

Preprint

SYNTHETIC CONTROL AS ONLINE LINEAR REGRESSION

JIAFENG CHEN

Harvard Business School and Department of Economics, Harvard University

This paper notes a simple connection between synthetic control and online learning. Specifically, we recognize synthetic control as an instance of *Follow-The-Leader* (FTL). Standard results in online convex optimization then imply that, even when outcomes are chosen by an adversary, synthetic control predictions of counterfactual outcomes for the treated unit perform almost as well as an oracle weighted average of control units' outcomes. Synthetic control on differenced data performs almost as well as oracle weighted difference-in-differences, potentially making it an attractive choice in practice. We argue that this observation further supports the use of synthetic control estimators in comparative case studies.

KEYWORDS: Synthetic control, Online convex optimization, Difference-in-differences, Regret.

1. INTRODUCTION

Synthetic control ([Abadie and Gardeazabal, 2003](#), [Abadie et al., 2015](#)) is an increasingly popular method for causal inference among policymakers, private institutions, and social scientists alike. In parallel, there is a rapidly growing methodological literature providing statistical guarantees for synthetic control methods.¹ Existing results for synthetic control—and for modifications thereof—are typically derived under a low-rank linear factor model

Jiafeng Chen: jiafengchen@g.harvard.edu

This paper would not have been possible without the generous encouragement and guidance of Isaiah Andrews. I also thank the editor, two anonymous referees, Alberto Abadie, Susan Athey, Xiaohong Chen, Avi Feller, Bruno Ferman, Wayne Gao, Edward Glaeser, Sukjin Han, Patrick Kline, Scott Kominers, Apoorva Lal, Namrata Narain, Ross Mattheis, David Ritzwoller, Brad Ross, Jonathan Roth, Bas Sanders, Suproteem Sarkar, Jesse Shapiro, Neil Shephard, Jann Spiess, Elie Tamer, Jaume Vives-i-Bastida, Davide Viviano, Chris Walker, and Jennifer Walsh for helpful discussions. I am also grateful for comments from participants of the Harvard Graduate Student Workshop in Econometrics, the 2022 Synthetic Control Workshop at the Princeton University Center for Statistics and Machine Learning, and the 2022 NBER Labor Studies Summer Institute Meeting.

¹See the review by [Abadie \(2021\)](#) as well as the special section on synthetic control methods in the *Journal of the American Statistical Association* ([Abadie and Cattaneo, 2021](#)).

or a vector autoregressive model of the outcomes (see, among others, [Abadie et al., 2010](#), [Ben-Michael et al., 2019, 2021](#), [Ferman and Pinto, 2021](#), [Viviano and Bradic, 2019](#)).² While these statistical guarantees formally hold under these outcome models, a number of authors have expressed optimism that the synthetic control method is robust to these modeling assumptions.³

On the other hand, in empirical settings where synthetic control is commonly applied—where the treated unit is an aggregate entity like a country or a U.S. state—plausible outcome modeling may be challenging. [Manski and Pepper \(2018\)](#), in studying the effect of gun laws in the United States using state-level crime rates, provocatively ask, “what random process should be assumed to have generated the existing United States, with its realized state-year crime rates?” Granted, the low-rank linear factor model is a general class of data-generating processes and may even arise under finer-grained models on the individual outcomes contained in the aggregate data ([Shi et al., 2022](#)). But to pessimists and skeptics, perhaps even such a model is implausible for the settings considered by many synthetic control studies. Indeed, if practitioners were willing to fully commit to an outcome model, perhaps they should estimate the outcome model directly—e.g., use factor model-based methods ([Bai and Ng, 2002](#), [Bai, 2003](#), [Xu, 2017](#), [Athey et al., 2021](#))—instead of using synthetic control?

As a result, existing methodological results seem to leave practitioners in a somewhat awkward position. On the one hand, synthetic control is intuitively appealing, and it is conjectured to have good properties under a variety of outcome models. On the other hand, perhaps existing outcome models that have so far proved sufficiently analytically tractable are not always compelling in common empirical settings. To address this tension, this paper

²Notably, like this paper, [Bottmer et al. \(2021\)](#) consider a design-based framework which conditions on the outcomes and considers randomness arising solely from assignment of the treated unit or the treatment time period.

³For instance, [Ben-Michael et al. \(2019\)](#) write, “Outcome modeling can also be sensitive to model misspecification, such as selecting an incorrect number of factors in a factor model. Finally, [... synthetic control] can be appropriate under multiple data generating processes (e.g., both the autoregressive model and the linear factor model) so that it is not necessary for the applied researcher to take a strong stand on which is correct.” [Abadie and Vives-i-Bastida \(2021\)](#) write, “Synthetic controls are intuitive, transparent, and produce reliable estimates for a variety of data generating processes.”

provides a few theoretical results and offers a novel interpretation of synthetic control methods. In particular, we seek guarantees for synthetic control that do not rely on any outcome model. Consequently, our results complement existing, model-based ones.

It is unlikely that nontrivial guarantees on the performance of synthetic control exist without any structure on the outcomes. However, we can derive guarantees of synthetic control’s performance *relative* to a class of alternatives, such as weighted matching or weighted difference-in-differences (DID) estimators, which practitioners may otherwise choose. Our first main result shows that, on average over *hypothetical treatment timings*, synthetic control predictions are never much worse than the predictions made by any weighted matching estimator. Our second main result shows that the same is true for synthetic control on differenced data versus any weighted DID estimator. These results imply that if there is a weighted matching or DID estimator that performs well, synthetic control likewise performs well. To be clear, these *regret guarantees* average over hypothetical treatment timings, which can be interpreted as expected loss under random treatment timing, a design-based assumption.

Taken together, our results provide reassurances for practitioners, as they offer justifications for synthetic control that do not rely on particular statistical models of the outcomes. At least on average over hypothetical treatment timings, regardless of outcomes, variations of synthetic control are competitive against common estimators, such as weighted matching and weighted DID estimators. Additionally, our second result introduces a novel version of synthetic control that is competitive against DID. Since DID is extremely popular in practice (Currie et al., 2020) and is thus a natural benchmark, this version of synthetic control may be particularly attractive.

We derive our results by casting prediction with panel data as an instance of *online convex optimization*, and by recognizing synthetic control as an online regression algorithm known as *Follow-The-Leader* (FTL, a name coined by Kalai and Vempala, 2005).⁴ Regret guarantees on FTL in the online convex optimization literature translate directly to guarantees

⁴For an introduction to online convex optimization, see Hazan (2019), Orabona (2019), Cesa-Bianchi and Lugosi (2006), and Shalev-Shwartz (2011).

for synthetic control against a class of alternative estimators. Since most results in online convex optimization have been derived under an adversarial model—where an imagined adversary generates the data—these results translate to guarantees on synthetic control without any structure on the outcome process.

This paper is perhaps closest to [Viviano and Bradic \(2019\)](#). They propose an ensemble scheme to aggregate predictions from multiple predictive models, which can include synthetic control, interactive fixed effects models, and random forests. Using results from the online learning literature, [Viviano and Bradic \(2019\)](#)’s ensemble scheme has the no-regret property, making the ensemble predictions competitive against the predictions of any fixed predictive model in the ensemble. Under sampling processes that yield good performance for some predictive model in the ensemble, [Viviano and Bradic \(2019\)](#) then derive performance guarantees for the ensemble learner. In contrast, we study synthetic control directly in the worst-case setting, and connect corresponding worst-case results to guarantees on statistical risk in a design-based framework. We show that synthetic control algorithms *themselves* are no-regret online algorithms and are in fact competitive against a wide class of matching or DID estimators.

[Section 2](#) sets up the notation and the decision protocol and presents our main results for synthetic control. [Section 3](#) presents several extensions that show alternative guarantees on modifications of synthetic control; in particular, we show that synthetic control on differenced data is competitive against a class of difference-in-differences estimators. [Section 4](#) concludes the paper.

2. SETUP AND MAIN RESULTS

Consider a simple setup for synthetic control, following [Doudchenko and Imbens \(2016\)](#). There are T time periods and $N + 1$ units. To simplify convergence rate expressions, we assume $T > N$ unless noted otherwise, but this assumption is not strictly necessary for our results. Let unit 0 be the only treated unit, first treated at some time $S \in \{1, \dots, T\} \equiv [T]$. The other N units are referred to as control units. Since we observe the treated potential outcomes for the treated unit after S , estimating causal effects for unit 0 amounts to

predicting the unobserved, post- S untreated potential outcomes of this unit. Thus, we focus on untreated potential outcomes.

Let the full panel of untreated potential outcomes be \mathbf{Y} with representative entry y_{it} , where (i) $\mathbf{Y}_{1:s} = (y_{0t}, \dots, y_{Nt})_{t=1}^s$ collects all untreated potential outcomes until and including time s , and (ii) $\mathbf{y}_t = (y_{1t}, \dots, y_{Nt})'$ is the vector of control unit outcomes at time t . Additionally, we let $\mathbf{y}(1) = (y_1(1), \dots, y_T(1))'$ denote the treated potential outcomes of unit 0, which are only observable for times $t \geq S$. Similarly, we let $\mathbf{y}(0) = (y_{01}, \dots, y_{0T})'$ denote the untreated potential outcomes of unit 0, which are observable for $t < S$. The analyst is tasked with predicting y_{0S} from observed data, which typically consist of pre-treatment outcomes of unit 0 and outcomes of untreated units. Like the main analysis in [Doudchenko and Imbens \(2016\)](#), we do not consider covariates extensively, though [Section 3.3](#) considers matching on covariates as a form of regularization.⁵

Synthetic control ([Abadie and Gardeazabal, 2003](#), [Abadie et al., 2010](#)), in its basic form, chooses some convex weights $\hat{\theta}_S$ that minimize past prediction errors

$$\hat{\theta}_S \in \arg \min_{\theta \in \Theta} \sum_{t=1}^{S-1} (y_{0t} - \theta' \mathbf{y}_t)^2, \quad (1)$$

where $\Theta \equiv \{(\theta_1, \dots, \theta_N) \in \mathbb{R}^N : \theta_i \geq 0, 1' \theta = 1\}$ is the simplex. For a one-step-ahead forecast for y_{0S} , synthetic control outputs the weighted average $\hat{y}_S \equiv \hat{\theta}'_S \mathbf{y}_S$, and forms the treatment effect estimate $\hat{\tau}_S \equiv y_S(1) - \hat{y}_S$.

Theoretical guarantees for treatment effect estimates $\hat{\tau}_S$ often rely on statistical models of the outcomes \mathbf{Y} . While synthetic control has good performance under a range of outcome models, one may still doubt whether these models are plausible—and whether the underlying repeated sampling thought experiments are appropriate—in the spirit of comments by [Manski and Pepper \(2018\)](#). In contrast to the usual outcome modeling approach, we instead

⁵To extend our analysis to cases with covariates, at a minimum, we can interpret \mathbf{Y} as the residuals of the untreated potential outcomes against some fixed regression function of the covariates, i.e. $y_{it} = y_{it}^* - h_t(x_i)$, for fixed h_t (perhaps estimated from auxiliary data), outcomes y_{it}^* , and covariate vectors x_i . The residualization is similar to [Section 5.5 in Doudchenko and Imbens \(2016\)](#) and expression (16) in [Abadie \(2021\)](#), but is stronger due to h_t being fixed for different adversarial choices of \mathbf{Y} . Our results apply so long as these residuals obey the boundedness assumption $\|\mathbf{Y}\|_\infty \leq 1$ that we impose later.

consider a worst-case setting where the outcomes are generated by an adversary.⁶ Doing so has the appeal of giving decision-theoretic justification for methods while being entirely agnostic towards the data-generating process. Since a dizzying range of reasonable data-generating models and identifying assumptions are possible in panel data settings—yet perhaps none are unquestionably realistic—this worst-case view is valuable, and worst-case guarantees can be comforting.

In particular, we assume an adversary picks the outcomes \mathbf{Y} —or, equivalently, we derive results that hold uniformly over $\{\mathbf{Y} : \|\mathbf{Y}\|_\infty \leq 1\}$. Specifically, we consider the following protocol between an analyst and an adversary:

(P1) The analyst commits to a class of linear prediction rules $\hat{y}_t \equiv f(\mathbf{y}_t; \theta_t(\mathbf{Y}_{1:t-1})) = \theta_t' \mathbf{y}_t$, parametrized by some $\theta_t \in \Theta$ that may be chosen as a function of the past data $\mathbf{Y}_{1:t-1}$. We refer to the maps $\sigma \equiv \{\theta_t(\cdot) : t \in [T]\}$ as the analyst’s *strategy*. This means that if the treatment time S is equal to t , then the analyst reports \hat{y}_t as their prediction for the untreated potential outcome at the first period after treatment.

(P2) The adversary chooses the matrix of outcomes \mathbf{Y} . In order to obtain nontrivial bounds, we assume that the adversary cannot choose arbitrarily large outcomes, and without further loss of generality, we assume $\|\mathbf{Y}\|_\infty \leq 1$. Since we are interested in the worst case, the adversary may choose \mathbf{Y} with knowledge of σ .

(P3) The analyst suffers loss equal to squared prediction error at time S : i.e., $\ell(\hat{y}_S, y_{0S}) \equiv (\hat{y}_S - y_{0S})^2$.

Under such a protocol, the analyst’s average squared loss, averaging over *hypothetical values of S* , is

$$\frac{1}{T} \sum_{S=1}^T (y_{0S} - \hat{y}_S)^2 = \frac{1}{T} \sum_{S=1}^T (y_{0S} - \theta_S' \mathbf{y}_S)^2 = \mathbb{E}_{S \sim \text{Unif}[T]} [(y_{0S} - \hat{y}_S)^2]. \quad (2)$$

Most results in this paper are guarantees in terms of the decision criterion (2) for synthetic control, where synthetic control (1) is viewed as a particular strategy σ under (P1) to (P3).

⁶The adversarial framework, popular in online learning, dates to the works of [Hannan \(1958\)](#) and [Blackwell \(1956\)](#).

As the second equality in (2) indicates, under an additional assumption that treatment timing is *uniformly random*, $S \sim \text{Unif}[T]$, the average loss over hypothetical treatment timings is equal to the expected squared loss over S . This additional assumption is a design-based perspective (Doudchenko and Imbens, 2016, Bottmer et al., 2021) on the panel causal inference problem. This perspective enables us to interpret average prediction loss over hypothetical treatment timings as expected prediction loss under the random treatment time S . The latter can in turn be thought of as design-based risk. Uniformly random assignment of S is restrictive, but we shall relax this requirement in Section 3.1 and Appendix B.1.⁷

We now make clear the connection with online convex optimization (see Section 1.1 in Hazan, 2019). Online convex optimization works with the following general protocol. Time t increments sequentially for T periods, and at time t :

(O1) An online player chooses some $\theta_t \in \Theta$, where $\Theta \subset \mathbb{R}^d$ is a bounded convex set. The choice θ_t may depend on the loss functions $\{\ell_s : s < t\}$ chosen by the adversary in the past.

(O2) After θ_t is chosen, an adversary chooses a loss function $\ell_t : \Theta \rightarrow \mathbb{R}$ from some given set of loss functions, which may be further parametrized. These loss functions are constrained to be convex and bounded but can otherwise be quite general. They are often further constrained in order to obtain specific regret results.

(O3) The player suffers loss $\ell_t(\theta_t)$ and observes $\ell_t(\cdot)$.⁸ The player may update their decision θ_{t+1} based on $\ell_1(\cdot), \dots, \ell_t(\cdot)$.

At the end of the game, the online player suffers total loss $\sum_{t=1}^T \ell_t(\theta_t)$.

Our setup of the panel prediction protocol, (P1) to (P3), is then an instance of online convex optimization, (O1) to (O3). To see this, the most important step is to recognize that the analyst's loss (2) is analogous to the online player's loss, and therefore to think of the analyst as making sequential decisions where \mathbf{Y} is sequentially revealed to them. This change in perspective relies on (i) our choice of decision criterion (2) and (ii) the fact

⁷The protocol (P1) to (P3) easily generalizes when we replace $f(\mathbf{y}_t, \theta_t)$ with any known scalar function and $\ell(\cdot, \cdot)$ with any loss function, so long as $\theta \mapsto \ell(f(\mathbf{y}_t, \theta), y_{0t})$ is convex and bounded. Our results in Section 3.3 allow for general loss functions.

⁸A closely related setting where the player only observes $\ell_t(\theta_t)$ instead of the entire loss function $\ell_t(\cdot)$ is known as *bandit convex optimization* (see Chapter 6 in Hazan, 2019), of which the adversarial multi-armed bandit problem (Robbins, 1952, Bubeck and Cesa-Bianchi, 2012) is a special case.

that the analyst’s decisions $\theta_t(\cdot)$ only require outcomes prior to t . Indeed, by fixing \mathbf{Y} and considering the hypothetical values of $S = 1, \dots, T$ sequentially, we can treat the analyst as if they were solving an online problem and learning from data in the past—even though, for any particular value of S , they are only confronted with a static, offline problem. To be clear, we are not considering some online version of synthetic control; the connection to online convex optimization comes from considering hypothetical, unrealized values of S .

After viewing the analyst’s problem as an online problem, we may straightforwardly establish the remaining correspondences. First, note that the simplex Θ is convex and bounded. Second, note that we may imagine the adversary in the panel prediction game as picking loss functions $\ell_t(\cdot)$ of the form $\theta \mapsto (y_{0t} - \theta' \mathbf{y}_t)^2$, parametrized by the potential outcomes (y_{0t}, \mathbf{y}_t) . These loss functions are indeed convex in θ and bounded, since both θ and \mathbf{Y} are bounded. Finally, note that the average loss (2) is equal to $\frac{1}{T} \sum_{t=1}^T \ell_t(\theta_t)$, which is simply the total loss in the online protocol scaled by $\frac{1}{T}$.⁹

Having recognized our setup as an instance of online convex optimization, the main observation of this paper recognizes that synthetic control is an online learning algorithm known as *Follow-the-Leader* (FTL). FTL, under (O1) to (O3), is the algorithm that, when prompted for a decision in (O1), simply chooses θ_t to minimize past losses:^{10,11}

$$\theta_t \in \arg \min_{\theta \in \Theta} \sum_{s < t} \ell_s(\theta).$$

⁹It may be tempting to ask whether the same argument applies to “horizontal regression” (Athey et al., 2021), where one regresses y_{iS} on y_{i1}, \dots, y_{iS-1} , perhaps constraining the coefficients to some bounded, convex set. Since synthetic control can be viewed as a “vertical regression,” where one regresses y_{0t} on y_{1t}, \dots, y_{Nt} , it seems we may apply our argument to the transposed \mathbf{Y} matrix. Indeed, we may formulate analogous claims by replacing t with i , s with j , S with some randomly chosen unit $M \in [N]$, and T with N . However, a difficulty with this interpretation is that synthetic control (1) naturally only uses information in the past ($t < S$), but the analogous restriction in horizontal regression, $i < M$, for a randomly chosen treated unit $M \in [N]$, is much less natural.

¹⁰FTL is also known as fictitious play in game theory (Brown, 1951). The name “follow-the-leader,” coined by Kalai and Vempala (2005), is popular in the recent computer science literature. For an introduction to FTL and similar algorithms, see Chapter 5 in Hazan (2019) and Chapters 1 and 7 in Orabona (2019).

¹¹When there are multiple minima, the choice of θ_t does not affect our theoretical guarantees. Nevertheless, it seems sensible in practice to take the minimum that is smallest in some norm, e.g. $\|\cdot\|_2$.

OBSERVATION 1: *Synthetic control (1) is an instance of FTL applied to the panel prediction protocol (P1) to (P3).*

Standard online convex optimization results on *regret* then apply to synthetic control as well. Before introducing these results, let us define regret as the gap between the total loss of a strategy σ and the best fixed weights θ in hindsight:

$$\text{Regret}_T(\sigma; \mathbf{Y}) \equiv \sum_{t=1}^T \ell_t(\theta_t) - \min_{\theta \in \Theta} \sum_{t=1}^T \ell_t(\theta) \quad (3)$$

$$= \sum_{S=1}^T (y_{0S} - \theta'_S \mathbf{y}_S)^2 - \min_{\theta \in \Theta} \sum_{S=1}^T (y_{0S} - \theta'_S \mathbf{y}_S)^2 \quad (4)$$

$$= T \left(\mathbb{E}_S[(y_{0S} - \theta'_S \mathbf{y}_S)^2] - \min_{\theta \in \Theta} \mathbb{E}_S[(y_{0S} - \theta'_S \mathbf{y}_S)^2] \right) \quad (5)$$

$$\geq T \left(\mathbb{E}_S[(y_{0S} - \theta'_S \mathbf{y}_S)^2] - \mathbb{E}_S[(y_{0S} - \theta'_S \mathbf{y}_S)^2] \right) \text{ for any } \theta \in \Theta. \quad (6)$$

(4) observes that, in our setting, regret is the difference between total squared prediction error of a strategy σ and that of the best fixed weights θ chosen in hindsight, summing over hypothetical treatment times S . (5) interprets the sum of losses as T times the expected loss under random treatment timing. Finally, (6) observes that regret is an upper bound of the expected error gap between the strategy σ and any fixed weights θ . We refer to $\arg \min_{\theta \in \Theta} \sum_{S=1}^T (y_{0S} - \theta'_S \mathbf{y}_S)^2$ as the *oracle weighted match*—the best set of weights for a given realization of the data \mathbf{Y} .

Focusing on regret rather than loss shifts the goalposts from performance to *competition*, which is a more fruitful perspective in our adversarial setting. After all, we cannot hope to obtain meaningful loss control as the all-powerful adversary can make the analyst miserable. However, the crucial insight of regret analysis is that, for certain strategies σ , the adversary cannot simultaneously make the analyst suffer high loss while letting some fixed strategy θ perform well—in other words, if any fixed θ performs well, then σ performs almost as well

over time. Indeed, if regret is sublinear, i.e., $\text{Regret}_T \leq o(T)$,¹² then the strategy σ never performs much worse than any fixed weights θ , on average over hypothetical treatment timing S . In this case, we can interpret σ as a strategy that is *competitive* against the class of weighted matching estimators.

It may seem surprising that these no-regret strategies σ exist in the first place. We emphasize that σ can output different weights θ_t , chosen adaptively over time, while σ is compared to an oracle that uses the best fixed weights. As a result, σ can compensate for its lack of oracle access by changing its choices judiciously over time.

The main result of this paper shows that the regret of synthetic control under quadratic loss is logarithmic in T . The result follows from a direct application of Hazan et al. (2007)'s regret bound for FTL (Theorem 5 in their paper, reproduced as [Theorem A.1](#) in the appendix).

THEOREM 2.1: *With bounded outcomes $\|\mathbf{Y}\|_\infty \leq 1$, synthetic control (1), denoted σ , satisfies the regret bound¹³*

$$\text{Regret}_T(\sigma, \mathbf{Y}) \leq 16N(\log(\sqrt{NT}) + 1) = O(N \log T).$$

[Theorem 2.1](#) shows that the synthetic control strategy (1) achieves logarithmic regret—and as a result, the average difference between the losses of synthetic control and losses of the oracle weighted match vanishes quickly as a function of T .¹⁴ In particular, if there exists

¹²We mean $\text{Regret}_T \leq o(T)$ in the sense that $\limsup_{T \rightarrow \infty} \frac{1}{T} \text{Regret}_T \leq 0$, since it is possible for Regret_T to be negative. Following the online convex optimization literature, we sometimes refer to σ as no-regret if it has sublinear regret.

¹³We say $f(N, T) = O(g(N, T))$ for $g(N, T) > 0$ if, for any sequence $N_T < T$ and $T \rightarrow \infty$,

$$\limsup_{T \rightarrow \infty} \frac{f(N_T, T)}{g(N_T, T)} < \infty.$$

In the conclusion of [Theorem 2.1](#), the inequality does not require $T > N$. The assumption $T > N$ is only used for the simplification $16N(\log(\sqrt{NT}) + 1) = O(N \log N + N \log T) = O(N \log T)$. Of course, the regret bound is less interesting if $\limsup N \log T/T > 0$.

¹⁴Restricting θ to the simplex Θ —a debated choice in the synthetic control literature—is somewhat important for the dependence on N , in so far as the simplex is bounded in $\|\cdot\|_1$. This is a consequence of the assumption that the outcomes \mathbf{Y} are bounded in the dual norm $\|\cdot\|_\infty$, which implies a bound on $\theta' \mathbf{y}_t$ that is free of N, T . In contrast, if we let $\Theta = \{\theta : \|\theta\|_2 \leq D/2\}$ be an ℓ_2 -ball, then the regret bound worsens to $O(D^2 N^2 \log(T))$.

a weighted average of the untreated units' outcomes that tracks $\mathbf{y}(0)$ well, then the average one-step-ahead loss of synthetic control estimates is only worse by $O\left(\frac{N \log T}{T}\right)$.

On its own, [Theorem 2.1](#) is purely an optimization result; we now offer a few comments on its statistical implications. As a preview, under random treatment timing, [Theorem 2.1](#) implies that the *risk* of estimating the causal effect at time S for synthetic control is not too much higher than that for any weighted matching estimator. Indeed, if any weighted matching estimator performs well, then synthetic control achieves low risk as well. Our discussion below translates [Theorem 2.1](#) into guarantees on the expected loss at treatment time—expressing regret as [\(5\)](#)—which relies on the design assumption that S is randomly assigned. Nevertheless, we stress that we could view [Theorem 2.1](#) purely as guarantees of average loss over hypothetical timings S —expressing regret only as [\(4\)](#)—which does not require a treatment timing assumption.

We can interpret regret as a gap in the design-based risk of estimating treatment effects. Specifically, we can interpret the expected loss of predicting the untreated outcome as the risk of estimating the treatment effect:

$$\begin{aligned} \text{Risk}(\sigma, \mathbf{Y}, \mathbf{y}(1)) &\equiv \mathbb{E}_S [(\tau_S - \hat{\tau}_S(\sigma))^2] \\ &\equiv \mathbb{E}_S [((y_S(1) - y_{0S}) - (y_S(1) - \hat{y}_S))^2] \\ &= \mathbb{E}_S [(y_{0S} - \hat{y}_S)^2]. \end{aligned} \tag{7}$$

Hence, [\(5\)](#) and [\(7\)](#), combined with [Theorem 2.1](#), imply that the risk of using synthetic control is no more than $N \log T / T$ worse than the risk of the oracle weighted match,¹⁵ regardless of the potential outcomes $\mathbf{Y}, \mathbf{y}(1)$:

$$\text{Risk}(\sigma, \mathbf{Y}, \mathbf{y}(1)) - \min_{\theta \in \Theta} \text{Risk}(\theta, \mathbf{Y}, \mathbf{y}(1)) = \frac{1}{T} \text{Regret}_T(\sigma, \mathbf{Y}) = O\left(\frac{N \log T}{T}\right). \tag{8}$$

This observation connects regret on prediction of the untreated potential outcome with differences in the risk of estimating treatment effects. Roughly speaking, [\(8\)](#) shows that

¹⁵We slightly abuse notation and use θ to denote the strategy that outputs θ every period.

synthetic control estimates of one-step-ahead causal effects are competitive against that of any fixed weighted match, for any realization of $\mathbf{Y}, \mathbf{y}(1)$, on average over S .

Of course, since the guarantee (8) holds for every \mathbf{Y} , it continues to hold when we average over \mathbf{Y} and $\mathbf{y}(1)$, over a joint distribution P that respects the boundedness condition $\|\mathbf{Y}\|_\infty \leq 1$. In this sense, analyzing regret in the adversarial framework not only does not preclude statistical interpretations, but rather facilitates analysis in a wide range of outcome models.¹⁶ Formally, let \mathcal{P} be a family of distributions for $\mathbf{Y}, \mathbf{y}(1)$ such that $P(\|\mathbf{Y}\|_\infty \leq 1) = 1$ for all $P \in \mathcal{P}$. Under an outcome model P , we may understand $\text{Risk}(\sigma, \mathbf{Y}, \mathbf{y}(1))$ as *conditional risk* and $\mathbb{E}_P \text{Risk}(\sigma, \mathbf{Y}, \mathbf{y}(1))$ as *unconditional risk*. Then, (8) implies that¹⁷

$$\sup_{P \in \mathcal{P}} \mathbb{E}_P \left[\text{Risk}(\sigma, \mathbf{Y}, \mathbf{y}(1)) - \min_{\theta \in \Theta} \text{Risk}(\theta, \mathbf{Y}, \mathbf{y}(1)) \right] = O\left(\frac{N \log T}{T}\right). \quad (9)$$

Therefore, the unconditional risk of synthetic control is never much worse than the risk of the oracle weighted match

$$R_\Theta^* \equiv \mathbb{E}_P \left[\min_{\theta \in \Theta} \text{Risk}(\theta, \mathbf{Y}, \mathbf{y}(1)) \right].$$

Hence, if the data-generating process P guarantees that R_Θ^* is small, then synthetic control achieves low expected risk as well. Concretely speaking, this latter requirement is that, for most realizations of the data, had we observed all the potential outcomes, we could find a

¹⁶The technique of “online-to-batch conversion” in the online learning literature exploits this intuition to prove results in batch (i.i.d.) settings via results in online adversarial settings.

¹⁷Abernethy et al. (2009) show that a minimax theorem applies, and

$$\sup_P \inf_\sigma \mathbb{E}_P \left[\text{Risk}(\sigma, \mathbf{Y}, \mathbf{y}(1)) - \min_{\theta \in \Theta} \text{Risk}(\theta, \mathbf{Y}, \mathbf{y}(1)) \right] = \frac{1}{T} \inf_\sigma \sup_{\mathbf{Y}} \text{Regret}_T(\sigma, \mathbf{Y}).$$

Note that the \leq direction is immediate via the min-max inequality. This result shows that the worst-case optimal risk differences in a stochastic setting (i.e. the analyst knows P and responds to it optimally) is equal to minimax regret. In this sense, worst-case regret analysis is not by itself conservative for a stochastic setting—minimax regret is a tight upper bound for performance in stochastic settings.

weighted match that tracks the potential outcomes y_{01}, \dots, y_{0T} well, so that¹⁸

$$\mathbb{E}_P \left[\min_{\theta \in \Theta} \frac{1}{T} \sum_{t=1}^T (y_{0t} - \theta' \mathbf{y}_t)^2 \right] \approx 0.$$

In many empirical settings, it seems plausible that the oracle weighted match performs well.¹⁹ Abadie (2021) states the following intuition in many comparative case studies: “[T]he effect of an intervention can be inferred by comparing the evolution of the outcome variables of interest between the unit exposed to treatment and a group of units that are similar to the exposed unit but were not affected by the treatment.” More formally speaking, a well-fitting oracle weighted match also resembles—and implies—Abadie et al. (2010)’s assumption that there exists a perfect pre-treatment fit of the outcomes. When the oracle weighted match performs well, our regret guarantees imply a guarantee on the loss of the feasible synthetic control estimator, making it an attractive option for causal inference in comparative case studies.

Even if no weighted average of the untreated units tracks y_{0t} closely, synthetic control continues to enjoy the assurance that it performs almost as well as the best weighted match. Moreover, in the general online learning setup (O1) to (O3), this no-regret property cannot be attained without choosing θ_t in some data-dependent manner.²⁰ This observation rules out alternatives such as simple difference-in-differences, which does not aggregate the control units in a data-dependent manner. In contrast, in Section 3, we additionally show that synthetic control on differenced data performs almost as well as the best *weighted* difference-in-differences estimator, a popular class of estimators in practice.

¹⁸Also, observe that $\mathbb{E}_P[\min_{\theta \in \Theta} \frac{1}{T} \sum_{t=1}^T (y_{0t} - \theta' \mathbf{y}_t)^2] \leq \min_{\theta \in \Theta} \mathbb{E}_P[\frac{1}{T} \sum_{t=1}^T (y_{0t} - \theta' \mathbf{y}_t)^2]$, and thus the guarantee (9) is stronger in the sense that it allows the oracle θ to depend on the realization of the data.

¹⁹We recognize that under many data-generating models, there is unforecastable, idiosyncratic randomness in y_{0t} . As a result, there may not exist a synthetic match that perfectly tracks the *realized* series y_{0t} (even though such a match may exist that tracks various conditional expectations of y_{0t} quite well). In many such cases, since squared error can be orthogonally decomposed, risk differences for estimating y_{0t} are also risk differences for estimating conditional means μ_t of y_{0t} . We discuss these results in Appendix B.3.

²⁰See Appendix A.2 for a simple argument in a general setup with unspecified $\ell(\cdot)$. Since simple DID does not choose weights adaptively, it fails to control regret against the class of weighted DID estimators that we discuss in Section 3.

3. EXTENSIONS

3.1. *Non-uniform treatment timing*

The previous interpretations—in (5) and (7)—rely on interpreting average loss over hypothetical values of S as expected loss over S , which requires uniform treatment timing $S \sim \text{Unif}[T]$. Despite being plausible in certain settings and appearing elsewhere in the literature (Doudchenko and Imbens, 2016, Bottmer et al., 2021), this assumption is perhaps crude.²¹ To some extent this is inevitable: Since we are agnostic on the outcome generation process, it is unavoidable to make treatment timing assumptions in order to obtain non-trivial statistical results on estimation of causal quantities. Nevertheless, note that such an assumption is only necessary for interpreting average losses as expected losses. The *a priori* proposition that *it is reasonable to expect a causal estimator to predict well relative to some oracle, at least on average over hypothetical treatment timings*, strikes us as defensible. Accepting this dictum relieves us of any need to model treatment timing.

Even if we wish to maintain the interpretation of average loss as expected loss, we can relax the uniform treatment timing assumption. In this subsection, we show that if the treatment timing distribution is known, then a weighted version of synthetic control achieves logarithmic weighted regret. Moreover, even if the treatment timing distribution is non-uniform, unknown, and possibly chosen by the adversary, we continue to show that synthetic control performs well if some weighted average of untreated units predicts y_{0S} accurately. Both results have constants that worsen if the treatment timing distribution deviates far from $\text{Unif}[T]$.

Suppose the conditional distribution $(S \mid \mathbf{Y})$ is denoted by $\pi = (\pi_1, \dots, \pi_T)'$, which may depend on \mathbf{Y} . Note that, for a known π , we may apply the same argument in [Theorem 2.1](#)

²¹Doudchenko and Imbens (2016) discuss inference in synthetic control via randomization of the treatment timing in their Section 6.2. Bottmer et al. (2021) consider randomization of the treated period in their Assumption 2, though, in their setting, the treatment lasts only one period. We also note that the randomness *per se* of S conditional on \mathbf{Y} can be realistic, but that its distribution is uniform and known is restrictive.

to the following weighted synthetic control estimator:

$$\hat{\theta}_S^\pi \in \arg \min_{\theta \in \Theta} \sum_{t < S} \pi_t (y_{0t} - \theta' \mathbf{y}_t)^2, \quad (10)$$

by redefining the loss functions $\ell_t(\cdot)$. This argument shows that (10) achieves $\log T$ weighted regret, stated in the following corollary. Note that (10) implements FTL with loss functions $\ell_t(\theta) \equiv \pi_t (y_{0t} - \theta' \mathbf{y}_t)^2$, and hence the argument of Hazan et al. (2007) applies.

COROLLARY 3.1: *Suppose $S \sim \pi$, $\frac{1}{CT} \leq \pi_t \leq \frac{C}{T}$ for some C , and $\|\mathbf{Y}\|_\infty \leq 1$. Then weighted synthetic control (10), denoted σ_π , achieves weighted regret bound*

$$\begin{aligned} \text{Regret}_T(\sigma_\pi; \pi, \mathbf{Y}) &\equiv T \cdot \left(\mathbb{E}_{S \sim \pi} [(y_{0S} - \hat{\theta}'_S \mathbf{y}_S)^2] - \min_{\theta \in \Theta} \mathbb{E}_{S \sim \pi} [(y_{0S} - \theta' \mathbf{y}_S)^2] \right) \\ &\leq 16C^3 N \left[\log \frac{\sqrt{NT}}{C^2} + 1 \right] = O(C^3 N \log T). \end{aligned} \quad (11)$$

Corollary 3.1 shows that the weighted regret—a difference in π -expected loss—is logarithmic in T , thereby controlling the worst-case gap between weighted synthetic control and the oracle weighted match for the expected loss. Assuming a known π could be reasonable. With a known dynamic treatment regime, π can depend on $\mathbf{Y}_{1:S-1}$, but is known whenever the analyst is prompted for a prediction at time S .²² We can also interpret Corollary 3.1 as providing guarantees on differences in Bayes risk under the analyst’s prior $S \sim \pi$, independent of \mathbf{Y} .

Even when π is *unknown* and chosen by the adversary, we can bound the loss of unweighted synthetic control, so long as π is not too far from uniform.

²²Since the bound is for a fixed \mathbf{Y} , we can allow π to depend on \mathbf{Y} , so long as $\pi_t(\mathbf{Y})$ is known at time $t + 1$ so that the analyst can compute (10). This allows for Corollary 3.1 to be applied in the following example, which is a more realistic design-based setting. There is a known *dynamic treatment regime* (Chakraborty and Murphy, 2014) parametrizing the treatment hazard: That is,

$$P(S = t \mid S \geq t, \mathbf{Y}) = r_t(\mathbf{Y}_{1:t-1})$$

for some known $r_t(\cdot)$. Then $\pi_t(\mathbf{Y}) = P(S = t \mid \mathbf{Y}) = (1 - r_1) \cdots (1 - r_{t-1}) r_t$ is a function of $\mathbf{Y}_{1:t-1}$. We thank Davide Viviano for suggesting this extension.

COROLLARY 3.2: *Suppose $S \sim \pi$, $\pi_t \leq C/T$ for some C , and $\|\mathbf{Y}\|_\infty \leq 1$. Then synthetic control (1), denoted σ , achieves the following bound on the expected loss*

$$\mathbb{E}_{S \sim \pi} \left[(y_{0S} - \hat{\theta}'_S \mathbf{y}_S)^2 \right] \leq C \left(\min_{\theta \in \Theta} \frac{1}{T} \sum_{t=1}^T (y_{0t} - \theta' \mathbf{y}_t)^2 + \frac{1}{T} \text{Regret}_T(\sigma; \mathbf{Y}) \right), \quad (12)$$

where $\text{Regret}_T(\sigma; \mathbf{Y})$ is defined by (4). Hence, for any joint distribution Q of (\mathbf{Y}, S) where $Q(S = t | \mathbf{Y}) \leq C/T$ for all t , and $Q(\|\mathbf{Y}\|_\infty \leq 1) = 1$, we have the average loss bound

$$\mathbb{E}_Q[(y_{0S} - \hat{\theta}'_S \mathbf{y}_S)^2] \leq C \left(\mathbb{E}_Q \left[\min_{\theta \in \Theta} \frac{1}{T} \sum_{t=1}^T (y_{0t} - \theta' \mathbf{y}_t)^2 \right] + O\left(\frac{N \log T}{T}\right) \right). \quad (13)$$

The result (12) shows that, uniformly over all bounded \mathbf{Y} and bounded treatment distributions π , the expected squared error is bounded by the average loss of the oracle weighted match plus the regret, all scaled with a constant C that indexes how far π deviates from the uniform distribution. Under the same assumption that the oracle weighted match performs well on average, (12) continues to show that the treatment estimation risk of synthetic control is small. Since such a result is valid for all \mathbf{Y} and π , we may understand (12) as a bound that holds even in a setting where the adversary picks both \mathbf{Y} and π , with the restriction that $\pi_t \leq C/T$, but otherwise unrestricted in the dependence between \mathbf{Y} and π .

As before, since (12) is a guarantee uniformly over \mathbf{Y} , it is also a guarantee when we average over \mathbf{Y} under an outcome model, yielding (13). Again, (13) shows that *for any* joint distribution of the bounded outcomes and the treatment timing, the unconditional risk of synthetic control is small when the expected oracle conditional risk, $\mathbb{E}_Q[\min_{\theta \in \Theta} \frac{1}{T} \sum_{t=1}^T (y_{0t} - \theta' \mathbf{y}_t)^2]$, is small—so long as S has sufficient randomness conditional on \mathbf{Y} so that C is not too large.

So far, we have considered weighted averages of untreated units as the class of competing estimators. These competing estimators are matching estimators. However, a more common class of competing estimators in applications are difference-in-differences (DID) estimators. It turns out that synthetic control on preprocessed data has regret guarantees against a class of DID estimators, which we turn to in the next subsection.

3.2. Competing against DID

Section 2 shows that the original synthetic control estimator is competitive against a class of matching estimators that use weighted averages of untreated units as matches for the treated unit. However, in many applications in economics, matching estimators are much less popular than DID estimators, since the latter accounts for unobserved confounders that are additive and constant over time. In this subsection, we show that synthetic control on differenced data is competitive against a large class of DID estimators. Additionally, Appendix A.3 offers regret guarantees against other flavors of DID estimators.

In practice, a common DID specification is the following two-way fixed effects regression:

$$\min_{\mu_i, \alpha_t, \lambda} \sum_{i=0}^N \sum_{t=1}^S \left(y_{it}^{\text{obs}} - \mu_i - \alpha_t - \lambda \mathbb{1}[(i, t) = (0, S)] \right)^2,$$

where the observed outcome $y_{it}^{\text{obs}} = y_{it}$ for all $(i, t) \neq (0, S)$, and $y_{0S}^{\text{obs}} = y_S(1)$. This specification regresses the observed outcomes on unit and time fixed effects, and uses the estimated coefficient λ as an estimate of the treatment effect $y_S(1) - y_{0S}$. Implicitly, this regression uses the estimated fixed effects $\mu_0 + \alpha_S$ as a forecast for the unobserved y_{0S} . We consider a weighted generalization of this regression, a special case of the synthetic DID estimators in Arkhangelsky et al. (2021):^{23,24}

$$\min_{\mu_i, \alpha_t, \lambda} \sum_{i=0}^N \sum_{t=1}^S w_i (y_{it}^{\text{obs}} - \mu_i - \alpha_t - \lambda \mathbb{1}[(i, t) = (0, S)])^2 \quad w_0 = 1, \sum_{i=1}^N w_i = 1, w_i \geq 0. \quad (14)$$

For convex weights $w = (w_1, \dots, w_N)'$, denote by $\sigma_{\text{TWFE}}(w)$ the strategy that estimates (14) on the data $(\mathbf{Y}_{1:t-1}, \mathbf{y}_t)$ at time t ,²⁵ and outputs the estimated coefficients $\mu_0 + \alpha_t$ as a prediction for y_{0t} . By varying over $w \in \Theta$, we obtain a class of competing DID strategies,

²³The weight w_0 does not affect $\mu_0 + \alpha_S$ achieving the optimum in the least-squares problem, per the calculation in Appendix A.6. As a result, we normalize $w_0 = 1$. Moreover, specifically, (14) is a special case of synthetic DID, (1) in Arkhangelsky et al. (2021), with only unit-level weights and no time-level weights.

²⁴(14) is underdetermined if $S = 1$. The ensuing discussion assumes $\sum_{i=1}^N w_i y_{i1}$ is the weighted two-way fixed effects prediction for y_{01} .

²⁵The value of y_{0t} does not enter $\alpha_S + \mu_0$ since it is absorbed by the coefficient λ .

where conventional DID corresponds to picking uniform weights $w = (1/N, \dots, 1/N)'$. We calculate in [Appendix A.6](#) that the prediction that $\sigma_{\text{TWFE}}(w)$ makes is

$$\hat{y}_t(\sigma_{\text{TWFE}}(w)) = \frac{1}{t-1} \sum_{s=1}^{t-1} y_{0s} + w' \left(\mathbf{y}_t - \frac{1}{t-1} \sum_{s=1}^{t-1} \mathbf{y}_s \right) \quad t \geq 2,$$

which simply uses the outcome difference against historical averages of untreated units to forecast that of unit 0. Note that this strategy amounts to using a weighted match with weight w on the *differenced data*

$$\tilde{y}_{i1} = y_{i1} \quad \tilde{y}_{it} \equiv y_{it} - \frac{1}{t-1} \sum_{s=1}^{t-1} y_{is} \quad |\tilde{y}_{it}| \leq 2$$

to forecast the same differences of unit 0, \tilde{y}_{0t} . Therefore, we may apply [Theorem 2.1](#) and show the following regret bound.

THEOREM 3.3: *Consider synthetic control on the differenced data, where the analyst computes*

$$\hat{\theta}_t \in \arg \min_{\theta \in \Theta} \sum_{s < t} (\tilde{y}_{0s} - \theta' \tilde{\mathbf{y}}_s)^2$$

and predicts $\hat{y}_t = \frac{1}{t-1} \sum_{s < t} y_{0s} + \hat{\theta}_t' \tilde{\mathbf{y}}_t$. Here, $\tilde{y}_{it} = y_{it} - \frac{1}{t-1} \sum_{s < t} y_{is}$ is the difference against historical means, and $\tilde{\mathbf{y}}_t = (\tilde{y}_{1,t}, \dots, \tilde{y}_{N,t})'$. Then we have the following regret guarantee against the oracle σ_{TWFE} , whose weights are chosen ex post:

$$\sum_{t=1}^T (y_{0t} - \hat{y}_t)^2 - \min_{\theta \in \Theta} \sum_{t=1}^T (y_{0t} - \hat{y}_t(\sigma_{\text{TWFE}}(\theta)))^2 \leq CN \log T$$

for some constant C .

[Theorem 3.3](#) shows that synthetic control on differenced data controls regret against the class of DID estimators (14).²⁶ In particular, the class of DID benchmarks corresponds to weighted two-way fixed effects regressions, and synthetic control is competitive against any fixed weighting. In this sense, [Theorem 3.3](#) builds on the intuition that synthetic control is a generalization of DID ([Doudchenko and Imbens, 2016](#)) to show that a version of synthetic control performs as well as any weighted DID estimator. Again, if any weighted DID estimator performs well, then [Theorem 3.3](#) becomes a performance guarantee on synthetic control. Moreover, since (14) is a popular alternative for many practitioners—setting aside whether there is a weighted DID that performs well—[Theorem 3.3](#) shows that it is without much loss to use synthetic control in such settings instead. Since DID is more popular in practice than weighted matching, competitive performance against DID is a more relevant consideration, which suggests prioritizing synthetic control on differenced data \tilde{y}_{it} over classic synthetic control (1).²⁷

To the best of our knowledge, the difference scheme \tilde{y}_{it} has yet to be considered in the literature. We do note that since the resulting predictions are equivalent to a weighted two-way fixed effects regression, this proposed synthetic control scheme can be thought of as synthetic DID ([Arkhangelsky et al., 2021](#)) with weights chosen by constrained least-squares on \tilde{y}_{it} . We also note that \tilde{y}_{it} is slightly different from [Ferman and Pinto \(2021\)](#)’s demeaned synthetic control, which takes the difference $\dot{y}_{it} \equiv y_{it} - \frac{1}{t} \sum_{s=1}^t y_{is}$. In [Appendix A.3](#), we show that [Ferman and Pinto \(2021\)](#)’s demeaned synthetic control achieves logarithmic regret against a different class of DID estimators that we call static DID estimators.²⁸ Another

²⁶The benchmark class of DID estimators in [Theorem 3.3](#) output predictions in a sequential manner, in so far as the coefficients in the regression (14) depend on S . In contrast, [Proposition A.3](#) compares synthetic control against a class of static DID estimators that do not exhibit this feature.

²⁷This comment is with the caveat that the constant in [Theorem 3.3](#) is worse than that in [Theorem 2.1](#). It seems possible to further improve the guarantee in [Theorem 3.3](#), since in our proof, we solely use the implication $|\tilde{y}_{it}| \leq 2$ and do not restrict the adversary from choosing \tilde{y}_{it} where the implied $|y_{it}| > 1$. We leave such a refinement to future work.

Of course, this observation also implies that [Theorem 3.3](#) holds without bounded outcomes $\|\mathbf{Y}\|_\infty \leq 1$ and solely with bounded differences $\max_{i,t} |\tilde{y}_{it}| \leq 2$.

²⁸Under certain conditions, [Ferman and Pinto \(2021\)](#) (Proposition 3) show that the demeaned synthetic control in [Proposition A.3](#) dominates DID with uniform weighting $\theta_i = 1/N$. The results [Proposition A.3](#) and [Theorem 3.3](#) are in a similar flavor, and show that synthetic control is competitive against DID with any fixed weighting,

popular alternative is first-differencing (Abadie, 2021), which by similar arguments may be shown to control regret against a class of *two-period* weighted DID strategies that output $\hat{y}_t(\sigma_{2P-DID}(\theta)) \equiv y_{0t-1} + \theta'(\mathbf{y}_t - \mathbf{y}_{t-1})$ as successive predictions.

3.3. Regularization, covariates, and other extensions

Theorem 2.1 shows that synthetic control, as FTL, gives logarithmic regret when we consider quadratic loss. However, to some extent this bound is an artifact of using squared losses, whose curvature ensures that the FTL predictions do not move around excessively over time. If we replace the loss function with the absolute loss $|\hat{y} - y|$, then the regret may be linear in T —no better than that of the trivial prediction $\hat{y}_t \equiv 0$ (see Example 2.10 in Orabona, 2019).

Motivated by the lack of general sublinear regret guarantees in FTL, the online learning literature proposes a large class of algorithms called *Follow-The-Regularized-Leader* (FTRL), where regularization helps stabilize the FTL predictions. With linear prediction functions $f(\mathbf{y}; \theta) = \theta' \mathbf{y}$, such strategies take the form

$$\theta_t \in \arg \min_{\theta \in \Theta} \sum_{s < t} \ell(\theta' \mathbf{y}_s, y_{0s}) + \frac{1}{\eta} \Phi(\theta) \quad (15)$$

for some convex penalty Φ and regularization strength $1/\eta > 0$. Here, we let $\ell(\cdot, \cdot)$ denote a generic convex and bounded loss function, generalizing our previous framework. Many regularized variants of synthetic control have been proposed (among others, Chernozhukov et al., 2021, Doudchenko and Imbens, 2016, Hirshberg, 2021). These regularized estimators have the form (15), though most such estimators are based on quadratic loss.

OBSERVATION 2: *Regularized synthetic control with penalty $\Phi(\cdot)$ is an instance of FTRL, where $\ell(\cdot, \cdot)$ is typically quadratic loss.*

on average over random assignment of treatment time. Of course, Proposition A.3 and Theorem 3.3 are not generalizations of Ferman and Pinto (2021)'s result—for one, we consider average loss under random treatment timing, and Ferman and Pinto (2021) consider a fixed treatment time under an outcome model, with the number of pre-treatment periods tending to infinity.

Moreover, we can think of synthetic control with covariates as regularized synthetic control as well. With time-invariant covariates $\mathbf{x}_j = (x_{1j}, \dots, x_{Nj})'$ for $j = 1, \dots, J$, synthetic control may choose weights θ to additionally match the covariates (see, e.g., (7) in [Abadie, 2021](#)):

$$\hat{\theta}_{S,x} \in \arg \min_{\theta \in \Theta} \sum_{t < S} (y_{0t} - \theta' \mathbf{y}_t)^2 + \frac{1}{2\eta} \sum_{j=1}^J \eta_j (x_{0j} - \theta' \mathbf{x}_j)^2, \quad (16)$$

for some given η_j that indexes the importance of matching covariate j . Observe that, for fixed x_{0j}, \mathbf{x}_j , (16) is a special case of (15); in particular, (16) uses a quadratic penalty of the form

$$\Phi(\theta) = \frac{1}{2}(\mathbf{x} - \mathbf{X}\theta)' H(\mathbf{x} - \mathbf{X}\theta)$$

for some positive definite H , vector \mathbf{x} , and conformable matrix \mathbf{X} . Thus, under the assumption that the covariates x_{0j}, \mathbf{x}_j are fixed and not chosen by the adversary, we may analyze synthetic control with time-invariant covariates as a special case of FTRL.

Motivated by the importance of loss function curvature, we slightly generalize and consider regularized synthetic control estimators using generic loss functions. A standard result in online convex optimization (e.g. Corollary 7.9 in [Orabona \(2019\)](#), Theorem 5.2 in [Hazan \(2019\)](#)) shows that choices of η exist to obtain \sqrt{T} regret.²⁹ The conditions for this result are highly general, explaining the popularity of FTRL in online convex optimization. We specialize to a few choices of the penalty function Φ in the synthetic control setting; see [Theorem A.4](#) for a general statement.

THEOREM 3.4: *Consider regularized synthetic control (15), equivalently FTRL, with penalty function $\Phi(\theta)$ and θ restricted to the simplex Θ . Let $\ell(\theta' \mathbf{y}_t, y_{0t})$ be a convex loss function in θ , not necessarily quadratic, to be specified.*

1. *Consider the quadratic penalty $\Phi(\theta) = \frac{1}{2}(\mathbf{x} - \mathbf{X}\theta)' H(\mathbf{x} - \mathbf{X}\theta)$. Assume the Hessian $\nabla_{\theta\theta'} \Phi(\cdot) = \mathbf{X}' H \mathbf{X}$ is positive definite with minimum eigenvalue normalized to 1. Let $K = \sup_{\theta \in \Theta} \Phi(\theta) - \inf_{\theta \in \Theta} \Phi(\theta)$ be the range of $\Phi(\cdot)$. Then, for both squared loss $\ell(\hat{y}, y) =$*

²⁹This rate matches the lower bound for linear losses. See Chapter 5 of [Orabona \(2019\)](#).

$\frac{1}{2}(y - \hat{y})^2$ and linear loss $\ell(\hat{y}, y) = |y - \hat{y}|$, we have $\text{Regret}_T \leq 2\sqrt{2KNT}$ with the choice $\eta = \sqrt{K(2NT)^{-1}}$.

Moreover, if $\mathbf{x} = 0$ and $\mathbf{X} = H = I$, then $\Phi(\theta) = \frac{1}{2}\|\theta\|^2$ is the ridge penalty, for which we obtain $\text{Regret}_T \leq 2\sqrt{NT}$ with the choice $\eta = 1/\sqrt{4NT}$.

2. For the entropy penalty $\Phi(\theta) = \sum_i \theta_i \log \theta_i + \log(N)$, for both squared and linear losses, we have $\text{Regret}_T \leq 3\sqrt{T \log N}$ with the choice $\eta = \sqrt{(\log N)/T}$.

These results hold for any $N, T > 0$ and allow for $T \leq N$.

Naturally, these choices correspond to regularized variants of synthetic control. As we discuss above, quadratic penalties generalize ridge penalization (Hirshberg, 2021) and matching on covariates.³⁰ The entropy penalty, which is very natural when the parameters lie on the simplex, is a special case of the proposal in Robbins et al. (2017); the resulting regret bound has better dependence on N and obtains the no-regret property as long as $\frac{\log N}{T} \rightarrow 0$.³¹ For these guarantees, the choice of η does require knowledge on the total number of periods T . This may be relaxed via the “doubling trick” (see Shalev-Shwartz (2011), Section 2.3.1), if we allow for different regularization strengths η_S for different realizations of S .

We conclude this section by pointing out a few other extensions. First, another weakening of the uniform treatment timing requirement can be achieved by considering the maximal regret over subperiods of $[T]$, also known as *adaptive regret*. We show in Appendix B.1 that a modification to the synthetic control algorithm—which still outputs a weighted average of untreated units—achieves worst subperiod regret of order $\log T$. Such a result implies that if we *additionally let the adversary pick a subperiod* of length T' , and treatment is uniformly randomly assigned on this subperiod, then modified synthetic control is at most $\frac{\log T}{T'}$ -worse on expected loss than the oracle weighted match. Of course, this regret guarantee is meaningful only when the subperiod is sufficiently long, i.e., $T' \gg \log T$. Second, under

³⁰Ridge penalties are a special case of elastic net penalties proposed by (Doudchenko and Imbens, 2016). Theorem A.4 applies to elastic net penalties with nonzero ℓ_2 component as well.

Note that when $\mathbf{X} \in \mathbb{R}^{J \times N}$ represents pre-treatment covariates of the control units, $\mathbf{X}'H\mathbf{X}$ being positive definite requires that the dimension of the covariates is at least the number of control units.

³¹Interestingly, ℓ_1 -penalty (proposed by, e.g., Chernozhukov et al., 2021) alone is not strongly convex (See Section 9.1.2 of Boyd and Vandenberghe, 2004), and Theorem A.4 does not apply. However, Theorem A.4 only contains sufficient conditions, and so this alone is not a criticism of ℓ_1 -penalty.

a design-based framework on treatment timing, we can test sharp hypotheses of the form $H_0 : \mathbf{y}(1) - \mathbf{y}(0) = \mathbf{z}$ by leveraging symmetries induced by random treatment timing. We briefly discuss inference in [Appendix B.2](#).

4. CONCLUSION

This paper notes a simple connection between synthetic control methods and online convex optimization. Synthetic control is an instance of Follow-The-Leader, which are well-studied strategies in the online learning literature. We present standard regret bounds for FTL that apply to synthetic control, which have interpretations as bounds for expected regret under random treatment timing. These regret bounds translate to bounds on expected risk gap under outcome models and imply that synthetic control is competitive against a wide class of matching estimators. In cases where some weighted match of untreated units predict the unobserved potential outcomes, these results show that synthetic control achieves low expected loss. Moreover, the regret bounds can be adapted to be regret bounds against difference-in-differences strategies. Lastly, we draw an analogous connection between regularized synthetic control and Follow-the-Regularized-Leader, a popular class of strategies in online learning.

We now point out a few limitations of this paper and directions for future work. First and foremost, the approach we have taken in this paper is deliberately pessimistic. Living in fear of an adversary constrained solely by bounded outcomes is perhaps too paranoid for sound decision-making. For instance, this worst-case perspective is not particularly amenable to incorporating covariates, since matching on covariates is inherently based on the hope that the covariates are predictive of potential outcomes. Further constraining the adversary ([Rakhlin et al., 2011](#)) may be an interesting direction for future research. For instance, it may be fruitful to consider an adversary with a fixed budget for how much y_{0t}, \mathbf{y}_t deviate from $y_{0,t-1}, \mathbf{y}_{t-1}$. Constraining the adversary may also render covariates useful, even in a worst-case framework.

It may also be interesting to consider alternative online protocols. So far, we have considered a thought experiment where, before each step t , the analyst only has access to data $\mathbf{Y}_{1:t-1}$ to output a prediction function. In practice, the analyst typically does have access

to $\mathbf{y}_1, \dots, \mathbf{y}_T$. Alternative protocols have been considered in the online learning literature. One example is the Vovk–Azoury–Warmuth forecaster (See Section 7.10 in [Orabona, 2019](#)), where we assume the analyst additionally has access to \mathbf{y}_t before they are prompted for a prediction at time t . In this case, regularized strategies can also achieve $\log T$ regret. Additionally, [Bartlett et al. \(2015\)](#) consider the fixed design setting in which $\mathbf{y}_1, \dots, \mathbf{y}_T$ is fully accessible to the analyst before they are prompted for a prediction. [Bartlett et al. \(2015\)](#) give a simple and explicit minimax regret strategy for online linear regression, which we may adapt into a synthetic control estimator.

We have only considered regret on one-step-ahead prediction for y_{0S} , but synthetic control estimates are often extrapolated multiple time periods ahead in practice. In attempting to extend our results to k -step-ahead prediction, it is natural to consider $\check{y}_{it} = (y_{it}, \dots, y_{i,t+k})$, and to attempt a similar argument on \check{Y} . The chief difficulty in doing so is one of delayed feedback, where the analyst cannot update their time- S decision based on loss from times $1, \dots, S - 1$. That is, for k -step-ahead prediction, the analyst, viewed as an online player who is prompted for a forecast of $\check{y}_{0,S} = (y_{0S}, y_{0,S+1}, \dots, y_{0,S+k-1})$, does not have access to their prediction loss for $\check{y}_{0,S-1} = (y_{0,S-1}, y_{0S}, \dots, y_{0,S+k-2})$, since $y_{0,S+k-2}$ is not yet observed. As a result, unlike (O3) in the standard online convex optimization protocol, the analyst does not have access to $\ell_1(\cdot), \dots, \ell_{S-1}(\cdot)$ when making decisions θ_S —rendering our results here insufficient. That said, delayed feedback—where the online player only has knowledge of the loss function after k periods—is studied in online learning ([Weinberger and Ordentlich, 2002](#), [Korotin et al., 2018](#), [Flaspohler et al., 2021](#)), and we leave an exploration to future work.

REFERENCES

- ABADIE, ALBERTO (2021): “Using synthetic controls: Feasibility, data requirements, and methodological aspects,” *Journal of Economic Literature*, 59, 391–425. [[1](#), [5](#), [13](#), [20](#), [21](#)]
- ABADIE, ALBERTO AND MATIAS D CATTANEO (2021): “Introduction to the Special Section on Synthetic Control Methods,” *Journal of the American Statistical Association*, 116, 1713–1715. [[1](#)]
- ABADIE, ALBERTO, ALEXIS DIAMOND, AND JENS HAINMUELLER (2010): “Synthetic control methods for comparative case studies: Estimating the effect of California’s tobacco control program,” *Journal of the American statistical Association*, 105, 493–505. [[2](#), [5](#), [13](#), [37](#)]

- (2015): “Comparative politics and the synthetic control method,” *American Journal of Political Science*, 59, 495–510. [1]
- ABADIE, ALBERTO AND JAVIER GARDEAZABAL (2003): “The economic costs of conflict: A case study of the Basque Country,” *American economic review*, 93, 113–132. [1, 5]
- ABADIE, ALBERTO AND JAUME VIVES-I-BASTIDA (2021): “Synthetic Controls in Action,” Tech. rep. [2]
- ABERNETHY, JACOB, ALEKH AGARWAL, PETER L BARTLETT, AND ALEXANDER RAKHLIN (2009): “A stochastic view of optimal regret through minimax duality,” *arXiv preprint arXiv:0903.5328*. [12]
- ARKHANGELSKY, DMITRY, SUSAN ATHEY, DAVID A HIRSHBERG, GUIDO W IMBENS, AND STEFAN WAGER (2021): “Synthetic difference-in-differences,” *American Economic Review*, 111, 4088–4118. [17, 19]
- ATHEY, SUSAN, MOHSEN BAYATI, NIKOLAY DOUDCHENKO, GUIDO IMBENS, AND KHASHAYAR KHOSRAVI (2021): “Matrix completion methods for causal panel data models,” *Journal of the American Statistical Association*, 116, 1716–1730. [2, 8]
- BAI, JUSHAN (2003): “Inferential theory for factor models of large dimensions,” *Econometrica*, 71, 135–171. [2]
- BAI, JUSHAN AND SERENA NG (2002): “Determining the number of factors in approximate factor models,” *Econometrica*, 70, 191–221. [2]
- BARTLETT, PETER L, WOUTER M KOOLEN, ALAN MALEK, EIJI TAKIMOTO, AND MANFRED K WARMUTH (2015): “Minimax fixed-design linear regression,” in *Conference on Learning Theory*, PMLR, 226–239. [24]
- BEN-MICHAEL, ELI, AVI FELLER, AND JESSE ROTHSTEIN (2019): “Synthetic controls and weighted event studies with staggered adoption,” *arXiv preprint arXiv:1912.03290*. [2]
- (2021): “The augmented synthetic control method,” *Journal of the American Statistical Association*, 116, 1789–1803. [2]
- BLACKWELL, DAVID (1956): “An analog of the minimax theorem for vector payoffs,” *Pacific Journal of Mathematics*, 6, 1–8. [6]
- BOTTMER, LEA, GUIDO IMBENS, JANN SPIESS, AND MERRILL WARNICK (2021): “A Design-Based Perspective on Synthetic Control Methods,” *arXiv preprint arXiv:2101.09398*. [2, 7, 14, 37]
- BOYD, STEPHEN P AND LIEVEN VANDENBERGHE (2004): *Convex optimization*, Cambridge university press. [22]
- BROWN, GEORGE W (1951): “Iterative solution of games by fictitious play,” *Activity analysis of production and allocation*, 13, 374–376. [8]
- BUBECK, SÉBASTIEN AND NICOLO CESA-BIANCHI (2012): “Regret analysis of stochastic and nonstochastic multi-armed bandit problems,” *Foundations and Trends® in Machine Learning*, 5, 1–122. [7]
- CESA-BIANCHI, NICOLO AND GÁBOR LUGOSI (2006): *Prediction, learning, and games*, Cambridge university press. [3]
- CHAKRABORTY, BIBHAS AND SUSAN A MURPHY (2014): “Dynamic treatment regimes,” *Annual review of statistics and its application*, 1, 447–464. [15]

- CHERNOZHUKOV, VICTOR, KASPAR WÜTHRICH, AND YINCHU ZHU (2021): “An exact and robust conformal inference method for counterfactual and synthetic controls,” *Journal of the American Statistical Association*, 116, 1849–1864. [20, 22, 37]
- CURRIE, JANET, HENRIK KLEVEN, AND ESMÉE ZWIERS (2020): “Technology and big data are changing economics: Mining text to track methods,” in *AEA Papers and Proceedings*, vol. 110, 42–48. [3]
- DOUDCHENKO, NIKOLAY AND GUIDO W IMBENS (2016): “Balancing, regression, difference-in-differences and synthetic control methods: A synthesis,” Tech. rep., National Bureau of Economic Research. [4, 5, 7, 14, 19, 20, 22, 30]
- FERMAN, BRUNO AND CRISTINE PINTO (2021): “Synthetic controls with imperfect pretreatment fit,” *Quantitative Economics*, 12, 1197–1221. [2, 19, 20, 30]
- FLASPOHLER, GENEVIEVE E, FRANCESCO ORABONA, JUDAH COHEN, SOUKAYNA MOUATADID, MIRUNA OPRESCU, PAULO ORENSTEIN, AND LESTER MACKEY (2021): “Online learning with optimism and delay,” in *International Conference on Machine Learning*, PMLR, 3363–3373. [24]
- HANNAN, JAMES (1958): “Approximation to Bayes risk in repeated play,” in *Contributions to the Theory of Games (AM-39), Volume III*, Princeton University Press, 97–140. [6]
- HAZAN, ELAD (2019): “Introduction to online convex optimization,” *arXiv preprint arXiv:1909.05207*. [3, 7, 8, 21, 35, 36]
- HAZAN, ELAD, AMIT AGARWAL, AND SATYEN KALE (2007): “Logarithmic regret algorithms for online convex optimization,” *Machine Learning*, 69, 169–192. [10, 15, 27]
- HIRSHBERG, DAVID A (2021): “Least squares with error in variables,” *arXiv preprint arXiv:2104.08931*. [20, 22]
- KALAI, ADAM AND SANTOSH VEMPALA (2005): “Efficient algorithms for online decision problems,” *Journal of Computer and System Sciences*, 71, 291–307. [3, 8]
- KOROTIN, ALEXANDER, VLADIMIR V’ YUGIN, AND EVGENY BURNAEV (2018): “Aggregating strategies for long-term forecasting,” in *Conformal and Probabilistic Prediction and Applications*, PMLR, 63–82. [24]
- LEI, JING, MAX G’ SELL, ALESSANDRO RINALDO, RYAN J TIBSHIRANI, AND LARRY WASSERMAN (2018): “Distribution-free predictive inference for regression,” *Journal of the American Statistical Association*, 113, 1094–1111. [37]
- MANSKI, CHARLES F AND JOHN V PEPPER (2018): “How do right-to-carry laws affect crime rates? Coping with ambiguity using bounded-variation assumptions,” *Review of Economics and Statistics*, 100, 232–244. [2, 5]
- ORABONA, FRANCESCO (2019): “A modern introduction to online learning,” *arXiv preprint arXiv:1912.13213*. [3, 8, 20, 21, 24, 32]
- RAKHLIN, ALEXANDER, KARTHIK SRIDHARAN, AND AMBUJ TEWARI (2011): “Online learning: Stochastic, constrained, and smoothed adversaries,” *Advances in neural information processing systems*, 24. [23]
- ROBBINS, HERBERT (1952): “Some aspects of the sequential design of experiments,” *Bulletin of the American Mathematical Society*, 58, 527–535. [7]

- ROBBINS, MICHAEL W, JESSICA SAUNDERS, AND BEAU KILMER (2017): “A framework for synthetic control methods with high-dimensional, micro-level data: evaluating a neighborhood-specific crime intervention,” *Journal of the American Statistical Association*, 112, 109–126. [22]
- SHALEV-SHWARTZ, SHAI (2011): “Online learning and online convex optimization,” *Foundations and trends in Machine Learning*, 4, 107–194. [3, 22]
- SHI, CLAUDIA, DHANYA SRIDHAR, VISHAL MISRA, AND DAVID BLEI (2022): “On the Assumptions of Synthetic Control Methods,” in *International Conference on Artificial Intelligence and Statistics*, PMLR, 7163–7175. [2]
- VIVIANO, DAVIDE AND JELENA BRADIC (2019): “Synthetic learner: model-free inference on treatments over time,” *arXiv preprint arXiv:1904.01490*. [2, 4]
- WEINBERGER, MARCELO J AND ERIK ORDENTLICH (2002): “On delayed prediction of individual sequences,” *IEEE Transactions on Information Theory*, 48, 1959–1976. [24]
- XU, YIQING (2017): “Generalized synthetic control method: Causal inference with interactive fixed effects models,” *Political Analysis*, 25, 57–76. [2]

APPENDIX A: PROOFS AND ADDITIONAL RESULTS

A.1. Proofs of [Theorem 2.1](#) and [Corollaries 3.1](#) and [3.2](#)

We reproduce [Theorem 5](#) of [Hazan et al. \(2007\)](#) in our notation.

THEOREM A.1—[Theorem 5](#), [Hazan et al. \(2007\)](#): Assume that for all t , the function $\ell_t : \Theta \rightarrow \mathbb{R}$ can be written as

$$\ell_t(\theta) = g_t(v_t' \theta)$$

for a univariate convex function $g_t : \mathbb{R} \rightarrow \mathbb{R}$ and some vector $v_t \in \mathbb{R}^n$. Assume that for some $R, a, b > 0$, we have $\|v_t\|_2 \leq R$ and for all $\theta \in \Theta$, we have $|g_t'(v_t' \theta)| \leq b$ and $g_t''(v_t' \theta) \geq a$, for all t . Then FTL on ℓ_t satisfies the following regret bound:

$$\text{Regret}_T \leq \frac{2nb^2}{a} \left[\log \left(\frac{DRaT}{b} \right) + 1 \right]$$

where $D = \max_{x, y \in \Theta} \|x - y\|_2$ is the diameter of Θ .

PROOF OF [THEOREM 2.1](#): [Theorem 2.1](#) follows immediately from [Theorem 5](#) in [Hazan et al. \(2007\)](#), reproduced in our notation as [Theorem A.1](#). The proof of this theorem relies

solely on optimality of θ_t (and the associated first-order condition); thus, in the case of multiple minima when minimizing $\sum_{t=1}^s \ell_t(\theta)$, any particular sequence of minima $\{\theta_t\}$ satisfies the guarantee.

Since Θ is the simplex, we know

$$D = \max_{\theta_1, \theta_2 \in \Theta} \|\theta_1 - \theta_2\|_2 \leq \max_{\theta_1, \theta_2 \in \Theta} \|\theta_1 - \theta_2\|_1 \leq \max_{\theta_1, \theta_2 \in \Theta} \|\theta_1\|_1 + \|\theta_2\|_1 = 2.$$

We choose $g_t(x) = \frac{1}{2}(y_{0t} - x)^2$ with $g'_t(x) = x - y_{0t}$ and $g''_t(x) = 1$. (The scaling by $1/2$ means that we obtain a bound on $1/2$ times the regret.) The vectors $v_t = \mathbf{y}_t$, whose dimensions are $n = N$ and whose 2-norms are bounded by $R = \sqrt{N}$. Note that $|\mathbf{y}'_t \theta| = |v'_t \theta| \leq \|v_t\|_\infty \|\theta\|_1 \leq 1$. Hence, $|g'_t(v'_t \theta)| = |\mathbf{y}'_t \theta - y_{0t}| \leq |\mathbf{y}'_t \theta| + |y_{0t}| \leq 2 \equiv b$ and $g''_t(x) \geq 1 \equiv a$. To summarize, we have $R = \sqrt{N}$, $a = 1$, $b = 2$, $D = 2$, and $n = N$.

Plugging in, we have

$$\frac{1}{2} \text{Regret}_T \leq 8N(\log(\sqrt{NT}) + 1),$$

which rearranges into the claim. *Q.E.D.*

PROOF OF COROLLARY 3.1 : The proof for Corollary 3.1 follows similarly, now with

$$g_t(x) = \frac{T}{2} \pi_t (y_{0t} - x)^2 \quad g'_t(x) = T \pi_t (x - y_{0t}) \quad g''_t(x) = T \pi_t.$$

Note that, since $\frac{1}{CT} \leq \pi_t \leq \frac{C}{T}$, we can take $a = 1/C$ and $b = 2C$. Doing so yields the expression in Corollary 3.1. *Q.E.D.*

PROOF OF COROLLARY 3.2: For Corollary 3.2, and in particular (12), by $(1, \infty)$ -Hölder's inequality,

$$\sum_{t=1}^T \pi_t (y_{0t} - \hat{\theta}'_t \mathbf{y}_t)^2 \leq \left(\max_t \pi_t \right) \sum_{t=1}^T (y_{0t} - \hat{\theta}'_t \mathbf{y}_t)^2 \leq \frac{C}{T} \sum_{t=1}^T (y_{0t} - \hat{\theta}'_t \mathbf{y}_t)^2.$$

We then apply [Theorem 2.1](#) to bound $\sum_{t=1}^T (y_{0t} - \hat{\theta}'_t \mathbf{y}_t)^2 = \min_{\theta \in \Theta} \sum_{t=1}^T (y_{0t} - \theta' \mathbf{y}_t)^2 + \text{Regret}_T$.

(13) follows immediately from (12) by taking the expectation \mathbb{E}_Q , noting that

$$\begin{aligned} \mathbb{E}_Q[(y_{0S} - \hat{\theta}'_S y_S)^2] &= \mathbb{E}_Q \left[\sum_{t=1}^T \mathbb{1}(S=t) (y_{0t} - \hat{\theta}'_t y_t)^2 \right] \\ &= \mathbb{E} \left[\mathbb{E} \left[\sum_{t=1}^T \mathbb{1}(S=t) (y_{0t} - \hat{\theta}'_t y_t)^2 \mid \mathbf{Y} \right] \right] \\ &= \mathbb{E} \left[\sum_{t=1}^T Q(S=t \mid \mathbf{Y}) (y_{0t} - \hat{\theta}'_t y_t)^2 \right] \end{aligned}$$

We then apply (12) to complete the proof. Q.E.D.

A.2. Lack of regret control for fixed strategies

LEMMA A.2: *In the online convex optimization setup, suppose the class of loss functions available to the adversary satisfies the following property: There exists $\epsilon > 0$ such that for any $\theta \in \Theta$, there exists $\tilde{\theta}$ and ℓ_1, \dots, ℓ_T , for which $\ell_t(\tilde{\theta}) \leq \ell_t(\theta) - \epsilon$. Then, the regret of any fixed strategy that outputs $\theta_t = \theta$ for every period is at least ϵT .*

PROOF: Let $\ell_t, \tilde{\theta}$ be the sequence of loss functions and alternative satisfying the required property on the class of loss functions. Then $\text{Regret}_T(\theta) \geq \sum_t \ell_t(\theta) - \sum_t \ell_t(\tilde{\theta}) = \epsilon T$.

Q.E.D.

It is easy to see that the loss functions in the panel prediction problem are rich enough to satisfy the property in [Lemma A.2](#). Fix, say, $\epsilon < 0.0001$. For any θ , we can find $\tilde{\theta} \in \Theta$ where $\|\tilde{\theta} - \theta\|_1 \geq \sqrt{\epsilon}$. Then, there exists some \mathbf{y} , $\|\mathbf{y}\|_\infty \leq 1$ where

$$|(\tilde{\theta} - \theta)' \mathbf{y}| = \max_{\|\mathbf{y}\|_\infty \leq 1} |(\tilde{\theta} - \theta)' \mathbf{y}| = \|\tilde{\theta} - \theta\|_1 \geq \sqrt{\epsilon}$$

since $\|\cdot\|_1$ is the dual norm to $\|\cdot\|_\infty$. The adversary chooses $\mathbf{y}_t = \mathbf{y}$ for all $t \in [T]$ and $y_{0t} = \tilde{\theta}' \mathbf{y}_t$. Then $\ell_t(\tilde{\theta}) = 0$ but $\ell_t(\theta) \geq (\sqrt{\epsilon})^2 = \epsilon$.

A.3. Static DID regret control

We could consider affine predictors with bounded intercepts

$$f(\mathbf{y}_t; \theta_0, \theta_1) = \theta_0 + \theta_1' \mathbf{y}_t \quad \Theta = [-2, 2] \times \Delta^{N-1}.$$

This choice corresponds to variations of synthetic control proposed by [Doudchenko and Imbens \(2016\)](#) and [Ferman and Pinto \(2021\)](#) in efforts to mimic behavior of DID estimators.³² Our regret bound from [Theorem 2.1](#) generalizes immediately to the affine predictions, where the benchmark oracle the regret measures against is

$$\min_{(\theta_0, \theta_1) \in \Theta} \sum_{t=1}^T (y_{0t} - \theta_0 - \theta_1' \mathbf{y}_t)^2. \quad (17)$$

(17) simultaneously chooses the best intercept and the best set of convex weights in hindsight. Because (17) is limited to using the same intercept for prediction in each period, it is, in some sense, a *static* DID estimator.

[Theorem 2.1](#) can be adapted to show that synthetic control with an intercept is competitive against static DID.

PROPOSITION A.3: *Consider demeaned synthetic control, where the analyst outputs the prediction $\hat{y}_t = \hat{\theta}_{0t} + \hat{\theta}_t' \mathbf{y}_t$ by solving the least-squares problem*

$$\hat{\theta}_{0t}, \hat{\theta}_t = \arg \min_{\theta_0, \theta \in [-2, 2] \times \Delta^{N-1}} \sum_{s < t} (y_{0s} - \theta_0 - \theta' \mathbf{y}_s)^2.$$

Then, under bounded data $\|\mathbf{Y}\|_\infty \leq 1$, we have the following regret bound:

$$\sum_{t=1}^T (y_{0t} - \hat{y}_t)^2 - \min_{\theta_0, \theta \in [-2, 2] \times \Delta^{N-1}} \sum_{t=1}^T (y_{0t} - \theta_0 - \theta' \mathbf{y}_t)^2 \leq CN \log T$$

for some constant C .

³²Synthetic control with an intercept is equivalent to synthetic control with demeaned data $\{y_s - \frac{1}{t} \sum_{k \leq t} y_k : s = 1, \dots, t\}$ ([Ferman and Pinto, 2021](#)), since the constraint that $\theta_0 \in [-2, 2]$ does not bind.

PROOF: We define the loss as $\frac{1}{2}(x - y)^2$, which only affects the regret up to a factor of 2. [Proposition A.3](#) can be proved with [Theorem A.1](#). Note that the diameter of the parameter space $[-2, 2] \times \Delta^{N-1}$ can be bounded by $D = 2 \cdot \sqrt{2^2 + 1} = 2\sqrt{5}$. The 2-norm of the vector $v_t = [1, \mathbf{y}'_t]'$ is now bounded by $R = \sqrt{N + 1}$. The 1-norm of the parameter vector $\vartheta = [\theta_0, \theta']'$ is now bounded by $2 + 1 = 3$. Hence, $|v'_t \vartheta| \leq 3$. Hence, we may take $b = 3 + 1 = 4$ and $a = 1$. Plugging in, we obtain

$$\text{Regret}_T \leq 64N \left[\log \left(\frac{\sqrt{5}}{2} \sqrt{N + 1} T \right) + 1 \right] < CN \log T$$

for some C .

Q.E.D.

A.4. Proof of [Theorem 3.3](#)

Similarly to the proof of [Proposition A.3](#), suppose the adversary picks the differences $|\tilde{y}_{it}| \leq 2$, *without* the constraint that the resulting levels obey the restriction $\|\mathbf{Y}\|_\infty \leq 1$. An application of [Theorem A.1](#) shows that

$$\sum_{t=1}^T (\tilde{y}_{0t} - \hat{\theta}'_t \tilde{\mathbf{y}}_t)^2 - \min_{\theta \in \Theta} \sum_{t=1}^T (\tilde{y}_{0t} - \theta' \tilde{\mathbf{y}}_t)^2 \leq CN \log T$$

for some C , uniformly over $|\tilde{y}_{it}| \leq 2$, where $\hat{\theta}_t$ is the FTL strategy on the data \tilde{y}_{it} , which is exactly the synthetic control on the differenced data when \mathbf{Y} is chosen by the adversary.

Now, given any $\|\mathbf{Y}\|_\infty \leq 1$, we have that the corresponding differences \tilde{y}_{it} obey the above regret bound, since they are bounded by 2. Moreover, for both synthetic control ($\theta_t = \hat{\theta}_t$) and the oracle σ_{TWFE} ($\theta_t = \theta$), the prediction error of the data y_{0t} is equal to the prediction error on the differences:

$$y_{0t} - \hat{y}_t = \frac{1}{t-1} \sum_{s<t} y_{0s} + \tilde{y}_{0t} - \left(\frac{1}{t-1} \sum_{s<t} y_{0s} + \theta'_t \tilde{\mathbf{y}}_t \right) = \tilde{y}_{0t} - \theta'_t \tilde{\mathbf{y}}_t.$$

Hence, we may rewrite the above regret bound as the bound

$$\sum_{t=1}^T (y_{0t} - \hat{y}_t)^2 - \min_{\theta \in \Theta} \sum_{t=1}^T (y_{0t} - \hat{y}_t(\sigma_{\text{TWFE}}(\theta)))^2 \leq CN \log T.$$

A.5. Proof of [Theorem 3.4](#)

THEOREM A.4: *Assume that*

1. $\ell_t(\theta) \equiv \ell(\theta' \mathbf{y}_t, y_{0t})$ is convex in θ for any \mathbf{Y} .
2. The regularizer $\Phi(\theta)$ is 1-strongly convex in some norm $\|\cdot\|$. Normalize Φ such that its minimum over Θ is zero and maximum is $K < \infty$.
3. All subgradients $\nabla_{\theta} \ell_t(\theta)$ are bounded in the dual norm $\|\cdot\|_*$, uniformly over Θ, \mathbf{Y} :

$$\|\nabla_{\theta} \ell_t(\theta)\|_*^2 \leq G.$$

Then FTRL attains the regret bound

$$\text{Regret}_T \leq \frac{K}{\eta} + \frac{\eta TG}{2}.$$

We first reproduce Corollary 7.9 from [Orabona \(2019\)](#) in our notation. Consider an FTRL algorithm that regularizes according to

$$\theta_t \in \arg \min_{\theta} \sum_{s \leq t} \ell_s(\theta) + \frac{1}{\eta} \Phi(\theta).$$

This corresponds to choosing $\eta_t = \eta$, $\psi(x) = \Phi(x)$, and $\min_{\theta} \Phi(\theta) = 0$ in [Orabona \(2019\)](#).

THEOREM A.5—Corollary 7.9, [Orabona \(2019\)](#): *Let ℓ_t be a sequence of convex loss functions. Let $\Phi : \Theta \rightarrow \mathbb{R}$ be μ -strongly convex with respect to the norm $\|\cdot\|$. Then, FTRL guarantees*

$$\sum_{t=1}^T \ell_t(\theta_t) - \sum_{t=1}^T \ell_t(\theta) \leq \frac{\Phi(\theta)}{\eta} + \frac{\eta}{2\mu} \sum_{t=1}^T \|g_t\|_*^2$$

for all subgradients $g_t \in \partial \ell_t(\theta_t)$ and all $\theta \in \Theta$, where $\|\cdot\|_*$ is the dual norm of $\|\cdot\|$.

PROOF OF [THEOREM A.4](#): [Theorem A.4](#) then follows immediately where $\|g_t\|_*^2 \leq G$, $\Phi(\theta) \leq K$, and $\mu = 1$. *Q.E.D.*

PROOF OF [THEOREM 3.4](#): For both squared and absolute losses, we can bound the gradient of the loss function in terms of

$$\|\nabla_{\theta} \ell_t(\theta)\|_* = \|\nabla q(y - \hat{y}) \cdot \mathbf{y}_t\|_* = |\nabla q(y - \hat{y})| \|\mathbf{y}_t\|_* \leq 2 \sup_{\|\mathbf{y}\|_{\infty} \leq 1} \|\mathbf{y}\|_*$$

under any norm, where $q(t) = t^2/2$ or $q(t) = |t|$. This is because (i) for squared loss, the gradient $|\nabla f| = |y - \hat{y}|$ is bounded by 2 and (ii) for absolute loss, the subgradients $|\nabla f|$ are bounded by 1 and hence by 2. Hence, we should pick G to be $4 \sup_{\|\mathbf{y}\|_{\infty} \leq 1} \|\mathbf{y}\|_*^2$.

For the quadratic penalty assumed, it is 1-strongly convex with respect to $\|\cdot\|_2$ by the assumption that the minimum eigenvalue of its Hessian is 1. Thus the dual norm $\|\cdot\|_*$ is also the Euclidean norm, and we may take $G = 4N$. This yields the bound by [Theorem A.4](#), since

$$\frac{K}{\sqrt{K(2TN)^{-1}}} + \frac{4NT}{2} \sqrt{\frac{K}{2TN}} = 2\sqrt{2}\sqrt{NTK}.$$

Setting $K = 1/2$ yields the ridge penalty result.

The entropy penalty is 1-strongly convex with respect to $\|\cdot\|_1$.³³ Thus we may take $G = 4\|\mathbf{y}_t\|_{\infty}^2 = 4$. The maximum of entropy (shifted so that its minimum is zero) can take $K = \log N$. This yields the bound via [Theorem A.4](#). *Q.E.D.*

³³This is a well-known result in online convex optimization. To prove it, we first note that

$$\Phi(y) = \Phi(x) + \nabla \Phi(x)'(y - x) + D_{\text{KL}}(y||x),$$

where $\Phi(x) = \sum_i x_i \log x_i + C$, $D_{\text{KL}}(y||x) = \sum_i y_i \log(y_i/x_i)$, and x, y lie in the interior of the simplex. Pinsker's inequality then implies

$$\Phi(y) \geq \Phi(x) + \nabla \Phi(x)'(y - x) + \frac{1}{2}\|x - y\|_1^2.$$

This is exactly the definition of 1-strong convexity with respect to $\|\cdot\|_1$.

A.6. Two-way fixed effect calculation

Consider the TWFE regression with known, nonnegative weights $\sum_{i=1}^N w_i = 1$ and the normalization $w_0 = 1$:

$$\arg \min_{\mu_i, \alpha_t} \sum_{\substack{i,t:(i,t) \neq (0,S) \\ i \in \{0, \dots, N\} \\ t \in [S]}} w_i (y_{it} - \mu_i - \alpha_t)^2.$$

We may eliminate $(i, t) = (0, S)$ from the sum since $\lambda \mathbb{1}(i = 0, S = t)$ in (14) absorbs that term, leaving μ_i, α_t unaffected. Consider forecasting y_{0S} with $\mu_0 + \alpha_S$ that solves the above program. As a reminder, in this subsection, we show that the estimated $\mu_0 + \alpha_S$ takes the form of forecasting with weighted average on differenced data.

The first-order condition for μ_i takes the form

$$\sum_{t=1}^{S-1} y_{it} - \mu_i - \alpha_t + \mathbb{1}(i \neq 0)(y_{iS} - \mu_i - \alpha_t) = 0.$$

Hence,

$$\mu_i = \begin{cases} \bar{y}_i - \bar{\alpha} & i \neq 0 \\ \bar{y}_0 - \frac{S}{S-1} \bar{\alpha} + \frac{1}{S-1} \alpha_S & i = 0 \end{cases}$$

where $\bar{\alpha} = \frac{1}{S} \sum_{t=1}^S \alpha_t$ and \bar{y}_i is the sample mean of observations for unit i over time $1, \dots, S$, with the understanding that y_{0S} is not included for \bar{y}_0 . Hence, the forecast is $\mu_0 + \alpha_S = \bar{y}_0 + \frac{S}{S-1} (\alpha_S - \bar{\alpha})$.

Let us inspect the first-order condition for α_S :

$$\sum_{i=1}^N w_i (y_{iS} - \mu_i - \alpha_S) = \sum_{i=1}^N w_i (y_{iS} - \bar{y}_i + \bar{\alpha} - \alpha_S) = 0.$$

Rearrange to obtain that $\alpha_S - \bar{\alpha} = \sum_{i=1}^N w_i \left(\frac{S-1}{S} y_{iS} - \frac{1}{S} \sum_{t=1}^{S-1} y_{it} \right)$. Therefore, $\frac{S}{S-1}(\alpha_S - \bar{\alpha}) = \sum_{i=1}^N w_i \left(y_{iS} - \frac{1}{S-1} \sum_{t=1}^{S-1} y_{it} \right)$. Thus the forecast is

$$\mu_0 + \alpha_S = \frac{1}{S-1} \sum_{t=1}^{S-1} y_{0t} + \sum_{i=1}^N w_i \left(y_{iS} - \frac{1}{S-1} \sum_{t=1}^{S-1} y_{it} \right).$$

Note that arriving at this result does not use the fact that $w_0 = 1$. Hence, w_0 does not matter for $\mu_0 + \alpha_S$.

APPENDIX B: FURTHER EXTENSIONS

B.1. Adaptive regret

The online learning literature also has results for controlling the *adaptive regret*:

$$\text{AdaptiveRegret}_T = \sup_{1 \leq r < s \leq T} \sum_{t=r}^s \left\{ \ell_t(\theta_t) - \min_{\theta_{r,s}} \sum_{t=r}^s \ell_t(\theta_{r,s}) \right\}, \quad (18)$$

which is the worst regret over any subinterval of $[T]$. An upper bound of adaptive regret serves as an upper bound of the regret over any subperiod indexed by $r < s$. In particular, suppose we obtain a $O(\log T)$ upper bound on adaptive regret, then we obtain meaningful *average* regret upper bounds for all subperiods significantly longer than $O(\log T)$.

A simple meta-algorithm called *Follow The Leading History* (FLH) (Algorithm 31 in [Hazan, 2019](#)) serves as a wrapper for an online learning algorithm σ , such that

$$\text{AdaptiveRegret}_T(\text{FLH}(\sigma)) \leq \text{Regret}_T(\sigma) + O(\log T). \quad (19)$$

When applied to synthetic control, FLH takes the following form. We initialize $p_1^1 = 1$ and set $\alpha = \frac{1}{4}$. At each time t , when prompted to make a prediction about y_{0t} :

1. Consider the synthetic control estimated weights $\theta_t^1, \dots, \theta_t^t$, where θ_t^j is the synthetic control weights estimated based on data from *time horizons* $j, \dots, t-1$.
2. Output the weighted average $\theta_t = \sum_{j=1}^t p_t^j \theta_t^j$.

3. After receiving \mathbf{y}_t, y_{0t} (and hence receiving $\ell_t(\theta) = \frac{1}{2}(y_{0t} - \theta' \mathbf{y}_t)^2$), instantiate

$$p_{t+1}^i \leftarrow \frac{p_t^i e^{-\alpha \ell_t(\theta_t^i)}}{\sum_{j=1}^t p_t^j e^{-\alpha \ell_t(\theta_t^j)}} \quad 1 \leq i \leq t.$$

4. Set $p_{t+1}^{t+1} = \frac{1}{t+1}$ and further update

$$p_{t+1}^i \leftarrow \left(1 - \frac{1}{t+1}\right) p_{t+1}^i \quad 1 \leq i \leq t.$$

At each step, FLH applied to synthetic control continues to output a convex weighted average of control unit outcomes, making it a type of synthetic control algorithm. Theorem 10.5 in [Hazan \(2019\)](#) then implies the bound (19) for the above algorithm.³⁴ In a nutshell, FLH treats synthetic control predictions from different horizons as *expert predictions*, and applies a no-regret online learning algorithm to aggregate these expert predictions. We direct readers to [Hazan \(2019\)](#) for further intuitions about the algorithm.

Combined with [Theorem 2.1](#) for synthetic control, we find that the adaptive regret of FLH-synthetic control is of the same order $O(N \log T + N \log N)$. This means that the average regret over any subperiod of length T' is $O\left(\frac{N \log T + N \log N}{T'}\right)$, a meaningful bound for long subperiods $T' \gg N \log T$. In other words, in a protocol where the adversary *additionally* picks a subperiod of length T' , and nature subsequently samples a treatment timing uniformly randomly over the subperiod, FLH-synthetic control achieves expected regret bound of $O\left(\frac{N \log T + N \log N}{T'}\right)$. The adaptive regret bound thus partially relaxes the requirement for uniform treatment timing, and allows for expected regret control over random treatment timing on any subperiod.

³⁴The proof follows immediately since $\frac{1}{2}(y_{0t} - \theta' \mathbf{y}_t)^2$ is $\frac{1}{4}$ -exp-concave. That is,

$$\theta \mapsto \exp\left(-\frac{1}{4} \cdot \frac{1}{2}(y_{0t} - \theta' \mathbf{y}_t)^2\right)$$

is concave. This is because $-2 \leq y_{0t} - \theta' \mathbf{y}_t \leq 2$, and $g(x) = \exp(-\frac{1}{4} \cdot \frac{1}{2}x^2)$ is concave on $x \in [-2, 2]$. The Hessian of $\exp(-\frac{1}{4} \cdot \frac{1}{2}(y_{0t} - \theta' \mathbf{y}_t)^2)$ in θ is then $g''(y_{0t} - \theta' \mathbf{y}_t) \mathbf{y}_t \mathbf{y}_t'$, which is negative semidefinite.

B.2. A note on inference

Under the treatment assignment model $S \sim \text{Unif}[T]$, we may test the sharp null $H_0 : \mathbf{y}(1) = \mathbf{y}(0)$, leveraging symmetries arising from treatment assignment. This is similar in spirit to [Bottmer et al. \(2021\)](#), who consider design-based inference under random assignment of the treated unit. They compute the variance of the estimated treatment effect (for treated unit $M \sim \text{Unif}[N]$ at some fixed time S) under random assignment, holding the outcomes fixed, and propose an unbiased estimator. This is also similar in spirit to unit-randomization-based placebo tests ([Abadie et al., 2010](#)).

Let $y_t = y_{0t}$ for $t < S$ and let $y_t = \mathbf{y}_t(1)$ for $t \geq S$ be the observed time series of the treated unit. For any prediction \hat{y}_t that does not depend on S —not limited to synthetic control predictions—we may form the residuals $r_t = |y_t - \hat{y}_t|$. One (finite-sample) test of the sharp null rejects when r_S is at least the $\lceil T(1 - \alpha) \rceil^{\text{th}}$ order statistic of the sample $\{r_1, \dots, r_T\}$. Since, under the null, r_S is equally likely to equal any of $\{r_1, \dots, r_T\}$, the probability of it being the among largest $100\alpha\%$ is bounded by α . Similarly, if $S \sim \pi$ where $\pi_t \leq C/T$, a least-favorable test may be constructed by rejecting when $r_t \geq r_{(T - \lfloor T\alpha/C \rfloor)}$. Informally speaking, this test is more powerful when the predictions \hat{y}_t are better, and our regret guarantees are in this sense informative for inference. Moreover, note that this procedure is very similar to conformal inference ([Lei et al., 2018](#), [Chernozhukov et al., 2021](#)). Conformal intervals rely on the assumption that the data is exchangeable in the underlying sampling process. This symmetry is true here by virtue of assuming $S \sim \text{Unif}[T]$, since the treated period is equally likely to be any one.

The argument above does not use the regret result. From Markov’s inequality, we can control the probability for the prediction error to deviate far relative to its expectation

$$\mathbb{P}_{S \sim \text{Unif}[T]} [(y_{0S} - \hat{y}_S)^2 > c] \leq \frac{\mathbb{E}_S[\ell_S(\theta_S)]}{c} \leq \frac{1}{c} \left(\min_{\theta \in \Theta} \frac{1}{T} \sum_{i=1}^T \ell_t(\theta) + \frac{1}{T} \text{Regret}_T \right).$$

Under assumptions where the pre-treatment loss $\min_{\theta} \frac{1}{S-1} \sum_{t < S} \ell_t(\theta)$ is a consistent estimator for the oracle performance $\min_{\theta} \frac{1}{T} \sum_{i=1}^T \ell_t(\theta)$, the above observation allows for

predictive confidence intervals for the untreated outcome and confidence intervals of the treatment effect, which are valid over random treatment timing.

B.3. Risk interpretation under idiosyncratic errors

We consider another interpretation of (9). In many data-generating processes,

$$\mathbb{E}_P \left[\min_{\theta} \text{Risk}(\theta, \mathbf{Y}, \mathbf{y}(1)) \right]$$

may not be small, because the realized data \mathbf{Y} may contain certain unforecastable components. The purpose of this section is to leverage the decomposition

$$\mathbb{E}_P[(\hat{y}_{0t} - y_{0t})^2] = \mathbb{E}_P[\epsilon_t^2] + \mathbb{E}_P[(\hat{y}_{0t} - \mu_t)^2],$$

where $\epsilon_t = y_{0t} - \mu_t$ is some unforecastable component satisfying $\mathbb{E}_P[\epsilon_t \hat{y}_{0t}] = 0$. This decomposition breaks prediction errors into forecastable and unforecastable components. Because of this additive decomposition, under certain conditions on ϵ_t , we can interpret risk differences as regret on estimating the forecastable component μ_t (since $\mathbb{E}_P[\epsilon_t^2]$ cancels in the difference). We can also decompose risk into the oracle error on estimating μ_t , the regret against the oracle on estimating μ_t , and the variance of the unforecastable errors ϵ_t .

For a fixed θ , under uniform treatment timing we have that

$$\mathbb{E}_P[\text{Risk}(\theta, \mathbf{Y}, \mathbf{y}(1))] = \mathbb{E}_P[\mathbb{E}_S(y_{0S} - \mu_S)^2] + \mathbb{E}_P[\mathbb{E}_S(\theta' \mathbf{y}_S - \mu_S)^2]$$

for *some* mean component μ_t , possibly random, of the outcome process y_{0t} . For instance, we may take $\mu_t = \mathbb{E}_P[y_{0t} \mid \mathbf{Y}_{1:t-1}, \mathbf{y}_t]$. For this μ_t , we can also write

$$\mathbb{E}_P[\text{Risk}(\sigma, \mathbf{Y}, \mathbf{y}(1))] = \mathbb{E}_P[\mathbb{E}_S(y_{0S} - \mu_S)^2] + \mathbb{E}_P[\mathbb{E}_S(\hat{\theta}'_t \mathbf{y}_S - \mu_S)^2],$$

since $\hat{\theta}'_t \mathbf{y}_t$ depends solely on $\mathbf{Y}_{1:t-1}, \mathbf{y}_t$. We thus have the following implication of (9)

$$\mathbb{E}_P[\mathbb{E}_S(\hat{\theta}'_t \mathbf{y}_S - \mu_S)^2] - \min_{\theta \in \Theta} \mathbb{E}_P[\mathbb{E}_S(\theta' \mathbf{y}_S - \mu_S)^2] \leq \frac{1}{T} \sup_{\|\mathbf{Y}\|_\infty \leq 1} \text{Regret}_T(\sigma; \mathbf{Y}),$$

which says that the risk difference of estimating the conditional mean μ_t —the forecastable component of the outcome process—is upper bounded by the regret. As a corollary, if $P = P_T$ is a sequence of data-generating processes where, as $T \rightarrow \infty$,

$$\min_{\theta \in \Theta} \mathbb{E}_P[\mathbb{E}_S(\theta' \mathbf{y}_S - \mu_S)^2] \rightarrow 0,$$

then we obtain a consistency result for synthetic control, in that

$$\mathbb{E}_P[\mathbb{E}_S(\hat{\theta}'_t \mathbf{y}_S - \mu_S)^2] \rightarrow 0$$

as well.

Shifting from risk differences to risks themselves, this means that the treatment effect estimation risk for synthetic control admits the following upper bound

$$\begin{aligned} \mathbb{E}_P[\text{Risk}(\sigma, \mathbf{Y}, \mathbf{y}(1))] &\leq \min_{\theta \in \Theta} \mathbb{E}_P[\mathbb{E}_S(\theta' \mathbf{y}_S - \mu_S)^2] + \frac{1}{T} \sup_{\|\mathbf{Y}\|_\infty \leq 1} \text{Regret}_T(\sigma; \mathbf{Y}) \\ &\quad + \mathbb{E}_P[\mathbb{E}_S(y_{0S} - \mu_S)^2], \end{aligned}$$

where the first term is the best possible error on the forecastable component μ_t , the second term is the average regret, and the third term is the variance of the unforecastable component that cannot be improved upon. We think the first two terms are likely small, and the last term is unavoidable.

This argument also extends to non-uniformly random treatment timing. Suppose we have a joint distribution Q of $(\mathbf{Y}, \mathbf{y}(1), S)$ such that $\pi_t(\mathbf{Y}) = Q(S = t \mid \mathbf{Y}) \leq C/T$. Suppose further that $y_{0t} = \mu_t + \epsilon_t$, where $\mathbb{E}_Q[\epsilon_t \mid \mu_t, \pi_t, \mathbf{Y}_{1:t-1}, \mathbf{y}_t] = 0$ for some mean component

μ_t .³⁵ Then we have a similar decomposition of the risk of estimating the treatment effect at S :

$$\begin{aligned}
\mathbb{E}_Q[(y_{0S} - \hat{\theta}'_S \mathbf{y}_t)^2] &= \sum_{t=1}^T \mathbb{E}_Q[\pi_t(\mathbf{Y})(y_{0t} - \hat{\theta}'_S \mathbf{y}_t)^2] \\
&= \sum_{t=1}^T \mathbb{E}_Q[\pi_t(\mathbf{Y})(y_{0t} - \mu_t)^2] + \mathbb{E}_Q[\pi_t(\mathbf{Y})(\mu_t - \hat{\theta}'_t \mathbf{y}_t)^2] \\
&\quad + 2\mathbb{E}_Q[\pi_t \epsilon_t (\mu_t - \hat{\theta}'_t \mathbf{y}_t)] \\
&= \mathbb{E}_Q[\epsilon_S^2] + \mathbb{E}_Q[(\mu_S - \hat{\theta}'_S \mathbf{y}_S)^2] \quad (\text{Last term is zero}) \\
&\leq \mathbb{E}_Q[\epsilon_S^2] + \frac{C}{T} \sum_{t=1}^T \mathbb{E}_Q[(\mu_t - \hat{\theta}'_t \mathbf{y}_t)^2] \quad ((1, \infty)\text{-H\"older's inequality}) \\
&\leq \mathbb{E}_Q[\epsilon_S^2] + C \left(\min_{\theta \in \Theta} \frac{1}{T} \sum_{t=1}^T \mathbb{E}_Q[(\mu_t - \theta' \mathbf{y}_t)^2] \right. \\
&\quad \left. + \frac{1}{T} \sup_{\|\mathbf{Y}\|_\infty \leq 1} \text{Regret}_T(\sigma; \mathbf{Y}) \right).
\end{aligned}$$

The last right-hand side is equal to the variance of the unforecastable component ϵ_S plus C times the oracle risk on estimating the mean component, as well as $O(NT^{-1} \log T)$ regret. If the oracle risk for estimating the mean component is small, then synthetic control is close to optimal, and its risk on estimating the mean component $\mathbb{E}_Q[(\mu_S - \hat{\theta}'_S \mathbf{y}_S)^2]$ is also small.³⁶

³⁵We can take $\mu_t = \mathbb{E}[y_{0t} \mid \mathbf{y}_t, \mathbf{Y}_{1:t-1}]$ whenever $S \perp\!\!\!\perp \mathbf{Y}$ under Q .

³⁶Note that the bound

$$\mathbb{E}_Q[(y_{0S} - \hat{\theta}'_S \mathbf{y}_t)^2] \leq C \left(\mathbb{E}_Q[\epsilon_S^2] + \min_{\theta \in \Theta} \frac{1}{T} \sum_{t=1}^T \mathbb{E}_Q[(\mu_t - \theta' \mathbf{y}_t)^2] + \frac{1}{T} \sup_{\|\mathbf{Y}\|_\infty \leq 1} \text{Regret}_T(\sigma; \mathbf{Y}) \right)$$

is immediate and allows for $\mu_t = \mathbb{E}[y_{0t} \mid \mathbf{Y}_{1:t-1}, \mathbf{y}_t] = 0$, yet the scaled idiosyncratic risk $C\mathbb{E}_Q[\epsilon_S^2]$ may be large.