{s}, \widetilde{p}) = \sup_{\alpha \in \mathcal{V}_{I_0}} \frac{1}{n^{1/2}} \sum_{i=1}^n \frac{\alpha^\top(g_i x_i)}{(\alpha^\top \widehat{\Sigma}_n \alpha)^{1/2}},\tag{3}$$

here $\widehat{\Sigma}_n = \frac{1}{n} \sum_{i=1}^n x_i x_i^{\top}$. Now recall that $\mathcal{D} = \mathcal{I}_0 \cup \mathcal{I}_1$ for all $\mathcal{D} \neq \mathcal{S}$ with $|\mathcal{D}| = s$. To see the connection, first note that the above optimization problem can be viewed as the following:

$$\begin{split} \widehat{L}_{n} &= \max_{I_{1} \in \mathcal{J}_{I_{0}}} \max_{\alpha \in \mathcal{V}_{I_{0} \cup I_{1}}} \frac{\alpha_{\mathcal{D}}^{+}(\sum_{i=1}^{n} g_{i}x_{i,\mathcal{D}}/n^{1/2})}{\{\alpha_{\mathcal{D}}^{+}(\widehat{\Sigma}_{n,\mathcal{D}}\mathcal{D})\alpha_{\mathcal{D}}\}^{1/2}} \\ &= \max_{I_{1} \in \mathcal{J}_{I_{0}}} \left\{ (\sum_{i=1}^{n} g_{i}x_{i,\mathcal{D}}/n^{1/2})^{\top} \widehat{\Sigma}_{n,\mathcal{D}}^{-1} \sum_{i=1}^{n} g_{i}x_{i,\mathcal{D}}/n^{1/2}) \right\}^{1/2} \\ &= \max_{I_{1} \in \mathcal{J}_{I_{0}}} \{g^{\top} X_{\mathcal{D}} (X_{\mathcal{D}}^{\top} X_{\mathcal{D}})^{-1} X_{\mathcal{D}}^{\top} g\}^{1/2} \\ &= \max_{I_{1} \in \mathcal{J}_{I_{0}}} \|P_{I_{0} \cup I_{1}}g\|_{2}. \end{split}$$

Thus it is essential to study the asymptotic property of \hat{L}_n . Now we define the standardized version of \hat{L}_n as follows

$$L_n := L_n(\tilde{s}, \tilde{p}) = \sup_{\alpha \in \mathcal{V}_{I_0}} \frac{1}{n^{1/2}} \sum_{i=1}^n \alpha^\top(g_i x_i).$$

Let $\mathbf{Z} = (Z_1, \dots, Z_{\tilde{p}})$ be \tilde{p} -variate Gaussian random variable with covariance matrix $I_{\tilde{p} \times \tilde{p}}$ and define the random variable $T^* := T^*(\tilde{s}, \tilde{p}) = \sup_{\alpha \in \mathcal{V}_{I_0}} \alpha_{F_{I_0}}^\top \mathbf{Z}$.

Lemma B.1. There exists universal constants K_0, K_1 such that for any $\delta_1 \in (0, K_0K_1]$,

$$|L_n - T^*| \lesssim n^{-1} c_n^{1/2}(\tilde{s}, \tilde{p}) + K_0 K_1 n^{-3/2} c_n^2(\tilde{s}, \tilde{p}) + \delta_1$$
(4)

holds with probability at least $1 - C\Delta_n(s, \tilde{p}; \delta_1)$ where $c_n(s, \tilde{p}) = s \log(e\tilde{p}/s) \lor \log n$ and

$$\Delta_n(\tilde{s}, \tilde{p}; \delta_1) = (K_0 K_1)^3 \frac{\{\tilde{s}b_n(\tilde{s}, \tilde{p})\}^2}{\delta_1^3 n^{1/2}} + (K_0 K_1)^4 \frac{\{\tilde{s}b_n(\tilde{s}, \tilde{p})\}^5}{\delta_1^4 n}$$

with $b_n(\tilde{s}, \tilde{p}) = \log(\tilde{p}/\tilde{s}) \vee \log n$.

Lemma B.2. Assume that the sample size satisfies $n \ge C_1(K_0 \lor K_1)^4 c_n(\tilde{s}, \tilde{p})$. then with probability at least $1 - C_2 n^{-1/2} c_n^{1/2}(\tilde{s}, \tilde{p})$,

$$\left|\widehat{L}_{n} - L_{n}\right| \lesssim (K_{0} \vee K_{1})^{2} K_{0} K_{1} n^{-1/2} c_{n}(\tilde{s}, \tilde{p}),$$
(5)

where $c_n(\tilde{s}, \tilde{p}) = \tilde{s} \log(e\tilde{p}/\tilde{s}) \vee \log n$.

Proof of the above two lemmas are omitted as it is in the same line of the proofs of Fan, Shao and Zhou (2018). Now applying Lemma B.1 and B.2 with

$$\delta_1 = \delta_1(s, \tilde{p}) = (K_0 K_1)^{3/4} \min[1, n^{-1/8} \{ \tilde{s} b_n(\tilde{s}, \tilde{p})^{3/8} \}]$$

yields that with probability at least $1 - C(K_0K_1)^{3/4}n^{-1/8}\{sb_n(s, \tilde{p})\}^{7/8}$,

$$\left|\widehat{L}_n - T^*\right| \lesssim (K_0 K_1)^{3/4} n^{-1/8} \{\widetilde{s}b_n(\widetilde{s}, \widetilde{p})\}^{3/8}.$$

Together with Lemma 2.3 from Chernozhukov, Chetverikov and Kato (2014) we can conclude that

$$\sup_{t \in \mathbb{R}} \left| \mathbb{P}(\widehat{L}_n \le t) - \mathbb{P}(T^* \le t) \right| \lesssim C(K_0 K_1)^{3/4} n^{-1/8} \{ \tilde{s} b_n(\tilde{s}, \tilde{p}) \}^{7/8}.$$
(6)

Next by the definition of T^* it follows that

$$T^{*2} = \max_{I_1 \in \mathcal{J}_{I_0}} \| \mathbf{Z}_{I_0 \cup I_1} \|_2^2 = \sum_{j \in I_0} Z_j^2 + \max_{I_1 \in \mathcal{J}_{I_0}} \sum_{k \in I_1} Z_k^2$$

Let $W \sim N_{(p-s)}(0, I_{(p-s)})$ be a Gaussian vector independent of **Z**. Thus it follows that

$$T^{*2} \stackrel{d}{=} \sum_{j \in \mathcal{I}_0} Z_j^2 + \sum_{k=p-2s+|\mathcal{I}_0|+1}^{p-s} W_{(k:p-s)}^2 \leq \sum_{j \in \mathcal{I}_0} Z_j^2 + (s-|\mathcal{I}_0|) W_{(p-s:p-s)}^2.$$

From Equation (6) it also follows that

$$\sup_{t \ge 0} \left| \mathbb{P}(\widehat{L}_n^2 \le t) - \mathbb{P}(T^{*2} \le t) \right| \lesssim C(K_0 K_1)^{3/4} n^{-1/8} \{ \tilde{s} b_n(\tilde{s}, \tilde{p}) \}^{7/8}.$$
(7)

Now from the assumption, we have $r = 1 + 8\delta_0$. Assume that

$$s \le 0.5 \min\{\delta_0, \frac{2\delta_0^2}{\{(1+6\delta_0)^{1/2} + (1+4\delta_0)^{1/2}\}^2}\}\log p.$$
(8)

Hence $|\mathcal{I}_0| \le s - 1 \le 0.5\delta_0 \log p < \delta_0 \log p$. Thus we have,

$$\begin{split} & \mathbb{P}(T^{*2} > 2(1+4\delta_0)(s-|I_0|)\log p) \\ &= \mathbb{P}(\sum_{j \in I_0} Z_j^2 + (s-|I_0|)W_{(p-s:p-s)}^2 > 2(1+4\delta_0)(s-|I_0|)\log p) \\ &\leq \mathbb{P}\left(\sum_{j \in I_0} Z_j^2 > 4\delta_0(s-|I_0|)\log p\right) + \mathbb{P}\left((s-|I_0|)W_{(p-s:p-s)}^2 > 2(1+2\delta_0)(s-|I_0|)\log p\right) \\ &\stackrel{(a)}{\leq} \mathbb{P}\left(\frac{\sum_{j \in I_0} Z_j^2 - |I_0|}{|I_0|} > (4\delta_0\log p - |I_0|)/|I_0|\right) + \mathbb{P}\left(W_{(p-s:p-s)}^2 > 2(1+2\delta_0)\log p\right) \\ &\stackrel{(b)}{\leq} \mathbb{P}\left(\frac{\sum_{j \in I_0} Z_j^2 - |I_0|}{|I_0|} > (3\delta_0\log p)/|I_0|\right) + \mathbb{P}\left(W_{(p-s:p-s)}^2 > 2(1+2\delta_0)\log p\right) \\ &\lesssim \exp(-0.75\delta_0\log p) + (p-s)\mathbb{P}\left(W_1^2 > 2(1+2\delta_0)\log p\right) \\ &\lesssim p^{-0.75\delta_0} + C\frac{p^{-2\delta_0}}{\sqrt{\log p}}. \end{split}$$

Inequality (a) uses $s - |I_0| \ge 1$ and inequality (b) uses $|I_0| < s < \delta_0 \log p$ (Condition (8)). Also, the first probability bound in (b) follows from Equation (56) in Wainwright (2009) and the fact that

Variable selection with heterogeneous signals

 $(3\delta_0 \log p)/|\mathcal{I}_0| \ge 6 > 4$. The last inequality in (b) follows from tail bound of standard Gaussian distribution. Now define the event $\mathcal{E}_{\mathcal{I}_0} := \{\|P_{\mathcal{S}\mathcal{S}}\|_2^2/(s \log p) > 2(1+4\delta_0)R_{\mathcal{I}_0}\}$ where $R_{\mathcal{I}_0} := (s - |\mathcal{I}_0|)/s$. Recall that

$$\frac{\|P_{Sg}\|_{2}^{2}}{s\log p} \geq \frac{\left\|\frac{\sum_{j\in S\setminus I_{0}}X_{j}\beta_{j}+P_{S}E}{(1+\|\beta_{S\setminus I_{0}}\|_{2}^{2})^{1/2}}\right\|_{2}^{2}}{s\log p} \leq \frac{\left\{\frac{\left\|\sum_{j\in S\setminus I_{0}}X_{j}\beta_{j}\right\|_{2}}{(1+\|\beta_{S\setminus I_{0}}\|_{2}^{2})^{1/2}} - \frac{\|P_{S}E\|_{2}}{(1+\|\beta_{S\setminus I_{0}}\|_{2}^{2})^{1/2}}\right\}^{2}}{s\log p} = (T_{1}^{1/2} - T_{2}^{1/2})^{2}.$$
(9)

where

$$T_{1} := \frac{\left\|\sum_{j \in S \setminus I_{0}} X_{j} \beta_{j}\right\|_{2}^{2}}{(1 + \left\|\beta_{S \setminus I_{0}}\right\|_{2}^{2}) s \log p} \ge \frac{2rR_{I_{0}}}{1 + 2r(s - |I_{0}|)(\log p)/n} \frac{V_{n}}{n}$$

and $V_n := \frac{\|\sum_{j \in S \setminus I_0} X_j \beta_j\|_2^2}{\|\beta_{S \setminus I_0}\|_2^2}$ is an χ_n^2 random variable. Also we have

$$T_2 := \frac{\|P_{\mathcal{S}}E\|_2^2}{(1 + \|\beta_{\mathcal{S} \setminus I_0}\|_2^2)s\log p} \le V_s/(s\log p)$$

where $V_s := \|P_S E\|_2^2$ is an χ_s^2 random variable independent of X_S . Next, we state the following simple algebraic relationship:

$$(1+6\delta_0)^{1/2} - \frac{\delta_0}{(1+6\delta_0)^{1/2} + (1+4\delta_0)^{1/2}} \ge (1+4\delta_0)^{1/2}.$$

In light of Equation (9) and using the above algebraic inequality we have the following:

$$\mathcal{E}_{\mathcal{I}_0}^c \subseteq \{T_1 \le 2(1+6\delta_0)R_{\mathcal{I}_0}\} \bigcup \{T_2 \ge 2\delta_0^2 \{(1+6\delta_0)^{1/2} + (1+4\delta_0)^{1/2}\}^{-2}R_{\mathcal{I}_0}\}$$

Next, we have

$$\mathbb{P}(T_1 \le 2(1+6\delta_0)R_{I_0}) \le \mathbb{P}\left(\frac{V_n}{n} \le \frac{1+6\delta_0}{1+8\delta_0}(1+2rs\log p/n)\right)$$

Now choose large *n* such that $(1 + 6\delta_0)(1 + 2rs \log p/n)) < (1 + 7\delta_0)$. Then for large *n* we have,

$$\mathbb{P}(T_1 \le 2(1+6\delta_0)R_{I_0}) \le \mathbb{P}\left(|V_n/n-1| \ge \frac{\delta_0}{1+8\delta_0}\right) \le \exp\left\{-C^* \frac{\delta_0^2}{(1+8\delta_0)^2}n\right\}$$

where C^* is a universal constant. Now we analyze the quantity T_2 . We have the following inequalities:

$$\mathbb{P}(T_2 \ge \frac{2\delta_0^2}{\{(1+6\delta_0)^{1/2} + (1+4\delta_0)^{1/2}\}^2} R_{I_0})$$

$$\le \mathbb{P}(V_s/s \ge \frac{2\delta_0^2}{\{(1+6\delta_0)^{1/2} + (1+4\delta_0)^{1/2}\}^2} R_{I_0} \log p)$$

$$\leq \mathbb{P}(V_s \geq \frac{2\delta_0^2}{\{(1+6\delta_0)^{1/2} + (1+4\delta_0)^{1/2}\}^2} \log p)$$

$$\leq \mathbb{P}(|V_s/s - 1| \geq 0.5 \frac{2\delta_0^2}{\{(1+6\delta_0)^{1/2} + (1+4\delta_0)^{1/2}\}^2} (\log p)/s)$$

$$\leq \exp(-C' \frac{2\delta_0^2}{\{(1+6\delta_0)^{1/2} + (1+4\delta_0)^{1/2}\}^2} \log p) \quad (\text{Using Condition (8)})$$

$$= p^{-C' \frac{2\delta_0^2}{\{(1+6\delta_0)^{1/2} + (1+4\delta_0)^{1/2}\}^2}} (C' > 0 \text{ is universal constant}).$$

Ultimately it shows that $\mathbb{P}(\mathcal{E}_{I_0}^c) \lesssim \exp\left\{-C^* \frac{\delta_0^2}{(1+8\delta_0)^2}n\right\} + p^{-C' \frac{2\delta_0^2}{\{(1+6\delta_0)^{1/2}+(1+4\delta_0)^{1/2}\}^2}}$. Now we are ready to show that the error probability goes to 0.

$$\begin{split} &\mathbb{P}_{\beta}(\widehat{S}\neq S) \leq \sum_{I_{0}\in S} \mathbb{P}_{\beta}\left(\max_{I_{1}:I_{1}\in\mathcal{J}_{I_{0}}} \frac{\|P_{I_{0}\cup I_{1}}Y\|_{2}^{2}}{s\log p} > \frac{\|P_{S}Y\|_{2}^{2}}{s\log p}\right) \\ &\leq \sum_{k=0}^{s-1} \sum_{I_{0}:|I_{0}|=k} \mathbb{P}_{\beta}\left(\max_{I_{1}:I_{1}\in\mathcal{J}_{I_{0}}} \frac{\|P_{I_{0}\cup I_{1}}Y\|_{2}^{2}}{s\log p} > \frac{\|P_{S}Y\|_{2}^{2}}{s\log p}\right) \\ &\leq \sum_{k=0}^{s-1} \sum_{I_{0}:|I_{0}|=k} \mathbb{P}_{\beta}\left(\max_{I_{1}:I_{1}\in\mathcal{J}_{I_{0}}} \frac{\|P_{I_{0}\cup I_{1}}Y\|_{2}^{2}}{s\log p} > \frac{\|P_{S}Y\|_{2}^{2}}{s\log p}, \mathcal{E}_{I_{0}}\right) + \mathbb{P}(\mathcal{E}_{I_{0}}^{c}) \\ &\leq \sum_{k=0}^{s-1} \sum_{I_{0}:|I_{0}|=k} \mathbb{P}_{\beta}\left(\max_{I_{1}:I_{1}\in\mathcal{J}_{I_{0}}} \frac{\|P_{I_{0}\cup I_{1}}Y\|_{2}^{2}}{s\log p} > \frac{\|P_{S}Y\|_{2}^{2}}{s\log p}, \mathcal{E}_{I_{0}}\right) + \mathbb{P}(\mathcal{E}_{I_{0}}^{c}) \\ &\leq \sum_{k=0}^{s-1} \sum_{I_{0}:|I_{0}|=k} \mathbb{P}_{\beta}\left(\max_{I_{1}:I_{1}\in\mathcal{J}_{I_{0}}} \frac{\|P_{I_{0}\cup I_{1}}Y\|_{2}^{2}}{s\log p} > 2(1+4\delta_{0})R_{I_{0}}\right) + \mathbb{P}(\mathcal{E}_{I_{0}}^{c}) \\ &\leq \sum_{k=0}^{s-1} \sum_{I_{0}:|I_{0}|=k} \mathbb{P}_{\beta}\left(\max_{I_{1}:I_{1}\in\mathcal{J}_{I_{0}}} \frac{\|P_{I_{0}\cup I_{1}}Y\|_{2}^{2}}{s\log p} > 2(1+4\delta_{0})R_{I_{0}}\right) + \mathbb{P}(\mathcal{E}_{I_{0}}^{c}) \\ &\lesssim \sum_{k=0}^{s-1} \sum_{I_{0}:|I_{0}|=k} \mathbb{P}_{\beta}\left(\sum_{I_{1}:I_{1}\in\mathcal{J}_{I_{0}}} \frac{\|P_{I_{0}\cup I_{1}}Y\|_{2}^{2}}{s\log p} + \exp\left\{-C^{*}\frac{\delta_{0}^{2}}{(1+8\delta_{0})^{2}n\right\} + p^{-C'}\frac{2\delta_{0}^{2}}{((1+6\delta_{0})^{1/2}+(1+4\delta_{0})^{1/2})^{2}} + n^{-1/8}\{sb_{n}(s,p)\}^{7/8}\right] \\ &\lesssim \sum_{k=0}^{s-1} \sum_{I_{0}:|I_{0}|=k} \mathbb{P}_{\beta}\left(\sum_{I_{0}:I_{0}\in\mathcal{I}} \frac{\delta_{0}^{2}}{(1+8\delta_{0})^{2}}n\right) + p^{-C'}\frac{2\delta_{0}^{2}}{((1+6\delta_{0})^{1/2}+(1+4\delta_{0})^{1/2})^{2}} + n^{-1/8}\{sb_{n}(s,p)\}^{7/8}\right] \\ &\lesssim \sum_{k=0}^{s-1} \sum_{I_{0}:I_{0}\in\mathcal{I}} \frac{S_{0}}{1} + \exp\left\{-C^{*}\frac{\delta_{0}^{2}}{(1+8\delta_{0})^{2}}n\right\} + p^{-C'\frac{2\delta_{0}^{2}}{((1+6\delta_{0})^{1/2}+(1+4\delta_{0})^{1/2})^{2}}} + n^{-1/8}\{sb_{n}(s,p)\}^{7/8}\right]. \end{split}$$

Variable selection with heterogeneous signals

Inequality (a) uses $\tilde{s}b_n(\tilde{s}, \tilde{p}) \leq sb_n(s, p)$ for large p. Thus if

$$s \lesssim \left(\delta_0 \wedge \frac{2\delta_0^2}{\{(1+6\delta_0)^{1/2} + (1+4\delta_0)^{1/2}\}^2} \wedge \frac{k}{16}\right) \log p$$
$$= \left(\frac{2\delta_0^2}{\{(1+6\delta_0)^{1/2} + (1+4\delta_0)^{1/2}\}^2} \wedge \frac{k}{16}\right) \log p,$$

then error probability goes to 0 uniformly over $\beta \in \mathcal{M}_{s}^{a}$.

B.2. Model consistency of BSS for sub-Gaussian model

In this section, we will show that Theorem 4.1 also holds beyond the Gaussian model. We assume that the entries of the design matrix X are i.i.d. *mean-zero* and *sub-Gaussian* with *unit variance*. We also assume that the entries of E are also i.i.d. *mean-zero* and *sub-Gaussian* with *unit variance* and independent of X.

In this setup, the results of Theorem 3.1 in Fan, Shao and Zhou (2018) are also valid and the proof steps follow exactly the same steps as the proof of Theorem 4.1 until the introduction of the random variables V_n and V_s .

We note that Gaussianity was only used to characterize the distributions of V_n and V_s . In particular, due to Gaussianity, we have $V_n \sim \chi_n^2$ and $V_s \sim \chi_s^2$. However, under the sub-Gaussian case, V_n is the sum of *n* independent sub-Exponential random variables with *unit-mean*. Hence, the probability bound for T_1 shown in the original proof is also valid.

Next, for V_s , we can use Theorem 1.1 of Rudelson and Vershynin (2013). Note that, there exists a constant $K_{\psi_2} > 1$ such that $\|\varepsilon\|_{\psi_2} \le K_{\psi_2}$, where $\|\varepsilon\|_{\psi_2} := \inf_{t>0} \{t : \mathbb{E} \exp(\varepsilon^2/t^2) \le 2\}$. By Theorem 1.1 of Rudelson and Vershynin (2013), we can obtain the same probability bound for T_2 if $s \le (0.5/K_{\psi_2}^2) \min\{\delta_0, \frac{2\delta_0^2}{\{(1+6\delta_0)^{1/2}+(1+4\delta_0)^{1/2}\}^2}\}\log p$. Hence, the rest of the proof is verbatim to the proof in the Gaussian case.

B.3. Proof of Theorem 4.3

In this section, we will show that BSS fails to recover the exact support when r = 1. We highlight three main steps of the proof:

- 1. Convert the BSS problem in the problem of selecting the model with maximum spurious correlation.
- 2. Use results from Fan, Shao and Zhou (2018) to find the asymptotic distribution of the maximum spurious correlation statistics.
- 3. Use the exact form of the asymptotic distribution along with scaling and centering parameters to approximate the recovery probability.

Recall the linear model $Y = X\beta + E$ with $S := S(\beta)$ as the set of active features and $s = |S| = O(\log p)$. As r = 1, there exists $j_0 \in S$ such that $\beta_{j_0} = \{(2 \log p)/n\}^{1/2}$. WLOG, let us assume that $j_0 = 1$ and define $I_0 = S \setminus \{1\}$. In order for BSS to recover the exact support S, it is necessary that

$$\max_{j \notin S} \left\| P_{I_0 \cup \{j\}} Y \right\|_2^2 < \| P_S Y \|_2^2$$

$$\Leftrightarrow \max_{j \notin S} \left\| P_{I_0 \cup \{j\}} \tilde{Y} \right\|_2^2 < \left\| P_S \tilde{Y} \right\|_2^2 \quad \text{(where } \tilde{Y} = X_1 \beta_1 + E\text{)}$$
$$\Leftrightarrow \max_{j \notin S} \left\| P_j^{\dagger} \tilde{Y} \right\|_2^2 < \left\| P_1^{\dagger} \tilde{Y} \right\|_2^2,$$

where P_j^{\dagger} is the orthogonal projection operator onto the sub-space span{ \tilde{X}_j } for all $j \in S^c \cup \{1\}$, where $\tilde{X}_j = (I_n - P_{\tilde{I}_0})X_j$. Due to Gaussianity, it follows that $\|\tilde{X}_j\|_2^2 \sim \chi_{n-s+1}^2$. Now, note that

$$\begin{aligned} \left\| P_1^{\dagger} \tilde{Y} \right\|_2^2 &= \beta_1^2 \left\| \tilde{X}_1 \right\|_2^2 + 2\beta_1 \tilde{X}_1^{\top} E + \left\| P_1^{\dagger} E \right\|_2^2 \\ &= (2\log p) \frac{\left\| \tilde{X}_1 \right\|_2^2}{n} + 2(2\log p)^{1/2} \frac{\tilde{X}_1^{\top} E}{\sqrt{n}} + \left\| P_1^{\dagger} E \right\|_2^2 \end{aligned}$$

Next, define the events

$$\mathcal{E}_{1} := \left\{ \left| \frac{\|\tilde{X}_{1}\|_{2}^{2}}{n-s+1} - 1 \right| \le 1/\sqrt{\log p} \right\}, \quad \mathcal{E}_{2} := \left\{ \frac{\tilde{X}_{1}^{\top}E}{\|\tilde{X}_{1}\|_{2}} \le -1 \right\}, \quad \mathcal{E}_{3} := \left\{ \left\| P_{1}^{\dagger}E \right\|_{2}^{2} \le 16 \right\}.$$

When p > 4, by Bernstein's inequality, we have $\mathbb{P}(\mathcal{E}_1^c) \le e^{-c_1 \frac{n}{\log p}}$, where $c_1 > 0$ is a universal constant. Recall that $n \ge p^k$, which tells that $\lim_{p \to \infty} \mathbb{P}(\mathcal{E}_1^c) = 0$. Next, note that $\frac{\tilde{X}_1^{\top} E}{\|\tilde{X}_1\|_2} \sim \mathsf{N}(0, 1)$. Next, we introduce a useful lemma.

Lemma B.3 (Gordon (1941)). Let $\Phi(\cdot)$ denote the cumulative distribution function of standard Gaussian distribution. Then for all $x \ge 0$, the following inequalities are true:

$$\left(\frac{x}{1+x^2}\right)\frac{e^{-x^2/2}}{\sqrt{2\pi}} \le 1 - \Phi(x) \le \left(\frac{1}{x}\right)\frac{e^{-x^2/2}}{\sqrt{2\pi}}.$$

By the above lemma we can conclude $\mathbb{P}(\mathcal{E}_2^c) \leq 1 - \frac{e^{-1/2}}{2\sqrt{2\pi}}$. Finally, as $\left\|P_1^{\dagger}E\right\|_2^2 \sim \chi_1^2$, we have $\mathbb{P}(\mathcal{E}_3^c) \leq 2e^{-8}$. Since n + 4 > 4s for large p, we have the following under $\mathcal{E}_1 \cap \mathcal{E}_2$:

$$\frac{\tilde{X}_1^{\top} E}{\sqrt{n}} = \frac{\tilde{X}_1^{\top} E}{\left\|\tilde{X}_1\right\|_2} \times \frac{\left\|\tilde{X}_1\right\|_2}{\sqrt{n-s+1}} \times \sqrt{\frac{n-s+1}{n}} \le -\frac{\sqrt{3}\{1 - (\log p)^{-1/2}\}}{2}$$

Here we used the fact that $\sqrt{1 - (\log p)^{-1/2}} > 1 - (\log p)^{-1/2}$. We define the event $\mathcal{E} := \bigcap_{i=1}^{3} \mathcal{E}_i$. Then, we have

$$\begin{split} \mathbb{P}(\widehat{S}_{\text{best}} = S) &\leq \mathbb{P}\left(\max_{j \notin S} \left\| P_j^{\dagger} \widetilde{Y} \right\|_2^2 < \left\| P_1^{\dagger} \widetilde{Y} \right\|_2^2 \right) \\ &\leq \mathbb{P}\left(\max_{j \notin S} \left\| P_j^{\dagger} \widetilde{Y} \right\|_2^2 < \left\| P_1^{\dagger} \widetilde{Y} \right\|_2^2, \mathcal{E} \right) + \mathbb{P}(\mathcal{E}^c) \\ &\leq \mathbb{P}\left\{\max_{j \notin S} \left\| P_j^{\dagger} \widetilde{Y} \right\|_2^2 < 2\log p \left(1 + \frac{1}{\sqrt{\log p}}\right) - (6\log p)^{1/2} \left(1 - \frac{1}{\sqrt{\log p}}\right) + 16 \right\} + \mathbb{P}(\mathcal{E}^c). \end{split}$$

10

We further note that $g := \tilde{Y}/(1 + \beta_1^2)^{1/2}$ follows a standard isotropic Gaussian distribution. Using Theorem 3.1 of Fan, Shao and Zhou (2018) and the fact that $1 + \beta_1^2 > 1$, we get

$$\begin{split} &\mathbb{P}(\widehat{S}_{\text{best}} = S) \\ &\leq \mathbb{P}\left\{Z_{(p-s:p-s)}^{2} \leq 2\log p \left(1 + \frac{1}{\sqrt{\log p}}\right) - (6\log p)^{1/2} \left(1 - \frac{1}{\sqrt{\log p}}\right) + 16\right\} + \mathbb{P}(\mathcal{E}^{c}) + o(1) \\ &\leq \mathbb{P}\left\{Z_{(p-s:p-s)}^{2} - 2\log(p-s) + \log\log(p-s) \leq \underbrace{-\left(\sqrt{6} - 2\right)(\log p)^{1/2} + \log\log p + O(1)}_{:=t_{p}}\right\} \\ &+ \mathbb{P}(\mathcal{E}^{c}) + o(1). \end{split}$$

where $Z^2_{(p-s:p-s)}$ is the maximum order statistics of $\{Z^2_j\}_{j \in [p-s]}$ with $\{Z_j\}_{j \in [p-s]}$ being i.i.d. standard Gaussian. Finally from Remark 3.3 of Fan, Shao and Zhou (2018), we know

$$Z^{2}_{(p-s:p-s)} - 2\log(p-s) + \log\log(p-s) \xrightarrow{d} \Lambda,$$

where $\mathbb{P}(\Lambda \le t) = \exp(-\pi^{-1/2} \exp(-t/2))$. As $t_p \to -\infty$, we have

$$\lim_{p \to \infty} \mathbb{P}(\widehat{\mathcal{S}}_{\text{best}} = \mathcal{S}) \le 1 - \frac{e^{-1/2}}{2\sqrt{2\pi}} + 2e^{-8} + \lim_{p \to \infty} \mathbb{P}(\mathcal{E}_1^c) < 0.9.$$

In other words, when r = 1, i.e., $a = \{2 \times 1 \times (\log p)/n\}^{1/2}$, we have

$$\lim_{p \to \infty} \sup_{\beta \in \mathcal{M}_{s}^{a}} \mathbb{P}(\widehat{\mathcal{S}}_{\text{best}} \neq \mathcal{S}) > \frac{1}{10}.$$

B.4. Minimax 0-1 loss under AURWM regime

We first present a result form Wang, Wainwright and Ramchandran (2010) which gives us the necessary condition for asymptotic exact recovery.

Theorem B.4 (Wang, Wainwright and Ramchandran (2010)). Consider the model (1) with the design matrix $X \in \mathbb{R}^{n \times p}$ be drawn with i.i.d elements from any distribution with zero mean and unit variance. Let $a := \min_{i \in S(\beta)} |\beta_i|$, i.e., it denote the minimum signal strength of β . Define the function

$$f_m(p,s,a) := \frac{\log \binom{p-s+m}{m} - 1}{\frac{1}{2} \log \left(1 + ma^2 (1 - \frac{m}{p-s+m})\right)}, \quad 1 \le m \le s.$$

Then $n \ge \max\{f_1(p, s, a), \dots, f_s(p, s, a), s\}$ is necessary for asymptotic exact recovery.

In the light of Theorem B.4 the proof of Proposition 4.4 follows immediately. To see this note that if r < 1 then there exists $\alpha \in (0, 1)$ such that $r = 1 - \alpha$. Also recall that $a = \{2r(\log p)/n\}^{1/2}$, $s = O(\log p)$

and $n = \lfloor p^k \rfloor$. Note that $f_1(p, s, a)/n \sim \frac{1}{1-\alpha}$. This shows that asymptotically the necessary condition in above theorem is violated and hence $r \ge 1$ is necessary.

Appendix C: Results related to ETS

C.1. Proof of Theorem 5.2

We first briefly describe the main steps of the proof:

- 1. We first establish the ℓ_2 -error bound of $\hat{\beta}$, i.e., we show that $\|\hat{\beta} \beta\|_2 \le \epsilon^{1/2}$.
- 2. Next, we upper bound the 0-1 loss $\mathbb{P}_{\beta}(\hat{\eta} \neq \eta)$ by decomposing it across the coordinates.
- 3. We analyze each of the terms separately and use ℓ_2 -error bound along with Gaussian tail inequalities to establish model consistency.

Now we are ready for the main proof. Due to Assumption 5.1, there exists a sequence $\{\alpha_p\}_{p\geq 1} \subseteq [0,\infty)$ converging to 0 such that with probability $1 - \alpha_p$ the following is true: \mathcal{A} requires no more than $T(\epsilon, p, \beta)$ iterations to output $\hat{\beta}$ that satisfies $\|\hat{\beta} - \beta\|_2 \leq \epsilon^{1/2}$.

For notational brevity, we write η instead of η_{β} . Define the event $\mathcal{H} = \{\|\hat{\beta} - \beta\|_2 \le \epsilon^{1/2}\}$ and we have $\mathbb{P}(\mathcal{H}^c) \le \alpha_p$. Note that $\hat{\beta}$ is based on the subsample \mathcal{D}_1 .

Next, for algebraic convenience we again reparametrize δ as $8\delta_0$ and set $\epsilon = 6\delta_0$, $\varsigma = (1 + \epsilon)^{1/2}$. Now note that for any $\beta \in \mathcal{M}_a^s$, we have

$$\begin{split} \mathbb{P}_{\beta}(\hat{\eta} \neq \eta | \mathcal{D}_{1}) &\leq \sum_{j:\beta_{j}=0} \mathbb{P}(\hat{\eta}_{j} = 1, \mathcal{H} | \mathcal{D}_{1}) + \sum_{j:\beta_{j}\neq0} \mathbb{P}_{\beta}(\hat{\eta}_{j} \neq 1, \mathcal{H} | \mathcal{D}_{1}) + \mathbb{P}(\mathcal{H}^{c} | \mathcal{D}_{1}) \\ &= \sum_{j:\beta_{j}=0} \mathbb{P}_{\beta}(\left|\Delta_{j}\right| > \kappa_{\varsigma}(X_{j}^{(2)}), \mathcal{H} | \mathcal{D}_{1}) + \sum_{j:\beta_{j}\neq0} \mathbb{P}_{\beta}(\left|\Delta_{j}\right| \leq \kappa_{\varsigma}(X_{j}^{(2)}), \mathcal{H} | \mathcal{D}_{1}) + \mathbb{P}(\mathcal{H}^{c} | \mathcal{D}_{1}), \end{split}$$

Using the fact that conditionally on $\hat{\beta}$ and X_j , the random variable Δ_j has the same distribution as the random variable in (11) in the main paper, we conclude that for all $j \notin S(\beta)$,

$$\begin{split} \mathbb{P}(\eta_{j} = 1, \mathcal{H} \mid \mathcal{D}_{1}) &\leq \mathbb{P}\Biggl((1+\epsilon)^{1/2} \left| g_{j} \right| > \frac{a \left\| X_{j}^{(2)} \right\|_{2}}{2} + \frac{(1+\epsilon) \log p}{a \left\| X_{j}^{(2)} \right\|_{2}}, \mathcal{H} \mid \mathcal{D}_{1} \Biggr) \\ &= 2\mathbb{E}\left\{ \overline{\Phi}\Biggl(\frac{a \left\| X_{j}^{(2)} \right\|_{2}}{2(1+\epsilon)^{1/2}} + \frac{(1+\epsilon)^{1/2} \log p}{a \left\| X_{j}^{(2)} \right\|_{2}} \Biggr) \right\}. \end{split}$$

Here $\overline{\Phi}(\cdot)$ denotes the survival function of the standard Gaussian random variable. Now note that for each *j* we have $\|X_j^{(2)}\|_2^2 \stackrel{d}{=} V_{n_2}$, where V_{n_2} is a chi-squared random variable with n_2 degrees of freedom. Thus we have

$$\mathbb{P}(\eta_j = 1, \mathcal{H} \mid \mathcal{D}_1) \le 2\mathbb{E}\left\{\overline{\Phi}\left(\frac{aV_{n_2}^{1/2}}{2(1+\epsilon)^{1/2}} + \frac{(1+\epsilon)^{1/2}\log p}{aV_{n_2}^{1/2}}\right)\right\}.$$

Variable selection with heterogeneous signals

Analogous argument and the fact that $|\beta_j| \ge a$ for all $\beta_j \ne 0$, leads to the fact that for all $j \in S(\beta)$,

$$\mathbb{P}(\eta_j \neq 1, \mathcal{H} \mid \mathcal{D}_1) \le 2\mathbb{E}\left\{\overline{\Phi}\left(\max\left\{\frac{aV_{n_2}^{1/2}}{2(1+\epsilon)^{1/2}} - \frac{(1+\epsilon)^{1/2}\log p}{aV_{n_2}^{1/2}}, 0\right\}\right)\right\}$$

Now recall that $\epsilon = 6\delta_0$ and $\gamma \in (0, \frac{\delta_0}{1+8\delta_0})$. With this choice of tuning parameters it is easy to see that $r(1-\gamma)/(1+\epsilon) \ge \frac{1+7\delta_0}{1+6\delta_0} > 1$ and hence as $p \to \infty$ we have

$$W_{n_2} := \frac{1}{(\log p)^{1/2}} \left(\frac{aV_{n_2}^{1/2}}{2(1+\epsilon)^{1/2}} - \frac{(1+\epsilon)^{1/2}\log p}{aV_{n_2}^{1/2}} \right)$$
$$\xrightarrow{\mathbf{p}} \frac{1}{(2r)^{1/2}} \left\{ r \left(\frac{1-\gamma}{1+\epsilon} \right)^{1/2} - \left(\frac{1+\epsilon}{1-\gamma} \right)^{1/2} \right\} > 0.$$

The above display uses the fact that $n_2/n \to 1 - \gamma$ and $V_{n_2}/n_2 \xrightarrow{p} 1$ as $p \to \infty$. Next let is define the following quantity q:

$$q := q(\epsilon, \delta_0, \gamma) = \frac{1}{\{2(1+8\delta_0)\}^{1/2}} \left\{ (1+8\delta_0) \left(\frac{1-\gamma}{1+\epsilon}\right)^{1/2} - \left(\frac{1+\epsilon}{1-\gamma}\right)^{1/2} \right\}$$

Due to choice of ϵ and γ it is easy to show q > 0. Now define the event $G_{n_2} := \{W_{n_2} > q/2\}$. Before we proceed it is useful to note the following:

$$W_{n_2} = \frac{1}{(2r)^{1/2}} \left(\frac{r\{V_{n_2}/(n(1-\gamma))\}^{1/2}(1-\gamma)^{1/2}}{(1+\epsilon)^{1/2}} - \frac{(1+\epsilon)^{1/2}}{\{V_{n_2}/(n(1-\gamma))\}^{1/2}(1-\gamma)^{1/2}} \right).$$

Next define the function

$$H(u) := \frac{1}{(2+16\delta_0)^{1/2}} \left\{ u(1+8\delta_0) \left(\frac{1-\gamma}{1+\epsilon}\right)^{1/2} - \frac{1}{u} \left(\frac{1+\epsilon}{1-\gamma}\right)^{1/2} \right\}, \quad u > 0$$

As $r = 1 + 8\delta_0$ we have $W_{n_2} = H(\{V_{n_2}/(n(1-\gamma))\}^{1/2})$. It is also easy to see that $H(\cdot)$ is strictly increasing function on $(0,\infty)$ and H(1) = q. Hence $\lambda_{\delta_0} := H^{-1}(q/2) \in (0,1)$. Now $G_{n_2}^c = \{W_{n_2} \le q/2\} \subseteq \{H(\{V_{n_2}/(n(1-\gamma))\}^{1/2}) \le q/2\}$. Thus a straightforward calculation shows that

$$\mathbb{P}(G_{n_2}^c) \le \mathbb{P}\left(\frac{V_{n_2}}{n_2} \le \frac{n(1-\gamma)}{n_2}\lambda_{\delta_0}^2\right).$$

Choose *p* large enough such that $n_2/n > \lambda_{\delta_0}(1-\gamma)$ and hence we have,

$$\mathbb{P}(G_{n_2}^c) \le \mathbb{P}\left(\frac{V_{n_2}}{n_2} \le \lambda_{\delta_0}\right) \lesssim \exp(-K_{\delta_0}n_2),$$

where $K_{\delta_0} = (1 - \lambda_{\delta_0})^2 / 8$. Note that $\overline{\Phi}(t) \le e^{-t^2/2}$ for all t > 0. Using this fact we have the following:

$$\mathbb{P}(\eta_j = 1, \mathcal{H} \mid \mathcal{D}_1) \le \mathbb{E}\left[\exp\left\{-\left(1 + \frac{W_{n_2}^2}{2}\right)\log p\right\} \mathbb{1}_{G_{n_2}}\right] + \mathbb{P}(G_{n_2}^c) \lesssim p^{-(1+q^2/8)} + \exp(-K_{\delta_0}n_2),$$

for all $j \notin S(\beta)$. Similarly,

$$\mathbb{P}(\eta_j \neq 1, \mathcal{H} \mid \mathcal{D}_1) \lesssim p^{-q^2/8} + \exp(-K_{\delta_0} n_2), \quad \forall j \in \mathcal{S}(\beta).$$

Thus marginalizing out \mathcal{D}_1 and summing over all j we get,

$$\sup_{\beta \in \mathcal{M}_a^s} \mathbb{P}_{\beta}(\hat{\eta} \neq \eta) \lesssim p^{-q^2/8} \log p + p \exp(-K_{\delta_0} n_2) + \alpha_p.$$

The result follows by taking $p \to \infty$.

Remark C.1. Note that $q^2 = \Omega(\frac{\delta_0^2}{1+\delta_0^2})$ and it shows that the upper bound in the above display deteriorates as $\delta_0 \to 0$. Also, as δ_0 approaches 0, the term K_{δ_0} also approaches 0. Hence, the rate of decay worsens as $\delta_0 \to 0$, and ETS continues to lose statistical power.

C.2. Proof of Corollary 5.3

Similar to previous proofs, we reparametrize δ by $8\delta_0$ and set $\epsilon = 6\delta_0$, $\varsigma = (1 + \epsilon)^{1/2}$. Now note that it is enough to prove the following:

$$\lim_{p \to \infty} \inf_{\beta \in \mathcal{M}_{s}^{a}} \mathbb{P}_{\beta} \left(\max_{j \notin \mathcal{S}(\beta)} \left| \Delta_{j} \right| < \min_{j \in \mathcal{S}(\beta)} \left| \Delta_{j} \right| \right) \to 1$$

as $p \to \infty$. To this end first define the following quantity:

$$t_p := \frac{\left(\frac{2rn_2\log p}{n}\right)^{1/2}}{2} + \frac{\varsigma^2\log p}{\left(\frac{2rn_2\log p}{n}\right)^{1/2}}$$

We will show that $\lim_{p\to\infty} \inf_{\beta\in\mathcal{M}_s^a} \mathbb{P}_{\beta}(\min_{j\in\mathcal{S}(\beta)} |\Delta_j| > t_p, \max_{j\notin\mathcal{S}(\beta)} |\Delta_j| \le t_p) \to 1$ as $p \to \infty$. For convenience let us define the events $G_{\min} := \{\min_{j\in\mathcal{S}(\beta)} |\Delta_j| > t_p\}$ and $G_{\max} := \{\max_{j\notin\mathcal{S}(\beta)} |\Delta_j| \le t_p\}$. Let \mathcal{H} be the event as defined in Section C.1. First we will analyze $\mathbb{P}_{\beta}(G_{\min}^c)$. Note that $\mathbb{P}_{\beta}(G_{\min}^c) \le \mathbb{P}_{\beta}(G_{\min}^c \cap \mathcal{H}) + \mathbb{P}_{\beta}(\mathcal{H}^c)$. Now the second term goes to 0 uniformly over $\beta \in \mathcal{M}_s^a$. Also using Equation (11) under the event \mathcal{H} we get

$$\begin{split} \sup_{\beta \in \mathcal{M}_{\mathcal{S}}^{a}} & \mathbb{P}_{\beta}(G_{\min}^{c} \cap \mathcal{H}) \\ \leq \sup_{\beta \in \mathcal{M}_{\mathcal{S}}^{a}} & \mathbb{P}_{\beta}\left(\min_{j \in \mathcal{S}(\beta)} \left|\beta_{j}\|X_{j}^{(2)}\|_{2} + (1+\epsilon)^{1/2}g_{j}\right| \le t_{p}\right) \\ \leq \sup_{\beta \in \mathcal{M}_{\mathcal{S}}^{a}} & \mathbb{P}_{\beta}\left(\max_{j \in \mathcal{S}(\beta)} \frac{|g_{j}|}{(\log p)^{1/2}} \ge \frac{1}{(1+\epsilon)^{1/2}(\log p)^{1/2}} \left\{a\min_{j \in \mathcal{S}(\beta)} \|X_{j}^{(2)}\|_{2} - t_{p}\right\}\right) \\ \leq \sup_{\beta \in \mathcal{M}_{\mathcal{S}}^{a}} & \mathbb{P}_{\beta}\left(\max_{j \in \mathcal{S}(\beta)} \frac{|g_{j}|}{(\log p)^{1/2}} \ge \frac{1}{(1+\epsilon)^{1/2}(\log p)^{1/2}} \left\{a\min_{j \in [p]} \|X_{j}^{(2)}\|_{2} - t_{p}\right\}\right) \end{split}$$

where $\{g_j\}_{j \in \mathcal{S}(\beta)}$ are non i.i.d. standard Gaussian. Note that $|\mathcal{S}(\beta)| = O(\log p)$. Hence

$$\max_{j \in \mathcal{S}(\beta)} |g_j| = O_{\mathbb{P}}(\log \log p),$$

which tells that

$$\max_{j \in \mathcal{S}(\beta)} \frac{|g_j|}{(\log p)^{1/2}} \xrightarrow{\mathsf{p}} 0.$$

Also using lemma 3 from Fletcher, Rangan and Goyal (2009) we have

$$\frac{1}{(1+\epsilon)^{1/2} (\log p)^{1/2}} \left(a \min_{j \in [p]} \|X_j^{(2)}\|_2 - t_p \right) \xrightarrow{p} \frac{1}{(2r)^{1/2}} \left\{ r \left(\frac{1-\gamma}{1+\epsilon} \right)^{1/2} - \left(\frac{1+\epsilon}{1-\gamma} \right)^{1/2} \right\}.$$

The right-hand side of the above display is at least $q(\epsilon, \delta_0, \gamma)$ (defined in Section C.1) which is strictly positive. Again for compactness we use q instead of $q(\epsilon, \delta_0, \gamma)$. The above display motivates us to define the following event:

$$\mathcal{E}_p = \left\{ \frac{1}{(1+\epsilon)^{1/2} (\log p)^{1/2}} \left(a \min_{j \in [p]} \|X_j^{(2)}\|_2 - t_p \right) \ge q/2 \right\},\$$

and it follows that $\mathbb{P}(\mathcal{E}_p^c) \to 0$ as $p \to \infty$. This leads to the following inequality:

$$\sup_{\beta \in \mathcal{M}_{s}^{a}} \mathbb{P}_{\beta}(G_{\min}^{c} \cap \mathcal{H}) \leq \sup_{\beta \in \mathcal{M}_{s}^{a}} \mathbb{P}_{\beta}\left(\max_{j \in \mathcal{S}(\beta)} \frac{|g_{j}|}{(\log p)^{1/2}} \geq q/2\right) + \mathbb{P}(\mathcal{E}_{p}^{c})$$
$$\leq p^{-q^{2}/8} \log p + \mathbb{P}(\mathcal{E}_{p}^{c}) \to 0.$$

Thus we have $\sup_{\beta \in \mathcal{M}_s^a} \mathbb{P}_{\beta}(G_{\min}^c) \to 0$. Similarly it can be shown that $\sup_{\beta \in \mathcal{M}_s^a} \mathbb{P}_{\beta}(G_{\max}^c) \to 0$ as $p \to \infty$. These two claims together complete the proof.

C.3. Discussion on Remark 5.4

As $r > 1 + \delta_*$, by reparameterizing δ_* by 8δ , we have $r > 1 + 8\delta$. Now we are basically back to the setting of the proof of Theorem 5.1 and all of the proof steps are exactly the same as that of Theorem 5.1 with δ in place of δ_0 . This allows us to choose the threshold using the knowledge of δ_* , and we do not need the knowledge of a. In particular, one can construct the threshold $\kappa_{\varsigma}(X_i^{(2)})$ with $\varsigma = (1 + A_2 \delta_*)^{1/2}$, where A_2 is the same universal constant as described in Theorem 5.2

C.4. Discussion on examples of ETS

Solving ℓ_0 -regularized problem:

Proofs for ETS-IHT:

We first introduce some standard assumptions for analyzing ETS-IHT.

Definition C.2 (RSC property). A differentiable function $F : \mathbb{R}^p \to \mathbb{R}$ is said to satisfy restricted strong convexity (RSC) at sparsity level $s = s_1 + s_2$ with strong convexity constraint ℓ_s if the following holds for all θ_1, θ_2 s.t. $\|\theta_1\|_0 \le s_1$ and $\|\theta_2\|_0 \le s_2$:

$$F(\theta_1) - F(\theta_2) \ge \langle \theta_1 - \theta_2, \nabla_{\theta} F(\theta_2) \rangle + \frac{\ell_s}{2} \|\theta_1 - \theta_2\|_2^2.$$

Definition C.3 (RSS property). A differentiable function $F : \mathbb{R}^p \to \mathbb{R}$ is said to satisfy restricted strong smoothness (RSS) at sparsity level $s = s_1 + s_2$ with strong smoothness constraint L_s if the following holds for all θ_1, θ_2 s.t. $\|\theta_1\|_0 \le s_1$ and $\|\theta_2\|_0 \le s_2$:

$$F(\theta_1) - F(\theta_2) \le \langle \theta_1 - \theta_2, \nabla_{\theta} F(\theta_2) \rangle + \frac{L_s}{2} \|\theta_1 - \theta_2\|_2^2.$$

Now we quote an important theorem from Jain, Tewari and Kar (2014) which quantifies the suboptimality gap of Algorithm 2.

Theorem C.4 (Jain, Tewari and Kar (2014)). Let *F* has RSC and RSS parameters given by $\ell_{2\hat{s}+s}(F) = \alpha$ and $L_{2\hat{s}+\hat{\pi}}(F) = L$ respectively. Call Algorithm 2 with $\hat{s} \ge 32L^2\ell^{-2}s$ and h = 2/(3L). Also let $\hat{\beta} = \arg\min_{\theta, \|\theta\|_0 \le s} F(\theta)$. Then the iterate of Algorithm 2 for $t = O(L\ell^{-1}\log(F(\beta^{(0)})/\epsilon))$ satisfies:

$$F(\beta^{(t)}) - F(\hat{\beta}) \le \epsilon.$$

In our setup, the observations $\{x_i\}_{i=1}^n$ are coming from i.i.d. mean zero isotropic Gaussian distribution. Thus, lemma 6 from Agarwal, Negahban and Wainwright (2012) immediately tells that RSC and RSS at any sparsity level *m* hold for $f_{n_1}(\cdot)$ with probability at least $1 - \exp(-c_0n_1)$ with $\ell_m = \frac{1}{2} - c_1(m \log p)/n_1$ and $L_m = 2 + c_1(m \log p)/n_1$, where c_0, c_1 are universal constants. Now set $m = 2\hat{s} + s$ and recall that $n_1 \sim \gamma p^k$. If $n_1 > 4c_1(2\hat{s} + s) \log p$ then we have $\ell_m \ge 1/4$ and $L_m \le 9/4$, which means that $L_m/(9\ell_m) \le 1$. Thus to apply Theorem C.4 it is enough to choose $\hat{s} = 2592s$. Also by the assumption on *n* for large *p* we have $n_1 > 4c_1(2\hat{s} + s) \log p$. Let $f_{n_1}(\theta) := n_1^{-1} ||Y^{(1)} - X^{(1)}\theta||_2^2$ for $\theta \in \mathbb{R}^p$. Note that $f_{n_1}(0) = n_1^{-1} ||Y^{(1)}||_2^2 \stackrel{d}{=} (1 + ||\beta||_2^2)V_{n_1}/n_1$, where V_{n_1} is chisquare random variable with n_1 degrees of freedom. Also by Bernstein's type inequality it follows that $|(V_{n_1}/n_1) - 1| \le 1/2$ with probability at least $1 - \exp(-c_4n_1)$, where c_4 is a universal positive constant. Thus if $t = O(L_m \ell_m^{-1} \log((1 + ||\beta||_2^2)/\epsilon_0)) = O(\log p + \log((1 + ||\beta||_{\infty})/\epsilon_0))$, then we have $f_{n_1}(\beta^{(t)}) - f_{n_1}(\hat{\beta}) \le \epsilon_0$. Thus by Theorem 3 of Jain, Tewari and Kar (2014) it follows that with probability at least $1 - \exp(-c_0n_1) - \exp(-c_4n_1) - c_2p^{-c_3}(c_2, c_3$ are universal constants) we have

$$\|\beta^{(t)} - \beta\|_2 \le C \left(\frac{s\log p}{n_1}\right)^{1/2} + (8\epsilon_0)^{1/2} \le (9\epsilon_0)^{1/2}, \text{ for large } p.$$

C is a positive universal constant in the above inequality. If we set $\epsilon = 9\epsilon_0$, then Assumption 5.1 holds with α_p equal to $\exp(-c_0n_1) + \exp(-c_4n_1) + c_2p^{-c_3}$ and $T(\epsilon, p, \beta) = O(\log p + \log((1 + ||\beta||_{\infty})/\epsilon)))$.

Solving ℓ_1 -regularized problem:

We start by recalling the definition of the loss function $f_{n_1,\lambda}(\theta) = n_1^{-1} ||Y^{(1)} - X^{(1)}\theta||_2^2 + \lambda ||\theta||_1$, and define the minimizer of the loss function $\hat{\beta}_L := \arg \min_{\theta} f_{n_1,\lambda}(\theta)$. First, we will prove Proposition 5.5.

Proof of Proposition 5.5:

Let us assume

$$f_{n_1,\lambda}(\hat{\beta}) - f_{n_1,\lambda}(\hat{\beta}_L) \le \epsilon_0 < 1.$$

Now, we will establish the ℓ_2 -error rate between $\hat{\beta}$ and β . We write $\hat{b} = \hat{\beta} - \beta$. Since, $\hat{\beta}_L$ is optimal, we have

$$f_{n_1,\lambda}(\hat{\beta}) \le f_{n_1,\lambda}(\hat{\beta}_L) + \epsilon_0 \le f_{n_1,\lambda}(\beta) + \epsilon_0$$

Rearrangement of the above inequality yields

$$0 \le \frac{1}{n_1} \left\| X^{(1)} \hat{b} \right\|_2^2 \le \frac{2E^\top X^{(1)} \hat{b}}{n_1} + \lambda (\|\beta\|_1 - \|\hat{\beta}\|_1) + \epsilon_0.$$
(10)

Since, $||X_j^{(1)}||_2^2 \sim \chi_{n_1}^2$, we have

$$\mathbb{P}\left(\underbrace{\max_{j \in [p]} n_1^{-1/2} \left\| X_j^{(1)} \right\|_2 < \sqrt{4/3}}_{:=\mathcal{E}} \right) \ge 1 - 2p^{-2}, \text{ for large } p.$$

Using Gaussianity of $X_j^{(1)\top} E / ||X_j^{(1)}||_2$, we have

$$\mathbb{P}\left(\frac{1}{n_1} \left\| X^{(1)\top} E \right\|_{\infty} > 2\sqrt{\frac{\log p}{n_1}} \right) \le \mathbb{P}\left(\max_{j \in [p]} \left| \frac{X_j^{(1)\top} E}{\left\| X_j^{(1)} \right\|_2} \right| > \sqrt{3\log p} \right) + \mathbb{P}(\mathcal{E}^c)$$
$$\le 2p^{-0.5} + 2p^{-2}.$$

Setting $\lambda = 8\{(\log p)/n_1\}^{1/2}$, we have $\lambda \ge 2 \|X^{(1)\top}E\|_{\infty}/n_1$ with probability at least $1 - 2p^{-0.5} - 2p^{-2}$. Now, since β is *s*-sparse with support on S, we have

$$\|\beta\|_{1} - \|\hat{\beta}\|_{1} = \|\beta_{\mathcal{S}}\|_{1} - \|\beta_{\mathcal{S}} + \hat{b}_{\mathcal{S}}\|_{1} - \|\hat{b}_{\mathcal{S}^{c}}\| \le \|\hat{b}_{\mathcal{S}}\|_{1} - \|\hat{b}_{\mathcal{S}^{c}}\|_{1}$$

Substituting this in the basic inequality (10) and using $\lambda \ge 2 \|X^{(1)\top}E\|_{\infty}/n_1$, we get

$$0 \leq \frac{1}{n_1} \left\| X^{(1)} \hat{b} \right\|_2^2 \leq 2 \left\| \frac{E^\top X^{(1)}}{n_1} \right\|_{\infty} \left\| \hat{b} \right\|_1 + \lambda(\left\| \hat{b}_{\mathcal{S}} \right\|_1 - \left\| \hat{b}_{\mathcal{S}^c} \right\|_1) + \epsilon_0$$

$$\leq (\lambda/2) \left\| \hat{b} \right\|_1 + \lambda(\left\| \hat{b}_{\mathcal{S}} \right\|_1 - \left\| \hat{b}_{\mathcal{S}^c} \right\|_1) + \epsilon_0$$

$$\leq (\lambda/2) \{ 3 \left\| \hat{b}_{\mathcal{S}} \right\|_1 - \left\| \hat{b}_{\mathcal{S}^c} \right\|_1 \} + \epsilon_0.$$
(11)

Hence, we have

$$\|\hat{b}\|_{1}^{2} = (\|\hat{b}_{\mathcal{S}}\|_{1} + \|\hat{b}_{\mathcal{S}^{c}}\|_{1})^{2} \le (4\|\hat{b}_{\mathcal{S}}\|_{1} + (2\epsilon_{0}/\lambda))^{2} \le 32\|\hat{b}_{\mathcal{S}}\|_{1}^{2} + \frac{8\epsilon_{0}^{2}}{\lambda^{2}}.$$

Now by Theorem 7.16 of Wainwright (2019), we have the following with probability at least $1 - 2 \exp(-n_1/32)$:

$$\frac{\left\|X^{(1)}\theta\right\|_{2}^{2}}{n_{1}} \ge c_{1} \left\|\theta\right\|_{2}^{2} - c_{2} \frac{\log p}{n_{1}} \left\|\theta\right\|_{1}^{2} \quad \text{for every } \theta \in \mathbb{R}^{p},$$

where $c_1, c_2 > 0$ are universal constants. Using the above fact we have

$$\frac{\left\|X^{(1)}\hat{b}\right\|_{2}^{2}}{n_{1}} \ge c_{1}\left\|\hat{b}\right\|_{2}^{2} - 32c_{2}\frac{s\log p}{n_{1}}\left\|\hat{b}\right\|_{2}^{2} - 8c_{2}\frac{\log p}{n_{1}\lambda^{2}}\epsilon_{0}^{2} \ge \frac{c_{1}}{2}\left\|\hat{b}\right\|_{2}^{2} - \epsilon_{0}^{2},$$

when $32c_2s \log p/(n_1) < c_1/2$ and $8c_2 \log p/(n_1\lambda^2) < 1$. This is possible for large enough values of p and λ .

Case 1: If $(c_1/4) \|\hat{b}\|_2^2 > \epsilon_0^2$, then using (11), we get

$$\frac{c_1}{4} \left\| \hat{b} \right\|_2^2 \le \frac{3\lambda\sqrt{s}}{2} \left\| \hat{b} \right\|_2 + \epsilon_0.$$

This bound involves a quadratic form of $\|\hat{b}\|_2$; computing the roots of the quadratic form we get the following bound:

$$\|b\|_2 \leq \underbrace{\frac{6\lambda\sqrt{s}}{c_1}}_{=O(\sqrt{(s\log p)/n_1})} + \frac{2\sqrt{\epsilon_0}}{\sqrt{c_1}}.$$

Case 2: If $(c_1/4) \|\hat{b}\|_2^2 \le \epsilon_0^2$, then $\|\hat{b}\|_2 \le 2\epsilon_0/\sqrt{c_1} < 2\sqrt{\epsilon_0/c_1}$. The last inequality uses the fact that $\epsilon_0 < 1$.

Combining the bounds obtained in Case 1 and Case 2 and with probability at least $1-2p^{-0.5}-2p^{-2}-e^{-n_1/32}$ (which is $\ge 1-3p^{-0.5}$ for large p), we finally have $\|\hat{\beta}-\beta\|_2 \le (C_3\epsilon_0)^{1/2}$ for large enough p and C_3 being an absolute constant. Now, using the reparameterization $\epsilon = C_3\epsilon_0$, we have $\epsilon < C_3$, and the initial sub-optimality gap turns out to be ϵ/C_3 . This finishes the proof.

Proof for ETS-PICASSO:

First, we will show that the estimator generated by the PICASSO algorithm is a good estimate of β . Let $\hat{\beta}^{\text{picasso}}$ be the estimator obtained by applying PICASSO for minimizing $f_{n_1}(\theta; \lambda, \mathcal{D}_1)$. Now define the largest and smallest s_0 -sparse eigenvalues of $G := X^{(1)\top} X^{(1)} / n_1$ as:

$$\rho_+(s_0) := \max_{\nu: \|\nu\|_0 \le s} \frac{\nu^\top G \nu}{\|\nu\|_2^2}; \quad \text{and} \quad \rho_-(s_0) := \min_{\nu: \|\nu\|_0 \le s} \frac{\nu^\top G \nu}{\|\nu\|_2^2}.$$

18

Let $\tilde{s}_{\psi} = (484\psi^2 + 100\psi)s$, where $\psi > 0$ is constant. An application of union bound and Equation 4.22 in Vershynin (2018) yields that

$$\mathbb{P}\left\{\max_{\mathcal{D}:|\mathcal{D}|\leq s+2\tilde{s}_{2}}\left\|\frac{X_{\mathcal{D}}^{(1)\top}X_{\mathcal{D}}^{(1)}}{n_{1}}-\mathsf{I}_{p}\right\|_{\mathrm{op}}\lesssim\frac{\log p}{p^{k/2}}\right\}\geq1-2p^{-2},$$
(12)

for large values of p. This ensures that for large values of p, we have

$$0.99 \le \rho_{-}(s+2\tilde{s}_{2}) \le \rho_{+}(s+2\tilde{s}_{2}) \le 1.1.$$
(13)

Defining $\kappa := \rho_+(s + 2\tilde{s}_2)/\rho_-(s + 2\tilde{s}_2)$, we have $\kappa < 2$. Thus, Assumption 3.5 of Zhao, Liu and Zhang (2018) holds with $\tilde{s} = \tilde{s}_2 > \tilde{s}_{\kappa}$. Also, (12) shows that for large *p*

$$\mathbb{P}\left(\underbrace{\max_{j \in [p]} n_1^{-1/2} \left\| X_j^{(1)} \right\|_2 < \sqrt{4/3}}_{:=\mathcal{E}}\right) \ge 1 - 2p^{-2}.$$

Hence, we have

$$\mathbb{P}\left(\frac{1}{n_1} \left\| X^{(1)\top} E \right\|_{\infty} > 2\sqrt{\frac{\log p}{n_1}} \right) \le \mathbb{P}\left(\max_{j \in [p]} \left| \frac{X_j^{(1)\top} E}{\left\| X_j^{(1)} \right\|_2} \right| > \sqrt{3\log p} \right) + \mathbb{P}(\mathcal{E}^c)$$

$$\le 2p^{-0.5} + 2p^{-2}.$$

Hence, Assumption 3.1 of Zhao, Liu and Zhang (2018) holds with high probability when $\lambda_N = \lambda \ge 8\{\log p/n_1\}^{1/2}$, where *N* denotes the final iteration count of the outermost loop of PICASSO and λ_K denotes the regularization parameter at the *K*th iteration of the outer loop. Also, in this case, $N = O(\log(||X^{(1)^{\top}}Y^{(1)}/n_1||_{\infty}\sqrt{n_1/\log p}))$ which follows from the description of PICASSO (see Algorithm 3 in Zhao, Liu and Zhang (2018)). From triangle inequality, it follows that

$$\left\|\frac{X^{(1)\top}Y^{(1)}}{n_1}\right\|_{\infty} \le \|G\beta\|_{\infty} + \left\|X^{(1)\top}E/n_1\right\|_{\infty},$$

and $||G\beta||_{\infty} \le ||G||_{\infty,\infty} ||\beta||_{\infty} \le \sqrt{p} ||G||_{\text{op}} ||\beta||_{\infty}$. Note that

$$||G||_{\text{op}} = \frac{p}{n_1} \left\| \frac{X^{(1)} X^{(1)\top}}{p} \right\|_{\text{op}}.$$

Thus, applying Equation (4.22) in Vershynin (2018), we get $||G\beta||_{\infty} \leq (p^{1.5}/n) ||\beta||_{\infty}$ with probability at least $2 \exp(-n_1)$. Hence, we have $N = O(\log p)$ with probability at least $1 - 2p^{-0.5} - 2p^{-2} - \exp(n_1)$.

Now will make sure that Assumption 3.7 of Zhao, Liu and Zhang (2018) also holds. Before, going any further let us clarify some notations. At *K*th outer iteration of PICASSO, Zhao, Liu and Zhang (2018) denotes the inner and middle loop precision parameters as τ_K and δ_K , and the active set initialization parameter is φ . To be consistent with our notation we set $\epsilon_0 = \delta_K$ for every *K*. Also, we choose the

parameters in such a way so that

$$\epsilon_0 \le \min\{1/8, C_3\}, \quad \tau_K \le \frac{\epsilon_0}{\rho_+(s+2\tilde{s})} \sqrt{\frac{\rho_-(1)}{\rho_+(1)(s+2\tilde{s})}}, \quad \varphi \le 1/8,$$
 (14)

 C_3 is the same constant ad in Proposition 5.5. Using (13), one can set $\tau_K = O(\epsilon_0/\sqrt{\log p})$ which will be less than 1 for large *p*. So, under the above conditions, Assumption 3.7 in Zhao, Liu and Zhang (2018) holds. Hence, part (iii) of Theorem 3.12 in Zhao, Liu and Zhang (2018) tells that

$$f_{n_1}(\widehat{\beta}^{\text{picasso}};\lambda,\mathcal{D}_1) - f_{n_1}(\widehat{\beta}_L;\lambda,\mathcal{D}_1) \le \epsilon_0 \frac{500\lambda^2 s}{11}.$$

If $\lambda = 8\sqrt{(\log p)/n_1}$, then for large p

$$f_{n_1}(\hat{\beta}^{\mathsf{picasso}}; \lambda, \mathcal{D}_1) - f_{n_1}(\hat{\beta}_L; \lambda, \mathcal{D}_1) \le \epsilon_0 / C_3$$

Hence, due to Proposition 5.5, we have $\|\widehat{\beta}^{\text{picasso}} - \beta\|_2^2 \le \epsilon_0$. Also, using Theorem 3.12 and Lemma 3.13 of Zhao, Liu and Zhang (2018), we get that PICASSO needs no more than

$$T(\epsilon_0, \beta, p) = O\left((\log p + \log \|\beta\|_{\infty})(\log p)^3 \{\log \log p + \log(\epsilon_0^{-1})\}\right).$$

But the above facts are true when $\epsilon_0 \leq \min\{1/8, C_3\}$. In order to extend the above results to a bigger range of ϵ_0 we first define $C_{\epsilon} := \min\{1/8, C_3\}/2$. If $\epsilon \leq C_{\epsilon}$, then one can get the same result by setting δ_K and τ_K appropriately as prescribed in (14). If $\epsilon > C_{\epsilon}$, then setting $\epsilon_0 = C_{\epsilon}$ and using (14), we again get $\|\hat{\beta}^{\text{picasso}} - \beta\|_2^2 \leq \epsilon_0 < \epsilon$ within $O((\log p + \log \|\beta\|_{\infty})(\log p)^3 \{\log \log p + \log(C_{\epsilon}^{-1}))$ iterations. Thus, conditions in Assumption 5.1 is met with

$$T(\epsilon, p, \beta) = O\left((\log p + \log \|\beta\|_{\infty})(\log p)^{3} \{\log \log p + \log(\epsilon^{-1} \vee C_{\epsilon}^{-1})\right),$$

and with probability at least $1 - O(p^{-0.5})$.

Proof for ETS-PGH:

Let $\hat{\beta}^{pgh}$ be the solution obtained by minimizing $f_{n_1}(\theta; \lambda, \mathcal{D}_1)$ via PGH method. For the proof of this part we will use Theorem 3.2 of Xiao and Zhang (2013). To apply that theorem we need to make sure that Assumption 3.2 of Xiao and Zhang (2013) is satisfied. To avoid notational confusion, we use $\tilde{\gamma}$ and $\tilde{\delta}$ to denote the parameters γ, δ' considered in Xiao and Zhang (2013) respectively. Let $\tilde{\delta} = 0.1$ and $\tilde{\gamma} = 2$. By a similar argument as before, it can be shown that with probability at least $1 - 2p^2$

$$\kappa(G, s_0) := \frac{\rho_+(s_0)}{\rho_-(s_0)} \le \frac{1+\nu}{1-\nu},$$

where $s_0 = \lfloor 46(1 + \tilde{\gamma})s \rfloor$, $\nu = 0.1$ and p is sufficiently large. Also, we choose

$$\lambda = 8 \max\{2, \frac{\tilde{\gamma} + 1}{\tilde{\gamma}(1 - \tilde{\delta}) - (1 + \tilde{\delta})}\}\{\log p/n_1\}^{1/2} \ge 4 \max\{2, \frac{\tilde{\gamma} + 1}{\tilde{\gamma}(1 - \tilde{\delta} - (1 + \tilde{\delta}))}\} \left\|X^{(1)^{\mathsf{T}}} E/n_1\right\|_{\infty}.$$
 (15)

The last inequality of the above display shows that $\lambda \ge 8 \|E^{\top} X^{(1)} / n_1\|_{\infty}$ and it happens with the probability at least $1 - O(p^{-0.5})$. Hence, following the arguments of the second bullet point on page 10 of

Xiao and Zhang (2013), we can conclude that the Assumption 3.2 of Xiao and Zhang (2013) holds with $\tilde{s} = \lfloor 22(1 + \tilde{\gamma})s \rfloor$, $\gamma_{\text{inc}} = 1.2$ (see Xiao and Zhang (2013)), $L_{\min} = 1.32$ and λ . Using part 3 of Theorem 3.2 in Xiao and Zhang (2013), for a given precision level ϵ_0 , we have

$$f_{n_1}(\widehat{\beta}^{\mathsf{pgh}};\lambda,\mathcal{D}_1) - f_{n_1}(\widehat{\beta}_L;\lambda,\mathcal{D}_1) \le O(\epsilon_0 s \sqrt{(\log p)/n_1}) < \epsilon_0/C_3, \quad \text{for large } p.$$

 C_3 is the same universal constant as in Proposition 5.5. If $\epsilon_0 = \min\{1, C_3\}/2 =: C_4$, then $\|\widehat{\beta}^{pgh} - \beta\|_2^2 \le \epsilon_0$, and the total iteration complexity is

$$O\left(\log p \log \log p + \log \max\{1, (\lambda^2/\epsilon_0^2) \log p\}\right),\$$

which for large value of of p, the form $O((\log p + \log \|\beta\|_{\infty}) \log \log p + \log \max\{1, (1/\epsilon_0^2) \log p\})$ (as $\lambda < 1$). If $\epsilon_0 \le C_4$, then the order becomes

$$O((\log p + \log \|\beta\|_{\infty}) \log \log p + \log(1/\epsilon_0)).$$

Otherwise, i.e., if $\epsilon_0 > C_4$ one can set the tolerance level at $\epsilon = C_4$ and the overall order in that case is $O((\log p + \log \|\beta\|_{\infty}) \log \log p + \log(1/C_4))$. Thus, for a given tolerance level ϵ , the total iteration complexity is $O((\log p + \log \|\beta\|_{\infty}) \log \log p + \log(\epsilon^{-1} \vee C_4^{-1}))$. Thus, Assumption 5.1 holds with probability at least $1 - O(p^{-0.5})$ and $T(\epsilon, p, \beta) = O((\log p + \log \|\beta\|_{\infty}) \log \log p + \log(\epsilon^{-1} \vee C_4^{-1}))$.

References

- AGARWAL, A., NEGAHBAN, S. and WAINWRIGHT, M. J. (2012). Fast global convergence of gradient methods for high-dimensional statistical recovery. Ann. Statist. 40 2452–2482. https://doi.org/10.1214/12-AOS1032
- CHERNOZHUKOV, V., CHETVERIKOV, D. and KATO, K. (2014). Gaussian approximation of suprema of empirical processes. *Ann. Statist.* **42** 1564–1597.
- FAN, J., SHAO, Q.-M. and ZHOU, W.-X. (2018). Are discoveries spurious? Distributions of maximum spurious correlations and their applications. Ann. Statist. 46 989–1017. https://doi.org/10.1214/17-AOS1575
- FLETCHER, A. K., RANGAN, S. and GOYAL, V. K. (2009). Necessary and sufficient conditions for sparsity pattern recovery. *IEEE Trans. Inform. Theory* 55 5758–5772. https://doi.org/10.1109/TIT.2009.2032726
- GORDON, R. D. (1941). Values of Mills' ratio of area to bounding ordinate and of the normal probability integral for large values of the argument. *Ann. Math. Statist.* **12** 364–366.
- JAIN, P., TEWARI, A. and KAR, P. (2014). On Iterative Hard Thresholding Methods for High-dimensional M-Estimation. In Advances in Neural Information Processing Systems 27. https://proceedings.neurips.cc/paper/ 2014/hash/218a0aefd1d1a4be65601cc6ddc1520e-Abstract.html.
- RUDELSON, M. and VERSHYNIN, R. (2013). Hanson-Wright inequality and sub-gaussian concentration. *Electron. Commun. Probab.* 18 1–9. https://doi.org/10.1214/ECP.v18-2865
- VERSHYNIN, R. (2018). High-Dimensional Probability: An Introduction with Applications in Data Science. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press. https://doi.org/10.1017/9781108231596
- WAINWRIGHT, M. J. (2009). Information-theoretic limits on sparsity recovery in the high-dimensional and noisy setting. *IEEE Trans. Inform. Theory* 55 5728–5741. https://doi.org/10.1109/TIT.2009.2032816
- WAINWRIGHT, M. J. (2019). High-dimensional statistics: a non-asymptotic viewpoint. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press. https://doi.org/10.1017/9781108627771
- WANG, W., WAINWRIGHT, M. J. and RAMCHANDRAN, K. (2010). Information-theoretic limits on sparse signal recovery: dense versus sparse measurement matrices. *IEEE Trans. Inform. Theory* 56 2967–2979. https://doi.org/10.1109/TIT.2010.2046199
- XIAO, L. and ZHANG, T. (2013). A proximal-gradient homotopy method for the sparse least-squares problem. SIAM J. Optim. 23 1062–1091. https://doi.org/10.1137/12086999
- ZHAO, T., LIU, H. and ZHANG, T. (2018). Pathwise coordinate optimization for sparse learning: algorithm and theory. Ann. Statist. 46 180–218. https://doi.org/10.1214/17-AOS1547