

Article prepared for the *Journal of Economic Literature*

Using Synthetic Controls: Feasibility, Data Requirements, and Methodological Aspects

Alberto Abadie
MIT

June 2020

Abstract

Probably because of their interpretability and transparent nature, synthetic controls have become widely applied in empirical research in economics and the social sciences. This article aims to provide practical guidance to researchers employing synthetic control methods. The article starts with an overview and an introduction to synthetic control estimation. The main sections discuss the advantages of the synthetic control framework as a research design, and describe the settings where synthetic controls provide reliable estimates and those where they may fail. The article closes with a discussion of recent extensions, related methods, and avenues for future research. (*JEL* C21, C23)

1. Introduction

Synthetic control methods (Abadie and Gardeazabal, 2003; Abadie et al., 2010) have become widely applied in empirical research in economics and other disciplines. Under appropriate conditions, synthetic controls provide substantial advantages as a research design method in the social sciences. These advantages, I believe, explain the increasing popularity of synthetic control methods in empirical research. At the same time, the validity of synthetic control estimators depends on important practical requirements. Perfunctory applications that ignore the context of the empirical investigation and the characteristics of the data may miss the mark, producing

Alberto Abadie, Department of Economics, MIT, abadie@mit.edu. I would like to thank Matias Cattaneo, Brigham Frandsen, Bruno Ferman, Javier Gardeazabal, Jens Hainmueller, Jérémy L'Hour, Guido Imbens, Devavrat Shah, Kaspar Wüthrich, participants at the Conference on Synthetic Controls and Related Methods at MIT/IDSS, NeurIPS 2019, the Chamberlain Online Seminar, Steven Durlauf (the editor), and five referees for insightful comments, and Kevin Li for expert research assistance. NSF support through grant SES-1756692 is gratefully acknowledged. Some fragments of this article build on an internal document previously commissioned by the Inter-American Development Bank.

misleading estimates.

My goal with this article is to provide guidance for empirical practice to researchers employing synthetic control methods. With this goal in mind, I put special emphasis on feasibility, data requirements, contextual requirements, and methodological issues related to the empirical application of synthetic controls. Particularly important is characterizing the practical settings where synthetic controls may be useful and those where they may fail.

Section 2 briefly introduces the ideas behind the synthetic control methodology in the context of comparative case studies. Section 3 discusses some of the formal aspects of the synthetic control methodology that are of particular interest for empirical applications. Readers who are already familiar with the synthetic control methodology may only need to read Sections 3.3 to 3.5 in detail, and skim through Section 2 and the rest of Section 3 in order to acquaint themselves with terms and notation that will be employed in later sections. Sections 4 through 6 comprise the core of the article. Section 4 discusses the practical advantages of synthetic control estimators. Sections 5 and 6 discuss contextual and data requirements for synthetic control empirical studies. I discuss the validity of these requirements in applied settings and potential ways to adapt the research design when the requirements do not hold in practice. Section 7 describes robustness and diagnostic checks to evaluate the credibility of a synthetic control counterfactual and to measure the extent to which results are sensitive to changes in the study design. Section 8 discusses extensions and recent proposals. The final section contains conclusions and describes open areas for research on synthetic controls.

2. A Primer on Synthetic Control Estimators

In a recent *Journal of Economic Perspectives* survey on the econometrics of policy evaluation, Susan Athey and Guido Imbens describe synthetic controls as “arguably the most important innovation in the policy evaluation literature in the last 15 years” (Athey and Imbens, 2017). In the last few years, synthetic controls have been applied to study the effects of right-to-carry laws (Donohue et al., 2019), legalized prostitution (Cunningham and Shah, 2018), immigration policy (Bohn et al., 2014), corporate political connections (Acemoglu et al., 2016), taxation (Kleven et al., 2013), organized crime (Pinotti, 2015) and many other key policy issues. They

have also been adopted as the main tool for data analysis across different sides of the issues in recent prominent debates on the effects of immigration (Borjas, 2017; Peri and Yasenov, 2019) and minimum wages (Allegretto et al., 2017; Jardim et al., 2017; Neumark and Wascher, 2017; Reich et al., 2017). Synthetic controls are also applied outside economics: in the social sciences, biomedical disciplines, engineering, etc. (see, e.g., Heersink et al., 2017; Pieters et al., 2017). Outside academia, synthetic controls have found considerable coverage in the popular press (see, e.g., Guo, 2015; Douglas, 2018) and have been widely adopted by multilateral organizations, think tanks, business analytics units, governmental agencies, and consulting firms. For example, the synthetic control method plays a prominent role in the official evaluation of the effects of the massive Bill & Melinda Gates Foundation’s *Intensive Partnerships for Effective Teaching* program (Gutierrez et al., 2016).

Synthetic control methods were originally proposed in Abadie and Gardeazabal (2003) and Abadie et al. (2010) with the aim to estimate the effects of aggregate interventions, that is, interventions that are implemented at an aggregate level affecting a small number of large units (such as a cities, regions, or countries), on some aggregate outcome of interest. More recently, synthetic control methods have been applied to settings with a large number of units.¹ We will discuss this and other extensions in Section 8.

Consider a setting where one aggregate unit, such as a state, or a school district, is exposed to an event or intervention of interest. For example, Abadie et al. (2010) study the effect of a large tobacco-control program adopted in California in 1988; Bifulco et al. (2017) evaluate the effects of an educational program adopted in the Syracuse, NY, school district in 2008. In accordance with the program evaluation literature in economics, the terms “treated” and “untreated” will refer to units exposed and not exposed to the event or intervention of interest, respectively. I will use the terms “event”, “intervention”, and “treatment” interchangeably. Traditional regression analysis techniques require large samples and many observed instances of the event or intervention of interest and, as a result, they are often ill-suited to estimate the effects of infrequent events, such as policy interventions, on aggregate units. Economists have approached the estimation of

¹See, e.g., Acemoglu et al. (2016), Kreif et al. (2016), Abadie and L’Hour (2019), and Dube and Zipperer (2015).

the effects of large scale but infrequent interventions using time-series analysis and comparative case studies. Single-unit time-series analysis is an effective tool to study the short-term effects of policy interventions in cases when we expect short-term effects to be of a substantial magnitude.² However, the use of time series techniques to estimate medium and long-term effects of policy intervention is complicated by the presence of shocks to the outcome of interest, aside from the effect of the intervention. Comparative case studies are based on the idea that the effect of an intervention can be inferred by comparing the evolution of the outcome variables of interest between the unit exposed to treatment and a group of units that are similar to the exposed unit but were not affected by the treatment. This can be achieved whenever the evolution of the outcomes for the unit affected by the intervention and the comparison units is driven by common factors that induce a substantial amount of co-movement.

Comparative case studies have long been applied to the evaluation of large-scale events or aggregate interventions. For example, to estimate the effects of the massive arrival of Cuban expatriates to Miami during the 1980 Mariel Boatlift on native unemployment in Miami, Card (1990) compares the evolution of native unemployment in Miami at the time of the boatlift to the average evolution of native unemployment in four other cities in the US. Similarly, Card and Krueger (1994) use Pennsylvania as a comparison to estimate the effects of an increase in the New Jersey minimum wage on employment in fast food restaurants in New Jersey. A drawback of comparative case studies of this type is that the selection of the comparison units is not formalized and often relies on informal statements of affinity between the units affected by the event or intervention of interest and a set of comparison units. Moreover, when the units of observation are a small number of aggregate entities, like countries or regions, no single unit alone may provide a good comparison for the unit affected by the intervention.

The synthetic control method is based on the idea that, when the units of observation are a small number of aggregate entities, a combination of unaffected units often provides a more

²The literature on “interrupted time-series” is particularly relevant in the context of policy evaluation. See, for example, Cook and Campbell (1979) which discusses the limitations of this methodology if interventions are gradual rather than abrupt and/or if the causal effect of an intervention is delayed in time. Interrupted time-series methods are closely related to regression-discontinuity design techniques (see, e.g., Thistlethwaite and Campbell, 1960).

appropriate comparison than any single unaffected unit alone. The synthetic control methodology seeks to formalize the selection of the comparison units using a data driven procedure. As we will discuss later, this formalization also opens the door to a mode of quantitative inference for comparative case studies.

3. Formal Aspects of the Synthetic Control Method

3.1. The Setting

Suppose that we obtain data for $J + 1$ units: $j = 1, 2, \dots, J + 1$. Without loss of generality, we assume that the first unit ($j = 1$) is the treated unit, that is, the unit affected by the policy intervention of interest.³ The “donor pool”, that is, the set of potential comparisons, $j = 2, \dots, J + 1$ is a collection of untreated units, not affected by the intervention. We assume also that our data span T periods and that the first T_0 periods are before the intervention. For each unit, j , and time, t , we observe the outcome of interest, Y_{jt} . For each unit, j , we also observe a set of k predictors of the outcome, X_{1j}, \dots, X_{kj} , which may include pre-intervention values of Y_{jt} and which are themselves unaffected by the intervention. The $k \times 1$ vectors $\mathbf{X}_1, \dots, \mathbf{X}_{J+1}$ contain the values of the predictors for units $j = 1, \dots, J + 1$, respectively. The $k \times J$ matrix, $\mathbf{X}_0 = [\mathbf{X}_2 \cdots \mathbf{X}_{J+1}]$, collects the values of the predictors for the J untreated units. For each unit, j , and time period, t , we will define Y_{jt}^N to be the potential response without intervention. For the unit affected by the intervention, $j = 1$, and a post-intervention period, $t > T_0$, we will define Y_{1t}^I to be the potential response under the intervention.⁴ Then, the effect of the intervention of interest for the affected unit in period t (with $t > T_0$) is:

$$\tau_{1t} = Y_{1t}^I - Y_{1t}^N. \quad (1)$$

Because unit “one” is exposed to the intervention after period T_0 , it follows that for $t > T_0$ we have $Y_{1t} = Y_{1t}^I$. Simply put, for the unit affected by the intervention and a post-intervention

³The synthetic control framework can easily accommodate estimation with multiple treated units by fitting separate synthetic controls for each of the treated units. In practice, however, estimation with several treated units may carry some practical complications that are discussed in Section 8.

⁴ Y_{1t}^I and Y_{jt}^N are the potential outcomes of Rubin’s Model for Causal Inference (see, e.g., Rubin, 1974; Holland, 1986). To simplify notation, I exclude the start time of the intervention from the notation for Y_{1t}^I . Notice, however, that the value of Y_{1t}^I depends in general not only on when the intervention starts, but also other features of the intervention that are fixed in our analysis and, therefore, excluded from the notation.

period we observe the potential outcome under the intervention. The great policy evaluation challenge is to estimate Y_{1t}^N for $t > T_0$: how the outcome of interest would have evolved for the affected unit *in the absence of the intervention*. This is a counterfactual outcome, as the affected unit was, by definition, exposed to the intervention of interest after $t = T_0$. As equation (1) makes clear, given that Y_{1t}^I is observed, the problem of estimating the effect of a policy intervention is equivalent to the problem of estimating Y_{1t}^N . Notice also that equation (1) allows the effect of the intervention to change over time. This is crucial because intervention effects may not be instantaneous and may accumulate or dissipate as time after the intervention passes.

3.2. Estimation

Comparative case studies aim to reproduce Y_{1t}^N —that is, the value of the outcome variable that would have been observed for the affected unit in the absence of the intervention—using one unaffected unit or a small number of unaffected units that had similar characteristics as the affected unit at the time of the intervention. When the data consist of a few aggregate entities, such as regions or countries, it is often difficult to find a single unaffected unit that provides a suitable comparison for the unit affected by the policy intervention of interest. As mentioned above, the synthetic control method is based on the observation that a combination of units in the donor pool may approximate the characteristics of the affected unit substantially better than any unaffected unit alone. A synthetic control is defined as a weighted average of the units in the donor pool. Formally, a synthetic control can be represented by a $J \times 1$ vector of weights, $\mathbf{W} = (w_2, \dots, w_{J+1})'$. Given a set of weights, \mathbf{W} , the synthetic control estimators of Y_{1t}^N and τ_{1t} are, respectively:

$$\hat{Y}_{1t}^N = \sum_{j=2}^{J+1} w_j Y_{jt}, \quad (2)$$

and

$$\hat{\tau}_{1t} = Y_{1t} - \hat{Y}_{1t}^N. \quad (3)$$

To avoid extrapolation, the weights are restricted to be non-negative and to sum to one, so synthetic controls are weighted averages of the units in the donor pool. The requirement that weights should be non-negative and no greater than one can be relaxed at the cost of allowing

extrapolation. For example, Abadie et al. (2015) show that, in the context of estimating the effect of a policy intervention, there is a regression estimator that can be represented as a synthetic control with weights that are unrestricted except for that the sum of the weights is equal to one. By not restricting the weights to be in $[0, 1]$, regression allows extrapolation.⁵ Restricting synthetic control weights to be non-negative and sum to one generates synthetic controls that are weighted averages of the outcomes of units in the donor pool, with weights that are typically sparse (see Section 4). That is, only a small number of units in the donor pool contribute to the estimate of the counterfactual of interest, \widehat{Y}_{1t}^N , and the contribution of each unit is represented by its synthetic control weight. Because synthetic control weights define a weighted average and because they are sparse, the specific nature of a synthetic control counterfactual estimate is particularly transparent, relative to competing methods. Notice also that considering synthetic controls with weights that sum to one may be warranted only if the variables in the data are rescaled to correct for differences in size between units (e.g., per-capita income) or if such correction is not needed because the variables in the data do not scale with size (e.g., prices).

As an example, a synthetic control that assigns equal weights, $w_j = 1/J$, to each of the units in the control group results in the following estimator for τ_{1t} :

$$\widehat{\tau}_{1t} = Y_{1t} - \frac{1}{J} \sum_{j=2}^{J+1} Y_{jt}. \quad (4)$$

In this case, the synthetic control is the simple average of all the units in the donor pool. A population-weighted version is

$$\widehat{\tau}_{1t} = Y_{1t} - \sum_{j=2}^{J+1} w_j^{pop} Y_{jt}, \quad (5)$$

where w_j^{pop} is the population in unit j (e.g., at the time of the intervention) as a fraction of the total population in the donor pool. If, however, a single unit, m , in the donor pool is used as a comparison, then $w_m = 1$, $w_j = 0$ for $j \neq m$, and

$$\widehat{\tau}_{1t} = Y_{1t} - Y_{mt}. \quad (6)$$

⁵See Section 4 for details. Doudchenko and Imbens (2016), Ferman (2019) and Li (2019) discuss the role of weight restrictions as regularization devices. Doudchenko and Imbens (2016) and Chernozhukov et al. (2019a) propose alternative regularization procedures for synthetic controls based on the elastic net and the lasso, respectively.

For nearest-neighbor estimators, m is the index value that minimizes $\|\mathbf{X}_1 - \mathbf{X}_j\|$ over j for some norm $\|\cdot\|$.

Expressing the comparison unit as a synthetic control motivates the question of how the weights, w_2, \dots, w_{J+1} , should be chosen in practice. Abadie and Gardeazabal (2003) and Abadie et al. (2010) propose to choose w_2, \dots, w_{J+1} so that the resulting synthetic control best resembles the pre-intervention values for the treated unit of predictors of the outcome variable. That is, given a set of non-negative constants, v_1, \dots, v_k , Abadie and Gardeazabal (2003) and Abadie et al. (2010) propose to choose the synthetic control, $\mathbf{W}^* = (w_2^*, \dots, w_{J+1}^*)'$ that minimizes

$$\|\mathbf{X}_1 - \mathbf{X}_0 \mathbf{W}\| = \left(\sum_{h=1}^k v_h (X_{h1} - w_2 X_{h2} - \dots - w_{J+1} X_{hJ+1})^2 \right)^{1/2} \quad (7)$$

subject to the restriction that w_2, \dots, w_{J+1} are non-negative and sum to one.⁶ Then, the estimated treatment effect for the treated unit at time $t = T_0 + 1, \dots, T$ is

$$\hat{\tau}_{1t} = Y_{1t} - \sum_{j=2}^{J+1} w_j^* Y_{jt}. \quad (8)$$

The positive constants v_1, \dots, v_k in (7) reflect the relative importance of the synthetic control reproducing the values of each of the k predictors for the treated unit, X_{11}, \dots, X_{k1} . For a given set of weights, v_1, \dots, v_k , minimizing equation (7) can be easily accomplished using constrained quadratic optimization. That is, each potential choice of $\mathbf{V} = (v_1, \dots, v_k)$ produces a synthetic control, $\mathbf{W}(\mathbf{V}) = (w_2(\mathbf{V}), \dots, w_{J+1}(\mathbf{V}))'$, which can be determined by minimizing equation (7), subject to the restriction that the weights in $\mathbf{W}(\mathbf{V})$ are positive and sum to one.

Of course, a question remains about how to choose \mathbf{V} . A simple selector of v_h is the inverse of the variance of X_{h1}, \dots, X_{hJ+1} , which in effect rescales all rows of $[\mathbf{X}_1 : \mathbf{X}_0]$ to have unit variance. Alternatively, Abadie and Gardeazabal (2003) and Abadie et al. (2010) choose \mathbf{V} , such that the synthetic control $\mathbf{W}(\mathbf{V})$ minimizes the mean squared prediction error (MSPE) of

⁶For the sake of expositional simplicity, I discuss only the normalized Euclidean norm in equation (7). Of course, other norms are possible. Also, to avoid notational clutter, dependence of the norm in equation (7) from the weights v_1, \dots, v_k is left implicit in the notation.

this synthetic control with respect to Y_{1t}^N

$$\sum_{t \in \mathcal{T}_0} \left(Y_{1t} - w_2(\mathbf{V})Y_{2t} - \dots - w_{J+1}(\mathbf{V})Y_{J+1t} \right)^2,$$

for some set $\mathcal{T}_0 \subseteq \{1, 2, \dots, T_0\}$ of pre-intervention periods. Abadie et al. (2015) propose a related method to choose v_1, \dots, v_k via out-of-sample validation. The ideas behind out-of-sample validation selection of v_1, \dots, v_k are described next. The goal of the synthetic control is to approximate the trajectory that would have been observed for Y_{1t} and $t > T_0$ in the absence of the intervention. For that purpose, the synthetic control method selects a set of weights \mathbf{W} such that the resulting synthetic control resembles the affected unit before the intervention along the values of the variables X_{11}, \dots, X_{k1} . The question of choosing $\mathbf{V} = (v_1, \dots, v_k)$ boils down to assessing the relative importance of each of X_{11}, \dots, X_{k1} as a predictor of Y_{1t}^N . That is, the value v_h aims to reflect the relative importance of approximating the value of X_{h1} for predicting Y_{1t}^N in the post-intervention period, $t = T_0 + 1, \dots, T$. Because Y_{1t}^N is not observed for $t = T_0 + 1, \dots, T$, we cannot directly evaluate the relative importance of fitting each predictor to approximate Y_{1t}^N in the post-intervention period. However, Y_{1t}^N is observed for the pre-intervention periods $t = 1, 2, \dots, T_0$, so it is possible to use pre-intervention data to assess the predictive power on Y_{1t}^N of the variables X_{1j}, \dots, X_{kj} . This can be accomplished in the following manner.

1. Divide the pre-intervention periods into a initial *training* period and a subsequent *validation* period. For simplicity and concreteness, we will assume that T_0 is even and the training and validation periods span $t = 1, \dots, t_0$ and $t = t_0 + 1, \dots, T_0$, respectively, with $t_0 = T_0/2$. In practice, the lengths of the training and validation periods may depend on application-specific factors, such as the extent of data availability on outcomes in the pre-intervention and post-intervention periods, and the specific times when the predictors are measured in the data.
2. For every value \mathbf{V} , let $\tilde{w}_2(\mathbf{V}), \dots, \tilde{w}_{J+1}(\mathbf{V})$ be the synthetic control weights computed with training period data on the predictors. The mean squared prediction error of this synthetic

control with respect to Y_{1t}^N in the validation period is:

$$\sum_{t=t_0+1}^{T_0} \left(Y_{1t} - \tilde{w}_2(\mathbf{V})Y_{2t} - \dots - \tilde{w}_{J+1}(\mathbf{V})Y_{J+1t} \right)^2, \quad (9)$$

3. Minimize the mean squared prediction error in the previous equation with respect to \mathbf{V} .
4. Use the resulting \mathbf{V}^* and data on the predictors for the last t_0 periods before in the intervention, $t = T_0 - t_0 + 1, \dots, T_0$, to calculate $\mathbf{W}^* = \mathbf{W}(\mathbf{V}^*)$.⁷

To give sharpness to the discussion of the properties and practical implementation of synthetic control estimators, I will refer as a running example to an application in Abadie et al. (2015), which estimates the effect of the 1990 German reunification on per-capita GDP in West Germany. In this application, the intervention is the 1990 German reunification and the treated unit is the former West Germany. The donor pool consists a set of industrialized countries, and \mathbf{X}_1 and \mathbf{X}_0 collect pre-reunification values of predictors of economic growth. Figure 1(a) compares the trajectory of per-capita GDP before and after the reunification for West Germany and a simple average of the countries in the donor pool, for the years 1960-2003. This is the comparison in equation (4). Average per-capita GDP among the countries in the donor pool fails to reproduce the trajectory of per-capita GDP for West Germany even before the reunification takes place in 1990. Moreover, the restriction of parallel trends required for difference-in-differences models (see, e.g., Abadie, 2005; Angrist and Pischke, 2008) fails to hold in the pre-intervention data. Figure 1(b) reports the trajectory of per-capita GDP for West Germany and for a synthetic control calculated in the manner explained in this section. This figure shows that a weighted average of the countries in the donor pool is able to closely approximate the trajectory of per-capita GDP for West Germany before the German reunification.

Moreover, the synthetic control of Figure 1(b) closely reproduces the pre-reunification values of economic growth predictors for West Germany. In columns (1) to (4), Table 1 reports the value of economic growth predictors for West Germany, \mathbf{X}_1 , for the synthetic control, $\mathbf{X}_0\mathbf{W}^*$, for the

⁷As discussed in Klößner et al. (2018), cross-validation weights are not always unique. That is, minimization of equation (9) may not have a unique solution. In principle, this could be dealt with via penalization (e.g., adding a term $\gamma \sum_{h=1}^k v_m^2$ for some $\gamma > 0$ to equation (9), which favors dense sets of weights). In practice, however, researchers should aim to demonstrate that their results are not overly sensitive to particular choices of \mathbf{V} .

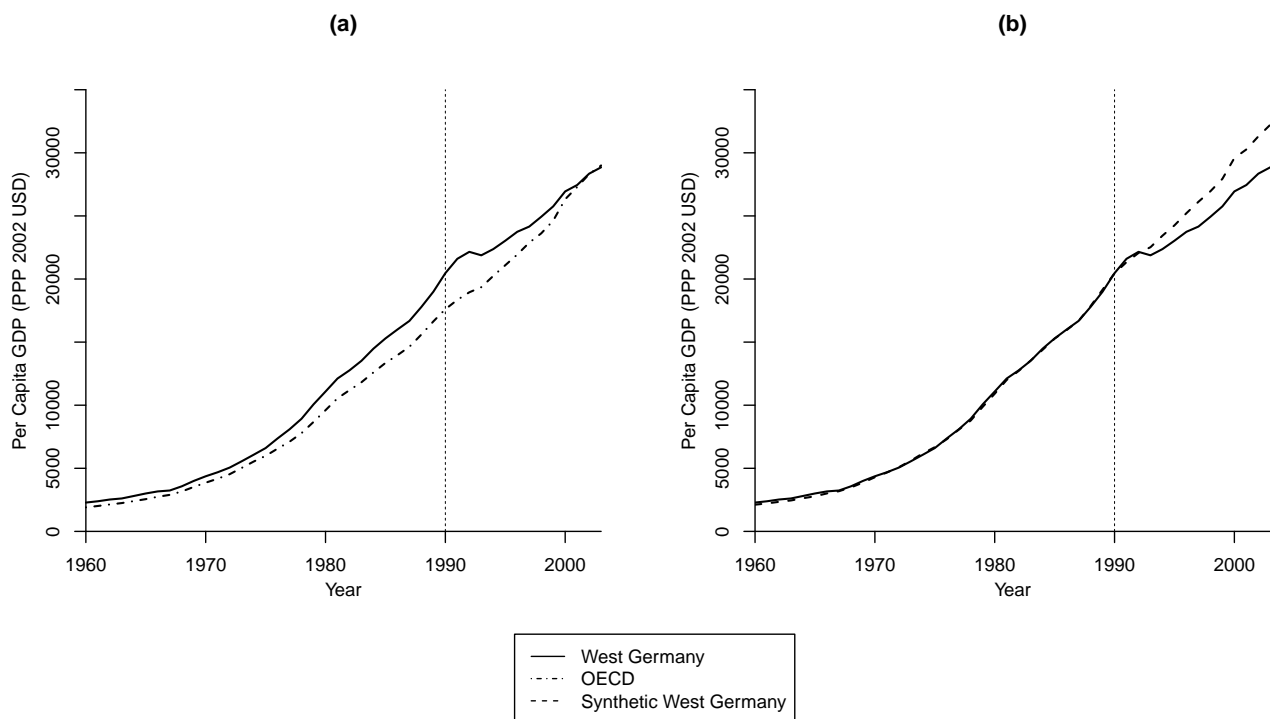


Figure 1: Synthetic control estimation in the German reunification example

Panel (a) compares the evolution of per capita GDP in West Germany to the evolution of per capita GDP for a simple average of OECD countries. In panel (b) the comparison is with a synthetic control calculated in the manner explained in Section 3.2. See Abadie et al. (2015) for details.

simple average of the units in the donor pool as in equation (4), and for the single comparison estimator of equation (6), where m is the index of the nearest neighbor in terms of the values of the predictors in \mathbf{X}_1 (see Table 1 note for details). The results in Table 1 illustrate the potential benefit in terms of the fit of the covariates from using synthetic control methods in studies with a few comparison units in the donor pool. While the simple average of the countries in the OECD sample and the nearest neighbor both fail to reproduce the economic growth predictors for West Germany prior to the reunification, a synthetic control provides a rather accurate approximation to the value of the predictors for West Germany.

	West Germany	Synthetic West Germany	OECD Average	Austria (Nearest Neighbor)
	(1)	(2)	(3)	(4)
GDP per-capita	15808.9	15802.24	13669.4	14817
Trade openness	56.8	56.9	59.8	74.6
Inflation rate	2.6	3.5	7.6	3.5
Industry share	34.5	34.5	34.0	35.5
Schooling	55.5	55.2	38.7	60.9
Investment rate	27.0	27.0	25.9	26.6

Note: First column reports \mathbf{X}_1 , second column reports $\mathbf{X}_0\mathbf{W}^*$, the third column reports a simple average of \mathbf{X}_j for the 16 OECD countries in the donor pool, and the last column reports the value of \mathbf{X}_j for the nearest neighbor of West Germany in terms of predictors values. GDP per capita, inflation rate, and trade openness are averages for 1981-1990 period. Industry share (of value added) is the average for 1981-1989. Schooling is the average for 1980 and 1985. Investment rate is averaged over 1980-1984. See Abadie et al. (2015) for variable definitions and sources. The nearest neighbor in column (4) minimizes the Euclidean norm of the pairwise differences between the values of the predictors for West Germany and for each of the countries in the donor pool, after rescaling the predictors to have unit variance.

Table 1: Economic growth predictor means before the German reunification

Table 2 relays the identities and contributions of each of the units in the donor pool to the synthetic control for West Germany. Austria carries the largest weight, with the U.S., Japan, Switzerland, and the Netherlands also contributing to the synthetic control with weights in

Australia	–
Austria	0.42
Belgium	–
Denmark	–
France	–
Greece	–
Italy	–
Japan	0.16
Netherlands	0.09
New Zealand	–
Norway	–
Portugal	–
Spain	–
Switzerland	0.11
United Kingdom	–
United States	0.22

Table 2: Synthetic control weights for West Germany

decreasing order. The rest of the countries in the donor pool do not contribute to the synthetic control for West Germany. As we will see later, the sparsity of the weights in Table 2 is typical of synthetic control estimators, and is a consequence of the geometric characteristics of the solution to the optimization problem that generates synthetic controls.

3.3. Bias bound

Abadie et al. (2010) study the bias properties of synthetic controls estimators for the cases when Y_{1t}^N is generated by (i) a linear factor model, or (ii) a vector autoregressive model.⁸ They show that, under some conditions, the synthetic control estimator is unbiased for a vector autoregressive model, and provide a bias bound for a linear factor model. Here, I will restrict the exposition to the linear factor model, which can be seen as a generalization of difference-in-differences. Consider the following linear factor model for Y_{jt}^N ,

$$Y_{jt}^N = \delta_t + \boldsymbol{\theta}_t \mathbf{Z}_j + \boldsymbol{\lambda}_t \boldsymbol{\mu}_j + \varepsilon_{jt}, \quad (10)$$

⁸Notice that the assumptions on the data generating process involve Y_{jt}^N , but not Y_{1t}^I . Since $Y_{1t}^I = Y_{1t}$ is observed, estimation of τ_{1t} for $t > T_0$ requires no assumptions on the process that generates Y_{1t}^I .

where δ_t is a time trend, \mathbf{Z}_j and $\boldsymbol{\mu}_j$ are vectors of observed and unobserved predictors of Y_{jt}^N , respectively, with coefficients $\boldsymbol{\theta}_t$ and $\boldsymbol{\lambda}_t$, and ε_{jt} are zero mean individual transitory shocks. In the time series literature in econometrics, $\boldsymbol{\theta}_t$ and $\boldsymbol{\lambda}_t$ are referred to as common factors, and \mathbf{Z}_j and $\boldsymbol{\mu}_j$ as factor loadings. δ_t is a common factor with constant loadings across units, while $\boldsymbol{\lambda}_t$ represents a set of common factors with varying loadings across units. A difference-in-differences/fixed effects panel model can be obtained from equation (10) by restricting $\boldsymbol{\lambda}_t$ to be time invariant, so $\boldsymbol{\lambda}_t = \boldsymbol{\lambda}$ (see Bai, 2009). This has the effect of restricting the mean outcomes of units with the same values for the observed predictors, $\mathbf{Z}_j = \mathbf{z}$, to follow parallel trends, $\delta_t + \boldsymbol{\theta}_t \mathbf{z} + \boldsymbol{\lambda} \boldsymbol{\mu}_j$. A linear factor model provides a useful extension to the difference-in-differences/fixed effects panel data models by allowing Y_{jt}^N to depend on multiple unobserved components, $\boldsymbol{\mu}_j$, with coefficients, $\boldsymbol{\lambda}_t$, that change in time. In contrast to difference-in-differences, the linear factor model does not impose parallel mean outcome trends for units with the same values for \mathbf{Z}_j .

Abadie et al. (2010) provide a characterization of the bias of the synthetic control estimator for the case when the synthetic control reproduces the characteristics of the treatment unit. Let \mathbf{X}_1 be the vector that includes \mathbf{Z}_1 and the pre-intervention outcomes for the treated unit, and let \mathbf{X}_0 be the matrix that collects the same variables for the untreated units. Suppose that $\mathbf{X}_1 = \mathbf{X}_0 \mathbf{W}^*$, that is, the synthetic control represented by \mathbf{W}^* is able to reproduce the characteristics of the treated unit (including the values of the pre-intervention outcomes). Then the bias of $\hat{\tau}_{it}$ is controlled by the ratio between the scale of the individual transitory shocks, ε_{it} and the number of pre-intervention periods, T_0 . The intuition behind this result is rather immediate. Under the factor model in equation (10), a synthetic control that reproduces the values \mathbf{Z}_1 and $\boldsymbol{\mu}_1$ would provide an unbiased estimator of the treatment effect for the treated. If $\mathbf{X}_1 = \mathbf{X}_0 \mathbf{W}^*$, then the synthetic control matches the value of \mathbf{Z}_1 . On the other hand, $\boldsymbol{\mu}_1$ is not observed, so it cannot be matched directly in the data. However, a synthetic control that reproduces the values of \mathbf{Z}_1 but fails to reproduce the values of $\boldsymbol{\mu}_1$ can only provide a close match for the pre-treatment outcomes if differences in the values of the individual transitory shocks between the treated and the synthetic control compensate for the differences in unobserved factor loadings. This is unlikely to happen when the scale of the transitory shocks, ε_{it} , is small or

the number of pre-treatment periods, T_0 , is large. In contrast, a small number of pre-intervention periods combined with enough variation in the unobserved transitory shocks may result in a close match for pre-treatment outcomes even if the synthetic control does not closely match the values of $\boldsymbol{\mu}_1$. This is a form of over-fitting and a potential source of bias.

In practice, the condition $\mathbf{X}_1 = \mathbf{X}_0\mathbf{W}^*$ is replaced by the approximate version $\mathbf{X}_1 \approx \mathbf{X}_0\mathbf{W}^*$. It is important to notice, however, that for any particular data-set there are not ex-ante guarantees on the size of the difference $\mathbf{X}_1 - \mathbf{X}_0\mathbf{W}^*$. When this difference is large, Abadie et al. (2010) recommend against the use of synthetic controls because of the potential for substantial biases. For the factor model in equation (10), obtaining a good fit $\mathbf{X}_1 \approx \mathbf{X}_0\mathbf{W}^*$ when \mathbf{X}_1 and \mathbf{X}_0 include pre-intervention outcomes typically requires that the variance of the transitory shock is small (see Ferman and Pinto, 2018). Moreover, because the bias bound depends inversely on T_0 , one could erroneously conclude that under the factor model in equation (10), the synthetic control estimator is unbiased as T_0 goes to infinity. However, the bias bound in Abadie et al. (2010) is derived under $\mathbf{X}_1 = \mathbf{X}_0\mathbf{W}^*$, and its practical relevance depends on the ability of the synthetic control to reproduce the trajectory of the outcome for the treated unit. Sizeable biases may persist as $T_0 \rightarrow \infty$, unless the quality of the fit, $\mathbf{X}_1 - \mathbf{X}_0\mathbf{W}^*$, is good. That is, the ability of a synthetic control to reproduce the trajectory of the outcome variable for the treated unit over an extended period of time, as in Figure 1(b), provides an indication of low bias. However, a large T_0 cannot drive down the bias if the fit is bad. In practice, synthetic controls may not perfectly fit the characteristics of the treated units. Section 7 discusses a backdating exercise that can often be used to obtain an indication of the size and direction of the bias arising from imperfect fit.

The risk of overfitting may also increase with the size of the donor pool, especially when T_0 is small. For any fixed T_0 , a larger J makes it easier to fit pre-treatment outcomes even when there are substantial discrepancies in factor loadings between the treated unit and the synthetic control. Consistent with this argument, the bias bound for $\hat{\tau}_{1t}$ derived in Abadie et al. (2010) depends positively on J . Under a factor model for Y_{it}^N , a large number of units in the donor pool may create or exacerbate the bias of the synthetic control estimator, especially if the values

of $\boldsymbol{\mu}_j$ in the donor pool greatly differ from $\boldsymbol{\mu}_1$.⁹ Moreover, the factor model in equation (10) should be interpreted only as an approximation to a more general (non-linear) process for Y_{it}^N . If the process that determines Y_{it}^N is non-linear in the attributes of the units, even a close fit by a synthetic control, which is a weighted average, could potentially result in large interpolation biases.

A practical implication of the discussion in the previous paragraph is that each of the units in the donor pool have to be chosen judiciously to provide a reasonable control for the treated unit. Including in the donor pool units that are regarded by the analyst to be unsuitable controls because of large discrepancies in the values of their observed attributes \mathbf{Z}_j or because of suspected large differences in the values of the unobserved attributes $\boldsymbol{\mu}_j$ relative to the treated unit is a recipe for bias.

There are other factors that contribute to the bias bound in Abadie et al. (2010). In particular, the value of the bound increases with the number on unobserved factors, that is, the number of components in $\boldsymbol{\mu}_j$. The dependence of the bias bound on the number of unobserved factors is relevant for the discussion on the choice of predictors for the synthetic control method in the next section.

3.4. Variable selection

A synthetic control provides a predictor of Y_{it}^N for $t > T_0$, the potential outcome without the intervention for the treated units in a post-intervention period. Like for any other prediction procedure, the choice of predictors (in \mathbf{X}_1 and \mathbf{X}_0 for synthetic control estimators) is a fundamental part of the estimation task. This section discusses variable selection in the synthetic control method. To aid the discussion of the different issues involved in variable selection for synthetic controls, I will employ the concepts and notation of the linear factor model framework of Section 3.3. Predictor variables in \mathbf{X}_1 and \mathbf{X}_0 typically include both pre-intervention values of the outcome variable as well as other predictors, \mathbf{Z}_j .

Pre-intervention values of the outcome variable, which are naturally available in panel data

⁹A large J may be beneficial in high-dimensional settings, as demonstrated in Ferman (2019), who shows that under certain conditions synthetic control estimators may asymptotically unbiased as $T_0 \rightarrow \infty$ and $J \rightarrow \infty$.

settings, play a crucial role in reproducing the unobserved factor loadings in $\boldsymbol{\mu}_j$ in the linear factor model of Section 3.3. They also arise organically as predictors for the synthetic control estimators under a vector auto-regression model for the process that generates the data (see Abadie et al., 2010). The credibility of a synthetic control estimator depends on its ability to track the trajectory of the outcome variable for the treated unit for an extended pre-intervention period. Provided that a good fit for pre-intervention outcomes is attained, the researcher has some flexibility in the way pre-intervention outcomes are incorporated in \mathbf{X}_1 and \mathbf{X}_0 . Consider the German reunification application of Section 3.2. As reported in Table 1, the set of predictors in \mathbf{X}_1 and \mathbf{X}_0 includes average per-capita GDP in 1981-1990, and no other pre-intervention outcome. Notice, however, that the resulting synthetic control is able to track the trajectory of per-capita GDP for West Germany for the entire 1960-1990 pre-intervention period. This happens because per-capita GDP figures for OECD countries strongly co-move in time across countries. This co-movement of the outcome variable of interest across the different units in the data is exactly what synthetic controls are designed to exploit. It makes it possible to match the entire trajectory of GDP per-capita for West Germany by fitting only the average level of GDP per capita in the 1981-1990 period. Given this premise, one potential advantage from using a summary measure of pre-reunification GDP per capita to calculate the synthetic control for West Germany (as opposed to, say, including all ten different annual values of GDP per capita for the 1981-1990 as predictors) resides in a higher sparsity of the resulting synthetic control. As will be discussed in Section 4 below, the number of units in the donor pool that carry positive weights in a synthetic control estimator is controlled by the number of predictors in \mathbf{X}_1 and \mathbf{X}_0 , and sparse synthetic controls (that is, synthetic controls made of a small number of comparison units) are easy to interpret and evaluate.

Part of the literature on synthetic controls emphasizes estimators that depend only on pre-intervention outcomes and ignore the information of other predictors, \mathbf{Z}_j . This reliance on pre-intervention outcomes only, while adopted in many cases for technical or expositional convenience, may create the mistaken impression that other predictors play a minor role in synthetic control estimators. Notice, however, that in equation (10), covariates excluded from \mathbf{Z}_j are mechanically

absorbed into $\boldsymbol{\mu}_j$, which increases the number of components of $\boldsymbol{\mu}_j$ and, therefore, the bound on the bias, as discussed in the previous section. By excluding \mathbf{Z}_j from the set of predictors in \mathbf{X}_1 and \mathbf{X}_0 , not only do we aim to implicitly match the values of $\boldsymbol{\mu}_j$ through their effects on the pre-intervention outcomes, but also the values of \mathbf{Z}_j .

Data driven methods for variable selection evaluate the predictive power of alternative sets of predictors. This can be done in the synthetic control method framework by considering alternative sets of variables for the procedure discussed in Section 3.2 to choose v_1, \dots, v_k . The procedure divides the pre-intervention periods into an initial training period and a subsequent validation period. Synthetic control weights are computed using data from the training period only. The validation period can then be used to evaluate the predictive power of the resulting synthetic control. This procedure can be used to select predictors or to evaluate the predictive power of a given set of predictors as in Section 7 below. Section 7 below discusses how to assess the robustness of the results to alternative sets of predictors.

Finally, it is worth noting that post-intervention outcomes are not used in the calculation of synthetic control weights. This property of synthetic control methods can be exploited to provide guarantees against the use of results to guide specification searches. This is because synthetic control weights can be calculated using pre-intervention data only in the design phase of the study, before post-intervention outcomes are observed or realized. Section 4 discusses this issue in more detail.

3.5. Inference

Abadie et al. (2010) propose a mode of inference for the synthetic control framework that is based on permutation methods. In its simpler version, the effect on the intervention is estimated separately for each of the units in the sample. Consider the case with a single treated unit, as in Section 3.1. A permutation distribution can be obtained by iteratively reassigning the treatment to the units in the donor pool and estimating “placebo effects” in each iteration. Then, the permutation distribution is constructed by pooling the effect estimated for the treated unit together with placebo effects estimated for the units in the donor pool. The effect of the treatment on the unit affected by the intervention is deemed to be significant when its magnitude

is extreme relative to the permutation distribution.

One potential complication with this procedure is that, even if a synthetic control is able to closely fit the trajectory of the outcome variable for the treated unit before the intervention, the same may not be true for all the units in the donor pool. For this reason, Abadie et al. (2010) propose a test statistic that measures the ratio of the post-intervention fit relative to the pre-intervention fit. For $0 \leq t_1 \leq t_2 \leq T$ and $j = \{1, \dots, J + 1\}$, let

$$R_j(t_1, t_2) = \left(\frac{1}{t_2 - t_1 + 1} \sum_{t=t_1}^{t_2} (Y_{jt} - \widehat{Y}_{jt}^N)^2 \right)^{1/2}, \quad (11)$$

where \widehat{Y}_{jt}^N is the outcome on period t produced by a synthetic control when unit j is coded as treated and using all other J units to construct the donor pool. This is the root mean squared prediction error (RMSPE) of the synthetic control estimator for unit j and time periods t_1, \dots, t_2 . The ratio between the post-intervention RMSPE and pre-intervention RMSPE for unit j is

$$r_j = \frac{R_j(T_0 + 1, T)}{R_j(1, T_0)}. \quad (12)$$

That is, r_j measures the quality of the fit of a synthetic control for unit j in the post-treatment period, relative to the quality of the fit in the pre-treatment period. Abadie et al. (2010) use the permutation distribution of r_j for inference. An alternative solution to the problem of poor pre-treatment fit in the donor pool is to base inference on the distribution $R_j(T_0 + 1, T)$ after discarding those placebo runs with $R_j(1, T_0)$ substantially larger than $R_1(1, T_0)$ (see Abadie et al., 2010).

A p -value for the inferential procedure based on the permutation distribution of r_j , as described above, is given by

$$p = \frac{1}{J + 1} \sum_{j=1}^{J+1} I_+(r_1 - r_j),$$

where $I_+(\cdot)$ is an indicator function which returns one for non-negative arguments and zero otherwise. While p -values are often used to summarize the results of testing procedures, the permutation distribution of the test statistics, r_j , or of the placebo gaps, $Y_{jt} - \widehat{Y}_{jt}^N$, are easy to report/visualize and provide additional information (e.g., on the magnitude of the differences between the estimated treatment effect on the treated unit and the placebo gaps in the donor

pool). P -values have also been used to construct confidence intervals for estimated treatment effects by inverting a test statistic (see, e.g., Firpo and Possebom, 2018).

Replacing $Y_{jt} - \widehat{Y}_{jt}^N$ in $R(T_0 + 1, T)$ with their positive or negative parts, $(Y_{jt} - \widehat{Y}_{jt}^N)^+$ or $(Y_{jt} - \widehat{Y}_{jt}^N)^-$, leads to one-sided inference. One-sided inference may result in a substantial gain of power.¹⁰ This is an important consideration in many comparative case study settings, where samples are considerably small. Alternative test statistics (see, e.g., Firpo and Possebom, 2018) could potentially be used to direct power to specific sets of alternatives.

As discussed in Abadie et al. (2010) this mode of inference reduces to classical randomization inference (Fisher, 1935) when the intervention is randomly assigned, a rather improbable setting, especially in contexts with aggregate units. More generally, this mode of inference evaluates significance relative to a benchmark distribution for the assignment process, one that is implemented directly in the data. Abadie et al. (2010) use a uniform benchmark, but one could easily depart from the uniform case. Firpo and Possebom (2018) propose a sensitivity analysis procedure that considers deviations from the uniform benchmark.

Because in most observational settings assignment to the intervention is not randomized, one could, in principle, adopt permutation schemes that incorporate information in the data on the assignment probabilities for the different units in the sample (as in, e.g., Rosenbaum, 1984). However, in many comparative case studies it is difficult to articulate the nature of a plausible assignment mechanism or even the specific nature of a placebo intervention. Consider, for example, the 1990 German reunification application in Abadie et al. (2015). In that context, it would be difficult to articulate the nature of the assignment mechanism or even describe placebo interventions (France would reunify with whom?) Moreover, even if a plausible assignment mechanism exists, estimation of the assignment mechanism is often hopeless, because many comparative case studies feature a single or a small number of treated units.

It is important to note that the availability of a well-defined procedure to select the comparison unit, like the one provided by the synthetic control method, makes the estimation of the effects of

¹⁰Notice that, in the presence of a treatment effect on the treated unit, permutations in which the treated unit contributes to the placebo synthetic control will tend to produce effects of the opposite sign to the effect on the treated unit, increasing the power of the one sided test.

placebo interventions feasible. Without a formal description of the procedure used to choose the comparison for the treated unit, it would be difficult to re-apply the same estimation procedure to the units in the donor pool. In this sense, the formalization of the choice of the comparison unit provided by the synthetic control method opens the door to a mode of quantitative inference in the context of comparative case studies.

Another important point to notice is that the permutation method described in this section does not attempt to approximate the sampling distributions of test statistics. Sampling-based statistical tests employ restrictions on the sampling mechanism (data generating process) to derive a distribution of a test statistic in a thought experiment where alternative samples could have been obtained from the sampling mechanism that generated the data. In a comparative case study framework, however, sampling-based inference is complicated—sometimes because of the absence of a well-defined sampling mechanism or data generating process, and sometimes because the sample is the same as the population. For example, in their study of the effect of terrorism on economic outcomes in Spain, Abadie and Gardeazabal (2003) employ a sample consisting of all Spanish regions. Here, sampling is not done at random from a well-defined superpopulation. As in classical randomization tests (Fisher, 1935), design-based inference takes care of these complications by conditioning on the sample and considering only the variation in the test statistic that is induced by the assignment mechanism (see, e.g., Abadie et al., 2020).¹¹

4. Why use synthetic controls?

In this section, I will describe some advantages of synthetic control estimators relative to alternative methods. For the sake of concreteness and because linear regression is arguably the most widely applied tool in empirical research in economics, I emphasize the differences between synthetic control estimators and linear regression estimators. However, much of the discussion applies more generally to other estimators of treatment effects.

A linear regression estimator of the effect of the treatment can easily be constructed using the panel data structure described in Section 3.1. Let \mathbf{Y}_0 be the $(T - T_0) \times J$ matrix of post-

¹¹In particular, this is in contrast to the bias bound calculations in Abadie et al. (2010), which are performed over the distribution of the individual transitory shocks, ε_{it} .

intervention outcomes for the units in the donor pool with (t, j) -element equal to Y_{j+1, T_0+t} . Let $\bar{\mathbf{X}}_1$ and $\bar{\mathbf{X}}_0$ be the result of augmenting \mathbf{X}_1 and \mathbf{X}_0 , respectively, with a row of ones. A regression-based estimator of the counterfactual Y_{1t}^N for $t > T_0$ is $\hat{\mathbf{B}}' \bar{\mathbf{X}}_1$, where $\hat{\mathbf{B}} = (\bar{\mathbf{X}}_0' \bar{\mathbf{X}}_0)^{-1} \bar{\mathbf{X}}_0' \mathbf{Y}'_0$. That is, the regression-based estimator is akin to a synthetic control, as it uses a linear combination, $\mathbf{Y}_0 \mathbf{W}^{reg}$, of the outcomes in the donor pool, with $\mathbf{W}^{reg} = \bar{\mathbf{X}}_0' (\bar{\mathbf{X}}_0' \bar{\mathbf{X}}_0)^{-1} \bar{\mathbf{X}}_1$, to reproduce the outcome of the treated unit in the absence of the intervention. Some advantages of synthetic controls relative to regression-based counterfactual are listed next.

No extrapolation. Synthetic control estimators preclude extrapolation, because synthetic control weights are non-negative and sum to one. It is easy to check that, like their synthetic control counterparts, the regression weights in \mathbf{W}^{reg} sum to one. Unlike the synthetic control weights, however, regression weights may be outside the $[0, 1]$ interval, allowing extrapolation outside of the support of the data (see Abadie et al., 2015, for details).¹² Table 3 reports regression weights for the German reunification example. In this application, the regression counterfactual utilizes negative values for four countries.

Transparency of the fit. Linear regression uses extrapolation to guarantee a perfect fit of the characteristics of the treated unit, $\bar{\mathbf{X}}_0 \mathbf{W}^{reg} = \bar{\mathbf{X}}_1$ (and, therefore, $\mathbf{X}_0 \mathbf{W}^{reg} = \mathbf{X}_1$) even when the untreated units are completely dissimilar in their characteristics to the treated unit. In contrast, synthetic controls make transparent the actual discrepancy between the treated unit and the convex combination of untreated units that provides the counterfactual of interest, $\mathbf{X}_1 - \mathbf{X}_0 \mathbf{W}^*$. This discrepancy is equal to the difference between columns (1) and (2) in Table 1. In addition, Figure 1(b) brings to light the fit of a synthetic control in terms of pre-intervention outcomes. That is, the information in Table 1 and Figure 1 makes clear the extent to which the observations in the donor pool can approximate the characteristics of the treated units by interpolation only. In some applications, comparisons like that of columns (1) and (2) of Table 1 may reveal that it is not possible to approximate the characteristics of the treated unit(s) using a weighted average of the units in the donor pool. In that case, Abadie et al. (2010, 2015) advise against using synthetic controls.

¹²See King and Zeng (2006) on the dangers of relying on extrapolation to estimate counterfactuals.

Safeguard against specification searches. In contrast to regression but similar to classical matching methods, synthetic controls do not require access to post-treatment outcomes in the design phase of the study, when synthetic controls are calculated. This implies that all the data analysis on design decisions like the identity of the units in the donor pool or the predictors in \mathbf{X}_1 and \mathbf{X}_0 can be made without knowing how they affect the conclusions of the study (see Rubin, 2007, for a related discussion regarding matching estimators). Moreover, synthetic control weights can be calculated and pre-registered/publicized before the post-treatment outcomes are realized, or before the actual intervention takes place. That is, pre-registration of synthetic control weights can play a role similar to pre-analysis plans in randomized control trials (see, e.g., Olken, 2015), providing a safeguard against specification searches and p -hacking.

Transparency of the counterfactual. Synthetic controls make explicit the contribution of each comparison unit to the counterfactual of interest. Moreover, because the synthetic control coefficients are proper weights and are sparse (more on sparsity below), they allow a simple and precise interpretation of the nature of the estimate of the counterfactual of interest. For the application to the effects of the German reunification in Table 2, the counterfactual for West Germany is given by a weighted average of Austria (0.42), Japan (0.16), the Netherlands (0.09), Switzerland (0.11) and the US (0.22) with weights in parentheses. Simplicity and transparency of the counterfactual allows the use of the expert knowledge to evaluate the validity of a synthetic control and the directions of potential biases. For instance, smaller neighboring countries of West Germany, such as Austria, Netherlands and Switzerland, have a substantial weight on the composition of the synthetic control of Table 2. If economic growth in these countries was negatively affected by the German reunification during the 1990-2003 period (perhaps because West Germany diverted demand and investment from these countries to East Germany) this would imply that Figure 1(b) estimates a lower bound on the magnitude (absolute value) of the negative effect of the German reunification on per capita GDP in West Germany.

Sparsity. As evidenced in the results of Tables 2 and 3 synthetic controls are sparse, but regression weights are not. As discussed above, sparsity plays an important role for the interpretation and evaluation of the estimated counterfactual. The sparsity of synthetic control weights has an

immediate geometric interpretation. Assume, for now, that \mathbf{X}_1 falls outside the convex hull of the columns of \mathbf{X}_0 . This is typical in empirical practice and a consequence of the curse of dimensionality. Assume also that the columns of \mathbf{X}_0 are in general position (that is, there is no set of m columns, with $2 \leq m \leq k + 1$, that fall into an $(m - 2)$ -dimensional hyperplane). Then, the synthetic control is unique and sparse—with the number of non-zero weights bounded by k —as it is the projection of \mathbf{X}_1 on the convex hull of the columns of \mathbf{X}_0 . Figure 2 provides a visual representation of the geometric interpretation of the sparsity property of synthetic control estimators. Only the control observations marked in red contribute to the synthetic control.

Notice that Table 1 indicates that the synthetic control for West Germany falls close to but outside the convex hull of the values of economic growth predictors in the donor pool (otherwise, columns (1) and (2) would be identical). As a result, the number of non-zero weights in Table 2 is not larger than the number of variables in Table 1. If desired, sparsity can be increased by imposing a bound on the density (number of non-zero weights) of \mathbf{W}^* in the calculation of synthetic controls (see Abadie et al., 2015).

In some cases, especially in applications with many treated units, the values of the predictors for some of the treated units may fall in the convex hull of the columns of \mathbf{X}_0 . Then, synthetic controls are not necessarily unique or sparse. That is, a minimizer of equation (7) may not be unique or sparse, although sparse solutions with no more than $k + 1$ non-zero weights always exist. A question is then how to choose among the typically infinite number of solutions to the minimization of equation (7). A modification of the synthetic control estimator in Abadie and L’Hour (2019) discussed in Section 8 addresses this problem and produces synthetic controls that are unique and sparse (provided that untreated observations are in general quadratic position, see Abadie and L’Hour, 2019, for details). In contrast, as shown in Table 3, regression estimators are typically not sparse.

It is important to notice that the role of sparsity in the context of synthetic control methods differs from the usual role that sparsity plays in other statistical methods like the lasso, where a sparsity-inducing regularization is employed to prevent overfitting, and where the interpretation of the lasso coefficients is often not at issue. Like for the lasso, the goal of synthetic controls

Australia	0.12
Austria	0.26
Belgium	0.00
Denmark	0.08
France	0.04
Greece	-0.09
Italy	-0.05
Japan	0.19
Netherlands	0.14
New Zealand	0.12
Norway	0.04
Portugal	-0.08
Spain	-0.01
Switzerland	0.05
United Kingdom	0.06
United States	0.13

Table 3: Regression weights for West Germany

is out-of-sample prediction; in particular, prediction of Y_{1t}^N for $t > T_0$. In contrast to the lasso, however, the identity and magnitude of non-zero coefficients constitute important information to interpret the nature of the estimate, and evaluate its validity and the potential for biases.

One of the greatest appeals of the synthetic control method resides, in my opinion, in the interpretability of the estimated counterfactuals, which results from the weighted average nature of synthetic control estimators and from the sparsity of the weights.

Despite the practical advantages of synthetic control methods, successful application of synthetic control estimators crucially depends on important contextual and data requirements, which are discussed in the next two sections.

5. Contextual Requirements

This section will discuss contextual requirements, that is, the conditions on the context of the investigation under which synthetic controls are appropriate tools for policy evaluation, as well as suitable ways to modify the analysis when these conditions do not perfectly hold. It is important,




Figure 2: Projecting \mathbf{X}_1 on the convex hull of \mathbf{X}_0

however, to point out that most of the requirements listed in this section pertain not only to synthetic control methods but also to any other type of comparative case study research design.

Size of the effect and volatility of the outcome. As previously discussed, the goal of comparative case studies is to estimate the effect of a policy intervention on the unit (e.g., state or region) exposed to an intervention of interest. That is, comparative case studies typically estimate the effect of an intervention on a single treated unit or on a small number of treated units. The nature of this exercise, which focuses on a single treated unit or on a small number of treated units, indicates that small effects will be indistinguishable from other shocks to the outcome of the affected unit, especially if the outcome variable of interest is highly volatile.¹³ As a result, the impact of “small” interventions with effects of a magnitude similar to the volatility of the outcome are difficult to detect. Even a large effect may be difficult to detect if the volatility of the outcome is also large. Outcome variables that include substantial random noise elevate the risk of overfitting, as explained in Section 3.3. In cases where substantial volatility is present in

¹³In studies that seek to estimate the average effect of an intervention that is observed in a large number of instances, the volatility of the outcome variable can often be reduced by averaging. In contrast, as explained above, comparative case studies in many cases focus on the effect of a single event or intervention.

the outcome of interest it is advisable to remove it via filtering in both the exposed unit as well as in the units in the donor pool before applying synthetic control techniques.¹⁴ Notice, however, that the challenge posed by volatility comes only from the fraction of it that is generated by unit-specific factors (e.g., the individual-specific transitory shocks, ε_{jt} , in equation (10)). Volatility generated by common factors affecting other units (e.g., the common factors λ_t in equation (10)) can be differentiated-out by choosing an appropriate synthetic control.

Availability of a comparison group. The very nature of comparative case studies implies that inference based on these methods will be faulty in the absence of a suitable comparison group. First and foremost, in order to have units available for the donor pool, it is important that not all units adopt interventions similar to the one under investigation during the period of the study. Units that adopt an intervention similar to the one adopted by the unit of interest should not be included in the donor pool because they are affected by the intervention, very much like the unit of interest. It is also important to eliminate from the donor pool any units that may have suffered large idiosyncratic shocks to the outcome of interest during the study period, if it is judged that such shocks would not have affected the outcome of the unit of interest in the absence of the intervention.¹⁵ Moreover, it is important to restrict the donor pool to units with characteristics that are similar to the affected unit. The reason is that, while the restrictions placed on the weights, \mathbf{W} , do not allow extrapolation, interpolation biases may still be important if the synthetic control matches the characteristics of the affected unit by averaging away large discrepancies between the characteristics of the affected unit and the characteristics of the units in the synthetic control. For the German reunification example, Abadie et al. (2015) restrict the donor pool to a set of OECD economies. Related to this point, Abadie and L’Hour (2019) propose adding to the objective function in equation (7) a set of penalty terms that depend on the discrepancies between the characteristics of the affected unit and the characteristics of the individual units included in the synthetic control (see Section 8 for details).

¹⁴For example, Amjad et al. (2018) propose singular value thresholding to de-noise data for synthetic controls.

¹⁵As an example, in their study of the effect of California’s tobacco control legislation, Abadie et al. (2010) discard from the donor pool several states that adopted large-scale tobacco programs or substantially increased taxes on tobacco during the sample period of the study.

No anticipation. As in any research design that exploits time variation in the outcome variable to estimate the effect of an intervention, synthetic control estimators may be biased if forward looking economic agents react in advance of the policy intervention under investigation, or if certain components of the intervention are put in place in advance of the formal implementation/enactment of the intervention. If there are signs of anticipation, it is advisable to backdate the intervention in the data set to a period before any anticipation effect can be expected, so the full extent of the effect of the intervention can be estimated. Notice that backdating the intervention in the data does not mechanically bias the estimator of the effect of the intervention even if some periods before the intervention are mistakenly recorded as post-intervention periods. The reason is that, as shown in equations (2) and (3), the synthetic control estimator does not restrict the time variation in the effect of the intervention. Therefore, periods barely affected by the intervention may show small or zero effects, while subsequent periods may produce a large estimated effect. This is in contrast with much of the practice using panel data models, where in many instances the effect of an intervention is restricted to be constant across post-intervention periods.

No interference. In the setup of Section 3.1, we defined the potential outcomes Y_{1t}^I and Y_{it}^N only in terms of the treatment status for unit 1 and unit i , respectively, at time t . This is the stable unit treatment value assumption in Rubin (1980), which implies that there is no interference across units. That is, units' outcomes are invariant to other units' treatments. In some instances, however, an intervention may have spillover effects on units that are not directly targeted by it. Assuming that such spillover effects do not exist is a strong restriction that must often be enforced in the design of the study or accounted for in the analysis of the results.

The assumption of no interference can be enforced in the design of a study by discarding from the donor pool those units with outcomes possibly affected by the intervention on the treated unit. Notice that there is a potential tension between this practice and the issues discussed in *Availability of a Comparison Group*. On the one hand, it is advisable to select for the donor pool units that are affected by the same regional economic shocks as the unit where the intervention happens. On the other hand, if spillover effects are substantial and affect units in close geograph-

ical proximity, those units may provide a biased estimate of the counterfactual outcome without intervention for the unit affected by the intervention. In cases when units potentially affected by spillover effects are discarded from the donor pool, the transparency of the fit of synthetic controls allows researchers to evaluate the reduction in the quality of the match between the characteristics of the treated unit and the characteristics of the synthetic control.

Potential spillover effects can also be accounted for in the analysis phase of a synthetic control study. If units affected by spillover effects are included in the synthetic control, the researcher should be aware of the potential direction of the bias of the resulting estimator. For example, Abadie et al. (2015) estimate the economic impact of the 1990 German reunification using a synthetic control of other OECD countries to approximate the trajectory of the counterfactual per capita GDP for West Germany in the absence of the unification. As explained above, if countries that compose the synthetic control for West Germany, like Austria, suffered from the negative effects of the German reunification, then we would expect the synthetic control estimator to be attenuated. That is, in this case, the synthetic control estimate would provide a lower bound on the magnitude of the causal effect of the German reunification on GDP per capita in West Germany. Notice that it is the transparency of the counterfactual and sparsity of the synthetic control counterfactual estimate that makes this exercise possible. In regression settings, like the one in Section 4, the weight of each unit in the counterfactual estimate is rarely computed in empirical practice, and non-sparse weight vectors often including negative components complicate the evaluation of potential biases.

Convex hull condition. Synthetic control estimates are predicated on the idea that a combination of unaffected units can approximate the pre-intervention characteristics of the affected unit. Once the synthetic control is constructed it should be checked that the differences in the characteristics of the affected unit and the synthetic control are small, that is:

$$X_{11} - w_2 X_{12} - \dots - w_J X_{1J+1} \approx 0, \dots, X_{k1} - w_2 X_{k2} - \dots - w_J X_{kJ+1} \approx 0.$$

In mathematical parlance, we need that $(X_{11}, X_{21}, \dots, X_{k1})$ falls close to the convex hull of the set of points $\{(X_{12}, X_{22}, \dots, X_{k2}), \dots, (X_{1J+1}, X_{2J+1}, \dots, X_{kJ+1})\}$. If the unit affected by the

intervention of interest is “extreme” in the value of a particular variable, such a value may not be closely approximated by a synthetic control.¹⁶

The fact that the value of a particular predictor for the treated cannot be closely approximated by the synthetic control may be less of a concern if the synthetic control closely tracks the trajectory of the outcome variable for the unit affected by the intervention during a hold-out validation period. In some cases, however, the unit affected by the intervention of interest may be extreme in the values of the outcome variable before the intervention and, as a result, there will not be a weighted average of untreated units that can approximate the trajectory of the outcome variable for the treated unit before the intervention. A potential way to proceed in those cases is to transform the outcome to time differences, $\Delta Y_{jt} = Y_{jt} - Y_{jt-1}$, or growth rates, $100 \times \Delta Y_{jt} / Y_{jt-1}$. Similarly, one could measure outcomes in differences with respect to pre-intervention means, $\tilde{Y}_{jt} = Y_{jt} - (1/T_0) \sum_{h=1}^{T_0} Y_{jh}$ (Ferman and Pinto, 2018). Consider the particular case where a synthetic control is calculated on the basis of all pre-intervention outcomes (equally-weighted). That is, each of the T_0 rows of \mathbf{X}_1 and \mathbf{X}_0 contains the outcome values in one of the pre-intervention periods for the treated unit and the donor pool, respectively, and all pre-intervention periods carry the same weight in the calculation of the synthetic control. For this particular case, measuring the outcomes in \mathbf{X}_1 and \mathbf{X}_0 in deviations with respect to the units’ pre-intervention means is equivalent to a proposal in Doudchenko and Imbens (2016), who measure outcomes in levels but allow for a constant shift in the synthetic control fit, $\mathbf{X}_1 - \alpha \mathbf{I}_{k \times 1} - \mathbf{X}_0 \mathbf{W}$ (with t -th row equal to $Y_{1t} - \alpha - w_2 Y_{2t} - \dots - w_{J+1} Y_{J+1t}$), where $\mathbf{I}_{k \times 1}$ is a vector of ones of the same dimension as \mathbf{X}_1 . As explained in Doudchenko and Imbens (2016), allowing for a constant shift between \mathbf{X}_1 and $\mathbf{X}_0 \mathbf{W}$ makes little sense in more general settings, when multiple covariates of different scales, instead of pre-intervention outcomes only, are included in \mathbf{X}_1 and \mathbf{X}_0 . Notice, however, that measuring outcomes in deviations with respect to their pre-intervention means instead of allowing for a constant shift between \mathbf{X}_1 and $\mathbf{X}_0 \mathbf{W}$ may still be useful to account for differences in the level of the outcomes across units, even if \mathbf{X}_1 and \mathbf{X}_0 include other predictors,

¹⁶For example, Abadie et al. (2015) find that because inflation levels were particularly low for West Germany before the reunification, the value of this variable cannot be closely reproduced by a synthetic control composed by other OECD countries. See Table 1.

aside from pre-intervention outcomes.

Transformations of the outcome variable like those in the previous paragraph may be useful in some cases because, as evidenced in the vast difference-in-differences literature, there are instances when a comparison group can reproduce the changes in the outcome variable for the unit of interest even if the level of the outcome variable cannot be reproduced. In other cases, however, credible counterfactuals require reproducing not only the trend of the outcome variable for the treated but also the level. For example, some formulations of the convergence hypothesis in economic growth imply that countries with different levels of per-capita GDP will tend to experience different growth rates on average, in the absence of an intervention. Similarly, non-linearities in labor earnings profiles over the life cycle imply that differences in the age distribution across populations will typically result in differences in the growth of labor earnings.¹⁷ More generally, there may not exist a combination of untreated units that provide a credible approximation to the treated units, and the conventional synthetic control estimator should not be used in that case.

It should also be noted that differencing the dependent variable may result in a substantial increase in the part of the variance of the outcome that is attributable to noise, potentially inducing an increase in bias. As an example, consider the linear factor model in equation (10). Differencing equation (10) we obtain

$$\Delta Y_{jt}^N = \Delta \delta_t + \Delta \boldsymbol{\theta}_t \mathbf{Z}_j + \Delta \boldsymbol{\lambda}_t \boldsymbol{\mu}_j + \Delta \varepsilon_{jt},$$

where $\Delta Y_{jt}^N = Y_{jt}^N - Y_{jt-1}^N$ with analogous expressions for $\Delta \delta_t$, $\Delta \boldsymbol{\theta}_t$, $\Delta \boldsymbol{\lambda}_t$, and $\Delta \varepsilon_{jt}$. Notice first that the differenced equation retains the linear factor structure. Notice also that differencing the outcome may help control the bias when the vectors of common factors $\boldsymbol{\theta}_t$ and $\boldsymbol{\lambda}_t$, or at least some of their components, vary little in time. In that case, the magnitudes of $\Delta \boldsymbol{\theta}_t$ and $\Delta \boldsymbol{\lambda}_t$ may be small even if the magnitudes of $\boldsymbol{\theta}_t$ and $\boldsymbol{\lambda}_t$ are large. This is the usual rationale for working with differenced outcomes and the basis for difference-in-differences estimators. There

¹⁷Notice that for non-linearities in the process that generates Y_{it}^N may require that, for each unit j contributing to the synthetic control, \mathbf{X}_j is reasonably close to \mathbf{X}_1 . Section 8 describes a synthetic control estimator with weights that penalize $\|\mathbf{X}_1 - \mathbf{X}_j\|$.

may be opposing forces at play, however. Suppose, in particular, that the idiosyncratic shocks, ε_{jt} , are independent or roughly independent in time. Then, the variance of $\Delta\varepsilon_{jt}$ is larger than the variance of ε_{jt} . Now, following the characterization of the bias in Section 3.3, a larger residual variance may result in a higher risk of overfitting and an increase in the bias of the synthetic control estimator.

Time horizon. The effect of some interventions may take time to emerge or to be of sufficient magnitude to be quantitatively detected in the data. An obvious but unsatisfying approach to this problem is to wait until the effects of the intervention run their course. A more proactive approach is to use surrogate outcomes or leading indicators of the outcome variable of interest.

6. Data Requirements

This section discusses data requirements for credible applications of synthetic controls. Like many of the contextual requirements in the previous section, the data requirements discussed here apply not only to synthetic control estimation but, more generally, to comparative case study methods.

Aggregate data on predictors and outcomes. From the previous discussion, it can be seen that the synthetic control method requires the availability of data on outcomes and predictors of the outcome for the unit or units exposed to the intervention of interest and a set of comparison units. Predictors and outcomes are often series routinely reported by government agencies, multilateral organizations and private entities. Examples of these types of outcomes are state-level crime rates in the U.S. (Donohue et al., 2019), country-level per-capita GDP (Abadie et al., 2015), and state-level cigarette consumption statistics in the U.S. (Abadie et al., 2010), which are routinely reported in publications produced or commissioned by the FBI, the World Bank and tobacco industry groups, respectively. Sometimes, when aggregate data do not exist, aggregates of micro-data are employed in comparative case studies. For example, in his study of the labor market effects of the Mariel Boatlift in Miami, Card (1990) uses micro-data from the Current Population Survey (CPS) to estimate aggregate values for wage rates and unemployment for workers in Miami and a set of four comparison cities before and after the Mariel Boatlift. Similarly, Bohn

et al. (2014) use data from the CPS to estimate the fraction of the population composed by non-citizens Hispanic by state in the U.S.

Sufficient pre-intervention information. The credibility of a synthetic control estimator depends in great part on its ability to steadily track the trajectory of the outcome variable for the affected unit before the intervention. As discussed in Section 3.3, Abadie, Diamond, and Hainmueller (2010) show that if the data generating process follows a linear factor model, then the bias of the synthetic control estimator is bounded by a function that is inversely proportional to the number of pre-intervention periods (provided that the synthetic control closely tracks the trajectory of the outcome variable for the affected unit during the pre-intervention periods). Therefore, when designing a synthetic control study, it is of crucial importance to collect information on the affected unit and the donor pool for a large pre-intervention window.

A caveat to the preference for a large number of pre-intervention periods is given by the possibility of structural breaks. Consider the linear factor model of equation (10). In this model, structural stability is represented by the restriction of constant factor loadings. Even if the model is a good representation of the distribution of the data at a relatively short time scale, its accuracy may suffer once we allow the number of periods to be large enough. Choosing v_1, \dots, v_k to up-weight the most recent measures (relative to the prediction window) included in \mathbf{X}_1 and \mathbf{X}_0 helps alleviate structural instability concerns.

With a small number of pre-intervention periods, close or even perfect fit of the predictor values for the treated unit may be spuriously attained, in which case the resulting synthetic control may fail to reproduce the trajectory of the outcome for the treated unit in the absence of the intervention. The severity of this problem can be diminished if powerful predictors of post-intervention values of Y_{jt}^N , aside from pre-intervention values of the outcome, are included in \mathbf{X}_j , reducing the residual variance and, as a result, the risk of overfitting.

Sufficient post-intervention information. This data requirement derives partly from the *Time horizon* contextual requirement in Section 5. The evaluation data must include outcome measures that are possibly affected by the intervention and are relevant for the policy decision or

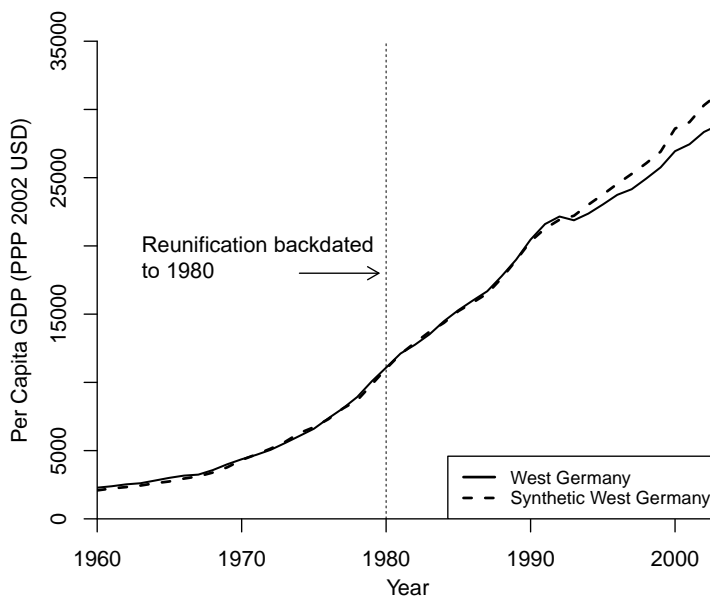


Figure 3: Backdating the 1990 German reunification application

scientific inquiry that is the object of the study. This may be problematic if the effect of an intervention is expected to arise gradually over time and if no forward looking measures of the outcome are available. Conversely, in some practical instances, the effect of an intervention may dissipate rapidly after showing substantial effectiveness for a few initial periods. Extensive post-intervention information allows a more complete picture of the effects of the intervention, in time and across the various outcomes of interest.

7. Robustness and Diagnostic Checks

The credibility of a synthetic control estimator depends on its ability to reproduce outcomes for the treated unit in the absence of the intervention. This section presents diagnostic checks that can be used to evaluate the credibility of synthetic control counterfactuals in actual applications, as well as robustness exercises to assess sensitivity of results to changes in the design of the study.

Backdating. The possibility of backdating was discussed in Section 5 as a way to address anticipation effects on the outcome variable before an intervention occurs. In the absence of anticipation effects, the same idea can be applied to assess the credibility of a synthetic control in concrete empirical applications. Figure 3 shows the result of estimating the effect of the 1990 German

reunification with the intervention backdated to 1980. Two important features of the results are as follows. First, as one would hope for, the synthetic control estimator closely tracks per-capita GDP for West Germany in the 1981-1990 period, before the start of the actual intervention. This is the “in-time placebo test” in Abadie et al. (2015) and similar to the “pre-program test” in Heckman and Hotz (1989). The absence of estimated effects prior to the intervention provides credibility of the synthetic control estimator, as it demonstrates that the synthetic control is able to reproduce the trajectory of the outcome variable for the treated unit before the intervention occurs. Second, a gap between per-capita GDP for West Germany and its synthetic control counterpart appears around the time of the German reunification, as in Figure 1(b). This is the case even when the intervention is ten-year backdated in the data and the procedure uses no information on the timing of the actual intervention. The shape and direction of the gap in Figure 3 is similar to that of Figure 1(b), albeit of a somewhat lower magnitude. The fact that the estimated effect of the German reunification appears shortly after 1990 even when the intervention is artificially ten-year backdated in the data provides credibility to the synthetic control estimator of the 1990 German reunification.

Robustness tests. Regardless of the estimation method employed in the analysis, the main conclusions of an empirical study should display some level of robustness with respect to changes in the study design. In the context of synthetic controls, two important ways the design of a study may influence results are (i) the choice of units in the donor pool, and (ii) the choice of predictors of the outcome variable. The first choice corresponds to the columns in \mathbf{X}_0 , and the second one corresponds to the rows in $[\mathbf{X}_1 : \mathbf{X}_0]$.

As an example of a robustness test, Figure 4 reports the results of a leave-one-out re-analysis of the German reunification data in Abadie et al. (2015), taking from the sample one-at-a-time each of the countries that contribute to the synthetic control in Table 2. All leave-one-out estimates closely track the per-capita GDP series for West Germany before 1990. The resulting estimates for the years after the reunification are all negative and centered around the result produced using the entire donor pool. The main conclusion of a negative estimate of the German reunification on per-capita GDP is robust to the exclusion of any particular country. In other examples, however,

results may not be as robust as those in Figure 4, and the scientific significance of the estimates should be evaluated with that information in mind. If the exclusion of a unit from the donor pool has a large effect on results without a discernible change in pre-intervention fit, this may warrant investigating if the change in the magnitude of the estimate is caused by the effects of other interventions or by particularly large idiosyncratic shocks on the outcome of the excluded untreated unit.




Figure 4: Leave-one-out estimates of the effect of the 1990 German reunification

8. Extensions and Related Methods

As the literature on synthetic control methods and related methods has greatly expanded in recent years, it has become increasingly difficult for researchers interested in applying these methods to figure out what is available where. In this section, I provide a brief guide to the recent contributions in the area. This represents only an incomplete snapshot of a literature that is rapidly evolving.

Multiple treated units. Several recent articles consider estimation and inference with synthetic controls for the case where there are multiple treated units. Notice that the presence of multiple

treated units does not give rise to additional conceptual challenges for the estimation of synthetic controls. Treatment effects can be estimated for each treated unit separately and aggregated in a second step if desired. However, the presence of multiple treated units creates some practical problems for estimation, as well as new challenges and opportunities for inference.

A potential complication with synthetic control estimation is that the minimizer of equation (7) subject to the weight constraints may not be unique, especially if the values of the predictors for a treated unit fall inside the convex hull of the values of the predictors for the donor pool. Suppose, for now, that only the first unit is treated. If \mathbf{X}_1 belongs to the convex hull of the columns of \mathbf{X}_0 , this implies that we can find \mathbf{W}^* such that $\mathbf{X}_1 = \mathbf{X}_0\mathbf{W}^*$. Moreover, the number of minimizers to equation (7) may be (and will typically be) infinite.¹⁸ That is, there may exist an infinite number of solutions to the problem of minimizing of equation (7), subject to the weight constraints, that perfectly reproduce \mathbf{X}_1 . An algorithm minimizing equation (7) subject to the weight constraints may select a solution, \mathbf{W}^* , with positive entries for units that are far away from the treated unit in the space of the predictors, even when an alternative solution exists based only on units with predictor values similar to \mathbf{X}_1 . This may, in turn, lead to large interpolation biases that remain unchecked under the illusion of perfect fit, $\mathbf{X}_1 = \mathbf{X}_0\mathbf{W}^*$.¹⁹

Even in moderate dimensions, k , the curse of dimensionality works to keep treated observations outside of the convex hull of the units in the donor pool. Therefore, in settings with one treated unit, multiplicity of solutions is rarely an issue, and if it arises it can often be easily addressed by restricting the donor pool to units with predictor values most similar to the values of the predictor for the unit exposed to the treatment. However, in settings with many treated and untreated units, multiplicity of solutions and how to choose among them become important issues for estimation. Moreover, large interpolation biases may also arise in settings where the predictor values for treated units fall outside the convex hull of the predictor values for the units in the donor pool, especially when the units contributing to synthetic controls are far away from

¹⁸If \mathbf{W}_1^* and \mathbf{W}_2^* are both minimizers of equation (7) subject to the weight constraints, then so is $a\mathbf{W}_1^* + (1 - a)\mathbf{W}_2^*$, for $a \in (0, 1)$. This implies that the number of solutions to the minimization equation (7) given the weight constraints can only be one or infinity.

¹⁹Notice, however, that perfect fit, $\mathbf{X}_1 = \mathbf{X}_0\mathbf{W}^*$, where \mathbf{W}^* minimizes equation (7) subject to the weight constraints, is indicative of the possibility of infinite solutions.

the treated units in the space of predictors. That is, there may be cases such that $\mathbf{X}_1 \approx \mathbf{X}_0 \mathbf{W}^*$ but where \mathbf{X}_j greatly differs from \mathbf{X}_1 , for some unit j contributing to the synthetic control.

To address these challenges, Abadie and L’Hour (2019) propose a synthetic control estimator that incorporates a penalty for pairwise matching discrepancies between the treated units and each of the units that contribute to their synthetic controls. Consider a setting with I treated units and J untreated units. We will index observations so that the treated units come first. That is, units $j = 1, \dots, I$ are treated and units $j = I + 1, \dots, I + J$ are untreated, with $I + J$ units in total. As in previous sections, \mathbf{X}_j is the vector of predictor values for unit j , and \mathbf{X}_0 is the matrix of the predictor values for the units in the donor pool. To simplify the exposition, I will abstract from the choice of v_1, \dots, v_k , which can be implemented (at a higher computational cost) using the procedure in Section 3.2. For $\lambda > 0$, the estimator in Abadie and L’Hour (2019) minimizes

$$\|\mathbf{X}_i - \mathbf{X}_0 \mathbf{W}\|^2 + \lambda \sum_{j=I+1}^{I+J} w_j \|\mathbf{X}_i - \mathbf{X}_j\|^2, \quad (13)$$

with respect to $\mathbf{W} = (w_{I+1}, \dots, w_{I+J})'$, for each treated unit, $i = 1, \dots, I$, subject to the constraints that the weights w_{I+1}, \dots, w_{I+J} are non-negative and sum to one.²⁰ The first term in equation (13) is the aggregate discrepancy between the predictor values for treated unit i and its synthetic control. The second term in equation (13) penalizes pairwise matching discrepancies between the predictor values for unit i and each of the units that contribute to its synthetic control, weighted by the magnitudes of their contributions. The penalty term is added to equation (13) with the aim of reducing interpolation biases. As $\lambda \rightarrow \infty$, the penalized estimator converges to one-to-one matching. As $\lambda \rightarrow 0$, the estimator uses an aggregate of pairwise matching discrepancies weighted by \mathbf{W} to select among all synthetic controls that attain the minimal value for $\|\mathbf{X}_i - \mathbf{X}_0 \mathbf{W}\|$. Values of λ between zero and infinity trade off aggregate fit of the synthetic control and pairwise fit of each of the units that contribute to it.

Abadie and L’Hour (2019) show that if $\lambda > 0$, then the minimizer of equation (13) is unique and sparse (provided that the columns of \mathbf{X}_0 are in general quadratic position, see Abadie and L’Hour (2019) for details). They also provide cross-validation techniques to select λ .

²⁰Although this is not reflected in the notation in equation (13), λ may depend on i .

Let $\mathbf{W}_i^* = (w_{iI+1}^*, \dots, w_{iI+J}^*)'$ be the solution to the minimization problem in equation (13). Then, the estimated treatment effect for $i = 1, \dots, I$ and $t = T_0 + 1, \dots, T$ is as in (8),

$$\hat{\tau}_{it} = Y_{it} - \sum_{j=I+1}^{I+J} w_{ij}^* Y_{jt}, \quad (14)$$

with average treatment effect given by

$$\hat{\tau}_t = \frac{1}{I} \sum_{i=1}^I \hat{\tau}_{it}.$$

In many instances, especially where the sample units are aggregates, like regions or countries, a weighted average (e.g., population-weighted, or GDP-weighted) treatment effect may be most relevant.

Dube and Zipperer (2015) and Abadie and L'Hour (2019) propose extensions of the permutation methods in Abadie et al. (2010) to the case with multiple treated units. They employ rank-based statistics on the permutation distribution of treatment effects, where the identity of the treated units is permuted at random in the data. In particular, Abadie and L'Hour (2019) propose the following simple generalization of the permutation test in Abadie et al. (2010). They consider a setting with I treated units and J untreated units. In each permutation $b = 1, \dots, B$, the identities of the I treated units are reassigned in the data among the $I + J$ units in the sample, and statistics $r_{b,1}, \dots, r_{b,I}$ are calculated for the units coded as treated in the permutation. These statistics could be (bias-corrected) synthetic control estimates of treatment effects, or rescaled versions that take into account the pre-intervention fit as in equation (12), or their absolute values, positive parts, or negative parts, depending on the context. Notice that when I is small relative to $I + J$, it may be possible to consider all possible treatment reassignments, in which case B is equal to $(I + J)$ -choose- I . If considering all possible treatment reassignments is computationally expensive, inference can be based on B random draws from all subsets of I units in the sample. Let $r_{0,1}, \dots, r_{0,I}$ be the same statistics calculated for the actual treated units. B permutation repetitions, in addition to the original sample values for treatment, produce $I \times (B + 1)$ statistics, $r_{0,1}, \dots, r_{0,I}, \dots, r_{B,1}, \dots, r_{B,I}$. Now, for each $b = 0, \dots, B$, one can calculate t_b equal to the sum of the ranks of $r_{b,1}, \dots, r_{b,I}$ within $r_{0,1}, \dots, r_{0,I}, \dots, r_{B,1}, \dots, r_{B,I}$.

The permutation inference in Abadie and L’Hour (2019) is based on the “extremeness” of the statistic t_0 within the permutation distribution t_0, t_1, \dots, t_B . Notice that, for $I = 1$ this mode of inference amounts to the permutation test in Abadie et al. (2010).

Hainmueller (2012) and Robbins et al. (2017) consider also settings with multiple treated units. Instead of producing a separate synthetic control for each treated unit, they calculate a single synthetic control to match aggregate values of the predictors between the treated and non-treated samples. As in the usual synthetic control estimator, the weights in Hainmueller (2012) and Robbins et al. (2017) are non-negative and sum to a pre-determined constant (typically equal to one, or to the number of treated units, depending on the scaling of the variables in the data-set). These estimators require that there is at least a convex combination units in the donor pool that exactly matches a pre-specified set of moments of the predictors for the treated units. Among the sets of weights that perfectly reproduce the moments for the treated sample, Hainmueller (2012) and Robbins et al. (2017) choose the one that minimizes a measure of discrepancy with respect to constant weights.

Bias correction. Another practical complication in settings with many treated units is that, even with a moderate k , the predictor values for some of the treated units may not be closely reproduced by a synthetic control, or may be closely reproduced only by combinations of units with large pairwise matching discrepancies in predictor values with respect to the treated unit. At the same time, including those ill-fitted units in the calculation of the aggregate effect may be important for the desired interpretation of the estimate (e.g., as an estimate of the average effect of the treatment on the treated). In that case, one could be concerned about the potential biases produced by matching discrepancies between the values of the predictors for the treated units and those for the respective synthetic controls.²¹ Bias-corrections play also an important role reducing regularization biases in inferential methods for regression-based variants of synthetic controls (see, e.g., Arkhangelsky et al., 2019; Chernozhukov et al., 2019b).

Abadie and L’Hour (2019) and Ben-Michael et al. (2019) propose modifications of the synthetic

²¹Related to this problem, Ferman and Pinto (2018) and Botosaru and Ferman (2019) study the properties of synthetic control estimators for case where the value of the predictors for a treated unit cannot be closely matched by a synthetic control.

control estimator along the lines of the bias-correction techniques of Rubin (1973), Quade (1982), and Abadie and Imbens (2011). They use regression-adjustments to attenuate the bias of synthetic control estimators in settings where the synthetic control counterfactual is constructed using untreated units with values of the predictors that do not closely reproduce the predictor values for the treated unit or units. For $t = T_0 + 1, \dots, T$, let $\hat{\mu}_{0t}$ be a sample regression function (parametric or nonparametric) estimated by regressing the untreated outcomes, $Y_{I+1,t}, \dots, Y_{I+J,t}$ on the values of the predictors for the untreated units, $\mathbf{X}_{I+1}, \dots, \mathbf{X}_{I+J}$. The bias-corrected synthetic control estimator for unit i is

$$\hat{\tau}_{it} = \left(Y_{it} - \sum_{j=I+1}^{I+J} w_{ij}^* Y_{jt} \right) - \sum_{j=I+1}^{I+J} w_{ij}^* \left(\hat{\mu}_{0t}(\mathbf{X}_i) - \hat{\mu}_{0t}(\mathbf{X}_j) \right). \quad (15)$$

The first term on the right-hand side of (15) is the synthetic control estimator in (14). The second term uses a regression adjustment to correct for discrepancies between the predictor values for the treated unit and the predictor values for the units that contribute to the synthetic control. Alternatively, the estimator in (15) can be expressed as

$$\hat{\tau}_{it} = (Y_{it} - \hat{\mu}_{0t}(\mathbf{X}_i)) - \sum_{j=I+1}^{I+J} w_{ij}^* (Y_{jt} - \hat{\mu}_{0t}(\mathbf{X}_j)). \quad (16)$$

Equation (16) provides an interpretation of the bias-corrected synthetic control estimator as a synthetic control estimator applied to regression residuals. The bias correction in equation (16) is related to the proposal in Doudchenko and Imbens (2016) to residualize the outcomes with respect to covariates before calculating synthetic controls.

A different avenue to evaluate the bias of synthetic control estimators, which was discussed in Section 7, is given by the availability of pre-intervention periods, when the effect of the treatment is not yet realized. In the absence of anticipation effects, estimates of treatment effects before the intervention are reflective of estimation biases. If biases are stable in time, estimates of those biases could be used to correct synthetic control estimates. Bias adjustments of this type, which are closely related to difference-in-differences methods, are proposed in Arkhangelsky et al. (2019) and Chernozhukov et al. (2019b).

Regression-based methods and extrapolation. Several articles have contributed regression-based estimators for synthetic controls. These procedures allow extrapolation by considering synthetic controls that are not convex combinations of the units in the donor pool. Doudchenko and Imbens (2016) consider an estimator that fits all pre-treatment outcomes for the treated, with weights that may be negative and may not sum to one, and allow for a constant shift in the level of the synthetic control estimator. They propose to use an elastic net—that is, a combination of lasso (L_1) and ridge (L_2) penalties—to regularize the weights. The counterfactual estimates for $t = T_0 + 1, \dots, T$ in Doudchenko and Imbens (2016) are

$$\widehat{Y}_{1t}^N = \widehat{\alpha} + \sum_{j=2}^{J+1} \widehat{w}_j Y_{jt}, \quad (17)$$

where $\widehat{\alpha}, \widehat{w}_2, \dots, \widehat{w}_{J+1}$ minimize

$$\sum_{t=1}^{T_0} \left(Y_{1t} - \alpha - \sum_{j=2}^{J+1} w_j Y_{jt} \right)^2 + \lambda_1 \left(\frac{1 - \lambda_2}{2} \sum_{j=2}^{J+1} w_j^2 + \lambda_2 \sum_{j=2}^{J+1} |w_j| \right), \quad (18)$$

with respect to $\alpha, w_2, \dots, w_{J+1} \in \mathbb{R}^{J+1}$ (unrestricted), and $\lambda_1 \geq 0$ and $0 \leq \lambda_2 \leq 1$ are regularization parameters selected by cross-validation. To incorporate additional predictors in their estimation procedure, Doudchenko and Imbens (2016) propose to use least squares in a first step to residualize the outcomes Y_{jt} for $j = 1, \dots, J + 1$ in equation (18) with respect to any other covariates. Chernozhukov et al. (2019a) consider different penalty terms, including lasso regularization (i.e., $\lambda_2 = 1$) in the context of an inferential procedure for synthetic controls. Regression estimators of this type are also related to the panel data approach to program evaluation estimator in Hsiao et al. (2012), where $\lambda_1 = 0$ and the parameters in equation (17) are estimated by unpenalized least squares. Li (2019) considers the same estimator as in Hsiao et al. (2012), but regularizes the weights $\widehat{w}_2, \dots, \widehat{w}_{J+1}$ to be non-negative.

Arkhangelsky et al. (2019) introduce a synthetic control estimator that weights not only the units in the control group, but also the pre-intervention time periods, to approximate the counterfactual of interest. The time weights in Arkhangelsky et al. (2019) play a similar role as the predictor weights, v_1, \dots, v_k , of Section 3.2. They reflect the importance of each of the individual

predictors, which in the leading version of the estimator of Arkhangelsky et al. (2019) are past outcome values.

Matrix completion/estimation methods. Amjad et al. (2018, 2019) and Athey et al. (2020) propose related methods that use tools from the matrix completion/matrix estimation literature. Suppose, as before, that unit $j = 1$ is the treated unit, and units $j = 2, \dots, J + 1$ are not treated. Amjad et al. (2018) posit a non-linear factor-structure model for the untreated, $Y_{jt}^N = f(\mu_j, \lambda_t) + \varepsilon_{jt}$, with $j = 2, \dots, J + 1$ and $t = 1, \dots, T$, where ε_{jt} is random noise. Their framework allows for the presence of missing values in the matrix $\{Y_{jt}^N\}$, with $j = 1, \dots, J + 1$, and $t = 1, \dots, T$. Using matrix estimation methods, in particular, singular value thresholding (see Chatterjee, 2015), Amjad et al. (2018) estimate a low-rank approximation, $\{\widehat{M}_{jt}\}$, to the matrix $\{M_{jt}\} = \{f(\mu_j, \lambda_t)\}$. The objects \widehat{M}_{jt} are used to de-noise the outcomes Y_{jt}^N and to impute missing values, if any. Then, synthetic controls are obtained as linear combinations of \widehat{M}_{jt} , with coefficients estimated by ridge regression of Y_{1t} on $\widehat{M}_{2t}, \dots, \widehat{M}_{J+1t}$ in the pre-intervention periods. The estimator in Amjad et al. (2018) does not incorporate covariates, using data on outcomes, Y_{jt} , only. Amjad et al. (2019) modifies the estimator in Amjad et al. (2018) to incorporate additional variables aside from the outcome of interest, under the assumption that all variables depend on common latent factors. Athey et al. (2020) postulate the model $Y_{jt}^N = M_{jt} + \varepsilon_{jt}$, for $j = 1, \dots, J + 1$ and $t = 1, \dots, T$, where ε_{jt} is again random noise. In their framework, missing entries in the matrix $\{Y_{jt}^N\}$, with $j = 1, \dots, J + 1$ and $t = 1, \dots, T$, arise naturally for the treated observation (or treated observations, if multiple units are treated) in the post-treatment periods. Athey et al. (2020) assume that the matrix $\{M_{jt}\}$, with $j = 1, \dots, J + 1$ and $t = 1, \dots, T$, is low-rank, which allows them to obtain an estimate, $\{\widehat{M}_{jt}\}$, via matrix completion techniques. The estimated counterfactual outcomes without the treatment for the treated are the values of \widehat{M}_{jt} such that Y_{jt}^N is missing. Extensions allow for models with covariates and the inclusion of time fixed effects and unit fixed effects (separate from the low-rank matrix, M_{jt}).

Inference. Several studies have proposed inferential tools for synthetic controls as alternatives to the permutation test in Section 3.5. Firpo and Possebom (2018) propose several generalizations of the permutation test in Section 3.5 and contribute confidence sets based on inverting the

results of these tests. In a repeated sampling framework for stationary data and large T_0 , Hahn and Shi (2017) propose to apply the end-of-sample instability test of Andrews (2003) to obtain an inferential procedure for synthetic control estimators. In the context of synthetic control estimators, the end-of-sample instability test of Andrews (2003) is related to the backdating ideas of Section 7. It compares the values of treatment effects computed for the $T - T_0$ post-intervention periods to the distribution of the of same values computed for every subset of $T - T_0$ consecutive pre-intervention periods. Related also to Andrews’s end-of-sample instability test, Chernozhukov et al. (2019a) devise a sampling-based inferential procedure for synthetic controls and related methods that employs permutations of regression residuals in the time dimension. In particular, Chernozhukov et al. (2019a) assume $Y_{1t}^N = P_t^N + u_t$, where u_1, \dots, u_T are stationary and weakly dependent with mean zero. Let $\tau_{T_0+1}, \dots, \tau_T$ be the effects of the treatment on the treated unit (unit one) at times $t = T_0 + 1, \dots, T$.²² The potential outcome under the intervention is $Y_{1t}^I = P_t^N + \tau_t + u_t$ for $t > T_0$. To simplify the exposition, assume that u_1, \dots, u_T are i.i.d. Then, the distribution of a function, $S(u_{T_0+1}, \dots, u_T)$, of the post-intervention values of u_t should be the same as the distribution of $S(u_{\pi(T_0+1)}, \dots, u_{\pi(T)})$, where $\pi(1), \dots, \pi(T)$ is a random permutation of $1, \dots, T$. Suppose for now that P_t^N is known. Then, under a null hypothesis, $\tau_{T_0+1} = a_{T_0+1}, \dots, \tau_T = a_T$, we can compute $u_t = Y_{1t} - P_t^N - a_t$, where $a_t = 0$ for $1 \leq t \leq T_0$. As a result, we can test the null hypothesis by comparing the value of $S(u_{T_0+1}, \dots, u_T)$ to its permutation distribution, that is, the distribution of $S(u_{\pi(T_0+1)}, \dots, u_{\pi(T)})$, which can be directly computed in the data. A feasible implementation of the test requires estimation of the residuals, u_1, \dots, u_T . In the context of the synthetic control method, Chernozhukov et al. (2019a) adopt the model $P_t^N = \sum_{j=2}^{J+1} w_j Y_{jt}$ with non-negative weights that sum to one, and $E[u_t Y_{jt}] = 0$ for $j = 2, \dots, J + 1$, and implement their test on constrained least squares residuals, $\hat{u}_1, \dots, \hat{u}_T$. The proposal in Chernozhukov et al. (2019a) differs from other synthetic control procedures in two important respects. First, while much of the literature on synthetic controls has adopted the linear factor model of Section 3.3 as a working model to understand the properties of synthetic control estimators, Chernozhukov et al. (2019a) adopt instead the

²²To be consistent with the notation for P_t^N and u_t and because only unit one is treated, here I drop the subscript indicating that identity of the treated unit from the notation for treatment effect.

restriction $E[u_t Y_{jt}] = 0$ for $j = 2, \dots, J+1$ to estimate P_t^N .²³ They show, however, that regardless of the validity of the model, their testing procedure remains valid as long as the estimated residuals, $\hat{u}_1, \dots, \hat{u}_T$ are exchangeable under the null hypothesis. Second, in contrast to other synthetic control procedures that compute the weights, w_2, \dots, w_{J+1} using pre-intervention data only, in the inferential procedure of Chernozhukov et al. (2019a) the synthetic control weights are estimated under the null hypothesis, $\tau_{T_0+1} = a_{T_0+1}, \dots, \tau_T = a_T$, using data on $Y_{it}^N = Y_{it} - a_t$ for all periods, including the periods after the intervention. For a similar set of models, Chernozhukov et al. (2019b) propose bias-corrected synthetic control estimation and confidence intervals for the mean value of the treatment effect over the post-intervention period, $(\tau_{T_0+1} + \dots + \tau_T)/(T - T_0)$ in settings when both T_0 and $T - T_0$ are large. Similar to difference-in-differences, their bias-correction procedure of Chernozhukov et al. (2019b) adjusts for differences in pre-intervention outcomes between the treated unit and the synthetic control. Confidence intervals are based on an asymptotically pivotal t -statistic and centered on the average of K -fold cross-fitted versions of the bias-corrected synthetic control estimate. Cattaneo et al. (2019) propose predictive intervals for synthetic control estimators and related methods. They consider a setting where data are stationary (perhaps after a transformation) and adopt the predictive model $Y_{1t}^N = \sum_{j=2}^{J+1} w_j Y_{jt} + u_t$, where the weights w_2, \dots, w_{J+1} are defined to minimize $E[(Y_{1t}^N - \tilde{w}_2 Y_{2t} - \dots - \tilde{w}_{J+1} Y_{J+1t})^2]$ with respect to $\tilde{w}_2, \dots, \tilde{w}_{J+1}$, for each $t = 1, \dots, T$ subject to weight constraints. Cattaneo et al. (2019) use pre-intervention data along with the assumption of stationarity to obtain bounds on $\sum_{j=2}^{J+1} w_j Y_{jt}$ and u_t for $t > T_0$, which they combine to obtain a predictive interval on Y_{1t}^N for $t > T_0$.

Other contributions. In this article, I have provided a brief description of selected strands of the literature on synthetic controls and related methods, starting with the canonical estimator in Sections 2 and 3, and describing some extensions and related methods in the current section.

²³To understand the differences between these two frameworks, notice that when the data are generated by the linear factor model of Section 3.3, and there is an unbiased synthetic control—that is, a synthetic control with weights, w_2, \dots, w_{J+1} that exactly reproduces \mathbf{Z}_1 and $\boldsymbol{\mu}_1$ —then, the restriction $E[u_t Y_{jt}] = 0$ for $j = 2, \dots, J+1$ and $u_t = Y_{1t}^N - \sum_{j=2}^{J+1} w_j Y_{jt}$ does not hold in general (see Ferman and Pinto, 2018). One exception is given by the results in Ferman (2019), which imply that $E[u_t Y_{jt}] = 0$ for $j = 2, \dots, J+1$ will approximately hold as $J \rightarrow \infty$ if there are weights, w_2, \dots, w_{J+1} that asymptotically recover Z_1 and μ_1 and are increasingly diluted among the units in the donor pool.

The literature is vast in its totality, however, and there are many noteworthy contributions I did not cover. They include Bai and Ng (2019), Brodersen et al. (2015), Gobillon and Magnac (2016), Gunsilius (2020), Kennedy-Shaffer et al. (2020), Viviano and Bradic (2019), and Xu (2017), among many others. Many of these are much recent proposals. Samartsidis et al. (2019) studies the performance of the canonical synthetic control estimator and of related methods in the context of the German reunification example of Section 3.2. As the set of methods on synthetic controls keeps expanding and enriching the applied econometrics toolkit, this is still a young literature and much remains to be done. I mention some open areas in the final section of this article.

9. Conclusions

Synthetic controls provide many practical advantages for the estimation of the effects of policy interventions and other events of interest. However, like for any other statistical procedure (and especially for those aimed at estimating causal effects), the credibility of the results depends crucially on the level of diligence exerted in the application of the method and on whether contextual and data requirements are met in the empirical application at hand. In this article, I emphasize the notion that mechanical applications of synthetic controls that do not take into account the context of the investigation or the nature of the data are risky enterprises. To this end, the article discusses the methodological underpinnings of synthetic control estimators and the conditions under which they provide suitable estimates of causal effects. It also describes how the analysis may be modified in the cases when those conditions do not hold. Finally, the article discusses some recent extensions that widen the applicability, robustness and flexibility of the method.

Open areas of related research abound, both methodological and empirical. Results on sampling-based inference, external validity, sensitivity to model restrictions, estimation with multiple interventions, and the identification of the channels through which the effect of an event or intervention operates, to mention a few, are scant or absent in the synthetic controls literature. An area of recent heightened interest regarding the use of synthetic controls is the design of experimental interventions in settings where the intervention of interest can only be applied

to one or a small number of aggregate units. In addition, existing results on robust and efficient computation of synthetic controls are scarce, and more research is needed on the computational aspects of this methodology. On the empirical side, many of the events and the policy interventions economists care about take place at an aggregate level, affecting entire aggregate units like school districts, cities, regions, or countries. This is exactly the setting synthetic controls were designed for, and potential applications of synthetic controls in economics are many.

References

- Abadie, A. (2005). Semiparametric difference-in-differences. *Review of Economic Studies* 72(1), 1–19.
- Abadie, A., S. Athey, G. W. Imbens, and J. M. Wooldridge (2020). Sampling-based versus design-based uncertainty in regression analysis. *Econometrica* 88(1), 265–296.
- Abadie, A., A. Diamond, and J. Hainmueller (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., A. Diamond, and J. Hainmueller (2015). Comparative politics and the synthetic control method. *American Journal of Political Science* 59(2), 495–510.
- Abadie, A. and J. Gardeazabal (2003). The economic costs of conflict: A case study of the Basque Country. *American Economic Review* 93(1), 113–132.
- Abadie, A. and G. W. Imbens (2011). Bias-corrected matching estimators for average treatment effects. *Journal of Business & Economic Statistics* 29(1), 1–11.
- Abadie, A. and J. L’Hour (2019). A penalized synthetic control estimator for disaggregated data. Working paper. Available at <https://sites.google.com/site/jeremyhour/research>.
- Acemoglu, D., S. Johnson, A. Kermani, J. Kwak, and T. Mitton (2016). The value of connections in turbulent times: Evidence from the United States. *Journal of Financial Economics* 121, 368–391.

- Allegretto, S., A. Dube, M. Reich, and B. Zipperer (2017). Credible research designs for minimum wage studies: A response to Neumark, Salas, and Wascher. *ILR Review* 70, 559–592.
- Amjad, M., D. Shah, and D. Shen (2018). Robust synthetic control. *Journal of Machine Learning Research* 19(22), 1–51.
- Amjad, M. J., V. Misra, D. Shah, and D. Shen (2019). mRSC: Multidimensional robust synthetic control. In *Proc. ACM Meas. Anal. Comput. Syst.*, Volume 3.
- Andrews, D. W. K. (2003). End-of-sample instability tests. *Econometrica* 71(6), 1661–1694.
- Angrist, J. D. and J.-S. Pischke (2008). *Mostly Harmless Econometrics: An Empiricist’s Companion*. Princeton University Press.
- Arkhangelsky, D., S. Athey, D. A. Hirshberg, G. W. Imbens, and S. Wager (2019). Synthetic difference in differences. Working Paper 25532. National Bureau of Economic Research.
- Athey, S., M. Bayati, N. Doudchenko, G. Imbens, and K. Khosravi (2020). Matrix completion methods for causal panel data models. Working paper.
- Athey, S. and G. W. Imbens (2017). The state of applied econometrics: Causality and policy evaluation. *The Journal of Economic Perspectives* 31, 23–32.
- Bai, J. (2009). Panel data models with interactive fixed effects. *Econometrica* 77(4), 1229–1279.
- Bai, J. and S. Ng (2019). Matrix completion, counterfactuals, and factor analysis of missing data. *arXiv e-prints*, arXiv:1910.06677.
- Ben-Michael, E., A. Feller, and J. Rothstein (2019). The augmented synthetic control method. Working paper.
- Bifulco, R., R. Rubenstein, and H. Sohn (2017). Using synthetic controls to evaluate the effect of unique interventions: The case of Say Yes to Education. *Evaluation Review* 41(6), 593–619.

- Bohn, S., M. Lofstrom, and S. Raphael (2014). Did the 2007 Legal Arizona Workers Act reduce the state’s unauthorized immigrant population? *Review of Economics and Statistics* 96(2), 258–269.
- Borjas, G. J. (2017). The wage impact of the Marielitos: A reappraisal. *ILR Review* 70, 1077–1110.
- Botosaru, I. and B. Ferman (2019). On the Role of Covariates in the Synthetic Control Method. *The Econometrics Journal*. Forthcoming.
- Brodersen, K. H., F. Gallusser, J. Koehler, N. Remy, S. L. Scott, et al. (2015). Inferring causal impact using bayesian structural time-series models. *The Annals of Applied Statistics* 9(1), 247–274.
- Card, D. (1990). The impact of the Mariel Boatlift on the Miami labor market. *ILR Review* 44, 245–257.
- Card, D. and A. B. Krueger (1994). Minimum wages and employment: A case study of the fast-food industry in New Jersey and Pennsylvania. *American Economic Review* 84, 772–793.
- Cattaneo, M. D., Y. Feng, and R. Titiunik (2019). Prediction intervals for synthetic control methods. *arXiv e-prints*, arXiv:1912.07120.
- Chatterjee, S. (2015). Matrix estimation by universal singular value thresholding. *The Annals of Statistics* 43(1), 177–214.
- Chernozhukov, V., K. Wüthrich, and Y. Zhu (2019a). An Exact and Robust Conformal Inference Method for Counterfactual and Synthetic Controls. *arXiv e-prints*, arXiv:1712.09089v6.
- Chernozhukov, V., K. Wüthrich, and Y. Zhu (2019b). Practical and robust t -test based inference for synthetic control and related methods. *arXiv e-prints*, arXiv:1812.10820v2.
- Cook, T. D. and D. T. Campbell (1979). *Quasi-Experimentation. Design & Analysis Issues for Field Settings*. Boston: Houghton Mifflin.

- Cunningham, S. and M. Shah (2018). Decriminalizing indoor prostitution: Implications for sexual violence and public health. *The Review of Economic Studies* 85(3), 1683–1715.
- Donohue, J. J., A. Aneja, and K. D. Weber (2019). Right-to-carry laws and violent crime: A comprehensive assessment using panel data and a state-level synthetic control analysis. *Journal of Empirical Legal Studies* 16(2), 198–247.
- Doudchenko, N. and G. W. Imbens (2016). Balancing, regression, difference-in-differences and synthetic control methods: A synthesis. NBER working paper no. 22791.
- Douglas, J. (2018). How an analysis of Basque terrorism helps economists understand Brexit. *The Wall Street Journal*, November 7.
- Dube, A. and B. Zipperer (2015). Pooling multiple case studies using synthetic controls: An application to minimum wage policies. IZA Discussion Papers 8944, Institute for the Study of Labor (IZA).
- Ferman, B. (2019). On the properties of the synthetic control estimator with many periods and many controls. *arXiv e-prints*, arXiv:1906.06665.
- Ferman, B. and C. Pinto (2018). Synthetic controls with imperfect pre-treatment fit. Working paper.
- Firpo, S. and V. Possebom (2018). Synthetic control method: Inference, sensitivity analysis and confidence sets. *Journal of Causal Inference* 6(2).
- Fisher, R. (1935). *The Design of Experiments*. Edinburgh: Oliver and Boyd.
- Gobillon, L. and T. Magnac (2016). Regional policy evaluation: Interactive fixed effects and synthetic controls. *Review of Economics and Statistics* 98(3), 535–551.
- Gunsilius, F. (2020). Distributional synthetic controls. *arXiv e-prints*, arXiv:2001.06118v3.
- Guo, J. (2015). Seriously, here’s one amazing math trick to learn what can’t be known. *Wonkblog*, *The Washington Post*, October 30.

- Gutierrez, I. A., G. Weinberger, and J. Engberg (2016). *Improving Teaching Effectiveness: Impact on Student Outcomes: The Intensive Partnerships for Effective Teaching Through 2013–2014*. Santa Monica, CA: RAND Corporation.
- Hahn, J. and R. Shi (2017). Synthetic control and inference. *Econometrics* 5(4), 52.
- Hainmueller, J. (2012). Entropy balancing for causal effects: A multivariate reweighting method to produce balanced samples in observational studies. *Political Analysis* 20(1), 25–46.
- Heckman, J. J. and V. J. Hotz (1989). Choosing among alternative nonexperimental methods for estimating the impact of social programs: The case of manpower training. *Journal of the American Statistical Association* 84(408), 862–874.
- Heersink, B., B. D. Peterson, and J. A. Jenkins (2017). Disasters and elections: Estimating the net effect of damage and relief in historical perspective. *Political Analysis* 25, 260–268.
- Holland, P. V. (1986). Statistics and causal inference. *Journal of the American Statistical Association* 81, 945–960.
- Hsiao, C., H. S. Ching, and S. K. Wan (2012). A panel data approach for program evaluation: Measuring the benefits of political and economic integration of Hong Kong with mainland China. *Journal of Applied Econometrics* 27, 705–740.
- Jardim, E., M. C. Long, R. Plotnick, E. V. Inwegen, J. Vigdor, and H. Wething (2017). Minimum wage increases, wages, and low-wage employment: Evidence from Seattle. National Bureau of Economic Research.
- Kennedy-Shaffer, L., V. de Gruttola, and M. Lipsitch (2020). Novel methods for the analysis of stepped wedge cluster randomized trials. *Statistics in Medicine* 39(7), 815–844.
- King, G. and L. Zeng (2006). The dangers of extreme counterfactuals. *Political Analysis* 14(2), 131–159.
- Kleven, H. J., C. Landais, and E. Saez (2013). Taxation and international migration of superstars: Evidence from the European football market. *American Economic Review* 103(5), 1892–1924.

- Klößner, S., A. Kaul, G. Pfeifer, and M. Schieler (2018). Comparative politics and the synthetic control method revisited: a note on abadie et al. (2015). *Swiss Journal of Economics and Statistics* 154(1), 11.
- Kreif, N., R. Grieve, D. Hangartner, A. J. Turner, S. Nikolova, and M. Sutton (2016). Examination of the synthetic control method for evaluating health policies with multiple treated units. *Health Economics* 25(12), 1514–1528.
- Li, K. T. (2019). Statistical inference for average treatment effects estimated by synthetic control methods. *Journal of the American Statistical Association*. Forthcoming.
- Neumark, D. and W. Wascher (2017). Reply to “Credible research designs for minimum wage studies”. *ILR Review* 70, 593–609.
- Olken, B. A. (2015). Promises and perils of pre-analysis plans. *Journal of Economic Perspectives* 29(3), 61–80.
- Peri, G. and V. Yasenov (2019). The labor market effects of a refugee wave: Synthetic control method meets the Mariel boatlift. *Journal of Human Resources* 54(2), 267–309.
- Pieters, H., D. Curzi, A. Olper, and J. Swinnen (2017). Effect of democratic reforms on child mortality: A synthetic control analysis. *The Lancet Global Health* 4, e627–e632.
- Pinotti, P. (2015). The economic costs of organised crime: Evidence from Southern Italy. *Economic Journal* 125, F203–F232.
- Quade, D. (1982). Nonparametric analysis of covariance by matching. *Biometrics* 38(3), 597–611.
- Reich, M., S. Allegretto, and A. Gødoy (2017). Seattle’s minimum wage experience 2015-16. CWED policy brief, Center on wage and employment dynamics, UC Berkeley.
- Robbins, M. W., J. Saunders, and B. Kilmer (2017). A framework for synthetic control methods with high-dimensional, micro-level data: Evaluating a neighborhood-specific crime intervention. *Journal of the American Statistical Association* 112, 109–126.

- Rosenbaum, P. R. (1984). Conditional permutation tests and the propensity score in observational studies. *Journal of the American Statistical Association* 79(387), 565–574.
- Rubin, D. B. (1973). The use of matched sampling and regression adjustments to remove bias in observational studies. *Biometrics* 29, 185–203.
- Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies,. *Journal of Educational Psychology* 66, 688–701.
- Rubin, D. B. (1980). Discussion of “Randomization analysis of experimental data in the Fisher randomization test” by Basu. *Journal of the American Statistical Association* 75(371), 591–593.
- Rubin, D. B. (2007). The design versus the analysis of observational studies for causal effects: Parallels with the design of randomized trials. *Statistics in Medicine* 26(1).
- Samartsidis, P., S. R. Seaman, A. M. Presanis, M. Hickman, and D. De Angelis (2019). Assessing the causal effect of binary interventions from observational panel data with few treated units. *Statistical Science*. Forthcoming.
- Thistlethwaite, D. L. and D. T. Campbell (1960). Regression-discontinuity analysis: An alternative to the ex-post facto experiment. *Journal of Educational Psychology* 51, 309–317.
- Viviano, D. and J. Bradic (2019). Synthetic learner: Model-free inference on treatments over time. *arXiv e-prints*, arXiv:1904.01490.
- Xu, Y. (2017). Generalized synthetic control method: Causal inference with interactive fixed effects models. *Political Analysis* 25(1), 57–76.