

Synthetic Controls for Experimental Design

Alberto Abadie
MIT

Jinglong Zhao
Boston University

December 2023

Abstract

This article studies experimental design in settings where the experimental units are large aggregate entities (e.g., markets), and only one or a small number of units can be exposed to the treatment. In such settings, randomization of the treatment may result in treated and control groups with very different characteristics at baseline, inducing biases. We propose a variety of synthetic control designs (Abadie, Diamond and Hainmueller, 2010, Abadie and Gardeazabal, 2003) as experimental designs to select treated units in non-randomized experiments with large aggregate units, as well as the untreated units to be used as a control group. Average potential outcomes are estimated as weighted averages of treated units, for potential outcomes with treatment—and control units, for potential outcomes without treatment. We analyze the properties of estimators based on synthetic control designs and propose new inferential techniques. We show that in experimental settings with aggregate units, synthetic control designs can substantially reduce estimation biases in comparison to randomization of the treatment.

1. Introduction

Consider the problem of a ride-sharing company choosing between two compensation plans for drivers (Doudchenko et al., n.d.; Jones and Barrows, 2019). The company can either keep the current compensation plan or adopt a new one with higher incentives. In order to estimate the effect of a change in compensation plans on profits, the company’s data science unit designs an experimental evaluation where the new plan is deployed at a small scale, say, in one of the local markets (cities) in the country. In this setting, a randomized control trial—or A/B test, where drivers in a local market are randomized into the new plan (active treatment arm) or the status-quo (control treatment arm)—is problematic. On the one hand, such an experiment

Alberto Abadie, Department of Economics, MIT, abadie@mit.edu. Jinglong Zhao, Questrom School of Business, Boston University, jinglong@bu.edu. The authors are grateful to Victor Chernozhukov, Guido Imbens, Rahul Mazumder, Jaume Vives-i-Bastida, Yinchu Zhu and seminar participants at Amazon.com, Brown, Harvard, and Princeton for helpful comments and discussions. Abadie gratefully acknowledges NSF funding (SES-1756692).

raises equity concerns, as drivers in the same local market but in different treatment arms obtain different compensations for the same jobs. On the other hand, if drivers in the active treatment arm respond to higher incentives by working longer hours, they will effectively steal business from drivers in the control arm of the experiment, which will result in biased experimental estimates.¹

One possible approach to this problem is to assign an entire local market to treatment, and use the rest of the local markets, which remain under the current compensation plan during the experimental periods, as potential comparison units. In this setting, using randomization to assign the active treatment allows ex-ante (i.e., pre-randomization) unbiased estimation of the effect of the active treatment. However, ex-post (i.e., post-randomization) biases can be large if, at baseline, the treated unit is different from the untreated units in the values of the features that affect the outcomes of interest.

As in the ride-sharing example where there is only one treated local market, large biases may arise more generally in randomized studies when either the treatment arm or the control arm contains a small number of units, so randomized treatment assignment may not produce treated and control groups that are similar in their features. In those cases, the fact that estimation biases would have averaged out over alternative treatment assignments is of little comfort to a researcher who is, in practice, limited to one assignment only.

To address these challenges, we propose the use of the synthetic control method (Abadie, Diamond and Hainmueller, 2010, Abadie and Gardeazabal, 2003) as an experimental design to select treated units in non-randomized experiments, as well as the untreated units to be used as a comparison group. We adopt the name *synthetic control designs* to refer to the resulting experimental designs.^{2,3}

¹A randomized evaluation across many markets is a potential solution to the problem of experimental interference between drivers. In practice, however, large-scale market-level randomized evaluations are often unfeasible. In the context of the ridesharing company example, large-scale market-level randomized evaluations (*i*) could be prohibitively expensive, (*ii*) could still raise substantial equity concerns, (*iii*) could create great disappointment among the large number of drivers in the treated cities if the program is rolled back after experimentation, and (*iv*) in some cases, the number of cities where the company operates could be too small for effective randomization.

²While we leave the “experimental” qualifier implicit in “synthetic control design”, it should be noted that the synthetic control designs proposed in this article differ from observational synthetic control designs (e.g., Abadie, Diamond and Hainmueller, 2010, Abadie and Gardeazabal, 2003, Doudchenko and Imbens, 2016), for which the identity of the treated unit(s) is taken as given.

³See, e.g., Abadie (2021), Amjad, Shah and Shen (2018), Arkhangelsky et al. (2021), Doudchenko and Imbens

In our framework, the choice of the treated unit (or treated units, if multiple treated units are desired) aims to accomplish two goals. First, it is often useful to select the treated units such that their features are representative of the features of an aggregate of interest, like an entire country market. The treatment effect for the treated units selected in this way may more accurately reflect the effect of the treatment on the aggregate of interest. Second, the treated units should not be idiosyncratic in the sense that their features cannot be closely approximated by the untreated units. Otherwise, the reliability of the estimate of the effect on the treated unit may be questionable. We show how to achieve these two objectives, whenever they are possible to achieve, using synthetic control methods.

While we are aware of the extensive use of synthetic control methods for experimental design in data science units, especially in the technology industry,⁴ the academic literature on this subject is at a nascent stage. There are, however, a few publicly available studies that are connected to this article. Aside to the present article, to our knowledge, Doudchenko et al. (n.d.) and Doudchenko et al. (2021) are the only other publicly available studies on the topic of experimental design with synthetic controls. The focus of Doudchenko et al. (n.d.) is on statistical power, which they calculate by simulation of the estimated effects of placebo interventions on historical (pre-experimental) data. That is, the selection of treated units is based on a measure of statistical power implied by the distribution of the placebo estimates for each unit. As a result, estimates based on the procedure in Doudchenko et al. (n.d.) target the effect of the treatment for the unit or units that are most closely tracked in the placebo distribution. In the same spirit, the target parameter in Doudchenko et al. (2021) is the treatment effect for a weighted average of treated units that can be closely matched in their pre-treatment outcomes by a weighted average of untreated units. In the present article, we aim to take a different perspective on the problem of unit selection in experiments with synthetic controls; one that takes into account the extent to which different sets of treated and control units approximate an aggregate causal effect of interest chosen by the analyst, such as the average treatment effect for

(2016) for background material on synthetic controls and related methods.

⁴See, in particular, Jones and Barrows (2019), which also provides the basis for the ride-sharing example above.

the relevant population.⁵ The inferential methods in the present article also differ from those in the related literature. In particular, Doudchenko et al. (2021) propose a permutation procedure for inference that requires that potential outcomes without the treatment are independent and identically distributed (i.i.d.) in time. In contrast, the inferential procedure proposed in the present article allows for time series dependence and non-stationarity in outcomes, which are pervasive features of time-series data.⁶ Another important difference between the present article and Doudchenko et al. (n.d.) and Doudchenko et al. (2021) is that Doudchenko et al. (n.d.) and Doudchenko et al. (2021) make use of pre-treatment outcomes only to select treated and control units, while our method allows the use of other observed features of the units. Agarwal, Shah and Shen (2021) propose synthetic interventions, a framework related to synthetic controls, and apply it to estimate treatment effect heterogeneity in an experimental setting with multiple treatments. Their work primarily focuses on the analyses of experimental data, but not on the design of experiments. Bottmer et al. (2021) is also related to the present article in the sense that they study synthetic control estimation in an experimental setting. Their article, however, considers only the case when the treatment is randomized, and is not concerned with issues of experimental design.

The research designs in this article are also related to ex-ante synthetic control designs for observational studies (see Abadie (2021), Kasy and Lehner (2023) and Chen (2023) for an online version of the problem) in that the synthetic weights are computed and can be pre-registered before post-intervention outcomes are realized. However, our methods significantly differ in one key aspect: they address the challenge that experimenters face when selecting specific units for exposure to the intervention of interest.

An alternative approach to control post-randomization bias involves stratifying units based on covariate values prior to treatment randomization within those strata. Stratification can significantly reduce post-randomization biases if units have similar covariate values within strata. However, traditional stratification methods do not adapt to the setting considered in this article,

⁵Consistent with the majority of literature on synthetic controls, our focus is primarily on average treatment effects. For an analysis of distributional effects using synthetic controls, see Gunsilius (2023).

⁶Our inferential approach is geared towards hypothesis testing. Firpo and Possebom (2018) and Cattaneo, Feng and Titiunik (2021) offer valuable insights into interval estimation with synthetic controls.

which features a limited number of large aggregate entities as units of analysis, and a single unit or a handful chosen for treatment. Because every stratum in stratified designs must have at least one unit randomized into treatment, the number of strata cannot exceed the desired number of treated units in the experiment. In the case of only one treated unit, we would be limited to a single stratum. This may lead to significant variation in units' characteristics within strata, reducing the appeal of stratification procedures. Additionally, a stratified design with a small number of strata/treated units might end up selecting a set of treated units that are not truly representative of the target population.

The rest of the article is organized as follows: Section 2 presents and discusses the synthetic control designs proposed in this article. Section 3 details the formal properties of estimators based on these synthetic control designs. In Section 4, we report the findings from an empirical validation of synthetic control designs, using sales data from a sample of Walmart stores. Section 5 discusses the results of simulation studies. Finally, Section 6 provides concluding remarks. The appendix contains proofs and supplemental materials.

2. Synthetic Control Designs

We consider a setting with T time periods and J units, which may represent J local markets as in the ride-sharing example in the previous section. Let T_0 be the number of pre-experimental periods, with $1 \leq T_0 < T$. At the end of period T_0 , a researcher designs an experiment to conduct during periods $T_0 + 1, T_0 + 2, \dots, T$. Using information available at T_0 , the experimenter aims to select the set of units that will be administered treatment (intervention) during the experimental periods.

To define causal parameters, we formally adopt a potential outcomes framework. For any $j \in \{1, \dots, J\}$ and any $t \in \{T_0 + 1, \dots, T\}$, let Y_{jt}^I be the potential outcome for unit j at time t when unit j is exposed to treatment starting at $T_0 + 1$. Similarly, for any $j \in \{1, \dots, J\}$ and any $t \in \{1, \dots, T\}$, let Y_{jt}^N be the potential outcome for unit j at time t under no treatment. In the ridesharing example, Y_{jt}^I and Y_{jt}^N could measure net income divided by market size, under

the active and the control treatment, respectively. Unit-level treatment effects are defined as

$$Y_{jt}^I - Y_{jt}^N,$$

for $j = 1, \dots, J$ and $t = T_0 + 1, \dots, T$. They represent the effect of switching unit j to the active treatment at time $T_0 + 1$ on the outcome of unit j at time $t > T_0$. We aim to estimate the average treatment effect

$$\tau_t = \sum_{j=1}^J f_j \cdot (Y_{jt}^I - Y_{jt}^N),$$

for $t = T_0 + 1, \dots, T$. In this expression, f_1, \dots, f_J represents a set of known positive weights that are relevant to the definition of the average. In the ride-sharing example of the previous section, f_j could represent the size of local market j as a share of the national market. Without loss of generality, and because it is the case in many applications, we can assume that the weights f_j sum to one,

$$\sum_{j=1}^J f_j = 1.$$

In the case when units are equally weighted, we set $f_j = 1/J$ for $j = 1, \dots, J$. We use the notation \mathbf{f} for a vector that collects the values of f_j for all the units, i.e., $\mathbf{f} = (f_1, \dots, f_J)$.

At time T_0 , in order to estimate the treatment effect τ_t for $t = T_0 + 1, \dots, T$, the experimenter chooses $\mathbf{w} = (w_1, \dots, w_J)$ and $\mathbf{v} = (v_1, \dots, v_J)$, such that

$$\sum_{j=1}^J w_j = 1,$$

$$\sum_{j=1}^J v_j = 1,$$

$$w_j \geq 0, v_j \geq 0, \text{ and } w_j v_j = 0, \forall j = 1, \dots, J.$$

Units with $w_j > 0$ are units that will be assigned to the intervention of interest from $T_0 + 1$ to T . These units are chosen to estimate average outcomes under the intervention of interest. Units with $w_j = 0$ constitute an untreated reservoir of potential control units (a “donor pool”). Among units with $w_j = 0$, those with $v_j > 0$ are used to estimate average outcomes under no intervention.

The first goal of the experimenter is to choose w_1, \dots, w_J such that

$$\sum_{j=1}^J w_j Y_{jt}^I = \sum_{j=1}^J f_j Y_{jt}^I, \quad (1)$$

for $t = T_0 + 1, \dots, T$. If equation (1) holds, a weighted average of outcomes for the units selected for treatment reproduces the average outcome with treatment for the entire population of J units. In practice, however, the choice of w_1, \dots, w_J cannot directly rely on matching the population average of Y_{jt}^I , as in equation (1). The quantities Y_{jt}^I are unobserved before time $T_0 + 1$, and will remain unobserved in the experimental periods for the units that are not exposed to the treatment. Instead, we aim to approximate equation (1) using predictors observed at T_0 of the values of $Y_{jT_0+1}^I, \dots, Y_{jT}^I$. Note also that it is not possible to use the weights $w_1 = f_1, \dots, w_J = f_J$, because it would leave no units in the donor pool, making the set of units with $v_j > 0$ empty and violating the second condition in (2).

The second goal of the experimenter is to choose v_1, \dots, v_J such that

$$\sum_{j=1}^J v_j Y_{jt}^N = \sum_{j=1}^J f_j Y_{jt}^N, \quad (2)$$

or, alternatively,

$$\sum_{j=1}^J v_j Y_{jt}^N = \sum_{j=1}^J w_j Y_{jt}^N. \quad (3)$$

If equations (2) or (3) hold, a weighted average of outcomes for the units in the donor pool reproduces the average outcome without treatment for the entire population of J units (equation (2)), or for the units selected for treatment (equation (3)). Like in the previous case with treated outcomes, it is not feasible to directly choose v_1, \dots, v_J so that equation (2) or (3) is satisfied. Instead, we propose a variety of methods to approximate either (2) or (3) based on predictors of $Y_{jT_0+1}^N, \dots, Y_{jT}^N$.

For the treated units, we define $Y_{jt} = Y_{jt}^N$ if $t = 1, \dots, T_0$, and $Y_{jt} = Y_{jt}^I$ if $t = T_0 + 1, \dots, T$. For the untreated units, we define $Y_{jt} = Y_{jt}^N$, for all $t = 1, \dots, T$. That is, Y_{jt} is the outcome observed for unit $j = 1, \dots, J$ at time $t = 1, \dots, T$. We say that

$$\sum_{j=1}^J w_j Y_{jt} \quad \text{and} \quad \sum_{j=1}^J v_j Y_{jt}$$

are the synthetic treated and synthetic control outcomes, respectively. The difference between these two quantities is

$$\tau_t(\mathbf{w}, \mathbf{v}) = \sum_{j=1}^J w_j Y_{jt} - \sum_{j=1}^J v_j Y_{jt},$$

for $t = T_0 + 1, \dots, T$. Suppose that equations (1) and (2) hold. Then, $\tau_t(\mathbf{w}, \mathbf{v})$ is equal to the average treatment effect, τ_t . If equation (3) holds instead, then $\tau_t(\mathbf{w}, \mathbf{v})$ is equal to the average effect of the treatment on the treated (\mathbf{w} -weighted),

$$\tau_t^T = \sum_{j=1}^J w_j \cdot (Y_{jt}^I - Y_{jt}^N)$$

(Doudchenko et al., 2021).

We choose $\mathbf{w} = (w_1, \dots, w_J)$ and $\mathbf{v} = (v_1, \dots, v_J)$ to match the pre-intervention values of predictors of the potential outcomes Y_{jt}^N and Y_{jt}^I for $t > T_0$.

Let \mathbf{X}_j be a column vector of pre-intervention features of unit j . We see the features in \mathbf{X}_j as predictors of the values of Y_{jt}^N and Y_{jt}^I in the experimental periods, in a sense that will be made precise in Section 3. We use the notation

$$\overline{\mathbf{X}} = \sum_{j=1}^J f_j \mathbf{X}_j.$$

That is, $\overline{\mathbf{X}}$ is the vector of population values for the predictors in \mathbf{X}_j . For any real vector \mathbf{x} , $\|\mathbf{x}\|$ is the Euclidean norm of \mathbf{x} , and $\|\mathbf{x}\|_0$ is the number of non-zero coordinates of \mathbf{x} . Let \underline{m} and \overline{m} be positive integers such that $1 \leq \underline{m} \leq \overline{m} \leq J - 1$. A simple selector of $\mathbf{w} = (w_1, \dots, w_J)$ and $\mathbf{v} = (v_1, \dots, v_J)$ is

$$\begin{aligned} \min_{\substack{w_1, \dots, w_J, \\ v_1, \dots, v_J}} & \left\| \overline{\mathbf{X}} - \sum_{j=1}^J w_j \mathbf{X}_j \right\|^2 + \left\| \overline{\mathbf{X}} - \sum_{j=1}^J v_j \mathbf{X}_j \right\|^2 \\ \text{s.t.} & \sum_{j=1}^J w_j = 1, \\ & \sum_{j=1}^J v_j = 1, \\ & w_j, v_j \geq 0, \quad \forall j = 1, \dots, J, \end{aligned}$$

$$\begin{aligned}
w_j v_j &= 0, \quad \forall j = 1, \dots, J, \\
\underline{m} &\leq \|\mathbf{w}\|_0 \leq \overline{m}.
\end{aligned} \tag{4}$$

The first term of the objective function in (4) measures the discrepancies between the population average of the features in \mathbf{X}_j (\mathbf{f} -weighted) and the averages of the features for units assigned to the treatment group (\mathbf{w} -weighted). The second term is analogous but with the second average taken over the units assigned to no intervention (\mathbf{v} -weighted). The first four constraints require that the weights in \mathbf{w} , as well as the weights in \mathbf{v} , are non-negative and sum to one. They also require that any unit selected for treatment cannot be utilized as a control unit — so, if $w_j > 0$, then $v_j = 0$. The last constraint allows a minimum and maximum number of units assigned to treatment. This restriction is of practical importance in a variety of contexts, especially when experimentation is costly and the experimenter is restricted in the number of units that may receive the treatment. We say that the design is *Unconstrained* if $\underline{m} = 1$ and $\overline{m} = J - 1$; otherwise, we say the design is *Constrained*. The last constraint in (4) is not the only conceivable restriction to the size or cost of the experiment. An explicit upper bound on the cost of an experiment would be given by $\mathbf{c}'\mathbf{d} \leq B$, where the j -th coordinate of \mathbf{c} is equal to the cost of assigning unit j to treatment, \mathbf{d} is a J -dimensional vector with ones at coordinates where $w_j > 0$, and zeros otherwise, and B is the experimenter's budget.

Let $\mathbf{w}^* = (w_1^*, \dots, w_J^*)$ and $\mathbf{v}^* = (v_1^*, \dots, v_J^*)$ be a solution to the optimization problem in (4). In practice, we do not require optimality of $(\mathbf{w}^*, \mathbf{v}^*)$, as long as $(\mathbf{w}^*, \mathbf{v}^*)$ is feasible and satisfies $\overline{\mathbf{X}} - \sum_{j=1}^J w_j^* \mathbf{X}_j \approx \mathbf{0}$ and $\overline{\mathbf{X}} - \sum_{j=1}^J v_j^* \mathbf{X}_j \approx \mathbf{0}$, where $\mathbf{0}$ is a vector of zeros of the same dimension as \mathbf{X}_j . Suppose that units with $w_j^* > 0$ are assigned to treatment in the experiment, and units with $w_j^* = 0$ are kept untreated. A synthetic control estimator of τ_t is $\widehat{\tau}_t = \tau_t(\mathbf{w}^*, \mathbf{v}^*)$, i.e.,

$$\widehat{\tau}_t = \sum_{j=1}^J w_j^* Y_{jt} - \sum_{j=1}^J v_j^* Y_{jt}. \tag{5}$$

This estimator is based on approximations to equations (1) and (2) that rely on \mathbf{X}_j , the observed predictors of the potential outcomes, Y_{jt}^N , and Y_{jt}^I . Note that for every solution to (4) with $\underline{m} \leq \|\mathbf{v}\|_0 \leq \overline{m}$, there is another solution that swaps the roles of the treated and the untreated in the experiment without changing the objective value.

In what follows, we take the weight selector in (4) as a starting point for synthetic control designs, and modify it in several respects. A second formulation of the synthetic control design is based on equations (1) and (3),

$$\begin{aligned}
& \min_{\substack{w_1, \dots, w_J, \\ v_1, \dots, v_J}} \left\| \bar{\mathbf{X}} - \sum_{j=1}^J w_j \mathbf{X}_j \right\|^2 + \beta \left\| \sum_{j=1}^J w_j \mathbf{X}_j - \sum_{j=1}^J v_j \mathbf{X}_j \right\|^2 \\
& \text{s.t.} \quad \sum_{j=1}^J w_j = 1, \\
& \quad \quad \sum_{j=1}^J v_j = 1, \\
& \quad \quad w_j, v_j \geq 0, \quad \forall j = 1, \dots, J, \\
& \quad \quad w_j v_j = 0, \quad \forall j = 1, \dots, J, \\
& \quad \quad \underline{m} \leq \|\mathbf{w}\|_0 \leq \bar{m}.
\end{aligned} \tag{6}$$

The parameter $\beta > 0$ reflects the trade-off between selecting treated units to fit the aggregate value of the predictors $\bar{\mathbf{X}}$, and selecting control units to fit the aggregate value of the treated units. A small value of β favors designs with treated units that closely match $\bar{\mathbf{X}}$. A large value of β , instead, favors designs with aggregate treated and aggregate control units that closely match each other. While it is possible to use data-driven selectors of β , the rule of thumb $\beta = 1$ provides a natural choice that equally weights the two terms in the objective function in (6). For this formulation of the synthetic control design, treatment assignment and estimation follow the same procedures as with the previous one. Large values for β produce estimators that target the \mathbf{w} -weighted average effect of the treatment on the treated, τ_t^T of Doudchenko et al. (2021). Small values of β prioritize estimation of the average treatment effect, τ_t .

In our third formulation of the synthetic control design, the experimenter selects a synthetic treated unit to match the average values of the characteristics in the population. However, unlike the design in (6), the experimenter chooses multiple synthetic controls, one for each unit that contributes to the synthetic treated unit. For any J -dimensional vector of non-negative coordinates, $\mathbf{w} = (w_1, \dots, w_J)$, let $\mathcal{J}_{\mathbf{w}}$ be the set of the indices with non-zero coordinates,

$\mathcal{J}_{\mathbf{w}} = \{j : w_j > 0\}$. Our next version of the synthetic control design is:

$$\begin{aligned}
& \min_{\substack{w_j, \forall j=1,2,\dots,J, \\ v_{ij}, \forall i,j=1,2,\dots,J}} \left\| \bar{\mathbf{X}} - \sum_{j=1}^J w_j \mathbf{X}_j \right\|^2 + \xi \sum_{j=1}^J w_j \left\| \mathbf{X}_j - \sum_{i=1}^J v_{ij} \mathbf{X}_i \right\|^2 \\
& \text{s.t.} \quad \sum_{j=1}^J w_j = 1, \\
& \quad w_j \geq 0, \quad \forall j = 1, \dots, J, \\
& \quad \sum_{i=1}^J v_{ij} = 1, \quad \forall j \in \mathcal{J}_{\mathbf{w}}, \\
& \quad v_{ij} = 0, \quad \forall i \in \mathcal{J}_{\mathbf{w}}, j = 1, \dots, J, \\
& \quad v_{ij} \geq 0, \quad \forall j \in \mathcal{J}_{\mathbf{w}}, i = 1, \dots, J, \\
& \quad v_{ij} = 0, \quad \forall j \notin \mathcal{J}_{\mathbf{w}}, i = 1, \dots, J, \\
& \quad \underline{m} \leq \|\mathbf{w}\|_0 \leq \bar{m}.
\end{aligned} \tag{7}$$

The parameter $\xi > 0$ arbitrates potential trade-offs between selecting treated units to fit the aggregate value of the predictors $\bar{\mathbf{X}}$, and selecting control units to fit the values of the predictors for the treated units. A small value of ξ favors experimental designs with treated units that closely match $\bar{\mathbf{X}}$. A large value of ξ , instead, favors designs where the values of the predictors for the treated units are closely matched by their respective synthetic controls.

Let $\{w_j^*, v_{ij}^*\}_{i,j=1,\dots,J}$ be a solution of the optimization problem in (7). As before, we do not strictly require optimality of $\{w_j^*, v_{ij}^*\}_{i,j=1,\dots,J}$, provided $\{w_j^*, v_{ij}^*\}_{i,j=1,\dots,J}$ is feasible and $\bar{\mathbf{X}} - \sum_{j=1}^J w_j^* \mathbf{X}_j \approx \mathbf{0}$ and $\mathbf{X}_j - \sum_{i=1}^J v_{ij}^* \mathbf{X}_i \approx \mathbf{0}$ for all j such that $w_j^* > 0$. Assign units with $w_j^* > 0$ to treatment in the experiment, and keep units with $w_j^* = 0$ untreated. Let

$$v_j^* = \sum_{i=1}^J w_i^* v_{ij}^*. \tag{8}$$

Then,

$$\begin{aligned}
\hat{\tau}_t &= \sum_{j=1}^J w_j^* Y_{jt} - \sum_{j=1}^J v_j^* Y_{jt} \\
&= \sum_{j=1}^J w_j^* \left(Y_{jt} - \sum_{i=1}^J v_{ij}^* Y_{it} \right).
\end{aligned}$$

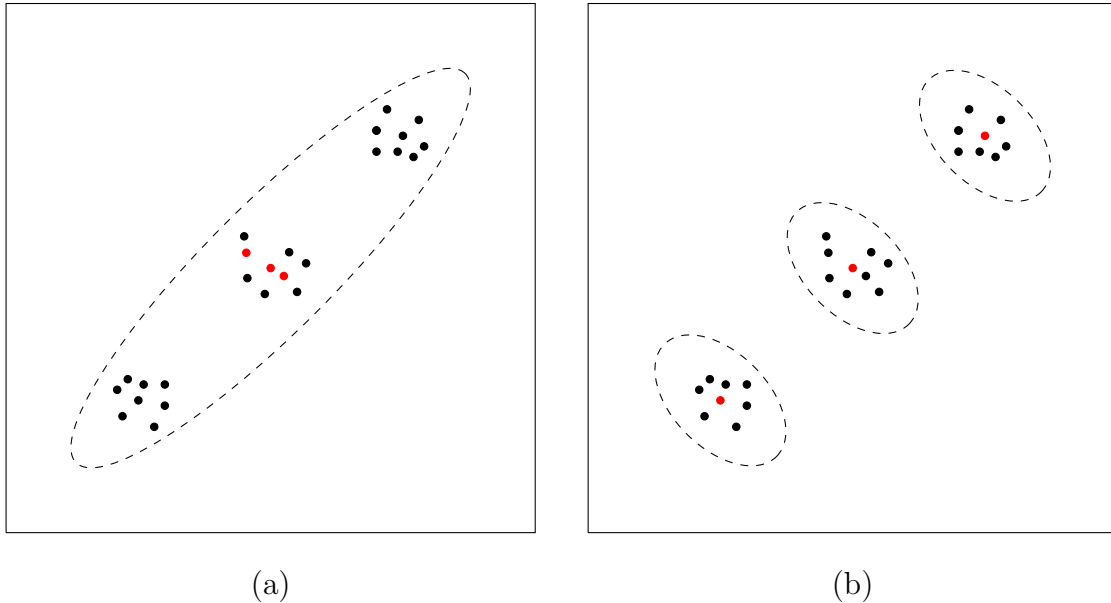


Figure 1: Clustering in a synthetic control design

Note: Panels (a) and (b) plot the values of the predictors in \mathbf{X}_j , which is bivariate in this simple example. Units assigned to treatment are drawn in red. In panel (a) we treat the entire sample as a single cluster. In panel (b) we divide the sample into three clusters and assign one unit in each cluster to the treatment.

Our next adjustment to the synthetic control design is motivated by settings where the units may be naturally divided in clusters with similar values in the predictors, $\mathbf{X}_1, \dots, \mathbf{X}_J$. For example, weather patterns, which may be highly dependent across cities in the same region (e.g., Northeast, Midwest, etc., in the US), may influence the seasonality of the demand for ride-sharing services. In those cases, it is natural to treat each cluster (each region, in our example) as a distinct experimental design to ameliorate interpolation biases. Figure 1 illustrates this point. Panels (a) and (b) depict identical samples in the space of the predictors. In this simple example, we have two predictors only, and their values for each of the units are represented by the coordinates of the points in the figure, which are the same in the two panels. Red dots represent units assigned to treatment. All other units are plotted as black dots. Panel (a) visualizes the result of treating the entire sample as one cluster. Three units are assigned to treatment. They closely reproduce the value of $\overline{\mathbf{X}}$, but they all fall in the same central cluster, far away from observations in other clusters. In panel (b), assignment to treatment takes into account

the clustered nature of the data, and one unit is treated per cluster. This provides a better approximation of the distribution of the predictor values for the entire sample, ameliorating concerns of interpolation biases.

Suppose we divide the set of J available units into K clusters. Let \mathcal{I}_k be the set of indices for the units in cluster k . The cluster mean is

$$\bar{\mathbf{X}}_k = \sum_{j \in \mathcal{I}_k} f_j \mathbf{X}_j / \sum_{j \in \mathcal{I}_k} f_j,$$

for each cluster $k = 1, \dots, K$. For each index $i = 1, \dots, J$, let $k(i)$ be the cluster to which unit i belongs, i.e., $i \in \mathcal{I}_{k(i)}$. A clustered version of the synthetic control design in (7) is given by:

$$\begin{aligned} \min_{\substack{w_j, \forall j=1,2,\dots,J, \\ v_{ij}, \forall i,j=1,2,\dots,J}} & \sum_{k=1}^K \left(\sum_{j \in \mathcal{I}_k} f_j \right) \left\{ \left\| \bar{\mathbf{X}}_k - \sum_{j \in \mathcal{I}_k} w_j \mathbf{X}_j \right\|^2 + \xi \sum_{j \in \mathcal{I}_k} w_j \left\| \mathbf{X}_j - \sum_{i,j \in \mathcal{I}_k} v_{ij} \mathbf{X}_i \right\|^2 \right\} \\ \text{s.t.} & \sum_{j \in \mathcal{I}_k} w_j = 1, \quad \forall k = 1, \dots, K, \\ & w_j \geq 0, \quad \forall j = 1, \dots, J, \\ & \sum_{i=1}^J v_{ij} = 1, \quad \forall j \in \mathcal{J}_w \\ & v_{ij} \geq 0, \quad \forall j \in \mathcal{J}_w, i = 1, \dots, J, \\ & v_{ij} = 0, \quad \forall j \notin \mathcal{J}_w, i = 1, \dots, J, \\ & v_{ij} = 0, \quad \forall i \in \mathcal{J}_w, j = 1, \dots, J, \\ & v_{ij} = 0, \quad \forall i, j, \text{ such that } k(i) \neq k(j), \\ & \underline{m} \leq \|\mathbf{w}\|_0 \leq \bar{m}. \end{aligned}$$

We conclude this section by discussing other possible extensions to the synthetic control design. First, it is well known that synthetic control estimators may not be unique. Lack of uniqueness is typical in settings where the values of the predictors that a synthetic control is targeting (i.e., $\bar{\mathbf{X}}$ in equation (4), or \mathbf{X}_j for a treated unit in equation (7)) fall inside the convex hull of the values of \mathbf{X}_j for the units in the donor pool. To address this issue we adapt the penalized estimator of Abadie and L'Hour (2021) to the synthetic control designs proposed in this article. The penalized synthetic control estimator of Abadie and L'Hour (2021) is unique

as long as predictor values for the units in the donor pool are in general quadratic position (see Abadie and L’Hour, 2021, for details). Moreover, penalized synthetic controls favor solutions where the synthetic units are composed of units that have predictor values, \mathbf{X}_j , similar to the target values. Applying the penalized synthetic control of Abadie and L’Hour (2021) to the objective function of (4), we obtain

$$\begin{aligned}
& \min_{\substack{w_1, \dots, w_J, \\ v_1, \dots, v_J}} \left\| \bar{\mathbf{X}} - \sum_{j=1}^J w_j \mathbf{X}_j \right\|^2 + \left\| \bar{\mathbf{X}} - \sum_{j=1}^J v_j \mathbf{X}_j \right\|^2 \\
& \quad + \lambda_1 \sum_{j=1}^J w_j \left\| \bar{\mathbf{X}} - \mathbf{X}_j \right\|^2 + \lambda_2 \sum_{j=1}^J v_j \left\| \bar{\mathbf{X}} - \mathbf{X}_j \right\|^2 \\
& \text{s.t.} \quad \sum_{j=1}^J w_j = 1, \\
& \quad \sum_{j=1}^J v_j = 1, \\
& \quad w_j, v_j \geq 0, \quad \forall j = 1, \dots, J, \\
& \quad w_j v_j = 0, \quad \forall j = 1, \dots, J, \\
& \quad \underline{m} \leq \|\mathbf{w}\|_0 \leq \bar{m}.
\end{aligned} \tag{9}$$

Here, λ_1 and λ_2 are positive constants that penalize discrepancies between the target values of the predictor $\bar{\mathbf{X}}$ and the values of the predictors for the units that contribute to their synthetic counterparts.⁷

Other types of penalization are possible. In particular, Doudchenko and Imbens (2016), Doudchenko et al. (2021) and others have proposed synthetic control estimators that use ridge or elastic net regularization on the synthetic control weights (e.g., on w_j and v_j in design (4)). The synthetic control designs proposed in this article can incorporate regularization on the weights.

Finally, Abadie and L’Hour (2021), Arkhangelsky et al. (2021), and Ben-Michael, Feller and Rothstein (2021) have proposed bias-correction techniques for synthetic control methods. Section OA.1 in the Online Appendix provides details on how to apply bias correction techniques

⁷See Abadie and L’Hour (2021) for details on penalized synthetic control estimators. The synthetic control design in (9) is a penalized version of (4). Section OA.1 in the Online Appendix discusses how to apply the Abadie and L’Hour penalty to the other synthetic designs proposed in this article.

in a synthetic control design.

3. Formal Results

We introduce an extension of the linear factor model commonly employed in the synthetic control literature, and use it to analyze the properties of estimators based on synthetic control designs.

Assumption 1 *Potential outcomes follow a linear factor model,*

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

$$Y_{jt}^I = v_t + \boldsymbol{\gamma}'_t \mathbf{Z}_j + \boldsymbol{\eta}'_t \boldsymbol{\mu}_j + \xi_{jt}, \quad (10b)$$

where \mathbf{Z}_j is a $(R \times 1)$ vector of observed covariates, $\boldsymbol{\theta}_t$ and $\boldsymbol{\gamma}_t$ are $(R \times 1)$ vectors of unknown parameters, $\boldsymbol{\mu}_j$ is a $(F \times 1)$ vector of unobserved covariates, $\boldsymbol{\lambda}_t$ and $\boldsymbol{\eta}_t$ are $(F \times 1)$ vectors of unknown parameters, and ϵ_{jt} and ξ_{jt} are unobserved random shocks.

Equation (10a) is the linear factor model for potential outcomes under no treatment that is commonly employed in the literature as a benchmark model to analyze the properties of synthetic control estimators (see, e.g., Abadie, Diamond and Hainmueller, 2010, Ferman, 2021). Equation (10b) extends the linear factor structure to potential outcomes under treatment. The reason is that, in contrast to synthetic control estimation with observational data, synthetic control designs require the choice of a treatment group in addition to the choice of a comparison group.

We employ the covariates in \mathbf{Z}_j as well as some pre-experimental values of the outcome variables Y_{jt} to construct the vectors of predictors, \mathbf{X}_j . In particular, let $\mathcal{E} \subseteq \{1, \dots, T_0\}$, let $T_{\mathcal{E}} = |\mathcal{E}|$, and let $\mathbf{Y}_j^{\mathcal{E}}$ be the $(T_{\mathcal{E}} \times 1)$ vector of $T_{\mathcal{E}}$ pre-experimental outcomes for unit j and time indices in \mathcal{E} . We define

$$\mathbf{X}_j = \begin{pmatrix} \mathbf{Y}_j^{\mathcal{E}} \\ \mathbf{Z}_j \end{pmatrix},$$

for $j = 1, \dots, J$. That is, the vector of predictors \mathbf{X}_j collects the covariates in \mathbf{Z}_j and the pre-experimental outcome values Y_{jt} for the “fitting periods” in \mathcal{E} . In practice, the values in \mathbf{X}_j are often scaled to make them independent of units of measurement, or to reflect the relative importance of each of the predictors (see, e.g., Abadie, 2021).

The next assumption gathers regularity conditions on model primitives.

Assumption 2

- (i) $F \leq T_{\mathcal{E}}$. Moreover, let $\boldsymbol{\lambda}_{\mathcal{E}}$ be the $(T_{\mathcal{E}} \times F)$ matrix with rows equal to the $\boldsymbol{\lambda}_t$'s indexed by \mathcal{E} . Let $\zeta_{\mathcal{E}}$ be the smallest eigenvalue of $\boldsymbol{\lambda}'_{\mathcal{E}}\boldsymbol{\lambda}_{\mathcal{E}}$. Then, $\underline{\zeta} = \zeta_{\mathcal{E}}/T_{\mathcal{E}} > 0$.
- (ii) For each $j = 1, \dots, J$, $\epsilon_{j1}, \dots, \epsilon_{jT}$ is a sequence of i.i.d. sub-Gaussian random variables with mean zero and variance proxy $\bar{\sigma}^2$. For any $j = 1, \dots, J$, $\xi_{jT_0+1}, \dots, \xi_{jT}$ is a sequence of i.i.d. sub-Gaussian random variables with mean zero, variance proxy $\bar{\sigma}^2$, and independent of $\epsilon_{j1}, \dots, \epsilon_{jT}$.

Assumption 2(i) is similar to conditions in Abadie, Diamond and Hainmueller (2010). Assumption 2(ii) is similar to conditions in Abadie, Diamond and Hainmueller (2010), Doudchenko and Imbens (2016), Chernozhukov, Wüthrich and Zhu (2021), and Arkhangelsky et al. (2021). Sub-Gaussianity is not strictly necessary, but it simplifies the form of our results. It can be relaxed by assuming bounded finite order moments (instead of bounding the entire moment generating function). Sub-Gaussianity is a relatively mild assumption. It holds for any Gaussian distribution, as well as any distribution with a bounded support. Distributions with heavy tails, such as the Cauchy distribution, are not sub-Gaussian. Notably, Assumption 2(ii) allows for dependence of ϵ_{jt} and ξ_{jt} across units.

Unless otherwise noted, all probability statements are over the joint distribution of ϵ_{jt} and ξ_{jt} and conditional on the values of the other components on the right-hand sides of equations (10a) and (10b). The next assumption pertains to the quality of the synthetic control fit. For concreteness, we focus on the base design in (4), and choose where $\boldsymbol{w}^* = (w_1^*, \dots, w_J^*)$ and $\boldsymbol{v}^* = (v_1^*, \dots, v_J^*)$ so that the synthetic treated and synthetic control units reproduce the average values of \boldsymbol{X}_j .

Assumption 3 *With probability one, (i)*

$$\sum_{j=1}^J w_j^* \boldsymbol{Z}_j = \sum_{j=1}^J v_j^* \boldsymbol{Z}_j = \sum_{j=1}^J f_j \boldsymbol{Z}_j,$$

and (ii)

$$\sum_{j=1}^J w_j^* \boldsymbol{Y}_j^{\mathcal{E}} = \sum_{j=1}^J v_j^* \boldsymbol{Y}_j^{\mathcal{E}} = \sum_{j=1}^J f_j \boldsymbol{Y}_j^{\mathcal{E}}.$$

Assumption 3 implies that the synthetic treated and control units defined by \mathbf{w}^* and \mathbf{v}^* provide a perfect fit for $\bar{\mathbf{X}}$. Assumption 3 is a strong restriction, which may only hold approximately in practice. The next assumption relaxes the perfect fit condition in Assumption 3.

Assumption 4 *There exists a positive constant $d > 0$, such that with probability one, (i)*

$$\left\| \sum_{j=1}^J w_j^* \mathbf{Z}_j - \sum_{j=1}^J f_j \mathbf{Z}_j \right\|_2^2 \leq Rd^2, \quad \left\| \sum_{j=1}^J v_j^* \mathbf{Z}_j - \sum_{j=1}^J f_j \mathbf{Z}_j \right\|_2^2 \leq Rd^2,$$

and (ii)

$$\left\| \sum_{j=1}^J w_j^* \mathbf{Y}_j^\mathcal{E} - \sum_{j=1}^J f_j \mathbf{Y}_j^\mathcal{E} \right\|_2^2 \leq T_\mathcal{E} d^2, \quad \left\| \sum_{j=1}^J v_j^* \mathbf{Y}_j^\mathcal{E} - \sum_{j=1}^J f_j \mathbf{Y}_j^\mathcal{E} \right\|_2^2 \leq T_\mathcal{E} d^2.$$

Let $\lambda_{t,f}$ be the f -th coordinate of $\boldsymbol{\lambda}_t$, and

$$\bar{\lambda} = \max_{\substack{t=1,\dots,T \\ f=1,\dots,F}} |\lambda_{t,f}|.$$

We define η_{tf} , θ_{tr} , γ_{tr} , $\bar{\eta}$, $\bar{\theta}$ and $\bar{\gamma}$ analogously, so $|\eta_{tf}| \leq \bar{\eta}$ for $t = T_0 + 1, \dots, T$, $f = 1, \dots, F$, $|\theta_{tr}| \leq \bar{\theta}$ for $t = 1, \dots, T$, $r = 1, \dots, R$, and $|\gamma_{tr}| \leq \bar{\gamma}$, for $t = T_0 + 1, \dots, T$, $r = 1, \dots, R$. The next theorem provides a bias bound for estimators based on synthetic control designs.

Theorem 1 *If Assumptions 1 – 3 hold, then for any $t \geq T_0 + 1$,*

$$|E[\hat{\tau}_t - \tau_t]| \leq \frac{\bar{\lambda}(\bar{\eta} + \bar{\lambda})F}{\underline{\zeta}} \sqrt{2 \log(2J)} \frac{\bar{\sigma}}{\sqrt{T_\mathcal{E}}}. \quad (13)$$

If Assumptions 1, 2, and 4 hold, then for any $t \geq T_0 + 1$,

$$|E[\hat{\tau}_t - \tau_t]| \leq \left((\bar{\gamma} + \bar{\theta})R + \frac{\bar{\lambda}(\bar{\eta} + \bar{\lambda})F}{\underline{\zeta}} (1 + \bar{\theta}R) \right) d + \frac{\bar{\lambda}(\bar{\eta} + \bar{\lambda})F}{\underline{\zeta}} \sqrt{2 \log(2J)} \frac{\bar{\sigma}}{\sqrt{T_\mathcal{E}}}. \quad (14)$$

Note that, while the factor model in equations (10a) and (10b) leave the sign and scale of $\boldsymbol{\lambda}_t$ and $\boldsymbol{\eta}_t$ free (e.g., multiplying $\boldsymbol{\lambda}_t$ and dividing $\boldsymbol{\mu}_t$ by the same non-zero constant does not change the value of $\boldsymbol{\lambda}'_t \boldsymbol{\mu}_t$), the value of the bound in Theorem 1 is invariant to changes in the sign or the scale of $\boldsymbol{\lambda}_t$ and $\boldsymbol{\eta}_t$. Moreover, the bound in (14) does not depend on the scale of \mathbf{Z}_j , because changing the scale of \mathbf{Z}_j leaves the product $\bar{\theta}d$ unchanged. The scale of Y_{jt} does affect

the bound in (14) because the treatment effect τ_t is measured in the same units as Y_{jt} . The results in Theorem 1 do not depend on the specific formulation of the synthetic control design (e.g., *Constrained* vs. *Unconstrained*).

The bias bounds (13) and (14) depend on the ratio between the scale of ϵ_{jt} , represented by $\bar{\sigma}$, and the number of fitting periods $T_{\mathcal{E}}$. Intuitively, the bias of the synthetic control estimator is small when a good fit in pre-experimental outcomes (Assumption 3) is obtained by implicitly fitting the values of the latent variables, μ_j . Overfitting happens when pre-experimental outcomes are instead fitted out of the variability in the individual transitory shocks, ϵ_{jt} . A small number of fitting periods $T_{\mathcal{E}}$ combined with large variability in ϵ_{jt} increases the risk of overfitting and, as a result, increases the bias bound. Similarly, for any fixed value of $T_{\mathcal{E}}$, the bias bound increases with J , reflecting the increased risk of over-fitting caused by increased variability in ϵ_{jt} over larger donor pools. Finally, the number of unobserved factors F enters the bound (13) linearly, which highlights the importance of including the observed predictors \mathbf{Z}_j — other than pre-experimental outcomes — in the vector of fitting variables \mathbf{X}_j . Under the factor model in equations (10a) and (10b), observed predictors not included in \mathbf{Z}_j are shifted to $\boldsymbol{\mu}_j$, increasing F and the magnitude of the bound.⁸

We next turn our attention to inference. We utilize a set of “blank periods,” $\mathcal{B} \subseteq \{1, \dots, T_0\} \setminus \mathcal{E}$, which comprise pre-experimental periods whose outcomes Y_{jt} have not been used to calculate \mathbf{w}^* or \mathbf{v}^* . Because pre-experimental periods that are not in \mathcal{E} or \mathcal{B} are not used in our procedure, without loss of generality we consider $\mathcal{B} = \{1, \dots, T_0\} \setminus \mathcal{E}$. We therefore assume that the number of elements of \mathcal{B} is $T_{\mathcal{B}} = |\mathcal{B}| = T_0 - T_{\mathcal{E}}$. We aim to test the null hypothesis:

For $t = T_0 + 1, \dots, T$, and $j = 1, \dots, J$,

$$Y_{jt}^I = \delta_t + \boldsymbol{\theta}'_t \mathbf{Z}_j + \boldsymbol{\lambda}'_t \boldsymbol{\mu}_j + \xi_{jt}, \quad (15)$$

⁸Shifting predictors from \mathbf{Z}_j to $\boldsymbol{\mu}_j$ changes the bias bound (13) in a more complex manner than what might be inferred from a cursory look at the bias formula. First, moving predictors from \mathbf{Z}_j to $\boldsymbol{\mu}_j$ also means shifting components of $\boldsymbol{\theta}_t$ to $\boldsymbol{\lambda}_t$, which can change the value of $\underline{\zeta}$. Poincaré’s separation theorem implies that $\underline{\zeta}$ cannot increase as a result of this shift. Moreover, moving predictors from \mathbf{Z}_j to $\boldsymbol{\mu}_j$ cannot decrease the values of $\bar{\lambda}$ and $\bar{\eta}$. Overall, the value of the bias bound in (13) cannot decrease and will typically increase by moving predictors from \mathbf{Z}_j to $\boldsymbol{\mu}_j$. This is not necessarily true for the bound in (14), because a shift of components from \mathbf{Z}_j to $\boldsymbol{\mu}_j$ decreases the value of R .

where ξ_{jt} has the same distribution as ϵ_{jt} .

Under the null hypothesis in (15), the distribution of Y_{jt}^I is the same as the distribution of Y_{jt}^N , for $t = T_0 + 1, \dots, T$, and $j = 1, \dots, J$. But the realized values of Y_{jt}^I and Y_{jt}^N may differ.

Recall from (5) that, for $t \in \{T_0 + 1, \dots, T\}$, a synthetic control estimator is defined as

$$\widehat{\tau}_t = \sum_{j=1}^J w_j^* Y_{jt} - \sum_{j=1}^J v_j^* Y_{jt}.$$

Let $\widehat{u}_t = \widehat{\tau}_t, \forall t \in \{T_0 + 1, \dots, T\}$ be the synthetic control estimator in the experimental periods. Similarly, for each $t \in \mathcal{B}$ in the blank periods, let

$$\widehat{u}_t = \sum_{j=1}^J w_j^* Y_{jt} - \sum_{j=1}^J v_j^* Y_{jt}.$$

Such \widehat{u}_t for $t \in \mathcal{B}$ are “placebo” treatment effects estimated for the blank periods. We study the properties of a test based on combinations from the set $\{\widehat{u}_t : t \in \mathcal{B} \cup \{T_0 + 1, \dots, T\}\}$.

We define Π as the set of all $(T - T_0)$ -combinations of $\mathcal{B} \cup \{T_0 + 1, \dots, T\}$. That is, for each $\pi \in \Pi$, π is a subset of indices from the blank periods and the experimental periods $\mathcal{B} \cup \{T_0 + 1, \dots, T\}$, such that $|\pi| = T - T_0$. The cardinality of Π is $|\Pi| = (T - T_\mathcal{E})! / ((T - T_0)!(T_0 - T_\mathcal{E})!)$. For each $\pi \in \Pi$, let $\pi(i)$ be the i^{th} smallest value in π , and

$$\widehat{\mathbf{e}}_\pi = (\widehat{u}_{\pi(1)}, \widehat{u}_{\pi(2)}, \dots, \widehat{u}_{\pi(T-T_0)}).$$

In addition, let $\widehat{\mathbf{e}} = (\widehat{u}_{T_0+1}, \dots, \widehat{u}_T) = (\widehat{\tau}_{T_0+1}, \dots, \widehat{\tau}_T)$. This is a vector of treatment effect estimates from the experimental periods. For any $(T - T_0)$ -dimensional vector $\mathbf{e} = (e_1, \dots, e_{T-T_0})$, we adopt the test statistic,

$$S(\mathbf{e}) = \frac{1}{T - T_0} \sum_{t=1}^{T-T_0} |e_t|. \quad (16)$$

Other choices of test statistics are possible, such as those based on an L_p -norm of \mathbf{e} (Chernozhukov, Wüthrich and Zhu, 2021) and one-sided versions of the resulting test statistics (i.e., with e_t or $-e_t$ replacing $|e_t|$ in equation (16)).

The p -value of a permutation test on (16) is

$$\hat{p} = \frac{1}{|\Pi|} \sum_{\pi \in \Pi} \mathbb{1}\{S(\hat{\mathbf{e}}_\pi) \geq S(\hat{\mathbf{e}})\} \quad (17)$$

Theorem 2 below shows that if $\boldsymbol{\lambda}_t$ are exchangeable random variables for $t \in \mathcal{B} \cup \{T_0+1, \dots, T\}$, then a test of the null hypothesis in (15) based on the p -value in (17) is exact.

Theorem 2 *Suppose that Assumptions 1 and 3(i) hold. Assume that $\{\boldsymbol{\lambda}_t\}_{t \in \mathcal{B} \cup \{T_0+1, \dots, T\}}$ is a sequence of exchangeable random variables independent of $\{\epsilon_{jt}\}_{t \in \mathcal{B} \cup \{T_0+1, \dots, T\}}$ and $\{\xi_{jt}\}_{t \in \{T_0+1, \dots, T\}}$. Assume also that for each $j = 1, \dots, J$, $\{\epsilon_{jt}\}_{t \in \mathcal{B} \cup \{T_0+1, \dots, T\}}$ and $\{\xi_{jt}\}_{t \in \{T_0+1, \dots, T\}}$ are two sequences of exchangeable random variables, respectively. Then under the null hypothesis (15), we have*

$$\alpha - \frac{1}{|\Pi|} \leq \Pr(\hat{p} \leq \alpha) \leq \alpha,$$

for any $\alpha \in [0, 1]$, where $\Pr(\hat{p} \leq \alpha)$ is taken over the distribution of $\{\xi_{jt}, \epsilon_{jt}, \boldsymbol{\lambda}_t\}$.

If Assumption 1 holds but Assumption 3(i) is violated, the same result holds if $\{(\boldsymbol{\theta}_t, \boldsymbol{\lambda}_t)\}_{t \in \mathcal{B} \cup \{T_0+1, \dots, T\}}$ is a sequence of exchangeable random variables independent of $\{\epsilon_{jt}\}_{t \in \mathcal{B} \cup \{T_0+1, \dots, T\}}$ and $\{\xi_{jt}\}_{t \in \{T_0+1, \dots, T\}}$, and if for each $j = 1, \dots, J$, $\{\epsilon_{jt}\}_{t \in \mathcal{B} \cup \{T_0+1, \dots, T\}}$ and $\{\xi_{jt}\}_{t \in \{T_0+1, \dots, T\}}$ are two sequences of exchangeable random variables, respectively. Here, $\Pr(\hat{p} \leq \alpha)$ is taken over the distribution of $\{\xi_{jt}, \epsilon_{jt}, \boldsymbol{\lambda}_t, \boldsymbol{\theta}_t\}$.

Note that, under the assumptions of Theorem 2, the potential outcome series Y_{jt}^N is allowed to be non-stationary through the term $\delta_t + \boldsymbol{\theta}'_t \mathbf{Z}_j$ in equation (10a). This is in contrast to a related result in Doudchenko et al. (2021), which requires that the potential outcomes Y_{jt}^N are i.i.d. over time. Exchangeability of $\boldsymbol{\lambda}_t$ is a strong restriction. Theorem OA.1 in Section OA.2 in the Online Appendix relaxes this restriction by showing that conditional on $\boldsymbol{\lambda}_t$ (i.e., without resorting to exchangeability of $\boldsymbol{\lambda}_t$), the p -value in (17) is still approximately valid for large $T_\mathcal{E}$.

In some settings, the number of possible combinations, $|\Pi|$, could be very large, making exact calculation of \hat{p} computationally expensive. In those instances, random samples of the combinations in Π can be used to approximate the p -value in equation (17).

The inferential technique proposed in this article is related to, but distinct from, the permutation methods in Abadie, Diamond and Hainmueller (2010), Chernozhukov, Wüthrich and Zhu

(2019), Chernozhukov, Wüthrich and Zhu (2021), Firpo and Possebom (2018), Lei and Candès (2021), and others. Inferential methods that reassign treatment across units (e.g., Abadie, Diamond and Hainmueller, 2010) are not appropriate for the designs of Section 2, which explicitly selects treated and control units to satisfy an optimality criterion.

Similar to Chernozhukov, Wüthrich and Zhu (2021), our method is based on rearrangements of estimated treatment effects across time periods. But unlike Chernozhukov, Wüthrich and Zhu (2021), which proposes permutations over all periods including the pre-intervention periods, our inferential method permutes only over the blank periods and post-intervention periods, which are not used to estimate the weights in the synthetic control design. Relative to Chernozhukov, Wüthrich and Zhu (2021), the generative models of equations (10a) and (10b), which allow for unobserved factors, and the finite sample nature of the results require a novel testing procedure that, similar to split conformal prediction methods (Lei et al., 2018, Vovk, Gammerman and Shafer, 2005), takes advantage of the availability of blank periods.

4. Empirical Illustration Using Walmart Data

In this section, we illustrate the applicability of the methods in this article using store-level data from Walmart (Prakash, 2023). The data set is a balanced panel of weekly sales for $J = 45$ Walmart stores and $T = 143$ weeks, between the week of February 5, 2010 and the week of October 26, 2012. We estimate the effect of a placebo intervention and show that, in the presence of a good pre-intervention fit, the methods of Section 3 produce point estimates that are close to zero and a test result that does not reject the sharp null hypothesis in (15).

We consider the design of a fictitious experiment across stores implemented on July 20, 2012 (week 129 in the data). Out of the $T_0 = 128$ pre-experimental weeks, we take the first $T_{\mathcal{E}} = 100$ weeks as the fitting period, and the last $(T_0 - T_{\mathcal{E}}) = 28$ weeks as the blank period. The number of weeks in the experimental period is $T - T_0 = 15$. The outcomes $\{Y_{jt}\}_{j=1,\dots,J,t=1,\dots,T}$ are weekly sales (units of revenue are undisclosed in the data). We use uniform weights $f_j = 1/J$ for $j = 1, \dots, J$, to average sales across all stores.

We compute synthetic treated and control units that apply the *Constrained* formulation in (4) with $\bar{m} = 2$. We adopt $\bar{m} = 2$ because using only one store for the synthetic treated fails




Figure 2: Synthetic Treatment Unit and Synthetic Control Unit, $\bar{m} = 2$

Note: The black solid line represents the synthetic treated outcome. The black dashed line represents the synthetic control outcome. The blue dashed lines are individual stores' sales.

to produce a good fit between the resulting synthetic treated and synthetic control units during the fitting period. Increasing \bar{m} to $\bar{m} = 3$ brings only marginal improvements in fit. Figures 2 and 3 report the results when $\bar{m} = 2$. Section OA.5 in the Online Appendix reports the results for $\bar{m} = 1$ and $\bar{m} = 3$.

Figure 2 reports the time-series of weekly sales for the synthetic treated (black solid line) and synthetic control (black dashed line) units, as well as for each individual store in the data set (blue dashed lines). Weekly sales for the synthetic treated and the synthetic control units closely follow each other during the fitting period. The gap between the two synthetic units remains small following the fitting period, which demonstrates good out-of-sample predictive power in the absence of an intervention.

Figure 3 reports the difference in weekly sales between the synthetic treated and the synthetic control units. The p -value of equation (17), calculated over the residuals of Figure 3, is equal to 0.933, which results in a failure to reject the null hypothesis (15).




Figure 3: Treatment Effect Estimate, when $\bar{m} = 2$.

Note: This figure reports the difference between the synthetic treated and synthetic control outcomes of Figure 2. For the experimental periods, this is the treatment effect estimate.

5. Simulation Study

In this section, we report the results of a simulation study that showcases the behavior of the estimators proposed in this article. We consider a setting with $J = 15$ units, $R = 7$ observable covariates and $F = 11$ unobservable covariates. We simulate data for $T = 30$ periods in total, with $T_0 = 25$ pre-experimental periods, and $T - T_0 = 5$ experimental or post-intervention periods. We compute weights during the first $T_{\mathcal{E}} = 20$ periods, and leave periods $t = 21, \dots, 25$ as blank periods.

For our baseline simulation design, we use the factor model in Assumption 1 to generate potential outcomes. For $t = 1, \dots, T$, we generate the series δ_t and v_t as small-to-large rearrangements of T i.i.d. Uniform $(0, 20)$ random variables. For $j = 1, \dots, J$, we set both \mathbf{Z}_j and $\boldsymbol{\mu}_j$ to be random vectors of i.i.d. Uniform $(0, 1)$ random variables. For $t = 1, \dots, T$, we set $\boldsymbol{\theta}_t$, $\boldsymbol{\gamma}_t$, $\boldsymbol{\lambda}_t$, and $\boldsymbol{\eta}_t$ to be random vectors of i.i.d. Uniform $(0, 10)$ random variables. Finally, for $j = 1, \dots, J$, and any $t = 1, \dots, T$, we set ϵ_{jt} and ξ_{jt} to be i.i.d. Normal $(0, \sigma^2)$ random variables, with $\sigma^2 = 1$. In Online Appendix OA.6 we present additional simulation results of alternative values of the noise parameter σ^2 .




Figure 4: Synthetic Treatment Unit and Synthetic Control Unit, $\sigma^2 = 1$

Note: The black solid line represents the synthetic treated outcome (w^* -weighted). The black dashed line represents the synthetic control outcome (v^* -weighted).

5.1. Results for a Single Simulation

Using the data generating process described above, we draw a single sample and conduct the synthetic control design in (4), with parameters $\underline{m} = 1$ and $\bar{m} = 14$, i.e., no constraint on the number of treated units. We report the results in Figures 4 and 5. In Figure 4, each blue dashed line represents an outcome trajectory Y_{jt} , for $t = 1, \dots, T$ and $j = 1, \dots, J$. The black solid line represents the trajectory of the synthetic treated unit $\sum_{j=1}^J w_j^* Y_{jt}$, for $t = 1, \dots, T$. The black dashed line represents the trajectory of the synthetic control unit $\sum_{j=1}^J v_j^* Y_{jt}$, for $t = 1, \dots, T$. The synthetic treated and synthetic control units closely track each other in the pre-experimental periods. They diverge in the experimental periods, when a treatment effect emerges as a result of the differences in the parameters of the data generating processes for Y_{jt}^N and Y_{jt}^I . Figure 5 reports the difference between the synthetic treated and the synthetic control outcomes. The inferential procedure of Section 3 produces p -value equal to 0.004 for the null hypothesis of no treatment effect in (15).




Figure 5: Treatment Effect Estimate, when $\sigma^2 = 1$.

Note: This figure reports the difference between the synthetic treated and synthetic control outcomes of Figure 4. For the experimental periods, this is the treatment effect estimate.

5.2. Performance Across Multiple Simulations

In this section, we compare the performance of different synthetic control designs over 1000 simulations that independently generate the model primitives (i.e., the factor loadings, covariates, and error terms) of Assumption 1. The data generating process is the same as in Section 5.1.

We consider five varieties of the synthetic control design:

1. *Unconstrained design:* This is the design in (4) without a cardinality constraint, so $\underline{m} = 1$ and $\overline{m} = J - 1 = 14$.
2. *Constrained design:* Same as the design in (4), but with $\underline{m} = 1$ and $\overline{m} = 1, \dots, 7$.
3. *Weakly-targeted design:* This is the design in (6). We vary β from 0.01 to 100.
4. *Unit-level design:* This is the design in (7), which fits a different synthetic control to each unit assigned to treatment. We vary ξ from 0.01 to 100.
5. *Penalized design:* This is the design in (9), with $\lambda = \lambda_1 = \lambda_2$. We vary λ from 0.01 to 100.

The *Constrained* design imposes sparsity in the synthetic treatment weights through a hard cardinality constraint specified by the integer \overline{m} . For the *Unit-level* design, large values of ξ generate sparsity in the synthetic treated weights. A sufficiently large value of ξ produces a

Unit-level design where the only treated unit is the unit that can be most closely fitted by a convex combination of the other units. For large values of λ , the *Penalized* design behaves like a one-to-one matching design, assigning all the weight to one treated and one control unit.

For the *Unit-level* design, synthetic control weights are aggregated as in (8). For the *Unconstrained* and *Penalized* designs, the synthetic treated and synthetic control weights can always be swapped without changing the objective values for their respective designs. For the *Constrained* design, the weights can be swapped when $\|\mathbf{v}^*\|_0 \leq \bar{m}$. When it is possible to swap synthetic treated and synthetic control weights, we choose the treated units so that the number of units with positive weights in \mathbf{w}^* is smaller than the number of units with positive weights in \mathbf{v}^* . When $\|\mathbf{w}^*\|_0 = \|\mathbf{v}^*\|_0$, we determine whether or not to swap following a specific rule described in Section B in the Appendix.

5.2.1. Average Treatment Effects

The first panel of Table 1 reports average treatment effects, τ_t , averaged over 1000 simulations. The second panel reports estimates of the average treatment effects, mean absolute error, mean squared error, and p -value, all averaged over 1000 simulations, as well as rejection rates. Mean absolute error (MAE) and mean squared error (MSE) are defined as

$$MAE = \frac{1}{T - T_0} \sum_{t=T_0+1}^T |\hat{\tau}_t - \tau_t|, \quad MSE = \frac{1}{T - T_0} \sum_{t=T_0+1}^T (\hat{\tau}_t - \tau_t)^2, \quad (18)$$

and the p -value is defined as in (17). Because the treatment effect is not equal to zero in the simulation of Table 1, smaller p -values and larger rejection rates reflect better performance of the testing procedure for a particular design.

In Table 1, the *Unconstrained* design has a strong relative performance. The performance of the *Constrained* design improves for larger \bar{m} , and is virtually identical to the performance of the *Unconstrained* design when $\bar{m} = 7$. The performance of *Weakly-targeted* and *Unit-level* designs is best when β and ξ take intermediate values. The *Penalized* design produces results similar to the *Unconstrained* design for small values of the penalization parameter λ .

Table 1: Average Treatment Effects (Averages over 1000 Simulations)

		τ_t								
		$t = 26$	$t = 27$	$t = 28$	$t = 29$	$t = 30$				
		-13.05	-10.92	-7.65	-5.17	-2.22				
		$\hat{\tau}_t$					MAE	MSE	\hat{p}	$\hat{p} < 0.05$
		$t = 26$	$t = 27$	$t = 28$	$t = 29$	$t = 30$				
<i>Unconstrained</i>		-13.03	-10.93	-7.64	-5.14	-2.17	0.96	1.45	0.016	0.920
<i>Constrained</i>	$\bar{m} = 1$	-12.89	-10.86	-7.58	-5.29	-2.32	2.93	13.29	0.065	0.642
	$\bar{m} = 2$	-13.04	-10.94	-7.65	-5.19	-2.21	1.69	4.48	0.029	0.815
	$\bar{m} = 3$	-13.04	-10.87	-7.64	-5.14	-2.09	1.28	2.56	0.020	0.888
	$\bar{m} = 4$	-13.05	-10.89	-7.66	-5.16	-2.13	1.10	1.92	0.017	0.915
	$\bar{m} = 5$	-12.98	-10.92	-7.64	-5.13	-2.17	1.02	1.65	0.017	0.915
	$\bar{m} = 6$	-13.05	-10.91	-7.61	-5.09	-2.16	0.97	1.48	0.016	0.916
	$\bar{m} = 7$	-13.04	-10.92	-7.64	-5.13	-2.18	0.94	1.41	0.016	0.926
<i>Weakly-targeted</i>	$\beta = 0.01$	-13.14	-10.97	-7.86	-5.24	-2.39	1.38	3.00	0.022	0.889
	$\beta = 0.1$	-13.06	-10.87	-7.68	-5.16	-2.24	1.04	1.73	0.016	0.934
	$\beta = 1$	-13.13	-10.90	-7.67	-5.20	-2.23	1.02	1.64	0.014	0.940
	$\beta = 10$	-13.15	-10.89	-7.72	-5.24	-2.26	1.11	1.93	0.014	0.941
	$\beta = 100$	-13.15	-10.92	-7.68	-5.23	-2.24	1.21	2.30	0.013	0.948
<i>Unit-level</i>	$\xi = 0.01$	-13.02	-10.85	-7.72	-5.11	-2.35	1.25	2.52	0.020	0.908
	$\xi = 0.1$	-13.08	-10.86	-7.76	-5.13	-2.28	1.07	1.85	0.016	0.920
	$\xi = 1$	-13.00	-10.82	-7.67	-5.09	-2.25	1.38	3.02	0.021	0.887
	$\xi = 10$	-13.05	-10.94	-7.73	-5.18	-2.21	2.23	8.02	0.031	0.816
	$\xi = 100$	-13.05	-10.89	-7.67	-5.10	-2.15	2.86	12.90	0.039	0.775
<i>Penalized</i>	$\lambda = 0.01$	-13.01	-10.91	-7.66	-5.13	-2.23	1.00	1.56	0.016	0.929
	$\lambda = 0.1$	-13.03	-10.86	-7.64	-5.12	-2.23	1.26	2.52	0.019	0.914
	$\lambda = 1$	-12.95	-10.85	-7.64	-5.06	-2.12	2.14	7.36	0.040	0.771
	$\lambda = 10$	-12.94	-10.74	-7.59	-5.16	-2.02	3.69	21.41	0.092	0.533
	$\lambda = 100$	-12.98	-10.68	-7.60	-5.26	-1.99	4.16	26.95	0.114	0.470

Note: Unless otherwise noted, all designs use $m = 1$ and $\bar{m} = 14$.

5.2.2. Average Treatment Effects on the Treated

The first five columns in Table 2 report averages of τ_t^T , the average effects of treatment on the treated units. These quantities depend on the weights for the treated units, which are different across different formulations of the synthetic control design. The next five columns report averages of $\hat{\tau}_t$. They are the same as in Table 1, yet we use them as estimators for τ_t^T in Table 2. The last two columns of Table 2 report averages across simulations of the mean absolute error and the mean squared error, defined as in (18) but with τ_t^T replacing τ_t .

The results in Table 2 are qualitatively similar to those for τ_t in Table 1, with one notable exception. As expected, for intermediate and large values of β , the *Weakly-targeted* design outperforms the other τ_t designs when the goal is to estimate τ_t^T . This is because in the *Weakly-targeted* design the synthetic control weights are targeted to τ_t^T (and more so as β becomes large).

Table 2: Average Treatment Effects on the Treated (Averages over 1000 Simulations)

	$\hat{\tau}_t$							$\hat{\tau}_t$							MAE^T	MSE^T	
	$t = 26$	$t = 27$	$t = 28$	$t = 29$	$t = 30$	$t = 26$	$t = 27$	$t = 28$	$t = 29$	$t = 30$	$t = 26$	$t = 27$	$t = 28$	$t = 29$	$t = 30$		
<i>Unconstrained</i>																	
<i>Constrained</i>																	
$\bar{m} = 1$	-13.05	-10.88	-7.65	-5.17	-2.24	-13.03	-10.93	-7.64	-5.14	-2.17	1.13	2.02					
$\bar{m} = 2$	-12.81	-10.74	-7.63	-5.12	-2.38	-12.89	-10.86	-7.58	-5.29	-2.32	3.04	14.62					
$\bar{m} = 3$	-13.10	-10.86	-7.69	-5.31	-2.33	-13.04	-10.94	-7.65	-5.19	-2.21	1.85	5.39					
$\bar{m} = 4$	-13.02	-10.91	-7.71	-5.22	-2.14	-13.04	-10.87	-7.64	-5.14	-2.09	1.47	3.39					
$\bar{m} = 5$	-13.03	-10.87	-7.66	-5.17	-2.12	-13.05	-10.89	-7.66	-5.16	-2.13	1.24	2.47					
$\bar{m} = 6$	-13.08	-10.87	-7.62	-5.17	-2.19	-12.98	-10.92	-7.64	-5.13	-2.17	1.19	2.25					
$\bar{m} = 7$	-13.11	-10.91	-7.59	-5.18	-2.21	-13.05	-10.91	-7.61	-5.09	-2.16	1.14	2.05					
$\bar{m} = 7$	-13.04	-10.89	-7.66	-5.18	-2.25	-13.04	-10.92	-7.64	-5.13	-2.18	1.12	1.98					
<i>Weakly-targeted</i>																	
$\beta = 0.01$	-13.02	-10.91	-7.64	-5.16	-2.22	-13.14	-10.97	-7.86	-5.24	-2.39	1.48	3.43					
$\beta = 0.1$	-13.03	-10.90	-7.66	-5.15	-2.22	-13.06	-10.87	-7.68	-5.16	-2.24	1.17	2.17					
$\beta = 1$	-13.05	-10.91	-7.62	-5.13	-2.25	-13.13	-10.90	-7.67	-5.20	-2.23	1.08	1.84					
$\beta = 10$	-13.05	-10.88	-7.65	-5.12	-2.22	-13.15	-10.89	-7.72	-5.24	-2.26	1.03	1.64					
$\beta = 100$	-13.04	-10.94	-7.65	-5.17	-2.21	-13.15	-10.92	-7.68	-5.23	-2.24	1.02	1.67					
<i>Unit-level</i>																	
$\xi = 0.01$	-13.04	-10.90	-7.67	-5.16	-2.24	-13.02	-10.85	-7.72	-5.11	-2.35	1.33	2.84					
$\xi = 0.1$	-13.04	-10.87	-7.68	-5.18	-2.26	-13.08	-10.86	-7.76	-5.13	-2.28	1.18	2.25					
$\xi = 1$	-13.11	-10.86	-7.72	-5.17	-2.28	-13.00	-10.82	-7.67	-5.09	-2.25	1.43	3.21					
$\xi = 10$	-13.06	-10.77	-7.62	-5.23	-2.19	-13.05	-10.94	-7.73	-5.18	-2.21	1.96	6.16					
$\xi = 100$	-13.05	-10.74	-7.59	-5.17	-2.21	-13.05	-10.89	-7.67	-5.10	-2.15	2.35	8.79					
<i>Penalized</i>																	
$\lambda = 0.01$	-13.05	-10.90	-7.63	-5.18	-2.24	-13.01	-10.91	-7.66	-5.13	-2.23	1.16	2.12					
$\lambda = 0.1$	-13.03	-10.81	-7.62	-5.11	-2.19	-13.03	-10.86	-7.64	-5.12	-2.23	1.38	3.02					
$\lambda = 1$	-13.02	-10.75	-7.64	-5.14	-2.25	-12.95	-10.85	-7.64	-5.06	-2.12	2.23	7.99					
$\lambda = 10$	-12.87	-10.86	-7.63	-5.22	-2.28	-12.94	-10.74	-7.59	-5.16	-2.02	3.80	23.09					
$\lambda = 100$	-12.85	-10.85	-7.62	-5.20	-2.30	-12.98	-10.68	-7.60	-5.26	-1.99	4.22	28.53					

Table 3: Average Treatment Effects, Nonlinear Model (Averages over 1000 Simulations)

		τ_t								
		$t = 26$	$t = 27$	$t = 28$	$t = 29$	$t = 30$				
		-13.18	-10.72	-7.96	-5.47	-2.43				
		$\hat{\tau}_t$					<i>MAE</i>	<i>MSE</i>	\hat{p}	$\hat{p} < 0.05$
		$t = 26$	$t = 27$	$t = 28$	$t = 29$	$t = 30$				
<i>Unconstrained</i>		-13.44	-10.91	-8.18	-5.85	-2.74	1.99	15.42	0.059	0.743
<i>Constrained</i>	$\bar{m} = 1$	-15.70	-13.18	-10.50	-7.76	-4.78	3.51	32.78	0.061	0.717
	$\bar{m} = 2$	-14.27	-11.86	-8.90	-6.44	-3.34	2.64	20.52	0.061	0.725
	$\bar{m} = 3$	-13.69	-11.38	-8.38	-5.95	-2.97	2.23	17.84	0.058	0.745
	$\bar{m} = 4$	-13.58	-11.09	-8.23	-5.89	-2.75	2.10	17.84	0.058	0.754
	$\bar{m} = 5$	-13.37	-10.97	-8.14	-5.79	-2.88	2.05	15.82	0.060	0.747
	$\bar{m} = 6$	-13.54	-11.03	-8.31	-5.86	-2.86	2.00	15.51	0.060	0.738
	$\bar{m} = 7$	-13.49	-10.94	-8.17	-5.86	-2.78	1.98	15.32	0.058	0.743
<i>Weakly-targeted</i>	$\beta = 0.01$	-11.67	-9.02	-6.37	-3.87	-1.00	2.59	20.13	0.116	0.604
	$\beta = 0.1$	-12.08	-9.60	-6.87	-4.30	-1.46	2.15	15.75	0.082	0.680
	$\beta = 1$	-12.52	-10.12	-7.34	-4.81	-1.93	1.96	13.01	0.057	0.758
	$\beta = 10$	-13.00	-10.56	-7.83	-5.27	-2.34	2.19	16.70	0.031	0.859
	$\beta = 100$	-13.25	-10.72	-8.01	-5.39	-2.58	2.45	20.81	0.024	0.876
<i>Unit-level</i>	$\xi = 0.01$	-11.76	-9.15	-6.51	-3.91	-1.15	2.57	20.54	0.118	0.593
	$\xi = 0.1$	-13.11	-10.59	-7.82	-5.15	-2.29	2.06	15.81	0.060	0.754
	$\xi = 1$	-13.74	-11.12	-8.42	-5.75	-2.84	2.37	19.70	0.029	0.850
	$\xi = 10$	-13.74	-11.20	-8.55	-5.89	-3.09	3.02	26.38	0.028	0.866
	$\xi = 100$	-13.79	-11.16	-8.54	-5.90	-3.08	3.20	28.37	0.029	0.863
<i>Penalized</i>	$\lambda = 0.01$	-13.40	-10.93	-8.32	-5.82	-2.83	1.97	16.13	0.056	0.757
	$\lambda = 0.1$	-13.33	-10.80	-8.12	-5.56	-2.65	2.07	13.81	0.045	0.779
	$\lambda = 1$	-13.32	-10.84	-8.15	-5.39	-2.60	3.08	27.44	0.056	0.738
	$\lambda = 10$	-13.39	-10.82	-7.95	-5.34	-2.58	3.85	37.08	0.103	0.595
	$\lambda = 100$	-13.35	-10.82	-8.00	-5.29	-2.57	4.10	41.23	0.117	0.562

Note: Unless otherwise noted, all designs use $m = 1$ and $\bar{m} = 14$.

5.2.3. Performance with Nonlinearities

We now study the behavior of the estimators based on synthetic control designs under deviations from the linear model in (10a) and (10b). We consider a nonlinear data generating process,

$$Y_{jt}^N = \delta_t + \exp(\boldsymbol{\theta}'_t \mathbf{Z}_j) + \exp(\boldsymbol{\lambda}'_t \boldsymbol{\mu}_j) + \epsilon_{jt},$$

$$Y_{jt}^I = v_t + \exp(\boldsymbol{\gamma}'_t \mathbf{Z}_j) + \exp(\boldsymbol{\eta}'_t \boldsymbol{\mu}_j) + \xi_{jt}.$$

The motivation to study a nonlinear model is that nonlinearities may induce interpolation biases, affecting the relative performance of the different designs. All parameter values are the same as in the simulation setup of section 5.1, except for the values of $\boldsymbol{\theta}_t$, $\boldsymbol{\gamma}_t$, $\boldsymbol{\lambda}_t$, and $\boldsymbol{\eta}_t$, which are chosen to be random vectors of i.i.d. Uniform (0, 3) random variables, instead of Uniform (0, 10), to control the magnitude of the exponential components in the nonlinear design.

Tables 3 and 4 report the results for τ_t and τ_t^T , respectively. In comparison to the results in Tables 1 and 2, we now see that the *Unit-level* and *Penalized* designs can easily match and in

Table 4: Average Treatment Effects on the Treated, Nonlinear Model (Averages over 1000 Simulations)

	τ_t^T							$\hat{\tau}_t$							MAE^T	MSE^T
	$t = 26$	$t = 27$	$t = 28$	$t = 29$	$t = 30$	$t = 26$	$t = 27$	$t = 28$	$t = 29$	$t = 30$	$t = 26$	$t = 27$	$t = 28$	$t = 29$		
<i>Unconstrained</i>																
<i>Constrained</i>																
$\bar{m} = 1$	-13.29	-10.65	-7.96	-5.47	-2.40	-13.44	-10.91	-8.18	-5.85	-2.74	2.31	20.07				
$\bar{m} = 2$	-13.41	-10.81	-8.01	-5.25	-2.61	-15.70	-13.18	-10.50	-7.76	-4.78	3.53	29.28				
$\bar{m} = 3$	-13.19	-10.88	-7.93	-5.42	-2.47	-14.27	-11.86	-8.90	-6.44	-3.34	2.86	26.06				
$\bar{m} = 4$	-13.17	-10.82	-8.07	-5.54	-2.44	-13.69	-11.38	-8.38	-5.95	-2.97	2.52	20.96				
$\bar{m} = 5$	-13.29	-10.82	-8.03	-5.59	-2.31	-13.58	-11.09	-8.23	-5.89	-2.75	2.44	20.98				
$\bar{m} = 6$	-13.23	-10.77	-7.88	-5.58	-2.30	-13.37	-10.97	-8.14	-5.79	-2.88	2.39	20.72				
$\bar{m} = 7$	-13.36	-10.71	-7.99	-5.50	-2.36	-13.54	-11.03	-8.31	-5.86	-2.86	2.32	19.91				
	-13.31	-10.68	-7.97	-5.46	-2.39	-13.49	-10.94	-8.17	-5.86	-2.78	2.30	19.91				
<i>Weakly-targeted</i>																
$\beta = 0.01$	-13.14	-10.67	-7.97	-5.47	-2.45	-11.67	-9.02	-6.37	-3.87	-1.00	2.68	21.06				
$\beta = 0.1$	-13.19	-10.66	-7.92	-5.47	-2.45	-12.08	-9.60	-6.87	-4.30	-1.46	2.31	17.03				
$\beta = 1$	-13.22	-10.76	-7.93	-5.46	-2.45	-12.52	-10.12	-7.34	-4.81	-1.93	1.91	10.68				
$\beta = 10$	-13.28	-10.70	-7.95	-5.49	-2.49	-13.00	-10.56	-7.83	-5.27	-2.34	1.33	3.86				
$\beta = 100$	-13.24	-10.74	-8.03	-5.43	-2.54	-13.25	-10.72	-8.01	-5.39	-2.58	1.14	2.56				
<i>Unit-level</i>																
$\xi = 0.01$	-13.20	-10.76	-7.98	-5.44	-2.48	-11.76	-9.15	-6.51	-3.91	-1.15	2.68	21.11				
$\xi = 0.1$	-13.35	-10.81	-8.07	-5.53	-2.55	-13.11	-10.59	-7.82	-5.15	-2.29	1.99	10.94				
$\xi = 1$	-13.40	-10.77	-8.05	-5.47	-2.56	-13.74	-11.12	-8.42	-5.75	-2.84	1.47	4.45				
$\xi = 10$	-13.35	-10.66	-7.99	-5.38	-2.71	-13.74	-11.20	-8.55	-5.89	-3.09	1.49	4.30				
$\xi = 100$	-13.39	-10.63	-7.95	-5.37	-2.69	-13.79	-11.16	-8.54	-5.90	-3.08	1.56	4.55				
<i>Penalized</i>																
$\lambda = 0.01$	-13.32	-10.69	-8.07	-5.49	-2.42	-13.40	-10.93	-8.32	-5.82	-2.83	2.25	18.48				
$\lambda = 0.1$	-13.17	-10.72	-8.09	-5.59	-2.55	-13.33	-10.80	-8.12	-5.56	-2.65	1.87	10.06				
$\lambda = 1$	-13.40	-10.80	-7.98	-5.35	-2.66	-13.32	-10.84	-8.15	-5.39	-2.60	1.97	8.89				
$\lambda = 10$	-13.41	-10.89	-8.00	-5.32	-2.64	-13.39	-10.82	-7.95	-5.34	-2.58	2.79	17.48				
$\lambda = 100$	-13.42	-10.87	-8.02	-5.26	-2.63	-13.35	-10.82	-8.00	-5.29	-2.57	3.06	21.54				

Table 5: Average Treatment Effects Under the Null Hypothesis (15) (Averages over 1000 Simulations)

		τ_t								
		$t = 26$	$t = 27$	$t = 28$	$t = 29$	$t = 30$				
		-0.01	0.00	0.00	0.02	-0.01				
		$\hat{\tau}_t$					<i>MAE</i>	<i>MSE</i>	\hat{p}	$\hat{p} < 0.05$
		$t = 26$	$t = 27$	$t = 28$	$t = 29$	$t = 30$				
<i>Unconstrained</i>		0.02	-0.05	0.01	0.03	0.05	1.09	1.90	0.498	0.049
<i>Constrained</i>	$\bar{m} = 1$	-0.04	-0.10	0.02	-0.19	0.09	3.00	14.36	0.504	0.035
	$\bar{m} = 2$	0.05	-0.09	-0.01	0.14	0.06	1.82	5.24	0.496	0.042
	$\bar{m} = 3$	-0.03	0.05	0.04	0.07	0.05	1.43	3.21	0.498	0.057
	$\bar{m} = 4$	-0.02	-0.03	-0.01	0.01	0.00	1.22	2.37	0.499	0.046
	$\bar{m} = 5$	0.06	-0.06	-0.02	0.04	0.02	1.15	2.09	0.489	0.050
	$\bar{m} = 6$	0.04	-0.01	-0.01	0.10	0.04	1.10	1.91	0.496	0.046
	$\bar{m} = 7$	0.01	-0.02	0.02	0.05	0.05	1.08	1.86	0.497	0.051
<i>Weakly-targeted</i>	$\beta = 0.01$	-0.13	-0.06	-0.23	-0.06	-0.18	1.46	3.34	0.492	0.045
	$\beta = 0.1$	-0.03	0.03	-0.03	0.02	-0.03	1.14	2.05	0.491	0.052
	$\beta = 1$	-0.08	0.01	-0.06	-0.03	-0.01	1.05	1.74	0.492	0.059
	$\beta = 10$	-0.10	-0.04	-0.09	-0.09	-0.08	1.00	1.56	0.484	0.047
	$\beta = 100$	-0.11	-0.01	-0.07	-0.05	-0.06	1.01	1.60	0.461	0.062
<i>Unit-level</i>	$\xi = 0.01$	0.02	0.06	-0.04	0.07	-0.13	1.31	2.76	0.491	0.050
	$\xi = 0.1$	-0.04	0.01	-0.08	0.06	-0.04	1.15	2.14	0.509	0.048
	$\xi = 1$	0.08	0.03	0.04	0.08	-0.01	1.41	3.11	0.494	0.053
	$\xi = 10$	-0.04	-0.17	-0.12	0.03	-0.10	1.95	6.13	0.495	0.048
	$\xi = 100$	-0.06	-0.15	-0.04	0.06	-0.07	2.35	8.86	0.494	0.051
<i>Penalized</i>	$\lambda = 0.01$	0.03	0.00	-0.03	0.03	-0.01	1.13	1.99	0.485	0.058
	$\lambda = 0.1$	0.02	-0.02	-0.03	-0.02	-0.06	1.35	2.92	0.492	0.052
	$\lambda = 1$	0.11	-0.08	-0.01	0.10	0.13	2.21	7.83	0.505	0.038
	$\lambda = 10$	-0.02	0.12	0.02	0.04	0.27	3.77	22.94	0.511	0.057
	$\lambda = 100$	-0.08	0.17	-0.01	-0.08	0.32	4.19	28.38	0.509	0.057

Note: Unless otherwise noted, all designs use $\underline{m} = 1$ and $\overline{m} = 14$.

some cases improve the performance of the *Unconstrained* design, especially for the estimation of τ_t^T . By fitting each treated unit with a unit-specific synthetic control, the *Unit-level* design can ameliorate interpolation biases induced by the aggregation of \mathbf{X}_j . The *Penalized* design selects synthetic treated and control units close to $\overline{\mathbf{X}}$ in the space of the predictors, which can reduce interpolation biases at the potential cost of lower precision for large values of λ (in which case, the *Penalized* design employs a small number of units in the synthetic treated and synthetic control). As in Table 4, when the parameter of interest is τ_t^T , the *Weakly targeted* design easily outperforms the unconstrained estimator for large values of β .

5.2.4. Test size

In this section, we generate the model primitives under the null hypothesis (15). That is, we employ a data generating process such that the values of the common factors and the distributions

Table 6: Treatment Effects Estimates Using Randomized Assignments and Difference-in-Means estimators (Averages over 1000 Simulations)

	τ_t					MAE	MSE
	$t = 26$	$t = 27$	$t = 28$	$t = 29$	$t = 30$		
	-13.05	-10.92	-7.65	-5.17	-2.22		
	$\hat{\tau}_t$					MAE	MSE
	$t = 26$	$t = 27$	$t = 28$	$t = 29$	$t = 30$		
$\bar{m} = 1$	-12.90	-10.65	-7.17	-4.90	-2.08	5.73	52.34
$\bar{m} = 2$	-12.87	-10.67	-7.33	-5.05	-1.98	4.54	32.58
$\bar{m} = 3$	-12.68	-10.51	-7.22	-4.87	-1.85	3.83	23.41
$\bar{m} = 4$	-12.99	-10.94	-7.78	-5.13	-2.16	3.38	17.72
$\bar{m} = 5$	-12.93	-10.75	-7.45	-5.14	-2.17	3.10	15.00
$\bar{m} = 6$	-12.91	-10.85	-7.53	-5.09	-2.12	3.08	15.78
$\bar{m} = 7$	-12.91	-10.66	-7.53	-5.02	-2.07	2.87	12.90

Note: In this table, \bar{m} stands for the maximum number of treated units allowed by randomization.

of the idiosyncratic error variables are unaffected by the intervention.

We report the simulation results in Table 5, which organizes information in the same way as in Table 1. Because the data are generated from the same distribution under treatment and under no treatment, the average treatment effects in Table 5 are close to zero. The same is true for the averages of $\hat{\tau}_t$ for all designs. Under the null hypothesis (15), the p -value should approximately follow a uniform distribution between zero and one. The results in Table 5 show good behavior of our testing procedure under the null hypothesis: average p -values and rejection rates are close to 0.5 and 0.05, respectively.

5.2.5. Comparison to Randomized Control Trials

Randomized treatment assignment produces ex-ante (pre-randomization) unbiased estimation of the average treatment effect. As we show below, however, ex-post (post-randomization) biases can be large, especially when only a small number of units are treated.

Table 6 reports performance measures of a randomized design for the simulation set-up of Section 5.2. For seven treated units, randomization of the treatment produces MAE and MSE equal to 2.87 and 12.90, respectively, which are substantially larger than the 0.94 and 1.41 values for the *Constrained* design in Table 1. For one treated unit only, the difference is even more stark; randomization of the treatment produces MAE and MSE equal to 5.73 and 52.34, substantially

larger than the 2.93 and 13.29 for the *Constrained* design in Table 1. These results underscore the appropriateness of synthetic control designs in experimental studies with a limited number of aggregate units.

6. Conclusions

Experimental design methods have largely been concerned with settings where a large number of experimental units are randomly assigned to a treatment arm, and a similarly large number of experimental units are assigned to a control arm. This focus on large samples and randomization has proven to be enormously useful in various classes of problems, but becomes inadequate when treating more than a few units is unfeasible, as is often the case in experimental studies with large aggregate units (e.g., markets). In that case, randomized designs may produce estimators that are substantially biased (post-randomization) relative to the average treatment effect or to the average treatment effect on the treated. Large biases can be expected when the unit or units assigned to treatment fail to approximate average outcomes under treatment for the entire population, or when the units in the control arm fail to approximate the outcomes that treated units would experience without treatment.

In this article we have applied synthetic control techniques, widely used in observational studies, to the design of experiments when treatment can only be applied to a small number of experimental units. The synthetic control design optimizes jointly over the identities of the units assigned to the treatment and the control arms, and over the weights that determine the relative contribution of those units to reproduce the counterfactuals of interest. We propose various designs aimed to estimate average treatment effects, analyze the properties of such designs and the resulting estimators, and devise inferential methods to test a null hypothesis of no treatment effects. In addition, we report result from an application to retail sales data and simulation results that demonstrate the applicability and computational feasibility of the methods proposed in this article. We show that synthetic control design can substantially outperform randomized designs in experimental settings with a small number of treated units.

Corporate researchers and academic investigators are often confronted with settings where interventions at the level of micro-units (i.e., customers, workers, or families) are unfeasible,

impractical or ineffective (see, e.g., Duflo, Glennerster and Kremer, 2007, Jones and Barrows, 2019). There is, in consequence, a wide range of potential applications of experimental design methods for large aggregate entities, like the ones proposed in this article.

References

- Abadie, Alberto.** 2021. “Using Synthetic Controls: Feasibility, Data Requirements, and Methodological Aspects.” *Journal of Economic Literature*, 59(2): 391–425.
- Abadie, Alberto, Alexis Diamond, and Jens 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, Alberto, and Javier Gardeazabal.** 2003. “The Economic Costs of Conflict: A Case Study of the Basque County.” *American Economic Review*, 93(1): 113–132.
- Abadie, Alberto, and Jérémy L’Hour.** 2021. “A Penalized Synthetic Control Estimator for Disaggregated Data.” *Journal of the American Statistical Association*, 116(536): 1817–1834.
- Agarwal, Anish, Devavrat Shah, and Dennis Shen.** 2021. “Synthetic Interventions.” *arXiv e-print*, arXiv:2006.07691.
- Amjad, Muhammad, Devavrat Shah, and Dennis Shen.** 2018. “Robust Synthetic Control.” *Journal of Machine Learning Research*, 19(22): 1–51.
- Arkhangelsky, Dmitry, Susan Athey, David A. Hirshberg, Guido W. Imbens, and Stefan Wager.** 2021. “Synthetic Difference-in-Differences.” *American Economic Review*, 111(12): 4088–4118.
- Ben-Michael, Eli, Avi Feller, and Jesse Rothstein.** 2021. “The Augmented Synthetic Control Method.” *Journal of the American Statistical Association*, 116(536): 1789–1803.
- Bobkov, SG, and GP Chistyakov.** 2014. “Bounds on the Maximum of the Density for Sums of Independent Random Variables.” *Journal of Mathematical Sciences*, 199(2): 100–106.

- Bottmer, Lea, Guido Imbens, Jann Spiess, and Merrill Warnick.** 2021. “A Design-Based Perspective on Synthetic Control Methods.” *arXiv e-prints*, arXiv:2101.09398.
- Cattaneo, Matias D., Yingjie Feng, and Rocio Titiunik.** 2021. “Prediction Intervals for Synthetic Control Methods.” *Journal of the American Statistical Association*, 116(536): 1865–1880.
- Chen, Jiafeng.** 2023. “Synthetic Control as Online Linear Regression.” *Econometrica*, 91(2): 465–491.
- Chernozhukov, Victor, Kaspar Wüthrich, and Yinchu Zhu.** 2019. “Distributional Conformal Prediction.” *arXiv preprint arXiv:1909.07889*.
- Chernozhukov, Victor, Kaspar Wüthrich, and Yinchu Zhu.** 2021. “An Exact and Robust Conformal Inference Method for Counterfactual and Synthetic Controls.” *Journal of the American Statistical Association*, 116(536): 1849–1864.
- Doudchenko, Nick, Khashayar Khosravi, Jean Pouget-Abadie, Sebastien Lahaie, Miles Lubin, Vahab Mirrokni, Jann Spiess, et al.** 2021. “Synthetic Design: An Optimization Approach to Experimental Design with Synthetic Controls.” *Advances in Neural Information Processing Systems*, 34.
- Doudchenko, Nikolay, and Guido W Imbens.** 2016. “Balancing, Regression, Difference-in-Differences and Synthetic Control Methods: A Synthesis.” National Bureau of Economic Research.
- Doudchenko, Nikolay, David Gilinson, Sean Taylor, and Nils Wernerfelt.** n.d. “Designing Experiments with Synthetic Controls.” https://mackinstitute.wharton.upenn.edu/wp-content/uploads/2020/03/Wernerfelt-Nils-Doudchenko-Nick-Gilinson-David-and-Taylor-Sean_Designing-Experiments-with-Synthetic-Controls.pdf.

- Duflo, Esther, Rachel Glennerster, and Michael Kremer.** 2007. “Using Randomization in Development Economics Research: A Toolkit.” *Handbook of development economics*, 4: 3895–3962.
- Ferman, Bruno.** 2021. “On the Properties of the Synthetic Control Estimator with Many Periods and Many Controls.” *Journal of the American Statistical Association*, 116(536): 1764–1772.
- Firpo, Sergio, and Vitor Possebom.** 2018. “Synthetic Control Method: Inference, Sensitivity Analysis and Confidence Sets.” *Journal of Causal Inference*, 6(2): 20160026.
- Gunsilius, F. F.** 2023. “Distributional Synthetic Controls.” *Econometrica*, 91(3): 1105–1117.
- Jones, Nick, and Sam Barrows.** 2019. “Synthetic Control and Alternatives to A/B Testing at Uber.” *Presented at PyData Amsterdam 2019*, <https://youtu.be/j5DoJV5S2Ao>.
- Kasy, Max, and Lukas Lehner.** 2023. “Employing the Unemployed of Marienthal: Evaluation of a Guaranteed Job Program.” https://maxkasy.github.io/home/files/papers/Jobguarantee_marienthal.pdf.
- Lei, Jing, Max G’Sell, Alessandro Rinaldo, Ryan J Tibshirani, and Larry Wasserman.** 2018. “Distribution-Free Predictive Inference for Regression.” *Journal of the American Statistical Association*, 113(523): 1094–1111.
- Lei, Lihua, and Emmanuel J Candès.** 2021. “Conformal Inference of Counterfactuals and Individual Treatment Effects.” *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 83(5): 911–938.
- Prakash, Sourav.** 2023. “Walmart Condensed Sales Data.” <https://www.kaggle.com/ds/3471234>.
- Rigollet, Philippe, and Jan-Christian Hütter.** 2019. “High Dimensional Statistics.” <http://www-math.mit.edu/~rigollet/PDFs/RigNotes17.pdf>.

Vovk, Vladimir, Alex Gammerman, and Glenn Shafer. 2005. *Algorithmic Learning in a Random World*. Springer Science & Business Media.

Appendix

A. Proofs

A.1. Proof of Theorem 1

Proof of Theorem 1. For any period $t = T_0 + 1, \dots, T$ we decompose $(\widehat{\tau}_t - \tau_t)$ as follows,

$$\begin{aligned}\widehat{\tau}_t - \tau_t &= \left(\sum_{j=1}^J w_j^* Y_{jt}^I - \sum_{j=1}^J v_j^* Y_{jt}^N \right) - \left(\sum_{j=1}^J f_j Y_{jt}^I - \sum_{j=1}^J f_j Y_{jt}^N \right) \\ &= \left(\sum_{j=1}^J w_j^* Y_{jt}^I - \sum_{j=1}^J f_j Y_{jt}^I \right) - \left(\sum_{j=1}^J v_j^* Y_{jt}^N - \sum_{j=1}^J f_j Y_{jt}^N \right).\end{aligned}\quad (\text{A.1})$$

The first term in (A.1) measures the difference between the synthetic treatment outcome and the aggregated treatment outcomes. The second term measures the difference between the synthetic control outcome and the aggregate control outcomes. We bound these two terms separately. From (10b), we obtain

$$\begin{aligned}\sum_{j=1}^J w_j^* Y_{jt}^I - \sum_{j=1}^J f_j Y_{jt}^I &= \boldsymbol{\gamma}'_t \left(\sum_{j=1}^J w_j^* \mathbf{Z}_j - \sum_{j=1}^J f_j \mathbf{Z}_j \right) \\ &\quad + \boldsymbol{\eta}'_t \left(\sum_{j=1}^J w_j^* \boldsymbol{\mu}_j - \sum_{j=1}^J f_j \boldsymbol{\mu}_j \right) + \left(\sum_{j=1}^J w_j^* \xi_{jt} - \sum_{j=1}^J f_j \xi_{jt} \right)\end{aligned}\quad (\text{A.2})$$

Similarly, using expression (10a), we obtain

$$\begin{aligned}\sum_{j=1}^J w_j^* \mathbf{Y}_j^\mathcal{E} - \sum_{j=1}^J f_j \mathbf{Y}_j^\mathcal{E} &= \boldsymbol{\theta}_\mathcal{E} \left(\sum_{j=1}^J w_j^* \mathbf{Z}_j - \sum_{j=1}^J f_j \mathbf{Z}_j \right) \\ &\quad + \boldsymbol{\lambda}_\mathcal{E} \left(\sum_{j=1}^J w_j^* \boldsymbol{\mu}_j - \sum_{j=1}^J f_j \boldsymbol{\mu}_j \right) + \left(\sum_{j=1}^J w_j^* \boldsymbol{\epsilon}_j^\mathcal{E} - \sum_{j=1}^J f_j \boldsymbol{\epsilon}_j^\mathcal{E} \right),\end{aligned}$$

where $\boldsymbol{\theta}_\mathcal{E}$ is the $(T_\mathcal{E} \times R)$ matrix with rows equal to the $\boldsymbol{\theta}_t$'s indexed by \mathcal{E} , and $\boldsymbol{\epsilon}_j^\mathcal{E}$ is defined analogously. Pre-multiplying by $\boldsymbol{\eta}'_t(\boldsymbol{\lambda}'_\mathcal{E} \boldsymbol{\lambda}_\mathcal{E})^{-1} \boldsymbol{\lambda}'_\mathcal{E}$ yields

$$\begin{aligned}\boldsymbol{\eta}'_t(\boldsymbol{\lambda}'_\mathcal{E} \boldsymbol{\lambda}_\mathcal{E})^{-1} \boldsymbol{\lambda}'_\mathcal{E} &\left(\sum_{j=1}^J w_j^* \mathbf{Y}_j^\mathcal{E} - \sum_{j=1}^J f_j \mathbf{Y}_j^\mathcal{E} \right) \\ &= \boldsymbol{\eta}'_t(\boldsymbol{\lambda}'_\mathcal{E} \boldsymbol{\lambda}_\mathcal{E})^{-1} \boldsymbol{\lambda}'_\mathcal{E} \boldsymbol{\theta}_\mathcal{E} \left(\sum_{j=1}^J w_j^* \mathbf{Z}_j - \sum_{j=1}^J f_j \mathbf{Z}_j \right)\end{aligned}$$

$$\begin{aligned}
& + \boldsymbol{\eta}'_t \left(\sum_{j=1}^J w_j^* \boldsymbol{\mu}_j - \sum_{j=1}^J f_j \boldsymbol{\mu}_j \right) \\
& + \boldsymbol{\eta}'_t (\boldsymbol{\lambda}'_\varepsilon \boldsymbol{\lambda}_\varepsilon)^{-1} \boldsymbol{\lambda}'_\varepsilon \left(\sum_{j=1}^J w_j^* \boldsymbol{\epsilon}_j^\varepsilon - \sum_{j=1}^J f_j \boldsymbol{\epsilon}_j^\varepsilon \right). \tag{A.3}
\end{aligned}$$

Equations (A.2) and (A.3) imply

$$\begin{aligned}
\sum_{j=1}^J w_j^* Y_{jt}^I - \sum_{j=1}^J f_j Y_{jt}^I & = (\boldsymbol{\gamma}'_t - \boldsymbol{\eta}'_t (\boldsymbol{\lambda}'_\varepsilon \boldsymbol{\lambda}_\varepsilon)^{-1} \boldsymbol{\lambda}'_\varepsilon \boldsymbol{\theta}_\varepsilon) \left(\sum_{j=1}^J w_j^* \mathbf{Z}_j - \sum_{j=1}^J f_j \mathbf{Z}_j \right) \\
& + \boldsymbol{\eta}'_t (\boldsymbol{\lambda}'_\varepsilon \boldsymbol{\lambda}_\varepsilon)^{-1} \boldsymbol{\lambda}'_\varepsilon \left(\sum_{j=1}^J w_j^* \mathbf{Y}_j^\varepsilon - \sum_{j=1}^J f_j \mathbf{Y}_j^\varepsilon \right) \\
& - \boldsymbol{\eta}'_t (\boldsymbol{\lambda}'_\varepsilon \boldsymbol{\lambda}_\varepsilon)^{-1} \boldsymbol{\lambda}'_\varepsilon \sum_{j=1}^J w_j^* \boldsymbol{\epsilon}_j^\varepsilon \\
& + \boldsymbol{\eta}'_t (\boldsymbol{\lambda}'_\varepsilon \boldsymbol{\lambda}_\varepsilon)^{-1} \boldsymbol{\lambda}'_\varepsilon \sum_{j=1}^J f_j \boldsymbol{\epsilon}_j^\varepsilon \\
& + \left(\sum_{j=1}^J w_j^* \xi_{jt} - \sum_{j=1}^J f_j \xi_{jt} \right). \tag{A.4}
\end{aligned}$$

If Assumption 3 holds, (A.4) becomes

$$\begin{aligned}
\sum_{j=1}^J w_j^* Y_{jt}^I - \sum_{j=1}^J f_j Y_{jt}^I & = -\boldsymbol{\eta}'_t (\boldsymbol{\lambda}'_\varepsilon \boldsymbol{\lambda}_\varepsilon)^{-1} \boldsymbol{\lambda}'_\varepsilon \sum_{j=1}^J w_j^* \boldsymbol{\epsilon}_j^\varepsilon \\
& + \boldsymbol{\eta}'_t (\boldsymbol{\lambda}'_\varepsilon \boldsymbol{\lambda}_\varepsilon)^{-1} \boldsymbol{\lambda}'_\varepsilon \sum_{j=1}^J f_j \boldsymbol{\epsilon}_j^\varepsilon \\
& + \left(\sum_{j=1}^J w_j^* \xi_{jt} - \sum_{j=1}^J f_j \xi_{jt} \right). \tag{A.5}
\end{aligned}$$

Only the first term on the right-hand side of (A.5) has a non-zero mean (because the weights, w_j^* , depend on the error terms $\boldsymbol{\epsilon}_j^\varepsilon$). Therefore,

$$\left| E \left[\sum_{j=1}^J w_j^* Y_{jt}^I - \sum_{j=1}^J f_j Y_{jt}^I \right] \right| = \left| E \left[\boldsymbol{\eta}'_t (\boldsymbol{\lambda}'_\varepsilon \boldsymbol{\lambda}_\varepsilon)^{-1} \boldsymbol{\lambda}'_\varepsilon \sum_{j=1}^J w_j^* \boldsymbol{\epsilon}_j^\varepsilon \right] \right|.$$

Using the same line of reasoning for the second term on the right-hand side of (A.1), we obtain

$$\left| E \left[\sum_{j=1}^J v_j^* Y_{jt}^N - \sum_{j=1}^J f_j Y_{jt}^N \right] \right| = \left| E \left[\boldsymbol{\lambda}'_t (\boldsymbol{\lambda}'_\varepsilon \boldsymbol{\lambda}_\varepsilon)^{-1} \boldsymbol{\lambda}'_\varepsilon \sum_{j=1}^J v_j^* \boldsymbol{\epsilon}_j^\varepsilon \right] \right|.$$

For any $t \geq T_0 + 1$ and $s \in \mathcal{E}$, under Assumption 2 (i), we apply Cauchy-Schwarz inequality and the eigenvalue bound on the Rayleigh quotient to obtain

$$\begin{aligned} (\boldsymbol{\eta}'_t(\boldsymbol{\lambda}'_{\mathcal{E}}\boldsymbol{\lambda}_{\mathcal{E}})^{-1}\boldsymbol{\lambda}_s)^2 &\leq (\boldsymbol{\eta}'_t(\boldsymbol{\lambda}'_{\mathcal{E}}\boldsymbol{\lambda}_{\mathcal{E}})^{-1}\boldsymbol{\eta}_t) (\boldsymbol{\lambda}'_s(\boldsymbol{\lambda}'_{\mathcal{E}}\boldsymbol{\lambda}_{\mathcal{E}})^{-1}\boldsymbol{\lambda}_s) \\ &\leq \left(\frac{\bar{\eta}^2 F}{T_{\mathcal{E}}\underline{\zeta}}\right) \left(\frac{\bar{\lambda}^2 F}{T_{\mathcal{E}}\underline{\zeta}}\right). \end{aligned}$$

Similarly,

$$(\boldsymbol{\lambda}'_t(\boldsymbol{\lambda}'_{\mathcal{E}}\boldsymbol{\lambda}_{\mathcal{E}})^{-1}\boldsymbol{\lambda}_s)^2 \leq \left(\frac{\bar{\lambda}^2 F}{T_{\mathcal{E}}\underline{\zeta}}\right)^2. \quad (\text{A.6})$$

Let

$$\bar{\epsilon}_{jt}^{\mathcal{E}} = \boldsymbol{\eta}'_t(\boldsymbol{\lambda}'_{\mathcal{E}}\boldsymbol{\lambda}_{\mathcal{E}})^{-1}\boldsymbol{\lambda}'_{\mathcal{E}}\boldsymbol{\epsilon}_j^{\mathcal{E}} = \sum_{s \in \mathcal{E}} \boldsymbol{\eta}'_t(\boldsymbol{\lambda}'_{\mathcal{E}}\boldsymbol{\lambda}_{\mathcal{E}})^{-1}\boldsymbol{\lambda}_s \epsilon_{js}.$$

Because $\bar{\epsilon}_{jt}^{\mathcal{E}}$ is a linear combination of independent sub-Gaussians with variance proxy $\bar{\sigma}^2$, it follows that $\bar{\epsilon}_{jt}^{\mathcal{E}}$ is sub-Gaussian with variance proxy $(\bar{\eta}\bar{\lambda}F/\underline{\zeta})^2\bar{\sigma}^2/T_{\mathcal{E}}$. Let $\mathcal{S} = \{\mathbf{w} \in \mathbb{R}^J : \sum_{j=1}^J w_j = 1\}$. Theorem 1.16 from Rigollet and Hütter (2019) implies

$$\begin{aligned} &\left| E \left[\sum_{j=1}^J w_j^* Y_{jt}^I - \sum_{j=1}^J f_j Y_{jt}^I \right] \right| \\ &= \left| E \left[\sum_{j=1}^J w_j^* \bar{\epsilon}_{jt}^{\mathcal{E}} \right] \right| \leq E \left[\max_{\mathbf{w} \in \mathcal{S}} \left| \sum_{j=1}^J w_j \bar{\epsilon}_{jt}^{\mathcal{E}} \right| \right] \leq \frac{\bar{\eta}\bar{\lambda}F}{\underline{\zeta}} \sqrt{2 \log(2J)} \frac{\bar{\sigma}}{\sqrt{T_{\mathcal{E}}}}. \end{aligned}$$

An analogous argument yields

$$\left| E \left[\sum_{j=1}^J v_j^* Y_{jt}^N - \sum_{j=1}^J f_j Y_{jt}^N \right] \right| \leq \frac{\bar{\lambda}^2 F}{\underline{\zeta}} \sqrt{2 \log(2J)} \frac{\bar{\sigma}}{\sqrt{T_{\mathcal{E}}}},$$

which completes the proof of the theorem.

Suppose now Assumption 4 holds (but Assumption 3 does not). To obtain a bound on the bias we need to bound the first two terms in (A.4). Recall that

$$\boldsymbol{\eta}'_t(\boldsymbol{\lambda}'_{\mathcal{E}}\boldsymbol{\lambda}_{\mathcal{E}})^{-1}\boldsymbol{\lambda}_s \leq \frac{\bar{\lambda}\bar{\eta}F}{T_{\mathcal{E}}\underline{\zeta}}.$$

Therefore, the absolute value of each element in vector $(\gamma'_t - \eta'_t(\lambda'_\varepsilon \lambda_\varepsilon)^{-1} \lambda'_\varepsilon \theta_\varepsilon)$ is bounded by $\bar{\gamma} + \bar{\theta} \frac{\bar{\lambda} \bar{\eta} F}{\underline{\zeta}}$. Cauchy–Schwarz inequality and Assumption 4 imply

$$\begin{aligned} & \left| (\gamma'_t - \eta'_t(\lambda'_\varepsilon \lambda_\varepsilon)^{-1} \lambda'_\varepsilon \theta_\varepsilon) \left(\sum_{j=1}^J w_j^* \mathbf{Z}_j - \sum_{j=1}^J f_j \mathbf{Z}_j \right) \right| \\ & \leq \left(\bar{\gamma} + \bar{\theta} \frac{\bar{\lambda} \bar{\eta} F}{\underline{\zeta}} \right) \sqrt{R} \left\| \sum_{j=1}^J w_j^* \mathbf{Z}_j - \sum_{j=1}^J f_j \mathbf{Z}_j \right\|_2 \\ & \leq \left(\bar{\gamma} + \bar{\theta} \frac{\bar{\lambda} \bar{\eta} F}{\underline{\zeta}} \right) R d, \end{aligned}$$

and

$$\left| \eta'_t(\lambda'_\varepsilon \lambda_\varepsilon)^{-1} \lambda'_\varepsilon \left(\sum_{j=1}^J w_j^* \mathbf{Y}_j^\varepsilon - \sum_{j=1}^J f_j \mathbf{Y}_j^\varepsilon \right) \right| \leq \frac{\bar{\lambda} \bar{\eta} F}{\underline{\zeta}} d.$$

Combining the last two displayed equations with (A.4), we have

$$\left| E \left[\sum_{j=1}^J w_j^* Y_{jt}^I - \sum_{j=1}^J f_j Y_{jt}^I \right] \right| \leq \left(\bar{\gamma} R + \frac{\bar{\lambda} \bar{\eta} F}{\underline{\zeta}} (1 + \bar{\theta} R) \right) d + \frac{\bar{\lambda} \bar{\eta} F}{\underline{\zeta}} \sqrt{2 \log(2J)} \frac{\bar{\sigma}}{\sqrt{T_\varepsilon}}.$$

An analogous derivation produces

$$\left| E \left[\sum_{j=1}^J v_j^* Y_{jt}^N - \sum_{j=1}^J f_j Y_{jt}^N \right] \right| \leq \left(\bar{\theta} R + \frac{\bar{\lambda}^2 F}{\underline{\zeta}} (1 + \bar{\theta} R) \right) d + \frac{\bar{\lambda}^2 F}{\underline{\zeta}} \sqrt{2 \log(2J)} \frac{\bar{\sigma}}{\sqrt{T_\varepsilon}},$$

which finishes the proof of the theorem. ■

A.2. Proof of Theorem 2

Proof of Theorem 2. Recall that

$$\hat{u}_t = \sum_{j=1}^J w_j^* Y_{jt} - \sum_{j=1}^J v_j^* Y_{jt},$$

for $t \in \mathcal{B} \cup \{T_0 + 1, \dots, T\}$. For $t \in \{T_0 + 1, \dots, T\}$, \hat{u}_t are the post-intervention estimates of the treatment effects; and for $t \in \mathcal{B}$, \hat{u}_t are the placebo treatment effects estimated for the blank periods. Let

$$u_t = \sum_{j=1}^J w_j^* \epsilon_{jt} - \sum_{j=1}^J v_j^* \epsilon_{jt}$$

for $t \in \mathcal{B}$, and

$$u_t = \sum_{j=1}^J w_j^* \xi_{jt} - \sum_{j=1}^J v_j^* \epsilon_{jt}$$

for $t \in \{T_0 + 1, \dots, T\}$. The null hypothesis (15) and the assumptions of Theorem 2 imply that $\{u_t\}_{t \in \mathcal{B} \cup \{T_0 + 1, \dots, T\}}$ is a sequence of exchangeable random variables. Additionally, Assumption 1 and the null hypothesis (15) imply

$$\hat{u}_t = \sum_{j=1}^J w_j^* Y_{jt} - \sum_{j=1}^J v_j^* Y_{jt} = \boldsymbol{\theta}'_t \sum_{j=1}^J (w_j^* - v_j^*) \mathbf{Z}_j + \boldsymbol{\lambda}'_t \sum_{j=1}^J (w_j^* - v_j^*) \boldsymbol{\mu}_j + u_t,$$

for $t \in \mathcal{B} \cup \{T_0 + 1, \dots, T\}$. The result of the theorem then follows from Theorem D.1 in Chernozhukov, Wüthrich and Zhu (2021). ■

B. Swapping Treated and Control Weights

Recall that when it is possible to swap synthetic treated and synthetic control weights, we choose the treated units so that the number of units with positive weights in \mathbf{w}^* is smaller than the number of units with positive weights in \mathbf{v}^* . When $\|\mathbf{w}^*\|_0 = \|\mathbf{v}^*\|_0$, we determine whether or not to swap using the following rule. For the *Unconstrained* design, we choose the treated group to be the one with the smallest index among the units with positive weights. We use the same procedure based on lowest index for *Constrained* with $\bar{m} = 7$ (highest value) and *Penalized* with $\lambda = 0.01$ (lowest value). Then, starting from $\bar{m} = 7$ and for smaller values of \bar{m} , we assign to the treated group the set of weights that is most similar to the weights obtained for $\|\mathbf{w}^*\|_0 \leq \bar{m} + 1$ (in terms of what units obtain positive weights). In those cases where the two sets of swappable weights for $\|\mathbf{w}^*\|_0 \leq \bar{m}$ are equally similar to the synthetic treated weights for $\|\mathbf{w}^*\|_0 \leq \bar{m} + 1$, we select the set of weights with the smallest index. We follow the analogous procedure for $\lambda > 0.01$, starting from smaller values of λ .

Online Appendix
Synthetic Controls for Experimental Design

Alberto Abadie and Jinglong Zhao

December 4, 2023

OA.1. Designs Based on Penalized and Bias-corrected Synthetic Control Methods

Consider the design problem in (7),

$$\underbrace{\left\| \bar{\mathbf{X}} - \sum_{j=1}^J w_j \mathbf{X}_j \right\|^2}_{(a)} + \xi \sum_{j=1}^J w_j \underbrace{\left\| \mathbf{X}_j - \sum_{i=1}^J v_{ij} \mathbf{X}_i \right\|^2}_{(b)}. \quad (\text{OA.1})$$

To apply the penalized synthetic control method of Abadie and L'Hour (2021) to this design, we replace the term (a) in (OA.1) with

$$\left\| \bar{\mathbf{X}} - \sum_{j=1}^J w_j \mathbf{X}_j \right\|^2 + \lambda_1 \sum_{j=1}^J w_j \|\bar{\mathbf{X}} - \mathbf{X}_j\|^2, \quad (\text{OA.2})$$

and the terms (b) with

$$\left\| \mathbf{X}_j - \sum_{i=1}^J v_{ij} \mathbf{X}_i \right\|^2 + \lambda_2 \sum_{i=1}^J v_{ij} \|\mathbf{X}_j - \mathbf{X}_i\|^2. \quad (\text{OA.3})$$

Here, λ_1 and λ_2 are positive constants that penalize discrepancies between the target values of the predictors ($\bar{\mathbf{X}}$ in (OA.2) and \mathbf{X}_j in (OA.3)) and the values of the predictors for the units that contribute to their synthetic counterparts.

All designs of Section 2 depend on terms akin to (a) and (b) in (OA.1). These terms can be adapted as in (OA.2) and (OA.3) to implement the penalized synthetic control design of Abadie and L'Hour (2021).

For all the designs in Section 2, the bias-corrected estimator of Abadie and L'Hour (2021) is

$$\widehat{\tau}_t^{BC} = \sum_{j=1}^J w_j^* (Y_{jt} - \widehat{\mu}_0(\mathbf{X}_j)) - \sum_{j=1}^J v_j^* (Y_{jt} - \widehat{\mu}_0(\mathbf{X}_j)),$$

where $t \geq T_0 + 1$ and the terms $\widehat{\mu}_{0t}(\mathbf{X}_j)$ are the fitted values of a regression of untreated outcomes, Y_{jt}^N , on unit's characteristics, \mathbf{X}_j . To avoid over-fitting biases, $\widehat{\mu}_{0t}(\mathbf{X}_j)$ can be cross-fitted for the untreated.

OA.2. Approximate Validity when λ_t are not Exchangeable

Recall that in Theorem 2 we have shown that when λ_t are exchangeable for $t \in \mathcal{B} \cup \{T_0 + 1, \dots, T\}$ the p -value in (17) is exact. In this section, we discuss the case when λ_t are not necessarily exchangeable. We show below in Theorem OA.1 that the p -value in (17) is approximately valid when $T_{\mathcal{E}}$ is large.

Theorem OA.1 *Suppose that Assumptions 1 – 3 hold, and that for $j = 1, \dots, J$, $t = 1, \dots, T$, ϵ_{jt} are continuously distributed with (a version of) the probability density function bounded by a constant $\kappa < \infty$. Then, under the null hypothesis (15), the p -values of equation (17) are approximately valid. In particular, there is an event \mathcal{C} , such that conditional on \mathcal{C} , for any $\alpha \in (0, 1]$, we have*

$$\alpha - 2z_2 - \frac{1}{|\Pi|} \leq \Pr(\hat{p} \leq \alpha) \leq \alpha + 2z_2,$$

and the event \mathcal{C} happens with probability equal or greater than

$$\Pr(\mathcal{C}) = 1 - 2J \exp\left(-\frac{z_1^2 \zeta^2}{8\bar{\sigma}^2 \bar{\lambda}^4 F^2} T_{\mathcal{E}}\right) - \frac{z_1}{z_2} \cdot 2e\sqrt{2J \min\{T - T_0, T_0 - T_{\mathcal{E}}\}} \kappa,$$

where z_1, z_2 are arbitrary positive constants.

A limitation of the result in Theorem OA.1 is that there are values of the parameters of the data generating for which the result of the theorem provides a tight bound on test size only for large values of $T_{\mathcal{E}}$. We prove Theorem OA.1 next.

OA.2.1. Definitions

First, define $T_p = \min\{T - T_0, T_0 - T_{\mathcal{E}}\}$. Next, recall that

$$\hat{u}_t = \sum_{j=1}^J w_j^* Y_{jt} - \sum_{j=1}^J v_j^* Y_{jt},$$

for $t \in \mathcal{B} \cup \{T_0 + 1, \dots, T\}$. For $t \in \{T_0 + 1, \dots, T\}$, \hat{u}_t are the post-intervention estimates of the treatment effects; and for $t \in \mathcal{B}$, \hat{u}_t are the placebo treatment effects estimated for the blank periods.

Let

$$u_t = \sum_{j=1}^J w_j^* \epsilon_{jt} - \sum_{j=1}^J v_j^* \epsilon_{jt} \quad (\text{OA.4})$$

for $t \in \mathcal{B}$, and

$$u_t = \sum_{j=1}^J w_j^* \zeta_{jt} - \sum_{j=1}^J v_j^* \epsilon_{jt} \quad (\text{OA.5})$$

for $t \in \{T_0 + 1, \dots, T\}$. For each $\pi \in \Pi$, similar to our definition of $\widehat{\mathbf{e}}_\pi$, define the $(T - T_0)$ -dimensional vector

$$\mathbf{e}_\pi = (u_{\pi(1)}, u_{\pi(2)}, \dots, u_{\pi(T-T_0)}).$$

In addition, let $\mathbf{e} = (u_1, \dots, u_{T-T_0}) = (\tau_{T_0+1}, \dots, \tau_T)$. It is useful to observe that, under the null hypothesis in (15), the random variables u_t for $t \in \mathcal{B} \cup \{T_0 + 1, \dots, T\}$ are independent and identically distributed.

Next, define the following two functions. Let

$$\widehat{F}(x) = \frac{1}{|\Pi|} \sum_{\pi \in \Pi} \mathbb{1} \{S(\widehat{\mathbf{e}}_\pi) < x\},$$

and

$$\widetilde{F}(x) = \frac{1}{|\Pi|} \sum_{\pi \in \Pi} \mathbb{1} \{S(\mathbf{e}_\pi) < x\}.$$

The proof of Theorem OA.1 proceeds in four steps. In step one, we define a high probability event, \mathcal{C}_1 , such that u_t and \widehat{u}_t are close to each other under \mathcal{C}_1 . In step two, we define a high probability event, \mathcal{C}_2 , such that many components of $\{S(\mathbf{e}_\pi)\}_{\pi \in \Pi}$ are well-separated from $S(\mathbf{e})$ under \mathcal{C}_2 . In step three, we show that, conditional on \mathcal{C}_1 and \mathcal{C}_2 , the ordering of $S(\mathbf{e}_\pi)$ and $S(\mathbf{e})$ will be the same as the ordering of $S(\widehat{\mathbf{e}}_\pi)$ and $S(\widehat{\mathbf{e}})$ for most $\pi \in \Pi$, which implies that $\widehat{F}(S(\widehat{\mathbf{e}}))$ and $\widetilde{F}(S(\mathbf{e}))$ are also close to each other. In step four, we conclude the proof by linking $\widehat{F}(S(\widehat{\mathbf{e}}))$ to the estimated p -value, and $\widetilde{F}(S(\mathbf{e}))$ to the nominal level α .

OA.2.2. Lemmas for the Proof of Theorem OA.1

Lemma OA.2 (Corollary 2, Bobkov and Chistyakov (2014)) *Let X_1, X_2, \dots, X_n be independent and continuously distributed random variables with densities $f_{X_1}, f_{X_2}, \dots, f_{X_n}$. For any $k \in \{1, 2, \dots, n\}$, let Λ_{X_k} be the smallest upper bound on the probability density f_{X_k} . Suppose that for any $k \in \{1, 2, \dots, n\}$, $\Lambda_{X_k} \leq \kappa$. Then for any a_1, a_2, \dots, a_n such that $\sum_{k=1}^n a_k^2 = 1$,*

$$\Lambda_{(a_1 X_1 + a_2 X_2 + \dots + a_n X_n)} \leq \sqrt{e} \kappa.$$

Lemma OA.3 *Let X be a continuously distributed random variable with a density f_X . Let Λ_X be the smallest upper bound on the probability density f_X .*

1. *The random variable $|X|$ has a density $f_{|X|}$ bounded by $\Lambda_{|X|} \leq 2\Lambda_X$;*
2. *For any constant $a \neq 0$, the random variable aX has a density, f_{aX} , bounded by $\Lambda_{aX} \leq \Lambda_X/a$;*

Proof of Lemma OA.3. To prove 1, note that for any $v \geq 0$,

$$f_{|X|}(v) = f_X(v) + f_X(-v) \leq 2\Lambda_X.$$

To prove 2, note that for any $v \geq 0$,

$$f_{aX}(v) = \frac{1}{|a|} f_X(v/a) \leq \frac{1}{|a|} \Lambda_X.$$

■

Lemma OA.4 *Recall that u_t is defined as (OA.4) and (OA.5), for the blank periods and the experimental periods, respectively. The probability density of u_t can be bounded by*

$$\Lambda_{u_t} \leq \frac{1}{2} \sqrt{eJ} \kappa.$$

Proof of Lemma OA.4. Fix any $t \in \mathcal{B} \cup \{T_0 + 1, \dots, T\}$. Using Lemma OA.2 (Bobkov and Chistyakov, 2014, Corollary 2), let there be J variables ϵ_{jt} for any $j \in \{1, 2, \dots, J\}$. For any

$j \in \{1, 2, \dots, J\}$, define $a_j = \frac{w_j^* - v_j^*}{\sqrt{\sum_{j=1}^J (w_j^* - v_j^*)^2}}$ such that $\sum_{j=1}^J a_j^2 = 1$. Using a_j , we can write u_t as $u_t = \sqrt{\sum_{j=1}^J (w_j^* - v_j^*)^2} \cdot \sum_{j=1}^J a_j \epsilon_{jt}$. When J is even,

$$\begin{aligned}
\Lambda_{u_t} &\leq \frac{1}{\sqrt{\sum_{j=1}^J (w_j^* - v_j^*)^2}} \cdot \Lambda_{\sum_{j=1}^J a_j \epsilon_{jt}} \\
&\leq \frac{1}{\sqrt{\sum_{j=1}^J (w_j^* - v_j^*)^2}} \cdot \sqrt{e\kappa} \\
&\leq \frac{1}{\sqrt{\sum_{j=1}^J (\frac{2}{J})^2}} \cdot \sqrt{e\kappa} \\
&= \frac{\sqrt{J}}{2} \sqrt{e\kappa},
\end{aligned}$$

where the first inequality is due to Lemma OA.3 Part 2; the second inequality is due to Lemma OA.2; the third inequality is due to convexity and Jensen's inequality, and the worst case is taken when $w_j^* = 2/J$ for one half of total units and $v_j^* = 2/J$ for the other half. When J is odd,

$$\begin{aligned}
\Lambda_{u_t} &\leq \frac{1}{\sqrt{\sum_{j=1}^J (w_j^* - v_j^*)^2}} \cdot \Lambda_{\sum_{j=1}^J a_j \epsilon_{jt}} \\
&\leq \frac{1}{\sqrt{\sum_{j=1}^J (w_j^* - v_j^*)^2}} \cdot \sqrt{e\kappa} \\
&\leq \frac{1}{\sqrt{\frac{J+1}{2} (\frac{2}{J+1})^2 + \frac{J-1}{2} (\frac{2}{J-1})^2}} \cdot \sqrt{e\kappa} \\
&= \sqrt{\frac{J^2 - 1}{J}} \cdot \frac{\sqrt{e\kappa}}{2}, \\
&\leq \frac{\sqrt{J}}{2} \sqrt{e\kappa},
\end{aligned}$$

where the first inequality is due to Lemma OA.3 Part 2; the second inequality is due to Lemma OA.2; the third inequality is due to convexity and Jensen's inequality, and the worst case is taken when $w_j^* = 2/(J+1)$ for $(J+1)/2$ of total units and $v_j^* = 2/(J-1)$ for the other $(J-1)/2$ of total units. ■

OA.2.3. Proof of Theorem OA.1

Proof of Theorem OA.1. (Step one.) Note that

$$\begin{aligned} \sum_{j=1}^J w_j^* \mathbf{Y}_j^\mathcal{E} - \sum_{j=1}^J v_j^* \mathbf{Y}_j^\mathcal{E} &= \boldsymbol{\theta}_\mathcal{E} \left(\sum_{j=1}^J w_j^* \mathbf{Z}_j - \sum_{j=1}^J v_j^* \mathbf{Z}_j \right) \\ &\quad + \boldsymbol{\lambda}_\mathcal{E} \left(\sum_{j=1}^J w_j^* \boldsymbol{\mu}_j - \sum_{j=1}^J v_j^* \boldsymbol{\mu}_j \right) + \left(\sum_{j=1}^J w_j^* \boldsymbol{\epsilon}_j^\mathcal{E} - \sum_{j=1}^J v_j^* \boldsymbol{\epsilon}_j^\mathcal{E} \right). \end{aligned}$$

Assumption 3 implies

$$\left(\sum_{j=1}^J w_j^* \boldsymbol{\mu}_j - \sum_{j=1}^J v_j^* \boldsymbol{\mu}_j \right) = -(\boldsymbol{\lambda}'_\mathcal{E} \boldsymbol{\lambda}_\mathcal{E})^{-1} \boldsymbol{\lambda}'_\mathcal{E} \left(\sum_{j=1}^J w_j^* \boldsymbol{\epsilon}_j^\mathcal{E} - \sum_{j=1}^J v_j^* \boldsymbol{\epsilon}_j^\mathcal{E} \right).$$

Under the null hypothesis (15), it follows that

$$\begin{aligned} \hat{u}_t &= \sum_{j=1}^J w_j^* Y_{jt} - \sum_{j=1}^J v_j^* Y_{jt} \\ &= -\boldsymbol{\lambda}'_t (\boldsymbol{\lambda}'_\mathcal{E} \boldsymbol{\lambda}_\mathcal{E})^{-1} \boldsymbol{\lambda}'_\mathcal{E} \sum_{j=1}^J (w_j^* - v_j^*) \boldsymbol{\epsilon}_j^\mathcal{E} + u_t, \end{aligned}$$

for $t \in \mathcal{B} \cup \{T_0 + 1, \dots, T\}$. We next define an event

$$\begin{aligned} \mathcal{C}_1 &= \left\{ \forall t \in \mathcal{B} \cup \{T_0 + 1, \dots, T\}, \left| \boldsymbol{\lambda}'_t (\boldsymbol{\lambda}'_\mathcal{E} \boldsymbol{\lambda}_\mathcal{E})^{-1} \boldsymbol{\lambda}'_\mathcal{E} \sum_{j=1}^J (w_j^* - v_j^*) \boldsymbol{\epsilon}_j^\mathcal{E} \right| \leq z_1 \right\} \\ &= \left\{ \max_{t \in \mathcal{B} \cup \{T_0 + 1, \dots, T\}} \left| \boldsymbol{\lambda}'_t (\boldsymbol{\lambda}'_\mathcal{E} \boldsymbol{\lambda}_\mathcal{E})^{-1} \boldsymbol{\lambda}'_\mathcal{E} \sum_{j=1}^J (w_j^* - v_j^*) \boldsymbol{\epsilon}_j^\mathcal{E} \right| \leq z_1 \right\}. \end{aligned}$$

Note that,

$$\begin{aligned} &\max_{t \in \mathcal{B} \cup \{T_0 + 1, \dots, T\}} \left| \boldsymbol{\lambda}'_t (\boldsymbol{\lambda}'_\mathcal{E} \boldsymbol{\lambda}_\mathcal{E})^{-1} \boldsymbol{\lambda}'_\mathcal{E} \sum_{j=1}^J (w_j^* - v_j^*) \boldsymbol{\epsilon}_j^\mathcal{E} \right| \\ &\leq \max_{t \in \mathcal{B} \cup \{T_0 + 1, \dots, T\}} \sum_{j=1}^J |w_j^* - v_j^*| \sum_{s \in \mathcal{E}} |\boldsymbol{\lambda}'_t (\boldsymbol{\lambda}'_\mathcal{E} \boldsymbol{\lambda}_\mathcal{E})^{-1} \boldsymbol{\lambda}_s| |\epsilon_{js}| \\ &\leq \sum_{j=1}^J |w_j^* - v_j^*| \sum_{s \in \mathcal{E}} \frac{\bar{\lambda}^2 F}{T_\mathcal{E} \zeta} |\epsilon_{js}|, \end{aligned}$$

where the first inequality is due to (A.6), and the second inequality is because $|w_j^* - v_j^*| \geq 0$ and $\epsilon_{js} \geq 0$. Therefore,

$$\begin{aligned} \Pr(\mathcal{C}_1) &\geq 1 - \Pr\left(\sum_{j=1}^J \frac{|w_j^* - v_j^*|}{2} \sum_{s \in \mathcal{E}} \frac{\bar{\lambda}^2 F}{T_{\mathcal{E}} \zeta} |\epsilon_{js}| > \frac{z_1}{2}\right) \\ &\geq 1 - \sum_{j=1}^J \Pr\left(\sum_{s \in \mathcal{E}} \frac{\bar{\lambda}^2 F}{T_{\mathcal{E}} \zeta} |\epsilon_{js}| > \frac{z_1}{2}\right) \\ &\geq 1 - 2J \exp\left(-\frac{z_1^2 \zeta^2}{8\bar{\sigma}^2 \bar{\lambda}^4 F^2} T_{\mathcal{E}}\right), \end{aligned}$$

where the second inequality follows from union bound, and the third inequality is the Chernoff bound for sub-Gaussian random variables.

(Step two.) Define $\tilde{z}_1 = 2z_1 > 0$. For each $k \in \{0, 1, 2, \dots, T - T_0\}$, we define the following sets of permutations. First, define $\Pi_0 = \{\pi_0\}$, where π_0 is defined as the set of post-intervention indices $\pi_0 = \{T_0 + 1, \dots, T\}$. Then, for any $k \in \{1, 2, \dots, T - T_0\}$, define

$$\Pi_k = \left\{ \pi \in \Pi \mid |\pi \setminus \pi_0| = k \right\}$$

to be the set of $(T - T_0)$ -combinations with exactly k many indices from the blank periods. Using the above definitions, we can decompose Π into

$$\Pi = \bigcup_{k=0}^{T_p} \Pi_k,$$

where $T_p = \min\{T - T_0, T_0 - T_{\mathcal{E}}\}$.

Then, for any $k \in \{1, 2, \dots, T_p\}$ and $\pi \in \Pi_k$, we focus on the following indicator

$$\mathbb{1}\left\{ \left| \sum_{t \in \pi \setminus \pi_0} |u_t| - \sum_{t \in \pi_0 \setminus \pi} |u_t| \right| \leq 2kz_1 \right\}.$$

The above indicator involves a total of $2k$ many $|u_t|$'s. Intuitively, it is obtained by canceling out common terms in $S(\mathbf{e}_\pi)$ and $S(\mathbf{e})$.

Below we focus on the properties of the sum of such indicators. First, focus on the probability density of $\left| \sum_{t \in \pi \setminus \pi_0} |u_t| - \sum_{t \in \pi_0 \setminus \pi} |u_t| \right|$. We have

$$\Lambda_{\left| \sum_{t \in \pi \setminus \pi_0} |u_t| - \sum_{t \in \pi_0 \setminus \pi} |u_t| \right|} \leq 2\Lambda_{\sum_{t \in \pi \setminus \pi_0} |u_t| - \sum_{t \in \pi_0 \setminus \pi} |u_t|}$$

$$\begin{aligned}
&\leq \frac{2}{\sqrt{2k}} \Lambda_{\sum_{t \in \pi \setminus \pi_0} \frac{1}{\sqrt{2k}} |u_t| - \sum_{t \in \pi_0 \setminus \pi} \frac{1}{\sqrt{2k}} |u_t|} \\
&\leq \frac{2}{\sqrt{2k}} \sqrt{e} \Lambda_{|u_t|} \\
&\leq \frac{2}{\sqrt{2k}} \sqrt{e} \sqrt{eJ\kappa} \\
&= \frac{\sqrt{2J}}{\sqrt{k}} e\kappa,
\end{aligned}$$

where the first inequality is due to Lemma OA.3-1; the second inequality is due to Lemma OA.2-2; the third inequality is due to Lemma OA.2; the last inequality is due to Lemma OA.4 and OA.3-1. We obtain

$$\Pr \left(\left| \sum_{t \in \pi \setminus \pi_0} |u_t| - \sum_{t \in \pi_0 \setminus \pi} |u_t| \right| \leq 2kz_1 \right) \leq 2e\sqrt{2Jk}z_1\kappa.$$

Next, due to Markov inequality, for any constant $z_2 > 0$, we have

$$\begin{aligned}
&\Pr \left(\sum_{k=1}^{T_p} \sum_{\pi \in \Pi_k} \mathbb{1} \left\{ \left| \sum_{t \in \pi \setminus \pi_0} |u_t| - \sum_{t \in \pi_0 \setminus \pi} |u_t| \right| \leq 2kz_1 \right\} \geq |\Pi|z_2 \right) \\
&\leq \frac{1}{|\Pi|z_2} \sum_{k=1}^{T_p} \sum_{\pi \in \Pi_k} E \left[\mathbb{1} \left\{ \left| \sum_{t \in \pi \setminus \pi_0} |u_t| - \sum_{t \in \pi_0 \setminus \pi} |u_t| \right| \leq 2kz_1 \right\} \right] \\
&\leq \frac{\sum_{k=1}^{T_p} |\Pi_k| 2e\sqrt{2Jk}z_1\kappa}{|\Pi|z_2}.
\end{aligned}$$

To conclude step two, define the event

$$\mathcal{C}_2 = \left\{ \sum_{k=1}^{T_p} \sum_{\pi \in \Pi_k} \mathbb{1} \left\{ \left| \sum_{t \in \pi \setminus \pi_0} |u_t| - \sum_{t \in \pi_0 \setminus \pi} |u_t| \right| \leq 2kz_1 \right\} < |\Pi|z_2 \right\}. \quad (\text{OA.6})$$

The probability that event \mathcal{C}_2 happens is at least

$$\Pr(\mathcal{C}_2) \geq 1 - \frac{\sum_{k=1}^{T_p} |\Pi_k| \sqrt{k} 2e\sqrt{2J}z_1\kappa}{|\Pi|z_2}.$$

(Step three.) Conditional on event \mathcal{C}_2 , fewer than $|\Pi|z_2$ of the absolute value terms in (OA.6) are such that $\left| \sum_{t \in \pi \setminus \pi_0} |u_t| - \sum_{t \in \pi_0 \setminus \pi} |u_t| \right| \leq 2kz_1$. For all the others, $\left| \sum_{t \in \pi \setminus \pi_0} |u_t| - \sum_{t \in \pi_0 \setminus \pi} |u_t| \right| > 2kz_1$.

Conditional on event \mathcal{C}_1 , we know that $|\hat{u}_t - u_t| \leq z_1$ for any $t \in \mathcal{B} \cup \{T_0 + 1, \dots, T\}$. So we have that $\sum_{t \in \pi \setminus \pi_0} |u_t| - \sum_{t \in \pi_0 \setminus \pi} |u_t| > 2kz_1$ implies

$$\begin{aligned}
S(\hat{\mathbf{e}}_\pi) - S(\hat{\mathbf{e}}) &= \frac{1}{T - T_0} \sum_{t \in \pi} |\hat{u}_t| - \frac{1}{T - T_0} \sum_{t \in \pi_0} |\hat{u}_t| \\
&= \frac{1}{T - T_0} \left(\sum_{t \in \pi \setminus \pi_0} |\hat{u}_t| - \sum_{t \in \pi_0 \setminus \pi} |\hat{u}_t| \right) \\
&\geq \frac{1}{T - T_0} \left(\sum_{t \in \pi \setminus \pi_0} (|u_t| - z_1) - \sum_{t \in \pi_0 \setminus \pi} (|u_t| + z_1) \right) \\
&> \frac{1}{T - T_0} (2kz_1 - 2kz_1) \\
&= 0,
\end{aligned}$$

where the first equality is due to definition $S(\mathbf{e}_\pi) = \frac{1}{T - T_0} \sum_{t \in \pi} |u_t|$. Similarly, \mathcal{C}_1 and $\sum_{t \in \pi \setminus \pi_0} |u_t| - \sum_{t \in \pi_0 \setminus \pi} |u_t| < -2kz_1$ implies

$$\begin{aligned}
S(\hat{\mathbf{e}}_\pi) - S(\hat{\mathbf{e}}) &= \frac{1}{T - T_0} \left(\sum_{t \in \pi \setminus \pi_0} |\hat{u}_t| - \sum_{t \in \pi_0 \setminus \pi} |\hat{u}_t| \right) \\
&\leq \frac{1}{T - T_0} \left(\sum_{t \in \pi \setminus \pi_0} (|u_t| + z_1) - \sum_{t \in \pi_0 \setminus \pi} (|u_t| - z_1) \right) \\
&< \frac{1}{T - T_0} (-2kz_1 + 2kz_1) \\
&= 0.
\end{aligned}$$

Combining both cases, we know that conditional on \mathcal{C}_1 and when $\left| \sum_{t \in \pi \setminus \pi_0} |u_t| - \sum_{t \in \pi_0 \setminus \pi} |u_t| \right| > 2kz_1$, the ordering of $S(\mathbf{e}_\pi)$ and $S(\mathbf{e})$ is the same as the ordering of $S(\hat{\mathbf{e}}_\pi)$ and $S(\hat{\mathbf{e}})$. As a result, for those π such that $\left| \sum_{t \in \pi \setminus \pi_0} |u_t| - \sum_{t \in \pi_0 \setminus \pi} |u_t| \right| > 2kz_1$, we have $\mathbb{1} \{S(\hat{\mathbf{e}}_\pi) \geq S(\hat{\mathbf{e}})\} = \mathbb{1} \{S(\mathbf{e}_\pi) \geq S(\mathbf{e})\}$. There are at most $|\Pi|z_2$ many π 's that contribute to the following summation,

$$\left| \sum_{k=1}^{T_p} \sum_{\pi \in \Pi_k} \left(\mathbb{1} \{S(\hat{\mathbf{e}}_\pi) \geq S(\hat{\mathbf{e}})\} - \mathbb{1} \{S(\mathbf{e}_\pi) \geq S(\mathbf{e})\} \right) \right| < |\Pi|z_2.$$

Note that $S(\hat{\mathbf{e}}_{\pi_0}) = S(\hat{\mathbf{e}})$ and $S(\mathbf{e}_{\pi_0}) = S(\mathbf{e})$, so $\mathbb{1} \{S(\hat{\mathbf{e}}_\pi) > S(\hat{\mathbf{e}})\} = \mathbb{1} \{S(\mathbf{e}_{\pi_0}) > S(\mathbf{e})\}$ is always

true. Combining π_0 we have

$$\begin{aligned} & \left| \sum_{\pi \in \Pi} \left(\mathbb{1}\{S(\hat{\mathbf{e}}_\pi) \geq S(\hat{\mathbf{e}})\} - \mathbb{1}\{S(\mathbf{e}_\pi) \geq S(\mathbf{e})\} \right) \right| \\ &= \left| \sum_{k=0}^{T_p} \sum_{\pi \in \Pi_k} \left(\mathbb{1}\{S(\hat{\mathbf{e}}_\pi) \geq S(\hat{\mathbf{e}})\} - \mathbb{1}\{S(\mathbf{e}_\pi) \geq S(\mathbf{e})\} \right) \right| \\ &< |\Pi| z_2. \end{aligned}$$

We conclude step three using the following block of inequalities. For any $\alpha \in (0, 1]$,

$$\begin{aligned} & \left| \Pr \left(1 - \hat{F}(S(\hat{\mathbf{e}})) \leq \alpha \right) - \Pr \left(1 - \tilde{F}(S(\mathbf{e})) \leq \alpha \right) \right| \\ &= \left| \Pr \left(1 - \frac{1}{|\Pi|} \sum_{\pi \in \Pi} \mathbb{1}\{S(\hat{\mathbf{e}}_\pi) < S(\hat{\mathbf{e}})\} \leq \alpha \right) - \Pr \left(1 - \frac{1}{|\Pi|} \sum_{\pi \in \Pi} \mathbb{1}\{S(\mathbf{e}_\pi) < S(\mathbf{e})\} \leq \alpha \right) \right| \\ &= \left| \Pr \left(\sum_{\pi \in \Pi} \mathbb{1}\{S(\hat{\mathbf{e}}_\pi) \geq S(\hat{\mathbf{e}})\} \leq \alpha |\Pi| \right) - \Pr \left(\sum_{\pi \in \Pi} \mathbb{1}\{S(\mathbf{e}_\pi) \geq S(\mathbf{e})\} \leq \alpha |\Pi| \right) \right| \\ &= \left| E \left[\mathbb{1} \left\{ \sum_{\pi \in \Pi} \mathbb{1}\{S(\hat{\mathbf{e}}_\pi) \geq S(\hat{\mathbf{e}})\} \leq \alpha |\Pi| \right\} \right] - E \left[\mathbb{1} \left\{ \sum_{\pi \in \Pi} \mathbb{1}\{S(\mathbf{e}_\pi) \geq S(\mathbf{e})\} \leq \alpha |\Pi| \right\} \right] \right| \\ &\leq E \left| \mathbb{1} \left\{ \sum_{\pi \in \Pi} \mathbb{1}\{S(\hat{\mathbf{e}}_\pi) \geq S(\hat{\mathbf{e}})\} \leq \alpha |\Pi| \right\} - \mathbb{1} \left\{ \sum_{\pi \in \Pi} \mathbb{1}\{S(\mathbf{e}_\pi) \geq S(\mathbf{e})\} \leq \alpha |\Pi| \right\} \right| \\ &\leq \Pr \left(\left| \alpha |\Pi| - \sum_{\pi \in \Pi} \mathbb{1}\{S(\mathbf{e}_\pi) \geq S(\mathbf{e})\} \right| \leq \left| \sum_{\pi \in \Pi} \left(\mathbb{1}\{S(\hat{\mathbf{e}}_\pi) \geq S(\hat{\mathbf{e}})\} - \mathbb{1}\{S(\mathbf{e}_\pi) \geq S(\mathbf{e})\} \right) \right| \right), \end{aligned}$$

where the second inequality is due to the following: $|\mathbb{1}\{a \leq c\} - \mathbb{1}\{b \leq c\}| \leq \mathbb{1}\{|c - b| \leq |a - b|\}$.

Conditional on events \mathcal{C}_1 and \mathcal{C}_2 , we obtain

$$\begin{aligned} & \left| \Pr \left(1 - \hat{F}(S(\hat{\mathbf{e}})) \leq \alpha \right) - \Pr \left(1 - \tilde{F}(S(\mathbf{e})) \leq \alpha \right) \right| \\ &\leq \Pr \left(\left| \alpha |\Pi| - \sum_{\pi \in \Pi} \mathbb{1}\{S(\mathbf{e}_\pi) \geq S(\mathbf{e})\} \right| < |\Pi| z_2 \right) \\ &\leq \frac{2|\Pi| z_2}{|\Pi|} \\ &= 2z_2, \end{aligned} \tag{OA.7}$$

where the last inequality is because $\sum_{\pi \in \Pi} \mathbb{1}\{S(\mathbf{e}_\pi) \geq S(\mathbf{e})\}$ is a discrete uniform distribution over $\{1, 2, \dots, |\Pi|\}$, and that there are at most $2|\Pi|z_2$ many integers centered around $\alpha|\Pi|$.

(Step four.) Note that, for any $\alpha \in (0, 1]$,

$$\alpha - \frac{1}{|\Pi|} \leq \Pr\left(1 - \tilde{F}(S(\mathbf{e})) \leq \alpha\right) \leq \alpha.$$

So conditional on events \mathcal{C}_1 and \mathcal{C}_2 , (OA.7) implies

$$\Pr\left(1 - \widehat{F}(S(\widehat{\mathbf{e}})) \leq \alpha\right) \leq \Pr\left(1 - \tilde{F}(S(\mathbf{e})) \leq \alpha\right) + 2z_2 \leq \alpha + 2z_2$$

and

$$\Pr\left(1 - \widehat{F}(S(\widehat{\mathbf{e}})) \leq \alpha\right) \geq \Pr\left(1 - \tilde{F}(S(\mathbf{e})) \leq \alpha\right) - 2z_2 \geq \alpha - 2z_2 - \frac{1}{|\Pi|}.$$

Combining both parts, conditional on $\mathcal{C} = \mathcal{C}_1 \cap \mathcal{C}_2$, we have

$$\alpha - 2z_2 - \frac{1}{|\Pi|} \leq \Pr(\widehat{p} \leq \alpha) = \Pr\left(1 - \widehat{F}(S(\widehat{\mathbf{e}})) \leq \alpha\right) \leq \alpha + 2z_2,$$

and \mathcal{C} happens with probability at least

$$\begin{aligned} \Pr(\mathcal{C}_1 \cap \mathcal{C}_2) &\geq (1 - \Pr(\mathcal{C}_1)) + (1 - \Pr(\mathcal{C}_2)) - 1 \\ &\geq 1 - 2J \exp\left(-\frac{z_1^2 \zeta^2}{8\bar{\sigma}^2 \lambda^4 F^2} T_{\mathcal{E}}\right) - \frac{\sum_{k=1}^{T_p} |\Pi_k| \sqrt{k}}{|\Pi|} \cdot \frac{z_1}{z_2} \cdot 2e\sqrt{2J\kappa}, \end{aligned}$$

which finishes the proof. ■

OA.3. Estimating the Average Effect of Treatment on the Treated Units

In Section 3 we have shown formal results of the bias bound in estimating the average treatment effect. In this section, we present similar results for estimating the average effect of treatment on the treated units. Similar to Assumption 3, we start with a perfect fit assumption.

Assumption 5 *With probability one, (i)*

$$\sum_{j=1}^J w_j^* \mathbf{Z}_j = \sum_{j=1}^J v_j^* \mathbf{Z}_j,$$

and (ii)

$$\sum_{j=1}^J w_j^* \mathbf{Y}_j^{\mathcal{E}} = \sum_{j=1}^J v_j^* \mathbf{Y}_j^{\mathcal{E}}.$$

In practice, Assumption 3 may only hold approximately. The next assumption accommodates settings with imperfect fit.

Assumption 6 *There exists a positive constant $d > 0$, such that with probability one,*

$$\left\| \sum_{j=1}^J w_j^* \mathbf{Z}_j - \sum_{j=1}^J v_j^* \mathbf{Z}_j \right\|_2^2 \leq Rd^2, \quad \left\| \sum_{j=1}^J w_j^* \mathbf{Y}_j^\mathcal{E} - \sum_{j=1}^J v_j^* \mathbf{Y}_j^\mathcal{E} \right\|_2^2 \leq T_\mathcal{E} d^2.$$

Using the above assumptions, we are able to provide the following bias bound.

Theorem OA.5 *If Assumptions 1, 2, and 5 hold, then for any $t \geq T_0 + 1$,*

$$|E[\widehat{\tau}_t^T - \tau_t^T]| \leq \frac{\bar{\lambda}^2 F}{\underline{\zeta}} 2\sqrt{2 \log(2J)} \frac{\bar{\sigma}}{\sqrt{T_\mathcal{E}}}.$$

If Assumptions 1, 2, and 6 hold, then for any $t \geq T_0 + 1$,

$$|E[\widehat{\tau}_t - \tau_t]| \leq \left(\bar{\theta}R + \frac{\bar{\lambda}^2 F}{\underline{\zeta}} (1 + \bar{\theta}R) \right) d + \frac{\bar{\lambda}^2 F}{\underline{\zeta}} 2\sqrt{2 \log(2J)} \frac{\bar{\sigma}}{\sqrt{T_\mathcal{E}}}.$$

We provide the following result on inference.

Theorem OA.6 *Suppose that Assumptions 1 and 5(i) hold. Assume that $\{\boldsymbol{\lambda}_t\}_{t \in \mathcal{B} \cup \{T_0+1, \dots, T\}}$ is a sequence of exchangeable random variables independent of $\{\epsilon_{jt}\}_{t \in \mathcal{B} \cup \{T_0+1, \dots, T\}}$ and $\{\xi_{jt}\}_{t \in \{T_0+1, \dots, T\}}$. Assume also that for each $j = 1, \dots, J$, $\{\epsilon_{jt}\}_{t \in \mathcal{B} \cup \{T_0+1, \dots, T\}}$ and $\{\xi_{jt}\}_{t \in \{T_0+1, \dots, T\}}$ are two sequences of exchangeable random variables, respectively. Then under the null hypothesis (15), we have*

$$\alpha - \frac{1}{|\Pi|} \leq \Pr(\widehat{p} \leq \alpha) \leq \alpha,$$

for any $\alpha \in [0, 1]$, where $\Pr(\widehat{p} \leq \alpha)$ is taken over the distributions of $\{\xi_{jt}, \epsilon_{jt}, \boldsymbol{\lambda}_t\}$.

If Assumption 1 holds but 5(i) is violated, the same result holds if $\{(\boldsymbol{\theta}_t, \boldsymbol{\lambda}_t)\}_{t \in \mathcal{B} \cup \{T_0+1, \dots, T\}}$ is a sequence of exchangeable random variables independent of $\{\epsilon_{jt}\}_{t \in \mathcal{B} \cup \{T_0+1, \dots, T\}}$ and $\{\xi_{jt}\}_{t \in \{T_0+1, \dots, T\}}$, and if for each $j = 1, \dots, J$, $\{\epsilon_{jt}\}_{t \in \mathcal{B} \cup \{T_0+1, \dots, T\}}$ and $\{\xi_{jt}\}_{t \in \{T_0+1, \dots, T\}}$ are two sequences of exchangeable random variables, respectively. Here, $\Pr(\widehat{p} \leq \alpha)$ is taken over the distributions of $\{\xi_{jt}, \epsilon_{jt}, \boldsymbol{\lambda}_t, \boldsymbol{\theta}_t\}$.

OA.3.1. Proofs of Theorem OA.5 and Theorem OA.6

Proof of Theorem OA.5. For any period $t = T_0 + 1, \dots, T$ we decompose $(\widehat{\tau}_t^T - \tau_t^T)$ as follows,

$$\widehat{\tau}_t^T - \tau_t^T = \sum_{j=1}^J w_j^* Y_{jt}^N - \sum_{j=1}^J v_j^* Y_{jt}^N.$$

From (10a), we obtain

$$\begin{aligned} \sum_{j=1}^J w_j^* Y_{jt}^N - \sum_{j=1}^J v_j^* Y_{jt}^N &= \boldsymbol{\theta}'_t \left(\sum_{j=1}^J w_j^* \mathbf{Z}_j - \sum_{j=1}^J v_j^* \mathbf{Z}_j \right) \\ &\quad + \boldsymbol{\lambda}'_t \left(\sum_{j=1}^J w_j^* \boldsymbol{\mu}_j - \sum_{j=1}^J v_j^* \boldsymbol{\mu}_j \right) + \left(\sum_{j=1}^J w_j^* \epsilon_{jt} - \sum_{j=1}^J v_j^* \epsilon_{jt} \right). \end{aligned} \quad (\text{OA.8})$$

Similarly, using expression (10a), we obtain

$$\begin{aligned} \sum_{j=1}^J w_j^* \mathbf{Y}_j^\mathcal{E} - \sum_{j=1}^J v_j^* \mathbf{Y}_j^\mathcal{E} &= \boldsymbol{\theta}_\mathcal{E} \left(\sum_{j=1}^J w_j^* \mathbf{Z}_j - \sum_{j=1}^J v_j^* \mathbf{Z}_j \right) \\ &\quad + \boldsymbol{\lambda}_\mathcal{E} \left(\sum_{j=1}^J w_j^* \boldsymbol{\mu}_j - \sum_{j=1}^J v_j^* \boldsymbol{\mu}_j \right) + \left(\sum_{j=1}^J w_j^* \boldsymbol{\epsilon}_j^\mathcal{E} - \sum_{j=1}^J v_j^* \boldsymbol{\epsilon}_j^\mathcal{E} \right), \end{aligned}$$

where $\boldsymbol{\theta}_\mathcal{E}$ is the $(T_\mathcal{E} \times R)$ matrix with rows equal to the $\boldsymbol{\theta}_t$'s indexed by \mathcal{E} , and $\boldsymbol{\epsilon}_j^\mathcal{E}$ is defined analogously. Pre-multiplying by $\boldsymbol{\lambda}'_t (\boldsymbol{\lambda}'_\mathcal{E} \boldsymbol{\lambda}_\mathcal{E})^{-1} \boldsymbol{\lambda}'_\mathcal{E}$ yields

$$\begin{aligned} &\boldsymbol{\lambda}'_t (\boldsymbol{\lambda}'_\mathcal{E} \boldsymbol{\lambda}_\mathcal{E})^{-1} \boldsymbol{\lambda}'_\mathcal{E} \left(\sum_{j=1}^J w_j^* \mathbf{Y}_j^\mathcal{E} - \sum_{j=1}^J v_j^* \mathbf{Y}_j^\mathcal{E} \right) \\ &= \boldsymbol{\lambda}'_t (\boldsymbol{\lambda}'_\mathcal{E} \boldsymbol{\lambda}_\mathcal{E})^{-1} \boldsymbol{\lambda}'_\mathcal{E} \boldsymbol{\theta}_\mathcal{E} \left(\sum_{j=1}^J w_j^* \mathbf{Z}_j - \sum_{j=1}^J v_j^* \mathbf{Z}_j \right) \\ &\quad + \boldsymbol{\lambda}'_t \left(\sum_{j=1}^J w_j^* \boldsymbol{\mu}_j - \sum_{j=1}^J v_j^* \boldsymbol{\mu}_j \right) \\ &\quad + \boldsymbol{\lambda}'_t (\boldsymbol{\lambda}'_\mathcal{E} \boldsymbol{\lambda}_\mathcal{E})^{-1} \boldsymbol{\lambda}'_\mathcal{E} \left(\sum_{j=1}^J w_j^* \boldsymbol{\epsilon}_j^\mathcal{E} - \sum_{j=1}^J v_j^* \boldsymbol{\epsilon}_j^\mathcal{E} \right). \end{aligned} \quad (\text{OA.9})$$

Equations (OA.8) and (OA.9) imply

$$\sum_{j=1}^J w_j^* Y_{jt}^N - \sum_{j=1}^J v_j^* Y_{jt}^N = (\boldsymbol{\theta}'_t - \boldsymbol{\lambda}'_t (\boldsymbol{\lambda}'_\mathcal{E} \boldsymbol{\lambda}_\mathcal{E})^{-1} \boldsymbol{\lambda}'_\mathcal{E} \boldsymbol{\theta}_\mathcal{E}) \left(\sum_{j=1}^J w_j^* \mathbf{Z}_j - \sum_{j=1}^J v_j^* \mathbf{Z}_j \right)$$

$$\begin{aligned}
& + \boldsymbol{\lambda}'_t(\boldsymbol{\lambda}'_\mathcal{E}\boldsymbol{\lambda}_\mathcal{E})^{-1}\boldsymbol{\lambda}'_\mathcal{E}\left(\sum_{j=1}^J w_j^* \mathbf{Y}_j^\mathcal{E} - \sum_{j=1}^J v_j^* \mathbf{Y}_j^\mathcal{E}\right) \\
& - \boldsymbol{\lambda}'_t(\boldsymbol{\lambda}'_\mathcal{E}\boldsymbol{\lambda}_\mathcal{E})^{-1}\boldsymbol{\lambda}'_\mathcal{E}\left(\sum_{j=1}^J w_j^* \boldsymbol{\epsilon}_j^\mathcal{E} - \sum_{j=1}^J v_j^* \boldsymbol{\epsilon}_j^\mathcal{E}\right) \\
& + \left(\sum_{j=1}^J w_j^* \epsilon_{jt} - \sum_{j=1}^J v_j^* \epsilon_{jt}\right). \tag{OA.10}
\end{aligned}$$

If Assumption 3 holds, (OA.10) becomes

$$\begin{aligned}
\sum_{j=1}^J w_j^* Y_{jt}^N - \sum_{j=1}^J v_j^* Y_{jt}^N & = -\boldsymbol{\lambda}'_t(\boldsymbol{\lambda}'_\mathcal{E}\boldsymbol{\lambda}_\mathcal{E})^{-1}\boldsymbol{\lambda}'_\mathcal{E}\left(\sum_{j=1}^J w_j^* \boldsymbol{\epsilon}_j^\mathcal{E} - \sum_{j=1}^J v_j^* \boldsymbol{\epsilon}_j^\mathcal{E}\right) \\
& + \left(\sum_{j=1}^J w_j^* \epsilon_{jt} - \sum_{j=1}^J v_j^* \epsilon_{jt}\right). \tag{OA.11}
\end{aligned}$$

Only the first term on the right-hand side of (OA.11) has a non-zero mean (because the weights w_j^* and v_j^* , depend on the error terms $\boldsymbol{\epsilon}_j^\mathcal{E}$). Therefore,

$$\begin{aligned}
\left|E\left[\sum_{j=1}^J w_j^* Y_{jt}^N - \sum_{j=1}^J v_j^* Y_{jt}^N\right]\right| & = \left|E\left[\boldsymbol{\lambda}'_t(\boldsymbol{\lambda}'_\mathcal{E}\boldsymbol{\lambda}_\mathcal{E})^{-1}\boldsymbol{\lambda}'_\mathcal{E}\left(\sum_{j=1}^J w_j^* \boldsymbol{\epsilon}_j^\mathcal{E} - \sum_{j=1}^J v_j^* \boldsymbol{\epsilon}_j^\mathcal{E}\right)\right]\right| \\
& \leq \left|E\left[\boldsymbol{\lambda}'_t(\boldsymbol{\lambda}'_\mathcal{E}\boldsymbol{\lambda}_\mathcal{E})^{-1}\boldsymbol{\lambda}'_\mathcal{E}\sum_{j=1}^J w_j^* \boldsymbol{\epsilon}_j^\mathcal{E}\right]\right| + \left|E\left[\boldsymbol{\lambda}'_t(\boldsymbol{\lambda}'_\mathcal{E}\boldsymbol{\lambda}_\mathcal{E})^{-1}\boldsymbol{\lambda}'_\mathcal{E}\sum_{j=1}^J v_j^* \boldsymbol{\epsilon}_j^\mathcal{E}\right]\right|.
\end{aligned}$$

For any $t \geq T_0 + 1$ and $s \in \mathcal{E}$, under Assumption 2 (i), we apply Cauchy-Schwarz inequality and the eigenvalue bound on the Rayleigh quotient to obtain

$$\left(\boldsymbol{\lambda}'_t(\boldsymbol{\lambda}'_\mathcal{E}\boldsymbol{\lambda}_\mathcal{E})^{-1}\boldsymbol{\lambda}_s\right)^2 \leq \left(\frac{\bar{\lambda}^2 F}{T_\mathcal{E}\underline{\zeta}}\right)^2.$$

Let

$$\bar{\boldsymbol{\epsilon}}_{jt}^\mathcal{E} = \boldsymbol{\lambda}'_t(\boldsymbol{\lambda}'_\mathcal{E}\boldsymbol{\lambda}_\mathcal{E})^{-1}\boldsymbol{\lambda}'_\mathcal{E}\boldsymbol{\epsilon}_j^\mathcal{E} = \sum_{s \in \mathcal{E}} \boldsymbol{\lambda}'_t(\boldsymbol{\lambda}'_\mathcal{E}\boldsymbol{\lambda}_\mathcal{E})^{-1}\boldsymbol{\lambda}_s \epsilon_{js}.$$

Because $\bar{\boldsymbol{\epsilon}}_{jt}^\mathcal{E}$ is a linear combination of independent sub-Gaussians with variance proxy $\bar{\sigma}^2$, we know $\bar{\boldsymbol{\epsilon}}_{jt}^\mathcal{E}$ is sub-Gaussian with variance proxy $(\bar{\lambda}^2 F/\underline{\zeta})^2 \bar{\sigma}^2/T_\mathcal{E}$. Let $\mathcal{S} = \{\mathbf{w} \in \mathbb{R}^J : \sum_{j=1}^J w_j = 1\}$ be the unit simplex. Theorem 1.16 from Rigollet and Hütter (2019) implies

$$\left|E\left[\sum_{j=1}^J w_j^* Y_{jt}^N - \sum_{j=1}^J v_j^* Y_{jt}^N\right]\right| \leq \left|E\left[\sum_{j=1}^J w_j^* \bar{\boldsymbol{\epsilon}}_{jt}^\mathcal{E}\right]\right| + \left|E\left[\sum_{j=1}^J v_j^* \bar{\boldsymbol{\epsilon}}_{jt}^\mathcal{E}\right]\right|$$

$$\begin{aligned}
&\leq E \left[\max_{\mathbf{w} \in \mathcal{S}} \left| \sum_{j=1}^J w_j \bar{\epsilon}_{jt}^{\mathcal{E}} \right| \right] + E \left[\max_{\mathbf{v} \in \mathcal{S}} \left| \sum_{j=1}^J v_j \bar{\epsilon}_{jt}^{\mathcal{E}} \right| \right] \\
&\leq \frac{\bar{\lambda}^2 F}{\underline{\zeta}} 2\sqrt{2 \log(2J)} \frac{\bar{\sigma}}{\sqrt{T_{\mathcal{E}}}},
\end{aligned}$$

which finishes the proof of the theorem.

Suppose now Assumption 6 holds (but Assumption 5 does not). To obtain a bound on the bias we need to bound the first two terms in (OA.10). Recall that

$$\boldsymbol{\lambda}'_t (\boldsymbol{\lambda}'_{\mathcal{E}} \boldsymbol{\lambda}_{\mathcal{E}})^{-1} \boldsymbol{\lambda}_s \leq \frac{\bar{\lambda}^2 F}{T_{\mathcal{E}} \underline{\zeta}}.$$

Therefore, the absolute value of each element in vector $(\boldsymbol{\theta}'_t - \boldsymbol{\lambda}'_t (\boldsymbol{\lambda}'_{\mathcal{E}} \boldsymbol{\lambda}_{\mathcal{E}})^{-1} \boldsymbol{\lambda}'_{\mathcal{E}} \boldsymbol{\theta}_{\mathcal{E}})$ is bounded by $\bar{\theta} \left(1 + \frac{\bar{\lambda}^2 F}{\underline{\zeta}}\right)$. Cauchy–Schwarz inequality and Assumption 6 imply

$$\begin{aligned}
&\left| (\boldsymbol{\theta}'_t - \boldsymbol{\lambda}'_t (\boldsymbol{\lambda}'_{\mathcal{E}} \boldsymbol{\lambda}_{\mathcal{E}})^{-1} \boldsymbol{\lambda}'_{\mathcal{E}} \boldsymbol{\theta}_{\mathcal{E}}) \left(\sum_{j=1}^J w_j^* \mathbf{Z}_j - \sum_{j=1}^J v_j^* \mathbf{Z}_j \right) \right| \\
&\leq \bar{\theta} \left(1 + \frac{\bar{\lambda}^2 F}{\underline{\zeta}}\right) \sqrt{R} \left\| \sum_{j=1}^J w_j^* \mathbf{Z}_j - \sum_{j=1}^J v_j^* \mathbf{Z}_j \right\|_2 \\
&\leq \bar{\theta} \left(1 + \frac{\bar{\lambda}^2 F}{\underline{\zeta}}\right) R d,
\end{aligned}$$

and

$$\left| \boldsymbol{\lambda}'_t (\boldsymbol{\lambda}'_{\mathcal{E}} \boldsymbol{\lambda}_{\mathcal{E}})^{-1} \boldsymbol{\lambda}'_{\mathcal{E}} \left(\sum_{j=1}^J w_j^* \mathbf{Y}_j^{\mathcal{E}} - \sum_{j=1}^J v_j^* \mathbf{Y}_j^{\mathcal{E}} \right) \right| \leq \frac{\bar{\lambda}^2 F}{\underline{\zeta}} d.$$

Combining the last two displayed equations with (OA.10), we have

$$\left| E \left[\sum_{j=1}^J w_j^* Y_{jt}^I - \sum_{j=1}^J f_j Y_{jt}^I \right] \right| \leq \left(\bar{\theta} R + \frac{\bar{\lambda}^2 F}{\underline{\zeta}} (1 + \bar{\theta} R) \right) d + \frac{\bar{\lambda}^2 F}{\underline{\zeta}} 2\sqrt{2 \log(2J)} \frac{\bar{\sigma}}{\sqrt{T_{\mathcal{E}}}},$$

which finishes the proof of the theorem. ■

Proof of Theorem OA.6. Recall that

$$\hat{u}_t = \sum_{j=1}^J w_j^* Y_{jt} - \sum_{j=1}^J v_j^* Y_{jt},$$

for $t \in \mathcal{B} \cup \{T_0 + 1, \dots, T\}$. For $t \in \{T_0 + 1, \dots, T\}$, \hat{u}_t are the post-intervention estimates of the treatment effects; and for $t \in \mathcal{B}$, \hat{u}_t are the placebo treatment effects estimated for the blank periods. Let

$$u_t = \sum_{j=1}^J w_j^* \epsilon_{jt} - \sum_{j=1}^J v_j^* \epsilon_{jt}$$

for $t \in \mathcal{B}$, and

$$u_t = \sum_{j=1}^J w_j^* \xi_{jt} - \sum_{j=1}^J v_j^* \epsilon_{jt}$$

for $t \in \{T_0 + 1, \dots, T\}$. The null hypothesis (15) and the assumptions of Theorem OA.6 imply that $\{u_t\}_{t \in \mathcal{B} \cup \{T_0 + 1, \dots, T\}}$ is a sequence of exchangeable random variables. Additionally, Assumption 1 and the null hypothesis (15) imply

$$\hat{u}_t = \sum_{j=1}^J w_j^* Y_{jt} - \sum_{j=1}^J v_j^* Y_{jt} = \boldsymbol{\theta}'_t \sum_{j=1}^J (w_j^* - v_j^*) \mathbf{Z}_j + \boldsymbol{\lambda}'_t \sum_{j=1}^J (w_j^* - v_j^*) \boldsymbol{\mu}_j + u_t,$$

for $t \in \mathcal{B} \cup \{T_0 + 1, \dots, T\}$. The result of the theorem follows now from Theorem D.1 in Chernozhukov, Wüthrich and Zhu (2021). ■

OA.4. Implementations of the Optimization Formulations

To computationally solve (4), i.e., the *Unconstrained* design, we propose two methods. The first method is by enumeration, which takes advantage of the objective function of (4) being separated between \mathbf{w} and \mathbf{v} . If we knew which units were to receive treatment and which units were to receive control, then we could decompose (4) into two classical synthetic control problems and solve both of them efficiently. We brute force enumerate all the possible combinations of the treatment units and control units. Because the two groups of treated and control units can be swapped, we only enumerate combinations such that the cardinality of the treated group is smaller than or equal to the cardinality of the control group. When the cardinality of the treated group is equal to the cardinality of the control group, we prioritize the treated group to be the one with the smallest index among the units with positive weights.

The second method solves a constrained optimization problem, by converting it into the canonical form of a Quadratic Constraint Quadratic Program (QCQP), which we detail below. The decision variables are w_j and $v_j, \forall j = 1, \dots, J$. For simplicity, we write it in a vector form $\tilde{\mathbf{W}} = (w_1, w_2, \dots, w_J, v_1, v_2, \dots, v_J)$.

Let M be the dimension of the predictors \mathbf{X}_j . Let X be an $M \times J$ matrix, each column of which is \mathbf{X}_j , which stands for the predictors of unit j .

Define $P^0 = \{P_{k,l}^0\}_{k,l=1,\dots,2J} \in \mathbb{R}^{2J \times 2J}$, such that P^0 has only two diagonal blocks, while the two off-diagonal blocks are zero. Define for any $k, l = 1, \dots, 2J$,

$$P_{k,l}^0 = \begin{cases} \sum_{i=1}^M X_{i,k} X_{i,l}, & k, l = 1, \dots, J; \\ \sum_{i=1}^M X_{i,(k-J)} X_{i,(l-J)}, & k, l = J+1, \dots, 2J; \\ 0, & \text{otherwise.} \end{cases}$$

Define $\mathbf{q}^0 \in \mathbb{R}^{2J}$, such that for any $k = 1, \dots, 2J$

$$\mathbf{q}_k^0 = \begin{cases} -2 \sum_{i=1}^M X_{i,k} \cdot \left(\sum_{j=1}^J f_j X_{i,j} \right), & k = 1, \dots, J; \\ -2 \sum_{i=1}^M X_{i,k-J} \cdot \left(\sum_{j=1}^J f_j X_{i,j} \right), & k = J+1, \dots, 2J. \end{cases}$$

Further define $\mathbf{e}_1 = (1, 1, \dots, 1, 0, 0, \dots, 0)'$ whose first J elements are 1 and last J elements 0; and $\mathbf{e}_2 = (0, 0, \dots, 0, 1, 1, \dots, 1)'$ whose first J elements are 0 and last J elements 1.

Finally, define $P^1 = \{P_{k,l}^1\}_{k,l=1,\dots,2J} \in \mathbb{R}^{2J \times 2J}$ such that P^1 only has non-zero values in the two off-diagonal blocks, i.e., for any $k, l = 1, \dots, 2J$,

$$P_{k,l}^1 = \begin{cases} 1, & k = l + J; \\ 1, & k = l - J; \\ 0, & \text{otherwise.} \end{cases}$$

Using the above notations we re-write the (non-convex) QCQP as follows,

$$\min \quad \tilde{\mathbf{W}}' P^0 \tilde{\mathbf{W}} + \mathbf{q}^0' \tilde{\mathbf{W}} \tag{OA.12}$$

$$\begin{aligned}
\text{s.t. } \mathbf{e}'_1 \tilde{\mathbf{W}} &= 1, \\
\mathbf{e}'_2 \tilde{\mathbf{W}} &= 1, \\
\tilde{\mathbf{W}}' P^1 \tilde{\mathbf{W}} &= 0, \\
\tilde{\mathbf{W}} &\geq \mathbf{0}.
\end{aligned}$$

The first computational method (enumeration) solves two synthetic control problems in each iteration. The synthetic control problems can be efficiently solved. We implement the synthetic control problem using the “lsei” function from “limSolve” package in R 4.0.2. For the second computational method (quadratic programming), the problem (OA.12) is implemented using Gurobi 9.0.2 in R 4.0.2. Since the QCQP is non-convex, the computation leads to some numerical errors up to 0.001 in finding the treated and control weights. So we round the treated and control weights to the nearest 2-digits in the implementation of the QCQP. Moreover, for all the weights that are less than or equal to 0.01, we trim the weights to zero. This is because smaller weights suffer from greater impacts of numerical errors, and that numerical errors could cause zero weights to be non-zero, thus having a non-negligible impact on the swapping rule.

To conclude, we compare the treated and control weights calculated from both methods. Both methods yield the same treated and control weights up to some negligible rounding error, while the first method takes longer computational time.

The other designs are computationally implemented using either one of the above two methods. The *Constrained* design is implemented using the enumeration method. In cases when the cardinality constraint \bar{m} is small, this brute force enumeration is very efficient. The *Weakly-targeted* design is implemented using the quadratic programming method. In the QCQP formulation, the objective function has both a different quadratic term P^0 and a different linear term \mathbf{q}^0 . The *Unit-Level* design is implemented using the enumeration method. The *Penalized* design is implemented using the quadratic programming method. In the QCQP formulation, the objective function has the same quadratic term P^0 and a different linear term \mathbf{q}^0 .

OA.5. Additional Illustrations Using Walmart Data

In this section, we present results for $\bar{m} = 1$ and $\bar{m} = 3$. Using only one treated unit ($\bar{m} = 1$) fails to produce a good fit between the treated and synthetic control unit in the fitting periods. For the case of $\bar{m} = 1$, Figures 6 and 7 reveal a substantial gap with a clear seasonal trend between the two synthetic units. Figures 8 and 9 report results for $\bar{m} = 3$. Increasing \bar{m} from $\bar{m} = 2$ to $\bar{m} = 3$ results in a minor improvement in fit, and leaves estimation results substantively unchanged.




Figure 6: Synthetic Treated Unit and Synthetic Control Unit, $\bar{m} = 1$

Note: The black solid line represents the synthetic treated outcome. The black dashed line represents the synthetic control outcome. The blue dashed lines are individual stores' sales.




Figure 7: Placebo Treatment Effects, $\bar{m} = 1$

Note: This figure reports the difference between the synthetic treated and synthetic control outcomes of Figure 6.




Figure 8: Synthetic Treatment Unit and Synthetic Control Unit, when $\bar{m} = 3$.

Note: The black solid line represents the synthetic treated outcome. The black dashed line represents the synthetic control outcome. The blue dashed lines are individual stores' sales.




Figure 9: Treatment Effect Estimate, when $\bar{m} = 3$.

Note: This figure reports the difference between the synthetic treated and synthetic control outcomes of Figure 8. For the experimental periods, this is the treatment effect estimate.

OA.6. Additional Simulation Results

In Section 5.1 the idiosyncratic shocks are i.i.d. Normal with variance $\sigma^2 = 1$. Figures 10 and 11 report results for $\sigma^2 = 5$ and $\sigma^2 = 10$, respectively. Figures 12 and 13 report differences between the outcomes for the synthetic treated and the synthetic control units for the same values for σ^2 . As the value of σ^2 increases, the quality of the post-treatment estimation and inference deteriorates, and the p -value for the null hypotheses of in (15) increases. The deterioration in pre-treatment fit in Figures 10 and 11 provides a diagnosis of the accuracy of the respective estimates.




Figure 10: Synthetic Treatment Unit and Synthetic Control Unit, when $\sigma^2 = 5$.




Figure 11: Synthetic Treatment Unit and Synthetic Control Unit, when $\sigma^2 = 10$.




Figure 12: Treatment Effect Estimate, when $\sigma^2 = 5$.




Figure 13: Treatment Effect Estimate, when $\sigma^2 = 10$.