



Série des Documents de Travail

n° 2020-17

**An alternative to synthetic control for models
with many covariates under sparsity**

Marianne BLÉHAUT¹
Xavier D'HAULTFOEUILLE²
Jérémy L' HOUR³
Alexandre B. TSYBAKOV⁴

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 CREST-ENSAE, marianne.blehaut@ensae.fr

² CREST-ENSAE, xavier.dhaultfoeuille@ensae.fr

³ CREST-ENSAE, jeremy.l.hour@ensae.fr

⁴ CREST-ENSAE, alexandre.tsybakov@ensae.fr

An alternative to synthetic control for models with many covariates under sparsity*

Marianne Bléhaut[†] Xavier D’Haultfoeuille[‡] Jérémy L’Hour[§]
Alexandre B. Tsybakov[¶]

May 26, 2020

Abstract

The synthetic control method is an econometric tool to evaluate causal effects when only one unit is treated. While initially aimed at evaluating the effect of large-scale macroeconomic changes with very few available control units, it has increasingly been used in place of more well-known microeconomic tools in a broad range of applications, but its properties in this context are unknown. This paper introduces an alternative to the synthetic control method, which is developed both in the usual asymptotic framework and in the high-dimensional scenario. We propose an estimator of average treatment effect that is doubly robust, consistent and asymptotically normal. It is also immunized against first-step selection mistakes. We illustrate these properties using Monte Carlo simulations and applications to both standard and potentially high-dimensional settings, and offer a comparison with the synthetic control method.

Keywords: treatment effect, synthetic control, covariate balancing, high-dimension.

1 Introduction

The synthetic control method (Abadie and Gardeazabal, 2003; Abadie et al., 2010, 2015) is one of the most recent additions to the empiricist’s toolbox, gaining popularity not only in economics, but also in political science, medicine, etc. It provides a sound methodology in many settings where only long aggregate panel data is available to the researcher. The method has been specifically developed in a context where a single sizeable unit such as a country, a state or a city undergoes a large-scale policy change (referred to as the treatment or intervention hereafter), while only a moderate number of control units (the donor pool) is available to construct a counterfactual through a synthetic unit. This unit is defined as a convex combination of units from the donor pool that best resembles the

*We thank participants at the 2016 North American and European Meetings of the Econometric Society, the 2017 IAAE Meeting and CREST internal seminars for their useful comments and discussions. We acknowledge funding from Investissements d’Avenir (ANR-11-IDEX-0003/Labex Ecodec/ANR-11-LABX-0047).

[†]CREST-ENSAE, marianne.blehaut@ensae.fr.

[‡]CREST-ENSAE, xavier.dhaultfoeuille@ensae.fr

[§]CREST-ENSAE, jeremy.l.hour@ensae.fr.

[¶]CREST-ENSAE, alexandre.tsybakov@ensae.fr.

treated unit before the intervention. Then the treatment effect is estimated from the difference in outcomes between the treated unit and its synthetic unit after the intervention takes place. In contexts such as those described above, the synthetic unit possesses several appealing properties (for more details on such properties, see the recent survey of Abadie, 2019). First, it does not lead to extrapolation outside the support of the data: because weights are non-negative, the counterfactual never takes a value outside of the convex hull defined by the donor pool. Second, one can assess simply its fit, making it easy to judge the quality of the counterfactual. Third, the synthetic unit is sparse: the number of control units receiving a non-zero weight is at most equal to the dimension of the matching variable plus one.

The method has still some limitations, in particular when applied to micro data, for which it was not initially intended. In such cases, the number of untreated units n_0 is typically greater than the dimension p of variables X used to construct the synthetic units. Then, as soon as the treated unit falls into the convex hull defined by the donor pool, the synthetic control solution is not uniquely defined (see in particular Abadie and L’Hour, 2019). Second, and still related to the fact that the method was not developed for micro data, there is, to the best of our knowledge, no asymptotic theory available for synthetic control yet. This means in particular, that inference cannot be conducted in a standard way. A third issue is related to variable selection. The standard synthetic control method, as advocated in Abadie et al. (2010), not only minimizes the norm $\|\cdot\|_V$ – defined for a vector a of dimension p and diagonal positive-definite matrix V , as $\|a\|_V = \sqrt{a^T V a}$ – between the characteristics of the treated and those of its synthetic unit under constraints, but also employs a bi-level optimization program over the weighting matrix V so as to obtain the best possible pre-treatment fit. Diagonal elements of V are interpreted as a measure of the predicting power of each characteristics for the outcome (see, e.g., Abadie et al., 2010; Abadie, 2019). This approach has been criticized for being unstable and yielding unreproducible results, see in particular Klößner et al. (2018).

We consider an alternative to the synthetic control that addresses these issues. Specifically, we consider a parametric form for the synthetic control weights, $W_i = h(X_i^T \beta_0)$, where we estimate the unknown parameter β_0 . This approach warrants the uniqueness of the solution in low-dimensional cases where $p < n_0$. With micro data, it may thus be seen as a particular solution of the synthetic control method. We show that the average treatment on the treated (ATT) parameter can be estimated with a two-step GMM estimator, where β_0 is computed in a first step so that the reweighted control group matches some features of the treated units. A key result is the double robustness of the estimator, as defined by Bang and Robins (2005). Specifically, we show that misspecifications in the synthetic control weights do not prevent valid inference if the outcome regression function is linear for the control group.

We then turn to the high-dimensional case where p is large, possibly greater than n_0 . This case actually corresponds to the initial set-up of the synthetic control method, and is therefore crucial to take into consideration. We depart from the synthetic control method by introducing an ℓ_1 penalization term in the minimization program used to estimate β_0 . We thus perform variable selection in a similar way as the Lasso, but differently from the synthetic control method, which

relies on the aforementioned optimization over V (leading to overweighting the variables that are good predictors of the outcome and underweighting the others).

We also study the asymptotic properties of our estimator. Building on double robustness, we construct an estimator that is immunized against first-step selection mistakes in the sense defined for example by Chernozhukov et al. (2015*b*) or Chernozhukov et al. (2018). This construction requires an extra step, which models the outcome regression function and provides a bias correction, a theme that has also been developed in Ben-Michael et al. (2018), Abadie and L’Hour (2019) and Arkhangelsky et al. (2019). We show that both in the low- and high-dimensional case, the estimator is consistent and asymptotically normal. Consequently, we develop inference based on asymptotic approximation, which can be used in place of permutation tests when randomization of the treatment is not warranted.

Apart from its close connection with the synthetic control method, the present paper is related to the literature on treatment effect evaluation through propensity score weighting and covariate balancing. Several efforts have been made to include balance between covariates as an explicit objective for estimation with or without relation to the propensity score (e.g. Hainmueller, 2012; Graham et al., 2012). Our paper is in particular related to that of Imai and Ratkovic (2014), who integrate propensity score estimation and covariate balancing in the same framework. We extend their paper by considering the case of high-dimensional covariates. Note that the covariate balancing idea is related to the calibration estimation in survey sampling, see in particular Deville and Särndal (1992).

It also partakes in the econometric literature addressing variable selection, and more generally the use of machine learning tools, when estimating a treatment effect, especially but not exclusively in a high-dimensional framework. The lack of uniformity for inference after a selection step has been raised in a series of papers by Leeb and Pötscher (2005, 2008*a,b*), echoing earlier papers by Leamer (1983) who put into question the credibility of many empirical policy evaluation results. One recent solution proposed to circumvent this post-selection conundrum is the use of double-selection procedures (Belloni and Chernozhukov, 2013; Farrell, 2015; Chernozhukov et al., 2015*a*; Chernozhukov et al., 2018). For example, Belloni et al. (2014) highlight the dangers of selecting controls exclusively in their relation to the outcome and propose a three-step procedure that helps selecting more controls and guards against omitted variable biases much more than a simple “post-single-selection” estimator, as it is usually done by selecting covariates based on either their relation with the outcome or with the treatment variable, but rarely both. Farrell (2015) extends the main approach of Belloni et al. (2014) by allowing for heterogeneous treatment effects, proposing an estimator that is robust to either model selection mistakes in propensity scores or in outcome regression. However, Farrell (2015) proves the root- n consistency of his estimator only when both models hold true (see Assumption 3 and Theorem 3 therein). Our paper is also related to the work of Athey et al. (2018), who consider treatment effect estimation under the assumption of a linear conditional expectation for the outcome equation. As we do, they also estimate balancing weights, to correct for the bias arising in this high-dimensional setting, but because of their linearity assumption, they do not require to estimate a propensity score. Their method is then somewhat

simpler than ours, but it does not enjoy the double-robustness property of ours.

Finally, a recent work by Bradic et al. (2019) parallel to ours¹ suggests a Lasso-type procedure assuming logistic propensity score and linear specification for both treated and untreated items. Their method is similar but still slightly different from ours. Also, the main focus in Bradic et al. (2019) is to prove asymptotic normality under possibly weaker conditions on the sparsity levels of the parameters. Namely, they allow for sparsity up to $o(\sqrt{n}/\log(p))$ or, in some cases up to $o(n/\log(p))$ where n is the sample size, when the eigenvalues of the population Gram matrix do not depend on n . Similar results can be developed for our method at the expense of much additional technical effort. We have opted not to pursue in this direction since it only helps to include relatively non-sparse models that are not of interest for the applications we have in mind.

The paper is organized as follows. Section 2 introduces the set-up and the identification strategy behind our estimator. Section 3 presents the estimator both in the low- and high-dimensional case and studies its asymptotic properties. Section 4 examines the finite sample properties of our estimator through a Monte Carlo experiment. Section 5 revisits LaLonde (1986)’s dataset to compare our procedure with other high-dimensional econometric tools and the effect of the large-scale tobacco control program of Abadie et al. (2010) for a comparison with synthetic control. Section 6 concludes. All proofs are in the Appendix.

2 Covariate Balancing Weights and Double Robustness

We are interested in the effect of a binary treatment, coded by $D = 1$ for the treated and $D = 0$ for the non-treated. We let $Y(0)$ and $Y(1)$ denote the potential outcome under no treatment and under the treatment, respectively. The observed outcome is then $Y = DY(1) + (1 - D)Y(0)$. We also observe a random vector $X \in \mathbb{R}^p$ of pre-treatment characteristics. The quantity of interest is the average treatment effect on the treated (ATT), defined as:

$$\theta_0 = \mathbb{E}[Y(1) - Y(0)|D = 1].$$

Here and in what follows, we assume that the random variables are such that all the considered expectations are finite. Since no individual is observed in both treatment states, identification of the counterfactual $\mathbb{E}[Y(0)|D = 1]$ is achieved through the following two ubiquitous conditions.

Assumption 2.1 (Nested Support) $\mathbb{P}[D = 1|X] < 1$ almost surely and $0 < \mathbb{P}[D = 1] < 1$.

Assumption 2.2 (Mean Independence) $\mathbb{E}[Y(0)|X, D = 1] = \mathbb{E}[Y(0)|X, D = 0]$.

Assumption 2.1, a version of the usual common support condition, requires that there exist control units for any possible value of the covariates in the population. Since the ATT is the parameter of interest, we are never reconstructing a counterfactual for control units so $\mathbb{P}[D = 1|X] > 0$ is

¹Our results have been presented as early as 2016 at the North American and European Summer Meetings of the Econometric Society (see https://www.econometricsociety.org/sites/default/files/regions/program_ESEM2016.pdf), and again in 2017 during the IAAE Meeting in Sapporo, see Bléhaut et al. (2017).

not required. Assumption 2.2 states that conditional on a set of observed confounding factors, the expected potential outcome under no treatment is the same for treated and control individuals. This assumption is a weaker form of the classical conditional independence assumption $(Y(0), Y(1)) \perp\!\!\!\perp D|X$.

In policy evaluation settings, the counterfactual is usually identified and estimated as a weighted average of non-treated unit outcomes:

$$\theta_0 = \mathbb{E}[Y(1)|D = 1] - \mathbb{E}[WY(0)|D = 0], \quad (2.1)$$

where W is a random variable called the weight. Popular choices for the weight are the following:

1. Linear regression: $W = \mathbb{E}[DX^T]\mathbb{E}[(1 - D)XX^T]^{-1}X$, also referred to as the Oaxaca-Blinder weight (Kline, 2011);
2. Propensity score: $W = P[D = 1|X]/(1 - P[D = 1|X])$;
3. Matching: $W = P(D = 1)f_{X|D=1}(X)/[P(D = 0)f_{X|D=0}(X)]$; ²
4. Synthetic controls: see Abadie et al. (2010).

In this paper, we propose another choice of weight W , which can be seen as a parametric alternative to the synthetic control. An advantage is that it is well-defined whether or not the number of untreated observations n_0 is greater than the dimension p of X , whereas the synthetic control estimator is not uniquely defined when $n_0 > p$. Formally, we look for weights W that:

- (i) satisfy a balancing condition as in the synthetic control method;
- (ii) are positive;
- (iii) depend only on the covariates;
- (iv) can be used whether $n_0 > p$ or $n_0 \leq p$ (high-dimensional regime).

Satisfying a balancing condition means that

$$\mathbb{E}[DX] = \mathbb{E}[W(1 - D)X]. \quad (2.2)$$

Up to a proportionality constant, this is equivalent to $\mathbb{E}[X|D = 1] = \mathbb{E}[WX|D = 0]$. In words, W balances the first moment of the observed covariates between the treated and the control group. The definition of the observable covariates X is left to the econometrician and can include transformation of the original covariates so as to match more features of their distribution. The idea behind such weights relies on the principle of ‘‘covariate balancing’’ as in, e.g., Imai and Ratkovic (2014). The following lemma shows that under Assumption 2.1 weights satisfying the balancing condition always exist.

²Assuming here that the conditional densities $f_{X|D=1}$ and $f_{X|D=0}$ exist. Of course, $P(D = 1)f_{X|D=1}(X)/[P(D = 0)f_{X|D=0}(X)] = P[D = 1|X]/(1 - P[D = 1|X])$, but the methods of estimation of W differ in the two cases.

Lemma 2.1 *If Assumption 2.1 holds, the propensity score weight $W = \mathbb{P}[D = 1|X]/(1 - \mathbb{P}[D = 1|X])$ satisfies the balancing condition (2.2).*

The proof of this lemma is straightforward by plugging the expression of W in Equation (2.2) and using the law of iterated expectations. Note that W is not a unique solution of (2.2). The linear regression weight $W = \mathbb{E}[DX^T]\mathbb{E}[(1 - D)XX^T]^{-1}X$ also satisfies the balancing condition but it can be negative and its use is problematic in high-dimensional regime.

Lemma 2.1 would suggest solving a binary choice model to obtain estimators of $\mathbb{P}[D = 1|X]$ and of the weight W as a first step, and then plugging W in (2.1) to estimate θ_0 . However, an inconsistent estimator of the propensity score leads to an inconsistent estimator of θ_0 and does not guarantee that the corresponding weight will achieve covariate balancing. Finally, estimation of the propensity score can be problematic when there are very few treated units. For these reasons, we consider another approach where estimation is based directly on balancing equations:

$$\mathbb{E}[(D - (1 - D)W)X] = 0. \tag{2.3}$$

An important advantage of this approach over the usual one based on the propensity score estimation through maximum likelihood is its double robustness (for a definition, see, e.g., Bang and Robins, 2005). Indeed, let W_1 denote the weights identified by (2.3) under a misspecified model on the propensity score. It turns out that if the balancing equations (2.3) hold for W_1 the estimated treatment effect will still be consistent provided that $\mathbb{E}[Y(0)|X]$ is linear in X . The formal result is given in Theorem 2.1 below.

Theorem 2.1 (Double Robustness) *Let Assumptions 2.1-2.2 hold and let $w : \mathbb{R}^p \rightarrow (0, +\infty)$ be a measurable function such that $\mathbb{E}[w(X)|Y] < \infty$, $\mathbb{E}[w(X)\|X\|_2] < \infty$, where $\|\cdot\|_2$ denotes the Euclidean norm. Assume the balancing condition*

$$\mathbb{E}[(D - (1 - D)w(X))X] = 0. \tag{2.4}$$

Then, for any $\mu \in \mathbb{R}^p$ the ATT θ_0 can be expressed as

$$\theta_0 = \frac{1}{\mathbb{P}(D = 1)} \mathbb{E}[(D - (1 - D)w(X))(Y - X^T\mu)] \tag{2.5}$$

in each of the following two cases:

1. $\mathbb{E}[Y(0)|X] = X^T\mu_0$ for some $\mu_0 \in \mathbb{R}^p$;
2. $P[D = 1|X] = w(X)/(1 + w(X))$.

In (2.5), the effect of X is taken out from Y in a linear fashion, while the effect of X on D is taken out by re-weighting the control group to obtain the same mean for X . Theorem 2.1 shows that an estimator based on (2.5) enjoys the double robustness property. Theorem 2.1 is similar to the result of Kline (2011) for the Oaxaca-Blinder estimator, which is obtained under the assumption that the propensity score follows specifically a log-logistic model in the propensity-score-well-specified

case. Theorem 2.1 is more general. It can be applied under parametric modeling of W as well as in nonparametric settings.

In this paper, we consider a parametric model for $w(X)$. Namely, we assume that $P[D = 1|X] = G(X^T \beta_0)$ for some unknown $\beta_0 \in \mathbb{R}^p$ and some known strictly increasing cumulative distribution function G . Then $w(X) = h(X^T \beta_0)$ with $h = G/(1 - G)$ and β_0 is identified by the balancing condition

$$\mathbb{E} [(D - (1 - D)h(X^T \beta_0)) X] = 0. \quad (2.6)$$

Clearly, h is a positive strictly increasing function, which implies that its primitive H is strictly convex. A classical example is to take G as the c.d.f. of the logistic distribution, in which case $h(u) = H(u) = \exp(u)$ for $u \in \mathbb{R}$. The strict convexity of H implies that β_0 is the unique solution of a strictly convex program:

$$\beta_0 = \arg \min_{\beta \in \mathbb{R}^p} \mathbb{E} [(1 - D)H(X^T \beta) - DX^T \beta]. \quad (2.7)$$

This program is well-defined, whether or not $P[D = 1|X] = G(X^T \beta_0)$. Note also that definitions (2.6) and (2.7) are equivalent provided that $\mathbb{E} [h(X^T \beta) \|X\|_2] < \infty$ for β in a vicinity of β_0 . Indeed, it follows from the dominated convergence theorem that, under this assumption and due to the fact that any convex function is locally Lipschitz, differentiation under the expectation sign is legitimate in (2.7).

We are now ready to state the main identification theorem justifying the use of ATT estimation methods developed below. It is a straightforward corollary of Theorem 2.1.

Theorem 2.2 (Parametric Double Robustness) *Let Assumptions 2.1-2.2 hold. Assume that $\beta_0 \in \mathbb{R}^p$ and a positive strictly increasing function h are such that $\mathbb{E} [h(X^T \beta_0) \|X\|_2] < \infty$, $\mathbb{E} [h(X^T \beta_0) |Y|] < \infty$ and condition (2.6) holds. Then, for any $\mu \in \mathbb{R}^p$, the ATT θ_0 satisfies*

$$\theta_0 = \frac{1}{\mathbb{P}(D=1)} \mathbb{E} [(D - (1 - D)h(X^T \beta_0)) (Y - X^T \mu)], \quad (2.8)$$

in each of two the following cases.

1. *There exists $\mu_0 \in \mathbb{R}^p$ such that $\mathbb{E}[Y(0)|X] = X^T \mu_0$.*
2. *$P[D = 1|X] = G(X^T \beta_0)$ with $G = h/(1 + h)$.*

At this stage, the parameter μ in Equation (2.8) does not play any role and can, for example, be zero. However, we will see below that in the high-dimensional regime, choosing μ carefully is crucial to obtain an “immunized” estimator of θ_0 that enjoys the desirable asymptotic properties.

3 A Parametric Alternative to Synthetic Control

We now assume to have a sample $(D_i, X_i, Y_i)_{i=1 \dots n}$ of i.i.d. random variables with the same distribution as (D, X, Y) .

3.1 Estimation with low-dimensional covariates

Consider first an asymptotic regime where the dimension p of the covariates is fixed, while the sample size n tends to infinity. We call it the low-dimensional regime. Define an estimator of β_0 via the empirical counterpart of (2.7):

$$\hat{\beta}_{\text{ld}} \in \arg \min_{\beta \in \mathbb{R}^p} \frac{1}{n} \sum_{i=1}^n [(1 - D_i)H(X_i^T \beta) - D_i X_i^T \beta]. \quad (3.1)$$

Next, we plug $\hat{\beta}_{\text{ld}}$ in the empirical counterpart of (2.8) to obtain the following estimator of θ_0 :

$$\tilde{\theta}_{\text{ld}} := \frac{1}{\frac{1}{n} \sum_{i=1}^n D_i} \left(\frac{1}{n} \sum_{i=1}^n [D_i - (1 - D_i)h(X_i^T \hat{\beta}_{\text{ld}})] Y_i \right).$$

Note that if X includes the intercept, $\tilde{\theta}_{\text{ld}}$ satisfies the desirable property of location invariance, namely it does not change if we replace all Y_i by $Y_i + c$, for any $c \in \mathbb{R}$.

Set $Z := (D, X, Y)$, $Z_i := (D_i, X_i, Y_i)$ and introduce the function

$$g(Z, \theta, (\beta, \mu)) := [D - (1 - D)h(X^T \beta)][Y - X^T \mu] - D\theta.$$

Then the estimator $\tilde{\theta}_{\text{ld}}$ satisfies

$$\frac{1}{n} \sum_{i=1}^n g(Z_i, \tilde{\theta}_{\text{ld}}, (\hat{\beta}_{\text{ld}}, 0)) = 0. \quad (3.2)$$

This estimator is a two-step GMM. It is consistent and asymptotically normal under mild regularity conditions, with asymptotic variance $\mathbb{E} [g^2(Z, \theta_0, (\beta_0, \mu_0))] / \mathbb{E}(D)^2$, where

$$\mu_0 = \mathbb{E}[h'(X^T \beta_0) X X^T | D = 0]^{-1} \mathbb{E}[h'(X^T \beta_0) X Y | D = 0]. \quad (3.3)$$

This can be shown by standard techniques (see, e.g., Section 6 in Newey and McFadden, 1994). Notice that since $h'(X^T \beta_0) > 0$ the vector μ_0 is the coefficient of the weighted population regression of Y on X for the control group. This observation is useful for the derivation of the “immunized” estimator in the high-dimensional case, to which we now turn.

3.2 Estimation with high-dimensional covariates

We now consider that p may grow with n , with possibly $p \gg n$. This can be of interest in several situations. First, in macroeconomic problems, n is actually small, and p may easily be of comparable size. For example, in the Tobacco control program application by Abadie et al. (2010) the control group size is limited due to the fact that the observational unit is the state but many pre-treatment outcomes are included among the covariates. Section 5.2 revisits this example. Second, researcher may want to consider a flexible form for the weights by including transformations of the covariates. For instance, one may want to interact categorical variables with other covariates or consider, e.g., different powers of continuous variables if one wants to allow for flexible non-linear effects. See Section 5.1 for an application considering such transformations. Third, one may want not only to

balance the first moments of the distribution of the covariates but also the second moments, the covariances, the third moments and so on to make the distribution more similar between the treated and the control group. In this case, high-dimensional settings seem to be of interest as well.

In high-dimensional regime, the GMM estimator in (3.1) is, in general, not consistent. We therefore propose an alternative Lasso-type method by adding in (3.1) an ℓ_1 penalization term:

$$\hat{\beta} \in \arg \min_{\beta \in \mathbb{R}^p} \left(\frac{1}{n} \sum_{i=1}^n [(1 - D_i)H(X_i^T \beta) - D_i X_i^T \beta] + \lambda \sum_{j=1}^p \psi_j |\beta_j| \right). \quad (3.4)$$

Here, $\lambda > 0$ is an overall penalty parameter set to dominate the noise in the gradient of the objective function and $\{\psi_j\}_{j=1, \dots, p}$ are covariate specific penalty loadings set as to grant good asymptotic properties. The penalty loadings can be adjusted using the algorithm presented in Appendix A.

This type of penalization offers several advantages. First, the program (3.4) has almost surely a unique solution when the entries of X have a continuous distribution, cf. Lemma 5 in Tibshirani (2013), which cannot be granted for its non-penalized version (3.1). Second, it yields a sparse solution in the sense that some entries of the vector of estimated coefficients are set exactly to zero if the penalty is large enough, which is not the case for estimators based on ℓ_2 penalization. The ℓ_0 -penalized estimator shares the same sparsity property but is very costly to compute, whereas (3.4) can be easily solved by computationally efficient methods, see, e.g., Hastie et al. (2009).

The use of covariate specific penalty loadings goes back to Bickel et al. (2009); the particular choice of penalty loadings that we consider below is inspired by Belloni et al. (2012). A drawback of penalizing by the ℓ_1 -norm is that it induces a bias in estimation of the coefficients. But this is not an issue here since we are ultimately interested in estimating θ_0 rather than β_0 . The solution $\hat{\beta}$ of (3.4) only plays the role of a pilot estimator.

The estimator $\hat{\beta}$ is consistent as n tends to infinity under assumptions analogous to those used in Bunea et al. (2007); Bickel et al. (2009) for the Lasso with quadratic loss, see Theorem 3.2 below. As in the low-dimensional case (cf. Section 3.1), one is then tempted to consider the plug-in estimator for the ATT based on Equation (2.8) with $\mu = 0$:

$$\tilde{\theta} = \frac{1}{\sum_{i=1}^n D_i} \sum_{i=1}^n [D_i - (1 - D_i)h(X_i^T \hat{\beta})] Y_i. \quad (3.5)$$

We refer to this estimator as the naive plug-in estimator. However, as mentioned above, the Lasso estimator $\hat{\beta}$ of the nuisance parameter β_0 is not asymptotically unbiased. In high-dimensional regime where p grows with n , naive plug-in estimators suffer from a regularization bias and may not be asymptotically normal with zero mean, as illustrated for example in Belloni et al. (2014); Chernozhukov et al. (2015b, 2018). Therefore, following the general approach of Chernozhukov et al. (2015b, 2018), we develop an immunized estimator that, at the first order, is insensitive to $\hat{\beta}$. We show that this estimator is asymptotically normal with mean zero and an asymptotic variance that does not depend on the properties of the pilot estimator $\hat{\beta}$. The idea is to choose parameter

μ in (2.8) such that the expected gradient of the estimating function $g(Z, \theta, (\beta, \mu))$ with respect to β is zero when taken at (θ_0, β_0) . This holds for $\mu = \mu_0$, where μ_0 satisfies

$$\mathbb{E} [(1 - D)h'(X^T \beta_0)(Y - X^T \mu_0)X] = 0. \quad (3.6)$$

Notice that if the corresponding matrix is invertible we get the low-dimensional solution (3.3). Clearly, μ_0 depends on unknown quantities and we need to estimate it. To this end, observe that Equation (3.6) corresponds to the first-order condition of a weighted least-squares problem,³ namely

$$\mu_0 = \arg \min_{\mu \in \mathbb{R}^p} \mathbb{E} [(1 - D)h'(X^T \beta_0)(Y - X^T \mu)^2]. \quad (3.7)$$

Since X is high-dimensional we cannot estimate μ_0 via the empirical counterpart of (3.7). Instead, we consider a Lasso-type estimator

$$\hat{\mu} \in \arg \min_{\mu \in \mathbb{R}^p} \left(\frac{1}{n} \sum_{i=1}^n [(1 - D_i)h'(X_i^T \hat{\beta})(Y_i - X_i^T \mu)^2] + \lambda' \sum_{j=1}^p \psi'_j |\mu_j| \right). \quad (3.8)$$

Here, similarly to (3.4), $\lambda' > 0$ is an overall penalty parameter set to dominate the noise in the gradient of the objective function and $\{\psi'_j\}_{j=1, \dots, p}$ are covariate-specific penalty loadings. Importantly, by estimating μ_0 we do not introduce, at least asymptotically, an additional source of variability since by construction, the gradient of the moment condition (2.8) (if we consider it as function of β rather than of β_0) with respect to (β, μ) vanishes at point (β_0, μ_0) .

Finally, the immunized ATT estimator is defined as

$$\hat{\theta} := \frac{1}{\sum_{i=1}^n D_i} \sum_{i=1}^n \left(D_i - (1 - D_i)h(X_i^T \hat{\beta}) \right) (Y_i - X_i^T \hat{\mu}).$$

Intuitively, the immunized procedure corrects the naive plug-in estimator in the case where the balancing program has “missed” a covariate that is very important to predict the outcome:

$$\hat{\theta} = \tilde{\theta} - \left[\frac{1}{n_1} \sum_{i: D_i=1}^n X_i - \frac{1}{n_1} \sum_{i: D_i=0}^n h(X_i^T \hat{\beta}) X_i \right]^T \hat{\mu},$$

where n_1 is the number of treated observations. This has a flavor of Frish-Waugh-Lowell partialling-out procedure for model selection as observed by Belloni et al. (2014) and further developed in Chernozhukov et al. (2015b).

To summarize, the estimation procedure in high-dimensional regime consists of the three following steps. Each step is computationally simple as it needs at most to minimize a convex function:

1. (*Balancing step.*) For a given level of penalty λ and positive covariate-specific penalty loadings $\{\psi_j\}_{j=1}^p$, compute $\hat{\beta}$ defined by

$$\hat{\beta} \in \arg \min_{\beta \in \mathbb{R}^p} \left(\frac{1}{n} \sum_{i=1}^n [(1 - D_i)H(X_i^T \beta) - D_i X_i^T \beta] + \lambda \sum_{j=1}^p \psi_j |\beta_j| \right). \quad (3.9)$$

³The assumptions under which we prove the results below guarantee that μ_0 defined here is unique. Extension to the case of multiple solutions can be worked out as well. It is technically more involved but in our opinion does not add much to the understanding of the problem.

2. (*Immunization step.*) For a given level of penalty λ' and covariate-specific penalty loadings $\{\psi'_j\}_{j=1}^p$, and using $\hat{\beta}$ obtained in the previous step, compute $\hat{\mu}$ defined by:

$$\hat{\mu} \in \arg \min_{\mu \in \mathbb{R}^p} \left(\frac{1}{n} \sum_{i=1}^n [(1 - D_i)h'(X_i^T \hat{\beta}) (Y_i - X_i^T \mu)^2] + \lambda' \sum_{j=1}^p \psi'_j |\mu_j| \right). \quad (3.10)$$

3. (*ATT estimation.*) Estimate the ATT using the immunized estimator:

$$\hat{\theta} = \frac{1}{\sum_{i=1}^n D_i} \sum_{i=1}^n \left[D_i - (1 - D_i)h(X_i^T \hat{\beta}) \right] (Y_i - X_i^T \hat{\mu}). \quad (3.11)$$

3.3 Asymptotic Properties

The current framework poses several challenges to achieving asymptotically valid inference. First, X can be high-dimensional since we allow for $p \gg n$ provided that sparsity conditions are met (see Assumption 3.1 below). Second, the ATT estimation is affected by the estimation of the nuisance parameters (β_0, μ_0) and we wish to neutralize their influence. Finally, the ℓ_1 -penalized estimators we use for β_0 and μ_0 are not conventional. The estimator of β_0 relies on a non-standard loss function and, to our knowledge, the properties of $\hat{\beta}$ that we need are not available in the literature, cf., e.g., van de Geer (2016) and the references therein. The estimator of μ_0 is close to the usual Lasso except for the weights that depend on $\hat{\beta}$. In general, discrepancy in the weights can induce an extra bias. Thus, it is not granted that such an estimator achieves properties close to the Lasso. We show below that it holds true under our assumptions. Our proof techniques may be of interest for other problems of similar type.

Let $\eta = (\beta, \mu)$ denote the vector of two nuisance parameters and recall that $Z = (D, X, Y)$. In what follows, we write for brevity $g(Z, \theta, \eta)$ instead of $g(Z, \theta, (\beta, \mu))$. In particular, for the value $\eta_0 := (\beta_0, \mu_0)$ we have

$$\mathbb{E}g(Z, \theta_0, \eta_0) = 0. \quad (3.12)$$

Hereafter, the notation $a \lesssim b$ means that $a \leq cb$ for some constant $c > 0$ independent of the sample size n . We denote by Φ and Φ^{-1} the cumulative distribution function and the quantile function of a standard normal random variable, respectively. We use the symbol $\mathbb{E}_n(\cdot)$ to denote the empirical average, that is $\mathbb{E}_n(a) = n^{-1} \sum_{i=1}^n a_i$ for $a = (a_1, \dots, a_n)$. Finally, for a vector $\delta = (\delta_1, \dots, \delta_p) \in \mathbb{R}^p$ and a subset $S \subseteq \{1, \dots, p\}$ we consider the restricted vector $\delta_S = (\delta_j \mathbb{I}(j \in S))_{j=1}^p$, where $\mathbb{I}(\cdot)$ denotes the indicator function, and we set $\|\delta\|_0 := \text{Card} \{1 \leq j \leq p : \delta_j \neq 0\}$, $\|\delta\|_1 := \sum_{j=1}^p |\delta_j|$, $\|\delta\|_2 := \sqrt{\sum_{j=1}^p \delta_j^2}$ and $\|\delta\|_\infty := \max_{j=1, \dots, p} |\delta_j|$.

We now state the assumptions used to prove the asymptotic results.

Assumption 3.1 (Sparsity restrictions) *The nuisance parameter is sparse in the following sense:*

$$\|\beta_0\|_0 \leq s_\beta, \quad \|\mu_0\|_0 \leq s_\mu$$

for some integers $s_\beta, s_\mu \in [1, p]$.

Assumption 3.2 (Conditions on function h) *Function h is increasing, twice continuously differentiable on \mathbb{R} and*

(i) *the second derivative h'' is Lipschitz on any compact subset of \mathbb{R} ,*

(ii) *either $\inf_{u \in \mathbb{R}} h'(u) \geq c_2$ or $\|\beta_0\|_1 \leq c_3$ where $c_2 > 0$ and $c_3 > 0$ are constants independent of n .*

We also need some conditions on the distribution of data. The random vectors $Z_i = (D_i, X_i, Y_i)$ are assumed to be i.i.d. copies of $Z = (D, X, Y)$ with $D \in \{0, 1\}$, $X \in \mathbb{R}^p$ and $Y \in \mathbb{R}$. Throughout the paper, we assume that Z depends on n , so that in fact we deal with a triangular array of random vectors. This dependence on n is needed for rigorously stating asymptotical results since we consider the setting where the dimension $p = p(n)$ is a function of n . Thus, in what follows Z is indexed by n but for brevity we typically suppress this dependence in the notation. On the other hand, all constants denoted by c (with various indices) and K appearing below are independent of n .

Assumption 3.3 (Conditions on the distribution of data) *The random vectors Z_i are i.i.d. copies of $Z = (D, X, Y)$ with $D \in \{0, 1\}$, $X \in \mathbb{R}^p$ and $Y \in \mathbb{R}$ satisfying (2.6), (3.6), (3.12) and the following conditions:*

(i) *There exist constants $K > 0$ and $c_1 > 0$ such that*

$$\max\{\|X\|_\infty, |X^T \beta_0|, |Y - X^T \mu_0|\} \leq K \text{ (a.s.)}, \quad \text{and} \quad 0 < \mathbb{P}(D = 1) < 1.$$

(ii) *Non-degeneracy conditions. There exists $c_1 > 0$ such that for all $j = 1, \dots, p$,*

$$\min \left\{ \begin{aligned} &\mathbb{E}((Y - X^T \mu_0 - \theta_0)^2 | D = 1), \quad \mathbb{E}(h^2(X^T \beta_0)(Y - X^T \mu_0)^2 | D = 0), \\ &\mathbb{E}((X^T e_j)^2 | D = 1), \quad \mathbb{E}(h^2(X^T \beta_0)(X^T e_j)^2 | D = 0), \\ &\mathbb{E}((h'(X^T \beta_0))^2 (Y - X^T \mu_0)^2 (X^T e_j)^2 | D = 0) \end{aligned} \right\} \geq c_1,$$

where e_j denotes the j th canonical basis vector in \mathbb{R}^p .

Assumption 3.4 (Condition on the population Gram matrix of the control group) *The population Gram matrix of the control group*

$$\Sigma := \mathbb{E}((1 - D)XX^T)$$

is such that

$$\min_{v \in \mathbb{R}^p \setminus \{0\}} \frac{v^T \Sigma v}{\|v\|_2^2} \geq \kappa_\Sigma, \tag{3.13}$$

where the minimal eigenvalue κ_Σ is a positive number.

Note that, in view of Assumption 3.3, κ_Σ is uniformly bounded: $\kappa_\Sigma \leq e_1^T \Sigma e_1 \leq K^2$.

Assumption 3.5 (Dimension restrictions) *The integers $p, s = \max(s_\beta, s_\mu) \in [1, p/2]$ and the value $\kappa_\Sigma > 0$ are functions of n satisfying the following growth conditions:*

$$(i) \frac{s^2 \log(p)}{\kappa_\Sigma^2 \sqrt{n}} \rightarrow 0 \text{ as } n \rightarrow \infty,$$

$$(ii) s/\kappa_\Sigma = o(p) \text{ as } n \rightarrow \infty,$$

$$(iii) \log(p) = o(n^{1/3}) \text{ as } n \rightarrow \infty.$$

Finally, we define the penalty loadings for estimation of nuisance parameters. The gradients of the estimating function with respect to the nuisance parameters are

$$\begin{aligned} \nabla_\mu g(Z, \theta, \eta) &= - [D - (1 - D)h(X^T \beta)] X, \\ \nabla_\beta g(Z, \theta, \eta) &= -(1 - D)h'(X^T \beta) [Y - X^T \mu] X. \end{aligned}$$

For each $i = 1, \dots, n$, we define the random vector $U_i \in \mathbb{R}^{2p}$ with entries corresponding to these gradients:

$$U_{i,j} := \begin{cases} - [D_i - (1 - D_i)h(X_i^T \beta_0)] X_{i,j} & \text{if } 1 \leq j \leq p, \\ -(1 - D_i)h'(X_i^T \beta_0) [Y_i - X_i^T \mu_0] X_{i,j} & \text{if } p + 1 \leq j \leq 2p, \end{cases}$$

where $X_{i,j}$ is the j th entry of X_i .

Assumption 3.6 (Penalty Loadings) *Let $c > 1$ and $\gamma \in (0, 2p)$ be such that $\log(1/\gamma) \lesssim \log(p)$ and $\gamma = o(1)$ as $n \rightarrow \infty$. The penalty loadings for estimation of β_0 satisfy*

$$\begin{aligned} \lambda &:= c\Phi^{-1}(1 - \gamma/2p)/\sqrt{n}, \\ \psi_{j,\max} \geq \psi_j &\geq \sqrt{n^{-1} \sum_{i=1}^n U_{i,j}^2} \quad \text{for } j = 1, \dots, p. \end{aligned}$$

The penalty loadings for estimation of μ_0 satisfy

$$\begin{aligned} \lambda' &:= 2c\Phi^{-1}(1 - \gamma/2p)/\sqrt{n}, \\ \psi'_{j,\max} \geq \psi'_j &\geq \sqrt{n^{-1} \sum_{i=1}^n U_{i,j+p}^2} \quad \text{for } j = 1, \dots, p. \end{aligned}$$

Here, the upper bounds on the loadings are

$$\begin{aligned} \psi_{j,\max} &= \sqrt{\frac{1}{n} \sum_{i=1}^n \max_{|u| \leq K} [(1 - D_i)h(u) - D_i]^2 X_{i,j}^2}, \\ \psi'_{j,\max} &= \sqrt{\frac{1}{n} \sum_{i=1}^n (1 - D_i) \max_{|u| \leq K} (|h'(u)|^2 [Y_i - u]^2) X_{i,j}^2} \end{aligned}$$

representing feasible majorants for $\sqrt{n^{-1} \sum_{i=1}^n U_{i,j}^2}$ under our assumptions.

The values $\sqrt{n^{-1} \sum_{i=1}^n U_{i,j}^2}$ depend on the unknown parameters. Thus, we cannot choose ψ_j, ψ'_j equal to these values but we can take them equal to the upper bounds $\psi_j = \psi_{j,\max}$ and $\psi'_j = \psi'_{j,\max}$. A more flexible iterative approach of choosing feasible loadings is discussed in Appendix A.

The following theorem constitutes the main asymptotic result of the paper.

Theorem 3.1 (Asymptotic Normality of the Immunized Estimator) *Let Assumptions 3.1 - 3.6 hold. Then the immunized estimator $\hat{\theta}$ defined in Equation (3.11) satisfies*

$$\hat{\sigma}^{-1} \sqrt{n}(\hat{\theta} - \theta_0) \xrightarrow{\mathcal{D}} \mathcal{N}(0, 1) \text{ as } n \rightarrow \infty,$$

where $\hat{\sigma}^2 := \frac{1}{n} \sum_{i=1}^n g^2(Z_i, \hat{\theta}, \hat{\eta}) \left(\frac{1}{n} \sum_{i=1}^n D_i \right)^{-2}$ is a consistent estimator of the asymptotic variance, and $\xrightarrow{\mathcal{D}}$ denotes convergence in distribution.

The proof is given in Appendix B.2. An important point underlying the root- n convergence and asymptotic normality of $\hat{\theta}$ is the fact that the expected gradient of g with respect to η is zero at η_0 . In granting this property, we follow the general methodology of estimation in the presence of high-dimensional nuisance parameters developed in Belloni et al. (2012, 2014, 2017); Chernozhukov et al. (2015b) among other papers by the same authors. The second important ingredient of the proof is to ensure that the estimator $\hat{\eta}$ converges fast enough to the nuisance parameter η_0 . Its rate of convergence is given in the following theorem.

Theorem 3.2 (Nuisance Parameter Estimation) *Under Assumptions 3.1-3.6, we have, with probability tending to 1 as $n \rightarrow \infty$,*

$$\|\hat{\beta} - \beta_0\|_1 \lesssim \frac{s_\beta}{\kappa_\Sigma} \sqrt{\frac{\log(p)}{n}}, \tag{3.14}$$

$$\|\hat{\mu} - \mu_0\|_1 \lesssim \frac{s}{\kappa_\Sigma} \sqrt{\frac{\log(p)}{n}}. \tag{3.15}$$

The proof is given in Appendix B.3. Note that the rate in (3.15) depends not only on the sparsity index s_μ of μ_0 but on the maximum $s = \max(s_\beta, s_\mu)$. This is natural since one should account for the accuracy of the preliminary estimator $\hat{\beta}$ used to obtain $\hat{\mu}$. Inspection of the proof shows that Theorem 3.2 remains valid under weaker assumptions, namely, we can replace Assumption 3.5 on s, p, κ_Σ by the condition (C.2). We also note that Assumption 3.5(i) in Theorem 3.1 can be modified to $(s/\sqrt{n}) \log(p) \rightarrow 0$ if κ_Σ is a constant independent of n as it is assumed, for example, in Bradic et al. (2019). Such a modification would require a substantially more involved proof but only improves upon considering relatively non-sparse cases. This does not seem of much added value for using the methods in practice when the sparsity index is typically small. Moreover, in the high-dimensional scenario we find it more important to specify the dependency of the growth conditions on κ_Σ .

4 Monte Carlo Simulations

The aim of this experiment is two-fold: illustrate the better properties of the immunized estimator over the naive plug-in, and compare it with other estimators. In particular, we compare it with a similar estimator proposed by Farrell (2015). We consider the following DGP. The p covariates are distributed as $X \sim \mathcal{N}(0, \Sigma)$, where the (i, j) element of the variance-covariance matrix satisfies $\Sigma_{i,j} = .5^{|i-j|}$. The treatment equation follows a logit model, $\Pr(D = 1|X) = \Lambda(X^T \gamma_0)$ with $\Lambda(u) = 1/(1 + \exp(-u))$. The potential outcomes satisfy

$$Y(d) = \exp(X^T \mu_0) + d(\zeta_0 X^T \gamma_0) + \varepsilon, \quad \varepsilon|D, X \sim \mathcal{N}(0, 1).$$


We assume the following form for the j th entry of γ_0 and μ_0 :

$$\gamma_{0j} = \begin{cases} \rho_\gamma (-1)^j / j^2 & \text{if } j \leq 10 \\ 0 & \text{otherwise} \end{cases}, \quad \mu_{0j} = \begin{cases} \rho_\mu (-1)^j / j^2 & \text{if } j \leq 10 \\ \rho_\mu (-1)^{j+1} / (p - j + 1)^2 & \text{if } j \geq p - 9 \\ 0 & \text{otherwise.} \end{cases}$$

We are thus in an strictly sparse setting for both equations where only ten covariates play a role in the treatment assignment and twenty in the outcome. Figure 1 depicts the precise pattern of the corresponding coefficients for $p = 30$. ρ_γ and ρ_μ are constants that fix the signal-to-noise ratio. More precisely, ρ_γ is set so that $R^2 = 0.3$ in the latent model for D , and ρ_μ is set so that $R^2 = 0.8$ in the model for $Y(0)$. Finally, we let $\zeta_0 = [V(\exp(X^T \mu_0)) / 5V(Y(0))]^{1/2}$. This implies that the variance of the individual treatment effect $\zeta_0 X^T \gamma_0$ is one fifth of the variance of $Y(0)$. In this set-up, the ATT satisfies $E[Y(1) - Y(0)|D = 1] = \zeta_0 E[Z \Lambda(Z)] / E[\Lambda(Z)]$, with $Z \sim \mathcal{N}(0, \gamma_0^T \Sigma \gamma_0)$. We compute it using Monte Carlo simulations.

We consider several estimators of the ATT. The first is the naive plug-in estimator defined in (3.5). Next, we consider our proposed estimator defined in (3.11), with $H(x) = h(x) = \exp(x)$. We also consider the estimator proposed by Farrell (2015). This estimator is also defined by (3.11), but $\hat{\beta}$ and $\hat{\mu}$ therein are obtained by a Logit Lasso and an unweighted Lasso regression, respectively. For these three estimators, the penalty loadings for the first-step estimators are set as in Appendix A. The last estimator, called the oracle hereafter, is our low-dimensional estimator defined by (3.2), where in the first step we only include the ten covariates affecting the treatment. For all estimators, we construct 95% confidence estimator on the ATT using the normal approximation and asymptotic variance estimators. We estimate the asymptotic variance of the naive plug-in estimator making as if we were in a low-dimensional setting. This means that this estimator would have an asymptotic coverage of 0.95 if p remained fixed.

Coefficient value



Notes: In this example, $\rho_\gamma = \rho_\mu = 1$. The central region of the graph represents the coefficients γ_0 and μ_0 associated with variables that do not play a role in either the equation equation or the outcome equation. The left region shows the coefficients associated with variables that are important for both equations. In the right region, only entries in μ_0 is different from 0, meaning that the variables determine the outcome equation but not the selection equation.

Figure 1: Sparsity patterns of γ_0 (crosses) and μ_0 (circles).

In our DGP, the variables X_j with $j \geq p - 9$ matter in the outcome equation but are irrelevant in the treatment assignment rule. Given that the propensity score is correctly specified, the four estimators are consistent. The oracle estimator should be the most accurate, as it incorporates the information on which covariates matter in the propensity score. Because the balancing program misses some covariates that are relevant for the outcome variable (namely, the X_j with $j \geq p - 9$), the naive plug-in estimator is expected to be asymptotically biased. The immunized procedure should correct for this bias.

Table 1 display the results. We consider several values of n and p that approximate a relatively high-dimensional setting. For every couple (n, p) , we report the root mean squared error (RMSE), the bias and the coverage rate of the confidence intervals associated to each estimator. Our estimator performs well in all settings, with a correct coverage rate and often the lowest RMSE over all estimators. The oracle has always a coverage rate close to 0.95 and a bias very close to 0, as one could expect, but it does not always exhibit the lowest RMSE. This is because, intuitively, the immunized estimator and that of Farrell (2015) trade off variance with some bias in their first steps, sometimes resulting in slightly lower RMSE on the final estimator. Note that the bias of the

final estimator, though asymptotically negligible, results in a slight undercoverage of the confidence intervals. Yet, even with $n = 500$ and $p = 1,000$, the coverage rate is still of 0.84, much higher than that of the naive plug-in estimator (0.485). This estimator exhibits a large bias and a low coverage rate even for $p = 50$ and $n = 2,000$. This shows the importance of correcting for the bias of the first-step estimator, even when p/n is quite small. Finally, the estimator of Farrell (2015) exhibit similar performances as those of the immunized estimator, though it displays a slightly larger RMSE and smaller coverage rate with this particular DGP.

Table 1: Monte-Carlo simulations

	$n = 500$			$n = 1,000$			$n = 2,000$		
	RMSE	Bias	CR	RMSE	Bias	CR	RMSE	Bias	CR
				$p = 50$					
Naive plug-in	0.312	0.264	0.62	0.228	0.195	0.601	0.169	0.144	0.608
Immunized	0.186	0.102	0.872	0.121	0.059	0.907	0.084	0.03	0.907
Farrell	0.197	0.117	0.857	0.132	0.075	0.885	0.093	0.046	0.887
Oracle	0.202	-0.017	0.929	0.143	-0.025	0.938	0.105	-0.03	0.935
				$p = 200$					
Naive plug-in	0.318	0.274	0.587	0.238	0.208	0.557	0.179	0.156	0.544
Immunized	0.185	0.108	0.883	0.125	0.066	0.897	0.085	0.037	0.911
Farrell	0.195	0.121	0.87	0.135	0.081	0.873	0.095	0.052	0.882
Oracle	0.194	-0.023	0.936	0.141	-0.029	0.946	0.108	-0.034	0.934
				$p = 500$					
Naive plug-in	0.339	0.297	0.532	0.247	0.216	0.522	0.185	0.165	0.494
Immunized	0.202	0.128	0.841	0.128	0.07	0.881	0.087	0.044	0.912
Farrell	0.212	0.141	0.81	0.14	0.086	0.854	0.098	0.059	0.866
Oracle	0.205	-0.019	0.927	0.146	-0.033	0.931	0.105	-0.032	0.938
				$p = 1,000$					
Naive plug-in	0.345	0.309	0.485	0.258	0.23	0.478	0.194	0.175	0.449
Immunized	0.199	0.133	0.835	0.135	0.082	0.862	0.09	0.051	0.885
Farrell	0.211	0.146	0.807	0.146	0.097	0.823	0.101	0.066	0.853
Oracle	0.194	-0.017	0.947	0.14	-0.022	0.943	0.103	-0.026	0.939

Notes: RMSE and CR stand respectively for root mean squared error and coverage rate. The nominal coverage rate is 0.95. The results are based on 10,000 simulations for each (n, p) . The naive plug-in and immunized estimators are defined in (3.5) and (3.11), respectively. “Farrell” is the estimator considered by Farrell (2015). “Oracle” is defined by (3.2), where in the first step we only include the ten covariates affecting the treatment.

5 Empirical Applications

5.1 Job Training Program

We revisit LaLonde (1986), who examines the ability of econometric methods to recover the causal effect of employment programs.⁴ This dataset was first built to assess the impact of the National Supported Work (NSW) program. The NSW is a transitional, subsidized work experience program targeted towards people with longstanding employment problems: ex-offenders, former drug addicts, women who were long-term recipients of welfare benefits and school dropouts. The quantity of interest is the average effect of the program for the participants on 1978 yearly earnings. The treated group gathers people who were randomly assigned to this program from the population at risk (with a sample size of $n_1 = 185$). Two control groups are available. The first one comes from the Panel Study of Income Dynamics (PSID) (sample size $n_0 = 2,490$). The second one comes from the experiment (sample size $n_0 = 260$) and is therefore directly comparable to the treated group. It provides us with a benchmark for the ATT. Hereafter, we use the group of participants and the PSID sample to compute our estimator and compare it with other competitors and the experimental benchmark.

To allow for a flexible specification, we follow Farrell (2015) by taking the raw covariates of the dataset (age, education, black, hispanic, married, dummy variable of no degree, income in 1974, income in 1975, dummy variables of no earnings in 1974 and in 1975), two-by-two-interactions between the continuous and dummy variables, two-by-two interactions between the dummy variables and powers up to degree 5 of the continuous variables. Continuous variables are linearly rescaled to $[0, 1]$. We end up with 172 variables to select from. The experimental benchmark for the ATT estimate is \$1,794, with a standard error of 671. We compare several estimators: the naive plug-in estimator, the immunized plug-in estimator, the doubly-robust estimator of Farrell (2015), the double-post-selection linear estimator of Belloni et al. (2014), and the plain OLS estimator including all the covariates. For the four penalized estimators, the penalty loadings on the first-step estimators are set as in Appendix A.

Table 2 displays the results. Our immunized estimator and that of Farrell (2015) give credible values for the ATT with respect to the experimental benchmark, with also similar standard errors.⁵ Notably, they are the only ones, out of five estimators, that display a significant, positive impact, as the experimental benchmark. The immunized estimator estimator offers a substantial improvement on bias and standard error over the naive plug-in estimator, in line with the evidence from the Monte Carlo experiment. Finally, the OLS estimator in Column (7) presents a benchmark of a

⁴For more discussion on the NSW program and the controversy regarding econometric estimates of causal effects based on nonexperimental data, see LaLonde (1986) and the subsequent contributions by Dehejia and Wahba (1999, 2002); Smith and Todd (2005).

⁵Farrell (2015)'s estimate shown in Table 2 differs from that displayed in Farrell's paper because contrary to him, we have not automatically included education, the dummy of no degree and the 1974 income in the set of theory pre-selected covariates. When doing so, the results are slightly better but not qualitatively different for this estimator. We chose not to do so as it would bias the comparison with other estimators, which do not include a set of pre-selected variables.

very simple model that does not use any selection at all. Not surprisingly, this estimator has a substantially larger standard error than the others.

Table 2: ATT estimates on the NSW program.

	Estimator		
	Experimental benchmark (1)	Naive plug-in (2)	Immunized estimator (3)
Point estimate	1,794.34	401.89	1,608.99
Standard error	(671.00)	(746.07)	(705.38)
95% confidence interval	[519; 3,046]	[-1,060; 1,864]	[226; 2991]
# variables in propensity score	none	9	9
# variables in outcome equation	none	none	12

	Farrell (2015) (4)	BCH (2014) (5)	OLS estimator (6)
Point estimate	1420.43	226.82	83.17
Standard error	(670.32)	(867.51)	(1,184.48)
95% confidence interval	[107; 2734]	[-1473; 1927]	[-2,238; 2,405]
# variables in propensity score	3	16	none
# variables in outcome equation	13	10	172


Notes: For details on the estimators, see the text. Standard errors and confidence intervals are based on the asymptotic distribution.

5.2 California Tobacco Control Program

Proposition 99 is one of the first and most ambitious large-scale tobacco control program, implemented in 1989 in California. It includes a vast array of measures, including an increase in cigarette taxation of 25 cents per pack, and a significant effort in prevention and education. In particular, the tax revenues generated by Proposition 99 were used to fund anti-smoking campaigns. Abadie et al. (2010) analyzes the impact of the law on tobacco consumption in California. Since this program was only enforced in California, it is a nice example where the synthetic control method applies whereas more standard public policy evaluation tools cannot be used. It is possible to reproduce a synthetic California by reweighting other states so as to imitate California's behavior.

For this purpose, Abadie et al. (2010) consider the following covariates: retail price of cigarettes, state log income per capita, percentage of the population between 15 and 24, per capita beer consumption (all 1980-1988 averages). Cigarette consumptions for the years between 1970 and 1975, 1980 and 1988 are also included. Using the same variables, we conduct the same analysis with our estimator. Figure 2 displays the estimated effect of Proposition 99 using the immunized

estimator.

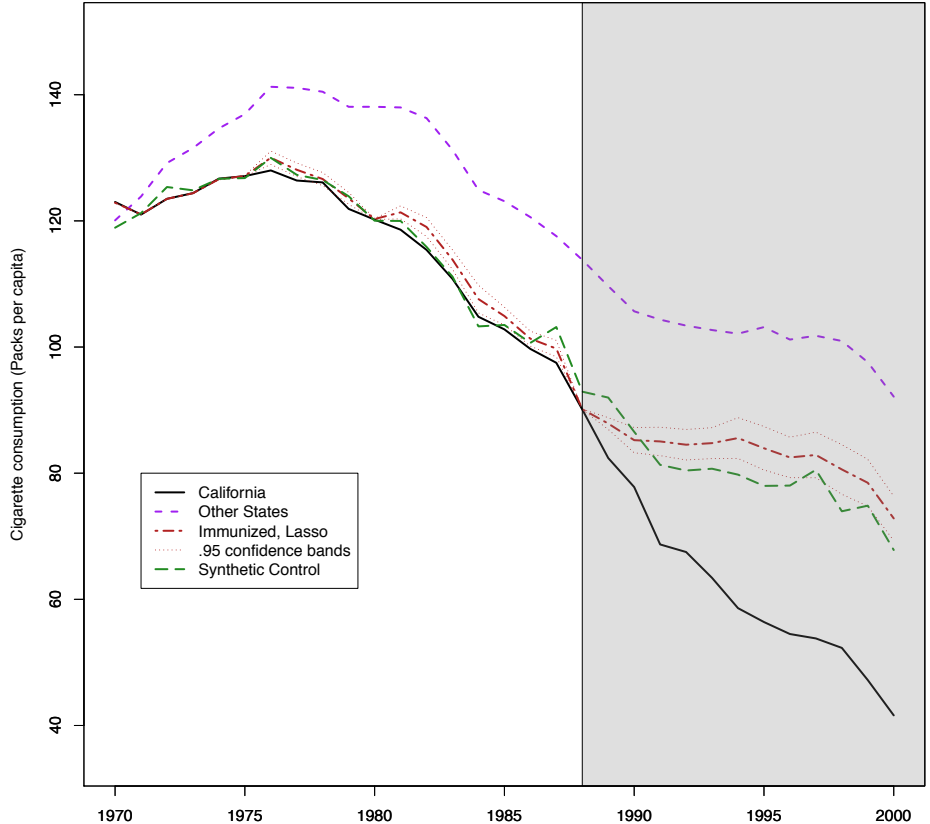


Notes: the shaded area represents the post-treatment period. The 95% confidence interval is based on the asymptotic approximation.

Figure 2: The effect of Proposition 99 on per capita tobacco consumption.

We find almost no effect of the policy over the pre-treatment period, giving credibility to the counterfactual employed. A steady decline takes place after 1988, and in the long-run, tobacco consumption is estimated to have decreased by about 30 packs per capita per year in California as a consequence of the policy. The variance is larger towards the end of the period because covariates are measured in the pre-treatment period and they become less relevant as predictors. Note also that by construction, including 1970 to 1975, 1980 and 1988 cigarette consumptions among the covariates yields an almost perfect fit at these dates because of the immunization procedure. The fit is not perfect, however, because of the shrinkage induced by the ℓ_1 -penalization.

Figure 3 displays a comparison between the immunized estimator and the synthetic control method. The dashed green line is the synthetic control counterfactual. It does not match exactly the plot of Abadie et al. (2010), in which the weights given to each predictors are optimized to fit best the outcome over the whole pre-treatment period. Instead, the green curve in Figure 3 optimizes the predictor weights using only the years 1970 through 1975, 1980 and 1988. This strategy brings a fairer comparison with our estimator that does not use California’s per capita tobacco consumption outside those dates to optimize the fit. In such a case, the years from 1976 to 1987, excluding 1980, can be used as sort of placebo tests.



Notes: The solid black line is California tobacco consumption as in the data. The dotted purple line is a simple average of other U.S. states. The dashed red line is the immunized estimator as presented in the paper along with the 95% confidence bands. The dashed green line is synthetic California.

Figure 3: Cigarette consumption in California, actual and counterfactual

Both our estimator and the synthetic control are credible counterfactuals, as they are able to closely match California pre-treatment tobacco consumption. They offer a sizable improvement over a sample average over the U.S. that did not implement any tobacco control program. Furthermore, even if our estimator gives a result relatively similar to the synthetic control, it displays a smoother pattern especially towards the end of the 1980s. The estimated treatment effect appears to be larger with the immunized estimate than with the synthetic control. However, it is hard to conclude that this difference is significant, because absent any asymptotic theory on the synthetic control estimator, it is unclear how one could make a test on the difference between the two. In fact, the availability of standard asymptotic approximation for confidence intervals is to the advantage of our method.

6 Conclusion

In this paper, we propose an estimator that makes a link between the synthetic control method, typically used with aggregated data and n_0 smaller than or of the same order as p , and treatment effect methods used for micro data for which $n_0 \gg p$. Our method accommodates both settings. In

the low-dimensional regime, it pins down one of the solutions of the synthetic control problem, which admits an infinity of solutions. In the high-dimensional regime, the estimator is a regularized and immunized version of the low-dimensional one and then differs from the synthetic control estimator. The simulations and applications suggest that it works well in practice.

In our study, we have focused on specific procedures based on ℓ_1 -penalization and proved that they achieve good asymptotic behavior in possibly high-dimensional regime under sparsity restrictions. Other types of estimators could be explored using these ideas. For example, in the high-dimensional regime, our strategy can be used with the whole spectrum of sparsity-related penalization techniques, such as group Lasso, fused Lasso, adaptive Lasso, Slope, among others.

References

- Abadie, A. (2019), ‘Using synthetic controls: Feasibility, data requirements, and methodological aspects’, *Journal of Economic Literature* **forthcoming**.
- Abadie, A., Diamond, A. and Hainmueller, J. (2010), ‘Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California’s Tobacco Control Program’, *Journal of the American Statistical Association* **105**(490), 493–505.
- Abadie, A., Diamond, A. and Hainmueller, J. (2015), ‘Comparative politics and the synthetic control method’, *American Journal of Political Science* **59**(2), 495–510.
- Abadie, A. and Gardeazabal, J. (2003), ‘The Economic Costs of Conflict: A Case Study of the Basque Country’, *American Economic Review* **93**(1), 113–132.
- Abadie, A. and L’Hour, J. (2019), A penalized synthetic control estimator for disaggregated data. Working Paper.
- Arkhangelsky, D., Athey, S., Hirshberg, D. A., Imbens, G. W. and Wager, S. (2019), Synthetic Difference in Differences, Working Paper wp2019_907, CEMFI.
- Athey, S., Imbens, G. W. and Wager, S. (2018), ‘Approximate residual balancing: De-biased inference of average treatment effects in high dimensions’. arXiv:1604.07125.
- Bang, H. and Robins, J. M. (2005), ‘Doubly robust estimation in missing data and causal inference models’, *Biometrics* **61**(4), 962–973.
- Bellec, P. C., Lecué, G. and Tsybakov, A. B. (2018), ‘Slope meets Lasso: Improved oracle bounds and optimality’, *Annals of Statistics* **46**(6B), 3603–3642.
- 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**(6), 2369–2429.
- Belloni, A. and Chernozhukov, V. (2013), ‘Least squares after model selection in high-dimensional sparse models’, *Bernoulli* **19**(2), 521–547.
- Belloni, A., Chernozhukov, V., Fernández-Val, I. and Hansen, C. (2017), ‘Program evaluation and causal inference with high-dimensional data’, *Econometrica* **85**(1), 233–298.
- Belloni, A., Chernozhukov, V. and Hansen, C. (2014), ‘Inference on treatment effects after selection among high-dimensional controls’, *The Review of Economic Studies* **81**(2), 608–650.
- Ben-Michael, E., Feller, A. and Rothstein, J. (2018), ‘The Augmented Synthetic Control Method’, *arXiv e-prints* p. arXiv:1811.04170.
- Bickel, P., Ritov, Y. and Tsybakov, A. B. (2009), ‘Simultaneous analysis of Lasso and Dantzig selector’, *Annals of Statistics* **37**(4), 1705–1732.

- Bléhaut, M., D'Haultfoeuille, X., L'Hour, J. and Tsybakov, A. B. (2017), 'A parametric generalization of the synthetic control method, with high dimension', *Unpublished, 2017 IAAE Meeting presentation* pp. 1–53.
URL: <https://editorialexpress.com/conference/IAAE2017/program/IAAE2017.html>
- Bradic, J., Wager, S. and Zhu, Y. (2019), 'Sparsity double robust inference of average treatment effects'. arXiv preprint arXiv:1905.00744.
- Bunea, F., Tsybakov, A. B. and Wegkamp, M. H. (2007), 'Sparsity oracle inequalities for the Lasso.', *Electronic Journal of Statistics* **1**, 169–194.
- 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', *The Econometrics Journal* **21**(1), C1–C68.
- Chernozhukov, V., Hansen, C. and Spindler, M. (2015a), 'Post-Selection and Post-Regularization Inference in Linear Models with Many Controls and Instruments', *American Economic Review* **105**(5), 486–90.
- Chernozhukov, V., Hansen, C. and Spindler, M. (2015b), 'Valid post-selection and post-regularization inference: An elementary, general approach', *Annual Review of Economics* **7**(1), 649–688.
- de la Peña, V. H., Lai, T. L. and Shao, Q.-M. (2009), *Self-Normalized Processes: Limit Theory and Statistical Applications*, 1st edn, Springer-Verlag Berlin Heidelberg.
- Dehejia, R. H. and Wahba, S. (1999), 'Causal effects in nonexperimental studies: Reevaluating the evaluation of training programs', *Journal of the American Statistical Association* **94**(448), pp. 1053–1062.
- Dehejia, R. H. and Wahba, S. (2002), 'Propensity score-matching methods for nonexperimental causal studies', *The Review of Economics and Statistics* **84**(1), 151–161.
- Deville, J.-C. and Särndal, C.-E. (1992), 'Calibration estimators in survey sampling', *Journal of the American statistical Association* **87**(418), 376–382.
- Farrell, M. H. (2015), 'Robust inference on average treatment effects with possibly more covariates than observations', *Journal of Econometrics* **189**(1), 1 – 23.
- Graham, B. S., Pinto, C. C. D. X. and Egel, D. (2012), 'Inverse Probability Tilting for Moment Condition Models with Missing Data', *Review of Economic Studies* **79**(3), 1053–1079.
- Hainmueller, J. (2012), 'Entropy balancing for causal effects: A multivariate reweighting method to produce balanced samples in observational studies', *Political Analysis* **20**(1), 25–46.
- Hastie, T., Tibshirani, R. and Friedman, J. (2009), *The Elements of Statistical Learning: Data mining, Inference and Prediction*, 2nd edn, Springer.

- Imai, K. and Ratkovic, M. (2014), ‘Covariate balancing propensity score’, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **76**(1), 243–263.
- Kline, P. (2011), ‘Oaxaca-blinder as a reweighting estimator’, *American Economic Review* **101**(3), 532–37.
- Klößner, S., Kaul, A., Pfeifer, G. and Schieler, M. (2018), ‘Comparative politics and the synthetic control method revisited: A note on abadie et al.(2015)’, *Swiss journal of economics and statistics* **154**(1), 11.
- LaLonde, R. J. (1986), ‘Evaluating the Econometric Evaluations of Training Programs with Experimental Data’, *American Economic Review* **76**(4), 604–20.
- Leamer, E. E. (1983), ‘Let’s take the con out of econometrics’, *The American Economic Review* **73**(1), 31–43.
- Leeb, H. and Pötscher, B. M. (2005), ‘Model Selection And Inference: Facts And Fiction’, *Econometric Theory* **21**(01), 21–59.
- Leeb, H. and Pötscher, B. M. (2008a), ‘Recent developments in model selection and related areas’, *Econometric Theory* **24**, 319–322.
- Leeb, H. and Pötscher, B. M. (2008b), ‘Sparse estimators and the oracle property, or the return of Hodges’ estimator’, *Journal of Econometrics* **142**(1), 201–211.
- Newey, W. K. and McFadden, D. (1994), Chapter 36 large sample estimation and hypothesis testing, in ‘Handbook of Econometrics’, Vol. 4 of *Handbook of Econometrics*, Elsevier, pp. 2111 – 2245.
- Rudelson, M. (2020), Personal communication.
- Rudelson, M. and Zhou, S. (2013), ‘Reconstruction from anisotropic random measurements’, *IEEE Transactions on Information Theory* **59**(6), 3434–3447.
- Smith, J. and Todd, P. (2005), ‘Does matching overcome LaLonde’s critique of nonexperimental estimators?’, *Journal of Econometrics* **125**(1-2), 305–353.
- Tibshirani, R. J. (2013), ‘The lasso problem and uniqueness.’, *Electronic Journal of Statistics* **7**, 1456–1490.
- van de Geer, S. A. (2016), *Estimating and Testing under Sparsity*, Springer.

A Algorithm for Feasible Penalty Loadings

Consider the ideal penalty loadings for estimation of β_0 defined as

$$\lambda := c\Phi^{-1}(1 - \gamma/2p)/\sqrt{n}$$

$$\bar{\psi}_j := \sqrt{\frac{1}{n} \sum_{i=1}^n [(1 - D_i)h(X_i^T \beta_0) - D_i]^2 X_{i,j}^2} \text{ for } j = 1, \dots, p,$$

and the ideal penalty loadings for estimation of μ_0 :

$$\lambda' := 2c\Phi^{-1}(1 - \gamma/2p)/\sqrt{n}$$

$$\bar{\psi}'_j := \sqrt{\frac{1}{n} \sum_{i=1}^n (1 - D_i)h'(X_i^T \beta_0)^2 [Y_i - X_i^T \mu_0]^2 X_{i,j}^2} \text{ for } j = 1, \dots, p.$$

Here $c > 1$ is an absolute constant, $\gamma > 0$ is a tuning parameter while β_0 and μ_0 are the true coefficients. We follow Belloni et al. (2014) and set $\gamma = .05$ and $c = 1.1$.

We first estimate the ideal penalty loadings $\{\bar{\psi}_j\}_{j=1}^p$ of the balancing step using the following algorithm. Set a small constant $\epsilon > 0$ and a maximal number of iterations k_0 .

1. Start by using a preliminary estimate $\beta^{(0)}$ of β_0 . For example, take $\beta^{(0)}$ with the entry corresponding to the intercept equal to $\log(\sum_{i=1}^n D_i / \sum_{i=1}^n (1 - D_i))$ and all other entries equal to zero. Then, for all $j = 1, \dots, p$, set

$$\tilde{\psi}_j^{(0)} = \sqrt{\frac{1}{n} \sum_{i=1}^n [(1 - D_i)h(X_i^T \beta^{(0)}) - D_i]^2 X_{i,j}^2}.$$

At step k , set $\tilde{\psi}_j^{(k)} = \sqrt{\frac{1}{n} \sum_{i=1}^n [(1 - D_i)h(X_i^T \beta^{(k)}) - D_i]^2 X_{i,j}^2}$, $j = 1, \dots, p$.

2. Estimate the model by the penalized balancing Equation (3.4) using the penalty level λ and penalty loadings found previously, to obtain $\hat{\beta}^{(k)}$.
3. Stop if $\max_{j=1, \dots, p} |\tilde{\psi}_j^{(k)} - \tilde{\psi}_j^{(k-1)}| \leq \epsilon$ or if $k > k_0$. Set $k = k + 1$ and go to step 1 otherwise.

Asymptotic validity of this approach is established analogously to (Belloni et al., 2012, Lemma 11). Estimation of the penalty loadings $\bar{\psi}'_j$ on the immunization step follows a similar procedure where we replace β_0 in the formula for $\bar{\psi}'_j$ by its estimator obtained on the balancing step.

B Proofs

B.1 Proof of Theorem 2.1

First, note that we have $\mathbb{E}[(1 - D)w(X)X] = \mathbb{E}[DX]$. As a result, for any $\mu \in \mathbb{R}^p$,

$$\mathbb{E}[(D - (1 - D)w(X))(Y - X^T \mu)] = \mathbb{E}[(D - (1 - D)w(X))Y].$$

Since $(1 - D)Y = (1 - D)Y(0)$ and $DY = DY(1)$, the value θ_0 satisfies the moment condition (2.8) if and only if

$$\mathbb{E}[w(X)Y(1 - D)] = \mathbb{E}[DY(0)].$$

By the Mean Independence assumption, $\mathbb{E}(D|X)\mathbb{E}(Y(0)|X) = \mathbb{E}(DY(0)|X)$. Thus,

$$\mathbb{E}[w(X)Y(1 - D)] = \mathbb{E}[\mathbb{E}(w(X)Y(0)(1 - D)|X)] = \mathbb{E}[w(X)(1 - \mathbb{E}(D|X))\mathbb{E}(Y(0)|X)]. \quad (\text{B.1})$$

We consider the two cases of the theorem separately.

1. In the linear case $\mathbb{E}(Y(0)|X) = X^T\mu_0$ we have

$$\begin{aligned} \mathbb{E}[w(X)Y(1 - D)] &= \mathbb{E}[w(X)(1 - D)X^T\mu_0] \\ &= \mathbb{E}[DX^T\mu_0] \\ &= \mathbb{E}[D\mathbb{E}(Y(0)|X)] \\ &= \mathbb{E}[DY(0)]. \end{aligned}$$

The first equality here is due to (B.1). The second equality follows from the fact that $\mathbb{E}[(1 - D)w(X)X] = \mathbb{E}[DX]$. The last equality uses the Mean Independence assumption.

2. Propensity score satisfies $P(D = 1|X) = w(X)/(1 + w(X))$. In this case, using (B.1) we have

$$\begin{aligned} \mathbb{E}[w(X)Y(1 - D)] &= \mathbb{E}[w(X)(1 - P(D = 1|X))\mathbb{E}(Y(0)|X)] \\ &= \mathbb{E}[P(D = 1|X)\mathbb{E}(Y(0)|X)] \\ &= \mathbb{E}[\mathbb{E}(D|X)\mathbb{E}(Y(0)|X)] \\ &= \mathbb{E}[DY(0)], \end{aligned}$$

where the last equality follows from the Mean Independence assumption. □

B.2 Proof of Theorem 3.1

Denote the observed data by $Z_i = (Y_i, D_i, X_i)$, and by π_0 the probability of being treated: $\pi_0 := \mathbb{P}(D = 1)$. The estimating moment function for θ_0 is $g(Z, \theta, \eta) := [D - (1 - D)h(X^T\beta)][Y - X^T\mu] - D\theta$. Recall that we define (θ_0, η_0) as the values satisfying:

$$\mathbb{E}g(Z, \theta_0, \eta_0) = 0.$$

All these quantities depend on the sample size n but for the sake of brevity we suppress this dependency in the notation except for the cases when we need it explicitly.

By the Taylor expansion and the linearity of the estimating function g in θ , there exists $t \in (0, 1)$ such that

$$\begin{aligned} \mathbb{E}_n[g(Z, \hat{\theta}, \hat{\eta})] &= \mathbb{E}_n[g(Z, \theta_0, \hat{\eta})] + \hat{\pi}(\theta_0 - \hat{\theta}) \\ &= \hat{\pi}(\theta_0 - \hat{\theta}) + \mathbb{E}_n[g(Z, \theta_0, \eta_0)] + (\hat{\eta} - \eta_0)^T \mathbb{E}_n[\nabla_{\eta}g(Z, \theta_0, \eta_0)] \\ &\quad + \frac{1}{2}(\hat{\eta} - \eta_0)^T \mathbb{E}_n[\nabla_{\eta}^2g(Z, \theta_0, \tilde{\eta})](\hat{\eta} - \eta_0), \end{aligned}$$

where $\tilde{\eta} := t\eta_0 + (1-t)\hat{\eta}$ and $\hat{\pi} = \frac{1}{n} \sum_{i=1}^n D_i$. The immunized estimator satisfies $\mathbb{E}_n[g(Z, \hat{\theta}, \hat{\eta})] = 0$. Thus, we obtain

$$\begin{aligned} \hat{\pi} \sqrt{n}(\hat{\theta} - \theta_0) &= \underbrace{\sqrt{n} \mathbb{E}_n [g(Z, \theta_0, \eta_0)]}_{:=I_1} + \underbrace{\sqrt{n}(\hat{\eta} - \eta_0)^T \mathbb{E}_n [\nabla_{\eta} g(Z, \theta_0, \eta_0)]}_{:=I_2} \\ &\quad + \underbrace{2^{-1} \sqrt{n}(\hat{\eta} - \eta_0)^T \mathbb{E}_n [\nabla_{\eta}^2 g(Z, \theta_0, \tilde{\eta})]}_{:=I_3} (\hat{\eta} - \eta_0). \end{aligned}$$

Now, to prove Theorem 3.1 we proceed as follows. First, we show that I_1 converges in distribution to a zero mean normal random variable with variance $\mathbb{E}g^2(Z, \theta_0, \eta_0)$ (Step 1 below), while I_2 and I_3 tend to zero in probability (Steps 2 and 3). This and the fact that $\hat{\pi} \rightarrow \pi_0$ (a.s.) imply that $\sqrt{n}(\hat{\theta} - \theta_0)$ is asymptotically normal with some positive variance, which in turn implies that $\hat{\theta} \rightarrow \theta_0$ in probability. Finally, using this property we prove that $\mathbb{E}_n [g^2(Z_i, \hat{\theta}, \hat{\eta})] \rightarrow \mathbb{E}g^2(Z, \theta_0, \eta_0)$ in probability (Step 4). Combining Steps 1 to 4 and using Slutsky lemma leads to the result of the theorem. Thus, to complete the proof of the theorem it remains to establish Steps 1 to 4.

Step 1

In this part, we write $g_{i,n} := g(Z_i, \theta_0, \eta_0)$ (making the dependence on n explicit in the notation). Recall that $\mathbb{E}g_{i,n} = 0$. We apply the Lindeberg-Feller central limit theorem for triangular arrays by checking a Lyapunov condition. Since $(g_{i,n})_{1 \leq i \leq n}$ are i.i.d., it suffices to prove that

$$\limsup_{n \rightarrow \infty} \frac{\mathbb{E}(g_{1,n}^{2+\delta})}{(\mathbb{E}g_{1,n}^2)^{1+\delta/2}} < \infty \quad (\text{B.2})$$

for some $\delta > 0$. Assumption 3.3(i) implies that $\mathbb{E}(g_{1,n}^{2+\delta})$ is bounded uniformly in n . Moreover,

$$\mathbb{E}(g_{1,n}^2) = \pi_0 \mathbb{E}((Y(1) - X^T \mu_0 - \theta_0)^2 | D = 1) + (1 - \pi_0) \mathbb{E}(h(X^T \beta_0)^2 (Y(0) - X^T \mu_0)^2 | D = 0).$$

Due to Assumption 3.3(ii) we have $\mathbb{E}(g_{1,n}^2) \geq c_1$, where $c_1 > 0$ does not depend on n . Thus, (B.2) holds. It follows that $(\mathbb{E}(g_{1,n}^2))^{-1/2} I_1$ converges in distribution to a standard normal random variable.

Step 2

Set $\psi_{j+p} = \psi'_j, j = 1, \dots, p$, and denote by Ψ a diagonal matrix of dimension $2p$ with diagonal elements $\psi_j, j = 1, \dots, 2p$. Let also $\bar{\Psi}$ be a diagonal matrix of dimension $2p$ with diagonal elements $\bar{\psi}_j = \sqrt{n^{-1} \sum_{i=1}^n U_{i,j}^2}, j = 1, \dots, 2p$. By Assumption 3.6 we have $\psi_j \geq \bar{\psi}_j$. Hence,

$$|I_2| \leq \|\Psi(\hat{\eta} - \eta_0)\|_1 \|\Psi^{-1} \sqrt{n} \mathbb{E}_n [\nabla_{\eta} g(Z, \theta_0, \eta_0)]\|_{\infty} \leq \|\Psi(\hat{\eta} - \eta_0)\|_1 \|\bar{\Psi}^{-1} \sqrt{n} \mathbb{E}_n [\nabla_{\eta} g(Z, \theta_0, \eta_0)]\|_{\infty}.$$

Here, the term $\|\bar{\Psi}^{-1} \sqrt{n} \mathbb{E}_n [\nabla_{\eta} g(Z, \theta_0, \eta_0)]\|_{\infty}$ is a maximum of self-normalized sums of variables $U_{i,j}$ and it can be bounded by using standard inequalities for self-normalized sums, cf. Lemma C.1 below. From the orthogonality conditions and Assumption 3.3 we have $\mathbb{E}(U_{i,j}) = 0, \mathbb{E}(U_{i,j}^2) \geq c_1$

and $\mathbb{E}(|U_{i,j}|^3) < \infty$ for any i and any j . By Lemma C.1 and the fact that $\Phi^{-1}(1-a) \leq \sqrt{-2\log(a)}$ for all $a \in (0, 1)$ we obtain that, with probability tending to 1 as $n \rightarrow \infty$,

$$\|\bar{\Psi}^{-1}\sqrt{n}\mathbb{E}_n[\nabla_\eta g(Z, \theta_0, \eta_0)]\|_\infty \leq \Phi^{-1}(1-\gamma/2p) \leq \sqrt{2\log(2p/\gamma)} \lesssim \sqrt{\log(p)}.$$

Next, inequalities (B.18) and (B.23) imply that $\|\Psi(\hat{\eta} - \eta_0)\|_1 \lesssim (s/\kappa_\Sigma)\sqrt{\log(p)/n}$ with probability tending to 1 as $n \rightarrow \infty$. Using these facts and the growth condition (i) in Assumption 3.5 we conclude that I_2 converges to 0 in probability as $n \rightarrow \infty$.

Step 3

Let $\mathbf{H} := \mathbb{E}_n[\nabla_\eta^2 g(Z, \theta_0, \tilde{\eta})] \in \mathbb{R}^{2p \times 2p}$ and let $h_{k,j}$ be the elements of matrix \mathbf{H} . We have

$$|I_3| \leq \frac{\sqrt{n}}{2} \|\hat{\eta} - \eta_0\|_1^2 \max_{1 \leq k, j \leq 2p} |h_{k,j}|. \quad (\text{B.3})$$

We now control the random variable $\max_{1 \leq k, j \leq 2p} |h_{k,j}|$. To do this, we first note that

$$\begin{aligned} \frac{\partial^2}{\partial \beta \partial \beta^T} g(Z, \theta, \eta) &= -(1-D)h''(X^T \beta) [Y - X^T \mu] X X^T, \\ \frac{\partial^2}{\partial \mu \partial \beta^T} g(Z, \theta, \eta) &= \frac{\partial^2}{\partial \beta \partial \mu^T} g(Z, \theta, \eta) = (1-D)h'(X^T \beta) X X^T, \\ \frac{\partial^2}{\partial \mu \partial \mu^T} g(Z, \theta, \eta) &= 0. \end{aligned}$$

It follows that

$$\max_{1 \leq k, j \leq 2p} |h_{k,j}| \leq \max \left(\max_{1 \leq k, j \leq p} |\tilde{h}_{k,j}|, \max_{1 \leq k, j \leq p} |\bar{h}_{k,j}| \right), \quad (\text{B.4})$$

where

$$\tilde{h}_{k,j} = \frac{1}{n} \sum_{i=1}^n (1-D_i) h''(X_i^T \tilde{\beta}) (Y_i - X_i \tilde{\mu}) X_{i,k} X_{i,j},$$

and

$$\bar{h}_{k,j} = \frac{1}{n} \sum_{i=1}^n (1-D_i) h'(X_i^T \tilde{\beta}) X_{i,k} X_{i,j}.$$

We now evaluate separately the terms $\max_{1 \leq k, j \leq p} |\tilde{h}_{k,j}|$ and $\max_{1 \leq k, j \leq p} |\bar{h}_{k,j}|$.

Note that $\tilde{h}_{k,j}$ can be decomposed as

$$\tilde{h}_{k,j} = \tilde{h}_{k,j,1} + \tilde{h}_{k,j,2} + \tilde{h}_{k,j,3} + \tilde{h}_{k,j,4},$$

where

$$\begin{aligned} \tilde{h}_{k,j,1} &= n^{-1} \sum_{i=1}^n (1-D_i) h''(X_i^T \beta_0) (Y_i - X_i^T \mu_0) X_{i,k} X_{i,j}, \\ \tilde{h}_{k,j,2} &= n^{-1} \sum_{i=1}^n (1-D_i) h''(X_i^T \beta_0) X_i^T (\mu_0 - \tilde{\mu}) X_{i,k} X_{i,j}, \\ \tilde{h}_{k,j,3} &= n^{-1} \sum_{i=1}^n (1-D_i) (h''(X_i^T \tilde{\beta}) - h''(X_i^T \beta_0)) (Y_i - X_i^T \mu_0) X_{i,k} X_{i,j}, \\ \tilde{h}_{k,j,4} &= n^{-1} \sum_{i=1}^n (1-D_i) (h''(X_i^T \tilde{\beta}) - h''(X_i^T \beta_0)) X_i^T (\mu_0 - \tilde{\mu}) X_{i,k} X_{i,j}. \end{aligned}$$

It follows from Assumption 3.3 that, for all k, j ,

$$|\tilde{h}_{k,j,2}| \leq C \|\mu_0 - \hat{\mu}\|_1. \quad (\text{B.5})$$

Here and in what follows we denote by C positive constants depending only on K that can be different on different appearances. Next, from Assumptions 3.3 and 3.2 (i) we obtain that, for all k, j ,

$$|\tilde{h}_{k,j,3}| \leq C \|\beta_0 - \hat{\beta}\|_1, \quad |\tilde{h}_{k,j,4}| \leq C \|\beta_0 - \hat{\beta}\|_1 \|\mu_0 - \hat{\mu}\|_1. \quad (\text{B.6})$$

From (B.5), (B.6), Theorem 3.2 and the growth condition (i) in Assumption 3.5 we find that, with probability tending to 1 as $n \rightarrow \infty$,

$$\max_{1 \leq k, j \leq p} (|\tilde{h}_{k,j,2}| + |\tilde{h}_{k,j,3}| + |\tilde{h}_{k,j,4}|) \lesssim (s/\kappa_\Sigma) \sqrt{\log(p)/n} \lesssim 1. \quad (\text{B.7})$$

Next, again from Assumption 3.3, we deduce that $|\mathbb{E}(\tilde{h}_{k,j,1})| \leq C$, while by Hoeffding's inequality

$$\mathbb{P}(|\tilde{h}_{k,j,1} - \mathbb{E}(\tilde{h}_{k,j,1})| \geq x) \leq 2 \exp(-Cnx^2), \quad \forall x > 0.$$

This and the union bound imply that there exists $C > 0$ large enough such that, with probability tending to 1 as $n \rightarrow \infty$,

$$\max_{1 \leq k, j \leq p} |\tilde{h}_{k,j,1}| \leq C(1 + \sqrt{\log(p)/n}) \lesssim 1. \quad (\text{B.8})$$

Finally, combining (B.7) and (B.8) we obtain that, with probability tending to 1 as $n \rightarrow \infty$,

$$\max_{1 \leq k, j \leq p} |\tilde{h}_{k,j}| \lesssim 1.$$

Quite similarly, we get that with probability tending to 1 as $n \rightarrow \infty$,

$$\max_{1 \leq k, j \leq p} |\bar{h}_{k,j}| \lesssim 1.$$

Thus, with probability tending to 1 as $n \rightarrow \infty$ we have $\max_{1 \leq k, j \leq 2p} |h_{k,j}| \lesssim 1$. On the other hand, $\|\hat{\eta} - \eta_0\|_1 \lesssim (s/\kappa_\Sigma) \sqrt{\log(p)/n}$ with probability tending to 1 as $n \rightarrow \infty$ due to Theorem 3.2. Using these facts together with (B.3) and the growth condition (i) in Assumption 3.5 we conclude that I_3 tends to 0 in probability as $n \rightarrow \infty$.

Step 4

We now prove that if $\hat{\theta} \rightarrow \theta_0$ in probability then $\frac{1}{n} \sum_{i=1}^n g^2(Z_i, \hat{\theta}, \hat{\eta}) \rightarrow \mathbb{E}g^2(Z, \theta_0, \eta_0)$ in probability as $n \rightarrow \infty$. We have

$$g(Z, \theta, \eta) = [D - (1 - D)h(X^T \beta)][Y - X^T \mu] - D\theta.$$

Theorem 3.2 and the growth condition (i) in Assumption 3.5 imply that $\|\beta_0 - \hat{\beta}\|_1$ is bounded by 1 on an event \mathcal{A}_n of probability tending to 1 as $n \rightarrow \infty$. Using Assumption 3.3 we deduce that, on the event \mathcal{A}_n , the values $X_i^T \beta_0$ and $X_i^T \hat{\beta}$ for all i belong to a subset of \mathbb{R} of diameter at most $2K$. On the other hand, Assumption 3.2 (i) implies that function h is bounded and Lipschitz on any

compact subset of \mathbb{R} . Therefore, using again Assumption 3.3 we find that on the event \mathcal{A}_n we have $|h(X_i^T \beta_0) - h(X_i^T \hat{\beta})| \leq C \|\beta_0 - \hat{\beta}\|_1$ for all i . It follows from this remark and from Assumption 3.3 that, on the event \mathcal{A}_n ,

$$\begin{aligned} |g(Z_i, \hat{\theta}, \hat{\eta}) - g(Z_i, \theta_0, \eta_0)| &\leq |\hat{\theta} - \theta_0| + |h(X_i^T \beta_0) - h(X_i^T \hat{\beta})| |Y - X^T \mu_0| + |h(X_i^T \hat{\beta})| |X_i^T (\mu_0 - \hat{\mu})| \\ &\leq |\hat{\theta} - \theta_0| + C \left(\|\beta_0 - \hat{\beta}\|_1 + \|\mu_0 - \hat{\mu}\|_1 + \|\mu_0 - \hat{\mu}\|_1 \|\beta_0 - \hat{\beta}\|_1 \right) := \zeta_n \end{aligned}$$

for all i . Also note that, due to Assumption 3.3, the random variables $g(Z_i, \theta_0, \eta_0)$ are a.s. uniformly bounded. Thus, using the equality $b^2 - a^2 = (b - a)^2 + 2a(b - a)$, $\forall a, b \in \mathbb{R}$, we get that, on the event \mathcal{A}_n ,

$$|g^2(Z_i, \hat{\theta}, \hat{\eta}) - g^2(Z_i, \theta_0, \eta_0)| \leq C(\zeta_n^2 + \zeta_n).$$

for all i . Using this fact together with Theorem 3.2, the growth condition (i) in Assumption 3.5 and the convergence $\hat{\theta} \rightarrow \theta_0$ in probability we find that $\frac{1}{n} \sum_{i=1}^n g^2(Z_i, \hat{\theta}, \hat{\eta}) - \frac{1}{n} \sum_{i=1}^n g^2(Z_i, \theta_0, \eta_0) \rightarrow 0$ in probability as $n \rightarrow \infty$. We conclude by applying the law of large numbers to the sum $\frac{1}{n} \sum_{i=1}^n g^2(Z_i, \theta_0, \eta_0)$. \square

B.3 Proof of Theorem 3.2

B.3.1 Proof of (3.14)

Recall that $\hat{\beta}$ is defined as:

$$\hat{\beta} \in \arg \min_{\beta \in \mathbb{R}^p} \left(\frac{1}{n} \sum_{i=1}^n [(1 - D_i) H(X_i^T \beta) - D_i X_i^T \beta] + \lambda \sum_{j=1}^p \psi_j |\beta_j| \right) \quad (\text{B.9})$$

with penalty loadings satisfying Assumption 3.6. Let $\Psi_d \in \mathbb{R}^{p \times p}$ be the diagonal matrix with diagonal entries ψ_1, \dots, ψ_p . Let also $\bar{\Psi}_d \in \mathbb{R}^{p \times p}$ be a diagonal matrix with diagonal entries $\bar{\psi}_j = \sqrt{n^{-1} \sum_{i=1}^n U_{i,j}^2}$, $j = 1, \dots, p$. We denote by $S_0 \subseteq \{1, \dots, p\}$ the set of indices of non-zero components of β_0 . By assumption, $\text{Card}(S_0) = \|\beta_0\|_0 \leq s_\beta$.

Step 1: Concentration Inequality

We first bound the sup-norm of the gradient of the objective function using Lemma C.1. Recall that for $1 \leq j \leq p$ we have $U_{i,j} = [(1 - D_i) h(X_i^T \beta_0) - D_i] X_{i,j}$. Set $\mathcal{S}_j := \frac{1}{\sqrt{n}} \sum_{i=1}^n U_{i,j} / \bar{\psi}_j$ and consider the event

$$\mathcal{B} := \left\{ \frac{1}{n} \max_{1 \leq j \leq p} \left| \sum_{i=1}^n \frac{U_{i,j}}{\bar{\psi}_j} \right| \leq \frac{\lambda}{c} \right\}.$$

By construction, the random variables $U_{i,j}$ are i.i.d., $\mathbb{E}(U_{i,j}) = 0$, $\mathbb{E}(U_{i,j}^2) \geq c_1$ and $\mathbb{E}(|U_{i,j}|^3) \leq C$ by Assumptions 3.3 and 3.2. Using these remarks, 3.6 and Lemma C.1 we obtain

$$\begin{aligned} \mathbb{P}(\mathcal{B}^C) &= \mathbb{P} \left(\frac{c}{\sqrt{n}} \max_{1 \leq j \leq p} |\mathcal{S}_j| > c \Phi^{-1}(1 - \gamma/2p) / \sqrt{n} \right) \\ &= \mathbb{P} \left(\max_{1 \leq j \leq p} |\mathcal{S}_j| > \Phi^{-1}(1 - \gamma/2p) \right) = o(1) \text{ as } n \rightarrow \infty. \end{aligned}$$

Step 2: Restricted Eigenvalue condition for the empirical Gram matrix

The empirical Gram matrix is

$$\hat{\Sigma} := \frac{1}{n} \sum_{i=1}^n (1 - D_i) X_i X_i^T = \frac{1}{n} \sum_{i=1}^n (1 - D_i)^2 X_i X_i^T.$$

We also recall the Restricted Eigenvalue (RE) condition (Bickel et al., 2009). For a non-empty subset $S \subseteq \{1, \dots, p\}$ and $\alpha > 0$, define the set:

$$\mathcal{C}[S, \alpha] := \{v \in \mathbb{R}^p : \|v_{S^C}\|_1 \leq \alpha \|v_S\|_1, v \neq 0\} \quad (\text{B.10})$$

where S^C stands for the complement of S . Then, for given $s \in \{1, \dots, p\}$ and $\alpha > 0$, the matrix $\hat{\Sigma}$ satisfies the $\text{RE}(s, \alpha)$ condition if there exists $\kappa(\hat{\Sigma}) > 0$ such that

$$\min_{\substack{S \subseteq \{1, \dots, p\}: \\ \text{Card}(S) \leq s}} \min_{v \in \mathcal{C}[S, \alpha]} \frac{v^T \hat{\Sigma} v}{\|v_S\|_2^2} \geq \kappa(\hat{\Sigma}). \quad (\text{B.11})$$

We now use Lemma C.2, stated and proved in Section C below. Note that Assumption 3.5 implies (C.2) therein and set $V_i = (1 - D_i)X_i$. Then, for any $s \in [1, p/2]$ and $\alpha > 0$, $\hat{\Sigma}$ satisfies the $\text{RE}(s, \alpha)$ condition, with $\kappa(\hat{\Sigma}) = c_* \kappa_\Sigma$ where $c_* \in (0, 1)$ is an absolute constant, with probability tending to 1 as $n \rightarrow \infty$.

Step 3: Basic inequality

At this step, we prove that with probability tending to 1 as $n \rightarrow \infty$, $\hat{\beta}$ satisfies the following inequality (further called the basic inequality):

$$\tau(\hat{\beta} - \beta_0)^T \hat{\Sigma}(\hat{\beta} - \beta_0) \leq 2\lambda \left(\|\Psi_d \beta_0\|_1 - \|\Psi_d \hat{\beta}\|_1 \right) + \frac{2\lambda}{c} \|\Psi_d(\hat{\beta} - \beta_0)\|_1, \quad (\text{B.12})$$

where $\tau > 0$ is a constant that does not depend on n .

By optimality of $\hat{\beta}$ we have

$$\frac{1}{n} \sum_{i=1}^n [\gamma_{\hat{\beta}}(X_i, D_i) - \gamma_{\beta_0}(X_i, D_i)] \leq \lambda \left(\|\Psi_d \beta_0\|_1 - \|\Psi_d \hat{\beta}\|_1 \right),$$

where $\gamma_\beta(X, D) := (1 - D)H(X^T \beta) - DX^T \beta$. Subtracting the inner product of the gradient $\nabla_\beta \gamma_{\beta_0}(X_i, D_i)$ and $\hat{\beta} - \beta_0$ on both sides we find

$$\begin{aligned} & \frac{1}{n} \sum_{i=1}^n [\gamma_{\hat{\beta}}(X_i, D_i) - \gamma_{\beta_0}(X_i, D_i) - ((1 - D_i)h(X_i^T \beta_0) - D_i)(\hat{\beta} - \beta_0)^T X_i] \\ & \leq \lambda \left(\|\Psi_d \beta_0\|_1 - \|\Psi_d \hat{\beta}\|_1 \right) - \frac{1}{n} \sum_{i=1}^n ((1 - D_i)h(X_i^T \beta_0) - D_i)(\hat{\beta} - \beta_0)^T X_i. \end{aligned} \quad (\text{B.13})$$

Using Taylor expansion we get that there exists $0 \leq t \leq 1$ such that

$$\begin{aligned} & \frac{1}{n} \sum_{i=1}^n \gamma_{\hat{\beta}}(X_i, D_i) - \gamma_{\beta_0}(X_i, D_i) - ((1 - D_i)h(X_i^T \beta_0) - D_i) (\hat{\beta} - \beta_0)^T X_i \\ &= \frac{1}{2} (\hat{\beta} - \beta_0)^T \left[\frac{1}{n} \sum_{i=1}^n (1 - D_i) X_i X_i^T h'(X_i^T \tilde{\beta}) \right] (\hat{\beta} - \beta_0), \end{aligned}$$

where $\tilde{\beta} = t\hat{\beta} + (1-t)\beta_0$. Plugging this into (B.13) and using the facts that $|\sum_i a_i b_i| \leq \|a\|_1 \|b\|_\infty$ and $\psi_j \geq \bar{\psi}_j$ we get that, on the event \mathcal{B} , which occurs with probability tending to 1 as $n \rightarrow \infty$:

$$\begin{aligned} & \frac{1}{2} (\hat{\beta} - \beta_0)^T \left[\frac{1}{n} \sum_{i=1}^n (1 - D_i) X_i X_i^T h'(X_i^T \tilde{\beta}) \right] (\hat{\beta} - \beta_0) \tag{B.14} \\ & \leq \lambda \left(\|\Psi_d \beta_0\|_1 - \|\Psi_d \hat{\beta}\|_1 \right) - \frac{1}{n} \sum_{i=1}^n ((1 - D_i)h(X_i^T \beta_0) - D_i) (\hat{\beta} - \beta_0)^T X_i \\ & \leq \lambda \left(\|\Psi_d \beta_0\|_1 - \|\Psi_d \hat{\beta}\|_1 \right) + \max_{1 \leq j \leq p} \left| \frac{1}{n} \sum_{i=1}^n \frac{X_{i,j}}{\bar{\psi}_j} ((1 - D_i)h(X_i^T \beta_0) - D_i) \right| \|\Psi_d(\hat{\beta} - \beta_0)\|_1 \\ & \leq \lambda \left(\|\Psi_d \beta_0\|_1 - \|\Psi_d \hat{\beta}\|_1 \right) + \frac{\lambda}{c} \|\Psi_d(\hat{\beta} - \beta_0)\|_1. \end{aligned}$$

By Assumption 3.2 we have $h' > 0$, which implies that the left-hand side of (B.14) is non-negative. Hence we have, under the event \mathcal{B} ,

$$0 \leq \lambda \left(\|\Psi_d \beta_0\|_1 - \|\Psi_d \hat{\beta}\|_1 \right) + (\lambda/c) \|\Psi_d(\hat{\beta} - \beta_0)\|_1.$$

which implies that

$$\|\Psi_d \hat{\beta}\|_1 \leq c_0 \|\Psi_d \beta_0\|_1 \tag{B.15}$$

where $c_0 = (c+1)/(c-1)$. By Assumption 3.6, we have $\max_j \psi_j \leq \max_j \psi_{j,\max} \leq \bar{\psi}$ where $\bar{\psi} > 0$ is a constant that does not depend on n . On the other hand, Assumption 3.3(ii) and the fact that the random variables $U_{i,j}^2$ are uniformly bounded implies that $\min_j \psi_j \geq \sqrt{c_1/2} := \underline{\psi}$ with probability tending to 1 as $n \rightarrow \infty$ (this follows immediately from Hoeffding's inequality, the union bound and the fact that $\log(p)/n \rightarrow 0$ due to Assumption 3.5(i)). These remarks and (B.15) imply that, with probability tending to 1 as $n \rightarrow \infty$,

$$\|\hat{\beta}\|_1 \leq c_0 \frac{\bar{\psi}}{\underline{\psi}} \|\beta_0\|_1. \tag{B.16}$$

We now use Assumption 3.2(ii). If $\|\beta_0\|_1 \leq c_3$, then $\|\hat{\beta}\|_1 \leq c_0(\bar{\psi}/\underline{\psi})c_3$ with probability tending to 1 as $n \rightarrow \infty$, so that $\min_{i=1,\dots,n} h'(X_i^T \tilde{\beta}) \geq h'(-K \max(1, c_0 \frac{\bar{\psi}}{\underline{\psi}})c_3) > 0$ where we have used Assumptions 3.3(i) and 3.2. Otherwise, given Assumption 3.2(ii), $h' \geq c_2$ on the whole real line so obviously $\min_{i=1,\dots,n} h'(X_i^T \tilde{\beta}) \geq c_2$. It follows that there exists $\tau > 0$ that does not depend on n such that, with probability tending to 1 as $n \rightarrow \infty$,

$$\tau v^T \hat{\Sigma} v \leq v^T \left[\frac{1}{n} \sum_{i=1}^n (1 - D_i) X_i X_i^T h'(X_i^T \tilde{\beta}) \right] v, \quad \forall v \in \mathbb{R}^p. \tag{B.17}$$

Using (B.17) with $v = \hat{\beta} - \beta_0$ and combining it with inequality (B.14) yields (B.12).

Step 4: Control of the ℓ_1 -error for $\hat{\beta}$

We prove that with probability tending to 1 as $n \rightarrow \infty$,

$$\|\Psi_d(\hat{\beta} - \beta_0)\|_1 \leq \left(\frac{c}{c-1}\right) \frac{4\bar{\psi}^2 \lambda s_\beta}{c_* \tau \kappa_\Sigma}, \quad \|\hat{\beta} - \beta_0\|_1 \leq \left(\frac{c}{c-1}\right) \frac{4\bar{\psi}^2 \lambda s_\beta}{\underline{\psi} c_* \tau \kappa_\Sigma}. \quad (\text{B.18})$$

It suffices to prove the first inequality in (B.18). The second inequality follows as an immediate consequence.

We will use the basic inequality (B.12). First, we bound $\|\Psi_d \beta_0\|_1 - \|\Psi_d \hat{\beta}\|_1$. By the triangular inequality,

$$\|\Psi_d \beta_{0, S_0}\|_1 - \|\Psi_d \hat{\beta}_{S_0}\|_1 \leq \|\Psi_d(\beta_{0, S_0} - \hat{\beta}_{S_0})\|_1.$$

Furthermore,

$$\begin{aligned} \|\Psi_d \beta_{0, S_0^c}\|_1 - \|\Psi_d \hat{\beta}_{S_0^c}\|_1 &= 2\|\Psi_d \beta_{0, S_0^c}\|_1 - \|\Psi_d \beta_{0, S_0^c}\|_1 - \|\Psi_d \hat{\beta}_{S_0^c}\|_1 \\ &\leq 2\|\Psi_d \beta_{0, S_0^c}\|_1 - \|\Psi_d(\beta_{0, S_0^c} - \hat{\beta}_{S_0^c})\|_1 \\ &\leq -\|\Psi_d(\beta_{0, S_0^c} - \hat{\beta}_{S_0^c})\|_1. \end{aligned}$$

The last inequality follows from the fact that $\|\beta_{0, S_0^c}\|_1 = 0$. Hence,

$$\begin{aligned} &\|\Psi_d \beta_0\|_1 - \|\Psi_d \hat{\beta}\|_1 + \frac{1}{c} \|\Psi_d(\hat{\beta} - \beta_0)\|_1 \\ &\leq \left(1 + \frac{1}{c}\right) \|\Psi_d(\hat{\beta}_{S_0} - \beta_{0, S_0})\|_1 - \left(1 - \frac{1}{c}\right) \|\Psi_d(\hat{\beta}_{S_0^c} - \beta_{0, S_0^c})\|_1. \end{aligned} \quad (\text{B.19})$$

Plugging this result in (B.12) we get, with probability tending to 1 as $n \rightarrow \infty$,

$$\begin{aligned} &(\hat{\beta} - \beta_0)^T \hat{\Sigma}(\hat{\beta} - \beta_0) \\ &\leq \frac{2\lambda}{\tau} \left[\left(1 + \frac{1}{c}\right) \|\Psi_d(\hat{\beta}_{S_0} - \beta_{0, S_0})\|_1 - \left(1 - \frac{1}{c}\right) \|\Psi_d(\hat{\beta}_{S_0^c} - \beta_{0, S_0^c})\|_1 \right], \end{aligned} \quad (\text{B.20})$$

and thus

$$\begin{aligned} (\hat{\beta} - \beta_0)^T \hat{\Sigma}(\hat{\beta} - \beta_0) + \frac{2\lambda}{\tau} \left(1 - \frac{1}{c}\right) \|\Psi_d(\hat{\beta} - \beta_0)\|_1 &\leq \frac{4\lambda}{\tau} \|\Psi_d(\hat{\beta}_{S_0} - \beta_{0, S_0})\|_1 \\ &\leq \frac{4\lambda \bar{\psi} \sqrt{s_\beta}}{\tau} \|\hat{\beta}_{S_0} - \beta_{0, S_0}\|_2, \end{aligned} \quad (\text{B.21})$$

where we have used the fact that $\text{Card}(S_0) \leq s_\beta$ due to Assumption 3.1. Recall that $c > 1$. From inequality (B.20) and the fact that $(\hat{\beta} - \beta_0)^T \hat{\Sigma}(\hat{\beta} - \beta_0) \geq 0$ we obtain a cone condition $\Psi_d(\hat{\beta} - \beta_0) \in \mathcal{C}[S_0, c_0]$ for $\Psi_d(\hat{\beta} - \beta_0)$, which in turn implies (with probability tending to 1 as $n \rightarrow \infty$) a cone condition $\hat{\beta} - \beta_0 \in \mathcal{C}[S_0, c_0 \bar{\psi} / \underline{\psi}]$ for $\hat{\beta} - \beta_0$. Therefore, using (B.11) (where we recall that $\kappa(\hat{\Sigma}) = c_* \kappa_\Sigma$), we obtain that, with probability tending to 1 as $n \rightarrow \infty$,

$$(\hat{\beta} - \beta_0)^T \hat{\Sigma}(\hat{\beta} - \beta_0) + \frac{2\lambda}{\tau} \left(1 - \frac{1}{c}\right) \|\Psi_d(\hat{\beta} - \beta_0)\|_1 \leq \frac{4\lambda \bar{\psi} \sqrt{s_\beta}}{\tau} \sqrt{\frac{(\hat{\beta} - \beta_0)^T \hat{\Sigma}(\hat{\beta} - \beta_0)}{c_* \kappa_\Sigma}}.$$

Using here the inequality $ab \leq (a^2 + b^2)/2, \forall a, b > 0$, we find that, with probability tending to 1 as $n \rightarrow \infty$,

$$\frac{1}{2}(\hat{\beta} - \beta_0)^T \hat{\Sigma}(\hat{\beta} - \beta_0) + \frac{2\lambda}{\tau} \left(1 - \frac{1}{c}\right) \|\Psi_d(\hat{\beta} - \beta_0)\|_1 \leq \frac{8\lambda^2 \bar{\psi}^2 s_\beta}{\tau^2 c_* \kappa_\Sigma}. \quad (\text{B.22})$$

which implies the first inequality in (B.18). Since $\lambda \lesssim \sqrt{\log(p)/n}$ the proof of (3.14) is complete.

B.3.2 Proof of (3.15)

Recall that $\hat{\mu}$ is defined as:

$$\hat{\mu} \in \arg \min_{\mu \in \mathbb{R}^p} \left(\frac{1}{n} \sum_{i=1}^n (1 - D_i) h'(X_i^T \hat{\beta}) (Y_i - X_i^T \mu)^2 + \lambda' \sum_{j=1}^p \psi'_j |\mu_j| \right).$$

Let $\Psi' \in \mathbb{R}^{p \times p}$ denote the diagonal matrix with diagonal entries ψ'_1, \dots, ψ'_p . We will prove that, with probability tending to 1 as $n \rightarrow \infty$,

$$\|\Psi'(\hat{\mu} - \mu_0)\|_1 \lesssim \frac{s}{\kappa_\Sigma} \sqrt{\frac{\log(p)}{n}}. \quad (\text{B.23})$$

Using an argument analogous to that after (B.15) we easily get that (B.23) implies (3.15).

Step 1: Concentration inequality

Define $V_{ij} := (1 - D_i) h'(X_i^T \beta_0) [Y_i - X_i^T \mu_0] X_{i,j}$, $\mathcal{S}'_j := \frac{1}{\sqrt{n}} \sum_{i=1}^n V_{ij} / \psi'_j$ and consider the event

$$\mathcal{B}' := \left\{ \frac{2}{n} \max_{1 \leq j \leq p} \left| \sum_{i=1}^n \frac{V_{ij}}{\psi'_j} \right| \leq \frac{\lambda'}{c} \right\}.$$

The random variables $V_{ij}, i = 1, \dots, n$, are i.i.d. and $\mathbb{E}(V_{ij}) = 0$, $\mathbb{E}(V_{ij}^2) \geq c_1$ and $\mathbb{E}(|V_{ij}|^3) < C$ for all i, j by Assumptions 3.2 and 3.3. Using Assumptions 3.1, 3.6, and Lemma C.1 we obtain

$$\begin{aligned} \mathbb{P}(\mathcal{B}'^C) &= \mathbb{P} \left(\frac{c}{\sqrt{n}} \max_{1 \leq j \leq p} |\mathcal{S}'_j| > c\Phi^{-1}(1 - \gamma/2p) / \sqrt{n} \right) \\ &= \mathbb{P} \left(\max_{1 \leq j \leq p} |\mathcal{S}'_j| > \Phi^{-1}(1 - \gamma/2p) \right) = o(1) \quad \text{as } n \rightarrow \infty. \end{aligned}$$

Step 2: Control of the ℓ_1 -error for $\hat{\mu}$

Introduce the notation $\gamma_{\beta, \mu}(Z_i) = (1 - D_i) h'(X_i^T \beta) (Y_i - X_i^T \mu)^2$. It follows from the definition of $\hat{\mu}$ that

$$\frac{1}{n} \sum_{i=1}^n [\gamma_{\hat{\beta}, \hat{\mu}}(Z_i) - \gamma_{\hat{\beta}, \mu_0}(Z_i)] \leq \lambda' (\|\Psi' \mu_0\|_1 - \|\Psi' \hat{\mu}\|_1).$$

Here

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n [\gamma_{\hat{\beta}, \hat{\mu}}(Z_i) - \gamma_{\hat{\beta}, \mu_0}(Z_i)] &= (\hat{\mu} - \mu_0)^T \left(\frac{1}{n} \sum_{i=1}^n (1 - D_i) h'(X_i^T \hat{\beta}) X_i X_i^T \right) (\hat{\mu} - \mu_0) \\ &\quad + \frac{2}{n} \sum_{i=1}^n (1 - D_i) h'(X_i^T \hat{\beta}) (Y_i - X_i^T \mu_0) X_i^T (\mu_0 - \hat{\mu}). \end{aligned}$$

Therefore,

$$\begin{aligned} &(\hat{\mu} - \mu_0)^T \left(\frac{1}{n} \sum_{i=1}^n (1 - D_i) h'(X_i^T \hat{\beta}) X_i X_i^T \right) (\hat{\mu} - \mu_0) \\ &\leq \lambda' (\|\Psi' \mu_0\|_1 - \|\Psi' \hat{\mu}\|_1) + \frac{2}{n} \sum_{i=1}^n (1 - D_i) h'(X_i^T \beta_0) (Y_i - X_i^T \mu_0) X_i^T (\hat{\mu} - \mu_0) + R_n, \end{aligned} \quad (\text{B.24})$$

where

$$\begin{aligned} R_n &= \frac{2}{n} \sum_{i=1}^n (1 - D_i) \left[h'(X_i^T \hat{\beta}) - h'(X_i^T \beta_0) \right] (Y_i - X_i^T \mu_0) X_i^T (\hat{\mu} - \mu_0) \\ &= \frac{2}{n} \sum_{i=1}^n (1 - D_i) (Y_i - X_i^T \mu_0) h''(X_i^T \tilde{\beta}) X_i^T (\hat{\beta} - \beta_0) X_i^T (\hat{\mu} - \mu_0) \end{aligned}$$

with $\tilde{\beta} = t\hat{\beta} + (1-t)\beta_0$ for some $t \in [0, 1]$. Introducing the matrix

$$A = \frac{2}{n} \sum_{i=1}^n (1 - D_i) (Y_i - X_i^T \mu_0) h''(X_i^T \tilde{\beta}) X_i X_i^T,$$

we can write

$$R_n = (\hat{\beta} - \beta_0)^T A (\hat{\mu} - \mu_0).$$

From (B.24) we deduce that on the event \mathcal{B}' that occurs with probability tending to 1 as $n \rightarrow \infty$,

$$\begin{aligned} &(\hat{\mu} - \mu_0)^T \left(\frac{1}{n} \sum_{i=1}^n (1 - D_i) h'(X_i^T \hat{\beta}) X_i X_i^T \right) (\hat{\mu} - \mu_0) \\ &\leq \lambda' (\|\Psi' \mu_0\|_1 - \|\Psi' \hat{\mu}\|_1) + \frac{\lambda'}{c} \|\Psi'(\hat{\mu} - \mu_0)\|_1 + (\hat{\beta} - \beta_0)^T A (\hat{\mu} - \mu_0). \end{aligned}$$

We now use (B.17) (noticing that $\hat{\beta} = \tilde{\beta}$ for $t = 1$) to obtain that, with probability tending to 1 as $n \rightarrow \infty$,

$$\begin{aligned} \tau (\hat{\mu} - \mu_0)^T \hat{\Sigma} (\hat{\mu} - \mu_0) &\leq \lambda' (\|\Psi' \mu_0\|_1 - \|\Psi' \hat{\mu}\|_1) + \frac{\lambda'}{c} \|\Psi'(\hat{\mu} - \mu_0)\|_1 \\ &\quad + (\hat{\beta} - \beta_0)^T A (\hat{\mu} - \mu_0). \end{aligned} \quad (\text{B.25})$$

Next, observe that, with probability tending to 1 as $n \rightarrow \infty$,

$$(\hat{\beta} - \beta_0)^T A (\hat{\mu} - \mu_0) - (\tau/2) (\hat{\mu} - \mu_0)^T \hat{\Sigma} (\hat{\mu} - \mu_0) \leq C (\hat{\beta} - \beta_0)^T \hat{\Sigma} (\hat{\beta} - \beta_0) \quad (\text{B.26})$$

where $C > 0$ is a constant that does not depend on n . To see this, set $u_i = (\hat{\beta} - \beta_0)^T X_i$ $v_i = (\hat{\mu} - \mu_0)^T X_i$ and $a_i = (Y_i - X_i^T \mu_0) h''(X_i^T \tilde{\beta})$. We have

$$\begin{aligned} (\hat{\beta} - \beta_0)^T A(\hat{\mu} - \mu_0) - (\tau/2)(\hat{\mu} - \mu_0)^T \hat{\Sigma}(\hat{\mu} - \mu_0) &= \frac{\tau}{2n} \sum_{i=1}^n (1 - D_i) \left(\frac{4a_i}{\tau} u_i v_i - v_i^2 \right) \\ &\leq \frac{1}{\tau n} \sum_{i=1}^n (1 - D_i) a_i^2 u_i^2. \end{aligned}$$

This implies (B.26) since (3.14) and Assumptions 3.3(i) and 3.2 guarantee that, with probability tending to 1 as $n \rightarrow \infty$, we have $\max_i |a_i| \leq C$ for a constant $C > 0$ that does not depend on n .

We also note that, due to (B.22), with probability tending to 1 as $n \rightarrow \infty$,

$$(\hat{\beta} - \beta_0)^T \hat{\Sigma}(\hat{\beta} - \beta_0) \lesssim \frac{\lambda^2 s_\beta}{\kappa_\Sigma} \lesssim \frac{s_\beta \log(p)}{n \kappa_\Sigma}. \quad (\text{B.27})$$

Combining (B.25), (B.26) and (B.27) we finally get that, with probability tending to 1 as $n \rightarrow \infty$,

$$\frac{\tau}{2} (\hat{\mu} - \mu_0)^T \hat{\Sigma}(\hat{\mu} - \mu_0) \leq \lambda' \left(\|\Psi' \mu_0\|_1 - \|\Psi' \hat{\mu}\|_1 + \frac{1}{c} \|\Psi'(\hat{\mu} - \mu_0)\|_1 \right) + \frac{\bar{c} s_\beta \log(p)}{n \kappa_\Sigma}, \quad (\text{B.28})$$

where $\bar{c} > 0$ is a constant that does not depend on n . Let $S_1 \subseteq \{1, \dots, p\}$ denote the set of indices of non-zero components of μ_0 . By assumption, $\text{Card}(S_1) = \|\mu_0\|_0 \leq s_\mu$. The same argument as in (B.19) (where we replace $\beta_0, \hat{\beta}, \Psi_d, S_0$ by $\mu_0, \hat{\mu}, \Psi', S_1$, respectively) yields

$$\begin{aligned} &\|\Psi' \mu_0\|_1 - \|\Psi' \hat{\mu}\|_1 + \frac{1}{c} \|\Psi'(\hat{\mu} - \mu_0)\|_1 \\ &\leq \left(1 + \frac{1}{c}\right) \|\Psi'(\hat{\mu}_{S_1} - \mu_{0,S_1})\|_1 - \left(1 - \frac{1}{c}\right) \|\Psi'(\hat{\mu}_{S_1^c} - \mu_{0,S_1^c})\|_1. \end{aligned}$$

This and (B.28) imply that, with probability tending to 1 as $n \rightarrow \infty$,

$$\frac{\tau}{2} (\hat{\mu} - \mu_0)^T \hat{\Sigma}(\hat{\mu} - \mu_0) + \lambda' \left(1 - \frac{1}{c}\right) \|\Psi'(\hat{\mu} - \mu_0)\|_1 \leq 4\lambda' \|\Psi'(\hat{\mu}_{S_1} - \mu_{0,S_1})\|_1 + \frac{\bar{c} s_\beta \log(p)}{n \kappa_\Sigma}, \quad (\text{B.29})$$

where we have used the fact that $c > 1$. We now consider two cases:

1. $\lambda' \|\Psi'(\hat{\mu}_{S_1} - \mu_{0,S_1})\|_1 \leq \frac{\bar{c} s_\beta \log(p)}{n \kappa_\Sigma}$. In this case, inequality (B.29) implies

$$\|\Psi'(\hat{\mu} - \mu_0)\|_1 \lesssim \frac{s_\beta \log(p)}{\lambda' n \kappa_\Sigma}$$

and (B.23) follows immediately since $\sqrt{\log(p)/n} \lesssim \lambda'$. Consequently, (3.15) holds in this case.

2. $\lambda' \|\Psi'(\hat{\mu}_{S_1} - \mu_{0,S_1})\|_1 > \frac{\bar{c} s_\beta \log(p)}{n \kappa_\Sigma}$. Then with probability tending to 1 as $n \rightarrow \infty$ we have

$$\frac{\tau}{2} (\hat{\mu} - \mu_0)^T \hat{\Sigma}(\hat{\mu} - \mu_0) + \lambda' \left(1 - \frac{1}{c}\right) \|\Psi'(\hat{\mu} - \mu_0)\|_1 \leq 5\lambda' \|\Psi'(\hat{\mu}_{S_1} - \mu_{0,S_1})\|_1.$$

This inequality is analogous to (B.21). In particular, it implies the cone condition, which now takes the form $\|\Psi'(\hat{\mu}_{S_0^c} - \mu_{0,S_0^c})\|_1 \leq \alpha \|\Psi'(\hat{\mu}_{S_0} - \mu_{0,S_0})\|_1$ with $\alpha = 5c/(c-1)$. Therefore,

we can use an argument based on the Restricted Eigenvalue condition, which is completely analogous to that after inequality (B.21) (we omit the details here). It leads to the following analog of (B.22):

$$(\hat{\mu} - \mu_0)^T \hat{\Sigma} (\hat{\mu} - \mu_0) + C\lambda' \|\Psi'(\hat{\mu} - \mu_0)\|_1 \lesssim \frac{(\lambda')^2 s_\mu}{\kappa_\Sigma}, \quad (\text{B.30})$$

where $C > 0$ is a constant that does not depend on n . Since $\lambda' \lesssim \sqrt{\log(p)/n}$ we get (B.23). Thus, the proof of (3.15) is complete. \square

C Auxiliary Lemmas

Lemma C.1 (Deviation of maximum of self-normalized sums) *Consider*

$$\mathcal{S}_j := \sum_{i=1}^n U_{i,j} \left(\sum_{i=1}^n U_{i,j}^2 \right)^{-1/2},$$

where $U_{i,j}$ are independent random variables across i with mean zero and for all i, j we have $\mathbb{E}[|U_{i,j}|^3] \leq C_1$, $\mathbb{E}[U_{i,j}^2] \geq C_2$ for some positive constants C_1, C_2 independent of n . Let $p = p(n)$ satisfy the condition $\log(p) = o(n^{1/3})$ and let $\gamma = \gamma(n) \in (0, 2p)$ be such that $\log(1/\gamma) \lesssim \log(p)$. Then,

$$\mathbb{P} \left(\max_{1 \leq j \leq p} |\mathcal{S}_j| > \Phi^{-1}(1 - \gamma/2p) \right) = \gamma(1 + o(1))$$

as $n \rightarrow \infty$.

Proof. We use a corollary of a result from de la Peña et al. (2009) given by Belloni et al. (2012, p.2409), which in our case can be stated as follows. Let \mathcal{S}_j and $U_{i,j}$ satisfy the assumptions of the present lemma. If there exist positive numbers $\ell > 0$, $\gamma > 0$ such that

$$0 < \Phi^{-1}(1 - \gamma/2p) \leq \frac{C_2^{1/2} n^{1/6}}{C_1^{1/3} \ell} - 1, \quad (\text{C.1})$$

then,

$$\mathbb{P} \left(\max_{1 \leq j \leq p} |\mathcal{S}_j| > \Phi^{-1}(1 - \gamma/2p) \right) \leq \gamma \left(1 + \frac{A}{\ell^3} \right),$$

where $A > 0$ is an absolute constant.

Now, since $\Phi^{-1}(1 - \gamma/2p) \leq \sqrt{2 \log(2p/\gamma)}$ and we assume that $\log(1/\gamma) \lesssim \log(p)$ and $\log(p) = o(n^{1/3})$ condition (C.1) is satisfied with $\ell = \ell(n) = (n^{1/3}/\log(p))^{1/4}$ for n large enough. Then $\ell(n) \rightarrow \infty$ as $n \rightarrow \infty$ and the lemma follows. \square

Lemma C.2 *Let $s \in [1, p/2]$ be an integer and $\alpha > 0$. Let $V \in \mathbb{R}^p$ be a random vector such that $\|V\|_\infty \leq M < \infty$ (a.s.), and set $\Sigma = \mathbb{E}(VV^T)$. Let Σ satisfy (3.13) and*

$$s/\kappa_\Sigma = o(p) \text{ as } n \rightarrow \infty, \text{ and } s \lesssim \frac{n\kappa_\Sigma^2}{\log(p) \log^3(n)}. \quad (\text{C.2})$$

Consider i.i.d. random vectors V_1, \dots, V_n with the same distribution as V . Then, for all n large enough with probability at least $1 - \exp(-C \log(p) \log^3(n))$ where $C > 0$ is a constant depending only on M the empirical matrix $\hat{\Sigma} = \frac{1}{n} \sum_{i=1}^n V_i V_i^T$ satisfies the $RE(s, \alpha)$ condition with $\kappa(\hat{\Sigma}) = c_* \kappa_\Sigma$ where $c_* \in (0, 1)$ is an absolute constant.

Proof. We set the parameters of Theorem 22 in Rudelson and Zhou (2013) as follows $s_0 = s$, $k_0 = \alpha$, and due to (3.13) we have, in the notation of that theorem, $K^2(s_0, 3k_0, \Sigma^{1/2}) \leq 1/\kappa_\Sigma$ and $\rho \geq \kappa_\Sigma$. Also note that $\|\Sigma^{1/2} e_j\|_2^2 = \mathbb{E}[(V^T e_j)^2] \leq M^2$, where e_j denotes the j th canonical basis vector in \mathbb{R}^p (this, in particular, implies that $\kappa_\Sigma \leq M^2$). Thus, the value d defined in Theorem 22 in Rudelson and Zhou (2013) satisfies $d \lesssim s/\kappa_\Sigma$ and condition $d \leq p$ holds true for n large enough due to (C.2). Next, note that condition $n \geq x \log^3(x)$ is satisfied for all $x \leq n/\log^3(n)$ and $n \geq 3$, so that the penultimate display formula in Theorem 22 of Rudelson and Zhou (2013) can be written as $d \log(p)/\rho \lesssim n/\log^3(n)$. Given the above bounds for d and ρ , we have a sufficient condition for this inequality in the form $s/\kappa_\Sigma^2 \lesssim n/\log^3(n)$, which is granted by (C.2). Thus, all the conditions of Theorem 22 in Rudelson and Zhou (2013) are satisfied and we find that, for all n large enough, with probability at least $1 - \exp(-C \log(p) \log^3(n))$ where $C > 0$ is a constant depending only on M we have

$$\min_{\substack{S \subseteq \{1, \dots, p\}: \\ \text{Card}(S) = s}} \min_{v \in \mathcal{C}[S, \alpha]} \frac{v^T \hat{\Sigma} v}{\|v_S\|_2^2} \geq (1 - 5\delta) \kappa_\Sigma, \quad (\text{C.3})$$

where $\delta \in (0, 1/5)$ (remark that there is a typo in Theorem 22 in Rudelson and Zhou (2013) that is corrected in (C.3): the last formula of that theorem should be $(1 - 5\delta) \|\Sigma^{1/2} u\|_2 \leq \frac{\|Xu\|_2}{\sqrt{n}} \leq (1 + 3\delta) \|\Sigma^{1/2} u\|_2$ where $0 < \delta < 1/5$ [Rudelson (2020)]). It remains to note that though at first glance (C.3) differs from (B.11) (in (C.3) we have $\text{Card}(S) = s$ rather than $\text{Card}(S) \leq s$), these two conditions are equivalent. Indeed, as shown in (Bellec et al., 2018, page 3607),

$$\bigcup_{S \subseteq \{1, \dots, p\}: \text{Card}(S) \leq s} \left\{ v \in \mathbb{R}^p : \|v_{S^c}\|_1 \leq \alpha \|v_S\|_1 \right\} = \left\{ v \in \mathbb{R}^p : \|v\|_1 \leq (1 + \alpha) \sum_{j=1}^s v_j^* \right\}$$

where $v_1^* \geq \dots \geq v_p^*$ denotes a non-increasing rearrangement of $|v_1|, \dots, |v_p|$. On the other hand,

$$\begin{aligned} \bigcup_{S \subseteq \{1, \dots, p\}: \text{Card}(S) = s} \left\{ v \in \mathbb{R}^p : \|v_{S^c}\|_1 \leq \alpha \|v_S\|_1 \right\} &\supseteq \left\{ v \in \mathbb{R}^p : \|v_{S_*^c(v)}\|_1 \leq \alpha \|v_{S_*(v)}\|_1 \right\} \\ &= \left\{ v \in \mathbb{R}^p : \|v\|_1 \leq (1 + \alpha) \sum_{j=1}^s v_j^* \right\} \end{aligned}$$

where $S_*(v)$ is the set of s largest in absolute value components of v . □