Efficient Difference-in-Differences and Event Study Estimators^{*}

Xiaohong Chen[†]

Pedro H. C. Sant'Anna[‡]

Haitian Xie[§]

June 21, 2025

Abstract

This paper investigates efficient Difference-in-Differences (DiD) and Event Study (ES) estimation using short panel data sets within the heterogeneous treatment effect framework, free from parametric functional form assumptions and allowing for variation in treatment timing. We provide an equivalent characterization of the DiD potential outcome model using sequential conditional moment restrictions on observables, which shows that the DiD identification assumptions typically imply nonparametric overidentification restrictions. We derive the semiparametric efficient influence function (EIF) in closed form for DiD and ES causal parameters under commonly imposed parallel trends assumptions. The EIF is automatically Neyman orthogonal and yields the smallest variance among all asymptotically normal, regular estimators of the DiD and ES parameters. Leveraging the EIF, we propose simple-to-compute efficient estimators. Our results highlight how to optimally explore different pre-treatment periods and comparison groups to obtain the tightest (asymptotic) confidence intervals, offering practical tools for improving inference in modern DiD and ES applications even in small samples. Calibrated simulations and an empirical application demonstrate substantial precision gains of our efficient estimators in finite samples.

^{*}Authors are listed in alphabetical order. We are grateful to Dmitry Arkhangelsky, Gregorio Caetano, Kevin Chen, Andrew Goodman-Bacon, Yingying Li, Jann Spiess, Liyang Sun, Sarah Vicol, Davide Viviano, Stefan Wager, Jeff Wooldridge, and the audience at several conferences (ESAM2024, CES2025, SETA2025, STATA Texas Empirical Micro Conference) and workshops (NYFed, Stanford, UPenn, BU, Emory, GSM/Peking Univ, City Univ of Hong Kong, Rice, U Carlos III de Madrid) for their helpful comments.

[†]Yale University. E-mail xiaohong.chen@yale.edu.

[‡]Emory University. E-mail: pedro.santanna@emory.edu

Peking University. E-mail: xht@gsm.pku.edu.cn

1 Introduction

Difference-in-Differences (DiD) and Event Study (ES) designs are among the most widely used empirical strategies in economics and related fields. For instance, recent data indicates that over 30% of 2024 NBER applied microeconomics working papers mention DiD or ES—more than any other causal inference method,¹ and their popularity is also rapidly expanding in empirical macroeconomics and finance (Goldsmith-Pinkham, 2024). Although recent methodological advances have improved the robustness of DiD estimators to treatment effect heterogeneity,² several empirically relevant econometric questions remain underexplored.

First, modern DiD and ES estimators are designed to accommodate rich treatment effect heterogeneity using short panel data. Unfortunately, most of them can have wide confidence intervals in applications with limited sample size.³ Whether the flexibility of modern DiD necessarily entails a meaningful loss in power compared to more restrictive estimators, such as two-way fixed effects (TWFE), remains an open question. Second, most heterogeneity-robust DiD implementations in short panels implicitly weight all pre-treatment periods equally or discard them entirely. These assumptions are generally unsupported by theory, data or subject-specific knowledge, and often yield imprecise estimators. Third, although it has become a common empirical practice to report several DiD and ES estimators,⁴ there is limited formal guidance on how to compare them, especially so when applied researchers are unwilling to impose strong, arbitrary restrictions on temporal dependence or treatment effect heterogeneity. Sometimes, even basic questions such as whether different DiD estimators target the same causal parameters or if they rely on similar identification assumptions can be difficult to answer.

In this paper, we address these challenges by developing a unified framework for semiparametrically efficient estimation of DiD and ES causal parameters under point-identification assumptions namely, parallel trends and no anticipation. We allow for several commonly-used DiD designs, including designs with a single treatment date and staggered treatment adoption, when covariates may or may not be important for identification. In particular, we (a) characterize the DiD potential outcome model in terms of equivalent restrictions on the joint distribution of observables, (b) derive the semiparametric efficient influence functions (EIF) for DiD and ES parameters and the efficient (i.e., smallest) variance bounds, (c) provide closed-form root-n asymptotically normal and efficient estimators as well as consistent estimators of the variances, and (d) show that semiparametric efficiency requires non-uniform weighting of pre-treatment periods and untreated cohorts. In addition, we highlight that the gains in efficiency/power can be empirically relevant in finite samples, even in designs without variation in treatment timing and covariates.

¹The NBER data used in Goldsmith-Pinkham (2024) is up to May 2024.

²See Roth, Sant'Anna, Bilinski and Poe (2023), de Chaisemartin and D'Haultfœuille (2023) and Baker, Callaway, Cunningham, Goodman-Bacon and Sant'Anna (2025) for recent reviews of DiD advances and a DiD practitioner's guide. See also Baker, Larcker and Wang (2022) for an overview with application in finance.

 $^{^{3}}$ See, e.g., Chiu, Lan, Liu and Xu (2023), Weiss (2024), and Lal (2025) for a discussion.

⁴See, e.g., Braghieri, Levy and Makarin (2022), Hansen and Wingender (2023), Arold (2024), and Mast (2024).

We start by providing an observationally equivalent characterization of the DiD potential outcome identification assumptions in terms of restrictions on the joint distribution of observables. This characterization links modern DiD designs to econometric models of sequential conditional moment restrictions with unknown functions of observables (Ai and Chen, 2012). It clarifies the informational content embedded in modern DiD identification assumptions, and enables our subsequent derivation of the semiparametric efficiency results for DiD models. Importantly, we show that DiD models are typically nonparametrically overidentified in the sense of Chen and Santos (2018). To the best of our knowledge, this is the first result to establish such nonparametric overidentification in a causal inference setting without imposing parametric or semiparametric functional-form restrictions on nuisance components.

Next, we derive the semiparametric efficient influence functions for DiD and ES parameters in closed-forms, under various parallel trends assumptions commonly employed in the DiD literature, including settings with and without variation in treatment timing and cases where parallel trends hold only after conditioning on observed covariates. By definition, an EIF for a causal parameter has mean zero and its second moment is the semiparametric efficient variance bound, which is the smallest asymptotic variance across all possible root-n consistent and asymptotically normal estimators under the DiD design identification conditions. Moreover, we provide closed-form estimators based on the EIFs to achieve the efficient variance bounds. To our knowledge, these are the first semiparametric efficiency results in DiD setups with multiple periods and without parametric functional-form assumptions. Importantly, these efficiency results only explore the DiD identification conditions and do not involve additional hard-to-justify restrictions on treatment effect heterogeneity (e.g., homoskedasticity) or the serial correlation of the outcomes.

Our efficient estimators for the DiD and ES are computed using sample EIFs, which are automatically Neyman orthogonal moments.⁵ The efficient estimators aggregate over multiple comparison groups and pre-treatment periods using analytically derived optimal weights that are proportional to the (conditional) covariances of outcome changes.

According to our semiparametric efficient variance bounds, DiD and ES estimators that treat all pre-treatment periods and comparison groups as equally informative or only use a particular comparison group and the last pre-treatment period as baseline, are generally inefficient, even in setups with a single treatment date. To achieve semiparametric efficiency, it is important to take into account that different untreated groups and pre-treatment periods are not equally informative, and that the best way to aggregate these information to gain precision involves constructing weights that depend on (conditional) covariance terms related to the outcome changes from different pre-treatment periods t_{pre} to a given post-treatment period t_{post} and different comparison groups. We introduce visualization tools that clarify how different baseline periods and comparison groups contribute to the

⁵For nonparametrically overidentified models (Chen and Santos, 2018), not all Neyman orthogonal moment-based doubly robust causal estimators are semiparametric efficient. Only the ones corresponding to the EIFs are semiparametric efficient.

estimator's efficiency, illustrating the structure implied by the semiparametric theory. Importantly, we stress that the geometry of these weights arises from our semiparametric efficiency results.⁶

Our new semiparametric efficient variance bounds offer a rigorous benchmark for evaluating a broad class of DiD and ES estimators, including TWFE and those developed by de Chaisemartin and D'Haultfœuille (2020), Callaway and Sant'Anna (2021), Sun and Abraham (2021), Borusyak, Jaravel and Spiess (2024), Wooldridge (2021), Gardner (2021), among others. Our proposed EIF-based estimators attain the efficient bounds, thereby dominating the existing estimators in terms of asymptotic efficiency. Simulation studies confirm the theoretical rankings.

We illustrate the practical relevance of our framework through empirically calibrated simulations and a real-data application. Our simulations build on DiD designs from Arkhangelsky et al. (2021) (for single-treatment) and Baker et al. (2022) (for staggered-adoption), calibrated to CPS and Compustat data, respectively. In both settings, our semiparametric efficient estimators deliver markedly lower root mean square error and narrower confidence intervals than alternative popular DiD estimators—often exceeding 40% gains in precision, with no loss in bias performance. We further revisit Dobkin, Finkelstein, Kluender and Notowidigdo (2018)'s real-data analysis of hospitalization and out-of-pocket medical expenses, using the data compiled by Sun and Abraham (2021). In this application, alternative estimators would often require sample sizes at least 30% larger to achieve precision comparable to ours. These findings reinforce the central insight of the paper: aligning estimation strategies with the informational structure of the DiD model identification assumptions can yield substantial improvements in precision across a range of DiD settings.

The rest of the paper is as follows. Section 2 formally introduces the general DiD designs and the parameters of interest. Section 3 derives the semiparametric EIFs and the efficient variance bounds. Section 4 proposes our semiparametric efficient estimation. Sections 5 and 6 present the simulation studies and real-data application, respectively. Section 7 concludes with discussions of future extensions. Appendix A introduces Hausman-type as well as incremental Sargan tests for the overidentification restrictions. Appendix B extends the semiparametric efficiency results to an instrumented DiD setting. Appendix C presents the proofs of the theoretical results.

Related Literature: This paper contributes to the popular DiD literature that accommodates treatment effect heterogeneity (see the references mentioned above). We provide the first semiparametric efficiency bounds for DiD and ES estimators in settings with multiple periods and varying forms of the parallel trends assumption, including covariate-conditional and staggered adoption designs. There are some recent work on efficient estimation under some extra conditions in DiD literature. Relative to Sant'Anna and Zhao (2020), who studies a just-identified two-period, two-group DiD model, we generalize to more empirically relevant overidentified settings with multiple periods and staggered adoption. In contrast to Borusyak et al. (2024), Wooldridge (2021), and Harmon (2023),

⁶For large n and large T panel data with single treatment date design, Arkhangelsky, Athey, Hirshberg, Imbens and Wager (2021) synthetic DiD exploits information from pre-treatment periods via a researcher-specified weighting criterion. Their weight is not for semiparametric efficiency consideration, however. See Remark 3.1 for details.

our efficiency results do not rely on auxiliary assumptions such as functional form, homoskedasticity, absence of serial correlation, or other time-series dependence restrictions. These assumptions are convenient but typically hard to justify on theoretical or empirical grounds, particularly in settings with heterogeneous treatment effects and dynamic responses.

Our semiparametric efficiency result builds on Ai and Chen (2012), as we characterize the DiD potential outcome model in terms of sequential conditional moment restrictions with unknown functions of observables.⁷ Without conditioning variables, the DiD potential outcome model can be equivalently transformed into Hansen's overidentified unconditional generalized moment restrictions (Hansen, 1982), and the semiparametric efficient variance bound of Chamberlain (1987) would also be applicable. Even in the case without covariates, our closed-form EIF expressions lead to simpler estimation procedure and provide new insights into how an efficient DiD estimator should weight different pre-treatment and comparison groups—to the best of our knowledge, these insights are new to the literature.

2 Framework, causal parameters, and estimands

We start by describing our general setup. We consider a setting with T periods indexed by $t \in \mathcal{T} = \{1, 2, \ldots, T\}$. At any given time t > 1, units can start receiving a binary treatment, and different units can start receiving treatment at different points. We focus on setups where treatment is an absorbing state, so once a unit is treated, it remains treated until t = T. Let $D_{i,t}$ be an indicator for whether unit i is treated by period t and let $G_i = \min\{t : D_{i,t} = 1\}$ be a "group" variable (or cohort) that indicates the first period at which unit i has received treatment. If unit i does not receive treatment by t = T, we set $G_i = \infty$ and refer to these units as "never treated". Since treatment is an absorbing state, $D_{i,t} = 1$ for all $t \ge G_i$. Let \mathcal{G} denote the support of G and let $\mathcal{G}_{trt} = \mathcal{G} \setminus \{\infty\}$ be the support of G among "eventually-treated" units. Without loss of generality, we assume that a "never-treated" group always exists. If all units are eventually treated, we drop all the data from when the last cohort is treated, so the last-treated cohort becomes the "never-treated" cohort, and T here denotes the number of available periods in the subset of the data that we will use in our analysis. Finally, we also assume that a vector of pre-treatment covariates X_i , whose support is denoted by $\mathcal{X} \subseteq \mathbb{R}^d$ is available. Throughout the paper, we adopt a "large-n, fixed-T" short panel data regime, as is typical in most DiD setups.

In this paper, we adopt the Neyman-Rubin potential outcome framework, indexing each potential outcome by the entire treatment path. Let $\mathbf{0}_s$ and $\mathbf{1}_s$ denote s-dimensional vectors of zeros and ones, respectively. We denote the potential outcome of unit i in period t if they were first treated at time g by $Y_{i,t}(\mathbf{0}_{g-1}, \mathbf{1}_{T-g+1})$, and denote by $Y_{i,t}(\mathbf{0}_T)$ their "never-treated" potential outcome. To simplify notation, we can explore that treatment is an absorbing state and index potential outcomes

⁷Chamberlain (1992) and Ai and Chen (2003) contain unknown functions of observables, but without sequential moments.

by treatment starting time: $Y_{i,t}(g) = Y_{i,t}(\mathbf{0}_{g-1}, \mathbf{1}_{T-g+1})$ and $Y_{i,t}(\infty) = Y_{i,t}(\mathbf{0}_T)$. In practice, though, we observe $Y_{i,t} = \sum_{g \in \mathcal{G}} 1\{G_i = g\}Y_{i,t}(g)$, where $1\{A\}$ is the indicator function that takes value one if A is true, and zero otherwise. Henceforth, we write $G_g = 1\{G = g\}$.

We assume that we observe a random sample of $(Y_{t=1}, \ldots, Y_{t=T}, X', G)'$.

Assumption S (Random Sampling). $\{(Y_{i,t=1},\ldots,Y_{i,t=T},X'_i,G_i)'\}_{i=1}^n$ is a random sample from $(Y_{t=1},\ldots,Y_{t=T},X',G)'$.

We also impose overlap assumptions to avoid "extrapolation" and ensure that our identification arguments are nonparametrically valid.

Assumption O (Overlap). For each $g \in \mathcal{G}$, $\mathbb{E}[G_g|X] \in (0,1)$ almost surely (a.s.).

We maintain the following no-anticipation condition throughout the paper.

Assumption NA (No-anticipation). For every $g \in \mathcal{G}_{trt}$, and every pre-treatment periods t < g, $\mathbb{E}[Y_{i,t}(g)|G = g, X] = \mathbb{E}[Y_{i,t}(\infty)|G = g, X]$ almost surely.

Assumption NA states that, on average, units do not change their behavior before treatment starts, which essentially requires that the eventually treated units do not anticipate the treatment taking place. If treatment is announced in advance and units are expected to change their behavior due to it, this assumption suggests that researchers define the treatment date as the time of treatment announcements and not the time of treatment take-up. See Malani and Reif (2015) for a discussion.

2.1 Causal parameters of interest

Many counterfactuals may be of interest in our context with potential variation in treatment timing. One particular family of causal parameters that have been popular in empirical work is the ATT(g, t) defined by

$$ATT(g,t) := \mathbb{E}\left[Y_t(g) - Y_t(\infty) | G = g\right], \text{ for } t \ge g.$$

$$(2.1)$$

Note that ATT(g, t) measures the Average Treatment Effect at time t of starting treatment at time g versus not starting it, among the units that indeed started treatment at time g. Let $CATT(g, t, X) := \mathbb{E}[Y_t(g) - Y_t(\infty)|G = g, X]$ denote the conditional ATT(g, t) given covariates X.

In setups with a unique treatment date, there is a single treated group, G = g with $g \neq \infty$, so tracking ATT(g,t) over time for this treated group provides a measure of how treatment effects evolve with elapsed treatment time. This is usually referred to as event studies. With multiple groups defined by their treatment timing, ATT(g,t) still provides information about treatment effect dynamics for each treated cohort $g \in \mathcal{G}_{trt}$. However, in such cases, researchers often want to summarize the average causal effects using weighted averages of ATT(g,t) among units that started treatment e periods ago, with $e \ge 0$. A natural aggregation uses the size of each treated cohort as weights, leading to the event study parameter that we denote by ES(e),

$$ES(e) := \mathbb{E}\left[ATT\left(G, G+e\right) \middle| G+e \in [2, T]\right] = \sum_{g \in \mathcal{G}_{trt}} \mathbb{P}(G=g | G+e \in [2, T]) ATT(g, g+e).$$
(2.2)

One may also want to further aggregate the event study coefficients to recover a scalar summary measure. Let \mathcal{E} denote the support of post-treatment event time E = t - G, $t \ge G$, and let N_E denote its cardinality. Then,

$$ES_{\text{avg}} := \frac{1}{N_E} \sum_{e \in \mathcal{E}} ES(e), \qquad (2.3)$$

provides a simple average of all post-treatment event study coefficients.

Under Assumptions S, O and NA only, one cannot identify ATT(g,t) for the post-treatment periods $t \ge g$, as $\mathbb{E}[Y_t(\infty)|G = g]$ is not observed or identified. DiD procedure tackles this problem by restricting the evolution of average untreated potential outcomes across the treated and comparison groups, a parallel trends (PT) assumption. Depending on the treatment timing and covariates, different versions of PT assumptions have been imposed in the literature. We discuss these below.

2.2 Parallel trends and DiD estimands

Two natural types of (conditional) PT have been adopted in the literature: one that effectively restricts trends of $Y_t(\infty)$ only in post-treatment periods and uses never-treated units as the relevant comparison group, and one that restricts trends of $Y_t(\infty)$ in both pre and post-treatment periods, and allow researchers to use any untreated group as a comparison group. We formally state these two PT assumptions below by allowing them to hold only after conditioning on covariates X, i.e., by allowing for covariate-specific trends. If covariates are unavailable or do not play an important role, one can take X = 1 *a.s.* as a special case to resort to an unconditional DiD analysis.

Assumption PT-Post (Parallel Trends in the post-treatment periods for all treated groups). For each $t \in \{2, ..., T\}$ and $g \in \mathcal{G}_{trt}$ such that $t \ge g$,

$$\mathbb{E}[Y_t(\infty) - Y_{t-1}(\infty)|G = g, X] = \mathbb{E}[Y_t(\infty) - Y_{t-1}(\infty)|G = \infty, X] a.s.$$

Assumption PT-All (Parallel Trends for all groups and periods). For each $t \in \{2, ..., T\}$ and $(g, g') \in \mathcal{G}_{trt} \times \mathcal{G}$,

$$\mathbb{E}[Y_t(\infty) - Y_{t-1}(\infty)|G = g, X] = \mathbb{E}[Y_t(\infty) - Y_{t-1}(\infty)|G = g', X] \ a.s.$$

Assumption PT-Post only imposes PT in the post-treatment periods, $t \ge g$, and imposes that, conditional on covariates, the average untreated potential outcome evolution of any treated group $g \in \mathcal{G}_{trt}$ and the never-treated group is the same. This assumption has been effectively used by Callaway and Sant'Anna (2021) and Sun and Abraham (2021) when constructing DiD estimators using nevertreated units as the comparison group. In setups with a single treatment date without covariates, this is the assumption implicitly used in TWFE event study regressions (Jacobson, LaLonde and Sullivan, 1993)

$$Y_{i,t} = \alpha_t + \eta_i + \sum_{e \neq -1} \mathbb{1}[E_{i,t} = e]\beta_e + \varepsilon_{i,t}, \qquad (2.4)$$

where $Y_{i,t}$ is the outcome for unit *i* period *t*, α_t and η_i are time and unit fixed effects often described as capturing time-invariant or common time-varying shocks, $E_{i,t} = t - G_i$ is the time relative to treatment (event-time), $\varepsilon_{i,t}$ is an idiosyncratic error term. The summation runs over all possible values of $E_{i,t}$ among eventually-treated units except for event-time -1.

Under Assumption PT-Post (and Assumptions S, O, NA), the ATT(g,t)'s are identified as

$$ATT(g,t) = \mathbb{E}[Y_t - Y_{g-1}|G = g] - \mathbb{E}[\mathbb{E}[Y_t - Y_{g-1}|G = \infty, X]|G = g].$$
(2.5)

In this case, $t_{\text{pre}} = g - 1$ is the only reliable baseline period because the PT condition does not necessarily hold for the pre-treatment periods. In addition, Assumption PT-Post prevents one from using untreated groups other than the never-treated as the comparison group. In such cases, even if more data on pre-treatment outcomes are available, or the relative size of the never-treated units is small compared to other untreated groups, this information cannot be used to improve the precision of DiD and ES estimators. But this raises the questions: When is PT plausible in post-treatment periods but not in pre-treatment periods? When does economic theory suggest we can only use a specific untreated group as a valid comparison group? The answers to these questions are probably application-dependent. But if we want to leverage more data to potentially get more informative inferences, we must strengthen Assumption PT-Post, leading us to Assumption PT-All.

Assumption PT-All states that, conditional on covariates, the average untreated potential outcome follows the same path in all treatment groups and periods. It is perhaps the most popular identification assumption used in the literature, as invoked by de Chaisemartin and D'Haultfœuille (2020), Sun and Abraham (2021), Gardner (2021), Marcus and Sant'Anna (2021), Borusyak et al. (2024), and Harmon (2023) in the unconditional form, and Callaway and Sant'Anna (2021), Wooldridge (2021), and Lee and Wooldridge (2023) in the conditional form. In setups with a single treatment date and no covariates, this assumption is arguably the one that justifies using a "static" TWFE regression specification

$$Y_{i,t} = \alpha_t + \eta_i + \beta D_{i,t} + \varepsilon_{i,t}, \qquad (2.6)$$

where $\beta = ES_{avg}$ is the parameter of interest. Unlike Assumption PT-Post, Assumption PT-All imposes a parallel pre-trends condition across all time periods.

As Assumption PT-All allows one to use different pre-treatments as baseline periods and leverage different sets of untreated units as a comparison group, the setup becomes much richer. The following lemma highlights how to combine comparison groups and pre-treatment periods to identify the ATT(g,t)'s, hinting that the DiD model becomes overidentified. Let $\mathcal{H}^{g,t} = \{(g', t_{\text{pre}}, t'_{\text{pre}}) \in \mathcal{G} \times \mathcal{T} \times \mathcal{T} : g > t'_{\text{pre}}, g' > \max\{t_{\text{pre}}, t'_{\text{pre}}\}\}.$

Lemma 2.1. Let Assumptions S, O, NA, and PT-All hold. Then, for every group $(g, g') \in \mathcal{G}_{trt} \times \mathcal{G}_{trt}$

and time periods $(t, t_{pre}, t'_{pre}) \in \mathcal{T} \times \mathcal{T} \times \mathcal{T}$ such that $t \ge g$, $g > t'_{pre}$, and $g' > \max\{t_{pre}, t'_{pre}\}$, with probability one,

$$CATT(g,t,X) = \underbrace{\mathbb{E}[Y_t - Y_{t'_{pre}}|G = g, X]}_{:=m_{g,t,t'_{pre}}(X)} - \left(\underbrace{\mathbb{E}[Y_t - Y_{t_{pre}}|G = \infty, X]}_{:=m_{\infty,t,t_{pre}}(X)} + \underbrace{\mathbb{E}[Y_{t_{pre}} - Y_{t'_{pre}}|G = g', X]}_{:=m_{g',t_{pre},t'_{pre}}(X)}\right), \quad (2.7)$$

and, as a consequence,

$$ATT(g,t) = \mathbb{E}\left[Y_t - Y_{t'_{pre}} \middle| G = g\right] - \mathbb{E}\left[\left(m_{\infty,t,t_{pre}}(X) + m_{g',t_{pre},t'_{pre}}(X)\right) \middle| G = g\right].$$
(2.8)

More generally, for any covariate-specific weights $w_{g',t_{pre},t'_{pre}}^{att(g,t)}(X)$ summing to one, we have that

$$ATT(g,t) = \mathbb{E}\bigg[\sum_{(g',t_{pre},t'_{pre})\in\mathcal{H}^{g,t}} w_{g',t_{pre},t'_{pre}}^{att(g,t)}(X) \left[m_{g,t,t'_{pre}}(X) - \left(m_{\infty,t,t_{pre}}(X) + m_{g',t_{pre},t'_{pre}}(X)\right)\right]\bigg|G = g\bigg].$$

Lemma 2.1 leverages several restrictions on the indexes worth explaining. First, the restriction $t \ge g$ means we identify ATT(g,t)'s in group g post-treatment periods. The restriction $g > t'_{\text{pre}}$ suggests using any of the group's g pre-treatment period as a baseline, while the restriction $g' > \max\{t_{\text{pre}}, t'_{\text{pre}}\}\$ means that we are using the eventually treated cohort g' that is treated after periods $t_{\text{pre}}, t'_{\text{pre}}\$ as part of our effective comparison group. Interestingly, these restrictions allow $t_{\text{pre}}\$ to be a post-treatment period for group g, as long as it is a pre-treatment period for cohort g'. Figure 1 illustrates why this is possible in our setup using a stylized example with four groups and ten time periods—essentially, when parallel trends hold in all pre-treatment periods, one can "subtract back" some "excessive" corrections using the never-treated units.

Figure 1: Illustration of how we can use t_{pre} in post-treatment period for group g



ATT(3,5) using $t'_{\rm pre} = 1$, $t_{\rm pre} = 7$, and with $G \in \{\infty, 8\}$ as comparison group

Notes: Illustrative example of a setup with 10 time periods and four different treatment groups, where one is interested in learning about ATT(3,5). Dark blue means units are treated (post-treatment periods), while light blue denotes pre-treatment periods. The hashed area denotes the "active" comparison group. The illustration highlights that one can use a combination of two not-yet-treated cohorts, G = 8 and $G = \infty$, as a comparison group, as well as leverage $t'_{\rm pre} = 1$ and $t_{\rm pre} = 7$ as "baseline periods". Using a post-treatment $t_{\rm pre} = 7$ as a baseline period is possible because the never-treated comparison group bridges this back to pre-treatment periods, since parallel trends hold in all periods for all groups.

Overall, Lemma 2.1 indicates that one can use any group g's pre-treatment period as a baseline, combine never-treated units with any not-yet-treated cohort g' > g to form comparison groups, and

use pre-treatment data for eventually-treated cohort g' when identifying an ATT(g, t). In practice, there are infinitely many ways one can combine these to identify the ATT(g, t)'s. In practice, however, one very relevant question is how to combine these different estimands to asymptotically obtain the most precise DiD estimator for the ATT(g, t)'s, or to obtain the most precise event study parameters ES(e) as defined in (2.2)? We discuss these points in the next section.

Remark 2.1. As Lemma 2.1 indicates, any covariate-specific weights $w_{g',t_{\text{pre}},t'_{\text{pre}}}^{\text{att}(g,t)}(X)$ summing to one can be used to identify a ATT(g,t). That is, the result in Lemma 2.1 allows for non-convex weights. Different from the DiD literature that discusses limitations of simple two-way fixed effects specifications (Goodman-Bacon, 2021, Borusyak et al., 2024, de Chaisemartin and D'Haultfœuille, 2020, Sun and Abraham, 2021), non-convex weights are *not* a concern in our context. Under the identifications in Lemma 2.1, the DiD model is overidentified, and (2.7) is *homogeneous* across baseline periods and comparison groups. Thus, issues related to an estimand not being weakly causal (Blandhol, Bonney, Mogstad and Torgovitsky, 2022) due to negative weights are not binding in our context.

3 Semiparametric Efficiency Bound for DiD and ES

In this Section, we derive the semiparametric efficiency bound for the ATT(g,t)'s under Assumption PT-Post and under Assumption PT-All. We also provide analogous results for the event study parameters $ES(e), e \ge 0$. These semiparametric efficiency bounds characterize the lowest possible asymptotic variance attainable by any regular, asymptotically linear estimator given the identification assumptions. Based on the efficient influence function, we derive closed-form expressions for estimands that can be used to construct efficient estimators under minimal smoothness assumptions. These results provide a benchmark for evaluating the asymptotic efficiency of existing DiD and ES estimators in the literature.

To build intuition, we first focus on the canonical case of a single treatment date. This simpler setting provides a transparent view of the key ideas and allows us to isolate the informational content of each pre-treatment period. We then generalize the results to the more complex case of staggered treatment timing. When covariates are uninformative for identification, the results further simplify to the unconditional DiD setup by setting X = 1.

3.1 DiD setups with a single treatment date

Consider the setting with a single treatment date, resulting in two groups: treated (G = g) and never-treated $(G = \infty)$. In this case, the event study parameters ES(e) are equivalent to ATT(g, t)'s with t = g + e. We proceed under either Assumption PT-Post or Assumption PT-All, noting that the former is a special case of the latter. With two groups, we have $G_{\infty} = 1 - G_g$ and $p_{\infty}(X) = 1 - p_g(X)$, with $p_g(X) = \mathbb{E}[G_g|X]$ being the propensity score, i.e., the probability of belonging to the treated group given covariates X. Under Assumptions S, O, NA, and PT-All, Lemma 2.1 implies that for any post-treatment period $t \ge g$, any pre-treatment period $t_{\text{pre}} \in \{1, \ldots, g-1\}$, and any weight $w_{t_{\text{pre}}}^{\text{att}(g,t)}(\cdot)$ summing to one,

$$ATT(g,t) = \mathbb{E}\left[Y_t - Y_{t_{\text{pre}}} \middle| G = g\right] - \mathbb{E}\left[\mathbb{E}\left[Y_t - Y_{t_{\text{pre}}} \middle| G = \infty, X\right] \middle| G = g\right]$$
$$= \mathbb{E}\left[\sum_{t_{\text{pre}}=1}^{g-1} w_{t_{\text{pre}}}^{\text{att}(g,t)}(X) \left[m_{g,t,t_{\text{pre}}}(X) - m_{\infty,t,t_{\text{pre}}}(X)\right] \middle| G = g\right].$$

This expression shows how information from multiple pre-treatment periods can be used to identify our target parameters, but it does not yet indicate *how* to weigh each period to gain precision. To this end, the next step is to characterize all the information content in our identification assumptions. In the next lemma, we establish an equivalent representation of our identification assumptions that impose restrictions on potential outcomes using sequential conditional moment restrictions on observable variables.

Lemma 3.1 (Moment-restrictions for overidentified DiD with single treatment time).

Suppose there are only two treatment groups, such that $G \in \mathcal{G} = \{g, \infty\}$, with $g \ge 2$. The family of probability distributions of $(Y_{t=1} \cdots, Y_{t=T}, X', G)$ satisfying Assumptions S, O, NA, and PT-All are observationally equivalent to the family of probability distributions of $(Y_{t=1} \cdots, Y_{t=T}, X', G)$ satisfying Assumptions S, O, and the following set of moment restrictions: for all $t \in \{g, \ldots, T\}$, with probability one,

$$\begin{split} \mathbb{E}[G_g(ATT(g,t) - CATT(g,t,X))] &= 0, \\ \mathbb{E}\left[CATT(g,t,X) - \frac{G_g(Y_t - Y_{g-1})}{p_g(X)} + \frac{G_{\infty}(Y_t - Y_{g-1})}{p_{\infty}(X)} \middle| X\right] &= 0, \\ \mathbb{E}\left[\frac{G_g(Y_{t_{pre}} - Y_1)}{p_g(X)} - \frac{G_{\infty}(Y_{t_{pre}} - Y_1)}{p_{\infty}(X)} \middle| X\right] &= 0, \text{ for all } 2 \leqslant t_{pre} \leqslant g-1, \\ \mathbb{E}[G_g - p_g(X)|X] &= 0. \end{split}$$

Lemma 3.1 lays the groundwork for deriving the efficient influence function and the semiparametric efficiency bound. Beyond that, it succinctly presents all the identifying power of our identification assumptions and highlights that one can construct many different DiD estimators based on it: it is just a matter of choosing an estimation method suitable for moment restrictions. Examples include (two-step) generalized methods of moments (GMM; see, e.g., Hansen, 1982, Ackerberg, Chen, Hahn and Liao, 2014), generalized empirical likelihood (Newey and Smith, 2004), and minimum distance estimators (Ai and Chen, 2003, 2007, 2012).

A natural question is whether the choice of the estimation procedure matters in terms of asymptotic efficiency. As discussed in the previous section, the DiD model characterized in Lemma 3.1 is overidentified in the sense of Chen and Santos (2018). As such, not every nonparametric estimation procedure is asymptotically semiparametrically efficient, implying that it is relevant to characterize the minimum asymptotic variance for any regular, asymptotic linear (RAL) DiD estimator, and to consider estimators that achieve this bound. Although different efficiency tools are available in the literature, many of them are not suitable for our DiD model in Lemma 3.1. For instance, the DiD model in Lemma 3.1 is based on sequential conditional moment restrictions with unknown nonparametric functions, making the efficiency results in Chamberlain (1992) and Ai and Chen (2003) not directly applicable. In addition, note that the conditional average treatment effects among the treated units in period t in our DiD model, CATT(g, t, X), is given by

$$CATT(g,t,X) = \mathbb{E}\left[Y_t - Y_{t_{\text{pre}}} | G = g, X\right] - \mathbb{E}\left[Y_t - Y_{t_{\text{pre}}} | G = \infty, X\right] \text{ for any } 1 \le t_{\text{pre}} \le g - 1.$$

This suggests that our first-step nuisance function is overidentified, making the efficiency results of Ackerberg et al. (2014) based on two-step GMM unsuitable for our context. On the other hand, we can build on Ai and Chen (2012), as their models are well-suited to our setup. Nonetheless, applying their results still requires substantial additional work: we need to first solve a calculus variations problem and then unravel several nested matrix inversions before arriving at the concise expression provided in Theorems 3.1 and 3.2.

In what follows, we present our semiparametric efficiency result under Assumption PT-All, and then present as a special case the semiparametric efficiency under Assumption PT-Post. Let $\pi_g = \mathbb{E}[G_g]$ denote the population proportion of treated units, and denote the group-specific conditional expectations of the evolution of outcomes from period t_{pre} to period t by

$$m_{g,t,t_{\rm pre}}(X) := \mathbb{E}[Y_t - Y_{t_{\rm pre}} | G = g, X].$$
 (3.1)

Let **1** denote a column vector of ones with the appropriate length varying according to the context. Let $W = (Y_{t=1}, \ldots, Y_{t=T}, X', G_g)$ denote the available random variables, and

$$\widetilde{Y}_{g,t,t_{\rm pre}} = \frac{1}{\pi_g} \left(G_g - \frac{p_g(X)}{p_\infty(X)} G_\infty \right) \left(Y_t - Y_{t_{\rm pre}} - m_{\infty,t,t_{\rm pre}}(X) \right).$$
(3.2)

For $1 \leq t_{\text{pre}} \leq g - 1$, let

$$\mathbb{IF}_{t_{\rm pre}}^{att(g,t)} = \widetilde{Y}_{g,t,t_{\rm pre}} - \frac{G_g}{\pi_g} ATT(g,t), \qquad (3.3)$$

and denote the g-1 column vector that stacks of these g-1 influence functions by $\mathbb{IF}^{att(g,t)} = (\mathbb{IF}_1^{att(g,t)}, \mathbb{IF}_2^{att(g,t)}, \cdots, \mathbb{IF}_{g-1}^{att(g,t)})'$. Finally, denote the $(g-1) \times (g-1)$ conditional covariance matrix of $\mathbb{IF}^{att(g,t)}$ as $V_{gt}(X) = \text{Cov}(\mathbb{IF}^{att(g,t)}|X)$, and let $V_{gt}^*(X)$ be a matrix of the same dimension as $V_{gt}(X)$ with the (j, k)-th element being

$$\frac{1}{p_g(X)}\operatorname{Cov}(Y_t - Y_j, Y_t - Y_k | G = g, X) + \frac{1}{1 - p_g(X)}\operatorname{Cov}(Y_t - Y_j, Y_t - Y_k | G = \infty, X).$$
(3.4)

Theorem 3.1. Suppose there are only two treatment groups, such that $G \in \mathcal{G} = \{g, \infty\}$, with $g \ge 2$. Under Assumptions S, O, NA and PT-All, the efficient influence function of a $ATT(g,t), t \ge g$ is given by

$$\mathbb{EIF}^{att(g,t)} = \frac{\mathbf{1}' V_{gt}(X)^{-1}}{\mathbf{1}' V_{gt}(X)^{-1} \mathbf{1}} \mathbb{IF}^{att(g,t)} = \frac{\mathbf{1}' V_{gt}^*(X)^{-1}}{\mathbf{1}' V_{gt}^*(X)^{-1} \mathbf{1}} \mathbb{IF}^{att(g,t)}.$$

Assuming that the second moment of the efficient influence function is finite, the semiparametric

efficiency bound for any regular, asymptotically linear estimator for the ATT(g,t) is given by

$$V_{eff} = \frac{1}{\pi_g^2} \left(\mathbb{E}[G_g(CATT(g, t, X) - ATT(g, t))^2] + \mathbb{E}\left[\frac{p_g(X)^2 |V_{gt}^*(X)|}{\sum_{j,j'=1}^{g-1} (-1)^{j+j'} |V_{gt,jj'}^*(X)|}\right] \right),$$

where $V_{gt,jj'}^*$ denotes the minor matrix obtained by removing the *j*th row and *j'*th column of V_{gt}^* .⁸

In the proof of Theorem 3.1, we show that including one post-treatment at a time and focusing on one ATT(g,t) at a time leads to the same efficiency bound as including all post-treatments and estimating all ATT(g,t)'s. The main advantage of working with one ATT(g,t) at a time is that it avoids dealing with a much larger variance-covariance matrix, which would preclude us from getting more intuition about efficiently aggregating each pre-treatment period. When it comes to estimation, this simpler characterization will allow us to avoid dealing with a much larger number of conditional moment restrictions that would lead to more computationally challenging procedures. Moreover, as a side result, this characterization allows us to present the semiparametric efficiency bound for the just-identified DiD setup when one imposes Assumption PT-Post instead of Assumption PT-All in Corollary 3.1 as a corollary of Theorem 3.1.

Corollary 3.1. Suppose there are only two treatment groups, such that $G \in \mathcal{G} = \{g, \infty\}$, with $g \ge 2$. Under Assumptions S, O, NA and PT-Post, the efficient influence function of a $ATT(g,t), t \ge g$, is equal to $\mathbb{IF}_{g-1}^{att(g,t)}$, and the semiparametric efficiency bound reduces to $V_{eff,j-id} = \mathbb{E}\left[\left(\widetilde{Y}_{g,t,g-1} - \frac{G_g}{\pi_g}ATT(g,t)\right)^2\right].$

Theorem 3.1 and Corollary 3.1 have several interesting implications. First, suppose we only use the pre-treatment period $t_{\text{pre}} = g - 1$ as a baseline, so PT holds only for post-treatment periods (Assumption PT-Post). Corollary 3.1 shows that the efficient influence function is $\mathbb{IF}_{g,g-1}^{g,t}$ under Assumption PT-Post. This is effectively the same as dropping data from all periods but time t(post-treatment) and time g - 1 (pre-treatment), transforming the multi-period DiD model into a two-period DiD setup. Thus, under Assumption PT-Post, Corollary 3.1 can be understood as generalizing the semiparametric efficiency bound result of Sant'Anna and Zhao (2020) from the much simpler 2-group-2-period DiD setup to the multi-period DiD setup with a single treatment date. Importantly, Corollary 3.1 highlights that pre-treatment data beyond the last pre-treatment period cannot be used to improve asymptotic efficiency, which can be hard to motivate empirically.

When one can leverage multiple pre-treatment periods (Assumption PT-All), Theorem 3.1 highlights that the efficient influence function, in this case, is an optimally weighted average of all influence functions based on different baseline periods, where the optimal set of weights is computed by minimizing the conditional variance given X. Notably, the optimal way to aggregate pretreatment periods for estimating ATT(g,t) will generally differ from when estimating $ATT(g,t_{pre})$, with $t_{pre} \neq t$, with weights generally being non-uniform across pre-treatment periods and covariate

⁸In the formula of the efficiency bound, we follow the convention in matrix theory to define the determinant of an empty matrix to be one.

strata. Of course, there are special cases where weighing all pre-treatment information equally is optimal. This is the case when outcome changes are (conditionally) uncorrelated across periods for both the treated and comparison group and have constant (conditional) variances across pre-treatment periods for both treatment groups (Wooldridge, 2021, Borusyak et al., 2024). Outside this scenario, it is hard to see a setup where equal weights for all pre-treatment periods are asymptotically optimal. Our semiparametric efficiency results do not impose any of these strong, hard-to-empirically motivate assumptions. As highlighted in Lemma 3.1, we do not impose additional restrictions beyond the identification assumptions, therefore being able to capture much richer notions of heterogeneity.

The efficient weights in Theorem 3.1 also vary with covariates. This is intuitive, as the optimal way to combine pre-treatment periods may vary across covariate strata. For instance, it can be that for a particular partition of the covariate space, more recent pre-treatment periods are "more informative", while for another partition, it can be that all pre-treatment periods are equally informative. The flexibility of the weights varying with X is meant to capture exactly this. It is also worth mentioning that even when the covariance terms in (3.4) are trivial functions of X, a type of homoscedasticity assumption within each treatment group, $V_{gt}^*(X)$ and, therefore, the optimal weights would still generally depend on X via the propensity score $p_g(X)$.⁹ Thus, Theorem 3.1 stresses the importance of choosing covariate-specific weights to achieve semiparametric efficiency.

The efficient influence function in Theorem 3.1 also informs how one can construct DiD estimators that achieve the semiparametric efficiency bound by effectively weighting different pre-treatment periods. More specifically, we can leverage the fact that the efficient influence function has mean zero to get the following DiD estimand:

$$ATT(g,t) = \mathbb{E}\left[\frac{\mathbf{1}'V_{gt}^{*}(X)^{-1}}{\mathbf{1}'V_{gt}^{*}(X)^{-1}\mathbf{1}}\widetilde{Y}_{g,t}\right],$$
(3.5)

where $\tilde{Y}_{g,t} = (\tilde{Y}_{g,t,1}, \ldots, \tilde{Y}_{g,t,g-1})'$ is a $(g-1) \times 1$ column vector of transformed outcomes. Thus, the DiD estimand ATT(g,t) uses a combination of pre-treatment periods as "effective" baselines, where the weights for each of these pre-treatment periods are given by $\mathbf{1}'V_{gt}^*(X)^{-1}/\mathbf{1}V_{gt}^*(X)^{-1}\mathbf{1}$. This expression highlights that, indeed, "more informative" baseline periods get weighted more than "less informative" periods, where the relevant notion of "informativeness" is given by the inverse of $V_{gt}^*(X)$ with elements defined in (3.4).

Remark 3.1. To our knowledge, the synthetic DiD procedure of Arkhangelsky et al. (2021) is the only other panel-data method that weights pre-treatment periods in a non-uniform manner. However, many important differences between our procedure and theirs are worth stressing. First, and in contrast to Arkhangelsky et al. (2021), our non-uniform weights are solved in closed-form for the semiparametric efficiency purpose. Second, our weights can vary with covariates, while covariates play no major role in their synthetic DiD procedures. On the other hand, our results depend on the

⁹It may be the case that the ratio of conditional covariance and the propensity score in (3.4) is constant, even when all these terms are not trivial functions of the covariates. Nonetheless, this seems like a knife-edge case requiring a very specific covariance structure, which is unrealistic for most applications.

plausibility of parallel trends, while synthetic DiD procedures can accommodate some deviations of parallel trends. Finally, while we focus on setups with a fixed number of periods, Arkhangelsky et al. (2021) inferential procedures are tailored for applications with a diverging number of pre-treatment periods.

Remark 3.2. Researchers often want to summarize the average treatment effects across all periods into a scalar parameter,

$$ES_{\text{avg}} = \frac{1}{N_E} \sum_{e \in \mathcal{E}} ES(e)$$

It follows from Theorem 3.1 and the Delta method that one can form efficient estimators for ES_{avg} in this two-group DiD setup using

$$ES_{\text{avg}} = \frac{1}{T - g + 1} \sum_{t=g}^{T} \mathbb{E} \left[\frac{\mathbf{1}' V_{gt}(X)^{-1}}{\mathbf{1}' V_{gt}(X)^{-1} \mathbf{1}} \widetilde{Y}_{g,t} \right],$$
(3.6)

with ATT(g, t)'s as defined in (3.5). As evident from (3.6), the efficient way to combine pre-treatment information also varies with covariates, and it is unlikely to use a uniform weighting scheme in many applications with any serial correlation.

Remark 3.3. From Lemma 3.1, Theorem 3.1, and Corollary 3.1, it is clear that our DiD model is over-identified. A natural consequence of these results is that we can directly assess the validity of the identification assumption via a Hausman-type test. One can also sequentially assess whether using data from periods further away from the treatment date is warranted via incremental Sargan tests. We discuss these testing procedures in Appendix A.

3.2 DiD setups with staggered treatment designs

We now extend our semiparametric efficiency bound results for ATT and ES estimators from the previous section to the staggered setups and provide guidelines on the most precise RAL DiD and ES estimators. We start studying the overidentified staggered DiD model for ATT(g,t), and later discuss the just-identified one as a special case. As we discuss below, the efficiency of ES(e), $e \ge 0$, follows from the efficiency of ATT(g,t)'s and the Delta method.

As staggered treatment times imply the existence of multiple treatment groups, we now have that $G_{\infty} = 1 - \sum_{g \in \mathcal{G}_{trt}} G_g$ and $p_{\infty}(X) = 1 - \sum_{g \in \mathcal{G}_{trt}} p_g(X)$, with $p_g(X) = \mathbb{E}[G_g|X]$ being a generalized propensity score, i.e., the probability of belonging to the group that is treated for the first time in period g given covariates X.

The following lemma characterizes the information content of our staggered DiD identification assumptions based on potential outcomes using sequential conditional moment restrictions on observable variables.

Lemma 3.2 (Moment-restrictions for overidentified staggered DiD). The family of probability distributions of $(Y_{t=1}, \dots, Y_{t=T}, X', G)$ satisfying Assumptions S, O, NA, and PT-All are observational equivalent to the family of probability distributions of $(Y_{t=1} \cdots, Y_{t=T}, X', G)$ satisfying Assumptions S, O and the following set of moment restrictions: for all $g, g' \in \mathcal{G}_{trt} \times \mathcal{G}_{trt}$ and post-treatment periods $t \in \{g, \ldots, T\}$, with probability one,

$$\mathbb{E}[\pi_{g} - G_{g}] = 0,$$

$$\mathbb{E}[G_{g}(ATT(g, t) - CATT(g, t, X))] = 0,$$

$$\mathbb{E}\left[CATT(g, t, X) - \frac{G_{g}(Y_{t} - Y_{g-1})}{p_{g}(X)} + \frac{G_{\infty}(Y_{t} - Y_{g-1})}{p_{\infty}(X)} \Big| X\right] = 0,$$

$$\mathbb{E}\left[\frac{G_{g'}(Y_{t_{pre}} - Y_{1})}{p_{g'}(X)} - \frac{G_{\infty}(Y_{t_{pre}} - Y_{1})}{p_{\infty}(X)} \Big| X\right] = 0, \text{ for all } 2 \leq t_{pre} \leq g' - 1, \qquad (3.7)$$

$$\mathbb{E}[G_{g} - p_{g}(X)|X] = 0.$$

The first two (unconditional) sets of moment conditions in Lemma 3.2 define each cohort's relative size, π_g , and the ATT(g,t) as a functional of the conditional ATT(g,t). The third and fourth sets of (conditional) moment restrictions define the conditional ATT(g,t) given covariates, CATT(g,t,X), and the set of overidentification restrictions implied by our Assumptions S, O, NA, and PT-All. By combining the third and fourth conditional moment restriction, it is clear that any pre-treatment period $t_{\rm pre} < g$ and multiple comparison groups can be used to identify the CATT(g,t,X), and therefore, the ATT(g,t). This is what essentially rationalizes Lemma 2.1. The last set of conditional moment restrictions defines the generalized propensity scores, $p_g(X)$. Finally, also note that, from Lemma 3.2, we can construct the event study parameters $ES(e), e \ge 0$, as

$$ES(e) = \sum_{g \in \mathcal{G}_{trt,e}} \frac{\pi_g}{\sum_{g \in \mathcal{G}_{trt,e}} \pi_g} ATT(g, g + e),$$
(3.8)

where, for $e \ge 0$, $\mathcal{G}_{trt,e} = \{g \in \mathcal{G}_{trt} : g + e \le T\}$ denotes the set of eventually-treated groups that have data for *e*-periods after treatment started.

Similarly to Lemma 3.1, Lemma 3.2 presents all the identifying power of our identification assumptions in the staggered DiD setup and highlights that one can construct many different DiD estimators based on it. The relevant question is how to fully explore the empirical content of the moment conditions in Lemma 3.2 to form DiD and ES estimators with appealing semiparametric efficiency guarantees. Here, in contrast to the setup in Lemma 3.1, we can use different pre-treatment periods and different comparison groups when constructing DiD and ES estimators. Hence, the setup is more interesting, though it also requires additional notation to characterize how the "most precise" DiD and ES estimators should look.

Let the "generated outcome" for a given ATT(g, t) be defined as

$$\widetilde{Y}_{g',t_{\rm pre}}^{\rm att(g,t)} = \frac{G_g}{\pi_g} (Y_t - Y_1 - m_{\infty,t,t_{\rm pre}}(X) - m_{g',t_{\rm pre},1}(X)) - \frac{p_g(X)}{p_{\infty}(X)} \frac{G_{\infty}}{\pi_g} (Y_t - Y_{t_{\rm pre}} - m_{\infty,t,t_{\rm pre}}(X)) - \frac{p_g(X)}{p_{g'}(X)} \frac{G_{g'}}{\pi_g} (Y_{t_{\rm pre}} - Y_1 - m_{g',t_{\rm pre},1}(X)), \quad (3.9)$$

and, for $g' \in \mathcal{G}_{trt}$ and $1 \leq t_{pre} \leq g' - 1$, let

$$\mathbb{IF}_{g',t_{\text{pre}}}^{att(g,t)} = \widetilde{Y}_{g',t_{\text{pre}}}^{\text{att}(g,t)} + \frac{G_g}{\pi_g} ATT(g,t).$$
(3.10)

Note that $\tilde{Y}_{g',t_{\text{pre}}}^{\text{att}(g,t)}$ uses data from the earliest pre-treatment and from t_{pre} . It also leverages observations from the "never-treated" group G_{∞} and from g' as the comparison group. When g' = g, though, it only uses data from G_{∞} as a comparison group, and, therefore, (3.9) and (3.10) reduce to (3.2) and (3.3), respectively.

Next, we collect all noncollinear $\mathbb{IF}_{g',t_{\mathrm{pre}}}^{att(g,t)}$ to characterize the "relevant" set of influence functions to be used in the construction of the efficient influence function. For g = g', let $\mathbb{IF}_{g}^{att(g,t)} = (\mathbb{IF}_{g,1}^{att(g,t)}, \mathbb{IF}_{g,2}^{att(g,t)}, \cdots, \mathbb{IF}_{g,g-1}^{att(g,t)})'$, while for each $g' \neq g$, we do not include the very first time period and let $\mathbb{IF}_{g',2}^{att(g,t)} = (\mathbb{IF}_{g',3}^{att(g,t)}, \mathbb{IF}_{g',3}^{att(g,t)}, \cdots, \mathbb{IF}_{g',g'-1}^{att(g,t)})'$. Denote the vector stacking all these vectors together by $\mathbb{IF}_{\mathrm{stg}}^{att(g,t)}$

$$\mathbb{IF}_{stg}^{att(g,t)} = (\mathbb{IF}_{g'}^{att(g,t),'}, g' \in \mathcal{G}_{trt})', \qquad (3.11)$$

and let its conditional covariance matrix be defined as $\Omega_{gt}(X) = \text{Cov}(\mathbb{IF}_{stg}^{att(g,t)}|X)$. Let $\Omega_{gt}^*(X)$ be a matrix of the same dimension as Ω_{gt} with (j, k)-th element given by

$$\frac{1}{p_g(X)} \operatorname{Cov}(Y_t - Y_1, Y_t - Y_1 | G = g, X) + \frac{1}{p_{\infty}(X)} \operatorname{Cov}(Y_t - Y_{t'_j}, Y_t - Y_{t'_k} | G = \infty, X)
- \frac{1\{g = g'_j\}}{p_g(X)} \operatorname{Cov}(Y_t - Y_1, Y_{t'_j} - Y_1 | G = g, X) - \frac{1\{g = g'_k\}}{p_g(X)} \operatorname{Cov}(Y_t - Y_1, Y_{t'_k} - Y_1 | G = g, X)
+ \frac{1\{g_j = g'_k\}}{p_{g'_j}(X)} \operatorname{Cov}(Y_{t'_j} - Y_1, Y_{t'_k} - Y_1 | G = g'_j, X),$$
(3.12)

where (g'_s, t'_s) is the value that g' and t_{pre} takes in the s-th entry of $\mathbb{IF}^{att(g,t)}_{stg}$. Finally, let $q_{g,e} = \mathbb{P}(G = g|G + e \in [2, T])$.

The next Theorem establishes the semiparametric efficiency bound for ATT(g,t) and ES(e) parameters under PT-All and staggered treatment adoption. We present the semiparametric efficiency bound for the just-identified DiD setup when one imposes Assumption PT-Post instead of Assumption PT-All in Corollary 3.2.

Theorem 3.2. Suppose that treatment adoption is staggered over time. Under Assumptions S, O, NA and PT-All, the efficient influence function of π_g and ATT(g,t), $t \ge g$, are given by

$$\mathbb{EIF}^{\pi_g} = G_g - \pi_g,$$

$$\mathbb{EIF}_{stg}^{att(g,t)} = \frac{\mathbf{1}'\Omega_{gt}(X)^{-1}}{\mathbf{1}'\Omega_{gt}(X)^{-1}\mathbf{1}} \mathbb{IF}_{stg}^{att(g,t)} = \frac{\mathbf{1}'\Omega_{gt}^*(X)^{-1}}{\mathbf{1}'\Omega_{gt}^*(X)^{-1}\mathbf{1}} \mathbb{IF}_{stg}^{att(g,t)}$$

Consequently, the efficient influence function of a ES(e), $e \ge 0$, is given by

$$\mathbb{EIF}_{stg}^{es(e)} = \sum_{g \in \mathcal{G}_{trt,e}} \left(q_{g,e} \mathbb{EIF}_{stg}^{att(g,g+e)} + \frac{ATT(g,g+e)}{\sum_{g' \in \mathcal{G}_{trt,e}} \pi_{g'}} \left((G_g - \pi_g) - q_{g,e} \sum_{s \in \mathcal{G}_{trt,e}} (G_s - \pi_s) \right) \right).$$

The semiparametric efficiency bounds are obtained as the second moments of the efficient influence functions, assuming that these second moments are finite.

Corollary 3.2. Suppose that treatment adoption is staggered over time. Under Assumptions S, O, NA and PT-Post, the model for ATT(g,t), $t \ge g$, is nonparametrically just-identified, and the efficient influence function of a ATT(g,t), $t \ge g$, is equal to $\mathbb{IF}_{g,g-1}^{att(g,t)}$. In such cases, the efficient influence function of a ES(e), $e \ge 0$, is given by

$$\mathbb{EIF}_{j\text{-}id,stg}^{es(e)} = \sum_{g \in \mathcal{G}_{trt,e}} \left(q_{g,e} \mathbb{IF}_{g,g-1}^{att(g,g+e)} + \frac{ATT(g,g+e)}{\sum_{g' \in \mathcal{G}_{trt,e}} \pi_{g'}} \left((G_g - \pi_g) - q_{g,e} \sum_{s \in \mathcal{G}_{trt,e}} (G_s - \pi_s) \right) \right).$$

Theorem 3.2 and Corollary 3.2 have some practical implications. First, when one is only willing to assume parallel trends for post-treatment periods, as in Assumption PT-Post, Corollary 3.2 shows that this is essentially the same problem as in the single treatment setup and that ATT(g,t) is just-identified: one can only use g-1 period as baseline, and the never-treated group as comparison group.

When one can leverage multiple pre-treatment periods as in Assumption PT-All, Theorem 3.2 highlights that the efficient influence function is an optimally weighted average of all influence functions based on different baseline periods and different comparison groups. Theorem 3.2 also shows that the optimal way to pool information across periods and comparison groups is likely to be non-uniform, highlighting that not every comparison group and pre-treatment period has the same "relevance" when it comes to inference. Importantly, these weights are allowed to vary with the data-generating process and are derived as a consequence of our semiparametric efficiency considerations. As no currently available staggered DiD and ES estimator shares these weighting schemes, they do not generally achieve the semiparametric efficiency bound, suggesting that we can potentially conduct more informative inferences.

Theorem 3.2 also provides a "blueprint" on how one can construct DiD and ES estimators that enjoy attractive semiparametric efficiency asymptotic guarantees. Similarly to the single-treatment date, we can explore the efficient influence function to get the following DiD ATT(g,t) estimands:

$$ATT(g,t) = ATT_{\rm stg}(g,t) := \mathbb{E}\left[\frac{\mathbf{1}'\Omega_{gt}^*(X)^{-1}}{\mathbf{1}'\Omega_{gt}^*(X)^{-1}\mathbf{1}}\widetilde{Y}^{\rm att(g,t)}\right],\tag{3.13}$$

where $\tilde{Y}^{\text{att}(g,t)}$ is a column vector of all (non-collinear) generated outcomes $\tilde{Y}^{\text{att}(g,t)}_{g',t_{\text{pre}}}$ that leveraged different (g', t_{pre}) pairs. The efficiency weights are given by $\mathbf{1}'\Omega^*_{gt}(X)^{-1}/\mathbf{1}\Omega^*_{gt}(X)^{-1}\mathbf{1}$. We recommend plotting the expected value of these weights so that one can have a better understanding of how each pre-treatment period and comparison group is leveraged for efficiency considerations. We do this in our simulations and empirical application.

Based on (3.13), one can straightforwardly form estimands for event study parameters $ES(e), e \ge 0$, by plugging (3.13) into (3.8),

$$ES(e) = ES_{stg}(e) := \sum_{g \in \mathcal{G}_{trt,e}} \frac{\pi_g}{\sum_{g \in \mathcal{G}_{trt,e}} \pi_g} ATT_{stg}(g, g + e)$$
(3.14)

An analogous procedure follows for ES_{avg} .

Remark 3.4 (Alternative Parallel Trends). In this section, we present semiparametric efficiency results for two types of parallel trends: Assumption PT-Post that uses never-treated units as the comparison group and does not restrict pre-treatment trends, and Assumption PT-All that imposes parallel trends in all periods and across all groups. We view these two PT assumptions as the ends of a spectrum, and we recognize that other types of parallel trend assumptions are also possible. It is straightforward to adjust our results in Theorem 3.2 to cover those intermediate cases, too. For instance, the PT assumption involving not-yet-treated units in Callaway and Sant'Anna (2021) restricts pre-treatment trends only once the first group is treated at period g_{\min} . In that particular case, data from period t = 1 until $t = g_{\min} - 2$ are not informative for any ATT(g, t), so we could drop them entirely from the analysis and apply Theorem 3.2 using the rest of the retained data. Gormley and Matsa (2011), Deshpande and Li (2019), Fadlon and Nielsen (2021), Cengiz, Dube, Lindner and Zipperer (2019), and Baker et al. (2022) discuss other parallel trends over a set of pre- and post-treatment periods that allow one to different untreated cohorts as a comparison group. Effectively, all one needs to do to cover these intermediate cases is to adjust the $\mathbb{IF}_{stg}^{att(g,t)}$ in (3.11) only to include moment equations that are justified by the given version of the parallel trends assumption.

4 Semiparametric efficient estimation and inference

In this section, we discuss how one can estimate and make inferences about the ATT(g,t)'s and ES(e)'s by leveraging the EIF-based estimands in (3.13) and (3.14). The estimator and inference procedures for DiD setups without covariates are straightforward, as they follow from standard plugin arguments and can be seen as a special case of our proposal. The results from DiD setups with a single treatment time follow as special cases, too.

We first discuss estimation for ATT(g,t)'s. The EIF-based estimand $ATT_{stg}(g,t)$ in (3.13) naturally suggests a two-step estimation procedure for ATT(g,t)'s, where one first estimate the nuisance parameters, $m(X) := ((m_{\infty,t,t'_{pre}}(X), t'_{pre} < t), (m_{g',t'_{pre},1}, g' \in \mathcal{G}_{trt}, t'_{pre} < g'))$ and $p_{ratio}(X) := (p_g(X)/p_{g'}(X), (g,g') \in \mathcal{G} \times \mathcal{G})$, the conditional covariance matrix $\Omega_{gt}^*(X)$ (or $\Omega_{gt}(X)$), and then use their fitted values to form a plug-in estimators fo ATT(g,t) based on (3.13).

When one uses parametric working models for the nuisance parameters, it is easy to show that the resulting estimator is doubly robust in the sense that it remains consistent for the ATT(g,t) as long as the working models for either (but not necessarily both) $p_{ratio}(X)$ or m(X) are correctly specified, regardless of the weighting scheme adopted; see, e.g., Sant'Anna and Zhao (2020, Theorem 1) and Callaway and Sant'Anna (2021, Theorem 2).¹⁰ Inference, in this case, follows from delta-method arguments.

¹⁰In our case, as our estimator averages across different working models for the propensity score and regression adjustment for untreated units, it is possible to refine the notion of double robustness in the sense that our ATT(g,t)estimator remains consistent if *weighted averages* of functionals of $p_{ratio}(X)$ or of m(X) are consistently estimated. As we advocate for an efficiency-oriented weighting scheme, we would not require it to hold for all possible weighted averages (which would essentially require that m(X) or $p_{ratio}(X)$ to be correctly specified).

In practice, however, parametric models may be too restrictive, and their misspecification can still lead to asymptotic biases. Thus, in what follows, we describe a flexible nonparametric estimation procedure that can leverage modern and traditional estimators for the nuisance functions.

First, recall that m(X) involves conditional expectations of changes in outcomes over time among some untreated units, given covariates, i.e., it involves terms like $m_{g',t'_{\text{pre}},t_{\text{pre}}}(X) = \mathbb{E}[Y_{t'_{\text{pre}}} - Y_{t_{\text{pre}}}|G = g', X]$ with $g' > \max\{t_{\text{pre}}, t'_{\text{pre}}\}$. This is nothing more than a nonparametric regression problem, and we can use a variety of estimators for it, including sieve-based (Chen, 2007), kernel-based (Cattaneo, Jansson and Ma, 2018), or recent machine learning methods such as random forests, lasso, ridge, deep neural nets, boosted trees, and the ensemble of these methods (Chernozhukov, Chetverikov, Demirer, Duflo, Hansen, Newey and Robins, 2018). Denote these estimators as $\hat{m}(X)$.

Analogously, $p_{\text{ratio}}(X)$ involves terms like ratios of $\mathbb{E}[G_g|X]$, and each of these conditional expectations can also be estimated using these same estimation procedures.¹¹ Note that modeling each propensity score separately and then constructing their ratios can potentially lead to instabilities associated with estimated propensity scores close to zero. As the propensity scores enter into $ATT_{\text{stg}}(g,t)$ (3.13) in a ratio format, it suffices to directly estimate $p_g(X)/p_{g'}(X)$, which is not bounded between zero and one, can be more flexibly modeled, and can lead to more stable procedures.

Towards this end, note that $p_g(X)/p_{g'}(X)$ can be found as the unique solution to the minimization problem:

$$\frac{p_g(X)}{p_{g'}(X)} = \operatorname*{arg\,min}_{r_{g,g'}} \mathbb{E}\left[\left(r_{g,g'}(X) - \frac{p_g(X)}{p_{g'}(X)}\right)^2 p_{g'}(X)\right] = \operatorname*{arg\,min}_{r_{g,g'}} \mathbb{E}\left[r_{g,g'}(X)^2 G_{g'} - 2r_{g,g'}(X)G_g\right].$$
(4.1)

This equivalence allows the use of $\mathbb{E}\left[r_{g,g'}(X)^2 G_{g'} - 2r_{g,g'}(X)G_g\right]$ as a (convex) loss function to estimate $r_{g,g'}(X) := p_g(X)/p_{g'}(X)$, and different estimation procedures can be used. An easy-to-implement estimator for $r_{g,g'}(X)$ is the sieve-based estimator $\hat{r}_{g,g'}(X) = \left(\psi^K(X)'\hat{\beta}_K\right)$, where $\psi^K(x)$ is a K-dimensional vector of flexible transformations of the X such as (tensor products of) cubic B-splines,

$$\widehat{\beta}_{K} = \underset{\beta_{K}}{\operatorname{arg\,min}} \mathbb{E}_{n} \left[G_{g'} \left(\psi^{K}(X)' \beta_{K} \right)^{2} - 2G_{g} \left(\psi^{K}(X)' \beta_{K} \right) \right], \tag{4.2}$$

and we avoid indexing $\hat{\beta}_K$ and β_K by (g, g') to simplify the notation burden. The ((g, g')-specific) sieve index K can be selected using some information criteria as follows:

$$\widehat{K} = \underset{K}{\operatorname{arg\,min}} 2\mathbb{E}_{n} \left[G_{g'} \left(\psi^{K}(X)' \widehat{\beta}_{K} \right)^{2} - 2G_{g} \left(\psi^{K}(X)' \widehat{\beta}_{K} \right) \right] + \frac{C_{n}K}{n},$$

where the information criterion corresponds to the Akaike information criterion (AIC) when $C_n = 2$, or the Bayesian information criterion (BIC) when $C_n = \log(n)$. In the appendix, we show the consistency of the estimator based on \hat{K} following Chen and Liao (2014). Similarly to (4.1), we can

¹¹It is often desirable to impose that $p_g(X)$ is bounded between zero and one and choose estimators that respect this constraint, such as multinomial series logit/probit and local multinomial logit/probit estimators. See, e.g., Chen (2007) for a discussion of likelihood-based sieve estimators, and Staniswalis (1989) for kernel-based ones.

estimate $s_{g'} := 1/p_{g'}(X)$ by solving the empirical analog of the minimization problem

$$s_{g'}(X) = \operatorname*{arg\,min}_{a_g} \mathbb{E}\left[a_g(X)^2 G_g - 2a_g(X)\right].$$

For instance, one can use sieve-estimators analogous to (4.2) for this task.

To estimate the conditional covariance terms of Ω_{gt}^* as defined in (3.12), we propose using a Nadaraya-Watson-type estimator based on kernel smoothing. Let Ker be a kernel function on the covariates space and h > 0 a bandwidth. Denote $K_h(\cdot) = Ker(\cdot/h)/h$. For $x = X_i$, we use $\widehat{Cov}(Y_t - Y_{t_{\text{pre}}}, Y_t - Y_{t'_{\text{pre}}}|G = g', X = x)$ as an estimator for each the covariance terms $Cov(Y_t - Y_{t_{\text{pre}}}, Y_t - Y_{t'_{\text{pre}}}|G = g', X = x)$, where $\widehat{Cov}(Y_t - Y_{t_{\text{pre}}}, Y_t - Y_{t'_{\text{pre}}}|G = g', X = x)$, where $\widehat{Cov}(Y_t - Y_{t_{\text{pre}}}, Y_t - Y_{t'_{\text{pre}}}|G = g', X = x)$ is define as

$$\frac{\sum_{i':G_{i'}=g'} K_h(X_{i'}-x)(Y_{i',t}-Y_{i',t_{\text{pre}}}-\hat{m}_{g',t,t_{\text{pre}}}(X_{i'}))(Y_{i',t}-Y_{i',t'_{\text{pre}}}-\hat{m}_{g',t,t'_{\text{pre}}}(X_{i'}))}{\sum_{i':G_{i'}=g'} K_h(X_{i'}-x)}$$

Based on these estimators for the propensity score ratios and the conditional covariance terms, we estimate Ω_{gt}^* by $\hat{\Omega}_{gt}^*$ with each (j, k)-th element given by the plug-in estimators of (3.12).

Based on these estimators for the nuisance functions and weights, we can estimate ATT(g,t) by¹²

$$\widehat{ATT}_{\text{stg}}(g,t) = \mathbb{E}_n \left[\frac{\mathbf{1}' \widehat{\Omega}_{gt}^*(X)^{-1}}{\mathbf{1}' \widehat{\Omega}_{gt}^*(X)^{-1} \mathbf{1}} \widehat{\widetilde{Y}}_{\text{stg}}^{\text{att}(g,t)} \right],$$
(4.3)

where $\hat{\widetilde{Y}}_{stg}^{att(g,t)}$ the estimated analog of $\widetilde{Y}_{stg}^{att(g,t)}$, with each $\hat{\widetilde{Y}}_{g',t_{pre}}^{att(g,t)}$ given by

$$\widehat{\widehat{Y}}_{g',t_{\rm pre}}^{\rm att(g,t)} = \frac{G_g}{\widehat{\pi}_g} (Y_t - Y_1 - \widehat{m}_{\infty,t,t_{\rm pre}}(X) - \widehat{m}_{g',t_{\rm pre},1}(X)) \\
- \widehat{r}_{g,\infty}(X) \frac{G_\infty}{\widehat{\pi}_g} (Y_t - Y_{t_{\rm pre}} - \widehat{m}_{\infty,t,t_{\rm pre}}(X)) - \widehat{r}_{g,g'}(X) \frac{G_{g'}}{\widehat{\pi}_g} (Y_{t_{\rm pre}} - Y_1 - \widehat{m}_{g',t_{\rm pre},1}(X)) \Big), \quad (4.4)$$

and $\hat{\pi}_g = \mathbb{E}_n[G_g].$

The estimator for the event study parameter is

$$\widehat{ES}(e) = \sum_{g \in \mathcal{G}_{trt}} \frac{\widehat{\pi}_g}{\sum_{g' \in \mathcal{G}_{trt,e}} \widehat{\pi}_{g'}} \widehat{ATT}_{stg}(g, g+e).$$
(4.5)

The following theorem establishes the large sample properties of our proposed DiD and ES estimators, and highlights that they achieve the semiparametric efficiency bound.

Theorem 4.1. Let Assumptions O, NA and PT-All and the regularity conditions listed in Assumption C.1 in the Appendix hold. Then, our proposed nonparametric estimator for $\widehat{ATT}_{stg}(g,t)$ is consistent, asymptotically normal, and achieves the semiparametric efficiency bound, i.e., as $n \to \infty$,

$$\begin{split} \sqrt{n}(\widehat{ATT}_{stg}(g,t) - ATT(g,t)) &= \frac{1}{\sqrt{n}} \sum_{i=1}^{n} \frac{\mathbf{1}' \widehat{\Omega}_{gt}^*(X_i)^{-1}}{\mathbf{1}' \widehat{\Omega}_{gt}^*(X_i)^{-1} \mathbf{1}} \mathbb{IF}_{stg}^{att(g,t)}(W_i) + o_p(1) \\ &\stackrel{d}{\to} N(0, V_{eff}), \end{split}$$

with $V_{eff} = \operatorname{Var}\left(\frac{\mathbf{1}'\hat{\Omega}_{gt}^{*}(X)^{-1}}{\mathbf{1}'\hat{\Omega}_{gt}^{*}(X)^{-1}\mathbf{1}} \mathbb{IF}_{stg}^{att(g,t)}(W)\right)$ being the semiparametric efficiency bound in Theorem 3.2.

 $^{^{12}}$ It is straightforward to consider variants of our estimator using sample-splitting arguments, too.

Consequently, $\widehat{ES}(e)$ is also asymptotically normal and semiparametrically efficient.

Theorem 4.1 provides the pointwise asymptotic normality results for each (g, t) pair, with $t \ge g$. In practice, researchers are often interested in making inferences about multiple groups and periods to better understand treatment effect dynamics and heterogeneity. In such cases, to avoid multipletesting problems, it is recommended to construct simultaneous confidence bands. We omit the details of this procedure as it is standard.¹³ To compute standard errors, one can leverage a multiplier bootstrap procedure or take the square root of the average of the estimated EIF squared divided by the sample size.

Remark 4.1. Although the results in Theorem 4.1 focus on the large sample properties of our efficient DiD estimators, $\widehat{ATT}_{stg}(g,t)$, in its proof, we also establish consistency and asymptotically normality for related estimators for the ATT(g,t)'s that replaces the estimators for the efficient weights $\mathbf{1}'\hat{\Omega}_{gt}^*(X)^{-1}/\mathbf{1}'\hat{\Omega}_{gt}^*(X)^{-1}\mathbf{1}$ with any consistent estimator for any weights w(x) that sum up to one for each covariate value x. Interestingly, we do not need to require rate conditions for these weights, allowing one to use many different weights as discussed in Lemma 2.1.

Remark 4.2. As our proposed estimators are based on efficient influence functions, they satisfy the Neyman orthogonality condition by construction. This implies that it is relatively straightforward to use modern machine learning estimators for the nuisance parameter without incurring a loss of asymptotic semiparametric efficiency in large samples.¹⁴

4.1 Efficient estimation when covariates are not present

When Assumptions NA and PT-All hold unconditionally, the efficiency results can be simplified by dropping X from the conditioning set. This leads to an even simpler construction of the efficient estimator. In particular, (3.13) reduces to

$$ATT(g,t) = \frac{\mathbf{1}'(\Omega_{gt}^*)^{-1}}{\mathbf{1}'(\Omega_{qt}^*)^{-1}\mathbf{1}}\widetilde{Y}^{\operatorname{att}(g,t)},$$

where each entry $\widetilde{Y}^{\rm att(g,t)}_{g',t_{\rm pre}}$ of $\widetilde{Y}^{\rm att(g,t)}$ is given by

$$\sum_{g',t_{\rm pre}}^{\text{fatt}(g,t)} = \mathbb{E}[Y_t - Y_1 | G = g] - (\mathbb{E}[Y_t - Y_{t_{\rm pre}} | G = \infty] + \mathbb{E}[Y_{t_{\rm pre}} - Y_1 | G = g']).$$

The covariance matrix Ω_{qt}^* has (j, k)-th element given by

$$\frac{1}{\pi_g} \operatorname{Cov}(Y_t - Y_1, Y_t - Y_1 | G = g) + \frac{1}{\pi_\infty} \operatorname{Cov}(Y_t - Y_{t'_j}, Y_t - Y_{t'_k} | G = \infty)$$
$$-\frac{\mathbf{1}\{g = g'_j\}}{\pi_g} \operatorname{Cov}(Y_t - Y_1, Y_{t'_j} - Y_1 | G = g) - \frac{\mathbf{1}\{g = g'_k\}}{\pi_g} \operatorname{Cov}(Y_t - Y_1, Y_{t'_k} - Y_1 | G = g)$$

¹³See, e.g., Theorems 2 and 3, Algorithm 1, and Corollaries 1 and 2 in Callaway and Sant'Anna (2021).

¹⁴Many machine learning procedures involve Neyman orthogonal moments with sample-splitting, which can lead to a loss of precision in DiD and ES estimation in small samples. This phenomenon is not unique to DiD; see, e.g., Chen, Chen and Tamer (2023) for a discussion.

$$+\frac{\mathbf{1}\{g_j = g'_k\}}{\pi_{g'_j}}\operatorname{Cov}(Y_{t'_j} - Y_1, Y_{t'_k} - Y_1 | G = g'_j),$$

where the indices (g'_s, t'_s) are as defined in (3.12).

In estimation, one can replace the expectations and covariances for each group with the corresponding within-group sample means and sample covariances. This procedure does not involve choosing tuning parameters or estimating conditional expectations.

When covariates are absent, the efficient GMM estimator also attains the semiparametric efficiency bound. Our DiD and ES estimators match that efficiency yet are substantially simpler to compute: they are available in closed form, whereas the optimal GMM approach entails solving a high-dimensional optimization problem that is often not practical. In addition, the construction of our efficient estimator provides new insights about how to weigh different pre-treatment and comparison groups to achieve efficiency—it is only a matter of analyzing the efficient weights $\mathbf{1}'(\Omega_{qt}^*)^{-1}/\mathbf{1}'(\Omega_{qt}^*)^{-1}\mathbf{1}$.

5 Monte Carlo simulations

So far, we have motivated and established that our efficient DiD and ES estimators have attractive efficiency properties that allow researchers to get more precise inference procedures. In this section, we aim to see how these properties play out in realistic empirical settings. We build on Arkhangelsky et al. (2021) and Baker et al. (2022) and consider two sets of simulation studies calibrated to datasets people have used for DiD. The first set of simulations builds on Arkhangelsky et al. (2021). It considers a single treatment date and leverages Current Population Survey (CPS) data for constructing a variety of outcomes and treatment assignments for a panel of US states. The second set of simulations builds on Baker et al. (2022). It explores Compustat data to calibrate outcomes and treatment assignments in staggered treatment setups with treatment timing varying across states. To simplify exposure and match the simulation designs in Arkhangelsky et al. (2021) and Baker et al. (2022), covariates play no important role in these Monte Carlo setups.

In our first set of simulations, we consider ES_{avg} as in (2.3) as the parameter of interest and compare our efficient DiD plug-in estimator based on (3.6) (EDiD), traditional OLS-based TWFE estimator for β based on (2.6) (TWFE), average of post-treatment event-study TWFE estimators β_e based on (2.4) (DTWFE), and the Arkhangelsky et al. (2021) Synthetic DiD estimator (SDiD). We note that the DTWFE estimator would be the efficient estimator if parallel trends were to hold only in post-treatment periods.¹⁵

In our second set of simulations, we also consider ES_{avg} as in (2.3) as the parameter of interest and compare our efficient (plug-in) estimator \widehat{ES}_{avg} based on (4.5) (EDiD), the average of post-

¹⁵In this non-staggered setup, the OLS estimate of β based on (2.6) coincides with the post-treatment average of the event-studies coefficients of the estimators proposed by Wooldridge (2021), Gardner (2021), and Borusyak et al. (2024). Similarly, the post-treatment average of the OLS estimates of β_e in (2.4) coincides with Callaway and Sant'Anna (2021), Sun and Abraham (2021), and de Chaisemartin and D'Haultfœuille (2020) estimators.

treatment event-study estimates based on (i) Callaway and Sant'Anna (2021) and Sun and Abraham (2021) DiD estimators using the "never-treated" as a comparison group (CS-SA), (ii) Callaway and Sant'Anna (2021) and de Chaisemartin and D'Haultfœuille (2020) DiD estimators using the "notyet-treated" as a comparison group (CS-dCDH), and (iii) Wooldridge (2021), Gardner (2021), and Borusyak et al. (2024) "imputation" estimators (BJS-G-W). We do not include staggered synthetic DiD estimators in the comparison as we are not aware of a paper describing and establishing the statistical properties of synthetic DiD estimators for ES_{avg} . In this exercise, we note that the CS-SA estimator would be the efficient one if parallel trends were to hold only for post-treatment periods and the never-treated group was the only valid comparison group.

We compare all these estimators regarding bias, root-mean-squared-error (RMSE) relative to our efficient DiD estimator, 95% empirical coverage, and 95% confidence interval length relative to our efficient DiD. We consider analytical and bootstrapped standard errors for coverage and length of confidence intervals but use the Gaussian critical values. For each data generating process (DGP), we consider 1,000 Monte Carlo experiments, and for each experiment, we use 300 bootstrap draws.¹⁶

5.1 Simulations based on CPS with single treatment date

Our first set of simulations builds on Arkhangelsky et al. (2021) and explores CPS data to construct an empirically motivated data-generating process. We differ from Arkhangelsky et al. (2021) in some aspects: (i) we consider heterogeneous treatment effects across units; (ii) consistent with our DiD regime with "large n" and fixed T, we consider short panels with T = 7—four pre-treatment and three post-treatment periods; (iii) we do not limit the maximum number of treated units in a given simulation; and (iv) all our outcomes are measured in log to avoid violating support restrictions in the simulation. Apart from these differences, the construction of our data-generating process follows from Arkhangelsky et al. (2021), which we describe below for completeness.

As in Arkhangelsky et al. (2021), we start designing our simulations using data on wages for women with positive wages in the March outgoing rotation groups in the CPS from 1979 to 2018. We then take logs and average the observations by state-year cells, leaving us with aggregate data from 50 states in 40 years. This will serve as the baseline dataset for our simulation designs.

Next, we consider a baseline specification for our untreated potential outcomes, $Y_{i,t}(\infty)$. We consider the same interactive fixed-effects specification as Arkhangelsky et al. (2021),

$$Y_{i,t}(\infty) = \gamma_i \upsilon_t + \varepsilon_{i,t},\tag{5.1}$$

with γ_i and υ_t being 4-dimensional vectors of latent unit and time factors, and $\varepsilon_{i,t}$ being a mean-zero Gaussian error term that follows an AR(2) process. This specification does not impose a two-way fixed effects structure on $Y_{i,t}(\infty)$ and allows for correlation over times within each state; accounting for serial correlation is very important when conducting valid inference using DiD and other panel

¹⁶For Synthetic DiD, Arkhangelsky et al. (2021) have not proposed a plug-in analytical standard error. As such, we do not report plug-in analytical standard errors from it.

data methods, see, e.g., Wooldridge (2003) and Bertrand, Duflo and Mullainathan (2004).

To construct a realistic set of simulations, we use the CPS aggregated data to estimate the γ_i 's, υ_t 's, and the variance-covariance matrix of the error terms. In matrix notation, the interactive fixed-effects model (5.1) is given by $Y(\infty) = L + E$, where $L = \Gamma \Upsilon'$, and we estimate L as

$$L := \arg\min_{L:rank(L)=4} \sum_{i,t} (Y_{i,t}^* - L_{i,t})^2,$$

where $Y_{i,t}^*$ is the log wage in state *i* in year *t* in the CPS data. To help with interpretation, we decompose *L* as unit-and-time fixed effects, *F*, and an interactive term *M*, with

$$F_{i,t} = \alpha_t + \eta_i = \frac{1}{T} \sum_{l=1}^T L_{i,l} + \frac{1}{N} \sum_{j=1}^N L_{j,t} - \frac{1}{NT} \sum_{i,t}^{N \times T} L_{i,t}$$
$$M_{i,t} = L_{i,t} - F_{i,t}.$$

We estimate the variance-covariance matrix of the residuals $\varepsilon_{i,t}$ by fitting an AR(2) model to the residuals of $Y_{i,t}^* - L_{i,t}$, and then compute the implied variance-covariance matrix of the error term for state i, Σ . Like Arkhangelsky et al. (2021), we assume that the error terms are independent across units and Σ is constant across states.

Next, we describe the treatment assignment process. In this set of simulations, units can start treatment in t = 2009, so we only have two groups: $G_i = 2009$ (treated units) and $G_i = \infty$ (untreated units). To simulate whether a state is treated, we follow Arkhangelsky et al. (2021) and consider the treatment status $D_{i,t} = 1\{t \ge 2009\}1\{G_i = 2009\}$ with

$$1\{G_i = 2009\} \sim \text{Bernoulli} (\pi_i),$$

$$\pi_i = P(G_i = 2009 | \eta_i, M_i) = \frac{\exp(\phi_\eta \eta_i + \phi_M M_i)}{1 + \exp(\phi_\eta \eta_i + \phi_M M_i)}.$$

We choose ϕ_{η} and ϕ_M as the coefficient estimates from a logistic regression of an observed binary characteristic of the state *i* on η_i and M_i .¹⁷ We consider three different characteristics relating treatment groups to minimum wage laws (part of our baseline specification), abortion rights, and gun control laws. We also consider a completely random treatment assignment. Unlike Arkhangelsky et al. (2021), though, we do not restrict the maximum number of treated units; only the minimum number of treated units to be two.¹⁸

Based on these parameters, we simulate untreated and treated potential outcomes $Y_{i,t}(\infty)$ and $Y_{i,t}(2009)$ as

$$Y_{i,t}(\infty) = \alpha_t + \eta_i + M_{i,t} + \varepsilon_{i,t}$$
(5.2)

$$Y_{i,t}(2009) = Y_{i,t}(\infty) + \tau_i \{ t \ge 2009 \}$$
(5.3)

where $\varepsilon_i = (\varepsilon_{i,t=1}, \ldots, \varepsilon_{i,t=T})'$ have a multivariate Gaussian distribution with mean zero and variance

¹⁷These "unit factors" are the first four left singular vectors from single value decomposition of the outcome Y, as described in Arkhangelsky et al. (2021).

¹⁸If we have less than two treated states in a simulation, we randomly select ten out of the 50 states to be treated completely at random. This is similar to Arkhangelsky et al. (2021), though their minimum number of treated states was one.

 Σ , and τ_i have a Gaussian distribution with mean zero and variance one. Note that we have a heterogeneous treatment effects model as the treatment effect varies across states. This differs from Arkhangelsky et al. (2021), who sets $\tau_i = 0$ for all states. However, we stress that ATT(2009, t) = 0 for all periods. The observed outcomes for unit *i* in time *t* are $Y_{i,t} = 1\{G_i = 2009\}Y_{i,t}(2009) + 1\{G_i = \infty\}Y_{i,t}(\infty)$. To respect our DiD setup with "large *n* and fixed *T*", we restrict T = 7 and keep data from t = 2005 until t = 2011.

Note that we do not impose that the parallel trends assumption holds exactly across periods. If this assumption is violated, we should see a bias in estimates that rely on this assumption, namely our efficient DiD estimator and the DTWFE in (2.4). If parallel trends in post-treatment periods is violated, we should also see a bias in the TWFE estimates based on (2.6).

Like Arkhangelsky et al. (2021), we consider different choices of treatment assignment and outcome variables. In addition, we also consider settings where we drop different components of the data generating process for the outcome, such as assuming that the error terms are serially independent ("No Corr"), that there is no interactive component M ("No M"), no additive two-way fixed effects ("No F"), or there is L ("Only noise").¹⁹ For the alternative outcome variables, we consider *log* of hours worked and *log* of unemployment rate; Arkhangelsky et al. (2021) consider these variables in levels, but that leads to some negative values during the simulation for these variables, violating their natural support restriction.

Finally, we consider two setups. We simulate data from the 50 states and seven periods for the first one. In this case, we expect that inference procedures based on asymptotic approximations may be challenging given the limited (effective) sample size of 50. In the second setup, we increase the number of cross-sectional units by drawing with replacement 200 states (each with its own error term, treatment assignment, and treatment effect). Here, we expect inference procedures based on asymptotic approximations to be better than in the setup with n = 50, though we acknowledge that n = 200 is still reasonably small.

Tables 1 present the bias and relative RMSE of the four different estimators for ES_{avg} when n = 50 and n = 200. At a high level, we have found that all estimators are (nearly) unbiased when n = 50 and that the bias becomes even smaller when n = 200. This suggests that, in these DGPs, the assumption that parallel trends hold in all periods is a good approximation. In terms of RMSE, however, there is a significant variation across estimators, with our efficient DiD estimator performing the best in nearly all DGPs, with Synthetic DiD coming as second, and the DTWFE estimator based on (2.4) being the least precise. The result for DTWFE is expected, though, as this estimator does not leverage data from any pre-treatment period other than t = 2008; all the other estimators leverage three additional pre-treatment periods. It is also important to stress the large RMSE gains that our efficient DiD estimator enjoys compared to the synthetic DiD estimator. For

¹⁹We omit results for DGPs in which there is no error term ("No noise"). We do it because the RMSE and the length of 95% confidence interval for our efficient estimators are close to zero, making it hard to report relative performance measures.

		Relative RMSE					Bias $(\times 10)$				
	Sample size	EDiD	TWFE	DTWFE	SDiD	Ε	DiD	TWFE	DTWFE	SDiD	
1. Baseline	50	1	3.57	12.46	1.53		0.01	0.60	2.58	0.00	
	200	1	2.32	3.37	1.95		0.00	-0.01	0.00	0.00	
Outcome Model											
2. No corr	50	1	3.52	12.33	1.45		0.02	0.59	2.46	0.00	
	200	1	2.27	3.33	1.95	-	0.01	0.00	-0.01	0.00	
3. No M	50	1	3.67	13.04	1.49	-	0.02	0.61	2.45	0.03	
	200	1	2.17	3.00	1.68	-	0.01	0.02	0.00	0.01	
4. No F	50	1	1.47	1.88	1.42		0.00	-0.02	-0.04	-0.02	
	200	1	1.64	2.18	1.63		0.00	0.00	-0.01	0.00	
5. Only noise	50	1	1.26	1.67	1.25		0.02	-0.02	-0.01	-0.01	
	200	1	1.47	1.87	1.48		0.00	-0.01	-0.01	-0.01	
Treatment Assignment											
6. Gun law	50	1	7.27	18.23	4.67		0.00	-0.08	-0.23	-0.11	
	200	1	8.70	12.94	6.05		0.00	0.03	0.02	0.02	
7. Abortion	50	1	6.99	17.19	4.77	-	0.01	0.55	1.75	0.33	
	200	1	8.04	12.64	5.23		0.00	-0.01	-0.05	0.00	
8. Random	50	1	4.13	15.30	1.53	-	0.01	-0.07	-0.22	0.01	
	200	1	2.56	3.74	2.17		0.00	0.01	0.00	0.01	
Outcome Variable											
9. Ln Hours	50	1	1.01	1.92	0.95	-	0.34	0.12	1.53	0.02	
	200	1	1.24	1.95	1.21		0.01	0.05	0.02	0.07	
10. Ln U-rate	50	1	0.82	1.44	0.82		0.73	-0.24	-0.36	-0.26	
	200	1	1.03	1.53	1.01		0.03	0.00	0.01	0.00	

Table 1: Simulation results for CPS data: Relative RMSE and Bias

Notes: Simulation results for CPS data based on 1,000 Monte Carlo experiments. The DGP is summarized in Section 5. The baseline case uses state minimum wage laws to simulate treatment assignment and generates outcomes using the DGPs described in Section 5. In subsequent settings, we omit parts of the DGP (rows 2–5), consider different distributions for the treatment variable (rows 6-8), and different distributions for the outcome variable (rows 9-10). The dataset has 7 times periods and either n = 50 or n = 200. Relative RMSE is reported using our efficient DiD estimator as benchmark. All results for bias are multiplied by 10 for readability.

instance, when n = 50, the RMSE of the synthetic DiD estimator is nearly 60% larger than that of our efficient DiD estimator in the baseline model. When treatment assignment is based on gun laws or abortion rights, their RMSE is approximately 5 times larger than ours when n = 50, and this difference grows further when n = 200. On the other hand, we notice that when the outcome of interest is the log unemployment rate, the synthetic DiD and the TWFE estimators have smaller RMSE than our efficient estimator. However, this difference reduces when n = 200.

Table 2 present the empirical coverage of 95% confidence intervals for ES_{avg} when n = 50 and when n = 200. We consider both nonparametric clustered-bootstrapped standard errors and analytical standard errors. Our simulation results show that using bootstrapped standard errors to construct confidence intervals yields very good coverage properties for all estimators across all considered DGPs, even in the challenging setup with n = 50. On the other hand, our results also reveal that inference based on analytical standard errors with our efficient DiD estimator is anticonservative

		Bootstrap				Analytical					
	Sample size	EDiD	TWFE	DTWFE	SDiD		EDiD	TWFE	DTWFE	SDiD	
1. Baseline	50	0.94	0.92	0.93	0.93		0.80	0.92	0.92	-	
	200	0.96	0.97	0.97	0.97		0.94	0.98	0.97	-	
Outcome Model											
2. No corr	50	0.92	0.94	0.94	0.93		0.80	0.93	0.92	-	
	200	0.93	0.97	0.97	0.96		0.93	0.97	0.97	-	
3. No M	50	0.94	0.93	0.93	0.94		0.82	0.92	0.92	-	
	200	0.93	0.97	0.97	0.96		0.91	0.96	0.97	-	
4. No F	50	0.94	0.94	0.94	0.95		0.82	0.94	0.92	-	
	200	0.95	0.95	0.96	0.96		0.94	0.95	0.96	-	
5. Only noise	50	0.93	0.93	0.92	0.95		0.79	0.92	0.90	-	
	200	0.94	0.95	0.95	0.95		0.93	0.94	0.95	-	
Treatment Assignment											
6. Gun law	50	0.95	0.94	0.94	0.97		0.94	0.94	0.94	-	
	200	0.94	0.96	0.96	0.97		0.94	0.96	0.96	-	
7. Abortion	50	0.94	0.92	0.91	0.93		0.92	0.93	0.91	-	
	200	0.95	0.98	0.98	0.97		0.95	0.97	0.98	-	
8. Random	50	0.94	0.94	0.94	0.98		0.93	0.94	0.94	-	
	200	0.95	0.93	0.93	0.94		0.95	0.94	0.93	-	
Outcome Variable											
9. Ln Hours	50	0.93	0.94	0.93	0.95		0.77	0.94	0.92	-	
	200	0.94	0.95	0.96	0.96		0.93	0.95	0.95	-	
10. Ln U-rate	50	0.93	0.91	0.92	0.92		0.79	0.90	0.91	-	
	200	0.95	0.94	0.95	0.95		0.92	0.94	0.94	-	

Table 2: Simulation results for CPS data: Coverage

Notes: Simulation results for CPS data based on 1,000 Monte Carlo experiments. The DGP is summarized in Section 5. The dataset has 7 times periods and either n = 50 or n = 200. "Bootstrap" results use 300 nonparametric clustered bootstrap replications to compute bootstrapped clustered standard errors for all methods. "Analytical" results use asymptotic approximations for computing standard errors for all methods but synthetic DiD, as analytical standard errors are not available. All the other details are the same as in Table 1.

when n = 50, though the results improve when n = 200. Thus, in setups with very limited sample sizes, we recommend using cluster bootstrap standard errors.

Finally, Table 3 presents the length of 95% confidence intervals for all the estimators, using bootstrap and analytical standard errors. Since it is only reasonable to compare the length of confidence intervals across methods that control size, we focus our attention on the bootstrap results. Here, we note that our efficient DiD estimator tends to have substantially shorter confidence intervals than other available DiD estimators. For instance, in the baseline DGP, the bootstrapped confidence interval for the synthetic DiD estimator is more than 40% larger than our efficient DiD estimator when n = 50, and approximately two times larger than our efficient DiD estimator when n = 200; these gains are even larger when compared to TWFE and DTWFE estimators. When treatment assignment is based on gun laws across states, the length of bootstrapped confidence intervals of all other considered estimators is more than four times larger than those based on our efficient estimator. We reach a qualitatively similar conclusion when using abortion rights as treatment. On the other hand, when the outcome of interest is the log unemployment rate, we note that bootstrap inference using synthetic DiD or traditional TWFE estimators is more precise than our efficient DiD estimator. However, this difference shrinks when n = 200.

		Bootstrap					Analytical					
	Sample size	EDiD	TWFE	DTWFE	SDiD		EDiD	TWFE	DTWFE	SDiD		
1. Baseline	50	1.00	3.63	12.58	1.43		1.00	5.44	18.40	-		
	200	1.00	2.46	3.60	2.08		1.00	2.65	3.86	-		
Outcome Model												
2. No corr	50	1.00	3.67	12.78	1.36		1.00	5.46	18.57	-		
	200	1.00	2.47	3.58	2.11		1.00	2.67	3.85	-		
3. No M	50	1.00	3.56	12.90	1.32		1.00	5.32	18.83	-		
	200	1.00	2.33	3.39	1.84		1.00	2.52	3.64	-		
4. No F	50	1.00	1.42	1.85	1.44		1.00	2.14	2.72	-		
	200	1.00	1.58	2.18	1.57		1.00	1.71	2.35	-		
5. Only noise	50	1.00	1.15	1.55	1.23		1.00	1.74	2.28	-		
	200	1.00	1.31	1.74	1.34		1.00	1.41	1.87	-		
Treatment Assignment												
6. Gun law	50	1.00	6.41	16.94	4.33		1.00	7.00	18.12	-		
	200	1.00	8.86	13.18	5.92		1.00	8.99	13.33	-		
7. Abortion	50	1.00	6.40	16.44	3.84		1.00	7.14	18.00	-		
	200	1.00	8.53	13.30	5.44		1.00	8.76	13.53	-		
8. Random	50	1.00	3.93	14.71	2.13		1.00	4.31	15.71	-		
	200	1.00	2.36	3.49	2.03		1.00	2.40	3.53	-		
Outcome Variable												
9. Ln Hours	50	1.00	1.05	2.01	1.03		1.00	1.58	2.95	-		
	200	1.00	1.37	2.07	1.33		1.00	1.48	2.21	-		
10. Ln U-rate	50	1.00	0.75	1.34	0.79		1.00	1.13	1.98	-		
	200	1.00	0.99	1.51	0.99		1.00	1.07	1.62	-		

Table 3: Simulation results for CPS data: Relative length of confidence interval

Notes: Simulation results for CPS data based on 1,000 Monte Carlo experiments. The data generating process is summarized in Section 5. The dataset has 7 times periods and either n = 50 or n = 200. "Bootstrap" results use 300 nonparametric clustered bootstrap replications to compute bootstrapped clustered standard errors for all methods. "Analytical" results use asymptotic approximations for computing standard errors for all methods but synthetic DiD, as analytical standard errors are not available. Length of Confidence intervals are measured relative to our efficient DiD estimator. All the other details are the same as in Table 1.

In practice, we anticipate that empirical researchers may also be interested in better understanding the contribution of each DiD component that uses a different pre-treatment period as a baseline period. In other words, one may be interested in understanding the "contribution" of each pre-treatment period to form efficient estimators. Figure 2 plots this in a heatmap-style for four different DGPs, using a single simulation draw. As evident from this, the optimal way to aggregate pre-treatment periods to gain efficiency varies with the DGP, and the further it deviates from uniform weights, the higher the gains are compared to TWFE. This highlights that, in general, we can substantially improve upon TWFE.

Another point worth stressing is related to negative efficiency weights. In our context, negative



Figure 2: Contribution of pre-treatment periods for the Efficient DiD estimator for ES_{avg}

Notes: The figure displays the contribution of each pre-treatment period when constructing the efficient DiD estimator for ES_{avg} based on a single draw of the data-generating process. Sample size n = 50 and T = 7, with four pre-treatment periods and three post-treatment periods. The weight color scale is the same across all sub-figures except Figure 2c.

efficiency weights do not lead to concerns related to treatment effect sign preservation as discussed in Goodman-Bacon (2021), Borusyak et al. (2024), and de Chaisemartin and D'Haultfœuille (2020), for example. The reason for this is that our overidentification conditions imply that the ATT(g,t)'s are *homogeneous* across baseline periods, allowing us to potentially non-convex sums across different baseline periods. As illustrated in Figure 2c and backed up by the results in Tables 1 and 3, this should not be interpreted as a concern. Negative efficiency weights arise simply as a consequence of the dependence structure of the changes in outcome in treatment and comparison groups, and de facto leveraging them can lead to substantial gains in precision, as illustrated in DGP 6.

Altogether, these simulation results highlight the excellent finite sample properties of our efficient DiD estimator in realistic DGPs: they tend to have much smaller RMSE and shorter confidence intervals than other available estimators. Importantly, the gains in precision can be very substantial.

5.2 Simulations based on Compustat with staggered treatment

Our second set of simulations builds on Baker et al. (2022) and considers DGPs with staggered treatment adoption calibrated to Compustat panel data. We differ from Baker et al. (2022) in some aspects: (i) consistent with our DiD framework with "large n" and fixed T, we consider a setup with n = 400 firms and follow them for T = 11 years; (ii) we consider that error terms are serially correlated with varying autocorrelation parameter ρ to allow for richer outcome dynamics and assess its impact of the performance of different estimators; (iii) we follows a one-step treatment assignment in which firms are assigned to different cohorts. Apart from these differences, the construction of our data-generating process follows that of Baker et al. (2022), which allows for three treatment dates, with all units eventually being treated, and treatment effects evolving dynamically with heterogeneous trends. We describe the DGPs below for completeness.

As in Baker et al. (2022), we begin with a sample of all firms in Compustat over the 36-year period from 1980 to 2015 that are U.S. incorporated, non-financial, and contain at least five observations. This leaves us with a sample of 12,020 different firms. Using this unbalanced panel, we compute returns on assets (ROA) and then decompose it into year fixed effects, firm/unit fixed effects, and residuals:

$$\{\widehat{\eta}_i\}_{i=1}^{12,020}, \ \{\widehat{\alpha}_t\}_{t=1}^{36}, \ \{\widehat{\varepsilon}_{i,t}\}_{i,t=1}^{176,622}, \text{ from } ROA_{i,t} = \eta_i + \alpha_t + \varepsilon_{i,t}.$$

The empirical (marginal) distribution of these terms serves as the benchmark distribution in our Monte Carlo simulations. We consider three possible treatment dates, G = 5, G = 8, and G = 11, with each firm randomly assigned to each treatment group with probability 1/3.

Based on these features, we simulate potential outcomes $Y_{i,t}(5)$, $Y_{i,t}(8)$, $Y_{i,t}(11)$, and $Y_{i,t}(\infty)$ as

$$Y_{i,t}(\infty) = \widetilde{\alpha}_t + \widetilde{\eta}_i + \widetilde{\varepsilon}_{i,t}$$
(5.4)

$$Y_{i,t}(5) = Y_{i,t}(\infty) + 0.5 \times \sigma_{ROA} 1\{t \ge 5\} (t - 5 + 1)$$
(5.5)

$$Y_{i,t}(8) = Y_{i,t}(\infty) + 0.3 \times \sigma_{ROA} 1\{t \ge 8\} (t - 8 + 1)$$
(5.6)

$$Y_{i,t}(11) = Y_{i,t}(\infty) + 0.1 \times \sigma_{ROA} 1\{t \ge 11\} (t - 11 + 1), \qquad (5.7)$$

where $\tilde{\eta}_i$ and $\tilde{\alpha}_t$ are drawn from the empirical distribution of $\{\hat{\eta}_i\}_{i=1}^{12,020}$ and $\{\hat{\alpha}_t\}_{t=1}^{36}$, respectively, $\tilde{\varepsilon}_{i,t} = \rho \tilde{\varepsilon}_{i,t-1} + \tilde{u}_{i,t}$, with $\tilde{u}_{i,t}$ being drawn from the empirical distribution of $\{\hat{\varepsilon}_{i,t}\}_{i=1}^{176,622}$, and σ_{ROA} refers to the sample standard deviation of ROA (0.309). To assess the impact of serial correlation on our results, we consider $\rho \in \{-1.1, -1, -0.5, 0, 0.5, 1, 1.1\}$.

The observed outcomes are given by $Y_{i,t} = \sum_{g \in \{5,8,11\}} Y_{i,t}(g) \times 1\{G_i = g\}$. As we consider a setup where all units are eventually treated, we are unable to identify average treatment effects in period t = 11. Therefore, we effectively drop the data from the last period (t = 11), and cohort G = 11 acts as "never-treated" units. In our setup, we have that ATT(5,t) = 0.154(t-4) for t = 5..., 10, and ATT(8,t) = 0.093(t-7) for t = 8,9,10. As a consequence, we have that for e = 0,1,...,5, ES(e) is equal to 0.123, 0.247, 0.370, 0.617, 0.772, and 0.926, respectively, and $ES_{avg} = 0.509$.

Table 4 presents the bias and relative RMSE of the four different estimators for ES_{avg} . As expected, all estimators are unbiased. However, in terms of RMSE, there is a good amount of variation across estimators depending on the residual serial correlation. In all scenarios, our proposed efficient DiD estimator yields the lowest RMSE, which supports our theoretical results. When the errors are spherical ($\rho = 0$), our efficient estimator performs very similarly to BJS-G-W. This is aligned with the efficiency results derived under these additional assumptions on the error term in Borusyak et al. (2024) and Wooldridge (2021). Outside this specific setup, though, the BJS-G-W estimator has no efficiency guarantees. Outside of our efficient estimators, our simulations also suggest that it is generally not possible to rank other estimators in terms of efficiency, as their performance depends on the specific DGP.

Table 5 presents the empirical coverage of 95% confidence intervals for ES_{avg} . All estimators display correct coverage across all considered DGPs. In terms of inferential precision measured by

		Relativ	e RMSE		Bias $(\times 10)$						
	EDiD	BJS-G-W	CS-SA	CS-dCDH	EDiD	BJS-G-W	CS-SA	CS-dCDH			
$\rho = 0$	1	1.01	1.61	1.55	0.00	0.00	0.00	0.00			
$\rho = 0.5$	1	1.04	1.26	1.22	-0.01	-0.01	0.00	0.00			
$\rho = 1$	1	1.18	1.08	1.04	0.00	0.01	0.01	0.01			
$\rho = 1.1$	1	1.47	1.23	1.18	-0.03	-0.02	-0.02	-0.02			
$\rho = -0.5$	1	1.04	2.31	2.20	0.00	0.00	0.00	0.00			
$\rho = -1$	1	1.58	3.22	3.37	0.00	0.00	0.00	0.00			
$\rho = -1.1$	1	2.23	3.37	3.77	0.00	0.00	0.00	-0.01			

Table 4: Simulation results for Compustat data: Relative RMSE and Bias

Notes: Simulation results for Compustat data based on 1,000 Monte Carlo experiments. The DGP is summarized in Section 5. In all setups, we randomly draw 400 firms and follow them for 11 periods. The baseline case has error terms that are randomly drawn from the residual distribution as described in Section 5. In subsequent settings, we consider error terms with different degrees of serial correlation. The parameter of interest is ES_{avg} . Relative RMSE is reported using our efficient DiD estimator as a benchmark. All results for bias are multiplied by 10 for readability.

Table 5: Simulation results for Compustat data: Coverage and relative length of confidence interval

		Cov	erage		Rel Length of CI						
	EDiD	BJS-G-W	CS-SA	CS-dCDH	EDiD	BJS-G-W	CS-SA	CS-dCDH			
$\rho = 0$	0.93	0.93	0.94	0.94	1	1.01	1.62	1.55			
$\rho = 0.5$	0.95	0.95	0.95	0.95	1	1.05	1.27	1.23			
$\rho = 1$	0.94	0.96	0.95	0.95	1	1.24	1.11	1.07			
$\rho = 1.1$	0.94	0.96	0.95	0.95	1	1.54	1.27	1.22			
$\rho = -0.5$	0.93	0.93	0.94	0.95	1	1.04	2.35	2.25			
$\rho = -1$	0.93	0.94	0.95	0.95	1	1.64	3.33	3.48			
$\rho = -1.1$	0.95	0.95	0.94	0.94	1	2.22	3.38	3.78			

Notes: Simulation results for Compustat data based on 1,000 Monte Carlo experiments. The DGP is summarized in Section 5. In all setups, we randomly draw 400 firms and follow them for 11 periods. The baseline case has error terms that are randomly drawn from the residual distribution as described in Section 5. In subsequent settings, we consider error terms with different degrees of serial correlation. The parameter of interest is ES_{avg} . Coverage is computed using empirical coverage associated with 95% confidence intervals based on asymptotic approximations (analytical standard errors). Relative length of confidence intervals are reported using our efficient DiD estimator as benchmark.

the length of the 95% confidence intervals, the results once again highlight the practical gains of using our efficient DiD estimator. The relative size of these gains varies with the degree of serial correlation in the data, as expected and consistent with our theoretical results.

As we discussed before, an appealing feature of our proposal is that we can visualize the contribution of the treatment group, the effective comparison group, and the pre-treatment period used as baseline for the construction of our efficient DiD estimator for ES_{avg} . Figure 3 plots this in a heatmap-style for three different DGPs, using a single simulation draw. Overall, we can see that the treatment group G = 5 "contributes more" to the \widehat{ES}_{avg} , which is again expected, as this group

Figure 3: Contribution of treatment and comparison groups, and pre-treatment periods for the Efficient DiD estimator for ES_{avg}



Notes: The figure displays the contribution of each pre-treatment period, comparison group, and treatment group when constructing the efficient DiD estimator for ES_{avg} based on a single draw of the data-generating process. Sample size n = 400 and T = 11, and staggered treatment adoption. The weight color scale varies across DGPs.

has six available post-treatment periods; group G = 8 has only three. Across DGPs, we can also see that the efficient weights vary across comparison groups and pre-treatment periods. When there is no serial correlation, most of the weights come from when G = g' = 5, indicating that the efficient estimator is leveraging the "never-treated" units as a leading comparison group. In this case, we can also see that all pre-treatment periods have similar weighting schemes, again highlighting the appeal of exploring various pre-treatment periods as a baseline. When the error term follows a unit root process, we observe that the efficient estimator allocates almost all the weights to the last available pre-treatment periods, and earlier pre-treatment periods are less important for efficiency considerations. When the errors are negatively serially correlated with $\rho = -1$, we see that the weights are differently spread than in the previous case. Overall, these heatmap plots illustrate that the optimal way to aggregate information across treatment groups, comparison groups, and pre-treatment periods varies from DGP to DGP, and being able to (asymptotically) achieve the semiparametric efficiency bound without needing to know these additional features of the data can be very attractive. Similarly to the setup with a single treatment date, negative efficiency weights provide no reason to be concerned about bias in our context. These efficiency weights appear as a consequence of the semiparametric efficiency results, and leveraging their structure can lead to practical gains in precision.

6 Empirical Illustration

We illustrate how our efficient DiD and ES estimators can be used by estimating the effect of hospitalization on out-of-pocket medical spending using publicly available survey data from the Health and Retirement Study (HRS) from the replication package of Dobkin et al. (2018). Like Dobkin et al. (2018), we explore the variation in the timing of hospitalization observed in the HRS to estimate the different treatment effects of interest. Dobkin et al. (2018) analyzes many other outcomes of interest and also leverage (not publicly available) hospitalization data linked to credit reports.

To conduct our illustrative analysis, we follow the same sample selection steps used by Sun and Abraham (2021), who also revisited Dobkin et al. (2018). The sample selection closely follows Dobkin et al. (2018), and restricts attention to non-pregnancy-related hospital admissions and to adults who are hospitalized at ages 50-59. Like Sun and Abraham (2021), we restrict our analysis to a subsample of individuals who appear throughout waves 7-11 (approximately spanning 2004-2012), so we maintain a balanced panel with 652 individuals in 5 waves, t = 7, ..., 11. Given the data construction, all individuals in the sample are hospitalized by t = 11, but were not hospitalized in wave t = 7. Thus, as we do not have a valid comparison group for wave t = 11, we also drop data from that period.

Overall, we have 4 treatment groups: individuals (first) hospitalized in t = 8 ($G_i = 8$, 252 individuals), t = 9, ($G_i = 9$, 176 individuals), t = 10 ($G_i = 10$, 163 individuals), and those hospitalized in t = 11 ($G_i = \infty$, 65 individuals). Note that as we are only using data up to wave 10, we can re-label the units treated in wave t = 11 as the "never-treated" group, so we match our notation.

As discussed in detail in Sun and Abraham (2021), (unconditional) parallel trends (in all periods and groups) and no-anticipation assumptions are plausible in this setting, and one should a priori expect that the effect of hospitalization is heterogeneous across individuals and periods. Figure 4 plots the evolution of the average out-of-pocket medical spending across all treatment groups, and also provides suggestive evidence in favor of the plausibility of parallel trends assumptions across all periods and groups. Thus, the "heterogeneous robust" DiD and ES estimators for staggered treatment adoption proposed by de Chaisemartin and D'Haultfœuille (2020), Callaway and Sant'Anna (2021), Sun and Abraham (2021), Gardner (2021), Wooldridge (2021), and Borusyak et al. (2024) are wellmotivated. When parallel trends is plausible across all periods and groups, we should expect all modern DiD estimators to provide similar point estimates.



Figure 4: Evolution of average out-of-pocket medical spending across different treatment groups

Notes: Plot of evolutions of average out-of-pocket medical spending across four cohorts defined by the time they were first hospitalized. The waves of the first hospitalizations are displayed in dashed vertical lines. For each cohort, points before the vertical lines are pre-treatment periods, and after the dashed vertical lines are post-treatment periods. We use the same balanced panel data as Sun and Abraham (2021).

In this context, we are interested in different treatment effect parameters. As we discussed before, a natural starting point is to analyze the six different post-treatment ATT(g,t)'s that measure the evolution of ATTs since hospitalization for each group. For instance, ATT(8,8) would give the average treatment effect of hospitalization on out-of-pocket medical spending in time t = 8, among individuals in group g = 8, i.e., those who were hospitalized in time 8; ATT(8,9) and ATT(8,10)would provide the average treatment effect for these same individuals but in periods t = 9 (one year after hospitalization) and t = 10 (two years after hospitalization). The interpretation of the other ATT(g,t)'s is analogous. We may also want to summarize all these different ATT(g,t)'s into more aggregated summary measures to ease interpretations and potentially gain precision. A natural way to summarize the effects is to look at the event-study parameters ES(e) as in (2.2)—which provides an average treatment effect by length of treatment exposure e—and their averages, ES_{avg} , as in (2.3).

Table 6 reports the estimates of all these target parameters and displays their standard errors clustered at individual level in parenthesis. We report estimates using different methods: our efficient DiD procedure (EDiD), the DiD estimator of Callaway and Sant'Anna (2021) and Sun and Abraham (2021) that uses the never-treated units as the comparison group (CS-SA), the one of Callaway and Sant'Anna (2021) and de Chaisemartin and D'Haultfœuille (2020) that uses the not-yet-treated units as the comparison group (CS-dCDH), and the imputation DiD estimator of Borusyak et al. (2024), Gardner (2021), and Wooldridge (2021) (BJS-G-W). Several overall patterns arise. First, note that the point estimates for each target parameter are similar across all the different DiD estimators. This is expected, as parallel trends across all periods and groups and no-anticipation are plausible in this application. We also note that instantaneous treatment effects tend to be larger than later effects, across all cohorts; for instance, estimates of ATT(8, 8) tend to be larger than ATT(8, 9)and ATT(8, 10), and ATT(9, 9) larger than ATT(9, 10). This is well captured in the event-study

Table 6: Estimates for the effect of hospitalization on out-of-pocket medical spending

Estimator	ATT(8,8)	ATT(8,9)	ATT(8, 10)	ATT(9,9)	ATT(9, 10)	ATT(10, 10)	ES(0)	ES(1)	ES(2)	ES_{avg}
EDiD	$3072 \\ (806)$	$1112 \\ (637)$	$1038 \\ (817)$	$3063 \\ (690)$	$90 \\ (641)$	2908 (894)	3024 (486)	$692 \\ (471)$	$1038 \\ (816)$	$1585 \\ (521)$
CS-SA	$2826 \\ (1035)$	825 (909)	$800 \\ (1008)$	3031 (702)	107 (651)	$3092 \\ (995)$	$2960 \\ (539)$	$530 \\ (585)$	$800 \\ (1008)$	$1430 \\ (647)$
CS-dCDH	$3029 \\ (913)$	$1248 \\ (861)$	$800 \\ (1008)$	$3324 \\ (959)$	$107 \\ (651)$	$3092 \\ (995)$	$3134 \\ (536)$	$779 \\ (570)$	$800 \\ (1008)$	$1571 \\ (566)$
BJS-G-W	$3029 \\ (916)$	1285 (767)	$ \begin{array}{c} 1021 \\ (851) \end{array} $	3239 (862)	77 (729)	2758 (957)	$3017 \\ (555)$	788 (587)	$ \begin{array}{l} 1021 \\ (851) \end{array} $	1609 (582)

Notes: This table reports four different sets of estimates for the dynamic effects of hospitalization on out-of-pocket medical spending. Each column is a different target parameter related to ATT(g, t), ES(e) or ES_{avg} . Standard errors clustered at the individual level are reported in parentheses. The sample includes observations from wave t = 7 to t = 10, for 652 individuals. We leverage data from Dobkin et al. (2018) and follow the same sample construction steps as Sun and Abraham (2021).

estimates, ES(0), ES(1), and ES(2), indicating that the effect of hospitalization on out-of-pocket medical spending is mostly transitory. This is in line with the findings in Dobkin et al. (2018).

Another overall pattern that emerges from the results in Table 6 relates to the inference precision of these estimators, as measured by their standard deviations. In line with our theoretical results, our efficient DiD estimators have the smallest standard errors across all target parameters. Outside our efficient DiD procedure, though, no other estimator uniformly dominates the remaining ones in terms of length of confidence intervals. In fact, CS-SA, CS-dCDH, and BJS-G-W take turns into leading to the second most precise estimator across target parameters, with CS-SA being "second best" for ATT(9,9) and ATT(9,10), CS-dCDH coming second for ATT(8,8), ATT(9,10), ES(0), ES(1), and ES_{avg} , and BJS-G-W for ATT(8,9), ATT(8,10), ATT(10,10), and ES(2). This highlights that relying on additional assumptions, such as homoskedasticity and restrictions on serial correlation, to justify efficiency gains may be practically hard. Our proposed efficient DiD estimator does not rely on these conditions and, in fact, delivers more precise inference procedures.

Table / Relative efferency of estimators for the effect of hegoitalizat		
TALLE I DETAILVE EURIEUR VILLESLUHALUNS DI LUE EUELL DE HOSDILAUZAL	on on medical	spending
Table 1. Relative enterency of estimators for the enter of hospitalizat	.on on moulou	sponding

Estimator	ATT(8,8)	ATT(8,9)	ATT(8, 10)	ATT(9,9)	ATT(9, 10)	ATT(10, 10)	ES(0)	ES(1)	ES(2)	ES_{avg}
EDiD	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
CS-SA	1.65	2.04	1.52	1.04	1.03	1.24	1.23	1.54	1.52	1.54
CS-dCDH	1.28	1.83	1.52	1.93	1.03	1.24	1.21	1.46	1.52	1.18
BJS-G-W	1.29	1.45	1.09	1.56	1.29	1.15	1.30	1.55	1.09	1.25

Notes: This table reports the estimated asymptotic relative efficiency (ARE) of four different sets of estimates for the dynamic effects of hospitalization on out-of-pocket medical spending, with our efficient DiD estimator as the benchmark. Each column provides the ARE for a different target parameter related to ATT(g,t), ES(e) or ES_{avg} . The ARE is computed as the square of the ratio of standard errors of the DiD estimator relative to the standard error of our efficient DiD estimator from Table 6.

But one may wonder: do these efficiency gains matter in practice? To answer this practically relevant question, we display in Table 7 estimates of the asymptotic relative efficiency (ARE) of our

proposed efficient DiD estimator with respect to the other available DiD estimators.²⁰ Heuristically speaking, the ARE provides a relative measure of sample size needed for other DiD estimators to achieve the same precision as our efficient DiD estimator. The results in Table 7 highlight that the gains in ARE of our efficient DiD estimator tend to be large. For example, for ATT(8,8), CS-SA, CS-dCDH, and BJS-G-W DiD estimators would respectively need a sample size with 65%, 28%, and 29% more hospitalized individuals than our EDiD estimator to achieve the same precision as EDiD. For ATT(8,9), the gains are even larger, with CS-SA, CS-dCDH, and BJS-G-W DiD requiring 104%, 83%, and 45% more hospitalized individuals than our EDiD estimator to equalize precision. These ARE gains remain large for the event-study coefficients, and their average.

Figure 5: Contribution of comparison groups and pre-treatment periods for the Efficient DiD estimator for ATT(g, t)'s



Notes: Weights attached to different ATT(g, t) estimators that leverage different comparison groups and baseline periods in our application. The weights sum up to one in each row and arise as a consequence of our semiparametric efficient results in Theorem 3.2.

To gain further insights into how our efficient DiD and ES estimators achieve efficiency gains, in Figure 5 we report the weights attached to the different (effective) comparison groups and pretreatment periods for each ATT(g,t) of interest. Overall, there are four different (non-collinear) $\hat{Y}_{g',t_{\text{pre}}}^{\text{att}(g,t)}$ as defined in (4.4) that one can use to estimate each ATT(g,t), and the weights attached to them sum up to one, i.e., the weights in Figure 5 sum up to one in each row. If we were to weigh each of these uniformly, each of them would have a 0.25 weight. In practice, we can see that our efficient DiD estimators never use these uniform weights, and, in fact, these weights vary according

²⁰For any parameter η of a distribution F, and for estimators $\hat{\eta}_1$ and $\hat{\eta}_2$ approximately $N(\eta, V_1/n)$ and $N(\eta, V_2/n)$, respectively, the asymptotic relative efficiency of $\hat{\eta}_2$ with respect to $\hat{\eta}_1$ is given by V_1/V_2 ; see, e.g., Section 8.2 in van der Vaart (1998).

to the parameter of interest. This, by itself, highlights the appeal of our procedure to "adapt" to these different scenarios with the researcher needing to take a stand on additional hard-to-justify assumptions.

It is also interesting to analyze these weights separately. We start with ATT(8,8). For this parameter, the efficient estimator puts weights of 0.29 on the DiD estimator that uses only the never-treated (in our case, cohort treated in wave 11) as a comparison group (G = g' = 5), 0.54 on the DiD estimator that only uses cohort G = 9 as a comparison group (g' = 9, t' = 8), and 0.03 on the DiD estimator that uses cohort G = 10 as the comparison group. Both of these use period t = 7 as the baseline. Interestingly, the efficient DiD estimator for ATT(8,8) puts 0.13 weight on the DiD estimator that uses "never-treated" and G = 10 cohort as comparison group and leverages the first period as pre-treatment periods as well as $t_{\rm pre} = 9$, illustrating the "bridging" phenomenon discussed in the stylized example in Figure 1. Here, note that $t_{\rm pre} = 9$ is a post-period for G = 5 but a pre-treatment period for group G = 10. The flexibility of our efficient DiD estimator to leverage this additional information from the data leads to the efficiency gains highlighted in Table 7.

For cohort G = 8 in period t = 9, we note that the efficient DiD estimator weights the DiD estimator based on never-treated units (i.e., the CS-SA estimator) by 0.30, the DiD estimator that leverages never-treated and G = 9 as comparison groups by 0.49 (0.48), and the DiD estimator that leverages only G = 10 as a comparison group by 0.25. An interpretation for the weights for ATT(8, 10) is analogous. For ATT(9, 9) and ATT(9, 10), note that the efficient DiD estimator puts nearly 0.90 weight on the CS-SA DiD estimator, explaining why its performance is similar to theirs (but better than the other alternative DiD estimators). For ATT(10, 10), our efficient estimator put 0.58 weight on the CS-SA estimator, 0.04, and 0.19 weight on the DiD estimator that uses never-treated as comparison and period eight and seven as baseline period, respectively, and the remaining 0.20 weight on the DiD estimator that combines never-treated and the cohort treated in wave 9 as comparison groups and use first two periods as baseline periods. Together, these flexible weighting schemes demonstrate that our efficient DiD estimators are flexible and capable of extracting information from the data to provide the gains in precision that our theoretical results indicate.

We conclude the session by discussing how one can use our results to assess the stability of our results. In Figure 6, we plot our efficient ES estimators with their 95% confidence intervals in blue, and in gray, we plot all possible event-study estimators that attach unit weight to the different ATT(g,t) estimators in Figure 5 that uses a particular (g',t_{pre}) pair. For ES(0), that are 64 possible estimators $(4 \times 4 \times 4)$ that leverage different combinations of ATT(8,8), ATT(9,9), and ATT(10,10). For ES(1), we have 16 possible estimators, and for ES(2), we have four possible estimators. As you can see in Figure 6, all these different combinations lead to very stable event-study estimators, indicating the plausibility of Assumption PT-All in our context. Figure 6 can be interpreted as an adaptation of the specification curve analysis (Simonsohn, Simmons and Nelson, 2020) to our DiD context.

Figure 6: Assessing the stability of our event-study estimators



Notes: The plot displays our efficient event study estimators for event time zero, one, and two, and their 95% confidence intervals computed using analytical standard errors in blue. In gray, we display all possible event-study estimators that attach unit weight to the different ATT(g,t) estimators in Figure 5 that uses a particular (g', t_{pre}) pair. For ES(0), that are 64 possible estimators $(4 \times 4 \times 4)$ that leverage different combinations of ATT(8,8), ATT(9,9), and ATT(10,10). For ES(1), we have 16 possible estimators, and for ES(2), we have four possible estimators.

7 Conclusion

Empirical researchers routinely face a variety of alternative DiD and ES estimators to choose from. Existing approaches often overlook the fact that different pre-treatment periods and comparison groups may have varying identification power, leading them to treat all pre-treatment periods as equally informative or to use only the last pre-treatment period as a baseline. The semiparametric efficiency framework developed in this paper provides a principled approach to leveraging information from the pre-treatment period to form DiD and ES estimators with minimum variance under the maintained identification assumptions, while remaining agnostic to parametric functional form or parametric error structure. As our simulations and empirical illustration demonstrate, the gains in precision can be substantial, addressing the concern that heterogeneous robust DiD estimators may not be informative compared to more traditional estimators that assume treatment effect heterogeneity away.

In practice, practitioners may want to assess if there is any statistical evidence against using any pre-treatment period as the baseline in their DiD/ES procedure, or against using any specific available comparison group (when treatment is staggered). We present in Appendix A how to construct Hausman-type tests for overidentification specification test of DiD models, or how to construct alternative estimators that aim to trade off bias and variance (Armstrong, Kline and Sun, 2024). We also discuss visualization schemes to informally quantify the sensitivity of the results concerning the usage of different pre-treatment baseline periods, and different comparison groups, or to "eyeball" the magnitude of potential model violations.

Our framework opens avenues for several extensions. It would be useful to study semiparametric efficient estimation for nonlinear DiD-type setups such as the Changes-in-Changes model (Athey and Imbens, 2006), setups where treatment can turn on and off over multiple periods (de Chaisemartin and D'Haultfœuille, 2020, 2024), and setups with unbalanced panel or repeated cross-sections (Call-

away and Sant'Anna, 2021)—in this latter case, we expect new sources of overidentification related to event-study weights. Finally, it is important to stress that all results derived in this paper leveraged the large-n, fixed-T panel data regime, as parallel trends are more likely to be satisfied when the number of periods is small. In setups where parallel trends are not plausible, we recommend that researchers use alternative estimators that do not rely on parallel trends assumptions, e.g., Arkhangelsky et al. (2021), Viviano and Bradic (2021), and Imbens and Viviano (2023). Developing semiparametric efficient estimators in such frameworks is left for future research.

References

- Ackerberg, Daniel, Xiaohong Chen, Jinyong Hahn, and Zhipeng Liao, "Asymptotic Efficiency of Semiparametric Two-step GMM," The Review of Economic Studies, 2014, 81 (3), 919–943.
- Ai, Chunrong and Xiaohong Chen, "Efficient estimation of models with conditional moment restrictions containin unknown functions," *Econometrica*, 2003, 71 (6), 1795–1843.
- _ and __, "Estimation of possibly misspecified semiparametric conditional moment restriction models with different conditioning variables," *Journal of Econometrics*, 2007, 141 (1), 5–43.
- **and** __, "The semiparametric efficiency bound for models of sequential moment restrictions containing unknown functions," *Journal of Econometrics*, 2012, *170* (2), 442–457.
- Arkhangelsky, Dmitry, Susan Athey, David A. Hirshberg, Guido W. Imbens, and Stefan Wager, "Synthetic Difference-in-Differences," American Economic Review, 2021, 111 (12), 4088–4118.
- Armstrong, Timothy B., Patrick Kline, and Liyang Sun, "Adapting to Misspecification," arXiv:2305.14265v4, 2024.
- Arold, Benjamin W, "Evolution vs. Creationism in the Classroom: The Lasting Effects of Science Education," The Quarterly Journal of Economics, 2024, p. qjae019.
- Athey, Susan and Guido W. Imbens, "Identification and Inference in Nonlinear Difference-in-Differences Models," *Econometrica*, 2006, 74 (2), 431–497.
- Baker, Andrew, Brantly Callaway, Scott Cunningham, Andrew Goodman-Bacon, and Pedro
 H. C. Sant'Anna, "Difference-in-Differences Designs: A Practitioner's Guide," arXiv:2503.13323, 2025.
- __, David F. Larcker, and Charles C. Y. Wang, "How much should we trust staggered difference-indifferences estimates?," Journal of Financial Economics, 2022, 144 (2), 370–395.
- Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan, "How Much Should We Trust Differences-In-Differences Estimates?," *The Quarterly Journal of Economics*, 2004, 119 (1), 249–275.
- Blandhol, Christine, John Bonney, Magne Mogstad, and Alexander Torgovitsky, "When is TSLS Actually LATE?," Working Paper 29709, National Bureau of Economic Research August 2022.
- Borusyak, Kirill, Xavier Jaravel, and Jann Spiess, "Revisiting Event Study Designs: Robust and Efficient Estimation," *Review of Economic Studies*, 2024, *91* (6), 3253–3285.
- Braghieri, Luca, Ro'ee Levy, and Alexey Makarin, "Social Media and Mental Health," American Economic Review, 2022, 112 (11), 3660–93.

- Callaway, Brantly and Pedro H. C. Sant'Anna, "Difference-in-Differences with multiple time periods," Journal of Econometrics, 2021, 225 (2), 200–230.
- Cattaneo, Matias D, Michael Jansson, and Xinwei Ma, "On the effect of bias estimation on coverage accuracy in nonparametric inference," *Journal of the American Statistical Association*, 2018, 113 (522), 767–779.
- Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer, "The Effect of Minimum Wages on Low-Wage Jobs," *The Quarterly Journal of Economics*, 2019, 134 (3), 1405–1454.
- Chamberlain, Gary, "Asymptotic efficiency in estimation with conditional moment restrictions," *Journal* of Econometrics, 1987, 34 (3).
- _, "Efficiency bounds for semiparametric regression," *Econometrica*, 1992, 60 (3), 567–596.
- Chen, Jiafeng, Xiaohong Chen, and Elie Tamer, "Efficient estimation of average derivatives in NPIV models: Simulation comparisons of neural network estimators," *Journal of Econometrics*, 2023, 235 (2), 1848–1875.
- Chen, Xiaohong, "Large Sample Sieve Estimation of Semi-Nonparametric Models," in James J. Heckman and Edward E. Leamer, eds., *Handbook of Econometrics*, Vol. 6B, Amsterdam: Elsevier, 2007, chapter 76, pp. 5549–5632.
- _ and Andres Santos, "Overidentification in Regular Models," *Econometrica*, 2018, 86 (5).
- _ and Zhipeng Liao, "Sieve inference on irregular parameters," Journal of Econometrics, 2014, 182 (1), 70–86.
- Chernozhukov, Victor, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins, "Double/debiased machine learning for treatment and structural parameters," *The Econometrics Journal*, 2018, 21 (1), C1—C68.
- Chiu, Albert, Xingchen Lan, Ziyi Liu, and Yiqing Xu, "What To Do (and Not to Do) with Causal Panel Analysis under Parallel Trends: Lessons from A Large Reanalysis Study," *Working Paper*, 2023.
- de Chaisemartin, Clément and Xavier D'Haultfœuille, "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects," *American Economic Review*, 2020, 110 (9), 2964–2996.
- and ____, "Two-way fixed effects and differences-in-differences with heterogeneous treatment effects: a survey," *Econometrics Journal*, 2023, 26 (3), C1—C30.
- and __, "Difference-in-Differences Estimators of Intertemporal Treatment Effects," The Review of Economics and Statistics, 2024, Forthcoming, 1–45.
- Deshpande, Manasi and Yue Li, "Who Is Screened Out? Application Costs and the Targeting of Disability Programs," American Economic Journal: Economic Policy, 2019, 11 (4), 213–248.
- Dobkin, Carlos, Amy Finkelstein, Raymond Kluender, and Matthew J. Notowidigdo, "The economic consequences of hospital admissions," *American Economic Review*, 2018, 108 (2), 308–352.
- Fadlon, Itzik and Torben Heien Nielsen, "Family Labor Supply Responses to Severe Health Shocks: Evidence from Danish Administrative Records[†]," American Economic Journal: Applied Economics, 2021, 13 (3), 1–30.
- Gardner, John, "Two-stage differences in differences," Working Paper, 2021.

Goldsmith-Pinkham, Paul, "Tracking the Credibility Revolution across Fields," arXiv:2405.20604, 2024.

- Goodman-Bacon, Andrew, "Difference-in-differences with variation in treatment timing," Journal of Econometrics, 2021, 225 (2), 254–277.
- Gormley, Todd A. and David A. Matsa, "Growing Out of Trouble? Corporate Responses to Liability Risk," *Review of Financial Studies*, 2011, 24 (8), 2781–2821.
- Hansen, Casper Worm and Asger Mose Wingender, "National and Global Impacts of Genetically Modified Crops," *American Economic Review: Insights*, June 2023, 5 (2), 224–40.
- Hansen, Lars Peter, "Large Sample Properties of Generalized Method of Moments Estimators," Econometrica, 1982, 50 (4), 1029–1054.
- Harmon, Nikolaj A., "Difference-in-Differences and Efficient Estimation of Treatment Effects," Working Paper, 2023.
- Imbens, Guido W. and Davide Viviano, "Identification and Inference for Synthetic Controls with Confounding," *arXiv:2312.00955*, 2023.
- Jacobson, Louis S., Robert J. LaLonde, and Daniel G. Sullivan, "Earnings losses of displaced workers," American Economic Review, 1993, 83 (4), 685–709.
- Lal, Apoorva, "When can we get away with using the two-way fixed effects regression?," arXiv:2503.05125 [econ.EM], 2025.
- Lee, Soo Jeong and Jeffrey M. Wooldridge, "A Simple Transformation Approach to Difference-in-Differences Estimation for Panel Data," *Working Paper*, 2023. Available at SSRN: http://dx.doi.org/ 10.2139/ssrn.4516518.
- Malani, Anup and Julian Reif, "Interpreting pre-trends as anticipation: Impact on estimated treatment effects from tort reform," *Journal of Public Economics*, 2015, 124, 1–17.
- Marcus, Michelle and Pedro H. C. Sant'Anna, "The role of parallel trends in event study settings: An application to environmental economics," *Journal of the Association of Environmental and Resource Economists*, 2021, 8 (2), 235–275.
- Mast, Evan, "Warding off Development: Local Control, Housing Supply, and NIMBYs," The Review of Economics and Statistics, 2024, 106 (3), 671–680.
- Newey, Whitney K. and Richard J. Smith, "Higher order properties of GMM and generalized empirical likelihood estimators," *Econometrica*, 2004, 72 (1), 219–255.
- Roth, Jonathan, Pedro H. C. Sant'Anna, Alyssa Bilinski, and John Poe, "What's Trending in Difference-in-Differences? A Synthesis of the Recent Econometrics Literature," *Journal of Econometrics*, 2023, 235 (2), 2218–2244.
- Sant'Anna, Pedro H. C. and Jun Zhao, "Doubly robust difference-in-differences estimators," Journal of Econometrics, 2020, 219 (1), 101–122.
- Simonsohn, Uri, Joseph P. Simmons, and Leif D. Nelson, "Specification curve analysis," Nature Human Behaviour, 2020, 4, 1208–1214.
- Staniswalis, Joan G, "The kernel estimate of a regression function in likelihood-based models," Journal of the American Statistical Association, 1989, 84 (405), 276–283.
- Sun, Liyang and Sarah Abraham, "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects," *Journal of Econometrics*, 2021, 225 (2), 175–199.

van der Vaart, Aad W., Asymptotic Statistics, Cambridge: Cambridge University Press, 1998.

- Viviano, Davide and Jelena Bradic, "Dynamic covariate balancing: estimating treatment effects over time," June 2021. arXiv:2103.01280 [econ, math, stat].
- Weiss, Amanda, "How Much Should We Trust Modern Difference-in-differences Estimates?," OSF Preprints, 2024.
- Wooldridge, Jeffrey M, "Cluster-Sample Methods in Applied Econometrics," American Economic Review P&P, 2003, 93 (2), 133–138.
- ____, "Two-Way Fixed Effects, the Two-Way Mundlak Regression, and Difference-in-Differences Estimators," Working Paper, 2021, pp. 1–89.