

Belloni, Alexandre; Chernozhukov, Victor; Chetverikov, Denis; Hansen, Christian Bailey; Kato, Kengo

Working Paper

High-dimensional econometrics and regularized GMM

cemmap working paper, No. CWP35/18

Provided in Cooperation with:

Institute for Fiscal Studies (IFS), London

Suggested Citation: Belloni, Alexandre; Chernozhukov, Victor; Chetverikov, Denis; Hansen, Christian Bailey; Kato, Kengo (2018) : High-dimensional econometrics and regularized GMM, cemmap working paper, No. CWP35/18, Centre for Microdata Methods and Practice (cemmap), London,
<https://doi.org/10.1920/wp.cem.2018.3518>

This Version is available at:

<https://hdl.handle.net/10419/189751>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

High-dimensional econometrics and regularized GMM

Alexandre Belloni
Victor Chernozhukov
Denis Chetverikov
Christian Hansen
Kengo Kato

The Institute for Fiscal Studies
Department of Economics, UCL

cemmap working paper CWP35/18

HIGH-DIMENSIONAL ECONOMETRICS AND REGULARIZED GMM

BY ALEXANDRE BELLONI, VICTOR CHERNOZHUKOV, DENIS CHETVERIKOV, CHRISTIAN HANSEN,
AND KENGO KATO

ABSTRACT. This chapter presents key concepts and theoretical results for analyzing estimation and inference in high-dimensional models. High-dimensional models are characterized by having a number of unknown parameters that is not vanishingly small relative to the sample size. We first present results in a framework where estimators of parameters of interest may be represented directly as approximate means. Within this context, we review fundamental results including high-dimensional central limit theorems, bootstrap approximation of high-dimensional limit distributions, and moderate deviation theory. We also review key concepts underlying inference when many parameters are of interest such as multiple testing with family-wise error rate or false discovery rate control. We then turn to a general high-dimensional minimum distance framework with a special focus on generalized method of moments problems where we present results for estimation and inference about model parameters. The presented results cover a wide array of econometric applications, and we discuss several leading special cases including high-dimensional linear regression and linear instrumental variables models to illustrate the general results.

CONTENTS

1. Introduction	2
1.1. Inference with Many Approximate Means	3
1.2. Inference with Many Parameters and Moments	7
2. Inference with Many Approximate Means	10
2.1. Setting	10
2.2. CLT, Bootstrap, and Moderate Deviations	10
2.3. Simultaneous Confidence Intervals.	17
2.4. Multiple Testing with FWER Control	20
2.5. Multiple Testing with FDR Control	22

2.6. Inference on Functionals of Many Approximate Means	24
3. Estimation with Many Parameters and Moments	33
3.1. Regularized Minimum Distance Estimation Problem	33
3.2. Double/De-Biased RGMM	47
3.3. Construction of Approximate Mean Estimators for θ via Orthogonal Moments	47
3.4. Testing Parameters α with Nuisance Parameters $\eta = \theta$ via Neyman-Orthogonal Scores	49
3.5. Analysis of DRGMM	51
4. Bibliographical Notes and Open Problems	57
Appendix A. Fundamental Tools	64
A.1. Tool Set 1: Moderate Deviation Inequality for Self-normalized Sums	64
A.2. Tool Set 2: Maximal and Deviation Inequalities	64
A.3. Tool Set 3: High-dimensional Central Limit and Bootstrap Theorem, Gaussian Comparison and Anti-Concentration Inequalities	65
Appendix B. Proofs for Section 2	69
Appendix C. Proofs for Section 3	85
Appendix D. Technical Lemmas	94
References	98

1. INTRODUCTION

In this chapter, we review some of the main ideas and concepts from the literature on estimation and inference in high dimensions. High-dimensional models naturally arise in many contexts. First, empirical researchers may want to build more flexible models in an effort to approximate real phenomena better. Second, they may want to use more “flexible” controls to make conditional exogeneity more plausible in an effort to (more plausibly) identify causal/structural effects. Third, researchers may want to analyze policy effects

on very high-dimensional outcomes and/or across many groups. Fourth, researchers may wish to leverage high-dimensional exclusion restrictions (“many instruments”) in an effort to pin down structural parameters better. These and other contexts motivate the set of methods and results we overview in this chapter. In addition to providing an overview of useful tools, we develop some new results in order to make existing results more useful for applications in econometrics. We note that, since the literature on high-dimensional estimation and inference is large, we have opted to review only *some* of the main results from this literature. In this regard, our exposition complements other reviews, e.g. [9], [64], [33], [54]. For a textbook-level treatment, we refer an interested reader, for example, to [32], [69], [74], and [108].

High-dimensionality typically refers to a setting where the number of parameters in a model is non-negligible compared to the sample size available. The presence of a large number of parameters often necessitates us to design estimation and inference methods that are different from those used in classical, low-dimensional, settings. High-dimensional models have always been of interest in econometrics and have recently been gaining in popularity. The recent interest in these models is due to both the availability of rich, modern data sets and to advances in the analysis of high-dimensional settings, such as the emergence of high-dimensional central limit theorems and regularization and post-regularization methods for estimation and inference.

1.1. Inference with Many Approximate Means. We split the chapter into two parts. In the first part, we consider inference using the *Many Approximate Means* (MAM) framework. In particular, we assume that we have a potentially high-dimensional vector of parameters

$$\theta_0 := (\theta_{01}, \dots, \theta_{0p})' \in \mathbb{R}^p$$

and its estimator

$$\hat{\theta} := (\hat{\theta}_1, \dots, \hat{\theta}_p)' \in \mathbb{R}^p,$$

having an approximately linear form,

$$\sqrt{n}(\hat{\theta} - \theta_0) = \frac{1}{\sqrt{n}} \sum_{i=1}^n Z_i + r_n, \quad (1.1)$$

where Z_1, \dots, Z_n are independent zero-mean random vectors in \mathbb{R}^p , sometimes referred to as “influence functions”, and $r_n \in \mathbb{R}^p$ is a vector of linearization errors that are asymptotically negligible; see the next section for the formal requirement. The vectors Z_1, \dots, Z_n are either directly observable or can be consistently estimated. Here, we allow for the case $p \gg n$.

This framework is rather general and covers, in particular, the case of testing multiple means with $r_n = 0$. More generally, this framework covers multiple linear and non-linear

M -estimators and also accommodates many de-biased estimators; see, e.g., [75] for explicit conditions giving rise to linearization (1.1) in low-dimensional settings and [19] for conditions in the high-dimensional settings with $p \gg n$.

While conceptually easy-to-understand, the MAM framework allows us to present fundamental concepts in high-dimensional settings:

1. Simultaneous inference.
2. Inference with False Discovery Rate control.
3. Estimation based on ℓ_1 -regularization.

The first concept here includes simultaneous confidence interval construction for all (or some) components of the vector $\theta_0 = (\theta_{01}, \dots, \theta_{0p})'$. As we explain in the next section, constructing *simultaneous* confidence intervals is especially important in the high-dimensional settings and we explain how to construct such intervals. This concept also includes multiple testing with family-wise error rate (FWER) control, where we simultaneously test hypotheses about different components of the vector θ_0 and we want to make sure that the probability of at least one false null rejection does not exceed the pre-specified level α . The second concept includes multiple testing with false discovery rate (FDR) control, where we simultaneously test hypotheses about different components of θ_0 and we want to make sure that the fraction of falsely rejected null hypotheses among all rejected null hypotheses does not exceed the pre-specified level α , at least in expectation. FDR control is more liberal than FWER control, so procedures with FDR control typically have larger power than those with FWER control. This higher power may be particularly important, for example, in geno-economics, where procedures with FWER control often fail to find any association between the outcome variables and genes; see Example 3 below for the details. The third concept includes estimation of linear functionals of the vector θ_0 . We will show that estimating such functionals sometimes requires forms of regularization, and we will explain the details of ℓ_1 -regularization. This discussion will prepare us for the more ambitious problems arising in the second part of the chapter.

To perform the tasks described above, we will use some fundamental tools:

- I. High-Dimensional Central Limit Theorem (with $p \gg n$)
- II. Moderate Deviation Theorem (Central Limit Theorem over Tail Areas)
- III. Regularization (focusing on ℓ_1 -type regularization)

One of the main goals of the first part of this chapter will therefore be to provide statements and discussion of these key tools in a simple but interesting framework. Outside of being useful in the MAM framework, these tools play an important role in the theory of high-dimensional estimation and inference more generally.

We next review some simple motivating examples that fall into the MAM framework.

Example 1 (Randomized Control Trials with Many Outcome Variables). Consider a randomized control trial with n participants, where each participant $i = 1, \dots, n$ is randomly assigned to either the treatment group ($D_i = 1$) or the control group ($D_i = 0$). Let γ denote the probability of being assigned to the treatment group and suppose that for each participant i , we observe a large number of outcome variables represented by a vector $Y_i = (Y_{i1}, \dots, Y_{ip})' \in \mathbb{R}^p$, which is often the case in practice. For each outcome variable $j = 1, \dots, p$, we then can estimate the average treatment effect

$$\theta_{0j} = E[Y_{ij} \mid D_i = 1] - E[Y_{ij} \mid D_i = 0]$$

by

$$\hat{\theta}_j = \frac{1}{n} \sum_{i=1}^n \left(\frac{D_i Y_{ij}}{\gamma} - \frac{(1 - D_i) Y_{ij}}{1 - \gamma} \right).$$

Clearly, this setting falls into the MAM framework with

$$Z_{ij} = \left(\frac{D_i Y_{ij}}{\gamma} - E[Y_{ij} \mid D_i = 1] \right) - \left(\frac{(1 - D_i) Y_{ij}}{1 - \gamma} - E[Y_{ij} \mid D_i = 0] \right)$$

for all $i = 1, \dots, n$ and $j = 1, \dots, p$. We note also that the MAM framework covers many other, more complicated, treatment effect settings, beyond randomized control trials, since treatment effect estimators are often asymptotically linear; e.g. see [78], [76], [1], and [44], among many others. ■

Example 2 (Randomized Control Trials with Many Groups). Consider the same randomized control trial as in the previous example and suppose that for each participant $i = 1, \dots, n$, we have only one outcome variable $Y_i \in \mathbb{R}$ but we observe several discrete covariates, represented by a vector $X_i = (X_{i1}, \dots, X_{id})' \in \mathbb{R}^d$. For simplicity, we can assume that each covariate is binary, so that $X_{ik} \in \{0, 1\}$ for all $k = 1, \dots, d$. In this case, the vector X_i can take $p = 2^d$ different values, denoted by x_1, \dots, x_p , and we can split all participants into p groups depending on their values of X_i . Suppose, also for simplicity, that each group consists of $\bar{n} = n/p$ participants, i.e. all groups are equal in size. We then can estimate group-specific average treatment effects

$$\theta_{0j} = E[Y_i \mid X_i = x_j, D_i = 1] - E[Y_i \mid X_i = x_j, D_i = 0]$$

by

$$\hat{\theta}_j = \frac{1}{\bar{n}} \sum_{i: X_i = x_j} \left(\frac{D_i Y_i}{\gamma} - \frac{(1 - D_i) Y_i}{1 - \gamma} \right),$$

and this setting again falls into the MAM framework, with n replaced by \bar{n} . Note also that since the number of groups, $p = 2^d$, is exponential in the number of covariates, d , it is likely that $p \gg \bar{n}$ or at least $p \sim n$, making our analysis in this chapter particularly relevant.

To give a specific example of an experiment with many groups, consider the Tennessee Student Teacher Achievement Ratio (STAR) project conducted from 1985-89 and studied, e.g. in [83] among many others. In this project, over 11000 students from kindergarten to

third grade in 79 schools were randomly assigned into small (13 to 17 students) or regular (22 to 25 students) classes. Classroom teachers were also randomly assigned to classes. Different student achievements were subsequently measured over many years. The project also collected many demographic variables characterizing students, teachers, and schools. For example, available data include gender (male or female) and race (white, black, asian, hispanic, native american, or other) for both students and teachers, month of birth (from Jan to Dec) for students, and educational achievement (associate, bachelor, master, master+, specialist, or doctoral) and years of teaching (from 0-42) for teachers. All these characteristics can be used to form a large number of groups of student-teacher pairs. ■

Example 3 (Genoeconomics). In genoeconomics, a field that combines genetics and economics, researchers are interested in studying how genes affect economic behavior. This field is of interest because genetic information, for example, can provide direct measures of preferences of economic agents and can serve as a source of exogenous variation. The vast majority of the humane genome is the same among all humans, and the differences occur “only” in around 52 million SNPs (single-nucleotide polymorphisms). Most SNPs take only three values, (0, 1, 2), and modern technologies allow measuring the values of many, if not all, of these SNPs with minimal costs. The datasets in genoeconomics, therefore, often take the following form: We have a random sample of n humans, where n is of order of hundreds or thousands, and for each human $i = 1, \dots, n$, we have an outcome variable $Y_i \in \mathbb{R}$ and the vector of SNP values, $X_i = (X_{i1}, \dots, X_{ip})' \in \mathbb{R}^p$, where p can be of order of thousands or even millions. To measure the association between the outcome variable and the SNP $j = 1, \dots, p$, we can use the slope coefficient θ_j in the linear regression

$$Y_i = \alpha_j + \theta_{0j}X_{ij} + \epsilon_{ij}, \quad E[\epsilon_{ij} | X_{ij}] = 0,$$

which can be estimated by

$$\hat{\theta}_j = \frac{\sum_{i=1}^n (X_{ij} - \bar{X}_j)(Y_i - \bar{Y})}{\sum_{i=1}^n (X_{ij} - \bar{X}_j)^2},$$

where $\bar{X}_j = \sum_{i=1}^n X_{ij}/n$ and $\bar{Y} = \sum_{i=1}^n Y_i/n$. Clearly, this setting falls into the MAM framework with

$$Z_{ij} = \frac{(X_{ij} - E[X_{ij}])(Y_i - E[Y_i])}{\text{Var}(X_{ij})}$$

for all $i = 1, \dots, n$ and $j = 1, \dots, p$. We refer the reader to [25] for more detailed discussion of genoeconomics. ■

Example 4 (Structural Models with Many Parameters). The examples above outline simple cases where parameters are estimated either by sample means (Examples 1 and 2) or by quantities that can be easily approximated by sample means (Example 3). In structural econometrics, we often use more sophisticated estimators of structural parameters, such as GMM. In the second part of the chapter, we therefore develop a high-dimensional regularized GMM estimator. This could be of interest, for example, in demand elasticity estimation, where the elasticity parameter θ_{0j} varies across product groups (or product characteristics) $j = 1, \dots, p$. We show that it is possible to construct asymptotically

unbiased estimators of these parameters using the double/de-biased regularized GMM approach. These estimators are asymptotically linear and fall into the MAM framework. We therefore can use inferential tools developed for the MAM framework to construct simultaneous confidence intervals and conduct multiple hypothesis testing using various approaches we discuss in the first part of the chapter. ■

1.2. Inference with Many Parameters and Moments. In the second part of the chapter, we study estimation and inference in the high-dimensional GMM setting, where both the number of moment equations and the dimensionality of the parameter of interest may be large. Specifically, we consider a random vector $X \in \mathbb{R}^{d_x}$, a vector of parameters $\theta \in \mathbb{R}^p$, and a vector-valued score function $g(X, \theta)$ mapping $\mathbb{R}^{d_x} \times \mathbb{R}^p$ into \mathbb{R}^m , for some $m \geq p$. For the moment function

$$g(\theta) := E[g(X, \theta)], \quad (1.2)$$

we assume that the true parameter value θ_0 satisfies

$$g(\theta_0) = 0. \quad (1.3)$$

We are then interested in estimating θ_0 and carrying out inference on θ_0 using a random sample X_1, \dots, X_n from the distribution of X . We allow for the case $m \gg n$ and $p \gg n$.

We develop a Regularized GMM estimator (RGMM) of θ_0 and study its properties under various structural assumptions, such as sparsity or approximate sparsity of θ_0 . This novel estimator extends the Dantzig selector of Candès and Tao in [36] that was developed specifically for estimating linear mean regression models.

To gain intuition behind the RGMM estimator, we also consider a general minimum distance estimation problem, where the parameter θ_0 is known to satisfy (1.3) but the function $g(\theta)$ does not necessarily take the form (1.2). Assuming that an estimator of $g(\theta)$ is available, we formulate a Regularized Minimum Distance (RMD) estimator of θ_0 and develop its properties under easily-interpretable high-level conditions. Specializing these conditions for the GMM setting then allows us to derive properties of the RGMM estimator under relatively low-level conditions.

Like other estimators developed for high-dimensional models, such as Lasso, the RGMM estimator is suitable for coping high-dimensionality of the problem but has a complicated asymptotic distribution, making inference based on this estimator problematic. We therefore also develop a Double/Debiased RGMM estimator (DRGMM) that is asymptotically linear, and thus fits into the MAM framework, reemphasizing the role of the MAM framework, and reducing the problem of inference on θ_0 to our analysis in the first part of the paper. Importantly, for our DRGMM estimator, we also consider a version with the optimal weighting matrix.

Example 5 (Linear Mean Regression Model). One of the simplest examples falling into the GMM framework is the linear mean regression model,

$$Y = W'\theta_0 + \epsilon, \quad E[\epsilon | W] = 0, \quad (1.4)$$

where $Y \in \mathbb{R}$ is an outcome variable, $W \in \mathbb{R}^p$ a vector of covariates, $\epsilon \in \mathbb{R}$ noise, and $\theta_0 \in \mathbb{R}^p$ a parameter of interest. This model fits into the GMM framework with

$$g(X, \theta) = (Y - W'\theta)W, \quad X = (W, Y),$$

which corresponds to the most widely studied case in the literature on high-dimensional models, e.g. the Lasso estimator of Tibshirani [107] and the Dantzig Selector of Candès and Tao [36]. More generally, we can take a vector-valued function $t: \mathbb{R}^p \rightarrow \mathbb{R}^m$ with $m \geq p$, consider a vector $t(W)$ of transformations of W , and set

$$g(X, \theta) = (Y - W'\theta)t(W), \quad X = (W, Y).$$

By considering a sufficiently rich vector of functions t and using optimally-weighted GMM, we can construct an estimator that achieves the semiparametric efficiency bound.

In this example, as well as in Examples 6 and 7 below, we are often interested in a *low*-dimensional sub-vector of θ_0 corresponding to the covariates of interest in the vector W but sometimes the whole vector θ_0 or some *high*-dimensional sub-vector of θ_0 is of interest as well. For example, suppose that $W = (P', D \cdot P)'$, where $D \in \{0, 1\}$ is a binary treatment variable and $P \in \mathbb{R}^{p/2}$ is a high-dimensional vector of controls, so that the model (1.4) becomes

$$Y = P'\alpha_0 + D \cdot P'\beta_0 + \epsilon, \quad E[\epsilon | P, D] = 0$$

with $\theta_0 = (\alpha_0', \beta_0')'$, where both α_0 and β_0 are $(p/2)$ -dimensional vectors of parameters. Assuming that D is randomly assigned conditional on P then implies that $P'\beta_0$ is the Conditional Average Treatment Effect (CATE) for the outcome variable Y and that β_0 is the vector of derivatives of the CATE with respect to the vector of controls P . Thus, the whole vector β_0 may be of interest. ■

Example 6 (Linear IV Regression Model). Consider the linear IV regression model

$$Y = W'\theta_0 + \epsilon, \quad E[\epsilon | Z] = 0,$$

where we use the same notation as above and, in addition, $Z \in \mathbb{R}^{dz}$ is a vector of instruments. This model fits into the GMM framework with

$$g(X, \theta) = (Y - W'\theta)Z, \quad X = (W, Z, Y),$$

or, more generally,

$$g(X, \theta) = (Y - W'\theta)t(Z), \quad X = (W, Z, Y),$$

where $t: \mathbb{R}^{dz} \rightarrow \mathbb{R}^m$ is a vector-valued function. The case where p is small and m is large (larger than the sample size n) was originally studied in [7] and the case where both p and m are large is considered in [54], [68], [17], and [70]. One of the novel parts of our analysis is that we can allow for the optimal weighting of the moment conditions. ■

Example 7 (Nonlinear IV Regression Model). More generally, consider a non-linear IV regression model

$$\mathbb{E}[f(Y, W'\theta_0) \mid Z] = 0,$$

where we use the same notation as in Example 6 and, in addition, $f: \mathbb{R}^2 \rightarrow \mathbb{R}$ is some function. This can be of interest, for example, in the analysis of discrete choice models where Y is binary (or, more generally, discrete). In the same fashion as above, this model fits into the GMM framework with

$$g(X, \theta) = f(Y, W'\theta)Z, \quad X = (W, Z, Y)$$

or, more generally,

$$g(X, \theta) = f(Y, W'\theta)t(Z),$$

where $t: \mathbb{R}^{dz} \rightarrow \mathbb{R}^m$ is a vector-valued function. ■

Notation. In what follows, all models and probability measures P can be indexed by the sample size n , so that models and their dimensions can change with n , allowing dimensionality to increase with n . We use “wp $\rightarrow 1$ ” to abbreviate the phrase “with probability that converges to 1”, and we use arrows \rightarrow_P and \rightsquigarrow_P to denote convergence in probability and in distribution, respectively. The symbol \sim means “distributed as”. The notation $a \lesssim b$ means $a = O(b)$ and $a \lesssim_P b$ means $a = O_P(b)$. We also use the notation $a \vee b = \max\{a, b\}$ and $a \wedge b = \min\{a, b\}$. For any $a \in \mathbb{R}$, $\lfloor a \rfloor$ denotes the largest integer that is smaller than or equal to a , and $\lceil a \rceil$ denotes the smallest integer that is larger than or equal to a . For a positive integer m , $[m] = \{1, \dots, m\}$.

Next, for any vector $x = (x_1, \dots, x_p)'$, we denote the ℓ_1 and ℓ_2 norms of x by $\|x\|_1 = \sum_{j=1}^p |x_j|$ and $\|x\|_2 = (\sum_{j=1}^p x_j^2)^{1/2}$, respectively. The ℓ_0 -“norm” of x , $\|x\|_0$, denotes the number of non-zero components of the vector x . Moreover, for any vector $x = (x_1, \dots, x_p)'$ in \mathbb{R}^p and any set of indices $T \subset \{1, \dots, p\}$, we use $x_T = (x_{T_1}, \dots, x_{T_p})'$ to denote the vector in \mathbb{R}^p such that $x_{T_j} = x_j$ for $j \in T$ and $x_{T_j} = 0$ for $j \in T^c$, where $T^c = \{1, \dots, p\} \setminus T$. For any matrix A of p columns, we use $\|A\|$ to denote the operator norm of A : $\|A\| = \sup_{x \in \mathbb{R}^p, \|x\|_2=1} \|Ax\|_2$.

The transpose of a column vector x is denoted by x' . For a differentiable map $\mathbb{R}^d \ni x \mapsto f(x) \in \mathbb{R}^k$, we use $\partial_{x'} f$ to denote the $k \times d$ Jacobian matrix $\partial f / \partial x' = (\partial f_i / \partial x_j)_{1 \leq i \leq k, 1 \leq j \leq d}$, and we correspondingly use the expression $\partial_{x'} f(x_0)$ to denote $\partial_{x'} f(x) \big|_{x=x_0}$, etc. When we have an event A whose occurrence depends on two independent random vectors, X and Y , we use $P_X(A)$ to denote the probability of A with respect to the distribution of X , holding Y fixed. For given Z_1, \dots, Z_n , we use the notation $Z_1^n = (Z_1, \dots, Z_n)$. We use Φ and ϕ to denote the cdf and pdf of the standard normal distribution.

Finally, we use standard empirical process theory notation. In particular, $\mathbb{E}_n[\cdot]$ abbreviates the average $n^{-1} \sum_{i=1}^n [\cdot]$ over index $i = 1, \dots, n$, e.g. $\mathbb{E}_n[f(z_i)]$ denotes $n^{-1} \sum_{i=1}^n f(z_i)$. Also, if Z is a random vector with law P and support \mathcal{Z} , $(Z_i)_{i \in [n]}$ is a random sample from

the distribution of Z , and \mathcal{F} is a class of functions $f: \mathcal{Z} \rightarrow \mathbb{R}$, then $\mathbb{G}_n f := \mathbb{G}_n f(Z) := n^{-1/2} \sum_{i=1}^n (f(Z_i) - \mathbb{E}[f(Z)])$ for all $f \in \mathcal{F}$.

2. INFERENCE WITH MANY APPROXIMATE MEANS

2.1. Setting. Suppose that we have a parameter $\theta_0 = (\theta_{01}, \dots, \theta_{0p})' \in \mathbb{R}^p$ and an estimator $\hat{\theta} = (\hat{\theta}_1, \dots, \hat{\theta}_p)' \in \mathbb{R}^p$ of this parameter that has an approximately linear form:

$$\sqrt{n}(\hat{\theta} - \theta_0) = \frac{1}{\sqrt{n}} \sum_{i=1}^n Z_i + r_n, \quad (2.1)$$

where Z_1, \dots, Z_n are independent zero-mean random vectors in \mathbb{R}^p , sometimes called the “influence functions,” and $r_n = (r_{n1}, \dots, r_{np})' \in \mathbb{R}^p$ is a vector of linearization errors that are small in the sense that

$$\|r_n\|_\infty = o_P\left(1/\sqrt{\log(pn)}\right), \quad (2.2)$$

with a more precise requirement provided in Condition A. The vectors Z_1, \dots, Z_n may not be directly observable, and we assume some estimators $\hat{Z}_1, \dots, \hat{Z}_n$ of these vectors are available in this case. In this section, we are interested in carrying out different types of inference on θ_0 . We are primarily interested in the case where p is larger or much larger than n , but the results below apply when p is smaller than n as well. Throughout the chapter, we refer to this setting as the *Many Approximate Means* (MAM) framework.

In this section, we review results from the literature on the high-dimensional Central Limit Theorem (CLT), high-dimensional bootstrap theorems, moderate deviations for self-normalized sums, simultaneous confidence intervals, multiple testing with the Family-Wise Error Rate (FWER) control, and multiple testing with the False Discovery Rate (FDR) control. All results to be reviewed below exist in the literature for the case of many *exact* means, where the approximation errors are not present, $r_n = 0$, and the vectors Z_i are observed. We extend these results to allow for many *approximate* means and also for unobservable but estimable vectors Z_i , i.e. we extend the results to cover the MAM framework. This extension is important because many estimators we work with are *asymptotically* linear but do not have to be linear in finite samples.

At the end of this section, we also consider the problem of estimating linear functionals of θ_0 , which motivates such concepts as sparsity and ℓ_1 -regularization and prepares us for the discussion in the second part of the chapter.

2.2. CLT, Bootstrap, and Moderate Deviations. To perform inference on θ_0 , we first need to develop a distributional approximation for

$$\sqrt{n}(\hat{\theta} - \theta_0).$$

When p is fixed and n gets large, $\sqrt{n}(\hat{\theta} - \theta_0)$ converges in distribution to a zero-mean Gaussian random vector by a classical CLT but here we are interested in the case with $p \rightarrow \infty$ or even $p/n \rightarrow \infty$ as $n \rightarrow \infty$ making classical CLTs inapplicable. We therefore rely on the high-dimensional CLT developed in [45, 49, 50]. To state the result, and also to extend it to allow for the MAM framework, we will use the following regularity conditions. Let $(B_n)_{n \geq 1}$, $(\delta_n)_{n \geq 1}$, and $(\beta_n)_{n \geq 1}$ be given sequences of constants satisfying $B_n \geq 1$, $\delta_n \searrow 0$, and $\beta_n \searrow 0$. Here, B_n is allowed to grow to infinity as n gets large.

Condition M. (i) $n^{-1} \sum_{i=1}^n \mathbb{E}[Z_{ij}^2] \geq 1$ for all $j \in [p]$ and (ii) $n^{-1} \sum_{i=1}^n \mathbb{E}[|Z_{ij}|^{2+k}] \leq B_n^k$ for all $j \in [p]$ and $k = 1, 2$.

Since r_n is asymptotically negligible, in the sense that (2.2) holds, it follows from (2.1) that, for all $j \in [p]$, the asymptotic variance of $\sqrt{n}(\hat{\theta}_j - \theta_{0j})$ is equal to $n^{-1} \sum_{i=1}^n \mathbb{E}[Z_{ij}^2] \geq 1$. Thus, the first part of Condition M requires that this variance is bounded away from zero. Such a condition precludes existence of super-efficient estimators and is typically imposed even in classical settings, where p is small relative to n . The second part of Condition M imposes the mild requirement that $n^{-1} \sum_{i=1}^n \mathbb{E}[|Z_{ij}|^3]$ and $n^{-1} \sum_{i=1}^n \mathbb{E}[|Z_{ij}|^4]$ do not increase too quickly with n .

Condition E. Either of the following moment bounds holds:

- E.1 $\mathbb{E}[\exp(|Z_{ij}|/B_n)] \leq 2$ for all $i \in [n]$ and $j \in [p]$, and $\left(\frac{B_n^2 \log^7(pn)}{n}\right)^{1/6} \leq \delta_n$, or
 E.2 $\mathbb{E}[\max_{j \in [p]} (|Z_{ij}|/B_n)^4] \leq 1$ for all $i \in [n]$, and $\left(\frac{B_n^4 \log^7(pn)}{n}\right)^{1/6} \leq \delta_n$.

The first part of Condition E.1 requires that Z_{ij} 's have light tails. In particular, under Condition E.1, the tails have to be sub-exponential:

$$\mathbb{P}(|Z_{ij}| > x) = \mathbb{P}(\exp(|Z_{ij}|/B_n) > \exp(x/B_n)) \leq 2 \exp(-x/B_n), \quad \text{for all } x > 0,$$

by Markov's inequality. In fact, Lemma 2.2.1 in [113] shows that if

$$\mathbb{P}(|Z_{ij}| > x) \leq 2 \exp(-x/C_n), \quad \text{for all } x > 0,$$

for some $C_n > 0$, then $\mathbb{E}[\exp(|Z_{ij}|/B_n)] \leq 2$ holds with $B_n = 3C_n$. Thus, the first part of Condition E.1 is equivalent to all Z_{ij} 's having sub-exponential tails. The first part of Condition E.2, on the other hand, allows for heavy-tails but imposes some moment conditions on $\max_{j \in [p]} |Z_{ij}|$. Conditions E.1 and E.2 are therefore non-nested. The second parts of Conditions E.1 and E.2 impose restrictions on how fast B_n and p can grow. Note that we never impose E.1 and E.2 simultaneously.

Condition A. (i) The linearization errors obey $\mathbb{P}(\max_{j \in [p]} |r_{nj}| > \delta_n / \log^{1/2}(pn)) \leq \beta_n$, and (ii) the estimates of the influence functions obey $\mathbb{P}(\max_{j \in [p]} \mathbb{E}_n[(\hat{Z}_{ij} - Z_{ij})^2] > \delta_n^2 / \log^2(pn)) \leq \beta_n$.

The first part of Condition A requires that the approximation errors in the vector r_n are asymptotically negligible, and clarifies (2.2). Note that if (2.2) holds, then it is rather standard to show that there exist *some* sequences of positive constants $(\delta_n)_{n \geq 1}$ and $(\beta_n)_{n \geq 1}$ satisfying $\delta_n \searrow 0$ and $\beta_n \searrow 0$ such that the first part of Condition A holds. The second part of Condition A requires the estimators \hat{Z}_{ij} of Z_{ij} to be sufficiently precise. Again, if \hat{Z}_{ij} 's satisfy

$$\max_{j \in [p]} \sqrt{\mathbb{E}_n[(\hat{Z}_{ij} - Z_{ij})^2]} = o_P(1/\log(pn)), \quad (2.3)$$

then there exist *some* $(\delta_n)_{n \geq 1}$ and $(\beta_n)_{n \geq 1}$ satisfying $\delta_n \searrow 0$ and $\beta_n \searrow 0$ such that the second part of Condition A holds.

In order to state a key CLT result, let \mathcal{A} be the class of all (closed) rectangles in \mathbb{R}^p , i.e. sets A of the form

$$A = \left\{ w = (w_1, \dots, w_p)' \in \mathbb{R}^p : w_{l_j} \leq w_j \leq w_{r_j} \text{ for all } j \in [p] \right\},$$

where $w_l = (w_{l_1}, \dots, w_{l_p})'$ and $w_r = (w_{r_1}, \dots, w_{r_p})'$ are two vectors such that $w_{l_j} \leq w_{r_j}$ for all $j \in [p]$. (Here, both w_{l_j} and w_{r_j} can take values of $-\infty$ or $+\infty$.) Denote $V := n^{-1} \sum_{i=1}^n \mathbb{E}[Z_i Z_i']$, and let $N(0, V)$ be a zero-mean Gaussian random vector in \mathbb{R}^p with covariance matrix V . The following theorem establishes the Gaussian approximation for the distribution of $\sqrt{n}(\hat{\theta} - \theta_0)$, which extends Proposition 2.1 in [50] to allow for many *approximate* means.

Theorem 2.1 (CLT for Many Approximate Means). *Under Conditions M, E, and A, the distribution of $\sqrt{n}(\hat{\theta} - \theta_0)$ over rectangles is approximately Gaussian:*

$$\sup_{A \in \mathcal{A}} \left| \mathbb{P}(\sqrt{n}(\hat{\theta} - \theta_0) \in A) - \mathbb{P}(N(0, V) \in A) \right| \leq C(\delta_n + \beta_n), \quad (2.4)$$

where C is a universal constant.

It is useful to note that Theorem 2.1 allows p to be larger or much larger than n . For example, the theorem implies that if Z_i 's are i.i.d zero-mean random vectors with each component bounded in absolute value by some constant C (independent of n) almost surely and the variance of each component bounded from below by one, then

$$\sup_{A \in \mathcal{A}} \left| \mathbb{P}(\sqrt{n}(\hat{\theta} - \theta_0) \in A) - \mathbb{P}(N(0, V) \in A) \right| \rightarrow 0 \text{ as } n \rightarrow \infty \quad (2.5)$$

as long as $\log^7 p = o(n)$ and (2.2) and (2.3) hold. Thus, Theorem 2.1 shows that Gaussian approximation over rectangles is possible even if p is *exponentially* large in n .

Note, however, that the Gaussian approximation here is stated only for the probability of $\sqrt{n}(\hat{\theta} - \theta_0)$ hitting rectangles $A \in \mathcal{A}$. The same Gaussian approximation may not hold if we look at more general classes of sets, e.g. all (Borel measurable) convex sets. In fact, it is known that if we replace the class of all rectangles \mathcal{A} in (2.5) by the class of all Borel

measurable convex sets, then we must assume that $p = o(n^{1/3})$, meaning $p \ll n$, in order to satisfy (2.5); see discussion on p. 2310 of [50]. On the other hand, as we will see below, the class of all rectangles is large enough to make Theorem 2.1 useful in many applications. See also Remark 2.2 below on how we can extend the class of rectangles and still allow for $p \gg n$.

The Gaussian approximation result of Theorem 2.1 is useful as applications below indicate, but does not immediately give a practical distributional approximation since the covariance matrix V is typically unknown. We therefore also consider bootstrap approximations. In particular, we consider the Gaussian (or multiplier) and empirical (or non-parametric) types of bootstrap. For the Gaussian bootstrap, let $e = (e_1, \dots, e_n)'$ be a vector consisting of i.i.d. $N(0, 1)$ random variables independent of the data yielding the estimator $\hat{\theta}$. A Gaussian bootstrap draw of the estimator $\hat{\theta}$ is then defined as

$$\sqrt{n}(\hat{\theta}^* - \hat{\theta}) := \frac{1}{\sqrt{n}} \sum_{i=1}^n e_i \hat{Z}_i \quad \text{or} \quad \hat{\theta}^* = \hat{\theta} + \frac{1}{n} \sum_{i=1}^n e_i \hat{Z}_i. \quad (2.6)$$

Alternatively, letting $e = (e_1, \dots, e_n)'$ be a vector following the multinomial distribution with parameters n and success probabilities $1/n, \dots, 1/n$ independent of the data, we can define an empirical bootstrap draw of the estimator $\hat{\theta}$ as

$$\sqrt{n}(\hat{\theta}^* - \hat{\theta}) := \frac{1}{\sqrt{n}} \sum_{i=1}^n (e_i - 1) \hat{Z}_i \quad \text{or} \quad \hat{\theta}^* = \hat{\theta} + \frac{1}{n} \sum_{i=1}^n (e_i - 1) \hat{Z}_i. \quad (2.7)$$

Equivalently, the empirical bootstrap draw of $\hat{\theta}$ can be constructed as

$$\sqrt{n}(\hat{\theta}^* - \hat{\theta}) = \frac{1}{\sqrt{n}} \sum_{i=1}^n (\hat{Z}_i^* - \bar{\hat{Z}}) \quad \text{or} \quad \hat{\theta}^* = \hat{\theta} + \frac{1}{n} \sum_{i=1}^n (\hat{Z}_i^* - \bar{\hat{Z}}), \quad (2.8)$$

where $\hat{Z}_1^*, \dots, \hat{Z}_n^*$ are an i.i.d. sample from the empirical distribution of $\hat{Z}_1, \dots, \hat{Z}_n$, and $\bar{\hat{Z}} = n^{-1} \sum_{i=1}^n \hat{Z}_i$. Indeed, the latter expression (2.8) reduces to the former expression (2.7) by setting each e_i as the number of times that \hat{Z}_i is “redrawn” in the bootstrap sample, and the vector $e = (e_1, \dots, e_n)'$ then follows the multinomial distribution with parameters n and success probabilities $1/n, \dots, 1/n$ independent of the data.

The following theorems show that the distribution of the bootstrap draw with respect to e approximates the Gaussian distribution given in Theorem 2.1.

Theorem 2.2 (Gaussian Bootstrap for Many Approximate Means). *Under Conditions M, E, and A, the distribution of a $N(0, V)$ vector over rectangles can be approximated by the Gaussian bootstrap: there exists a universal constant C such that for the Gaussian bootstrap draw $\hat{\theta}^*$ given in (2.6),*

$$\sup_{A \in \mathcal{A}} \left| \mathbb{P}_e(\sqrt{n}(\hat{\theta}^* - \hat{\theta}) \in A) - \mathbb{P}(N(0, V) \in A) \right| \leq C\delta_n$$

holds with probability at least $1 - \beta_n - n^{-1}$ in the case of E.1 and $1 - \beta_n - (\log n)^{-2}$ in the case of E.2.

Theorem 2.3 (Empirical Bootstrap for Many Approximate Means). *Assume that Conditions M, E, and A are satisfied. In addition, assume that $\mathbb{P}(\max_{i \in [n]} \max_{j \in [p]} |\hat{Z}_{ij} - Z_{ij}| > 1) \leq \beta_n$. Then the distribution of a $N(0, V)$ vector over rectangles can be approximated by the empirical bootstrap: there exists a universal constant C such that for the empirical bootstrap draw $\hat{\theta}^*$ given in (2.7),*

$$\sup_{A \in \mathcal{A}} \left| \mathbb{P}_e(\sqrt{n}(\hat{\theta}^* - \hat{\theta}) \in A) - \mathbb{P}(N(0, V) \in A) \right| \leq C\delta_n$$

holds with probability at least $1 - 2\beta_n - n^{-1}$ in the case of E.1 and

$$\sup_{A \in \mathcal{A}} \left| \mathbb{P}_e(\sqrt{n}(\hat{\theta}^* - \hat{\theta}) \in A) - \mathbb{P}(N(0, V) \in A) \right| \leq C\delta_n \log^{1/3} n \quad (2.9)$$

holds with probability at least $1 - 2\beta_n - (\log n)^{-2}$ in the case of E.2.

Remark 2.1 (Comparison of Gaussian and Empirical Bootstraps). Comparing Theorems 2.2 and 2.3 suggests that the Gaussian bootstrap may be more accurate than the empirical bootstrap. However, it is important to remember that both theorems only give upper bounds on the distributional approximation errors, and so such a conjecture may or may not be valid. In fact, there is some evidence that the empirical bootstrap may be more accurate than the Gaussian bootstrap because the former is able to better match higher-order moments of Z_i 's; see [58]. ■

Remark 2.2 (Sparsely Convex Sets). We note that Theorems 2.1-2.3 can be extended to allow for somewhat more general classes of sets, beyond the class of rectangles. In particular, these theorems can be extended to allow for classes of *sparsely convex sets*. For an integer $s > 0$, we say that $A \subset \mathbb{R}^p$ is an s -sparsely convex set if there exists an integer $Q > 0$ and convex sets $A_q \subset \mathbb{R}^p$, $q \in [Q]$, such that $A = \bigcap_{q \in [Q]} A_q$ and the indicator function of each A_q , $w \mapsto 1\{w \in A_q\}$, depends on at most s components of its argument $w = (w_1, \dots, w_p)'$. Each rectangle, for example, is clearly a 1-sparsely convex set. An example of 2-sparsely convex set is

$$A = \left\{ w = (w_1, \dots, w_p)' \in \mathbb{R}^p : \max_{j, k \in [p]; j \neq k} (w_j^2 + w_k^2) \leq x \right\}, \quad x \geq 0.$$

Theorems 2.1-2.3 can be extended to allow \mathcal{A} to be the class of all s -sparsely convex sets as long as s is not too large; we refer to [50] for details in the case of many exact means and leave the case of many approximate means to future work. ■

Remark 2.3 (Weakening Condition E.2). We also note that the second part of Condition E above can be slightly weakened and generalized. In particular, it can be replaced by the following condition:

E.2' For some $q \in [4, \infty)$, $E [\max_{j \in [p]} (|Z_{ij}|/B_n)^q] \leq 1$ for all $i \in [n]$, and $\left(\frac{B_n^2 \log^7(pn)}{n}\right)^{1/6} + \left(\frac{B_n^2 \log^3(pn)}{n^{1-2/q}}\right)^{1/3} \leq \delta_n$.

If we use this alternative version of Condition E, Theorems 2.1-2.3 still hold but the constant C in these theorems then depend on q , whenever E.2' is used. The same remark also applies to all theorems below where Condition E is used. ■

Figure 1 illustrates Theorems 2.1, 2.2, and 2.3 for rectangles A of a particular type:

$$A = \left\{ w = (w_1, \dots, w_p)' \in \mathbb{R}^p : -x \leq w_j \leq x \text{ for all } j \in [p] \right\}, \quad x \geq 0.$$

Specifically, Figure 1 plots

$$P(\sqrt{n}\|\hat{\theta} - \theta_0\|_\infty \leq x) \text{ against } P(\|N(0, V)\|_\infty \leq x)$$

and

$$P(\sqrt{n}\|\hat{\theta} - \theta_0\|_\infty \leq x) \text{ against } P_e(\sqrt{n}\|\hat{\theta}^* - \hat{\theta}\|_\infty \leq x)$$

as x varies from 0 to ∞ for different values of n and p and a distribution of Z_i 's motivated by the problem of selecting the regularization parameter of the RMD estimator in Section 3, where $\hat{\theta}^*$ is either the Gaussian or the empirical bootstrap draw. The figure indicates that both Gaussian and bootstrap approximations in Theorems 2.1, 2.2, and 2.3 are rather precise.

Another useful result for inference in high-dimensional settings is a moderate deviation theorem for self-normalized sums, which we present below. This result typically leads to conservative inference but requires very weak moment conditions. In particular, it does not require Condition E.

Theorem 2.4 (Moderate Deviations for Many Approximate Means). *Assume that Conditions M and A are satisfied. Also, let $\bar{C} > 0$ be some constant and assume that $(2\bar{C})^3 B_n \log^{3/2}(pn)/\sqrt{n} \leq \delta_n \leq 1$. Then there exist constants n_0 and C depending only on \bar{C} such that*

$$\left| P\left(\frac{\sqrt{n}(\hat{\theta}_j - \theta_{0j})}{(\mathbb{E}_n[\hat{Z}_{ij}^2])^{1/2}} > x\right) - (1 - \Phi(x)) \right| \leq C\left((1 - \Phi(x))\delta_n + \beta_n + (pn)^{-1}\right) \quad (2.10)$$

for all $n \geq n_0$ and $1 \leq x \leq \bar{C} \log^{1/2}(pn)$. In addition,

$$P\left(\max_{j \in [p]} \frac{\sqrt{n}(\hat{\theta}_j - \theta_{0j})}{(\mathbb{E}_n[\hat{Z}_{ij}^2])^{1/2}} > \Phi^{-1}(1 - \alpha/p)\right) \leq \alpha + C\left(\alpha\delta_n + \beta_n + (pn)^{-1}\right) \quad (2.11)$$

for all $n \geq n_0$ and α such that $1 \leq \Phi^{-1}(1 - \alpha/p) \leq \bar{C} \log^{1/2}(pn)$.

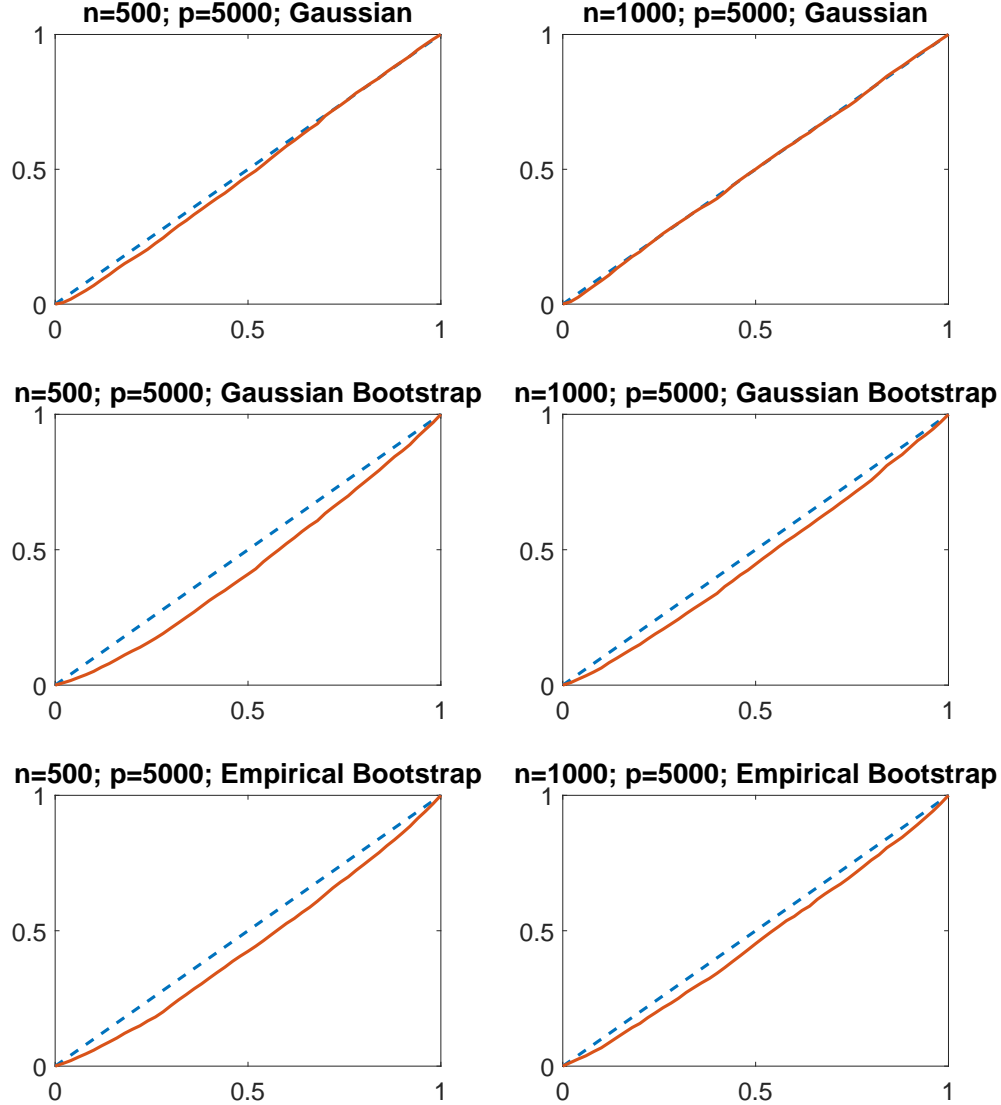


FIGURE 1. P-P plots comparing the distribution of $\max_{j \in [p]} \sqrt{n} |\hat{\theta}_j - \theta_{0j}|$ with its Gaussian, Gaussian bootstrap, and empirical bootstrap approximations in the example motivated by the problem of selecting the regularization parameter of the RMD estimator in Section 3. Here, Z_{ij} 's are generated as $Z_{ij} = W_{ij} \varepsilon_i$ with $\varepsilon_i \sim t(4)$, (the t -distribution with four degrees of freedom), and W_{ij} 's are non-stochastic (simulated once using $U[0, 1]$ distribution independently across i and j). We assume that Z_{ij} 's are observed and set $\hat{Z}_{ij} = Z_{ij}$ for all i and j . The dashed line is 45° . To generate bootstrap approximations, we use randomly selected sample of Z_{ij} 's. The figure indicate that all three approximations are good, and the quality of the approximation is particularly good for the tail probabilities, which is most relevant for practical applications.

Since for any $x > 0$, we have $1 - \Phi(x) \leq \exp(-x^2/2)$ by Proposition 2.5 in [60], it follows that $\Phi^{-1}(1 - 1/pn) \leq \sqrt{2 \log(pn)}$, and so setting $\alpha = 1/n$ in (2.11) gives the following corollary of Theorem 2.4:

Corollary 2.1 (Maximal Inequality for Many Approximate Means). *Assume that Conditions M and A are satisfied. Also, assume that $(2\sqrt{2})^3 B_n \log^{3/2}(pn)/\sqrt{n} \leq \delta_n \leq 1$. Then there exist universal constants n_0 and C such that for all $n \geq n_0$,*

$$\mathbb{P} \left(\max_{j \in [p]} \frac{\sqrt{n}(\hat{\theta}_j - \theta_{0j})}{(\mathbb{E}_n[\hat{Z}_{ij}^2])^{1/2}} > \sqrt{2 \log(pn)} \right) \leq C(\beta_n + n^{-1}).$$

In particular,

$$\frac{|\sqrt{n}(\hat{\theta}_j - \theta_{0j})|}{(\mathbb{E}_n[\hat{Z}_{ij}^2])^{1/2}} = O_P(\sqrt{\log(pn)}) \quad (2.12)$$

uniformly over $j \in [p]$.

If Z_{ij} 's are all bounded, or at least sub-Gaussian, it is straightforward to show by combining the union bound and exponential inequalities, such as those of Hoeffding or Bernstein, that

$$\frac{|\sqrt{n}(\hat{\theta}_j - \theta_{0j})|}{(V_{jj})^{1/2}} = O_P(\sqrt{\log(pn)}) \quad (2.13)$$

uniformly over $j \in [p]$. Comparing (2.12) and (2.13) now reveals an interesting feature of Theorem 2.4 and Corollary 2.1: replacing the *true value* V_{jj} of the asymptotic variance of $\sqrt{n}(\hat{\theta}_j - \theta_{0j})$ by an *estimator* $\mathbb{E}_n[\hat{Z}_{ij}^2]$ allows us to obtain the same bound, $O_P(\sqrt{\log(pn)})$, for the normalized version of $\sqrt{n}(\hat{\theta}_j - \theta_{0j})$ without imposing strong moment conditions on the data, such as boundedness, since Corollary 2.1 only assumes four finite moments of the Z_{ij} 's (via Condition M). Results of this form were used previously by [7] in the theory of high-dimensional estimation via Lasso to allow for noise with heavy tails. Also, [46] used such results to develop computationally efficient tests of many moment inequalities for heavy-tailed data.

2.3. Simultaneous Confidence Intervals. When only one θ_{0j} is of interest, it follows from standard arguments that

$$\frac{\sqrt{n}(\hat{\theta}_j - \theta_{0j})}{(\mathbb{E}_n[\hat{Z}_{ij}^2])^{1/2}} \rightsquigarrow_P N(0, 1)$$

as long as $r_{nj} = o_P(1)$ and $\mathbb{E}_n[(\hat{Z}_{ij} - Z_{ij})^2] = o_P(1)$. We can thus, e.g., construct a two-sided confidence interval for θ_{0j} with asymptotic coverage $1 - \alpha$ for some $\alpha \in (0, 1)$ as

$$CS_j(1 - \alpha) = \left[\hat{\theta}_j - (\mathbb{E}_n[\hat{Z}_{ij}^2])^{1/2} \frac{z_{\alpha/2}}{\sqrt{n}}; \hat{\theta}_j + (\mathbb{E}_n[\hat{Z}_{ij}^2])^{1/2} \frac{z_{\alpha/2}}{\sqrt{n}} \right],$$

where $z_{\alpha/2}$ denotes the $(1 - \alpha/2)$ -quantile of $N(0, 1)$, i.e., $1 - \Phi(z_{\alpha/2}) = \alpha/2$. However, the confidence intervals above are too optimistic when many components θ_{0j} of the parameter vector θ_0 are of interest, and it is likely that one or several θ_{0j} 's will fall out of their respective confidence intervals. Therefore, to obtain valid inferential statements, we need to carry out a multiplicity adjustment to explicitly take into account that many θ_{0j} are of interest. In this subsection, we demonstrate how to perform this adjustment and construct *simultaneous* confidence intervals for multiple components of θ_0 using Theorems 2.1–2.4.

The following quantity will play an important role in our analysis:

$$\lambda(1 - \alpha) := (1 - \alpha) \text{ quantile of } \|\sqrt{n}W(\hat{\theta} - \theta_0)\|_\infty,$$

where $W := \text{diag}(w_1, \dots, w_p)$ is a diagonal, potentially unknown, weighting matrix. For concreteness, for all $j \in [p]$, we often set $w_j = V_{jj}^{-1/2}$, which normalizes each $\sqrt{n}w_j(\hat{\theta}_j - \theta_{0j})$ to have asymptotic variance one, or $w_j = 1$, which simplifies the analysis.

If we knew $\lambda(1 - \alpha)$ and W , we would be able to use confidence intervals

$$CS = \prod_{j \in [p]} CS_j, \quad CS_j := \left[\hat{\theta}_j - \frac{\lambda(1 - \alpha)}{w_j \sqrt{n}}; \hat{\theta}_j + \frac{\lambda(1 - \alpha)}{w_j \sqrt{n}} \right],$$

since they clearly satisfy the desired coverage condition,

$$\mathbb{P}(\theta_{0j} \in CS_j \text{ for all } j \in [p]) = 1 - \alpha + o(1). \quad (2.14)$$

In practice, however, $\lambda(1 - \alpha)$ is typically unknown and has to be estimated from the data. To this end, let $\hat{W} := \text{diag}(\hat{w}_1, \dots, \hat{w}_p)$ be an estimator of W and let

$$\hat{\lambda}(1 - \alpha) := (1 - \alpha) \text{ quantile of } \|\sqrt{n}\hat{W}(\hat{\theta}^* - \hat{\theta})\|_\infty \mid \hat{W}, (\hat{Z}_i)_{i=1}^n,$$

where $\hat{\theta}^*$ is obtained via the Gaussian bootstrap, (2.6). The case where $\hat{\theta}^*$ is obtained via the empirical bootstrap, (2.7), can be analyzed similarly. We then can define feasible confidence intervals as

$$CS = \prod_{j \in [p]} CS_j, \quad CS_j := \left[\hat{\theta}_j - \frac{\hat{\lambda}(1 - \alpha)}{\hat{w}_j \sqrt{n}}; \hat{\theta}_j + \frac{\hat{\lambda}(1 - \alpha)}{\hat{w}_j \sqrt{n}} \right]. \quad (2.15)$$

Below, we will show that these feasible confidence intervals still satisfy (2.14), under certain regularity conditions. To prove this claim, we impose the following condition:

Condition W. (i) For some $C_W \geq 1$, the diagonal elements of the matrix $W = \text{diag}(w_1, \dots, w_p)$ satisfy $C_W^{-1}V_{jj}^{-1/2} \leq w_j \leq C_W$ for all $j \in [p]$, and (ii) the estimator $\hat{W} = \text{diag}(\hat{w}_1, \dots, \hat{w}_p)$ of the matrix $W = \text{diag}(w_1, \dots, w_p)$ satisfies $\mathbb{P}(\max_{j \in [p]} |\hat{w}_j - w_j|^2 (1 + \mathbb{E}_n[\hat{Z}_{ij}^2]) > \delta_n^2 / \log^2(pn)) \leq \beta_n$.

This condition holds trivially with $C_W = 1$ if we set $\hat{w}_j = w_j = 1$ for all $j \in [p]$ (recall that by Condition M, we have $V_{jj} \geq 1$ for all $j \in [p]$). Also, we show that Conditions M, E, and A imply Condition W, with possibly different δ_n and β_n , if we set $w_j = V_{jj}^{-1/2}$ and $\hat{w}_j = (\mathbb{E}_n[\hat{Z}_{ij}^2])^{-1/2}$ for all $j \in [p]$ as a part of the proof of Theorem 2.8 below.

The key observation that allows us to show that the confidence intervals (2.15) satisfy the desired coverage condition (2.14) will be to show that the bootstrap quantile function $\hat{\lambda}(1 - \alpha)$, as well as the original quantile function $\lambda(1 - \alpha)$, can be approximated by the Gaussian quantile function,

$$\lambda^g(1 - \alpha) := (1 - \alpha) \text{ quantile of } \|WN(0, V)\|_\infty.$$

It is this place, where Theorems 2.1-2.3 play a key role. Formally, we have the following results.

Theorem 2.5 (Quantile Comparison). *Assume that Conditions M, E, A, and W are satisfied. Then there exists a constant C depending only on C_W such that for $\epsilon_n := C(\delta_n + \beta_n)$,*

$$\lambda^g(1 - \alpha - \epsilon_n) \leq \lambda(1 - \alpha) \leq \lambda^g(1 - \alpha + \epsilon_n). \quad (2.16)$$

In addition,

$$\lambda^g(1 - \alpha - \epsilon_n) \leq \hat{\lambda}(1 - \alpha) \leq \lambda^g(1 - \alpha + \epsilon_n), \quad (2.17)$$

holds with probability at least $1 - 2\beta_n - n^{-1}$ in the case of E.1 and $1 - 2\beta_n - (\log n)^{-2}$ in the case of E.2. Moreover, for any $a \in (0, 1)$,

$$\lambda^g(1 - a) \leq \bar{\sigma} \Phi^{-1}(1 - a/(2p)) \leq \bar{\sigma} \sqrt{2 \log(2p/a)}, \quad (2.18)$$

where $\bar{\sigma} := \max_{j \in [p]} (w_j V_{jj}^{1/2})$.

Theorem 2.6 (Simultaneous Confidence Intervals). *Assume that Conditions M, E, A, and W are satisfied. Then the confidence intervals (2.15) satisfy the desired coverage condition (2.14). Moreover, with probability $1 - o(1)$, uniformly over $j \in [p]$, the maximum of weighted radii of these confidence intervals is bounded from above by*

$$\sup\{\|W(\hat{\theta} - \theta)\|_\infty : \theta \in CS\} \leq (1 + o(1)) \lambda^g(1 - \alpha + \epsilon_n) / \sqrt{n},$$

where CS is defined in (2.15).

Remark 2.4 (Simultaneous Confidence Intervals via Moderate Deviations). The construction of the simultaneous confidence intervals above relies upon the bootstrap approximation $\hat{\lambda}(1 - \alpha)$ of the quantile function $\lambda(1 - \alpha)$. Alternatively, we can use the moderate deviation theorem for this purpose. Specifically, set $w_j = V_{jj}^{-1/2}$ and $\hat{w}_j = (\mathbb{E}_n[\hat{Z}_{ij}^2])^{-1/2}$ for all $j \in [p]$. Then it follows from Theorem 2.4 that $\Phi^{-1}(1 - \alpha/2p)$ can be used as a good upper bound for $\lambda(1 - \alpha)$. Therefore, using the same arguments as those in the proof of Theorem 2.6, we can show that, under certain regularity conditions allowing for $p \gg n$,

the confidence intervals

$$CS = \prod_{j \in [p]} CS_j, \quad CS_j := \left[\hat{\theta}_j - \frac{\Phi^{-1}(1 - \alpha/2p)}{(\mathbb{E}_n[\hat{Z}_{ij}^2])^{-1/2}\sqrt{n}}; \hat{\theta}_j + \frac{\Phi^{-1}(1 - \alpha/2p)}{(\mathbb{E}_n[\hat{Z}_{ij}^2])^{-1/2}\sqrt{n}} \right].$$

satisfies the desired coverage condition (2.14). ■

2.4. Multiple Testing with FWER Control. In this subsection, we are interested in simultaneously testing hypotheses about different components of θ_0 . For concreteness, for each $j \in [p]$, we consider testing

$$H_j : \theta_{0j} \leq \bar{\theta}_{0j} \text{ against } H'_j : \theta_{0j} > \bar{\theta}_{0j}$$

for some given value $\bar{\theta}_{0j}$, where H_j is the null and H'_j is the alternative. The results below also apply for testing $H_j : \theta_{0j} = \bar{\theta}_{0j}$ against $H'_j : \theta_{0j} \neq \bar{\theta}_{0j}$ with obvious modifications of test statistics and critical values.

Since we are interested in testing these hypotheses simultaneously for all $j \in [p]$, we seek a procedure that would reject *at least* one true null hypothesis with probability not larger than $\alpha + o(1)$, uniformly over a large class of data-generating processes and, in particular, uniformly over the set of true null hypotheses. In the literature, procedures with this property are said to have *strong control of the Family-Wise Error Rate* (FWER).

More formally, let \mathcal{P} be a set of probability measures for the distribution of the data corresponding to different data generating processes, and let $P \in \mathcal{P}$ be the true probability measure. Each null hypothesis H_j is equivalent to $P \in \mathcal{P}_j$ for some subset \mathcal{P}_j of \mathcal{P} . Let $\mathcal{W} := \{1, \dots, p\}$ and for $w \subset \mathcal{W}$ denote $\mathcal{P}^w := (\cap_{j \in w} \mathcal{P}_j) \cap (\cap_{j \notin w} \mathcal{P}_j^c)$ where $\mathcal{P}_j^c := \mathcal{P} \setminus \mathcal{P}_j$. In words, \mathcal{P}_j is the set of probability measures corresponding to the j th null hypothesis H_j being true and \mathcal{P}^w is the set of probability measures such that all null hypotheses H_j with $j \in w$ are true and all null hypotheses H_j with $j \notin w$ are false.

Corresponding to this notation, strong FWER control means

$$\sup_{w \subset \mathcal{W}} \sup_{P \in \mathcal{P}^w} P_P \{\text{reject at least one } H_j \text{ for } j \in w\} \leq \alpha + o(1) \quad (2.19)$$

where P_P denotes the probability distribution generated by the probability measure of the data P . We seek a procedure that satisfies (2.19).

We consider three different (but related) procedures: the Bonferroni, Bonferroni-Holm, and Romano-Wolf procedures. All three procedures will be based on the t -statistics,

$$t_j := \frac{\sqrt{n}(\hat{\theta}_j - \bar{\theta}_{0j})}{(\mathbb{E}_n[\hat{Z}_{ij}^2])^{1/2}}, \quad j \in [p]. \quad (2.20)$$

For each $j \in [p]$, the *Bonferroni* procedure rejects H_j if $t_j > \Phi^{-1}(1 - \alpha/p)$. It is clear why this procedure satisfies (2.19): For any $w \subset \mathcal{W}$ and any $P \in \mathcal{P}^w$,

$$\max_{j \in w} t_j \leq \max_{j \in w} \frac{\sqrt{n}(\hat{\theta}_j - \theta_{0j})}{(\mathbb{E}_n[\hat{Z}_{ij}^2])^{1/2}} \leq \max_{j \in [p]} \frac{\sqrt{n}(\hat{\theta}_j - \theta_{0j})}{(\mathbb{E}_n[\hat{Z}_{ij}^2])^{1/2}},$$

and under the conditions of Theorem 2.4,

$$\mathbb{P}_P \left(\max_{j \in [p]} \frac{\sqrt{n}(\hat{\theta}_j - \theta_{0j})}{(\mathbb{E}_n[\hat{Z}_{ij}^2])^{1/2}} > \Phi^{-1}(1 - \alpha/p) \right) \leq \alpha + o(1),$$

where $o(1)$ does not depend on (w, P) . This result is formally stated in Theorem 2.7.

The Bonferroni procedure is a one-step method, which determines the hypotheses H_j to be rejected in just one step. This procedure can be improved by employing multi-step methods. In particular, so-called *stepdown* methods also have strong FWER control but may reject some H_j 's that are not rejected by the Bonferroni procedure, thus yielding important power improvements. We will consider the following form of the stepdown methods:

- (1) For a subset $w \subset \mathcal{W}$, let $c_{1-\alpha, w}$ be a (potentially conservative) estimator of the $(1-\alpha)$ quantile of $\max_{j \in w} t_j$. On the first step, let $w(1) = \mathcal{W}$ and reject all hypotheses H_j satisfying $t_j > c_{1-\alpha, w(1)}$. If no null hypothesis is rejected, then stop. If some H_j 's are rejected, let $w(2)$ be the set of all null hypotheses that were not rejected on the first step and move to (2).
- (2) On step $l \geq 2$, reject all hypotheses H_j for $j \in w(l)$ satisfying $t_j > c_{1-\alpha, w(l)}$. If no null hypothesis is rejected, then stop. If some H_j 's are rejected, let $w(l+1)$ be the subset of all null hypotheses j among $w(l)$ that were not rejected and proceed to the next step.

Here, we obtain the *Bonferroni-Holm* procedure, suggested in [77], by setting

$$(BH) \quad c_{1-\alpha, w} = \Phi^{-1}(1 - \alpha/|w|)$$

for all $w \subset \mathcal{W} = \{1, \dots, p\}$, where $|w|$ denotes the number of elements in w , and we obtain the *Romano-Wolf* procedure, suggested in [103], by setting

$$(RW) \quad c_{1-\alpha, w} = (1 - \alpha) \text{ quantile of } \max_{j \in w} \frac{\sqrt{n}(\hat{\theta}^* - \hat{\theta})}{(\mathbb{E}_n[\hat{Z}_{ij}^2])^{1/2}} \mid (\hat{Z}_i)_{i=1}^n,$$

where $\hat{\theta}^*$ is a bootstrap version of $\hat{\theta}$. In what follows, we maintain that $\hat{\theta}^*$ is obtained with the Gaussian bootstrap, (2.6), though the case of the empirical bootstrap can also be considered. To show that the Bonferroni-Holm and Romano-Wolf procedures have the strong FWER control (2.19), we will use the moderate deviation result in Theorem 2.4 and the high-dimensional CLT and bootstrap results in Theorems 2.1 and 2.2, respectively.

In [103], Romano and Wolf proved the following general result regarding the strong FWER control of the stepdown methods: If the critical values $c_{1-\alpha,w}$ satisfy

$$c_{1-\alpha,w'} \leq c_{1-\alpha,w''} \quad \text{whenever } w' \subset w'' \text{ and} \quad (2.21)$$

$$\sup_{w \subset \mathcal{W}} \sup_{P \in \mathcal{P}^w} \mathbb{P}_P \left(\max_{j \in w} t_j > c_{1-\alpha,w} \right) \leq \alpha + o(1), \quad (2.22)$$

then the stepdown method described above satisfies (2.19). Indeed, let w be the set of true null hypotheses and suppose that the method rejects at least one of these hypotheses. Let l be the step when the method rejects a true null hypothesis for the first time, and let H_{j_0} be this hypothesis. Clearly, we have $w(l) \supset w$. It then follows from (2.21) that

$$\max_{j \in w} t_j \geq t_{j_0} > c_{1-\alpha,w(l)} \geq c_{1-\alpha,w}.$$

Combining these inequalities with (2.22) yields (2.19).

We now formally establish strong FWER control of the three procedures described above.

Theorem 2.7 (Strong FWER Control by Bonferroni and Bonferroni-Holm Procedures). *Let \mathcal{P} be a class of probability measures for the distribution of the data such that Conditions M and A are satisfied with the same B_n , δ_n , and β_n for all $P \in \mathcal{P}$ and assume that $(2\sqrt{2})^3 B_n \log^{3/2}(pn)/\sqrt{n} \leq \delta_n$ for all $n \geq 1$. Then both Bonferroni and Bonferroni-Holm procedures have the strong FWER control property (2.19).*

Theorem 2.8 (Strong FWER Control by Romano-Wolf Procedure). *Let \mathcal{P} be a class of probability measures for the distribution of the data such that Conditions M, E, and A are satisfied with the same B_n , δ_n , and β_n for all $P \in \mathcal{P}$. Then the Romano-Wolf procedure with the Gaussian bootstrap critical values $c_{1-\alpha,w}$ described above has the strong FWER control property (2.19).*

Remark 2.5 (Comparison of Procedures with Strong FWER Control). Clearly, the Bonferroni procedure is the first step of the Bonferroni-Holm procedure. Thus, the set of hypotheses rejected by the Bonferroni-Holm procedure always contains all the hypotheses rejected by the Bonferroni procedure, which implies that the Bonferroni-Holm procedure is a more powerful version of the Bonferroni procedure. Since both procedures have strong FWER control, it follows that the former is preferable among two procedures. Regarding the comparison between the Bonferroni-Holm and Romano-Wolf procedures, both procedures have their own advantages. In particular, the Bonferroni-Holm procedure requires weaker conditions, as it does not require Condition E; but the Romano-Wolf procedure is typically more powerful because it uses non-conservative bootstrap critical values $c_{1-\alpha,w}$. ■

2.5. Multiple Testing with FDR Control. As in the previous subsection, here we are again interested in testing

$$H_j : \theta_j \leq \bar{\theta}_{0j} \text{ against } H'_j : \theta_j > \bar{\theta}_{0j} \quad (2.23)$$

simultaneously for all $j = 1, \dots, p$. In this subsection, however, we seek a procedure with a different type of control. In particular, we look for procedures controlling the *the False Discovery Rate* (FDR).

When p is large or very large, the requirement of controlling the FWER, which we had in the previous subsection, is often considered too stringent. This is because large values of p lead to large critical values $c_{1-\alpha, \mathcal{W}}$ used by the Bonferroni-Holm and Romano-Wolf procedures in the first step; and, when p is very large, it often turns out that these procedures simply fail to reject any null hypothesis. In [26], Benjamini and Hochberg offer an alternative, weaker, requirement: Instead of controlling the FWER, they suggest to control the False Discovery Rate (FDR), which is defined as the expected proportion of falsely rejected null hypotheses among all rejected null hypotheses.

To define the FDR formally, let $\mathcal{H}_0 \subset \{1, \dots, p\}$ be the set of indices j corresponding to the true null hypotheses H_j , and suppose that we have a procedure for testing (2.23) simultaneously for all $j \in [p]$ that rejects null hypotheses H_j with $j \in \mathcal{H}_R \subset \{1, \dots, p\}$. Thus, the total number of rejected null hypotheses is $|\mathcal{H}_R|$, and the total number of falsely rejected null hypotheses is $|\mathcal{H}_0 \cap \mathcal{H}_R|$. We then can define the *False Discovery Proportion* (FDP) by

$$FDP := \frac{|\mathcal{H}_0 \cap \mathcal{H}_R|}{|\mathcal{H}_R| \vee 1}.$$

Here, we put $|\mathcal{H}_R| \vee 1$ instead of $|\mathcal{H}_R|$ in the denominator to avoid division by zero since it may happen that all null hypotheses are accepted, so that $|\mathcal{H}_R| = 0$, and we want to define $FDP = 0$ when $|\mathcal{H}_R| = 0$. The FDR is then defined as the expected value of the FDP:

$$FDR = E[FDP].$$

We would like to have a procedure with the FDR being as small as possible and with large power in terms of rejecting false null hypotheses; so, for $\alpha \in (0, 1)$, we seek a procedure controlling the FDR in the sense that

$$FDR \leq \alpha + o(1).$$

Like in the previous subsection, we can also say that we have a strong control of FDR over a set \mathcal{P} of probability measures for the distribution of the data corresponding to different data generating processes if

$$\sup_{P \in \mathcal{P}} FDR_P \leq \alpha + o(1), \tag{2.24}$$

where we use index P to emphasize that the FDR depends on the distribution of the data represented by the probability measure P .

Next, we describe the *Benjamini-Hochberg* procedure, suggested in [26], for testing (2.23) simultaneously for all $j \in [p]$ with the FDR control. Recall the t -statistics t_j defined in (2.20), and let $t_{(1)} \geq \dots \geq t_{(p)}$ be the ordered sequence of t_j 's. Also, define $t_{(0)} := +\infty$.

Then let

$$\hat{k} := \max \left\{ j = 0, 1, \dots, p : 1 - \Phi(t_{(j)}) \leq \alpha j/p \right\}.$$

The Benjamini-Hochberg procedure rejects all H_j with $t_j \geq t_{(\hat{k})}$. For simplicity, to prove that this procedure controls the FDR, we will assume in this subsection that the random vectors $Z_i, i \in [n]$, are i.i.d., all having the same distribution as that of some random vector $Z^0 = (Z_1^0, \dots, Z_p^0)$. In addition, we will impose the following conditions:

Condition C. For some $0 < r < 1$ and $0 < \rho < (1 - r)/(1 + r)$, (i) the correlation between any two components of Z^0 is bounded in absolute value by r : $\max_{1 \leq j < k \leq p} |\mathbb{E}[Z_j^0 Z_k^0]| \leq r$, and (ii) for all $j \in [p]$, we have $|\mathbb{E}[Z_j^0 Z_k^0]| \geq (\log p)^{-3}$ for at most p^ρ different values of $k \in [p]$.

Condition F. (i) There exists $\mathcal{H} = \mathcal{H}_n \subset \{1, \dots, p\}$ such that $|\mathcal{H}| \geq \log \log p$ and $\sqrt{n}(\theta_{0j} - \bar{\theta}_{0j})/V_{jj}^{1/2} \geq 2\sqrt{\log p}$ for all $j \in \mathcal{H}$; (ii) for some $\gamma \in (0, 1)$, the number of false null hypotheses H_j , say p_1 , satisfied $p_1 \leq \gamma p$.

Conditions C and F are adapted from [89]. Condition C restricts dependence between Z_1^0, \dots, Z_p^0 and essentially corresponds to Condition (C1) in [89]. Condition C implies that each t -statistic t_j is highly correlated with at most p^ρ other t -statistics and is only weakly correlated with all other t -statistics. Condition F is a technical condition that restricts the number of false null hypotheses to be not too small and not too large.

Theorem 2.9 (Control of False Discovery Rate). *Let \mathcal{P} be a class of probability measures for the distribution of the data such that Conditions M, A, C, and F are satisfied with the same $B_n, \delta_n, \beta_n, r, \rho$, and γ for all $P \in \mathcal{P}$. Also, assume that $B_n = B$ for some constant B (independent of n) and all $n \geq 1$ and that uniformly over $P \in \mathcal{P}$, $p \rightarrow \infty$ and $\log p = o(n^\zeta)$ for some $\zeta < 3/23$ as $n \rightarrow \infty$. Then the Benjamini-Hochberg procedure has the FDR control property (2.24).*

Theorem 2.9 extends Theorem 4.1 in [89] to allow for many *approximate* means. The proof of this theorem relies critically on the moderate deviation result in Theorem 2.4. It is worth mentioning that [89] finds a phase transition phenomenon that the Benjamini-Hochberg method can not control the FDR when $\log p \geq c_0 n^{1/3}$ for some constant $c_0 > 0$; see Corollary 2.1 in [89].

2.6. Inference on Functionals of Many Approximate Means. We next consider the problem of estimating linear functionals of the vector θ_0 , say $a'\theta_0$, where $a \in \mathbb{R}^p$ is a vector of loadings. These functionals are of interest in many settings. For instance, as we discussed in Example 5 in the Introduction, the conditional average treatment effects may take the form of such functionals. We consider two cases separately: $\|a\|_1 = 1$ and $\|a\|_2 = 1$. As we will see, the analysis of the first case is straightforward and the results discussed so far immediately apply in this case. On the other hand, we will see that the second case is substantially more complicated, and treating it will require introducing some form of

regularization. It is this second case which will help us to develop intuition for the results to be discussed in the second part of the chapter.

2.6.1. *Inference on Functionals of Many Means without Regularization.* Inference using the MAM framework extends to the case where we are interested in linear functionals of θ_0 of the form

$$a'\theta_0, \text{ with } \|a\|_1 = 1.$$

Indeed, by Hölder's inequality,

$$\sup_{\|a\|_1=1} |a'(\hat{\theta} - \theta_0)| = \|\hat{\theta} - \theta_0\|_\infty,$$

which can be small even if p is much larger than n . For instance, in the “ideal noise model” where

$$\hat{\theta} - \theta_0 \sim N(0, I_p/n), \quad (2.25)$$

with I_p denoting the $p \times p$ identity matrix, we have that

$$\|\hat{\theta} - \theta_0\|_\infty \lesssim_P \sqrt{\log p/n},$$

so the error $\|\hat{\theta} - \theta_0\|_\infty$ is small if $\log p$ is much smaller than n .

Moreover, we have for a large collection of functionals $a'_k\theta_0$, indexed by a_k with $\|a_k\|_1 = 1$ and $k \in [p]$ that under (2.2),

$$\sqrt{n}a'_k(\hat{\theta} - \theta_0) = \frac{1}{\sqrt{n}} \sum_{i=1}^n a'_k Z_i + o_P(1/\sqrt{\log(pn)}) \quad \text{uniformly in } k \in [p].$$

It follows that $(a'_k\hat{\theta}_0)_{k=1}^p$ are approximate means themselves, so we can use the “many approximate means approach” for the inference on these functionals as well.

2.6.2. *Inference on Functionals of Many Means with Regularization.* Here we again consider inference on linear functionals $a'\theta_0$. However, in contrast to the discussion above, we assume that $\|a\|_2 = 1$ instead of $\|a\|_1 = 1$. Using only the condition that $\|a\|_2 = 1$, the previous “reduction back to many means” is not possible. Indeed,

$$\sup_{\|a\|_2=1} |a'(\hat{\theta} - \theta_0)| = \|\hat{\theta} - \theta_0\|_2,$$

so estimating all such functionals uniformly well will require having $\|\hat{\theta} - \theta_0\|_2$ small, which is typically not possible. For example, even in the ideal noise model (2.25), we have

$$\|\hat{\theta} - \theta_0\|_2^2 = \chi^2(p)/n = (p/n)(1 + O_P(1/\sqrt{p})),$$

which diverges to infinity if p increases more rapidly than n .

We therefore need to replace the estimator $\hat{\theta}$ by another estimator, say $\tilde{\theta}$, which performs favorably even when p is large relative to n in the sense that

$$\|\tilde{\theta} - \theta_0\|_2^2 \text{ is small.}$$

In order to achieve this goal, we will need two key ingredients. First, we will need to assume that θ_0 has some structure which can help to reduce the complexity of the estimation. Second, we will need to construct an estimator $\tilde{\theta}$ that is employing this structure to reduce the estimation error. In what follows, we shall rely on approximate sparsity as the structure on θ_0 and make use of ℓ_1 -norm regularization to obtain the estimator $\tilde{\theta}$.

A simple structure that provides useful intuition is the *exact sparsity structure with separation from zero*. Under this structure, θ_0 has s large components of size bigger than the maximal estimation error $\rho = \|\hat{\theta} - \theta_0\|_\infty$ and all remaining components *exactly* equal to zero. If θ_0 has this structure and the maximal estimation error r is known, we can use a simple estimator $\tilde{\theta} = (\tilde{\theta}_1, \dots, \tilde{\theta}_p)'$ with

$$\tilde{\theta}_j = \begin{cases} \hat{\theta}_j, & \text{if } |\hat{\theta}_j| > \rho, \\ 0, & \text{if } |\hat{\theta}_j| \leq \rho, \end{cases}$$

for all $j \in [p]$. This estimator might be referred to as a “selection-based estimator.” In practice, ρ will typically be unknown, but its distribution can be estimated by the bootstrap, as discussed in Section 2.2. We can then replace ρ in the definition of $\tilde{\theta}$ by an estimate of the $(1 - \alpha)$ -quantile of its distribution for some $\alpha = \alpha_n \rightarrow 0$. Figure 2 illustrates this selection-based estimator.

Assuming exact sparsity with separation structure provides a substantial dimension reduction that greatly eases the task of learning θ_0 as long as s is small. For example, in the ideal noise model, letting $T \subset \{1, \dots, p\}$ denote the set of indices of non-zero components of the vector θ_0 , the use of $\tilde{\theta}$ defined in the previous paragraph under this structure leads to the following bound:

$$\|\tilde{\theta} - \theta_0\|_2 \leq \sqrt{\sum_{j \in T} |\hat{\theta}_j - \theta_{0j}|^2} \leq \sqrt{\|N(0, I_s/n)\|^2} = \sqrt{\chi^2(s)/n} \lesssim_P \sqrt{s/n}, \quad (2.26)$$

which can be small provided that s is much smaller than n . However, the exact sparsity with separation structure seems unrealistic as a model for real econometric applications. A structure in which all parameters are either exactly zero or magically align themselves to be larger in magnitude than ρ seems extremely unintuitive and unlikely to correspond to most sensible economic models. We will not work with this structure further, though we will consider an exact sparsity structure with *no separation* since it helps to convey some of the main ideas of the theory of high-dimensional estimation.

More generally, we will consider an *approximately* sparse structure where the coefficients, sorted in non-increasing order in terms of absolute size, smoothly decline in magnitude

Selection-Based Estimation in ES Model




FIGURE 2. Illustration of the Selection-Based Estimator in Exactly Sparse Models. The true values θ_j are shown as blue points. The approximate mean estimators $\hat{\theta}_j$ are shown as red circles. The selection-based mean estimators $\tilde{\theta}_j$ are shown in blue circles. Here $\theta_j = 50/j^{3/2}$ for $j = 1, \dots, 8$ and $\theta_j = 0$ for $j > 8$, $\hat{\theta}_j = \theta_j + N(0,1)$, and $\tilde{\theta}_j = 1(|\hat{\theta}_j| > \lambda)\hat{\theta}_j$ with $\lambda = \Phi^{-1}(1 - \alpha/2p)$ with $\alpha = .1$. Note that the θ_j 's do not magically stay sufficiently far from zero for model selection to work perfectly, invalidating the (rather incredible) perfect model selection story very frequently used in the classical and modern literature. Note that the resulting selection-based estimator still performs well from the estimation point of view. Selection-based estimators of this kind have been known as post-model selection estimators. These selection-based estimators indeed perform well in exactly and approximately sparse models as shown by Belloni and Chernozhukov [10] who study the properties of the post-Lasso estimators in more general models.

towards zero. For this structure, we expect to achieve a rate of convergence similar to that in (2.26). Intuitively, under an approximately sparse structure, the unregularized estimator $\hat{\theta}$ still informs us about the components θ_{0j} that can not be distinguished from zero. We can thus use this information to set the regularized estimator $\tilde{\theta}_j$ for such components to zero. In such cases, we will be making an error by rounding those coefficients to zero, but the error will be negligible as long as the coefficients decrease to zero sufficiently quickly. We develop these results formally below.

For a given estimator $\hat{\theta}$, we consider the following ℓ_1 -regularization procedure:

$$\tilde{\theta} = \arg \min_{\theta \in \mathbb{R}^p} \|\theta\|_1 : \quad \|\hat{\theta} - \theta\|_\infty \leq \lambda, \quad (2.27)$$

where we minimize the ℓ_1 norm of the coefficients subject to them deviating from the initial estimates, $\hat{\theta}$, in the ℓ_∞ norm by at most λ . In (2.27), λ is the regularization parameter which controls the shrinkage in the estimator $\tilde{\theta}$. At one extreme, setting $\lambda = 0$ results in no regularization and yields $\tilde{\theta} = \hat{\theta}$. At the other extreme, setting $\lambda = \infty$ produces maximal regularization and results in $\tilde{\theta} = 0$. In general, we aim to set λ to be of the order of the estimation error $\|\hat{\theta} - \theta_0\|_\infty$ as in (2.29) below. Figure 3 illustrates the regularized estimator $\tilde{\theta}$.

To establish properties of the estimator $\tilde{\theta}$ in (2.27), observe that the optimization problem in (2.27) separates into p independent problems: For each $j \in [p]$,

$$\tilde{\theta}_j = \min_{\theta_j \in \mathbb{R}} |\theta_j| : \quad |\hat{\theta}_j - \theta_j| \leq \lambda. \quad (2.28)$$

It follows that the explicit solution $\tilde{\theta}$ of the optimization problem in (2.27) is given by

$$\tilde{\theta}_j = (|\hat{\theta}_j| - \lambda)_+ \text{sign}(\hat{\theta}_j), \quad j \in [p],$$

where for any $a \in \mathbb{R}$, we use $(a)_+$ to denote $a1\{a > 0\}$. This solution is known as the soft-thresholded estimator. In this simple setting, it also coincides with the well-known Lasso and Dantzig selector estimators.

To carry out (2.27), we need to choose the regularization parameter λ . We will assume that λ is chosen so that

$$\lambda \geq (1 - \alpha)\text{-quantile of } \|\hat{\theta} - \theta_0\|_\infty. \quad (2.29)$$

As follows from the discussion above, we can approximate such a λ either via self-normalized moderate deviations,

$$\lambda = n^{-1/2} \Phi^{-1}(1 - \alpha/2p) \max_{j \in [p]} (\mathbb{E}_n[\hat{Z}_{ij}^2])^{1/2}, \quad (2.30)$$

or via the bootstrap as outlined in Section 2.2. In the ideal noise model (2.25), we can also choose λ as

$$\lambda = n^{-1/2} \Phi^{-1}(1 - \alpha/2p) \leq \sqrt{2 \log(2p/\alpha)/n}. \quad (2.31)$$

Regularized Estimation in AS Model




FIGURE 3. Illustration of the Regularized Estimator in an Approximately Sparse Models. The true values θ_j are shown as black points. The approximate mean estimators $\hat{\theta}_j$ are shown as red circles. The ℓ_1 -regularized mean estimators $\tilde{\theta}_j$ are shown in blue circles. Here $\theta_j = 50/j^{3/2}$, $\hat{\theta}_j = \theta_j + N(0, 1)$, and $\tilde{\theta}_j = (|\hat{\theta}_j| - \lambda)_+ \text{sign}(\hat{\theta}_j)$ with $\lambda = \Phi^{-1}(1 - \alpha/2p)$ with $\alpha = .1$.

We analyze the estimator $\tilde{\theta}$ in (2.27) under three different conditions.

Condition ES. The parameter θ_0 is exactly sparse: There exists $T \subset \{1, \dots, p\}$ with cardinality s such that $\theta_{0j} \neq 0$ only for $j \in T$.

Condition AS. The parameter θ_0 is approximately sparse: For some $A > 0$ and $a > 1/2$, the non-increasing rearrangement $(|\theta_{0j}|^*)_{j \in [p]}$ of absolute values of coefficients $(|\theta_{0j}|)_{j \in [p]}$ obeys

$$|\theta_{0j}|^* \leq Aj^{-a}, \quad j \in [p].$$

Condition DM. The parameter θ_0 has bounded ℓ_1 norm: $\|\theta_0\|_1 \leq K$ for some $K > 0$.

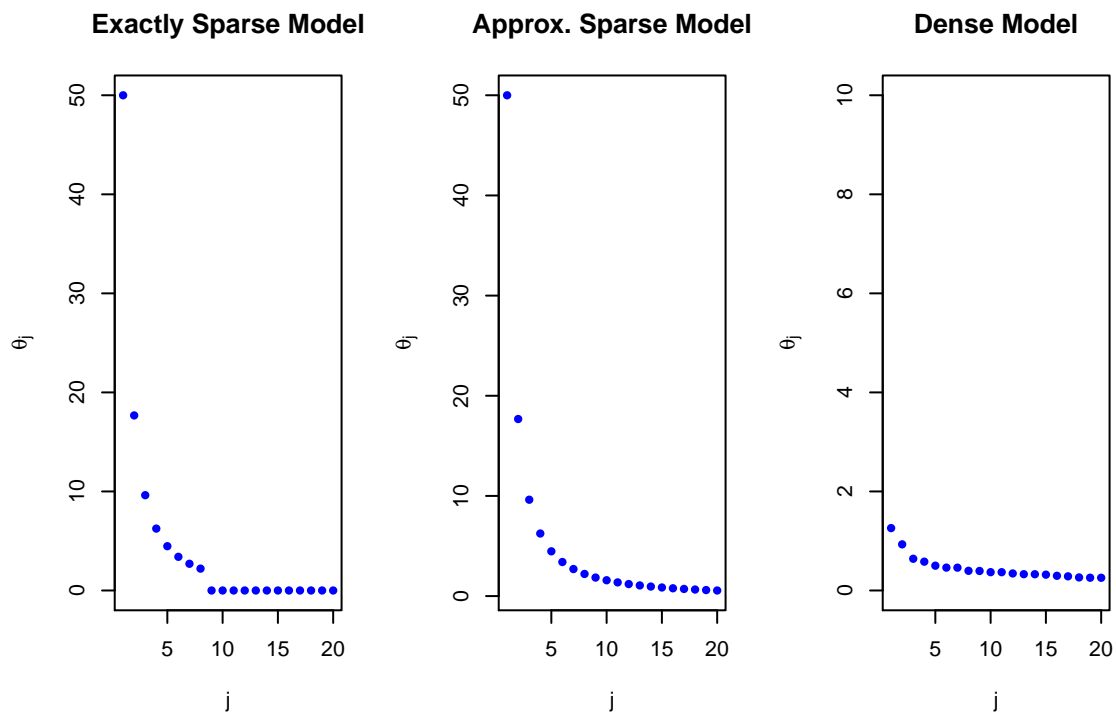


FIGURE 4. A visual illustration of the ES, AS, and DM Models. The true values θ_j are shown as blue points. In the Exactly Sparse (ES) Model, $\theta_j = 50/j^{3/2}$ for $j = 1, \dots, 8$ and $\theta_j = 0$ for $j > 8$. In the Approximately Sparse (AS) Model, $\theta_j = 10/j^{3/2}$. In the Dense Model (DM), $\theta_j = 10/(2p) + [10/(2p)]v_j$, where $\{v_j\}_{j=1}^p$ is a non-decreasing rearrangement of p i.i.d. draws of standard exponential variables and $\|\theta\|_1 \leq 10$.

Figure 4 illustrates these conditions. Speaking informally, ES can be thought of as a special case of AS, and AS can be thought of a special case of DM when $a > 1$. More formally, Condition ES implies Condition AS with any (a, A) such that $A \geq \|\theta_0\|_\infty s^a$; and, as long as $a > 1$, Condition AS implies Condition DM with any $K \geq Aa/(a - 1)$. While Conditions ES and AS require θ_0 to be either sparse or approximately sparse, Condition DM allows θ_0 to be “dense” – to have many elements that are all of similar, small size – such that neither sparsity nor approximate sparsity holds. When working with Condition AS, it will be convenient to denote $s = \lceil (A/\lambda)^{1/a} \rceil$, which can be thought of as the “effective” dimension of the approximately sparse θ_0 .

Condition AS on approximate sparsity can also be compared with conditions typically imposed in the literature on nonparametric series estimation, e.g. [95], where the *unordered* sequence of coefficients $(\theta_{0j})_{j \in [p]}$ is often required to obey $|\theta_{0j}| \leq Aj^{-a}$, which is referred to as a smoothness condition. The approximate sparsity condition thus can be considered as a relaxation of the smoothness condition.

We now establish a bound on the estimation error $\tilde{\theta} - \theta_0$ in the ℓ_q norm, where $q \geq 1$.

Theorem 2.10 (Estimation Bounds for ℓ_1 -Regularized Many Means). *Suppose that (2.29) holds. Then with probability at least $1 - \alpha$, we have $|\tilde{\theta}_j| \leq |\theta_{0j}|$ and $|\tilde{\theta}_j - \theta_{0j}| \leq 2\lambda$ for all $j \in [p]$ and*

- (i) *under Condition DM: $\|\tilde{\theta} - \theta_0\|_q \leq 2K^{1/q}\lambda^{1-1/q}$ if $q \geq 1$;*
- (ii) *under Condition ES: $\|\tilde{\theta} - \theta_0\|_q \leq 2s^{1/q}\lambda$ if $q \geq 1$, and $\tilde{\theta}_j = 0$ for all $j \notin T$;*
- (iii) *under Condition AS: $\|\tilde{\theta} - \theta_0\|_q \leq C_{a,q}s^{1/q}\lambda$ if $q > a^{-1}$, where $s = \lceil (A/\lambda)^{1/a} \rceil$,*

where $C_{a,q}$ is a constant depending only on a and q .

Corollary 2.2 (Estimation Bounds for ℓ_1 -Regularized Many Means). *Suppose that (2.29) holds with $\alpha \rightarrow 0$ and that $\lambda \lesssim \sqrt{\log p/n}$. Then*

- (i) *under Condition DM: $\|\tilde{\theta} - \theta_0\|_q \lesssim_P K^{1/q}(\log p/n)^{(q-1)/(2q)}$ if $q \geq 1$;*
- (ii) *under Condition ES: $\|\tilde{\theta} - \theta_0\|_q \lesssim_P s^{1/q}\sqrt{\log p/n}$ if $q \geq 1$, and $\tilde{\theta}_j = 0$ for all $j \notin T$;*
- (iii) *under Condition AS: $\|\tilde{\theta} - \theta_0\|_q \lesssim_P C_{a,q}s^{1/q}\sqrt{\log p/n}$ if $q > a^{-1}$, where $s = \lceil (A/\lambda)^{1/a} \rceil$,*

where $C_{a,q}$ is a constant depending only on a and q .

Note that the condition $\lambda \lesssim \sqrt{\log p/n}$ in Corollary 2.2 can easily be satisfied in the ideal noise model and in many other models, as long as $\max_{j \in [p]} \mathbb{E}_n[\hat{Z}_{ij}^2] \lesssim 1$; see (2.30) and (2.31).

Corollary 2.2 illustrates the power of regularization in high-dimensional settings when θ_0 has some structure. In particular, Conditions ES and AS both imply that the estimation

error of $\tilde{\theta}$ in the ℓ_2 norm satisfies

$$\|\tilde{\theta} - \theta_0\|_2 \lesssim_P \sqrt{s \log p/n},$$

where s is the effective dimension in the case of Condition AS as long as λ is chosen appropriately. Hence, the ambient dimension p affects the rate only through a log factor, and the effective dimension appears in a very natural form through the ratio $\sqrt{s/n}$. Thus, under either of these two conditions, we have consistency in the ℓ_2 norm as long as $s \log p/n$ tends to zero. Under Condition DM, the estimation error satisfies

$$\|\tilde{\theta} - \theta_0\|_2 \lesssim_P \sqrt{K}(\log p/n)^{1/4},$$

which tends to zero as long as $K^2 \log p/n \rightarrow 0$.

Proof of Theorem 2.10. It follows from (2.29) that with probability at least $1 - \alpha$, we have $\|\hat{\theta} - \theta_0\|_\infty \leq \lambda$, in which case

$$|\tilde{\theta}_j| \leq |\theta_{0j}| \text{ and } |\tilde{\theta}_j - \theta_{0j}| \leq |\tilde{\theta}_j - \hat{\theta}_j| + |\hat{\theta}_j - \theta_{0j}| \leq 2\lambda, \quad \text{for all } j \in [p] \quad (2.32)$$

by (2.28). This gives the first asserted claim.

To prove the second claim, we assume that (2.32) holds. Then, under Condition DM, $\|\theta_0\|_1 \leq K$, and so

$$\|\tilde{\theta} - \theta_0\|_q^q \leq \|\tilde{\theta} - \theta_0\|_1 \|\tilde{\theta} - \theta_0\|_\infty^{q-1} \leq (2K)(2\lambda)^{q-1} \leq 2^q K \lambda^{q-1}$$

by (2.32), which gives (i).

Further, under Condition ES, (2.32) implies that $|\tilde{\theta}_j| \leq |\theta_{0j}| = 0$ for all $j \in T^c$, and so

$$\|\tilde{\theta} - \theta_0\|_q^q = \|(\tilde{\theta} - \theta_0)_T\|_q^q \leq s \|\tilde{\theta} - \theta_0\|_\infty^q \leq s(2\lambda)^q,$$

which gives (ii).

Finally, consider the case of AS and assume, without loss of generality, that components of θ_0 are decreasing in absolute values, $|\theta_{01}| \geq \dots \geq |\theta_{0p}|$, so that $|\theta_{0j}| \leq A j^{-a}$ for all $j \in [p]$. Then, denoting $\bar{T} := \{j \in [p] : A j^{-a} > \lambda\}$, it follows from the triangle inequality and (2.32) that

$$\|\tilde{\theta} - \theta_0\|_q \leq \|(\tilde{\theta} - \theta_0)_{\bar{T}}\|_q + \|(\tilde{\theta} - \theta_0)_{\bar{T}^c}\|_q \leq \|(\tilde{\theta} - \theta_0)_{\bar{T}}\|_q + 2\|(\theta_0)_{\bar{T}^c}\|_q. \quad (2.33)$$

Here, $|\bar{T}| = s - 1$, and so

$$\|(\tilde{\theta} - \theta_0)_{\bar{T}}\|_q < s^{1/q} \|\tilde{\theta} - \theta_0\|_\infty \leq 2s^{1/q} \lambda. \quad (2.34)$$

Also,

$$\|(\theta_0)_{\bar{T}^c}\|_q^q \leq \sum_{j=1}^p (A j^{-a} \mathbf{1}\{A j^{-a} \leq \lambda\})^q = A^q \sum_{j=1}^p j^{-aq} \mathbf{1}\{j \geq (A/\lambda)^{1/a}\} \leq \frac{2^{aq} s \lambda^q}{aq - 1}, \quad (2.35)$$

where the last inequality is established by replacing the sum by an integral, as formally shown in Lemma D.1 in Appendix D. Combining (2.33), (2.34), and (2.35) gives (iii) and completes the proof of the theorem. ■

3. ESTIMATION WITH MANY PARAMETERS AND MOMENTS

3.1. Regularized Minimum Distance Estimation Problem. In this section, we consider the minimum distance estimation problem. We develop a Regularized Minimum Distance (RMD) estimator and study its properties.

Suppose that we have a target moment function $\theta \mapsto g(\theta)$, mapping $\Theta \subset \mathbb{R}^p$ to \mathbb{R}^m , and its empirical version $\theta \mapsto \hat{g}(\theta)$, also mapping $\Theta \subset \mathbb{R}^p$ to \mathbb{R}^m , where both p and m may be large. Assume that θ_0 is the unique solution of the following equation:

$$g(\theta_0) = 0.$$

We are interested in estimating θ_0 using the empirical version $\theta \mapsto \hat{g}(\theta)$ of the function $\theta \mapsto g(\theta)$.

We define the RMD estimator $\hat{\theta}$ as a solution of the optimization problem

$$\min_{\theta \in \Theta} \|\theta\|_1 : \|\hat{g}(\theta)\|_\infty \leq \lambda. \tag{3.1}$$

where λ is a regularization parameter. We will choose λ so that a solution of (3.1) exists with large probability; and in the event that $\|\hat{g}(\theta)\|_\infty > \lambda$ for all $\theta \in \Theta$ so that the optimization problem (3.1) has no solution, we can set $\hat{\theta}$ to be equal to any particular element of Θ . In the linear mean regression model, the RMD estimator reduces to the Dantzig Selector proposed by Candès and Tao in [36].

Let $\alpha \in (0, 1)$ be a constant, which should be thought of as some small number. We will assume that the regularization parameter λ satisfies the following condition:

Condition L. *The regularization parameter λ is such that*

$$\|\hat{g}(\theta_0)\|_\infty \leq \lambda, \text{ with probability at least } 1 - \alpha. \tag{3.2}$$

When $\hat{g}_j(\theta_0) \sim N(0, \sigma^2/n)$ for all $j \in [m]$, setting $\lambda = n^{-1/2} \sigma \Phi^{-1}(1 - \alpha/(2m))$ is sufficient to satisfy Condition L. This Gaussianity of the moment conditions occurs, for example, in the high-dimensional linear regression model with homoscedastic Gaussian noise under appropriate normalization of the covariates; e.g. see [9]. More generally, we can use the self-normalization method or the bootstrap to approximate λ as discussed in Section 2, provided that a preliminary estimator of θ_0 is available.

The key consequence of Condition L is that θ_0 is feasible in the optimization problem (3.1) with probability at least $1 - \alpha$, in which case a solution $\hat{\theta}$ of this optimization problem exists and, by optimality, satisfies $\|\hat{\theta}\|_1 \leq \|\theta_0\|_1$. As we discuss below, this property is crucial to handle high-dimensional models.

Denote

$$\mathcal{R}(\theta_0) := \{\theta \in \Theta : \|\theta\|_1 \leq \|\theta_0\|_1\},$$

which we sometimes refer to as the restricted set. Also, let there be some sequences of positive constants $(\epsilon_n)_{n \geq 1}$ and $(\delta_n)_{n \geq 1}$ satisfying $\epsilon_n \searrow 0$ and $\delta_n \searrow 0$. To establish properties of the RMD estimator $\hat{\theta}$, we will use the following high-level conditions:

Condition EMC. *The empirical moment function concentrates around the target moment function:*

$$\sup_{\theta \in \mathcal{R}(\theta_0)} \|\hat{g}(\theta) - g(\theta)\|_\infty \leq \epsilon_n \text{ with probability at least } 1 - \delta_n.$$

Condition MID. *The target moment function obeys the following identifiability condition:*

$$\{\|g(\theta) - g(\theta_0)\|_\infty \leq \epsilon, \theta \in \mathcal{R}(\theta_0)\} \text{ implies } \|\theta - \theta_0\|_\ell \leq r(\epsilon; \theta_0, \ell),$$

for all $\epsilon > 0$, where $\epsilon \mapsto r(\epsilon; \theta_0, \ell)$ is a weakly increasing rate function, mapping $[0, \infty)$ to $[0, \infty)$ and depending on the true value θ_0 and the semi-norm of interest ℓ .

Conditions EMC and MID encode the key blocks we need for the results. A large part of this section will be devoted to the verification of these conditions in particular examples. Condition EMC will be verified using empirical process methods, where contraction inequalities play a big role. Condition MID encodes both local and global identification of θ_0 . It states that if $g(\theta)$ is close to $g(\theta_0)$ in the ℓ_∞ norm and θ is weakly smaller than θ_0 in terms of the ℓ_1 norm, then θ is close to θ_0 in the semi-norm of interest ℓ . As we explain below, the validity of this condition depends on the interplay between the structure of g , the semi-norm ℓ , and the structure of θ_0 . To appreciate the latter point, we note that if $\theta_0 = 0$, then Condition MID always holds with $r(\epsilon; 0, \ell) = 0$ for all semi-norms ℓ . This means that $\theta_0 = 0$ is identified in the restricted set $\mathcal{R}(\theta_0)$ under no other assumptions on g . The structure of g and θ_0 will begin to play an important role when $\theta_0 \neq 0$, as we discuss below.

Using these conditions, we can immediately obtain the following elementary but important result on the properties of the RMD estimator:

Proposition 3.1 (Bounds on Estimation Error of RMD Estimator). *Assume that Conditions L, EMC, and MID are satisfied for some semi-norm ℓ . Then with probability at least $1 - \alpha - \delta_n$, the RMD estimator $\hat{\theta}$ obeys*

$$\|\hat{\theta} - \theta_0\|_\ell \leq r(\lambda + \epsilon_n; \theta_0, \ell). \quad (3.3)$$

Proof of Proposition 3.1. Consider the event that $\lambda \geq \|\hat{g}(\theta_0)\|_\infty$, $\hat{\theta} \in \mathcal{R}(\theta_0)$, and $\|\hat{g}(\hat{\theta}) - g(\hat{\theta})\|_\infty \leq \epsilon_n$. By the union bound and Condition EMC, this event occurs with probability at least $1 - \alpha - \delta_n$ since Condition L implies that with probability at least $1 - \alpha$, we have $\lambda \geq \|\hat{g}(\theta_0)\|_\infty$ and $\hat{\theta} \in \mathcal{R}(\theta_0)$. On this event, we have, by the definition of the RMD estimator, $\|\hat{g}(\hat{\theta})\|_\infty \leq \lambda$, and so

$$\|g(\hat{\theta})\|_\infty \leq \|g(\hat{\theta}) - \hat{g}(\hat{\theta})\|_\infty + \|\hat{g}(\hat{\theta})\|_\infty \leq \epsilon_n + \lambda \quad (3.4)$$

by the triangle inequality. In turn, (3.4) implies (3.3) via Condition MID since $g(\theta_0) = 0$, which gives the asserted claim. ■

Example 8 (Regularized GMM). An important special case of the minimum distance estimation problem is the Generalized Method of Moments (GMM) estimation problem, which corresponds to

$$g(\theta) = AEg(X, \theta), \quad \hat{g}(\theta) = \hat{A}\hat{E}_n g(X, \theta),$$

where $(x, \theta) \mapsto g(x, \theta)$ is a measurable score function, mapping $\mathbb{R}^{d_x} \times \Theta$ to \mathbb{R}^m , A is a positive definite weighting matrix, and \hat{A} is an estimator of this matrix. In practice, a simple choice of A is the diagonal weighting matrix,

$$A^2 = \text{diag}(\text{Var}(\mathbb{E}_n g(X, \tilde{\theta})))^{-1},$$

where $\tilde{\theta}$ is a guess or a preliminary estimator of θ_0 . The analysis in this section will be suitable for the case where the weighting matrix is known, i.e. $\hat{A} = A$, and so, without much loss of generality, we set $A = I$ to simplify the exposition. We turn to other choices and estimation of A in Section 3.2, where we consider optimal inference for individual components of θ_0 . ■

In what follows, let

$$G := (\partial/\partial\theta')g(\theta)|_{\theta=\theta_0}$$

be the Jacobian matrix. This matrix plays an important role in encoding the local information about θ_0 . For convenience, for all $j \in [m]$, we will use $G_j(\theta)$ to denote the j th row of the matrix $G(\theta)$. Regarding the semi-norm ℓ , we shall be focusing mainly on the ℓ_1 norm $\|\cdot\|_1$, ℓ_2 norm $\|\cdot\|_2$, and the Fisher norm $\|\cdot\|_F$, i.e. $\ell \in \{\ell_1, \ell_2, F\}$, where we define the Fisher norm by

$$\|\delta\|_F := |\delta'(G'G)^{1/2}\delta|^{1/2}, \quad \delta \in \mathbb{R}^p.$$

To make Proposition 3.1 operational, we need to verify Conditions EMC and MID and provide suitable choices of ϵ_n , δ_n , and the rate function r . We will do so separately in linear and non-linear models.

3.1.1. *Linear Case.* Here we consider the case where $\theta \mapsto g(\theta)$ is linear:

$$g(\theta) = G\theta + g(0). \quad (3.5)$$

Two lead examples of this case are the linear mean regression model and the linear instrumental variables (IV) model:

Example 9 (Linear regression models). The linear mean regression model,

$$Y = W'\theta_0 + \epsilon, \quad E\epsilon W = 0, \quad (3.6)$$

corresponds to (3.5) with $g(\theta) = E[(Y - W'\theta)W]$, $G = -EWW'$, and $g(0) = EYW$. The linear IV regression model,

$$Y = W'\theta_0 + \epsilon, \quad E\epsilon Z = 0, \quad (3.7)$$

corresponds to (3.5) with $g(\theta) = E[(Y - W'\theta)Z]$, $G = -EZW'$, and $g(0) = EYZ$. For both models, we will be able to show that the RMD estimator $\hat{\theta}$ has a fast rate of convergence as long as θ_0 is either exactly sparse or approximately sparse under simple, intuitive conditions on the matrix G . ■

We first discuss Condition MID. We consider three cases as in Section 2:

- exactly sparse model: θ_0 obeys Condition ES;
- approximately sparse model: θ_0 obeys Condition AS;
- dense model: θ_0 obeys Condition DM.

In the dense model, where $\|\theta_0\|_1 \leq K$, we can immediately deduce that whenever G is symmetric and non-negative definite, Condition MID holds with $\ell = F$ and

$$r(\epsilon; \theta_0, F) \leq \sqrt{2K\epsilon}.$$

Indeed, to prove this inequality, observe that for any $\theta \in \mathcal{R}(\theta_0)$ satisfying $\|g(\theta) - g(\theta_0)\|_\infty \leq \epsilon$, we have

$$\begin{aligned} \|\theta - \theta_0\|_F^2 &= |(\theta - \theta_0)'G(\theta - \theta_0)| \leq \|\theta - \theta_0\|_1 \|G(\theta - \theta_0)\|_\infty \\ &= \|\theta - \theta_0\|_1 \|g(\theta) - g(\theta_0)\|_\infty \leq 2K\epsilon, \end{aligned}$$

since $\theta \in \mathcal{R}(\theta_0)$ implies that $\|\theta - \theta_0\|_1 \leq \|\theta\|_1 + \|\theta_0\|_1 \leq 2\|\theta_0\|_1 \leq 2K$ via the triangle inequality. This result will imply “slow” rates of convergence of the estimator $\hat{\theta}$, though these slow rates will be sufficient in some applications.

In the exactly sparse or approximately sparse models, we proceed as follows. Define a modulus of continuity

$$k(\theta_0, \ell) := \inf_{\theta \in \mathcal{R}(\theta_0): \|\theta - \theta_0\|_\ell > 0} \|G(\theta - \theta_0)\|_\infty / \|\theta - \theta_0\|_\ell,$$

which we can also call an identifiability factor. Whenever $k(\theta_0, \ell) \neq 0$, we immediately obtain that Condition MID holds with

$$r(\epsilon; \theta_0, \ell) \leq k(\theta_0, \ell)^{-1} \epsilon.$$

There exist methods in the literature to show that $k(\theta_0, \ell) \neq 0$ and to bound $k(\theta_0, \ell)$ from below in the exactly sparse model. When we consider the approximately sparse model, we will first sparsify θ_0 to $\theta_0(\epsilon) = (\theta_{01}(\epsilon), \dots, \theta_{0p}(\epsilon))'$ with

$$\theta_{0j}(\epsilon) := \begin{cases} \theta_{0j} + \text{sign}(\theta_{0j})\Delta/(s-1), & \text{if } Aj^{-a} > \epsilon, \\ 0, & \text{if } Aj^{-a} \leq \epsilon, \end{cases}$$

where $s := \lceil (A/\epsilon)^{1/a} \rceil$ and $\Delta := \sum_{j=1}^p |\theta_{0j}| 1\{Aj^{-a} \leq \epsilon\}$. We will then provide a bound on $k^{-1}(\theta_0, \ell)\epsilon$ in terms of $k^{-1}(\theta_0(\epsilon), \ell)\epsilon$ and the approximation error $\|\theta_0 - \theta_0(\epsilon)\|_\ell$. Note that we assume that the components of θ_0 are decreasing in absolute values, $|\theta_{01}| \geq \dots \geq |\theta_{0p}|$ in this construction, which is without loss of generality because we do not use this information in the estimation. We also inflate components of $\theta_{0j}(\epsilon)$ with $Aj^{-a} > \epsilon$ by using $\theta_{0j} + \text{sign}(\theta_{0j})\Delta/(s-1)$ instead of θ_{0j} to make sure that $\theta \in \mathcal{R}(\theta_0)$ implies $\theta \in \mathcal{R}(\theta_0(\epsilon))$, which will be important in the verification of Condition MID. Finally, we assume that $A > \epsilon$ to make sure that $s - 1 \geq 1$.

To present one possible lower bound for $k(\theta_0, \ell)$ in the exactly sparse model, we introduce some further notation. For $J \subset \{1, \dots, m\}$ and $H \subset \{1, \dots, p\}$, let $G_{J,H} = (G_{j,h})_{j \in J, h \in H}$ be the submatrix of G consisting of all rows $j \in J$ and all columns $h \in H$ of G . Also, for an integer $l \geq s$, define the l -sparse smallest and l -sparse largest singular values of G by

$$\sigma_{\min}(l) := \min_{|H| \leq l} \max_{|J| \leq l} \sigma_{\min}(G_{J,H}), \quad \sigma_{\max}(l) := \max_{|H| \leq l} \max_{|J| \leq l} \sigma_{\max}(G_{J,H}), \quad (3.8)$$

respectively, where $\sigma_{\min}(G_{J,H})$ and $\sigma_{\max}(G_{J,H})$ are the smallest and the largest singular values of the matrix $G_{J,H}$. We then have the following lower bound on $k(\theta_0, \ell)$, which is an immediate consequence of Theorem 1 in [17]:

Lemma 3.1 (Lower Bound for $k(\theta_0, \ell)$ in Exactly Sparse Model). *Under Condition ES, there exists a universal constant $c \geq 1$ such that if $\sigma_{\min}^2(l)/\sigma_{\max}^2(l) \geq c\mu_n$ and $\sigma_{\max}^2(l) \geq 1$ for $l = \lceil s/\mu_n \rceil$ and some $\mu_n \in (0, 1/c)$, then*

$$k(\theta_0, \ell_q) \geq s^{-1/q} \mu_n, \quad q \in \{1, 2\}. \quad (3.9)$$

In Lemma 3.1, the quantity μ_n appears as a key factor determining the modulus of continuity $k(\theta_0, \ell)$. When μ_n is bounded away from zero, we say that we have a strongly identified model. Otherwise, i.e. when $\mu_n > 0$ is drifting towards zero, we say that we have a non-strongly identified model. The quantity μ_n in turn can be easily bounded in linear regression models:

Example 9 (Linear Regression Models, Continued). In the linear mean regression model (3.6), assume that all eigenvalues of $G = -EWW'$ are bounded in absolute values from above and away from zero uniformly over n . Then for any integer $l \geq s$, we have

$$\sigma_{\min}(l) = \min_{|H| \leq l} \max_{|J| \leq l} \sigma_{\min}(G_{J,H}) \geq \min_{|H| \leq l} \sigma_{\min}(G_{H,H}) \geq \sigma_{\min}(G)$$

and, similarly,

$$\sigma_{\max}(l) = \max_{|H| \leq l} \max_{|J| \leq l} \sigma_{\max}(G_{J,H}) \leq \sigma_{\max}(G)$$

by the standard properties of eigenvalues of symmetric matrices. Hence, we can choose μ_n in Lemma 3.1 to be bounded away from zero, which means that the linear mean regression model is strongly identified with $k(\theta_0, \ell_q)$ satisfying $k(\theta_0, \ell_q)^{-1} \leq Cs^{1/q}$ for some constant $C > 0$.

In the linear IV regression model (3.7), assume first that there exist constants $\bar{\sigma} \geq \underline{\sigma} > 0$ such that for all $s/\log n \leq l \leq s \log n$ and any combination of l covariates W_{H_1}, \dots, W_{H_l} from the vector $W = (W_1, \dots, W_p)'$, there exists a combination of l instruments Z_{J_1}, \dots, Z_{J_l} from the vector $Z = (Z_1, \dots, Z_m)'$ such that the matrix $E[(Z_{J_1}, \dots, Z_{J_l})'(W_{H_1}, \dots, W_{H_l})]$ has singular values bounded from above by $\bar{\sigma}$ and from below by $\underline{\sigma}$, which means that Z_{J_1}, \dots, Z_{J_l} are strong instruments for the covariates W_{H_1}, \dots, W_{H_l} . Then again we can choose μ_n in Lemma 3.1 to be bounded away from zero, which again implies that we have a strongly identified model with $k(\theta_0, \ell_q)$ satisfying $k(\theta_0, \ell_q)^{-1} \leq Cs^{1/q}$. On the other hand, if for some covariate W_j , we only have weak instruments, μ_n will drift towards zero, and we obtain a non-strongly identified model. ■

In the analysis below, we use the bound (3.9) as a starting point. In particular, letting there be a sequence of constants $(L_n)_{n \geq 1}$ satisfying $L_n \geq 1$ for all $n \geq 1$, we impose the following assumption:

Condition LID. For $q \in \{1, 2\}$, either of the following conditions hold: (a) θ_0 obeys Condition ES and the bound $k(\theta_0, \ell_q) \geq s^{-1/q} \mu_n$ holds, or (b) θ_0 obeys Condition AS, the bound $k(\theta_0(\epsilon), \ell_q) \geq s^{-1/q} \mu_n$ holds for $s = \lceil (A/\epsilon)^{1/a} \rceil$, and $\|G_j\|_1 \leq L_n$ for each $j \in [m]$, where G_j denotes the j th row of G .

We use the notation LID as a shorthand for $\text{LID}(\theta_0, G)$ since the condition is indexed by both the parameter of interest θ_0 and the matrix G . We will make the dependence explicit when we invoke the condition to other parameters.

We can now verify Condition MID:

Lemma 3.2 (Bounds on Rate Function in Linear Models). (i) Under Condition LID(a), Condition MID holds with

$$r(\epsilon; \theta_0, \ell_q) \leq \epsilon s^{1/q} \mu_n^{-1}, \quad q \in \{1, 2\}.$$

(ii) Under Condition LID(b), Condition MID holds with

$$r(\epsilon; \theta_0, \ell_q) \leq C_{a,q} \left(L_n \epsilon s^{1/q} \mu_n^{-1} + \epsilon s^{1/q} \right), \quad q \in \{1, 2\},$$

as long as $a > 1$ and $A > \epsilon$, where $C_{a,q}$ is a constant depending only on a and q .

Proof of Lemma 3.2. The first claim is immediate from the definition of the identifiability factor $k(\theta_0, \ell_q)$. To show the second claim, assume that Condition AS holds, fix $q \in \{1, 2\}$, and take any $\theta \in \mathcal{R}(\theta_0)$ such that $\|g(\theta) - g(\theta_0)\|_\infty = \|G(\theta - \theta_0)\|_\infty \leq \epsilon$. By the triangle inequality,

$$\|\theta - \theta_0\|_q \leq \|\theta - \theta_0(\epsilon)\|_q + \|\theta_0(\epsilon) - \theta_0\|_q. \quad (3.10)$$

Below, we bound the two terms on the right-hand side of this inequality.

To bound $\|\theta_0(\epsilon) - \theta_0\|_q$, we have

$$\Delta = \sum_{j=1}^p |\theta_{0j}| \mathbf{1}\{Aj^{-a} \leq \epsilon\} \leq \sum_{j=1}^p Aj^{-a} \mathbf{1}\{Aj^{-a} \leq \epsilon\} \leq \frac{2^a \epsilon s}{a-1} \quad (3.11)$$

by Condition AS and Lemma D.1 in Appendix D. Thus,

$$\begin{aligned} \|\theta_0(\epsilon) - \theta_0\|_q^q &\leq \sum_{j=1}^p (\Delta/(s-1))^{q-1} \mathbf{1}\{j < (A/\epsilon)^{1/a}\} + \sum_{j=1}^p (Aj^{-a})^{q-1} \mathbf{1}\{j \geq (A/\epsilon)^{1/a}\} \\ &\leq (s-1)(\Delta/(s-1))^q + \frac{2^{aq} \epsilon^q s}{aq-1} \leq C_{a,q}^q \epsilon^q s. \end{aligned}$$

To bound $\|\theta - \theta_0(\epsilon)\|_q$, we have

$$\begin{aligned} \|G(\theta - \theta_0(\epsilon))\|_\infty &\leq \|G(\theta - \theta_0)\|_\infty + \|G(\theta_0 - \theta_0(\epsilon))\|_\infty \\ &\leq \epsilon + \|G(\theta_0 - \theta_0(\epsilon))\|_\infty \leq \epsilon + \max_{j \in [m]} \|G_j\|_1 \|\theta_0 - \theta_0(\epsilon)\|_\infty \leq C_{a,q} L_n \epsilon, \end{aligned}$$

where the last inequality follows from (3.11) and Condition AS. Therefore, since $\theta \in \mathcal{R}(\theta_0)$ implies $\theta \in \mathcal{R}(\theta_0(\epsilon))$, we have by Condition LID that

$$\|\theta - \theta_0(\epsilon)\|_q \leq C_{a,q} L_n \epsilon s^{1/q} \mu_n^{-1}.$$

Combining the bounds on $\|\theta_0(\epsilon) - \theta_0\|_q$ and $\|\theta - \theta_0(\epsilon)\|_q$ above with (3.10) gives the second asserted claim. \blacksquare

Remark 3.1 (Bounding $k(\theta_0, F)$ via $k(\theta_0, \ell_1)$). Note that once we have a lower bound on $k(\theta_0, \ell_1)$, we can always use it to obtain a lower bound on $k(\theta_0, F)$, whenever G is symmetric. Indeed, for any $\theta \in \mathbb{R}^p$ such that $\theta \neq \theta_0$, we have

$$(\theta - \theta_0)' G(\theta - \theta_0) \leq \|\theta - \theta_0\|_1 \|G(\theta - \theta_0)\|_\infty.$$

Rearranging this expression gives

$$\frac{\|G(\theta - \theta_0)\|_\infty}{\sqrt{(\theta - \theta_0)' G(\theta - \theta_0)}} \geq \sqrt{\frac{\|G(\theta - \theta_0)\|_\infty}{\|\theta - \theta_0\|_1}},$$

which implies that $k(\theta_0, F) \geq k^{1/2}(\theta_0, \ell_1)$. ■

We next verify Condition EMC. Let there be a sequence of constants $(\ell_n)_{n \geq 1}$ satisfying $\ell_n \geq 1$ for all $n \geq 1$. Consider the following condition:

Condition ELM. *The empirical moment function is linear,*

$$\hat{g}(\theta) = \hat{G}\theta + \hat{g}(0),$$

and with probability at least $1 - \delta_n$, we have

$$\max_{j \in [m]} \|\hat{G}_j - G_j\|_\infty \vee \|\hat{g}(0) - g(0)\|_\infty \leq \ell_n / \sqrt{n}.$$

Example 8 (Regularized GMM, Continued). When specialized to GMM problems, Condition ELM means that the score function is linear,

$$g(X, \theta) = G(X)\theta + g(X, 0),$$

where $G(X) := (\partial/\partial\theta')g(X, \theta)|_{\theta=\theta_0}$, and with probability at least $1 - \delta_n$, we have

$$\max_{j \in [m]} \|\mathbb{G}_n G_j(X)\|_\infty \vee \|\mathbb{G}_n g(X, 0)\|_\infty \leq \ell_n.$$

This form will be useful below to verify Condition ELM in the linear regression models. ■

Condition ELM is plausible. It is implied by many sufficient conditions based on self-normalized moderate deviations and high-dimensional central limit theorems, as reviewed in Section 2. In particular, it is possible to choose

$$\ell_n \propto \sqrt{\log(p \vee m \vee n)}. \tag{3.12}$$

in many cases, as illustrated in the case of linear regression models below. We highlight the slow growth of ℓ_n with respect to the number of parameters p and the number moment functions m , which is critical to allow the analysis to handle high-dimensional models.

Lemma 3.3 (Empirical Moment Concentration). *If Conditions DM and ELM are satisfied, then Condition EMC holds with*

$$\epsilon_n = n^{-1/2} \ell_n (K + 1).$$

Proof of Lemma 3.3. Conditions DM and ELM imply that with probability at least $1 - \delta_n$,

$$\begin{aligned} \sup_{\theta \in \mathcal{R}(\theta_0)} \|\hat{g}(\theta) - g(\theta)\|_\infty &\leq \sup_{\theta \in \mathcal{R}(\theta_0)} \|(\hat{G} - G)\theta + \hat{g}(0) - g(0)\|_\infty \\ &\leq \max_{j \in [m]} \|\hat{G}_j - G_j\|_\infty \|\theta_0\|_1 + \|\hat{g}(0) - g(0)\|_\infty \leq n^{-1/2} \ell_n (K + 1), \end{aligned}$$

which gives the asserted claim. ■

Example 9 (Linear Regression Models, Continued). Here, we verify Condition ELM for linear regression models under primitive conditions. Since the mean regression model is a special case of the IV regression model, we only consider the latter. In the linear IV model, we have $g(0) = E[YZ]$ and $G = -E[ZW']$. Suppose that $Z \in \mathbb{R}^m$ and $W \in \mathbb{R}^p$ are such that the following moment condition holds for some $\sigma > 0$:

$$\max_{j \in [m], k \in [p]} \left(E[Y^4] + E[Z_j^4] + E[W_k^4] \right) \leq \sigma^2.$$

Also, suppose that

$$E \left[\|(Y, Z', W')'\|_\infty^8 \right] \leq M_n^4$$

for some M_n , possibly growing to infinity, but such that $n^{-1/2} M_n^2 \log(p \vee m \vee n) \leq \sigma^2$. Let $(Y_i, Z_i, W_i)_{i=1}^n$ be a random sample from the distribution of (Y, Z, W) . Then by Hölder's inequality,

$$\sum_{i=1}^n E[(Y_i Z_{ij})^2] \leq \sum_{i=1}^n (E[Y_i^4] E[Z_{ij}^4])^{1/2} \leq n\sigma^2$$

for all $j \in [m]$. Hence, by Lemma A.3,

$$\begin{aligned} E \left[\max_{j \in [p]} \left| \sum_{i=1}^n (Y_i Z_{ij} - E[Y_i Z_{ij}]) \right| \right] &\leq A \left(\sigma \sqrt{n \log p} + \sqrt{E \left[\max_{i \in [n]} \|Y_i Z_i\|_\infty^2 \right] \log p} \right) \\ &\leq A \left(\sigma \sqrt{n \log p} + n^{1/4} M_n \log p \right) \leq 2A\sigma \sqrt{n \log p} \end{aligned}$$

for some universal constant $A > 0$. Therefore, applying Lemma A.2 with $s = 4$, $t = \sigma \sqrt{3n \log n}$, and σ^2 replaced by $n\sigma^2$ shows that there exist universal constants $c, C > 0$ such that with probability at least $1 - c/\log^2 n$,

$$\|\mathbb{G}_n g(X, 0)\|_\infty \leq C\sigma \sqrt{\log(p \vee n)},$$

By the same argument, again with probability at least $1 - c/\log^2 n$, we also have

$$\max_{j \in [m]} \|\mathbb{G}_n G_j(X)\|_\infty \leq C\sigma \sqrt{\log(p \vee m \vee n)}.$$

Condition ELM thus holds with $\delta_n = 2c/\log^2 n$ and

$$\ell_n = C\sigma \sqrt{\log(p \vee m \vee n)},$$

which is in accord with (3.12). ■

Summarizing the results in Proposition 3.1, Lemma 3.2, and Lemma 3.3, we obtain the following theorem:

Theorem 3.1 (Bounds on Estimation Error of RMD Estimator, Linear Case). *In the linear case (3.5), assume that Conditions L, DM, LID, and ELM are satisfied and $\lambda \leq \ell_n/\sqrt{n}$. Then with probability at least $1 - \alpha - \delta_n$,*

$$\|\hat{\theta} - \theta_0\|_q \leq \frac{\ell_n (K + 2) s^{1/q}}{\mu_n \sqrt{n}}, \quad q \in \{1, 2\} \tag{3.13}$$

in the case of LID(a) (exactly sparse model). Similarly, as long as $a > 1$ and $A > n^{-1/2}\ell_n(K+1)$,

$$\|\hat{\theta} - \theta_0\|_q \leq \frac{C_{a,q}(L_n + \mu_n)\ell_n(K+2)s^{1/q}}{\mu_n\sqrt{n}}, \quad q \in \{1, 2\} \quad (3.14)$$

in the case of LID(b) (approximately sparse model), where $C_{a,q}$ is a constant depending only on a and q .

Theorem 3.1 implies ℓ_q -rates of convergence of the RMD estimator and characterizes how the sparsity of θ_0 impacts these rates. The impact of the overall number of coefficients is bounded by the factor ℓ_n which typically grows logarithmically with the number of coefficients p and the number of moment conditions m . This is effectively the impact of not knowing the support of θ_0 . This implies that we can achieve consistent estimators even if $(p \vee m) \gg n$ since ℓ_n can grow much slower than root- n .

Remark 3.2 (Improving Theorem 3.1). The right-hand sides of the bounds in Theorem 3.1 depend linearly on K , which may be suboptimal if K is increasing with n . This dependence can be improved by reiterating the argument in Proposition 3.1 one more time. In particular, we can introduce a doubly restricted set

$$\bar{\mathcal{R}}(\theta_0) := \{\theta \in \Theta: \|\theta\|_1 \leq \|\theta_0\|_1 \text{ and } \|\theta - \theta_0\|_1 \leq \gamma_n\},$$

where γ_n denotes the right-hand side of either (3.13) or (3.14), depending on which part of Condition LID is imposed. Then by Theorem 3.1 and Condition L, $\hat{\theta} \in \bar{\mathcal{R}}(\theta_0)$ with probability at least $1 - \alpha - \delta_n$, and we can replace the set $\mathcal{R}(\theta_0)$ in the proof of Proposition 3.1 by $\bar{\mathcal{R}}(\theta_0)$. In turn, modifying slightly the proof of Lemma 3.3, which provides the value of ϵ_n to be used in the proof of Proposition 3.1, we can write

$$\begin{aligned} \epsilon_n &= \sup_{\theta \in \bar{\mathcal{R}}(\theta_0)} \|\hat{g}(\theta) - g(\theta)\|_\infty \leq \sup_{\theta \in \bar{\mathcal{R}}(\theta_0)} \|(\hat{G} - G)\theta_0 + (\hat{G} - G)(\theta - \theta_0) + \hat{g}(0) - g(0)\|_\infty \\ &\leq \|(\hat{G} - G)\theta_0\|_\infty + \max_{j \in [m]} \|\hat{G}_j - G_j\|_\infty \gamma_n + \|\hat{g}(0) - g(0)\|_\infty, \end{aligned}$$

which is typically of order $n^{-1/2}\ell_n$ (as long as $\gamma_n \rightarrow 0$ and $\|(\hat{G} - G)\theta_0\|_\infty$ is of order $n^{-1/2}\ell_n$) and thus can be substantially smaller than $\epsilon_n = n^{-1/2}\ell_n(K+1)$ derived in the proof of Lemma 3.3. Using these arguments in the proof of Proposition 3.1 may lead to improved bounds in Theorem 3.1 but we omit the formal statements for brevity of the chapter. ■

3.1.2. Restricted Non-Linear Case. We would like to find some useful conditions for nonlinear models, where the rates of convergence of the RMD estimator will be similar to what we have in the linear case.

Condition NLID. Assume that Condition LID holds and that the target moment function g satisfies a restricted non-linearity condition around $\theta = \theta_0$; namely

$$\{\|g(\theta) - g(\theta_0)\|_\infty \leq \epsilon, \theta \in \mathcal{R}(\theta_0)\} \text{ implies } \|G(\theta - \theta_0)\|_\infty/2 \leq \epsilon$$

for all $\epsilon \leq \epsilon^*$, where ϵ^* is a tolerance parameter, measuring the degree of the linearity of the problem, with $\epsilon^* = \infty$ in the linear case.

This condition is the gradient version of restricted convexity for ℓ_1 -penalized M-estimators in [8] and [93].

Example 10. As a reference case, we can take nonlinear models with

$$\epsilon^* = c\mu_n^2/(s^2)$$

for some constant $c > 0$. This case arises, for example, in nonlinear moment condition models of the sort $E\rho(Y, W'\theta_0)Z = 0$, where $\rho: \mathbb{R}^2 \rightarrow \mathbb{R}$ is a well-behaved residual function (e.g. as would arise in nonlinear regression and nonlinear IV regression). This nonlinearity creates an additional requirement on the effective sparsity of θ_0 , namely

$$s^2 \ll n/\mu_n,$$

which arises when we analyze estimation. ■

Lemma 3.4 (Bounds on Rate Function in Non-Linear Models). (i) Under Conditions LID(a) and NLID, Condition MID holds for all $\epsilon \leq \epsilon^*$ with

$$r(\epsilon; \theta_0, \ell_q) \leq 2\epsilon s^{1/q} \mu_n^{-1}, \quad q \in \{1, 2\}.$$

(ii) Under Conditions LID(b) and NLID, Condition MID holds for all $\epsilon \leq \epsilon^*$ with

$$r(\epsilon; \theta_0, \ell_q) \leq C_{a,q} \left(L_n \epsilon s^{1/q} \mu_n^{-1} + \epsilon s^{1/q} \right), \quad q \in \{1, 2\}$$

as long as $a > 1$ and $A > \epsilon$, where $C_{a,q}$ is a constant depending only on a and q .

This lemma can be proven using the same argument as that leading to Lemma 3.2, so we omit the proof.

Next, let there be sequences of positive constants $(B_{1n})_{n \geq 1}$ and $(B_{2n})_{n \geq 1}$. We consider an example of a sufficient condition that allows us to bound the empirical error in estimating θ_0 .

Condition ENM. (i) The target and empirical moment functions have the form $g(\theta) = E[g(X, \theta)]$ and $\hat{g}(\theta) = \mathbb{E}_n[g(X, \theta)]$, respectively, where $g(X, \theta) = (g_1(X, \theta), \dots, g_m(X, \theta))'$ is a vector of score functions, corresponding to the RGMM example. (ii) The score functions have the index form:

$$g_j(X, \theta) = \tilde{g}_j(X, Z_{u(j)}(X)' \vartheta_{u(j)}), \quad j \in [m], \quad \theta = (\vartheta'_1, \dots, \vartheta'_u)' \in \mathbb{R}^p, \quad (3.15)$$

where \tilde{g}_j is a measurable map from $\mathbb{R}^{d_x} \times \mathbb{R}$ to \mathbb{R} for all $j \in [m]$, Z_u is a measurable map from \mathbb{R}^{d_x} to \mathbb{R}^{p_u} for all $u \in [\bar{u}]$, $u(j) \in [\bar{u}]$ for all $j \in [m]$, ϑ_u is p_u -dimensional subvector of θ for all $u \in [\bar{u}]$, and $p = p_1 + \dots + p_{\bar{u}}$. (iii) The score functions are Lipschitz in the second argument, namely

$$|\tilde{g}_j(X, v) - \tilde{g}_j(X, \tilde{v})| \leq L_j(X)|v - \tilde{v}|, \quad \text{for all } (v, \tilde{v}) \in \mathbb{R}^2, \quad (3.16)$$

with probability one, where L_j is a measurable map from \mathbb{R}^{d_x} to \mathbb{R}_+ for all $j \in [m]$. (iv) Finally, we have

$$\text{Var}(g_j(X, \theta) - g_j(X, \theta_0)) \leq B_{1n}^2 \quad (3.17)$$

for all $\theta \in \mathcal{R}(\theta_0)$ and $j \in [m]$ and with probability at least $1 - \delta_n/6$,

$$\max_{j \in [m], k \in [p_{u(j)}]} \mathbb{E}_n[L_j^2(X) Z_{u(j)k}^2(X)] \leq B_{2n}^2, \quad \|n^{-1/2} \mathbb{G}_n(g(X, \theta_0))\|_\infty \leq n^{-1/2} \ell_n. \quad (3.18)$$

In many applications, the number of indices $Z_u(X)' \vartheta_u$ is small. As examples, the common class of single-index models clearly use just one index, $Z_u(X)' \vartheta_u = Z(X)' \theta$; and we would have two indexes - the supply index $Z_1(X)' \vartheta_1$ and the demand index $Z_2(X)' \vartheta_2$ - if we estimate linear supply and demand equations. The Lipschitz condition allows us to use Ledoux-Talagrand type contraction inequalities for bounding the error. Condition ENM, though plausible in a number of applications, is strong. Its chief appeal is in immediately providing useful bounds on the empirical error. One could also obtain bounds on the empirical error through the use of maximal inequalities that carefully exploit the geometry and entropy properties of the set of functions $\{g(X, \theta) : \theta \in \mathcal{R}(\theta_0)\}$.

Lemma 3.5 (Empirical Moment Concentration). *Assume that Conditions DM and ENM are satisfied. Then Condition EMC holds with*

$$\epsilon_n = n^{-1/2}(\tilde{\ell}_n + \ell_n),$$

where $\tilde{\ell}_n = C(B_{1n} + B_{2n}K \log^{1/2}(pn/\delta_n))$ and C is a universal constant.

Next, we demonstrate how Condition ENM can be verified in particular examples. Specifically, we consider logistic regression and nonlinear IV regression models.

Example 11 (Logistic Regression Model). Let $Y \in \{0, 1\}$ be a binary outcome of interest and $W = (W_1, \dots, W_p)' \in \mathbb{R}^p$ a vector of covariates linked by a logistic model, namely

$$\mathbb{E}[Y | W] = \Lambda(W' \theta_0), \quad \Lambda(t) = \exp(t) / \{1 + \exp(t)\}, \quad t \in \mathbb{R}.$$

The vector of score functions associated with this model is

$$g(X, \theta) = W(Y - \Lambda(W' \theta)), \quad X = (Y, W)'$$

and so $g_j(X, \theta) = \tilde{g}_j(X, W' \theta)$, where

$$\tilde{g}_j(X, t) = W_j(Y - \Lambda(t)), \quad t \in \mathbb{R}, \quad j \in [p].$$

Suppose that for some $\sigma > 0$,

$$\max_{j \in [p]} \mathbb{E}[W_j^4] \leq \sigma^4, \quad \frac{\log(pn)}{\sqrt{n}} (\mathbb{E}[\|W\|_\infty^4])^{1/2} \leq \sigma^2, \quad \frac{\log(pn)}{\sqrt{n}} (\mathbb{E}[\|W\|_\infty^8])^{1/2} \leq \sigma^4. \quad (3.19)$$

Then (3.16) holds with $L_j(X) = |W_j|$ since Λ is 1-Lipschitz. Moreover, since $|\Lambda(t) - \Lambda(\tilde{t})| \leq 1$ for any $(t, \tilde{t}) \in [0, 1]^2$,

$$\text{Var}(g_j(X, \theta) - g_j(X, \theta_0)) \leq \mathbb{E}[W_j^2] \leq \sigma^2$$

for any $\theta \in \mathbb{R}^p$ by the first inequality in (3.19). Therefore, (3.17) holds with any $B_{1n}^2 \geq \sigma^2$. Also, by the first and third inequalities in (3.19), it follows from Lemmas A.4 and A.5, where the latter is applied with $s = 2$ and $t = n\sigma^4$, that

$$\max_{j \in [p], k \in [p]} \mathbb{E}_n[W_j^2 W_k^2] \leq \max_{j \in [p]} \mathbb{E}_n[W_j^4] \leq C_1 \sigma^4$$

with probability at least $1 - c_1 / \log^2 n$, where c_1, C_1 are some universal constants. Finally, by the same arguments as those in Example 9, it follows from the first and second inequalities in (3.19) that

$$\|\mathbb{G}_n(g(X, \theta_0))\|_\infty \leq C_2 \sigma \sqrt{\log(pn)}$$

with probability at least $1 - c_2 / \log^2 n$, where c_2, C_2 are some universal constants. Conclude, by the union bound, that (3.18) holds probability at least $1 - \delta_n/6$ as long as we set

$$B_{2n}^2 \geq C_1 \sigma^4, \quad \delta_n = 6(c_1 + c_2) / \log^2 n, \quad \ell_n = C_2 \sigma \sqrt{\log(pn)}.$$

We can thus establish all requirements of Condition ENM. \blacksquare

Example 7 (Nonlinear IV Regression Model, Continued). Consider the model

$$\mathbb{E}[f(Y, W' \theta_0) \mid Z] = 0,$$

where $Y \in \mathbb{R}$ is an outcome variable, $W \in \mathbb{R}^p$ is a vector of endogenous covariates, $Z \in \mathbb{R}^m$ is a vector of instruments, $f: \mathbb{R}^2 \rightarrow \mathbb{R}$ is some known function, and $\theta_0 \in \mathbb{R}^p$ is a vector of parameters of interest. A vector of score functions associated with this model is

$$g(X, \theta) = Z f(Y, W' \theta), \quad X = (Y, W', Z')',$$

and so $g_j(X, \theta) = \tilde{g}_j(X, \theta)$, where

$$\tilde{g}_j(X, t) = Z_j f(Y, t), \quad t \in \mathbb{R}, \quad j \in [m].$$

Suppose that the function $f: \mathbb{R}^2 \rightarrow \mathbb{R}$ is Lipschitz in its second argument:

$$|f(Y, t) - f(Y, \tilde{t})| \leq \gamma(Y) |t - \tilde{t}|, \quad \text{for all } t, \tilde{t} \in \mathbb{R},$$

with probability one. Suppose also that for some $\sigma > 0$,

$$\max_{j \in [m]} \mathbb{E}[Z_j^2 f^2(Y, W' \theta_0)] \leq \sigma^2, \quad \frac{\log(mn)}{\sqrt{n}} (\mathbb{E}[\|Z\|_\infty^4 f^4(Y, W' \theta_0)])^{1/2} \leq \sigma^2, \quad (3.20)$$

$$\max_{j \in [m], k \in [p]} \mathbb{E}[\gamma^2(Y) Z_j^2 W_k^2] \leq \sigma^4, \quad \frac{\log(pmn)}{\sqrt{n}} (\mathbb{E}[\gamma^4(Y) \|Z\|_\infty^4 \|W\|_\infty^4])^{1/2} \leq \sigma^4. \quad (3.21)$$

Then (3.16) holds with $L_j(X) = Z_j \gamma(Y)$ for all $j \in [m]$. Also, for any $j \in [m]$ and $\theta \in \mathcal{R}(\theta_0)$,

$$\begin{aligned} |g_j(X, \theta) - g_j(X, \theta_0)| &\leq \gamma(Y) |Z_j W'(\theta - \theta_0)| \\ &\leq \gamma(Y) |Z_j| \|W\|_\infty \|\theta - \theta_0\|_1 \leq 2\gamma(Y) |Z_j| \|W\|_\infty \|\theta_0\|_1, \end{aligned}$$

and so (3.17) holds for all

$$B_{1n}^2 \geq 4 \|\theta_0\|_1^2 \max_{j \in [m]} \mathbb{E}[\gamma^2(Y) |Z_j|^2 \|W\|_\infty^2]. \quad (3.22)$$

Further, like in Example 11, by (3.21), it follows from Lemmas A.4 and A.5 that

$$\max_{j \in [m], k \in [p]} \mathbb{E}_n[\gamma^2(Y)Z_j^2W_k^2] \leq C_1\sigma^4$$

with probability at least $1 - c_1/\log^2(n)$, and by (3.20), it follows from Lemmas A.2 and A.3 that

$$\|\mathbb{G}_n(g(X, \theta_0))\|_\infty \leq C_2\sigma\sqrt{\log(mn)}$$

with probability at least $1 - c_2/\log^2(n)$, where c_1, C_1, c_2 and C_2 are universal constants. Conclude, by the union bound, that (3.18) holds with probability at least $1 - \delta_n/6$ as long as we set

$$B_{2n}^2 \geq C_1\sigma^4, \quad \delta_n = 6(c_1 + c_2)/\log^2 n, \quad \ell_n = C_2\sigma\sqrt{\log(mn)}.$$

We have thus verified all assumptions of Condition ENM. Note also that the conditions we give here are sufficient but sometimes are not necessary. For example, if we assume that the function $f: \mathbb{R}^2 \rightarrow \mathbb{R}$ is bounded in absolute value by a constant C , then (3.17) holds for all

$$B_n^2 \geq 4C \max_{j \in [m]} \mathbb{E}[Z_j^2].$$

Depending on the setting, this bound can be better than (3.22). ■

The following result is an immediate corollary of Proposition 3.1 and Lemmas 3.4 and 3.5.

Theorem 3.2 (Bounds on Empirical Error for Non-Linear RGMM). *Consider the non-linear case and assume that Conditions L, DM, LID, NLID, and ENM are satisfied. Also, assume that λ is chosen so that $\lambda \leq n^{-1/2}(\tilde{\ell}_n + \ell_n)$ and that the side condition $n^{-1/2}(\tilde{\ell}_n + \ell_n) \leq \epsilon^*/2$ holds. Then with probability at least $1 - \alpha - \delta_n$,*

$$\|\hat{\theta} - \theta_0\|_q \leq \frac{2(\tilde{\ell}_n + \ell_n)s^{1/q}}{\mu_n\sqrt{n}}, \quad q \in \{1, 2\},$$

in the case of LID(a) (exactly sparse model); and, as long as $a > 1$, and $A > n^{-1/2}\ell_n(K + 1)$,

$$\|\hat{\theta} - \theta_0\|_q \leq \frac{C_{a,q}(L_n + \mu_n)(\tilde{\ell}_n + \ell_n)s^{1/q}}{\mu_n\sqrt{n}}, \quad q \in \{1, 2\}$$

in the case of LID(b) (approximately sparse model), where $C_{a,q}$ is a constant depending only on a and q .

Theorem 3.2 shows that, under sparsity conditions, the RGMM estimator can be consistent for θ_0 in the ℓ_q -norm in nonlinear models. Importantly, the dependence of the convergence rates on the total number of parameters and moment conditions is controlled by ℓ_n and $\tilde{\ell}_n$, which typically grow logarithmically with p and m as shown in Examples 11 and 7. Thus consistency is possible even for high-dimensional models when the number of parameters exceeds the sample size. These results also highlight the different rates of convergence for different norms of interest. In particular, the RGMM estimator has good rates

of convergence in the ℓ_1 and ℓ_2 norms. However, the rate of convergence of the RGMM estimator in the max-norm is not optimal in many cases of interest; and additional tools, and estimators, are needed to obtain good estimators for that case.

3.2. Double/De-Biased RGMM. Section 3.1 focused on rates of convergence of regularized minimum distance estimators. We now turn to providing inferential statements for parameters of interest where we will leverage the obtained rates of convergence of the RGMM estimator. Our development will emphasize inference via the Neyman orthogonality principle. Specifically, we aim to construct moment equations $M(\alpha; \eta) = 0$ for the target parameters $\alpha \in \mathbb{R}^{p_1}$ given the nuisance parameter $\eta \in \mathbb{R}^p$ such that the true value α_0 of the parameter α obeys

$$M(\alpha_0; \eta_0) = 0, \quad (3.23)$$

where η_0 is the true value of the nuisance parameter and such that the equations are first-order insensitive to local perturbations of the nuisance parameter η around the true value:

$$\left. \partial_{\eta'} M(\alpha_0, \eta) \right|_{\eta=\eta_0} = 0. \quad (3.24)$$

We refer to the latter property as the Neyman orthogonality condition.

Given M , estimation of or inference about α_0 will then be based on some estimator \hat{M} of M , where we plug-in a potentially biased, regularized estimator $\hat{\eta}$ in place of the unknown η_0 . The role of the Neyman orthogonality condition, (3.24), is precisely to mitigate the impact of the use of such biased estimators (or other non-regular estimators) of η_0 on the estimation of α_0 . In many settings, basing estimation and inference for α_0 on estimating equations with the Neyman orthogonality property will allow us

- to construct a high-quality estimator of α_0 with an approximately Gaussian distribution or
- to carry out high-quality inference on α_0 via hypothesis testing (in cases where α_0 is weakly identified).

3.3. Construction of Approximate Mean Estimators for θ via Orthogonal Moments. Here we take the target parameter and the nuisance parameter to be the same, namely

$$\alpha = \theta, \quad \eta = \theta.$$

The estimator for the nuisance parameter will be $\hat{\eta} = \hat{\theta}$, the RGMM estimator from the previous subsections.

We can construct the Neyman orthogonal equations $M(\alpha; \eta)$ for the pair (α, η) as follows. First we define an optimal moment selection matrix:

$$\gamma_0 = G' \Omega^{-1},$$

where $\Omega = \mathbb{E}g(X, \theta_0)g(X, \theta_0)'$ and $G = (\partial/\partial\theta')g(\theta)|_{\theta=\theta_0}$. This $p \times m$ moment selection matrix can be used to collapse m -dimensional moment equations $g(\theta_0) = 0$ to p -dimensional moment equations $\gamma_0 g(\theta_0) = 0$. We will focus on the optimal moment selection matrix, although, in principle, sub-optimal moment selection matrices could be used in practice as well. For example, estimation of the m by m matrix Ω and its inverse might be a limiting factor in some settings, and one might consider using $\text{diag}(\Omega)$ instead as its inverse is trivial to compute.

Given the moment selection matrix, we define

$$M(\alpha; \eta) = \gamma_0 G(\alpha - \eta) + \gamma_0 g(\eta). \quad (3.25)$$

We then have that at the true parameter values

$$M(\theta_0; \theta_0) = 0,$$

and the Neyman orthogonality condition holds:

$$\partial_{\eta'} M(\theta_0; \eta) \Big|_{\eta=\theta_0} = -\gamma_0 G + \gamma_0 G = 0.$$

Heuristically, if we somehow knew γ_0 , G , and $\eta = \theta_0$, we could define an "oracle" linear estimator of $\alpha_0, \bar{\theta}$, as the root of

$$\bar{M}(\bar{\theta}; \theta_0) = \gamma_0 G(\bar{\theta} - \theta_0) + \gamma_0 \hat{g}(\theta_0) = 0;$$

that is

$$\sqrt{n}(\bar{\theta} - \theta_0) = -(\gamma_0 G)^{-1} \gamma_0 \sqrt{n} \hat{g}(\theta_0).$$

This estimator is linear, so it obeys

$$\sqrt{n}(\bar{\theta} - \theta_0) \approx_d N(0, V), \quad V = (G' \Omega^{-1} G)^{-1},$$

over the sets in \mathcal{A} , the class of all rectangles in \mathbb{R}^p , under the CLT conditions of Section 2. The variance matrix V here is the optimal variance matrix for GMM.

The above construction is infeasible because of the many unknowns, including the true value of the target parameter, appearing in it. To make construction feasible, we can plug-in estimators corresponding to the unknowns: instead:

- (1) plug-in the RGMM estimator $\hat{\theta}$ in place of θ_0 ,
- (2) plug-in an estimator $\hat{G} = \partial_{\theta'} \hat{g}(\hat{\theta})$ for G (or a regularized version);
- (3) plug-in an estimator $\hat{\Omega} = \mathbb{E}_n g(X, \hat{\theta})g(X, \hat{\theta})'$ for Ω (or a regularized version);
- (4) plug-in a regularized estimator $\hat{\gamma}$ for $\gamma_0 := G' \Omega^{-1}$, such that $\hat{\gamma}$ is well-behaved;
- (5) plug-in a regularized estimator $\hat{\mu}$ for $\mu_0 := (\gamma_0 G)^{-1}$, such that $\hat{\mu}$ is well-behaved.

Specific choices of estimators $\hat{\mu}$ and $\hat{\gamma}$ will be discussed later. Given estimators of all unknown quantities, we define a "two-step" estimator of the target parameter by solving

the estimated Neyman-orthogonal equation:

$$\hat{M}(\theta; \hat{\theta}) = \hat{\mu}^{-1}(\theta - \hat{\theta}) + \hat{\gamma}\hat{g}(\hat{\theta}) = 0.$$

The resulting solution to this equation, $\check{\theta}$, provides an estimator of the target parameter which we refer to as the double/debiased regularized GMM (DRGMM) estimator:

$$\sqrt{n}(\check{\theta} - \hat{\theta}) = -\hat{\mu}\hat{\gamma}\sqrt{n}\hat{g}(\hat{\theta}) \text{ or equivalently } \check{\theta} = \hat{\theta} - \hat{\mu}\hat{\gamma}\hat{g}(\hat{\theta}) \quad (3.26)$$

By exploiting the Neyman orthogonality property and further assumptions on the problem, we can show that this estimator approximates the infeasible “oracle” estimator defined above in the sense that

$$\sqrt{n}(\check{\theta} - \theta_0) = \sqrt{n}(\bar{\theta} - \theta_0) + o_P(1/\sqrt{\log p}).$$

Hence, the DRGMM estimator is also approximately linear and is therefore an approximate mean, so we have

$$\sqrt{n}(\check{\theta} - \theta_0) \approx_d N(0, V), \quad V = (G'\Omega^{-1}G)^{-1}, \quad (3.27)$$

over the class of all rectangles in \mathbb{R}^p under the conditions of Theorem 2.1 in Section 2. Indeed, given this construction, we are back to the MAM framework. We can thus use the inferential tools from Section 2 for immediate construction of simultaneous confidence bands and hypothesis testing with control of FWER or FDR. Inference done in this way will be optimal in the sense that the variance matrix V can not be generally improved by using any other moment selection matrix $\bar{\gamma}$ in place of γ_0 . Optimality may also be attained in other semi-parametric senses, which we do not discuss.

3.4. Testing Parameters α with Nuisance Parameters $\eta = \theta$ via Neyman-Orthogonal Scores. Here we take the target parameter, α , and the nuisance parameter, η , to be the different:

$$\alpha = \alpha, \quad \eta = \theta.$$

The true value of the parameter is given by $(\alpha'_0, \eta'_0)'$ and solves

$$g(\alpha_0, \theta_0) = 0.$$

Here, we are thinking of a situation where η_0 is strongly identified and can be well-estimated by RGMM while α_0 is only weakly or partially identified. We would thus like to use a robust testing approach to test values of α_0 and then invert to construct a confidence set for α_0 .

The estimator for the nuisance parameter will be $\hat{\eta} = \hat{\theta}$, the RGMM estimator from Section 3.1. We can construct the Neyman orthogonal equations $M(\alpha; \eta)$ for the pair (α, η) as follows. First, we define a $p' \times m$ moment selection matrix for α_0 ,

$$\xi_0, \text{ e.g. } \xi_0 = I \text{ or } \xi_0 = G_\alpha\Omega^{-1}, \quad G_\alpha = \partial_{\alpha'}g(\alpha_0, \theta_0),$$

where $p' \geq \dim(\alpha)$. The latter matrix will be optimal when α_0 is strongly identified.

Given the moment selection matrix, we define

$$M(\alpha; \theta) = (\xi_0 - \xi_0 G(G' \Omega^{-1} G)^{-1} G \Omega^{-1}) g(\alpha, \theta) = (\xi_0 - \xi_0 G \mu_0 \gamma_0) g(\alpha, \theta) \quad (3.28)$$

as in [54]. Using this estimating equation, we have that, at the true values of the parameters,

$$M(\alpha_0; \theta_0) = 0$$

using $g(\alpha_0, \theta_0) = 0$ and that the Neyman orthogonality condition holds:

$$\partial_{\eta'} M(\theta_0; \theta) \Big|_{\theta=\theta_0} = (\xi_0 - \xi_0 G(G' \Omega^{-1} G)^{-1} G \Omega^{-1}) G = 0.$$

Heuristically, if we knew θ_0 and all the extra parameters used in forming M , we could use the oracle Neyman-orthogonal score for testing α_0 :

$$\sqrt{n} \bar{M}(\alpha_0; \theta_0) = (\xi_0 - \xi_0 G \mu_0 \gamma_0) \sqrt{n} \hat{g}(\alpha_0, \theta_0).$$

This quantity is clearly linear in $\sqrt{n} \hat{g}(\alpha_0, \theta_0)$, so it obeys

$$\sqrt{n} \bar{M}(\alpha_0; \theta_0) \approx_d N(0, V_M), \quad V_M = (\xi_0 - \xi_0 G \mu_0 \gamma_0) \Omega (\xi_0 - \xi_0 G \mu_0 \gamma_0)',$$

over the sets in \mathcal{A} , the class of all rectangles in \mathbb{R}^p , under the CLT conditions of Section 2.

The above construction is again clearly infeasible. As outlined in Section 3.3, we can make the construction feasible by plugging in estimators for the various missing unknowns. Specifically, we will

- (1) plug-in the RGMM estimator $\hat{\theta}$ in place of θ_0 ,
- (2) plug-in an estimator $\hat{G} = \partial_{\theta'} \hat{g}(\hat{\theta})$ for G (or a regularized version);
- (3) plug-in an estimator $\hat{\Omega} = \mathbb{E}_n g(X, \hat{\theta}) g(X, \hat{\theta})'$ for Ω (or a regularized version);
- (4) plug-in a regularized estimator $\hat{\gamma}$ for $\gamma_0 := G' \Omega^{-1}$, such that $\hat{\gamma}$ is well-behaved;
- (5) plug-in a regularized estimator $\hat{\mu}$ for $\mu_0 := (\gamma_0 G)^{-1}$, such that $\hat{\mu}$ is well-behaved;
- (6) plug-in a regularized estimator $\hat{\xi}$ for ξ_0 .

We discuss estimation of $\hat{\mu}$, $\hat{\gamma}$, and $\hat{\xi}$ further in Section 3.5.

Given plug-in estimates of the unknown quantities, we define the Neyman-orthogonal score function for testing α_0 :

$$\sqrt{n} \hat{M}(\alpha_0; \hat{\theta}) = \sqrt{n} (\hat{\xi}_0 - \hat{\xi}_0 \hat{G} \hat{\mu}_0 \hat{\gamma}_0) \hat{g}(\alpha_0, \hat{\theta}).$$

By exploiting Neyman orthogonality property and further assumptions on the problem, we can show that this feasible score approximates the infeasible ‘‘oracle’’ score,

$$\sqrt{n} \hat{M}(\alpha_0; \hat{\theta}) = \sqrt{n} \bar{M}(\alpha_0; \hat{\theta}) + o_P(1/\sqrt{\log p}).$$

Hence, the feasible score is approximately linear and is therefore an approximate mean. We then have

$$\sqrt{n} \hat{M}(\alpha_0; \hat{\theta}) \approx_d N(0, V_M), \quad (3.29)$$

over the class of all rectangles in \mathbb{R}^p under the conditions of Theorem 2.1 in Section 2. We are thus back within the setting outlined in Section 2 and may use the inferential tools from Section 2 for construction of simultaneous confidence bands and hypothesis testing with control of FWER or FDR. Given that we can provide valid inferential statements based on (3.29) for any α_0 , we can invert to obtain confidence regions.

3.5. Analysis of DRGMM. In order to analyze the estimator (3.26), we can write, using elementary expansions and some algebra,

$$\sqrt{n}(\check{\theta} - \theta_0) = -\mu_0\gamma_0\sqrt{n}\hat{g}(\theta_0) + r, \quad (3.30)$$

where

$$r = r_1 + r_2 + r_3 \quad \left\{ \begin{array}{l} r_1 = \sqrt{n}(I - \hat{\mu}\hat{\gamma}\hat{G})(\hat{\theta} - \theta_0) \\ r_2 = \sqrt{n}(\hat{\mu}\hat{\gamma}(\hat{G} - \tilde{G}))(\hat{\theta} - \theta_0). \\ r_3 = \sqrt{n}(\hat{\mu}\hat{\gamma} - \mu_0\gamma_0)\hat{g}(\theta_0) \end{array} \right. \quad (3.31)$$

In (3.31), $\tilde{G} = \{-\partial_{\theta'}\hat{g}_k(\hat{\theta}_k^*)\}_{k=1}^m$ denotes a $m \times p$ matrix with rows $-\partial_{\theta'}\hat{g}_k(\hat{\theta}_k^*)$, $k \in [m]$, where each row is evaluated at a point $\hat{\theta}_k^*$ on the line between $\hat{\theta}$ and θ_0 .

Note that because of Neyman orthogonality property we expect the term r_1 to be small, in fact if we knew (μ_0, γ_0, G) the term would vanish. In linear moment models, $\hat{G} = \tilde{G}$, so that the second term vanishes, $r_2 = 0$. The last term can also vanish under mild conditions. We analyze the structure of these remainder terms in the lemma given below.

To fix ideas, we record a trivial proposition.

Proposition 3.2 (Approximate Linearity and Normality of DRGMM). *If the remainder term r obeys the conditions of Section 2, then DRGMM is an approximate mean estimator:*

$$\sqrt{n}(\check{\theta} - \theta_0) = \frac{1}{\sqrt{n}} \sum_{i=1}^n Z_i + r, \quad Z_i = -\mu_0\gamma_0 g(X_i, \theta_0).$$

Each component of the estimator is approximately normally distributed and satisfies self-normalized moderate deviations results of Section 2, provided the Z_i 's obey the regularity conditions of Section 2. The distribution of $\sqrt{n}(\check{\theta} - \theta_0)$ over rectangles can be approximated by bootstrapping the scores

$$\hat{Z}_i = \hat{\mu}\hat{\gamma}g(X_i, \hat{\theta}),$$

provided these obey the regularity conditions of Section 2. Consequently, the results on simultaneous inference with FWER control or pointwise testing with FDR control apply.

A crucial step to using Proposition 3.2 is to ensure the approximation error r is small. The following lemma is useful for thinking about estimators of γ_0 and μ_0 which are suitably well-behaved to obtain small approximation errors. Of course, the lemma only suggests one possible direction, and there are other strategies for estimating γ_0 and μ_0 to explore. In what follows, we use v_j to denote the j th row of some matrix v .

Lemma 3.6. *We have that the approximation errors r_1 , r_2 and r_3 as defined in (3.31) satisfy*

$$\begin{aligned} \|r_1\|_\infty &\leq \bar{r}_1 = \sqrt{n} \|I - \hat{\mu} \hat{\gamma} \hat{G}\|_\infty \|\hat{\theta} - \theta_0\|_1 \\ \|r_2\|_\infty &\leq \bar{r}_2 = \sqrt{n} \max_{j \in [p]} \|\hat{\mu}_j\|_1 \max_{j \in [p]} \|\hat{\gamma}_j\|_1 \|\hat{G} - \tilde{G}\|_\infty \|\hat{\theta} - \theta_0\|_1 \\ \|r_3\|_\infty &\leq \bar{r}_3 = \max_{j \in [p]} \|\hat{\mu}_j\|_1 \max_{j \in [p]} \|\hat{\gamma}_j - \gamma_{0j}\|_1 \|\sqrt{n} \hat{g}(\theta_0)\|_\infty \\ &\quad + \max_{j \in [p]} \|\hat{\mu}_j - \mu_{0j}\|_1 \max_{j \in [p]} \|\gamma_{0j}\|_1 \|\sqrt{n} \hat{g}(\theta_0)\|_\infty \end{aligned}$$

with simplifications occurring in the linear case, when $\tilde{G} = \hat{G}$, and in the case where the moment selection matrix is known, $\hat{\gamma} = \gamma_0$.

Proof of Lemma 3.6. In what follows we denote by e_j the coordinate vector, with 1 in the j -th position and 0 in the other positions. To obtain the first bound, we have that by Hölder's inequality

$$\|r_1\|_\infty \leq \sqrt{n} \max_{j \in [p]} \|e'_j - \hat{\mu}_j \hat{\gamma} \hat{G}\|_\infty \|\hat{\theta} - \theta_0\|_1 \leq \bar{r}_1.$$

The second bound follows from multiple applications of the Hölder's inequality

$$\begin{aligned} \|r_2\|_\infty &\leq \|\hat{\mu} \hat{\gamma} (\hat{G} - \tilde{G})\|_\infty \sqrt{n} \|\hat{\theta} - \theta_0\|_1 \\ &\leq \max_{j,k \in [p]} \|\hat{\mu}_j\|_1 \|\hat{\gamma}(\hat{G} - \tilde{G})\|_{\cdot,k} \sqrt{n} \|\hat{\theta} - \theta_0\|_1 \\ &= \max_{j,k \in [p]} \|\hat{\mu}_j\|_1 \max_{l \in [p]} \|\hat{\gamma}_l\|_1 \|\hat{G} - \tilde{G}\|_{l,k} \sqrt{n} \|\hat{\theta} - \theta_0\|_1 \\ &\leq \max_{j \in [p]} \|\hat{\mu}_j\|_1 \max_{l \in [p]} \|\hat{\gamma}_l\|_1 \|\hat{G} - \tilde{G}\|_\infty \sqrt{n} \|\hat{\theta} - \theta_0\|_1. \end{aligned}$$

Finally, we bound r_3 . It follows from Hölder's inequality and the triangle inequality that

$$\begin{aligned} \|r_3\|_\infty &\leq \max_{j \in [p]} |(\hat{\mu}_j \hat{\gamma} - \mu_{0j} \gamma_0) \sqrt{n} \hat{g}(\theta_0)| \\ &\leq \max_{j \in [p]} |\hat{\mu}_j (\hat{\gamma} - \gamma_0) \sqrt{n} \hat{g}(\theta_0)| + \max_{j \in [p]} |(\hat{\mu}_j - \mu_{0j}) \gamma_0 \sqrt{n} \hat{g}(\theta_0)| \\ &\leq \max_{j \in [p]} \|\hat{\mu}_j\|_1 \|(\hat{\gamma} - \gamma_0) \sqrt{n} \hat{g}(\theta_0)\|_\infty + \max_{j \in [p]} \|\hat{\mu}_j - \mu_{0j}\|_1 \|\gamma_0 \sqrt{n} \hat{g}(\theta_0)\|_\infty \\ &\leq \max_{j \in [p]} \|\hat{\mu}_j\|_1 \max_{j \in [p]} |(\hat{\gamma} - \gamma_0)_j \cdot \sqrt{n} \hat{g}(\theta_0)| \\ &\quad + \max_{j \in [p]} \|\hat{\mu}_j - \mu_{0j}\|_1 \max_{j \in [p]} |(\gamma_0)_j \cdot \sqrt{n} \hat{g}(\theta_0)| \\ &\leq \max_{j \in [p]} \|\hat{\mu}_j\|_1 \max_{j \in [p]} \|(\hat{\gamma} - \gamma_0)_j\|_1 \|\sqrt{n} \hat{g}(\theta_0)\|_\infty \\ &\quad + \max_{j \in [p]} \|\hat{\mu}_j - \mu_{0j}\|_1 \max_{j \in [p]} \|(\gamma_0)_j\|_1 \|\sqrt{n} \hat{g}(\theta_0)\|_\infty. \end{aligned}$$

■

A plausible approach is then to construct estimators $\hat{\gamma}$, $\hat{\mu}$, and $\hat{\theta}$ such that the upper bounds \bar{r}_1 , \bar{r}_2 , and \bar{r}_3 given in Lemma 3.6 approach zero sufficiently fast. A natural choice of $\hat{\theta}$ is given by the RGMM estimator discussed in Section 3.1. We can also obtain estimators \hat{G} and $\hat{\Omega}$ by plug-in expressions. We now turn to estimating γ_0 and μ_0 .

First, we consider one potential estimator for γ_0 . We do not attempt to use a standard plug-in estimate since $\hat{\Omega}$ will not be full rank in high-dimensional settings, so its inverse will be ill-posed even if Ω^{-1} is well-behaved. Instead, we define the estimator $\hat{\gamma}$ as the solution to the program:

$$\min_{\gamma \in \mathbb{R}^{p \times m}} \sum_{j \in [p]} \|\gamma_j\|_1 : \quad \|\gamma_j \hat{\Omega} - (\hat{G}')_j\|_\infty \leq \lambda_j^\gamma, \quad j \in [p], \quad (3.32)$$

where λ^γ is a vector of regularization parameters. By allowing $\lambda^\gamma > 0$, the constraint in (3.32) requires solving the inverse problem only approximately, which is needed to handle the rank deficiency of $\hat{\Omega}$.

We proceed similarly in our proposed estimator for μ_0 . We define the estimator $\hat{\mu}$ as the solution to the following program:

$$\min_{\mu \in \mathbb{R}^{p \times m}} \sum_{j \in [p]} \|\mu_j\|_1 : \quad \|\mu_j \hat{\gamma} \hat{G} - e'_j\|_\infty \leq \lambda_j^\mu, \quad j \in [p] \quad (3.33)$$

where e_j is a coordinate vector with 1 in the j -th position and 0 elsewhere and λ^μ is a vector of regularization parameters. Again, the use of positive regularization parameters $\lambda^\mu > 0$ allows us to work with approximate solutions which are needed to cope with the high-dimensionality.

We now summarize an algorithm for constructing the estimator $\check{\theta}$.

Algorithm for DRGMM.

- Step 1. Compute the RGMM estimator $\hat{\theta}$.
- Step 2. Use the plug-in rules $\hat{G} = \partial_{\theta'} \hat{g}(\hat{\theta})$ and $\hat{\Omega} = \mathbb{E}_n g(X, \hat{\theta}) g(X, \hat{\theta})'$.
- Step 3. Obtain the estimator $\hat{\gamma}$ as defined in (3.32).
- Step 4. Obtain the estimator $\hat{\mu}$ as defined in (3.33).
- Step 5. Update the initial RGMM estimator $\check{\theta} = \hat{\theta} - \hat{\mu} \hat{\gamma} \hat{g}(\hat{\theta})$.

We note that the regularized problems (3.32) and (3.33) can be cast as linear programming problems and can be solved separately by row $j \in [p]$. Both features are convenient from a computational perspective.

Next we proceed to analyze the properties of the estimators. The following lemma provides high-level conditions on the estimators of G and Ω and on the penalty choices to derive the needed ℓ_1 -rates of convergence for the rows of $\hat{\gamma}$ and $\hat{\mu}$.

Lemma 3.7. *Let $n^{1/2} \|\hat{\Omega} - \Omega\|_\infty \leq \ell_n^\Omega$ and $n^{1/2} \|\hat{G} - G\|_\infty \leq \ell_n^G$ with probability $1 - \delta_n$ and suppose that $\max_{j \in [p]} \|\gamma_{0j}\|_1 \leq K$. Let the penalty parameters satisfy $n^{1/2} \lambda_j^\gamma \geq K \ell_n^\Omega + \ell_n^G$ and $\lambda_j^\gamma \leq n^{-1/2} \ell_n$ for $j \in [p]$. Suppose that Condition LID(γ_{0j}, Ω) holds for each $j \in [p]$. Then, with*

probability $1 - 3\delta_n$ we have

$$\max_{j \in [p]} \|\hat{\gamma}_j - \gamma_{0j}\|_1 \leq \frac{C_{a,1} s \ell_n (2 + K)(L_n + \mu_n)}{\mu_n \sqrt{n}}.$$

Suppose that $\max_{j \in [p]} \|\mu_{0j}\|_1 \leq K$. Let the penalty parameters in (3.33) satisfy $n^{1/2} \lambda_j^\mu \geq 2K^2 \ell_n^G + K^3 \ell_n^\Omega + K^2 \max_{j \in [m]} n^{1/2} \lambda_j^\gamma$ and $\lambda_j^\mu \leq n^{-1/2} \ell_n'$ for $j \in [p]$. Suppose Condition LID($\mu_{0j}, G' \Omega^{-1} G$) holds for each $j \in [p]$. Then with probability $1 - \delta_n$ we have

$$\max_{j \in [p]} \|\hat{\mu}_j - \mu_{0j}\|_1 \leq \frac{C_{a,1} s \ell_n' (2 + K)(L_n + \mu_n)}{\mu_n \sqrt{n}}$$

Lemma 3.7 builds upon the theory of RGMM with linear score functions established in Section 3.1. As expected, condition LID is assumed to hold for the different Jacobian matrices. Lemma 3.7 also highlights sufficient conditions on how the penalty parameters should be chosen. Moreover, it assumes that the estimators $\hat{\Omega}$ and \hat{G} have good rates of convergence in the ℓ_∞ -norm.

The following lemmas complement the result of Lemma 3.7 by providing conditions and explicit bounds on the rates of convergence for the estimators of Ω and G in the linear and non-linear case.

Lemma 3.8 (Linear Score). *Consider the case of linear score, $g(\theta) = G\theta + g(0)$ where $G_{kj} = \mathbb{E}[G_{kj}(X)]$, $g_k(0) = \mathbb{E}[g_k(X, 0)]$. Let $\Omega = \mathbb{E}[g(X, \theta_0)g(X, \theta_0)']$ and $\hat{\Omega} = \mathbb{E}_n[g(X, \hat{\theta})g(X, \hat{\theta})']$. Suppose that:*

- (i) $\max_{k \in [m], j \in [p]} \mathbb{E}[G_{kj}^2(X)] \leq \sigma^2$, $\max_{k \in [m]} \mathbb{E}[g_k^2(0)] \leq \sigma^2$;
- (ii) $n^{-1/2} \mathbb{E}[\max_{i \in [n]} \|G(X_i)\|_\infty^2] \leq \delta_n \log^{-1/2}(2m)$;
- (iii) with probability $1 - \delta_n$ we have $\max_{k \in [m]} \mathbb{E}_n[\{G_k(X)'(\hat{\theta} - \theta_0)\}^2] \leq \Delta_{2n}^2$;
- (iv) $c \leq \max_{k \in [m]} \mathbb{E}[g_k^4(X, \theta_0)] \leq C$, $n^{-1/2} \mathbb{E}[\max_{i \in [n]} \|g(X, \theta_0)\|_\infty^4] \leq \delta_n \log^{-1/2}(2m)$.

Then, with probability $1 - C\delta_n$ we have

$$\begin{aligned} \|\hat{G} - G\|_\infty &\leq C\sigma \sqrt{n^{-1} \log(2m)} \\ \|\hat{\Omega} - \Omega\|_\infty &\leq C' \sqrt{n^{-1} \log(2m)} + C\Delta_{2n} + \Delta_{2n}^2 \end{aligned}$$

The moment assumptions in Lemma 3.8 are quite standard and allow for $m \gg n$. Requirement (iii) relies on the rate of convergence of $\hat{\theta}$ which impacts the estimation of Ω only. We note that there are examples in which we can bypass this term such as the homoskedastic linear instrumental variable case discussed in Theorem 3.3.

Next we provide conditions to derive bounds on the estimation error of G and Ω for the non-linear case. The conditions will assume a Lipschitz condition on the score and its derivative as stated in (3.15) and (3.16) of condition ENM.

Lemma 3.9 (Non-Linear Score). *Let the score $(g_k, k \in [m])$ and $(\partial_j g_k, j \in [p], k \in [m])$, $m \geq p$, satisfy conditions (3.15) and (3.16), with $(L_k(X), Z_k(X))_{k=1}^m$ and $(\tilde{L}_{kj}(X), \tilde{Z}_{kj}(X))_{k \in [m], j \in [p]}$ respectively. Suppose further that:*

(i) *for all $j \in [m], l \in [p]$, $\mathbb{E}[L_j^2(X)Z_{jl}^2(X)] \leq B_n^2$, $\mathbb{E}[|L_j(X)Z'_j v|^2 \{1 + g_k^2(X, \theta_0)\}] \leq C\|v\|^2$;*

and for each $k, j \in [m], l \in [p]$, $\mathbb{E}[\tilde{L}_{kj}^2(X)\tilde{Z}_{kjl}^2(X)] \leq B_n^2$, and $\mathbb{E}[|\tilde{L}_{kj}(X)\tilde{Z}'_{kj} v|^2] \leq C\|v\|^2$;

(ii) *with probability $1 - \delta_n$ we have $\max_{k, j \in [m], l \in [p]} \mathbb{E}_n[L_j^2(X)Z_{kl}^2 \{1 + g_k^2(X, \theta_0)\}] \leq B_n^2$;*

(iii) *with probability $1 - \delta_n$ we have $\|\hat{\theta} - \theta_0\|_\ell \leq \Delta_{\ell n}$ for $\ell \in \{1, 2\}$;*

(iv) $\mathbb{E}[g_k^4(X, \theta_0)] \leq C$, $k \in [m]$, $n^{-1/2}\mathbb{E}[\max_{i \in [n]} \|g(X_i, \theta_0)\|_\infty^4] \leq \delta_n \wedge \log^{-1/2} m$;

(v) $\mathbb{E}[G_{kj}^2(X, \theta_0)] \leq C$, $k \in [m], j \in [p]$; $n^{-1/2}\mathbb{E}[\max_{i \in [n]} \|G(X_i, \theta_0)\|_\infty^2] \leq \delta_n \wedge \log^{-1/2} m$.

Then, with probability $1 - C'\delta_n$ we have

$$\begin{aligned} \|\hat{G} - G\|_\infty &\leq C'\sqrt{n^{-1}\log(2m)} + C'B_n\Delta_{1n}\sqrt{n^{-1}\log(m/\delta_n)} + C'\Delta_{2n} \\ \|\hat{G} - \tilde{G}\|_\infty &\leq C'\sqrt{n^{-1}\log(2m)} + C'B_n\Delta_{1n}\sqrt{n^{-1}\log(m/\delta_n)} + C'\Delta_{2n} \\ \|\hat{\Omega} - \Omega\|_\infty &\leq C'\sqrt{n^{-1}\log(2m)} + C'B_n\Delta_{1n}\sqrt{n^{-1}\log(mp/\delta_n)} + 2B_n^2\Delta_{1n}^2 + 2C\Delta_{2n}. \end{aligned}$$

The moment conditions are quite standard. The bounds depend on the ℓ_1 and ℓ_2 -rates of convergence of the initial RGMM estimator $\hat{\theta}$.

The following theorems provide results that builds upon the RGMM estimator discussed in Section 3.1 and builds upon the previous lemmas to deliver the approximate linear expansion (3.30).

We begin with a result for the homoskedastic linear instrumental variables model. Let $Y = W'\theta_0 + \epsilon$ with $\mathbb{E}\epsilon Z = 0$, $\mathbb{E}\epsilon^2 Z Z' = \sigma^2 \mathbb{E} Z Z'$. Then, using the moment function $g(\theta) = \mathbb{E}[(Y - W'\theta)Z]$, we have $G = -\mathbb{E} Z W'$, $g(0) = \mathbb{E} Y Z$, and $\Omega = \sigma^2 \mathbb{E} Z Z'$. In this homoskedastic setting, we will compute $\hat{\gamma}$ in (3.32) with $\hat{\Omega} = \mathbb{E}_n Z Z'$.

Theorem 3.3 (Homoskedastic Linear IV Model). *Consider the homoskedastic high-dimensional linear instrumental variable model. Suppose:*

(1) $\|\theta_0\|_1 \leq K$, $\max_{j \in [p]} \|\gamma_{0j}\|_1 \leq K$ and $\max_{j \in [p]} \|\mu_{0j}\|_1 \leq K$;

(2) Conditions $LID(\theta_0, G)$, $LID(\gamma_{0j}, \Omega)$ and $LID(\mu_{0j}, G'\Omega^{-1}G)$ hold for $j \in [p]$;

(3) $\max_{k \in [m], j \in [p]} \mathbb{E}[Z_k^2 W_j^2] + \mathbb{E}[Y^2 Z_k^2] \leq C$, $c \leq \max_{k \in [m]} \mathbb{E}[\epsilon^4 Z_k^4] \leq C$;

(4) $n^{-1/2}\{\mathbb{E}[\max_{i \in [n]} \|Z_i W_i'\|_\infty^2] + \mathbb{E}[\max_{i \in [n]} \|\epsilon_i Z_i\|_\infty^4]\} \leq \delta_n \log^{-1/2}(2m)$;

(5) $K + L_n + \sigma^2 + \mu_n^{-1} \leq C$.

Then for $\bar{\lambda} = C'(1 + \sigma)\sqrt{n^{-1}\log(2mn)}$ for some fixed C' sufficiently large, setting $\lambda_j^\gamma = \frac{1}{2}\lambda_j^\mu = \bar{\lambda}$, with probability $1 - C\delta_n$ we have

$$\sqrt{n}(\hat{\theta} - \theta_0) = \frac{(G'\Omega^{-1}G)^{-1}G'\Omega^{-1}}{\sqrt{n}} \sum_{i=1}^n \epsilon_i Z_i + r, \quad \text{with } \|r\|_\infty \leq C u_n$$

provided that $n^{-1}s^2 \log^2(pmn) \leq u_n^2$.

Theorem 3.3 derives an approximate linear representation for the estimator that immediately allows us to construct simultaneous confidence regions for all parameters under the conditions of Section 2. The result exploits the homoskedasticity and bypasses the need to estimate σ^2 . In turn this allows the representation to hold under the mild sparsity requirement of $n^{-1}s^2 \log^2(pmn) \leq u_n^2$.

Under more stringent requirements, the next result considers the non-linear case where the functions g_k satisfies a Lipschitz condition; see condition ENM in Section 3.1.2.

Theorem 3.4 (Non-Linear Case). *Suppose the following conditions hold:*

- (1) $\|\theta_0\|_1 \leq K$, $\max_{j \in [p]} \|\gamma_{0j}\|_1 \leq K$ and $\max_{j \in [p]} \|\mu_{0j}\|_1 \leq K$;
- (2) Conditions $LID(\theta_0, G)$, $LID(\gamma_{0j}, \Omega)$ and $LID(\mu_{0j}, G'\Omega^{-1}G)$ hold for $j \in [p]$;
- (3) Condition ENM holds for the score $(g_k, k \in [m])$ with $(L_k(X), Z_k(X))_{k=1}^m$;
- (4) Condition ENM holds for $(\partial_j g_k, j \in [p], k \in [m])$, $m \geq p$, with $(\tilde{L}_{kj}(X), \tilde{Z}_{kj}(X))_{k \in [m], j \in [p]}$;
- (5) for all $j \in [m], l \in [p]$, $E[L_j^2(X)Z_{jl}^2(X)] \leq B_n^2$, $E[|L_j(X)Z_j'v|^2\{1 + g_k^2(X, \theta_0)\}] \leq C\|v\|^2$; and for each $k, j \in [m], l \in [p]$, $E[\tilde{L}_{kj}^2(X)\tilde{Z}_{kjl}^2(X)] \leq B_n^2$, and $E[|\tilde{L}_{kj}(X)\tilde{Z}_{kjl}'v|^2] \leq C\|v\|^2$;
- (6) with probability $1 - \delta_n$ we have $\max_{k, j \in [m], l \in [p]} \mathbb{E}_n[L_j^2(X)Z_{kl}^2\{1 + g_k^2(X, \theta_0)\}] \leq B_n^2$;
- (7) $E[g_k^4(X, \theta_0)] \leq C$, $k \in [m]$, $n^{-1/2}E[\max_{i \in [n]} \|g(X_i, \theta_0)\|_\infty^4] \leq \delta_n \wedge \log^{-1/2} m$;
- (8) $E[G_{kj}^2(X, \theta_0)] \leq C$, $k \in [m], j \in [p]$; $n^{-1/2}E[\max_{i \in [n]} \|G(X_i, \theta_0)\|_\infty^2] \leq \delta_n \wedge \log^{-1/2} m$.
- (9) $B_n + K + L_n + \mu_n^{-1} \leq C$.

For $\bar{a} \geq 0$ and $C' \geq 1$, let $\bar{\lambda} = C'n^{-\frac{1}{2}+\bar{a}}\Phi^{-1}(1 - (pmn)^{-1})$. Then, setting $\lambda_j^\gamma = \frac{1}{2}\lambda_j^\mu = \bar{\lambda}$, with probability $1 - C\delta_n$ we have

$$\sqrt{n}(\check{\theta} - \theta_0) = -\mu_0\gamma_0\hat{g}(\theta_0) + r, \quad \text{with } \|r\|_\infty \leq Cu_n$$

provided that $n^{-1+2\bar{a}}s^2 \log^2(pmn) \leq u_n^2$ and $\bar{\lambda} \geq Cn^{-1/2}s^{1/2} \log^{1/2}(2m)$ for some large $C > 0$.

Theorem 3.4 provides one approach to constructing estimators with suitable linearization based on the RGMM estimator and the estimators (3.32) and (3.33) for the nuisance parameters. Under suitable choice of penalty parameters, the requirement $n^{-1}s^3 \log^2(pmn) \leq u_n$ where $u_n = o(\log^{-1/2}(p))$ suffices to ensure that approximation errors do not distort the asymptotic coverage of (rectangular) confidence regions. We note that the derivation of practical choices of penalty parameters has drawn considerable attention in the literature. Although Theorem 3.4 allows us to postulate a choice $\bar{\lambda}$ that allows us to cover a class of s -sparse models, it would be of interest to obtain adaptive rules that are theoretically valid and practical. See Remark 3.3 below for some initial discussion.

Remark 3.3 (Practical Considerations for the DRGMM Estimator). We note that the penalty choices discussed in Lemma 3.7 rely on using upper bounds of some unknown quantities. Relying on upper bounds is common in deriving theoretical results, though using upper bounds in finite-samples may result in overpenalization and a deterioration in performance. An alternative approach attempts to make the estimators adaptive to the relevant unknown quantities. For example, we can define an alternative estimator for γ_0 as the

solution to the following optimization problem

$$\min_{\gamma \in \mathbb{R}^{p \times m}} \sum_{j \in [p]} \|\gamma_j\|_1 : \quad \|\gamma_j \hat{\Omega} - (\hat{G}')_j\|_\infty \leq \|\gamma_j\|_1 n^{-1/2} \ell_n^\Omega + n^{-1/2} \ell_n^G, \quad j \in [p]. \quad (3.34)$$

The optimization problem in (3.34) can still be written as a linear programming problem after adding additional variables, [23], [20], and [18]. The benefit of using (3.34) is that it avoids trying to guess K . Similarly, we can have

$$\min_{\mu \in \mathbb{R}^{p \times m}} \sum_{j \in [p]} \|\mu_j\|_1 : \quad \|\mu_j \hat{\gamma} \hat{G} - e'_j\|_\infty \leq \|\mu_j\|_1 \lambda_j^\mu, \quad j \in [p] \quad (3.35)$$

where $\lambda_j^\mu = 2 \max_{j \in [p]} \|\hat{\gamma}_j\|_1 n^{-1/2} \ell_n^G + \max_{j \in [p]} \|\hat{\gamma}_j\|_1^2 n^{-1/2} \ell_n^\Omega + \max_{j \in [p]} \|\hat{\gamma}_j\|_1 \|\lambda^\gamma\|_\infty$. This approach is justified by Hölder's inequality. Other approaches motivated by self-normalization moderate deviation theory leads to non-linear problems that can be handled via conic programming in special cases; see [24] and [17].

4. BIBLIOGRAPHICAL NOTES AND OPEN PROBLEMS

The literature on CLT with increasing dimensions is quite broad, and we refer the reader to Appendix I in [45] for an extensive review on this topic prior to the publication of [45]. In the following discussion, we mainly focus on the development after [45]. The high-dimensional CLT and bootstrap results in this chapter, namely Theorems 2.1–2.3, build upon Proposition 2.1, Corollary 4.2, and Proposition 4.3, respectively, in [50], which, in turn, improves on the results of their earlier paper [45]. Both papers use some important technical tools, such as anti-concentration inequalities and Gaussian comparison theorems, obtained in [49]. [116] provides a helpful exposition of these results targeting mathematically oriented graduate students. [45] also provides several useful applications of these results, including the choice of the regularization parameter for the Dantzig selector, specification testing with a parametric model under the null and a nonparametric one under the alternative, and multiple testing with FWER control. Another useful application is testing many moment inequalities, where the number of moment inequalities is larger than the sample size, which is studied in [46] in detail.

There are several extensions of high-dimensional CLT and bootstrap results of [45, 50]. [58] show that Condition E, which is imposed in Theorems 2.1–2.3, can be slightly improved if we are only concerned with inference based on the empirical bootstrap or if we use the multiplier bootstrap with Gaussian weights e_i replaced by weights satisfying

$$\mathbb{E}[e_i] = 0, \quad \mathbb{E}[e_i^2] = 1, \quad \mathbb{E}[e_i^3] = 1, \quad i \in [n]. \quad (4.1)$$

In particular, they show that the term $\log^7(pn)/n$ in Condition E can be replaced by the term $\log^5(pn)/n$. To obtain their results, they circumvent the Gaussian approximation and work directly with the bootstrap approximation. [123, 46, 122] develop time series extensions of the high-dimensional CLT. [42, 43] develop extensions of the high-dimensional

CLT and bootstrap theorems to U -statistics and randomized incomplete U -statistics, respectively ([42] focuses on the second order case). [81] studies Gaussian approximation to a high-dimensional vector of smooth Wiener functionals by combining the techniques developed in [45, 49, 50] and Malliavin calculus.

An important feature of Theorems 2.1–2.3 is that they provide distributional approximation results for the class of *rectangles*. There are also many related results in the literature if we are interested in other classes of sets. For example, [28, 29] show that a result like (2.5) with \mathcal{A} being the class of all *convex sets* is possible under certain moment conditions if $p = o(n^{2/7})$. More formally, [29] proves the following: Let Z_1, \dots, Z_n be zero-mean independent random vectors in \mathbb{R}^p and suppose that the covariance matrix of $S_n^Z = n^{-1/2} \sum_{i=1}^n Z_i$, $V = n^{-1} \sum_{i=1}^n \mathbb{E}[Z_i Z_i']$, is invertible; then

$$\sup_{A \subset \mathbb{R}^p: \text{convex}} |\mathbb{P}(S_n^Z \in A) - \mathbb{P}(N(0, V) \in A)| \leq \frac{Kp^{1/4}}{n^{3/2}} \sum_{i=1}^n \mathbb{E}[\|V^{-1/2} Z_i\|_2^3], \quad (4.2)$$

where K is a universal constant. In the simplest case where $V = I$ and $\|Z_i\|_2 \leq C\sqrt{p}$ for all $i \in [n]$ and some constant (independent of n), the right-hand side on (4.2) is $O(p^{7/4}n^{-1/2})$, which is $o(1)$ if $p = o(n^{2/7})$. This result is straightforward to use for inference since the covariance matrix V can be accurately estimated as long as $p \ll n$. Recently, [120] improves on the Bentkus condition in the case where Z_1, \dots, Z_n are i.i.d. with identity covariance matrix, and such that $\|Z_i\|_2 \leq C\sqrt{p}$ for all $i \in [n]$ and some constant C ; under these conditions [120] shows that the left-hand side of (4.2) is approaching zero provided that $p = o(n^{2/5})$ up to log factors. [126] shows that the multiplier bootstrap inference over all *centered balls* with weights e_i satisfying (4.1) is possible if $p = o(n^{1/2})$. See also [91] for some early results regarding the multiplier bootstrap with weights e_i satisfying (4.1).

Theorem 2.4 on moderate deviations of self-normalized sums extends the results of [80] to show that the moderate deviation inequality holds, up to some corrections, even if the original random variables in the denominator of the self-normalized sum are replaced by suitable estimators. This is particularly helpful when we allow for many approximate means as opposed to many exact means; see [7] for an application. Textbook-level treatment of the theory of self-normalized sums can be found in [57].

Theorem 2.6 on simultaneous confidence intervals is a rather simple application of Theorems 2.1–2.3. Similar results formulated in terms of particular applications can be found for example in [19], [11], and [17]. Simultaneous confidence intervals also constitute an important research topic in the literature on nonparametric estimation and inference. Useful references on this literature are provided in [48].

Theorems 2.7 and 2.8 on multiple testing with FWER control build upon [103] with the key difference that we allow for $p \rightarrow \infty$, and in particular $p/n \rightarrow \infty$, as $n \rightarrow \infty$. A closely related analog of these theorems can be found in [45] but our conditions here

are somewhat weaker than those in [45]. [118] explains importance of FWER control in multiple testing. A clean textbook-level treatment of multiple testing with FWER control can be found in [87]. A useful discussion of multiple testing in experimental economics can be found in [88]. [104] draws a connection between multiple testing with the FWER control and constructing sets covering the identified sets with a prescribed probability in partially identified models and provide a stepdown procedure for constructing such sets. Note also that our formulation of the Bonferroni-Holm procedure is slightly different but equivalent to the commonly used formulation, e.g. in [87]. We have changed the formulation to facilitate the comparison between the Bonferroni-Holm and Romano-Wolf procedures.

Theorem 2.9 on multiple testing with FDR control generalizes the results of [89] to allow for many approximate means. In turn, [89] generalizes the original results of [26] on the Benjamini-Hochberg procedure to allow for the unknown distribution of the data and also to allow for some dependence between the t -statistics. Moreover, [89] uses moderate deviation for self-normalized sums theory to allow for testing in ultra-high dimensions. Theorem 2.9 is also closely related to the results in [90], who considers multiple testing with FDR control for a specific setting: variable selection in a high-dimensional regression model. A useful discussion of multiple testing with FDR control and other types of control can be found in [101]. A textbook-level treatment is provided in [69].

Our Condition C for the analysis of the Benjamini-Hochberg procedure requires that each t -statistic is correlated with a relatively small set of other t -statistics. The procedure, however, remains valid under the so-called positive regression dependency condition; see [27] for details. There also has been a lot of research about related procedures; see e.g. [102]. A radically different procedure, which also allows for FDR control but which we did not consider in this chapter, is the knockoff filter of [5]. This alternative procedure is designed specifically for the variable selection problem in the linear mean regression model and requires the number of covariates to be smaller than the sample size but does not restrict dependence between covariates in any way. See also [35] where the knockoff filter is modified to allow for the high-dimensional regression model, with the number of covariates exceeding the sample size, in exchange for some other conditions.

We present our high-dimensional estimation results within the context of ℓ_1 -regularized minimum distance estimation focusing on the case where the parameter vector exhibits an approximately sparse structure. The estimation results clearly build upon fundamental work for ℓ_1 -penalized regression of [66] and [107]. This initial work has been expanded in many directions; see, for example, the textbook treatment of [32] as well as [36], [30], [8], [21], [7], [14], and [12] for results most closely related to the approach taken in this review.

More generally, providing methods for estimating high-dimensional models has been an active area of research for quite some time, and there is a large collection of methods available within the literature. [73] and [61] provide useful textbook introductions to a

wide array of methods that are useful in high-dimensional contexts. Developing new techniques for estimation in high-dimensional settings is also still an active area of research, so the list of methods available to researchers continues to expand. Further exploring the use of these procedures in economic applications and the impact of their use on inference about structural parameters seems like a useful avenue to pursue.

Methods for obtaining valid inferential statements following regularization in high-dimensional settings has been an active area of research in the recent statistics and econometrics literature. Early work on inference in high-dimensional settings focused on the exact sparsity structure with separation (from zero) discussed in Section 2.6.2; see, e.g., [62] for an early paper or [63] for a more recent review. A consequence of sparsity with strong separation from zero is that model selection does not impact the asymptotic distribution of the parameters estimated in the selected model, under regularity conditions. This property allows one to do inference using standard approximate distributions for the parameters of the selected model ignoring that model selection was done. While convenient, inferential results obtained relying on this structure may perform very poorly in more realistic approximately sparse structures as was noted in a series of papers; see, for example, [85] and [86].

The more recent work on inference about model parameters following the use of regularization, including the procedure outlined in this chapter, has focused on providing procedures that remain valid without maintaining exact sparsity with separation. As noted in Section 3.2, a key element in obtaining valid inferential statements is the use of estimating equations that satisfy the Neyman orthogonality condition. This idea dates at least to [96] who used the idea of projecting the score that identifies the parameter of interest onto the ortho-complement of the tangent space for nuisance parameters in the construction of the $C(\alpha)$, or orthogonal score, statistic. This idea also plays a key role in semiparametric and targeted learning theory; see, for example, [2], [94], [112], [105], and [111].

Within the high-dimensional context, much of the work on inference focuses on inference for prespecified low-dimensional parameters in the presence of high-dimensional nuisance parameters when ℓ_1 regularization or variable selection methods are used to estimate the nuisance parameters. [15] considers inference about parameters on a low-dimensional set of endogenous variables following selection of instruments from a high-dimensional set using lasso in a homoscedastic, Gaussian IV model. Their approach relies on the fact that the moment condition underlying IV estimation is Neyman orthogonal. These ideas were further developed in the context of providing uniformly valid inference about the parameters on endogenous variables in the IV context with many instruments to allow non-Gaussian heteroscedastic disturbances in [7]. [14], which to our knowledge provides the first formal statement of the Neyman orthogonality condition in the high-dimensional setting, covers inference on the parametric components of the partially linear model and average treatment effects. See also [13], [65], [82], [16], and [12], among others,

for further applications and generalizations explicitly making use of Neyman orthogonal estimating equations. As noted above, Neyman orthogonal estimating equations are closely related to Neyman's $C(\alpha)$ -statistic. The use of $C(\alpha)$ statistics for testing and estimation with high-dimensional approximately sparse models was first explored in the context of quantile regression in [19] and in the context of high-dimensional generalized linear models by [22]. Other uses of $C(\alpha)$ -statistics or close variants include those in [114], [99], [119], and [98]. Finally, a different strand of the literature has focused on ex-post "de-biasing" of estimators to enable valid inference as opposed to directly basing estimation and inference on orthogonal estimating equations. While seemingly distinct, the de-biasing approach is the same as approximately solving orthogonal estimating equations; see, for example, discussion in [54]. Important seminal contributions following the de-biasing approach are [124], [109], and [79].

Rather than focus on a pre-specified low-dimensional parameter, one may also do inference for high-dimensional parameters. [117] and [92] use sample splitting to provide procedures for multiple inference in high-dimensional settings that can control FWER and FDR under strong conditions. [97] also consider the construction of confidence sets for the entire parameter vector in a sparse high-dimensional regression using sample splitting ideas. [109] suggest using Bonferroni-Holm in conjunction with the de-sparsified lasso using a limiting distribution derived under homoscedastic Gaussian errors. As discussed in this chapter, the high-dimensional CLT and bootstrap results of [45] are broadly applicable for inference about high-dimensional parameters. [19] provides an early use of these results for construction of a simultaneous confidence rectangle for many target parameters within a rich class of models estimated using orthogonal estimating equations; see also [55] which implements inference for high-dimensional treatment or structural effects following [19]. [59] and [121] also consider bootstrap inference for many parameters estimated with debiased estimators in high-dimensional models building on [45]. More recently, [11] extends [19] to provide valid inference for many functional parameters, and [18] consider inference for many parameters in a high-dimensional linear model with errors in variables. Finally, [34], [127], [128], and [72] consider different approaches which allow testing hypotheses about functionals that may involve the entire high-dimensional parameter vector within different high-dimensional contexts.

There is also a rapidly growing body of research focused on learning economically interesting parameters that uses different high-dimensional methods and/or aims to provide reliable inferential statement under relatively weaker conditions. Data-adaptive estimation of nuisance functions, with an emphasis on using high-dimensional methods, is advocated in the targeted learning literature under the nomenclature "super learner" though many formal results in this literature are obtained in low-dimensional settings; see, e.g., [110], [111], and [125]. [3] is an important example that uses tree-based methods and sample splitting for estimating and performing inference about heterogeneous treatment effects. [115] considers estimation and inference for heterogeneous treatment effects using a variant of random forests with formal results established in a low-dimensional context. In

a high-dimensional setting, [44] use Neyman orthogonal estimating equations and sample splitting to provide a generic procedure for inference about low-dimensional parameters under weak conditions that allow for the use of wide variety of high-dimensional, machine learning methods. Sample splitting and orthogonal estimating equations are also employed in [53], which considers inference for high-dimensional conditional heterogeneous treatment effects. These ideas are also extended in [52] which provides inference for a variety of useful functionals of heterogeneous treatment effects, such as the best linear predictor of the conditional average treatment effect function, estimated via generic high-dimensional methods using sample splitting while accounting for uncertainty introduced from the sample splits under very mild conditions. [17] provides multipurpose inference methods for high-dimensional causal effects which cover endogenous treatments. [4] use a reweighting after regression adjustment via lasso which allows valid inference for the average treatment effect to be performed under very weak conditions on the propensity score as long as treatment and control conditional mean functions are linear. [56] considers inference for linear functionals of conditional expectations with estimated Riesz representers under very weak conditions. Note that among applications of [56] is estimation of average treatment effects, and that the conditions of [56] also impose weak assumptions on the propensity score. The advantage of [56]’s approach over that in [4] is that the former explicitly allows the tradeoff in the rate of estimating the inverse propensity score with the rate of estimating the regression function, allowing misspecification of both functions. The regularity conditions are also substantively weaker, for example allowing regressions functions be completely non-sparse when the inverse propensity score is well-approximated.

Most of the work in the recent literature on high-dimensional estimation and inference relies on approximate sparsity to provide dimension reduction and the corresponding use of sparsity-based estimators. Dense models are appealing in many settings and may be usefully employed in more moderate-dimensional settings. The many weak-instrument regime, popular in econometrics since at least [6], provides one such example. Approaches which provide valid inference for structural parameters within this setting can all be viewed as making use of regularization to avoid dramatic overfitting in the relationship between endogenous variables and the many available instruments. See, for example, [41], [100], [37], and [71] for approaches that explicitly use regularized first-stage estimation within a dense model framework. [39] and [40] consider inference for a low-dimensional set of coefficients in a linear model with number of variables proportional to but smaller than the sample size in a framework allowing for the coefficients on the nuisance variables to be dense. Within the same framework, [38] extend this work to address inference about parameters estimated using two-step procedures where the first step is a linear regression with many variables.

We conclude these bibliographic notes by noting that the references to the high-dimensional literature provided above are necessarily selective. The literature on high-dimensional estimation and inference is large and rapidly expanding, and it is impractical to give more

than a cursory overview highlighting a few examples. The goal of this review is to provide readers with a few key papers in several areas and a taste of existing results.

APPENDIX A. FUNDAMENTAL TOOLS

A.1. Tool Set 1: Moderate Deviation Inequality for Self-normalized Sums.

Lemma A.1. Let ξ_1, \dots, ξ_n be independent mean-zero random variables with $n^{-1} \sum_{i=1}^n \mathbb{E}[\xi_i^2] \geq 1$ and $\mathbb{E}[|\xi_i|^{2+\nu}] < \infty$ for all $i \in [n]$ where $0 < \nu \leq 1$. Let $S_n := \sum_{i=1}^n \xi_i$, $V_n^2 := \sum_{i=1}^n \xi_i^2$, and $D_{n,\nu} := (n^{-1} \sum_{i=1}^n \mathbb{E}[|\xi_i|^{2+\nu}])^{1/(2+\nu)}$. Then uniformly in $0 \leq x \leq n^{2(2+\nu)} / D_{n,\nu}$,

$$\left| \frac{\mathbb{P}(S_n/V_n \geq x)}{1 - \Phi(x)} - 1 \right| \leq K n^{-\nu/2} D_{n,\nu}^{2+\nu} (1+x)^{2+\nu},$$

where K is a universal constant.

Proof. See Theorem 7.4 in [57] or the original paper, [80]. Note that the formulation in these sources requires that $n^{-1} \sum_{i=1}^n \mathbb{E}[\xi_i^2] = 1$ but it is trivial to show that we can instead assume that $n^{-1} \sum_{i=1}^n \mathbb{E}[\xi_i^2] \geq 1$. ■

A.2. Tool Set 2: Maximal and Deviation Inequalities.

Lemma A.2. Let X_1, \dots, X_n be independent random vectors in \mathbb{R}^p . In addition, define $\sigma^2 := \max_{j \in [p]} \sum_{i=1}^n \mathbb{E}[X_{ij}^2]$. Then for every $s > 1$ and $t > 0$,

$$\begin{aligned} \mathbb{P} \left(\max_{j \in [p]} \left| \sum_{i=1}^n (X_{ij} - \mathbb{E}[X_{ij}]) \right| \geq 2\mathbb{E} \left[\max_{j \in [p]} \left| \sum_{i=1}^n (X_{ij} - \mathbb{E}[X_{ij}]) \right| \right] + t \right) \\ \leq e^{-t^2/(3\sigma^2)} + \frac{K_s}{t^s} \sum_{i=1}^n \mathbb{E} \left[\max_{j \in [p]} |X_{ij} - \mathbb{E}[X_{ij}]|^s \right], \end{aligned}$$

where K_s is a constant depending only on s .

Proof. See Lemma E.2 in [50]. ■

Lemma A.3. Let X_1, \dots, X_n be independent random vectors in \mathbb{R}^p with $p \geq 2$. Define $M := \max_{i \in [n]} \max_{j \in [p]} |X_{ij}|$ and $\sigma^2 := \max_{j \in [p]} \sum_{i=1}^n \mathbb{E}[X_{ij}^2]$. Then

$$\mathbb{E} \left[\max_{j \in [p]} \left| \sum_{i=1}^n (X_{ij} - \mathbb{E}[X_{ij}]) \right| \right] \leq K(\sigma \sqrt{\log p} + \sqrt{\mathbb{E}[M^2]} \log p),$$

where K is a universal constant.

Proof. See Lemma 8 in [49]. ■

Lemma A.4. Let X_1, \dots, X_n be independent random vectors in \mathbb{R}^p with $p \geq 2$ such that $X_{ij} \geq 0$ for all $i \in [n]$ and $j \in [p]$. Define $M := \max_{i \in [n]} \max_{j \in [p]} X_{ij}$. Then for any $s \geq 1$ and $t > 0$,

$$\mathbb{P} \left(\max_{j \in [p]} \sum_{i=1}^n X_{ij} \geq 2\mathbb{E} \left[\max_{j \in [p]} \sum_{i=1}^n X_{ij} \right] + t \right) \leq K\mathbb{E}[M^s]/t^s.$$

where K is a constant depending only on s .

Proof. See Lemma E.4 in [50]. ■

Lemma A.5. Let X_1, \dots, X_n be independent random vectors in \mathbb{R}^p with $p \geq 2$ such that $X_{ij} \geq 0$ for all $i \in [n]$ and $j \in [p]$. Define $M := \max_{i \in [n]} \max_{j \in [p]} X_{ij}$. Then

$$\mathbb{E} \left[\max_{j \in [p]} \sum_{i=1}^n X_{ij} \right] \leq K \left(\max_{j \in [p]} \mathbb{E} \left[\sum_{i=1}^n X_{ij} \right] + \mathbb{E}[M] \log p \right),$$

where K is a universal constant.

Proof. See Lemma 9 in [49]. ■

Lemma A.6. Let $(Y_1, \dots, Y_p)^T$ be a Gaussian random vector with $\mathbb{E}[Y_j] = 0$ and $\mathbb{E}[Y_j^2] \leq 1$ for all $j \in [p]$. For $\alpha \in (0, 1)$, let $c(\alpha)$ denote the $(1 - \alpha)$ quantile of the distribution of $\max_{j \in [p]} Y_j$. Then $c(\alpha) \leq \sqrt{2 \log p} + \sqrt{2 \log(1/\alpha)}$.

Proof. By the Borell-Sudakov-Tsirel'son inequality (Theorem A.2.1 in [113]), for every $r > 0$,

$$\mathbb{P} \left(\max_{j \in [p]} Y_j \geq \mathbb{E} \left[\max_{j \in [p]} Y_j \right] + r \right) \leq e^{-r^2/2},$$

which implies that

$$c(\alpha) \leq \mathbb{E} \left[\max_{j \in [p]} Y_j \right] + \sqrt{2 \log(1/\alpha)}. \tag{A.1}$$

In addition, by Proposition A.3.1 in [106],

$$\mathbb{E} \left[\max_{j \in [p]} Y_j \right] \leq \sqrt{2 \log p} \tag{A.2}$$

Combining (A.1) and (A.2) leads to the desired result. ■

A.3. Tool Set 3: High-dimensional Central Limit and Bootstrap Theorem, Gaussian Comparison and Anti-Concentration Inequalities. In Theorems A.1–A.3 below, we will follow the following setting: Let X_1, \dots, X_n be independent zero-mean random vectors

in \mathbb{R}^p , and consider $S_n^X = n^{-1/2} \sum_{i=1}^n X_i$. Denote by $V = n^{-1} \sum_{i=1}^n E[X_i X_i']$ the covariance matrix of S_n^X (which we assume to exist). In addition, let \mathcal{A} denote the class of all rectangles in \mathbb{R}^p , i.e. sets of the form

$$A = \left\{ w = (w_1, \dots, w_p)' \in \mathbb{R}^p : w_{lj} \leq w_j \leq w_{rj} \text{ for all } j \in [p] \right\},$$

where $w_l = (w_{l1}, \dots, w_{lp})'$ and $w_r = (w_{r1}, \dots, w_{rp})'$ are p -dimensional vectors with components in $(\mathbb{R} \cup \{-\infty\} \cup \{+\infty\})^p$.

Theorem A.1 (High-dimensional CLT). *Assume that for some constants $b > 0$ and $B \geq 1$,*

$$\frac{1}{n} \sum_{i=1}^n E[X_{ij}^2] \geq b \quad \text{and} \quad \frac{1}{n} \sum_{i=1}^n E[|X_{ij}|^{2+k}] \leq B \quad \text{for all } j \in [p] \text{ and } k = 1, 2. \quad (\text{A.3})$$

Then the following claims hold: (i) if

$$E[\exp(|X_{ij}|/B)] \leq 2 \text{ for all } i \in [n] \text{ and } j \in [p], \quad (\text{A.4})$$

then

$$\sup_{A \in \mathcal{A}} |\mathbb{P}(S_n^X \in A) - \mathbb{P}(N(0, V) \in A)| \leq K_1 \left(\frac{B^2 \log^7(pn)}{n} \right)^{1/6}, \quad (\text{A.5})$$

where K_1 is a constant depending only on b ; (ii) if for some constant $q \in [3, \infty)$,

$$E \left[\max_{j \in [p]} (|X_{ij}|/B)^q \right] \leq 1 \text{ for all } i \in [n], \quad (\text{A.6})$$

then

$$\sup_{A \in \mathcal{A}} |\mathbb{P}(S_n^X \in A) - \mathbb{P}(N(0, V) \in A)| \leq K_2 \left(\left(\frac{B^2 \log^7(pn)}{n} \right)^{1/6} + \left(\frac{B^2 \log^3(pn)}{n^{1-2/q}} \right)^{1/3} \right), \quad (\text{A.7})$$

where K_2 is a constant depending only on b and q .

Proof. See Proposition 2.1 in [50]. ■

Theorem A.2 (Multiplier Bootstrap). *Let e_1, \dots, e_n be i.i.d. $N(0, 1)$ random variables that are independent of X_1, \dots, X_n and consider $S_n^{eX} = n^{-1/2} \sum_{i=1}^n e_i X_i$. Assume that (A.3) holds for some constants $b > 0$ and $B \geq 1$. Recall $X_1^n = (X_1, \dots, X_n)$. Then the following claims hold: (i) if (A.4) is satisfied, then for any constant $\beta \in (0, e^{-1})$ with probability at least $1 - \beta$,*

$$\sup_{A \in \mathcal{A}} |\mathbb{P}(S_n^{eX} \in A \mid X_1^n) - \mathbb{P}(N(0, V) \in A)| \leq K_1 \left(\frac{B^2 \log^5(pn) \log^2(1/\beta)}{n} \right)^{1/6},$$

where K_1 is a constant depending only on b ; (ii) if (A.6) is satisfied for some constant $q \in [4, \infty)$, then for any constant $\beta \in (0, e^{-1})$ with probability at least $1 - \beta$,

$$\begin{aligned} & \sup_{A \in \mathcal{A}} |\mathbb{P}(S_n^{eX} \in A \mid X_1^n) - \mathbb{P}(N(0, V) \in A)| \\ & \leq K_2 \left(\left(\frac{B^2(\log p + \log(1/\beta)) \log^4 p}{n} \right)^{1/6} + \left(\frac{B^2(\log p + \beta^{-2/q}) \log^2 p}{n^{1-2/q}} \right)^{1/3} \right), \end{aligned}$$

where K_2 is a constant depending only on b and q .

Proof. The claim under (A.4) is proven in Corollary 4.2 of [50]. (Note that there, $S_n^{eX} = n^{-1/2} \sum_{i=1}^n e_i X_i$ is replaced by $n^{-1/2} \sum_{i=1}^n e_i (X_i - \bar{X}_n)$, where $\bar{X}_n := n^{-1} \sum_{i=1}^n X_i$ is the sample mean, but since X_i 's have mean zero, this change is not essential).

The claim under (A.6) improves upon Corollary 4.2 of [50]. To prove the asserted claim, note that, by Remark 4.1 in [50], there exists a constant $K > 0$ depending only on b such that on the event $\Delta_n \leq \bar{\Delta}_n$,

$$\sup_{A \in \mathcal{A}} |\mathbb{P}(S_n^{eX} \in A \mid X_1^n) - \mathbb{P}(N(0, V) \in A)| \leq K \bar{\Delta}_n^{1/3} \log^{2/3} p, \quad (\text{A.8})$$

where

$$\Delta_n := \max_{1 \leq j, k \leq p} \left| \frac{1}{n} \sum_{i=1}^n (X_{ij} X_{ik} - \mathbb{E}[X_{ij} X_{ik}]) \right|.$$

Also, as in the proof of Proposition 4.1 in [50], case (E.2), for all $t > 0$,

$$\mathbb{P} \left(\Delta_n > C \left(\left(\frac{B^2 \log p}{n} \right)^{1/2} + \frac{B^2 \log p}{n^{1-2/q}} \right) + t \right) \leq \exp \left(-\frac{nt^2}{3B^2} \right) + \frac{cB^q}{t^{q/2} n^{q/2-1}},$$

where $c, C > 0$ are constants depending only on b and q . Thus, setting

$$t = \bar{C} \left(\left(\frac{B^2 \log(1/\beta)}{n} \right)^{1/2} + \frac{B^2}{\beta^{2/q} n^{1-2/q}} \right)$$

for sufficiently large \bar{C} , which can be chosen to depend only on b and q , shows that with probability at least $1 - \beta$,

$$\Delta_n \leq (C + \bar{C}) \left(\left(\frac{B^2(\log p + \log(1/\beta))}{n} \right)^{1/2} + \frac{B^2(\log p + \beta^{-2/q})}{n^{1-q/2}} \right).$$

The asserted claim follows by combining this bound with (A.8). \blacksquare

Theorem A.3 (Empirical Bootstrap). Let X_1^*, \dots, X_n^* be i.i.d. random vectors from the empirical distribution of X_1, \dots, X_n and consider $S_n^{X^*} = n^{-1/2} \sum_{i=1}^n (X_i^* - \bar{X}_n)$, where $\bar{X}_n = n^{-1} \sum_{i=1}^n X_i$ is the sample mean. Assume that (A.3) holds for some constants $b > 0$ and $B \geq 1$.

Then the following claims hold: (i) if (A.4) is satisfied, then for any constant $\beta \in (n^{-2}, e^{-1})$ with probability at least $1 - \beta$,

$$\sup_{A \in \mathcal{A}} |\mathbb{P}(S_n^{X^*} \in A \mid X_1^n) - \mathbb{P}(N(0, V) \in A)| \leq K_1 \left(\frac{B^2 \log^7(pn)}{n} \right)^{1/6},$$

where K_1 is a constant depending only on b ; (ii) if (A.6) is satisfied for some constant $q \in [4, \infty)$, then for any constant $\beta \in (n^{-2}, e^{-1})$ with probability at least $1 - \beta$,

$$\begin{aligned} & \sup_{A \in \mathcal{A}} |\mathbb{P}(S_n^{X^*} \in A \mid X_1^n) - \mathbb{P}(N(0, V) \in A)| \\ & \leq K_2 \left(\left(\frac{B^2 \log^7(pn)}{n} \right)^{1/6} + \left(\frac{B^2 \log^3(pn)}{\beta^{2/q} n^{1-2/q}} \right)^{1/3} \right), \end{aligned}$$

where K_2 is a constant depending only on b and q .

Proof. This is Proposition 4.3 in [50]. ■

Lemma A.7 (Gaussian Comparison). Let $Y = (Y_1, \dots, Y_p)'$ and $Z = (Z_1, \dots, Z_p)'$ be mean-zero Gaussian random vectors in \mathbb{R}^p with covariance matrices $\Sigma^Y = (\Sigma_{jk}^Y)_{1 \leq j, k \leq p}$ and $\Sigma^Z = (\Sigma_{jk}^Z)_{1 \leq j, k \leq p}$, respectively. Let $\Delta := \max_{j, k \in [p]} |\Sigma_{jk}^Y - \Sigma_{jk}^Z|$ and let \mathcal{A} be the class of all rectangles in \mathbb{R}^p . Suppose that $\min_{j \in [p]} \Sigma_{jj}^Y \vee \min_{j \in [p]} \Sigma_{jj}^Z \geq \underline{\sigma}^2$ for some constant $\underline{\sigma} > 0$. Then

$$\sup_{A \in \mathcal{A}} |\mathbb{P}(Y \in A) - \mathbb{P}(Z \in A)| \leq C \Delta^{1/3} \log^{2/3} p,$$

where C is a constant that depends only on $\underline{\sigma}$.

Proof. The proof is implicit in the proof of Theorem 4.1 in [50]. This Gaussian comparison inequality here improves slightly on the original Gaussian comparison inequality in [49]. ■

Lemma A.8 (Anti-concentration Inequality). Let $Z = (Z_1, \dots, Z_p)'$ be a mean-zero Gaussian random vector in \mathbb{R}^p with $\sigma_j^2 := \mathbb{E}[Z_j^2] > 0$ for $j \in [p]$. Denote $\underline{\sigma} := \min_{j \in [p]} \sigma_j$. Then for all $\epsilon > 0$ and $x = (x_1, \dots, x_p)'$ in \mathbb{R}^p , we have

$$\mathbb{P}(Z \leq x + \epsilon) - \mathbb{P}(Z \leq x) \leq \frac{\epsilon}{\underline{\sigma}} (\sqrt{2 \log p} + 2), \tag{A.9}$$

where $x + \epsilon = (x_1 + \epsilon, \dots, x_p + \epsilon)'$.

Proof. This is Nazarov's inequality stated in Lemma A.1 in [50]. For a detailed proof of Lemma A.8, see [51]. ■

APPENDIX B. PROOFS FOR SECTION 2

Proof of Theorem 2.1. Take any rectangle $A \in \mathcal{A}$. Then there exist vectors $w_l = (w_{l1}, \dots, w_{lp})'$ and $w_r = (w_{r1}, \dots, w_{rp})'$ in \mathbb{R}^p such that

$$A = \left\{ w = (w_1, \dots, w_p)' \in \mathbb{R}^p : w_{lj} \leq w_j \leq w_{rj} \text{ for all } j \in [p] \right\}.$$

Also, denote $\Delta_n := \delta_n / \sqrt{\log p}$ and let

$$A^+ := \left\{ w = (w_1, \dots, w_p)' \in \mathbb{R}^p : w_{jl} - \Delta_n \leq w_j \leq w_{jr} + \Delta_n \text{ for all } j \in [p] \right\},$$

so that $A^+ \in \mathcal{A}$ is also a rectangle and $A \subset A^+$. Then, by Condition A and linearization (2.1),

$$\mathbb{P}(\sqrt{n}(\hat{\theta} - \theta_0) \in A) \leq \left(\frac{1}{\sqrt{n}} \sum_{i=1}^n Z_i \in A^+ \right) + \beta_n. \quad (\text{B.1})$$

In addition, under Conditions M and E, it follows from Theorem A.1 that

$$\mathbb{P} \left(\frac{1}{\sqrt{n}} \sum_{i=1}^n Z_i \in A^+ \right) \leq \mathbb{P}(N(0, V) \in A^+) + K_1 \delta_n, \quad (\text{B.2})$$

where K_1 is a constant that depends only on q .

Further, for any vector $a = (a_1, \dots, a_p)' \in \mathbb{R}^p$ and any number $b \in \mathbb{R}$, let $a + b$ and $a - b$ be vectors in \mathbb{R}^p defined by $(a + b)_j = a_j + b$ and $(a - b)_j = a_j - b$ for all $j \in [p]$, respectively. Note that $w = (w_1, \dots, w_p)' \in A^+ \setminus A$ implies that either $w_j \leq w_{jr} + \Delta_n$ for all j but $w_j > w_{jr}$ for some j or $w_j \geq w_{jl} - \Delta_n$ for all j but $w_j < w_{jl}$ for some j . Hence, by the union bound,

$$\begin{aligned} \mathbb{P}(N(0, V) \in A^+ \setminus A) &\leq \mathbb{P}(N(0, V) \leq w_r + \Delta_n) - \mathbb{P}(N(0, V) \leq w_r) \\ &\quad + \mathbb{P}(N(0, V) \geq w_l - \Delta_n) - \mathbb{P}(N(0, V) \geq w_l). \end{aligned}$$

In turn, by the anti-concentration inequality, Lemma A.8,

$$\mathbb{P}(N(0, V) \leq w_r + \Delta_n) - \mathbb{P}(N(0, V) \leq w_r) \leq K_2 \Delta_n \sqrt{\log p} \leq K_2 \delta_n$$

and

$$\begin{aligned} &\mathbb{P}(N(0, V) \geq w_l - \Delta_n) - \mathbb{P}(N(0, V) \geq w_l) \\ &= \mathbb{P}(-N(0, V) \leq -w_l + \Delta_n) - \mathbb{P}(-N(0, V) \leq -w_l) \leq K_2 \delta_n \end{aligned}$$

for some universal constant K_2 . Therefore,

$$\mathbb{P}(N(0, V) \in A^+) \leq \mathbb{P}(N(0, V) \in A) + 2K_2 \delta_n.$$

Combining this inequality with (B.1) and (B.2) gives

$$\mathbb{P}(\sqrt{n}(\hat{\theta} - \theta_0) \in A) \leq \mathbb{P}(N(0, V) \in A) + \beta_n + (K_1 + 2K_2) \delta_n, \quad (\text{B.3})$$

which is the upper bound for $P(\sqrt{n}(\hat{\theta} - \theta_0) \in A)$. To establish the lower bound, one can use the same argument with the rectangle

$$A^- := \left\{ w = (w_1, \dots, w_p)' \in \mathbb{R}^p : w_{lj} + \Delta_n \leq w_j \leq w_{rj} - \Delta_n \text{ for all } j \in [p] \right\},$$

and obtain

$$P(\sqrt{n}(\hat{\theta} - \theta_0) \in A) \geq P(N(0, V) \in A) - \beta_n - (K_1 + 2K_2)\delta_n. \quad (\text{B.4})$$

Combining (B.3) and (B.4) gives the asserted claim with $C := 1 + K_1 + 2K_2$. \blacksquare

Proof of Theorem 2.2. First, we consider the case of E.1. Take any rectangle $A \in \mathcal{A}$. Then there exist vectors in \mathbb{R}^p , $w_l = (w_{l1}, \dots, w_{lp})'$ and $w_r = (w_{r1}, \dots, w_{rp})'$, such that

$$A = \left\{ w = (w_1, \dots, w_p)' \in \mathbb{R}^p : w_{lj} \leq w_j \leq w_{rj} \text{ for all } j \in [p] \right\}.$$

Also, denote $\Delta_n := \delta_n / \sqrt{\log(pn)}$. Now, observe that conditional on $(Z_i, \hat{Z}_i)_{i=1}^n$, for all $j \in [p]$, the random variable $n^{-1/2} \sum_{i=1}^n e_i(\hat{Z}_{ij} - Z_{ij})$ is zero-mean Gaussian with variance $n^{-1} \sum_{i=1}^n (\hat{Z}_{ij} - Z_{ij})^2$. Hence, by the Borell-Sudakov-Tsirel'son inequality, for any $\alpha \in (0, 1)$,

$$P_e \left(\max_{j \in [p]} \left| \frac{1}{\sqrt{n}} \sum_{i=1}^n e_i(\hat{Z}_{ij} - Z_{ij}) \right| > \max_{j \in [p]} \sqrt{\mathbb{E}_n[(\hat{Z}_{ij} - Z_{ij})^2]} \left(\sqrt{2 \log p} + \sqrt{2 \log(1/\alpha)} \right) \right)$$

is bounded from above by 2α . Therefore, setting $\alpha = \delta_n/2 \geq 1/(2n)$ and using Condition A shows that there exists a universal constant K_1 such that with probability at least $1 - \beta_n$,

$$P_e \left(\max_{j \in [p]} \left| \frac{1}{\sqrt{n}} \sum_{i=1}^n e_i(\hat{Z}_{ij} - Z_{ij}) \right| > K_1 \Delta_n \right) \leq \delta_n. \quad (\text{B.5})$$

Hence, letting

$$A^+ := \left\{ w = (w_1, \dots, w_p)' \in \mathbb{R}^p : w_{lj} - K_1 \Delta_n \leq w_j \leq w_{rj} + K_1 \Delta_n \text{ for all } j \in [p] \right\},$$

so that $A^+ \in \mathcal{A}$ is also a rectangle and $A \subset A^+$, implies that with probability at least $1 - \beta_n$,

$$P_e \left(\frac{1}{\sqrt{n}} \sum_{i=1}^n e_i \hat{Z}_i \in A \right) \leq P_e \left(\frac{1}{\sqrt{n}} \sum_{i=1}^n e_i Z_i \in A^+ \right) + \delta_n. \quad (\text{B.6})$$

Further, by Theorem A.2, there exists a universal constant K_2 such that

$$P_e \left(\frac{1}{\sqrt{n}} \sum_{i=1}^n e_i Z_i \in A^+ \right) \leq P(N(0, V) \in A^+) + K_2 \delta_n \quad (\text{B.7})$$

holds with probability at least $1 - (2n)^{-1}$. Also, like in the proof of Theorem 2.1,

$$P(N(0, V) \in A^+) \leq P(N(0, V) \in A) + K_1 K_3 \Delta_n \sqrt{\log p} \leq P(N(0, V) \in A) + K_1 K_3 \delta_n,$$

where K_3 is a universal constant. Combining this inequality with (B.6) and (B.7) and noting that $\sqrt{n}(\hat{\theta}^* - \hat{\theta}) = n^{-1/2} \sum_{i=1}^n e_i \hat{Z}_i$ shows that with probability at least $1 - \beta_n - (2n)^{-1}$,

$$P_e(\sqrt{n}(\hat{\theta}^* - \hat{\theta}) \in A) \leq P(N(0, V) \in A) + C\delta_n, \quad (\text{B.8})$$

where $C = 1 + K_2 + K_1 K_3$, which is the upper bound for $P_e(\sqrt{n}(\hat{\theta}^* - \hat{\theta}) \in A)$. To establish the lower bound, one can use the same argument with the rectangle

$$A^- := \left\{ w = (w_1, \dots, w_p)' \in \mathbb{R}^p : w_{lj} + K_1 \Delta_n \leq w_j \leq w_{rj} - K_1 \Delta_n \text{ for all } j \in [p] \right\},$$

and obtain that with probability at least $1 - \beta_n - (2n)^{-1}$,

$$P_e(\sqrt{n}(\hat{\theta}^* - \hat{\theta}) \in A) \geq P(N(0, V) \in A) - C\delta_n. \quad (\text{B.9})$$

Combining (B.8) and (B.9) and noting that the same event (B.5) is used to establish both (B.8) and (B.9) gives the asserted claim for the Gaussian bootstrap in the case of E.1.

In the case of E.2, we proceed through the same steps but we note that (B.7) only holds with probability at least $1 - (\log n)^{-2}/2$. This completes the proof of the theorem. \blacksquare

Proof of Theorem 2.3. Note that we can assume without loss of generality that $\delta_n \leq 1$ since otherwise the claims of the theorem are trivial. Therefore, since $(B_n^2 \log^7(pn)/n)^{1/6} \leq \delta_n$ by Condition E and $B_n \geq 1$, it follows that $\log^7(pn)/n \leq 1$, and so $\delta_n \geq \log^3(pn)/\sqrt{n}$. This bound will be useful later in the proof.

Now, conditional on $(Z_i, \hat{Z}_i)_{i=1}^n$, the random vector

$$\frac{1}{\sqrt{n}} \sum_{i=1}^n (e_i - 1)(\hat{Z}_i - Z_i)$$

is equal in distribution to

$$\frac{1}{\sqrt{n}} \sum_{i=1}^n \left((\hat{Z}_i - Z_i)^* - \frac{1}{n} \sum_{l=1}^n (\hat{Z}_l - Z_l) \right),$$

where $(\hat{Z}_i - Z_i)^*$, $i \in [n]$, are i.i.d. random vectors from the empirical distribution of $\hat{Z}_1 - Z_1, \dots, \hat{Z}_n - Z_n$. Also, by Condition A,

$$\max_{j \in [p]} (\mathbb{E}_n[(\hat{Z}_{ij} - Z_{ij})^2])^{1/2} \leq \delta_n / \log(pn) \quad (\text{B.10})$$

holds with probability at least $1 - \beta_n$. In addition, by the assumption of the theorem,

$$\max_{i \in [n]} \max_{j \in [p]} \left| \hat{Z}_{ij} - Z_{ij} - \frac{1}{n} \sum_{l=1}^n (\hat{Z}_{lj} - Z_{lj}) \right| \leq 2 \quad (\text{B.11})$$

holds with probability at least $1 - \beta_n$. Therefore, applying the union bound and Bernstein's inequality (cf. [113], Lemma 2.2.9 or [31], p.36) on the intersection of events (B.10) and

(B.11) shows that with probability at least $1 - 2\beta_n$, for any $t > 0$,

$$P_e \left(\max_{j \in [p]} \left| \frac{1}{\sqrt{n}} \sum_{i=1}^n (e_i - 1)(\hat{Z}_i - Z_i) \right| > t \right) \leq 2p \exp \left(-\frac{nt^2/2}{n\delta_n^2/\log^2(pn) + 2\sqrt{nt}/3} \right).$$

Hence, since $\delta_n \geq \log^3(pn)/\sqrt{n}$, there exists a universal constant $K_1 > 0$ such that for $\Delta_n := \delta_n/\sqrt{\log(pn)}$, with probability at least $1 - 2\beta_n$,

$$P_e \left(\max_{j \in [p]} \left| \frac{1}{\sqrt{n}} \sum_{i=1}^n (e_i - 1)(\hat{Z}_i - Z_i) \right| > K_1 \Delta_n \right) \leq \delta_n. \quad (\text{B.12})$$

The rest of the proof proceeds through the same steps as those in the proof of Theorem 2.2 with (B.12) replacing (B.5) and applying Theorem A.3 instead of A.2. Note that in the case of E.2, we obtain $\delta_n \log^{1/3} n$ on the right-hand side of the bound (2.9) instead of δ_n , as we had for the Gaussian bootstrap in Theorem 2.2, because the bound in Theorem A.3 is slightly worse than that in Theorem A.2. \blacksquare

Proof of Theorem 2.4. We split the proof into three steps.

Step 1. Here we show that there exists n_1 depending only on \bar{C} such that for all $n \geq n_1$,

$$P \left(\frac{1}{n} \sum_{i=1}^n Z_{ij}^2 \leq \frac{1}{2} \text{ for some } j \in [p] \right) \leq \frac{1}{pn}. \quad (\text{B.13})$$

To prove (B.13), we will use the following inequality: If X_1, \dots, X_n are independent non-negative random variables, then for any $t > 0$,

$$P \left(\sum_{i=1}^n (X_i - \mathbb{E}[X_i]) \leq -t \right) \leq \exp \left(-\frac{t^2}{2 \sum_{i=1}^n \mathbb{E}[X_i^2]} \right), \quad (\text{B.14})$$

which can be found in Exercise 2.9 of [31]. Specifically, since

$$\frac{(2\bar{C})^3 B_n \log^{3/2}(pn)}{\sqrt{n}} \leq \delta_n \leq 1,$$

it follows that there exists n_1 depending only on \bar{C} such that for all $n \geq n_1$,

$$\frac{16B_n^2 \log(pn)}{n} = \frac{16B_n^2 \log^3(pn)}{n \log^2(pn)} \leq \frac{16}{(2\bar{C})^6 \log^2(pn)} \leq 1. \quad (\text{B.15})$$

Then for all $n \geq n_1$ and $j = 1, \dots, p$,

$$\begin{aligned}
 \mathbb{P} \left(\frac{1}{n} \sum_{i=1}^n Z_{ij}^2 \leq \frac{1}{2} \right) &\leq \mathbb{P} \left(\frac{1}{n} \sum_{i=1}^n Z_{ij}^2 \leq \frac{1}{2n} \sum_{i=1}^n \mathbb{E}[Z_{ij}^2] \right) = \mathbb{P} \left(\sum_{i=1}^n Z_{ij}^2 \leq \frac{1}{2} \sum_{i=1}^n \mathbb{E}[Z_{ij}^2] \right) \\
 &= \mathbb{P} \left(\sum_{i=1}^n (Z_{ij}^2 - \mathbb{E}[Z_{ij}^2]) \leq -\frac{1}{2} \sum_{i=1}^n \mathbb{E}[Z_{ij}^2] \right) \leq \exp \left(-\frac{(\sum_{i=1}^n \mathbb{E}[Z_{ij}^2])^2}{8 \sum_{i=1}^n \mathbb{E}[Z_{ij}^4]} \right) \\
 &\leq \exp \left(-\frac{n^2}{8nB_n^2} \right) \leq \exp(-2 \log(pn)) = (pn)^{-2}, \tag{B.16}
 \end{aligned}$$

where the inequality in the first line follows from Condition M, the inequality in the second line from (B.14) with $X_i = Z_{ij}^2$ and $t = 2^{-1} \sum_{i=1}^n \mathbb{E}[Z_{ij}^2]$, the first inequality in the third line from Condition M, and the second inequality in the third line from (B.15). Combining (B.16) with the union bound gives (B.13).

Step 2. Here we show that for all $n \geq n_1$,

$$\mathbb{P} \left(\max_{j \in [p]} \left| \frac{(\sum_{i=1}^n \hat{Z}_{ij}^2)^{1/2}}{(\sum_{i=1}^n Z_{ij}^2)^{1/2}} - 1 \right| > \frac{\delta_n \sqrt{2}}{\log(pn)} \right) \leq \beta_n + (pn)^{-1}. \tag{B.17}$$

To show (B.17), note that for all $n \geq n_1$, the left-hand side of (B.17) is equal to

$$\begin{aligned}
 &\mathbb{P} \left(\max_{j \in [p]} \frac{|(\sum_{i=1}^n \hat{Z}_{ij}^2)^{1/2} - (\sum_{i=1}^n Z_{ij}^2)^{1/2}|}{(\sum_{i=1}^n Z_{ij}^2)^{1/2}} > \frac{\delta_n \sqrt{2}}{\log(pn)} \right) \\
 &\leq \mathbb{P} \left(\max_{j \in [p]} \frac{(\sum_{i=1}^n (\hat{Z}_{ij} - Z_{ij})^2)^{1/2}}{(\sum_{i=1}^n Z_{ij}^2)^{1/2}} > \frac{\delta_n \sqrt{2}}{\log(pn)} \right) \\
 &\leq \mathbb{P} \left(\max_{j \in [p]} \left(\frac{1}{n} \sum_{i=1}^n (\hat{Z}_{ij} - Z_{ij})^2 \right)^{1/2} > \frac{\delta_n}{\log(pn)} \right) + (pn)^{-1} \leq \beta_n + (pn)^{-1},
 \end{aligned}$$

where the second line follows from the triangle inequality, and the third one from Step 1 and Condition A. This gives (B.17).

Step 3. Here we complete the proof. Fix $1 \leq x \leq \bar{C} \log^{1/2}(pn)$ and denote

$$\Delta_n(x) := KB_n(1+x)^3/\sqrt{n} \leq K(2\bar{C})^3 B_n \log^{3/2}(pn)/\sqrt{n} \leq K\delta_n, \tag{B.18}$$

where K is a universal constant from Lemma A.1. Now, fix $j \in [p]$ and observe that for all $n \geq n_1$,

$$\begin{aligned}
 \mathbb{P} \left(\frac{n(\hat{\theta}_j - \theta_{0j})}{(\sum_{i=1}^n \hat{Z}_{ij}^2)^{1/2}} > x \right) &\leq \mathbb{P} \left(\frac{n(\hat{\theta}_j - \theta_{0j})}{(\sum_{i=1}^n Z_{ij}^2)^{1/2}} > x - x \left| \frac{(\sum_{i=1}^n \hat{Z}_{ij}^2)^{1/2}}{(\sum_{i=1}^n Z_{ij}^2)^{1/2}} - 1 \right| \right) \\
 &\leq \mathbb{P} \left(\frac{n(\hat{\theta}_j - \theta_{0j})}{(\sum_{i=1}^n Z_{ij}^2)^{1/2}} > x - \frac{\sqrt{2}x\delta_n}{\log(pn)} \right) + \beta_n + (pn)^{-1}, \tag{B.19}
 \end{aligned}$$

where the second inequality follows from Step 2. Further, by linearization (2.1),

$$n(\hat{\theta}_j - \theta_{0j}) = \sum_{i=1}^n Z_{ij} + \sqrt{nr}r_{nj},$$

and by Condition A,

$$\mathbb{P}\left(\max_{j \in [p]} |r_{nj}| > \delta_n / \sqrt{\log(pn)}\right) \leq \beta_n.$$

Hence, for all $n \geq n_1$, the probability in (B.19) is bounded from above by

$$\mathbb{P}\left(\frac{\sum_{i=1}^n Z_{ij}}{(\sum_{i=1}^n Z_{ij}^2)^{1/2}} > x - \frac{\sqrt{2}x\delta_n}{\log(pn)} - \frac{\sqrt{2}\delta_n}{\sqrt{\log(pn)}}\right) + \beta_n + (pn)^{-1}, \quad (\text{B.20})$$

where we used Step 1 to bound $\sum_{i=1}^n Z_{ij}^2$. Next, since $\delta_n \leq 1$ and $x \geq 1$, there exists universal n_2 such that for all $n \geq n_2$, we have $2\sqrt{2}\delta_n \leq \log(pn)$ and $x \geq 2\sqrt{2}\delta_n / \sqrt{\log(pn)}$. Then for all $n \geq n_2$,

$$0 \leq x - \frac{\sqrt{2}x\delta_n}{\log(pn)} - \frac{\sqrt{2}\delta_n}{\sqrt{\log(pn)}} \leq x \leq \bar{C} \log^{1/2}(pn) \leq n^{1/6} / B_n^{1/3},$$

and so by Lemma A.1, for all $n \geq n_2$, the probability in (B.20) is bounded from above by

$$\left(1 - \Phi\left(x - \frac{\sqrt{2}x\delta_n}{\log(pn)} - \frac{\sqrt{2}\delta_n}{\sqrt{\log p}}\right)\right) (1 + \Delta_n(x)). \quad (\text{B.21})$$

Further, for any $y, z \geq 0$, we have $\Phi(y+z) - \Phi(y) \leq z\phi(y)$, where ϕ is the pdf of the standard normal distribution, and for any $y \geq 1$, $\phi(y) \leq 2y(1 - \Phi(y))$; see Proposition 2.5 in [60]. Hence, denoting $\gamma_n := \sqrt{2}x\delta_n / \log(pn) + \sqrt{2}\delta_n / \sqrt{\log(pn)}$, we have for all $n \geq n_2$ that

$$\begin{aligned} -\Phi\left(x - \frac{\sqrt{2}x\delta_n}{\log(pn)} - \frac{\sqrt{2}\delta_n}{\sqrt{\log(pn)}}\right) &= -\Phi(x - \gamma_n) \leq -\Phi(x) + \gamma_n\phi(x - \gamma_n) \\ &\leq -\Phi(x) + \gamma_n\phi(x)e^{x\gamma_n} \leq -\Phi(x) + 2x\gamma_n e^{x\gamma_n}(1 - \Phi(x)). \end{aligned}$$

Therefore, for all $n \geq n_2$, the expression in (B.21) is bounded from above by

$$(1 - \Phi(x))(1 + 2x\gamma_n e^{x\gamma_n})(1 + \Delta_n(x)).$$

Thus, using

$$x\gamma_n = \frac{\sqrt{2}x^2\delta_n}{\log(pn)} + \frac{\sqrt{2}x\delta_n}{\sqrt{\log(pn)}} \leq \sqrt{2}(\bar{C}^2 + \bar{C})\delta_n \leq \sqrt{2}(\bar{C}^2 + \bar{C}),$$

and (B.18) shows that there exists a constant C depending only on \bar{C} such that for all $n \geq n_0 := n_1 \vee n_2$,

$$\mathbb{P}\left(\frac{n(\hat{\theta}_j - \theta_{0j})}{(\sum_{i=1}^n \hat{Z}_{ij}^2)^{1/2}} > x\right) - (1 - \Phi(x)) \leq C\left(\delta_n(1 - \Phi(x)) + \beta_n + (pn)^{-1}\right).$$

This provides the upper bound in (2.10). To establish the lower bound, we use the same argument and note that for all $n \geq n_2$, we have

$$0 \leq x + \frac{\sqrt{2}x\delta_n}{\log(pn)} + \frac{\sqrt{2}\delta_n}{\sqrt{\log(pn)}} \leq 2x \leq 2\bar{C} \log^{1/2}(pn) \leq n^{1/6}/B_n^{1/3},$$

which ensures that we can use Lemma A.1. This completes the proof of (2.10).

To prove (2.11), we proceed like in the beginning of this step to show that for all $n \geq n_0$ and all $1 \leq x \leq \bar{C} \log^{1/2}(pn)$,

$$\mathbb{P} \left(\max_{j \in [p]} \frac{n(\hat{\theta}_j - \theta_{0j})}{(\sum_{i=1}^n \hat{Z}_{ij}^2)^{1/2}} > x \right) \leq \mathbb{P} \left(\max_{j \in [p]} \frac{\sum_{i=1}^n Z_{ij}}{(\sum_{i=1}^n Z_{ij}^2)^{1/2}} > x \right) + C(\beta_n + (pn)^{-1}).$$

In addition, by the union bound,

$$\mathbb{P} \left(\max_{j \in [p]} \frac{\sum_{i=1}^n Z_{ij}}{(\sum_{i=1}^n Z_{ij}^2)^{1/2}} > x \right) \leq p \max_{j \in [p]} \mathbb{P} \left(\frac{\sum_{i=1}^n Z_{ij}}{(\sum_{i=1}^n Z_{ij}^2)^{1/2}} > x \right),$$

and, by the arguments above, for all $n \geq n_0$, $j \in [p]$, and $1 \leq x \leq \bar{C} \log^{1/2}(pn)$,

$$\mathbb{P} \left(\frac{\sum_{i=1}^n Z_{ij}}{(\sum_{i=1}^n Z_{ij}^2)^{1/2}} > x \right) \leq (1 - \Phi(x))(1 + 2x\gamma_n e^{x\gamma_n})(1 + \Delta_n(x)) \leq 1 - \Phi(x) + C\delta_n(1 - \Phi(x)),$$

so that

$$\mathbb{P} \left(\max_{j \in [p]} \frac{n(\hat{\theta}_j - \theta_{0j})}{(\sum_{i=1}^n \hat{Z}_{ij}^2)^{1/2}} > x \right) \leq p(1 - \Phi(x)) + C(p\delta_n(1 - \Phi(x)) + \beta_n + (pn)^{-1}).$$

Setting $x = \Phi(1 - \alpha/p)$ here gives (2.11) and completes the proof of the theorem. \blacksquare

Proof of Theorem 2.5. Without loss of generality, we can assume that $\delta_n \leq 1$ since otherwise the result is trivial (by choosing C large enough). Also, there exists a universal constant n_1 such that $\log^4(pn) \geq (2\sqrt{2})^6$ for all $n \geq n_1$, and so under Condition E, for all $n \geq n_1$,

$$\frac{(2\sqrt{2})^6 B_n^2 \log^3(pn)}{n} \leq \frac{(2\sqrt{2})^6 \delta_n^6}{\log^4(pn)} \leq \frac{(2\sqrt{2})^6 \delta_n^2}{\log^4(pn)} \leq \delta_n^2.$$

In addition, let n_0 be a universal constant from Corollary 2.1. Clearly, it suffices to prove the result for $n \geq n_0 \vee n_1$ since the result for $n < n_0 \vee n_1$ is trivial. Then, by Corollary 2.1,

$$\mathbb{P} \left(\max_{j \in [p]} \frac{|\sqrt{n}(\hat{\theta}_j - \theta_{0j})|}{(\mathbb{E}_n[\hat{Z}_{ij}^2])^{1/2}} > \sqrt{2 \log(pn)} \right) \leq 2C(\beta_n + n^{-1}),$$

where C is a universal constant, and so with probability at least $1 - \beta_n - 2C(\beta_n + n^{-1})$, for all $j \in [p]$,

$$\begin{aligned} |\sqrt{n}(\hat{w}_j - w_j)(\hat{\theta}_j - \theta_{0j})| &= \frac{|\sqrt{n}(\hat{\theta}_j - \theta_{0j})|}{(\mathbb{E}_n[(\hat{Z}_{ij})^2])^{1/2}} \times |\hat{w}_j - w_j| (\mathbb{E}_n[(\hat{Z}_{ij})^2])^{1/2} \\ &\leq \sqrt{2 \log(pn)} \times (\delta_n / \log(pn)) = \sqrt{2} \delta_n / \sqrt{\log(pn)}, \end{aligned}$$

where we used Condition W. Also, by Conditions A and W, with probability at least $1 - \beta_n$, for all $j \in [p]$,

$$\left| \sqrt{n} w_j (\hat{\theta}_j - \theta_{0j}) - \frac{1}{\sqrt{n}} \sum_{i=1}^n w_j Z_{ij} \right| = w_j |r_{nj}| \leq C_W |r_{nj}| \leq C_W \delta_n / \sqrt{\log(pn)}.$$

Conclude that

$$\sqrt{n} \hat{W}(\hat{\theta} - \theta_0) = \frac{1}{\sqrt{n}} \sum_{i=1}^n W Z_i + \bar{r}_n, \quad (\text{B.22})$$

where $\bar{r}_n = (\bar{r}_{n1}, \dots, \bar{r}_{np})'$ is such that

$$\mathbb{P}(\|\bar{r}_n\|_\infty > (C_W + \sqrt{2}) \delta_n / \sqrt{\log(pn)}) \leq 2\beta_n + 2C(\beta_n + n^{-1}). \quad (\text{B.23})$$

Next, by the triangle inequality,

$$(\mathbb{E}_n[(\hat{w}_j \hat{Z}_{ij} - w_j Z_{ij})^2])^{1/2} \leq |\hat{w}_j - w_j| (\mathbb{E}_n[(\hat{Z}_{ij})^2])^{1/2} + w_j (\mathbb{E}_n[(\hat{Z}_{ij} - Z_{ij})^2])^{1/2},$$

and so by Conditions A and W,

$$\mathbb{P}\left(\max_{j \in [p]} \mathbb{E}_n[(\hat{w}_j \hat{Z}_{ij} - w_j Z_{ij})^2] > (C_W + 1)^2 \delta_n^2 / \log^2(pn)\right) \leq 2\beta_n. \quad (\text{B.24})$$

Thus, applying Theorem 2.1 for linearization (B.22), with (B.23) and (B.24) playing the role of Condition A, shows that

$$\sup_{A \in \mathcal{A}} \left| \mathbb{P}(\sqrt{n} \hat{W}(\hat{\theta} - \theta_0) \in A) - \mathbb{P}(WN(0, V) \in A) \right| \leq (C/2)(\delta_n + \beta_n + n^{-1}) \leq C(\delta_n + \beta_n) \quad (\text{B.25})$$

for some constant C depending only on C_W , and applying Theorem 2.2 shows that

$$\sup_{A \in \mathcal{A}} \left| \mathbb{P}_e(\sqrt{n} \hat{W}(\hat{\theta}^* - \hat{\theta}) \in A) - \mathbb{P}(WN(0, V) \in A) \right| \leq C \delta_n \quad (\text{B.26})$$

holds with probability at least $1 - 2\beta_n - n^{-1}$ in the case of E.1 and $1 - 2\beta_n - (\log n)^{-2}$ in the case of E.2 since the proof of Theorem 2.2 only requires the second part of Condition A.

We are now able to prove (2.16) and (2.17). To prove (2.16), denoting $\epsilon_n := C(\delta_n + \beta_n)$, we have by (B.25) that

$$\begin{aligned} \mathbb{P}(\|\sqrt{n} \hat{W}(\hat{\theta} - \theta_0)\|_\infty \leq \lambda^g (1 - \alpha - \epsilon_n)) &\leq \mathbb{P}(\|N(0, V)\|_\infty \leq \lambda^g (1 - \alpha - \epsilon_n)) + \epsilon_n \\ &\leq 1 - \alpha - \epsilon_n + \epsilon_n = 1 - \alpha, \end{aligned}$$

and, similarly,

$$P(\|\sqrt{n}\hat{W}(\hat{\theta} - \theta_0)\|_\infty \leq \lambda^g(1 - \alpha + \epsilon_n)) \geq 1 - \alpha.$$

Thus,

$$\lambda^g(1 - \alpha - \epsilon_n) \leq \lambda(1 - \alpha) \leq \lambda^g(1 - \alpha + \epsilon_n),$$

which gives (2.16). To prove (2.17), we use the same argument but apply (B.26) instead of (B.25).

Finally, we prove (2.18). Let ξ be a $N(0, 1)$ random variable. Then for any $a \in (0, 1)$,

$$\begin{aligned} P(\|WN(0, V)\|_\infty > \bar{\sigma}\Phi^{-1}(1 - a/(2p))) &\leq \sum_{j=1}^p P(|w_j V_{jj}^{1/2} \xi| > \bar{\sigma}\Phi^{-1}(1 - a/(2p))) \\ &\leq \sum_{j=1}^p P(|\xi| > \Phi^{-1}(1 - a/(2p))) = a, \end{aligned}$$

which implies that $\lambda^g(1 - a) \leq \bar{\sigma}\Phi^{-1}(1 - a/(2p))$. To prove that $\Phi^{-1}(1 - a/(2p)) \leq (2 \log(2p/a))^{1/2}$, we apply the inequality $1 - \Phi(x) \leq \exp(-x^2/2)$, which holds for all $x > 0$ as discussed in Proposition 2.5 of [60], with $x = (2 \log(2p/a))^{1/2}$. This completes the proof of the lemma. ■

Proof of Theorem 2.6. By Theorem 2.5 and its proof,

$$\begin{aligned} P(\|\sqrt{n}\hat{W}(\hat{\theta} - \theta_0)\|_\infty \leq \hat{\lambda}(1 - \alpha)) &\leq P(\|\sqrt{n}\hat{W}(\hat{\theta} - \theta_0)\|_\infty \leq \lambda^g(1 - \alpha + \epsilon_n)) + o(1) \\ &= P(\|WN(0, V)\|_\infty \leq \lambda^g(1 - \alpha + \epsilon_n)) + o(1) = 1 - \alpha + \epsilon_n + o(1) = 1 - \alpha + o(1), \end{aligned}$$

and, similarly,

$$\begin{aligned} P(\|\sqrt{n}\hat{W}(\hat{\theta} - \theta_0)\|_\infty \leq \hat{\lambda}(1 - \alpha)) &\geq P(\|\sqrt{n}\hat{W}(\hat{\theta} - \theta_0)\|_\infty \leq \lambda^g(1 - \alpha - \epsilon_n)) + o(1) \\ &= P(\|WN(0, V)\|_\infty \leq \lambda^g(1 - \alpha - \epsilon_n)) + o(1) = 1 - \alpha - \epsilon_n + o(1) = 1 - \alpha + o(1). \end{aligned}$$

Hence,

$$P(\|\sqrt{n}\hat{W}(\hat{\theta} - \theta)\|_\infty \leq \hat{\lambda}(1 - \alpha)) = 1 - \alpha + o(1),$$

which implies that the confidence intervals in (2.15) satisfy (2.14), which is the first asserted claim. To prove the second asserted claim, simply note that by Theorem 2.5, with probability $1 - o(1)$, $\hat{\lambda}(1 - \alpha) \leq \lambda^g(1 - \alpha + \epsilon_n)$. This completes the proof of the theorem. ■

Proof of Theorem 2.7. Since the set of null hypotheses H_j rejected by the Bonferroni-Holm procedure always contains the set of null hypotheses rejected by the Bonferroni procedure, it suffices to prove the result for the Bonferroni-Holm procedure. To do so, note that whenever $w' \subset w''$, it follows that $|w'| \leq |w''|$, and so

$$c_{1-\alpha, w'} = \Phi^{-1}(1 - \alpha/|w'|) \leq \Phi^{-1}(1 - \alpha/|w''|) = c_{1-\alpha, w''}.$$

Thus, (2.21) is satisfied, and we only need to verify (2.22). To this end, take any $w \subset \mathcal{W}$ and $P \in \mathcal{P}^w$. Denote $\bar{p} := |w|$. By discussion immediately after Theorem 2.4,

$$\Phi^{-1}(1 - \alpha/\bar{p}) \leq (2 \log(\bar{p}/\alpha))^{1/2} \leq (2 \log(\bar{p}n))^{1/2}$$

for all $n \geq 1/\alpha$. Together with $(2\sqrt{2})^3 B_n \log^{3/2}(\bar{p}n)/\sqrt{n} \leq \delta_n$, this makes it possible to apply Theorem 2.4, (2.11), with $\bar{C} = \sqrt{2}$, to show that

$$\begin{aligned} & \mathbb{P}_P \left(\max_{j \in w} \frac{\sqrt{n}(\hat{\theta}_j - \bar{\theta}_{0j})}{(\frac{1}{n} \sum_{i=1}^n \hat{Z}_{ij}^2)^{1/2}} > \Phi^{-1}(1 - \alpha/\bar{p}) \right) \\ & \leq \mathbb{P}_P \left(\max_{j \in w} \frac{\sqrt{n}(\hat{\theta}_j - \theta_{0j})}{(\frac{1}{n} \sum_{i=1}^n \hat{Z}_{ij}^2)^{1/2}} > \Phi^{-1}(1 - \alpha/\bar{p}) \right) \leq \alpha + C(\alpha\delta_n + \beta_n + n^{-1}) \end{aligned} \quad (\text{B.27})$$

for some universal constant C and all $n \geq n_0 \vee n_1$, where n_0 is a universal constant appearing in Theorem 2.4 and n_1 is such that for all $n \geq n_1$, we have $\delta_n \leq 1$ (recall that $\delta_n \searrow 0$). Since the right-hand side of (B.27) and $n_0 \vee n_1$ do not depend on (w, P) and $C(\alpha\delta_n + \beta_n + n^{-1}) = o(1)$, (2.22) holds, and the asserted claim follows. \blacksquare

Proof of Theorem 2.8. Since (2.21) holds trivially, it suffices to prove (2.22). To do so, we will use the following lemma.

Lemma B.1. *For $j \in [p]$, denote $\hat{w}_j := (n^{-1} \sum_{i=1}^n \hat{Z}_{ij}^2)^{-1/2}$ and $w_j := V_{jj}^{-1/2}$. Then Conditions M, E, and A imply that there exists universal constants $n_0 \in \mathbb{N}$ and $C \geq 1$ such that for all $n \geq n_0$, Condition W holds with $C_W = 1$ and with δ_n and β_n replaced by $C\delta_n$ and $\beta_n + \delta_n + (pn)^{-1}$, respectively.*

Proof. The first part of Condition W holds trivially with $C_W = 1$. To prove the second part, we will assume, without loss of generality, that $\delta_n \leq 1$. Observe that by the triangle inequality and Condition M, for all $j \in [p]$,

$$\begin{aligned} |\hat{w}_j - w_j|(\mathbb{E}_n[\hat{Z}_{ij}^2])^{1/2} &= \frac{|(\frac{1}{n} \sum_{i=1}^n \hat{Z}_{ij}^2)^{1/2} - (\frac{1}{n} \sum_{i=1}^n \mathbb{E}[Z_{ij}^2])^{1/2}|}{(\frac{1}{n} \sum_{i=1}^n \mathbb{E}[Z_{ij}^2])^{1/2}} \\ &\leq \frac{|(\frac{1}{n} \sum_{i=1}^n \hat{Z}_{ij}^2)^{1/2} - (\frac{1}{n} \sum_{i=1}^n Z_{ij}^2)^{1/2}|}{(\frac{1}{n} \sum_{i=1}^n \mathbb{E}[Z_{ij}^2])^{1/2}} + \frac{|(\frac{1}{n} \sum_{i=1}^n Z_{ij}^2)^{1/2} - (\frac{1}{n} \sum_{i=1}^n \mathbb{E}[Z_{ij}^2])^{1/2}|}{(\frac{1}{n} \sum_{i=1}^n \mathbb{E}[Z_{ij}^2])^{1/2}} \\ &\leq \left(\frac{1}{n} \sum_{i=1}^n (\hat{Z}_{ij} - Z_{ij})^2 \right)^{1/2} + \left| \frac{1}{n} \sum_{i=1}^n Z_{ij}^2 - \frac{1}{n} \sum_{i=1}^n \mathbb{E}[Z_{ij}^2] \right|. \end{aligned} \quad (\text{B.28})$$

The first term in (B.28) is bounded from above by $\delta_n/\log(pn)$ with probability at least $1 - \beta_n$ uniformly over all $j \in [p]$ by Condition A. Also, by Condition E and Lemma A.3,

$$\mathbb{E} \left[\max_{j \in [p]} \left| \frac{1}{n} \sum_{i=1}^n Z_{ij}^2 - \frac{1}{n} \sum_{i=1}^n \mathbb{E}[Z_{ij}^2] \right| \right] \leq K_1 \left(\sqrt{\frac{B_n^2 \log p}{n}} + \left(\mathbb{E} \left[\max_{i \in [n]} \max_{j \in [p]} Z_{ij}^4 \right] \right)^{1/2} \frac{\log p}{n} \right),$$

where K_1 is a universal constant. In addition, in the case of E.1, for some universal constant K_2 ,

$$\left(\mathbb{E} \left[\max_{i \in [n]} \max_{j \in [p]} Z_{ij}^4 \right] \right)^{1/2} \leq K_2 B_n^2 \log^2(pn),$$

and in the case of E.2,

$$\left(\mathbb{E} \left[\max_{i \in [n]} \max_{j \in [p]} Z_{ij}^4 \right] \right)^{1/2} \leq B_n^2 \sqrt{n}.$$

Thus, in both cases, since $\delta_n \leq 1$, by Condition E,

$$\mathbb{E} \left[\max_{j \in [p]} \left| \frac{1}{n} \sum_{i=1}^n Z_{ij}^2 - \frac{1}{n} \sum_{i=1}^n \mathbb{E}[Z_{ij}^2] \right| \right] \leq C \delta_n^2 / \log(pn),$$

for some universal constant C , so that by Markov's inequality, with probability at least $1 - \delta_n$,

$$\max_{j \in [p]} \left| \frac{1}{n} \sum_{i=1}^n Z_{ij}^2 - \frac{1}{n} \sum_{i=1}^n \mathbb{E}[Z_{ij}^2] \right| \leq C \delta_n / \log(pn).$$

Hence,

$$\mathbb{P} \left(\max_{j \in [p]} |\hat{w}_j - w_j|^2 \mathbb{E}_n[\hat{Z}_{ij}^2] > (1 + C)^2 \delta_n^2 / \log^2(pn) \right) \leq \beta_n + \delta_n.$$

Moreover, by Conditions E and A, as in Step 1 of the proof of Theorem 2.4, there exists a universal constant n_0 such that for all $n \geq n_0$,

$$\left(\frac{1}{n} \sum_{i=1}^n \hat{Z}_{ij}^2 \right)^{1/2} \geq \left(\frac{1}{n} \sum_{i=1}^n Z_{ij}^2 \right)^{1/2} - \left(\frac{1}{n} \sum_{i=1}^n (\hat{Z}_{ij} - Z_{ij})^2 \right)^{1/2} \geq \frac{1}{\sqrt{2}} - \frac{\delta_n}{\log(pn)} \geq \frac{1}{2}$$

for all $j \in [p]$ with probability at least $1 - \beta_n - (pn)^{-1}$. Thus, for all $n \geq n_0$,

$$\max_{j \in [p]} |\hat{w}_j - w_j|^2 (1 + \mathbb{E}_n[\hat{Z}_{ij}^2]) \leq 5 \max_{j \in [p]} |\hat{w}_j - w_j|^2 \mathbb{E}_n[\hat{Z}_{ij}^2] \leq 5(1 + C)^2 \delta_n^2 / \log^2(pn)$$

with probability at least $1 - \beta_n - \delta_n - (pn)^{-1}$, where we have β_n instead of $2\beta_n$ since the same event $\max_{j \in [p]} (\mathbb{E}_n[(\hat{Z}_{ij} - Z_{ij})^2])^{1/2} \leq \delta_n / \log(pn)$ is used twice. This gives the asserted claim. \blacksquare

Getting back to the proof of Theorem 2.8, take any $w \subset \mathcal{W}$ and $P \in \mathcal{P}^w$. Also, for $j \in [p]$, denote $\hat{w}_j := (n^{-1} \sum_{i=1}^n \hat{Z}_{ij}^2)^{-1/2}$, $w_j := V_{jj}^{-1/2}$, and

$$\bar{t}_j := \frac{\sqrt{n}(\hat{\theta}_j - \theta_{0j})}{\left(\frac{1}{n} \sum_{i=1}^n \hat{Z}_{ij}^2 \right)^{1/2}} = \sqrt{n} \hat{w}_j (\hat{\theta}_j - \theta_{0j}).$$

Then

$$\max_{j \in w} t_j \leq \max_{j \in w} \bar{t}_j,$$

and Condition W holds by Lemma B.1. Now, under conditions M, E, A, and W, we can proceed like in the proof of Theorems 2.5 and 2.6, with probabilities like $P_P(\max_{j \in w} \bar{t}_j \leq c_{1-\alpha, w})$ replacing probabilities like $P_P(\max_{1 \leq j \leq p} |\bar{t}_j| \leq c_{1-\alpha, w})$, to obtain

$$P_P \left(\max_{j \in w} \bar{t}_j \leq c_{1-\alpha, w} \right) \geq 1 - \alpha + o(1),$$

where the term $o(1)$ depends only on (δ_n, β_n) and on whether E.1 or E.2 is used. In particular, the term $o(1)$ does not depend on (w, P) . This gives (2.22) and completes the proof of the theorem. \blacksquare

Proof of Theorem 2.9. In this proof, all convergence results hold uniformly over $P \in \mathcal{P}$, and so we fix $P \in \mathcal{P}$, and drop the index P , i.e. we write, for example, P instead of P_P . Also, without loss of generality, we assume that $|\mathcal{H}| \leq \sqrt{\log p}$ (if $|\mathcal{H}| > \sqrt{\log p}$, we can redefine \mathcal{H} by keeping only the first $\sqrt{\log p}$ elements of it). We split the proof into six steps.

Step 1. Here we show that

$$P \left(\sum_{j=1}^p 1\{t_j > \sqrt{2 \log p}\} \geq |\mathcal{H}| \right) \rightarrow 1. \quad (\text{B.29})$$

To show (B.29), note that by Chebyshev's inequality, for all $j \in \mathcal{H}$ and $\epsilon > 0$,

$$P \left(\left| \frac{1}{n} \sum_{i=1}^n Z_{ij}^2 - V_{jj} \right| > \epsilon \right) \leq \frac{\sum_{i=1}^n E[Z_{ij}^4]}{\epsilon^2 n^2} \leq \frac{B^2}{\epsilon^2 n}.$$

Setting $\epsilon = V_{jj} \delta_n$ and recalling that by Condition M, $V_{jj} \geq 1$ for all $j \in \mathcal{H}$ gives

$$P \left(\left| \frac{\sum_{i=1}^n Z_{ij}^2}{n V_{jj}} - 1 \right| > \delta_n \right) \leq \frac{B^2}{V_{jj}^2 \delta_n^2 n} \leq \frac{B^2}{\delta_n^2 n},$$

and so by the union bound,

$$P \left(\max_{j \in \mathcal{H}} \left| \frac{\sum_{i=1}^n Z_{ij}^2}{n V_{jj}} - 1 \right| > \delta_n \right) \leq \frac{B^2 \sqrt{\log p}}{\delta_n^2 n} = o(1).$$

Thus, since $|x/y - 1| \leq |(x/y)^2 - 1|$ for all $x, y > 0$, it follows that

$$P \left(\max_{j \in \mathcal{H}} \left| \frac{(\sum_{i=1}^n Z_{ij}^2)^{1/2}}{(n V_{jj})^{1/2}} - 1 \right| > \delta_n \right) \leq P \left(\max_{j \in \mathcal{H}} \left| \frac{\sum_{i=1}^n Z_{ij}^2}{n V_{jj}} - 1 \right| > \delta_n \right) \leq \frac{B^2 \sqrt{\log p}}{\delta_n^2 n} = o(1).$$

Also, by Step 2 of the proof of Theorem 2.4,

$$P \left(\max_{j \in \mathcal{H}} \left| \frac{(\sum_{i=1}^n \hat{Z}_{ij}^2)^{1/2}}{(\sum_{i=1}^n Z_{ij}^2)^{1/2}} - 1 \right| > \delta_n \right) = o(1).$$

Therefore, recalling that $\hat{T}_j = n|\hat{\theta}_j|/(\sum_{i=1}^n \hat{Z}_{ij}^2)^{1/2}$,

$$\begin{aligned} & \mathbb{P}\left(\sum_{j=1}^p 1\{t_j > \sqrt{2\log p}\} \geq |\mathcal{H}|\right) \\ & \geq \mathbb{P}\left(\sum_{j \in \mathcal{H}} 1\{t_j > \sqrt{2\log p}\} \geq |\mathcal{H}|\right) = \mathbb{P}\left(t_j > \sqrt{2\log p} \text{ for all } j \in \mathcal{H}\right) \\ & \geq \mathbb{P}\left(\sqrt{n}(\hat{\theta}_j - \bar{\theta}_{0j})/V_{jj}^{1/2} > \sqrt{2\log p}(1 + \delta_n)^2 \text{ for all } j \in \mathcal{H}\right) - o(1). \end{aligned} \quad (\text{B.30})$$

Next, since $\delta_n \rightarrow 0$ and $\sqrt{n}(\theta_{0j} - \bar{\theta}_{0j})/V_{jj}^{1/2} > \sqrt{4\log p}$ for all $j \in \mathcal{H}$ by Condition L, it follows that the probability in (B.30) is bounded from below by

$$\begin{aligned} & \mathbb{P}\left(\sqrt{n}|\hat{\theta}_j - \theta_{0j}|/V_{jj}^{1/2} \leq (1/2)\sqrt{\log p} \text{ for all } j \in \mathcal{H}\right) \\ & \geq \mathbb{P}\left(\left|(nV_{jj})^{-1/2} \sum_{i=1}^n Z_{ij}\right| \leq (1/2)\sqrt{\log p} - \delta_n/\sqrt{\log p} \text{ for all } j \in \mathcal{H}\right) - o(1) \\ & \geq \mathbb{P}\left(\left|(nV_{jj})^{-1/2} \sum_{i=1}^n Z_{ij}\right| \leq (1/3)\sqrt{\log p} \text{ for all } j \in \mathcal{H}\right) - o(1) \\ & \geq 1 - \sqrt{\log p}/((1/3)\sqrt{\log p})^2 - o(1) = 1 - 9/\sqrt{\log p} = 1 - o(1), \end{aligned}$$

where the second line follows from Conditions A and M, the third from $\delta_n \rightarrow 0$, and the fourth from Chebyshev's inequality, the union bound, and the fact that $p \rightarrow \infty$. This gives (B.29) and completes the first step.

Step 2. Here we show that for any $j = 1, \dots, p$, the inequality $t_j \geq t_{(\hat{k})}$ holds if and only if $\hat{T}_j \geq \hat{t}$, where

$$\hat{t} := \inf \left\{ t \in \mathbb{R} : 1 - \Phi(t) \leq \frac{\alpha \max\{\sum_{j=1}^p 1\{t_j \geq t\}, 1\}}{p} \right\}. \quad (\text{B.31})$$

To prove this equivalence, we consider two cases separately: $\hat{k} = 0$ and $\hat{k} \geq 1$. First, suppose that $\hat{k} = 0$. Then for all $l \in [p]$,

$$1 - \Phi(t_{(l)}) > \frac{\alpha l}{p} = \frac{\alpha \max\{\sum_{j=1}^p 1\{t_j \geq t_{(l)}\}, 1\}}{p}.$$

Therefore, denoting $t_{(p+1)} := -\infty$, it follows that for all $t \leq t_{(1)}$, there exists $l \in [p]$ such that $t_{(l+1)} < t \leq t_{(l)}$ and

$$1 - \Phi(t) \geq 1 - \Phi(t_{(l)}) > \frac{\alpha \max\{\sum_{j=1}^p 1\{t_j \geq t_{(l)}\}, 1\}}{p} = \frac{\alpha \max\{\sum_{j=1}^p 1\{t_j \geq t\}, 1\}}{p}.$$

Hence, $\hat{t} > t_{(1)}$; so for all $j \in [p]$, we have $t_j < t_{(\hat{k})}$ and $t_j < \hat{t}$, so that equivalence holds.

Next, suppose that $\hat{k} \geq 1$. Then, by the same argument as above, $\hat{t} > t_{(\hat{k}+1)}$. Also,

$$1 - \Phi(t_{(\hat{k})}) \leq \frac{\alpha \hat{k}}{p} = \frac{\alpha \max\{\sum_{j=1}^p 1\{t_j \geq t_{(\hat{k})}\}, 1\}}{p}$$

by definition of \hat{k} , and so $\hat{t} \leq t_{(\hat{k})}$. Therefore, $t_{(\hat{k}+1)} < \hat{t} \leq t_{(\hat{k})}$. Thus, for all $j \in [p]$, the inequality $t_j \geq t_{(\hat{k})}$ holds if and only if $t_j \geq \hat{t}$, so that equivalence holds again. This completes the second step.

Step 3. Here we show that \hat{t} satisfies

$$\frac{pG(\hat{t})}{\max\{\sum_{j=1}^p 1\{t_j \geq \hat{t}\}, 1\}} = \alpha \quad (\text{B.32})$$

and

$$P(\hat{t} \leq G^{-1}(\alpha|\mathcal{H}|/p)) \rightarrow 1, \quad (\text{B.33})$$

where $G(t) := 1 - \Phi(t)$ for all $t \geq 0$. Indeed, (B.32) follows immediately from (B.31) since Φ is continuous. To prove (B.33), note that for all $x > 0$, $1 - \Phi(x) \leq e^{-x^2/2}$ by Proposition 2.5 in [60]. Therefore, as long as $|\mathcal{H}| \geq 1/\alpha$, which holds for all n large enough,

$$\alpha|\mathcal{H}|/p \geq 1/p = \exp\left(-\frac{(\sqrt{2\log p})^2}{2}\right) \geq \left(1 - \Phi(\sqrt{2\log p})\right) = G(\sqrt{2\log p}). \quad (\text{B.34})$$

Hence, since G is strictly decreasing,

$$G^{-1}(\alpha|\mathcal{H}|/p) \leq \sqrt{2\log p}. \quad (\text{B.35})$$

Therefore, setting $t = G^{-1}(\alpha|\mathcal{H}|/p)$, it follows from Step 1 that $\text{wp} \rightarrow 1$,

$$1 - \Phi(t) = G(t) = \frac{\alpha|\mathcal{H}|}{p} \leq \frac{\alpha \max\{\sum_{j=1}^p 1\{t_j \geq t\}, 1\}}{p},$$

and so by definition of \hat{t} , $\text{wp} \rightarrow 1$, $\hat{t} \leq t = G^{-1}(\alpha|\mathcal{H}|/p)$. This gives (B.33) and completes the third step.

Step 4. Here we show that for any sequence of constants $(\gamma_n)_{n \geq 1}$ such that $\gamma_n \rightarrow 0$,

$$\sup_{0 \leq t \leq (2\log p)^{1/2}} \left| 1 - \frac{G(t + \gamma_n/\sqrt{\log p})}{G(t)} \right| = o(1). \quad (\text{B.36})$$

To prove (B.36), note that it is immediate that

$$\sup_{0 \leq t \leq 1} \left| 1 - \frac{G(t + \gamma_n/\sqrt{\log p})}{G(t)} \right| = o(1).$$

Also, as long as $|\gamma_n|/\sqrt{\log p} \leq 1$, which holds for all n large enough, uniformly over $1 < t \leq (2\log p)^{1/2}$,

$$\begin{aligned} |G(t) - G(t + \gamma_n/\sqrt{\log p})| &= |\Phi(t + \gamma_n/\sqrt{\log p}) - \Phi(t)| = \frac{|\gamma_n|(\phi(t) \vee \phi(t + \gamma_n/\sqrt{\log p}))}{\sqrt{\log p}} \\ &= \frac{(1 + o(1))|\gamma_n|\phi(t)}{\sqrt{\log p}} \leq \frac{2(1 + o(1))t|\gamma_n|G(t)}{\sqrt{\log p}} = o(1)G(t), \end{aligned}$$

where the last inequality follows from the fact that for all $x \geq 1$, $\phi(x) \leq 2x(1 - \Phi(x)) = 2xG(x)$; see Proposition 2.5 in [60]. This gives (B.36).

Step 5. Here we show that

$$\sup_{0 \leq t \leq G^{-1}(\alpha|\mathcal{H}|/p)} \frac{\sum_{j \in \mathcal{H}_0} 1\{t_j \geq t\}}{p_0 G(t)} \leq 1 + o_p(1), \quad (\text{B.37})$$

where $p_0 := p - p_1$ is the number of true null hypotheses. To prove (B.37), for $j \in [p]$, let $\bar{t}_j := \sqrt{n}(\hat{\theta}_j - \theta_{0j})/(\mathbb{E}_n[\hat{Z}_{ij}^2])^{1/2}$, so that $\bar{t}_j \geq t_j$ for all $j \in \mathcal{H}_0$. We will show below that

$$\sup_{0 \leq t \leq G^{-1}(\alpha|\mathcal{H}|/p)} \left| \frac{\sum_{j \in \mathcal{H}_0} 1\{\bar{t}_j \geq t\}}{p_0 G(t)} - 1 \right| = o_p(1), \quad (\text{B.38})$$

which implies (B.37). To prove (B.38), for all $j \in [p]$, denote

$$T_j := \frac{\sum_{i=1}^n Z_{ij}}{(\sum_{i=1}^n Z_{ij}^2)^{1/2}} \text{ and } \bar{T}_j := \frac{\frac{1}{\sqrt{n}} \sum_{i=1}^n Z_{ij}}{\left(\frac{1}{n} \sum_{i=1}^n Z_{ij}^2 - \left(\frac{1}{n} \sum_{i=1}^n Z_{ij}\right)^2\right)^{1/2}},$$

so that

$$\bar{T}_j = \frac{T_j}{\sqrt{1 - T_j^2/n}}.$$

Equation (13) on page 2016 of [89] shows that for any sequence of positive constants $(d_n)_{n \geq 1}$ such that $d_n \rightarrow \infty$ and $d_n = o(p)$ as $n \rightarrow \infty$,

$$\sup_{0 \leq t \leq G_\kappa^{-1}(d_n/p)} \left| \frac{\sum_{j \in \mathcal{H}_0} 1\{\bar{T}_j \geq t\}}{p_0 G_\kappa(t)} - 1 \right| = o_p(1), \quad (\text{B.39})$$

where G_κ is some function such that $G_\kappa(t) \geq G(t)$ for all $t \in \mathbb{R}$ and, given that $\log^3 p/n = o(1)$, $G_\kappa(t) = G(t)(1 + o(1))$ uniformly over $0 \leq t \leq \sqrt{2 \log p}$ (in fact, [89] studied the case with absolute values, $|\bar{T}_j|$, but their argument works for \bar{T}_j instead of $|\bar{T}_j|$ as well). Therefore, since we have $G^{-1}(\alpha|\mathcal{H}|/(2p)) \leq \sqrt{2 \log p}$ for all n large enough, which can be established by the same arguments as those leading to (B.35),

$$\sup_{0 \leq t \leq G^{-1}(\alpha|\mathcal{H}|/(2p))} \left| \frac{\sum_{j \in \mathcal{H}_0} 1\{\bar{T}_j \geq t\}}{p_0 G(t)} - 1 \right| = o_p(1). \quad (\text{B.40})$$

Next, using the inequality $1 - \Phi(x) \leq \exp(-x^2/2)$ with $x = (4 \log p)^{1/2}$ gives $\Phi^{-1}(1 - 1/p^2) \leq (4 \log p)^{1/2}$. Therefore, using Theorem 2.4, (2.11), with $\alpha = 1/p = o(1)$ shows that

$$\max_{j \in [p]} |\bar{t}_j| = O_p(\sqrt{\log p}) \text{ and } \max_{1 \leq j \leq p} |T_j| = O_p(\sqrt{\log p}).$$

Thus,

$$\begin{aligned} |\bar{t}_j - T_j| &= \left| \frac{n(\hat{\theta}_j - \theta_{0j})}{(\sum_{i=1}^n \hat{Z}_{ij}^2)^{1/2}} - \frac{\sum_{i=1}^n Z_{ij}}{(\sum_{i=1}^n Z_{ij}^2)^{1/2}} \right| \\ &\leq |\bar{t}_j| \times \left| 1 - \frac{(\sum_{i=1}^n \hat{Z}_{ij}^2)^{1/2}}{(\sum_{i=1}^n Z_{ij}^2)^{1/2}} \right| + \frac{|\sqrt{nr}r_{nj}|}{(\sum_{i=1}^n Z_{ij}^2)^{1/2}} = o_p(1/\sqrt{\log p}) \end{aligned}$$

uniformly over $j \in [p]$ by Steps 1 and 2 in the proof of Theorem 2.4 and Condition A. Also,

$$|\bar{T}_j - T_j| = |T_j| \times \left| \frac{1}{\sqrt{1 - T_j^2/n}} - 1 \right| \leq |T_j| \times \left| \frac{1}{1 - T_j^2/n} - 1 \right| = \frac{|T_j|^3/n}{1 - T_j^2/n} = o_p(1/\sqrt{\log p})$$

uniformly over $j \in [p]$ since $\log^4 p/n = o(1)$. Hence, there exists a sequence of positive constants $(\gamma_n)_{n \geq 1}$ such that $\text{wp} \rightarrow 1$,

$$\max_{j \in [p]} |\bar{t}_j - \bar{T}_j| \leq \gamma_n / \sqrt{\log p}.$$

Further, by Step 4 and (B.35), for all $0 \leq t \leq G^{-1}(\alpha|\mathcal{H}|/p)$ and all n large enough,

$$G(t + \gamma_n / \sqrt{\log p}) = G(t)(1 + o(1)) \geq \alpha|\mathcal{H}|/(2p),$$

and so

$$t + \gamma_n / \sqrt{\log p} \leq G^{-1}(\alpha|\mathcal{H}|/(2p)).$$

Hence, uniformly over $0 \leq t \leq G^{-1}(\alpha|\mathcal{H}|/p)$, $\text{wp} \rightarrow 1$,

$$\frac{\sum_{j \in \mathcal{H}_0} 1\{\bar{t}_j \geq t\}}{p_0 G(t)} \geq \frac{\sum_{j \in \mathcal{H}_0} 1\{\bar{T}_j \geq t + \gamma_n / \sqrt{\log p}\}}{p_0 G(t)} = \frac{\sum_{j \in \mathcal{H}_0} 1\{\bar{T}_j \geq t + \gamma_n / \sqrt{\log p}\}}{p_0 G(t + \gamma_n / \sqrt{\log p})(1 + o(1))}$$

and

$$\frac{\sum_{j \in \mathcal{H}_0} 1\{\bar{t}_j \geq t\}}{p_0 G(t)} \leq \frac{\sum_{j \in \mathcal{H}_0} 1\{\bar{T}_j \geq t - \gamma_n / \sqrt{\log p}\}}{p_0 G(t)} = \frac{\sum_{j \in \mathcal{H}_0} 1\{\bar{T}_j \geq t - \gamma_n / \sqrt{\log p}\}}{p_0 G(t - \gamma_n / \sqrt{\log p})(1 + o(1))}.$$

Therefore, $\text{wp} \rightarrow 1$,

$$\sup_{0 \leq t \leq G^{-1}(\alpha|\mathcal{H}|/p)} \left| \frac{\sum_{j \in \mathcal{H}_0} 1\{\bar{t}_j \geq t\}}{p_0 G(t)} - 1 \right| \leq \sup_{0 \leq t \leq G^{-1}(\alpha|\mathcal{H}|/(2p))} 2 \times \left| \frac{\sum_{j \in \mathcal{H}_0} 1\{\bar{T}_j \geq t\}}{p_0 G(t)} - 1 \right|.$$

Combining this bound with (B.40) gives (B.38) and so (B.37).

Step 6. Here we complete the proof. By (B.33) in Step 3 and Step 5,

$$\frac{\sum_{j \in \mathcal{H}_0} 1\{t_j \geq \hat{t}\}}{p_0 G(\hat{t})} \leq 1 + o_p(1),$$

and so

$$\sum_{j \in \mathcal{H}_0} 1\{t_j \geq \hat{t}\} \leq p_0 G(\hat{t})(1 + o_p(1)). \quad (\text{B.41})$$

Therefore,

$$FDP = \frac{\sum_{j \in \mathcal{H}_0} 1\{t_j \geq \hat{t}\}}{\max\{\sum_{j=1}^p 1\{t_j \geq \hat{t}\}, 1\}} \leq \frac{p_0 G(\hat{t})(1 + o_p(1))}{p G(\hat{t})/\alpha} = \frac{\alpha p_0}{p} + o_p(1),$$

where the first equality follows from Step 2 and the second from (B.41) above and (B.32) in Step 3. Finally, since FDP is bounded between 0 and 1, it follows that

$$FDR = E[FDP] \leq \alpha p_0/p + o(1) \leq \alpha + o(1).$$

The asserted claim follows. \blacksquare

APPENDIX C. PROOFS FOR SECTION 3

Proof of Lemma 3.1. The proof follows immediately from Lemma D.2 in Appendix D. \blacksquare

Proof of Lemma 3.5. By the triangle inequality,

$$\sup_{\theta \in \mathcal{R}(\theta_0)} \|\hat{g}(\theta) - g(\theta)\|_\infty \leq \sup_{\theta \in \mathcal{R}(\theta_0)} \|\hat{g}(\theta) - g(\theta) - \hat{g}(\theta_0)\|_\infty + \|\hat{g}(\theta_0)\|_\infty,$$

where the second term on the right-hand side satisfies

$$\|\hat{g}(\theta_0)\|_\infty = \|n^{-1/2} \mathbb{G}_n(g(X, \theta_0))\|_\infty \leq n^{-1/2} \ell_n \quad (\text{C.1})$$

with probability $1 - \delta_n/6$ by Condition ENM. Thus, it remains to bound

$$\sup_{\theta \in \mathcal{R}(\theta_0)} \|\hat{g}(\theta) - g(\theta) - \hat{g}(\theta_0)\|_\infty.$$

To do so, we apply Lemma D.3 in Appendix D. Denote

$$h_j(X, t) = \tilde{g}_j(X, Z_{u(j)}(X)' \vartheta_{0u(j)} + t) - \tilde{g}_j(X, Z_{u(j)}(X)' \vartheta_{0u(j)}), \quad j \in [m],$$

$$\eta = \theta - \theta_0 = ((\vartheta_1 - \vartheta_{01})', \dots, (\vartheta_{\bar{u}} - \vartheta_{0\bar{u}})')',$$

and $\Delta = \{\theta - \theta_0 : \theta \in \mathcal{R}(\theta_0)\}$. Then $\|\Delta\|_1 = \sup_{\eta \in \Delta} \|\eta\|_1 \leq 2K$ by Condition DM and

$$\sup_{\theta \in \mathcal{R}(\theta_0)} \|\hat{g}(\theta) - g(\theta) - \hat{g}(\theta_0)\|_\infty = n^{-1/2} \max_{j \in [m]} \sup_{\eta \in \Delta} |\mathbb{G}_n(h_j(X, Z_{u(j)}(X)'(\vartheta_{u(j)} - \vartheta_{0u(j)})))|.$$

Also, by Condition ENM, the functions $h_j(X, t)$ have the contractive structure as required in Lemma D.3 and

$$\sup_{\eta \in \Delta, j \in [m]} \text{Var}(h_j(X, Z_{u(j)}(X)'(\vartheta_{u(j)} - \vartheta_{0u(j)}))) \leq B_{1n}^2.$$

Finally, with probability at least $1 - \delta_n/6$,

$$\max_{j \in [m], k \in [p_z]} \mathbb{E}_n[L_j^2(X) Z_{u(j)k}^2(X)] \leq B_{2n}^2,$$

again by Condition ENM. Therefore, applying Lemma D.3 shows that

$$\mathbb{P} \left(\max_{j \in [m]} \sup_{\eta \in \Delta} |\mathbb{G}_n(h_j(X, Z_{u(j)}(X))'(\vartheta_{u(j)} - \vartheta_{0u(j)}))| > t \right) \leq 5\delta_n/6$$

for all $t \geq \{4B_{1n}\} \vee \{16\sqrt{2}B_{2n}K \log^{1/2}(48pm/\delta_n)\} = \tilde{\ell}_n$. Combining this bound with (C.1) gives the asserted claim. \blacksquare

Proof of Lemma 3.7. We will verify conditions L and ELM to invoke Theorem 3.1 since Conditions LID(γ_{0j}, Ω) and DM ($\|\gamma_{0j}\|_1 \leq K$) are assumed. By definition of the estimator $\hat{\gamma}_j$ in (3.32), it follows that $\hat{\gamma}_j$ is an RGMM estimator with a linear score function since

$$g(\gamma_j) = \gamma_j \Omega + (G')_j \quad \text{and} \quad \hat{g}(\gamma_j) = \gamma_j \hat{\Omega} + (\hat{G}')_j.$$

In this case we have that ELM holds by considering $\hat{G} := \hat{\Omega}$, $\hat{g}(0) := (\hat{G}')_j$, and $\ell_n = \ell_n^\Omega \vee \ell_n^G$ by assumption.

Next we verify that the choice of penalty satisfies Condition L for some $\alpha = \delta_n$. Note that with probability $1 - \delta_n$

$$\begin{aligned} \|\gamma_{0j} \hat{\Omega} - (\hat{G}')_j\|_\infty &\leq \|\gamma_{0j} \hat{\Omega} - (G')_j\|_\infty + n^{-1/2} \ell_n^G \\ &\leq \|\gamma_{0j} \Omega - (G')_j\|_\infty + \|\gamma_{0j}\|_1 \|\hat{\Omega} - \Omega\|_\infty + n^{-1/2} \ell_n^G \\ &\leq 0 + Kn^{-1/2} \ell_n^\Omega + n^{-1/2} \ell_n^G \end{aligned}$$

so that γ_{0j} is feasible for all $j \in [p]$ if $\lambda_j^\gamma \geq Kn^{-1/2} \ell_n^\Omega + n^{-1/2} \ell_n^G$. Thus the result follows from Theorem 3.1.

For the second result, we will verify conditions L, DM, LID($\mu_{0j}, G' \Omega^{-1} G$) and ELM to invoke Theorem 3.1. Conditions DM and LID($\mu_{0j}, G' \Omega^{-1} G$) are assumed. By the definition of the estimator $\hat{\mu}_j$, given in (3.33), we have that $\hat{\mu}_j$ is a RGMM estimator associated with a linear score as we can write

$$g(\mu_j) = \mu_j \gamma_0 G - e'_j \quad \text{and} \quad \hat{g}(\mu_j) = \mu_j \hat{\gamma} \hat{G} - e'_j.$$

To verify Condition ELM, note that $g(0) = \hat{g}(0) = -e_j$ and by assumption $\|\mu_{0j}\|_1 \leq K$ so we need to bound $\|\hat{\gamma} \hat{G} - \gamma_0 G\|_\infty$. Using that $\gamma_0 G = G' \gamma'_0$ since $\gamma_0 = G' \Omega^{-1}$, we have with probability $1 - 2\delta_n$ that

$$\begin{aligned} \|\hat{\gamma} \hat{G} - \gamma_0 G\|_\infty &\leq \|\hat{\gamma}(\hat{G} - G)\|_\infty + \|\hat{\gamma} G - G' \gamma'_0\|_\infty \\ &\leq \max_{j \in [p]} \|\hat{\gamma}_j\|_1 \|\hat{G} - G\|_\infty + \|\hat{\gamma} \Omega \Omega^{-1} G - G' \Omega^{-1} G\|_\infty \\ &\leq \max_{j \in [p]} \|\gamma_{0j}\|_1 \|\hat{G} - G\|_\infty + \|\hat{\gamma}(\Omega - \hat{\Omega}) \gamma'_0\|_\infty + \|\hat{\gamma} \hat{\Omega} \gamma'_0 - G' \gamma'_0\|_\infty \\ &\leq K \|\hat{G} - G\|_\infty + \|\hat{\gamma}(\Omega - \hat{\Omega})\|_\infty \max_{j \in [p]} \|\gamma_{0j}\|_1 + \|\hat{\gamma} \hat{\Omega} \gamma'_0 - G' \gamma'_0\|_\infty \\ &\leq K \|\hat{G} - G\|_\infty + K^2 \|\Omega - \hat{\Omega}\|_\infty + \|(\hat{G}' - G') \gamma'_0\|_\infty + \|\hat{\gamma} \hat{\Omega} \gamma'_0 - \hat{G}' \gamma'_0\|_\infty \\ &\leq K \|\hat{G} - G\|_\infty + K^2 \|\Omega - \hat{\Omega}\|_\infty + K \|\hat{G}' - G'\|_\infty + K \|\hat{\gamma} \hat{\Omega} - \hat{G}'\|_\infty \\ &\leq Kn^{-1/2} \ell_n^G + K^2 n^{-1/2} \ell_n^\Omega + Kn^{-1/2} \ell_n^G + K \max_{j \in [p]} \lambda_j^\gamma =: r^* \end{aligned}$$

where we used that $\|\hat{\gamma}_j\|_1 \leq \|\gamma_{0j}\|_1$ for all $j \in [p]$ with probability $1 - \delta_n$ and that $\{\|\hat{G} - G\|_\infty \leq n^{-1/2}\ell_n^G$ and $\|\hat{\Omega} - \Omega\|_\infty \leq n^{-1/2}\ell_n^\Omega\}$ occurs with probability $1 - \delta_n$. Thus Condition ELM holds with $\ell_n = n^{1/2}r^*$.

Next we verify Condition L. We have that with probability at least $1 - 2\delta_n$

$$\|\mu_0 \hat{\gamma} \hat{G} - I\|_\infty \leq \|\mu_0 \gamma_0 G - I\|_\infty + \|\mu_0 (\hat{\gamma} \hat{G} - \gamma_0 G)\|_\infty \leq 0 + \max_{j \in [p]} \|\mu_{0j}\|_1 \|\hat{\gamma} \hat{G} - \gamma_0 G\|_\infty \leq Kr^*.$$

Therefore, if $\lambda_j^\mu \geq Kr^*$ we have that $(\mu_{0j})_{j \in [p]}$ is feasible for the optimization problem (3.33) with probability at least $1 - 2\delta_n$ (i.e., Condition L holds with $\alpha = 2\delta_n$). Then the result follows by Theorem 3.1. \blacksquare

Proof of Lemma 3.8. Let $M_n = \mathbb{E}[\max_{i \in [n]} \|G(X_i, \theta_0)\|_\infty^2] \vee \mathbb{E}[\max_{i \in [n]} \|g(X_i, \theta_0)\|_\infty^4]$.

To bound $\|\hat{G} - G\|_\infty$ we apply Lemma C.1(4) with $t = \delta_n^{-1}$, $\bar{q} = 2$. Indeed, under the assumed condition $n^{-1/2}M_n\{\delta_n^{-1} + \log(2m)\} \leq c$, with probability $1 - 2\delta_n$ we have

$$\max_{j \in [m], k \in [p]} |\mathbb{G}_n G_{jk}(X)| \leq C\sqrt{\sigma^2 \log(2m)} \quad \text{and} \quad \max_{k \in [m]} |\mathbb{G}_n g_k(X, 0)| \leq C\sqrt{\sigma^2 \log(2m)}.$$

Next we establish the bound on $\|\hat{\Omega} - \Omega\|_\infty$. Using the triangle inequality we have

$$\begin{aligned} n^{1/2}\|\hat{\Omega} - \Omega\|_\infty &= \max_{j \in [m], k \in [m]} n^{1/2} |\mathbb{E}_n[g_j(X, \hat{\theta})g_k(X, \hat{\theta})] - \mathbb{E}[g_j(X, \theta_0)g_k(X, \theta_0)]| \\ &\leq \max_{j \in [m], k \in [m]} n^{1/2} |\mathbb{E}_n[\{g_j(X, \hat{\theta}) - g_j(X, \theta_0)\}\{g_k(X, \hat{\theta}) - g_k(X, \theta_0)\}]| \\ &\quad + 2 \max_{j \in [m], k \in [m]} n^{1/2} |\mathbb{E}_n[g_j(X, \theta_0)\{g_k(X, \hat{\theta}) - g_k(X, \theta_0)\}]| \\ &\quad + \max_{j \in [m], k \in [m]} |\mathbb{G}_n(g_j(X, \theta_0)g_k(X, \theta_0))| \end{aligned} \tag{C.2}$$

To bound the last term of the right-hand-side (RHS) in (C.2) we apply Lemma C.1(4) with $q = 2$ and $t = \delta_n$, so that with probability $1 - \delta_n$

$$\begin{aligned} \max_{k, j \in [m]} |\mathbb{G}_n(g_k(X, \theta_0)g_j(X, \theta_0))| &\leq C \max_{k \in [m]} \mathbb{E}[g_k^4(X, \theta_0)]^{1/2} \sqrt{\log(2m)} \\ &\quad + Cn^{-1/2} \mathbb{E}[\max_{i \in [n]} \|g(X, \theta_0)\|_\infty^4] \{\delta_n^{-1} + \log(m)\} \\ &\leq C' \sqrt{\log(2m)} \end{aligned}$$

under the growth condition (iv).

To bound the first term of the RHS in (C.2) we note that by linearity of the score, $g(X, \hat{\theta}) - g(X, \theta_0) = G(X)'(\hat{\theta} - \theta_0)$, and applying the Cauchy-Schwarz inequality we have

$$\begin{aligned} n^{1/2} \mathbb{E}_n[\{g_j(X, \hat{\theta}) - g_j(X, \theta_0)\}\{g_k(X, \hat{\theta}) - g_k(X, \theta_0)\}] &\leq n^{1/2} \max_{k \in [m]} \mathbb{E}_n[\{G_k(X)'(\hat{\theta} - \theta_0)\}^2] \\ &\leq n^{1/2} \Delta_{2n}^2 \end{aligned}$$

where the last inequality holds with probability $1 - \delta_n$ by condition (iii).

To bound the remaining term in the RHS in (C.2), again using linearity, the Cauchy-Schwarz inequality, and condition (iii), with probability $1 - \delta_n$

$$\begin{aligned} |\mathbb{E}_n[g_j(X, \theta_0)\{g_k(X, \hat{\theta}) - g_k(X, \theta_0)\}]| &= |\mathbb{E}_n[g_j(X, \theta_0)G_k(X)](\hat{\theta} - \theta_0)| \\ &\leq \max_{j \in [m]} \{\mathbb{E}_n[g_j^2(X, \theta_0)]\}^{1/2} \max_{k \in [m]} \mathbb{E}_n[\{G_k(X)(\hat{\theta} - \theta_0)\}^2]^{1/2} \\ &\leq \max_{j \in [m]} \{\mathbb{E}_n[g_j^2(X, \theta_0)]\}^{1/2} \Delta_{2n} \end{aligned}$$

To bound $\max_{j \in [m]} \{\mathbb{E}_n[g_j^2(X, \theta_0)]\}^{1/2}$, we will apply Lemma C.1(3). We have

$$\begin{aligned} \mathbb{E}[\max_{j \in [m]} |\mathbb{E}_n[g_j^2(X, \theta_0)] - \mathbb{E}[g_j^2(X, \theta_0)]|] &\leq n^{-1} \mathbb{E}[\max_{i \in [n]} \|g(X, \theta_0)\|_\infty^2] \log(2m) \\ &\quad + n^{-1/2} \mathbb{E}[\max_{i \in [n]} \|g(X, \theta_0)\|_\infty^2]^{1/2} \log^{1/2}(2m) \\ &\leq C\delta_n \end{aligned}$$

under $n^{-1/2} \mathbb{E}[\max_{i \in [n]} \|g(X, \theta_0)\|_\infty^2]^{1/2} \leq \delta_n \log^{-1/2}(2m)$ assumed in condition (iv). Thus, by Markov's inequality we have that with probability $1 - \delta_n$

$$\max_{j \in [m]} \mathbb{E}_n[g_j^2(X, \theta_0)] \leq \max_{j \in [m]} \mathbb{E}[g_j^2(X, \theta_0)] + C.$$

■

Proof of Lemma 3.9. Recall that $\partial_j g_k(X, \theta) = G(X, \theta)$.

We first bound $\|\hat{G} - G\|_\infty$. We have that

$$\begin{aligned} \max_{j \in [p], k \in [m]} n^{1/2} |\hat{G}_{kj} - G_{kj}| &= \max_{j \in [p], k \in [m]} n^{1/2} |\mathbb{E}_n[\partial_j g_k(X, \hat{\theta})] - \mathbb{E}[\partial_j g_k(X, \theta_0)]| \\ &\leq \max_{j \in [p], k \in [m]} |\mathbb{G}_n(\partial_j g_k(X, \hat{\theta}) - \partial_j g_k(X, \theta_0))| \\ &\quad + \max_{j \in [p], k \in [m]} |\mathbb{G}_n(\partial_j g_k(X, \theta_0))| \\ &\quad + \max_{j \in [p], k \in [m]} n^{1/2} |\mathbb{E}[\partial_j g_k(X, \hat{\theta}) - \partial_j g_k(X, \theta_0)]| \end{aligned} \tag{C.3}$$

To control the second term in the RHS of (C.3) we apply Lemma C.1(4) with $q = 2$ and $t = \delta_n$, so that with probability $1 - \delta_n$

$$\max_{j \in [p], k \in [m]} |\mathbb{G}_n(\partial_j g_k(X, \theta_0))| = \max_{j \in [p], k \in [m]} |\mathbb{G}_n(G_{kj}(X, \theta_0))| \leq C\sqrt{\log(2m)}$$

under condition (v).

To bound the last term in the RHS of (C.3), we have

$$\begin{aligned}
\max_{j \in [p], k \in [m]} n^{1/2} |\mathbb{E}[\partial_j g_k(X, \hat{\theta}) - \partial_j g_k(X, \theta_0)]| \\
&\leq_{(1)} \max_{j \in [p], k \in [m]} n^{1/2} \mathbb{E}[|\tilde{L}_{kj}(X) \tilde{Z}_{jk}(X)'(\hat{\theta} - \theta_0)|] \\
&\leq_{(2)} \max_{j \in [p], k \in [m]} n^{1/2} \mathbb{E}[|\tilde{L}_{kj}(X) \tilde{Z}_{jk}(X)'(\hat{\theta} - \theta_0)|^2]^{1/2} \\
&\leq_{(3)} C n^{1/2} \|\hat{\theta} - \theta_0\|_2 \\
&\leq_{(4)} C n^{1/2} \Delta_{2n}
\end{aligned}$$

where (1) holds by the ENM condition $|G_{kj}(X, \tilde{\theta}) - G_{kj}(X, \theta)| \leq L_{kj}(X) |\tilde{Z}'_{kj}(\tilde{\theta} - \theta)|$, (2) holds by monotonicity of L_q norms, (3) holds by condition (i), and (4) holds with probability $1 - \delta_n$ under condition (iii).

Finally, to control the first term in the RHS of (C.3), we have by Lemma D.3 in Appendix D with $B_{1n} = B_n \Delta_{1n}$ and $B_{2n} = B_n$ so that

$$\max_{j \in [p], k \in [m]} |\mathbb{G}_n(\partial_j g_k(X, \hat{\theta}) - \partial_j g_k(X, \theta_0))| \leq C B_n \Delta_{1n} \log^{1/2}(m/\delta_n)$$

with probability $1 - 7\delta_n$ where we used that $\|\hat{\theta} - \theta_0\|_1 \leq \Delta_{1n}$ with probability at least $1 - \delta_n$ by condition (iii).

To bound $\|\hat{G} - \tilde{G}\|_\infty$ we note that

$$\|\hat{G} - \tilde{G}\|_\infty \leq \|\hat{G} - G\|_\infty + \|\tilde{G} - G\|_\infty$$

and that Lemma D.3 in Appendix D allows for $\bar{m} = m$ and different $\tilde{\theta}_j^*$, $j \in [m]$. Since for every $j \in [m]$ we have $\|\tilde{\theta}_j^* - \theta_0\|_1 \leq \|\hat{\theta} - \theta_0\|_1$ and $\|\tilde{\theta}_j^* - \theta_0\|_2 \leq \|\hat{\theta} - \theta_0\|_2$, the bound we derived on $\|\hat{G} - G\|_\infty$ applies to $\|\tilde{G} - G\|_\infty$ as well.

Next we establish the bound on $\|\hat{\Omega} - \Omega\|_\infty$. By the triangle inequality, we have

$$\begin{aligned}
n^{1/2} \|\hat{\Omega} - \Omega\|_\infty &= \max_{j \in [m], k \in [m]} n^{1/2} |\mathbb{E}_n[g_j(X, \hat{\theta})g_k(X, \hat{\theta})] - \mathbb{E}[g_j(X, \theta_0)g_k(X, \theta_0)]| \\
&\leq \max_{j \in [m], k \in [m]} |\mathbb{G}_n(g_j(X, \hat{\theta})g_k(X, \hat{\theta}) - g_j(X, \theta_0)g_k(X, \theta_0))| \\
&\quad + \max_{j \in [m], k \in [m]} |\mathbb{G}_n(g_j(X, \theta_0)g_k(X, \theta_0))| \\
&\quad + \max_{j \in [m], k \in [m]} n^{1/2} |\mathbb{E}[g_j(X, \hat{\theta})g_k(X, \hat{\theta}) - g_j(X, \theta_0)g_k(X, \theta_0)]|
\end{aligned} \tag{C.4}$$

To control the first term of the RHS in (C.4), further apply the triangle inequality to obtain

$$\begin{aligned}
&\max_{j \in [m], k \in [m]} |\mathbb{G}_n(g_j(X, \hat{\theta})g_k(X, \hat{\theta}) - g_j(X, \theta_0)g_k(X, \theta_0))| \\
&= \max_{j \in [m], k \in [m]} |\mathbb{G}_n(\{g_j(X, \hat{\theta}) - g_j(X, \theta_0)\}g_k(X, \hat{\theta}) - g_j(X, \theta_0)\{g_k(X, \theta_0) - g_k(X, \hat{\theta})\})| \\
&\leq \max_{j \in [m], k \in [m]} |\mathbb{G}_n(\{g_j(X, \hat{\theta}) - g_j(X, \theta_0)\}\{g_k(X, \hat{\theta}) - g_k(X, \theta_0)\})| \\
&\quad + 2 \max_{j \in [m], k \in [m]} |\mathbb{G}_n(\{g_j(X, \hat{\theta}) - g_j(X, \theta_0)\}g_k(X, \theta_0))| \\
&= (I) + 2(II)
\end{aligned}$$

Therefore, we obtain with probability $1 - C\delta_n$

$$\begin{aligned}
(I) &\leq_{(1)} \max_{j \in [m], k \in [m]} n^{1/2} \mathbb{E}_n[\{g_j(X, \hat{\theta}) - g_j(X, \theta_0)\}^2]^{1/2} \mathbb{E}_n[\{g_k(X, \hat{\theta}) - g_k(X, \theta_0)\}^2]^{1/2} \\
&\quad + \max_{j \in [m], k \in [m]} n^{1/2} \mathbb{E}[\{g_j(X, \hat{\theta}) - g_j(X, \theta_0)\}^2]^{1/2} \mathbb{E}[\{g_k(X, \hat{\theta}) - g_k(X, \theta_0)\}^2]^{1/2} \\
&\leq_{(2)} \max_{j \in [m], k \in [m]} n^{1/2} \mathbb{E}_n[\{L_j(X)Z_j'(\hat{\theta} - \theta_0)\}^2]^{1/2} \mathbb{E}_n[\{L_k(X)Z_k'(\hat{\theta} - \theta_0)\}^2]^{1/2} \\
&\quad + \max_{j \in [m], k \in [m]} n^{1/2} \mathbb{E}[\{L_j(X)Z_j'(\hat{\theta} - \theta_0)\}^2]^{1/2} \mathbb{E}[\{L_k(X)Z_k'(\hat{\theta} - \theta_0)\}^2]^{1/2} \\
&\leq_{(3)} \max_{j \in [m]} n^{1/2} \{\|\mathbb{E}_n[L_j^2(X)Z_jZ_j']\|_\infty + \|\mathbb{E}[L_j^2(X)Z_jZ_j']\|_\infty\} \|\hat{\theta} - \theta_0\|_1^2 \\
&\leq_{(4)} 2n^{1/2} B_n^2 \Delta_{1n}^2
\end{aligned}$$

where (1) follows from triangle inequality and Cauchy-Schwarz inequality, (2) follows from the ENM condition $|g_j(X, \hat{\theta}) - g_j(X, \theta_0)| \leq L_j(X)|Z_j'(\hat{\theta} - \theta_0)|$, (3) follows from $|v'Mv| \leq \|v\|_1 \|M\|_\infty$, and (4) follows from $\|\hat{\theta} - \theta_0\|_1 \leq \Delta_{1n}$ with probability $1 - \delta_n$ by condition (iii), $\|\mathbb{E}[L_j^2(X)Z_jZ_j']\|_\infty \leq \max_{k \in [p]} \mathbb{E}[L_j^2(X)Z_{jk}^2] \leq B_n^2$ by condition (i) and $\|\mathbb{E}_n[L_j^2(X)Z_jZ_j']\|_\infty \leq \max_{k \in [p]} \mathbb{E}_n[L_j^2(X)Z_{jk}^2] \leq B_n^2$ by condition (ii) with probability $1 - \delta_n$.¹

Next we will apply the contraction principle in Lemma D.3 in Appendix D to control (II) with $\tilde{L}_{jk}(X) := |L_j(X)g_k(X, \theta_0)|$. By condition (ii) and because for any θ such that $\|\theta - \theta_0\|_\ell \leq \Delta_{\ell n}$, $\ell \in \{\ell_1, \ell_2\}$,

$$\begin{aligned}
\text{Var}(\mathbb{G}_n(\{g_j(X, \theta) - g_j(X, \theta_0)\}g_k(X, \theta_0))) &\leq \text{Var}(\mathbb{G}_n(\{L_j(X)Z_j'(\theta - \theta_0)\}g_k(X, \theta_0))) \\
&\leq C\|\theta - \theta_0\|^2 \wedge B_n^2\|\theta - \theta_0\|_1^2 \\
&\leq C\Delta_{2n}^2 \wedge B_n^2\Delta_{1n}^2
\end{aligned}$$

by condition (i), we have that the condition of Lemma D.3 hold. Therefore, by Lemma D.3 we have with probability $1 - 7\delta_n$

$$(II) \leq C' \{B_n \Delta_{1n} \vee B_n \Delta_{1n} \log^{1/2}(mp/\delta_n)\}$$

where we used that $\|\hat{\theta} - \theta_0\|_1 \leq \Delta_{1n}$ with probability $1 - \delta_n$ by condition (iii).

To control the last term of the RHS in (C.4), we use that

$$\begin{aligned}
&|\mathbb{E}[g_j(X, \hat{\theta})g_k(X, \hat{\theta}) - g_j(X, \theta_0)g_k(X, \theta_0)]| \\
&\leq |\mathbb{E}[g_j(X, \hat{\theta})\{g_k(X, \hat{\theta}) - g_k(X, \theta_0)\}]| + |\mathbb{E}[\{g_j(X, \hat{\theta}) - g_j(X, \theta_0)\}g_k(X, \theta_0)]| \\
&\leq B_n^2\|\hat{\theta} - \theta_0\|_1^2 + 2 \max_{j, k \in [m]} |\mathbb{E}[\{g_j(X, \hat{\theta}) - g_j(X, \theta_0)\}g_k(X, \theta_0)]| \\
&\leq B_n^2\Delta_{1n}^2 + 2C\Delta_{2n}.
\end{aligned}$$

To bound the second term of the RHS in (C.4) we use Lemma C.1(4) so that with probability $1 - \delta_n$

$$\begin{aligned}
\max_{k, j \in [m]} |\mathbb{G}_n(g_k(X, \theta_0)g_j(X, \theta_0))| &\leq C \max_{k \in [m]} \mathbb{E}[g_k^4(X, \theta_0)]^{1/2} \sqrt{\log(2m)} \\
&\quad + Cn^{-1/2} \mathbb{E}[\max_{i \in [n]} \|g(X_i, \theta_0)\|_\infty^4] \{\delta_n^{-1} + \log(m)\} \\
&\leq C' \sqrt{\log(2m)}
\end{aligned}$$

¹An alternative bound can be obtained by $\{\max_{\text{eig}}(\mathbb{E}_n[L_j^2(X)ZZ']) \vee \max_{\text{eig}}(\mathbb{E}[L_j^2(X)ZZ'])\}n^{1/2}\|\hat{\theta} - \theta_0\|_2^2$.

under the growth condition (iv). ■

Proof of Theorem 3.3. We will establish that $\bar{r}_1 + \bar{r}_2 + \bar{r}_3 = o_P(\log^{-1/2} p)$. Throughout the proof we assume $s \geq 1$.

Because we are considering a special case, a few simplifications occur. Because this is a linear case, we can take $\tilde{G} = \hat{G}$ so that $\bar{r}_2 = 0$. Moreover, due to the homoskedastic setting, we note that $\mu_0\gamma_0$ is independent of σ^2 so we do not need to estimate σ^2 in the construction of $\hat{\mu}$ and $\hat{\gamma}$.

By Lemma C.1(4) with probability $1 - \delta_n$ we have

$$\begin{aligned} \|\sqrt{n}\hat{g}(\theta_0)\|_\infty &= \max_{k \in [m]} |\mathbb{G}_n(g_k(X, \theta_0))| \\ &\leq C \max_{k \in [m]} \mathbb{E}[g_k^2(X, \theta_0)]^{1/2} \sqrt{\log(2m)} \\ &\quad + Cn^{-1/2} \mathbb{E}[\max_{i \in [n]} \|g(X_i, \theta_0)\|_\infty^2] \{\delta_n^{-1} + \log(2m)\} \\ &\leq C' \sqrt{\log(2m)} \end{aligned}$$

under the growth condition $n^{-1/2} \mathbb{E}[\max_{i \in [n]} \|g(X, \theta_0)\|_\infty^2] \{\delta_n^{-1} + \log(2m)\} \leq c \log^{1/2}(2m)$. Thus Condition L holds for Step 1 under the proposed choice of penalty parameter. Moreover, since the score is linear, we have that Condition ELM holds since $\hat{G} = -\mathbb{E}_n ZW'$ and $\hat{g}(0) = \mathbb{E}_n YW$; and we have by Lemma 3.8 that with probability $1 - C\delta_n$

$$\|\hat{G} - G\|_\infty \leq C\sqrt{\log(2m)}, \quad \text{and} \quad \|\hat{g}(0) - g(0)\|_\infty \leq C\sqrt{\log(2m)}$$

under Conditions (3) and (4).

Since Conditions DM ($\|\theta_0\|_1 \leq K$) and LID(θ_0, G) hold by assumption, and by Theorem 3.1 with $\ell_n = C' \sqrt{\log(2m)}$, we have for $q \in \{1, 2\}$ that

$$\|\hat{\theta} - \theta_0\|_q \leq \Delta_{qn} := Cn^{-1/2} s^{1/q} \ell_n \{(1+L)\mu_n^{-1} + C\} \leq C'n^{-1/2} s^{1/q} \log^{1/2}(2m) \quad (\text{C.5})$$

where we used that $\mu_n \geq c$ and $L \leq C$.

Moreover by Lemma 3.8 we have that with probability $1 - C\delta_n$

$$\|\hat{G} - G\|_\infty \leq C\sqrt{\log(2m)}, \quad \text{and} \quad \|\mathbb{E}_n ZZ' - \mathbb{E}ZZ'\|_\infty \leq C\sqrt{\log(2m)}$$

under where we used that we do not need to estimate σ .

In order to apply Lemma 3.7, recall the choices of penalty

$$\bar{\lambda} = \lambda_j^\gamma = \frac{1}{2} \lambda_j^\mu = n^{-1/2} C(1 \vee K^2)(1+K)\Phi^{-1}(1 - (mpn)^{-1}), \quad (\text{C.6})$$

so that we consider $\ell_n = n^{1/2} \bar{\lambda} \geq C(1 \vee K^2)(1+K)\sqrt{\log(2m)}$ for some constant $C > 0$ chosen sufficiently large so they satisfy the requirements on the penalty choices of Lemma 3.7. Therefore, Lemma 3.7 yields with probability $1 - C\delta_n$ that

$$\max_{j \in [p]} \|\hat{\gamma}_j - \sigma^2 \gamma_{0j}\|_1 \leq C\bar{\lambda} s \{1 + \mu_n^{-1}\} \quad \text{and} \quad \|\hat{\gamma}_j\|_1 \leq \|\sigma^2 \gamma_{0j}\|_1 \leq CK \quad \text{for all } j \in [p]$$

$$\max_{j \in [p]} \|\hat{\mu}_j - \sigma^{-2} \mu_{0j}\|_1 \leq C' \bar{\lambda} s \{1 + \mu_n^{-1}\} \quad \text{and} \quad \|\hat{\mu}_j\|_1 \leq \|\sigma^{-2} \mu_{0j}\|_1 \leq CK \quad \text{for all } j \in [p]$$

Now we are in position to bound \bar{r}_1, \bar{r}_2 and \bar{r}_3 . By Lemma 3.6 and using the definition of the estimators (3.33) we have

$$\begin{aligned} \bar{r}_1 &= \sqrt{n} \|I - \hat{\mu} \hat{\gamma} \hat{G}\|_\infty \|\hat{\theta} - \theta_0\|_1 \leq \sqrt{n} \bar{\lambda} \Delta_{1n} \\ &\leq C(1 \vee K^2)(1 + K) \Phi^{-1}(1 - (pmn)^{-1}) \Delta_{1n} \\ \bar{r}_2 &= 0 \\ \bar{r}_3 &\leq KC \bar{\lambda} s \{1 + \mu_n^{-1}\} \|\sqrt{n} \hat{g}(\theta_0)\|_\infty \\ &\leq KC(1 \vee K^2)(1 + K) n^{-1/2} \Phi^{-1}(1 - (pmn)^{-1}) s \log^{1/2}(2m) \end{aligned}$$

Under $K \leq C, \mu_n \geq c$, and $n^{-1/2} s \log(2pmn) \leq u_n$, with probability $1 - C\delta_n$ we have

$$\bar{r}_1 + \bar{r}_2 + \bar{r}_3 \leq C u_n. \quad \blacksquare$$

Proof of Theorem 3.4. We will establish that $\bar{r}_1 + \bar{r}_2 + \bar{r}_3 = o_P(\log^{-1/2} p)$. Throughout the proof we assume $s \geq 1$.

By Lemma C.1(4) with probability $1 - \delta_n$ we have

$$\begin{aligned} \|\sqrt{n} \hat{g}(\theta_0)\|_\infty &= \max_{k \in [m]} |\mathbb{G}_n(g_k(X, \theta_0))| \\ &\leq C \max_{k \in [m]} \mathbb{E}[g_k^2(X, \theta_0)]^{1/2} \sqrt{\log(2m)} \\ &\quad + C n^{-1/2} \mathbb{E}[\max_{i \in [n]} \|g(X, \theta_0)\|_\infty^2] \{\delta_n^{-1} + \log(2m)\} \\ &\leq C' \sqrt{\log(2m)} \end{aligned}$$

under the growth condition $n^{-1/2} \mathbb{E}[\max_{i \in [n]} \|g(X, \theta_0)\|_\infty^2] \{\delta_n^{-1} + \log(2m)\} \leq c \log^{1/2}(2m)$.

By Theorem 3.2 with $\ell_n = C' \sqrt{\log(2m)}$, we have for $q \in \{1, 2\}$ that

$$\|\hat{\theta} - \theta_0\|_q \leq \Delta_{qn} := C n^{-1/2} s^{1/q} \ell_n \{(1 + L) \mu_n^{-1} + C\} \leq C' n^{-1/2} s^{1/q} \log^{1/2}(2m) \quad (\text{C.7})$$

where we used that $\mu_n \geq c$ and $L \leq C$.

Moreover by Lemma 3.9 we have that

$$\|\hat{G} - G\|_\infty \leq C \Delta_{2n}, \quad \|\hat{G} - \tilde{G}\|_\infty \leq C \Delta_{2n} \quad \text{and} \quad \|\hat{\Omega} - \Omega\|_\infty \leq C \Delta_{2n}$$

under $B_n \Delta_{1n} \log^{1/2}(m/\delta_n) \leq C n^{1/2} \Delta_{2n}$ (implied by $B_n s^{1/2} \log(mn) \leq C n^{1/2} \log(2m)$), and $s \geq 1$.

In order to apply Lemma 3.7, recall the choices of penalty

$$\bar{\lambda} = \lambda_j^\gamma = \frac{1}{2} \lambda_j^\mu = n^{-1/2 + \bar{a}} \Phi^{-1}(1 - (mpn)^{-1}) \quad (\text{C.8})$$

so that we consider $\ell_n = n^{1/2}\bar{\lambda} \geq C(1 \vee K^2)(1 + K)n^{1/2}\Delta_{2n}$ for some constant $C > 0$ chosen sufficiently large so they satisfy the requirements on the penalty choices of Lemma 3.7. Therefore, Lemma 3.7 yields with probability $1 - C\delta_n$ that

$$\max_{j \in [p]} \|\hat{\gamma}_j - \gamma_{0j}\|_1 \leq C\bar{\lambda}s\{1 + \mu_n^{-1}\} \quad \text{and} \quad \|\hat{\gamma}_j\|_1 \leq \|\gamma_{0j}\|_1 \leq K \quad \text{for all } j \in [p]$$

$$\max_{j \in [p]} \|\hat{\mu}_j - \mu_{0j}\|_1 \leq C'\bar{\lambda}s\{1 + \mu_n^{-1}\} \quad \text{and} \quad \|\hat{\mu}_j\|_1 \leq \|\mu_{0j}\|_1 \leq K \quad \text{for all } j \in [p]$$

Now we are in position to bound \bar{r}_1, \bar{r}_2 and \bar{r}_3 . By Lemma 3.6 and using the definition of the estimators (3.33) we have

$$\begin{aligned} \bar{r}_1 &= \sqrt{n}\|I - \hat{\mu}\hat{\gamma}\hat{G}\|_\infty \|\hat{\theta} - \theta_0\|_1 \leq \sqrt{n}\bar{\lambda}\Delta_{1n} \\ &\leq n^{\bar{a}}\Phi^{-1}(1 - (pmn)^{-1})\Delta_{1n} \\ \bar{r}_2 &= \sqrt{n} \max_{j \in [p]} \|\hat{\mu}_j\|_1 \max_{j \in [p]} \|\hat{\gamma}_j\|_1 \|\hat{G} - \tilde{G}\|_\infty \|\hat{\theta} - \theta_0\|_1 \\ &\leq CK^2 n^{1/2} \Delta_{2n} \Delta_{1n} \\ \bar{r}_3 &\leq KC\bar{\lambda}s\{1 + \mu_n^{-1}\} \|\sqrt{n}\hat{g}(\theta_0)\|_\infty \\ &\leq KC\bar{n}^{-1/2+\bar{a}}\Phi^{-1}(1 - (pmn)^{-1})s \log^{1/2}(2m) \end{aligned}$$

Under $K + \mu_n^{-1} \leq C$, and $n^{-1/2+\bar{a}}s \log(2pmn) \leq u_n$, with probability $1 - C\delta_n$ we have

$$\bar{r}_1 + \bar{r}_2 + \bar{r}_3 \leq Cu_n. \quad \blacksquare$$

Lemma C.1. Let $X_i, i = 1, \dots, n$, be independent random vectors in \mathbb{R}^p , $p \geq 3$. Define $\bar{m}_k^k := \max_{j \in [p]} \frac{1}{n} \sum_{i=1}^n \mathbb{E}[|X_{ij}|^k]$ and $M_k^k \geq \mathbb{E}[\max_{i \leq n} \|X_i\|_\infty^k]$. Then we have the following bounds:

$$\begin{aligned} (1) \quad \mathbb{E} \left[\max_{j \in [p]} \frac{1}{n} \sum_{i=1}^n |X_{ij}| \right] &\leq CM_1 n^{-1} \log p + C\bar{m}_1 \\ (2) \quad \mathbb{E} \left[\max_{j \in [p]} \frac{1}{n} \left| \sum_{i=1}^n X_{ij} - \mathbb{E}[X_{ij}] \right| \right] &\leq C\bar{m}_2 \sqrt{n^{-1} \log p} + CM_2 n^{-1} \log p \\ (3) \quad \mathbb{E} \left[\max_{j \in [p]} \frac{1}{n} \left| \sum_{i=1}^n |X_{ij}|^k - \mathbb{E}[|X_{ij}|^k] \right| \right] &\leq \frac{CM_k^k \log p}{n} + C \sqrt{\frac{M_k^k \bar{m}_k^k \log p}{n}} \end{aligned}$$

for some universal constant C . Moreover, for $\bar{q} \geq 2$, we have with probability $1 - t^{-\bar{q}/2}$

$$(4) \quad \max_{j \in [p]} |\mathbb{G}_n(X_j)| \leq C\bar{m}_2 \sqrt{\log p} + n^{-1/2} C_{\bar{q}} \{M_{\bar{q}} t^{1/2} + M_2(t + \log p)\}.$$

Proof. The proofs of the first inequalities are given in Lemmas 8 and 9 of [49]. The proof of the last inequality can be found in [24].

The last result follows from the second inequality and Theorem 5.1 in [47] with $\alpha = 1$. Indeed we have with probability $1 - t^{\bar{q}/2}$

$$\max_{j \in [p]} |\mathbb{G}_n(X_j)| \leq 2 \left\{ \bar{m}_2 \sqrt{\log p} + n^{-1/2} M_2 \log p \right\} + K(\bar{q}) \left[(\bar{m}_2 + n^{-1/2} M_{\bar{q}}) \sqrt{t} + n^{-1/2} M_2 t \right]$$

The result follows by collecting the terms. \blacksquare

APPENDIX D. TECHNICAL LEMMAS

Lemma D.1. For any $q > 1/a$,

$$A^q \sum_{j=1}^p j^{-aq} \mathbf{1}\{j \geq (A/\lambda)^{1/a}\} \leq \frac{2^{aq} s \lambda^q}{aq - 1},$$

where $s := \lceil (A/\lambda)^{1/a} \rceil$.

Proof of Lemma D.1. We consider two cases separately: $A \leq \lambda$ and $A > \lambda$. When $A \leq \lambda$, we have

$$\sum_{j=1}^p j^{-aq} \mathbf{1}\{j \geq (A/\lambda)^{1/a}\} \leq \sum_{j=1}^p j^{-aq} = 1 + \sum_{j=2}^p j^{-aq} \leq 1 + \int_1^\infty x^{-aq} dx = 1 + \frac{1}{aq - 1},$$

so that

$$A^q \sum_{j=1}^p j^{-aq} \mathbf{1}\{j \geq (A/\lambda)^{1/a}\} \leq \left(1 + \frac{1}{aq - 1}\right) \lambda^q = \left(1 + \frac{1}{aq - 1}\right) s \lambda^q \leq \frac{2^{aq} s \lambda^q}{aq - 1}.$$

When $A > \lambda$, we have $s = \lceil (A/\lambda)^{1/a} \rceil \geq 2$, and

$$\sum_{j=1}^p j^{-aq} \mathbf{1}\{j \geq (A/\lambda)^{1/a}\} = \sum_{j=s}^p j^{-aq} \leq \int_{s-1}^\infty x^{-aq} dx = \frac{(s-1)^{1-aq}}{aq - 1},$$

so that

$$A^q \sum_{j=1}^p j^{-aq} \mathbf{1}\{j \geq (A/\lambda)^{1/a}\} \leq \frac{s}{aq - 1} \left(\frac{A}{(s-1)^a} \right)^q \leq \frac{2^{aq} s \lambda^q}{aq - 1}$$

since $2(s-1) \geq (A/\lambda)^{1/a}$. Conclude that the asserted claim holds in both cases. \blacksquare

Lemma D.2 (General Lower Bound for $k(\theta_0, \ell)$ in Exactly Sparse Models). Recall the definition of l -sparse smallest and l -sparse largest singular values $\sigma_{\min}(l)$ and $\sigma_{\max}(l)$ in (3.8) and assume that Condition ES is satisfied. Then

$$k(\theta_0, \ell_q) \geq \max_{l \geq s} \frac{\sigma_{\min}(l)}{\sqrt{l}} \left(1 - \frac{\sigma_{\max}(l)}{\sigma_{\min}(l)} \sqrt{\frac{4s}{l}} \right) \frac{s^{1/2-1/q}}{2 + 4\sqrt{s/l}}, \quad q \in \{1, 2\}. \quad (\text{D.1})$$

Proof of Lemma D.2 By Condition ES, there exists $T \subset \{1, \dots, p\}$ such that $|T| = s$ and $\theta_{0j} = 0$ for all $j \in T^c$. Also, for any $\theta \in \mathcal{R}(\theta_0)$, we have $\|\theta_T\|_1 + \|\theta_{T^c}\|_1 = \|\theta\|_1 \leq \|\theta_0\|_1$, and so $\|\theta_{T^c}\|_1 \leq \|\theta_0\|_1 - \|\theta_T\|_1 \leq \|(\theta - \theta_0)_T\|_1$ by the triangle inequality. Hence, $k(\theta_0, \ell_q) \geq \kappa_q^G(s, 1)$ for $\kappa_q^G(s, 1)$ defined in front of Theorem 1 in [17]. Thus, given that by Theorem 1 in [17], $\kappa_q^G(s, 1)$ is bounded from above by the right-hand side of (D.1), the asserted claim follows. \blacksquare

Lemma D.3 (Maximal Inequality Based on Contraction Principle). *Let $(X_i)_{i=1}^n$ be independent random vectors with common support \mathcal{X} , and let \mathbb{G}_n be the corresponding empirical process. Consider the maximum over suprema of empirical processes with contractive structure:*

$$\max_{j \in [m]} \sup_{\eta \in \Delta} |\mathbb{G}_n(h_j(X, Z_{u(j)}(X)'v_{u(j)}))|,$$

where (1) $u(j) \in [\bar{u}]$ for all $j \in [m]$; (2) v_u is a parameter vector in \mathbb{R}^{p_u} for all $u \in [\bar{u}]$; (3) Δ is a set in \mathbb{R}^p with $p = p_1 + \dots + p_{\bar{u}}$; (4) $\eta = (v'_1, \dots, v'_{\bar{u}})'$; (5) for all $j \in [m]$, the link function h_j is a measurable map from $\mathcal{X} \times \mathbb{R}$ to \mathbb{R} , is pointwise Lipschitz: $|h_j(x, t) - h_j(x, s)| \leq L_j(x)|t - s|$ for all $x \in \mathcal{X}$ and $t, s \in \mathbb{R}$, and is passing through the origin $h_j(x, 0) = 0$, for all $x \in \mathcal{X}$; (6) for all $j \in [m]$, the function L_j is a measurable map from \mathcal{X} to \mathbb{R} ; and (7) for all $u \in [\bar{u}]$, the function Z_u is a measurable map from \mathcal{X} to \mathbb{R}^{p_u} .

Suppose that

$$\sup_{\eta \in \Delta, j \in [m]} \mathbb{E}_n \text{Var}(h_j(X, Z_{u(j)}(X)'v_{u(j)})) \leq B_{1n}^2, \quad \max_{j \in [m], k \in [p_{u(j)}]} \mathbb{E}_n [L_j^2(X) Z_{u(j)k}^2(X)] \leq B_{2n}^2$$

with probability at least $1 - \delta_n$. Then for all $t \geq 4B_{1n}$,

$$\mathbb{P} \left(\sup_{\eta \in \Delta, j \in [m]} |\mathbb{G}_n(h_j(X, Z_{u(j)}(X)'v_{u(j)}))| > t \right) \leq 4\delta_n + 8pm \exp \left(-\frac{t^2}{128B_{2n}^2 \|\Delta\|_{1,v}^2} \right),$$

where $\|\Delta\|_{1,v} := \max_{u \in [\bar{u}]} \sup_{\eta \in \Delta} \|v_u\|_1$. Thus,

$$\mathbb{P} \left(\sup_{\eta \in \Delta, j \in [m]} |\mathbb{G}_n(h_j(X, Z_{u(j)}(X)'v_{u(j)}))| > t \right) \leq 5\delta_n$$

as long as $t \geq \{4B_{1n}\} \vee \{8\sqrt{2}B_{2n}\|\Delta\|_{1,v} \log^{1/2}(8pm/\delta_n)\}$.

Proof of Lemma D.3. The proof is a variant of the argument given by [8]. Since the second asserted claim follows immediately from the first one, it suffices to prove the first one. To do so, we split the proof into two steps.

Step 1. We first invoke a symmetrization lemma for probabilities. Let $(\sigma_i)_{i=1}^n$ be i.i.d. copies of the Rademacher random variable σ , which takes on values $\{-1, 1\}$ with equal probabilities. Fix any t such that

$$t^2 \geq 16B_{1n}^2 \geq 16 \sup_{\eta \in \Delta, j \in [m]} \mathbb{E}_n \text{Var}(h_j(X, Z_{u(j)}(X)'v_{u(j)})).$$

Then by Lemma 2.3.7 in [113] and Chebyshev's inequality,

$$\begin{aligned} (\star) &:= \mathbb{P} \left(\sup_{\eta \in \Delta, j \in [m]} |\mathbb{G}_n(h_j(X, Z_{u(j)}(X)'v_{u(j)}))| > t \right) \\ &\leq 4\mathbb{P} \left(\sup_{\eta \in \Delta, j \in [m]} |\mathbb{G}_n(\sigma h_j(X, Z_{u(j)}(X_i)'v_{u(j)}))| > t/4 \right). \end{aligned}$$

Next, let

$$\mathcal{A} := \sup_{\eta \in \Delta, j \in [m]} |\mathbb{G}_n(\sigma h_j(X, Z_{u(j)}(X)'v_{u(j)}))|$$

and define the event $\Omega := \{\max_{j \in [m], k \in [p_{u(j)}]} \mathbb{E}_n[L_j^2(X)Z_{u(j)k}^2(X)] \leq B_{2n}^2\}$. Then

$$\mathbb{P}(\mathcal{A} > t/4) \leq \mathbb{P}(\mathcal{A} > t/4 \mid \Omega) + \mathbb{P}(\Omega^c) \leq \mathbb{E}[\mathbb{P}(\mathcal{A} > t/4 \mid (X_i)_{i=1}^n, \Omega)] + \delta_n,$$

where we used that $\mathbb{P}(\Omega^c) \leq \delta_n$ by assumption.

Now, by Markov's inequality, for $\psi := t/(16B_{2n}^2\|\Delta\|_{1,v}^2)$,

$$\begin{aligned} \mathbb{P}(\mathcal{A} > t/4 \mid (X_i)_{i=1}^n, \Omega) &\leq \exp(-\psi t/4) \mathbb{E}[\exp(\psi \mathcal{A}) \mid (X_i)_{i=1}^n, \Omega] \\ &\leq 2pm \exp(-\psi t/4) \exp(2\psi^2 B_{2n}^2 \|\Delta\|_{1,v}^2) \\ &= 2pm \exp(-t^2/\{128B_{2n}^2\|\Delta\|_{1,v}^2\}), \end{aligned} \tag{D.2}$$

where (D.2) is established in Step 2 below. Therefore,

$$(\star) \leq 4\delta_n + 8pm \exp\left(-\frac{t^2}{128B_{2n}^2\|\Delta\|_{1,v}^2}\right),$$

which is the first asserted claim. It remains to establish (D.2).

Step 2. Here, we bound $\mathbb{E}[\exp(\psi \mathcal{A}) \mid (X_i)_{i=1}^n, \Omega]$ and establish (D.2). In this step, we will condition throughout on $\{(X_i)_{i=1}^n, \Omega\}$ but omit explicit notation for this conditioning to keep the notation lighter.

We have

$$\begin{aligned} &\mathbb{E} \left[\exp \left(\psi \sup_{\eta \in \Delta, j \in [m]} |\mathbb{G}_n(\sigma h_j(X, Z_{u(j)}(X)'v_{u(j)}))| \right) \right] \\ &= \mathbb{E} \left[\exp \left(\frac{\psi}{\sqrt{n}} \sup_{\eta \in \Delta, j \in [m]} \left| \sum_{i=1}^n \sigma_i h_j(X_i, Z_{u(j)}(X_i)'v_{u(j)}) \right| \right) \right] \\ &\leq m \max_{j \in [m]} \mathbb{E} \left[\exp \left(\frac{2\psi}{\sqrt{n}} \sup_{\eta \in \Delta} \left| \sum_{i=1}^n \sigma_i L_j(X_i) Z_{u(j)}(X_i)'v_{u(j)} \right| \right) \right] \\ &= m \max_{j \in [m]} \mathbb{E} \left[\exp \left(2\psi \sup_{\eta \in \Delta} |\mathbb{G}_n(\sigma L_j(X) Z_{u(j)}(X)'v_{u(j)})| \right) \right] \end{aligned}$$

by Lemma D.4 in this appendix with

$$h_{ij}(t) = h_j(X_i, t), \quad \gamma_{ij} = L_j(X_i), \quad G(t) = \exp(\psi t / \sqrt{n}), \quad i \in [n], \quad j \in [m], \quad t \in \mathbb{R}$$

and

$$T_j = \{t = (t_1, \dots, t_n)' \in \mathbb{R}^n : t_i = Z_{u(j)}(X_i)' v_{u(j)} \text{ for all } i \in [n] \text{ and some } \eta \in \Delta\}, \quad j \in [m].$$

Further, for all $j \in [m]$,

$$\begin{aligned} & \mathbb{E} \left[\exp \left(2\psi \sup_{\eta \in \Delta} |\mathbb{G}_n(\sigma L_j(X) Z_{u(j)}(X)' v_{u(j)})| \right) \right] \\ & \leq \mathbb{E} \left[\exp \left(2\psi \max_{k \in [p_{u(j)}]} |\mathbb{G}_n(\sigma L_j(X) Z_{u(j)k}(X))| \|\Delta\|_{1,v} \right) \right] \\ & \leq p \max_{k \in [p_{u(j)}]} \mathbb{E} \left[\exp \left(2\psi |\mathbb{G}_n(\sigma L_j(X) Z_{u(j)k}(X))| \|\Delta\|_{1,v} \right) \right] \end{aligned}$$

by Hölder's inequality and the observation that $\exp(t) > 0$ for all $t \in \mathbb{R}$. Moreover, for all $j \in [m]$ and $k \in [p_{u(j)}]$,

$$\begin{aligned} & \mathbb{E} \left[\exp \left(2\psi |\mathbb{G}_n(\sigma L_j(X) Z_{u(j)k}(X))| \|\Delta\|_{1,v} \right) \right] \\ & \leq 2 \mathbb{E} \left[\exp \left(2\psi \mathbb{G}_n(\sigma L_j(X) Z_{u(j)k}(X)) \|\Delta\|_{1,v} \right) \right] \\ & \leq 2 \exp \left(2\psi^2 \mathbb{E}_n [L_j(X)^2 Z_{u(j)k}(X)^2] \|\Delta\|_1^2 \right) \leq 2 \exp \left(2\psi^2 B_{2n}^2 \|\Delta\|_{1,v}^2 \right), \end{aligned}$$

where the first inequality follows from observing that $\mathbb{G}_n(\sigma L_j(X) Z_{u(j)k}(X))$'s are symmetrically distributed, the second from sub-Gaussianity of $\mathbb{G}_n(\sigma L_j(X) Z_{u(j)k}(X))$ (see the proof of Lemma 2.2.7 in [113], for example), and the third from the definition of Ω (recall that we implicitly condition on Ω). Combining the inequalities above gives (D.2) and completes the proof of the lemma. \blacksquare

Lemma D.4 (Contraction Principle for Maxima). *For $i \in [n]$ and $j \in [m]$, let the functions $h_{ij} : \mathbb{R} \rightarrow \mathbb{R}$ be such that*

$$|h_{ij}(t) - h_{ij}(s)| \leq \gamma_{ij} |t - s|, \text{ for all } (t, s) \in \mathbb{R}^2 \text{ and } h_{ij}(0) = 0.$$

Also, let $G : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ be a non-decreasing convex function and for each $j \in [m]$, let T_j be a bounded set in \mathbb{R}^n . Then

$$\mathbb{E} \left[G \left(\sup_{j \in [m], t \in T_j} \left| \sum_{i=1}^n \sigma_i h_{ij}(t_i) \right| \right) \right] \leq m \max_{j \in [m]} \mathbb{E} \left[G \left(2 \sup_{t \in T_j} \left| \sum_{i=1}^n \sigma_i \gamma_{ij} t_i \right| \right) \right],$$

where $(\sigma_i)_{i=1}^n$ are i.i.d. Rademacher random variables, i.e. random variables taking values -1 and $+1$ with probability $1/2$ each.

Proof of Lemma D.4. Since $G(x) \geq 0$ for all $x \in \mathbb{R}_+$, we have

$$\begin{aligned} \mathbb{E} \left[G \left(\sup_{j \in [m], t \in T_j} \left| \sum_{i=1}^n \sigma_i h_{ij}(t_i) \right| \right) \right] &\leq \mathbb{E} \left[\sum_{j=1}^m G \left(\sup_{t \in T_j} \left| \sum_{i=1}^n \sigma_i h_{ij}(t_i) \right| \right) \right] \\ &\leq m \max_{j \in [m]} \mathbb{E} \left[G \left(\sup_{t \in T_j} \left| \sum_{i=1}^n \sigma_i h_{ij}(t_i) \right| \right) \right]. \end{aligned}$$

Also, by Theorem 4.12 in [84], for each $j \in [m]$,

$$\mathbb{E} \left[G \left(\sup_{t \in T_j} \left| \sum_{i=1}^n \sigma_i h_{ij}(t_i) \right| \right) \right] \leq \mathbb{E} \left[G \left(2 \sup_{t \in T_j} \left| \sum_{i=1}^n \sigma_i \gamma_{ij} t_i \right| \right) \right].$$

Combining these inequalities gives the asserted claim. \blacksquare

REFERENCES

- [1] Abadie, A. (2005). Semiparametric difference-in-differences estimators. *Review of Economic Studies* **72** 1-19.
- [2] Andrews, D. (1994). Asymptotics for semiparametric econometric models via stochastic equicontinuity. *Econometrica* **62** 43-72.
- [3] Athey, S. and Imbens, G. (2016). Recursive partitioning for heterogeneous causal effects. *Proceedings of the National Academy of Sciences* **113** 7353-7360.
- [4] Athey, S., Imbens, G. and Wager, S. (2017). Approximate residual balancing: Debiased inference of average treatment effects in high dimensions. *Journal of the Royal Statistical Society, Series B*, forthcoming.
- [5] Barber, R. and Candès, E. (2015). Controlling the false discovery rate via knockoffs. *Annals of Statistics* **43** 2055-2085.
- [6] Bekker, P. (1994). Alternative approximations to the distributions of instrumental variables estimators. *Econometrica* **63** 657-681.
- [7] Belloni, A., Chen, D., Chernozhukov, V. and Hansen, C. (2012). Sparse models and methods for optimal instruments with an application to eminent domain. *Econometrica* **80** 2369-2429.
- [8] Belloni, A. and Chernozhukov, V. (2011a). ℓ_1 -penalized quantile regression in high-dimensional sparse models. *Annals of Statistics* **39** 82-130.
- [9] Belloni, A. and Chernozhukov, V. (2011b). High-dimensional sparse econometric models: an introduction. *Springer Lecture Notes*.
- [10] Belloni, A. and Chernozhukov, V. (2013). Least squares after model selection in high-dimensional sparse models. *Bernoulli* **2** 521-547.
- [11] Belloni, A., Chernozhukov, V., Chetverikov, D. and Wei, Y. (2017). Uniformly valid post-regularization confidence regions for many functional parameters in Z-estimation framework. arXiv:1512.07619.

- [12] Belloni, A., Chernozhukov, V., Fernández-Val, I. and Hansen, C. (2017). Program evaluation and causal inference with high-dimensional data. *Econometrica* **85** 233-298.
- [13] Belloni, A., Chernozhukov, V. and Hansen, C. (2010). Inference for high-dimensional sparse econometric models. *Advances in Economics and Econometrics. 10th World Congress of the Econometric Society. August 2010 III.* 245-295.
- [14] Belloni, A., Chernozhukov, V. and Hansen, C. (2014). Inference on treatment effects after selection amongst high-dimensional controls. *Review of Economic Studies* **81** 608-650.
- [15] Belloni, A., Chernozhukov, V. and Hansen, C. (2010). LASSO methods for Gaussian instrumental variables models. ArXiv:1012.1297.
- [16] Belloni, A., Chernozhukov, V., Hansen, C., and Kozbur, D. (2016). Inference in high-dimensional panel models with an application to gun control. *Journal of Business and Economic Statistics*. **34** 590-605.
- [17] Belloni, A., Chernozhukov, V., Hansen, C., and Newey, W. (2017). Simultaneous confidence intervals for high-dimensional linear models with many endogenous variables. ArXiv:1712.08102.
- [18] Belloni, A., Chernozhukov, V., and Kaul, A. (2017). Confidence bands for coefficients in high-dimensional linear models with error-in-variables. ArXiv:1703.00469.
- [19] Belloni, A., Chernozhukov, V. and Kato, K. (2015). Uniform post-selection inference for least absolute deviation regression and other Z-estimation problems. *Biometrika* **102** 77-94.
- [20] Belloni, A., Chernozhukov, V., Kaul, A., Rosenbaum, M., and Tsybakov A. (2017). Pivotal estimation via self-normalization for high-dimensional linear models with error-in-variables. arXiv:1708.08353.
- [21] Belloni, A., Chernozhukov, V., and Wang L. (2011). Square-root-lasso: Pivotal recovery of sparse signals via conic programming. *Biometrika* **98**, 791-806.
- [22] Belloni, A., Chernozhukov, V., and Wei, Y. (2016). Post-selection inference for generalized linear models with many controls. *Journal of Business and Economic Statistics*. **34** 606-619.
- [23] Belloni, A., Rosenbaum, M., and Tsybakov A. (2016). An $\{\ell_1, \ell_2, \ell_\infty\}$ -regularization approach to high-dimensional errors-in-variables models. *Electronic Journal of Statistics* **10** 1729-1750.
- [24] Belloni, A., Rosenbaum, M., and Tsybakov A. (2016). Linear and conic programming approaches to high-dimensional errors-in-variables models. *Journal of the Royal Statistical Society, Series B*, forthcoming.
- [25] Benjamin, D., Cesarini, D., Chabris, C., Glaeser, E., and Laibson, D. (2012). The promises and pitfalls of geno-economics. *Annual Review of Economics* **4** 627-662.
- [26] Benjamini, Y. and Hochberg, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society. Series B* **57** 289-300.
- [27] Benjamini, Y. and Yekutieli, D. (2001). The control of the false discovery rate in multiple testing under dependency. *Annals of Statistics* **29** 1165-1188.

- [28] Bentkus, V. (2003). On the dependence of the Berry-Esseen bound on dimension. *Journal of Statistical Planning and Inference* **113** 385-402.
- [29] Bentkus, V. (2005). A Lyapunov-type bound in \mathbb{R}^d . *Theory Probab. Appl.* **49** 311-323.
- [30] Bickel, P., Ritov, Y. and Tsybakov A. (2009). Simultaneous analysis of Lasso and Dantzig selector. *Annals of Statistics* **37** 1705-1732.
- [31] Boucheron, A., Lugosi, M. and Massart, P. (2013). *Concentration Inequalities: A Nonasymptotic Theory of Independence*. Oxford University Press.
- [32] Bühlmann, P. and van de Geer, S. (2011). *Statistics for High-Dimensional Data: Methods, Theory, and Applications*. Springer.
- [33] Cai, T. and Sun, W. (2017). Large-scale global and simultaneous inference: estimation and testing in very high dimensions. *Annual Review of Economics* **9** 411-439.
- [34] Cai, T. and Guo, Z. (2016) Accuracy assessment for high-dimensional linear regression. arXiv:1603.03474.
- [35] Candès, E., Fan, Y., Janson, L. and Lv, J. (2017). Panning for gold: model-X knockoffs for high-dimensional controlled variable selection. arXiv:1610.02351.
- [36] Candès, E. and Tao, E. (2007). The Dantzig selector: statistical estimation when p is much larger than n . *Annals of Statistics* **35** 2313-2351.
- [37] Carrasco, M. (2012). A regularization approach to the many instruments problem. *Journal of Econometrics* **170** 383-398.
- [38] Cattaneo, M., Jansson, M. and Ma, X. (2017). Two-step estimation and inference with possibly many included covariates. University of Michigan working paper.
- [39] Cattaneo, M., Jansson, M. and Newey, W. (2016). Alternative asymptotics and the partially linear model with many regressors. *Econometric Theory*, forthcoming.
- [40] Cattaneo, M., Jansson, M. and Newey, W. (2017). Inference in linear regression models with many covariates and heteroskedasticity. *Journal of the American Statistical Association*, forthcoming.
- [41] Chamberlain, G. and Imbens, G. (2004). Random effects estimators with many instrumental variables. *Econometrica* **72** 295-306.
- [42] Chen, X. (2017). Gaussian and bootstrap approximations for high-dimensional U-statistics and their applications. *Annals of Statistics*, to appear.
- [43] Chen, X. and Kato, K. (2017). Randomized incomplete U -statistics in high dimensions. arXiv:1712:00771.
- [44] Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W. and Robins, J. (2018). Double/debiased machine learning for treatment and structural parameters. *Econometrics Journal* **21** C1-C68.
- [45] Chernozhukov, V., Chetverikov, D. and Kato, K. (2013a). Gaussian approximations and multiplier bootstrap for maxima of sums of high-dimensional random vectors. *Annals of Statistics* **41** 2786-2819.
- [46] Chernozhukov, V., Chetverikov, D. and Kato, K. (2013b). Testing many moment inequalities. arXiv:1312.7614.
- [47] Chernozhukov, V., Cheterikov, D., and Kato, K. (2014a). Gaussian approximation of suprema of empirical processes. *Annals of Statistics* **42** 1564-1597.

- [48] Chernozhukov, V., Cheterikov, D., and Kato, K. (2014b). Anti-concentration and honest, adaptive confidence bands. *Annals of Statistics* **42** 1787-1818.
- [49] Chernozhukov, V., Chetverikov, D. and Kato, K. (2015). Comparison and anti-concentration bounds for maxima of Gaussian random vectors. *Probability Theory and Related Fields* **162** 47-70.
- [50] Chernozhukov, V., Chetverikov, D. and Kato, K. (2017a). Central limit theorems and bootstrap in high dimensions. *The Annals of Probability* **45** 2309-2352.
- [51] Chernozhukov, V., Chetverikov, D. and Kato, K. (2017b). Detailed proof of Nazarov's inequality. arXiv:1711.10696.
- [52] Chernozhukov, V., Demirer, M., Duflo, E. and Fernández-Val, I. (2018). Generic machine learning inference on heterogeneous treatment effects in randomized experiments. ArXiv:1712.04802.
- [53] Chernozhukov, V., Goldman, M., Semenova, V. and Taddy, M. (2108). Orthogonal machine learning for demand estimation: High dimensional causal inference in dynamic panels. ArXiv:1712.09988.
- [54] Chernozhukov, V., Hansen, C. and Spindler M. (2015). Valid post-selection and post-regularization inference: An elementary, general approach. *Annual Review of Economics* **7**, 649-688.
- [55] Chernozhukov, V., Hansen, C. and Spindler M. (2016). hdm: High-dimensional metrics. ArXiv:1608.00354.
- [56] Chernozhukov, V., Newey, W. and Robins, J. (2018). Double/de-biased machine learning using regularized Riesz representers. arXiv:1802.08667.
- [57] de la Peña, V., Lai, T. and Shao, Q.M. (2009). *Self-Normalized Processes*. Springer.
- [58] Deng, H. and Zhang, C.-H. (2017). Beyond Gaussian approximation: bootstrap for maxima of sums of independent random vectors. arXiv:1705.09528.
- [59] Dezeure, R., Bühlmann, P. and Zhang, C.-H. (2016). High-dimensional simultaneous inference with the bootstrap. arXiv:1606.03940.
- [60] Dudley, R. (2014). *Uniform Central Limit Theorems*. Cambridge University Press.
- [61] Efron, B. and Hastie, T. (2016). *Computer Age Statistical Inference*. Cambridge University Press.
- [62] Fan, J. and Li, R. (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American Statistical Association* **96** 1348-1360.
- [63] Fan, J. and Lv, J. (2010). A selective overview of variable selection in high dimensional feature space. *Statistica Sinica* **20** 101-148.
- [64] Fan, J., Lv, J. and Qi, L. (2011). Sparse high-dimensional models in economics. *Annual Review of Economics* **3**, 291-317.
- [65] Farrell, M. (2015). Robust inference on average treatment effects with possibly more covariates than observations. *Journal of Econometrics*. **189** 1-23.
- [66] Frank, I. and Friedman, J. (1993). A statistical view of some chemometrics regression tools. *Technometrics* **35**, 109-135.
- [67] Frölich, M. (2007). Nonparametric IV estimation of local average treatment effects with covariates. *Journal of Econometrics* **139** 35-75.

- [68] Gautier, E. and Tsybakov, A. (2014). High-dimensional instrumental variables regression and confidence sets. ArXiv:1105.2454.
- [69] Giraud, C. (2015). *Introduction to high-dimensional statistics*. Chapman and Hall/CRC Press.
- [70] Gold, D., Lederer, J. and Tao, J. (2017). Inference for high-dimensional instrumental variables regression. arXiv:1708.05499.
- [71] Hansen, C. and Kozbur, D. (2014). Instrumental variables estimation with many weak instruments using regularized JIVE. *Journal of Econometrics* **182** 290-308.
- [72] Hansen, C., Kozbur, D. and Misra, S. (2017). Targeted undersmoothing. ArXiv:1706.07328.
- [73] Hastie, T., Tibshirani, R. and Friedman, J. (2009). *Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer.
- [74] Hastie, T., Tibshirani, R. and Wainwright, M. (2015). *Statistical Learning with Sparsity: the Lasso and Generalizations*. Chapman and Hall/CRC Press.
- [75] He, X. and Shao Q.-M. (2000). On parameters of increasing dimensions. *Journal of Multivariate Analysis* **73** 120-135.
- [76] Hirano, K., Imbens, R. and Ridder, G. (2003). Efficient estimation of average treatment effects using the estimated propensity score. *Econometrica* **71** 1161-1189.
- [77] Holm, S. (1979). A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics* **6** 65-70.
- [78] Imbens, G. and Rubin, D. (2015). *Causal inference for statistics, social, and biomedical sciences*. Cambridge University Press.
- [79] Javanmard, A. and Montanari, A. (2014). Confidence intervals and hypothesis testing for high-dimensional regression. *Journal of Machine Learning Research* **15** 2869-2909.
- [80] Jing, B.-Y., Shao, Q.-M. and Wang, Q. (2003). Self-normalized Cramér-type large deviations for independent random variables. *Annals of Probability* **31** 2167-2215.
- [81] Koike, Y. (2017). Gaussian approximation of maxima of Wiener functionals and its application to high-frequency data. arXiv:1709.00353.
- [82] Kozbur, D. (2015). Inference in additively separable models with a high-dimensional set of conditioning variables. arXiv:1503.05436.
- [83] Krueger, A. (1999). Experimental estimates of education production functions. *The Quarterly Journal of Economics* **114** 497-532.
- [84] Ledoux, M. and Talagrand, M. (1991). *Probability in Banach Spaces*. Springer.
- [85] Leeb, H. and Pötscher, B. (2008). Recent developments in model selection and related areas. *Econometric Theory* **24** 319-322.
- [86] Leeb, H. and Pötscher, B. (2008). Sparse estimators and the oracle property, or the return of Hodges' estimator. *Journal of Econometrics* **142** 201-211.
- [87] Lehmann, E. and Romano, J. (2005). *Testing statistical hypotheses*. Springer.
- [88] List, J., Shaikh, A. and Xu, Y. (2015). Multiple hypothesis testing in experimental economics. *NBER working paper 21875*.
- [89] Liu, W. and Shao, Q.-M. (2014). Phase transition and regularized bootstrap in large-scale t -tests with false discovery rate control. *Annals of Statistics* **42** 2003-2025.

- [90] Liu, W. and Luo, S. (2014). Hypothesis testing for high-dimensional regression models. *Working paper*.
- [91] Mammen, E. (1993). Bootstrap and wild bootstrap for high dimensional linear models. *The Annals of Statistics* **21** 255-285.
- [92] Meinhausen, N., Meier, L., and Bühlmann, P. (2009). p-values for high-dimensional regression. *Journal of the American Statistical Association* **104** 1671-1681.
- [93] Negahban, S., Ravikumar, P., Wainwright, M., and Yu, B. (2014). A unified framework for high-dimensional analysis of M-estimators with decomposable regularizers. *Statistical Science* **27** 538-557.
- [94] Newey, W. (1994). The asymptotic variance of semiparametric estimators. *Econometrica* **62** 1349-1382.
- [95] Newey, W. (1997). Convergence rates and asymptotic normality for series estimators. *Journal of Econometrics* **79** 147-168.
- [96] Neyman, J. (1959). Optimal asymptotic tests of composite statistical hypotheses, in U. Grenander, ed. *Probability and Statistics, the Harald Cramer Volume*. Wiley.
- [97] Nickl, R. and van de Geer, S. (2013). Confidence sets in sparse regression. *Annals of Statistics* **41** 2852-2876.
- [98] Ning, Y. and Liu, H. (2017). A general theory of hypothesis tests and confidence regions for sparse high dimensional models. *Annals of Statistics*. **45** 158-195.
- [99] Ning, Y., Zhao, T., and Liu, H. (2015). A likelihood ratio framework for high dimensional semiparametric regression. arXiv:1412.2295.
- [100] Okui, R. (2011). Instrumental variable estimation in the presence of many moment conditions. *Journal of Econometrics* **165** 70-86.
- [101] Romano, J., Shaikh, A. and Wolf, M. (2008a). Formalized data snooping based on generalized error rates. *Econometric Theory* **24** 404-447.
- [102] Romano, J., Shaikh, A. and Wolf, M. (2008b). Control of the false discovery rate under dependence using the bootstrap and subsampling. *Test* **17** 417-442.
- [103] Romano, J. and Wolf, M. (2005). Exact and approximate stepdown methods for multiple hypothesis testing. *Journal of American Statistical Association* **100** 94-108.
- [104] Romano, J. and Shaikh, A. (2010). Inference for the identified set in partially identified econometric models. *Econometrica* **78** 169-211.
- [105] Scharfstein, D., Rotnitzky, A. and Robins, J. (1999). Rejoinder to "Adjusting for non-ignorable drop-out using semiparametric non-response models". *Journal of the American Statistical Association* **94** 1135-1146.
- [106] Talagrand, M. (2011). *Mean Field Models for Spin Glasses*. Springer.
- [107] Tibshirani, R. (1996). Regression shrinkage and selection via the Lasso. *Journal of the Royal Statistical Society, Series B* **58** 267-288.
- [108] van de Geer, S. (2016). *Estimation and Testing under Sparsity*. Springer.
- [109] van de Geer, S., Bühlmann, P., Ritov, Y. and Dezeure, R. (2014). On asymptotically optimal confidence regions and tests for high-dimensional models. *Annals of Statistics* **42** 1166-1202.
- [110] van der Laan, M., Polley, M. and Hubbard, A. (2007). Super learner. *Statistical Applications in Genetics and Molecular Biology* **6**. Retrieved 24 February 2017 from

- doi:10.2202/1544-6115.1309.
- [111] van der Laan, M. and S. Rose (2011). *Targeted Learning: Causal Inference for Observational and Experimental Data*. Springer.
 - [112] van der Vaart, A. W. (1998). *Asymptotic Statistics*. Cambridge University Press.
 - [113] van der Vaart, A. and Wellner, J. (1996). *Weak Convergence and Empirical Processes: With Applications to Statistics*. Springer.
 - [114] Voorman, A., Shojaie, A., and Witten, D. (2014). Inference in high dimensions with the penalized score test. arXiv:1401.2678.
 - [115] Wager, S. and Athey, S. (2017). Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, forthcoming.
 - [116] Wasserman, L. (2014). Stein's method and the bootstrap in low and high dimensions: a tutorial. *Working paper*.
 - [117] Wasserman, L. and Roeder, K. (2009). High-dimensional variable selection. *Annals of Statistics* **37** 2178-2201.
 - [118] White, H. (2000). A reality check for data snooping. *Econometrica* **68** 1097-1126.
 - [119] Yang, Z., Ning, Y., and Liu, H. (2015). On semiparametric exponential family graphical models. ArXiv:1412.8697.
 - [120] Zhai, A. (2017). A high-dimensional CLT in \mathcal{W}_2 distance with near optimal convergence rate. *Probability Theory and Related Fields*, to appear.
 - [121] Zhang, X. and Cheng, G. (2017a). Simultaneous inference for high-dimensional linear models. *Journal of the American Statistical Association* **518** 757-768.
 - [122] Zhang, X. and Cheng, G. (2017b). Gaussian approximation for high dimensional vector under physical dependence. *Bernoulli*, to appear.
 - [123] Zhang, D. and Wu, W. (2017). Gaussian approximation for high-dimensional time series. *Annals of Statistics* **45** 1895-1919.
 - [124] Zhang, C.-H. and Zhang, S. (2014). Confidence intervals for low dimensional parameters in high dimensional linear models. *Journal of the Royal Statistical Society, Series B* **76** 217-242.
 - [125] Zheng, W. and van der Laan, M. (2011). Cross-validated targeted minimum-loss-based estimation. In *Targeted Learning* 459-474. Springer.
 - [126] Zhilova, M. (2016). Non-classical Berry-Esseen inequality and accuracy of the weighted bootstrap. arXiv:1611.0268.
 - [127] Zhu, Y. and Bradic, J. (2017). Linear hypothesis testing in dense high-dimensional linear models. *Journal of the American Statistical Association*, forthcoming.
 - [128] Zhu, Y. and Bradic, J. (2017). A projection pursuit framework for testing general high-dimensional hypothesis. arXiv:1705:01024.