Synthetic Control As Online Linear Regression

Jiafeng Chen*
Department of Economics, Harvard University
Harvard Business School
jiafengchen@g.harvard.edu

April 7, 2022

Abstract

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. 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
JEL codes: C1, C44

*I thank Alberto Abadie, Isaiah Andrews, Avi Feller, Bruno Ferman, Wayne Gao, Ed Glaeser, Sukjin Han, Pat Kline, Scott Kominers, Namrata Narain, David Ritzwoller, Brad Ross, Jonathan Roth, Suproteem Sarkar, Jesse Shapiro, Neil Shephard, Elie Tamer, Jaume Vives, Davide Viviano, Chris Walker, Jennifer Walsh for discussions and encouragement. I also thank participants of the Harvard Graduate Workshop in Econometrics and participants of the Synthetic Control Workshop at the Princeton University Center for Statistics and Machine Learning for helpful comments.

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 guarantees for synthetic control are typically derived under a linear factor model 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 the guarantees formally hold under these outcome models, there is a wide sense of 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 U.S. state or a country—plausible outcome modeling may be quite challenging. Manski and Pepper (2018), in studying the effect of right-to-carry 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?" The linear factor model is indeed a general class of data-generating process, but in this pessimistic view, perhaps even such a model is implausible for the settings considered by many synthetic control studies.

As a result, existing methodological results seem to leave practitioners in a bit of a bind. On the one hand, synthetic control is intuitively appealing, and it is widely believed to have good properties under a variety of outcome models. On the other hand, perhaps existing outcome models that have proved sufficiently tractable for deriving statistical guarantees are not sufficiently plausible in common empirical settings. This paper contributes a few theoretical results that address this tension, as well as providing a novel interpretation of synthetic control methods. In particular, we seek guarantees for synthetic control that do not rely on any outcome model.

It is unlikely that nontrivial guarantees on the *performance* of synthetic control exist

[1] See 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).

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

[3] For 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."

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 time, 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. These *regret* guarantees bounds relative performance on average over time, which can be interpreted as expected loss under a design assumption of random treatment timing.

Our results arise by casting prediction with panel data as an instance of *online convex optimization*,[4] and by recognizing that synthetic control as an online regression algorithm known as *Follow-The-Leader* (FTL). Regret guarantees on FTL in the online convex optimization literature translate directly to guarantees of synthetic control against a class of alternative estimators. Since most results in online convex optimization are 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 effect 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 prediction 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 guarantees in the worst-case setting to guarantees on statistical risk in a design-based framework. We make the point that synthetic control algorithms *themselves* are no-regret online algorithms, and are in fact competitive against a wide class of matching or difference-in-differences estimators.

This paper proceeds as follows. Section 2 sets up the notation and the decision protocol, and presents our main results for synthetic control. Section 3 presents a number of 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.

[4] See Hazan (2019), Orabona (2019), Cesa-Bianchi and Lugosi (2006), and Shalev-Shwartz (2011) for introductions to online learning.

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, where we assume $T > N$. Unit 0 is first treated at some time $\tau \in \{1, \ldots, T\} \equiv [T]$. Estimating causal effects for the treated unit amounts to predicting the unobserved potential outcome of this unit, and so we focus on untreated potential outcomes. Let the full panel of untreated potential outcomes be $\mathbf{Y}$, where (i) $\mathbf{Y}_i = (y_{i1}, \ldots, y_{iT})'$ is the time series of unit i, (ii) $\mathbf{Y}_t = (y_{0t}, \ldots, y_{Nt})'$ is the vector of outcomes at time t, and (iii) $\mathbf{y}_t = (y_{1t}, \ldots, y_{Nt})'$ is the vector of control unit outcomes at time t. We let $\mathbf{y}(1) = [y_1(1), \ldots, y_T(1)]'$ denote the treated outcomes of unit 0, which is only observable for times $t \geq \tau$. The analyst is tasked with predicting $y_{0,\tau}$ from observed data, which typically consist of pre-treatment outcomes of unit 0 and outcomes of untreated units. Similar to the main analysis in Doudchenko and Imbens (2016), we do not consider covariates.

Synthetic control (Abadie and Gardeazabal, 2003; Abadie *et al.*, 2010), in its basic form, chooses convex weights $\hat{\theta}_\tau$ to minimize past prediction errors

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

where $\Theta \equiv \{(\theta_1, \ldots, \theta_N) \colon \theta_i \geq 0, 1'\theta = 1\}$ is the simplex. For a one-step-ahead forecast for $y_{0,\tau}$, synthetic control outputs the weighted average $\hat{y}_\tau = \hat{\theta}_\tau' \mathbf{y}_\tau$, and forms the treatment effect estimate $\hat{\mathrm{TE}}_\tau = y_\tau(1) - \hat{y}_\tau$.

Theoretical guarantees for treatment effect estimates generated by (1) 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, in the spirit of comments by Manski and Pepper (2018). In contrast with the usual outcome modeling approach, we instead consider a worst-case setting where the outcomes are generated by an adversary.[5] 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 $\mathbf{Y}$. Specifically, we consider the following protocol between an analyst and an adversary:

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

1. 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 $\theta \in \Theta$ that may depend on the past.[6] We refer to the maps $\sigma \equiv (\theta_t(\cdot) \colon t \in [T])$ as the agent's *strategy.*

2. The adversary chooses the matrix of outcomes $\mathbf{Y}$. We assume that the adversary cannot choose arbitrarily large outcomes, and without further loss of generality, assume $\|\mathbf{Y}\|_\infty \le 1$. Since we are interested in the worst case, the adversary may choose $\mathbf{Y}$ with knowledge of σ.

3. The treatment time τ is sampled uniformly from $\{1, \ldots, T\}$, regardless of the adversary's choices $\mathbf{Y}$.

4. The analyst suffers loss $\ell(\hat{y}_{0,\tau}, y_{0,\tau}) \equiv (\hat{y}_\tau - y_{0,\tau})^2$.

Under such a protocol, the analyst's expected squared loss, over random treatment timing, is the average loss

$$\mathbb{E}_\tau \left[(\hat{y}_\tau - y_{0,\tau})^2\right] = \frac{1}{T}\sum_{t=1}^T (y_{0,t} - \hat{y}_t)^2 = \frac{1}{T}\sum_{t=1}^T (y_{0,t} - \theta_t'\mathbf{y}_t)^2. \tag{2}$$

We view the random sampling of treatment timing as a design-based perspective (Doudchenko and Imbens, 2016; Bottmer *et al.*, 2021) on the panel causal inference problem, which enables us to interpret average prediction loss over time as expected prediction loss of the treatment time τ. Of course, uniformly random assignment is restrictive, but we shall relax this requirement in Section 2.1 and Appendix B.[7]

We now make clear the connection with online convex optimization. Online convex optimization works with the following general protocol. At time t, an online player chooses some $\theta_t \in \Theta$, where $\Theta \subset \mathbb{R}^d$ is a bounded convex set. After θ is revealed, an adversary chooses a (convex) cost function $\ell_t : \Theta \to \mathbb{R}$, and the player suffers $\ell_t(\theta_t)$ and observes $\ell_t(\cdot)$. At the end of the game, the total loss suffered by the online player is $\sum_{t=1}^T \ell_t(\theta_t)$.

Our setup of the panel prediction protocol is then an instance of online convex optimization. To wit, first, in both protocols, the player makes decisions θ_t in a sequential manner; restriction to the simplex makes Θ convex. Second, the adversary in the panel prediction game may be thought of as picking loss functions $\ell_t(\cdot)$ of the form $\theta \mapsto \frac{1}{2}(y_{0,t} - \mathbf{y}_t'\theta)^2$, indexed by $(y_{0,t}, \mathbf{y}_t)$. Third, the expected loss in the panel prediction game, (2), is equal to the average loss $\frac{1}{T}\sum_{t=1}^T \ell_t(\theta_t)$, which is simply the total loss scaled by $1/T$.

[6] For simplicity, we do not consider randomized predictions, but randomized algorithms admit a similar analysis.

[7] Our decision framework 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 $\theta \mapsto \ell(f(\mathbf{y}_t, \theta), y_{0,t})$ is convex and bounded.

The main observation of this paper recognizes that synthetic control is an online learning algorithm known as *Follow-the-Leader* (FTL). FTL is the algorithm that simply chooses θ_t to minimize past losses:[8]

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

Observation 1. FTL is equivalent to synthetic control (1) under the panel prediction protocol.

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:

$$\begin{aligned}
\text{Regret}_T(\sigma; \mathbf{Y}) &\equiv \sum_{t=1}^{T} \ell_t(\theta_t) - \min_{\theta \in \Theta} \sum_{t=1}^{T} \ell_t(\theta) && (3)\\
&= \sum_{t=1}^{T} (y_{0,t} - \mathbf{y}_t'\theta_t)^2 - \min_{\theta \in \Theta} \sum_{t=1}^{T} (y_{0,t} - \mathbf{y}_t'\theta)^2 && (4)\\
&= T\left(\mathbb{E}_\tau[(y_{0,\tau} - \mathbf{y}_\tau'\theta_\tau)^2] - \min_{\theta \in \Theta} \mathbb{E}_\tau[(y_{0,\tau} - \mathbf{y}_\tau'\theta)^2]\right) && (5)\\
&\geq T\left(\mathbb{E}_\tau[(y_{0,\tau} - \mathbf{y}_\tau'\theta_\tau)^2] - \mathbb{E}_\tau[(y_{0,\tau} - \mathbf{y}_\tau'\theta)^2]\right) \text{ for any } \theta \in \Theta. && (6)
\end{aligned}$$

(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. (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 $\arg\min_{\theta \in \Theta} \sum_{t=1}^{T} (y_{0,t} - \mathbf{y}_t'\theta)^2$ as the *oracle weighted match*.

Focusing on regret rather than loss shifts the goalpost from performance to *competition*, which is a more fruitful perspective in the adversarial framework. 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. $\text{Regret}_T = o(T)$,[9] then the strategy σ never performs much worse as any fixed weights θ on average. In this case, we can interpret σ as a strategy that is *competitive* against the class of weighted matches. It can seem quite surprising that these no-regret strategies σ exist in the first place; we emphasize that σ can

[8] Also known as fictitious play in game theory (Brown, 1951).

[9] We sometimes refer to σ as no-regret if it has sublinear regret.

output different weights θ_t, chosen adaptively over time, while it is compared to fixed weights θ.

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

Theorem 2.1 (Theorem 5, Hazan *et al.* 2007)**.** *Under the setting of Observation 1 with bounded outcomes* $\|\mathbf{Y}\|_\infty \le 1$, *synthetic control* (1) *satisfies the regret bound*[10]

$$\text{Regret}_T(\sigma, \mathbf{Y}) \le 16N(\log(2\sqrt{N}T) + 1) = O(N \log T).$$

Theorem 2.1 shows that the unregularized synthetic control estimator (1) achieves logarithmic regret—and as a result, the average difference between the synthetic control losses and losses of the oracle weighted match vanishes quickly as a function of T.[11] In particular, if there exists a weighted average of the untreated unit's outcomes that track $\mathbf{Y}_0$ 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, Theorem 2.1 implies that the *risk* of estimating the causal effect at time τ for synthetic control is not too much higher than that for any weighted matching estimator; in particular, if any weighted matching estimator performs well, then synthetic control achieves low risk of estimating the causal effect. Our ensuing discussion interprets the guarantees as guarantees on the expected loss at treatment time (expressing regret as (5)), which relies on a design assumption that τ is randomly assigned. Nevertheless, we stress that we could view purely as guarantees of average loss over time (expressing regret only as (4)), which does not require a treatment timing assumption.

We can interpret regret as a gap in the *risk* of estimating treatment effects. In particular, we can interpret the expected loss of predicting the untreated outcome as the risk of estimating

[10] We say $f(N,T) = O(g(N,T))$ if, for any sequence $N_T < T$ and $T \to \infty$,

$$\limsup_{T\to\infty} \left| \frac{f(N_T, T)}{g(N_T, T)} \right| < \infty.$$

[11] The restriction of Θ as the simplex—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 $\mathbf{y}_t'\theta$ that is free of N, T. In contrast, if we let $\Theta = \{\theta : \|\theta\|_2 \le D/2\}$ be an ℓ_2-ball, then the regret bound worsens to $O(D^2 N^2 \log(T))$.

the treatment effect:

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

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,[12] regardless of realized outcomes $\mathbf{Y}$:

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

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

Of course, since the guarantee (8) holds for every $\mathbf{Y}$, it continues to hold when we average over $\mathbf{Y}, \mathbf{y}(1)$, over a joint distribution P that respects the boundedness condition $\|\mathbf{Y}\|_\infty \le 1$. 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.[13] Formally, let $\mathcal{P}$ be a family of distributions for $\mathbf{Y}, \mathbf{y}(1)$ such that $P(\|\mathbf{Y}\|_\infty \le 1) = 1$ 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 that[14]

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

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

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

[13]The 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.

[14]Abernethy *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 $\le$ 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, regret analysis is not by itself conservative for a stochastic setting—minimax regret is a tight upper bound for performance in stochastic settings.

the best oracle weighted match

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

Hence, if the data-generating process P guarantees that *if 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 can find a weighted match that tracks the potential outcomes $y_{0,1}, \ldots, y_{0,T}$ well, so that[15]

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

In many empirical settings, it seems plausible that the oracle weighted match performs well.[16] 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, the 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 $y_{0,t}$ closely, synthetic control continues to enjoy the assurance that it performs almost as well as the best weighted match. Moreover, this no-regret property cannot be attained without choosing θ_t in some data-dependent manner.[17] This observation rules out alternatives such as simple difference-in-differences, which does not aggregate in a data-dependent manner. 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 very popular class of estimators in practice.

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

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

[17] Roughly speaking, a strategy that picks some fixed $\theta \in \Theta$ suffers from an adversary who picks the data such that $\ell_t(\theta) - \ell_t(\tilde{\theta}) \ge c > 0$, for some $\tilde{\theta} \in \Theta$. Under such a configuration, the constant strategy would suffer $O(T)$ total regret.

2.1 Non-uniform treatment timing

The previous interpretations rely on interpreting average loss over time as expected loss over τ ((5) and (7)), which requires uniform treatment timing $\tau \sim \mathrm{Unif}[T]$. Despite being plausible in certain settings and appearing elsewhere in the literature (Doudchenko and Imbens, 2016; Bottmer *et al.*, 2021), this assumption is perhaps not entirely palatable.[18] Since we are being completely agnostic on the outcome generation process, it is unavoidable to make treatment timing assumptions in order to obtain nontrivial results on estimation of causal quantities. Nevertheless, note that such an assumption is only necessary to interpret average losses as expected losses. The a priori position that *it is reasonable to expect a causal estimator to predict well relative to a set of oracles, at least on average over time*, strikes us as defensible. Accepting this dictum relieves us of any need of treatment timing modeling.

Even if we wish to maintain the interpretation of total 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 distribution is non-uniform, unknown, and possibly chosen by the adversary, we continue to show that synthetic control performs well as long as some weighted average of untreated units predicts $y_{0,\tau}$ accurately.

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

$$\hat{\theta}_t^{\pi} \in \arg\min_{\theta \in \Theta} \sum_{s<t} \pi_s (y_{0,s} - \mathbf{y}_s' \theta)^2. \tag{10}$$

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

Corollary 2.2. *Suppose $\tau \sim \pi$, $\frac{1}{CT} \le \pi_t \le \frac{C}{T}$ for some C, and $\|\mathbf{Y}\|_\infty \le 1$. Then weighted synthetic control* (10)*, denoted σ_π, achieves weighted regret bound*

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

Corollary 2.2 shows that the weighted regret—a difference in π-expected loss—is logarith-

[18] We note that the randomness *per se* of τ conditional on $\mathbf{Y}$ can be realistic, but that its distribution is uniform and known is restrictive.

mic in T, thereby controlling the worst-case gap between weighted synthetic control and the best oracle weighted match for the expected loss.

Assuming a known π could be reasonable. With a known dynamic treatment regime, π can depend on $\mathbf{Y}$, but is known whenever the analyst is prompted for a prediction.[19] We can also interpret Corollary 2.2 as providing guarantees on gaps in Bayes risk under the analyst's prior $\tau \sim \pi$, independent of $\mathbf{Y}$.

Even when π is *unknown*, we can bound the loss of unweighted synthetic control.

Corollary 2.3. *Suppose $\tau \sim \pi$, $\pi_t \le C/T$ for some C, and $\|\mathbf{Y}\|_\infty \le 1$. Then synthetic control* (1) *achieves the loss bound*

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

Hence, for any joint distribution Q of $(\mathbf{Y}, \tau)$ where $Q(\tau = t \mid \mathbf{Y}) \le C/T$ for all t, and $Q(\|\mathbf{Y}\|_\infty \le 1) = 1$, we have the average loss bound

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

The result (12) shows that, uniformly over all bounded $\mathbf{Y}$ and bounded treatment distribution 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 $\pi_t \le C/T$, but otherwise unrestricted in the dependence of $\mathbf{Y}$ 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, $\mathbb{E}_Q\left[\min_{\theta\in\Theta}\frac{1}{T}\sum_{t=1}^{T}(y_{0,t} - \theta'\mathbf{y}_t)^2\right]$,

[19] Since 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 2.2 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,

$$\mathrm{P}(\tau = t \mid \tau \ge t, \mathbf{Y}) = r_t(\mathbf{Y}_{1:t-1})$$

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

is small—so long as τ has sufficient randomness conditional on $\mathbf{Y}$.

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

3 Extensions

3.1 Competing against DID

The previous section 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 section, we show that synthetic control on differenced data is competitive against a large class of DID estimators.

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

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

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 $y_\tau(1) - y_{0,\tau}$. Implicitly, this regression uses the estimated fixed effects $\mu_0 + \alpha_\tau$ as a forecast for the unobserved $y_{0,\tau}$.

We consider a weighted generalization of this regression:[20]

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

For convex weights $w = (w_1, \ldots, w_N)$, 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,[21] and outputs the estimated coefficients $\mu_0 + \alpha_t$ as a prediction for $y_{0,t}$. By varying over $w \in \Theta$, we obtain a class of competing DID strategies, where canonical DID corresponds to picking uniform weights $w = (1/N, \ldots, 1/N)'$. We

[20]The weight w_0 does not affect $\mu_0 + \alpha_\tau$ achieving the optimum in the least-squares problem, per the calculation in Appendix A.5. As a result, we normalize $w_0 = 1$.

[21]The value of $y_{0,t}$ does not enter $\alpha_\tau + \mu_0$ since it is absorbed by the coefficient λ.

calculate in Appendix A.5 that the prediction that $\sigma_{\mathrm{TWFE}}(w)$ makes is

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

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*

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

to forecast the same difference of unit 0, $\tilde{y}_{0,t}$. Therefore, we may apply Theorem 2.1 and show the following regret bound.

Theorem 3.1. *Consider synthetic control on the differenced data, where the analyst computes*

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

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

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

for some constant C*.*

Theorem 3.1 shows that synthetic control on differenced data controls regret against a class of DID estimators (14) that may output predictions in a sequential manner as well, in contrast to a similar result in Proposition A.2. 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 3.1 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 3.1 becomes a performance guarantee on synthetic control. Moreover, since (14) is a popular alternative in many empirical settings—setting aside whether there is a weighted DID that performs well—Theorem 3.1 shows that it is without (too much) loss to use synthetic control in such settings instead.[22]

To the best of our knowledge, the difference scheme $\tilde{y}_{it}$ has yet to be considered in the literature. We note that $\tilde{y}_{it}$ is slightly different from Ferman and Pinto (2021)'s demeaned synthetic control, which takes the difference $\check{y}_{it} \equiv y_{it} - \frac{1}{t}\sum_{s=1}^{t} y_{is}$. In Appendix A.2, 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. Another 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 $\hat{y}_t(\sigma_{\text{DID}}(\theta)) \equiv y_{0,t-1} + \theta' (\mathbf{y}_t - \mathbf{y}_{t-1})$ as successive predictions.

3.2 Regularization and other extensions

Theorem 2.1 shows that synthetic control / 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, such strategies take the form

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

for some convex penalty Φ. 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 FTRL, where $\ell(\cdot,\cdot)$ is typically quadratic loss.

[22]Under certain conditions, Ferman and Pinto (2021) (Proposition 3) show that the demeaned synthetic control in Proposition A.2 dominates DID with uniform weighting $\theta_i = 1/N$. The results Proposition A.2 and Theorem 3.1 are in a similar flavor, and show that synthetic control is competitive against DID with any fixed weighting, on average over random assignment of treatment time. Of course, Proposition A.2 and Theorem 3.1 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.

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.[23] 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.3 for a general statement.

Theorem 3.2. *Consider regularized synthetic control / FTRL with penalty function* $\Phi(\theta)$*. Let* $\ell(\mathbf{y}_t'\theta, y_{0,t})$ *be a convex loss function (in* θ*), not necessarily quadratic.*

1. *For the ridge penalty* $\Phi(\theta) = \frac{1}{2}\|\theta\|^2$*, for both squared loss* $\frac{1}{2}(y-\hat{y})^2$ *and linear loss* $|y-\hat{y}|$*, we have* $\text{Regret}_T \le 3\sqrt{NT}$ *with the choice* $\eta = 1/\sqrt{NT}$*.*

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

Naturally, these choices correspond to regularized variants of synthetic control. Quadratic penalties correspond to ridge penalization in Hirshberg (2021), and is a special case of elastic net penalty proposed by Doudchenko and Imbens (2016).[24] 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).[25] Additionally, Theorem 3.2 gives guidance on choosing the regularization strength η for different estimators, which depends only on N, T, and the bound on $\mathbf{Y}$.[26]

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 $[T]$, or *adaptive regret*. We show in Appendix B 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 a design-based framework on treatment timing, we can test sharp hypotheses of the form

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

[24] Theorem A.3 applies to elastic net penalties with nonzero ℓ_2 component as well.

[25] Interestingly, 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.3 does not apply. However, Theorem A.3 only contains sufficient conditions, and so this alone is not a criticism of ℓ_1-penalty.

[26] The dependence on T may be relaxed via the "doubling trick" (see Shalev-Shwartz (2011) section 2.3.1), if we allow for regularization that depends on τ.

$H_0 : [y_1(1) - y_{0,1}, \dots, y_T(1) - y_{0,T}] = [z_1, \dots, z_T]$ by leveraging symmetries induced by random treatment timing. We discuss inference in Appendix C.

4 Conclusion

This paper notes a very simple connection between synthetic control methods and online learning. Synthetic control is an instance of Follow-The-Leader strategies, which are well-studied 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. Under conditions 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. 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. 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 also 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 $\mathbf{y}_{1:T}$ is fully accessbile 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.

Likewise, we have only considered regret on one-step-ahead prediction for $y_{0,\tau}$, but synthetic control estimates are often extrapolated multiple time periods ahead in practice. Multi-step-ahead prediction seems closely related to the problem of delayed feedback in online learning (Weinberger and Ordentlich, 2002; Korotin *et al.*, 2018; Flaspohler *et al.*, 2021), where the learner is prompted for the prediction of $y_{0,t+D}$ at time t, and $y_{0,t+D}$ is revealed only at time $t + D$.

A Proofs and additional results

A.1 Proofs of Theorem 2.1 and Corollaries 2.2 and 2.3

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 $\ell_t : \Theta \to \mathbb{R}$ can be written as*

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

for a univariate convex function $g_t : \mathbb{R} \to \mathbb{R}$ 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 $g_t''(v_t'\theta) \geq a$. Then FTL on ℓ_t satisfies the following regret bound:

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

where $D = \text{diameter}(\Theta) = \max_{x,y\in\Theta}\|x - y\|_2$.

Theorem 2.1 follows immediately from Theorem 5 in Hazan *et al.* (2007), reproduced in our notation as Theorem A.1. Since Θ is the simplex, we know $D \leq 2$. We choose $g_t(x) = \frac{1}{2}(y_{0,t} - x)^2$ with $g_t'(x) = x - y_{0,t}$ 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 are $\mathbf{y}_t$, whose dimensions are $n = N$ and whose 2-norms are bounded by $R = \sqrt{N}$. Note that $|v_t'\theta| \leq \|v_t\|_\infty \|\theta\|_1 \leq 1$. Hence $|g_t'(v_t'\theta)| \leq 2 \equiv b$ and $g_t''(x) \geq 1 \equiv a$. Plugging in to obtain

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

which rearranges into the claim.

The proof for Corollary 2.2 follows similarly. 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 2.2.

For Corollary 2.3, 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

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

We then apply (12) to complete the proof.

A.2 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 estimators proposed by Doudchenko and Imbens (2016) and Ferman and Pinto (2021) in efforts to mimick behavior of DID estimators.[27] Our regret bound from Theorem 2.1 generalizes immediately to the affine predictions, where the benchmark oracle that the regret measures against is

$$\min_{(\theta_0, \theta_1) \in \Theta} \sum_{t=1}^T (y_{0,t} - \theta_0 - \mathbf{y}_t' \theta_1)^2. \tag{16}$$

(16) simultaneously chooses the best intercept and the best set of convex weights in hindsight. Because (16) 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.2. *Consider demeaned synthetic control, where the analyst outputs the prediction* $\hat{y}_t = \hat{\theta}_{0t} + \mathbf{y}_t' \hat{\theta}_t$ *via solving the least-squares problem*

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

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

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

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

for some constant C.

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.2 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 $|v_t'\vartheta| \le 3$. Hence we may take $b = 3+1 = 4$ and $a = 1$. Plugging in to obtain

$$\text{Regret}_T \le 32N\left[\log\left(\frac{\sqrt{5}}{2}\sqrt{N+1}T\right)+1\right] < CN\log T$$

for some C. □

The set of competitors for synthetic control in Proposition A.2 is constrained to use the same θ_0 in making predictions for each time period, and this may be a limitation.

A.3 Proof of Theorem 3.1

Similarly to the proof of Proposition A.2, suppose the adversary picks the differences $|\tilde{y}_{it}| \le 2$, *without* the constraint that the differences obey the restriction $\|\mathbf{Y}\|_\infty \le 1$. An application of Theorem A.1 shows that

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

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

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

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

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

A.4 Proof of Theorem 3.2

Theorem A.3. *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. *The gradients* $\nabla_\theta\ell_t(\theta)$ *are bounded in the dual norm* $\|\cdot\|_*$*, uniformly over* $\Theta,\mathbf{Y}$*:*

$$\|\nabla_\theta\ell_t(\theta)\|_*^2\le G.$$

Then FTRL attains the regret bound

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

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

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

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

Theorem A.4 (Corollary 7.9, Orabona 2019)**.** *Let* ℓ_t *be a sequence of convex loss functions. Let* $\Phi:\Theta\to\mathbb{R}$ *be* μ*-strongly convex w.r.t. the norm* $\|\cdot\|$*. Then, FTRL guarantees*

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

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

Theorem A.3 then follows immediately where $\|g_t\|_*^2\le G$, $\Phi(\theta)\le K$, and $\mu=1$.

A.4.1 Proof of Theorem 3.2

For both squared and absolute losses, we can bound the gradient of the loss function in terms of

$$\|\nabla_\theta \ell_t(\theta)\| \le |y_{0,t} - \mathbf{y}_t'\theta| \|\mathbf{y}_t\| \le 2\|\mathbf{y}_t\|$$

under any norm. Hence we should pick G to bound $4\|\mathbf{y}_t\|_*^2$

For the ridge penalty $\Phi(\theta) = \frac{1}{2}\|\theta\|_2^2$, it is 1-strongly convex with respect to $\|\cdot\|_2$. Thus we take $G = 4N$. The ridge penalty is bounded by $K = 1$ on the simplex. This yields the bound via Theorem A.3.

The entropy penalty is 1-strongly convex with respect to $\|\cdot\|_1$. 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.3.

A.5 Two-way fixed effect calculation

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

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

We may eliminate $(i,t) = (0,\tau)$ from the sum since $\lambda \mathbb{1}(i = 0, \tau = t)$ absorbs that term, leaving μ_i, α_t unaffected. Consider forecasting $y_{0,\tau}$ with $\mu_0 + \alpha_\tau$ that solves the above program.

The first-order condition for μ_i takes the form

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

Hence

$$\mu_i = \begin{cases} \overline{y}_{it} - \overline{\alpha} & i \neq 0 \\ \overline{y}_{0t} - \frac{\tau}{\tau-1}\overline{\alpha} + \frac{1}{\tau-1}\alpha_\tau & i = 0 \end{cases}$$

where $\overline{\alpha} = \frac{1}{\tau}\sum_{t=1}^{\tau} \alpha_t$ and $\overline{y}_{it}$ is the sample mean of observations for unit i, with the understanding that $y_{0\tau}$ is not observed.

Hence the forecast is

$$\mu_0 + \alpha_\tau = \overline{y}_{0t} + \frac{\tau}{\tau - 1}(\alpha_\tau - \overline{\alpha}).$$

Let us inspect the first-order condition for α_τ:

$$\sum_{i=1}^{N} w_i(y_{i\tau} - \mu_i - \alpha_\tau) = \sum_{i=1}^{N} w_i(y_{i\tau} - \overline{y}_{i\tau} + \overline{\alpha} - \alpha_\tau) = 0.$$

Hence

$$\alpha_\tau - \overline{\alpha} = \sum_{i=1}^{N} w_i \left(\frac{\tau - 1}{\tau} y_{i\tau} - \frac{1}{\tau} \sum_{t=1}^{\tau-1} y_{it} \right).$$

Therefore,

$$\frac{\tau}{\tau - 1}(\alpha_\tau - \overline{\alpha}) = \sum_{i=1}^{N} w_i \left(y_{i\tau} - \frac{1}{\tau - 1} \sum_{t=1}^{\tau-1} y_{it} \right).$$

Thus the forecast is

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

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

B Adaptive regret

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

$$\text{AdaptiveRegret}_T = \sup_{1 \le r < s \le T} \sum_{t=r}^{s} \left\{ \ell_t(\theta_t) - \min_{\theta_{r,s}} \sum_{t=r}^{s} \ell_t(\theta_{r,s}) \right\}, \tag{17}$$

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

$$\text{AdaptiveRegret}_T(\text{FLH}(\sigma)) \le \text{Regret}_T(\sigma) + O(\log T). \tag{18}$$

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 $y_{0,t}$:

1. Consider the synthetic control estimated weights $\theta_t^1, \ldots, \theta_t^t$, where θ_t^j is the synthetic

control weights estimated based on data from *time horizons* $j, \ldots, t-1$.

2. Output the weighted average $\theta_t = \sum_{j=1}^t p_t^j \theta_t^j$.

3. After receiving $\mathbf{y}_t, y_{0,t}$ (and hence receiving $\ell_t(\theta) = \frac{1}{2}(y_{0,t} - \theta' \mathbf{y}_t)^2$), 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 (18) for the above algorithm.[28] 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.

C A note on inference

Under the treatment assignment model $\tau \sim \text{Unif}[T]$, we may test the sharp null $H_0 : \mathbf{y}(1) = \mathbf{Y}_0$. Let $y_t = y_{0,t}$ for $t < \tau$ and let $y_t = \mathbf{y}_t(1)$ for $t \geq \tau$ be the observed time series of the treated unit. For any prediction $\hat{y}_t$ that does not depend on τ—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_τ is at least the $\lceil T(1-\alpha) \rceil^{\text{th}}$ order statistic of the sample $\{r_1, \ldots, r_T\}$.

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

Since, under the null, r_τ is equally likely to equal any of $\{r_1, \ldots, r_T\}$, the probability of it being the largest α is bounded by α. Similarly, if $\tau \sim \pi$ where $\pi_t \le C/T$, a least-favorable test may be constructed by rejecting when $r_t \ge r_{(T-\lfloor T\alpha/C \rfloor)}$. 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, from Markov's inequality, we can control the probability for the prediction error to deviate far relative to its expectation

$$\mathrm{P}_{\tau \sim \mathrm{Unif}[T]}\left[(y_{0,\tau} - \hat{y}_\tau)^2 > c\right] \le \frac{\mathbb{E}_\tau[\ell_\tau(\theta_\tau)]}{c} \le \frac{1}{c}\left(\min_{\theta \in \Theta} \frac{1}{T}\sum_{i=1}^T \ell_t(\theta) + \frac{1}{T}\mathrm{Regret}_T\right).$$

Under assumptions where the pre-treatment loss $\min_\theta \frac{1}{\tau-1}\sum_{t<\tau} \ell_t(\theta)$ 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.

D Risk interpretation under idiosyncratic errors

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

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

may not be small. Nevertheless, for a fixed θ, under uniform treatment timing we have that

$$\mathbb{E}_P[\mathrm{Risk}(\theta, \mathbf{Y}, \mathbf{y}(1))] = \mathbb{E}_P[\mathbb{E}_\tau(y_{0,\tau} - \mu_\tau)^2] + \mathbb{E}_P[\mathbb{E}_\tau(\theta'\mathbf{y}_\tau - \mu_\tau)^2]$$

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

$$\mathbb{E}_P[\mathrm{Risk}(\sigma, \mathbf{Y}, \mathbf{y}(1))] = \mathbb{E}_P[\mathbb{E}_\tau(y_{0,\tau} - \mu_\tau)^2] + \mathbb{E}_P[\mathbb{E}_\tau(\hat{\theta}_t'\mathbf{y}_\tau - \mu_\tau)^2],$$

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)

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

which says that the risk difference of estimating the conditional mean μ_t is upper bounded by the regret. If $P = P_T$ is a sequence of data-generating processes where, as $T \to \infty$,

$$\min_{\theta \in \Theta} \mathbb{E}_P[\mathbb{E}_\tau(\theta' \mathbf{y}_\tau - \mu_\tau)^2] \to 0,$$

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

$$\mathbb{E}_P[\mathbb{E}_\tau(\hat{\theta}_t' \mathbf{y}_\tau - \mu_\tau)^2] \to 0$$

as well.

From a risk perspective, this means that the treatment effect estimation risk for synthetic control admits the following upper bound

$$\mathbb{E}_P[\text{Risk}(\sigma, \mathbf{Y}, \mathbf{y}(1))] \le \min_{\theta \in \Theta} \mathbb{E}_P[\mathbb{E}_\tau(\theta' \mathbf{y}_\tau - \mu_\tau)^2] + \frac{1}{T} \sup_{\|\mathbf{Y}\|_\infty \le 1} \text{Regret}_T(\sigma; \mathbf{Y}) + \mathbb{E}_P[\mathbb{E}_\tau(y_{0,\tau} - \mu_\tau)^2],$$

where the first two terms are likely small, and the last term represents the risk incurred by unforecastable randomness in $y_{0,t}$.

In general, suppose we have a joint distribution Q of $(\mathbf{Y}, \mathbf{y}(1), \tau)$ such that $\pi_t(\mathbf{Y}) = Q(\tau = t \mid \mathbf{Y}) \le C/T$. Suppose further that $y_{0,t} = \mu_t + \epsilon_t$, 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.[29] Then we have a similar decomposition of the risk of estimating the treatment effect at τ:

$$\begin{aligned}
\mathbb{E}_Q[(y_{0,\tau} - \hat{\theta}_\tau' \mathbf{y}_t)^2] &= \sum_{t=1}^T \mathbb{E}_Q[\pi_t(\mathbf{Y})(y_{0,t} - \hat{\theta}_\tau' \mathbf{y}_t)^2] \\
&= \sum_{t=1}^T \mathbb{E}_Q\left[\pi_t(\mathbf{Y})(y_{0,t} - \mu_t)^2\right] + \mathbb{E}_Q[\pi_t(\mathbf{Y})(\mu_t - \hat{\theta}_t' \mathbf{y}_t)^2] + 2\mathbb{E}_Q[\pi_t \epsilon_t (\mu_t - \hat{\theta}_t' \mathbf{y}_t)] \\
&= \mathbb{E}_Q[\epsilon_\tau^2] + \mathbb{E}_Q[(\mu_\tau - \hat{\theta}_\tau' \mathbf{y}_\tau)^2] \\
&\le \mathbb{E}_Q[\epsilon_\tau^2] + \frac{C}{T} \sum_{t=1}^T \mathbb{E}_Q[(\mu_t - \hat{\theta}_t' \mathbf{y}_t)^2] \\
&\le \mathbb{E}_Q[\epsilon_\tau^2] + C \left(\min_{\theta \in \Theta} \frac{1}{T} \sum_{t=1}^T \mathbb{E}_Q[(\mu_t - \theta' \mathbf{y}_t)^2] + \frac{1}{T} \sup_{\|\mathbf{Y}\|_\infty \le 1} \text{Regret}_T(\sigma; \mathbf{Y}) \right).
\end{aligned}$$

The last right-hand side is equal to the variance of the unforecastable component ϵ_τ 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

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

optimal, and its risk on estimating the mean component $\mathbb{E}_Q[(\mu_\tau - \hat{\theta}_\tau' \mathbf{y}_\tau)^2]$ is also small.

Note that the bound

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

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

References

Abadie, A. (2021). Using synthetic controls: Feasibility, data requirements, and methodological aspects. *Journal of Economic Literature*, **59** (2), 391–425.

— and Cattaneo, M. D. (2021). Introduction to the special section on synthetic control methods.

—, Diamond, A. and Hainmueller, J. (2010). Synthetic control methods for comparative case studies: Estimating the effect of california's tobacco control program. *Journal of the American statistical Association*, **105** (490), 493–505.

—, — and — (2015). Comparative politics and the synthetic control method. *American Journal of Political Science*, **59** (2), 495–510.

— and Gardeazabal, J. (2003). The economic costs of conflict: A case study of the basque country. *American economic review*, **93** (1), 113–132.

— and Vives-i-Bastida, J. (2021). *Synthetic Controls in Action*. Tech. rep.

Abernethy, J., Agarwal, A., Bartlett, P. L. and Rakhlin, A. (2009). A stochastic view of optimal regret through minimax duality. *arXiv preprint arXiv:0903.5328*.

Bartlett, P. L., Koolen, W. M., Malek, A., Takimoto, E. and Warmuth, M. K. (2015). Minimax fixed-design linear regression. In *Conference on Learning Theory*, PMLR, pp. 226–239.

Ben-Michael, E., Feller, A. and Rothstein, J. (2019). Synthetic controls and weighted event studies with staggered adoption. *arXiv preprint arXiv:1912.03290*.

—, — and — (2021). The augmented synthetic control method. *Journal of the American Statistical Association*, **116** (536), 1789–1803.

Blackwell, D. (1956). An analog of the minimax theorem for vector payoffs. *Pacific Journal of Mathematics*, **6** (1), 1–8.

Bottmer, L., Imbens, G., Spiess, J. and Warnick, M. (2021). A design-based perspective on synthetic control methods. *arXiv preprint arXiv:2101.09398*.

Boyd, S. P. and Vandenberghe, L. (2004). *Convex optimization*. Cambridge university press.

Brown, G. W. (1951). Iterative solution of games by fictitious play. *Activity analysis of production and allocation*, **13** (1), 374–376.

Cesa-Bianchi, N. and Lugosi, G. (2006). *Prediction, learning, and games*. Cambridge university press.

Chakraborty, B. and Murphy, S. A. (2014). Dynamic treatment regimes. *Annual review of statistics and its application*, **1**, 447–464.

Chernozhukov, V., Wüthrich, K. and Zhu, Y. (2021). An exact and robust conformal inference method for counterfactual and synthetic controls. *Journal of the American Statistical Association*, **116** (536), 1849–1864.

Doudchenko, N. and Imbens, G. W. (2016). *Balancing, regression, difference-in-differences and synthetic control methods: A synthesis*. Tech. rep., National Bureau of Economic Research.

Ferman, B. and Pinto, C. (2021). Synthetic controls with imperfect pretreatment fit. *Quantitative Economics*, **12** (4), 1197–1221.

Flaspohler, G. E., Orabona, F., Cohen, J., Mouatadid, S., Oprescu, M., Orenstein, P. and Mackey, L. (2021). Online learning with optimism and delay. In *International Conference on Machine Learning*, PMLR, pp. 3363–3373.

Hannan, J. (1958). Approximation to bayes risk in repeated play. In *Contributions to the Theory of Games (AM-39), Volume III*, Princeton University Press, pp. 97–140.

Hazan, E. (2019). Introduction to online convex optimization. *arXiv preprint arXiv:1909.05207*.

—, Agarwal, A. and Kale, S. (2007). Logarithmic regret algorithms for online convex optimization. *Machine Learning*, **69** (2-3), 169–192.

Hirshberg, D. A. (2021). Least squares with error in variables. *arXiv preprint arXiv:2104.08931*.

Korotin, A., V'yugin, V. and Burnaev, E. (2018). Aggregating strategies for long-term forecasting. In *Conformal and Probabilistic Prediction and Applications*, PMLR, pp. 63–82.

Manski, C. F. and Pepper, J. V. (2018). How do right-to-carry laws affect crime rates? coping with ambiguity using bounded-variation assumptions. *Review of Economics and Statistics*, **100** (2), 232–244.

Orabona, F. (2019). A modern introduction to online learning. *arXiv preprint arXiv:1912.13213*.

Robbins, M. W., Saunders, J. and Kilmer, B. (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** (517), 109–126.

Shalev-Shwartz, S. (2011). Online learning and online convex optimization. *Foundations and trends in Machine Learning*, **4** (2), 107–194.

Viviano, D. and Bradic, J. (2019). Synthetic learner: model-free inference on treatments over time. *arXiv preprint arXiv:1904.01490*.

Weinberger, M. J. and Ordentlich, E. (2002). On delayed prediction of individual sequences. *IEEE Transactions on Information Theory*, **48** (7), 1959–1976.