Electronic Journal of Statistics
Vol. 16 (2022) 5955–5979
ISSN: 1935-7524
https://doi.org/10.1214/22-EJS2088

Uncertainty quantification for robust variable selection and multiple testing

Eduard Belitser[1] **and Nurzhan Nurushev**[2]

[1] *VU Amsterdam*
e-mail: e.n.belitser@vu.nl

[2] *Rabobank*
e-mail: nurushevn@gmail.com

Abstract: We study the problem of identifying the set of *active* variables, termed in the literature as *variable selection* or *multiple hypothesis testing*, depending on the pursued criteria. For a general *robust setting* of non-normal, possibly dependent observations and a generalized notion of *active set*, we propose a procedure that is used simultaneously for the both tasks, variable selection and multiple testing. The procedure is based on the *risk hull minimization* method, but can also be obtained as a result of an empirical Bayes approach or a penalization strategy. We address its quality via various criteria: the Hamming risk, FDR, FPR, FWER, NDR, FNR, and various *multiple testing risks*, e.g., MTR=FDR+NDR; and discuss a weak optimality of our results. Finally, we introduce and study, for the first time, the *uncertainty quantification problem* in the variable selection and multiple testing context in our robust setting.

MSC2020 subject classifications: Primary 62C15.
Keywords and phrases: Multiple testing, robust setting, variable selection, uncertainty quantification.

Received December 2021.

1. Introduction

We are concerned with the problem of identifying the set of *active* (or *significant*) variables. This task appears in a wide variety of applied fields as genomics, functional MRI, neuro-imaging, astrophysics, among others. Such data is typically available on a large number of observation units, which may or may not contain a signal; the signal, when present, may be relatively faint and is dispersed across different observation units in an unknown fashion (i.e., the *sparsity pattern* is unknown to the observer). A prototypical application is GWAS (*genome-wide association studies*), where millions of genetic factors are examined for their potential influence on phenotypic traits. Although the number of tested genomic locations sometimes exceeds 10^5 or even 10^6, it is often believed that only a small set of genetic locations have tangible influences on the outcome of the disease or the trait of interest. This is well modeled by the stylized assumption of signal sparsity.

arXiv: 2109.09239

Depending on pursued criteria the problem is termed in the literature as either *variables selection*, (also termed as *recovery of the sparsity pattern*), or *multiple testing problem*. Commonly, the problems of variable selection and multiple testing are studied separately in the literature, although there are conceptual similarities and connections between them. In fact, a variable selection method determines the corresponding multiple testing procedure and vice versa, the difference lies merely in different criteria for inference procedures.

1.1. The observations model

Suppose we observe a high-dimensional ($\mathbb{R}^n$-valued) vector $X = (X_1, \ldots, X_n) \sim \mathbb{P}_\theta$ such that $(X-\theta)/\sigma$ satisfies some condition that in a way prevents too much dependence between coordinates of X (see Condition (A1) below), where $\theta = (\theta_1, \ldots, \theta_n) \in \mathbb{R}^n$ is an unknown high-dimensional *signal*. Actually, $\mathbb{P}_\theta = \mathbb{P}^n_{\theta,\sigma}$, but we will omit the dependence on n and σ in the sequel. In other words, we observe

$$X_i = \theta_i + \sigma\xi_i, \qquad i \in [n] \triangleq \{1, \ldots, n\}, \tag{1}$$

where $\xi = (\xi_1, \ldots, \xi_n) \triangleq (X-\theta)/\sigma$ is the "noise" vector and $\sigma > 0$ is the known "noise intensity". We emphasize that we pursue a general *distribution-free* setting: the distribution of ξ is arbitrary, satisfying only Condition (A1) below. The purpose of introducing σ is that certain extra information can be converted into a smaller noise intensity σ, a "more informative" model. For example, suppose we originally observed X_{ij} with $\mathbb{E}_\theta X_{ij} = \theta_i$ and $\mathrm{Var}_\theta(X_{ij}) = 1$, such that $(X_{ij}, j \in [m])$ are independent for each $i \in [n]$, for some $m = m_n \to \infty$ as $n \to \infty$. By taking $X_i = \frac{1}{m_n}\sum_{j=1}^{m_n} X_{ij}$, we obtain the model (1) with $\sigma^2 = \frac{1}{m_n} \to 0$ as $n \to \infty$. We are interested in non-asymptotic results, which imply asymptotic ones if needed. Possible asymptotic regime is high-dimensional setup $n \to \infty$, the leading case in the literature for high dimensional models. Another possible asymptotics is $\sigma \to 0$, accompanying $n \to \infty$ or on its own.

The general goal is (for now, loosely formulated) to select the *active* coordinates $I_*(\theta) \subseteq [n]$ of the signal θ, based on the data X. In the sequel, we will need to properly formalize the notion of *active set* $I_*(\theta)$. In particular, in this paper we let $I_*(\theta)$ be not necessarily the support of θ, $S(\theta) = \{i \in [n] : \theta_i \neq 0\}$. The main motivation for this is that we may want to qualify some relatively small (but non-zero) coordinates of θ as "inactive", with the threshold depending on the number of such coordinates. On the other hand, if the non-zero coordinates θ_i are allowed to be arbitrarily close (relative to the noise intensity σ) to zero, then it becomes impossible to recover the signal support. So, even when relaxing the notion of active set, there are still principal limitations as no method should be able to distinguish between $|\theta_i| \asymp \sigma$ and $\theta_i = 0$. These limitations will be quantified by establishing an appropriate lower bound. To make the problem feasible, one needs either to impose some kind of *strong signal condition* on θ (typically done in the literature on variable selection), or somehow adjust (relax) the criterion that measures the procedure quality (typically done in the literature on the multiple testing). For example, certain procedures can control more tolerant criteria like FDR or NDR without any condition, but, as we show below, their sum can be controlled again only under some strong signal condition.

1.2. Notation

Denote the probability measure of X from the model (1) by $\mathbb{P}_\theta$, and by $\mathbb{E}_\theta$ the corresponding expectation. For the notational simplicity, we skip the dependence on σ and n of these quantities and many others. Denote by $1\{s \in S\} = 1_S(s)$ the indicator function of the set S, by $|S|$ the cardinality of the set S, the difference of sets $S\backslash S_0 = \{s \in S : s \notin S_0\}$. Let $[k] = \{1, \ldots, k\}$, $k \in \mathbb{N}$. For $I \subseteq [n]$ define $I^c = [n]\backslash I$. For two nonnegative sequences (a_l) and (b_l), $a_l \lesssim b_l$ means $a_l \leq cb_l$ for all l (its range should be clear from the context) with some absolute $c > 0$, and $a_l \asymp b_l$ means that $a_l \lesssim b_l$ and $b_l \lesssim a_l$. The symbol $\triangleq$ will refer to equality by definition, $\Phi(x) = \mathbb{P}(Z \leq x)$ for $Z \sim \mathrm{N}(0,1)$. Throughout we assume the conventions: $|\varnothing| = 0$, $\sum_{I\in\varnothing} a_I = 0$ for any $a_I \in \mathbb{R}$ and $0\log(a/0) = 0$ (hence $(a/0)^0 = 1$) for any $a > 0$, in all the definitions whenever $0/0$ occurs we set by default $0/0 = 0$. Introduce the function $\ell(x) = \ell_q(x) \triangleq x\log(qn/x)$, $x, q > 0$, increasing in $x \in [0, n]$ for all $q \geq e$. Finally, introduce the ordered $\theta_1^2, \ldots, \theta_n^2$: $\theta_{[1]}^2 \geq \theta_{[2]}^2 \geq \ldots \geq \theta_{[n]}^2$, define additionally $\theta_{[0]}^2 = \infty$ and $\theta_{[n+1]}^2 = 0$.

1.3. Variable selection and multiple testing in the literature

The most studied situation in the literature is the particular case of (1):

$$X_i \overset{\text{ind}}{\sim} \mathrm{N}(\theta_i, \sigma^2), \quad i \in [n]; \quad I_*(\theta) = S(\theta) \triangleq \{i \in [n] : \theta_i \neq 0\}, \tag{2}$$

where the support $S(\theta)$ plays the role of active coordinates of θ, and θ is assumed to be *sparse* in the sense that $\theta \in \ell_0[s] = \{\theta \in \mathbb{R}^n : |S(\theta)| \leq s\}$ with $s = s_n = o(n)$ as $n \to \infty$.

Considering the situation (2) for now, there is a huge literature on the active set recovery problem studied from various perspectives. For an $I \subseteq [n]$,

$$\eta_I = (\eta_i(I), i \in [n]) = \big(1\{i \in I\}, i \in [n]\big) \in \{0,1\}^n$$

is the binary representation of $I \subseteq [n]$. For a measurable $\check{I} = \check{I}(X)$, let $\check{\eta} = \eta_{\check{I}}$ be some data dependent selector, which is supposed to estimate $\eta_* = \eta_*(\theta) = \eta_{I_*}$, for the "true" *active set* $I_* = I_*(\theta)$ which will be defined later on.

The historically first approach to assess the quality of $\check{I}$ is via the probability of wrong recovery $\mathbb{P}_\theta(\check{I} \neq I_*(\theta))$. Another approach is based on the Hamming distance between $\check{\eta}$ and η_* which is determined by the number of positions at which $\check{\eta}$ and η_* differ:

$$|\check{\eta} - \eta_*| \triangleq \sum_{i=1}^n |\check{\eta}_i - \eta_i(I_*)| = \sum_{i=1}^n 1\{\check{\eta}_i \neq \eta_i(I_*)\} = |\check{I}\backslash I_*| + |I_*\backslash \check{I}|.$$

Then a measure of the quality of $\check{I}$ is the expected Hamming loss called *Hamming risk*:

$$R_H(\check{I}, I_*) = \mathbb{E}_\theta|\eta_{\check{I}} - \eta_{I_*}| = \mathbb{E}_\theta|\check{I}\backslash I_*| + \mathbb{E}_\theta|I_*\backslash\check{I}| = R_{FP}(\check{I}, I_*) + R_{FN}(\check{I}, I_*),$$

where R_{FP} and R_{FN} are the *false positives* and *false negatives* terms (in a way, Type I and Type II errors), respectively. Note that $\mathrm{P}_\theta\big(\check{I} \neq I_*\big) = \mathrm{P}_\theta\big(|\eta_{\check{I}} - \eta_{I_*}| \geq 1\big) \leq \mathbb{E}_\theta|\eta_{\check{I}} - \eta_{I_*}| = R_H(\check{I}, I_*)$, which means that the approach based on the Hamming risk can imply results for $\mathrm{P}_\theta(\check{I} \neq I_*)$.

As is already well understood in many papers in related situations, in order to be able to recover $S(\theta)$, the non-zero signal $\theta_{S(\theta)} = (\theta_i, i \in S(\theta))$ has to satisfy some sort of *strong signal condition*. If $\theta \in \ell_0[s]$ with polynomial sparsity parametrization $s = n^\beta$, $\beta \in (0, 1)$ in the normal model (2), [12] and [2] (see further references therein) express this condition in the form of $\theta_i^2 \gtrsim \sigma^2 \log(n)$, $i \in S(\theta)$, for the signal to be detectable. Actually, [12] studied an idealized chi-squared model $Y_i \overset{\text{ind}}{\sim} \chi^2_\nu(\lambda_i)$, $i \in [n]$, where $\chi^2_\nu(\lambda_i)$ is a chi-square distributed random variable with ν degrees of freedom and non-centrality parameter λ_i. Squaring the both sides of (2), we arrive at the above chi-squared model with non-centrality parameters $\lambda_i = \theta_i^2$ and $\nu = 1$.

The case of polynomial sparsity $s = n^\beta$ has been well investigated, especially in the normal model (2), with active coordinates $\theta_i^2 \gtrsim \sigma^2 \log(n)$, $i \in S(\theta)$. The situation with arbitrary signal sparsity s is more complex: assuming $s = s_n \to \infty$ as $n \to \infty$, the right scaling for the active coordinates of $\theta \in \ell_0[s]$ becomes essentially $\theta_i^2 \gtrsim \sigma^2 \log(n/s)$, $i \in S(\theta)$. A recent important reference on this topic is [8], see also [18], [9]. More on this is in Section 4.

Inference on the active set can also be looked at from the multiple testing perspective. In classical multiple testing problem for the situation (2), one considers the following sequence of tests:

$$H_{0,i} : \theta_i = 0 \quad \text{versus} \quad H_{1,i} : \theta_i \neq 0, \quad i \in [n].$$

To connect to variable selection, notice that in multiple testing language, a variable selector $\check{I} = \check{I}(X) \subseteq [n]$ gives the multiple testing procedure which rejects the corresponding null hypothesis $H_{0,i}$, $i \in \check{I}$, whereas $I_*(\theta)$ encodes which null hypothesis do not hold.

Now we give some definitions for the multiple testing framework. The convention $0/0 = 0$ is used in below definitions. In multiple testing framework, one is typically interested in controlling (up to some prescribed level) of a type I error. The most popular one is the *False Discovery Rate* (FDR):

$$\mathrm{FDR}(\check{I}) = \mathrm{FDR}(\check{I}, I_*) = \mathbb{E}_\theta \tfrac{|\check{I}\backslash I_*|}{|\check{I}|},$$

the averaged proportion of errors among the selected variables. In multiple testing terminology, this is the expected ratio of incorrect rejections to total rejections. This criterion, introduced in [6], has become very popular because it is "tolerant" and "scalable" in the sense that the more rejections are possible, the more false positives are allowed. It also delivers an adaptive signal estimator, see [1]. Besides FDR, we study other known multiple testing criteria: *False Positive Rate* (FPR)

$$\mathrm{FPR}(\check{I}) = \mathrm{FPR}(\check{I}, I_*) = \mathbb{E}_\theta \tfrac{|\check{I} \backslash I_*|}{n-|I_*|},$$

Non-Discovery Rate (NDR)

$$\mathrm{NDR}(\check{I}) = \mathrm{NDR}(\check{I}, I_*) = \mathbb{E}_\theta \tfrac{|I_* \backslash \check{I}|}{|I_*|},$$

False Non-Discovery Rate (FNR) (introduced by [15])

$$\mathrm{FNR}(\check{I}) = \mathrm{FNR}(\check{I}, I_*) = \mathbb{E}_\theta \tfrac{|I_* \backslash \check{I}|}{n-|\check{I}|},$$

k-Family-Wise Error Rate (k-FWER) (introduced in [19]) and *k-Family-Wise Non-Discovery Rate* (k-FWNR)

$$\text{k-FWER}(\check{I}, I_*) = \mathbb{P}_\theta(|\check{I} \backslash I_*| \geq k), \quad \text{k-FWNR}(\check{I}, I_*) = \mathbb{P}_\theta(|I_* \backslash \check{I}| \geq k).$$

The k-FWER is the probability of rejecting at least k true null hypotheses. The case $k = 1$ reduces to the usual FWER. The FDR, FPR and k-FWER have the flavor of type I error as they deal with the control of false positives, NDR, FNR and k-FWNR have the flavor of type II error as they deal with the control of false negatives.

In the multiple testing setting, most of theoretical studies rely on the fact that the null distribution is exactly known. In practice, it is often unreasonable to assume this, instead, the null distribution is commonly (e.g., in genomics) implicitly defined as the “background noise” of the measurements, and is adjusted via some pre-processing steps. The issue of finding an appropriate null distribution has been popularized in a series of papers by Efron, (see [13]-[14] and further references therein), where the concept of *empirical null distribution* was introduced. A recent reference on this topic is [21], see also references therein. Our robust setting actually aligns well with the fact that the null hypothesis distribution is unknown, in fact, we avoid the problem of estimating the null distribution and obtain results that are robust over a certain (rather general) family of null distributions.

1.4. Multiple testing risk, strong signal condition again

In multiple testing setting, controlling type I error only is clearly not enough to characterize the quality of the procedures. For example, taking $\check{I} = \varnothing$ gives the perfect FDR control: $\mathrm{FDR}(\varnothing, I_*) = 0$, but this is of course an unreasonable procedure. One needs to control also some type II error, for example, NDR, or FNR. Again, if considered as the only criterion, the NDR can be easily controlled simply by taking another unreasonable procedure $\check{I} = [n]$, yielding

$\mathrm{NDR}([n], I_*) = 0$. It is thus instructive to control these errors together, see [2], [20], [22], [10] and further references in these papers.

Introduce the *multiple testing risks* (MTR) as all possible sums of the probabilities of type I and type II errors. The first MTR is the sum of the FDR and the NDR:

$$\mathrm{MTR}_1(\check{I}) = \mathrm{MTR}_1(\check{I}, I_*) = \mathrm{FDR}(\check{I}, I_*) + \mathrm{NDR}(\check{I}, I_*).$$

The other MTR's are defined similarly: they are always combination of type I and type II errors. Apart from MTR_1, these are

$$\mathrm{MTR}_2 = \mathrm{FDR} + \mathrm{FNR}, \quad \mathrm{MTR}_3 = \mathrm{FPR} + \mathrm{NDR}, \quad \mathrm{MTR}_4 = \mathrm{FPR} + \mathrm{FNR}.$$

Relating the MTR to the Hamming risk R_H, notice that, the both MTR and R_H, although being different, always combine some sort of Types I and II errors, in essence controlling the *false positives* and *false negatives* simultaneously. Clearly, the MTR is a more mild criterion as it is always a proportion: $\mathrm{MTR} \ll R_H$.

It is desirable that MTR is as small as possible, e.g., converging to zero as $n \to \infty$. However, as is shown in related settings in [18], [2], [20], [22], [8], [10], this is in general impossible, which is not surprising because the same kind of principal limitation occurs for the Hamming risk in the case of variable selection. Again some sort of *strong signal condition* is unavoidable. The Hamming risk R_H is a more severe quality measure than MTR, so its convergence to zero should occur under a more severe strong/sparse signal condition. In some papers in related settings this difference is referred to as *exact* and *approximate* recovery of the active set. Yet another type of recovery, the so called *almost full recovery* is studied in [8], this is the convergence of $R_H/|I_*|$ to zero.

1.5. The scope

In this paper, we generalize the standard normal setting (2) to the more general setting (1) in that we pursue the *robust* inference in the sense that the distribution of the error vector ξ is unknown, but assumed to satisfy only certain condition, Condition (A1) below. In particular, the ξ_i's can be non-normal, not identically distributed, of non-zero mean, and even dependent; their distribution may depend on θ.

Next, we generalize the notion of active set $I_*(\theta)$ that is now not necessarily the sparsity class $\ell_0(s)$ and not necessarily with all large non-zero coordinates. Basically, for the problem of determining the active set to be well defined, the parameter θ has to possess a set $I_*(\theta)$ of *distinct* active coordinates, which is ensured by the condition $\theta \in \Theta(K)$. We elaborate on the meaning of $\Theta(K)$ in Section 3.2. Then the *sparsity* is expressed by $|I_*(\theta)|$ and the *strong signal* condition by the fact that $\theta \in \Theta(K)$ for some (sufficiently large) $K > 0$. The "goodness"' of θ is determined by the combination of the two components: how *sparse* and how *strong* the signal θ is. By looking at this combination, we exposes the so called *phase transition* effect, separating the impossibility and possibility to recover the active set, discussed in Section 4.2.

Finally, in this paper we address the new problem of *uncertainty quantification* (UQ) for the active set $\eta_* = \eta_*(\theta) = \eta(I_*(\theta))$, this is to be distinguished from the uncertainty quantification for the parameter θ. For the Hamming loss $|\cdot|$ on $\{0,1\}^n$, a confidence ball is $B(\hat{\eta},\hat{r}) = \{\eta \in \{0,1\}^n : |\hat{\eta}-\eta| \le \hat{r}\}$, where the center $\hat{\eta} = \hat{\eta}(X) : \mathbb{R}^n \mapsto \{0,1\}^n$ and radius $\hat{r} = \hat{r}(X) : \mathbb{R}^n \mapsto \mathbb{R}_+ = [0,+\infty]$ are measurable functions of the data X. The goal is to construct such a confidence ball $B(\hat{\eta}, C\hat{r})$ that for any $\alpha_1, \alpha_2 \in (0,1]$ and some function $r(\eta_*(\theta))$, $r : \mathbb{R}^n \to \mathbb{R}_+$, there exist $C, c > 0$ such that

$$\sup_{\theta\in\Theta_{\mathrm{cov}}} \mathbb{P}_\theta\big(\eta_*(\theta) \notin B(\hat{\eta}, C\hat{r})\big) \le \alpha_1, \qquad \sup_{\theta\in\Theta_{\mathrm{size}}} \mathbb{P}_\theta\big(\hat{r} \ge c r(\eta_*(\theta))\big) \le \alpha_2, \tag{3}$$

for some $\Theta_{\mathrm{cov}}, \Theta_{\mathrm{size}} \subseteq \mathbb{R}^n$. The function $r(\eta_*(\theta))$, called the *radial rate*, is a benchmark for the effective radius of the confidence ball $B(\hat{\eta}, C\hat{r})$. The first expression in (3) is called *coverage relation* and the second *size relation*. To the best of our knowledge, there are no results (3) on uncertainty quantification with the Hamming loss for the active set $\eta_*(\theta)$. It is desirable to find the smallest $r(\eta_*(\theta))$, the biggest Θ_{cov} and Θ_{size} such that (3) holds and $r(\eta_*(\theta)) \asymp R(\Theta_{\mathrm{size}})$, where $R(\Theta_{\mathrm{size}})$ is the optimal rate in estimation problem for $\eta_*(\theta)$. We derive some UQ results for the proposed selector $\hat{I}$.

Typically, the so called *deceptiveness* issue is pertinent to the UQ problem, meaning that the confidence set of the optimal size and high coverage can only be constructed for non-deceptive parameters (in particular, Θ_{cov} cannot be the whole set $\mathbb{R}^n$). Being non-deceptive is expressed by imposing some condition on the parameter; for example, the EBR (excessive bias restriction) condition $\Theta_{\mathrm{cov}} = \Theta_{\mathrm{ebr}} \subset \mathbb{R}^n$, see [3]–[5]. Interestingly, there is no deceptiveness issue as such for our UQ problem. An intuition behind this is as follows: the problem of active set recovery is more difficult than the UQ-problem in a sense that solving the former problem implies solving the latter. Then the condition $\theta \in \Theta(K)$ for the parameter to have distinct active coordinates implies also that the parameter is non-deceptive. In our case, we will have $\Theta_{\mathrm{cov}} = \Theta_{\mathrm{size}} = \Theta(K)$ for some $K > 0$.

1.6. Organization of the rest of the paper

In Section 2 we describe the criteria and procedures for variable selection and multiple testing, and introduce the generalized notion of active coordinates. In Section 3 we present the main results of the paper. In Section 4 we discuss a weak optimality of our results and a phase transition effect. The proofs of the theorems are collected in Section 5, which, despite generality of the setting and results, we could keep completely self-contained and relatively compact.

2. Preliminaries

2.1. Criterion for selecting active variables

Consider for the moment the estimation problem of θ when we use the projection estimators $\hat{\theta}(I) = (X_i 1\{i \in I\}, i \in [n])$, $I \subseteq [n]$. The quadratic loss of $\hat{\theta}(I)$ gives a theoretical criterion

$$\mathcal{C}_1^{\mathrm{th}}(I) = \|\hat{\theta}(I) - \theta\|^2 = \sum_{i \in I^c} \theta_i^2 + \sigma^2 \sum_{i \in I} \xi_i^2.$$

The best choice of I would be the one minimizing $\mathcal{C}_1^{\mathrm{th}}(I)$, which is the same as minimizing $\mathcal{C}_1^{\mathrm{th}}(I) + n\sigma^2$. However, neither θ nor ξ are observed. Substituting unbiased estimator $X_i^2 - \sigma^2$ instead of θ_i^2, $i \in I^c$, leads to the quantity

$$\mathcal{C}_2^{\mathrm{th}}(I) = \sum_{i \in I^c} X_i^2 + \sigma^2 |I| + \sigma^2 \sum_{i \in I} \xi_i^2 = \|X - \hat{\theta}(I)\|^2 + \sigma^2 |I| + \sigma^2 \sum_{i \in I} \xi_i^2$$

to minimize with respect to $I \subseteq [n]$, however still not usable in view of the term $\sigma^2 \sum_{i \in I} \xi_i^2$. If instead of $\sigma^2 \sum_{i \in I} \xi_i^2$, we use its expectation $\sigma^2 |I|$, we arrive essentially at Mallows's C_p-criterion (and AIC in the normal case)

$$\mathcal{C}_{\mathrm{Mallows}}(I) = \|X - \hat{\theta}(I)\|^2 + 2\sigma^2 |I|.$$

However, it is well known that the C_p-criterion leads to overfitting. An intuitive explanation is that using the expectation $\sigma^2 |I|$ as penalty in $\mathcal{C}_{\mathrm{Mallows}}(I)$ is too optimistic to control oscillations of its stochastic counterpart $\sigma^2 \sum_{i \in I} \xi_i^2$.

The next idea is to use some quantity $p(I)$ (instead of $|I|$) that majorizes $\sum_{i \in I} \xi_i^2$ in the more strict sense that for some $K, H_0, \alpha > 0$ and all $M \geq 0$

$$\mathbb{P}_\theta\Big(\sup_{I \subseteq [n]} \big(\sum_{i \in I} \xi_i^2 - K p(I) \big) \geq M \Big) \leq H_0 e^{-\alpha M}. \tag{4}$$

Remark 1. This idea is borrowed from the *risk hull minimization method* developed by Golubev in several papers; see [11] and references therein.

Thus, in view (4), using $K\sigma^2 p(I)$ instead of $\sigma^2 \sum_{i \in I} \xi_i^2$, we obtain a more adequate criterion $\mathcal{C}_3(I) = \|X - \hat{\theta}(I)\|^2 + \sigma^2 |I| + K\sigma^2 p(I)$. Since typically $|I| \lesssim p(I)$, the second term can be absorbed into the third, and we finally derive the criterion

$$\mathcal{C}(I) = \|X - \hat{\theta}(I)\|^2 + K\sigma^2 p(I), \tag{5}$$

for sufficiently large constant $K > 0$. It remains to determine $p(I)$, preferably smallest possible, for which (4) holds. We state a condition used throughout.

Condition (A1). For some $p_0(I)$ such that $p_0(I) \leq C_\xi |I|$, for some positive C_ξ, H_ξ, α_ξ and any $M \geq 0$,

$$\sup_{\theta \in \mathbb{R}^n} \mathbb{P}_\theta\Big(\sum_{i \in I} \xi_i^2 \geq p_0(I) + M \Big) \leq H_\xi e^{-\alpha_\xi M}, \quad I \subseteq [n]. \tag{A1}$$

If the distribution of ξ does not depend on θ (in some important specific cases), there is no supremum over $\theta \in \mathbb{R}^n$. For independent ξ_i's, (A1) holds with $p_0(I) \asymp |I|$, so that indeed $p_0(I) \le C_\xi |I|$, $I \subseteq [n]$, for some $C_\xi > 0$.

For fixed $C_\xi, H_\xi, \alpha_\xi > 0$, Condition (A1) with $p_0(I) = C_\xi |I|$ describes a class of possible measures $\mathbb{P}_\theta$ from (1), in the sequel denoted by $\mathcal{P}_A$:

$$\mathcal{P}_A = \{\mathbb{P}_\theta : \mathbb{P}_\theta \text{ satisfies (A1)}\}. \tag{6}$$

Remark 2. Condition (A1) is of course satisfied for independent normals $\xi_i \overset{\text{ind}}{\sim} \mathrm{N}(0,1)$ and for bounded (arbitrarily dependent) ξ_i's. Recall that the ξ_i's are not necessarily of zero mean, but for normals Condition (A1) is the weakest if $\mathbb{E}_\theta \xi_i = 0$. In a way, Condition (A1) prevents too much dependence, but it still allows some interesting cases of dependent ξ_i's. For example, one can show (in the same way as in [5]) that this conditions is fulfilled for ξ_i's that follow an autoregressive model AR(1) with normal white noise.

Let $\eta = \sup_{I\subseteq[n]} \big(\sum_{i\in I} \xi_i^2 - C_\xi |I| - \alpha_\xi^{-1}\big[|I| + \log\binom{n}{|I|}\big]\big)$. By (A1) and the union bound, it is not difficult to derive

$$\begin{aligned}
\mathbb{P}_\theta(\eta \ge M) &\le \sum_{I\subseteq[n]} \mathbb{P}_\theta\Big(\sum_{i\in I} \xi_i^2 \ge C_\xi |I| + \alpha_\xi^{-1}\big[|I| + \log\tbinom{n}{|I|}\big] + M\Big) \\
&\le H_\xi e^{-\alpha_\xi M} \sum_{I\subseteq[n]} e^{-|I|}\tbinom{n}{|I|}^{-1} = H_\xi e^{-\alpha_\xi M} \sum_{k=0}^{n} \sum_{I\subseteq[n]:|I|=k} e^{-k}\tbinom{n}{k}^{-1} \\
&= H_\xi e^{-\alpha_\xi M} \sum_{k=0}^{n} e^{-k} \le \tfrac{H_\xi e}{e-1} e^{-\alpha_\xi M} = H_0 e^{-\alpha_\xi M}.
\end{aligned}$$

where $H_0 = H_\xi/(1-e^{-1})$. As $\binom{n}{k} \le (\frac{en}{k})^k$, $k \in [n]$, we have

$$C_\xi |I| + \alpha_\xi^{-1}\big[|I| + \log\tbinom{n}{|I|}\big] \le (C_\xi + \alpha_\xi^{-1})|I| + \alpha_\xi^{-1}|I| \log\tfrac{en}{|I|} = \alpha_\xi^{-1}|I|\log\big(\tfrac{q_\xi n}{|I|}\big),$$

where $q_\xi = e^{C_\xi \alpha_\xi + 2}$. The last two displays imply the following relation (that we will need later): for appropriate $M_\xi > 0$ and $q = e^2$

$$\sup_{\theta\in\mathbb{R}^n} \sum_{I\subseteq[n]} \mathbb{P}_\theta\Big(\sum_{i\in I} \xi_i^2 \ge M_\xi |I| \log\big(\tfrac{qn}{|I|}\big) + M\Big) \le H_0 e^{-\alpha_\xi M}. \tag{7}$$

Although we will use the relation (7) only for $q = e^2$, it is also implied by (A1) for any $q > 1$ with appropriately chosen $M_\xi = M_\xi(q)$, certainly if $M_\xi(q) \ge \alpha_\xi^{-1}\big(\frac{\log q_\xi}{\log q} + 1\big)$.

Thus, we obtain that, under (A1), the criterion (4) is satisfied with $p(I) = \ell_q(|I|) = |I| \log\big(\frac{qn}{|I|}\big)$. According to (5), this motivates the definition of the so called *preselector*

$$\tilde{I} = \tilde{I}(K) = \arg\min_{I\subseteq[n]} \Big\{ \sum_{i\in I^c} X_i^2 + K\sigma^2 p(I)\Big\}, \tag{8}$$

where $p(I) = p_q(I) \triangleq \ell_q(|I|) = |I| \log\big(\frac{qn}{|I|}\big)$, for some $K > 0$ and $q = e^2$. If $\tilde{I}$ is not unique, take, say, the one with the biggest sum $\sum_{i\in\tilde{I}}(n-i)$.

Actually, an interesting interplay between constants K and q is possible, making certain constant in the proofs sharper. But we fix the second constant $q = e^2$ in (8) for the sake of mathematical succinctness.

Remark 3. Since $\ell'(x) = \log(\frac{en}{x})$, then for $1 \le x \le y \le n$, $(y-x)\log(\frac{qn}{ey}) \le \ell(y) - \ell(x) \le (y-x)\log(\frac{qn}{ex})$. By the definition (8) of $\tilde{I}$, we have that for each $i \in \tilde{I}$, with $|\tilde{I}| \ge 1$,

$$X_i^2 \ge K\sigma^2\big[|\tilde{I}|\log(\tfrac{qn}{|\tilde{I}|}) - |\tilde{I}-1|\log\big(\tfrac{qn}{|\tilde{I}|-1}\big)\big] \ge K\sigma^2\log\big(\tfrac{qn}{e|\tilde{I}|}\big),$$

and, similarly, for each $i' \in \tilde{I}^c$ with $|\tilde{I}| \le n-1$,

$$X_{i'}^2 \le K\sigma^2\big[(|\tilde{I}|+1)\log(\tfrac{qn}{|\tilde{I}|+1}) - |\tilde{I}|\log\big(\tfrac{qn}{|\tilde{I}|}\big)\big] \le K\sigma^2\log\big(\tfrac{qn}{e|\tilde{I}|}\big).$$

Next, define the *selector* $\hat{\eta} = (\hat{\eta}_i, i \in [n])$ (and respectively $\hat{I}$) of significant coordinates as

$$\hat{\eta}_i = \hat{\eta}_i(K) = 1\big\{X_i^2 \ge K\sigma^2\log\big(\tfrac{qn}{|\tilde{I}|}\big)\big\}, \quad \hat{I} = \hat{I}(K) = \{i \in [n] : \hat{\eta}_i = 1\}. \tag{9}$$

Notice that the selector is always a subset of the preselector: $\hat{I}(K) \subseteq \tilde{I}(K)$.

Remark 4. As we have already mentioned, this procedure can be related to the *risk hull minimization* (RHM) method developed by Golubev. Almost the same procedures can be derived as a result of the empirical Bayes approach with appropriately chosen prior, or as a result of the penalization strategy with appropriately chosen penalty, see [5]. It is interesting that several methodologies deliver akin procedures.

2.2. The notion of general active set

Suppose we consider an arbitrary θ and would still like somehow divide all the entries of θ into the groups of *active* and *inactive* coordinates. The support set $S(\theta) = \{i \in [n] : \theta_i \neq 0\}$ as active group and the rest as inactive is one traditional way of doing this. The idea is to soften this so that smallish, but nonzero, coordinates θ_i be possibly assigned to the inactive group.

For $\theta \in \mathbb{R}^n$ and $A \ge 0$, define the *active set* $I_*(A,\theta) = I_*(A,\theta,\sigma^2)$ as follows: with $q = e^2$,

$$I_*(A,\theta) = \arg\min_{I\subseteq[n]} r^2(I,\theta), \quad \text{where} \quad r^2(I,\theta) \triangleq \sum_{i\in I^c}\theta_i^2 + A\sigma^2|I|\log(\tfrac{qn}{|I|}) \tag{10}$$

and $I_*(A,\theta)$ is with the smallest sum $\sum_{i\in I_*(A,\theta)} i$ if the minimum is not unique. Denote for brevity $I_* = I_*(A,\theta)$. By the definition (10),

$$r^2(\theta) \triangleq r^2(I_*,\theta) \le r^2(I,\theta) \quad \text{for any} \quad I \subseteq [n].$$

The last relation and the same argument as in Remark 3 lead to the following two relations. If $i \in I_*$ with $|I_*| \geq 1$, then

$$\theta_i^2 \geq A\sigma^2\big[|I_*|\log(\tfrac{qn}{|I_*|}) - (|I_*|-1)\log(\tfrac{qn}{|I_*|-1})\big)\big] \geq A\sigma^2\log(\tfrac{en}{|I_*|}). \tag{11}$$

Conversely, if $i \in I_*^c$ with $|I_*| \leq n-1$, then

$$\theta_i^2 \leq A\sigma^2\big[(|I_*|+1)\log(\tfrac{qn}{|I_*|+1}) - |I_*|\log(\tfrac{qn}{|I_*|})\big)\big] \leq A\sigma^2\log\big(\tfrac{en}{|I_*|\vee 1}\big). \tag{12}$$

Also notice that I_* depends on the product $A\sigma^2$ rather than just on A. If we consider the asymptotic regime $\sigma^2 \to 0$, it is instructive to fix the product $A\sigma^2$ so that $A \to \infty$, which can be interpreted as if θ satisfies more and more stringent strong signal condition.

Remark 5. To get an idea what $I_*(A,\theta)$ means, suppose in (10) we had $A\sigma^2|I|\log(qn)$ instead of $A\sigma^2|I|\log(\frac{qn}{|I|})$. Then active coordinates would have been all $i \in [n]$ corresponding to "large" $\theta_i^2 \geq A\sigma^2\log(qn)$. The definition (10) does kind of the same, but the requirement for being active becomes slightly more lenient if there are more "large" coordinates. The function $x\log(qn/x)$ is increasing (in fact, for all $q \geq e$) for $x \in [1,n]$ slightly slower than $x\log(qn)$, creating the effect of "borrowing strength" via the number of active coordinates: the more such coordinates, the less stringent the property of being active becomes.

The family $\mathcal{I} = \mathcal{I}(\theta) = \{I_k^{\mathrm{vsp}}(\theta),\, k = 0,1,\ldots,n\}$, with $I_k^{\mathrm{vsp}}(\theta) = \{i \in [n] : \theta_i^2 \geq \theta_{[k]}^2\}$, is called *variable selection path* (VSP). It consists of at most $n+1$ embedded sets:

$$\varnothing = I_0^{\mathrm{vsp}}(\theta) \subseteq I_1^{\mathrm{vsp}}(\theta) \subseteq \ldots \subseteq I_n^{\mathrm{vsp}}(\theta) \subseteq I_{n+1}^{\mathrm{vsp}} \triangleq [n].$$

Clearly, $I_*(A,\theta) = I^{\mathrm{vsp}}_{|I_*(A,\theta)|}(\theta)$, $\theta \in \mathbb{R}^n$. If some of $\theta_{[k]}$ coincide, the corresponding sets $I_k^{\mathrm{vsp}}(\theta)$ in the variable selection path $\mathcal{I}$ merge. Accounting for these merges, notice that the true support $S(\theta) = I_n^{\mathrm{vsp}}(\theta)$ is the last set in the variable selection path $\mathcal{I}(\theta)$.

Remark 6. We state some further properties of the active set I_* and the variable selection path $\mathcal{I}$.

(a) The family $\{I_*(A,\theta), A \geq 0\}$ reproduces the variable selection path $\mathcal{I}$

$$\{I_*(A,\theta), A \geq 0\} = \mathcal{I}(\theta).$$

(b) For any $0 \leq A_1 \leq A_2$ and any $\theta \in \mathbb{R}^n$, we have

$$\varnothing \subseteq I_*(A_2,\theta) \subseteq I_*(A_1,\theta) \subseteq S(\theta) \subseteq [n].$$

(c) In view of (11) and (12), we have that, if $\theta_i^2 \geq A\sigma^2\log(qn/|S(\theta)|)$ for all $i \in S(\theta)$ and some $A > 0$, then $I_*(A',\theta) = S(\theta)$ for all $A' \leq A$.

3. Main results

In this section we present the main results.

3.1. Control of the preselector

First, we establish the results on over-size and under-size control of the preselector $\tilde{I} = \tilde{I}(K)$ defined by (8). Recall $\ell(x) = x\log(\frac{qn}{x})$, $x \geq 0$, $q = e^2$.

Theorem 1. *Let $\tilde{I} = \tilde{I}(K)$ be defined by (8), $I_*(A,\theta)$ be defined by (10), let H_0 be from (7). Then for any sufficiently large K_0 (e.g., $K_0 > 2M_\xi$) there exist $A_0 > 0$ and constants $M_0, \alpha_0 > 0$ such that for any $M \geq 0$,*

$$\sup_{\theta\in\mathbb{R}^n} e^{\alpha_0\ell(|I_*(A_0,\theta)|)}\mathbb{P}_\theta\big(|\tilde{I}(K_0)| \geq M_1|I_*(A_0,\theta)| + M\big) \leq H_0 e^{-\alpha_0 M/2}. \tag{i}$$

For any $K_1 > 0$, $\delta \in [0,1)$, there exist $A_1, \alpha_1, \alpha_1' > 0$ (depending on δ, K_1) such that for any $M \geq 0$

$$\sup_{\theta\in\mathbb{R}^n} e^{\alpha_1'\ell(|I_*(A_1,\theta)|)}\mathbb{P}_\theta\big(|\tilde{I}(K_1)| \leq \delta|I_*(A_1,\theta)| - M\big) \leq H_0 e^{-\alpha_1 M}. \tag{ii}$$

Remark 7. We can obtain other formulations of properties (i) and (ii).

- (a) Property (i): for any A_0 there exist (sufficiently large) K_0 and $M_1, \alpha_0 > 0$ (depending on A_0) such that for any $M \geq 0$, (i) holds.
- (b) Property (ii): for any $A_1 > 0$ there exist (sufficiently small) $K_1 > 0$ and $\delta \in [0,1]$ such that for any $M \geq 0$, (ii) holds.

Suppose we fix some sufficiently large $K > 0$ such that there exists A_0 for which (i) is fulfilled and consider the corresponding preliminary selector $\tilde{I}(K)$. For this K, the properties (i) and (ii) of Theorem 1 provide *separately* over-size and under-size control of $\tilde{I}$, with some $A_0(K)$ and $A_1(K)$, respectively. By analyzing the proof, we see that always $A_0(K) \leq A_1(K)$. Hence $I_*(A_1,\theta) \subseteq I_*(A_0,\theta)$ (as it should be), forming a *shell* $I_*(A_0,\theta)\backslash I_*(A_1,\theta)$.

3.2. Set of signals with distinct active coordinates

In the light of Theorem 1, *uniform* constants $A_0(K) \leq A_1(K)$ exist such that (i) and (ii) are fulfilled. We can say informally that $\tilde{I}$ "lives in an inflated shell" between $I_*(A_1(K),\theta)$ and $I_*(A_0(K),\theta)$, which can be thought of as *indifference zone* for the selector $\tilde{I} = \tilde{I}(K)$. For a general $\theta \in \mathbb{R}^n$, the sets $I_*(A_1(K),\theta) \subseteq I_*(A_0(K),\theta)$ can be far apart and the selector may have too much room to vary, so that this θ can be thought of as not having distinct active and inactive coordinates. An idea is to determine a set of θ's for which active coordinates are "identifiable".

This motivates the following definition of the *set* $\Theta(K)$ *of signals with distinctive active and inactive coordinates* as follows:

$$\Theta(K) = \{\theta \in \mathbb{R}^n : I_*(A_1(K),\theta) = I_*(A_0(K),\theta)\}. \tag{13}$$

In what follows, denote for brevity $I_* = I_*(A_1(K), \theta)$ and $I^* = I_*(A_0(K), \theta)$.

The quantities I_* and I^* depend on θ. The constants $A_0(K)$ and $A_1(K)$, not depending on θ, exist in view of Theorem 1. Hence, imposing $I_* = I^*$ determines a subset of $\mathbb{R}^n$ for which we can provide simultaneous control for oversizing and undersizing of $\tilde{I}$ by Theorem 1, i.e., there is no indifference zone. This heuristically means that $\Theta(K)$ describes a set of signals with distinctive active and inactive coordinates, so that the active coordinates are well defined.

The above definition (13) is still somewhat implicit, so we elucidate its meaning. We have $\Theta(K) = \cup_{I \in \mathcal{I}} \Theta_I(K)$, where

$$\Theta_I(K) = \Big\{ \theta \in \mathbb{R}^n : \theta_{[i]}^2 \geq A_1(K)\sigma^2 \log(\tfrac{en}{|I|}), i = 1, \dots, |I|;$$
$$\theta_{[j]}^2 \leq A_0(K)\sigma^2 \log(\tfrac{en}{|I|}), j = |I|+1, \dots, n \Big\},$$

$\theta_{[i]}^2$'s are ordered θ_i^2's: $\theta_{[1]}^2 \geq \ldots \geq \theta_{[n]}^2$. In words, $\theta \in \Theta_I(K)$ is such that $|I|$ biggest values of θ_i^2's are all bigger than $A_1(K)\sigma^2 \log(\frac{en}{|I|})$ and others do not exceed $A_0(K)\sigma^2 \log(\frac{en}{|I|})$. Then $I_* = I^* = I$ for this θ. Basically, there must be a sufficient "gap" between the active group I of coordinates and the nonactive one I^c. The magnitude of this gap depends on the size of the active group: the smaller the active group, the bigger the gap must be.

In particular, this notion generalizes the traditional strong signal requirement. Indeed, if $\theta_i^2 \geq A_1(K)\sigma^2 \log(\frac{qn}{|S(\theta)|})$ for $i \in S(\theta)$, then $I_* = I_*(A_1(K), \theta) = S(\theta) = I_*(A_0(K), \theta) = I^*$, so that $\theta \in \Theta(K)$. In fact, $I_*(A, \theta) = S(\theta)$ for all $A \in [0, A_1(K)]$.

The constants $A_0(K), A_1(K)$ evaluated in the proof of Theorem 1 are uniform over $\mathbb{P}_\theta \in \mathcal{P}_A$ defined by (6). These constants are of course far from being optimal as we use relatively rough bounds in the proof. We can extend the set $\Theta(K)$ by insisting on the best possible choices of those constants. Precisely, fix some $K, M_1, \delta > 0$, and define $A_0'(K)$ to be the biggest constant for which (i) holds with some $M_1' \leq M_1$ and $A_1'(K)$ be the smallest constant for which (ii) is fulfilled with some $\delta' \geq \delta$. Then the resulting set

$$\Theta''(K) = \{\theta \in \mathbb{R}^n : I_*(A_1'(K), \theta) = I_*(A_0'(K), \theta)\}$$

can be used instead of $\Theta(K)$ in all relevant claims, while $\Theta(K) \subseteq \Theta''(K)$.

On the other hand, for any $A_0' \leq A_0(K)$ and $A_1' \geq A_1(K)$,

$$\Theta'(A_0', A_1') = \{\theta \in \mathbb{R}^n : I_*(A_1', \theta) = I_*(A_0', \theta)\} \subseteq \Theta(K). \tag{14}$$

For a particular $\theta \in \Theta(K)$, there may be $A_0' < A_0(K)$ and $A_1' > A_1(K)$ such that $\theta \in \Theta'(A_0', A_1')$. The parameters A_0', A_1 reflect the strength of signal $\theta \in \Theta'(A_0', A_1')$: the bigger the difference between $A_0' < A_1'$, the stronger the signal θ. This yields yet another interpretation of the set $\Theta(K)$ as the "minimal" signal strength condition on $\theta \in \Theta(K)$. In particular, in view of property (c) from Remark 6, it follows that for $\theta \in \Theta'(0, A_1(K))$

$$I_*(A_1(K), \theta) = S(\theta) \quad \text{and} \quad \Theta'(0, A_1(K)) \subseteq \Theta(K). \tag{15}$$

The main role of $\Theta'(A'_0, A'_1)$ is to show the transition from the "hard" definition $S(\theta)$ of active set towards the "soft" definition $I_*(A_1(K), \theta)$.

3.3. Results on active set recovery and multiple testing

Now we establish the control of all the introduced quality measures (FPR, NDR, the Hamming rate, etc.) for the proposed active set selector $\hat{I}$.

Theorem 2. *Let $\hat{\eta}$ and $\hat{I} = \hat{I}(K)$ be defined by (9); $I^* = I_*(A_0(K), \theta)$, $I_* = I_*(A_1(K), \theta)$ and $\Theta(K)$ be defined by (13) for sufficiently large $K > 0$, and let $\eta_* = \eta_{I_*}$. Then there exist constants $H_1, H_2, \alpha_2, \alpha_3 > 0$ such that, uniformly in $\theta \in \mathbb{R}^n$,*

$$\mathrm{FPR}(\hat{I}, I^*) = \mathbb{E}_\theta \tfrac{|\hat{I} \backslash I^*|}{n - |I^*|} \leq H_1 (\tfrac{n}{|I^*| \vee 1})^{-\alpha_2}, \tag{16}$$

$$\mathrm{NDR}(\hat{I}, I_*) = \mathbb{E}_\theta \tfrac{|I_* \backslash \hat{I}|}{|I_*|} \leq H_2 (\tfrac{n}{|I_*|})^{-\alpha_3}. \tag{17}$$

Moreover, there exist $H_3, \alpha_4 > 0$ such that, uniformly in $\theta \in \Theta(K)$,

$$R_H(\hat{I}, I_*) = \mathbb{E}_\theta |\hat{\eta} - \eta_*| = \mathbb{E}_\theta \big(|\hat{I} \backslash I_*| + |I_* \backslash \hat{I}| \big) \leq H_3 n (\tfrac{n}{|I_*| \vee 1})^{-\alpha_4}. \tag{18}$$

From the above theorem, the next corollary follows immediately. It describes control of k-FWER, k-FWNR, the probability of wrong discovery, and the so called *almost full recovery* (relation (19) below).

Corollary 1. *For any $k = 1, \ldots, n$ and uniformly in $\theta \in \mathbb{R}^n$,*

$$\text{k-FWER}(\hat{I}, I^*) = \mathbb{P}_\theta (|\hat{I} \backslash I^*| \geq k) \leq \tfrac{H_1}{k} (n - I^*) (\tfrac{n}{|I^*| \vee 1})^{-\alpha_2},$$
$$\text{k-FWNR}(\hat{I}, I_*) = \mathbb{P}_\theta (|I_* \backslash \hat{I}| \geq k) \leq \tfrac{H_2}{k} I_* (\tfrac{n}{|I_*|})^{-\alpha_3}.$$

Uniformly in $\theta \in \Theta(K)$, $\mathbb{P}_\theta(\hat{I} \neq I_*) \leq H_3 n (\frac{n}{|I_*|})^{-\alpha_4}$ and

$$\tfrac{R_H(\hat{I}, I_*)}{|I_*|} = \tfrac{1}{|I_*|} \mathbb{E}_\theta \big(|\hat{I} \backslash I_*| + |I_* \backslash \hat{I}| \big) \leq H_3 (\tfrac{n}{|I_*| \vee 1})^{-(\alpha_4 - 1)}. \tag{19}$$

The following result establishes the control of $\mathrm{FDR}(\hat{I})$ and $\mathrm{FNR}(\hat{I})$.

Theorem 3. *With the same notation as in Theorem 2, there exist constants $H_5, H_6, \alpha_5, \alpha_6 > 0$ such that, uniformly in $\theta \in \Theta(K)$,*

$$\mathrm{FDR}(\hat{I}, I_*) = \mathbb{E}_\theta \tfrac{|\hat{I} \backslash I_*|}{|\hat{I}|} \leq H_5 (\tfrac{n}{|I_*| \vee 1})^{-\alpha_5}, \tag{20}$$

$$\mathrm{FNR}(\hat{I}, I_*) = \mathbb{E}_\theta \tfrac{|I_* \backslash \hat{I}|}{n - |\hat{I}|} \leq H_6 (\tfrac{n}{|I_*|})^{-\alpha_6}. \tag{21}$$

Theorems 2 and 3 imply the next corollary.

Corollary 2. *For some constants $H_7, \alpha_7 > 0$, uniformly in $\theta \in \Theta(K)$,*

$$\mathrm{MTR}_l(\hat{I}, I_*) \leq H_7 (\tfrac{n}{|I_*| \vee 1})^{-\alpha_7}, \quad l = 1, \ldots, 4.$$

Remark 8. Recall the family of measures $\mathcal{P}_A$ defined by (6) and introduce

$$\mathcal{P}_\Theta = \mathcal{P}_\Theta(K) = \{\mathbb{P}_\theta \in \mathcal{P}_A : \theta \in \Theta(K)\}. \tag{22}$$

Clearly $\mathcal{P}_\Theta \subseteq \mathcal{P}_A$. In all the theorems and corollaries so far, all the claims uniform over $\theta \in \mathbb{R}^n$ hold uniformly also over $\mathcal{P}_A$ and all the claims uniform over $\theta \in \Theta(K)$ hold uniformly also over $\mathcal{P}_\Theta$.

Notice that FPR and NDP are controlled uniformly in $\theta \in \mathbb{R}^n$, whereas all the other quantities only in $\theta \in \Theta(K)$. As we already mentioned, the uniform control of either just Type I error or just Type II error is not much of a value, because this can always be achieved. It is a combination of the two types errors that one should try to control. The most natural choices of such combinations are the Hamming risk and the MTR's, studied in the present paper. Another possible direction in obtaining interesting results is simultaneous control of Type I error (say, FDR) and some estimation risk (or posterior convergence rate in case of Bayesian approach). Such a route is investigated in [10]. We should mention that $\hat{I}$ could also be derived as a result of empirical Bayes approach with appropriately chosen prior, and similar results could be derived on optimal estimation and posterior convergence rate.

3.4. Asymptotic regimes

Let us finally discuss possible asymptotic regimes. First, we note that asymptotics $n \to \infty$ is not well defined, unless we describe how the true signal $\theta \in \mathbb{R}^n$ itself evolves with n. Assume that $\theta \in \mathbb{R}$ evolves with $n \in \mathbb{N}$ in such a way that $\frac{|S(\theta)|}{n} \le \kappa < 1$. Then from the definition (10) of active coordinates I_*, it follows that $\frac{|I_*|}{n} \le \kappa$, but it could also $\frac{|I_*|}{n} \to 0$. This basically means that the signal is not getting "less sparse", in fact it can become "more sparse", making all the three criterions closer to zero. On the other hand, what can happen in this situation is that the signal is "getting lost" by spreading it over the bigger amount of coordinates. If, under growing dimension, we want the signal still to contain a certain portion of active coordinates, we need to make those coordinates more prominent, i.e., to strengthen the strong signal condition. The active coordinates should be increasing in magnitude when dimension is growing: by (11), $\theta_i^2 \ge A\sigma^2 \log(\frac{en}{|I_*|})$, $i \in I_*$. This can be relaxed by decreasing σ^2. Indeed, if $\sigma^2 \to 0$, then the strong signal condition is easier to satisfy. The strong signal condition is also easier to satisfy if $|I_*|$ increases, but then all the rates in the theorems (and their corollaries) deteriorate. This complies with the heuristics that "massive" active sets are more difficult to recover.

It is also interesting to reveal what affects the powers α_i in Theorems 2 and 3 and their corollaries. In the present formulation, the α_i's depend only on K (also via $A_0 = A_0(K)$ and $A_1 = A_1(K)$) and the constants from Condition (A1). Going through the proofs, we see that the bigger the differences between $A_1(K)$ and K and between K and $A_0(K)$, the bigger the α_i's. Basically, these quantities are responsible for the strong signal condition. To make it more evident, assume

$\theta \in \Theta'(A_0', A_1')$ defined by (14), setting for simplicity $A_0' = 0$. Consider the asymptotic regime $A_1' \to \infty$ in the strong signal condition $\theta_i^2 \geq A_1' \sigma^2 \log(\frac{en}{|I_*|})$ for the active coordinates $i \in I_*$. One can obtain that all α_i's in the claims are some multiple of A_1', so that one can interpret the asymptotic regime $A_1' \to \infty$ as the strong signal condition becoming more and more stringent, which in turn leads to $\alpha_i \to \infty$ in Theorems 2 and 3 and their corollaries.

3.5. Quantifying uncertainty for the variable selector

Here we construct confidence ball $B(\hat{\eta}, \hat{r})$ with optimal properties. Let $B(\hat{\eta}, \hat{r}) = \{\eta \in \{0,1\}^n : |\hat{\eta} - \eta| \leq \hat{r}\}$, $\hat{\eta} = \hat{\eta}(K)$ and $\tilde{I} = \tilde{I}(K)$ be given by (8). Define

$$\hat{r} = \hat{r}(\tilde{I}) = n(\tfrac{n}{|\tilde{I}| \vee 1})^{-\alpha_4'}, \tag{23}$$

for some α_4' such that $0 < \alpha_4' < \alpha_4$, with α_4 from Theorem 2.

The following theorem describes the coverage and size properties of the confidence ball based on $\hat{\eta}$ and $\hat{r}$.

Theorem 4. *With the same notation as in Theorem 2, let $\hat{r}$ be defined by (23) and $r_* = r_*(\theta) = n(\frac{n}{|I_*| \vee 1})^{-\alpha_4'}$. Then there exist constants $M_1', H_7, H_8, \alpha_8, \alpha_9$ such that, uniformly in $\theta \in \Theta(K)$,*

$$\mathbb{P}_\theta\big(\eta_* \notin B(\hat{\eta}, \hat{r})\big) \leq H_7(\tfrac{n}{|I_*| \vee 1})^{-\alpha_8},$$
$$\mathbb{P}_\theta\big(\hat{r} \geq M_1' r_*\big) \leq H_8(\tfrac{n}{|I_*| \vee 1})^{-\alpha_9}.$$

According to the UQ-framework (3), we have $\Theta_{\mathrm{cov}} = \Theta_{\mathrm{size}} = \Theta(K)$ for some $K > 0$, $r(\eta_*(\theta)) = r_*(\theta) = n(\frac{n}{|I_*| \vee 1})^{-\alpha_4'}$. In view of the lower bound results of the next section, the radius $\hat{r}$ is *optimal*, in a weak sense.

As we mentioned in Section 1.5, typically the so called *deceptiveness* issue emerges in UQ problems. But in this case, interestingly, there is no deceptiveness issue as such for our UQ problem. A heuristic explanation is as follows: the problem of active set recovery is already more difficult than the UQ-problem in a sense that solving the former problem implies solving the latter. Basically, the condition $\theta \in \Theta(K)$ for the parameter to have distinct active coordinates implies also that the parameter is non-deceptive.

4. Discussion: Weak optimality, phase transition

4.1. Lower bounds

Define

$$\Theta_s(a) = \{\theta \in \ell_0[s] : |\theta_i| \geq a, i \in S(\theta)\}.$$

Let $\Theta_s^+(a)$ be the version of $\Theta_s(a)$ when we put $\theta_i \geq a$ instead of $|\theta_i| \geq a$ in the definitions. Clearly, $\Theta_s^+(a) \subset \Theta_s(a)$. For $\theta \in \Theta_s(a)$, the traditional active set is $I_*(\theta) = S(\theta)$. To ensure strict separation from the inactive set, one typically imposes $a \geq \bar{a}_n > 0$ for appropriate $\bar{a}_n$. Later we will also make sure that $\Theta_s(\bar{a}_n) \subseteq \Theta'(0, A_1') \subseteq \Theta(K)$, so that a lower bound over the set $\Theta_s(\bar{a}_n)$ will imply a lower bound over the set of interest $\Theta(K)$.

The minimax lower bound over the class $\Theta_s(a)$ for the problem of the recovery of the active set $I_*(\theta) = S(\theta)$ in the Hamming risk for the normal means model was derived by [8]. Precisely, under the normality assumption $\xi_i \overset{\text{ind}}{\sim} \mathrm{N}(0,1)$, Theorem 2.2 from [8] states: for any $s < n$, $s' \in (0, s]$,

$$r_H(\Theta_s^+(a)) \triangleq \inf_{\check{\eta}} \sup_{\theta \in \Theta_s^+(a)} \mathbb{E}_\theta |\check{\eta} - \eta_{S(\theta)}| \geq s'\Psi_+(s,a) - 4s' \exp\big\{ -\tfrac{(s-s')^2}{2s}\big\},$$

where $\Psi_+(s,a) = (\frac{n}{s} - 1)\Phi\big(-\frac{a}{2\sigma} - \frac{\sigma}{a}\log(\frac{n}{s}-1)\big) + \Phi\big(-\frac{a}{2\sigma} + \frac{\sigma}{a}\log(\frac{n}{s}-1)\big)$. If $a^2 \leq 2\sigma^2 \log(\frac{n}{s}-1)$ then, by taking $s' = s/2$ in the above display, we get the lower bound

$$r_H(\Theta_s^+(a)) \geq \tfrac{s}{2}\Phi(0) - 2se^{-s/8} = s\big(\tfrac{1}{4} - 2e^{-s/8}\big) > 0.085s \tag{24}$$

for $s \geq 20$. Expectedly, if $a^2 \leq 2\sigma^2 \log(\frac{n}{s}-1)$, it is impossible to achieve even consistency. On the other hand, if $a^2 > 2\sigma^2 \log(\frac{n}{s}-1)$ and $n/s \geq 2.7$, then

$$\Phi\big(-\tfrac{a}{2\sigma} - \tfrac{\sigma}{a}\log(\tfrac{n}{s}-1)\big) \geq \Phi(-a/\sigma) \geq (2/\pi)^{1/2} e^{-4a^2/\sigma^2}.$$

Assuming further $\frac{a^2}{\sigma^2} \lesssim s$ (this implies $s \gtrsim \log n$) and taking again $s' = s/2$,

$$\begin{aligned} r_H(\Theta_s^+(a)) &\geq \tfrac{(n-s)}{2}\Phi\big(-\tfrac{a}{2\sigma} - \tfrac{\sigma}{a}\log(\tfrac{n}{s}-1)\big) - 2se^{-s/8} \\ &\geq C_1(n-s)e^{-C_2a^2/\sigma^2} - C_3e^{-C_4 s} \geq C_5(n-s)e^{-C_6a^2/\sigma^2}. \end{aligned} \tag{25}$$

Assume that $a^2/\sigma^2 = A\log(\frac{en}{s})$ for some $A > 2$ and $\log n \lesssim s \leq n/2.7$. Then, under the normality assumption $\xi_i \overset{\text{ind}}{\sim} \mathrm{N}(0,1)$, it follows from (25) that for some $c_1, c_2 > 0$ (depending only on A, in fact $c_2 = C_6 A$)

$$\begin{aligned} \inf_{\check{I}} \sup_{\theta\in\Theta_s(a)} \mathbb{E}_\theta\big(|\check{I}\backslash S(\theta)| + |S(\theta)\backslash \check{I}|\big) &= \inf_{\check{I}} \sup_{\theta\in\Theta_s(a)} \mathbb{E}_\theta |\eta_{\check{I}} - \eta_{S(\theta)}| \\ &= r_H(\Theta_s(a)) \geq r_H(\Theta_s^+(a)) \geq c_1 n(n/s)^{-c_2}. \end{aligned} \tag{26}$$

Notice that the lower bound (26) is over rather "small" set $\Theta_s(a)$ which actually makes it strong. In view of (24) and (26), even in the simplest normal model with $\theta \in \Theta_s(a)$, we need a strong signal condition. For example, if $a^2 = \bar{a}_n^2 = A\sigma^2\log(\frac{en}{s})$ then we need $A > 2$ to avoid inconsistency in recovering the active set $S(\theta)$. According to the terminology from [8], if $r_H(\Theta_s(a)) \to 0$ as $n \to \infty$, the *exact recovery* of the active set $S(\theta)$ takes place; and *almost full recovery* occurs if $r_H(\Theta_s(a))/s \to 0$ as $n \to \infty$ (assuming that $s > 0$).

The above lower bound (26) reveals some sort of *phase transition*. Indeed, the almost full recovery (the mildest criterion) can occur if $s \ll n$ and the constant A is sufficiently large, so that $c_2 > 1$, or, if $s \asymp n$ (but $s/n \leq c < 1$) and

$c_2 \to \infty$. The exact recovery is even more difficult to fulfill, it can only occur if $n(n/s)^{-c_2} \to 0$ as $n \to \infty$. This is determined by the combination of two factors, the constant c_2 and the order of parameter s. The parameter s describes the sparsity of the signal θ and the constant c_2 depends on A (is a multiple of A) which expresses the *signal strength*.

By (15) (or, by property (c) from Remark 6), it follows that $\Theta_s(a(K)) \subseteq \Theta'(0, A_1(K)) \subseteq \Theta(K)$ for $a^2(K) = A_1(K)\sigma^2 \log(\frac{qn}{|S(\theta)|})$. Property (b) from Remark 6 implies also that $I_* = I_*(A_1(K), \theta) \subseteq S(\theta)$ for all $\theta \in \mathbb{R}^n$. The last two facts and (26) allow us to derive the following lower bound:

$$\inf_{\check{I}} \sup_{\theta \in \Theta(K)} \mathbb{E}_\theta \frac{|\check{I} \backslash I^*| + |I_* \backslash \check{I}|}{n(|I_*|/n)^{c_2}} \geq \inf_{\check{I}} \sup_{\theta \in \Theta'(0, A_1(K))} \mathbb{E}_\theta \frac{|\check{I} \backslash I^*| + |I_* \backslash \check{I}|}{n(|S(\theta)|/n)^{c_2}}$$
$$\geq \inf_{\check{I}} \sup_{\theta \in \Theta_s(a(K))} \mathbb{E}_\theta \frac{|\check{I} \backslash S(\theta)| + |S(\theta) \backslash \check{I}|}{n(|S(\theta)|/n)^{c_2}} = \frac{r_H(\Theta_s(a(K)))}{ne^{c_2 \log(s/n)}} \geq c_1, \tag{27}$$

where the distribution $\mathbb{P}_\theta$ is taken to be the product normal as in (2), $\log n \lesssim s \leq n/2.7$ and $A_1(K) > 2$. The normalizing factor for the Hamming risk is thus $ne^{c_2 \log(|I_*|/n)}$. Bounds for the MTR's can be derived similarly.

Recall that the bound (27) is in the regime $A_1(K) > 2$. Otherwise (when $A_1(K) \leq 2$), we have by (24) that $\inf_{\check{I}} \sup_{\theta \in \Theta(K)} \mathbb{E}_\theta \frac{|\check{I} \backslash I^*| + |I_* \backslash \check{I}|}{|I_*|} \geq 0.085$.

Remark 9. Notice that the measure $\mathbb{P}_\theta$ in the above lower bounds is the product normal measure (2). The same lower bound trivially holds in (27) when instead of $\sup_{\theta \in \Theta(K)}$ we take $\sup_{\mathbb{P}_\theta \in \mathcal{P}_\Theta}$, where $\mathcal{P}_\Theta$ is defined by (22). This is because the normal measures in (27) (parametrized by $\theta \in \Theta(K)$) are all contained in the family $\mathcal{P}_\Theta$ by the definition (22) of $\mathcal{P}_\Theta$.

4.2. Phase transition

Relating the lower bound (27) with the results of the previous section, we claim that the selector $\hat{I}$ (and the corresponding $\hat{\eta}$) is *optimal* in the following weak sense:

$$c_1 \leq \inf_{\check{I}} \sup_{\theta \in \Theta(K)} \frac{R_H(\check{I}, I_*)/n}{e^{c_2 \log(|I_*|/n)}}, \quad \sup_{\theta \in \Theta(K)} \frac{R_H(\hat{I}, I_*)/n}{e^{\alpha_4 \log((|I_*| \vee 1)/n)}} \leq H_3. \tag{28}$$

In view of Remark 9, the above relations hold also if instead of $\sup_{\theta \in \Theta(K)}$ we take $\sup_{\mathbb{P}_\theta \in \mathcal{P}_\Theta}$, where $\mathcal{P}_\Theta$ is defined by (22). Bounds for the MTR's can also be derived.

Admittedly, the optimality in (28) is in a weak sense as it is up to constants c_2, α_4, which differ in general. But this is the best we can achieve under the general robust setting in this paper. Constant c_2 in the lower bound is established for the normal sub-model (2) of our more general model (1), whereas α_4 is obtained uniformly for the general model (1) under Condition (A1) (thus determined by the constants from Condition (A1)).

Remark 10. If we want to match precisely the upper and lower bounds, we would need to specify the error distribution (or severely restrict the choice). This problem seemingly interesting and challenging does not align with the main focus of the present paper, the robust setting.

However, even these, not precisely matched, lower and upper bounds in (28) can demonstrate some sort of *phase transition* phenomenon, also in the general setting (1). Precisely, the minimax Hamming r_H risk (and MTR) can be either close to zero or not, depending on the combination of the signal sparsity $|I_*|$ and signal magnitude (how strong the signal is), reflected by the constants c_2, α_4. The dependence of the normalizing factor on sparsity $|I_*|$ is only through the ratio $n/|I_*|$.

The "informativeness" of the model (how "bad" the noise is) plays a role as well in that it determines how large K must be in the set $\Theta(K)$ for the upper bound in (28) to hold, which depends on the constants from Condition (A1). It also determines the value of the constant α_4 in (28), i.e., the signal strength of the active coordinates.

We can further distill the role of the signal strength by considering, say, $\Theta'(0, A_1')$ (with $A_1' \geq A_1(K)$) instead of $\Theta(K)$ in (28):

$$c_1 \leq \inf_{\check{I}} \sup_{\theta \in \Theta'(0,A_1)} \frac{R_H(\check{I}, I_*)/n}{e^{c_2 \log(|I_*|/n)}}, \qquad \sup_{\theta \in \Theta'(0,A_1)} \frac{R_H(\hat{I}, I_*)/n}{e^{\alpha_4 \log((|I_*| \vee 1)/n)}} \leq H_3.$$

In this case, $I_* = S(\theta)$ and then the both constant c_2 and α_4 are just (different) multiples of A_1'. This clearly demonstrates how the signal strength parameter A_1' enters the convergence rate for recovery of the active set $S(\theta)$, also making the relation between lower and upper bound more explicit.

5. Proofs of the theorems

Proof of Theorem 1. For any $a, b \in \mathbb{R}$, $(a+b)^2 \leq 2a^2+2b^2$, hence also $-(a+b)^2 \leq -a^2/2 + b^2$. Using these elementary inequalities, the definition (8) of $\tilde{I}(K)$, we derive that, for any $I, I_0 \subseteq [n]$,

$$\begin{aligned}
&\mathbb{P}_\theta(\tilde{I}(K) = I) \leq \mathbb{P}_\theta\Big(\sum_{i \in I^c} X_i^2 + K\sigma^2 \ell(|I|) \leq \sum_{i \in I_0^c} X_i^2 + K\sigma^2 \ell(|I_0|)\Big) \\
&= \mathbb{P}_\theta\Big(\sum_{i \in I \setminus I_0} \tfrac{X_i^2}{\sigma^2} - \sum_{i \in I_0 \setminus I} \tfrac{X_i^2}{\sigma^2} \geq K[\ell(|I|) - \ell(|I_0|)]\Big) \\
&\leq \mathbb{P}_\theta\Big(\sum_{i \in I \setminus I_0} (\tfrac{2\theta_i^2}{\sigma^2} + 2\xi_i^2) - \sum_{i \in I_0 \setminus I} (\tfrac{\theta_i^2}{2\sigma^2} - \xi_i^2) \geq K(\ell(|I|) - \ell(|I_0|))\Big) \\
&= \mathbb{P}_\theta\Big(\sum_{i \in I \setminus I_0} 2\xi_i^2 + \sum_{i \in I_0 \setminus I} \xi_i^2 \geq \sum_{i \in I_0 \setminus I} \tfrac{\theta_i^2}{2\sigma^2} - \sum_{i \in I \setminus I_0} \tfrac{2\theta_i^2}{\sigma^2} + K(\ell(|I|) - \ell(|I_0|))\Big).
\end{aligned}$$

In particular, for any I, I_0 such that $I_0 \subseteq I$, we have

$$\mathbb{P}_\theta(\tilde{I}(K) = I) \leq \mathbb{P}_\theta\Big(\sum_{i \in I \setminus I_0} \xi_i^2 \geq \tfrac{K}{2}(\ell(|I|) - \ell(|I_0|)) - \sum_{i \in I \setminus I_0} \tfrac{\theta_i^2}{\sigma^2}\Big), \tag{29}$$

and for any I, I_0 such that $I \subseteq I_0$, we have

$$\mathbb{P}_\theta(\tilde{I}(K) = I) \le \mathbb{P}_\theta\Big(\sum_{i\in I_0\backslash I} \xi_i^2 \ge \sum_{i\in I_0\backslash I} \tfrac{\theta_i^2}{2\sigma^2} + K(\ell(|I|) - \ell(|I_0|))\Big). \tag{30}$$

Now we prove (i). For brevity, denote for now $I_* = I_*(A_0) = I_*(A_0,\theta)$. If $A_0|I\backslash I_*|\log(\frac{qn}{|I\cup I_*|}) < \sum_{i\in I\backslash I_*}\frac{\theta_i^2}{\sigma^2}$ would hold for some $I \subseteq [n]$, then

$$\begin{aligned} r^2_{A_0}(I\cup I_*,\theta) &= \sum_{i\notin I\cup I_*} \theta_i^2 + A_0\sigma^2|I\cup I_*|\log(\tfrac{qn}{|I\cup I_*|}) \\ &\le \sum_{i\notin I\cup I_*} \theta_i^2 + A_0\sigma^2|I\backslash I_*|\log(\tfrac{qn}{|I\cup I_*|}) + A_0\sigma^2|I_*|\log(\tfrac{qn}{|I_*|}) \\ &< \sum_{i\notin I\cup I_*} \theta_i^2 + \sum_{i\in I\backslash I_*}\theta_i^2 + A_0\sigma^2|I_*|\log(\tfrac{qn}{|I_*|}) \\ &= \sum_{i\notin I_*}\theta_i^2 + A_0\sigma^2|I_*|\log(\tfrac{qn}{|I_*|}) = r^2_{A_0}(\theta), \end{aligned}$$

which contradicts the definition (10) of $I_* = I_*(A_0,\theta)$. Hence,

$$\sum_{i\in I\backslash I_*}\tfrac{\theta_i^2}{\sigma^2} \le A_0|I\backslash I_*|\log(\tfrac{qn}{|I\cup I_*|}) \le A_0|I|\log(\tfrac{qn}{|I|}) = A_0\ell(|I|).$$

Using this and (29) with $I_0 = I_*\cap I$ (so that $I\backslash I_0 = I\backslash I_*$) yields

$$\begin{aligned} \mathbb{P}_\theta(\tilde{I}(K_0) = I) &\le \mathbb{P}_\theta\Big(\sum_{i\in I\backslash I_0}\xi_i^2 \ge (\tfrac{K_0}{2} - A_0)\ell(|I|) - \tfrac{K_0}{2}\ell(|I_0|)\Big) \\ &\le \mathbb{P}_\theta\Big(\sum_{i\in I}\xi_i^2 \ge (\tfrac{K_0}{2} - A_0)\ell(|I|) - \tfrac{K_0}{2}\ell(|I_*|)\Big). \end{aligned}$$

Let $\mathcal{J} = \{I\subseteq[n] : \ell(|I|) \ge M_0\ell(|I_*(A_0)|) + M\}$, with $M_0 = K_0/(2C_0)$ where $C_0 = K_0/2 - A_0 - M_\xi > 0$, which holds for any $K_0 > 2(A_0 + M_\xi)$. The last display and (7) imply that, for any $\theta\in\mathbb{R}^n$, with $\alpha_0 = C_0\alpha_\xi$,

$$\begin{aligned} \mathbb{P}_\theta\big(\tilde{I}(K_0)\in\mathcal{J}\big) &= \sum_{I\in\mathcal{J}}\mathbb{P}_\theta(\tilde{I}(K_0) = I) \\ &\le \sum_{I\in\mathcal{J}}\mathbb{P}_\theta\Big(\sum_{i\in I}\xi_i^2 \ge (\tfrac{K_0}{2} - A_0)\ell(|I|) - \tfrac{K_0}{2}\ell(|I_*|)\Big) \\ &= \sum_{I\in\mathcal{J}}\mathbb{P}_\theta\Big(\sum_{i\in I}\xi_i^2 \ge M_\xi\ell(|I|) + C_0\ell(|I|) - \tfrac{K_0}{2}\ell(|I_*|)\Big) \\ &\le \sum_{I\subseteq[n]}\mathbb{P}_\theta\Big(\sum_{i\in I}\xi_i^2 \ge M_\xi\ell(|I|) + C_0M\Big) \le H_0e^{-\alpha_0M}. \end{aligned} \tag{31}$$

Let $M_1' = 4(M_0+1)$. If $|I| \ge M_1'|I_*| + M$ and $|I_*| \le qn/(M_1')^2$, then

$$\ell(|I|) = |I|\log\big(\tfrac{qn}{|I|}\big) \ge \tfrac{1}{2}\ell(|I|) + \tfrac{1}{2}|I| \ge \tfrac{M_1'}{2}|I_*|\log\big(\tfrac{qn}{M_1'|I_*|}\big) + \tfrac{M}{2}$$

$$\geq \tfrac{M_1'}{4}|I_*|\log\big(\tfrac{qn}{|I_*|}\big)+\tfrac{M}{2}=(M_0+1)\ell(|I_*|)+\tfrac{M}{2}.$$

Hence, for $|I_*|\leq qn/M_1^2$, by using (31),

$$\begin{aligned}\mathbb{P}_\theta\big(|\tilde I(K_0)|\geq M_1'|I_*|+M\big)&\leq \mathbb{P}_\theta\big(\ell(|\tilde I(K_0)|)\geq (M_0+1)\ell(|I_*|)+\tfrac{M}{2}\big)\\ &\leq H_0e^{-\alpha_0\ell(|I_*|)-\alpha_0 M/2}.\end{aligned}$$

If $|I_*|>qn/(M_1')^2$ and $M_1''=(M_1')^2/q$, then we trivially obtain

$$\mathbb{P}_\theta\big(|\tilde I(K_0)|\geq M_1''|I_*|+M\big)\leq \mathbb{P}_\theta\big(|\tilde I(K_0)|\geq M_1''|I_*|\big)\leq \mathbb{P}_\theta\big(|\tilde I(K_0)|>n\big)=0.$$

Hence the choice $M_1=\max\{M_1',M_1''\}$ ensures (i).

Next, we prove the assertion (ii). For the rest of the proof, denote for brevity $I_*=I_*(A_1)$. Define $\mathcal{T}=\{I\in\mathcal{I}:|I|\leq\delta|I_*|-M\}$, $\delta\in[0,1)$. Hence, for any $I\in\mathcal{T}$,

$$\begin{aligned}\ell(|I_*\cup I|)&=|I_*\cup I|\log(\tfrac{qn}{|I_*\cup I|})\leq |I_*|\log(\tfrac{qn}{|I_*\cup I|})+|I|\log(\tfrac{qn}{|I_*\cup I|})\\ &\leq(1+\delta)|I_*|\log(\tfrac{qn}{|I_*|})-M=(1+\delta)\ell(|I_*|)-M.\end{aligned}\tag{32}$$

Next, by (11) and the fact that $|I|\leq\delta|I_*|$, we obtain that for any $I\in\mathcal{T}$,

$$\sum_{i\in I_*\backslash I}\tfrac{\theta_i^2}{\sigma^2}\geq|I_*\backslash I|A_1\log(\tfrac{en}{|I_*|})\geq A_1(1-\delta)|I_*|\log(\tfrac{en}{|I_*|})\geq\tfrac{A_1(1-\delta)}{2}\ell(|I_*|).\tag{33}$$

Denote for brevity $C_{A_1}=\frac{A_1(1-\delta)}{4}-K_1(1+\delta)$. Using the relation (30) with $I_0=I_*\cup I$ (so that $I_0\backslash I=I_*\backslash I$), the relations (32), (33), and (7), we derive that

$$\begin{aligned}\mathbb{P}_\theta(\tilde I(K_1)\in\mathcal{T})&=\sum_{I\in\mathcal{T}}\mathbb{P}_\theta(\tilde I(K_1)=I)\\
&\leq\sum_{I\in\mathcal{T}}\mathbb{P}_\theta\Big(\sum_{i\in I_*\backslash I}\xi_i^2\geq\sum_{i\in I_*\backslash I}\tfrac{\theta_i^2}{2\sigma^2}+K_1(\ell(|I|)-\ell(|I_*\cup I|))\Big)\\
&\leq\sum_{I\in\mathcal{T}}\mathbb{P}_\theta\Big(\sum_{i\in I_*\backslash I}\xi_i^2\geq\big(\tfrac{A_1(1-\delta)}{4}-K_1(1+\delta)\big)\ell(|I_*|)+K_1[\ell(|I|)+M]\Big)\\
&=\sum_{I\in\mathcal{T}}\mathbb{P}_\theta\Big(\sum_{i\in I_*\backslash I}\xi_i^2\geq C_{A_1}\ell(|I_*|)+K_1\ell(|I|)+K_1M\Big)\\
&\leq\sum_{I\in\mathcal{T}}\mathbb{P}_\theta\Big(\sum_{i\in I_*\backslash I}\xi_i^2\geq M_\xi\ell(|I_*|)+(C_{A_1}-M_\xi)\ell(|I_*|)+K_1M\Big)\\
&\leq\sum_{I\subseteq[n]}\mathbb{P}_\theta\Big(\sum_{i\in I}\xi_i^2\geq M_\xi\ell(|I|)+(C_{A_1}-M_\xi)\ell(|I_*|)+K_1M\Big)\\
&\leq H_0e^{-\alpha_\xi M-\alpha_\xi(C_{A_1}-M_\xi)\ell(|I_*|)}=H_0e^{-\alpha_1M-\alpha_1'\ell(|I_*|)},\end{aligned}$$

where $\alpha_1=\alpha_\xi K_1$, $\alpha_1'=\alpha_\xi(C_{A_1}-M_\xi)$ and A_1 is assumed to be so large that $C_{A_1}>M_\xi$. This proves (ii). $\square$

Proof of Theorem 2. First we prove (16). Denote for brevity $A_0 = A_0(K)$ and recall that $I^* = I_*(A_0, \theta)$. Consider the case $|I^*| \geq 1$ and let $B = \{|\tilde{I}| > M_1|I^*|\}$. By (12), $\theta_i^2 \leq A_0\sigma^2 \log(\frac{qn}{|I^*|\vee 1})$ for all $i \in (I^*)^c$. Using this, Condition (A1) with $I = \{i\}$, (9), and property (i) of Theorem 1, we derive

$$\begin{aligned}
\mathrm{FPR}(\hat{I}, I^*) &= \mathbb{E}_\theta \tfrac{|\hat{I}\backslash I^*|}{n-|I^*|} = \tfrac{1}{n-|I^*|} \sum_{i\in(I^*)^c} \mathbb{E}_\theta \hat{\eta}(X_i)(1_B + 1_{B^c}) \\
&\leq \mathbb{P}_\theta(B) + \tfrac{1}{n-|I^*|} \sum_{i\in(I^*)^c} \mathbb{P}_\theta\big(\tfrac{2\theta_i^2}{\sigma^2} + 2\xi_i^2 \geq K \log(\tfrac{qn}{M_1|I^*|})\big) \\
&\leq \mathbb{P}_\theta(B) + \tfrac{1}{n-|I^*|} \sum_{i\in(I^*)^c} \mathbb{P}_\theta\big(\xi_i^2 \geq (\tfrac{K}{2} - A_0)\log(\tfrac{qn}{|I^*|}) - \tfrac{K}{2}\log M_1\big) \\
&\leq H_0 e^{-\alpha_0 \ell(|I^*|)} + \tfrac{H_\xi}{n-|I^*|} \sum_{i\in(I^*)^c} e^{-\alpha_\xi(\frac{K}{2}-A_0-C_\xi)\log(\frac{qn}{|I^*|}) + \frac{\alpha_\xi K}{2}\log M_1} \\
&\leq H_0 e^{-\alpha_0 \ell(|I^*|)} + H'(\tfrac{n}{|I^*|})^{-\alpha'},
\end{aligned} \tag{34}$$

for K and A_0 such that $\frac{K}{2} - A_0 > C_\xi$ and the property (i) of Theorem 1 can be applied, where $\alpha' = \alpha_\xi(\frac{K}{2} - A_0 - C_\xi)$ and $H' = H_\xi e^{\alpha_\xi K \log\sqrt{M_1}}/q^{\alpha'}$.

Now we consider the case $|I^*| = 0$ (i.e., $I^* = \varnothing$). Reasoning similarly to (34), now with $B = \{|\tilde{I}| > \log n\}$, we derive

$$\begin{aligned}
\mathrm{FPR}(\hat{I}, I^*) &\leq \mathbb{P}_\theta(B) + \tfrac{1}{n} \sum_{i\in[n]} \mathbb{P}_\theta\big(\tfrac{2\theta_i^2}{\sigma^2} + 2\xi_i^2 \geq K\log(\tfrac{qn}{\log n})\big) \\
&\leq \mathbb{P}_\theta(B) + \tfrac{1}{n} \sum_{i\in[n]} \mathbb{P}_\theta\big(\xi_i^2 \geq (\tfrac{K}{2} - A_0)\log(qn) - \tfrac{K}{2}\log\log n\big) \\
&\leq H'' n^{-\alpha''}.
\end{aligned} \tag{35}$$

The relations (34) and (35) establish (16).

Next, we proof the assertion (17). If $I_* = \varnothing$, the claim follows, assume $|I_*| \geq 1$. Denote $B_\delta = \{|\tilde{I}| \leq \delta|I_*|\}$. Using Condition (A1), the definition (9), (11), property (ii) of Theorem 1, and the fact that $(a+b)^2 \geq 2a^2/3 - 2b^2$ for any $a, b \in \mathbb{R}$, we have that, for $I_* = I_*(A_1)$ and $\hat{I} = \hat{I}(K)$,

$$\begin{aligned}
\mathrm{NDR}(\hat{I}, I_*) &= \frac{1}{|I_*|}\mathbb{E}_\theta|I_*\backslash\hat{I}| = \frac{1}{|I_*|}\sum_{i\in I_*}\mathbb{E}_\theta(1-\hat{\eta}(X_i)) \\
&= \frac{1}{|I_*|}\sum_{i\in I_*}\mathbb{P}_\theta\Big(|\theta_i + \sigma\xi_i| < \sigma\big[K\log(\tfrac{qn}{|\tilde{I}|})\big]^{1/2}\Big) \\
&\leq \frac{1}{|I_*|}\sum_{i\in I_*}\mathbb{P}_\theta\Big(\tfrac{2}{3}\theta_i^2 - 2\sigma^2\xi_i^2 < \sigma^2 K\log(\tfrac{qn}{|\tilde{I}|})\Big) \\
&\leq \frac{1}{|I_*|}\sum_{i\in I_*}\mathbb{P}_\theta\Big(\xi_i^2 > \tfrac{A_1}{3}\log(\tfrac{en}{|I_*|}) - \tfrac{K}{2}\log(\tfrac{qn}{|\tilde{I}|})\Big)
\end{aligned}$$

$$\leq \mathbb{P}_\theta(B_\delta) + \frac{1}{|I_*|}\sum_{i\in I_*}\mathbb{P}_\theta\Big(\xi_i^2 > \tfrac{A_1}{6}\log(\tfrac{qn}{|I_*|}) - \tfrac{K}{2}\log(\tfrac{qn}{|\tilde{I}|}), B_\delta^c\Big)$$
$$\leq H_0 e^{-\alpha_1'\ell(|I_*|)} + \frac{1}{|I_*|}\sum_{i\in I_*}\mathbb{P}_\theta\big(\xi_i^2 > (\tfrac{A_1}{6} - \tfrac{K}{2})\log(\tfrac{qn}{|I_*|}) + \tfrac{K}{2}\log\delta\big)$$
$$\leq H_0 e^{-\alpha_1'\ell(|I_*|)} + C_2(\tfrac{qn}{I_*})^{-\alpha_3'} \leq H_2(\tfrac{n}{|I_*|})^{-\alpha_3}, \tag{36}$$

for sufficiently large A_1 (such that $\frac{A_1}{6} - \frac{K}{2} > M_\xi$ and the property (ii) of Theorem 1 can be applied), where $\alpha_3' = \alpha_\xi(\frac{A_1}{3} - \frac{K}{2} - M_\xi)$ and $C_2 = H_\xi e^{-\alpha_\xi K\log\sqrt{\delta}}$. The relation (17) is proved.

Since $I^* = I_*(A_0(K)) = I_*(A_1(K)) = I_*$ for $\theta \in \Theta(K)$, the assertion (18) follows from the relations (34) and (36): uniformly in $\theta \in \Theta(K)$,

$$\mathbb{E}_\theta\big(|\hat{I}\backslash I_*| + |I_*\backslash\hat{I}|\big) \leq (n - |I_*|)H_1(\tfrac{n}{|I_*|\vee 1})^{-\alpha_2} + |I_*|H_2(\tfrac{n}{|I_*|})^{-\alpha_3}$$
$$\leq H_3 n(\tfrac{n}{|I_*|\vee 1})^{-\alpha_4}.$$

□

Proof of Theorem 3. First, we proof assertion (20). Introduce the event $B_\delta = \{|\tilde{I}| < \delta|I_*|\}$. We argue along the same lines as in (34) and (35) for the two cases $|I_*| > 0$ and $|I_*| = 0$. For the case $|I_*| > 0$, we use (ii) of Theorem 1, (16) of Theorem 2 and the fact that $I^* = I_*$ for $\theta \in \Theta(K)$ to derive

$$\mathrm{FDR}(\hat{I}) = \mathbb{E}_\theta\tfrac{|\hat{I}\backslash I_*|}{|\hat{I}|} = \mathbb{E}_\theta\big[\tfrac{|\hat{I}\backslash I^*|}{|\hat{I}|}(1_{B_\delta^c} + 1_{B_\delta})\big] \leq \tfrac{1}{\delta|I_*|}\mathbb{E}_\theta|\hat{I}\backslash I^*| + \mathbb{P}_\theta(B_\delta)$$
$$\leq \tfrac{1}{\delta|I_*|}H_1 n\big(\tfrac{n}{|I^*|\vee 1}\big)^{-\alpha_2} + H_0 e^{-\alpha_1'\ell(|I_*|)} \leq H_5(\tfrac{n}{|I_*|})^{-\alpha_6}.$$

The case $|I_*| = 0$ is handled as follows. If $|\hat{I}| = 0$, the claim holds. Assume $|\hat{I}| \geq 1$, then

$$\mathrm{FDR}(\hat{I}) \leq \mathbb{E}_\theta|\hat{I}\backslash I^*| \leq H_1 n n^{-\alpha_2} = H_1 n^{-(\alpha_2 - 1)}.$$

Next, we prove assertion (21). Introduce the event $B = \{|\tilde{I}| > M_1|I^*|\}$. Consider two cases: the case $|I^*| \geq n/(2M_1)$ and the case $|I^*| < n/(2M_1)$. Suppose $|I^*| \geq n/(2M_1)$, then $\mathrm{FNR}(\hat{I}) = \mathbb{E}_\theta\frac{|I_*\backslash\hat{I}|}{n-|\hat{I}|} \leq 1 \leq 2M_1\frac{|I^*|}{n} = 2M_1\frac{|I_*|}{n}$ and (21) holds, as $I^* = I_*$ for $\theta \in \Theta(K)$.

Now suppose $|I^*| < n/(2M_1)$. Then using the same reasoning as in (36), Theorem 1 and the fact that $I^* = I_*$ for $\theta \in \Theta(K)$, we again obtain (21):

$$\mathrm{FNR}(\hat{I}) = \mathbb{E}_\theta\tfrac{|I_*\backslash\hat{I}|}{n-|\hat{I}|} = \mathbb{E}_\theta\big[\tfrac{|I_*\backslash\hat{I}|}{(n-|\hat{I}|)}(1_{B^c} + 1_B)\big]$$
$$\leq \tfrac{1}{n - M_1|I^*|}\mathbb{E}_\theta|I_*\backslash\hat{I}| + \mathbb{P}_\theta(B) \leq H_6(\tfrac{n}{|I_*|})^{-\alpha_7}.$$

□

Proof of Theorem 4. We first establish the coverage property. Recall that $\hat{r} = \hat{r}(\tilde{I}) = n(\frac{n}{|\tilde{I}|\vee 1})^{-\alpha_4'}$ and by (18) of Theorem 2,

$$\sup_{\theta\in\Theta(K)} \mathbb{E}_\theta|\hat{\eta}-\eta_*| \le H_3 n(\tfrac{n}{|I_*|\vee 1})^{-\alpha_4}. \tag{37}$$

Denote $B_\delta = \{|\tilde{I}| \le \delta |I_*|\}$. Now, by using the Markov inequality, property (ii) of Theorem 1 and (37), we obtain, uniformly in $\theta \in \Theta(K)$,

$$\begin{aligned}\mathbb{P}_\theta\big(\eta_* \notin B(\hat{\eta},\hat{r})\big) &= \mathbb{P}_\theta\big(|\eta_*-\hat{\eta}| > \hat{r}, B_\delta^c\big) + \mathbb{P}_\theta(B_\delta)\\ &\le \mathbb{P}_\theta(B_\delta) + \mathbb{P}_\theta\big(|\eta_*-\hat{\eta}| > n(\tfrac{n}{(\delta|I_*|)\vee 1})^{-\alpha_4'}\big)\\ &\le H_0 e^{-\alpha_1'\ell(|I_*|)} + H_3\delta^{\alpha_4'}(\tfrac{n}{|I_*|\vee 1})^{-(\alpha_4-\alpha_4')} \le H_7(\tfrac{n}{|I_*|\vee 1})^{-\alpha_7},\end{aligned}$$

which proves the coverage property.

It remains to prove the size property. Let $B = \{|\tilde{I}| \ge M_1|I^*|\}$. For any $M_1' > M_1^{\alpha_4'} \vee 1$, we have that, uniformly in $\theta \in \Theta(K)$ (so that $I^* = I_*$),

$$\mathbb{P}_\theta\big(\hat{r} \ge M_1' r_*, B^c\big) \le \mathbb{P}_\theta\Big(n\big(\tfrac{n}{M_1|I^*|\vee 1}\big)^{-\alpha_4'} \ge M_1' n\big(\tfrac{n}{|I_*|\vee 1}\big)^{-\alpha_4'}\Big) = 0.$$

Using this, the property (i) of Theorem 1 and the fact that $I^* = I_*$ for $\theta \in \theta(K)$, we derive that, uniformly in $\theta \in \Theta(K)$,

$$\begin{aligned}\mathbb{P}_\theta\big(\hat{r} \ge M_1' r_*\big) &\le \mathbb{P}_\theta(B) + \mathbb{P}_\theta\big(\hat{r} \ge M_1' r_*, B^c\big)\\ &= \mathbb{P}_\theta(B) \le H_0 e^{-\alpha_0\ell(|I^*|)} \le H_8(\tfrac{n}{|I^*|\vee 1})^{-\alpha_8},\end{aligned}$$

yielding the size property. ∎

References

[1] Abramovich, F., Benjamini, Y., Donoho, D. L. and Johnstone, I. M. (2006). Adapting to unknown sparsity by controlling the false discovery rate. *Ann. Statist.* 34, 584–653. MR2281879

[2] Arias-Castro, E. and Chen, S. (2017). Distribution-free multiple testing. *Electron. J. Statist.* 11, 1983–2001. MR3651021

[3] Belitser, E. (2017). On coverage and local radial rates of credible sets. *Ann. Statist.* 45, 1124–1151. MR3662450

[4] Belitser, E. and Nurushev, N. (2019). General framework for projection structures. ArXiv:1904.01003.

[5] Belitser, E. and Nurushev, N. (2020). Needles and straw in a haystack: robust empirical Bayes confidence for possibly sparse sequences. *Bernoulli* 26, 191–225. MR4036032

[6] Benjamini, Y. and Hochberg, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J. Roy. Statist. Soc. Ser. B.* 57, 289–300. MR1325392

[7] Blanchard, G. and Roquain, E. (2009). Adaptive false discovery rate control under independence and dependence. *J. Mach. Learn. Res.* 10, 2837–2871. MR2579914

[8] Butucea, C., Ndaoud, M., Stepanova, N. and Tsybakov, A. (2018). Variable selection with Hamming loss. *Ann. Statist.* 46, 1837–1875. MR3845003

[9] Butucea, C. and Stepanova, N. (2017). Adaptive variable selection in nonparametric sparse additive models. *Electron. J. Stat.* 11, 2321–2357. MR3656494

[10] Castillo, I. and Roquain, É. (2018). On spike and slab empirical Bayes multiple testing. *Ann. Statist.* 48, 2548–2574. MR4152112

[11] Cavalier, L. and Golubev, Yu. (2006). Risk hull method and regularization by projection of ill-posed inverse problems. *Ann. Statist.* 34, 1653–1677. MR2283712

[12] Donoho, D. and Jin, J. (2004). Higher criticism for detecting sparse heterogeneous mixtures. *Ann. Statist.* 32, 962–994. MR2065195

[13] Efron, B. (2008). Microarrays, empirical Bayes and the two-groups model. *Statist. Sci.* 23, 1–22. MR2431866

[14] Efron, B. (2009). Empirical Bayes estimates for large-scale prediction problems. *J. Am. Stat. Assoc.* 104, 1015–1028. MR2562003

[15] Genovese, C. and Wasserman, L. (2002). Operating characteristics and extensions of the false discovery rate procedure. *J. Roy. Statist. Soc. Ser. B.* 64, 499–517. MR1924303

[16] Ingster, Y. (1997). Some problems of hypothesis testing leading to infinitely divisible distribution. *Math. Methods Statist.* 6, 47–69. MR1456646

[17] Ingster, Y. (1999). Minimax detection of a signal for ℓ_n^p-balls. *Math. Methods Statist.* 7, 401–428. MR1680087

[18] Ingster, Y. and Stepanova, N. (2014). Adaptive variable selection in nonparametric sparse regression. *J. Math. Sci.* 199, 184–201. MR3032218

[19] Lehmann, E. and Romano, J. (2005). Generalizations of the familywise error rate. *Ann. Statist.* 33, 1138–1154. MR2195631

[20] Rabinovich, M., Ramdas, A., Jordan, M. and Wainwright M. (2020). Optimal rates and tradeoffs in multiple testing. *Statist. Sinica* 30, 741–762. MR4214160

[21] Roquain, É. and Verzelen, N. (2022). False discovery rate control with unknown null distribution: is it possible to mimic the oracle? *Ann. Statist.* 50, 1095–1123. MR4404913

[22] Salomond, J.-B. (2017). Risk quantification for the thresholding rule for multiple testing using Gaussian scale mixtures. ArXiv:1711.08705.

[23] Sarkar, S. K., T. Zhou, T. and D. Ghosh, D. (2008). A general decision theoretic formulation of procedures controlling FDR and FNR from a Bayesian perspective. *Statist. Sinica* 18, 925–945. MR2440399