

Comparative Politics and the Synthetic Control Method Revisited: A Note on Abadie et al. (2015)

Stefan Klößner, Saarland University*
Ashok Kaul, Saarland University
Gregor Pfeifer, University of Hohenheim
Manuel Schieler, Saarland University

This Version: January 31, 2017
(First Version: January 13, 2016)

Abstract

Recently, Abadie et al. (2015) have expanded synthetic control methods by the so-called cross-validation technique. We find that their results are not being reproduced when alternative software packages are used or when the variables' ordering within the dataset is changed. We show that this failure stems from the cross-validation technique relying on non-uniquely defined predictor weights. For several of the robustness analyses of Abadie et al. (2015), we find the amount of the resulting ambiguity to be substantial. Applying unflawed, standard synthetic control methods reveals that the authors' results are particularly driven by a specific control country, the United States.

*Corresponding author: Stefan Klößner, Saarland University, Statistics & Econometrics, C3 1, D-66123, Germany; S.Kloessner@mx.uni-saarland.de; Phone 0049 681 302 3179.

1 Introduction

As a tool for policy evaluation, Abadie and Gardeazabal (2003) have introduced so-called synthetic control methods (SCM). For estimating the development of the treated unit in the absence of the treatment, the basic idea of SCM is to find suitable donor weights which describe how the treated unit is synthesized by a weighted mix of unaffected control units. In this context, 'suitable' means that treated and synthetic unit should resemble each other as closely as possible prior to the treatment, both with respect to the outcome of interest and with respect to so-called economic predictors. The latter are variables of predictive power for explaining the outcome. The data-driven SCM approach searches for optimal predictor weights in order to grant more importance to economic predictors with better predictive power. Properties of the SCM estimator, like (asymptotic) unbiasedness, have been developed by Abadie et al. (2010), while Gardeazabal and Vega-Bayo (2016) find that the SCM estimator performs well as compared to alternative panel approaches.

Over the last few years, many studies have applied SCM across several fields, e.g., Acemoglu et al. (2016) (political connections), Cavallo et al. (2013) (natural disasters), Gobillon and Magnac (2016) (enterprise zones), or Kleven et al. (2013) (taxation of athletes). Recently, the SCM approach has been expanded by Abadie et al. (2015) (German Reunification) to incorporate *cross-validation*: the predictor weights, whose data in the training period (first part of the pre-treatment period) are used to find optimal donor weights for synthesizing the treated unit, are selected such that the out-of-sample error in the validation period (second part of the pre-treatment period) is minimized.

When measuring the effect of the 1990 reunification on Germany's GDP per capita using the software package R, Abadie et al. (2015) found the following predictor weights: 44.2% (GDP per capita), 24.5% (investment rate), 13.4% (trade openness), 10.7% (amount of schooling), 7.2% (inflation rate), and 0.1% (industry share of value added). These predictor weights led to Germany being synthesized by Austria (42%), the United States (22%), Japan (16%), Switzerland (11%), and the Netherlands (9%).

When trying to replicate these results using the software package Stata, however, we found different predictor weights: 84.5% (GDP), 4.5% (investment), 5.1% (trade), 4.2% (schooling), 0.5% (inflation), and 1.2% (industry). The corresponding synthetic Germany was slightly different from the one obtained by Abadie et al. (2015): it consisted of Austria (43%), the United States (22%), Japan (15%), Switzerland (11%), and the Netherlands (9%).¹ We had sorted the countries alphabetically, while Abadie et al. (2015) had used a different ordering.² Although in theory, the ordering should have no effect on the estimation results (neither should the respective software package), we recalculated all weights using the ordering that had been used by Abadie et al. (2015). Surprisingly, we got yet *another* set of predictor weights: 71.0% (GDP), 11.1% (investment), 7.9% (trade), 6.4% (schooling), 2.7% (inflation), and 0.9% (industry). The corresponding weights for the countries synthesizing Germany were much closer to, but still different from the values found by Abadie et al. (2015).³

Closer inspection shows that the failure to reproduce the results of Abadie et al. (2015) stems from the newly introduced cross-validation technique. All the above mentioned predictor weights are equivalent solutions of the cross-validation approach, meaning that this technique is (in most applications) not well-defined since the predictor weights are not uniquely defined. As the cross-validation technique allows many different equivalent predictor weights, the results obtained by Abadie et al. (2015) are arbitrary in the sense that the authors could have obtained different results if they had used other software or organized the data differently.

We therefore investigate the corresponding ambiguity by conducting large-scaled Monte Carlo studies. The variation of the estimated post-treatment development of West German GDP is very small, with all estimates being significantly above Germany's actual GDP. Concerning several robustness studies of Abadie et al. (2015), however, we find quite large amounts of ambiguity, in particular for the so-called in-space placebo and leave-one-out studies. Thus,

¹The seemingly small differences of the countries' weights delivered by Abadie et al. (2015) (also called ADH subsequently) and Stata are actually more pronounced, cf. the 'W Weights Main' entries of columns 'ADH' and 'Alph.' in Table 1 below.

²The data file used by Abadie et al. (2015) implicitly orders the countries by the U.S., the UK, Austria, Belgium, Denmark, France, West Germany, Italy, the Netherlands, Norway, Switzerland, Japan, Greece, Portugal, Spain, Australia, New Zealand.

³See the 'W Weights Main' entries of columns 'ADH' and 'Orig.' in Table 1 below.

using the cross-validation cannot be recommended and standard synthetic control methods should be applied instead. When doing so, we confirm the main result of Abadie et al. (2015), indicating a significant drop in West German GDP due the reunification. In contrast to Abadie et al. (2015), however, detecting such a significant gap crucially hinges on including U.S. data.

The remainder of the paper unfolds as follows: Section 2, describes the synthetic control method with and without cross-validation. In Section 3, we elaborate on the reasons why the cross-validation technique is typically not well-defined, while the standard SCM approach does not suffer from this problem. We then analyze the extent to which the results of Abadie et al. (2015) are prone to ambiguity and compare them to those under the standard synthetic control approach. Section 5 concludes.

2 Synthetic Control Methods

In the following, we describe how synthetic control methods work both with and without the cross-validation technique. Many additional explanations, in particular on how to select potential comparison units and predictor variables, are provided in Abadie et al. (2015).⁴

For the synthetic control method, we have two types of data: the variable of interest, often denoted by the letter Y , and predictor variables, usually denoted by X . These are considered both for a unit that has at some point in time been 'treated', usually denoted by the subscript $_1$, and for so-called donor units. The latter are units not too different from the first one, but unaffected from the treatment, and denoted by the subscript $_0$. In the example discussed throughout this paper, the treated unit is Germany which has been reunified in 1990, the variable of interest is GDP per capita, and predictors are GDP per capita, a measure for trade openness, the inflation rate, the industry share of value added, the amount of schooling attained, and the investment rate. The donor units consist of sixteen OECD countries⁵ for which the synthetic control method determines non-negative so-called donor

⁴As we need a very formal representation of the SCM techniques for later use, we provide a rather mathematical translation of the verbal description and R-code given by Abadie et al. (2015).

⁵For more details on variables as well as donor choice, see Abadie et al. (2015, p. 509).

weights W : these weights describe to what extent each donor country is used to produce a 'synthetic' (i.e., counterfactual) Germany. Thereby, the weights should be such that synthetic Germany mimics actual Germany as well as possible with respect to the (pre-treatment) predictor variables. For the example at hand, this means that the differences between actual and synthetic Germany with respect to GDP per capita, trade openness, inflation rate, industry share, schooling, and investment rate should be as small as possible. As we have six predictors ($k = 6$), operationalizing the last statement requires introducing some weighting scheme for the predictors. These non-negative so-called predictor weights are usually denoted by v_m or V , and the cross-validation technique introduced in Abadie et al. (2015) is a new method to determine these weights. To this end, the pre-treatment period is divided into two parts, a training and a validation period: for the case of the German reunification, the training period is 1971-1980, while the validation period is 1981-1990.

In the training period, one makes use of the $(k \times J)$ -matrix $X_0^{(\text{train})}$ and the k -dimensional vector $X_1^{(\text{train})}$, containing time averages of the predictors' data for the donor units and the treated unit, respectively.⁶ For any predictor weights $V = (v_1, \dots, v_k)$, the donor weights $W_{(\text{train})}^*(V)$ in the training period are defined as the minimizer of $\sum_{m=1}^k v_m \left(X_{1m}^{(\text{train})} - X_{0m}^{(\text{train})} W \right)^2$ with respect to J -dimensional non-negative donor weights W summing to unity, i.e., as the solution of

$$\min_W \sum_{m=1}^k v_m \left(X_{1m}^{(\text{train})} - X_{0m}^{(\text{train})} W \right)^2 \quad \text{s.t.} \quad W \geq 0, \mathbb{1}'W = 1, \quad (1)$$

where $\mathbb{1}$ denotes the vector of ones, while $X_{1m}^{(\text{train})}$ and $X_{0m}^{(\text{train})}$ denote the m -th component and row of $X_1^{(\text{train})}$ and $X_0^{(\text{train})}$, respectively.

In the validation period, one uses the $(L \times J)$ -matrix $Y_0^{(\text{valid})}$ and the L -dimensional vector $Y_1^{(\text{valid})}$, containing the variable of interest's data for the validation period.⁷ The cross-validation defines predictor weights $V^* = (v_1^*, \dots, v_k^*)$ as those predictor weights that minimize

⁶ J denotes the number of donor units used to synthesize the treated unit, for the application under consideration, the German reunification, the $J = 16$ donor units are given by the U.S., the UK, Austria, Belgium, Denmark, France, Italy, the Netherlands, Norway, Switzerland, Japan, Greece, Portugal, Spain, Australia, and New Zealand.

⁷For the application at hand, annual data from 1981 to 1990 are used, resulting in $L = 10$.

the out-of-sample error $\|Y_1^{(\text{valid})} - Y_0^{(\text{valid})}W_{(\text{train})}^*(V)\|^2$ over V , i.e., V^* is the solution of⁸

$$\min_V \|Y_1^{(\text{valid})} - Y_0^{(\text{valid})}W_{(\text{train})}^*(V)\|^2 \text{ s.t. } V \geq 0, \mathbb{1}'V = 1, \quad (2)$$

where we have normalized the predictor weights to sum to unity.⁹

These predictor weights V^* are then used to determine $W_{(\text{main})}^*$ as the minimizer of $\sum_{m=1}^k v_m^* \left(X_{1m}^{(\text{valid})} - X_{0m}^{(\text{valid})}W \right)^2$, i.e., as the solution of

$$\min_W \sum_{m=1}^k v_m^* \left(X_{1m}^{(\text{valid})} - X_{0m}^{(\text{valid})}W \right)^2 \text{ s.t. } W \geq 0, \mathbb{1}'W = 1, \quad (3)$$

where the $(k \times J)$ -matrix $X_0^{(\text{valid})}$ and the k -dimensional vector $X_1^{(\text{valid})}$ contain time averages of the predictors' data for the validation period, which for the application at hand ranges from 1981 to 1990.

Thus, the synthetic control method with cross-validation is a two-step procedure: first, in the so-called 'training' step, V^* is determined by minimizing the cross-validation criterion, thereby making use of 'training' weights $W_{(\text{main})}^*(V)$ as defined by Equation (1). Then, in the second, so-called 'main' step, these predictor weights V^* are used to determine the 'main' donor weights $W_{(\text{main})}^*(V^*)$ by Equation (3). These 'main' donor weights $W_{(\text{main})}^*(V^*)$ are then employed for synthesizing the treated unit.

In contrast, the standard synthetic control method consists of only one step, not distinguishing between a training and validation period. Instead, all pre-treatment data are used, in our application those from 1971 to 1990, to build the following quantities: the $(k \times J)$ -matrix X_0 and the k -dimensional vector X_1 , containing time averages of the predictors' data for the donor units and the treated unit, respectively, as well as the $(\tilde{L} \times J)$ -matrix Y_0 and the \tilde{L} -dimensional vector Y_1 , containing the variable of interest's pre-treatment data for the donor units and the

⁸Abadie et al. (2015, p. 502): 'Intuitively, the cross-validation technique selects the weights v_m that minimize out-of-sample prediction errors.'

⁹ As obviously $W_{(\text{train})}^*(\alpha V) = W_{(\text{train})}^*(V)$ for all predictor weights V and positive constants $\alpha > 0$, one may assume without loss of generality that predictor weights are always scaled such that their components sum to unity, see, e.g., (Abadie and Gardeazabal, 2003, p. 128), (Abadie et al., 2015, Footnote 5, p. 497f).

treated unit, respectively.¹⁰ For given predictor weights $V = (v_1, \dots, v_k)$, the standard SCM approach defines the donor weights $W^*(V)$ as the minimizer of $\sum_{m=1}^k v_m (X_{1m} - X_{0m}W)^2$, i.e., as the solution of

$$\min_W \sum_{m=1}^k v_m (X_{1m} - X_{0m}W)^2 \quad \text{s.t. } W \geq 0, \mathbb{1}'W = 1. \quad (4)$$

Optimal predictor weights V^* are then determined by minimizing the in-sample error¹¹, i.e., as the solution of

$$\min_V \|Y_1 - Y_0 W^*(V)\|^2 \quad \text{s.t. } V \geq 0, \mathbb{1}'V = 1. \quad (5)$$

The donor weights $W^*(V^*)$ are then used for synthesizing the treated unit.

3 Well-Definedness of Synthetic Control Methods

A crucial insight as to why the cross-validation technique of Abadie et al. (2015) is not well-defined is the fact that, typically, there is no *unique* minimizer of the out-of-sample error $\|Y_1^{(\text{valid})} - Y_0^{(\text{valid})} W_{(\text{train})}^*(V)\|^2$, thus, Equation (2) does not define V^* unambiguously. The reason is that the mapping $W_{(\text{train})}^*$ defined by Equation (1) is often not injective — it regularly happens that $W_{(\text{train})}^*(\tilde{V})$ and $W_{(\text{train})}^*(V)$ coincide although \tilde{V} and V are different after scaling. Less formally, it is often the case that different predictor weights lead to the same 'training' weights. The problem of the cross-validation approach is that such different predictor weights \tilde{V} and V , although scaled and entailing identical $W_{(\text{train})}^*(\tilde{V}) = W_{(\text{train})}^*(V)$, typically lead to different $W_{(\text{main})}^*(\tilde{V}) \neq W_{(\text{main})}^*(V)$ in the 'main' step in Equation (3).

Actually, this is the reason behind the diverging results described above: all predictor weights given earlier, those found by Abadie et al. (2015) as well as our results obtained using Stata with two different W orderings for the donor countries, are equivalent solutions of

¹⁰ For the German reunification, annual data for the pre-treatment time span 1971-1990 are used, resulting in $\tilde{L} = 20$.

¹¹ There exists another method to determine predictor weights, the so-called regression-based method, which however is rarely used in practice.

Equation (2). This can be seen from the 'W Weights Training' rows of columns 'ADH', 'Orig.', and 'Alph.' of Table 1 below — all these different predictor weights produce the same $W_{(\text{train})}^*$ in Equation (1), leading to identical out-of-sample errors. However, although these different predictor weights V are equivalent with respect to Equation (1), i.e., produce the same donor weights $W_{(\text{train})}^*$ and therefore the same synthetic Germany in the training period, they are *not* equivalent with respect to Equation (3). More specifically, the corresponding donor weights $W_{(\text{main})}^*$ for the main application do not coincide (cf. the 'W Weights Main' rows of columns 'ADH', 'Orig', and 'Alph.' of Table 1). Overall, therefore, one obtains different synthetic versions for the treated unit, leading to different estimates for the post-treatment development of the treated unit in absence of the intervention, and potentially to diverging conclusions about the effect of the intervention. Thus, in the end, the cross-validation technique introduced by Abadie et al. (2015) is not properly defined, typically leading to ambiguous estimates of the treatment effect.

While $W_{(\text{train})}^*$ *in general* is not injective, it depends on the respective application whether or not there exist several different predictor weights minimizing the out-of-sample error, resulting in Equation (2) defining a *unique* vector of predictor weights or a set thereof. In some applications, there might be an up to scaling unique minimizer, making the cross-validation technique well-defined, while in other applications, there might exist many different minimizers. In the latter case, it is not clear how large the set of these minimizers will be. In the appendix, we therefore elaborate on a heuristic rule of thumb that allows to get an idea about the amount of ambiguity. It turns out that the decisive quantity in this context is the difference $k - \alpha$ between the number of predictors used, k , and the number of donor units that obtain positive weights in the training period, $\alpha := \#\{j : W_{(\text{train}),j}^* > 0\}$. If the difference $k - \alpha$ is positive, the predictor weights will typically not be uniquely defined by the cross-validation technique, with a generically increasing amount of ambiguity the larger $k - \alpha$. In case of the German reunification, six predictors (GDP, trade openness, inflation, industry share, schooling, investment rate), are used, but only five donor units obtain positive weights

in the training period (the U.S., Austria, Switzerland, Japan, Australia, cf. Table 1), thus $k - \alpha = 6 - 5 = 1 > 0$. Consequently, there exist many solutions for determining the predictor weights by the cross-validation technique, but the amount of ambiguity is expected to be rather small.

Finally, we briefly discuss why the standard synthetic control method does not suffer from similar problems. For the standard method, V^* as defined by Equation (5) in general consists of many solutions, in complete analogy to the cross-validation method and Equation (2). Again, the reason is that the mapping W^* defined by Equation (4) is often not injective: it regularly happens that $W^*(\tilde{V})$ equals $W^*(V)$ although \tilde{V} and V are different after scaling. However, in contrast to the cross-validation method, this is not a problem: the only quantity needed to synthesize the treated unit and estimate treatment effects are the donor weights $W^*(V^*)$ — which are indeed unique, even if V^* is not.

4 The Economic Cost of the 1990 German Reunification

4.1 Ambiguity of Results Using the Cross-Validation Technique

For the upcoming analysis, we retrieved from the AJPS Data Archive on Dataverse (<https://dataverse.harvard.edu/dataverse/ajps>) both the data and all code of Abadie et al. (2015). We also followed Abadie et al. (2015) by using R (R Core Team (2014)) in combination with package Synth (Abadie et al. (2011)). In particular, we first ran the code supplied by Abadie et al. (2015), storing all the results, especially the results for the donor units' weights in the training period. We then conducted large-scale Monte Carlo studies, searching for predictor weights that also lead to these donor weights in the training period, i.e., 'training-equivalent' predictor weights which also minimize Equation (2). These were then used to calculate the corresponding donor weights for the main period, GDP estimates, and follow-up quantities.¹²

¹²More precisely, we simulated values for v_1, \dots, v_k by independent draws from the Cauchy distribution, solved Equation (1) and checked whether $W_{(\text{train})}^*(v_1, \dots, v_k)$ was up to four digits equal to the 'training' weights given in Table 1. If this was the case, we computed the corresponding 'main' W weights and the other follow-up quantities like GDP estimates, etc., and stored the corresponding V weights for later use. The whole

		ADH	Orig.	Alph.	Min.	Max.
V Weights	GDP per capita	44.2	71.0	84.5	38.5	87.7
	Trade openness	13.4	7.9	5.1	4.4	14.6
	Inflation rate	7.2	2.7	0.5	0.0	8.1
	Industry share	0.1	0.9	1.2	0.0	1.3
	Schooling	10.7	6.4	4.2	3.7	11.6
	Investment rate	24.5	11.1	4.5	2.9	27.2
W Weights Main	USA	21.9	22.1	22.0	20.1	22.1
	UK	0.1	0.0	0.0	0.0	1.8
	Austria	41.8	42.2	43.1	41.6	47.1
	Netherlands	9.0	9.2	8.5	0.6	9.4
	Norway	0.1	0.0	0.0	0.0	1.9
	Switzerland	11.1	10.9	10.8	10.8	13.7
	Japan	15.5	15.7	15.4	9.9	15.7
W Weights Training	USA	13.5	13.5	13.5	13.5	13.5
	Austria	50.7	50.7	50.7	50.7	50.7
	Switzerland	16.6	16.6	16.6	16.6	16.6
	Japan	14.6	14.6	14.6	14.6	14.6
	Australia	4.5	4.5	4.5	4.5	4.5

Table 1: Results (predictor weights V , donor weights W for main application and in training period) obtained in different ways: 'ADH' stands for the results of Abadie et al. (2015), 'Orig.' are results from Stata with the same ordering of donors as in the code of Abadie et al. (2015), 'Alph.' denotes results from Stata with donors sorted alphabetically, 'Min.' and 'Max.' denote minimal and maximal values, respectively, found under the condition that the corresponding predictor weights V lead to identical donor weights W in the training period. All numbers are given in %, suppressing donors with weight less than 1%.

Table 1 summarizes the results for the predictor weights V and donor weights W : the columns 'ADH', 'Orig.', and 'Alph.' contain the results obtained by Abadie et al. (2015), by Stata using the same ordering of donor countries as did Abadie et al. (2015), and by Stata using the donor countries in alphabetical order, respectively. The 'Min.' and 'Max.' columns contain the smallest and largest values obtained in our Monte Carlo study, respectively. We find the weights of some predictors to vary substantially. For instance, the V weight of GDP can take values between 38.5% and 87.7%, while the inflation rate may be almost irrelevant with a weight of nearly zero, but also taken into considerable account when its weight is 8.1%. For the composition of synthetic Germany, we find similar ambiguity: the weight of Austria varies between 41.6% and 47.1%, the Netherlands can be essentially unimportant with a weight of only 0.6%, but also contribute 9.4% to synthetic Germany. In some cases, Germany is even

procedure was repeated for several million draws of the V weights.

synthesized by six instead of five countries, when the UK or Norway are attributed small but positive weights, respectively.

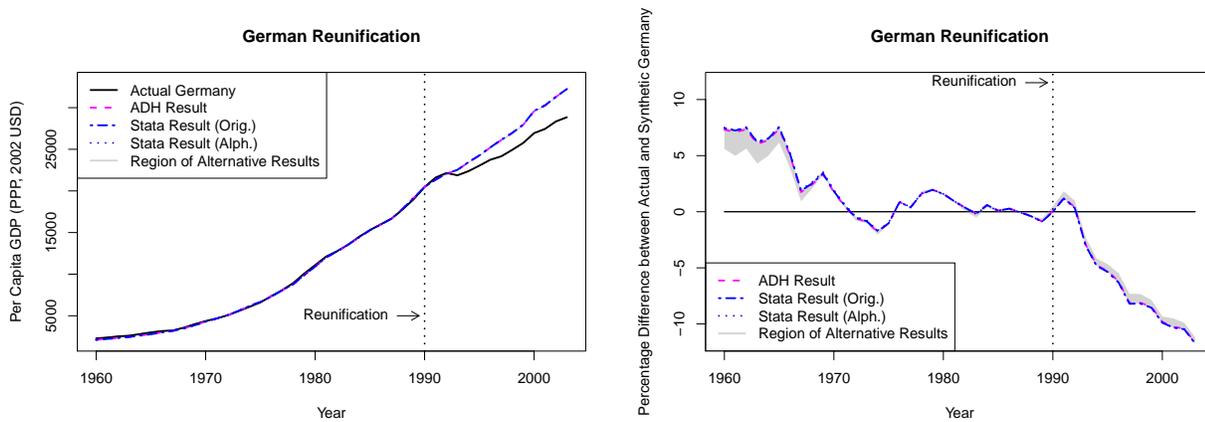


Figure 1: GDP timelines of actual Germany and various versions of synthetic Germany. Left: purchasing power parity (PPP)-adjusted GDP in 2002 U.S. dollars. Right: percentage differences between actual Germany and synthetic Germany. 'ADH' denotes the results obtained by Abadie et al. (2015), 'Stata Result (Orig.)' and 'Stata Result (Alph.)' stand for the results we obtained using Stata, using the same ordering of the donors as did Abadie et al. (2015) ('Orig.'), and alphabetical ordering ('Alph.'), respectively.

Figure 1 (the left part of which corresponds to Figure 3 of Abadie et al. (2015)) shows the timelines of GDP for actual Germany as well as several versions of synthetic Germany. From the original timelines, differences between the various versions of synthetic Germany are barely visible. This is due to the fact that German GDP per capita rose from roughly 3,000 U.S. dollars in 1960 to almost 30,000 U.S. dollars in 2003, necessitating a large scale for the plot which renders small differences between timelines essentially invisible. We therefore accompany these timelines by what might be called a 'relative gap plot', namely the percentage difference between actual and synthetic Germany. The grey area, which displays the range of ambiguity due to different but equivalent results, now becomes visible, but overall it is quite small, indicating that although donor weights are ambiguous, the conclusion with respect to a gap in German GDP after the reunification remains valid.

We now turn our attention to the in-space placebo study which artificially reassigns the reunification to all donor countries, thus treating Germany as a donor country while at the same time one of the donor countries takes the role of the treated unit. To evaluate the results, one calculates the ratios of post-treatment differences between actual and synthetic GDP

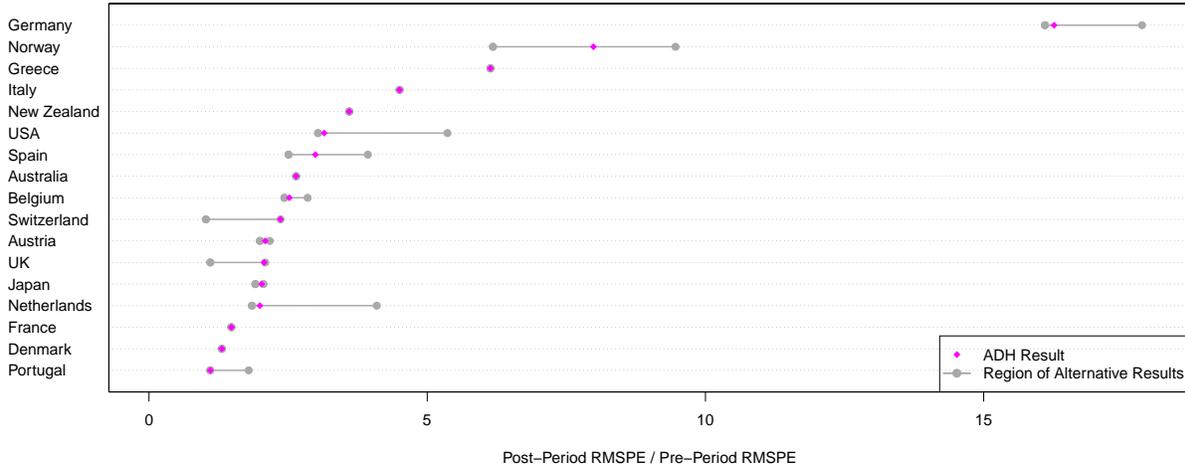


Figure 2: Ratios of post-treatment over pre-treatment root mean square prediction error (RMSPE) for in-space placebos. 'ADH' denotes the results obtained by Abadie et al. (2015).

values over corresponding pre-treatment differences. The results are displayed in Figure 2 which, as a special case, contains Figure 5 of Abadie et al. (2015). We find rather large ranges of ratios for some countries (Germany, Norway, the U.S., Spain, Switzerland, the UK, the Netherlands), and small to (almost) no ranges of ratios for other countries.¹³ In line with our heuristic rule of thumb (in detail in the appendix), the countries with large ranges are characterized by rather small numbers of donor countries contributing in the training period: Switzerland (one donor country), the U.S., Portugal, Spain (two donor countries), the UK, the Netherlands, Japan (three donor countries), Norway (four donor countries). Overall, notwithstanding the significant ambiguity of these ratios, the ratio for Germany is by far the largest, indicating that the reunification had a significant impact on German GDP per capita.

Figure 3 (the left part of which corresponds to Figure 6 of Abadie et al. (2015)) shows the results for the case when the U.S. data is removed from the sample — a so-called 'leave-one-out' analysis which, in the original Abadie et al. (2015) study, backs up the main finding. However, the grey area indicating the range of ambiguity now is very large and the original result found by Abadie et al. (2015) is extreme under all equivalent results: all other possible

¹³Note that ranges might be underestimated as these were calculated from the extensive, yet limited Monte Carlo study that we conducted.

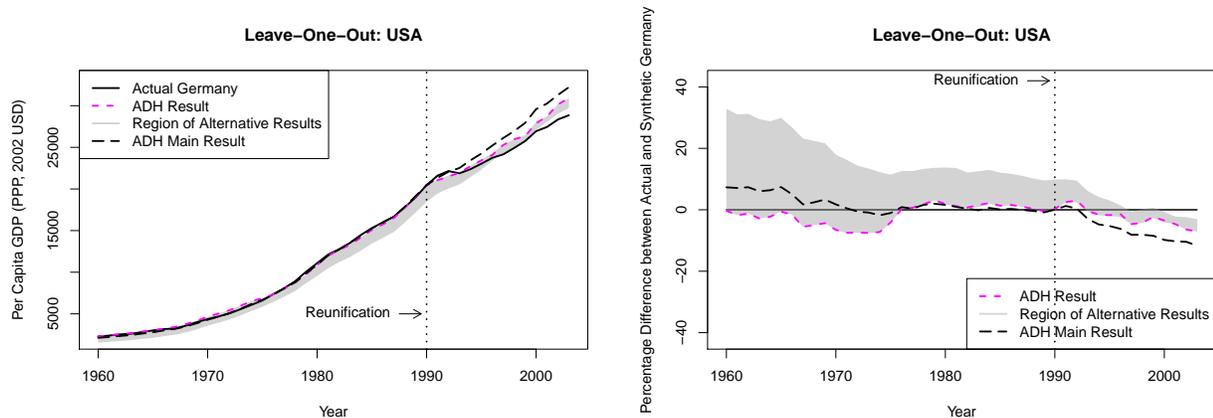


Figure 3: GDP timelines of actual Germany and various versions of synthetic Germany, leaving out U.S. data. Left: purchasing power parity (PPP)-adjusted GDP in 2002 U.S. dollars. Right: percentage differences between actual Germany and synthetic Germany. 'ADH result' denotes the corresponding result obtained by Abadie et al. (2015), 'ADH Main Result' their result obtained for complete data.

results show smaller post-treatment gaps between actual and synthetic Germany's GDP per capita, raising the question whether the gap in GDP due to the reunification crucially hinges on the U.S. acting as a donor country synthesizing Germany. The relative gap plot of Figure 3 also shows that the gap of approximately 7% in 2003 is not larger than the pre-treatment approximation error of about 7% in the early 1970's, strengthening the doubts whether there is still a significant gap in German GDP per capita after the reunification when the U.S. data is removed. Correspondingly, the ratio of post-treatment over pre-treatment differences is only 4.64, quite a small value as compared to the ratios of the in-space placebo study. Therefore, in contrast to what Abadie et al. (2015) find, it seems that including the U.S. data is essential for obtaining a significant gap in GDP between actual and synthetic Germany.

Overall, due to the not uniquely defined predictor weights, using the cross-validation technique leads to considerable ambiguity with respect to several robustness studies as conducted by Abadie et al. (2015).¹⁴ Therefore, using this technique cannot be recommended.

¹⁴Results for other robustness studies left out here to save space are available upon request.

4.2 Results Using Standard SCM

As an unflawed alternative to applying the cross-validation technique, we will now use the standard synthetic control method to analyze the reunification’s effect on West Germany’s GDP. Applying this method, we find Germany to be synthesized by Austria (62.6%), the United States (16.0%), Switzerland (13.1%), Japan (6.8%), and the Netherlands (1.5%). The corresponding timelines for GDP per capita are displayed in Figure 4, the results for the in-space placebo study can be found in the left part of Figure 5. These results are very similar to those obtained when using the cross-validation technique: after the reunification, Germany suffered from a significant loss in GDP per capita which amounted to roughly 11% in 2003.

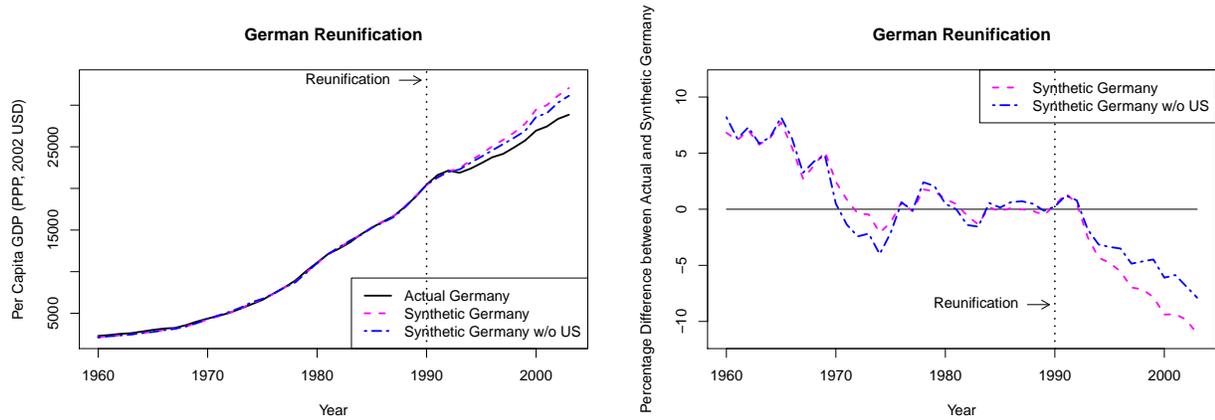


Figure 4: GDP timelines of actual Germany and two versions of synthetic Germany. Left: purchasing power parity (PPP)-adjusted GDP in 2002 U.S. dollars. Right: percentage differences between actual Germany and synthetic Germany. 'Synthetic Germany' stands for the results obtained when using standard SCM, 'Synthetic Germany w/o US' is the corresponding result when the U.S. data is removed from the sample.

Figure 4 as well as the right part of Figure 5 show the results after removing the U.S. data from the sample. The gap between actual and synthetic Germany reduces to approximately 8%, and the ratio of post-treatment differences to pre-treatment differences shrinks from 14.9 to 8.2, which is much smaller than the ratio for Norway (12.7). Therefore, the U.S. data is essential for detecting a significant gap in German GDP per capita caused by the reunification.

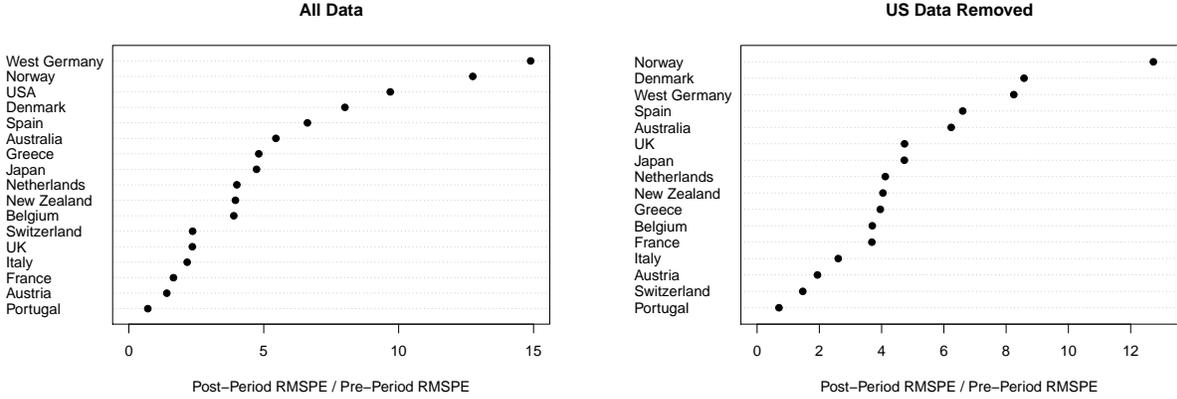


Figure 5: Ratios of post-treatment over pre-treatment root mean square prediction error (RMSPE) for in-space placebos, using standard SCM. Left: results when using all data. Right: results after removing U.S. data from the sample.

5 Conclusion

The synthetic control method is an important tool in policy evaluation which has been expanded by Abadie et al. (2015), who introduce the cross-validation technique for selecting predictor weights. In this paper, we have shown that this technique is flawed because it hinges on predictor weights which in many applications will not be uniquely defined. When using synthetic control methods in combination with cross-validation, one might therefore arrive at ambiguous results and conclusions.

As far as theory is concerned, we derive a heuristic rule of thumb which relates non-uniqueness of the predictor weights to the difference between the number of predictors and the number of donor units that synthesize the unit of interest in the training period. If this difference is positive, which is the case in most applications, predictor weights based on cross-validation are typically not uniquely defined and the ambiguity with respect to this non-uniqueness usually becomes larger the more this difference increases.

Empirically, examining the German reunification using the data of Abadie et al. (2015), we find that the amount of ambiguity is rather small as far as the main application is concerned. With respect to several robustness studies, however, the ambiguity implied by the predictors' non-uniqueness is significant, in particular for the leave-one-out and in-space placebo studies.

The failure of synthetic control methods with cross-validation is no failure of synthetic control methods as such. One can simply stick to the standard synthetic control method without cross-validation since it does not contain a second estimation step for which the predictor weights' uniqueness is crucial. When doing so for the example of the German reunification, we mostly confirm the results of Abadie et al. (2015) — there is a significant gap in German GDP due to the reunification. With respect to robustness, however, we find, in contrast to Abadie et al. (2015), that this result crucially depends on the U.S. data being included in the estimation. After removing the U.S. data from the sample, the estimated gap in GDP after the German reunification becomes much smaller and is no longer significant according to the in-space placebo study.

References

- Abadie, A., Diamond, A., and Hainmueller, J. (2010). Synthetic control methods for comparative case studies: Estimating the effect of california's tobacco control program. *Journal of the American Statistical Association*, 105(490):493–505.
- Abadie, A., Diamond, A., and Hainmueller, J. (2011). Synth: An R package for synthetic control methods in comparative case studies. *Journal of Statistical Software*, 42(13):1–17.
- Abadie, A., Diamond, A., and Hainmueller, J. (2015). Comparative politics and the synthetic control method. *American Journal of Political Science*, 59(2):495–510.
- Abadie, A. and Gardeazabal, J. (2003). The economic costs of conflict: A case study of the basque country. *The American Economic Review*, 93(1):113–132.
- Acemoglu, D., Johnson, S., Kermani, A., Kwak, J., and Mitton, T. (2016). The value of connections in turbulent times: Evidence from the united states. *Journal of Financial Economics*, 121(2):368 – 391.
- Cavallo, E., Galiani, S., Noy, I., and Pantano, J. (2013). Catastrophic Natural Disasters and Economic Growth. *The Review of Economics and Statistics*, 95(5):1549–1561.
- Gardeazabal, J. and Vega-Bayo, A. (2016). An empirical comparison between the synthetic control method and Hsiao et al.'s panel data approach to program evaluation. *Journal of Applied Econometrics*, pages n/a–n/a.
- Gobillon, L. and Magnac, T. (2016). Regional policy evaluation: Interactive fixed effects and synthetic controls. *The Review of Economics and Statistics*, 98(3):535–551.

Kleven, H. J., Landais, C., and Saez, E. (2013). Taxation and international migration of superstars: Evidence from the european football market. *The American Economic Review*, 103(5):1892–1924.

R Core Team (2014). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.

A Theory on Predictor Weights by Cross-Validation

Let V^* be a solution to Equation (2) and $W_{(\text{train})}^*(V^*)$ be the corresponding minimizer of Equation (1). We denote by $\mathcal{V} := \{V : W_{(\text{train})}^*(V) = W_{(\text{train})}^*(V^*)\}$ the set of all predictor weights V that lead to the same 'training' weights as V^* . Obviously, \mathcal{V} is closed under multiplication with positive constants, allowing to assume without loss of generality that predictor weights are always scaled such that the components of V sum to unity.¹⁵ The cross-validation technique is typically not well-defined if predictor weights $\tilde{V} \in \mathcal{V}$ exist that are different from V^* after scaling. Then, \tilde{V} is also an optimizer of the out-of-sample error, but leading to 'main' weights $W_{(\text{main})}^*(\tilde{V})$ which typically do not coincide with the corresponding 'main' weights belonging to V^* :¹⁶ $W_{(\text{main})}^*(\tilde{V}) \neq W_{(\text{main})}^*(V^*)$. Thus, well-definedness of the cross-validation technique crucially hinges on \mathcal{V} being a singleton after scaling. Furthermore, the larger \mathcal{V} , the more different weights for synthesizing, $W_{(\text{main})}^*(\tilde{V})$ for $\tilde{V} \in \mathcal{V}$, will typically exist, and the larger the amount of ambiguity of the cross-validation approach usually will be.

To develop a rule of thumb which sheds some light on how large \mathcal{V} and thus the resulting ambiguity of the cross-validation technique are, we state the following Lemma.

Lemma 1. *For any given predictor weights V , an optimizer W^* of Equation (1) must fulfil the following conditions:¹⁷*

1. for all j running through the components of W^* , with e_j denoting the j -th unit vector:

$$d_j(W^*, V) := \sum_{m=1}^k v_m \left(X_{1m}^{(\text{train})} - X_{0m}^{(\text{train})} W^* \right) X_{0m}^{(\text{train})} (W^* - e_j) \geq 0, \quad (6)$$

2. $d_j(W^*, V) = 0$ for all j with $W_j^* > 0$.

¹⁵See, e.g., (Abadie and Gardeazabal, 2003, p. 128), (Abadie et al., 2015, Footnote 5, p. 497f).

¹⁶Theoretically, it is possible that $W_{(\text{main})}^*(\tilde{V})$ and $W_{(\text{main})}^*(V^*)$ coincide. However, that would be quite a coincidence.

¹⁷One may prove that these conditions are not only necessary, but also sufficient for W^* being a minimizer of Equation (1).

Proof. For every j , consider $f_j(\delta) := \sum_{m=1}^k v_m \left(X_{1m}^{(\text{train})} - X_{0m}^{(\text{train})} ((1 - \delta)W^* + \delta e_j) \right)^2$. The derivative of f_j at $\delta = 0$, $f'_j(0) = 2 \sum_{m=1}^k v_m \left(X_{1m}^{(\text{train})} - X_{0m}^{(\text{train})} W^* \right) X_{0m}^{(\text{train})} (W^* - e_j) = 2 d_j(W^*, V)$, must be non-negative, as otherwise the convex combination $(1 - \delta)W^* + \delta e_j$ would for small positive δ yield a smaller value in (1) than does W^* . For j with $W_j^* > 0$, the vector $(1 - \delta)W^* + \delta e_j$ will have non-negative components summing to unity even for negative δ that are small enough in absolute value. Therefore, $f'_j(0)$ must vanish in that case, as otherwise $(1 - \delta)W^* + \delta e_j$ for small negative δ would yield a smaller value than W^* in (1). \square

Fixing $W^* := W_{(\text{train})}^*(V^*)$, Lemma 1 states the conditions V must fulfil to belong to \mathcal{V} : $d_j(W^*, V) \geq 0$ for all j with $W_j^* = 0$, and $d_j(W^*, V) = 0$ for all j with $W_j^* > 0$. As $d_j(W^*, V)$ is a linear function of v_1, \dots, v_k , the conditions for V to belong to \mathcal{V} thus consist of linear equations and inequalities: for the k unknown quantities v_1, \dots, v_k , we have α linear equations and $J - \alpha$ linear inequalities, with J denoting the number of donor units, and $\alpha := \{j : W_j^* > 0\}$, the number of donor units which obtain positive weights in the 'training' period. As a rule of thumb, we thus have the following:¹⁸

Rule of Thumb 1. *The cross-validation method is typically not well-defined if the difference $k - \alpha$ between the number of economic predictors (k) and the number of donor units with positive W weight in the 'training' period (α) is positive. The larger the difference $k - \alpha$, the larger is typically the ambiguity induced by the cross-validation technique.*

Finally, notice that \mathcal{V} is a convex set, as the conditions in Lemma 1 are linear in the V weights. In particular, this entails that as soon as \mathcal{V} contains predictor weights \tilde{V} different from V^* after scaling, \mathcal{V} contains infinitely many elements due to its convexity, with its dimension typically increasing with $k - \alpha$.

¹⁸It might be possible to strengthen the results of the rule of thumb to obtain a rigorous mathematical statement. This, however, is beyond the scope of this paper.