



Munich Personal RePEc Archive

Growth Effects of European Monetary Union: A Synthetic Control Approach

Lucke, Bernd

University of Hamburg

3 November 2022

Online at <https://mpra.ub.uni-muenchen.de/125280/>

MPRA Paper No. 125280, posted 12 Jul 2025 08:30 UTC

Growth Effects of European Monetary Union: A New Synthetic Control Approach

Bernd Lucke¹

Department of Economics

University of Hamburg

March 2024

Abstract

How did European Monetary Union (EMU) affect economic growth? Standard synthetic control methods (SCM) construct counterfactuals from pre-treatment data. But non-EU controls often diverge *post*-treatment. I propose a new method where SC weights are derived from consistent parameter estimates in *post*-treatment time and use controls which differ by EMU- but not by EU-membership. I find that EMU has benefited regions with competitive, export-oriented companies, while it had sizable detrimental growth effects on most French, Italian and Greek regions. Over 18 years, the latter seem to have had a loss in income between 15% and 30% relative to the non-EMU counterfactual.

Keywords: European Monetary Union, synthetic control methods

JEL: C21, E65, F33

¹ Prof. Dr. Bernd Lucke, von Melle Park 5, 20146 Hamburg. Tel. +49-40-42838-3996. Email: bernd.lucke@uni-hamburg.de

1. Introduction

European Monetary Union (EMU), inaugurated in 1999 with the introduction of the Euro, represents a crucial economic policy milestone. Navigating through challenges like the global financial and European sovereign debt crises, the Euro has, until recently, upheld its commitment to ensuring price stability in the Eurozone. However, its impact on economic growth remains uncertain.

Despite over two decades of EMU, scholarly assessments of its effects on real per-capita growth are sparse. Instead, numerous studies have explored EMU's influence on intermediate variables, including trade, capital accumulation and FDI, financial and political integration, real exchange rate volatility, international price elasticities, and institutional development². Not all results are conclusive, and the presumed impact on economic growth is often asserted rather than rigorously demonstrated through econometric methods.

Only few studies directly address EMU's effect on economic growth. Conti (2014) utilizes a differences-in-differences (DiD) framework for seventeen European countries, finding positive EMU effects on GDP per capita, but to a lesser extent for countries with high initial debt levels. Kalaitzoglou and Durgheu (2016) report ambiguous results, attributing them to easier access to finance and a tendency to overborrow. Dreyer and Schmid (2016), using panel GMM methods, conclude that Eurozone membership had no significant effect on growth from 1999 to 2012. Ioannatos (2018) reaches a similar conclusion for a sample covering the Eurozone until 2016.

Another group of studies, akin to this paper in targeting GDP per capita and employing the synthetic control method, includes Fernandez and Garcia Perea (2015), Verstegen et al. (2017), Puzzello and Gomis-Porqueras (2018), Gasparotti and Kullas (2019), and Gabriel and

² The discussion paper version of this article gives an in-depth review of the related literature.

Pessoa (2020). All in all, there is little consensus among these authors, a negative impact on Italy's GDP being the only seemingly robust finding.

The disagreement in results may stem from these studies using country-level data. This is problematic because idiosyncratic policies or large idiosyncratic shocks after the treatment date may make a control country unsuitable for determining the counterfactual. In fact, many OECD countries often used as controls (e. g. Japan, Mexico, the US or Turkey) diverged after 1999 from the Eurozone for reasons completely unrelated to the introduction of the Euro. This is likely to induce sizable biases of the growth effects estimated for the Euro introduction.

To overcome this, I propose two innovations: First, a new synthetic control (SC) method utilizing *post*-treatment structural information in the control regions to determine SC weights. This is in stark contrast to the standard approach of Abadie et al. (2010, henceforth ADH), where SC weights are functions of *pre*-treatment information only. But policies and institutions may change over time, which is why parameters in the underlying common-factors model are usually specified as time-dependent. Hence, SC weights for the counterfactual should be estimated in *post*-treatment time (from regions unaffected by treatment). Below, I show under which conditions this is possible.

Second, I confine the set of controls to units which are part of the EU Common Market (but not the Eurozone). Therefore, all controls share the EU's economic, legal and institutional framework with the treated unit. But they differ in respect to monetary policy – and this is what is intended.

ARDECO (Annual Regional Database of the European Commission) data are used in this study. In the NUTS3 classification, the cross-section dimension comprises 1027 regional units, providing real GDP per capita and a host of other structural information. As compared to country-level controls, the wealth of NUTS3 observations increases the likelihood of matching Eurozone regions with a “synthetic twin” outside the Eurozone (but within the Common Market).

The paper's organization is as follows: Section 2 presents the Rubin causal model in terms of a static factor model for potential outcomes. Section 3 reviews the standard ADH approach, and Section 4 outlines an alternative method. Section 5 provides information on the regional data available in the ARDECO dataset. Section 6 presents and interprets results for GDP growth of Eurozone NUTS3 regions vis-à-vis the estimated counterfactual. Section 7 concludes.

2. A Rubin causal factor model

Synthetic control methods are based on the Rubin causal model (cf. Rubin (1974), Holland (1986)) which defines the causal effect of a certain treatment, e. g. a policy intervention, as the difference between the potential outcome under treatment and the potential outcome under non-treatment. Modern synthetic control methods (SCM) (e. g. Abadie et al. (2010), Abadie (2021), Abadie and L'Hour (2021)) aim at constructing the counterfactual as a weighted mean of the outcomes of control units unaffected by the treatment. This weighted mean is called the synthetic control (SC).

Mathematically, the standard ADH-approach involves solving an optimistic bilevel minimization problem, cf. Malo et al. (2020). This type of problem is mathematically complex and a number of unwelcome issues (which initially went unnoticed), have been detected and discussed in the more recent literature, e. g. Kaul et al. (2015), Ferman and Pinto (2016), Becker and Klößner (2018), Ben-Michael et al. (2021), Chernozhukov et al. (2021).

2.1 The model

Suppose we observe a balanced panel of $K = J + L$ units over $T = T_0 + T_1$ periods of time. L units have been assigned to a policy intervention (the “treatment”) from period $T_0 + 1$ onward. Conditional on covariates (predictors), the assignment was random. J units are unaffected by the treatment. In line with much of the literature, I will first lay out the model as if L were equal

to one. Thus, unit 1 will denote the treated unit. Later I will argue that the total of L units will convey important information helpful to quantify the treatment effect for “unit 1”.

The observed outcome for unit i in period t is denoted y_{it} , $i = 1, \dots, K$. Denote the pre-treatment outcomes by $y_i^{pre} := (y_{1i} \ \dots \ y_{T_0 i})'$ and collect the pre-treatment outcomes of the controls in the $T_0 \times J$ matrix $Y_0^{pre} := (y_2^{pre} \ \dots \ y_{J+1}^{pre})$. For the treated units, define the $T_0 \times L$ matrix Y_1^{pre} accordingly and let the pre-treatment outcome of unit 1, y_1^{pre} , be the first column of Y_1^{pre} .

In post-treatment time, I assume that for each unit i and period $t > T_0$ the *potential* outcome in the case of non-treatment y_{it}^N is a linear function of an R -dimensional random vector z_i of observable predictors and of an F -dimensional random vector λ_t of unobservable shocks, $F \geq T_1$. Both z_i and λ_t are unaffected by treatment. In equation (1), their deterministic coefficient vectors $\theta_t \in \mathbb{R}^R$ and $\mu_t \in \mathbb{R}^F$ are, respectively, time- or unit-specific:

$$y_{it}^N = z_i' \theta_t + \lambda_t' \mu_i \quad \forall i = 1, \dots, K \quad \forall t > T_0 \quad (1)$$

Collect all covariates of the control units in the $R \times J$ matrix $Z_0 := (z_2 \ \dots \ z_{J+1})$ and define the $R \times L$ matrix Z_1 for the predictors of the treated units accordingly. I assume

A1: Properties of predictors

- a) Non-redundancy: $E(z_i z_i')$ is nonsingular $\forall i$.
- b) Completeness: $E(\lambda_t | Z_0, Z_1) = 0 \quad \forall t > T_0$.

In order to allow for a common growth component in the y_{it}^N 's, assume that the predictors contain a constant term $z_{1i} = 1 \quad \forall i$.

Note that I have defined (1) for $t > T_0$ only. Extending (1) to hold for $1 \leq t \leq T_0$ would be problematic: Since predictors z_i are constant over time, (1) would imply that z_i determines all

pre-treatment outcomes, even the outcomes of the initial period 1. Hence, z_i would need to be known already in period 1. This would disallow pre-treatment outcomes later than period 1 as predictors and, more generally, would greatly limit the predictive power of z_i for the treatment period if the time span T_0 prior to treatment is substantial. If, by contrast, (1) is defined for $t > T_0$ only, predictors may involve (functions of) pre-treatment outcomes y_{si} , $s \leq T_0$.

Potential outcomes in the case of treatment y_{ii}^{Tr} differ from y_{ii}^N by a treatment effect α_{ii} which is unit- and time-specific:

$$y_{ii}^{Tr} = \alpha_{ii} + y_{ii}^N, \quad \forall i \quad \forall t > T_0 \quad (2)$$

The treatment status of unit i is given by a binary variable d_{ii} such that

$$y_{ii} = d_{ii} y_{ii}^{Tr} + (1 - d_{ii}) y_{ii}^N, \quad d_{ii} \in \{0, 1\} \quad \forall i \quad (3)$$

Finally, some matrix notation: Collect all factor loadings of the control units in the deterministic $F \times J$ matrix $M_0 := (\mu_2 \dots \mu_{J+1})$. For the treated units, define the $F \times L$ matrix M_1 accordingly. In post-treatment time, let $\Theta^{post} := (\theta_{T_0+1} \dots \theta_T)'$ denote predictor coefficients and $\Lambda^{post} := (\lambda_{T_0+1} \dots \lambda_T)'$ the unobserved shocks. The product $H_0^{post} := \Lambda^{post} M_0$ is a $T_1 \times J$ random matrix with typical element $\eta_{it} := \mu_i' \lambda_t$. Define $H_1^{post} := \Lambda^{post} M_1$ accordingly. By A1, $E(H_0^{post} | Z_0, Z_1) = 0$ and $E(H_1^{post} | Z_0, Z_1) = 0$.

Collect outcomes in $y_i^{post} := (y_{(T_0+1)i} \dots y_{Ti})' \quad \forall i$, and $Y_0^{post} := (y_2^{post} \dots y_{J+1}^{post})$ for all controls. For the treated units, define the $T_1 \times L$ matrix Y_1^{post} accordingly and let y_1^{post} be the first column of Y_1^{post} . The treatment effects α_{ii} are organized in the $T_1 \times L$ matrix A with typical column α_i , $i = 1, \dots, L$, and typical row α_t' , $t > T_0$.

Since $y_{ii} = y_{ii}^N$ for all control units, we have

$$Y_0^{post} = \Theta^{post} Z_0 + \Lambda^{post} M_0 =: \Theta^{post} Z_0 + H_0^{post} \quad (4)$$

while for the treated unit 1, (1), (2) and (3) yield

$$y_1^{post} = \alpha_1 + \Theta^{post} z_1 + \Lambda^{post} \mu_1 =: \alpha_1 + \Theta^{post} z_1 + \eta_1^{post} \quad (5)$$

More generally, the outcomes of all L treated units are given by

$$Y_1^{post} = \Theta^{post} Z_1 + \Lambda^{post} M_1 + A \quad (6)$$

which, setting $H_1^{A,post} := H_1^{post} + A$, yields the treatment analogue of the control equation (4)

$$Y_1^{post} = \Theta^{post} Z_1 + H_1^{A,post}. \quad (7)$$

Treatment effects A are unobserved and may be correlated with predictors Z_1 or post-treatment shocks Λ^{post} . To account for this, suppose

$$A = B_A Z_1 + \Lambda^{post} C_A + E_A \quad (8)$$

Here, B_A and C_A are deterministic coefficient matrices of dimensions $T_1 \times R$ and $T_1 \times F$, respectively, while E_A is a $T_1 \times L$ matrix of zero-mean shocks uncorrelated with both Z_1 and Λ^{post} . Note that E_A is allowed to display correlation across units and time.

2.2 Identification

The right-hand side of (1) involves three unobservables, θ_t , λ_t and μ_t . In general, many observationally equivalent choices $\tilde{\theta}_t$, $\tilde{\lambda}_t$ and $\tilde{\mu}_t$ exist since

$$y_{it}^N = z_i' \theta_t + \mu_i' \lambda_t = z_i' (\theta_t - G_1 \lambda_t) + (z_i' G_1 G_2 + \mu_i' G_2) G_2^{-1} \lambda_t =: z_i' \tilde{\theta}_t + \tilde{\mu}_i' \tilde{\lambda}_t$$

for suitably chosen nonzero matrices G_1 and G_2 . But assumption A1 and the Law of Iterated

Expectations imply $E(z_i \eta_{it}) = E(E(z_i \mu_i' \lambda_t | Z_1, Z_0)) = E(z_i \mu_i' E(\lambda_t | Z_1, Z_0)) = 0_R$. Since

$E(z_i z_i')$ is non-singular we can premultiply (1) by z_i to get $z_i y_{it}^N = z_i z_i' \theta_t + z_i \eta_{it}$, take

expectations and solve for $\theta_t = E(z_i z_i')^{-1} E(y_{it}^N z_i)$. Hence, θ_t and η_{it} are uniquely identified

under A1 and the OLS estimate of θ_t is consistent with respect to the cross section dimension.

A1 does not separately identify μ_i and λ_i , though. But any identification will do, since we can be agnostic about the “true” shocks. The interest is in the treatment effects, not in “true” shocks. I impose the familiar principal components assumption

A2: Orthonormal shocks

In $H^{post} = \Lambda^{post} M$, the shocks in Λ^{post} are orthogonal, satisfying $E(\Lambda^{post} \Lambda^{post'}) = I_{T_1}$.

Moreover, MM' is a $T_1 \times T_1$ diagonal matrix.

This identification is unique up to sign changes in columns and rows. See the discussion paper version for additional technical conditions to achieve unique identification.

3. The Standard SC-Approach

The standard synthetic control approach has been popularized by Abadie and Gardeazabal (2003) and Abadie, Diamond and Hainmueller (ADH) (2010). The key idea is that the counterfactual y_{it}^N , $t > T_0$, can be approximated by a weighted average of the observed contemporaneous outcomes of the control units. Formally, if $y_1^{N,post} := (y_{(T_0+1)1}^N \cdots y_{T1}^N)'$, the ADH approach aims at finding $w^* \in \Delta_J := \{w \in \mathbb{R}^J \mid \iota_J' w = 1 \wedge w_i \geq 0 \ \forall i = 1, \dots, J\}$ such that

$$y_1^{N,post} \approx Y_0^{post} w^* \quad (9)$$

Here, ι_J is a $J \times 1$ vector of ones.

To find w^* , ADH's approach relies on y_1^{pre} , Y_0^{pre} and on the predictors z_1, Z_0 , where the latter may also include functions of some or all of the pre-treatment outcomes. ADH assume that (1) also holds for $t = 1, \dots, T_0$. Let us ignore that this may disallow endogenous variables (e. g. y_i^{pre} 's) as predictors. Then, using the notation $\Theta^{pre}, \Lambda^{pre}, H_0^{pre}$ for the pre-treatment analogues of $\Theta^{post}, \Lambda^{post}, H_0^{post}$ it follows that for any vector of weights $w \in \Delta_J$ we have

$$y_1^{pre} - Y_0^{pre} w = \Theta^{pre} (z_1 - Z_0 w) + \Lambda^{pre} (\mu_1 - M_0 w) \quad (10)$$

Since not all predictors may be equally informative for potential outputs, let $v \in \Delta_R$ be a vector of predictor weights and let $V := \text{diag}(v)$ be the corresponding diagonal $R \times R$ matrix.

ADH propose to solve the following optimistic bilevel minimization problem:

$$\min_{v \in \Delta_R, w \in \Delta_J} L_{out}(v, w) := \frac{1}{T_0} (y_1^{pre} - Y_0^{pre} w)' (y_1^{pre} - Y_0^{pre} w) \quad (11)$$

s. t.

$$w \in \Psi(v) := \underset{w \in \Delta_J}{\text{argmin}} L_{in}(v, w) := (z_1 - Z_0 w)' V (z_1 - Z_0 w), \quad V = \text{diag}(v)$$

This formulation is due to Malo et al. (2020).

Let $\Phi_1^Z := \{w \in \Delta_J \mid z_1 = Z_0 w\}$ and $\Phi_1^{Y^{pre}} := \{w \in \Delta_J \mid y_1^{pre} = Y_0^{pre} w\}$. ADH assume that the intersection $\Phi_1^Z \cap \Phi_1^{Y^{pre}}$ is nonzero. Under this (and some additional) assumptions ADH show that $Y_0^{post} w^0$ is an asymptotically unbiased estimator of $y_1^{N, post}$ when the number of pre-treatment observations T_0 approaches infinity.³

Unfortunately, the ADH approach is problematic in multiple regards: First, if λ_t contains an idiosyncratic shock which affects the treated unit, then, asymptotically, $\Phi_1^{Y^{pre}} = \emptyset$ with probability 1 and the ADH estimator may not be asymptotically unbiased, cf. Ferman and Pinto (2016). Second, V is not identified at the optimum, since Φ_1^Z is non-empty by assumption and, therefore, the optimal value of L_{in} is zero for any matrix V . Third, suboptimal weights w are chosen if all pre-treatment outcomes are included in the matrix of predictors: Then, a solution to the bilevel problem (11) is given by any $w \in \Phi_1^{Y^{pre}}$ along with $v = T_0^{-1} (\iota_{T_0}' \quad 0_R)'$, i. e. the choice of w is solely driven by the pre-treatment outcomes and all other predictors have no

³ Some papers (e. g. Malo et al. (2020)) state that ADH prove the „consistency“ of the SC-estimator. This is not true. ADH's proof shows asymptotic unbiasedness. ADH do not claim that the variance of $y_1^{N, post} - Y_0^{post} w^*$ converges to zero.

impact at all. This result is due to Kaul et al. (2015). See Lucke (2022) for further problems which arise in model (1) if potential outcomes are autocorrelated.

4. An alternative synthetic control approach

The overarching aim of synthetic control analysis is a good estimate of the counterfactual $y_1^{N,post}$. For this, let us focus on the synthetic control error $\varepsilon_1^{post}(w) := y_1^{N,post} - Y_0^{post}w$, $w \in \Delta_J$, i. e. the difference between the counterfactual of the treated unit and its synthetic control over the treatment period. I will estimate w by minimizing its mean squared prediction error

$$MSPE_1^{post}(w) := \frac{1}{T_1} E\left(\varepsilon_1^{post}(w)' \varepsilon_1^{post}(w) \middle| I_0\right)$$

where $I_0 := \{z_1, Z_0\}$ is the relevant information set. Optimal weights would be the solution of

$$\min_{w \in \Delta_J} MSPE_1^{post}(w) \tag{12}$$

Solving (12) has not yet been pursued in the literature presumably because $\varepsilon_1^{post}(w)$ depends on the unknown counterfactual $y_1^{N,post}$. But closer analysis reveals that the only unobservables required for the solution of (12) are $(\Theta^{post}, \mu_1, M_0)$. Estimating Θ^{post} consistently is straightforward and (μ_1, M_0) can be derived from decomposing specific regression residuals which asymptotically coincide with their true population counterparts. Moreover, problem (12) is a standard constrained quadratic minimization problem, which, numerically, is much easier to solve than the optimistic bilevel problem (11).

While (12) is quite different from the bilevel problem (11), under the assumptions imposed by ADH, a solution to (11) is also a solution to (12). This is stated in

Proposition 1:

Suppose that for all periods $t = 1, \dots, T$ the data are generated by (1), (2) and (3). If assumptions A1, A2 and the ADH-assumption $\Phi_1^Z \cap \Phi_1^{Y^{pre}} \neq \emptyset$ hold, then any $w^+ \in \Delta_J$ which solves the bilevel problem (11) is also a solution to $\min_{w \in \Delta_J} MSPE_1^{post}(w)$.

Proof: See Appendix.

Regrettably, the non-emptiness assumption $\Phi_1^Z \cap \Phi_1^{Y^{pre}} \neq \emptyset$ is rarely met in practical applications. It appears reasonable to forego this assumption and instead focus on directly minimizing (12). This would have the additional advantage that (1) need not hold in pre-treatment time so that a much larger set of predictors (including pre-treatment outcomes and other endogenous variables) may be used.

Proposition 2 establishes the existence of a constrained quadratic minimization problem equivalent to (12).

Proposition 2:

Suppose the data are generated by (1), (2) and (3). Suppose further that assumptions A1 and A2 hold. Then, solving $\min_{w \in \Delta_J} MSPE_1^{post}(w)$ is equivalent to solving

$$\min_{w \in \Delta_J} w' E \left(Z_0' \Theta^{post} \Theta^{post} Z_0 + M_0' M_0 \right) w - 2E \left(z_1' \Theta^{post} \Theta^{post} Z_0 + \mu_1' M_0 \right) w \quad (13)$$

Proof: See Appendix.

Problem (13) depends on the unknown parameters $(\Theta^{post}, \mu_1, M_0)$. Proposition 3 shows how consistent estimates of Θ^{post} and B_A can be found. Knowledge of B_A will allow a bias correction for those estimators of Θ^{post} which also incorporate observations from treated units. The consistency of such bias-corrected estimators is established in Proposition 4. (μ_1, M_0) can be computed from the residuals obtained from this group of bias-corrected, consistent estimators of Θ^{post} .

Proposition 3:

Suppose the data are generated by the factor model (1), (2) and (3). Assume that assumption A1 holds and that $M_{Z_0 Z_0} := p \lim_{J \rightarrow \infty} J^{-1} Z_0 Z_0'$ and $M_{Z_1 Z_1} := p \lim_{L \rightarrow \infty} L^{-1} Z_1 Z_1'$ are both finite and invertible. Moreover, assume that treatment effects conform with (8). Let $\hat{\Theta}_0^{post} := Y_0^{post} Z_0' (Z_0 Z_0')^{-1}$ and $\tilde{\Theta}_1^{post} := Y_1^{post} Z_1' (Z_1 Z_1')^{-1}$ be OLS-estimators of Θ^{post} in (4) and (7), respectively. Then

a) the estimator $\hat{\Theta}_0^{post}$ is J -consistent for Θ^{post} , i. e. $p \lim_{J \rightarrow \infty} \hat{\Theta}_0^{post} = \Theta^{post}$.

b) the estimator $\tilde{\Theta}_1^{post} - \hat{\Theta}_0^{post}$ is J, L -consistent for B_A , i. e. $p \lim_{\substack{J \rightarrow \infty \\ L \rightarrow \infty}} (\tilde{\Theta}_1^{post} - \hat{\Theta}_0^{post}) = B_A$

Proof: See Appendix.

Part b) of Proposition 3 says that the difference between the inconsistent estimator $\tilde{\Theta}_1^{post}$ and the consistent estimator $\hat{\Theta}_0^{post}$ is consistent for B_A , the matrix which accounts for correlation between the treatment effects A and the predictors Z_1 . Note that $\tilde{\Theta}_1^{post} - \hat{\Theta}_0^{post}$ is equal to the OLS estimator $\hat{B}_A := \hat{H}_1^{A, post} Z_1' (Z_1 Z_1')^{-1}$, where $\hat{H}_1^{A, post} := Y_1^{post} - \hat{\Theta}_0^{post} Z_1$.

Let us now study just one treated unit, i. e. unit 1, along with all J control units. For this purpose, let $\hat{A}_Z := \hat{B}_A Z_1$ and let \hat{a}_1 be the first column of \hat{A}_Z . Augment Y_0^{post} by y_1^{post} corrected for the estimated treatment effect, i. e. $Y_{10}^{post} := \begin{pmatrix} (y_1^{post} - \hat{a}_1) & Y_0^{post} \end{pmatrix}$.

Similarly, let $Z_{10} := \begin{pmatrix} z_1 & Z_0 \end{pmatrix}$ and $M_{10} := \begin{pmatrix} \mu_1 & M_0 \end{pmatrix}$. We can now jointly estimate equations (4) for the controls and (5) for the treated unit. Alternatively, using $Z := \begin{pmatrix} Z_1 & Z_0 \end{pmatrix}$, $M := \begin{pmatrix} M_1 & M_0 \end{pmatrix}$ and $Y^{post} := \begin{pmatrix} (Y_1^{post} - \hat{A}_z) & Y_0^{post} \end{pmatrix}$ we may jointly estimate (4) and (6), i. e. use all treated units simultaneously. A third approach would consist in estimating Θ^{post} just from

(6), i. e. from all treated units. All three approaches yield consistent estimates of Θ^{post} as is shown in

Proposition 4:

Suppose the data are generated by the factor model (1), (2) and (3). Assume that assumption A1 holds and that $M_{Z_0 Z_0} := p \lim_{J \rightarrow \infty} J^{-1} Z_0 Z_0'$ and $M_{Z_1 Z_1} := p \lim_{L \rightarrow \infty} L^{-1} Z_1 Z_1'$ are both finite and invertible. Moreover, assume that treatment effects are linear as in (8). Then

a) the estimator $\hat{\Theta}_{10}^{post} := Y_{10}^{post} Z_{10}' (Z_{10} Z_{10}')^{-1}$ is J -consistent for Θ^{post} ,

b) the estimator $\hat{\Theta}^{post} := Y^{post} Z' (Z Z')^{-1}$ is J, L -consistent for Θ^{post} ,

c) the estimator $\hat{\Theta}_1^{post} := (Y_1^{post} - \hat{A}_Z) Z_1' (Z_1 Z_1')^{-1}$ is J, L -consistent for Θ^{post} ,

i. e. $p \lim_{J \rightarrow \infty} \hat{\Theta}_{10}^{post} = p \lim_{\substack{J \rightarrow \infty \\ L \rightarrow \infty}} \hat{\Theta}^{post} = p \lim_{\substack{J \rightarrow \infty \\ L \rightarrow \infty}} \hat{\Theta}_1^{post} = \Theta^{post}$.

Proof: See Appendix.

Note that models (4) and (6) involve the same parameters Θ^{post} . Proposition 4 can be used to test the validity of this restriction by a simple F -Test.

While all estimators in Proposition 4 have the same probability limit, their finite sample properties may be different. There is no clear guidance which estimator should be preferred in synthetic control analysis, except for the fact that $\hat{\Theta}_{10}^{post}$ requires slightly weaker assumptions. Therefore, I will use $\hat{\Theta}_{10}^{post}$ in the following.

Estimating the control equations (4) jointly with at least one equation in (6) is convenient because under identification A2 the residuals, e. g. $\hat{H}_{10}^{post} = \hat{\Lambda}^{post} \begin{pmatrix} \hat{\mu}_1 & \hat{M}_0 \end{pmatrix}$, can be decomposed to yield unique estimates of μ_1 and M_0 . These estimates allow us to solve (13). Note that \hat{H}_{10}^{post} approaches the true matrix product $\Lambda^{post} M_{10}$ asymptotically and that its decomposition is unique under A2.

The optimal SCM-weights w^* which solve (13) are used to construct the synthetic control

$$\hat{y}_1^{N,post} := Y_0^{post} w^* \text{ and the estimated causal effects of treatment } \hat{\alpha}_1 := y_1^{Tr,post} - \hat{y}_1^{N,post}.$$

5. Data

To assess growth effects of EMU, I use regional data on per-capita GDP growth from the Annual Regional Database of the European Commission (ARDECO). The time span is 1980-2018⁴, the level of disaggregation is NUTS3, the lowest level available. Regions from formerly socialist economies are excluded, since no data are available prior to 1990 (and often later) and because growth in these regions is strongly driven by post-socialist catch-up growth – which would very likely be a major cofounder in the control group.

Hence, the analysis is confined to the former EU-15, i. e. regions from Austria, Belgium, West Germany, Denmark, Greece, Spain, Finland, France, Ireland, Italy, Luxembourg, Netherlands, Portugal, Sweden and the United Kingdom (UK). Of these countries, Denmark, Sweden and the UK kept their national currencies, while the others formed the Eurozone. The Euro became effective in 1999. Thus “treatment” starts in 1999 and extends through all subsequent years.

Denmark did not join the Euro, but it joined its predecessor, the Exchange Rate Mechanism II, in which it maintained a fixed exchange rate to the Euro over the treatment period. Therefore, to construct a synthetic control, I exclusively utilize regions from those two EU-15 countries that maintained and actively experienced exchange rate flexibility from 1999 to 2018, i. e. Sweden and the United Kingdom. These two countries consist of exactly 200 regions at NUTS3 level⁵. This is the control set (or donor pool). Since some of the British regions underwent a territorial redefinition in the year 2000 which left major traces in their

⁴ 2018 is the end date since no data are available for the United Kingdom after this date.

⁵ I delete region UKN15 (Mid and East Antrim) from the donor pool because this region was redefined with sizable increases in territory in 2015. Hence the growth rate of GDP in 2015 is artificially great (22%) which disqualifies UKN15 as a control region in treatment time.

recorded GDP growth rate of this year, I discard the first two treatment years 1999 and 2000 from the analysis of treatment effects. All results below refer to the eighteen post-treatment years 2001-2018⁶.

While comparing countries like Italy and Greece to Sweden or the UK on a national level may seem ambitious, the regional level offers higher chances of finding similar economic structures. Some regions, whether in northern or southern countries, share common traits such as a focus on agriculture, fisheries, traditional industries, tourism, or other services, as well as similar strengths and weaknesses in infrastructure. If the prevailing perspective is that Sweden and the UK are more competitive than Italy and Greece, it implies that a higher proportion of regions in the former countries house highly competitive companies compared to the latter. However, this does not imply that Sweden and the UK lack regions with traditional, underdeveloped, or rusty characteristics, akin to many regions in Italy and Greece. For the construction of synthetic controls, it suffices to have some overlap, i. e. just some regions in the control group which are similar to many regions in the treatment group.

ARDECO provides data on nominal GDP per capita on NUTS3 level. I divide by the national consumer price index and take log differences to have the growth rate of real GDP as outcome variable. Moreover, I use the following set of 23 covariates (plus a constant):

The first three predictors are simple statistics of pre-treatment outcomes: The log of regional real GDP in 1998, the log of the 1980-1998 average level of real GDP, and the 1980-1998 average growth rate of real GDP. To account for sectoral activities I use the 1998 shares of gross value added (GVA) for Agriculture, Industry, Construction, Trade, and Financial and Business Services. For the same sectors and the residual sector, I use the 1998 shares of labor costs in sectoral GVA as predictors. All these data are from ARDECO.

⁶ Thus, Greece is a regular Eurozone country for the full treatment period.

A third party (Cambridge Econometrics) was charged by the European Commission to compile sectoral capital stock data at NUTS2 level. This data has not yet been made available through ARDECO, but Cambridge Econometrics kindly supplied it to me. The sectoral disaggregation of capital stocks is the same as for GVA at NUTS3 level, so I broke down the sectoral NUTS2 capital stock data proportional to the GVA shares in order to construct measures of *total* capital stocks at NUTS3 level. The logs of their 1998 values are used as predictors.

Data from the EU census 2011 provide information on dwellings, their age and the number of flats per dwelling on NUTS3 and NUTS2 level. Dwellings are indicative of infrastructure, as inhabitants need roads, gas, electricity, telecommunication and so on. I employ the logarithm of dwellings built before 2001 as a proxy for infrastructure, utilizing both the log average age of the dwellings and the ratio of dwellings built between 1991 and 2000 over all dwellings as distinct indicators of infrastructure modernity. Additionally, I incorporate the ratio of dwellings with three or more flats over all dwellings as an indicator of the degree of urbanization in the NUTS3 region.

Finally, ARDECO also supplies some simple statistics like the population in 1998, the area in square kilometers and dummy variables indicating whether a NUTS3 region is a mountain region or a region bordering on some other country. These variables (population and area in logs) are also included in the set of predictors.

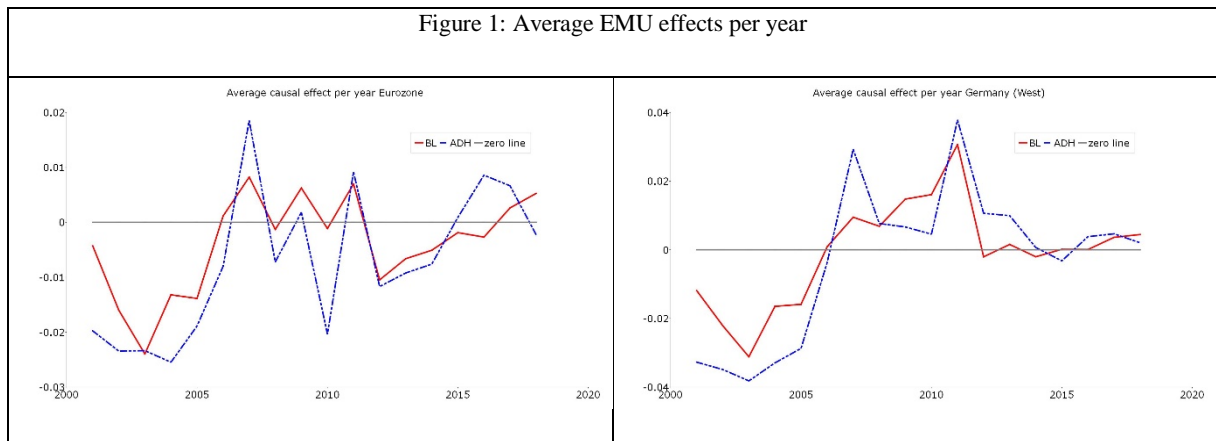
6. Synthetic Controls for EMU regions

This section describes and interprets the results obtained from estimating synthetic controls for EMU regions 2001-2018. I use two competing methods to estimate the synthetic controls: The standard ADH method in its “canonical” form (e. g. Chernozhukov (2021)) where all pre-treatment outcomes are used as predictors (and all other predictors are therefore redundant, cf. Kaul et al. (2015)). This approach minimizes the mean squared prediction error

for *pre*-treatment outcomes, cf. ADH and Abadie (2021). Alternatively, I use the approach outlined in Section IV which constructs the counterfactuals from the factor model (1) estimated on data in *post*-treatment time. In both approaches I use just one treated unit at a time along with all controls. I denote the results produced by both approaches by ADH and BL, respectively.

A necessary condition for the validity of the BL approach is the equality of the Θ^{post} matrices in the control and in the treatment group. Suppose the true parameters in the control group are Θ_0^{post} and the true parameters in the treatment group are Θ_1^{post} . Using Propositions 3 and 4c) we derive consistent estimates $\hat{\Theta}_0^{post}$ and $\hat{\Theta}_1^{post}$ so that we can test $H_0 : \Theta_0^{post} = \Theta_1^{post}$ with a simple F -test. The resulting F -statistic is 0.577 with associated P-value of 1.000, i. e. the null is easily accepted.

Let us now have a look at the average growth effect (across all Eurozone regions) over time. The left panel of Figure 1 displays for each year from 2001 to 2018 the cross section average of all the estimated causal effects of Eurozone membership.



BL estimates seem to have less variance than ADH estimates. Apart from this, ADH and BL estimates broadly suggest the same conclusion: There were, on average, sizable losses in GDP growth in the years prior to the financial crisis. During and immediately after the financial crisis, there may have been a weakly positive effect (BL) or no clear effect at all

(ADH) from the existence of a common currency. This came to an end and turned deeply into negative territory when the Eurozone's sovereign debt crisis unfolded in 2011.

Interestingly, much of the first two developments may have been driven by Germany, the EU's largest economy, cf. Figure 1, right panel. Here we find strongly negative EMU effects prior to the financial crisis and strongly positive effects between 2007 and 2011 – with not much of an effect after this date.

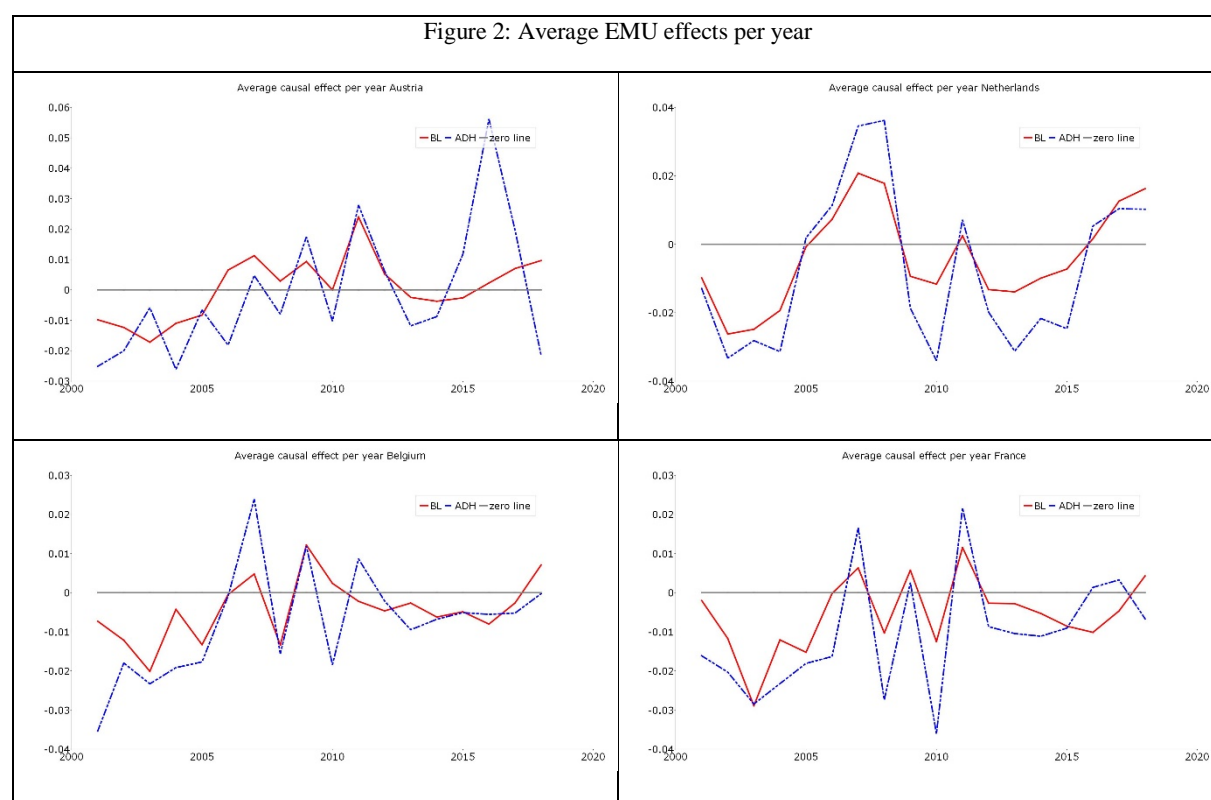
To understand this development, it is useful to go back to the late 1990s, when the European Monetary System (EMS) was headed for permanently fixed exchange rates and a common currency in 1999. In these pre-Euro years and in the early 2000s, Germany suffered from exceptionally low GDP growth. Sinn (2003, 2014) even dubbed Germany “the laggard of Europe” and argued that the (anticipated) introduction of the Euro was at the heart of the phenomenon:

In his view, prior to EMU, the EU's southern periphery was poor in capital due to inflation and depreciation risk. Hence, the marginal product of capital was high. When EMU approached, inflation and depreciation risks were reduced. Attracted by high yields, investors shifted capital from Germany and other northern EU countries to the southern periphery, competing down its interest rates to almost the German level. As a result, German companies lost their competitive edge in terms of access to cheaper credit and also suffered from weaker aggregate demand since capital and investors were moving south. Consequently, German growth was weaker than it had been under the German Mark. The causal effects estimated for Germany in the early 2000s are well in line with this reasoning, cf. the right panel of Figure 1.

But, second, the tide turned with the financial crisis. Investors worried about financial institutions and debt sustainability in the south, moving back capital to Europe's core economies. Germany, in particular, was seen as a safe haven so that interest spreads vis-à-vis Germany greatly increased in the runup to the 2010-2013 sovereign debt crisis. As a consequence, Germany recovered quickly from the financial crisis, whereas the Eurozone's

southern countries got deeper into trouble. Thus, European capital market integration propelled by EMU was quite beneficial for Germany in the crisis years. Between 2007 and 2011 the positive effects for German regions may actually have outweighed the detrimental effects for southern European regions so that the Eurozone average is estimated to have had a positive growth effect from the Euro.

Similar effects (negative before and positive during the financial crisis) can be seen for Austrian and Dutch regions, cf. the upper panels of Figure 2. By contrast, core EU countries Belgium and France (lower panel of Figure 2) did not see the negative effects prior to the financial crisis being offset during the crisis. It seems that capital flows have shied away from these countries more than from Germany, Austria and the Netherlands.

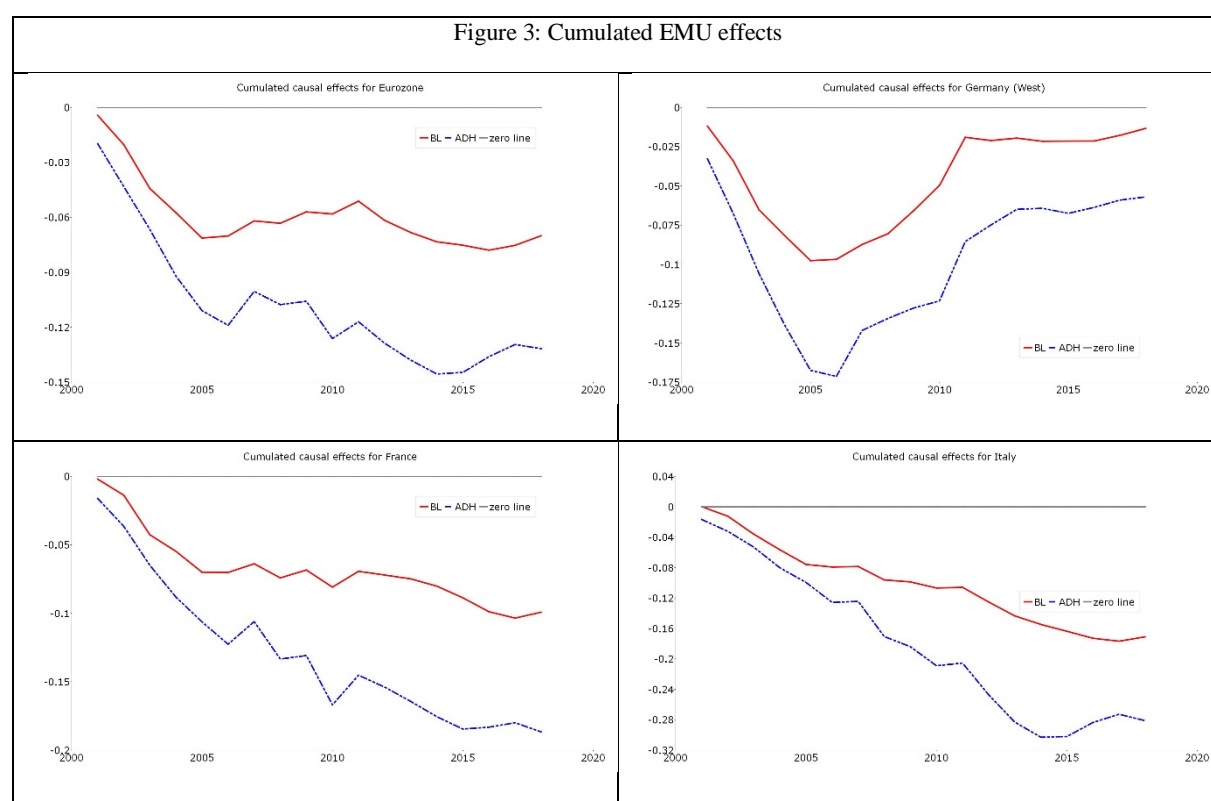


In the years from 2012 onward, the negative causal effects in southern Europe became so strong that even the aggregate Eurozone effects turned into negative territory (with a slow recovery over the next several years). These negative aggregate effects are clearly not driven by Germany, which, cf. Figure 1, experienced neither great harm nor great benefit from the

common currency in these years. Hence, focusing on the total Eurozone average is of limited value, since growth in European regions was quite differently affected by the common currency.

In Figure 3, we present the *cumulative* growth effect in the Eurozone and in the three largest Eurozone economies: Germany, France and Italy. Depicted is the total estimated gain or loss in per capita income at each point in time between 2001 and 2018.


According to these results, both ADH and BL estimates suggest that the net impact of EMU was negative for Eurozone growth. Both methods suggest that greatest losses in GDP growth occurred in the initial years of EMU (approximately until 2006) – perhaps some sort of phase-in costs of EMU. While ADH estimates suggest further losses up to 2015, BL estimates seem to stabilize from 2007 onward at a cumulated loss of 7% of GDP. The same may be true for ADH estimates at a later date and with a cumulative loss in growth almost twice as large. Neither method suggests that the Eurozone would have made good on forgone growth in years after 2018. Of course, Covid is a great obstacle for any extended analysis.



Generally, BL estimates are more cautious than their ADH counterparts. This can also be seen in the country-specific cumulated effects for Germany, France and Italy. For these countries, the Eurozone's heterogeneity is eye-catching. For Germany, the initially accumulated loss of nearly 9% of GDP in 2005 is nearly completely recovered in the subsequent years (upper right panel of Figure 3). For France, by contrast, both methods find that cumulated losses keep increasing even after the initial phase-in of EMU (lower left panel of Figure 3) – unlike the stabilization possibly under way for the Eurozone as a whole. For Italy (lower right panel of Figure 3), there is weak evidence of stabilization after 2015 or later, but only after a long and very negative development which peaked at a cumulative loss in GDP of roughly 15% (BL) or even 30% (ADH). As a rule of thumb, it seems fair to say that Italy lost about 1% of GDP growth for each year of EMU membership relative to a counterfactual with a national currency.

While Figures 1-3 give an impression of average treatment effects (ATTs) across regions for different points in time, we will now turn to ATTs across time for different regions. Clearly, it is suggestive to look at all regions in a particular country.

An interesting case is Belgium, cf. Figure 4. Here, the first region on the x-axis is Brussels, the next 22 regions are located in Flanders, and regions 24-44 are part of mostly French-speaking Wallonia. It is well known that companies in Wallonia are, on average, less competitive than companies in Flanders. For instance, OECD (2020) presents evidence that regional productivity in Wallonia is quite a bit lower than in Flanders and that Flanders was quicker and more successful in transitioning from manufacturing to services. Both ADH and BL estimates of the causal effects of EMU suggest that, by and large, the less competitive Wallonian regions have suffered more from EMU than the Flemish regions. This is quite plausible, as under EMU, issues related to competitiveness cannot be addressed through a depreciation of the national currency.



Note that SCs are estimated separately for each treated region. No information is supplied on the geographic location of a treated region or on the country to which it belongs or which language is spoken. The computer does not know which region is Flemish or Wallonian, French, Italian or German. Yet the results are very much in line with the view that Wallonian companies were in need of a currency devaluation to improve their competitiveness and that this is less so for Flemish companies. They also imply that, under prevailing exchange rates, average German companies did not face a competitiveness issue, while the French and Italian economies might have significantly benefited from a devaluation.

To illustrate this, the graphs in Figure 5 present the regions of each country sorted in a way that the estimated causal effects of EMU decrease from left to right. For instance, for Germany (upper left panel of Figure 5) this indicates that roughly one-third of German regions have benefited from EMU (slightly less according to ADH, slightly more according to BL). For almost all regions, the size of the average effects is in the range between $\pm 2.5\%$ of EMU-induced growth. In fact, 94% of the BL-estimates are in the $\pm 1\%$ interval and 76% in the $\pm 0.5\%$ interval. As these represent average values per year of EMU membership, even a 0.5% difference in growth rates accumulates to nearly 10% of GDP over 18 years. Hence, many of these effects exert a significant impact on economic well-being. However, many ATTs are relatively small, and although quite a few exhibit strong negative values, there appear to be just as many positive ones of approximately the same magnitude. Hence, in balance the German

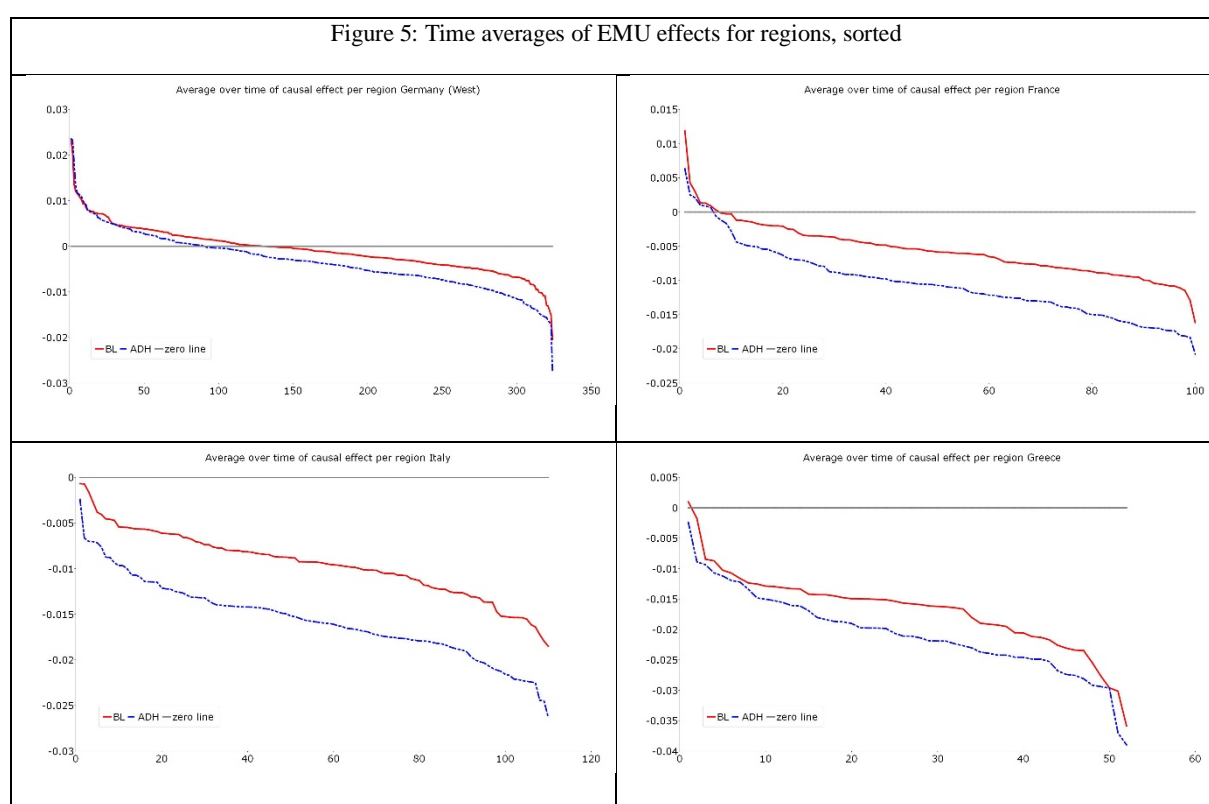
results suggest that Germany was only mildly negatively affected by EMU – a result in line with Figure 3, upper right panel.

Interestingly, Germany's top positive EMU-effects are associated with the German car industry. According to both methods, the highest positive effect is found for Ingolstadt, the headquarter and main production site of luxury brand manufacturer Audi AG. The second largest effect is found for Wolfsburg, the headquarter and main production site of Volkswagen AG. Still according to both methods, the third largest effect is found for Coburg, home to the large automotive supplier Brose Fahrzeugteile SE & Co KG. The region Dingolfing-Landau ranks fourth in ADH and eighth in BL, it hosts the main production site of the BMW Group.

Since German car manufacturers are strongly export-oriented, these results may be due to an undervalued Euro (compared to a counterfactual German mark) and, hence, favorable export conditions for German exporters in general. In fact, further analysis shows that it is not just the car industry which drives the most positive EMU-results in Germany. The top ten German regions also host headquarters or major production sites of companies active in mechanical engineering (e. g. Kaeser Kompressoren SE) or health technology (e. g. Siemens Healthineers AG), typically with strong positions on world markets. In line with this interpretation, observe that German net exports as a share of GDP have risen to record levels of between six and eight percent since the introduction of the Euro (up from about two percent prior to 1999).

A completely different picture emerges for the 100 French regions, cf. Figure 5, upper right panel. In continental France, both ADH and BL identify just four regions which have benefited from EMU – three more are overseas territories. Ignoring the latter, the positive effects in European France are limited to just three regions in the Greater Paris area plus (very marginally positive) the Département Hautes Garonne, where the French-German aerospace giant Airbus SE is located at Toulouse (production) and Blagnac (headquarter).

In stark contrast to the German results, both methods find that regions with major production sites of French carmakers have suffered from EMU. This is true for various regions in which Renault SA and Groupe PSA SA⁷, maker of well-known brands like Peugeot, Citroen and Opel, maintain some of their larger plants. It is also true for tyre producer Michelin SCA, e. g. at its main site in Clermont-Ferrand. Moving from cars to trains, French giant Alstom SA had one of its two major plants, now owned by General Electric, in the Territoire de Belfort, the region with the worst EMU effect in France (BL: -1.6%, ADH: -2.1%).




In general, there is no indication that French companies benefited from EMU by successfully competing on export markets. This may be just the flip side of the opposite conclusion suggested by the German results. Quite possibly, the Euro was undervalued from a German perspective while it was overvalued from a French (and many other Eurozone countries') perspective. The historical fact that, prior to EMU, the German Mark had typically

⁷ Today part of Stellantis N.V.

appreciated against the French Franc and the Italian Lira over long horizons, is quite in line with this reasoning.

Turning to Italy, there is not a single region for which either method comes up with a positive EMU-estimate. Rather, most regions seem to have suffered a loss in economic growth of at least half a percentage point per year and for many it may have been substantially more: Even the more cautious BL-estimates are below -1% per year for almost half of the regions and ADH estimates are lower than -1.5% for more than half of the regions. In other words: Results suggest that it would have been a Pareto-improvement if Italy had retained the Lira rather than join the Eurozone.

The uniformly negative EMU-experience of Italian regions is surprising because Italy is home to a large number of small and medium sized enterprises (SMEs) with good world market access and to some well-established large companies particularly in mechanical engineering, automotives (e. g. Fiat SpA), Textiles and fashion (e. g. Benneton Group, Gucci, Versace SpA) and foodstuffs (e. g. Barilla, Ferrero SpA). Much of this industrial base is concentrated in the north of Italy where it produces about 70% of Italian merchandise exports, far more than companies in Italy's middle, south or on its large islands, cf. Italian Trade Agency (2017). One might, therefore, expect that the effects of EMU are distributed unevenly across Italy, possibly resulting in less unfavorable outcomes in the northern regions. But as the maps in Figure 6 show, this is not the case. By and large, Italy's north does not do any better than the rest of the country, including its Mezzogiorno regions which are often perceived as rural and underdeveloped. This, once again, implies that companies competing internationally face particular challenges when operating in a Eurozone country that, if it had retained its national currency, would probably have depreciated against, for example, the German Mark.



Deeply red in Figure 6 is Greece. The distribution of its ATTs is shown in the lower right panel of Figure 5. Again, virtually all EMU-effects are estimated as negative. This is not surprising since Greece was the ailing Eurozone member where the sovereign debt crisis had started in 2010. The estimated EMU-effects in Greece are even more devastating than in Italy, with all but three (ADH) or four (BL) regions losing on average more than one percentage point of growth each year over 18 years of EMU – and many losing more than 1.5 or even more than 2 percentage points.

Overall, the maps in Figure 6 show that ADH and BL estimates are highly correlated. Across the whole Eurozone the correlation is 0.93, and while BL estimates are usually closer to zero, they are also less volatile: Their standard deviation is 0.7%, while ADH estimates have a standard deviation of 0.9% across the Eurozone. But note that these numbers are only indicative. They do not allow for any assessment of “significance” since the causal effects are not independent observations. Quite to the contrary, as the examples of France, Italy and Greece show.

Figure 6 indicates that Greece, Italy, France and the Wallonian part of Belgium endured particularly adverse effects by EMU, while Germany and Austria experienced mixed results

with some regions benefiting from EMU and at least as many recording mildly detrimental effects. Spain and Portugal (in the media often referred to as on equal footing with Italy or Greece) seem to have fared better than the latter two countries and better than France, although clearly worse than Germany and Austria with their net benefit just slightly below zero.

For most SC studies it is not possible to establish statistical significance in the formal sense of the word. So, are the results economically meaningful in any sense or is the true effect of EMU just zero and all results are coincidental? To answer this question, it is useful to look at (West) Germany and Italy. Suppose that the estimated ATTs are just random fluctuations around a true EMU effect of zero. Then it would be hard to explain why in Italy - a country of similar size and at least as much heterogeneity - not a single regional ATT is estimated as positive.

One might object that even if the estimated ATTs are zero mean random events they could be correlated across regions. In particular so if regions are small and in proximity to each other. Hence, one would expect regional clusters of similar ATTs also under the null of zero EMU effects. But would such a regional correlation structure extend over a whole country the size of Italy or most of France if it were just zero mean random effects which we measure? The example of Germany speaks strongly against this hypothesis: There is regional correlation for instance in the Southeast with many positive ATTs, but they do not stretch out beyond Bavaria and some neighboring regions in Baden-Württemberg. The traditional industry state of North Rhine-Westphalia in the middle western part of Germany, for instance, does not correlate greatly with Bavaria and actually has quite a few regions with negative ATTs. Turin in the north and Apulia in the south of Italy should be similarly disconnected if ATTs were just random. But they are not.

Moreover, the differences between the westernmost part of France and the bordering German regions are telling. Historically, this part of France, in particular Alsace and Lorraine, had great cultural influence from Germany and parts of the now French regions used to belong

to Germany in earlier times. But the estimated ATTs are sharply different to the East and to the West of the German-French border. This can hardly be explained by random events and it also speaks against the hypothesis of regional correlation of random events simply because of proximity and the resulting economic ties. It is much more probable that the sharp differences in ATTs are due to institutional settings which differ in France and in Germany and which tend to make the former country less successful in a currency union than the latter. Education, market regulation and wage setting behavior are among the candidates which have frequently been discussed to explain why a common monetary policy affects countries quite differently.

7. Conclusions

European Monetary Union (EMU) is one of the most important pillars of European integration. But after more than 20 years of EMU it is still not clear if EMU has benefited economic growth. The loss of exchange rate flexibility has certainly made it harder for governments to respond optimally to asymmetric shocks and, simultaneously, has disadvantaged companies which had trouble to withstand the competitive pressure of the Common Market. There is substantial evidence that such companies cluster in certain regions or even countries of the Eurozone and, therefore, the net effect of EMU on GDP may well be negative in just parts of the Eurozone, possibly on sub-country level.

In terms of regional data, the introduction of the Euro can be thought of as a natural experiment where treatment is, conditional on predictors, assigned as good as randomly. For not the economic characteristics of a certain region determined if the region joined the Eurozone. Rather, the decision was due to the will of the country's parliament and was probably the result of a complex consideration of economic, financial and political issues relevant for the country as a whole. Sometimes, e. g. in the UK, parliament was opposed to EMU, sometimes the people, in a referendum, opposed the will of the government (e. g. in Sweden). Since it is likely that some regions in countries which declined to join the Eurozone were, in economic

characteristics, not much different from regions in other countries which decided to join the Euro, this naturally suggests the former as potential control units in a major policy experiment.

Therefore, synthetic control methods may be well suited to quantify the growth effects of EMU. I have used the “canonical” ADH (2010) approach to the construction of synthetic controls along with a novel method which estimates the counterfactuals from the large sample of control regions in *post*-treatment time. While this is a great difference in underlying data, results do not differ by much. Essentially, both methods lead to the conclusion that most regions in France, Italy and Greece would have had much stronger growth had their country not joined the Eurozone. Regions with competitive industries, e. g. German car manufacturers, however, may have benefited greatly from an undervaluation of the common currency.

Hence, there is no uniformly positive or negative effect of EMU on regional growth, neither across regions nor over time. But there are many regions in the Eurozone which seemingly were disfavored by the common currency and these regions seem to cluster in certain countries like France, Italy, Greece or the French speaking part of Belgium. Also, increased capital mobility due to EMU has likely had heterogeneous effects across the Eurozone. In particular, in times of major exogenous shocks like the 2008-2009 financial crisis, it seems that EMU along with Capital Union cushioned a fairly competitive country like Germany even better than a national currency would have done, whereas regions in less advanced countries were adversely affected and would have fared better had their country still had exchange rate flexibility to counter a negative shock.

References:

Abadie, A., (2021): *Using Synthetic Controls: Feasibility, Data Requirements, and Methodological Aspects*, Journal of Economic Literature 59, pp. 391–425.

- 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, pp. 493–505.
- Abadie, A., and Gardeazabal, J., (2003): *The Economic Costs of Conflict: A Case Study of the Basque Country*, American Economic Review 93, pp. 113-132.
- Abadie, A., and L'Hour, J., (2021): *A Penalized Synthetic Control Estimator for Disaggregated Data*, Journal of the American Statistical Association 116, pp. 1817-1834.
- Athey, S., and Imbens, G. W., (2017): *The state of applied econometrics: Causality and policy evaluation*, Journal of Economic Perspectives 31, pp. 3–32.
- Bai, J., (2003): *Inferential Theory for Factor Models of Large Dimensions*, Econometrica 71, pp. 135-171.
- Bai, J., and Ng, S. (2002): *Determining the Number of Factors in Approximate Factor Models*, Econometrica 70, pp. 191-221.
- Becker, M., and Klößner, S., (2018): *Fast and reliable computation of generalized synthetic controls*, Econometrics and Statistics 5, pp. 1–19.
- Ben-Michael, E., Feller, A., and Rothstein, J., (2021): *The Augmented Synthetic Control Method*, Journal of the American Statistical Association, pp. 1–34, DOI: 10.1080/01621459.2021.1929245.
- Chernozhukov, V., Wüthrich, K., and Zhu, Y., (2021): *An Exact and Robust Conformal Inference Method for Counterfactual and Synthetic Controls*, arXiv:1712.09089v10 [econ.EM] 20 May 2021.
- Conti, M., (2014): *The Introduction of the Euro and Economic Growth: Some Panel Data Evidence*, Journal of Applied Economics 17, pp. 199-212.
- Dreyer, J. K., and Schmid, P., (2016): *Does the Eurozone enhance Economic Growth? EU and EZ Growth Effects following the Introduction of the Euro*, Proceedings of the 15th International Conference on Finance and Banking, Silesian University.

European Commission (2022): *Cohesion in Europe towards 2050*, 8th Report on Economic, Social and Territorial Cohesion, Brussels.

Ferman, B., and Pinto, C., (2016): *Revisiting the Synthetic Control Estimator*, MPRA Paper No. 73982, <https://mpra.ub.uni-muenchen.de/73982/>.

Fernández, C., and Garcia Perea, P., (2015): *The Impact of the Euro on Euro Area GDP Per Capita*, Banco de Espana Working Paper No. 1530, <https://ssrn.com/abstract=2690211> or <http://dx.doi.org/10.2139/ssrn.2690211>

Gabriel, R. D., and Pessoa, A. S., (2020): *Adopting the Euro: A Synthetic Control Approach*, <https://ssrn.com/abstract=3563044> or <http://dx.doi.org/10.2139/ssrn.3563044>

Garcia-Penalosa, C. (2010): Income distribution, economic growth and European integration. *The Journal of Economic Inequality*, 8(3): pp. 277–292. <https://doi.org/10.1007/s10888-010-9140-5>

Gasparotti, A., and Kullas, M., (2019): *20 Years of the Euro: Winners and Losers*, cep study, Centrum für Europäische Politik, Freiburg.

Holland, P. W., (1986): *Statistics and Causal Inference*, *Journal of the American Statistical Association* 81, pp. 945-960.

Empirica, 40(1): pp. 197–214. <https://doi.org/10.1007/s10663-011-9182-3>

Italian Trade Agency (2017): *Italy in the World Economy, Report 2016-2017*, Rome, <https://www.italiancompaniesworld.com/documents/ice-annual-report-italy-in-the-world-market.html>.

Ioannatos, P. E., (2018): *Has the Euro promoted Eurozone's growth?*, *Journal of Economic Integration* 33, pp. 1388-1411.

Kalaitzoglou, I. and Durgheu, B. (2016): *Financial and economic growth in Europe: Is the Euro beneficial for all countries?* *Journal of Economic Integration*, 31(2), pp. 414–471.

Kaul, A., Klößner, S., Pfeifer, G., and Schieler, M., (2015): *Synthetic Control Methods: Never Use All Pre-Intervention Outcomes Together with Covariates*, MPRA Paper 83790.

Kleven, H. J., Landais, C., & Saez, E., (2013): *Taxation and international migration of superstars: Evidence from the European football market*, American Economic Review 103, pp. 1892–1924.

Kutan, A. M. and Yigit, T. M. (2007). European integration, productivity growth and real convergence. European Economic Review, 51(6): 1370–1395.
<https://doi.org/10.1016/j.euroecorev.2006.11.001>

Lucke, B., (2022): *Consistent Estimation of Optimal Synthetic Control Weights*, Discussion Paper, University of Hamburg.

Malo, P., Eskelinen, J., Zhou, X., and Kuosmanen, T., (2020): *Computing Synthetic Controls Using Bilevel Optimization*, MPRA Paper No. 104085, <https://mpra.ub.uni-muenchen.de/104085/>

OECD (2020): *OECD Regions and Cities at a Glance 2020*, OECD Publishing, Paris, <https://doi.org/10.1787/959d5bA0-en>.

Puzzello, L., and Gomis-Porqueras, P., (2018): *Winners and Losers from the Euro*, European Economic Review 108, pp. 129-152.

Rubin, D. B., (1974): *Estimating Causal Effects of Treatments in Randomized and Nonrandomized Studies*, Journal of Educational Psychology 66, pp. 688-701.

Sinn, H.-W., (2003): *The Laggard of Europe*, cesifo Forum vol 4, Special Issue 1.

Sinn, H.-W., (2014): *The Euro Trap*, Oxford University Press, ISBN: 978-0-19-870213-9.

Verstegen, L., van Groezen, B., and Meijdam, L., (2017): *Benefits of EMU Participation: Estimates Using the Synthetic Control Method*, Tilburg University Discussion Paper 2017-032.

Appendix: Proofs

Proof of Proposition 1:

Clearly, any $w^+ \in \Phi_1^Z \cap \Phi_1^{Y^{pre}}$ implies $z_1 = Z_0 w^+$ and $y_1^{pre} = Y_0^{pre} w^+$. Hence, $L_{out}(v, w) = L_{in}(v, w) = 0$ for all choices of V . Since both problems are quadratic, this is the global minimum, i. e. w^+ solves (11).

Moreover, by (10), $\Lambda^{pre}(\mu_1 - M_0 w^+) = y_1^{pre} - Y_0^{pre} w^+ = 0$. This must be true for all possible shocks Λ^{pre} , so that we have $\mu_1 = M_0 w^+$. It follows that for $w = w^+$ the synthetic control error

$$\varepsilon_1^{post}(w) = y_1^{N, post} - Y_0^{post} w = \Theta^{post}(z_1 - Z_0 w) + \Lambda^{post}(\mu_1 - M_0 w)$$

is identically equal to zero, $\varepsilon_1^{post}(w^+) \equiv 0$. Thus, w^+ solves (12).

qed.

Proof of Proposition 2:

Multiplying the objective in (12) by the constant factor T_1 we have

$$\begin{aligned} T_1 MSPE_1^{post} &= E\left(\varepsilon_1^{post}(w)' \varepsilon_1^{post}(w) \middle| I_0\right) = E\left((y_1^N - Y_0^{post} w)'(y_1^N - Y_0^{post} w) \middle| I_0\right) \\ &= E\left(y_1^N' y_1^N - 2y_1^N'(\Theta^{post} Z_0 + \Lambda^{post} M_0)w + w'(\Theta^{post} Z_0 + \Lambda^{post} M_0)'(\Theta^{post} Z_0 + \Lambda^{post} M_0)w \middle| I_0\right) \\ &= E\left(y_1^N' y_1^N \middle| I_0\right) + w' Z_0' \Theta^{post} \Theta^{post} Z_0 w \\ &\quad + w' E\left(M_0' M_0 + Z_0' \Theta^{post} \Lambda^{post} M_0 + M_0' \Lambda^{post} \Theta^{post} Z_0 \middle| I_0\right) w \\ &\quad - 2E\left((\Theta^{post} z_1 + \Lambda^{post} \mu_1)'(\Theta^{post} Z_0 + \Lambda^{post} M_0) \middle| I_0\right) w \\ &= E\left(y_1^N' y_1^N \middle| I_0\right) + w' Z_0' \Theta^{post} \Theta^{post} Z_0 w + w' M_0' M_0 w \\ &\quad - 2E\left(z_1' \Theta^{post} \Theta^{post} Z_0 + z_1' \Theta^{post} \Lambda^{post} M_0 + \mu_1' \Lambda^{post} \Theta^{post} Z_0 + \mu_1' M_0 \middle| I_0\right) w \\ &= E\left(y_1^N' y_1^N \middle| I_0\right) + w' (Z_0' \Theta^{post} \Theta^{post} Z_0 + M_0' M_0) w - 2(z_1' \Theta^{post} \Theta^{post} Z_0 + \mu_1' M_0) w \end{aligned}$$

The first term, $E(y_1^N' y_1^N \middle| I_0)$, is independent of w . Therefore, (12) is equivalent to solving

$$\min_{w \in \Delta_J} w' E\left(Z_0' \Theta^{post} \Theta^{post} Z_0 + M_0' M_0\right) w - 2E\left(z_1' \Theta^{post} \Theta^{post} Z_0 + \mu_1' M_0\right) w$$

qed.

Proof of Proposition 3:

a) By A1 and the Law of Iterated Expectations we have

$$E(H_0^{post}) = E(E(H_0^{post} | z_1, Z_0)) = E(E(\Lambda^{post} | z_1, Z_0) M_0) = 0.$$

Hence, A1 along with the Weak Law of Large numbers implies $p \lim_{J \rightarrow \infty} J^{-1} H_0^{post} Z_0' = 0$. Using

(4) we get

$$p \lim_{J \rightarrow \infty} Y_0^{post} Z_0' (Z_0 Z_0')^{-1} = \Theta^{post} + p \lim_{J \rightarrow \infty} J^{-1} H_0^{post} Z_0' \quad p \lim_{J \rightarrow \infty} (J^{-1} Z_0 Z_0')^{-1} = \Theta^{post}$$

b) Observe that

$$\begin{aligned} \hat{\Theta}_1^{post} - \hat{\Theta}_0^{post} &= Y_1^{post} Z_1' (Z_1 Z_1')^{-1} - \hat{\Theta}_0^{post} \\ &= (\Theta^{post} Z_1 + \Lambda^{post} M_1 + A) Z_1' (Z_1 Z_1')^{-1} - \hat{\Theta}_0^{post} \\ &= \Theta^{post} - \hat{\Theta}_0^{post} + (\Lambda^{post} M_1 + B_A Z_1 + \Lambda^{post} C_A + E_A) Z_1' (Z_1 Z_1')^{-1} \\ &= B_A + \Theta^{post} - \hat{\Theta}_0^{post} + (\Lambda^{post} (M_1 + C_A) + E_A) Z_1' (Z_1 Z_1')^{-1} \\ &= B_A + \underbrace{\Theta^{post} - \hat{\Theta}_0^{post}}_{\xrightarrow[p \lim_{J \rightarrow \infty}]{} 0} + \left(\underbrace{L^{-1} \Lambda^{post} (M_1 + C_A) Z_1}_{\xrightarrow[p \lim_{L \rightarrow \infty}]{} 0} + \underbrace{L^{-1} E_A Z_1}_{\xrightarrow[p \lim_{L \rightarrow \infty}]{} 0} \right) \underbrace{(L^{-1} Z_1 Z_1')^{-1}}_{\xrightarrow[p \lim_{L \rightarrow \infty}]{} \text{finite}} \xrightarrow[p \lim_{J, L \rightarrow \infty}]{} B_A \end{aligned}$$

The convergence of the first term in parentheses follows from assumption A1 along with the Weak Law of Large Numbers.

qed.

Proof of Proposition 4:

a) A regression of Y_{10}^{post} on Z_{10} yields

$$\begin{aligned} \hat{\Theta}_{10}^{post} - \Theta^{post} &= \left((y_1^{post} - \hat{a}_1) \quad Y_0 \right) (z_1 \quad Z_0)' (Z_{10} Z_{10}')^{-1} - \Theta^{post} \\ &= \left((\Theta^{post} z_1 + \Lambda^{post} \mu_1 + (\alpha_1 - \hat{a}_1)) z_1' + (\Theta^{post} Z_0 + \Lambda^{post} M_0) Z_0' \right) (z_1 z_1' + Z_0 Z_0')^{-1} - \Theta^{post} \\ &= \underbrace{J^{-1} \Lambda^{post} (\mu_1 z_1' + M_0 Z_0')}_{\xrightarrow[p \lim_{J \rightarrow \infty}]{} 0} \underbrace{(J^{-1} (z_1 z_1' + Z_0 Z_0'))^{-1}}_{\xrightarrow[p \lim_{J \rightarrow \infty}]{} M_{Z_0 Z_0}^{-1}} + \underbrace{J^{-1} (\alpha_1 - \hat{a}_1) z_1'}_{\xrightarrow[p \lim_{J \rightarrow \infty}]{} 0} \underbrace{(J^{-1} (z_1 z_1' + Z_0 Z_0'))^{-1}}_{\xrightarrow[p \lim_{J \rightarrow \infty}]{} z_1' M_{Z_0 Z_0}^{-1}} \xrightarrow[p \lim_{J \rightarrow \infty}]{} 0 \end{aligned}$$

Observe that for given (fixed) L , $p \lim_{J \rightarrow \infty} (\alpha_1 - \hat{a}_1)$ is nonzero but finite, cf. the proof of

Proposition 3. Therefore, even for fixed L , we have $p \lim_{J \rightarrow \infty} J^{-1} (\alpha_1 - \hat{a}_1) = 0$.

b) Suppose that $L = \kappa J$, $\kappa \in \mathbb{R}^+$, as J approaches infinity. A regression of Y^{post} on Z yields

$$\begin{aligned} \hat{\Theta}^{post} - \Theta^{post} &= \left((Y_1^{post} - \hat{B}_A Z_1)' \quad Y_0' \right) (Z_1' \quad Z_0')^{-1} - \Theta^{post} \\ &= \left((\Theta^{post} Z_1 + \Lambda^{post} (M_1 + C_A) + (B_A - \hat{B}_A) Z_1 + E_A) Z_1' + (\Theta^{post} Z_0 + \Lambda^{post} M_0) Z_0' \right) (Z_1 Z_1' + Z_0 Z_0')^{-1} - \Theta^{post} \\ &= \left(\Lambda^{post} ((M_1 + C_A) Z_1' + M_0 Z_0') + (B_A - \hat{B}_A) Z_1 Z_1' + E_A Z_1' \right) (Z_1 Z_1' + Z_0 Z_0')^{-1} \\ &= \kappa \left(\underbrace{L^{-1} \Lambda^{post} ((M_1 + C_A) Z_1' + M_0 Z_0')}_{\xrightarrow[p \lim_{J, L \rightarrow \infty}]{}} 0 + \underbrace{(B_A - \hat{B}_A)}_{\xrightarrow[p \lim_{J, L \rightarrow \infty}]{}} \underbrace{L^{-1} Z_1 Z_1'}_{\xrightarrow[p \lim_{L \rightarrow \infty}]{}} M_{Z_1 Z_1} + \underbrace{L^{-1} E_A Z_1'}_{\xrightarrow[p \lim_{L \rightarrow \infty}]{}} 0 \right) \left(\underbrace{\kappa L^{-1} Z_1 Z_1'}_{\xrightarrow[p \lim_{L \rightarrow \infty}]{}} \kappa M_{Z_1 Z_1} + \underbrace{J^{-1} Z_0 Z_0'}_{\xrightarrow[p \lim_{J \rightarrow \infty}]{}} M_{Z_0 Z_0} \right)^{-1} \xrightarrow[p \lim_{J, L \rightarrow \infty}]{}} 0 \end{aligned}$$

where $p \lim_{J, L \rightarrow \infty} \hat{B}_A = B_A$ follows from Proposition 3.

c) A regression of $Y_1^{post} - \hat{A}_z$ on Z_1 yields

$$\begin{aligned} \hat{\Theta}_1^{post} - \Theta^{post} &= (Y_1^{post} - \hat{A}_z)' Z_1' (Z_1 Z_1')^{-1} - \Theta^{post} \\ &= (Y_1^{post} - \hat{B}_A Z_1)' (Z_1 Z_1')^{-1} - \Theta^{post} \\ &= (\Theta^{post} Z_1 + \Lambda^{post} M_1 + (A - \hat{B}_A Z_1)) Z_1' (Z_1 Z_1')^{-1} - \Theta^{post} \\ &= (\Theta^{post} Z_1 + \Lambda^{post} (M_1 + C_A) + (B_A - \hat{B}_A) Z_1 + E_A) Z_1' (Z_1 Z_1')^{-1} - \Theta^{post} \\ &= \underbrace{(B_A - \hat{B}_A)}_{\xrightarrow[p \lim_{J, L \rightarrow \infty}]{}} 0 + \left(\underbrace{L^{-1} \Lambda^{post} (M_1 + C_A) Z_1'}_{\xrightarrow[p \lim_{L \rightarrow \infty}]{}} 0 + \underbrace{L^{-1} E_A Z_1'}_{\xrightarrow[p \lim_{L \rightarrow \infty}]{}} 0 \right) \underbrace{(L^{-1} Z_1 Z_1')^{-1}}_{\xrightarrow[p \lim_{L \rightarrow \infty}]{}} M_{Z_1 Z_1}^{-1} \xrightarrow[p \lim_{J, L \rightarrow \infty}]{}} 0 \end{aligned}$$

where $p \lim_{J, L \rightarrow \infty} \hat{B}_A = B_A$ follows from Proposition 3.

qed.