

Uncertainty Quantification in Synthetic Controls with Staggered Treatment Adoption*

Matias D. Cattaneo[†] Yingjie Feng[‡] Filippo Palomba[§] Rocio Titiunik[¶]

October 14, 2024

Abstract

We propose principled prediction intervals to quantify the uncertainty of a large class of synthetic control predictions (or estimators) in settings with staggered treatment adoption, offering precise non-asymptotic coverage probability guarantees. From a methodological perspective, we provide a detailed discussion of different causal quantities to be predicted, which we call *causal predictands*, allowing for multiple treated units with treatment adoption at possibly different points in time. From a theoretical perspective, our uncertainty quantification methods improve on prior literature by (i) covering a large class of causal predictands in staggered adoption settings, (ii) allowing for synthetic control methods with possibly nonlinear constraints, (iii) proposing scalable robust conic optimization methods and principled data-driven tuning parameter selection, and (iv) offering valid uniform inference across post-treatment periods. We illustrate our methodology with an empirical application studying the effects of economic liberalization on real GDP per capita for Sub-Saharan African countries. Companion general-purpose software packages are provided in `Python`, `R`, and `Stata`.

Keywords: causal inference, synthetic controls, staggered treatment adoption, prediction intervals, non-asymptotic inference.

*We thank the co-Editor, Peter Hull, and three reviewers for their comments. We also thank Alberto Abadie, Simon Freyaldenhoven, and Bartolomeo Stellato for many insightful discussions. Cattaneo and Titiunik gratefully acknowledge financial support from the National Science Foundation (SES-2019432 and SES-2241575), Cattaneo gratefully acknowledges financial support from the National Institute of Health (R01 GM072611-16), and Feng gratefully acknowledges the financial support from the National Natural Science Foundation of China (NSFC) under grants 72203122, 72133002, and 72250064.

[†]Department of Operations Research and Financial Engineering, Princeton University.

[‡]School of Economics and Management, Tsinghua University.

[§]Department of Economics, Princeton University.

[¶]Department of Politics, Princeton University.

Contents

1	Introduction	1
1.1	Related Literature	3
1.2	Paper Organization	4
2	The Effect of Liberalization on GDP for Sub-Saharan African Countries	5
3	Basic Framework	7
4	Uncertainty Quantification	12
4.1	In-Sample Uncertainty Quantification	17
4.2	Out-of-Sample Uncertainty Quantification	21
4.3	Theoretical Justification	24
4.4	Simultaneous Prediction Intervals	25
5	Empirical Results	27
6	Theoretical Foundations	31
6.1	Discussion of Conditions (i)–(iv)	38
6.2	Defining Constraints in Simulation	40
6.3	Adjustment for Nonlinear Constraints	41
7	Conclusion	42

1 Introduction

The synthetic control (SC) method was introduced by [Abadie and Gardeazabal \(2003\)](#), and since then many extensions and generalizations have been proposed (see [Abadie, 2021](#), and references therein). The method is now part of the standard toolkit for program evaluation and treatment effect analysis ([Abadie and Cattaneo, 2018](#)), offering a complement to traditional difference-in-differences, event studies, and other panel data approaches for causal inference with longitudinal aggregate data and few treated units. Most of the synthetic control literature concentrates on identification, as well as on prediction or point estimation of treatment effects, under different causal inference frameworks and algorithmic implementations. In contrast, principled uncertainty quantification of synthetic control predictions or estimators in general settings remains mostly unexplored, despite its importance for empirical work.

Following [Cattaneo, Feng and Titiunik \(2021\)](#), who proposed non-asymptotic prediction intervals for synthetic controls with a single-treated-unit, we employ a causal inference framework where potential outcomes are assumed to be random and develop novel prediction intervals to quantify the uncertainty of a large class of synthetic control predictions (or estimators) in settings with staggered treatment adoption and multiple treated units. Our contributions include establishing precise non-asymptotic coverage probability guarantees for our novel prediction intervals, introducing scalable robust optimization implementations for possibly nonlinear constraints in the synthetic controls construction, developing principled tuning parameter selection, and proposing valid joint inference methods across time. Inferential procedures with non-asymptotic probability guarantees are valuable because synthetic control applications often have small sample sizes, impeding the applicability of asymptotic approximations. Conceptually, the proposed prediction intervals capture two sources of uncertainty: one coming from the construction of the synthetic control weights with pre-treatment data, and the other generated by the irreducible sampling variability introduced by the post-treatment outcomes. Our prediction intervals also take into account potential misspecification errors explicitly and enjoy other robustness properties due to their non-asymptotic, generic construction. For example, our methods allow for nonlinear regularization in the synthetic controls construction, which accommodates L2 penalization, L1/L2-combined penalization, and other variants thereof. These nonlinear penalization schemes are better suited for application where the

covariates exhibit collinearity/codependency, as opposed to L1 penalization schemes, which are better suited for covariate shrinkage/selection.

To motivate our methodological work, we begin by introducing an empirical application in Section 2, and we use this application throughout to illustrate our contributions. Following [Billmeier and Nannicini \(2013\)](#), we investigate the effect of economic liberalization in the 1990s on real GDP per capita for Sub-Saharan African countries. This application includes multiple adoption times and multiple treated units. Furthermore, we consider the special case of outcome-only synthetic controls to improve the exposition and carefully account for the fact that the outcome variable in this application is non-stationary and co-integrated. We also discuss more complex empirical analyses in the supplemental appendix.

Our first contribution is methodological in nature due to the complexity added by the staggered treatment adoption setup, which allows for (but does not require) the existence of multiple treated units changing from control to treatment status at possibly different points in time. In Section 3, we introduce a basic causal inference framework that is motivated by our empirical application and specifically tailored to synthetic control methods with staggered treatment adoption. Using this framework, we define different causal quantities to be predicted in the context of synthetic controls, which we refer to as *causal predictands*, and explain how prediction methods are implemented.

Section 4 then discusses uncertainty quantification in the context of our empirical application and basic causal inference framework. Motivated by our empirical application and recent advances in the synthetic controls literature (see [Abadie, 2021](#), and references therein), our analysis focuses exclusively on incorporating in-sample and out-of-sample uncertainty quantification for outcome-only synthetic control methods with non-stationary data and non-linear constraints. In this setting, we present novel prediction intervals with precise non-asymptotic guarantees for synthetic controls with staggered adoption, under easy-to-interpret sufficient conditions. We also discuss scalable, robust conic optimization implementations of our methods ([Boyd and Vandenberghe, 2004](#)), data-driven selection of tuning parameters, and valid joint inference across time.

The main empirical results are presented in Section 5. Our findings indicate that the economic liberalization in the 1990s did not have a positive economic impact on emerging Sub-Saharan African countries. This finding is in line with prior empirical results ([Billmeier and Nannicini, 2013](#)). The Supplemental Appendix provides additional empirical evidence supporting our main

findings, including a re-analysis using alternative synthetic control specifications, placebo treatment dates, donor pool constructions, and other related implementations under different assumptions.

While the basic framework developed in Sections 3 and 4 is sufficient for our empirical application, Section 6 offers a substantive generalization of our methods that allows for general misspecification error, multiple covariate features, and cross-equation re-weighting, among other features, when constructing the synthetic control weights with pre-treatment data. This general framework also provides foundational results for our empirical work, including a theoretical justification for the methods presented in Sections 3 and 4.

1.1 Related Literature

We contribute to developing prediction/estimation and inference methods for synthetic control settings with staggered treatment adoption. Putting aside generic linear factor model or matrix completion methods, [Ben-Michael, Feller and Rothstein \(2022\)](#), [Powell \(2022\)](#), and [Shaikh and Toulis \(2021\)](#) appear to be the only prior papers that have studied staggered treatment adoption for synthetic controls explicitly. The first paper focuses on prediction/estimation in settings where the pre-treatment fit is poor and develops penalization methods to improve the performance of the canonical synthetic control method. [Ben-Michael, Feller and Rothstein \(2022\)](#) also suggest employing a bootstrap method for assessing uncertainty, but no formalization is provided guaranteeing its (asymptotic) validity. [Powell \(2022\)](#) employs a standard parametric panel data model to discuss estimation and inference methods for a single common treatment effect that is valid in large samples. [Shaikh and Toulis \(2021\)](#) focus on uncertainty quantification employing a parametric duration model and propose a permutation-based inferential method under a symmetry assumption. Our paper complements these prior contributions by (i) developing a general causal inference framework for synthetic control methods with staggered treatment adoption, and (ii) offering nonparametric inference methods with demonstrable non-asymptotic coverage guarantees and allowing for misspecification in the construction of the synthetic control weights. We also propose novel scalable robust conic optimization implementations, principled tuning parameter selection methods, and a valid joint inference procedure across multiple time periods.

Our quantification of uncertainty via non-asymptotic prediction intervals follows [Cattaneo, Feng and Titiunik \(2021\)](#). These inference methods are motivated by [Vovk \(2012\)](#) and are most closely

related to prior work by [Chernozhukov, Wüthrich and Zhu \(2021a\)](#); [Chernozhukov, Wüthrich and Zhu \(2021b\)](#) on conformal prediction intervals. Relative to this literature, our contributions include: (i) allowing for a large class of causal predictands in staggered adoption settings (prior work covered only the canonical single treated unit case); (ii) covering a large class of synthetic control predictions with possibly nonlinear constraints (prior work allowed only linear constraints); (iii) developing scalable robust optimization implementations and proposing principled data-driven tuning parameter selection (prior work did not provide guidance on these issues); and (iv) introducing valid uniform inference methods across post-treatment periods (absent in prior work).

There are a few other, conceptually different, recent proposals to quantify uncertainty and conduct inference in the synthetic controls literature. For example, [Li \(2020\)](#) study correctly specified linear factor models, [Masini and Medeiros \(2021\)](#) study high-dimensional penalization methods, [Agarwal, Shah, Shen and Song \(2021\)](#) investigate matrix completion methods, and [Shi, Miao, Hu and Tchetgen \(2023\)](#) develop inference methods using a proximal causal inference framework. All these methods rely on asymptotic approximations, in most cases employing standard Gaussian critical values that assume away misspecification errors and other small sample issues. Our work complements these contributions by providing prediction intervals with non-asymptotic coverage guarantees ([Wainwright, 2019](#)). There is also a literature in econometrics on event studies that is loosely related to synthetic controls with staggered adoption: see, for example, [Freyaldenhoven et al. \(2019, Forthcoming\)](#), [Miller \(2023\)](#), and references therein. Finally, all the inferential methods mentioned so far contrast with the original method proposed by [Abadie, Diamond and Hainmueller \(2010\)](#), which relies on design-based permutation of treatment assignment assuming that the potential outcomes are non-random.

1.2 Paper Organization

Section 2 introduces our running empirical application. Section 3 presents the basic causal inference framework for outcome-only synthetic controls with staggered treatment adoption, and Section 4 discusses non-asymptotic uncertainty quantification in that context. Section 5 presents our empirical results. Section 6 gives a self-contained discussion of our most general framework and theoretical results. The Supplemental Appendix includes all proofs, additional empirical results, and other technical details omitted to improve the exposition.

We provide general-purpose software implementing our results in `Python`, `R`, and `Stata`, including detailed documentation and additional replication materials. This software is discussed in our companion article [Cattaneo, Feng, Palomba and Titiunik \(2024\)](#), where we addressed several implementation issues related to numerical optimization and tuning parameter selection. To complement the illustration, in Section S.6 of the Supplemental Appendix, we provide more details on how to prepare the data to analyze staggered treatment adoption using synthetic control methods using our companion software. In addition, in the Supplemental Appendix, Section S.2.3 demonstrates how to obtain tighter bounds for uncertainty quantification when the weighting matrix has a block diagonal structure, and Section S.5 shows how to reformulate the synthetic control problem as a scalable conic optimization problem to improve computational efficiency. For ease of understanding, Table S.1 in the Supplemental Appendix summarizes the notation used throughout the paper.

2 The Effect of Liberalization on GDP for Sub-Saharan African Countries

During the second half of the twentieth century, many countries around the world launched programs of (external) economic liberalization, booming from 22% in 1960 to 73% in the early 2000s ([Wacziarg and Welch, 2008](#)). In the last thirty years, scholars have investigated the social and economic consequences of such liberalization programs, often reaching conflicting conclusions (see, e.g., [Levine and Renelt, 1992](#); [Sachs, Warner, Åslund and Fischer, 1995](#); [DeJong and Ripoll, 2006](#)).

The impact of liberalization policies on economic welfare has been traditionally investigated with cross-country analyses (e.g. [Sachs, Warner, Åslund and Fischer, 1995](#)) and individual case studies (e.g. [Bhagwati and Srinivasan, 2001](#)). More recently, scholars have turned to synthetic control methods in the hope of employing a causal inference methodology that allows for the presence of time-varying unobservable confounders. Employing the synthetic control framework originally developed in [Abadie and Gardeazabal \(2003\)](#), [Billmeier and Nannicini \(2013\)](#) analyzed the effects of liberalization in four continents: Africa, Asia, North America, and South America. They used a pre-existing dataset of economic variables (previously used in [Giavazzi and Tabellini, 2005](#)) which includes 180 countries, covers the period 1963–2000 and contains an indicator for economic liberalization originally defined in [Sachs, Warner, Åslund and Fischer \(1995\)](#) and updated

in [Wacziarg and Welch \(2008\)](#) (hereafter, the Sachs-Warner indicator). More details on the data and the definition of economic liberalization can be found in Section S.6 of the Supplemental Appendix.

[Billmeier and Nannicini \(2013\)](#) studied the effect of economic liberalization—as measured by the Sachs-Warner indicator—on GDP per capita, to understand whether the adoption of liberalization programs affected economic growth. Building on their study of Sub-Saharan Africa, we illustrate how our formal inference framework can be an essential tool to aggregate results and draw general conclusions from synthetic control analyses under staggered treatment adoption by multiple treated units. The ambitious scope of the analysis in [Billmeier and Nannicini \(2013\)](#) resulted in a large number of synthetic control results. In Sub-Saharan Africa alone, they studied sixteen episodes of liberalization that occurred at ten different periods. The results exhibited considerable heterogeneity, which the authors summarized by grouping the effects into four categories according to two dimensions: whether the effect was positive and whether there was evidence that the effect was robust rather than “coincidental” (a notion of statistical significance). The four categories were (a) countries with a positive and strongly robust effect, (b) countries with a positive and somewhat robust effect, (c) countries with a positive but non-robust effect, and (d) countries with a null effect. They concluded that “Botswana is the only country clearly in the first group, offering a truly convincing success story” (p. 995).

[Billmeier and Nannicini \(2013\)](#) adjudicated the robustness or statistical significance of each of the sixteen effects with the Fisherian-type placebo test developed by [Abadie et al. \(2010\)](#), which was the main inference tool available at the time. Our new inference framework allows us to reconsider the evidence for Sub-Saharan Africa by providing formal tools to perform joint inference and draw a general conclusion about the effects of liberalization by considering the evidence altogether rather than piecemeal.

The data starts in the 1960s, which marks the beginning of Africa’s post-colonial era. Since independence from their colonial rulers until the late 1980s, many countries in Sub-Saharan Africa adopted neo-patrimonial political regimes based on the concentration of power in a single individual and the resulting cultivation of personalistic politics, widespread clientelistic networks, and the use of the resources of the state to achieve political legitimacy ([Bratton and Van de Walle, 1997](#)).

The 1980s put high stress on these regimes, as economic conditions deteriorated. Negative economic growth, increased debt services as a proportion of exports, currency overvaluation, low

commodity prices, low investment, and limited access to credit created an adverse economic environment that these non-democratic regimes were ill-equipped to handle. International finance organizations such as the World Bank and the International Monetary Fund incentivized programs of structural economic adjustment, with disputed success (Callaghy and Ravenhill, 1994).

The adverse economic conditions led to popular protests in the early 1990s, which marked the beginning of a wave of transitions to democracy throughout the continent. Starting in the early 1990s, many countries started a process of political liberalization, partly influenced by the fall of the Berlin Wall in November of 1989 and the subsequent collapse of the communist regimes of the Soviet Union (Bratton and Van de Walle, 1997).


As illustrated in Figure 1, most countries in Sub-Saharan Africa adopted the economic liberalization treatment between the late 1980s and the early or mid-1990s which, as just explained, was a period of major political and economic transition in Africa and beyond. This means that, for many countries in our sample, the economic liberalization treatment occurred nearly simultaneously with other major political and economic shocks that influenced the World’s geo-political equilibrium. Given these potential confounders, it is crucial to exploit the staggered adoption of the treatment to make inferences. While the worldwide political changes between the late 1980s and early 1990s affected all countries at the same time, the concrete measures of economic liberalization captured by the Sachs index were introduced in different years for different countries. Our proposed synthetic control methods allow us to define causal predictands and perform joint inferences that leverage the staggered introduction of the treatment, while controlling for worldwide shocks affecting all units.

3 Basic Framework

We set up a basic synthetic control framework that matches the empirical application in Section 2 and suffices to communicate key ideas of our proposed method. A more general and technically more involved framework is deferred to Section 6.

Suppose that we observe N countries (“units”) for T years (“periods”). Countries are indexed by $i = 1, \dots, N$, and years are indexed by $t = 1, \dots, T$. For each country i in year t , we can observe the GDP per capita Y_{it} (“outcome”) and a time T_i that indicates when country i adopted economic

Figure 1: *Staggered Treatment Adoption in Africa.*



Notes: A country i in year t is defined to be “closed” or “open” using to the Sachs index. Thick vertical lines—corresponding to 1987 and 1991—delimit the three main liberalization waves of [Billmeier and Nannicini \(2013\)](#).

liberalization (“treatment”). Assume that a country remains economically liberalized whenever $t \geq T_i$. (In our empirical application this is always verified.) Without loss of generality, countries are ordered in the adoption times so that $1 \leq T_1 \leq T_2 \leq \dots \leq T_N \leq \infty$, with $T_i = \infty$ denoting that country i remains untreated throughout the observation period. Let $\mathcal{N} = \{i : T_i = \infty\}$ be the group of “never-liberalized” countries and $\mathcal{E} = \{i : T_i < \infty\}$ the group of “ever-liberalized” countries. Define $J_0 = |\mathcal{N}|$ and $J_1 = |\mathcal{E}|$, where we use $|\mathcal{A}|$ to denote the number of elements in \mathcal{A} for any finite set \mathcal{A} .

We are interested in the effect of economic liberalization on a country’s GDP per capita. Let $Y_{it}(s)$ denote the potential GDP per capita of country i in year t that would be observed had country i adopted the economic liberalization treatment in year s , for $s = 1, \dots, T, \infty$, and assume $Y_{it}(s) = Y_{it}(\infty)$ whenever $t < s$. Implicitly, these simplifications impose two standard assumptions: no spillovers (the potential outcomes of country i depend only on i ’s adoption time) and no anticipation (a country’s potential outcomes prior to the treatment are equal to the outcomes it would have had if it had never been treated). The observed GDP per capita Y_{it} can be written as

$$Y_{it} = Y_{it}(\infty)\mathbf{1}(t < T_i) + Y_{it}(T_i)\mathbf{1}(t \geq T_i).$$

That is, whenever the economic liberalization has not been adopted, we always observe the potential GDP per capita under the “never-treated” status; otherwise the potential GDP per capita corresponding to the adoption time T_i is observed.

Our primary goal is to make a prediction of the missing counterfactual GDP per capita $Y_{it}(\infty)$ for an ever-liberalized country in the post-treatment period; we then use this unit-level predictand as the basis of a variety of other aggregate predictands.

For simplicity, we take all never-treated units in \mathcal{N} as “donors” for SC prediction, and let all years prior to the economic liberalization, $\{t : 1 \leq t \leq T_i - 1\}$, be the training period used to obtain the SC weights for each ever-liberalized country $i \in \mathcal{E}$. The idea of SC is to find a vector of weights on the donor countries such that the weighted average of pre-treatment GDP per capita of donor countries matches that of the treated country as closely as possible, subject to some regularization constraints. Let $\|\cdot\|_1$ and $\|\cdot\|_2$ denote the usual L1 and L2 vector norms respectively, and $(v_j : j \in \mathcal{A})$ denote a vector consisting of all v_j ’s with $j \in \mathcal{A}$ for a set \mathcal{A} . In our empirical analysis below, the SC weights are obtained via L1-L2 constrained least squares:

$$\widehat{\mathbf{w}}^{[i]} = \arg \min_{\mathbf{w} \in \mathcal{W}^{[i]}} \sum_{t=1}^{T_i-1} \left(Y_{it} - \mathbf{Y}'_{\mathcal{N}t} \mathbf{w} \right)^2 \quad \text{for each } i \in \mathcal{E}, \quad (3.1)$$

where $\mathbf{Y}_{\mathcal{N}t} = (Y_{jt} : j \in \mathcal{N})$ is the vector of GDP per capita for donor countries at time t , and $\mathcal{W}^{[i]} = \{\mathbf{w} \in \mathbb{R}^{J_0} : \|\mathbf{w}\|_1 \leq Q_1^{[i]}, \|\mathbf{w}\|_2 \leq Q_2^{[i]}\}$ for some constants $Q_1^{[i]} > 0$ and $Q_2^{[i]} > 0$. The choices of $Q_1^{[i]}$ and $Q_2^{[i]}$ are discussed in Section S.6.2.2 of the Supplemental Appendix. The prediction of

country i 's counterfactual GDP per capita k periods after the economic liberalization is given by

$$\widehat{Y}_{i(T_i+k)}(\infty) = \mathbf{Y}'_{\mathcal{N}(T_i+k)} \widehat{\mathbf{w}}^{[i]}, \quad \text{for each } i \in \mathcal{E} \text{ and } 0 \leq k \leq T - T_i.$$

In what follows, we define a variety of causal predictands and construct their SC predictions accordingly.

Example TSUS (Time-specific unit-specific predictand). The first predictand we consider is the effect of the intervention for a specific unit in a specific time period, which is the primary causal predictand of interest in classical SC analysis with only one treated unit.

For each ever-treated unit $i \in \mathcal{E}$, we can define the TSUS predictand in each post-treatment period:

$$\tau_{ik} := Y_{i(T_i+k)}(T_i) - Y_{i(T_i+k)}(\infty), \quad 0 \leq k \leq T - T_i.$$

In our empirical application, this predictand captures the effect of economic liberalization for a specific country, measured k years after the adoption of the liberalization policy. Given the prediction $\widehat{Y}_{i(T_i+k)}(\infty)$, we predict τ_{ik} by

$$\widehat{\tau}_{ik} := Y_{i(T_i+k)}(T_i) - \widehat{Y}_{i(T_i+k)}(\infty).$$

See Figure 2(a) for a graphical representation of $\widehat{\tau}_{ik}$. ┘

Example TAUS (Time-averaged unit-specific predictand). When there are multiple post-treatment periods, scholars may be interested in the effect on a treated unit averaged across all periods. This defines the TAUS predictand, which takes the average over time of the TSUS predictands for an ever-treated unit i :

$$\tau_i := \frac{1}{T - T_i + 1} \sum_{k=0}^{T-T_i} \tau_{ik}.$$

In our empirical application, this is the effect of economic liberalization for a specific country, averaged over the entire post-liberalization period. Given the TSUS prediction for country i in

period k , we predict τ_i by

$$\widehat{\tau}_i := \frac{1}{T - T_i + 1} \sum_{k=0}^{T-T_i} \widehat{\tau}_{ik}.$$

See Figure 2(b) for a graphical representation of $\widehat{\tau}_i$. ┘

Example TSUA (Time-specific unit-averaged predictand). When there are multiple treated units, it is also of interest to define a predictand that captures the average effect of the intervention across a given group of units, at a single period in time. Let $\mathcal{Q} \subseteq \{1, 2, \dots, T\}$ be a set of adoption times. The TSUA predictand is defined as

$$\tau_{\mathcal{Q}k} := \frac{1}{Q} \sum_{i: T_i \in \mathcal{Q}} \tau_{ik}, \quad Q = |\{i : T_i \in \mathcal{Q}\}|.$$

In our empirical application, this is the average effect of economic liberalization for the group of countries with adoption times in \mathcal{Q} , measured k years after liberalization. In later empirical analysis we consider the treatment effect for countries that liberalized in three different waves—before 1987, between 1987 and 1991, and after 1991. Given the TSUS predictions, we predict $\tau_{\mathcal{Q}k}$ by

$$\widehat{\tau}_{\mathcal{Q}k} := \frac{1}{Q} \sum_{i: T_i \in \mathcal{Q}} \widehat{\tau}_{ik}.$$

See Figure 2(c) for a graphical representation of $\widehat{\tau}_{\mathcal{Q}k}$. ┘

Example TAUA (Time-averaged unit-averaged predictand). Finally, when there are multiple treated units and multiple post-treatment periods, we may be interested in a predictand that captures the overall average effect of the intervention.

We define the TAUA predictand as the average of the TSUS predictands across all treated units and over L post-treatment periods:

$$\tau_{..} := \frac{1}{LJ_1} \sum_{k=1}^L \sum_{i \in \mathcal{E}} \tau_{ik}.$$

In our empirical application, this is the average effect of economic liberalization for all ever-liberalized countries over L years after the policy adoption. We assume all ever-treated units are observed at least L periods after the adoption for some $L \geq 0$. However, since the observation

ends at time T , some ever-treated units may have to be excluded in this definition as L varies. Given the TSUS predictions, we predict $\tau_{..}$ by

$$\hat{\tau}_{..} := \frac{1}{LJ_1} \sum_{k=1}^L \sum_{i \in \mathcal{E}} \hat{\tau}_{ik}.$$

┘

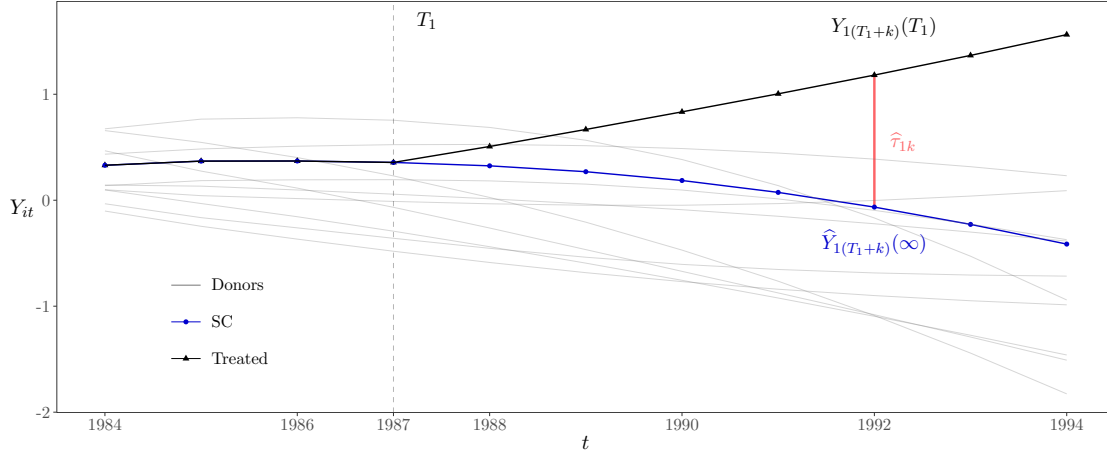
We conclude this section with a final note on the nature of the predictands defined above. The potential outcomes, treatment adoption times, and individual effects are all viewed as *random* quantities. We assume that there is only a fixed (possibly small) number of ever-treated units and time periods, which is often the case in synthetic control analysis and accommodates our empirical application. Thus, the various predictands defined above are also random quantities in general, which is why we prefer referring to them as “predictands” rather than as treatment “effects”. However, we do occasionally use the term “effect” or “predicted effect” to emphasize its randomness and to maintain consistency with the term used for analogous quantities defined in the SC literature under a fixed, non-random potential outcomes framework. In classical large-sample causal analysis, target parameters are often probability or ergodic (non-random) limits of the average effects above as $Q \rightarrow \infty$, $J_1 \rightarrow \infty$, and/or $T - T_i \rightarrow \infty$. Although our results are also valid in such large-sample settings, we develop statistical inference methods based on prediction intervals that describe a region where a new realization of a random causal predictand of interest is likely to be observed, rather than the usual confidence intervals giving a region in the parameter space for a non-random parameter of interest.

4 Uncertainty Quantification

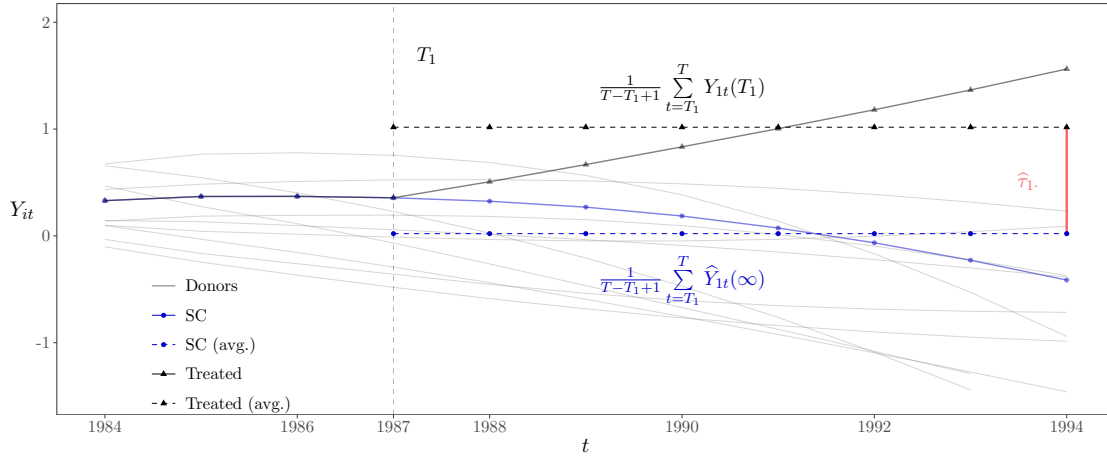
Let τ denote any of the four causal predictands defined above. Our goal is to construct a (random) prediction interval \mathcal{I} such that, with some high probability $1 - \pi$ over a conditioning set \mathcal{H} , \mathcal{I} covers τ with a pre-specified probability $1 - \alpha$ given \mathcal{H} :

$$\mathbb{P}\left\{\mathbb{P}[\tau \in \mathcal{I} \mid \mathcal{H}] \geq 1 - \alpha\right\} \geq 1 - \pi. \tag{4.1}$$

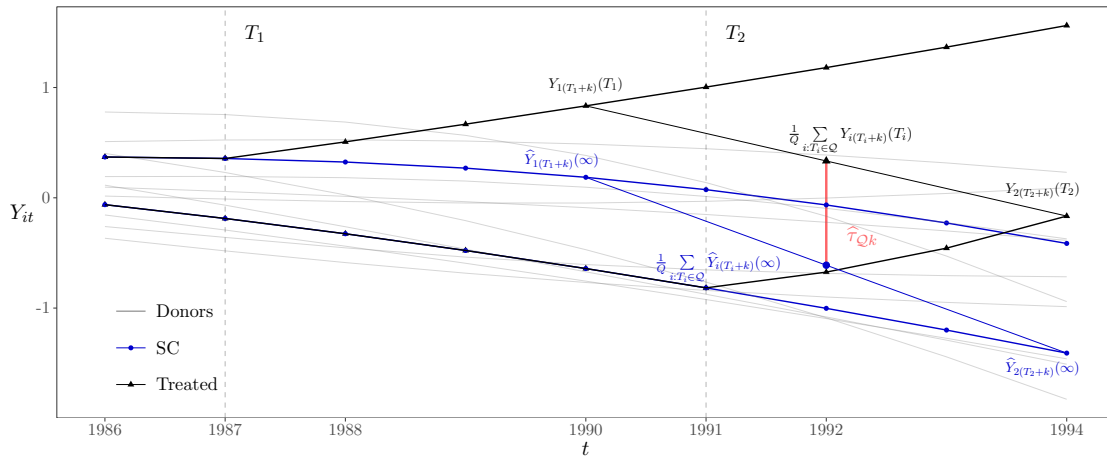
Figure 2: Graphical Representation of the Predictands.



(a) Time-specific unit-specific (TSUS) predictand.



(b) Time-averaged unit-specific (TAUS) predictand.



(c) Time-specific unit-averaged (TSUA) predictand.

Generally, the choice of the conditioning set \mathcal{H} determines the uncertainty that would not be taken into account by the prediction interval. If we do not condition on anything, technically, \mathcal{H} is a trivial σ -field, and \mathcal{I} reduces to an unconditional prediction interval. Since we regard synthetic controls as a regression-based out-of-sample forecasting problem, where it is standard to condition on the “covariates” (outcomes of donor countries in this context), we focus on the uncertainty from the ever-treated countries. Thus, we let $\mathcal{H} = \{Y_{it} : i \in \mathcal{N}, 1 \leq t \leq T\}$. Also, in practice one needs to set a desired (conditional) coverage level $(1 - \alpha)$, say 90%, whereas the probability loss π over \mathcal{H} is a “small” number that only needs to be theoretically characterized. In this paper, all results are valid if the training period is long enough ($T_i - 1$ is large), with the associated probability $1 - \pi$ characterized precisely. Thus, we say that the conditional prediction interval offers finite-sample probability guarantees. Our results imply that $\pi \rightarrow 0$ as $T_i \rightarrow \infty$, but no limits or asymptotic arguments are used in this paper.

To better understand the sources of uncertainty of SC predictions, we define the target quantity of the SC weights (conditional on \mathcal{H}) that is analogous to (3.1):

$$\mathbf{w}_0^{[i]} = \arg \min_{\mathbf{w} \in \mathcal{W}^{[i]}} \mathbb{E} \left[\sum_{t=1}^{T_i-1} \left(Y_{it} - \mathbf{Y}'_{\mathcal{N}t} \mathbf{w} \right)^2 \middle| \mathcal{H} \right]. \quad (4.2)$$

Thus, we can write

$$Y_{it}(\infty) = \mathbf{Y}'_{\mathcal{N}t} \mathbf{w}_0^{[i]} + u_{it}, \quad i \in \mathcal{E}, 1 \leq t \leq T, \quad (4.3)$$

where u_{it} is the corresponding pseudo-true residual relative to the conditioning set \mathcal{H} . Then, for each ever-treated country $i \in \mathcal{E}$, we have the following decomposition of the counterfactual outcome prediction:

$$\widehat{Y}_{i(T_i+k)}(\infty) - Y_{i(T_i+k)}(\infty) = \mathbf{Y}'_{\mathcal{N}(T_i+k)} (\widehat{\mathbf{w}}^{[i]} - \mathbf{w}_0^{[i]}) - u_{i(T_i+k)},$$

where $\mathbf{Y}'_{\mathcal{N}(T_i+k)} (\widehat{\mathbf{w}}^{[i]} - \mathbf{w}_0^{[i]})$ captures the *in-sample uncertainty* from the SC weights construction using the pre-treatment information, and $u_{i(T_i+k)}$ captures the *out-of-sample uncertainty* from the stochastic error in a specific post-treatment period. Notice that in-sample uncertainty quantification is necessary in this scenario since the conditioning set $Y_{it} \notin \mathcal{H}$ for $i \in \mathcal{E}$.

Accordingly, a similar decomposition can be performed for the prediction $\hat{\tau}$ of τ :

$$\hat{\tau} - \tau = \text{InErr}(\tau) + \text{OutErr}(\tau),$$

where $\text{InErr}(\tau)$ and $\text{OutErr}(\tau)$ denote the in-sample error and the out-of-sample error associated with the prediction of τ respectively. Specific expressions of such errors for each treatment effect prediction are given below:

$$\begin{aligned} \text{InErr}(\tau_{ik}) &= -\mathbf{Y}'_{\mathcal{N}(T_i+k)}(\widehat{\mathbf{w}}^{[i]} - \mathbf{w}_0^{[i]}), & \text{OutErr}(\tau_{ik}) &= u_{i(T_i+k)}, \\ \text{InErr}(\tau_{i\cdot}) &= -\frac{1}{T - T_i + 1} \sum_{k=0}^{T-T_i} \mathbf{Y}'_{\mathcal{N}(T_i+k)}(\widehat{\mathbf{w}}^{[i]} - \mathbf{w}_0^{[i]}), & \text{OutErr}(\tau_{i\cdot}) &= \frac{1}{T - T_i + 1} \sum_{k=0}^{T-T_i} u_{i(T_i+k)}, \\ \text{InErr}(\tau_{\mathcal{Q}k}) &= -\frac{1}{Q} \sum_{i:T_i \in \mathcal{Q}} \mathbf{Y}'_{\mathcal{N}(T_i+k)}(\widehat{\mathbf{w}}^{[i]} - \mathbf{w}_0^{[i]}), & \text{OutErr}(\tau_{\mathcal{Q}k}) &= \frac{1}{Q} \sum_{i:T_i \in \mathcal{Q}} u_{i(T_i+k)}, \\ \text{InErr}(\tau_{\cdot\cdot}) &= -\frac{1}{LJ_1} \sum_{k=1}^L \sum_{i \in \mathcal{E}} \mathbf{Y}'_{\mathcal{N}(T_i+k)}(\widehat{\mathbf{w}}^{[i]} - \mathbf{w}_0^{[i]}), & \text{OutErr}(\tau_{\cdot\cdot}) &= \frac{1}{LJ_1} \sum_{k=1}^L \sum_{i \in \mathcal{E}} u_{i(T_i+k)}. \end{aligned}$$

The target SC weights $\mathbf{w}_0^{[i]}$ and the residual u_{it} need to be understood in context. In principle, the weights $\mathbf{w}_0^{[i]}$ represent a “stable” cross-sectional relationship among treated and donor units that can be learned in the training period and translated to the post-treatment period, which is the common feature of all SC methods. In our empirical application, the outcome Y_{it} is GDP per capita, which is well known to be a non-stationary time series. Then, the idea of SC appears to be more applicable when GDP per capita sequences of different countries are cointegrated, where the “stable relationship” is given by the cointegrating vector, and the remainder u_{it} is a stationary series. Thus, we make the following assumption on the data generating process:

Assumption 1 (Data generating process). *Assume that for each $j \in \mathcal{N}$, $Y_{jt} = Y_{j(t-1)} + v_{jt}$, for each $i \in \mathcal{E}$, $Y_{it} = \mathbf{Y}'_{\mathcal{N}t} \mathbf{w}_0^{[i]} + u_{it}$, and $(\mathbf{u}'_t, \mathbf{v}'_t)$ is i.i.d. over t , where $\mathbf{u}_t = (u_{1t}, \dots, u_{J_1t})'$ and $\mathbf{v}_t = (v_{(J_1+1)t}, \dots, v_{Nt})'$. Assume \mathbf{v}_t is sub-Gaussian, and u_{it} is sub-Gaussian conditional on \mathcal{H} with parameter σ_{it} for each $i \in \mathcal{E}$.*

The outcomes of donor units are assumed to follow the simple unit root process with sub-Gaussian errors; other non-stationary patterns can also be accommodated at the cost of more technicalities. The conditional sub-Gaussianity of u_{it} , as precisely defined in Section 4.2, enables

us to characterize the tail probability of u_{it} exceeding a given bound, which is useful for quantifying out-of-sample uncertainty, but can also be replaced with weaker moment conditions. We emphasize that Assumption 1 is motivated by our empirical application and only used to illustrate the key idea of our method; our theory is more general and can be applied to other types of data, e.g., weakly dependent time series satisfying certain mixing conditions while allowing for a certain degree of model misspecification. See more discussion in Sections 4.1 and 6 below. Also, we view (4.3) as a *predictive* model, and the cointegrated system is merely one way to justify the SC approach; in the Supplemental Appendix Section S.3, we briefly discuss an alternative justification that assumes a linear factor model, and explain how to interpret the in-sample and out-of-sample uncertainty in that scenario.

We employ an intuitive strategy to construct prediction intervals for any causal predictand τ defined before. Specifically, if we can find random intervals $[\underline{M}_{\text{in}}(\tau), \overline{M}_{\text{in}}(\tau)]$ and $[\underline{M}_{\text{out}}(\tau), \overline{M}_{\text{out}}(\tau)]$ that (conditionally) cover the in-sample error $\text{InErr}(\tau)$ and the out-of-sample error $\text{OutErr}(\tau)$, respectively, with certain probabilities, i.e.,

$$\begin{aligned} \mathbb{P}\left\{\mathbb{P}\left[\underline{M}_{\text{in}}(\tau) \leq \text{InErr}(\tau) \leq \overline{M}_{\text{in}}(\tau) \mid \mathcal{H}\right] \geq 1 - \alpha_{\text{in}}\right\} &\geq 1 - \pi_{\text{in}} \quad \text{and} \\ \mathbb{P}\left\{\mathbb{P}\left[\underline{M}_{\text{out}}(\tau) \leq \text{OutErr}(\tau) \leq \overline{M}_{\text{out}}(\tau) \mid \mathcal{H}\right] \geq 1 - \alpha_{\text{out}}\right\} &\geq 1 - \pi_{\text{out}}, \end{aligned}$$

then, by a union bound we have

$$\mathbb{P}\left\{\mathbb{P}\left[\widehat{\tau} - \overline{M}_{\text{in}}(\tau) - \overline{M}_{\text{out}}(\tau) \leq \tau \leq \widehat{\tau} - \underline{M}_{\text{in}}(\tau) - \underline{M}_{\text{out}}(\tau) \mid \mathcal{H}\right] \geq 1 - \alpha_{\text{in}} - \alpha_{\text{out}}\right\} \geq 1 - \pi_{\text{in}} - \pi_{\text{out}}.$$

That is, the prediction interval $\mathcal{I}(\tau) := [\widehat{\tau} - \overline{M}_{\text{in}}(\tau) - \overline{M}_{\text{out}}(\tau), \widehat{\tau} - \underline{M}_{\text{in}}(\tau) - \underline{M}_{\text{out}}(\tau)]$ achieves $(1 - \alpha_{\text{in}} - \alpha_{\text{out}})$ coverage probability conditional on \mathcal{H} , which holds with probability at least $(1 - \pi_{\text{in}} - \pi_{\text{out}})$ over \mathcal{H} . In practice, we can set, for example, $\alpha_{\text{in}} = \alpha_{\text{out}} = \alpha/2$ to achieve the desired coverage level $1 - \alpha$, and our theory will precisely characterize $\pi_{\text{in}} + \pi_{\text{out}}$ and guarantee it is small at least when the training period is long enough. Bounding two errors separately makes the resultant prediction interval for τ conservative, but it clearly differentiates the contribution of the SC training procedure and the out-of-sample prediction procedure to the final inference.

4.1 In-Sample Uncertainty Quantification

We begin with the in-sample uncertainty quantification for the SC prediction of the time-specific counterfactual outcome $Y_{i(T_i+k)}(\infty)$ for each ever-treated unit, and then extend it to the four causal predictands defined before. Importantly, we take the decomposition (4.3) as a predictive model with no structural meaning, which does not have to be a “correctly specified” mean regression model. That is, we allow the outcomes of the donor countries $\mathbf{Y}_{\mathcal{N}t}$ to be possibly correlated with the SC residual u_{it} , i.e., $\boldsymbol{\gamma}^{[i]} := \mathbb{E}[\mathbf{Y}_{\mathcal{N}t}u_{it}|\mathcal{H}] \neq \mathbf{0}$, which violates the standard identification assumption for mean regression models. Accommodating this possibility is important, especially in high-dimensional scenarios, since SC constraints in $\mathcal{W}^{[i]}$ are usually imposed to make SC predictions feasible, and probably more stable, rather than being grounded in prior knowledge of the true data generating process.

Our proposed uncertainty quantification strategy only makes use of three basic facts: (i) $\widehat{\mathbf{w}}^{[i]}$ is the optimizer of the sample-based SC problem (3.1); (ii) $\mathbf{w}_0^{[i]}$ is the optimizer of the population-based SC problem (4.2); and (iii) the L1-L2 constraint set $\mathcal{W}^{[i]}$ used in our empirical application is convex. Actually, (i) and (ii) are true by definition of SC, and (iii) are true for all commonly used constraints, such as simplex, L1 constraint, L2 constraint, or some combinations thereof. Consequently, it can be shown that the following inequality holds (deterministically):

$$(\widehat{\mathbf{w}}^{[i]} - \mathbf{w}_0^{[i]})' \widehat{\mathbf{Q}}^{[i]} (\widehat{\mathbf{w}}^{[i]} - \mathbf{w}_0^{[i]}) - 2(\widehat{\boldsymbol{\gamma}}^{[i]} - \boldsymbol{\gamma}^{[i]})' (\widehat{\mathbf{w}}^{[i]} - \mathbf{w}_0^{[i]}) \leq 0 \quad \text{for each } i \in \mathcal{E},$$

where $\widehat{\mathbf{Q}}^{[i]} = \sum_{t=1}^{T_i-1} \mathbf{Y}_{\mathcal{N}t} \mathbf{Y}'_{\mathcal{N}t}$ and $\widehat{\boldsymbol{\gamma}}^{[i]} = \sum_{t=1}^{T_i-1} \mathbf{Y}_{\mathcal{N}t} u_{it}$. Therefore, the following bounds on the in-sample error of the prediction $\widehat{Y}_{i(T_i+k)}(\infty)$ hold:

$$\inf_{\boldsymbol{\delta} \in \mathcal{M}_{\widehat{\boldsymbol{\gamma}}^{[i]}}^{[i]}} \mathbf{Y}'_{\mathcal{N}(T_i+k)} \boldsymbol{\delta} \leq \mathbf{Y}'_{\mathcal{N}(T_i+k)} (\widehat{\mathbf{w}}^{[i]} - \mathbf{w}_0^{[i]}) \leq \sup_{\boldsymbol{\delta} \in \mathcal{M}_{\widehat{\boldsymbol{\gamma}}^{[i]}}^{[i]}} \mathbf{Y}'_{\mathcal{N}(T_i+k)} \boldsymbol{\delta},$$

where $\mathcal{M}_{\widehat{\boldsymbol{\gamma}}^{[i]}}^{[i]} = \{\boldsymbol{\delta} \in \Delta^{[i]} : \boldsymbol{\delta}' \widehat{\mathbf{Q}}^{[i]} \boldsymbol{\delta} - 2(\widehat{\boldsymbol{\gamma}}^{[i]} - \boldsymbol{\gamma}^{[i]})' \boldsymbol{\delta} \leq 0\}$ and $\Delta^{[i]} = \{\boldsymbol{\delta} : \boldsymbol{\delta} + \mathbf{w}_0^{[i]} \in \mathcal{W}^{[i]}\}$. Conditional on \mathcal{H} , $\widehat{\boldsymbol{\gamma}}^{[i]}$ is random, making the feasible set $\mathcal{M}_{\widehat{\boldsymbol{\gamma}}^{[i]}}^{[i]}$, and thus the resulting upper and lower bounds, stochastic as well. We can employ a normal distributional approximation of $\widehat{\boldsymbol{\gamma}}^{[i]}$ and, for instance, set the lower and upper bounds to $(\alpha_{\text{in}}/2)$ -quantile of $\inf_{\boldsymbol{\delta} \in \mathcal{M}_{\widehat{\boldsymbol{\gamma}}^{[i]}}^{[i]}} \mathbf{Y}'_{\mathcal{N}(T_i+k)} \boldsymbol{\delta}$ and $(1 - \alpha_{\text{in}}/2)$ -quantile

of $\sup_{\boldsymbol{\delta} \in \mathcal{M}_{\mathbf{G}}^{[i]}} \mathbf{Y}'_{\mathcal{N}(T_i+k)} \boldsymbol{\delta}$, respectively, conditional on \mathcal{H} , where $\mathcal{M}_{\mathbf{G}}^{[i]} = \{\boldsymbol{\delta} \in \Delta^{[i]} : \boldsymbol{\delta}' \widehat{\mathbf{Q}}^{[i]} \boldsymbol{\delta} - 2\mathbf{G}^{[i]'} \boldsymbol{\delta} \leq 0\}$, $\mathbf{G}^{[i]} | \mathcal{H} \sim \mathbf{N}(\mathbf{0}, \boldsymbol{\Sigma}^{[i]})$ and $\boldsymbol{\Sigma}^{[i]} = \mathbb{V}[\widehat{\boldsymbol{\gamma}}^{[i]} | \mathcal{H}]$. Such an interval has at least $(1 - \alpha_{\text{in}})$ conditional coverage probability for the in-sample error $\mathbf{Y}'_{\mathcal{N}(T_i+k)} (\widehat{\mathbf{w}}^{[i]} - \mathbf{w}_0^{[i]})$.

However, the above bounds cannot be directly implemented in practice because the normalized constraint set $\Delta^{[i]}$ depends on the unknown pseudo-true value $\mathbf{w}_0^{[i]}$, and the normal vector $\mathbf{G}^{[i]}$ depends on the unknown covariance matrix $\boldsymbol{\Sigma}^{[i]}$. We thus propose a feasible simulation-based strategy allowing for *possibly nonlinear* constraints, with unknown quantities replaced with plug-in approximations thereof.

First, we need a feasible constraint set $\Delta^{[i]\star}$ used in simulation that is “locally” similar to $\Delta^{[i]}$. Intuitively, we require that every point in the original infeasible constraint set $\Delta^{[i]}$ be *close* to the feasible constraint set $\Delta^{[i]\star}$ in simulation. Consequently, searching for an upper (or lower) bound within the infeasible set $\Delta^{[i]}$ can be replaced by doing so within the feasible set $\Delta^{[i]\star}$. In our empirical application, $\Delta^{[i]}$ is the L1-L2 constraint, which can be decomposed into a sequence of linear constraints on the weights vector $\mathbf{w} = (w_j : j \in \mathcal{N})$ of the form $\sum_{j \in \mathcal{N}} a_j w_j \leq Q_1^{[i]}$ for some $a_j \in \{+1, -1\}$, plus the L2 constraint $\sum_{j \in \mathcal{N}} w_j^2 \leq (Q_2^{[i]})^2$. Let $m_\ell^{[i]}(\mathbf{w}) \leq 0$ denote a generic inequality among them and d_{\leq} be the total number of constraints. We first use a parameter $\varrho^{[i]} = C^{[i]} \frac{\log T_0}{T_0}$ as a bound on the deviation of $\widehat{\mathbf{w}}^{[i]}$ from $\mathbf{w}_0^{[i]}$, where $C^{[i]}$ is some constant and can be chosen in a data-dependent way as described in Section 6.2. Then, we introduce a sequence of tuning parameters $\varrho_\ell^{[i]}$'s, where $\varrho_\ell^{[i]} = J_0 \varrho^{[i]}$ if $m_\ell^{[i]}(\mathbf{w}) \leq 0$ is a linear constraint, and $\varrho_\ell^{[i]} = 2 \|\widehat{\mathbf{w}}^{[i]}\|_1 \varrho^{[i]}$ if $m_\ell^{[i]}(\mathbf{w}) \leq 0$ is the L2 constraint. Importantly, the order of $\varrho_\ell^{[i]}$ accommodates the non-stationary data in our application and makes the condition on the “closeness” between $\Delta^{[i]}$ and $\Delta^{[i]\star}$ required in our general Theorem 1 hold. Now, given the parameters $\varrho_\ell^{[i]}$'s, we enforce the ℓ -th constraint to be binding in the simulation, i.e., $m_\ell^{[i]}(\mathbf{w}) \leq m_\ell^{[i]}(\widehat{\mathbf{w}}^{[i]})$, if $m_\ell^{[i]}(\widehat{\mathbf{w}}^{[i]}) > -\varrho_\ell^{[i]}$; otherwise, we keep the original constraint $m_\ell(\mathbf{w})^{[i]} \leq 0$. More discussion on methods to construct $\Delta^{[i]\star}$ is available in Section 6.2.

Second, we need an estimator $\widehat{\boldsymbol{\Sigma}}^{[i]}$ of the covariance matrix $\boldsymbol{\Sigma}^{[i]}$. A variety of well-established heteroskedasticity/serial-correlation-robust estimators can be used. We require $\widehat{\boldsymbol{\Sigma}}^{[i]}$ to be a “good” approximation of $\boldsymbol{\Sigma}^{[i]}$ in the sense of condition (iii) in Corollary 1 below. This allows us to approximate the infeasible normal distribution $\mathbf{N}(\mathbf{0}, \boldsymbol{\Sigma}^{[i]})$ by $\mathbf{N}(\mathbf{0}, \widehat{\boldsymbol{\Sigma}}^{[i]})$, which can be simulated using the data.

Once $\Delta^{[i]\star}$ and $\widehat{\Sigma}^{[i]}$ are available, we can simply draw random vectors from $\mathbf{N}(\mathbf{0}, \widehat{\Sigma}^{[i]})$ conditional on the data, and then set

$$\underline{M}_{\text{in}} = \underline{\mathfrak{c}}^*(\alpha_{\text{in}}/2) - \varepsilon_{\Delta} \quad \text{and} \quad \overline{M}_{\text{in}} = \overline{\mathfrak{c}}^*(1 - \alpha_{\text{in}}/2) + \varepsilon_{\Delta}, \quad (4.4)$$

where

$$\begin{aligned} \underline{\mathfrak{c}}^*(\alpha_{\text{in}}/2) &= (\alpha_{\text{in}}/2)\text{-quantile of } \inf_{\boldsymbol{\delta} \in \mathcal{M}_{\mathbf{G}}^{[i]\star}} \mathbf{Y}'_{\mathcal{N}(T_i+k)} \boldsymbol{\delta} \\ \overline{\mathfrak{c}}^*(1 - \alpha_{\text{in}}/2) &= (1 - \alpha_{\text{in}}/2)\text{-quantile of } \sup_{\boldsymbol{\delta} \in \mathcal{M}_{\mathbf{G}}^{[i]\star}} \mathbf{Y}'_{\mathcal{N}(T_i+k)} \boldsymbol{\delta}, \end{aligned}$$

conditional on the data, with $\mathcal{M}_{\mathbf{G}}^{[i]\star} = \{\boldsymbol{\delta} \in \Delta^{[i]\star} : \boldsymbol{\delta}' \widehat{\mathbf{Q}}^{[i]} \boldsymbol{\delta} - 2(\mathbf{G}^{[i]\star})' \boldsymbol{\delta} \leq 0\}$ and $\mathbf{G}^{[i]\star} | \text{Data} \sim \mathbf{N}(\mathbf{0}, \widehat{\Sigma}^{[i]})$. The constant ε_{Δ} , intuitively, measures the local “distance” between the two constraint sets $\Delta^{[i]}$ and $\Delta^{[i]\star}$ and is used to adjust for nonlinear constraints such as the L2 constraint imposed in our SC construction. However, in other applications with linear constraints (e.g., simplex and lasso) only, this adjustment is unnecessary so that we can set $\varepsilon_{\Delta} = 0$. In our empirical application, we set $\varepsilon_{\Delta} = C(\log T_0)^2 T_0^{-2} \|\mathbf{Y}_{\mathcal{N}(T_i+k)}\|_1$ for some constant $C > 0$, which accommodates the non-stationarity of the data and applies to the counterfactual outcome prediction problem; a general data-dependent strategy for selecting ε_{Δ} is proposed in Section 6.3. Under some mild regularity conditions, $[\underline{M}_{\text{in}}, \overline{M}_{\text{in}}]$ is a feasible prediction interval that achieves the desired conditional coverage of the counterfactual outcome $Y_{i(T_i+k)}$.

Now, we are ready to construct bounds on the in-sample error for each treatment effect prediction.

Example TSUS (Time-specific unit-specific predictand, continued). Since the in-sample error of $\widehat{\tau}_{ik}$ is the opposite of that of $\widehat{Y}_{i(T_i+k)}$, we simply set $\underline{M}_{\text{in}}(\tau_{ik}) = -\overline{M}_{\text{in}}$ and $\overline{M}_{\text{in}}(\tau_{ik}) = -\underline{M}_{\text{in}}$. \lrcorner

Example TAUS (Time-averaged unit-specific predictand, continued). Since the in-sample error of $\widehat{\tau}_{i\cdot}$ is (the opposite of) average in-sample error of $\widehat{Y}_{i(T_i+k)}$ over time, we can set $\underline{M}_{\text{in}}(\tau_{i\cdot}) = \underline{\mathfrak{c}}^*(\alpha_{\text{in}}/2) - \varepsilon_{\Delta}(\tau_{i\cdot})$ and $\overline{M}_{\text{in}}(\tau_{i\cdot}) = \overline{\mathfrak{c}}^*(1 - \alpha_{\text{in}}/2) + \varepsilon_{\Delta}(\tau_{i\cdot})$ where

$$\underline{\mathfrak{c}}^*(\alpha_{\text{in}}/2) = (\alpha_{\text{in}}/2)\text{-quantile of } \inf_{\boldsymbol{\delta} \in \mathcal{M}_{\mathbf{G}}^{[i]\star}} -\frac{1}{T - T_i + 1} \sum_{k=0}^{T-T_i} \mathbf{Y}'_{\mathcal{N}(T_i+k)} \boldsymbol{\delta},$$

$$\bar{\mathbf{c}}^*(1 - \alpha_{\text{in}}/2) = (1 - \alpha_{\text{in}}/2)\text{-quantile of } \sup_{\boldsymbol{\delta} \in \mathcal{M}_{\mathbf{G}}^{[i]*}} -\frac{1}{T - T_i + 1} \sum_{k=0}^{T-T_i} \mathbf{Y}'_{\mathcal{N}(T_i+k)} \boldsymbol{\delta},$$

conditional on the data, and $\varepsilon_{\Delta}(\tau_i) = C(\log T_0)^2 T_0^{-2} \left\| \frac{1}{T - T_i + 1} \sum_{k=0}^{T-T_i} \mathbf{Y}_{\mathcal{N}(T_i+k)} \right\|_1$ for some constant $C > 0$. \lrcorner

Example TSUA (Time-specific unit-averaged predictand, continued). Since the in-sample error of $\hat{\tau}_{\mathcal{Q}k}$ is (the opposite of) average in-sample error of $\hat{Y}_{i(T_i+k)}$ across multiple treated units, we can set $\underline{M}_{\text{in}}(\tau_{\mathcal{Q}k}) = \underline{\mathbf{c}}^*(\alpha_{\text{in}}/2) - \varepsilon_{\Delta}(\tau_{\mathcal{Q}k})$ and $\bar{M}_{\text{in}}(\tau_{\mathcal{Q}k}) = \bar{\mathbf{c}}^*(1 - \alpha_{\text{in}}/2) + \varepsilon_{\Delta}(\tau_{\mathcal{Q}k})$ where

$$\begin{aligned} \underline{\mathbf{c}}^*(\alpha_{\text{in}}/2) &= (\alpha_{\text{in}}/2)\text{-quantile of } \inf_{\boldsymbol{\delta} \in \mathcal{M}_{\mathbf{G}}^*} -\frac{1}{Q} \sum_{i: T_i \in \mathcal{Q}} \mathbf{Y}'_{\mathcal{N}(T_i+k)} \boldsymbol{\delta}^{[i]}, \\ \bar{\mathbf{c}}^*(1 - \alpha_{\text{in}}/2) &= (1 - \alpha_{\text{in}}/2)\text{-quantile of } \sup_{\boldsymbol{\delta} \in \mathcal{M}_{\mathbf{G}}^*} -\frac{1}{Q} \sum_{i: T_i \in \mathcal{Q}} \mathbf{Y}'_{\mathcal{N}(T_i+k)} \boldsymbol{\delta}^{[i]}, \end{aligned}$$

conditional on the data, $\boldsymbol{\delta} = (\boldsymbol{\delta}^{[i]} : T_i \in \mathcal{Q}) \in \mathbb{R}^{QJ}$, $\mathcal{M}_{\mathbf{G}}^* = \{\boldsymbol{\delta} \in \times_{i: T_i \in \mathcal{Q}} \Delta^{[i]*} : \sum_{i: T_i \in \mathcal{Q}} [\boldsymbol{\delta}^{[i]'} \hat{\mathbf{Q}}^{[i]} \boldsymbol{\delta}^{[i]'} - 2(\mathbf{G}^{[i]*})' \boldsymbol{\delta}^{[i]'}] \leq 0\}$, and $\varepsilon_{\Delta}(\tau_{\mathcal{Q}k}) = C(\log T_0)^2 T_0^{-2} \left\| \frac{1}{Q} \sum_{i: T_i \in \mathcal{Q}} \mathbf{Y}_{\mathcal{N}(T_i+k)} \right\|_1$ for some constant $C > 0$. In this case, the in-sample uncertainty depends on SC weights for multiple treated units, so we aggregate the previously describe quadratic constraint for each $\boldsymbol{\delta}^{[i]}$ into one single quadratic constraint for the entire vector $\boldsymbol{\delta}$. Note that with multiple treated units, we should draw the *long* vector $\mathbf{G}^* := (\mathbf{G}^{[i]*} : T_i \in \mathcal{Q})$ conditional on the data from $\mathbf{N}(\mathbf{0}, \hat{\boldsymbol{\Sigma}})$ with $\hat{\boldsymbol{\Sigma}}$ an estimate of $\boldsymbol{\Sigma} = \mathbb{V}[\hat{\boldsymbol{\gamma}} | \mathcal{H}]$ for $\hat{\boldsymbol{\gamma}} = (\hat{\boldsymbol{\gamma}}^{[i]} : T_i \in \mathcal{Q})$. Therefore, the correlation structure among different treated units is implicitly captured. \lrcorner

Example TAUA (Time-averaged unit-averaged predictand, continued). Since the in-sample error of $\hat{\tau}_{\cdot}$ is (the opposite of) average in-sample error of $\hat{Y}_{i(T_i+k)}$ across multiple treated units over time, we set $\underline{M}_{\text{in}}(\tau_{\cdot}) = \underline{\mathbf{c}}^*(\alpha_{\text{in}}/2) - \varepsilon_{\Delta}(\tau_{\cdot})$ and $\bar{M}_{\text{in}}(\tau_{\cdot}) = \bar{\mathbf{c}}^*(1 - \alpha_{\text{in}}/2) + \varepsilon_{\Delta}(\tau_{\cdot})$ where

$$\begin{aligned} \underline{M}_{\text{in}}(\tau_{\cdot}) &= (\alpha_{\text{in}}/2)\text{-quantile of } \inf_{\boldsymbol{\delta} \in \mathcal{M}_{\mathbf{G}}^*} -\frac{1}{LJ_1} \sum_{k=1}^L \sum_{i \in \mathcal{E}} \mathbf{Y}'_{\mathcal{N}(T_i+k)} \boldsymbol{\delta}^{[i]} \\ \bar{M}_{\text{in}}(\tau_{\cdot}) &= (1 - \alpha_{\text{in}}/2)\text{-quantile of } \sup_{\boldsymbol{\delta} \in \mathcal{M}_{\mathbf{G}}^*} -\frac{1}{LJ_1} \sum_{k=1}^L \sum_{i \in \mathcal{E}} \mathbf{Y}'_{\mathcal{N}(T_i+k)} \boldsymbol{\delta}^{[i]}, \end{aligned}$$

conditional on the data, $\mathcal{M}_{\mathbf{G}}^* = \{\boldsymbol{\delta} \in \times_{i: i \in \mathcal{E}} \Delta^{[i]*} : \sum_{i \in \mathcal{E}} [\boldsymbol{\delta}^{[i]'} \hat{\mathbf{Q}}^{[i]} \boldsymbol{\delta}^{[i]'} - 2(\mathbf{G}^{[i]*})' \boldsymbol{\delta}^{[i]'}] \leq 0\}$ and

$\varepsilon_{\Delta}(\tau_{\cdot}) = C(\log T_0)^2 T_0^{-2} \|\frac{1}{LJ_1} \sum_{k=1}^L \sum_{i \in \mathcal{E}} \mathbf{Y}_{\mathcal{N}(T_i+k)}\|_1$ for some constant $C > 0$. As in the previous example, we draw random vectors $\mathbf{G}^{\star} := (\mathbf{G}^{[i]\star} : i \in \mathcal{E})$ conditional on the data from $\mathbf{N}(\mathbf{0}, \widehat{\boldsymbol{\Sigma}})$ with $\widehat{\boldsymbol{\Sigma}}$ an estimate of $\boldsymbol{\Sigma} = \mathbb{V}[\widehat{\boldsymbol{\gamma}} | \mathcal{H}]$ for $\widehat{\boldsymbol{\gamma}} = (\widehat{\boldsymbol{\gamma}}^{[i]} : i \in \mathcal{E})$. \lrcorner

Remark 1 (Alternative Bounds). For the unit-averaged TSUA and TAUA predictands, we combine the quadratic constraints from optimization into a single constraint on the entire (centered) SC weights vector $\boldsymbol{\delta}$. This ‘‘aggregation’’ strategy is consistent with the view of taking (3.1) as a multiple-equation regression problem, as detailed in Section 6. By contrast, we could also retain the individual quadratic constraints $\boldsymbol{\delta}^{[i]\prime} \widehat{\mathbf{Q}}^{[i]} \boldsymbol{\delta}^{[i]} - 2(\mathbf{G}^{[i]\star})' \boldsymbol{\delta}^{[i]} \leq 0$, for $T_i \in \mathcal{Q}$ or $i \in \mathcal{E}$, each restricting a subvector $\boldsymbol{\delta}^{[i]}$ of $\boldsymbol{\delta}$. This alternative strategy applies to the special case of SC analysis in which SC weights for each treated unit are constructed separately, yielding possibly tighter bounds on in-sample errors; see more detailed discussion in Supplemental Appendix Section S.2.3. In Table S.2 of the Supplemental Appendix, we also illustrate the potential improvement of this method in our empirical application in terms of interval length. \lrcorner

Remark 2 (Scalable Optimization Implementations). The proposed bounds on in-sample errors are suprema and infima of linear functions subject to some linear or quadratic constraints on parameters. In Section S.5 of the Supplemental Appendix, we show such problems can be recast as conic optimization programs, which gives massive speed improvement in practice. \lrcorner

4.2 Out-of-Sample Uncertainty Quantification

To bound the out-of-sample error, we propose an easy-to-implement approach based on non-asymptotic concentration inequalities. Recall that by Assumption 1, $u_{it} - \mathbb{E}[u_{it} | \mathcal{H}]$ is conditional-on- \mathcal{H} sub-Gaussian with parameter σ_{it} , which implies that for any $\varepsilon > 0$,

$$\mathbb{P}\left(|u_{it} - \mathbb{E}[u_{it} | \mathcal{H}]| \geq \varepsilon \mid \mathcal{H}\right) \leq 2 \exp\left(-\frac{\varepsilon^2}{2\sigma_{it}^2}\right).$$

Consequently, for the out-of-sample error $u_{i(T_i+k)}$ of the counterfactual outcome prediction $\widehat{Y}_{i(T_i+k)}$, we can set

$$\begin{aligned}\underline{M}_{\text{out}} &= -\mathbb{E}[u_{i(T_i+k)}|\mathcal{H}] - \sqrt{2\sigma_{i(T_i+k)}^2 \log(2/\alpha_{\text{out}})}, \quad \text{and} \\ \overline{M}_{\text{out}} &= -\mathbb{E}[u_{i(T_i+k)}|\mathcal{H}] + \sqrt{2\sigma_{i(T_i+k)}^2 \log(2/\alpha_{\text{out}})},\end{aligned}\tag{4.5}$$

which yields a prediction interval $[\underline{M}_{\text{out}}, \overline{M}_{\text{out}}]$ that covers $-u_{i(T_i+k)}$ with at least $(1 - \alpha_{\text{out}})$ conditional coverage probability. We emphasize that the sub-Gaussianity assumption is one of many possibilities. The above strategy could be applied using other concentration inequalities requiring weaker moment conditions, though the resulting prediction intervals may be wider.

In practice, one could first construct pre-treatment residuals $\widehat{u}_{it} = Y_{it} - \mathbf{Y}'_{\mathcal{N}t} \widehat{\mathbf{w}}^{[i]}$ for each $i \in \mathcal{E}$ and $t < T_i$, and then estimate the conditional moments of u_{it} employing various parametric or nonparametric regression of \widehat{u}_{it} . Such estimates can then be translated into the necessary estimates of $\mathbb{E}[u_{i(T_i+k)}|\mathcal{H}]$ and $\sigma_{i(T_i+k)}^2$ for constructing $\underline{M}_{\text{in}}$ and $\overline{M}_{\text{out}}$. The unknown conditional moments could also be set using external information, or tabulated across different values to assess the sensitivity of the resulting prediction intervals.

We can apply the same idea to bound the out-of-sample errors of the four treatment effect predictions.

Example TSUS (Time-specific unit-specific predictand, continued). Since the out-of-sample error in this case is u_{it} , we simply set $\underline{M}_{\text{out}}(\tau_{ik}) = \mathbb{E}[u_{i(T_i+k)}|\mathcal{H}] - [2\sigma_{i(T_i+k)}^2 \log(2/\alpha_{\text{out}})]^{1/2}$ and $\overline{M}_{\text{out}}(\tau_{ik}) = \mathbb{E}[u_{i(T_i+k)}|\mathcal{H}] + [2\sigma_{i(T_i+k)}^2 \log(2/\alpha_{\text{out}})]^{1/2}$. \lrcorner

Example TAUS (Time-averaged unit-specific predictand, continued). Since $u_{it} - \mathbb{E}[u_{it}|\mathcal{H}]$ is condition-on- \mathcal{H} sub-Gaussian with parameter σ_{it} , it can be shown that the out-of-sample error $\text{OutErr}(\tau_i)$, as the average of u_{it} over time, is also sub-Gaussian, satisfying that for any $\varepsilon > 0$,

$$\mathbb{P}(|\text{OutErr}(\tau_i) - \mathbb{E}[\text{OutErr}(\tau_i)|\mathcal{H}]| \geq \varepsilon | \mathcal{H}) \leq 2 \exp\left(-\frac{\varepsilon^2}{2\sigma_i^2}\right), \quad \sigma_i := \frac{1}{T - T_i + 1} \sum_{t=T_i}^T \sigma_{it}.$$

Therefore, we can set

$$\underline{M}_{\text{out}}(\tau_i) = \frac{1}{T - T_i + 1} \sum_{k=0}^{T-T_i} \mathbb{E}[u_{i(T_i+k)}|\mathcal{H}] - \sqrt{2\sigma_i^2 \log(2/\alpha_{\text{out}})}, \quad \text{and}$$

$$\bar{M}_{\text{out}}(\tau_{i\cdot}) = \frac{1}{T - T_i + 1} \sum_{k=0}^{T-T_i} \mathbb{E}[u_{i(T_i+k)} | \mathcal{H}] + \sqrt{2\sigma_i^2 \log(2/\alpha_{\text{out}})}.$$

This result holds regardless of the dependence structure of u_{it} , but if u_{it} is indeed independent over t , the same result holds with $\sigma_i := \frac{1}{T-T_i+1} (\sum_{t=T_i}^T \sigma_{it}^2)^{1/2}$, leading to an improved bound. \lrcorner

Example TSUA (Time-specific unit-averaged predictand, continued). In this scenario, the out-of-sample error $\text{OutErr}(\tau_{Qk})$ is a cross-sectional average of u_{it} at different times. The uncertainty quantification strategy outlined previously in Example **TAUS** can still be applied, with the caveat that it is uncommon in SC analysis to assume u_{it} is stationary and/or independent over i . By contrast, it is reasonable to assume u_{it} is stationary and/or independent (at least weakly dependent) over time. Therefore, we employ the concentration inequality that holds under a general dependence structure and set

$$\begin{aligned} \underline{M}_{\text{out}}(\tau_{Qk}) &= \frac{1}{Q} \sum_{i:T_i \in Q} \mathbb{E}[u_{i(T_i+k)} | \mathcal{H}] - \sqrt{2\sigma_{Qk}^2 \log(2/\alpha_{\text{out}})}, \quad \text{and} \\ \bar{M}_{\text{out}}(\tau_{Qk}) &= \frac{1}{Q} \sum_{i:T_i \in Q} \mathbb{E}[u_{i(T_i+k)} | \mathcal{H}] + \sqrt{2\sigma_{Qk}^2 \log(2/\alpha_{\text{out}})} \end{aligned}$$

with $\sigma_{Qk} := \frac{1}{Q} \sum_{i:T_i \in Q} \sigma_{i(T_i+k)}$. \lrcorner

Example TAUA (Time-averaged unit-averaged predictand, continued). Since the adoption time T_i may be heterogeneous across i , the out-of-sample error $\text{OutErr}(\tau_{\cdot})$ is an average of out-of-sample errors of different units in different periods, which is (conditionally) sub-Gaussian as well. Then, we set

$$\begin{aligned} \underline{M}_{\text{out}}(\tau_{\cdot}) &= \frac{1}{LJ_1} \sum_{k=1}^L \sum_{i \in \mathcal{E}} \mathbb{E}[u_{i(T_i+k)} | \mathcal{H}] - \sqrt{2\sigma_{\cdot}^2 \log(2/\alpha_{\text{out}})}, \quad \text{and} \\ \bar{M}_{\text{out}}(\tau_{\cdot}) &= \frac{1}{LJ_1} \sum_{k=1}^L \sum_{i \in \mathcal{E}} \mathbb{E}[u_{i(T_i+k)} | \mathcal{H}] + \sqrt{2\sigma_{\cdot}^2 \log(2/\alpha_{\text{out}})} \end{aligned}$$

with $\sigma_{\cdot} := \frac{1}{LJ_1} \sum_{k=1}^K \sum_{i \in \mathcal{E}} \sigma_{i(T_i+k)}$. \lrcorner

In addition to the concentration-based approach described above, other strategies, including location-scale models and quantile regression, were proposed in [Cattaneo, Feng and Titiunik \(2021\)](#)

for out-of-sample uncertainty quantification. We briefly review them in Supplemental Appendix Section S.2.1.

4.3 Theoretical Justification

The following corollary, as a special case of our more general results presented in Section 6, provides a theoretical justification of the proposed method, which closely matches our empirical application. Let $T_0 = T_1 - 1$, and $\tilde{\mathbf{Q}} = \frac{1}{T_0} \sum_{t=1}^{T_0} \tilde{\mathbf{G}}(t/T_0) \tilde{\mathbf{G}}(t/T_0)'$ where $\tilde{\mathbf{G}}$ is a mean-zero Brownian motion on $[0, 1]$ with variance $\mathbb{E}[\mathbf{v}_t \mathbf{v}_t']$. We use $s_{\min}(\mathbf{A})$ and $s_{\max}(\mathbf{A})$ to denote the minimum and maximum singular values of a matrix \mathbf{A} , respectively. Also, for simplicity, in the following corollary we define $\Sigma = \mathbb{V}[\hat{\gamma} | \mathcal{H}]$ for $\hat{\gamma} = (\hat{\gamma}^{[i]} : i \in \mathcal{E})$ and let $\hat{\Sigma}$ be an estimator of Σ , as in Example [TAUA](#).

Corollary 1. *Let Assumption 1 hold. In addition, suppose that with probability over \mathcal{H} at least $1 - \pi_0$, the following conditions hold:*

- (i) $\min_{1 \leq t \leq T} s_{\min}(\mathbb{E}[\mathbf{u}_t \mathbf{u}_t' | \mathcal{H}]) > \eta$ for some constant $\eta > 0$;
- (ii) $(\log T_0)^{-1/5} \leq s_{\min}(\tilde{\mathbf{Q}}) \leq s_{\max}(\tilde{\mathbf{Q}}) \leq (\log T_0)^{1/5}$;
- (iii) $\mathbb{P}(\|\hat{\Sigma} - \Sigma\| \leq \epsilon_{\Sigma,1}^* | \mathcal{H}) \geq 1 - \epsilon_{\Sigma,2}^*$ with $\epsilon_{\Sigma,1}^* \leq \frac{T_0^2 \eta}{8\sqrt{d}(\log T_0)^{1/5}}$;
- (iv) For all $1 \leq \ell \leq d_{\leq}$, $\varrho_{\ell}^{[i]} < |m_{\ell}^{[i]}(\mathbf{w}_0^{[i]})| - (\log T_0) T_0^{-1}$ if $m_{\ell}^{[i]}(\mathbf{w}_0^{[i]}) \neq 0$.

Then, for any causal prediction $\hat{\tau} \in \{\hat{\tau}_{ik}, \hat{\tau}_i, \hat{\tau}_{Qk}, \hat{\tau}_{..}\}$, when T_0 is large enough,

$$\mathbb{P}\left\{\mathbb{P}(\tau \in [\hat{\tau} - \bar{M}_{\text{in}}(\tau) - \bar{M}_{\text{out}}(\tau), \hat{\tau} - \underline{M}_{\text{in}}(\tau) - \underline{M}_{\text{out}}(\tau)] | \mathcal{H}) \geq 1 - \alpha_{\text{in}} - \alpha_{\text{out}} - \epsilon\right\} \geq 1 - \pi,$$

where $\epsilon = 2\mathfrak{C}_{\epsilon}(\log T_0)^2 T_0^{-1/2} + 4\epsilon_{\Sigma,1}^* \sqrt{d}(\log T_0)^{1/5} T_0^{-2} \eta^{-1} + \epsilon_{\Sigma,2}^* + 3T_0^{-1}$, $\pi = \pi_0 + \frac{\mathfrak{C}_{\pi}}{T_0}$, and \mathfrak{C}_{ϵ} and \mathfrak{C}_{π} are some constants characterized in the proof.

The additional technical conditions imposed in this corollary are intuitive: (i) is a mild restriction used to guarantee the variance matrix Σ is non-degenerate; (ii) is used to characterize the upper and lower bounds on the matrices $\hat{\mathbf{Q}}^{[i]}$ and can be shown to hold with high probability when T_0 is large; (iii) is a high-level condition on the convergence of the variance estimator $\hat{\Sigma}$ and can be verified on a case-by-case basis; and (iv) guarantees that the non-binding constraints can be differentiated from the binding ones under our thresholding rule, making the ‘‘closeness’’ requirement for the constraint sets imposed in Theorem 1 satisfied. If these conditions hold with high probability (π_0 is

small), then this corollary shows that the proposed prediction intervals can achieve approximately $(1 - \alpha_{\text{in}} - \alpha_{\text{out}})$ conditional coverage probability, with the probability loss ϵ decreasing as T_0 grows and $\epsilon_{\Sigma,1}^*$ and $\epsilon_{\Sigma,2}^*$ get small (the variance estimator $\widehat{\Sigma}$ becomes more “precise”).

4.4 Simultaneous Prediction Intervals

So far we have focused on constructing prediction intervals that have high coverage of the desired predictands or treatment effects, in particular, the TSUS predictand, which captures the effect of the intervention for a given unit in *each* post-treatment period. In some applications, it might be appealing to construct prediction intervals that have high *simultaneous* coverage in multiple post-treatment periods, usually termed *simultaneous prediction intervals*. They can be employed to test, for example, whether the largest (or smallest) treatment effect across different periods is significantly different from zero.

Specifically, for a particular treated unit $i \in \mathcal{E}$, we aim to construct a sequence of intervals \mathcal{I}_k for $0 \leq k \leq L$ for some $L \leq T - T_i$ such that

$$\mathbb{P}\left\{\mathbb{P}[\tau_{ik} \in \mathcal{I}_k, \text{ for all } 0 \leq k \leq L \mid \mathcal{H}] \geq 1 - \alpha\right\} \geq 1 - \pi.$$

As described before, the uncertainty of the predicted TSUS effect $\widehat{\tau}_{ik}$ comes from the in-sample error $\text{InErr}(\tau_{ik}) = -Y'_{\mathcal{N}(T_i+k)}(\widehat{\mathbf{w}}^{[i]} - \mathbf{w}_0^{[i]})$ and the out-of-sample error $\text{OutErr}(\tau_{ik}) = u_{i(T_i+k)}$.

Regarding the in-sample error, the following is an immediate generalization of the bound described in Section 4.1, which enjoys simultaneous coverage in multiple periods:

$$\underline{M}_{\text{in}}(\tau_{ik}) = \underline{\mathfrak{c}}^*(\alpha_{\text{in}}/2) - \bar{\epsilon}_{\Delta} \quad \text{and} \quad \bar{M}_{\text{in}}(\tau_{ik}) = \bar{\mathfrak{c}}^*(1 - \alpha_{\text{in}}/2) + \bar{\epsilon}_{\Delta}, \quad (4.6)$$

where

$$\begin{aligned} \underline{\mathfrak{c}}^*(\alpha_{\text{in}}/2) &= (\alpha_{\text{in}}/2)\text{-quantile of } \inf_{\boldsymbol{\delta} \in \mathcal{M}_{\mathbf{G}}^{[i]*}, 0 \leq k \leq L} -\mathbf{Y}'_{\mathcal{N}(T_i+k)} \boldsymbol{\delta} \\ \bar{\mathfrak{c}}^*(1 - \alpha_{\text{in}}/2) &= (1 - \alpha_{\text{in}}/2)\text{-quantile of } \sup_{\boldsymbol{\delta} \in \mathcal{M}_{\mathbf{G}}^{[i]*}, 0 \leq k \leq L} -\mathbf{Y}'_{\mathcal{N}(T_i+k)} \boldsymbol{\delta}, \end{aligned}$$

and the constant $\bar{\epsilon}_{\Delta}$, as before, is a small tuning parameter used to adjust for nonlinear constraints

and can be set to zero if all constraints are linear. Given the adjustment $\varepsilon_\Delta(\tau_{ik})$ for each post-treatment period (see Example TSUS in Section 4.1), we can set $\bar{\varepsilon}_\Delta = \max_{0 \leq k \leq L} \varepsilon_\Delta(\tau_{ik})$. It can be shown that this generally guarantees that $[\underline{M}_{\text{in}}(\tau_{ik}), \bar{M}_{\text{in}}(\tau_{ik})]$ can simultaneously cover the in-sample error $\text{InErr}(\tau_{ik})$ for all $0 \leq k \leq L$, with $\mathbb{P}(\cdot|\mathcal{H})$ -probability at least $1 - \alpha_{\text{in}}$, up to some small loss.

Regarding the out-of-sample error, an easy-to-implement strategy analogous to that described in Section 4.2 is to adjust the bounds on the out-of-sample error $\text{OutErr}(\tau_{ik}) = u_{i(T_i+k)}$ based on maximal inequalities. Recall that each $u_{i(T_i+k)} - \mathbb{E}[u_{i(T_i+k)}|\mathcal{H}]$, $0 \leq k \leq L$, is conditional sub-Gaussian with parameter $\sigma_{i(T_i+k)}$ (but is not necessarily independent over k). Then,

$$\mathbb{P}\left(\max_{0 \leq k \leq L} |u_{i(T_i+k)} - \mathbb{E}[u_{i(T_i+k)}|\mathcal{H}]| \geq \varepsilon \mid \mathcal{H}\right) \leq 2 \sum_{k=0}^L \exp\left(-\frac{\varepsilon^2}{2\sigma_{i(T_i+k)}^2}\right).$$

If $\sigma_{i(T_i+k)} \leq \sigma_i$ for all $0 \leq k \leq T - T_i$, then one can set $\underline{M}_{\text{out}} = \mathbb{E}[u_{i(T_i+k)}|\mathcal{H}] - \varepsilon$ and $\bar{M}_{\text{out}} = \mathbb{E}[u_{i(T_i+k)}|\mathcal{H}] + \varepsilon$ with $\varepsilon = \sqrt{2\sigma_i^2 \log(2(L+1)/\alpha_2)}$. It can be shown that $[\underline{M}_{\text{out}}, \bar{M}_{\text{out}}]$ can simultaneously cover $\text{OutErr}(\tau_{ik})$ for all $0 \leq k \leq L$, with $\mathbb{P}(\cdot|\mathcal{H})$ -probability at least $1 - \alpha_{\text{out}}$. Compared with prediction intervals with validity for each period constructed the same way, these simultaneous prediction intervals are slightly wider due to the additional factor $\sqrt{\log(L+1)}$. In practice, one only needs to estimate the conditional mean and variance of u_{it} using the pre-treatment residuals; flexible parametric or non-parametric estimation methods can be used.

Again, the sub-Gaussianity assumption can be relaxed by using other concentration inequalities requiring weaker moment conditions, though the resulting simultaneous prediction intervals may be wider. Also, there are other strategies to construct prediction intervals that simultaneously cover multiple out-of-sample errors, though they are computationally more cumbersome and usually require more stringent conditions. See Supplemental Appendix Section S.2.2 for a brief discussion.

The idea outlined above to achieve simultaneous coverage is general and can also be used to, for example, construct prediction intervals that simultaneously cover the TSUS predictands for multiple treated units rather than for multiple post-treatment periods. In our empirical application, we construct simultaneous prediction intervals for time-averaged effects across different economies; see details in Section 5.

5 Empirical Results

We use our framework to evaluate the effect of economic liberalization on (log) real GDP per capita in the sixteen countries that adopted economic liberalization in Sub-Saharan Africa: Benin, Botswana, Cabo Verde, Cameroon, Gambia, Ghana, Guinea, Guinea-Bissau, Ivory Coast, Mali, Mauritius, Niger, South Africa, Uganda, and Zambia.

We emphasize that when the target predictand does not involve an average across countries (e.g., TSUS and TAUS), we include not-yet-treated countries in the donor pool. Including not-yet-treated units in other scenarios would raise issues, as their pre-treatment features would appear both in the conditioning set \mathcal{H} and in \mathbf{A} , thereby invalidating the attempt of quantifying their contribution to prediction uncertainty.

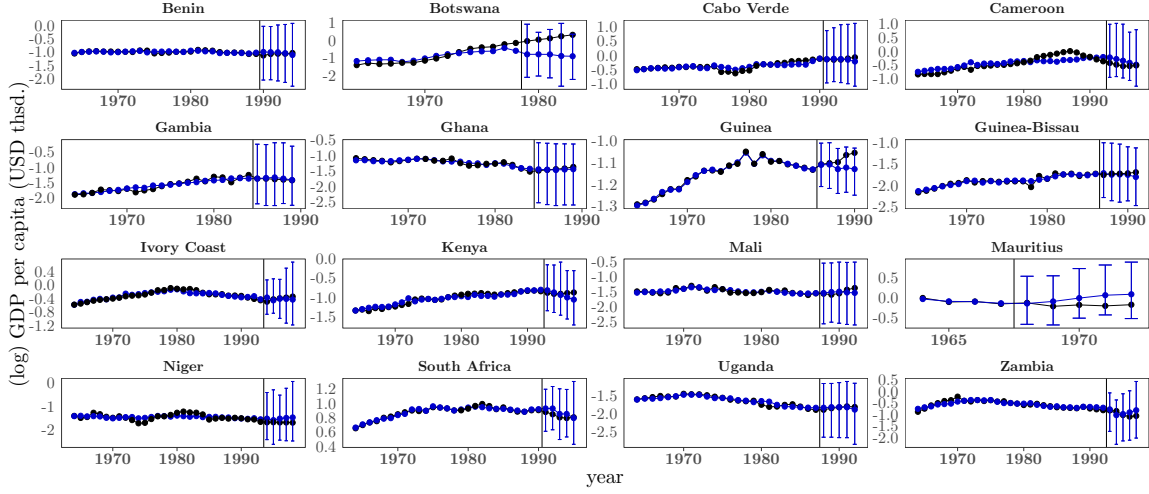
Overall, our point predictions suggest that (external) liberalization episodes in Sub-Saharan Africa had no impact on GDP per capita. Although we find that the point prediction for Botswana is large and positive, our predictands and uncertainty characterization indicate that this predicted effect cannot be distinguished from zero. Our new evidence thus suggests that Botswana’s success story was possibly the cause rather than the consequence of economic liberalization. In what follows, we report the predictands described in Section 3, Examples [TSUS](#), [TAUS](#), and [TSUA](#).

In all our results, we compute the weights using an L1-L2 constraint and use pre-treatment GDP as the only feature. In the Supplemental Appendix Section S.7, we also report the results using (i) a simplex-type constraint, (ii) a ridge-type constraint, (iii) more than one feature, (iv) using placebo treatment dates, and (v) leaving one of the donors out at a time. All the implementation details are reported in Supplemental Appendix Section S.6.

TSUS predicted effects in every period after liberalization (τ_{ik}). We first analyze the predicted treatment effect for every individual country that adopts liberalization in each post-liberalization period (up to five years after adoption), which is an example of the [TSUS](#) predictand. In Figure 3, we show the predicted synthetic control outcomes (panel (a)) with the corresponding 90% prediction intervals, and the computed weights $\widehat{\mathbf{w}}^{[i]}$, $i \in \mathcal{E}$ (panel (b)) for our sixteen countries.


In most countries, the realized trajectory of GDP per capita (black lines) lies roughly on top of the synthetic one (blue lines), suggesting that in the absence of the liberalization event, real income per capita would have been approximately equal. Looking at the 90% prediction intervals (blue

Figure 3: *Time-specific unit-specific (TSUS) predicted effects in every period, $\hat{\tau}_{ik}$.*



→ Treated → Synthetic Control

(a) $Y_{it}(T_i)$ and $\hat{Y}_{it}(\infty)$



(b) $\hat{w}_j^{[i]}, i \in \mathcal{E}, j \in \mathcal{N}$

Notes: Panel (a): TSUS prediction for every country in each of five periods after treatment. Blue bars report 90% prediction intervals. In-sample uncertainty is quantified using 200 simulations, whereas out-of-sample uncertainty is quantified using sub-Gaussian bounds. Panel (b): each dot represents the weight that the donor (row) gets in forming the synthetic control for the treated unit (column). When there is no dot, it means that the unit was not part of the donor pool for the treated unit in question.

vertical bars), we can see that in all cases the distance between the actual GDP series and the counterfactual one is not different from zero with high probability for almost all units and periods.

For Botswana, the realized trajectory of GDP per capita (black lines) lies above the synthetic one, suggesting a positive treatment effect. However, the 90% prediction intervals show that this effect cannot be distinguished from zero with high probability.

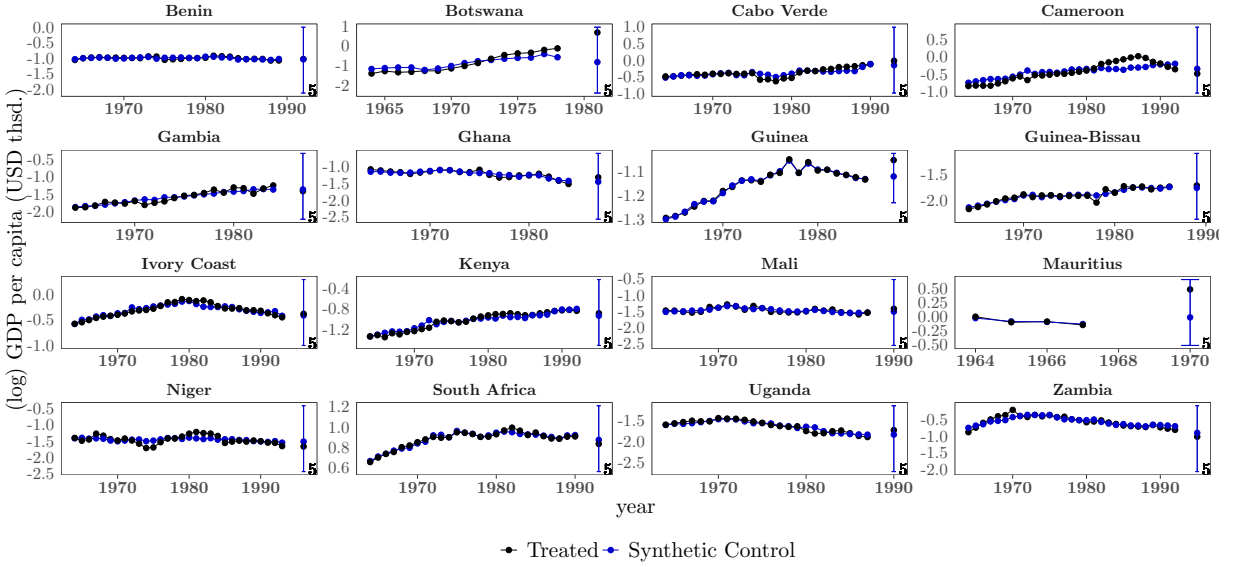
TAUS predicted effects, averaged over five years (τ_i). The second causal predictand of interest is the effect for each of the sixteen African countries we study, averaged over the five periods following the liberalization treatment in each country (up to the year 2000). This is an example of the [TAUS](#) predictand. Figure 4 shows that in all countries the liberalization episode seems to have a negligible effect on real GDP per capita. By looking at the prediction intervals, we can see that no treated unit shows a negative average effect with high probability.

Although Botswana's predictand is the highest, once uncertainty is quantified this effect cannot be distinguished from zero with high probability.


TSUA predicted effects, averaged over countries that liberalized in each of three waves: before 1987, between 1987 and 1991, and after 1991 ($\tau_{Q_1k}, \tau_{Q_2k}, \tau_{Q_3k}$). In interpreting the evidence for each individual country, [Billmeier and Nannicini \(2013\)](#) consider the hypothesis that liberalization only led to economic growth for countries that liberalized early. Our framework allows us to group countries according to the era in which they liberalized and consider their joint trajectory in a formal way. To do so, we use the three waves considered by [Billmeier and Nannicini \(2013\)](#) and study three different predictands that average over all countries that liberalized during each wave: the TSUA effect for all countries that liberalized before 1986 (Botswana, Gambia, Ghana, and Guinea), τ_{Q_1k} with $Q_1 = \{t : t < 1987\}$; the TSUA effect for all countries that liberalized between 1987 and 1991 (Benin, Cabo Verde, Guinea-Bissau, Mali, South Africa, and Uganda), τ_{Q_2k} with $Q_2 = \{t : 1987 \leq t \leq 1991\}$; and the TSUA effect for all countries that liberalized between 1992 and 1994 (Burkina Faso, Burundi, Cameroon, Ethiopia, Ivory Coast, Mozambique, Niger, Tanzania, and Zambia), τ_{Q_3k} with $Q_3 = \{t : 1991 < t \leq 1994\}$. The predicted effect is calculated for every post-liberalization period (up to five years), each representing a specific example of the [TSUA](#) predictand.

Our evidence does not support their hypothesis. The results, presented in Figure 5, show that, with high probability, we can conclude that countries that liberalized followed a trajectory similar to the counterfactual under no liberalization in the three waves. Although the trajectory of the

Figure 4: Time-averaged unit-specific (TAUS) predicted effects, averaged over five years, $\hat{\tau}_i$.



(a) $Y_{it}(T_i)$ and $\hat{Y}_{it}(\infty)$



(b) $\hat{w}_j^{[i]}, i \in \mathcal{E}, j \in \mathcal{N}$

Notes: Panel (a): TAUS prediction for every country averaged over the five periods following treatment (up to the year 2000). Blue bars report 90% prediction intervals. In-sample uncertainty is quantified using 200 simulations, whereas out-of-sample uncertainty is quantified using sub-Gaussian bounds. Panel (b): each dot represents the weight that the donor (row) gets in forming the synthetic control for the treated unit (column). When there is no dot, it means that the unit was not part of the donor pool for the treated unit in question.

early adopters is above the counterfactual, Figure 5(a) shows that the average effect for the early treated countries starts to diverge from the counterfactual trajectory of the synthetic control even before the treatment is adopted. Even if spurious, the observed divergence cannot be distinguished from zero once we use our framework to account for uncertainty.

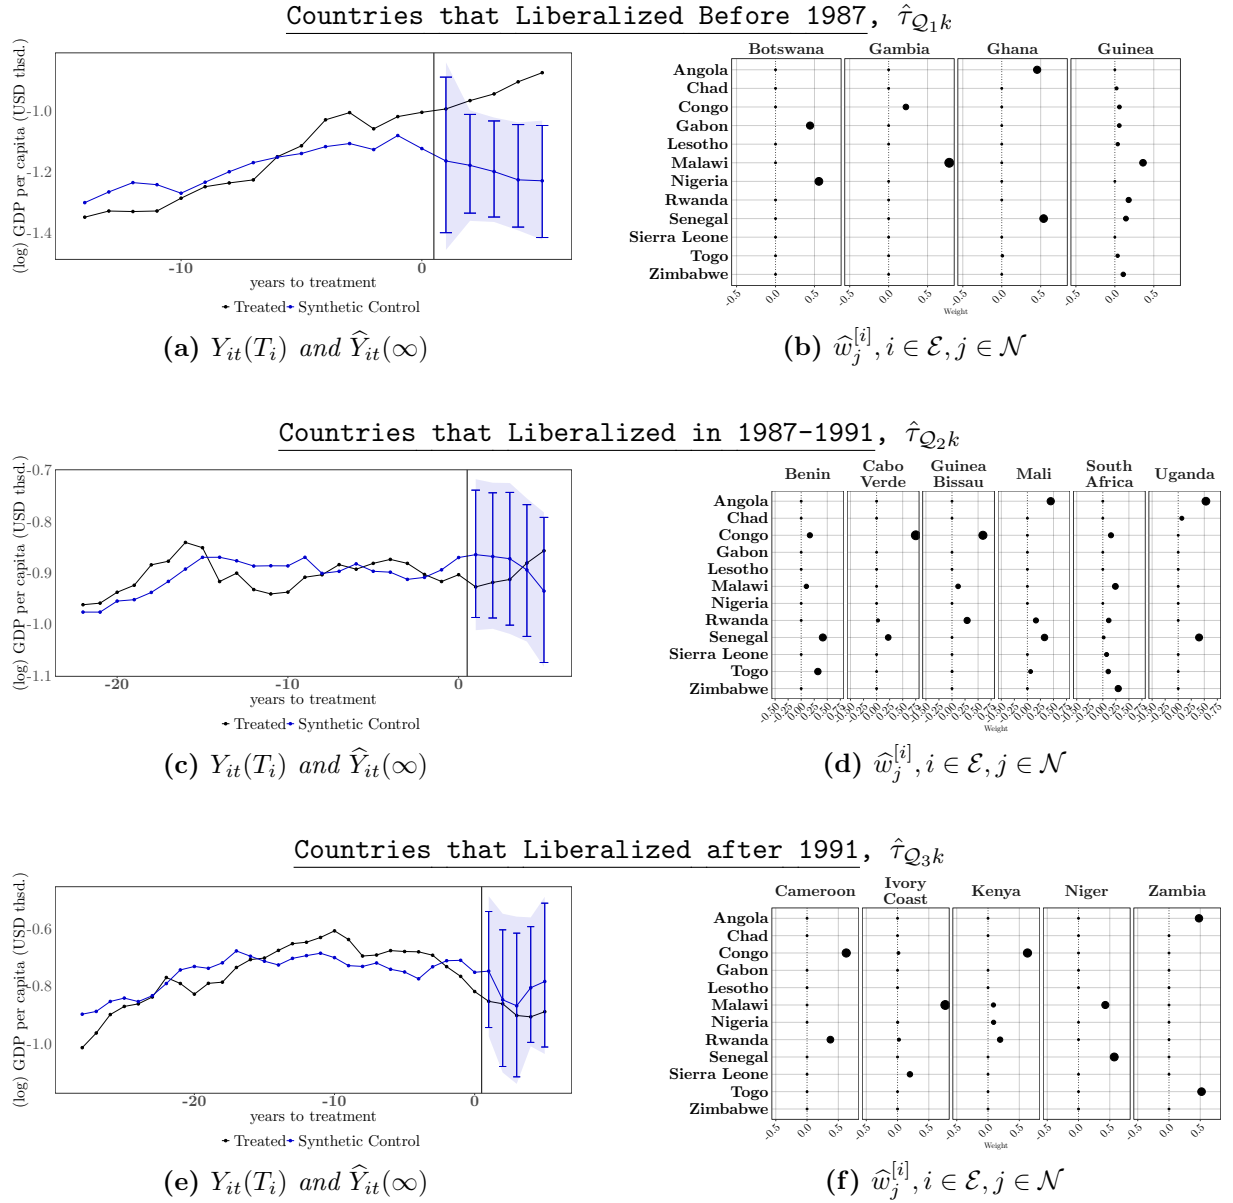
TSUA predicted effects, averaged over all liberalized countries ($\tau_{\mathcal{E}k}$). In this last and fourth exercise, we focus on a popular causal predictand: the effect in every period after treatment averaged over all the treated countries. This is yet another example of the [TSUA](#) predictand, with $\mathcal{Q} = \mathcal{E}$. This predictand thus averages over all countries that liberalized, regardless of when they did so. We report this predictand in every year after the adoption of liberalization (up to five years) which occurs at different times for different countries. Figure 6 reports the results. To compute this predictand, we first construct a synthetic control for each treated unit (see panel (b)) and then we pool all the synthetic controls together in every post-treatment period to get a single prediction and a single prediction interval. Panel (a) and panel (b) show that pooling across the 16 African countries that embarked on liberalization programs helps reduce the uncertainty surrounding the synthetic trajectory. Indeed, this TSUA predictand shows very similar trajectories of the synthetic real income and the actual real income, which are indistinguishable from each other with high probability for the average treated country.

6 Theoretical Foundations

This section presents a general framework that can accommodate more flexible specifications in SC analysis. As in Section 3, we still consider the case with J_0 never-treated units and J_1 ever-treated units that adopt a treatment at possibly different times. However, now we assume that a user may want to obtain SC weights by matching on M features (denoted by a subscript $l = 1, \dots, M$ below) with additional covariates adjustment, rather than relying solely on pre-treatment outcomes.

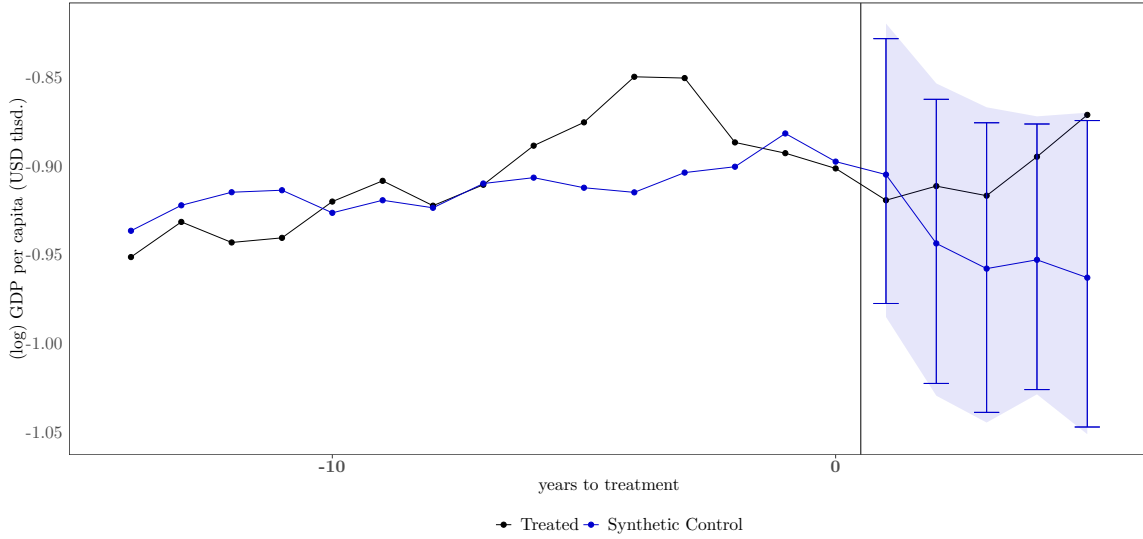
Specifically, let $\mathbf{A}_l^{[i]} = (a_{1,l}^{[i]}, \dots, a_{T_{i0},l}^{[i]})' \in \mathbb{R}^{T_{i0}}$ be the l -th feature of the treated unit i measured in T_{i0} (user-specified) pre-treatment periods. For each feature l and each treated unit i , there exist $J_0 + K$ variables that are used to predict or match the T_{i0} -dimensional vector $\mathbf{A}_l^{[i]}$. These $J_0 + K$ variables are separated into two groups denoted by $\mathbf{B}_l^{[i]} = (\mathbf{B}_{1,l}^{[i]}, \mathbf{B}_{2,l}^{[i]}, \dots, \mathbf{B}_{J_0,l}^{[i]}) \in \mathbb{R}^{T_{i0} \times J_0}$ and $\mathbf{C}_l^{[i]} = (\mathbf{C}_{1,l}^{[i]}, \dots, \mathbf{C}_{K,l}^{[i]}) \in \mathbb{R}^{T_{i0} \times K}$, respectively. More precisely, for each $j = 1, \dots, J_0$, $\mathbf{B}_{j,l}^{[i]} =$

Figure 5: Time-specific unit-averaged (TSUA) predicted effects in each period, averaged over three groups of countries.




Notes: TSUA prediction in every period after treatment (up to five years), averaged over all countries that liberalized in each of three waves: before 1987 (Botswana, Gambia, Ghana, and Guinea), between 1987 and 1991 (Benin, Cabo Verde, Guinea-Bissau, Mali, South Africa, and Uganda), and after 1991 (Burkina Faso, Burundi, Cameroon, Ethiopia, Ivory Coast, Mozambique, Niger, Tanzania, and Zambia). Blue bars report 90% prediction intervals, whereas blue-shaded areas report 90% simultaneous prediction intervals. In-sample uncertainty is quantified using 200 simulations, whereas out-of-sample uncertainty is quantified using sub-Gaussian bounds. Panel (b): each dot represents the weight that the donor (row) gets in forming the synthetic control for the treated unit (column). When there is no dot, it means that the unit was not part of the donor pool for the treated unit in question.

Figure 6: *Time-specific unit-averaged (TSUA) predicted effect, averaged over all treated units, $\widehat{\tau}_{\mathcal{E}k}$.*



(a) $Y_{it}(T_i)$ and $\widehat{Y}_{it}(\infty)$



(b) $\widehat{w}_j^{[i]}, i \in \mathcal{E}, j \in \mathcal{N}$

Notes: Panel (a): TSUA prediction in every period after treatment, averaged over all the treated countries. Blue bars report 90% prediction intervals, whereas blue-shaded areas report 90% simultaneous prediction intervals. In-sample uncertainty is quantified using 200 simulations, whereas out-of-sample uncertainty is quantified using sub-Gaussian bounds. Panel (b): each dot represents the weight that the donor (row) gets in forming the synthetic control for the treated unit (column). When there is no dot, it means that the unit was not part of the donor pool for the treated unit in question.

$(b_{j1,l}^{[i]}, \dots, b_{jT_{i0},l}^{[i]})'$ corresponds to the l -th feature of the j -th unit in the donor pool measured in T_{i0} pre-treatment periods, and for each $k = 1, \dots, K$, $\mathbf{C}_{k,l}^{[i]} = (c_{k1,l}^{[i]}, \dots, c_{kT_{i0},l}^{[i]})'$ is another vector of control variables used to predict $\mathbf{A}_l^{[i]}$ over the same pre-intervention time span. Stacking the M equations (corresponding to M features) for each treated unit, we define

$$\underbrace{\mathbf{A}^{[i]}}_{T_{i0} \cdot M \times 1} = \begin{bmatrix} \mathbf{A}_1^{[i]} \\ \vdots \\ \mathbf{A}_M^{[i]} \end{bmatrix}, \quad \underbrace{\mathbf{B}^{[i]}}_{T_{i0} \cdot M \times J_0} = \begin{bmatrix} \mathbf{B}_1^{[i]} \\ \vdots \\ \mathbf{B}_M^{[i]} \end{bmatrix}, \quad \underbrace{\mathbf{C}^{[i]}}_{T_{i0} \cdot M \times K \cdot M} = \begin{bmatrix} \mathbf{C}_1^{[i]} & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \mathbf{C}_2^{[i]} & \dots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & \mathbf{C}_M^{[i]} \end{bmatrix}.$$

For instance, in the Supplemental Appendix we revisit our empirical application, where $\mathbf{A}^{[i]}$ contains the (log) GDP per capita and the investment-to-GDP ratio ($M = 2$) of an ever-liberalized economy i during the pre-liberalization period, and $\mathbf{B}^{[i]}$ contains the same two features of the donor economies used to match $\mathbf{A}^{[i]}$. For each feature $l = 1, 2$, $\mathbf{C}_l^{[i]}$ contains an intercept and a linear time trend ($K = 2$).

We search for a vector of weights $\mathbf{w} = (\mathbf{w}^{[1]'}, \dots, \mathbf{w}^{[J_1]'})' \in \mathcal{W} \subseteq \mathbb{R}^{J_0 J_1}$, which is common across the M features, and a vector of coefficients $\mathbf{r} = (\mathbf{r}^{[1]'}, \dots, \mathbf{r}^{[J_1]'})' \in \mathcal{R} \subseteq \mathbb{R}^{K M J_1}$, such that the linear combination of $\mathbf{B}^{[i]}$ and $\mathbf{C}^{[i]}$ matches $\mathbf{A}^{[i]}$ as closely as possible, for all $i \in \mathcal{E}$. The feasibility sets \mathcal{W} and \mathcal{R} capture the restrictions imposed. Analogously to (3.1), such SC weights are obtained via the following optimization problem: for some $(\tilde{T} \cdot M) \times (\tilde{T} \cdot M)$ symmetric weighting matrix \mathbf{V} with $\tilde{T} = \sum_{i=1}^{J_1} T_{0i}$,

$$\hat{\boldsymbol{\beta}} := (\hat{\mathbf{w}}', \hat{\mathbf{r}}')' \in \arg \min_{\mathbf{w} \in \mathcal{W}, \mathbf{r} \in \mathcal{R}} (\mathbf{A} - \mathbf{B}\mathbf{w} - \mathbf{C}\mathbf{r})' \mathbf{V} (\mathbf{A} - \mathbf{B}\mathbf{w} - \mathbf{C}\mathbf{r}) \quad (6.1)$$

where

$$\underbrace{\mathbf{A}}_{\tilde{T} \cdot M \times 1} = \begin{bmatrix} \mathbf{A}^{[1]} \\ \vdots \\ \mathbf{A}^{[J_1]} \end{bmatrix}, \quad \underbrace{\mathbf{B}}_{\tilde{T} \cdot M \times J_0 \cdot J_1} = \begin{bmatrix} \mathbf{B}^{[1]} & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \mathbf{B}^{[2]} & \dots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & \mathbf{B}^{[J_1]} \end{bmatrix}, \quad \underbrace{\mathbf{C}}_{\tilde{T} \cdot M \times K \cdot M \cdot J_1} = \begin{bmatrix} \mathbf{C}^{[1]} & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \mathbf{C}^{[2]} & \dots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & \mathbf{C}^{[J_1]} \end{bmatrix}.$$

Accordingly, we write $\hat{\mathbf{w}} = (\hat{\mathbf{w}}^{[1]'}, \dots, \hat{\mathbf{w}}^{[J_1]'})'$ where each $\hat{\mathbf{w}}^{[i]} = (\hat{w}_1^{[i]}, \dots, \hat{w}_{J_0}^{[i]})'$ is the SC weights on J_0 donor units that are used to predict the counterfactual of the treated unit i . Similarly, write $\hat{\mathbf{r}} = (\hat{\mathbf{r}}^{[1]'}, \dots, \hat{\mathbf{r}}^{[J_1]'})'$ and $\hat{\boldsymbol{\beta}} = (\hat{\boldsymbol{\beta}}^{[1]'}, \dots, \hat{\boldsymbol{\beta}}^{[J_1]'})'$.

Remark 3 (Weighting Matrix). As pointed out by [Ben-Michael et al. \(2022\)](#), with multiple treated units, the SC weights could be constructed in two ways: (i) optimizing the separate fit for each treated unit; and (ii) optimizing the pooled fit for the average of the treated units. These ideas can be accommodated by choosing a proper weighting matrix \mathbf{V} . For example, assume $T_{i0} = T_0$ for simplicity. Taking $\mathbf{V} = \mathbf{I}_{T_0 M J_1}$ yields

$$\widehat{\boldsymbol{\beta}} = \arg \min_{\mathbf{w} \in \mathcal{W}, \mathbf{r} \in \mathcal{R}} \sum_{i=1}^{J_1} \sum_{l=1}^M \sum_{t=1}^{T_0} \left(a_{t,l}^{[i]} - \mathbf{b}_{t,l}^{[i]'} \mathbf{w}_l^{[i]} - \mathbf{c}_{t,l}^{[i]'} \mathbf{r}_l^{[i]} \right)^2,$$

where $\mathbf{B}_l^{[i]} := (\mathbf{b}_{1,l}^{[i]}, \dots, \mathbf{b}_{T_0,l}^{[i]})'$ is the l -th feature of the J_0 donor units, and $\mathbf{C}_l^{[i]} := (\mathbf{c}_{1,l}^{[i]}, \dots, \mathbf{c}_{T_0,l}^{[i]})'$ is the additional K variables used to predict $\mathbf{A}_l^{[i]}$. The objective above is equivalent to minimizing the sum of squared errors of the pre-treatment fit for *each* treated unit and thus is termed “separate fit”. By contrast, consider the following weighting matrix: $\mathbf{V} = \frac{1}{J_1^2} \mathbf{1}_{J_1} \mathbf{1}_{J_1}' \otimes \mathbf{I}_{T_0 M}$ where \otimes denotes the Kronecker product operator. Then,

$$\widehat{\boldsymbol{\beta}} = \arg \min_{\mathbf{w} \in \mathcal{W}, \mathbf{r} \in \mathcal{R}} \sum_{l=1}^M \sum_{t=1}^{T_0} \left[\frac{1}{J_1} \sum_{i=1}^{J_1} \left(a_{t,l}^{[i]} - \mathbf{b}_{t,l}^{[i]'} \mathbf{w}_l^{[i]} - \mathbf{c}_{t,l}^{[i]'} \mathbf{r}_l^{[i]} \right) \right]^2.$$

In this case, the goal is to minimize the sum of squared *averaged* errors across all treated units, which is usually termed “pooled fit”. ┘

Given the SC weights, the predicted counterfactual outcome of each treated unit $i \in \mathcal{E}$ is

$$\widehat{Y}_{i(T_i+k)}(\infty) := \mathbf{x}_{T_i+k}^{[i]'} \widehat{\mathbf{w}}^{[i]} + \mathbf{g}_{T_i+k}^{[i]'} \widehat{\boldsymbol{\Gamma}}^{[i]} = \mathbf{p}_{T_i+k}^{[i]'} \widehat{\boldsymbol{\beta}}^{[i]}, \quad \mathbf{p}_{T_i+k}^{[i]} = (\mathbf{x}_{T_i+k}^{[i]'}, \mathbf{g}_{T_i+k}^{[i]'}), \quad k \geq 0,$$

where $\mathbf{x}_{T_i+k}^{[i]}$ is a vector of predictors of the donor units used to predict the counterfactual of the treated unit i measured k periods after the treatment, and $\mathbf{g}_{T_i+k}^{[i]}$ is a vector of predictors that correspond to the additional control variables specified in $\mathbf{C}^{[i]}$. In general, variables included in $\mathbf{x}_{T_i+k}^{[i]}$ and $\mathbf{g}_{T_i+k}^{[i]}$ need not be the same as those in $\mathbf{B}^{[i]}$ and $\mathbf{C}^{[i]}$.

Again, let τ denote any of the four causal predictands defined in [Section 3](#). Then, the prediction of τ can be constructed accordingly and uniformly expressed as

$$\widehat{\tau} = \mathbf{L}(\{Y_{it}\}) - \mathbf{p}'_{\tau} \widehat{\boldsymbol{\beta}}, \tag{6.2}$$

with $L(\{Y_{it}\})$ some linear combination of observed post-treatment outcomes and $\mathbf{p}'_{\tau}\widehat{\boldsymbol{\beta}}$ the prediction of the corresponding counterfactual. \mathbf{p}_{τ} here denotes a predictor vector associated with the predictand τ , whose specific expression in each case is as follows:

$$\begin{aligned}\mathbf{p}_{\tau_{ik}} &= \underbrace{(\mathbf{0}'_{J_0+KM}, \dots, \mathbf{0}'_{J_0+KM})}_{(i-1) \text{ vectors}}, \mathbf{p}_{T_i+k}^{[i]'}, \underbrace{(\mathbf{0}'_{J_0+KM}, \dots, \mathbf{0}'_{J_0+KM})}'_{(J_1-i) \text{ vectors}}, \\ \mathbf{p}_{\tau_i} &= \left(\underbrace{(\mathbf{0}'_{J_0+KM}, \dots, \mathbf{0}'_{J_0+KM})}_{(i-1) \text{ vectors}}, \frac{1}{T - T_i + 1} \sum_{t \geq T_i} \mathbf{p}_t^{[i]'}, \underbrace{(\mathbf{0}'_{J_0+KM}, \dots, \mathbf{0}'_{J_0+KM})}'_{(J_1-i) \text{ vectors}} \right)', \\ \mathbf{p}_{\tau_{\mathcal{Q}k}} &= \left(\frac{1}{Q} \mathbf{p}_{T_1+k}^{[1]'} \mathbf{1}(T_1 \in \mathcal{Q}), \frac{1}{Q} \mathbf{p}_{T_2+k}^{[2]'} \mathbf{1}(T_2 \in \mathcal{Q}), \dots, \frac{1}{Q} \mathbf{p}_{T_{J_1}+k}^{[J_1]'} \mathbf{1}(T_{J_1} \in \mathcal{Q}) \right)', \quad \text{and} \\ \mathbf{p}_{\tau..} &= \left(\frac{1}{LJ_1} \sum_{k=1}^L \mathbf{p}_{T_1+k}^{[1]'}, \frac{1}{LJ_1} \sum_{k=1}^L \mathbf{p}_{T_2+k}^{[2]'}, \dots, \frac{1}{LJ_1} \sum_{k=1}^L \mathbf{p}_{T_{J_1}+k}^{[J_1]'} \right)',\end{aligned}$$

where $\mathbf{0}_{J_0+KM}$ denotes a $(J_0 + KM)$ -dimensional vector of zeros.

Analogously to (4.2), the pseudo-true value of SC weights in this framework is defined by

$$\boldsymbol{\beta}_0 := (\mathbf{w}'_0, \mathbf{r}'_0)' = \arg \min_{\mathbf{w} \in \mathcal{W}, \mathbf{r} \in \mathcal{R}} \mathbb{E} \left[(\mathbf{A} - \mathbf{B}\mathbf{w} - \mathbf{C}\mathbf{r})' \mathbf{V} (\mathbf{A} - \mathbf{B}\mathbf{w} - \mathbf{C}\mathbf{r}) \middle| \mathcal{H} \right], \quad (6.3)$$

and then we can write

$$\mathbf{A} = \mathbf{B}\mathbf{w}_0 + \mathbf{C}\mathbf{r}_0 + \mathbf{U}, \quad \mathbf{w}_0 \in \mathcal{W}, \quad \mathbf{r}_0 \in \mathcal{R}, \quad (6.4)$$

where $\mathbf{U} = (\mathbf{u}^{[1]'}, \dots, \mathbf{u}^{[J_1]'})' \in \mathbb{R}^{\tilde{T}M}$ is the corresponding pseudo-true residual relative to the σ -field $\mathcal{H} = \{\mathbf{B}, \mathbf{C}, \mathbf{p}_{\tau}\}$.

As before, we differentiate the contribution of the in-sample error $\text{InErr}(\tau)$ and the out-of-sample error $\text{OutErr}(\tau)$ to the uncertainty of SC prediction of τ . For in-sample uncertainty, the optimization bounds used in Section 4.1 can be generalized to this setup. Specifically, let $d = J_0 + KM$, $\mathbf{Z} = (\mathbf{B}, \mathbf{C})$, $\widehat{\mathbf{Q}} = \mathbf{Z}'\mathbf{V}\mathbf{Z}$, $\widehat{\boldsymbol{\gamma}}' = \mathbf{U}'\mathbf{V}\mathbf{Z}$, $\boldsymbol{\gamma} = \mathbb{E}[\widehat{\boldsymbol{\gamma}} | \mathcal{H}]$, and $\Delta = \{\boldsymbol{\beta} - \boldsymbol{\beta}_0 \in \mathbb{R}^d : \boldsymbol{\beta} \in \mathcal{W} \times \mathcal{R}\}$. It follows from the optimality of $\widehat{\boldsymbol{\beta}}$ and $\boldsymbol{\beta}_0$ and the convexity of \mathcal{W} and \mathcal{R} that $\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0 \in \Delta$ and $(\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0)' \widehat{\mathbf{Q}} (\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0) - 2(\widehat{\boldsymbol{\gamma}} - \boldsymbol{\gamma})' (\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0) \leq 0$. Also, given the expression (6.2), $\text{InErr}(\tau)$ can be generally expressed as $-\mathbf{p}_{\tau}'(\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0)$. Then, a valid, though infeasible, prediction interval for the in-sample error $\text{InErr}(\tau)$ is $[\underline{\mathfrak{c}}(\alpha_{\text{in}}/2), \bar{\mathfrak{c}}(1 - \alpha_{\text{in}}/2)]$, where $\underline{\mathfrak{c}}(\alpha)$ denotes the α -quantile of

$\inf_{\delta \in \mathcal{M}_{\mathbf{G}}} -\mathbf{p}'_{\tau} \delta$ and $\bar{\mathbf{c}}(\alpha)$ denotes α -quantile of $\sup_{\delta \in \mathcal{M}_{\mathbf{G}}} -\mathbf{p}'_{\tau} \delta$ conditional on \mathcal{H} for any $\alpha \in (0, 1)$, with $\mathcal{M}_{\mathbf{G}} = \{\delta \in \Delta : \delta' \widehat{\mathbf{Q}} \delta - 2\mathbf{G}' \delta \leq 0\}$, $\mathbf{G} | \mathcal{H} \sim \mathbf{N}(\mathbf{0}, \mathbf{\Sigma})$ and $\mathbf{\Sigma} = \mathbb{V}[\widehat{\gamma} | \mathcal{H}]$. Analogously to (4.4), once a feasible variance estimator $\widehat{\mathbf{\Sigma}}$ of $\mathbf{\Sigma}$ and a feasible constraint set Δ^* that locally approximates the original Δ (in the sense of condition (iii) in Theorem 1 below) is available, we can set

$$\underline{M}_{\text{in}}(\tau) = \underline{\mathbf{c}}^*(\alpha_{\text{in}}/2) - \varepsilon_{\Delta}(\tau) \quad \text{and} \quad \overline{M}_{\text{in}}(\tau) = \overline{\mathbf{c}}^*(1 - \alpha_{\text{in}}/2) + \varepsilon_{\Delta}(\tau) \quad (6.5)$$

where $\underline{\mathbf{c}}^*(\alpha_{\text{in}}/2)$ is the $(\alpha_{\text{in}}/2)$ -quantile of $\inf_{\delta \in \mathcal{M}_{\mathbf{G}}^*} -\mathbf{p}'_{\tau} \delta$, $\overline{\mathbf{c}}^*(1 - \alpha_{\text{in}}/2)$ is the $(1 - \alpha_{\text{in}}/2)$ -quantile of $\sup_{\delta \in \mathcal{M}_{\mathbf{G}}^*} -\mathbf{p}'_{\tau} \delta$ conditional on the data, and $\varepsilon_{\Delta}(\tau)$, as in Section 4, is a tuning parameter used to adjust for nonlinear constraints, with $\mathcal{M}_{\mathbf{G}}^* = \{\delta \in \Delta^* : \delta' \widehat{\mathbf{Q}} \delta - 2(\mathbf{G}^*)' \delta \leq 0\}$ and $\mathbf{G}^* | \text{Data} \sim \mathbf{N}(\mathbf{0}, \widehat{\mathbf{\Sigma}})$.

For out-of-sample error, analogously to (4.5), if $\text{OutErr}(\tau)$ is assumed to be sub-Gaussian conditional on \mathcal{H} with parameter $\sigma_{\mathcal{H}}$, then we can set

$$\begin{aligned} \underline{M}_{\text{out}}(\tau) &= \mathbb{E}[\text{OutErr}(\tau) | \mathcal{H}] - \sqrt{2\sigma_{\mathcal{H}}^2 \log(2/\alpha_{\text{out}})} \quad \text{and} \\ \overline{M}_{\text{out}}(\tau) &= \mathbb{E}[\text{OutErr}(\tau) | \mathcal{H}] + \sqrt{2\sigma_{\mathcal{H}}^2 \log(2/\alpha_{\text{out}})}, \end{aligned} \quad (6.6)$$

Now, we present a general theorem that justifies the above method under high-level conditions and covers the results given in Section 4 as a special case. Let $\|\cdot\|_*$ be the dual norm of $\|\cdot\|$ for any generic ℓ_p vector norm $\|\cdot\|$ on \mathbb{R}^d with $p \geq 1$, $\|\cdot\|_{\mathbf{F}}$ the Frobenius matrix norm, and $\mathcal{B}(\mathbf{0}, \varepsilon)$ an ε -neighborhood around zero for some $\varepsilon > 0$. Define the distance between a point $\mathbf{a} \in \mathbb{R}^d$ and a set $\Lambda \subseteq \mathbb{R}^d$ by $\text{dist}(\mathbf{a}, \Lambda) = \inf_{\lambda \in \Lambda} \|\mathbf{a} - \lambda\|$.

Theorem 1. *Assume \mathcal{W} and \mathcal{R} are convex, $\widehat{\beta}$ in (6.1) and β_0 in (6.3) exist, $\mathcal{H} = \sigma(\mathbf{B}, \mathbf{C}, \mathbf{p}_{\tau})$, and $\underline{M}_{\text{in}}(\tau)$, $\overline{M}_{\text{in}}(\tau)$, $\underline{M}_{\text{out}}(\tau)$ and $\overline{M}_{\text{out}}(\tau)$ are specified as in (6.5) and (6.6). In addition, for some finite non-negative constants ϵ_{γ} , π_{γ} , ϖ_{δ}^* , ϵ_{δ}^* , π_{δ}^* , ϖ_{Δ}^* , ϵ_{Δ}^* , π_{Δ}^* , $\epsilon_{\gamma,1}^*$, $\epsilon_{\gamma,2}^*$ and π_{γ}^* , the following conditions hold:*

- (i) $\mathbb{P}[\mathbb{P}(\mathbf{p}'_{\tau}(\widehat{\beta} - \beta_0) \in [\underline{\mathbf{c}}(\alpha_0), \overline{\mathbf{c}}(1 - \alpha_0)] | \mathcal{H}) \geq 1 - 2\alpha_0 - \epsilon_{\gamma}] \geq 1 - \pi_{\gamma}$ for any $\alpha_0 \in (0, 1)$;
- (ii) $\mathbb{P}[\mathbb{P}(\sup\{\|\delta\| : \delta \in \mathcal{M}_{\mathbf{G}}\} \leq \varpi_{\delta}^* | \mathcal{H}) \geq 1 - \epsilon_{\delta}^*] \geq 1 - \pi_{\delta}^*$;
- (iii) $\mathbb{P}[\mathbb{P}(\sup_{\mathbf{a} \in \Delta \cap \mathcal{B}(0, \varpi_{\delta}^*)} \text{dist}(\mathbf{a}, \Delta^* \cap \mathcal{B}(0, \varepsilon)) \leq \varpi_{\Delta}^* | \mathcal{H}) \geq 1 - \epsilon_{\Delta}^*] \geq 1 - \pi_{\Delta}^*$;
- (iv) $\mathbb{P}[\mathbb{P}(\|\mathbf{\Sigma}^{-1/2} \widehat{\mathbf{\Sigma}} \mathbf{\Sigma}^{-1/2} - \mathbf{I}_d\|_{\mathbf{F}} \leq 2\epsilon_{\gamma,1}^* | \mathcal{H}) \geq 1 - \epsilon_{\gamma,2}^*] \geq 1 - \pi_{\gamma}^*$;

(v) $\text{OutErr}(\tau) - \mathbb{E}[\text{OutErr}(\tau)|\mathcal{H}]$ is sub-Gaussian conditional on \mathcal{H} with parameter $\sigma_{\mathcal{H}}$.

Then, for $\epsilon_{\gamma,1}^* \in [0, 1/4]$ and $\epsilon_{\Delta}(\tau) \geq \|\mathbf{P}_{\tau}\|_* \varpi_{\Delta}^*$,

$$\mathbb{P}\left\{\mathbb{P}(\tau \in [\hat{\tau} - \bar{M}_{\text{in}}(\tau) - \bar{M}_{\text{out}}(\tau), \hat{\tau} - \underline{M}_{\text{in}}(\tau) - \underline{M}_{\text{out}}(\tau)]|\mathcal{H}) \geq 1 - \alpha_{\text{in}} - \alpha_{\text{out}} - \epsilon\right\} \geq 1 - \pi,$$

where $\epsilon = \epsilon_{\gamma} + 2\epsilon_{\gamma,1}^* + \epsilon_{\gamma,2}^* + 2\epsilon_{\delta}^* + \epsilon_{\Delta}^*$ and $\pi = \pi_{\gamma} + \pi_{\gamma}^* + \pi_{\delta}^* + \pi_{\Delta}^*$.

Assumptions (i)–(iv) imposed in Theorem 1 are high-level, which are used for in-sample uncertainty quantification and can be verified in many practically relevant scenarios such as the cointegrated data considered in Section 4. We give more detailed discussion of each condition in Section 6.1 below. Assumption (v), as we emphasized before, is a moment condition used to showcase our out-of-sample uncertainty quantification strategy and can be relaxed by utilizing other appropriate concentration inequalities. Finally, recall that the constant $\epsilon_{\Delta}(\tau)$ used to adjust for nonlinear constraints is included in the in-sample error bounds $\underline{M}_{\text{in}}(\tau)$ and $\bar{M}_{\text{in}}(\tau)$, and a general strategy for choosing $\epsilon_{\Delta}(\tau)$ is discussed in Section 6.3.

6.1 Discussion of Conditions (i)–(iv)

In this section, we discuss the justification of the high-level conditions (i)–(iv) in more detail.

- **Condition (i).** This condition formalizes the idea of distributional approximation of $\hat{\gamma} - \gamma$ by a Gaussian vector \mathbf{G} . It can be verified under different primitive conditions, such as Assumption 1 that accommodates non-stationary data and is applicable to our empirical application. Lemma S.1 in the Supplemental Appendix provides a general way to verify condition (i) by assuming the pseudo-true residuals in \mathbf{U} are independent over t conditional on \mathcal{H} . In fact, (i) also holds when the errors are only weakly dependent (e.g., β -mixing) conditional on \mathcal{H} .
- **Condition (ii).** This is a mild condition on the concentration of $\boldsymbol{\delta} \in \mathcal{M}_{\mathbf{G}}$. The requirement $\boldsymbol{\delta}'\hat{\mathbf{Q}}\boldsymbol{\delta} - 2\mathbf{G}'\boldsymbol{\delta} \leq 0$ is usually known as the *basic inequality* in regression analysis; see, for example, [Wainwright \(2019, Chapter 7\)](#) for the case of lasso. The vector \mathbf{G} is (conditional) Gaussian by construction, making condition (ii) easy to verify based on well-known bounds for Gaussian distributions. This condition holds in a variety of empirically relevant settings, including

outcomes-only regression with i.i.d. data, multi-equation regression with weakly dependent data, and cointegrated outcomes and feature settings.

- **Condition (iii).** This is a high-level requirement on the “closeness” between Δ and Δ^* . We generalize the strategy for constructing Δ^* described in Section 4.1 to this more flexible setup, which can be shown to satisfy condition (iii) if the constraints specified in \mathcal{W} and \mathcal{R} are formed by smooth functions. Suppose that

$$\mathcal{W} \times \mathcal{R} = \left\{ \boldsymbol{\beta} \in \mathbb{R}^d : \mathbf{m}_=(\boldsymbol{\beta}) = \mathbf{0}, \mathbf{m}_\leq(\boldsymbol{\beta}) \leq \mathbf{0} \right\},$$

where $\mathbf{m}_=(\cdot) \in \mathbb{R}^{d_=}$ and $\mathbf{m}_\leq(\cdot) \in \mathbb{R}^{d_\leq}$ and $d_=$ and d_\leq denote the number of equality and inequality constraints in $\mathcal{W} \times \mathcal{R}$, respectively. Let the ℓ -th constraint in $\mathbf{m}_\leq(\cdot)$ be $m_{\leq,\ell}(\cdot)$. Given tuning parameters $\varrho_\ell > 0$, $\ell = 1, \dots, d_\leq$, let $\mathcal{A} = \{\ell_1, \dots, \ell_k\}$ denote the set of indices for the inequality constraints such that $m_{\text{in},\ell}(\widehat{\boldsymbol{\beta}}) > -\varrho_\ell$. We define

$$\Delta^* = \left\{ \boldsymbol{\beta} - \widehat{\boldsymbol{\beta}} : \mathbf{m}_=(\boldsymbol{\beta}) = \mathbf{0}, m_{\leq,\ell}(\boldsymbol{\beta}) \leq m_{\leq,\ell}(\widehat{\boldsymbol{\beta}}) \mathbf{1}(\ell \in \mathcal{A}), \ell = 1, \dots, d_\leq \right\}. \quad (6.7)$$

The following lemma verifies condition (iii) for this Δ^* .

Lemma 1. *Let $\|\cdot\|$ be the Euclidean norm for vectors and the spectral norm for matrices. Assume that with probability over \mathcal{H} at least $1 - \pi_\Delta^*$, the following conditions hold: (i) $\mathbb{P}(\|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0\| \leq \varpi_\delta^* | \mathcal{H}) \geq 1 - \epsilon_\Delta^*$; (ii) $\mathbf{m}(\cdot) = (\mathbf{m}_=(\cdot)', \mathbf{m}_\leq(\cdot)')'$ is twice continuously differentiable on $\mathcal{B}(\boldsymbol{\beta}_0, \varpi_\delta^*)$ with $\inf_{\boldsymbol{\beta} \in \mathcal{B}(\boldsymbol{\beta}_0, \varpi_\delta^*)} s_{\min}(\frac{\partial}{\partial \boldsymbol{\beta}} \mathbf{m}(\boldsymbol{\beta})) \geq c_{\min}$ for some constant $c_{\min} > 0$; and (iii) for all $1 \leq \ell \leq d_\leq$ and some $\mathbf{c} > 0$ specified in the proof, $\varrho_\ell > \mathbf{c} \varpi_\delta^*$ and $\varrho_\ell < |m_{\leq,\ell}(\boldsymbol{\beta}_0)| - \mathbf{c} \varpi_\delta^*$ if $m_{\leq,\ell}(\boldsymbol{\beta}_0) \neq 0$. Then, for Δ^* defined in (6.7), condition (iii) in Theorem 1 holds with $\varpi_\Delta^* = \mathfrak{C}(\varpi_\delta^*)^2$ for some constant $\mathfrak{C} > 0$.*

In this lemma, the tuning parameters ϱ_ℓ are introduced to guarantee that with high probability, we can correctly differentiate the binding inequality constraints from the other non-binding ones. Section 6.2 below provides more practical details about choosing ϱ_ℓ . Also, the concentration requirement for $\widehat{\boldsymbol{\beta}}$ specified in this lemma is mild. Since $\widehat{\boldsymbol{\beta}}$ satisfies the basic inequality $(\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0)' \widehat{\mathbf{Q}}(\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0) - 2(\widehat{\boldsymbol{\gamma}} - \boldsymbol{\gamma})'(\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0) \leq 0$, the concentration of $\widehat{\boldsymbol{\beta}}$ can be shown by combining

a distributional approximation of $\hat{\gamma} - \gamma$ by a Gaussian vector \mathbf{G} and the idea outlined in the previous discussion about condition (ii).

- **Condition (iv).** This is a requirement that $\hat{\Sigma}$ be a “good” approximation of the unknown covariance matrix Σ . Many standard covariance estimation strategies such as the family of well-known heteroskedasticity-consistent estimators can be utilized.

6.2 Defining Constraints in Simulation

Section 6.1 proposes a procedure for constructing the constraint set Δ^* in simulation, which relies on a sequence of tuning parameters ϱ_ℓ , $\ell = 1, \dots, d_\leq$, to determine which inequality constraints are binding. Now, we propose a feasible strategy to select ϱ_ℓ .

Suppose that each constraint imposed in the analysis only restricts the SC weights associated with a particular treated unit (so there is no cross-treated-unit restriction), which is the most common case in SC applications. Consider a generic inequality constraint $m_{\leq, \ell}(\boldsymbol{\beta}^{[i]}) \leq 0$ associated with the treated unit i . Our goal is to use ϱ_ℓ to check if this inequality constraint is binding. Using the Taylor expansion, if $m_{\leq, \ell}(\boldsymbol{\beta}_0^{[i]}) = 0$, then $m_{\leq, \ell}(\hat{\boldsymbol{\beta}}^{[i]}) \approx \frac{\partial}{\partial \boldsymbol{\beta}'} m_{\leq, \ell}(\boldsymbol{\beta}_0^{[i]}) (\hat{\boldsymbol{\beta}}^{[i]} - \boldsymbol{\beta}_0^{[i]})$. We use this formula to define a range within which $m_{\leq, \ell}(\hat{\boldsymbol{\beta}}^{[i]})$ deviates from $m_{\leq, \ell}(\boldsymbol{\beta}_0^{[i]})$, with high probability.

Note that the basic inequality $\boldsymbol{\delta}' \hat{\mathbf{Q}} \boldsymbol{\delta} - 2\mathbf{G}' \boldsymbol{\delta} \leq 0$ in condition (ii) of Theorem 1, along with the Gaussian approximation of $\hat{\gamma}$ in condition (i), implies a bound on the deviation of $\hat{\boldsymbol{\beta}}^{[i]}$ from $\boldsymbol{\beta}_0^{[i]}$: with high probability, $\|\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0\| \leq 2\|\mathbf{G}\|/s_{\min}(\hat{\mathbf{Q}})$. Motivated by this fact, we propose to use $\varrho^{[i]} = \mathcal{C}(\log T_0)^c/T_0^{1/2}$ as a (data-dependent) bound on $\|\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0\|$, where $c = 1/2$ if the data are i.i.d. or weakly dependent, and $c = 1$ if $\mathbf{A}^{[i]}$ and $\mathbf{B}^{[i]}$ form a cointegrated system, and \mathcal{C} is one of the following

$$\mathcal{C}_1 = \frac{\hat{\sigma}_u}{\min_{1 \leq j \leq J_0} \hat{\sigma}_{b_j}}, \quad \mathcal{C}_2 = \frac{\max_{1 \leq j \leq J_0} \hat{\sigma}_{b_j} \hat{\sigma}_u}{\min_{1 \leq j \leq J_0} \hat{\sigma}_{b_j}^2}, \quad \mathcal{C}_3 = \frac{\max_{1 \leq j \leq J_0} \hat{\sigma}_{b_j u}}{\min_{1 \leq j \leq J_0} \hat{\sigma}_{b_j}^2}, \quad (6.8)$$

where $\hat{\sigma}_{b_j, u}$ is the estimated (unconditional) covariance between the pseudo-true residual $\mathbf{u}^{[i]}$ and the j -th column of $\mathbf{B}^{[i]}$ (the features of the j -th control unit), and $\hat{\sigma}_u$ and $\hat{\sigma}_{b_j}$ are the estimated (unconditional) standard deviation of $\mathbf{u}^{[i]}$ and the j -th column of $\mathbf{B}^{[i]}$, respectively. If the synthetic control weights were constructed based on both stationary and non-stationary features, the non-

stationary components would govern the precision of the estimation. In such cases, one could ignore the stationary components and set $c = 1$. Then, an intuitive choice of ϱ_ℓ would be

$$\varrho_\ell = \left\| \frac{\partial}{\partial \boldsymbol{\beta}} m_{\leq, \ell}(\widehat{\boldsymbol{\beta}}^{[i]}) \right\|_1 \times \varrho^{[i]}. \quad (6.9)$$

If $m_{\leq, \ell}(\widehat{\boldsymbol{\beta}}^{[i]}) > -\varrho_\ell$, we let the ℓ -th constraint be binding in the simulation.

6.3 Adjustment for Nonlinear Constraints

When some constraints in Δ are nonlinear (e.g., ridge-type constraints), we introduce a constant $\varepsilon_\Delta(\tau)$ to adjust the bounds on the in-sample error; this constant depends on the distance ϖ_Δ^* between the localized constraint sets Δ and Δ^* specified in condition (iii) of Theorem 1. This adjustment is only necessary for nonlinear constraints; it is not needed when the constraints are linear in parameters (e.g., simplex or lasso constraints).

The distance between Δ and Δ^* typically depends on the first and second derivatives of the constraint functions $\mathbf{m}_{\leq}(\cdot)$. Again, we focus on the case in which there is no cross-treated-unit constraint. Let $\mathcal{S}^{[i]}$ be the set of indices for inequality constraints related to a generic treated unit i . Denote by $\mathbf{m}_{\leq, \mathcal{S}^{[i]}}(\cdot)$ the vector of constraint functions $m_{\leq, \ell}(\cdot)$ with $\ell \in \mathcal{S}^{[i]}$. We propose to set

$$\varepsilon_\Delta^{[i]}(\tau) = \|\mathbf{p}_\tau^{[i]}\|_1 \times \frac{\sqrt{|\mathcal{S}^{[i]}|}}{2} s_{\min}^{-1} \left(\frac{\partial}{\partial \boldsymbol{\beta}} \mathbf{m}_{\leq, \mathcal{S}^{[i]}}(\widehat{\boldsymbol{\beta}}^{[i]}) \right) \times \max_{\ell \in \mathcal{S}^{[i]}} s_{\max} \left(\frac{\partial}{\partial \boldsymbol{\beta} \partial \boldsymbol{\beta}'} m_{\leq, \ell}(\widehat{\boldsymbol{\beta}}^{[i]}) \right) \times (\varrho^{[i]})^2,$$

where $\mathbf{p}_\tau^{[i]}$ denotes the subvector of \mathbf{p}_τ that corresponds to $\boldsymbol{\beta}_0^{[i]}$. Denote by \mathcal{E}_τ the set of treated units to which the causal predictand τ is related. Then set

$$\varepsilon_\Delta(\tau) = \sum_{i \in \mathcal{E}_\tau} \varepsilon_\Delta^{[i]}(\tau).$$

When constraints are linear, the second derivative of $m_{\leq, \ell}(\cdot)$ is zero, and the above choice of $\varepsilon_\Delta(\tau)$ is exactly 0. In the special case of L2 constraints used in our empirical application, the above choice of $\varepsilon_\Delta^{[i]}(\tau)$ simplifies to $\varepsilon_\Delta^{[i]}(\tau) = \|\mathbf{p}_\tau^{[i]}\|_1 \times (2\|\widehat{\boldsymbol{\beta}}^{[i]}\|_2)^{-1} \times (\varrho^{[i]})^2$.

7 Conclusion

We developed prediction intervals to quantify the uncertainty of a large class of synthetic control predictions (or estimators) in settings with staggered treatment adoption. Because many synthetic control applications have a limited number of observations, our inference procedures are based on non-asymptotic concentration arguments. The construction of our prediction intervals is designed to capture two sources of uncertainty: the first is the construction or estimation of the synthetic control weights with pre-treatment data, and the second is the variability of the post-treatment outcomes. By combining both sources in a prediction interval, our procedure offers precise non-asymptotic coverage probability guarantees and allows researchers to implement sensitivity analyses to assess how robust the conclusions of the analysis are to various levels of uncertainty. Our framework is general, allowing for one or multiple treated units, simultaneous or staggered treatment adoption, linear or non-linear constraints, and stationary or non-stationary data. To enhance implementation, we also showed how to recast the methods as conic optimization programs and how to choose the necessary tuning parameters in a principled data-driven way. We illustrated our methods with an empirical application studying the effect of economic liberalization on real GDP per capita in Sub-Saharan African countries, motivated by the work of [Billmeier and Nannicini \(2013\)](#).

All our methods are implemented in Python, R, and Stata software, which is publicly available (<https://nppackages.github.io/scpi/>), and discussed in detail in our companion article [Cattaneo, Feng, Palomba and Titiunik \(2024\)](#) and in Section S.6 of the Supplemental Appendix.

References

- Abadie, A. (2021), “Using Synthetic Controls: Feasibility, Data Requirements, and Methodological Aspects,” *Journal of Economic Literature*, 59, 391–425.
- Abadie, A., and Cattaneo, M. D. (2018), “Econometric Methods for Program Evaluation,” *Annual Review of Economics*, 10, 465–503.
- Abadie, A., Diamond, A., and Hainmueller, J. (2010), “Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California’s Tobacco Control Program,” *Journal of the American Statistical Association*, 105, 493–505.

- Abadie, A., and Gardeazabal, J. (2003), “The Economic Costs of Conflict: A Case Study of the Basque Country,” *American Economic Review*, 93, 113–132.
- Agarwal, A., Shah, D., Shen, D., and Song, D. (2021), “On Robustness of Principal Component Regression,” *Journal of the American Statistical Association*, 116, 1731–1745.
- Ben-Michael, E., Feller, A., and Rothstein, J. (2022), “Synthetic Controls with Staggered Adoption,” *Journal of the Royal Statistical Society, Series B*, 84, 351–381.
- Bhagwati, J., and Srinivasan, T. N. (2001), “Outward-Orientation and Development: Are Revisionists Right?” in *Trade, Development and Political Economy*, Springer, pp. 3–26.
- Billmeier, A., and Nannicini, T. (2013), “Assessing Economic Liberalization Episodes: A Synthetic Control Approach,” *Review of Economics and Statistics*, 95, 983–1001.
- Boyd, S., and Vandenberghe, L. (2004), *Convex Optimization*, Cambridge University Press.
- Bratton, M., and Van de Walle, N. (1997), *Democratic experiments in Africa: Regime transitions in comparative perspective*, Cambridge university press.
- Callaghy, T. M., and Ravenhill, J. (1994), *Hemmed In: Responses to Africa’s Economic Decline*, Columbia University Press.
- Cattaneo, M. D., Feng, Y., Palomba, F., and Titiunik, R. (2024), “`scpi`: Uncertainty Quantification for Synthetic Control Methods,” *Journal of Statistical Software*.
- Cattaneo, M. D., Feng, Y., and Titiunik, R. (2021), “Prediction Intervals for Synthetic Control Methods,” *Journal of the American Statistical Association*, 116, 1865–1880.
- Chernozhukov, V., Wüthrich, K., and Zhu, Y. (2021a), “Distributional Conformal Prediction,” *Proceedings of the National Academy of Sciences*, 118, e2107794118.
- Chernozhukov, V., Wüthrich, K., and Zhu, Y. (2021b), “An Exact and Robust Conformal Inference Method for Counterfactual and Synthetic Controls,” *Journal of the American Statistical Association*, 116, 1849–1864.
- DeJong, D. N., and Ripoll, M. (2006), “Tariffs and Growth: An Empirical Exploration of Contingent Relationships,” *Review of Economics and Statistics*, 88, 625–640.

- Freyaldenhoven, S., Hansen, C., Pérez, J. P., and Shapiro, J. M. (Forthcoming), “Visualization, identification, and estimation in the linear panel event-study design,” *Advances in Economics and Econometrics: Twelfth World Congress*.
- Freyaldenhoven, S., Hansen, C., and Shapiro, J. M. (2019), “Pre-event trends in the panel event-study design,” *American Economic Review*, 109, 3307–3338.
- Giavazzi, F., and Tabellini, G. (2005), “Economic and Political Liberalizations,” *Journal of Monetary Economics*, 52, 1297–1330.
- Levine, R., and Renelt, D. (1992), “A Sensitivity Analysis of Cross-Country Growth Regressions,” *American Economic Review*, 82, 942–963.
- Li, K. T. (2020), “Statistical Inference for Average Treatment Effects Estimated by Synthetic Control Methods,” *Journal of the American Statistical Association*, 115, 2068–2083.
- Masini, R., and Medeiros, M. C. (2021), “Counterfactual Analysis with Artificial Controls: Inference, High Dimensions and Nonstationarity,” *Journal of the American Statistical Association*, 116, 1773–1788.
- Miller, D. L. (2023), “An introductory guide to event study models,” *Journal of Economic Perspectives*, 37, 203–230.
- Powell, D. (2022), “Synthetic Control Estimation Beyond Comparative Case Studies: Does the Minimum Wage Reduce Employment?” *Journal of Business & Economic Statistics*, 40, 1302–1314.
- Sachs, J. D., Warner, A., Åslund, A., and Fischer, S. (1995), “Economic Reform and the Process of Global Integration,” *Brookings Papers on Economic Activity*, 1995, 1–118.
- Shaikh, A. M., and Toulis, P. (2021), “Randomization Tests in Observational Studies with Staggered Adoption of Treatment,” *Journal of the American Statistical Association*, 116, 1835–1848.
- Shi, X., Miao, W., Hu, M., and Tchetgen, E. T. (2023), “Theory for Identification and Inference with Synthetic Controls: A Proximal Causal Inference Framework,” arXiv:2108.13935.

- Vovk, V. (2012), “Conditional Validity of Inductive Conformal Predictors,” in *JMLR: Workshop and Conference Proceedings 25, Asian Conference on Machine Learning*, pp. 475–490.
- Wacziarg, R., and Welch, K. H. (2008), “Trade Liberalization and Growth: New Evidence,” *World Bank Economic Review*, 22, 187–231.
- Wainwright, M. J. (2019), *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*, Cambridge University Press.