

Estimation in high-dimensional linear regression: Post-Double-Autometrics as an alternative to Post-Double-Lasso*

Sullivan Hué[†], Sébastien Laurent^{†‡§}, Ulrich Aiounou[†] & Emmanuel Flachaire[†]

March 7, 2025

Abstract

Post-Double-Lasso is becoming the most popular method for estimating linear regression models with many covariates when the purpose is to obtain an accurate estimate of a parameter of interest, such as an average treatment effect. However, this method can suffer from substantial omitted variable bias in finite sample. We propose a new method called Post-Double-Autometrics, which is based on Autometrics, and show that this method outperforms Post-Double-Lasso. Its use in a standard application of economic growth sheds new light on the hypothesis of convergence from poor to rich economies.

JEL Classification: C21, C52, C55.

Keywords: Treatment effect, High dimension, Lasso, Autometrics.

*We thank participants at the ANR MLEforRisk workshop in Orléans, 2024 French Conference in Applied Econometrics using Stata, 2024 Quantitative Finance and Financial Econometrics conference in Marseille, FinEML conference 2024 in Lugano and Autometrics workshop in Aix-en-Provence for their comments. We acknowledge research support by the French National Research Agency Grants ANR-17-EURE-0020 and ANR-21-CE26-0007-01 (project MLEforRisk) and by the Excellence Initiative of Aix-Marseille University - A*MIDEX.

[†]Aix-Marseille University (Aix-Marseille School of Economics), CNRS & EHESS.

[‡]Aix-Marseille Graduate School of Management – IAE, France.

[§]Institut universitaire de France (IUF), France.

1 Introduction

In high-dimensional linear regression, covariate selection becomes a key issue when the purpose is to obtain a good estimate of a parameter of interest, such as an average treatment effect. It is important not to miss relevant variables correlated to the outcome and the treatment variable (to avoid bias), and not to include irrelevant variables (to avoid loss of accuracy). For this purpose, Post-Double-Lasso (Belloni et al., 2014) is becoming the most popular method. This method allows the use of many covariates, to avoid omitted variable bias, while controlling overfitting issue and loss of accuracy with a Lasso variable selection method.

However, Wüthrich and Zhu (2023) demonstrate that this method can experience significant omitted variable bias in finite samples, even in straightforward settings favorable to Lasso. They also show that the performance of this method is very sensitive to the choice of the regularization parameter of the Lasso, and no specific choice of this parameter systematically outperforms the others.

In this paper, we consider an alternative variable selection method, Autometrics (Doornik et al., 2009), which is based on statistical inference. This method performs automatic selection of variables based on “Hendry’s theory of reduction” (see Chapter 6 of Hendry and Doornik, 2014) and an automatic general-to-specific model selection (Hendry, 2000). Its main hyperparameter, the target size α , is the significance level used for inference. Simulation results show that this parameter allows to determine approximatively the frequency of irrelevant explanatory variables included in the terminal model (Hendry and Doornik, 2014, Epprecht et al., 2021, Flachaire et al., 2024).

Firstly, we propose a new method called Post-Double-Autometrics, that extends Post-Double-Lasso by using Autometrics instead of Lasso to do variable selection. We show through extensive Monte Carlo experiments that the proposed method provides the best variable selection method and outperforms Post-Double-Lasso in many cases. To illustrate, Figure 1 provides a numerical illustration of this phenomenon in a simple simulation experiment, with $n = 400$ observations and $p = 210$ covariates of which only 10 are relevant. The error term is *i.i.d.* and the covariates are moderately correlated. The Lasso penalization parameter λ^{min} is obtained by cross-validation, and the Autometrics target size is 5%. The left panel shows the behavior of the Post-Double-Lasso estimator of a treatment effect whose coefficient is equal to 0 while the right panel shows the empirical distribution of our proposed Post-Double-Autometrics estimator. In this setting, the Post-Double-Lasso estimator is upward biased while the Post-Double-Autometrics estimator is approximately unbiased and behaves almost like the oracle estimator. More details on this simulation will be given in Section 3 and more particularly in Section 3.2.

Secondly, we revisit a standard economic growth application, on the convergence hypothesis that the poor catch up the rich countries, with the Barro and

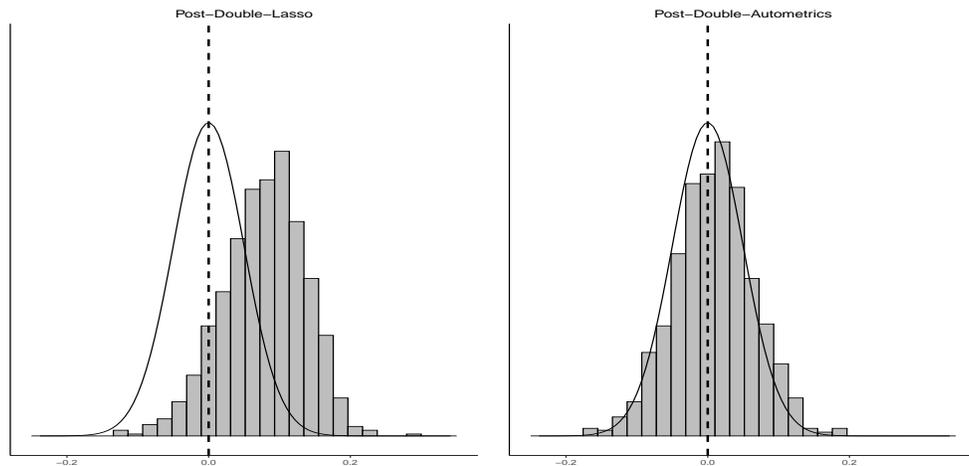


Figure 1: Empirical distribution of the Post-Double-Lasso (left panel) and Post-Double-Autometrics (right panel) estimation of the treatment effect. In both panels, the distribution of the oracle estimator corresponds to the solid line while the true value corresponds to the dotted vertical line.

Lee (1994) data consisting of a panel of countries for the period of 1960 to 1985. This hypothesis is often tested from a linear regression of growth rates on the initial level of GDP and additional characteristics of countries. Given that the number of covariates that describe such characteristics can be comparable to the sample size, covariate selection is an issue. Empirical results based on Lasso selection method support the convergence hypothesis (Belloni and Chernozhukov, 2011, Belloni et al., 2013). We show that this conclusion is sensitive to the choice of the penalization term, and that the convergence hypothesis is rejected with Post-Double-Autometrics whatever the target size. Further analysis shows that including or missing one specific covariate is most likely the source of support or lack of support for the convergence hypothesis.

The remainder of the paper is structured as follows. Section 2 presents the pitfalls of Lasso-based methods and introduce our new approach. Section 3 is devoted to Monte Carlo experiments and Section 4 to an empirical application. The paper concludes with Section 5.

2 Estimation in high-dimensional linear regression

In this section, we first present the framework, Post-Double-Lasso and finally Post-Double-Autometrics.

2.1 Framework

In accordance with the literature, we consider the following standard linear regression model

$$y_i = \mathbf{X}_i \boldsymbol{\beta} + d_i \delta + \epsilon_i, \quad (1)$$

$$d_i = \mathbf{X}_i \boldsymbol{\gamma} + \eta_i, \quad (2)$$

for $i = 1, \dots, n$, where y_i is the outcome, d_i is the treatment variable, $\mathbf{X}_i = (x_{1,i}, \dots, x_{p,i}) \in \mathbb{R}^p$ is a vector of p covariates, and ϵ_i and η_i are error terms. The parameter of interest δ is the average treatment effect. The covariates might include transformations of raw variables in order to capture non-linear effects (e.g., powers of raw variables, interaction terms).

In observational studies, the unconfoundedness condition is crucial to obtain unbiased estimation of the average treatment effect δ . This condition requires \mathbf{X} to contain all the confounder variables which explain both the outcome y and the treatment d . Otherwise, the estimation may suffer from substantial omitted variable bias (OVB) or confounding bias. In practice, it is difficult to account for all confounder variables. Here, we consider the strategy of including many covariates and using a variable selection method to select the relevant covariates (to make the estimate as efficient as possible).

2.2 Post-Double-Lasso

An early approach developed in the literature is Lasso, which performs estimation and variable selection simultaneously (Tibshirani, 1996). Considering the linear regression in (1), Lasso solves the following penalized regression problem

$$\hat{\boldsymbol{\theta}} = \arg \min_{\boldsymbol{\beta}} \left[\sum_{i=1}^n (y_i - \mathbf{Z}_i \boldsymbol{\theta})^2 + \lambda \sum_{j=1}^p \omega_j |\beta_j| \right], \quad (3)$$

where $\mathbf{Z}_i = (\mathbf{X}_i, d_i)$, $\boldsymbol{\theta} = (\boldsymbol{\beta}^\top, \delta)^\top$ and $\omega_j = 1 \forall j$. Estimation is feasible in a high dimensional setting, that is, when the number of covariates p is large and can exceed the number of observations n . The penalization restricts the magnitude of the coefficients, making some of them exactly zero with a sufficiently large λ . Several approaches have been proposed in the literature to set λ . For example, Bickel et al. (2009) proposed to set $\lambda^{bya} = 2\sigma\sqrt{2n^{-1}(1+\tau)\log p}$, where σ is the standard deviation of the error term, and $\tau > 0$. Belloni et al. (2012) proposed a penalization choice, denoted as λ^{bch} , to take into account heteroscedastic and clustered errors. Another popular approach is to set λ by cross-validation, with the value minimizing the prediction error, λ^{min} , or the value associated to the most parsimonious model within a 1-standard-error interval, λ^{1se} . See Homrighausen

and McDonald (2013, 2014) and Chetverikov et al. (2021) for theoretical results on the choice of λ by cross-validation.

The penalization shrinks all coefficients towards zero, introducing some bias into the coefficient estimates associated with relevant covariates. Since the interest is on the estimation of the parameter δ , it is recommended to use the Post-Lasso (Belloni and Chernozhukov, 2013) which relies on two steps. First, estimate (1) by Lasso to select a subset of variables, denoted \mathbf{X}^* . Second, estimate δ from the OLS regression of y on \mathbf{X}^* and d . To mitigate the bias introduced in the coefficients of relevant variables, the Adaptive-Lasso can also be used (Zou, 2006). It solves the penalized regression problem in (3), with weights $\omega_j = |\hat{\mu}_j|^{-\eta}$, where $\eta > 0$ and $\hat{\mu}_j$ is a consistent estimator of θ_j . Here, smaller weights are assigned to larger coefficients in the constraint and, thus, large non-zero coefficients shrink more slowly to zero as λ increases.

Belloni et al. (2014) show that Lasso usually fails to select some relevant variables in the first step, which may lead to a biased estimate of the parameter of interest δ . This typically occurs when some covariates in \mathbf{X} are correlated with d , but only have a small effect on y . To solve this issue, they proposed the Post-Double-Lasso. The standard linear regression model defined in (1)-(2) implies the following reduced form model:

$$y_i = \mathbf{X}_i \boldsymbol{\alpha} + \nu_i, \quad (4)$$

where $\boldsymbol{\alpha} = \boldsymbol{\beta} + \boldsymbol{\gamma} \delta$ and $\nu_i = \epsilon_i + \eta_i \delta$. In order to select relevant control variables, Post-Double-Lasso relies on the following two steps of Lasso

$$\hat{\boldsymbol{\alpha}} = \arg \min_{\boldsymbol{\alpha}} \left[\sum_{i=1}^n (y_i - \mathbf{X}_i \boldsymbol{\alpha})^2 + \lambda_1 \sum_{j=1}^p w_{1,j} |\alpha_j| \right], \quad (5)$$

$$\hat{\boldsymbol{\gamma}} = \arg \min_{\boldsymbol{\gamma}} \left[\sum_{i=1}^n (d_i - \mathbf{X}_i \boldsymbol{\gamma})^2 + \lambda_2 \sum_{j=1}^p w_{2,j} |\gamma_j| \right], \quad (6)$$

where $w_{1,j} = w_{2,j} = 1 \ \forall j$. Finally, δ is estimated from the OLS regression of y on d and the union of selected control variables, such as

$$\left(\hat{\boldsymbol{\beta}}, \hat{\delta} \right) = \arg \min_{\boldsymbol{\beta}, \delta} \left[\sum_{i=1}^n (y_i - \mathbf{X}_i \boldsymbol{\beta} - d_i \delta)^2 \right] \text{ s.t. } \beta_j = 0 \ \forall j \notin \{ \hat{\mathbf{I}}_{\mathbf{X}^*} \cup \hat{\mathbf{I}}_{\mathbf{X}^{**}} \}, \quad (7)$$

where $\hat{\mathbf{I}}_{\mathbf{X}^*} = \text{supp}(\hat{\boldsymbol{\alpha}}) = \{j : \hat{\alpha}_j \neq 0\}$ and $\hat{\mathbf{I}}_{\mathbf{X}^{**}} = \text{supp}(\hat{\boldsymbol{\gamma}}) = \{j : \hat{\gamma}_j \neq 0\}$. This method thus gives covariates omitted in the first step a chance to be selected in the second step, especially those correlated to d and likely to lead to an omitted variable bias. Similarly to the Post-Lasso, several choices can be considered for the penalization parameters λ_1 and λ_2 , and the Lasso can be replaced by the Adaptive-Lasso.

Although the Post-Double-Lasso is a substantial improvement over the Post-Lasso, Wüthrich and Zhu (2023) show that it can still lead to an OVB. Specifically, under certain assumptions, the authors theoretically demonstrate that the probability of not selecting any relevant control variables, given that at least one exists, becomes high when p is relatively large. This occurs when the magnitude of the control variable coefficients is not sufficiently large compared to the penalty parameter. In such cases, the Lasso method may fail to capture relevant variables in both steps of the Post-Double-Lasso procedure and induce an OVB on δ . Additionally, the authors show through simulations that the magnitude of the bias depends on the choice of the regularization parameters λ_1 and λ_2 , and that no single popular choice consistently leads to better results.

2.3 Post-Double-Autometrics

To address this OVB issue, we propose to use another variable selection method: Autometrics (Doornik et al., 2009). Autometrics is a popular variable selection model among econometricians because it relies on statistical inference to select the relevant variables. There exists at least six different software implementations of automatic “general-to-specific” (GETS) model selection (Matlab, OxMetrics, Scilab, STATA, EViews and R). Autometrics is one of them and is implemented in OxMetrics. See Sucarrat (2020) for more details on general-to-specific modelling and its implementations.

More specifically, Autometrics is an algorithm performing automatic variables selection with the GETS model selection approach (Hendry, 2000). As Lasso, this method starts from a generalized unrestricted model (GUM) which includes all potential relationships between the outcome y and the control variables \mathbf{X} , such as lag values, trends, step indicators or any other transformations of the raw variables aiming at capturing non-linearities. The algorithm then performs a battery of tests to remove insignificant variables from \mathbf{X} and to find a congruent parsimonious model. The final model is then estimated by OLS.¹ Using a multiple block-search algorithm, Autometrics can also be used with more observations than covariates.

Autometrics has three main advantages. First, unlike the Lasso performing variable selection based on the magnitude of coefficients (Wüthrich and Zhu, 2023), Autometrics relies on statistical inference. Therefore, Autometrics-based methods will select variables whenever their coefficients are significant, independently of the magnitude of their coefficients. Second, the tuning parameter of Autometrics is the targeted nominal size α . It allows the user to control the expected proportion of irrelevant variables included in the final model. This approach contrasts sharply with Lasso, where λ directly controls the number (or proportion) of variables se-

¹For more details on Autometrics, see Doornik et al. (2009) and Castle et al. (2023).

lected in the final model. Third, Autometrics may invalidate a reduction when the reduced model does not pass some pre-defined misspecification tests.

The target size α refers to the pre-specified significance level or threshold that guides the inclusion or exclusion of variables in the model during the selection process of Autometrics. Specifically, it determines the acceptable probability of incorrectly rejecting a true null hypothesis (i.e., including a variable that is not truly relevant). By setting a target size, Autometrics controls the false-positive rate of variable selection, balancing between overfitting (selecting too many irrelevant variables) and underfitting (failing to select relevant variables). In practical terms:

- A smaller target size (e.g., $\alpha = 0.01$) is more conservative, leading to fewer variables being selected as only those with strong statistical significance are retained.
- A larger target size (e.g., $\alpha = 0.05$) is less stringent, allowing more variables to be included, even if their significance is somewhat weaker.

The choice of target size directly impacts the number of control variables selected, which can influence both the results and the robustness of the model. However, unlike for Lasso, the tuning parameter of Autometrics is more straightforward and intuitive to select.

Lasso methods and Autometrics both start with a general model containing potentially many covariates, from which the goal is to select the most relevant ones. However, unlike Lasso, which estimates parameters through a single optimization process, Autometrics employs an iterative procedure, reducing the model using statistical tests and a tree search algorithm. As a result, deriving theoretical results for Autometrics is more challenging, and its properties are primarily assessed through Monte Carlo simulations, as shown in Hendry and Johansen (2011). Nonetheless, theoretical insights on similar general-to-specific model selection procedures and algorithms can be found in Johansen and Nielsen (2009), Hendry and Johansen (2011), Hendry and Johansen (2015), and Johansen and Nielsen (2016).

Autometrics allows to force the presence of variable d in Equation (1) in the terminal model. However, Autometrics is expected to miss relevant control variables with a low degree of non-centrality (i.e., true t -statistic) and, if these variables are strongly correlated with d , this can lead to an OVB for δ , similarly to Post-Lasso.

Therefore, following the approach of Belloni et al. (2014), we propose an extension of Autometrics, which we call Post-Double-Autometrics. Our method is similar to the Post-Double-Lasso described above, with Lasso being replaced by Autometrics in the first two steps. By giving covariates omitted in the first step a chance to be recovered in the second step, this method is expected to improve over the usual implementation of Autometrics. We expect Post-Double-Autometrics to address the OVB issue of Post-Double-Lasso, as Autometrics selects variables

based on statistical inference rather than coefficient magnitude. This means Autometrics will not overlook relevant variables, even if their coefficient magnitudes are small, as long as they are statistically significant.

Post-Double-Autometrics shares similarities with the indicator saturation method proposed by Hendry et al. (2006). Their approach consists in splitting the initial sample in two blocks, and applying a general-to-specific variable selection method on each block. A final variable selection is then applied on the union of variables selected in each block. However, the dependent variable considered in both blocks of the indicator saturation method is the same (y). Moreover, the main objective of the indicator saturation method is to detect outliers or breaks and not estimating an average treatment effect as in the Post-Double-Autometrics.

3 Simulations

In this section, we use Monte Carlo simulations to study the finite sample properties of different methods for estimating a parameter of interest in high-dimensional linear regression. We consider two single-selection methods, the naive Post-Lasso and Autometrics, and three double-selection methods, the benchmark Post-Double-Lasso, the extension Post-Double-Adaptive-Lasso, and our proposed Post-Double-Autometrics.² Monte Carlo simulations are performed on 1,000 replications.

The data generating process (DGP) corresponds to Equations (1)-(2), where $\mathbf{X}_i \stackrel{i.i.d}{\sim} \mathcal{N}(\mathbf{0}_p, \Sigma_{\mathbf{X}})$, $\Sigma_{\mathbf{X}}$ is a Toeplitz matrix with element (k, l) set to $\rho^{|k-l|}$ for $k, l = 1, \dots, p$, $\epsilon_i \stackrel{i.i.d}{\sim} \mathcal{N}(0, \sigma_\epsilon^2)$ and $\eta_i \stackrel{i.i.d}{\sim} \mathcal{N}(0, \sigma_\eta^2)$. The parameter ρ controls the degree of correlation between the covariates. The DGP is general enough to consider cases where variables are positively or negatively correlated (when ρ is respectively positive or negative) but also weakly or strongly correlated (i.e., when $|\rho|$ is respectively close to 0 or to 1).

To assess the variable selection properties of the different methods, we consider a few relevant variables ($\beta_j \neq 0, \gamma_j \neq 0$, for $j = 1, \dots, 10$) and many irrelevant variables ($\beta_j = \gamma_j = 0$, for $j = 11, \dots, p$). Instead of choosing arbitrary values for the non-zero coefficients, we calibrate their individual statistical significance as follows:

$$\beta_j = \psi_j^y \sqrt{\sigma_\epsilon^2 [\Sigma_{\mathbf{Z}}^{-1}]_{jj}}, \quad (8)$$

$$\gamma_j = \psi_j^d \sqrt{\sigma_\eta^2 [\Sigma_{\mathbf{X}}^{-1}]_{jj}}, \quad (9)$$

²For the adaptive-Lasso, we set $\eta = 1$ and $\hat{\mu}_j$ to the estimate of an OLS regression when $p + 1 < N/2$ and a ridge regression otherwise.

where ψ_j^y and ψ_j^d are the non-centrality parameters of the relevant covariates \mathbf{X}_j , $j = 1, \dots, 10$, in (1)-(2), $\mathbf{Z} = (\mathbf{X}, d)$, and

$$\Sigma_Z = \begin{pmatrix} \Sigma_{\mathbf{X}} & \Sigma_{\mathbf{X}}\gamma \\ \gamma'\Sigma_{\mathbf{X}} & \gamma'\Sigma_{\mathbf{X}}\gamma + \sigma_\eta^2 \end{pmatrix}.$$

The values of ψ_j^y and ψ_j^d correspond to the expected t -statistics for, respectively, the null hypotheses $H_0 : \beta_j = 0$ in Equation (1) and $H_0 : \gamma_j = 0$ in Equation (2). They allow to control the relevance of covariates, since the probability to reject a null hypothesis is high (low) for large (small) value of the non-centrality parameter. To simplify the notation, we will denote by ψ^y (resp. ψ^d) the non-centrality parameters of all non-zero coefficients β_j (resp. γ_j) of the 10 relevant covariates \mathbf{X}_j , $j = 1, \dots, 10$.

To study the finite sample performance of the different methods mentioned above, we report the average bias and root mean squared error (RMSE) of $\hat{\delta}$ as well as the proportion of selected relevant (potency) and irrelevant (gauge) variables included in the terminal model. High potency is important to avoid biased estimation, while for a given level of potency, a smaller gauge is expected to lead to more efficiency (i.e., lower variance of $\hat{\delta}$).

3.1 Independent covariates

In the first simulation, we set $n = 400$, $p = 210$, $\rho = 0$, i.e., $\Sigma_{\mathbf{X}} = I_p$, $\sigma_\epsilon^2 = \sigma_\eta^2 = 1$, and our parameter of interest $\delta = 0$. In this case, the covariates are orthogonal which is the most favorable case for all methods. Moreover, we use $\psi_j^y = 2.5$ and $\psi_j^d = 4$ for $j = 1, \dots, 10$, so that coefficients of the relevant variables have expected t -statistics equal to 2.5 for $H_0 : \beta_j = 0$ in Equation (1), and equal to 4 for $H_0 : \gamma_j = 0$ in Equation (2). The probability for these variables to be kept in the final model is therefore quite low in the first equation, and higher in the second equation. In this framework, double-selection methods are expected to outperform single-selection methods.

Table 1 shows bias and RMSE of the treatment effect $\hat{\delta}$, as well as the average proportion of relevant (potency) and irrelevant (gauge) variables. Note that for the Post-Double-Lasso and Post-Double-Autometrics the gauge and potency are computed on the union of \mathbf{X}^* and \mathbf{X}^{**} , and we consider several choices of tuning parameters λ and α .

First, double-selection methods provide substantial improvements over single selection methods, regardless of the choice of the variable selection method considered. The potency is always increased in the double-selection approaches, leading to lower bias and RMSE. Overall, the results suggest that single-selection methods are always biased in this setting, unlike double-selection methods. Therefore,

Model ($\psi^y = 2.5, \psi^d = 4$)	Bias	RMSE	Potency	Gauge
OLS	0.0010	0.0749	1.0000	1.0000
Post-Lasso, λ^{min}	0.1036	0.1298	0.3868	0.0532
Post-Lasso, λ^{lse}	0.1723	0.1802	0.0350	0.0024
Post-Lasso, λ^{bya}	0.1803	0.1856	0.0000	0.0000
Post-Lasso, λ^{bcch}	0.0833	0.1032	0.4857	0.0487
Post-Double-Lasso, λ^{min}	0.0011	0.0598	0.9999	0.2156
Post-Double-Lasso, λ^{lse}	0.0343	0.0698	0.7718	0.0192
Post-Double-Lasso, λ^{bya}	0.1660	0.1720	0.1000	0.0000
Post-Double-Lasso, λ^{bcch}	0.1690	0.1758	0.0454	0.0000
Post-Double-Adaptive-Lasso, λ^{min}	0.0017	0.0597	1.0000	0.3212
Post-Double-Adaptive-Lasso, λ^{lse}	0.0070	0.0564	0.9479	0.0858
Autometrics, $\alpha = 0.05$	0.0832	0.1139	0.4342	0.0364
Autometrics, $\alpha = 0.01$	0.1253	0.1451	0.2316	0.0088
Post-Double-Autometrics, $\alpha = 0.05$	0.0097	0.0556	0.9056	0.0821
Post-Double-Autometrics, $\alpha = 0.01$	0.0264	0.0603	0.8059	0.0204

Table 1: Bias, RMSE, Gauge and Potency of $\hat{\delta}$ with independent covariates ($\rho = 0$). Design: $\psi^y = 2.5, \psi^d = 4, n = 400, p = 210$.

we focus on double-selection methods in the following and leave single-selection methods out of the discussion.

Second, the performance of Post-Double-Lasso methods depends on the choice of the λ penalty term. Most of the relevant variables are missed (potency close to 0) with λ^{bya} and λ^{bcch} , leading to a treatment effect poorly estimated (large bias and RMSE). The bias is close to 0 for Post-Double-Lasso with λ^{min} as almost all relevant variables are selected (potency close to 1). However, λ^{min} selects too many (i.e., more than 20%) irrelevant variables which translates into a small loss of efficiency. Interestingly, Post-Double-Lasso with λ^{lse} selects very few irrelevant variables (gauge close to 2%) but also misses more than 20% of relevant variables which translated into a much larger bias. Results suggest that λ^{min} is the best choice for the penalty term of Post-Double-Lasso methods as it leads to an unbiased estimate of δ and a small RMSE.

Third, both the Post-Double-Adaptive-Lasso with λ^{min} and λ^{lse} are unbiased and have a RMSE slightly lower than Post-Double-Lasso with λ^{min} . Post-Double-Adaptive-Lasso with λ^{lse} performs particularly well because it has a high potency (close to 95%) and a gauge lower than 10%.

Fourth, Post-Double-Autometrics leads to large potency, and thus to small bias, regardless of the choice of α . Moreover, as Autometrics controls the proportion of irrelevant selected variables, while selecting most relevant variables, the gauge of Post-Double-Autometrics remains relatively small unlike Post-Double-

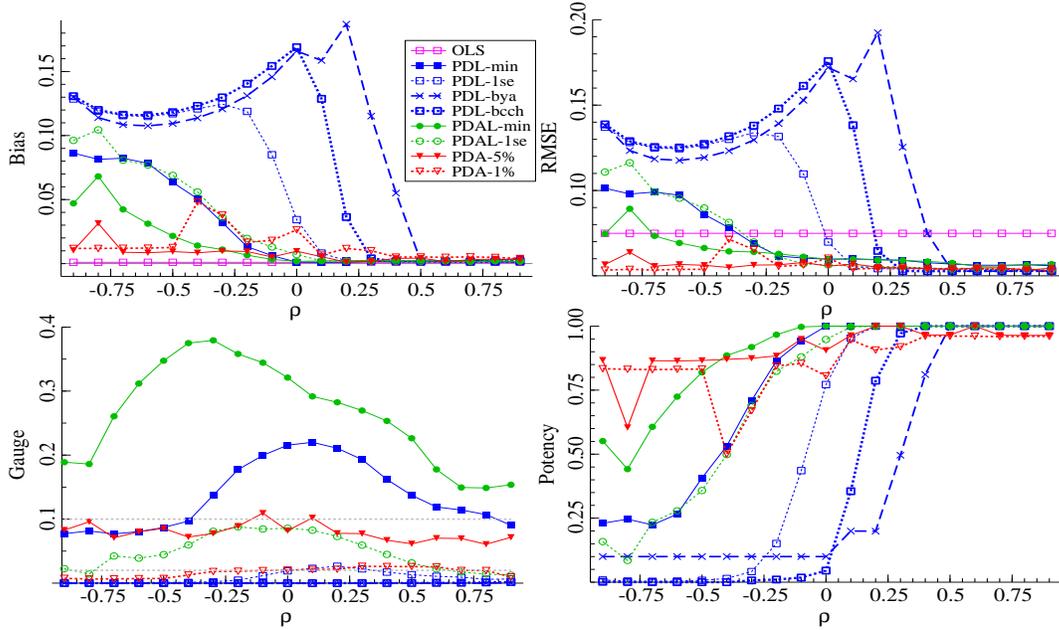


Figure 2: Bias, RMSE, Gauge and Potency of $\hat{\delta}$ with dependent covariates $\rho \in [-0.9; 0.9]$. Design: $\psi^y = 2.5$, $\psi^d = 4$, $n = 400$, $p = 210$.

(Adaptive-)Lasso with λ^{min} . The Post-Double-Autometrics method applies Autometrics successively to two different models and the expected gauge of each model is α . Therefore, if the number of redundant variables is very large (like in our simulation), approximately $2\alpha\%$ of redundant variables are expected to be retained when considering the union between \mathbf{X}^* and \mathbf{X}^{**} , which is exactly what we observe in Table 1.³

Finally, although OLS is unbiased, its RMSE is relatively large compared to those of unbiased double-selection methods. For example, the RMSE of OLS is about 35% larger than Post-Double-Autometrics with $\alpha = 0.05$.

To conclude, Post-Double-Autometrics with $\alpha = 0.05$ competes very well with Post-Double-Lasso with λ^{min} and Post-Double-Adaptive-Lasso with λ^{min} and λ^{1se} in this very simple and unrealistic framework of independent covariates.

3.2 Correlation and relevance of covariates

This section complements the previous simulation results by first allowing the covariates to be correlated and then by varying the non-centrality parameter ψ^d of the 10 relevant covariates in Equation (2).

Figure 2 displays bias and RMSE of $\hat{\delta}$ as well as the potency and gauge for

³Note that the gauge for the standard Autometrics method is approximately equal to α .

$\psi^y = 2.5$, $\psi^d = 4$, $n = 400$ and $p = 210$ dependent control variables, with $\rho \in [-0.9; 0.9]$, for OLS and double-selection methods, across 1,000 replications. The case $\rho = 0$ corresponds to the results in Table 1.⁴

The conclusions from Table 1 are confirmed, but some interesting new results are emerging. Indeed, Post-Double-Lasso methods are again found to be sensitive to the selection method of λ , Post-Double-Lasso with λ^{min} can lead to more accurate estimates than OLS when $\rho > -0.4$, and the Post-Double-Adaptive-Lasso improves the results compared to standard Lasso methods. Importantly, the performance of all Lasso-based methods deteriorates when ρ becomes strongly negative, i.e., the bias and RMSE increase while the potency largely decreases.⁵

In contrast, Post-Double-Autometrics provides reliable results regardless of α and ρ . The bias and RMSE of Post-Double-Autometrics remain close to 0 for all values of ρ , whereas the gauge is still close to 2α . The results thus suggest that Post-Double-Autometrics outperforms other estimation methods as it provides the best treatment effect estimation (lowest bias and RMSE) and the best variable selection (high potency and controlled gauge), while being insensitive to α and to the magnitude of the correlation between the covariates.

Finally, Figures 3-6 complement previous results by displaying respectively the bias, RMSE, potency and gauge as a function of the correlation coefficient, with $\rho \in [-0.9, 0.9]$, and of the non-centrality parameter ψ^d , with $\psi^d \in [1, 8]$, for $\psi^y = 2.5$, $n = 400$ and $p = 210$. The higher ψ^d is, the greater the probability that these variables will be selected, i.e., the more relevant they are. Therefore, the potency is expected to increase with ψ^d , resulting in lower values of the bias and RMSE. Note that the case $\rho = -0.6$ and $\psi^d = 4$ corresponds to the results displayed in Figure 1, while those of Figure 2 have been obtained for $\psi^d = 4$ and $\rho \in [-0.9, 0.9]$. For the sake of clarity, we focus on the best performing methods, i.e., Post-Double-Lasso for $\lambda = \{\lambda^{min}, \lambda^{lse}\}$, Post-Double-Adaptive-Lasso for $\lambda = \{\lambda^{min}, \lambda^{lse}\}$ and Post-Double-Autometrics for $\alpha = \{5\%, 1\%\}$.⁶ In the figures, these methods are denoted, respectively, *PDL-min*, *PDL-lse*, *PDAL-min*, *PDAL-lse*, *PDA-5%* and *PDA-1%*.

The bias in Lasso-based methods is influenced not only by the parameters λ and ρ , but also by the value of ψ^d . Even for large ψ^d , these methods remain biased when ρ is strongly negative, and the bias only disappears when $\psi^d \geq 7.5$. Similar patterns are observed for the RMSE. Furthermore, the results indicate that the gauge increases with ψ^d , reaching levels above 30% with λ^{min} . Comparable findings hold for the adaptive-Lasso methods, though with slightly better performance in terms of bias and RMSE. In summary, while Lasso-based double-selection methods

⁴Results for single-selection methods are available upon request.

⁵Hendry and Doornik (2014) found similar results, demonstrating that Lasso's effectiveness decreases significantly when variables are negatively correlated.

⁶Results for other methods are available upon request.

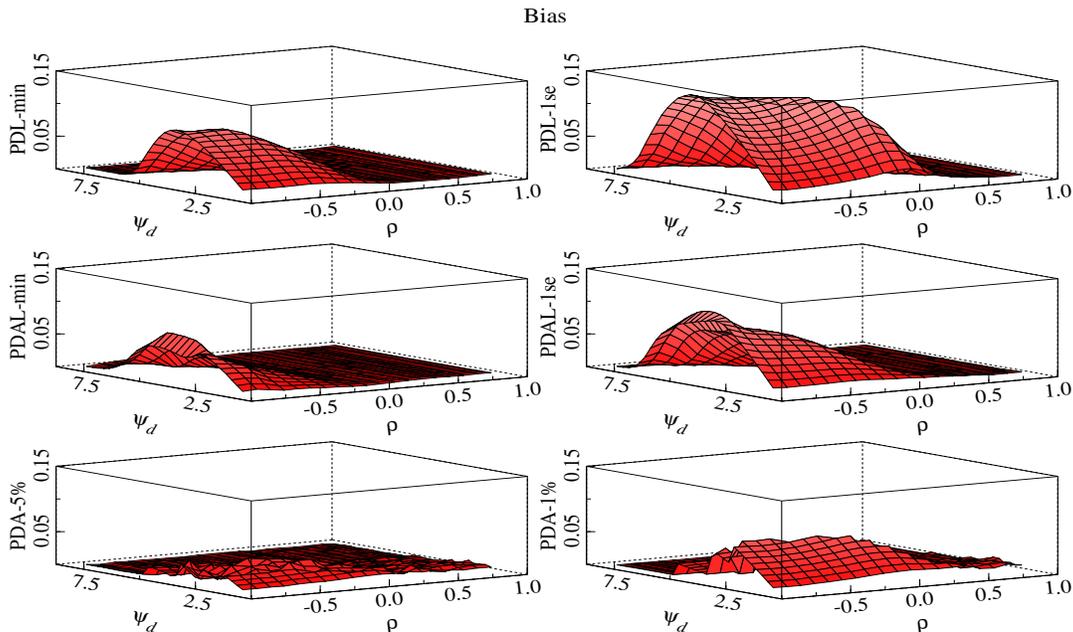


Figure 3: Bias of $\hat{\delta}$, with dependent covariates $\rho \in [-0.9; 0.9]$ and varying non-centrality measure $\psi^d \in [1; 8]$. Design: $\psi^y = 2.5$, $n = 400$, $p = 210$.

can yield accurate estimates of treatment effects, their performance depends on the choice of λ , the correlation among covariates (ρ), and the individual relevance of these covariates (ψ^d).

In contrast, the performance of Post-Double-Autometrics is remarkable in all cases. While the bias remains close to 0 for all values of ψ^d due to a potency close to 1, the gauge does not increase and is close to 2α regardless of the value of ψ^d . Moreover, the RMSE of Post-Double-Autometrics is the lowest among all methods, and especially for $\alpha = 0.05$. Results thus suggest that Post-Double-Autometrics outperforms other methods as it provides the most accurate treatment effect, regardless of the choice of α , the correlation between the covariates and the non-centrality parameter ψ^d .

3.3 More covariates than observations

Finally, we study the robustness of double-selection methods in a high-dimensional linear regression model with more variables than observations. More precisely, we set $n = 200$ and keep all other parameters unchanged, so that $p = 210 > n$. Figures 7-10 display respectively the bias, RMSE, potency and gauge as a function of the correlation coefficient, with $\rho \in [-0.9, 0.9]$, and $\psi^d \in [1, 8]$, for $\psi^y = 2.5$. Results are similar to those of Figures 3-6, but some differences can be observed.

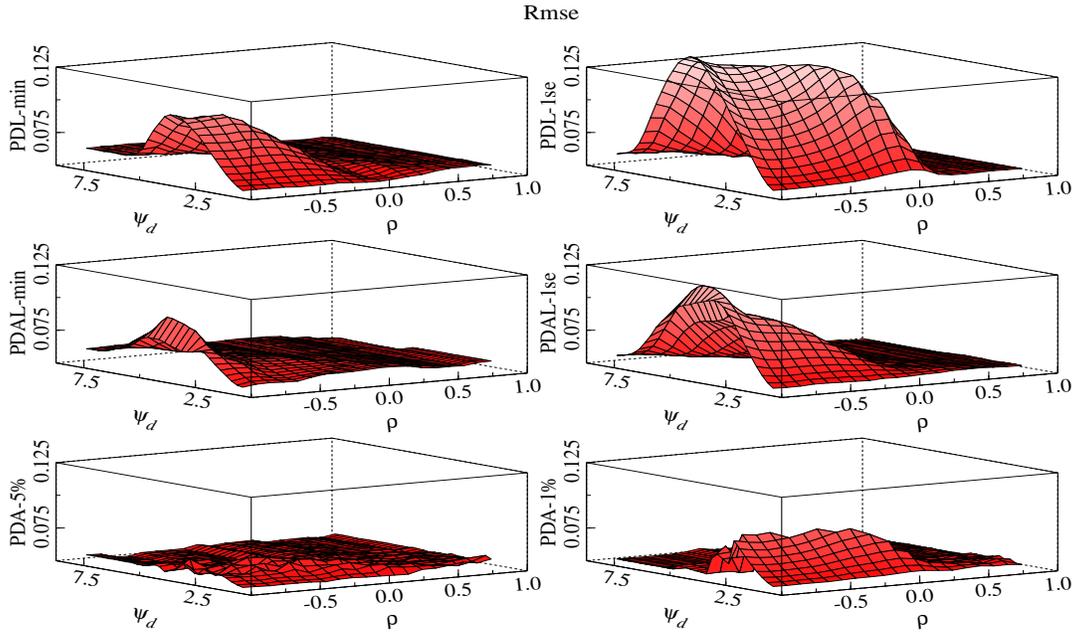


Figure 4: RMSE of $\hat{\delta}$, with dependent covariates $\rho \in [-0.9; 0.9]$ and varying non-centrality measure $\psi^d \in [1; 8]$. Design: $\psi^y = 2.5$, $n = 400$, $p = 210$.

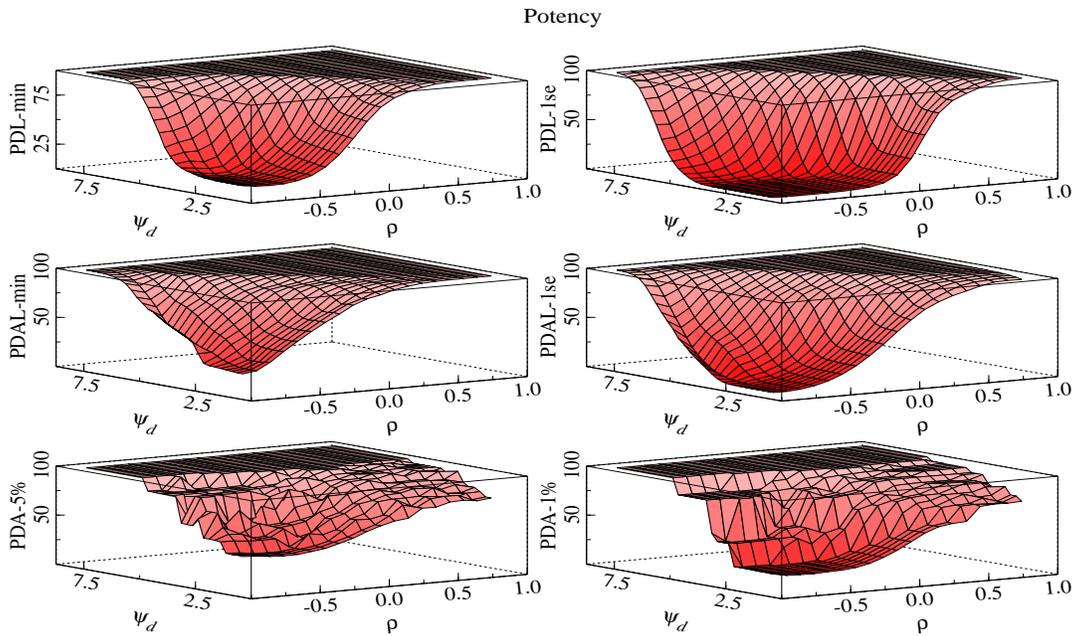


Figure 5: Potency of $\hat{\delta}$, with dependent covariates $\rho \in [-0.9; 0.9]$ and varying non-centrality measure $\psi^d \in [1; 8]$. Design: $\psi^y = 2.5$, $n = 400$, $p = 210$.

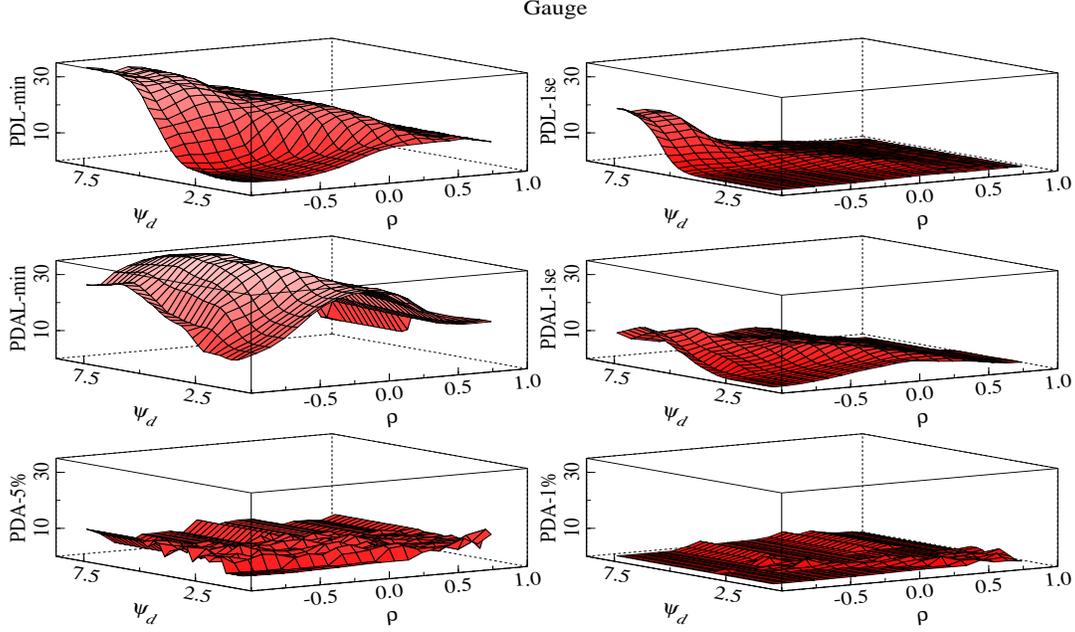


Figure 6: Gauge of $\hat{\delta}$, with dependent covariates $\rho \in [-0.9; 0.9]$ and varying non-centrality measure $\psi^d \in [1; 8]$. Design: $\psi^y = 2.5$, $n = 400$, $p = 210$.

First, the bias of Lasso-based methods is substantially larger than 0 for $p > n$ despite a larger potency. Indeed, as the sample size n decreased compared to the previous simulation but the non-centrality parameters ψ^y and ψ^d remained constant, the non-zeros coefficients in β and γ increased. Therefore, as Lasso selects variables based on the magnitude of the penalized coefficients (in absolute value), more relevant control variables are included in the terminal model, resulting in a larger potency. However, the effect of missed relevant control variables also increases (as β is larger), leading to a larger bias.

Second, the effect of the adaptive-Lasso over the Lasso is unclear for $p > n$. While the bias and RMSE of Post-Double-Adaptive-Lasso and Post Double-Lasso are similar for $\psi^d \leq 4$, they differ for larger values of ψ^d . Indeed, the bias and RMSE of Post-Double-Adaptive-Lasso increases with ψ^d , whereas they decrease towards 0 for Post-Double-Lasso when $\psi^d > 0$ as for the case $p < n$.

Third, Post-Double-Autometrics is still unbiased whatever the values of ρ and ψ^d . As Autometrics selects variables based on their significance, potency and bias remain similar. However, the gauge of Post-Double-Autometrics is slightly larger than 2α . Such result comes from the block-search performed by Autometrics when $p > n$, which leads to a small increase of the gauge.

To conclude, Post-Double-Autometrics also outperforms Post-Double(-Adaptive)-Lasso when $p > n$ and represents a reliable alternative when OLS is unfeasible.

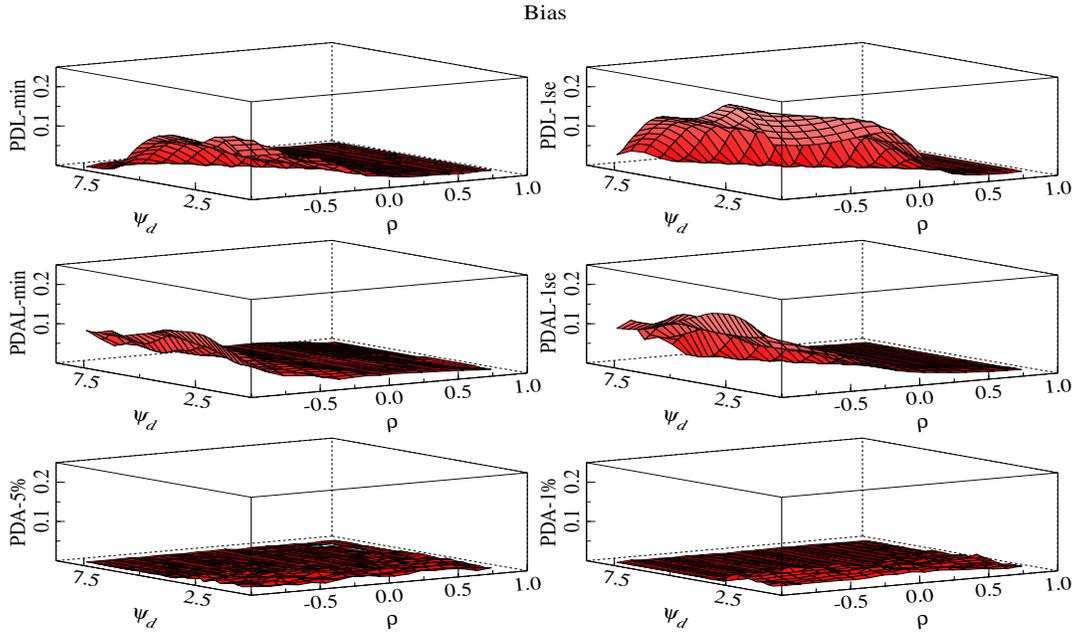


Figure 7: Bias of $\hat{\delta}$ when $p > n$, with dependent covariates $\rho \in [-0.9; 0.9]$ and varying non-centrality measure $\psi^d \in [1; 8]$. Design: $\psi^y = 2.5$, $n = 200$, $p = 210$.

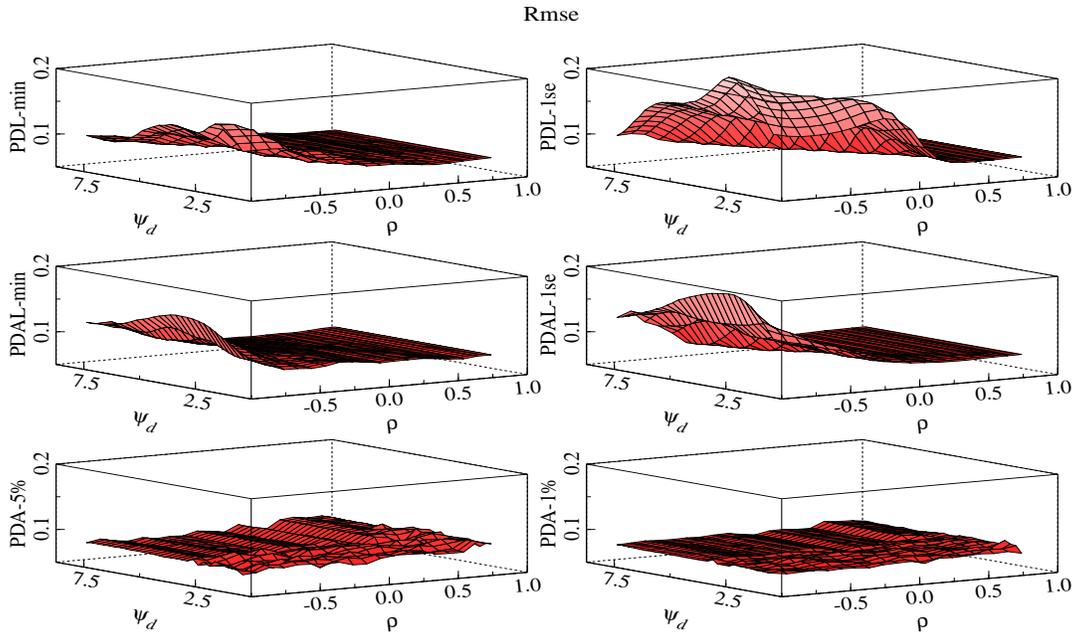


Figure 8: RMSE of $\hat{\delta}$ when $p > n$, with dependent covariates $\rho \in [-0.9; 0.9]$ and varying non-centrality measure $\psi^d \in [1; 8]$. Design: $\psi^y = 2.5$, $n = 200$, $p = 210$.

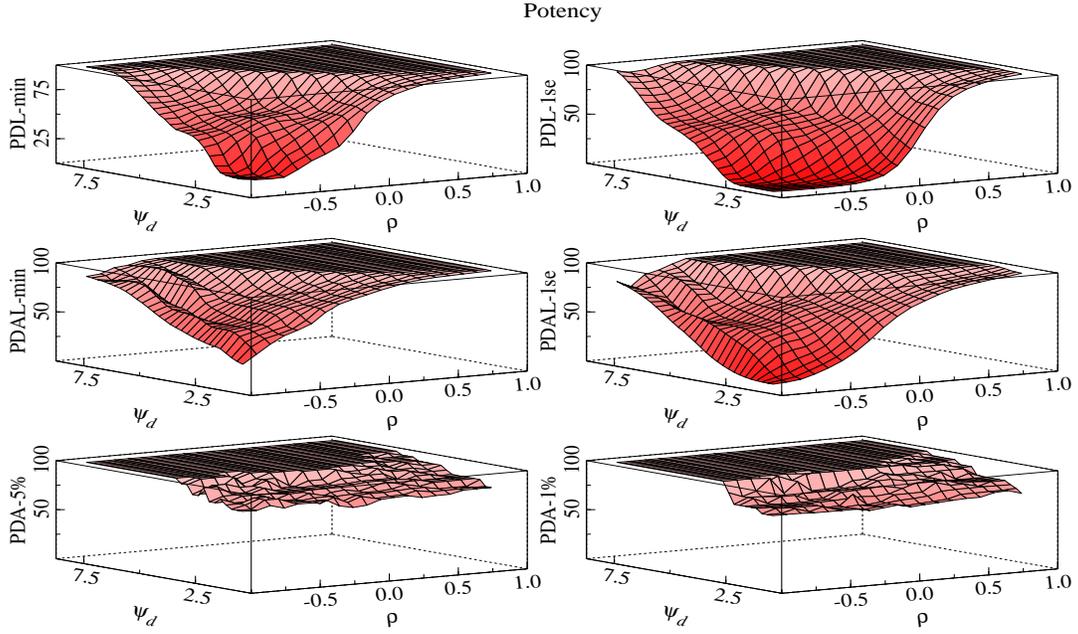


Figure 9: Potency of $\hat{\delta}$ when $p > n$, with dependent covariates $\rho \in [-0.9; 0.9]$ and varying non-centrality measure $\psi^d \in [1; 8]$. Design: $\psi^y = 2.5$, $n = 200$, $p = 210$.

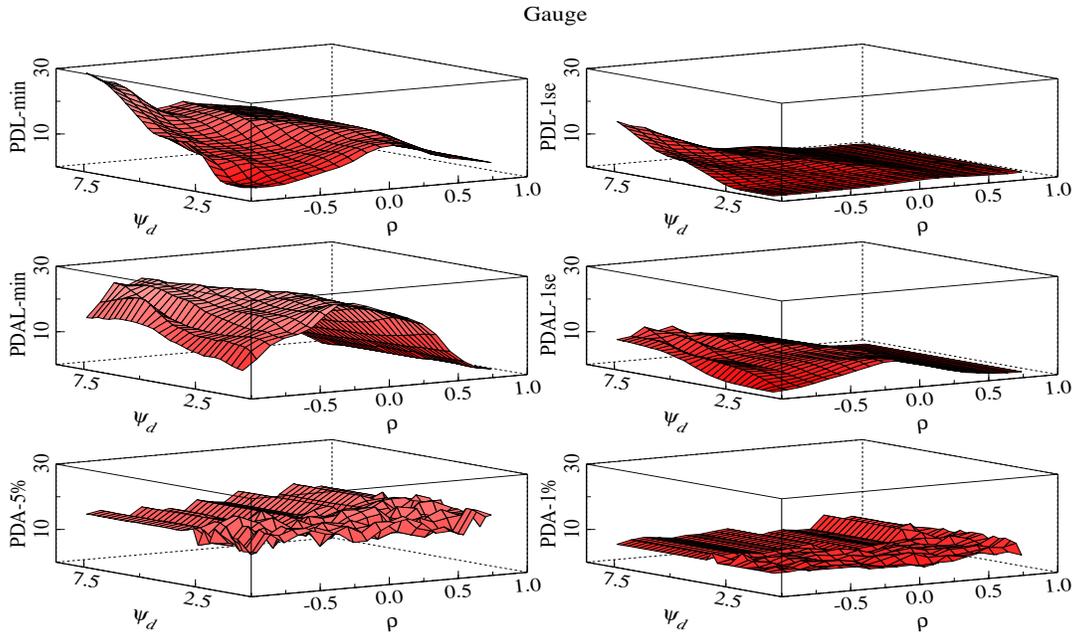


Figure 10: Gauge of $\hat{\delta}$ when $p > n$, with dependent covariates $\rho \in [-0.9; 0.9]$ and varying non-centrality measure $\psi^d \in [1; 8]$. Design: $\psi^y = 2.5$, $n = 200$, $p = 210$.

4 Empirical Application

In this section, we revisit an international economic growth example. In the empirical growth literature, estimating the effect of an initial level of GDP per capita on the growth rates of GDP per capita is a key issue. Indeed, an important finding of the Solow-Swan-Ramsey growth model is the hypothesis of convergence, stating that poorer countries should catch up with richer ones over time as they grow faster. Following this hypothesis, the effect of the initial level of GDP on its growth rate, for a given country, should be negative. To assess whether or not this assumption holds empirically, we use $n = 90$ complete observations extracted from the database of Barro and Lee (1994).⁷ The original dataset consists of a panel of 138 countries for the period of 1960 to 1985, and includes $p = 60$ control variables (\mathbf{X}), such as variables measuring the education, trade openness or mortality rate. The dependent variable (y) corresponds to the national growth rates in GDP per capita for the periods 1965-1975 and 1975-1985, and the variable of interest (d) is the real GDP per capita in 1965. See Table 2 for a complete description of the variables included in the dataset.

Many contributions in the literature investigated the convergence hypothesis. Using a simple bivariate regression model of growth rates on the initial level of GDP, Barro and Sala-i Martin (1995) showed that the convergence hypothesis is rejected. However, this estimate may be biased as the model does not take into account other countries' characteristics potentially affecting the growth rate, the initial level of GDP, or both. The literature thus focused on estimating this effect conditional on some characteristics (Levine and Renelt, 1992; Sala-I-Martin, 1997; Sala-i Martin et al., 2004). Given that the number of control variables is close to the number of observations, covariate selection is critical and the findings of these articles were severely criticized as they relied on ad-hoc procedures. More recently, using the same sample of 90 countries than in our application, Belloni and Chernozhukov (2011) and Belloni et al. (2013) shed light on the convergence hypothesis using Lasso-based approaches. The authors find that a negative and statistically significant effect of the initial level of GDP on its growth rate, supporting the convergence hypothesis of the Solow-Swan-Ramsey growth model. We revisit this application using the different methods presented in Section 2 to show whether or not our results are supporting those of the literature, and ultimately if the convergence hypothesis is verified.

Table 3 displays the results obtained for Post-Lasso and Post-Double-Lasso with $\lambda = \{\lambda^{bya}, \lambda^{min}, \lambda^{lse}, \lambda^{bch}\}$, Post-Adaptive-Lasso and Post-Double-Adaptive-Lasso with $\lambda = \{\lambda^{min}, \lambda^{lse}\}$, Autometrics and Post-Double-Autometrics with $\alpha = \{0.05, 0.01\}$, as well as OLS regressions with and without control variables. We

⁷The data is available in the R library `hdm` of Chernozhukov et al. (2016) using the command `data("GrowthData")`.

Name	Description	Dependent variable
Outcome	National growth rates in GDP per capita for the periods 1965-1975 and 1975-1985	<u>Variable of interest</u>
gdps465	Real GDP per capita (1980 international prices) in 1965	<u>Control variables</u>
bmp11	Black market premium Log (1+BMP)	
freeop	Measure of "Free trade openness"	
freetar	Measure of tariff restriction	
h65	Total gross enrollment ratio for higher education in 1965	
hm65	Male gross enrollment ratio for higher education in 1965	
hf65	Female gross enrollment ratio for higher education in 1965	
p65	Total gross enrollment ratio for primary education in 1965	
pm65	Male gross enrollment ratio for primary education in 1965	
pf65	Female gross enrollment ratio for primary education in 1965	
s65	Total gross enrollment ratio for secondary education in 1965	
sm65	Male gross enrollment ratio for secondary education in 1965	
sf65	Female gross enrollment ratio for secondary education in 1965	
fert65	Total fertility rate (children per woman) in 1965	
mort65	Infant Mortality Rate in 1965	
lifee065	Life expectancy at age 0 in 1965	
gpop1	Growth rate of population	
fert1	Total fertility rate (children per woman)	
mort1	Infant Mortality Rate (ages 0-1)	
invsh41	Ratio of real domestic investment (private plus public) to real GDP	
geetot1	Ratio of total nominal government expenditure on education to nominal GDP	
geerec1	Ratio of recurring nominal government expenditure on education to nominal GDP	
gde1	Ratio of nominal government expenditure on defense to nominal GDP	
govwb1	Ratio of nominal government "consumption" expenditure to nominal GDP (using current local currency)	
govsh41	Ratio of real government "consumption" expenditure to real GDP (Period average)	
gvxdxe41	Ratio of real government "consumption" expenditure net of spending on defense and on education to real GDP	
high65	Percentage of "higher school attained" in the total pop in 1965	
highm65	Percentage of "higher school attained" in the male pop in 1965	
highf65	Percentage of "higher school attained" in the female pop in 1965	
highc65	Percentage of "higher school complete" in the total pop	
highcm65	Percentage of "higher school complete" in the male pop	
highcf65	Percentage of "higher school complete" in the female pop	
human65	Average schooling years in the total population over age 25 in 1965	
humanm65	Average schooling years in the male population over age 25 in 1965	
humanf65	Average schooling years in the female population over age 25 in 1965	
hyr65	Average years of higher schooling in the total population over age 25	
hyrm65	Average years of higher schooling in the male population over age 25	
hyrf65	Average years of higher schooling in the female population over age 25	
no65	Percentage of "no schooling" in the total population	
nom65	Percentage of "no schooling" in the male population	
nof65	Percentage of "no schooling" in the female population	
pinstab1	Measure of political instability	
pop65	Total Population in 1965	
worker65	Ratio of total Workers to population	
pop1565	Population Proportion under 15 in 1965	
pop6565	Population Proportion over 65 in 1965	
sec65	Percentage of "secondary school attained" in the total pop in 1965	
secm65	Percentage of "secondary school attained" in male pop in 1965	
secf65	Percentage of "secondary school attained" in female pop in 1965	
secc65	Percentage of "secondary school complete" in the total pop in 1965	
seccm65	Percentage of "secondary school complete" in the male pop in 1965	
seccf65	Percentage of "secondary school complete" in female pop in 1965	
syr65	Average years of secondary schooling in the total population over age 25 in 1965	
syrm65	Average years of secondary schooling in the male population over age 25 in 1965	
syrf65	Average years of secondary schooling in the female population over age 25 in 1965	
teapri65	Pupil/Teacher Ratio in primary school	
teasec65	Pupil/Teacher Ratio in secondary school	
ex1	Ratio of export to GDP (in current international prices)	
im1	Ratio of import to GDP (in current international prices)	
xr65	Exchange rate (domestic currency per U.S. dollar) in 1965	
tot1	Terms of trade shock (growth rate of export prices minus growth rate of import prices)	

Table 2: Name and description of the variables of the application on growth data

Method	$\hat{\delta}$	Robust s.e.	90% CI	k^*
OLS without control variables	0.0013	0.0053	[-0.0074 ; 0.0101]	0
OLS with all control variables	-0.0094	0.0324	[-0.0627 ; 0.0439]	60
OLS, HCK standard errors	-0.0094	0.0354	[-0.0676 ; 0.0489]	60
Post-Lasso, λ^{bya}	0.0013	0.0053	[-0.0074 ; 0.0101]	0
Post-Lasso, λ^{min}	-0.0510	0.0138	[-0.0737; -0.0282]	20
Post-Lasso, λ^{1se}	-0.0347	0.0119	[-0.0543; -0.0151]	10
Post-Lasso, λ^{bcch}	-0.0174	0.0093	[-0.0328; -0.0020]	7
Post-Adaptive-Lasso, λ^{min}	-0.0429	0.0136	[-0.0652; -0.0205]	12
Post-Adaptive-Lasso, λ^{1se}	-0.0439	0.0120	[-0.0637; -0.0241]	7
Post-Double-Lasso, λ^{bya}	-0.0419	0.0167	[-0.0694; -0.0144]	7
Post-Double-Lasso, λ^{min}	-0.0314	0.0207	[-0.0655 ; 0.0027]	23
Post-Double-Lasso, λ^{1se}	-0.0563	0.0191	[-0.0876; -0.0249]	13
Post-Double-Lasso, λ^{bcch}	-0.0500	0.0158	[-0.0760; -0.0240]	7
Post-Double-Adaptive-Lasso, λ^{min}	-0.0446	0.0168	[-0.0723; -0.0169]	21
Post-Double-Adaptive-Lasso, λ^{1se}	-0.0538	0.0171	[-0.0819; -0.0258]	7
Autometrics, $\alpha = 0.05$	-0.0054	0.0118	[-0.0248 ; 0.0140]	12
Autometrics, $\alpha = 0.01$	-0.0430	0.0106	[-0.0606; -0.0255]	3
Post-Double-Autometrics, $\alpha = 0.05$	-0.0257	0.0188	[-0.0566 ; 0.0052]	14
Post-Double-Autometrics, $\alpha = 0.01$	-0.0143	0.0196	[-0.0465 ; 0.0179]	11

Table 3: Estimates, robust standard errors, 90% confidence intervals of the effect of the initial level of GDP δ , and number of selected variables k^* for the application on the growth data. Coefficients not significantly different from 0 at the 10% nominal level are displayed in bold.

also report the results of an OLS regression involving all control variables with the correction of Cattaneo et al. (2018) to control for high-dimension and heteroscedasticity. The results of the OLS methods without control variables suggest that the convergence hypothesis is rejected as the coefficient is not significant at the 10% nominal level, which is consistent with the finding of Barro and Sala-i Martin (1995). This result should be taken with caution as this model is subject to omitted variable bias due to the non-inclusion of any control variables. The inclusion of 60 control variables in the model (with or without correction of the standard error of $\hat{\delta}$) does not change the conclusion. However, this result should be interpreted with caution due to the high number of variables in the model, leading to a relatively large standard error.

To obtain more accurate estimates of δ , we rely on the methods presented in Section 2 to select the relevant control variables. The results show that the coefficient $\hat{\delta}$ is always negative, except for the two methods that do not select any covariates, i.e., OLS without control variables and Post-Lasso with λ^{bya} . Consistent with the findings of our Monte Carlo simulations, results of the Post-Lasso and Post-Double-Lasso are very sensitive to the choice of λ . While δ is significantly different from 0 at the 10% nominal level for both Post-Adaptive-Lasso and Post-Double-Adaptive-Lasso, the results for Post-Lasso and Post-Double-Lasso methods depend on the choice of λ . Specifically, δ is significantly different from 0 for all Post-Lasso and Post-Double-Lasso methods except for Post-Lasso with λ^{bya} and Post-Double-Lasso with λ^{min} .

Among the Post-Double-Lasso methods, λ^{min} selects more than 10 additional control variables compared with λ^{bya} , λ^{lse} and λ^{bcch} and out of these four methods is the only one allowing to reject the hypothesis of convergence (or equivalently to fail to reject the null hypothesis $H_0 : \delta = 0$). Naturally, decreasing the penalization parameter increases the number of selected control variables, which affects the statistical inference on δ .⁸ It echoes with the recommendation of Wüthrich and Zhu (2023) to check the sensibility of the results to the choice of λ by varying λ in a grid around the chosen ‘optimal’ value. Similar results are observed for the Post-Double-Adaptive-Lasso, even though δ remains significant for both values of λ .

Autometrics results differ between the two target sizes considered in the application. Specifically, while δ is significantly different from 0 for $\alpha = 0.01$ (with only 3 control variables selected), δ is not significantly different from 0 for $\alpha = 0.05$ (with 12 control variables selected).

In contrast, Post-Double-Autometrics provides consistent results for both values of α and suggests that the convergence hypothesis is not verified (i.e., $H_0 : \delta = 0$

⁸Similarly, Belloni and Chernozhukov (2011) show in their application that despite the fact that the number of selected control variables increases as λ decreases, their main conclusion remained unaffected by the choice of λ .

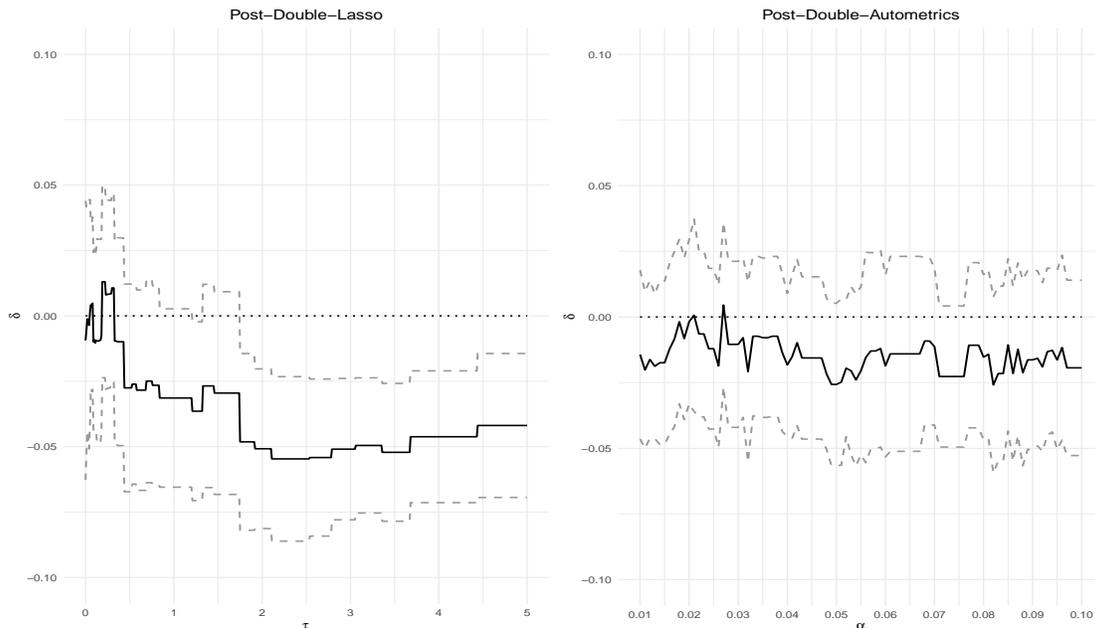


Figure 11: Estimate (solid line) and 90% confidence interval (dashed line) of δ with the Post-Double-Lasso and penalty term $\lambda = \tau\lambda^{min}$ for different values of τ (Left), and with Post-Double-Autometrics for different values of α (Right).

is not rejected). This conclusion does not support the previous findings of the literature obtained with Lasso-based methods.⁹ Moreover, the results obtained are more accurate than those of OLS methods with all control variables as the standard errors of $\hat{\delta}$ are much smaller.

To illustrate the (lack of) robustness of the Post-Double-Lasso, Figure 11 (left panel) displays the estimates and 90% confidence intervals of $\hat{\delta}$ with λ in the first two steps of the Post-Double-Lasso set to $\tau\lambda^{min}$ with $\tau \in (0, 5)$.¹⁰ The results show the sensitivity of the statistical inference on δ for Post-Double-Lasso to the value of λ . Indeed, while the coefficient is not significant for $\tau < 1.75$, it becomes negative and significant for $\tau \geq 1.75$. Results obtained for λ^{bya} , λ^{lse} , λ^{bch} are qualitatively similar and available upon request. Therefore, the conclusion of Post-Double-Lasso and Post-Double-Adaptive-Lasso regarding the convergence hypothesis strongly depends on the choice of λ .

On the other hand, the results of Post-Double-Autometrics are robust to the choice of α . Figure 11 (right panel) displays the estimate and 90% confidence

⁹Note that our results are also consistent with those of the Post-Double-Lasso with λ^{min} , which has been shown to be the overall best Lasso-based method in our previous section across several Monte Carlo simulation experiments.

¹⁰Note that when $\tau = 0$, the results are identical to those of the OLS method with all control variables, and when $\tau = 1$ the results are those displayed in Table 3.

interval of $\hat{\delta}$ for $\alpha \in [1\%, 10\%]$. The results show that the estimates and confidence intervals are stable across all considered values of α , and that the convergence hypothesis is never satisfied. In summary, unlike Belloni and Chernozhukov (2011) and Belloni et al. (2013), our results do not support the convergence hypothesis of the Solow-Swan-Ramsey growth model.

Finally, Table 4 displays the control variables selected by Post-Double-Autometrics, Post-Double-Lasso and Post-Double-Adaptive-Lasso. Among all control variables, the most important ones are the black market premium (*bmp1l*), which is a proxy of trade openness, the life expectancy at birth (*life065*), the percentage of workers within the population (*worker65*), and the ratio of export to GDP (*ex1*) as they are selected by several methods. Interestingly, the variable measuring the ratio of import to GDP (*im1*) is only selected by the three double-methods for which the convergence hypothesis is not verified, i.e., Post-Double-Autometrics with $\alpha = \{0.05, 0.01\}$ and Post-Double-Lasso with λ^{min} . This variable is likely the primary factor influencing the significance or non-significance of δ in the competing post-double methods, thereby contributing to the OVB. To validate this finding, we analyze the role of variable *im1* in the model.

First, Table 5 reports the *t*-statistic associated with the coefficient of *im1* in the two steps of Post-Double-Autometrics with $\alpha = \{0.05, 0.01\}$. The results confirm that *im1* is relevant, as it is selected in both steps. While the *t*-statistics are positive in the first step, they are negative in the second step. This suggests that, conditional on the other selected control variables, the correlation between *d* and *im1* is negative. Furthermore, the *t*-statistics are relatively small (ranging from 2 to 4 in absolute value). This mirrors previous Monte Carlo simulation findings, where we found that when the non-centrality parameters of relevant variables were small both in Equations (1) and (2), Post-Double-Autometrics selects relevant variables more frequently than Post-Double-Lasso.

Second, Table 6 shows the confidence intervals for the three methods, with and without variable *im1* in the final OLS regression. Interestingly, while the estimated effect of initial GDP, i.e., δ , is not significant when *im1* is included in the final OLS regression, it becomes significant when it is omitted. $\hat{\delta}$ remains negative, regardless of its significance, but the coefficient of *im1* is positive when included in the final model (0.274, 0.355, and 0.444 for Post-Double-Lasso with λ^{min} and Post-Double-Autometrics at $\alpha = \{0.05, 0.01\}$, respectively), leading to an insignificant δ . This suggests that omitting *im1* from the model introduces a negative bias. This finding highlights the importance of including *im1* and explains why models excluding this variable result in a significant coefficient for initial GDP. Contrary to Levine and Renelt (1992), we thus find that the share of imports is not substitutable to export share, but that each of these two variables should be included in the terminal model.

Finally, these three methods also select other variables in common, such as

	PDA		PDL				PDAL	
	$\alpha = 0.05$	$\alpha = 0.01$	λ^{bya}	λ^{min}	λ^{lse}	λ^{bcch}	λ^{min}	λ^{lse}
bmp11	✓	✓		✓	✓	✓	✓	
pf65	✓	✓						
lifee065	✓	✓	✓	✓	✓	✓	✓	✓
mort1	✓	✓					✓	
worker65	✓	✓		✓	✓		✓	✓
ex1	✓	✓		✓	✓		✓	
im1	✓	✓		✓				
mort65	✓							✓
gpop1	✓						✓	
no65	✓							
nof65	✓			✓	✓			
pop1565	✓			✓			✓	
secc65	✓							
seccm65	✓			✓				
sf65		✓	✓	✓		✓		
pop6565		✓	✓	✓	✓	✓	✓	✓
syr65		✓						
syrf65		✓						
freeop							✓	
fretar			✓	✓	✓	✓	✓	✓
h65							✓	
hm65			✓	✓	✓	✓	✓	✓
hf65				✓			✓	
pm65				✓			✓	
geetot1							✓	
geerec1				✓			✓	
gde1				✓			✓	
govwb1							✓	
govsh41				✓	✓			✓
gvxdxe41					✓		✓	
humanf65			✓	✓	✓	✓		
pinstab1				✓			✓	
seccf65				✓				
teapri65			✓	✓	✓			
teasec65				✓	✓			
xr65				✓				
tot1							✓	

Note: PDA, PDL and PDAL refer to Post-Double-Autometrics, Post-Double-Lasso and Post-Double-Adaptive-Lasso methods, respectively. The symbol “✓” indicates that a variable is selected by the corresponding method.

Table 4: Variables selected by Post-Double-Autometrics, Post-Double Lasso and Post-Double-Adaptive-Lasso methods.

Model	PDA, $\alpha = 0.05$	PDA, $\alpha = 0.01$
Step 1	2.6338	2.1340
Step 2	-2.5068	-3.8457

Table 5: t -statistics of the variable *im1* in the first two steps of the Post-Double-Autometrics method.

	with <i>im1</i>		without <i>im1</i>	
	$\hat{\delta}$	CI	$\hat{\delta}$	CI
PDL, λ^{min}	-0.0314	[-0.0655; 0.0027]	-0.0458	[-0.0732; -0.0183]
PDA, $\alpha = 0.05$	-0.0257	[-0.0566; 0.0052]	-0.0413	[-0.0713; -0.0113]
PDA, $\alpha = 0.01$	-0.0143	[-0.0465; 0.0179]	-0.0393	[-0.0654; -0.0132]

Table 6: Confidence interval with and without the variable *im1* included in the final OLS regression. Coefficients not significantly different from 0 at the 10% nominal level are displayed in bold.

variables related to the population not working (*pop1565* and *pop6565*) or to education (*nof65*, *seccm65* and *sf65*). The results thus suggest that the most important covariates allowing to reject the convergence hypothesis are mainly related to trade openness, infant life expectancy, proportion of worker in the population, exportation and importation (relative to the GDP), and education.¹¹ Furthermore, Post-Double-Lasso with λ^{min} selects several additional variables compared to Post-Double-Autometrics. These ones are mainly related to education, but also to tariff restriction, political instability, exchange rate and government expenditures. However, given that Post-Double-Lasso with λ^{min} selects on average many irrelevant variables (see previous section), most of these variables are probably irrelevant and could be dropped from the analysis.

5 Conclusion

In this paper, we have investigated the shortcomings of the Post-Double-Lasso method for estimating a parameter of interest, as an average treatment effect, in linear regression models with numerous covariates. While Post-Double-Lasso has gained popularity for its ability to mitigate omitted variable bias and manage overfitting through Lasso variable selection, we highlighted its vulnerabilities, particularly its sensitivity to the choice of the regularization parameter and the possibility of substantial omitted variable bias when some of the variables are negatively correlated.

As an alternative to Post-Double-Lasso, we introduced Post-Double-Autometrics, which leverages the Autometrics variable selection method, a method based on statistical inference to select relevant variables. Our findings demonstrate that Post-Double-Autometrics can effectively mitigate bias in estimating treatment effects by retaining relevant variables that might be overlooked by Lasso-based methods. Through extensive Monte Carlo simulations and an empirical illustration,

¹¹Note that one variable related to education (*pf65*) is only selected by both Post-Double-Autometrics methods.

we showed that Post-Double-Autometrics outperforms Post-Double-Lasso, yielding more reliable estimates of the treatment effect.

In this study, we adopt a fully parametric approach, assuming that the general model encompasses the true one — meaning that all relevant variables are included in Equations (1)-(2) prior to variable selection. In contrast, Chernozhukov et al. (2018) propose an alternative approach, Double Machine Learning, that extends Post-Double-Lasso by capturing non-linearities non-parametrically. A valuable extension of this work would be to adapt Post-Double-Autometrics using the recent GAMLA method proposed by Flachaire et al. (2024) that builds upon the Generalized Additive Model (Hastie and Tibshirani, 2017). This would allow the model to capture non-linearities while simultaneously selecting variables assumed to have a linear effect on the outcome variable. We leave this as a topic for future research.

References

- Barro, R. and Sala-i Martin, X. (1995). *Economic Growth*. Advanced Series in Economics Series. McGraw-Hill.
- Barro, R. J. and Lee, J.-W. (1994). Data set for a panel of 138 countries.
- 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. (2011). L1-penalized quantile regression in high-dimensional sparse models. *Annals of Statistics*, 39(1):82–130.
- 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., and Hansen, C. (2013). Inference for high-dimensional sparse econometric models. In Acemoglu, D., Arellano, M., and Dekel, E., editors, *Advances in Economics and Econometrics: Tenth World Congress*, pages 245–295. Cambridge University Press.
- Belloni, A., Chernozhukov, V., and Hansen, C. (2014). Inference on treatment effects after selection among high-dimensional controls. *Review of Economic Studies*, 81(2):608–650.
- Bickel, P. J., Ritov, Y., and Tsybakov, A. B. (2009). Simultaneous analysis of Lasso and Dantzig selector. *The Annals of Statistics*, 37(4):1705 – 1732.
- Castle, J., Doornik, J., and Hendry, D. (2023). Robust discovery of regression models. *Econometrics and Statistics*, 26:31–51.

- Cattaneo, M. D., Jansson, M., and Newey, W. K. (2018). Inference in linear regression models with many covariates and heteroscedasticity. *Journal of the American Statistical Association*, 113(523):1350–1361.
- 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. (2016). hdm: High-dimensional metrics. *R Journal*, 8(2):185–199.
- Chetverikov, D., Liao, Z., and Chernozhukov, V. (2021). On cross-validated lasso in high dimensions. *The Annals of Statistics*, 49(3):1300–1317.
- Doornik, J. A. et al. (2009). Autometrics. *Castle, and Shephard (2009)*, pages 88–121.
- Epprecht, C., Guégan, D., Veiga, A., and Correa da Rosa, J. (2021). Variable selection and forecasting via automated methods for linear models: LASO/adaLASSO and Autometrics. *Communications in Statistics - Simulation and Computation*, 50:103–122.
- Flachaire, E., Hué, S., Laurent, S., and Hacheme, G. (2024). Interpretable machine learning using partial linear models. *Oxford Bulletin of Economics and Statistics*, 86(3):519–540.
- Hastie, T. J. and Tibshirani, R. J. (2017). *Generalized additive models*. Routledge.
- Hendry, D., Johansen, S., and Santos, C. (2006). Selecting a regression saturated by indicators. *Unpublished paper, Economics Department, University of Oxford*.
- Hendry, D. F. (2000). *Econometrics: alchemy or science?: Essays in econometric methodology*. OUP Oxford.
- Hendry, D. F. and Doornik, J. A. (2014). *Empirical model discovery and theory evaluation: automatic selection methods in econometrics*. MIT Press.
- Hendry, D. F. and Johansen, S. (2011). The properties of model selection when retaining theory variables. *University of Copenhagen Discussion Paper*.
- Hendry, D. F. and Johansen, S. (2015). Model discovery and trygve haavelmo’s legacy. *Econometric Theory*, 31(1):93–114.
- Homrighausen, D. and McDonald, D. (2013). The lasso, persistence, and cross-validation. In *International conference on machine learning*, pages 1031–1039. PMLR.

- Homrighausen, D. and McDonald, D. J. (2014). Leave-one-out cross-validation is risk consistent for lasso. *Machine learning*, 97(1):65–78.
- Johansen, S. and Nielsen, B. (2009). An analysis of the indicator saturation estimator as a robust regression estimator. *Castle, and Shephard (2009)*, 1:1–36.
- Johansen, S. and Nielsen, B. (2016). Asymptotic theory of outlier detection algorithms for linear time series regression models. *Scandinavian Journal of Statistics*, 43(2):321–348.
- Levine, R. and Renelt, D. (1992). A sensitivity analysis of cross-country growth regressions. *The American economic review*, pages 942–963.
- Sala-i Martin, X., Doppelhofer, G., and Miller, R. I. (2004). Determinants of long-term growth: A bayesian averaging of classical estimates (bace) approach. *American Economic Review*, 94(4):813–835.
- Sala-I-Martin, X. X. (1997). I just ran two million regressions. *The American Economic Review*, 87(2):178–183.
- Sucarrat, G. (2020). User-specified general-to-specific and indicator saturation methods. *R Journal*, 12(2).
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 58(1):267–288.
- Wüthrich, K. and Zhu, Y. (2023). Omitted variable bias of lasso-based inference methods: A finite sample analysis. *Review of Economics and Statistics*, 105(4):982–997.
- Zou, H. (2006). The adaptive lasso and its oracle properties. *Journal of the American statistical association*, 101(476):1418–1429.

A Additional tables and figures for the empirical illustration

Method	$\hat{\delta}$	Robust s.e.	90% CI	k^*
Post-Elastic-Net, λ^{min}	-0.0502	0.0186	[-0.0808; -0.0197]	45
Post-Elastic-Net, λ^{lse}	-0.0447	0.0169	[-0.0724; -0.0169]	32
Post-Double-Elastic-Net, λ^{min}	-0.0094	0.0324	[-0.0627; 0.0439]	60
Post-Double-Elastic-Net, λ^{lse}	-0.0094	0.0324	[-0.0627; 0.0439]	60

Table 7: Estimates, robust standard errors, 90% confidence intervals of the effect of the initial level of GDP δ , and number of selected variables k^* of elastic-net based methods for the application on the growth data. Coefficients not significantly different from 0 at the 10% nominal level are displayed in bold.