

n° 2016-12

An $\{\ell_1, \ell_2, \ell_\infty\}$ Regularization Approach to High-Dimensional Errors-in-variables Models

A.Belloni¹

M.Rosembaum²

A.B.Tsybakov³

Les documents de travail ne reflètent pas la position du CREST et n'engagent que leurs auteurs.
Working papers do not reflect the position of CREST but only the views of the authors.

¹ The Fuqua School of Business (Duke University). E-mail: abn5@duke.edu

² Laboratoire de Probabilités et Modèles Aléatoires (Université Pierre et Marie Curie), CREST ENSAE-Paris Tech. E-mail: mathieu.rosenbaum@upmc.fr

³ CREST, ENSAE-Paris Tech. E-mail: alexandre.tsybakov@ensae.fr

An $\{\ell_1, \ell_2, \ell_\infty\}$ -Regularization Approach to High-Dimensional Errors-in-variables Models

Alexandre Belloni

*The Fuqua School of Business
Duke University
e-mail: abn5@duke.edu*

Mathieu Rosenbaum

*Laboratoire de Probabilités et Modèles Aléatoires,
Université Pierre et Marie Curie (Paris 6), and
Centre de Recherche en Economie et Statistique,
ENSAE-Paris Tech
e-mail: mathieu.rosenbaum@upmc.fr*

Alexandre B. Tsybakov

*Centre de Recherche en Economie et Statistique,
ENSAE-Paris Tech
e-mail: alexandre.tsybakov@ensae.fr*

Abstract: Several new estimation methods have been recently proposed for the linear regression model with observation errors in the design. Different assumptions on the data generating process have motivated different estimators and analysis. In particular, the literature considered (1) observation errors in the design uniformly bounded by some $\bar{\delta}$, and (2) zero-mean independent observation errors. Under the first assumption, the rates of convergence of the proposed estimators depend explicitly on $\bar{\delta}$, while the second assumption has been essentially applied when an estimator for the second moment of the observational error is available. This work proposes and studies two new estimators which, compared to other procedures for regression models with errors in the design, exploit an additional ℓ_∞ -norm regularization. The first estimator is applicable when both (1) and (2) hold but does not require an estimator for the second moment of the observational error. The second estimator is applicable under (2) and requires an estimator for the second moment of the observation error. Importantly, we impose no assumption on the accuracy of this pilot estimator, in contrast to the previously known procedures. As the recent proposals, we allow the number of covariates to be much larger than the sample size. We establish the rates of convergence of the estimators and compare them with the bounds obtained for related estimators in the literature. These comparisons show interesting insights on the interplay of the assumptions and the achievable rates of convergence.

1. Introduction

Several new estimation methods have been recently proposed for the linear regression model with observation errors in the design. Such problems arise in a variety of applications, see [7, 8, 10, 11]. In this work we consider the following regression model with observation errors in the design:

$$\begin{aligned}y &= X\theta^* + \xi, \\Z &= X + W.\end{aligned}$$

Here the random vector $y \in \mathbb{R}^n$ and the random $n \times p$ matrix Z are observed, the $n \times p$ matrix X is unknown, W is an $n \times p$ random noise matrix, and $\xi \in \mathbb{R}^n$ is a random noise vector. The vector of unknown parameters of interest is θ^* which is assumed to belong to a given convex subset Θ of \mathbb{R}^p characterizing some prior knowledge about θ^* (potentially $\Theta = \mathbb{R}^p$). Similarly to the recent literature on this topic, we consider the setting where the dimension p can be much larger than the sample size

n and the vector θ^* is s -sparse, which means that it has no more than s non-zero components.

The need for new estimators under errors in the design arises from the fact that standard estimators (e.g. Lasso and Dantzig selector) might become unstable, see [8]. To deal with this framework, various assumptions have been considered, leading to different estimators.

A classical assumption in the literature is a uniform boundedness condition on the errors in the design, namely,

$$|W|_\infty \leq \bar{\delta} \text{ almost surely,} \quad (1)$$

where $|\cdot|_q$ denotes the ℓ_q -norm for $1 \leq q \leq \infty$. Note that this assumption allows for various dependences between the errors in the design. In this setting, the Matrix Uncertainty selector (MU selector), which is robust to the presence of errors in the design, is proposed in [8]. The MU selector $\hat{\theta}^{MU}$ is defined as a solution of the minimization problem

$$\min\{|\theta|_1 : \theta \in \Theta, \left|\frac{1}{n}Z^T(y - Z\theta)\right|_\infty \leq \tau_1|\theta|_1 + \tau\}, \quad (2)$$

where the parameters τ_1 and τ depend on the level of the noises of W and ξ respectively. Under appropriate choices of these parameters and suitable assumptions on X , it was shown in [8] that with probability close to 1,

$$|\hat{\theta}^{MU} - \theta^*|_q \leq Cs^{1/q}\{\bar{\delta} + \bar{\delta}^2\}|\theta^*|_1 + Cs^{1/q}\sqrt{\frac{\log p}{n}}, \quad 1 \leq q \leq \infty. \quad (3)$$

Here and in what follows we denote by the same symbol C (or c') different positive constants that do not depend on θ^* , s , n , p , $\bar{\delta}$, but only on the variance parameters σ^2 and σ_*^2 (defined later). The result (3) implies consistency as the sample size n tends to infinity provided that the error in the design goes to zero sufficiently fast to offset $s^{1/q}|\theta^*|_1$, and the number of variables p and the sparsity s of θ^* do not grow too fast relative to the sample size n .

An alternative assumption considered in the literature is that the entries of the random matrix W are independent with zero mean, the values

$$\sigma_j^2 = \frac{1}{n} \sum_{i=1}^n \mathbb{E}(W_{ij}^2), \quad j = 1, \dots, p,$$

are finite, and data-driven estimators $\hat{\sigma}_j^2$ of σ_j^2 are available converging with an appropriate rate. This assumption motivated the idea to compensate the bias of using the observable $Z^T Z$ instead of the unobservable $X^T X$ in (2) thanks to the estimates of σ_j^2 . This compensated MU selector, introduced in [9] and denoted as $\hat{\theta}^C$, is defined as a solution of the minimization problem

$$\min\{|\theta|_1 : \theta \in \Theta, \left|\frac{1}{n}Z^T(y - Z\theta) + \widehat{D}\theta\right|_\infty \leq \tau_1|\theta|_1 + \tau\},$$

where \widehat{D} is the diagonal matrix with entries $\hat{\sigma}_j^2$ and $\tau_1 > 0$ and $\tau > 0$ are constants chosen according to the level of the noises and the accuracy of the $\hat{\sigma}_j^2$.

Rates of convergence of the compensated MU selector were established in [9]. Importantly, the compensated MU selector can be consistent as the sample size n increases even if the error in the design does not vanish. This is in contrast to the case of the MU selector, where the bounds are small only if the bound on the design error $\bar{\delta}$ is small. In particular, under regularity conditions, when θ^* is s -sparse, it is shown in [9] that with probability close to 1,

$$|\hat{\theta}^C - \theta^*|_q \leq Cs^{1/q}\sqrt{\frac{\log p}{n}}(|\theta^*|_1 + 1), \quad 1 \leq q \leq \infty. \quad (4)$$

Under the same alternative assumption, a conic programming based estimator $\hat{\theta}^{Conic}$ has been recently proposed and analyzed in [1]. The estimator $\hat{\theta}^{Conic}$ is defined as the first component of any solution of the optimization problem

$$\min_{(\theta, t) \in \mathbb{R}^p \times \mathbb{R}_+} \{|\theta|_1 + \lambda_2 t : \theta \in \Theta, \left| \frac{1}{n} Z^T (y - Z\theta) + \widehat{D}\theta \right|_\infty \leq \tau_2 t + \tau, |\theta|_2 \leq t\}, \quad (5)$$

where λ_2 , τ_2 and τ are some positive tuning constants. Akin to $\hat{\theta}^C$, this estimator compensates for the bias by using the estimators $\hat{\sigma}_j^2$ of σ_j^2 . However, it exploits a combination of ℓ_1 and ℓ_2 -norm regularization to be more adaptive. It was shown to attain a bound as in (4) and to be computationally feasible since it is cast as a tractable convex optimization problem (a second order cone programming problem). Moreover, under mild additional conditions, with probability close to 1, the estimator (5) achieves improved bounds of the form

$$|\hat{\theta}^{Conic} - \theta^*|_q \leq C s^{1/q} \sqrt{\frac{\log p}{n}} (|\theta^*|_2 + 1), \quad 1 \leq q \leq \infty, \quad (6)$$

provided that \widehat{D} converges to D in sup-norm with the rate $\sqrt{(\log p)/n}$. It is shown in [1] that the rate of convergence in (6) is minimax optimal in the considered model.

There have been other approaches to the errors-in-variables model, usually exploiting some knowledge about the vector θ^* , see [2, 3, 7, 10]. Assuming $|\theta^*|_1$ is known, [7] proposed an estimator $\hat{\theta}'$ defined as the solution of a non-convex program which can be well approximated by an iterative relaxation procedure. In the case where the entries of the regression matrix X are zero-mean subgaussian and θ^* is s -sparse, under appropriate assumptions, it is proved that for the error in ℓ_2 -norm ($q = 2$),

$$|\hat{\theta}' - \theta^*|_2 \leq C(\theta^*) s^{1/2} \sqrt{\frac{\log p}{n}} (|\theta^*|_2 + 1), \quad (7)$$

with probability close to 1. Here, the value $C(\theta^*)$ depends on θ^* , so that there is no guarantee that the estimator attains the optimal bound as in (6). Assuming that the sparsity s of θ^* is known and the non-zero components of θ^* are separated from zero in so that

$$|\theta_j^*| \geq C \sqrt{\frac{\log p}{n}} (|\theta^*|_2 + 1),$$

an orthogonal matching pursuit algorithm to estimate θ^* is introduced in [2, 3]. Focusing as in [7] on the particular case where the entries of the regression matrix X are zero-mean subgaussian, it is shown in [2, 3] that this last estimator satisfies a bound analogous to (6), as well as a consistent support recovery result, without requiring estimates of σ_j^2 .

The main purpose of this work is to show that an additional regularization term based on the ℓ_∞ -norm leads to improved rates of convergence in several situations. We propose two new estimators for θ^* . The first proposal is applicable under a new combination of the assumptions mentioned above. Namely, we assume that the components of the errors in the design are uniformly bounded by $\bar{\delta}$ as in (1), and that the rows of W are independent and with zero mean. However, we will neither assume that a data-driven estimator \widehat{D} is available, nor that specific features of θ^* are known (e.g. s or $|\theta^*|_1$). The estimator is defined as a solution of a regularized optimization problem which uses simultaneously ℓ_1 , ℓ_2 and ℓ_∞ regularization functions. It can be cast as a convex optimization problem and the solution can be easily computed. We study its rates of convergence in various norms in Section 3. One of the conclusions is that for $\bar{\delta} \gg \{(\log p)/n\}^{1/2}$, the new estimator has improved rates of convergence compared to the MU selector. Furthermore, note that the conic estimator $\hat{\theta}^{Conic}$ studied in [1] can also be applied. Indeed, our setting can be embedded into that of [1] with \widehat{D} being the identically zero $p \times p$

matrix, which means that we have an estimator of each σ_j^2 with an error bounded by $\bar{\delta}^2$. Comparing the bounds yields that the conic estimator $\hat{\theta}^{Conic}$ achieves the same rate as our new estimator if $\bar{\delta}$ is smaller than or of the order $\{(\log p)/n\}^{1/4}$. However, there is no bound for $\hat{\theta}^{Conic}$ available when $\bar{\delta} \gg \{(\log p)/n\}^{1/4}$.

The second estimator we propose applies to the same setting as in [1]. The idea of taking advantage of an additional ℓ_∞ -norm regularization can be used to improve the conic estimator $\hat{\theta}^{Conic}$ of [1] whenever the rate of convergence of the estimator \widehat{D} for σ_j^2 , $j = 1, \dots, p$, is slower than $\{(\log p)/n\}^{1/2}$. This motivated us to propose and analyze a modification of the conic estimator. We derive new rates of convergence that can lead to improvements. However, we acknowledge that in the case considered in [1], where the rate of convergence of \widehat{D} is $\{(\log p)/n\}^{1/2}$, there is no gain in the rates of convergence when using the additional ℓ_∞ -norm regularization. Therefore the proposed estimator is of particular interest in applications where it is costly and disruptive to generate precise measurements of X_i (the i -th row of X : the couple (X_i, Z_i) is observed in that case, and y_i is unobserved), while observations on (y_k, Z_k) are readily available. For example, interventions in experiments can be used to precisely measure a particular covariate of subjects but such interventions can invalidate the use of the outcome y_i for that particular subject. This is relevant in epidemiological research, where accurate measurements are commonly difficult to achieve without expensive interviews and tests which tend to impact the future behavior of the subject invalidating the associated outcome, see [4]. In such settings, we observe n observations of (y_k, Z_k) and n^* observations of the pair (X_i, Z_i) . When the (X_i, Z_i) are iid, the later can be typically used to construct estimates $\hat{\sigma}_j^2$ for σ_j^2 so that $|\widehat{D} - D|_\infty = b(\varepsilon) \leq C\sqrt{\{\log(p/\varepsilon)\}/n^*}$ with probability $1 - \varepsilon$. Nonetheless, given the sampling cost structure, it is often the case that $n^* \ll n$ and $b(\varepsilon) \gg \sqrt{(\log p)/n}$.

We emphasize that the new use of the ℓ_∞ -norm in the first order condition is precisely what drives our new rates of convergence. In the settings we are concerned with, the rates of convergence are dominated by the ‘‘crude’’ estimators \widehat{D} of D . The impact of using \widehat{D} instead of D is well controlled by $|(\widehat{D} - D)\theta^*|_\infty \leq b(\varepsilon)|\theta^*|_\infty$ since D is diagonal. Previous estimators that rely on ℓ_1 or ℓ_2 -regularization terms would not be able to fully exploit this structure and would not be sharp if $b(\varepsilon) \gg \sqrt{(\log p)/n}$. Furthermore, the ℓ_∞ -regularization allows us to achieve a convex formulation for the problem. We view the use of the ℓ_∞ -regularization as a new way to increase the adaptivity of an estimator that can be of independent interest in other applications.

The paper is organized as follows. Section 2 contains the notation, main assumptions and some preliminary lemmas needed to determine threshold constants in the algorithms. The definition and properties of our first estimator are given in Section 3 whereas those of our second procedure can be found in Section 4. Section 5 contains simulation results. Some auxiliary lemmas are relegated to an appendix.

2. Notation, assumptions, and preliminary lemmas

In this section, we introduce the assumptions which will be required to derive the rates of convergence of the proposed estimators. One set of conditions pertains to the design matrix and the second to the errors in the model. We also state preliminary lemmas related to the stochastic error terms. We start by introducing some notation.

2.1. Notation

Let $J \subset \{1, \dots, p\}$ be a set of integers. We denote by $|J|$ the cardinality of J . For a vector $\theta = (\theta_1, \dots, \theta_p)$ in \mathbb{R}^p , we denote by θ_J the vector in \mathbb{R}^p whose j -th component satisfies $(\theta_J)_j = \theta_j$ if $j \in J$, and $(\theta_J)_j = 0$ otherwise. For $\gamma > 0$, the random variable η is said to be *sub-gaussian with*

variance parameter γ^2 (or shortly γ -sub-gaussian) if, for all $t \in \mathbb{R}$,

$$\mathbb{E}[\exp(t\eta)] \leq \exp(\gamma^2 t^2/2).$$

A random vector $\zeta \in \mathbb{R}^p$ is said to be *sub-gaussian with variance parameter γ^2* if the inner products (ζ, v) are γ -sub-gaussian for any $v \in \mathbb{R}^p$ with $|v|_2 = 1$.

2.2. Design matrix

The performance of the estimators that we consider below is influenced by the properties of the Gram matrix

$$\Psi = \frac{1}{n} X^T X.$$

We will assume that:

(A1) *The matrix X is deterministic.*

In order to characterize the behavior of the design matrix, we set

$$m_2 = \max_{j=1, \dots, p} \frac{1}{n} \sum_{i=1}^n X_{ij}^2,$$

where X_{ij} are the elements of matrix X and we consider the sensitivity characteristics related to the Gram matrix Ψ . For $u > 0$, define the cone

$$C_J(u) = \{\Delta \in \mathbb{R}^p : |\Delta_{J^c}|_1 \leq u |\Delta_J|_1\},$$

where J is a subset of $\{1, \dots, p\}$. For $q \in [1, \infty]$ and an integer $s \in [1, p]$, the ℓ_q -sensitivity (cf. [5]) is defined as follows:

$$\kappa_q(s, u) = \min_{J: |J| \leq s} \left(\min_{\Delta \in C_J(u): |\Delta|_q = 1} |\Psi \Delta|_\infty \right).$$

Like in [5], we use here the sensitivities to derive the rates of convergence of estimators under sparsity. Importantly, as shown in [5], the approach based on sensitivities is more general than that based on the restricted eigenvalue or the coherence condition, see also [1, 6, 9]. In particular, under those conditions, we have $\kappa_q(s, u) \geq c s^{-1/q}$ for some constant $c > 0$, which implies the usual optimal bounds for the errors.

2.3. Disturbances

Next we turn to the error W in the design and the error ξ in the regression equation. We will make the following assumptions:

- (A2) *The elements of the random vector ξ are independent zero-mean sub-gaussian random variables with variance parameter σ^2 .*
- (A3) *The rows of the noise matrix W are independent zero-mean subgaussian random vectors with variance parameter σ_*^2 , and $\mathbb{E}[W_{ij}W_{ik}] = 0$ for all $1 \leq j < k \leq p$. Furthermore, W is independent of ξ .*

2.4. Bounds on the stochastic error terms

We now state some useful lemmas from [1] and [9] that provide bounds to various stochastic error terms that play a role in our analysis. We state them here because they introduce the thresholds δ_i, δ'_i that will be used in the definition of the estimators. In what follows, D is the diagonal matrix with diagonal elements $\sigma_j^2, j = 1, \dots, p$, and for a square matrix A , we denote by $\text{Diag}\{A\}$ the matrix with the same dimensions as A , the same diagonal elements, and all off-diagonal elements equal to zero.

Lemma 1. *Let $0 < \varepsilon < 1$ and assume (A1)-(A3). Then, with probability at least $1 - \varepsilon$ (for each event),*

$$\begin{aligned} \left| \frac{1}{n} X^T W \right|_\infty &\leq \delta_1(\varepsilon), & \left| \frac{1}{n} X^T \xi \right|_\infty &\leq \delta_2(\varepsilon), & \left| \frac{1}{n} W^T \xi \right|_\infty &\leq \delta_3(\varepsilon), \\ \left| \frac{1}{n} (W^T W - \text{Diag}\{W^T W\}) \right|_\infty &\leq \delta_4(\varepsilon), & \left| \frac{1}{n} \text{Diag}\{W^T W\} - D \right|_\infty &\leq \delta_5(\varepsilon), \end{aligned}$$

where

$$\begin{aligned} \delta_1(\varepsilon) &= \sigma_* \sqrt{\frac{2m_2 \log(2p^2/\varepsilon)}{n}}, & \delta_2(\varepsilon) &= \sigma \sqrt{\frac{2m_2 \log(2p/\varepsilon)}{n}}, \\ \delta_3(\varepsilon) &= \delta_5(\varepsilon) = \varpi(\varepsilon, 2p), & \delta_4(\varepsilon) &= \varpi(\varepsilon, p(p-1)), \end{aligned}$$

and for an integer N ,

$$\varpi(\varepsilon, N) = \max \left(\gamma_0 \sqrt{\frac{2 \log(N/\varepsilon)}{n}}, \frac{2 \log(N/\varepsilon)}{t_0 n} \right),$$

where γ_0, t_0 are positive constants depending only on σ, σ_* .

Lemma 2. *Let $0 < \varepsilon < 1$, $\theta^* \in \mathbb{R}^p$ and assume (A1)-(A3). Then, with probability at least $1 - \varepsilon$,*

$$\left| \frac{1}{n} X^T W \theta^* \right|_\infty \leq \delta'_1(\varepsilon) |\theta^*|_2,$$

where $\delta'_1(\varepsilon) = \sigma_* \sqrt{\frac{2m_2 \log(2p/\varepsilon)}{n}}$. In addition, with probability at least $1 - \varepsilon$,

$$\left| \frac{1}{n} (W^T W - \text{Diag}\{W^T W\}) \theta^* \right|_\infty \leq \delta'_4(\varepsilon) |\theta^*|_2,$$

where

$$\delta'_4(\varepsilon) = \max \left(\gamma_2 \sqrt{\frac{2 \log(2p/\varepsilon)}{n}}, \frac{2 \log(2p/\varepsilon)}{t_2 n} \right),$$

and γ_2, t_2 are positive constants depending only on σ_* .

The proofs of Lemmas 1 and 2 can be found in [9] and [1] respectively.

3. $\{\ell_1, \ell_2, \ell_\infty\}$ -MU selector

In this section, we define and analyze our first estimator. It can be seen as a compromise between the MU selector (2) and the conic estimator (5) achieved thanks to an additional ℓ_∞ -norm regularization. In the setting that we consider now, the estimate \hat{D} is not available but the rows of the design error matrix W are independent with mean 0, and its entries are uniformly bounded. Formally, in this section we make the following assumption:

(A4) *Almost surely, $|W|_\infty \leq \bar{\delta}$.*

Assumptions (A1)-(A4) imply the assumptions in [8]. However, they neither imply or are implied by the assumptions in [9]. That is, it is an intermediary set of conditions relative to the original assumptions for the MU selector in [8] and to those for the compensated MU selector in [9]. Importantly, we do not assume that there are some accurate estimators of σ_j^2 .

Similar to [8], consistent estimates require $\bar{\delta} \rightarrow 0$ as the sample size grows. The bound $\bar{\delta}$ does not necessarily scale with p as $\bar{\delta}$ is not derived from (A3). Assumption (A4) is motivated from sampling schemes where the precision with which the covariate X_{ij} is measured can be controlled by the practitioner to some degree and the technology allows $|Z_{ij} - X_{ij}| \leq \bar{\delta}$; in principle $\bar{\delta}$ could be made smaller if more expensive measurements were performed. Our finite sample analysis keeps track of the dependence of $\bar{\delta}$ explicitly allowing to capture its impact as a potential function of the sample size. A simple illustrative example is when the $W_{ij}, i = 1, \dots, n, j = 1, \dots, p$ are independent zero-mean random variables that are a.s. bounded by $\bar{\delta}$. Then the vector W_i is sub-gaussian with parameter $\bar{\sigma}_*^2 = \bar{\delta}^2$.

We consider the estimator $\hat{\theta}$ such that $(\hat{\theta}, \hat{t}, \hat{u}) \in \mathbb{R}^p \times \mathbb{R}_+ \times \mathbb{R}_+$ is a solution of the following minimization problem:

$$\min_{\theta, t, u} \{|\theta|_1 + \lambda_2 t + \lambda_\infty u : (\theta, t, u) \in \Theta, \left| \frac{1}{n} Z^T (y - Z\theta) \right|_\infty \leq \tau_2 t + \bar{\delta}^2 u + \tau, |\theta|_2 \leq t, |\theta|_\infty \leq u\}, \quad (8)$$

where $\lambda_2 > 0$ and $\lambda_\infty > 0$ are tuning constants and we allow for $(\theta, t, u) \in \Theta$ where Θ is a pre-specified convex set that contains $(\theta^*, |\theta^*|_2, |\theta^*|_\infty)$ and characterizes some prior knowledge (a trivial choice is $\Theta = \mathbb{R}^p \times \mathbb{R}_+ \times \mathbb{R}_+$). This estimator $\hat{\theta}$ will be further referred to as the $\{\ell_1, \ell_2, \ell_\infty\}$ -MU selector.

Remark 1 (Safeguard constraints). In order to further bound t and u , we can add constraints that exploit that $|\cdot|_q \leq |\cdot|_1$ for $q \geq 1$. Therefore, the constraints

$$\theta = \theta^+ - \theta^-, \quad \theta^+ \geq 0, \quad \theta^- \geq 0, \quad w = \sum_{j=1}^p \{\theta_j^+ + \theta_j^-\}, \quad u \leq t, \quad \text{and} \quad t \leq w$$

preserve the convexity of the optimization problem and can potentially yield additional performance. We note that our theoretical results allow for such safeguard constraints to be included in the estimation.

The estimator above attempts to mimic the conic estimator (5) without estimators $\hat{\sigma}_j^2$ for $\sigma_j^2, j = 1, \dots, p$. In order to make θ^* feasible for (8), the contribution of the unknown term $\frac{1}{n} \text{Diag}(W^T W) \theta^*$ needs to be bounded. This is precisely the role of the extra term $\bar{\delta}^2 u$ in the constraint since $|\theta|_\infty \leq u$ and $|\frac{1}{n} \text{Diag}(W^T W)|_\infty \leq \bar{\delta}^2$ almost surely. Note that the use of u and t instead of $|\theta|_\infty$ and $|\theta|_2$ in the constraint makes (8) a convex programming problem.

This new estimator exploits Assumptions (A2)-(A4) to achieve a rate of convergence that is intermediary relative to the rate of the MU selector and to that of the conic estimator.

Set $\tau_2 = \delta'_1(\varepsilon) + \delta'_4(\varepsilon)$ and $\tau = \delta_2(\varepsilon) + \delta_3(\varepsilon)$. Note that τ_2 and τ are of the order $\sqrt{(\log p)/n}$. The next theorem summarizes the performance of the estimator defined by solving (8).

Theorem 1. *Let Assumptions (A1)-(A4) hold. Assume that the true parameter θ^* is s -sparse and $(\theta^*, |\theta^*|_2, |\theta^*|_\infty)$ belongs to Θ . Let $0 < \varepsilon < 1$, $1 \leq q \leq \infty$ and $0 < \lambda_2, \lambda_\infty < \infty$. Let $\hat{\theta}$ be the $\{\ell_1, \ell_2, \ell_\infty\}$ -MU selector. If $\kappa_q(s, 1 + \lambda_2 + \lambda_\infty) \geq cs^{-1/q}$ for some constant $c > 0$, then with probability at least $1 - 7\varepsilon$,*

$$|\hat{\theta} - \theta^*|_q \leq Cs^{1/q} \sqrt{\frac{\log(c'p/\varepsilon)}{n}} (|\theta^*|_1 + 1) + Cs^{1/q} \bar{\delta}^2 |\theta^*|_1, \quad (9)$$

for some constants $C > 0$ and $c' > 0$ (here we set $s^{1/\infty} = 1$).

If in addition, $(1 + \lambda_2 + \lambda_\infty)\{\bar{\delta}^2 \lambda_\infty^{-1} + \lambda_2^{-1} \sqrt{\log(p/\varepsilon)/n}\} \leq c_1 \kappa_1(s, 1 + \lambda_2 + \lambda_\infty)$ for some small enough constant c_1 , then, with the same probability, we have

$$|\hat{\theta} - \theta^*|_q \leq C s^{1/q} \sqrt{\frac{\log(c'p/\varepsilon)}{n}} (|\theta^*|_2 + 1) + C s^{1/q} \bar{\delta}^2 |\theta^*|_\infty \quad (10)$$

for some constants $C > 0$ and $c' > 0$.

Proof. We proceed in three steps. Step 1 establishes initial relations and the fact that $\Delta = \hat{\theta} - \theta^*$ belongs to $C_J(1 + \lambda_2 + \lambda_\infty)$. Step 2 provides a bound on $|\frac{1}{n} X^T X \Delta|_\infty$. Step 3 establishes the rates of convergence stated in the theorem. We work on the event of probability at least $1 - 7\varepsilon$ where all the inequalities in Lemmas 1 and 2 are realized. Throughout the proof, $J = \{j : \theta_j^* \neq 0\}$. We often make use of the inequalities $|\theta|_\infty \leq |\theta|_2 \leq |\theta|_1$, $\forall \theta \in \mathbb{R}^p$.

Step 1. We first note that

$$\begin{aligned} |\frac{1}{n} Z^T (y - Z\theta^*)|_\infty &\leq |\frac{1}{n} Z^T \xi|_\infty + |\frac{1}{n} Z^T W \theta^*|_\infty \\ &\leq \delta_2(\varepsilon) + \delta_3(\varepsilon) + |\frac{1}{n} Z^T W \theta^*|_\infty \end{aligned} \quad (11)$$

with probability at least $1 - 2\varepsilon$ by Lemma 1. Next, Lemma 2 and the fact that, due to (1), we have $|\frac{1}{n} \text{Diag}(W^T W)|_\infty \leq \bar{\delta}^2$ imply

$$\begin{aligned} |\frac{1}{n} Z^T W \theta^*|_\infty &\leq |\frac{1}{n} X^T W \theta^*|_\infty + |\frac{1}{n} W^T W \theta^*|_\infty \\ &\leq |\frac{1}{n} X^T W \theta^*|_\infty + |\frac{1}{n} (W^T W - \text{Diag}(W^T W)) \theta^*|_\infty + |\frac{1}{n} \text{Diag}(W^T W) \theta^*|_\infty \\ &\leq \delta'_1(\varepsilon) |\theta^*|_2 + \delta'_4(\varepsilon) |\theta^*|_2 + \bar{\delta}^2 |\theta^*|_\infty. \end{aligned} \quad (12)$$

Combining (11) and (12) we get that $(\theta, t, u) = (\theta^*, |\theta^*|_2, |\theta^*|_\infty)$ is feasible for the problem (8), so that

$$|\hat{\theta}|_1 + \lambda_2 |\hat{\theta}|_2 + \lambda_\infty |\hat{\theta}|_\infty \leq |\hat{\theta}|_1 + \lambda_2 \hat{t} + \lambda_\infty \hat{u} \leq |\theta^*|_1 + \lambda_2 |\theta^*|_2 + \lambda_\infty |\theta^*|_\infty. \quad (13)$$

From (13) we easily obtain

$$|\hat{\theta}_{J^c}|_1 \leq (1 + \lambda_2 + \lambda_\infty) |\hat{\theta}_J - \theta^*|_1.$$

Arguments similar to (13) lead to

$$\hat{t} - |\theta^*|_2 \leq \frac{|\Delta|_1 + \lambda_\infty |\Delta|_\infty}{\lambda_2} \leq \frac{(1 + \lambda_\infty)}{\lambda_2} |\Delta|_1 \quad \text{and} \quad \hat{u} - |\theta^*|_\infty \leq \frac{|\Delta|_1 + \lambda_2 |\Delta|_2}{\lambda_\infty} \leq \frac{(1 + \lambda_2)}{\lambda_\infty} |\Delta|_1.$$

Step 2. We have

$$\begin{aligned} |\frac{1}{n} X^T X \Delta|_\infty &\leq |\frac{1}{n} Z^T X \Delta|_\infty + |\frac{1}{n} W^T X \Delta|_\infty \\ &\leq |\frac{1}{n} Z^T Z \Delta|_\infty + |\frac{1}{n} Z^T W \Delta|_\infty + |\frac{1}{n} W^T X \Delta|_\infty \\ &\leq |\frac{1}{n} Z^T (y - Z\theta^*)|_\infty + |\frac{1}{n} Z^T (y - Z\hat{\theta})|_\infty + |\frac{1}{n} Z^T W \Delta|_\infty + |\frac{1}{n} W^T X \Delta|_\infty. \end{aligned}$$

The results of Step 1 and of Lemmas 1 and 2 imply the following bounds

$$\begin{aligned} |\frac{1}{n} Z^T (y - Z\theta^*)|_\infty &\leq \tau_2 |\theta^*|_2 + \bar{\delta}^2 |\theta^*|_\infty + \tau, \\ |\frac{1}{n} Z^T (y - Z\hat{\theta})|_\infty &\leq \tau_2 \hat{t} + \bar{\delta}^2 \hat{u} + \tau \\ &\leq \tau_2 |\theta^*|_2 + \bar{\delta}^2 |\theta^*|_\infty + \tau + \{\tau_2(1 + \lambda_\infty)/\lambda_2 + \bar{\delta}^2(1 + \lambda_2)/\lambda_\infty\} |\Delta|_1, \\ |\frac{1}{n} W^T X \Delta|_\infty &\leq \delta_1 |\Delta|_1, \\ |\frac{1}{n} Z^T W \Delta|_\infty &\leq |\frac{1}{n} X^T W \Delta|_\infty + |\frac{1}{n} (W^T W - \text{Diag}(W^T W)) \Delta|_\infty + |\frac{1}{n} \text{Diag}(W^T W) \Delta|_\infty \\ &\leq \delta_1 |\Delta|_1 + \delta_4 |\Delta|_1 + \bar{\delta}^2 |\Delta|_\infty. \end{aligned}$$

These relations and the inequality $|\Delta|_\infty \leq |\Delta|_1$ yield that

$$|\frac{1}{n} X^T X \Delta|_\infty \leq 2\tau_2 |\theta^*|_2 + 2\bar{\delta}^2 |\theta^*|_\infty + 2\tau + (\bar{\delta}^2 \{(1 + \lambda_2 + \lambda_\infty)/\lambda_\infty\} + \{(1 + \lambda_\infty)/\lambda_2\} \tau_2 + 2\delta_1 + \delta_4) |\Delta|_1.$$

Step 3. Next note that $|\Delta|_1 \leq |\hat{\theta}|_1 + |\theta^*|_1 \leq (2 + \lambda_2 + \lambda_\infty)|\theta^*|_1$. Letting

$$\eta = \bar{\delta}^2 \{(1 + \lambda_2 + \lambda_\infty)/\lambda_\infty\} + \{(1 + \lambda_\infty)/\lambda_2\}\tau_2 + 2\delta_1 + \delta_4,$$

we have

$$|\frac{1}{n}X^T X \Delta|_\infty \leq 2\tau + (2\tau_2 + 2\bar{\delta}^2 + (2 + \lambda_2 + \lambda_\infty)\eta)|\theta^*|_1.$$

By the definition of the ℓ_q -sensitivity,

$$|\frac{1}{n}X^T X \Delta|_\infty \geq \kappa_q(s, 1 + \lambda_2 + \lambda_\infty)|\Delta|_q.$$

Now, (9) follows by combining the last two displays and the assumption on $\kappa_q(s, 1 + \lambda_2 + \lambda_\infty)$. To prove (10), we use that

$$\begin{aligned} |\frac{1}{n}X^T X \Delta|_\infty &\leq 2\tau_2|\theta^*|_2 + 2\bar{\delta}^2|\theta^*|_\infty + 2\tau + \eta|\Delta|_1 \\ &\leq 2\tau_2|\theta^*|_2 + 2\bar{\delta}^2|\theta^*|_\infty + 2\tau + \eta|\frac{1}{n}X^T X \Delta|_\infty / \kappa_1(s, 1 + \lambda_2 + \lambda_\infty). \end{aligned}$$

Under the condition that $(1 + \lambda_2 + \lambda_\infty)\{\bar{\delta}^2\lambda_\infty^{-1} + \lambda_2^{-1}\sqrt{\log(p/\varepsilon)/n}\} \leq c_1\kappa_1(s, 1 + \lambda_2 + \lambda_\infty)$ for c_1 small enough, by definition of η we have $\eta/\kappa_1(s, 1 + \lambda_2 + \lambda_\infty) \leq c'$ for some $0 < c' < 1$. Thus, we have

$$|\frac{1}{n}X^T X \Delta|_\infty \leq c(\tau_2|\theta^*|_2 + \bar{\delta}^2|\theta^*|_\infty + \tau),$$

which implies (10) in view of the definition of the ℓ_q -sensitivity and the assumption on $\kappa_q(s, 1 + \lambda_2 + \lambda_\infty)$. \square

Remark 2 (Relaxation of Assumption (A4)). We have stated Theorem 1 under Assumption (A4) to make the analysis streamlined with the previous literature, see [8]. However, inspection of the proofs shows that a more general condition can be used. The results of Theorem 1 hold with probability at least $1 - 7\varepsilon - \varepsilon'$ if instead of Assumption (A4) we require W to satisfy:

$$|\frac{1}{n}\text{Diag}(W^T W)|_\infty \leq \bar{\delta}^2,$$

with probability at least $1 - \varepsilon'$, for some $\varepsilon' > 0$.

Compared to [8], the results in Theorem 1 exploit the zero-mean condition on the noise matrix W . As in [8], the estimator is consistent as $\bar{\delta}$ goes to zero. In order to compare the rates in Theorem 1 with those for the MU selector, we recall that, by Theorem 3 in [8], the MU selector satisfies

$$|\hat{\theta}^{MU} - \theta^*|_q \leq C s^{1/q} \sqrt{\frac{\log(c'p/\varepsilon)}{n}} + C s^{1/q} (\bar{\delta} + \bar{\delta}^2) |\theta^*|_1$$

with probability close to 1. While both rates share some terms, a term of the order $s^{1/q}\bar{\delta}|\theta^*|_1$ appears only in the rate for the MU selector whereas a term of the order $s^{1/q}\sqrt{\log(c'p/\varepsilon)/n}|\theta^*|_1$ appears only for the $\{\ell_1, \ell_2, \ell_\infty\}$ -MU selector. Therefore, the improvement upon the original MU selector can be achieved when $\bar{\delta} \gg \sqrt{\log(c'p/\varepsilon)/n}$.

If the additional condition in the second part of Theorem 1 holds, we can use the bound (10) and a better accuracy is achieved by the proposed estimator. In particular, $|\theta^*|_1$ no longer drives the rate of convergence. The impact of $\bar{\delta}$ on this rate is in the term

$$s^{1/q}\bar{\delta}^2|\theta^*|_\infty \quad \text{instead of} \quad s^{1/q}(\bar{\delta} + \bar{\delta}^2)|\theta^*|_1 \quad (14)$$

for the MU selector. Furthermore, the rate of convergence of the new estimator also has a term of the form $|\theta^*|_2 s^{1/q}\sqrt{\log(c'p/\varepsilon)/n}$. Thus the new estimator obtains a better accuracy by exploiting additional assumptions together with the fact that $\bar{\delta}|\theta^*|_1$ is of larger order than $\sqrt{\log(c'p/\varepsilon)/n}|\theta^*|_2$,

which holds whenever $\bar{\delta} \gg \sqrt{\log(c'p/\varepsilon)/n}$. Finally, the impact of going from the ℓ_1 -norm to the ℓ_2 - or ℓ_∞ -norms is not negligible neither. For example, if all non-zero components of θ^* are equal to the same constant $a > 0$, we have $|\theta^*|_1 = sa$ while $|\theta^*|_2 = a\sqrt{s}$ and $|\theta^*|_\infty = a$. Then, the comparison in (14) is reduced to comparing

$$s^{1/q}\bar{\delta}^2 \quad \text{versus} \quad s^{1+1/q}(\bar{\delta} + \bar{\delta}^2),$$

featuring the maximum contrast between the two rates.

Finally, note that the conic estimator $\hat{\theta}^{Conic}$ studied in [1] can be also applied under the assumptions of this section. Indeed, our setting can be embedded into that of [1] with \hat{D} being the identically zero $p \times p$ matrix, which means that we have an estimator of each σ_j^2 with an error bounded by $b = \bar{\delta}^2$. The results in [1] assume $b \leq C\sqrt{(\log p)/n}$ but they do not apply to designs with b of larger order. Comparing the bound (10) in Theorem 1 to the bound (6) yields that the conic estimator $\hat{\theta}^{Conic}$ achieves the same rate as our new estimator whenever $\bar{\delta}$ is smaller than or of the order $\{(\log p)/n\}^{1/4}$. However, there is no bound for $\hat{\theta}^{Conic}$ available when $\bar{\delta} \gg \{(\log p)/n\}^{1/4}$.

4. $\{\ell_1, \ell_2, \ell_\infty\}$ -compensated MU selector

In this section, we discuss a modification of the conic estimator proposed in [1]. We introduce an additional ℓ_∞ -norm regularization to better adapt to the estimation error in \hat{D} . As discussed in the introduction, this is beneficial when the rate of convergence of \hat{D} to D is slower than $\sqrt{(\log p)/n}$, which is not covered by [1]. Here we consider the same assumptions as in [1] with the only difference that now we allow for any rate of convergence of \hat{D} to D . Thus, we replace Assumption (A4) by the following assumption on the availability of estimators for σ_j^2 , $j = 1, \dots, p$:

(A5) *There exist statistics $\hat{\sigma}_j^2$ and positive numbers $b(\varepsilon)$ such that for any $0 < \varepsilon < 1$, we have*

$$\mathbb{P} \left(\max_{j=1, \dots, p} |\hat{\sigma}_j^2 - \sigma_j^2| \geq b(\varepsilon) \right) \leq \varepsilon.$$

In what follows, we fix ε and set

$$\tau_2 = \delta'_1(\varepsilon) + \delta'_4(\varepsilon), \quad \tau = \delta_2(\varepsilon) + \delta_3(\varepsilon) \quad \text{and} \quad \tau_\infty = b(\varepsilon) + \delta_5(\varepsilon).$$

We are particularly interested in cases where τ_∞ is of larger order than $\sqrt{(\log p)/n}$. To define the estimator, we consider the following minimization problem:

$$\min_{\theta, t, u} \{ |\theta|_1 + \lambda_2 t + \lambda_\infty u : (\theta, t, u) \in \Theta, \left| \frac{1}{n} Z^T (y - Z\theta) + \hat{D}\theta \right|_\infty \leq \tau_2 t + \tau_\infty u + \tau, |\theta|_2 \leq t, |\theta|_\infty \leq u \}, \quad (15)$$

where $\lambda_2 > 0$ and $\lambda_\infty > 0$ are tuning constants and we allow for $(\theta, t, u) \in \Theta$ where Θ is a pre-specified convex set that contains $(\theta^*, |\theta^*|_2, |\theta^*|_\infty)$, see Remark 1.

Let $(\hat{\theta}, \hat{t}, \hat{u})$ be a solution of (15). We take $\hat{\theta}$ as estimator of θ^* and we call it the $\{\ell_1, \ell_2, \ell_\infty\}$ -compensated MU selector. The rates of convergence of this estimator are given in the next theorem.

Theorem 2. *Let Assumptions (A1)-(A3), and (A5) hold. Assume that the true parameter θ^* is s -sparse and $(\theta^*, |\theta^*|_2, |\theta^*|_\infty)$ belongs to Θ . Let $0 < \varepsilon < 1$ and $1 \leq q \leq \infty$. Suppose also that*

$$\kappa_q(s, 1 + \lambda_2 + \lambda_\infty) \geq cs^{-1/q} \quad (16)$$

for some constant $c > 0$ and that

$$s \leq c_1 \min \left\{ \frac{\lambda_2 \lambda_\infty}{(1 + \lambda_2 + \lambda_\infty)^3} \sqrt{\frac{n}{\log(p/\varepsilon)}}, \frac{\lambda_\infty}{(1 + \lambda_2 + \lambda_\infty)^2} \frac{1}{b(\varepsilon)} \right\}, \quad (17)$$

for some small enough constant $c_1 > 0$. Let $\hat{\theta}$ be the $\{\ell_1, \ell_2, \ell_\infty\}$ -compensated MU selector. Then, with probability at least $1 - 8\varepsilon$,

$$|\hat{\theta} - \theta^*|_q \leq C s^{1/q} \sqrt{\frac{\log(c'p/\varepsilon)}{n}} (|\theta^*|_2 + 1) + C s^{1/q} b(\varepsilon) |\theta^*|_\infty, \quad (18)$$

for some constants $C > 0$ and $c' > 0$ (here we set $s^{1/\infty} = 1$).

Under the same assumptions with $q = 1$, the prediction error admits the following bound, with the same probability:

$$\frac{1}{n} |X(\hat{\theta} - \theta^*)|_2^2 \leq C s \frac{\log(c'p/\varepsilon)}{n} (|\theta^*|_2 + 1)^2 + C s b^2(\varepsilon) |\theta^*|_\infty^2. \quad (19)$$

Proof. Throughout the proof, we assume that we are on the event of probability at least $1 - 8\varepsilon$ where the results of Lemmas 3, 4 and 5 in the Appendix hold. Property (26) in Lemma 4 implies that $\Delta = \hat{\theta} - \theta^*$ is in the cone $C_J(1 + \lambda_2 + \lambda_\infty)$, where $J = \{j : \theta_j^* \neq 0\}$. Therefore, by the definition of the ℓ_q -sensitivity and Lemma 5, we have

$$\kappa_q(s, 1 + \lambda_2 + \lambda_\infty) |\Delta|_q \leq \left| \frac{1}{n} X^T X \Delta \right|_\infty \leq \mu_0 + \mu_1 |\hat{\theta} - \theta^*|_1 + \mu_2 |\theta^*|_2 + \mu_\infty |\theta^*|_\infty,$$

where μ_0 and μ_2 are of the order $\sqrt{\frac{1}{n} \log(c'p/\varepsilon)}$, and μ_1 and μ_∞ are of the order $\sqrt{\frac{1}{n} \log(c'p/\varepsilon) + b(\varepsilon)}$. Using again (26), we have

$$\begin{aligned} |\Delta|_1 &= |\Delta_{J^c}|_1 + |\Delta_J|_1 \leq (2 + \lambda_2 + \lambda_\infty) |\Delta_J|_1 \\ &\leq (2 + \lambda_2 + \lambda_\infty) s^{1-1/q} |\Delta_J|_q \leq (2 + \lambda_2 + \lambda_\infty) s^{1-1/q} |\Delta|_q. \end{aligned}$$

It follows that

$$(\kappa_q(s, 1 + \lambda_2 + \lambda_\infty) - (2 + \lambda_2 + \lambda_\infty) \mu_1 s^{1-1/q}) |\Delta|_q \leq \mu_0 + \mu_2 |\theta^*|_2 + \mu_\infty |\theta^*|_\infty,$$

which implies, by (16),

$$(c - (2 + \lambda_2 + \lambda_\infty) \mu_1 s) s^{-1/q} |\Delta|_q \leq \mu_0 + \mu_2 |\theta^*|_2 + \mu_\infty |\theta^*|_\infty,$$

in view of the assumptions of the theorem. Recall that

$$\mu_1 \leq a \sqrt{\log(c'p/\varepsilon)/n} \{1 + (1 + \lambda_2) \lambda_\infty^{-1} + (1 + \lambda_\infty) \lambda_2^{-1}\} + ab(\varepsilon) \{1 + \lambda_\infty + \lambda_2\} / \lambda_\infty,$$

where $a > 0$ is a constant. Therefore, since we assume (17), and $\{1 + (1 + \lambda_2) \lambda_\infty^{-1} + (1 + \lambda_\infty) \lambda_2^{-1}\} \leq (1 + \lambda_\infty + \lambda_2)^2 / (\lambda_\infty \lambda_2)$, relation (18) follows if c_1 is small enough.

To prove (19), write first

$$\frac{1}{n} |X \Delta|_2^2 \leq \frac{1}{n} |X^T X \Delta|_\infty |\Delta|_1.$$

Next remark that from (18) with $q = 1$, we have

$$|\Delta|_1 \leq C s \sqrt{\frac{\log(c'p/\varepsilon)}{n}} (|\theta^*|_2 + 1) + C s b(\varepsilon) |\theta^*|_\infty.$$

Lemma 5 in the Appendix yields

$$\left| \frac{1}{n} X^T X \Delta \right|_\infty \leq \mu_0 + \mu_1 |\hat{\theta} - \theta^*|_1 + \mu_2 |\theta^*|_2 + \mu_\infty |\theta^*|_\infty. \quad (20)$$

Combining the above bound for $|\Delta|_1$ and (20), we get

$$\frac{1}{n} |X\Delta|_2^2 \leq C \frac{s \log(c'p/\varepsilon)}{n} (|\theta^*|_2 + 1)^2 + Csb^2(\varepsilon)|\theta^*|_\infty^2$$

since $\mu_1 s \leq C''$ for some constant $C'' > 0$ under our assumptions. This proves (19). \square

Theorem 2 generalizes the results in [1] to estimators \hat{D} that converge with rate $b(\varepsilon)$ of larger order than $\sqrt{(\log p)/n}$. At the same time, if $b(\varepsilon)$ is smaller than $\sqrt{(\log p)/n}$, both the conic estimator $\hat{\theta}^{Conic}$ of [1] and the $\{\ell_1, \ell_2, \ell_\infty\}$ -compensated MU selector achieve the same rate of convergence.

For such designs that condition (17) does not hold, the conclusions of Theorem 2 need to be slightly modified as shown in the next theorem.

Theorem 3. *Let Assumptions (A1)-(A3), and (A5) hold. Assume that the true parameter θ^* is s -sparse and $(\theta^*, |\theta^*|_2, |\theta^*|_\infty)$ belongs to Θ . Let $0 < \varepsilon < 1$ and $1 \leq q \leq \infty$. Let $\hat{\theta}$ be the $\{\ell_1, \ell_2, \ell_\infty\}$ -compensated MU selector. Then, with probability at least $1 - 8\varepsilon$,*

$$|\hat{\theta} - \theta^*|_q \leq \frac{C}{\kappa_q(s, 1 + \lambda_2 + \lambda_\infty)} \left\{ \sqrt{\frac{\log(c'p/\varepsilon)}{n}} (|\theta^*|_1 + 1) + b(\varepsilon)|\theta^*|_1 \right\},$$

for some constants $C > 0$ and $c' > 0$.

Proof. Again, throughout the proof, we assume that we are on the event of probability at least $1 - 8\varepsilon$ where the results of Lemmas 3, 4 and 5 in the Appendix hold. Property (26) in Lemma 4 implies that $\Delta = \hat{\theta} - \theta^*$ is in the cone $C_J(1 + \lambda_2 + \lambda_\infty)$, where $J = \{j : \theta_j^* \neq 0\}$. Since

$$|\Delta|_1 \leq |\hat{\theta}|_1 + |\theta^*|_1 \leq \{|\theta^*|_1 + \lambda_2|\theta^*|_2 + \lambda_\infty|\theta^*|_\infty\} + |\theta^*|_1 \leq (2 + \lambda_2 + \lambda_\infty)|\theta^*|_1,$$

we obtain

$$\left| \frac{1}{n} X^T X \Delta \right|_\infty \leq \mu_0 + \mu_1 |\Delta|_1 + \mu_2 |\theta^*|_2 + \mu_\infty |\theta^*|_\infty \leq \mu_0 + (\mu_1 + \mu_2 + \mu_\infty)(2 + \lambda_2 + \lambda_\infty) |\theta^*|_1.$$

Therefore

$$\kappa_q(s, 1 + \lambda_2 + \lambda_\infty) |\Delta|_q \leq \mu_0 + (\mu_1 + \mu_2 + \mu_\infty)(2 + \lambda_2 + \lambda_\infty) |\theta^*|_1,$$

which implies the result. \square

The result in Theorem 3 parallels the result for generic designs for the conic estimator [1]. Indeed, this result states that the additional ℓ_∞ -regularization does not worsen the guarantees obtained in [1]. For generic designs, our bounds do not achieve the previous dependence on $|\theta^*|_2$ and $|\theta^*|_\infty$ and, instead, the final dependence is on $|\theta^*|_1$.

5. Simulations

This section aims to illustrate the finite sample performance of the proposed estimators. We will focus on the $\{\ell_1, \ell_2, \ell_\infty\}$ -compensated MU selector only. We consider the following data generating process

$$y_i = x_i^T \theta^* + \xi_i, \quad z_i = x_i + w_i.$$

Here, ξ_i, w_i, x_i are independent and $\xi_i \sim \mathcal{N}(0, \sigma^2)$, $w_i \sim \mathcal{N}(0, \sigma_*^2 I_{p \times p})$, $x_i \sim \mathcal{N}(0, \Sigma)$ where $I_{p \times p}$ is the identity matrix and Σ is $p \times p$ matrix with elements $\Sigma_{ij} = \rho^{|i-j|}$. We consider the vector of unknown

parameters $\theta^* = 1.25(1, 1, 1, 1, 1, 0, \dots, 0)^T$. We set $\sigma = 0.128$, $\sigma_*^2 = 0.5$, and $\rho = 0.25$. We assume that σ is known and we set $\widehat{D} = D = \sigma_*^2 I_{p \times p}$. The penalty parameters are set as $\tau = \sigma \sqrt{\log(p/\varepsilon)/n}$, $b(\varepsilon) = \sigma_*^2 \sqrt{\log(p/\varepsilon)/n}$, for $\varepsilon = 0.05$.

In our first set of simulations, we illustrate the finite sample performance of the proposed estimator by setting $\lambda_2 = \lambda_\infty \in \{0.25, 0.5, 0.75, 1\}$. The $\{\ell_1, \ell_2, \ell_\infty\}$ -compensated MU selector will be denoted by $\{\ell_1, \ell_2, \ell_\infty\}$. We compare its performance with other recent proposals in the literature, namely the conic estimator (denoted as Conic (λ_2) for $\lambda_2 = 0.25, 0.5, 0.75, 1$), and the Compensated MU selector (cMU). We also provide the (infeasible) Dantzig selector which knows X (Dantzig X) and the Dantzig selector that uses only Z (Dantzig Z) as additional benchmark for the performance.

Method ($\lambda_2 = \lambda_\infty$)	$n = 300$ and $p = 10$			$n = 300$ and $p = 50$		
	Bias	RMSE	PR	Bias	RMSE	PR
Dantzig X	0.0265486	0.0321528	0.0349530	0.0301636	0.0349420	0.0386731
Dantzig Z	0.5892699	0.6218173	0.7118256	0.6032541	0.7246990	0.7526539
cMU	0.6002801	0.6526144	0.7375240	0.6684987	0.7074681	0.8148175
Conic (.25)	1.9261733	1.9567318	2.3165088	1.9952936	2.0190105	2.4085353
$\{\ell_1, \ell_2, \ell_\infty\}$ (.25)	1.7922416	1.8349666	2.1453927	1.9035308	1.9325326	2.2875796
Conic (.5)	0.3184083	0.4161670	0.4326569	0.3668194	0.4395404	0.4781078
$\{\ell_1, \ell_2, \ell_\infty\}$ (.5)	0.2137347	0.3505829	0.3382480	0.3489980	0.4491837	0.4605638
Conic (.75)	0.3179691	0.4158134	0.4322128	0.3668194	0.4395404	0.4781078
$\{\ell_1, \ell_2, \ell_\infty\}$ (.75)	0.2085334	0.3459298	0.3330411	0.2699453	0.3786945	0.3896168
Conic (1)	0.3179691	0.4158134	0.4322128	0.3661721	0.4390614	0.4773173
$\{\ell_1, \ell_2, \ell_\infty\}$ (1)	0.2078373	0.3455287	0.3324356	0.2483137	0.3691060	0.3736929

TABLE 1

Simulation results for 100 replications. For each estimator we provide average bias (Bias), average root-mean squared error (RMSE), and average prediction risk (PR).

Method ($\lambda_2 = \lambda_\infty$)	$n = 300$ and $p = 100$			$n = 300$ and $p = 300$		
	Bias	RMSE	PR	Bias	RMSE	PR
Dantzig X	0.0317776	0.0366155	0.0403419	0.0344617	0.0387848	0.0436396
Dantzig Z	0.6039890	0.8364059	0.7910512	0.6334052	1.0775665	0.8824695
cMU	0.6908240	0.7359536	0.8472447	0.7228791	0.7653174	0.8843476
Conic (.25)	2.0196204	2.0428152	2.4429977	2.0833543	2.0985979	2.5281871
$\{\ell_1, \ell_2, \ell_\infty\}$ (.25)	1.9363225	1.9646153	2.3321286	2.0016163	2.0247679	2.4181903
Conic (.5)	0.5032353	0.6479385	0.6390150	0.6809176	0.8886359	0.8367831
$\{\ell_1, \ell_2, \ell_\infty\}$ (.5)	0.4170439	0.5207218	0.5436218	0.4694103	0.5507253	0.5975351
Conic (.75)	0.3849631	0.4699933	0.5082582	0.4195124	0.4964321	0.5428568
$\{\ell_1, \ell_2, \ell_\infty\}$ (.75)	0.3250997	0.4312186	0.4512656	0.3869566	0.4747343	0.5104562
Conic(1)	0.3811186	0.4673239	0.5043246	0.4047078	0.4846393	0.5271225
$\{\ell_1, \ell_2, \ell_\infty\}$ (1)	0.2907918	0.4155573	0.4242000	0.3573025	0.4569624	0.4819208

TABLE 2

Simulation results for 100 replications. For each estimator we provide average bias (Bias), average root-mean squared error (RMSE), and average prediction risk (PR).

Tables 1 and 2 provide the performance of the proposed estimator when $\lambda_2 = \lambda_\infty$ and the performance of various benchmarks. As discussed in the literature, ignoring the error-in-variables issue can lead to worse performance as seen from the performance of Dantzig Z compared to the (infeasible) Dantzig X. The conic compensated estimator performs better than the compensated MU selector (cMU) when $\lambda_2 \in \{0.5, 0.75, 1\}$. The comparison of the proposed estimator and the conic estimator is easier to establish as we can parametrize them by λ_2 (as we set $\lambda_2 = \lambda_\infty$). In this case the conic estimator penalizes more aggressively the uncertainty of not knowing σ_j^2 . In essentially all cases¹ the proposed estimator yields improvements. The introduction of ℓ_∞ -norm regularization seems to alleviate regularization bias. Nonetheless, when setting $\lambda_2 = 0.25$ both the conic estimator and the proposed

¹The conic compensated estimator performs slightly better only with respect to RMSE in the case of $\lambda_2 = 0.5$. For all other parameters and metrics, the proposed estimator performed slightly better or substantially better.

estimator fail in the experiment. This failure occurs by not having enough penalty to control $t - |\theta|_2$ and $u - |\theta|_\infty$ which leads to a large right hand side $\tau_2 t + \tau_\infty u + \tau$ in the constraint

$$\left| \frac{1}{n} Z^T (y - Z\theta) + \widehat{D}\theta \right|_\infty \leq \tau_2 t + \tau_\infty u + \tau$$

in (15) and similarly the right hand side $\tau_2 t + \tau$ in (5). In turn, this leads to substantial regularization bias and therefore underfitting. In fact, detailed inspection of estimators in that case reveals that coefficients are very close to zero for both conic and the proposed estimator.

In the second set of simulations, we explore the performance of the proposed estimator for the case $\lambda_2 \neq \lambda_\infty$. Moreover, we also study a modified estimator that contains safeguard constraints. These constraints aim to mitigate the problem discussed above. The safeguard constraints are described in Remark 1. We denote by $\{\ell_1, \ell_2, \ell_\infty\}^*$ the estimator computed with the safeguards.

We consider the same design as before and we explore some combinations of values $(\lambda_2, \lambda_\infty) \in \{0.25, 0.5, 0.75, 1\} \times \{0.25, 0.5, 0.75, 1\}$ for both proposed estimators (with and without the safeguard constraints).

Method $(\lambda_2, \lambda_\infty)$	$n = 300$ and $p = 10$			$n = 300$ and $p = 50$		
	Bias	RMSE	PR	Bias	RMSE	PR
$\{\ell_1, \ell_2, \ell_\infty\} (1,1)$	0.2078373	0.3455287	0.3324356	0.2483137	0.3691060	0.3736929
$\{\ell_1, \ell_2, \ell_\infty\}^* (1,1)$	0.2078373	0.3455287	0.3324356	0.2483137	0.3691060	0.3736929
$\{\ell_1, \ell_2, \ell_\infty\} (1,.5)$	0.2534465	0.3997941	0.3725479	0.5214272	0.7086267	0.6514348
$\{\ell_1, \ell_2, \ell_\infty\}^* (1,.5)$	0.2392491	0.3623416	0.3569492	0.3980543	0.4729990	0.5112310
$\{\ell_1, \ell_2, \ell_\infty\} (.5,1)$	0.2077228	0.3455690	0.3322088	0.2448911	0.3690180	0.3723095
$\{\ell_1, \ell_2, \ell_\infty\}^* (.5,1)$	0.2077228	0.3455690	0.3322088	0.2448911	0.3690180	0.3723095
$\{\ell_1, \ell_2, \ell_\infty\} (.75,.75)$	0.2085334	0.3459298	0.3330411	0.2699453	0.3786945	0.3896168
$\{\ell_1, \ell_2, \ell_\infty\}^* (.75,.75)$	0.2085334	0.3459297	0.3330411	0.2699453	0.3786945	0.3896168
$\{\ell_1, \ell_2, \ell_\infty\} (.25,1)$	0.2078663	0.3458796	0.3322444	0.2439496	0.3684173	0.3715836
$\{\ell_1, \ell_2, \ell_\infty\}^* (.25,1)$	0.2078663	0.3458796	0.3322444	0.2439496	0.3684173	0.3715836
$\{\ell_1, \ell_2, \ell_\infty\} (.5,.5)$	0.2137347	0.3505829	0.3382480	0.3489980	0.4491837	0.4605638
$\{\ell_1, \ell_2, \ell_\infty\}^* (.5,.5)$	0.2137347	0.3505827	0.3382479	0.3382218	0.4225007	0.4490958
$\{\ell_1, \ell_2, \ell_\infty\} (.25,.5)$	0.2114159	0.3502938	0.3369809	0.3188151	0.4086438	0.4313163
$\{\ell_1, \ell_2, \ell_\infty\}^* (.25,.5)$	0.2114159	0.3502938	0.3369809	0.3188151	0.4086438	0.4313163
$\{\ell_1, \ell_2, \ell_\infty\} (.25,.25)$	1.7922416	1.8349666	2.1453927	1.9035308	1.9325326	2.2875796
$\{\ell_1, \ell_2, \ell_\infty\}^* (.25,.25)$	0.5477221	0.6050091	0.6780050	0.6151622	0.6574460	0.7535672

TABLE 3

Simulation results for 100 replications. For each estimator we provide average bias (Bias), average root-mean squared error (RMSE), and average prediction risk (PR).

Tables 3 and 4 show the performance for different values of λ_2 and λ_∞ . We note that these parameters seem to have different impact on the finite sample performance even if $\lambda_2 + \lambda_\infty$ is kept constant. Importantly, we observe that the addition of safeguard constraints virtually always leads to improvements although small (even zero sometimes) for most of the tested parameter values. In the case $\lambda_2 < \lambda_\infty$ using safeguard constraints makes almost no difference and overall performance of both estimators is better. In contrast, the estimators perform worse when $\lambda_2 > \lambda_\infty$ and the safeguard constraints lead to improvements. Finally, as expected, the safeguard constraints improve substantially the performance when $\lambda_2 = \lambda_\infty = 0.25$. In that case, the performance becomes comparable to that of the cMU estimator. Essentially, the safeguard constraints help to avoid severe underfitting. They are very helpful when the performance is below of what can be achieved. Nonetheless, we recommend to keep them in all cases as it does not impact negatively the estimator and the additional computational burden seems minimal.

In our third set of simulations, we consider a situation with a “crude” estimator of D . The design is the same as in our first set of simulations, however, in order to compute the estimator of D , we have independent observations of $(X_i, Z_i), i = 1, \dots, n^*$ where $n^* \leq n$. We consider $p \in \{50, 100\}$ and

Method ($\lambda_2, \lambda_\infty$)	$n = 300$ and $p = 100$			$n = 300$ and $p = 300$		
	Bias	RMSE	PR	Bias	RMSE	PR
$\{\ell_1, \ell_2, \ell_\infty\} (1,1)$	0.2907918	0.4155573	0.4242000	0.3573025	0.4569624	0.4819208
$\{\ell_1, \ell_2, \ell_\infty\}^* (1,1)$	0.2907918	0.4155573	0.4242000	0.3573084	0.4569653	0.4819268
$\{\ell_1, \ell_2, \ell_\infty\} (1,.5)$	0.6707248	0.8687948	0.8260765	1.0995021	1.2843061	1.3224733
$\{\ell_1, \ell_2, \ell_\infty\}^* (1,.5)$	0.4713115	0.5469680	0.5998890	0.5813572	0.6440090	0.7214057
$\{\ell_1, \ell_2, \ell_\infty\} (.5,1)$	0.2813842	0.4123716	0.4183949	0.3434854	0.4501870	0.4705890
$\{\ell_1, \ell_2, \ell_\infty\}^* (.5,1)$	0.2813842	0.4123716	0.4183949	0.3434113	0.4501763	0.4705578
$\{\ell_1, \ell_2, \ell_\infty\} (.75,.75)$	0.3250997	0.4312186	0.4512656	0.3869566	0.4747343	0.5104562
$\{\ell_1, \ell_2, \ell_\infty\}^* (.75,.75)$	0.3250997	0.4312186	0.4512656	0.3869382	0.4747230	0.5104387
$\{\ell_1, \ell_2, \ell_\infty\} (.25,1)$	0.2791982	0.4113275	0.4166136	0.3392525	0.4482318	0.4674070
$\{\ell_1, \ell_2, \ell_\infty\}^* (.25,1)$	0.2790174	0.4111578	0.4163974	0.3386830	0.4478506	0.4667958
$\{\ell_1, \ell_2, \ell_\infty\} (.5,.5)$	0.4170439	0.5207218	0.5436218	0.4694103	0.5507253	0.5975351
$\{\ell_1, \ell_2, \ell_\infty\}^* (.5,.5)$	0.3977208	0.4819926	0.5218939	0.4590288	0.5311414	0.5863702
$\{\ell_1, \ell_2, \ell_\infty\} (.25,.5)$	0.3726889	0.4645454	0.4983916	0.4324230	0.5082528	0.5569916
$\{\ell_1, \ell_2, \ell_\infty\}^* (.25,.5)$	0.3718829	0.4639374	0.497499	0.4357035	0.5115011	0.5608083
$\{\ell_1, \ell_2, \ell_\infty\} (.25,.25)$	1.9363225	1.9646153	2.3321286	2.0016163	2.0247679	2.4181903
$\{\ell_1, \ell_2, \ell_\infty\}^* (.25,.25)$	0.6365329	0.6854681	0.7851421	0.6669880	0.7127657	0.8198798

TABLE 4

Simulation results for 100 replications. For each estimator we provide average bias (Bias), average root-mean squared error (RMSE), and average prediction risk (PR).

$n^* = \{100, 250, 500\}$ and $n = 500$. Table 5 shows how the performance of the estimators change as less precise estimates of D are used. For example, consider the average prediction risk (PR) for $p = 50$ as n^* goes from 500 to 100 so that the quality of the estimator worsens. Table 5 shows that the cMU deteriorates by 0.51 (from 1.23 to 0.72), the conic deteriorates by 0.41 (from 0.82 to 0.41), and the proposed estimator deteriorates by 0.24 (from 0.55 to 0.31). Similarly when $p = 100$, we observe that cMU deteriorates by 0.52, the conic deteriorates by 0.47, and the proposed estimator deteriorates by 0.35. In both cases the findings are aligned with the theoretical results that the $\{\ell_1, \ell_2, \ell_\infty\}$ -based estimator is less sensitive to the use of a crude estimator for D .

Method	$n = 500, n^* = 100, p = 50$			$n = 500, n^* = 100, p = 100$		
	Bias	RMSE	PR	Bias	RMSE	PR
cMU	1.0175038	1.0400870	1.2283635	1.0641008	1.0909776	1.2845535
Conic (1)	0.6660582	0.7467607	0.8159505	0.7490445	0.8519112	0.9150378
$\{\ell_1, \ell_2, \ell_\infty\} (.5,1)$	0.4275867	0.5005111	0.5477009	0.5451567	0.6207584	0.6824172
Method	$n = 500, n^* = 250, p = 50$			$n = 500, n^* = 250, p = 100$		
	Bias	RMSE	PR	Bias	RMSE	PR
cMU	0.7520123	0.7795984	0.9150986	0.7884387	0.8152334	0.9568734
Conic (1)	0.4173097	0.4716240	0.5327973	0.4446488	0.4962641	0.5623790
$\{\ell_1, \ell_2, \ell_\infty\} (.5,1)$	0.2615103	0.3503973	0.3697514	0.2847444	0.3699376	0.3937119
Method	$n = 500, n^* = 500, p = 50$			$n = 500, n^* = 500, p = 100$		
	Bias	RMSE	PR	Bias	RMSE	PR
cMU	0.5922077	0.6226468	0.7269638	0.6258117	0.6564809	0.7636137
Conic (1)	0.3137644	0.3762319	0.4162300	0.3391774	0.3977718	0.4393967
$\{\ell_1, \ell_2, \ell_\infty\} (.5,1)$	0.1966452	0.3024311	0.3091859	0.2219782	0.3219536	0.3296119

TABLE 5

Simulation results for 100 replications. For each estimator we provide average bias (Bias), average root-mean squared error (RMSE), and average prediction risk (PR).

Appendix: Auxiliary lemmas

In what follows, we write for brevity $\delta_i = \delta_i(\varepsilon)$, $\delta'_i = \delta'_i(\varepsilon)$, and we set $\Delta = \hat{\theta} - \theta^*$, $J = \{j : \theta_j^* \neq 0\}$.

Lemma 3. Assume (A1)-(A3) and (A5). Then with probability at least $1 - 6\varepsilon$, the pair $(\theta, t, u) = (\theta^*, |\theta^*|_2, |\theta^*|_\infty)$ belongs to the feasible set of the minimization problem (15).

Proof. First, note that $Z^T(y - Z\theta^*) + n\widehat{D}\theta^*$ is equal to

$$\begin{aligned} & -X^T W\theta^* + X^T \xi + W^T \xi - (W^T W - \text{Diag}\{W^T W\})\theta^* \\ & - (\text{Diag}\{W^T W\} - nD)\theta^* + n(\widehat{D} - D)\theta^*. \end{aligned}$$

By definition of δ_i and b , with probability at least $1 - 4\varepsilon$, we have

$$|\frac{1}{n}X^T \xi|_\infty + |\frac{1}{n}W^T \xi|_\infty \leq \delta_2 + \delta_3 \quad (21)$$

$$|(\frac{1}{n}\text{Diag}\{W^T W\} - D)\theta^*|_\infty \leq |\frac{1}{n}\text{Diag}\{W^T W\} - D|_\infty |\theta^*|_\infty \leq \delta_5 |\theta^*|_\infty \quad (22)$$

$$|(\widehat{D} - D)\theta^*|_\infty \leq b(\varepsilon) |\theta^*|_\infty, \quad (23)$$

where in (22) and (23) we have used that the considered matrices are diagonal. Also, by Lemma 2, with probability at least $1 - 2\varepsilon$, we have

$$|\frac{1}{n}X^T W\theta^*|_\infty \leq \delta'_1 |\theta^*|_2 \quad (24)$$

$$|\frac{1}{n}(W^T W - \text{Diag}\{W^T W\})\theta^*|_\infty \leq \delta'_4 |\theta^*|_2. \quad (25)$$

Combining the decomposition of $Z^T(y - Z\theta^*) + n\widehat{D}\theta^*$ together with (21)-(25), we find that

$$|\frac{1}{n}Z^T(y - Z\theta^*) + \widehat{D}\theta^*|_\infty \leq \tau_2 |\theta^*|_2 + \tau_\infty |\theta^*|_\infty + \tau,$$

with probability at least $1 - 6\varepsilon$, which implies the lemma. \square

Lemma 4. *Let $\hat{\theta}$ be the $\{\ell_1, \ell_2, \ell_\infty\}$ -compensated MU-selector. Assume (A1)-(A3) and (A5). Then with probability at least $1 - 6\varepsilon$ (on the same event as in Lemma 3), we have*

$$|(\hat{\theta} - \theta^*)_{J^c}|_1 \leq (1 + \lambda_2 + \lambda_\infty) |(\hat{\theta} - \theta^*)_J|_1, \quad (26)$$

$$\hat{t} - |\theta^*|_2 \leq \{(1 + \lambda_\infty)/\lambda_2\} |\hat{\theta} - \theta^*|_1 \quad \text{and} \quad \hat{u} - |\theta^*|_\infty \leq \{(1 + \lambda_2)/\lambda_\infty\} |\hat{\theta} - \theta^*|_1. \quad (27)$$

Proof. Set $\Delta = \hat{\theta} - \theta^*$. On the event of Lemma 3, $(\theta^*, |\theta^*|_2, |\theta^*|_\infty)$ belongs to the feasible set of the minimization problem (5). Consequently,

$$|\hat{\theta}|_1 + \lambda_2 |\hat{\theta}|_2 + \lambda_\infty |\hat{\theta}|_\infty \leq |\hat{\theta}|_1 + \lambda_2 \hat{t} + \lambda_\infty \hat{u} \leq |\theta^*|_1 + \lambda_2 |\theta^*|_2 + \lambda_\infty |\theta^*|_\infty. \quad (28)$$

This implies

$$|\Delta_{J^c}|_1 + \lambda_2 |\Delta_{J^c}|_2 + \lambda_\infty |\Delta_{J^c}|_\infty \leq |\Delta_J|_1 + \lambda_2 |\Delta_J|_2 + \lambda_\infty |\Delta_J|_\infty \leq (1 + \lambda_2 + \lambda_\infty) |\Delta_J|_1,$$

and so

$$|\Delta_{J^c}|_1 \leq (1 + \lambda_2 + \lambda_\infty) |\Delta_J|_1.$$

and (26) follows. To prove (27), it suffices to note that (28) implies

$$\lambda_2 \hat{t} \leq |\theta^*|_1 - |\hat{\theta}|_1 + \lambda_2 |\theta^*|_2 + \lambda_\infty |\theta^*|_\infty - \lambda_\infty \hat{u} \leq |\hat{\theta} - \theta^*|_1 + \lambda_2 |\theta^*|_2 + \lambda_\infty |\hat{\theta} - \theta^*|_\infty$$

and the result follows since $|\hat{\theta}|_\infty \leq \hat{u}$ and $|\hat{\theta} - \theta^*|_\infty \leq |\hat{\theta} - \theta^*|_1$. Similar calculations yield the bound for \hat{u} . \square

Lemma 5. *Let $\hat{\theta}$ be the $\{\ell_1, \ell_2, \ell_\infty\}$ -compensated MU-selector. Assume (A1)-(A3) and (A5). Then, on a subset of the event of Lemma 3 having probability at least $1 - 8\varepsilon$, we have*

$$|\frac{1}{n}X^T X(\hat{\theta} - \theta^*)|_\infty \leq \mu_0 + \mu_1 |\hat{\theta} - \theta^*|_1 + \mu_2 |\theta^*|_2 + \mu_\infty |\theta^*|_\infty,$$

where $\mu_0 = \tau + \delta_2 + \delta_3$, $\mu_1 = 2\delta_1 + \delta_4 + \delta_5 + b(\varepsilon) + \{(1 + \lambda_\infty)/\lambda_2\}\tau_2 + \{(1 + \lambda_2)/\lambda_\infty\}\tau_\infty$, $\mu_2 = \tau_2 + \delta'_1$, $\mu_\infty = \tau_\infty + b(\varepsilon) + \delta'_4 + \delta_5$.

Note that μ_0 and μ_2 are of the order $\sqrt{\frac{1}{n} \log(c'p/\varepsilon)}$, and μ_1 and μ_∞ are of the order $\sqrt{\frac{1}{n} \log(c'p/\varepsilon)} + b(\varepsilon)$.

Proof. Throughout the proof, we assume that we are on the event of probability at least $1 - 6\varepsilon$ where inequalities (21) – (25) hold and $(\theta^*, |\theta^*|_2, |\theta^*|_\infty)$ belongs to the feasible set of the minimization problem (15). We have

$$\begin{aligned} |\frac{1}{n} X^T X \Delta|_\infty &\leq |\frac{1}{n} Z^T (Z \hat{\theta} - y) - \widehat{D} \hat{\theta}|_\infty + |(\frac{1}{n} Z^T W - D) \hat{\theta}|_\infty \\ &\quad + |(\widehat{D} - D) \hat{\theta}|_\infty + |\frac{1}{n} Z^T \xi|_\infty + |\frac{1}{n} W^T X \Delta|_\infty. \end{aligned}$$

Using the fact that $(\hat{\theta}, \hat{t}, \hat{u})$ belongs to the feasible set of the minimization problem (5) together with (27), we obtain

$$\begin{aligned} |\frac{1}{n} Z^T (Z \hat{\theta} - y) - \widehat{D} \hat{\theta}|_\infty &\leq \tau_2 \hat{t} + \tau_\infty \hat{u} + \tau \\ &\leq \{(1 + \lambda_\infty)/\lambda_2\} \tau_2 |\Delta|_1 + \tau_2 |\theta^*|_2 + \{(1 + \lambda_2)/\lambda_\infty\} \tau_\infty |\Delta|_1 + \tau_\infty |\theta^*|_\infty + \tau. \end{aligned}$$

Using that $\hat{\theta} = \theta^* + \Delta$, Assumption (A5) together with (23) yields that

$$\begin{aligned} |\frac{1}{n} X^T X \Delta|_\infty &\leq \{(1 + \lambda_\infty)/\lambda_2\} \tau_2 |\Delta|_1 + \tau_2 |\theta^*|_2 + \{(1 + \lambda_2)/\lambda_\infty\} \tau_\infty |\Delta|_1 + \tau_\infty |\theta^*|_\infty + \tau \\ &\quad + |(\frac{1}{n} Z^T W - D) \hat{\theta}|_\infty + |(\widehat{D} - D) \hat{\theta}|_\infty + |\frac{1}{n} Z^T \xi|_\infty + |\frac{1}{n} W^T X \Delta|_\infty \\ &\leq \{(1 + \lambda_\infty)/\lambda_2\} \tau_2 |\Delta|_1 + \tau_2 |\theta^*|_2 + \{(1 + \lambda_2)/\lambda_\infty\} \tau_\infty |\Delta|_1 + \tau_\infty |\theta^*|_\infty + \tau \\ &\quad + |(\frac{1}{n} Z^T W - D) \hat{\theta}|_\infty + b(\varepsilon) |\theta^*|_\infty + b(\varepsilon) |\Delta|_1 + \delta_2 + \delta_3 + |\frac{1}{n} W^T X \Delta|_\infty. \end{aligned}$$

Now remark that $|(\frac{1}{n} Z^T W - D) \hat{\theta}|_\infty \leq |(\frac{1}{n} Z^T W - D) \Delta|_\infty + |(\frac{1}{n} Z^T W - D) \theta^*|_\infty$. In view of Lemma 2 and (22), on the initial event of probability at least $1 - 6\varepsilon$,

$$\begin{aligned} &|(\frac{1}{n} Z^T W - D) \theta^*|_\infty \\ &\leq |(\frac{1}{n} (W^T W - \text{Diag}\{W^T W\}) \theta^*)|_\infty + |(\frac{1}{n} \text{Diag}\{W^T W\} - D) \theta^*|_\infty + |\frac{1}{n} X^T W \theta^*|_\infty \\ &\leq (\delta'_4 + \delta_5) |\theta^*|_\infty + \delta'_1 |\theta^*|_2. \end{aligned} \tag{29}$$

Moreover, we have

$$\begin{aligned} |(\frac{1}{n} Z^T W - D) \Delta|_\infty &\leq |\frac{1}{n} Z^T W - D|_\infty |\Delta|_1 \\ &\leq (|\frac{1}{n} (W^T W - \text{Diag}\{W^T W\})|_\infty + |\frac{1}{n} \text{Diag}\{W^T W\} - D|_\infty + |\frac{1}{n} X^T W|_\infty) |\Delta|_1. \end{aligned}$$

Therefore,

$$|(\frac{1}{n} Z^T W - D) \Delta|_\infty \leq (\delta_1 + \delta_4 + \delta_5) |\Delta|_1, \tag{30}$$

with probability at least $1 - 8\varepsilon$ (since we intersect the initial event of probability at least $1 - 6\varepsilon$ with the event of probability at least $1 - 2\varepsilon$ where the bounds δ_1 and δ_4 hold for the corresponding terms). Next, on the same event of probability at least $1 - 8\varepsilon$,

$$|\frac{1}{n} W^T X \Delta|_\infty \leq |\frac{1}{n} X^T W|_\infty |\Delta|_1 \leq \delta_1 |\Delta|_1. \tag{31}$$

To complete the proof, it suffices to plug (29) – (31) in the last inequality for $|\frac{1}{n} X^T X \Delta|_\infty$ and to obtain

$$\begin{aligned} |\frac{1}{n} X^T X \Delta|_\infty &\leq [2\delta_1 + \delta_4 + \delta_5 + b(\varepsilon) + \{(1 + \lambda_\infty)/\lambda_2\} \tau_2 + \{(1 + \lambda_2)/\lambda_\infty\} \tau_\infty] |\Delta|_1 \\ &\quad + \{\tau_2 + \delta'_1\} |\theta^*|_2 + \{\tau_\infty + b(\varepsilon) + \delta'_4 + \delta_5\} |\theta^*|_\infty + \tau + \delta_2 + \delta_3. \end{aligned}$$

□

Acknowledgement

The work of A.B. Tsybakov was supported by GENES, by the French National Research Agency (ANR) under the grants ANR -11- IDEX-0003-02, Labex ECODEC (ANR - 11-LABEX-0047), IPANEMA (ANR-13-BSH1-0004-02), and by the "Chaire Economie et Gestion des Nouvelles Données", under the auspices of Institut Louis Bachelier, Havas-Media and Paris-Dauphine.

References

- [1] A. Belloni, M. Rosenbaum and A. B. Tsybakov (2014) Linear and Conic Programming Approaches to High-Dimensional Errors-in-variables Models. [arxiv:1408.0241](#)
- [2] Y. Chen and C. Caramanis (2012) Orthogonal matching pursuit with noisy and missing data: Low and high-dimensional results. [arxiv:1206.0823](#)
- [3] Y. Chen and C. Caramanis (2013) Noisy and missing data regression: Distribution-oblivious support recovery. *Proc. of International Conference on Machine Learning (ICML)*.
- [4] L. S. Freedman, D. Midthune, R. J. Carroll, and V. Kipnis (2008) A comparison of regression calibration, moment reconstruction and imputation for adjusting for covariate measurement error in regression. *Statistics in Medicine*, **27**, 5195–5216.
- [5] E. Gautier and A. B. Tsybakov (2011) High-dimensional instrumental variables regression and confidence sets. [arxiv:1105.2454](#)
- [6] E. Gautier and A. B. Tsybakov (2013) Pivotal estimation in high-dimensional regression via linear programming. In: *Empirical Inference – Festschrift in Honor of Vladimir N. Vapnik*, B.Schölkopf, Z. Luo, V. Vovk eds., 195 - 204. Springer, New York e.a.
- [7] P-L. Loh and M. J. Wainwright (2012) High-dimensional regression with noisy and missing data: Provable guarantees with non-convexity. *Annals of Statistics*, **40**, 1637–1664.
- [8] M. Rosenbaum and A.B. Tsybakov (2010) Sparse Recovery under Matrix Uncertainty. *Annals of Statistics*, **38**, 2620–2651.
- [9] M. Rosenbaum and A.B. Tsybakov (2013) Improved Matrix Uncertainty Selector. In: *From Probability to Statistics and Back: High-Dimensional Models and Processes – A Festschrift in Honor of Jon A. Wellner*, M.Banerjee et al. eds. *IMS Collections*, vol.9, 276–290, Institute of Mathematical Statistics.
- [10] Ø. Sørensen, A. Frigessi and M. Thoresen (2012) Measurement error in Lasso: Impact and likelihood bias correction. [arxiv:1210.5378](#)
- [11] Ø. Sørensen, A. Frigessi and M. Thoresen (2014) Covariate Selection in High-Dimensional Generalized Linear Models with Measurement Error. [arxiv:1407.1070](#)