Dynamic covariate balancing: estimating treatment effects over time

Davide Viviano* Jelena Bradic†

Draft version July, 2022
First version: March, 2021

Abstract

This paper discusses the problem of estimation and inference on the effects of time-varying treatment. We propose a method for inference on the effects treatment *histories*, introducing a dynamic covariate balancing method combined with penalized regression. Our approach allows for (i) treatments to be assigned based on arbitrary past information, with the propensity score being unknown; (ii) outcomes and time-varying covariates to depend on treatment trajectories; (iii) high-dimensional covariates; (iv) heterogeneity of treatment effects. We study the asymptotic properties of the estimator, and we derive the parametric $n^{-1/2}$ convergence rate of the proposed procedure. Simulations and an empirical application illustrate the advantage of the method over state-of-the-art competitors.

Keywords: Causal Inference, Dynamic Treatments, High Dimensions, Treatment Effects, Panel Data.
JEL Code: C10.

*Stanford Graduate School of Business, 655 Knight Way, Stanford, CA 94305. Email: dviviano@stanford.edu.

†Department of Mathematics and Halicioğlu Data Science Institute, University of California San Diego, La Jolla, CA, 92093. Email: jbradic@ucsd.edu.

1 Introduction

This paper studies the estimation and inference of the average effect of treatment trajectories (i.e., treatment history) for observational studies with n independent units observed over T periods. Empirical examples include studying the effect of public health insurance (Finkelstein et al., 2012), negative political advertisements (Blackwell, 2013), or the long or short-run effects of minimum wages on employment.

We consider a setting where time-varying covariates and outcomes depend on past treatment assignments, and treatments are assigned sequentially based on arbitrary past information. Two alternative procedures can be considered for this case. First, researchers may explicitly model how treatment effects propagate over each period through time-varying covariates and intermediate outcomes. This approach is prone to large estimation error and misspecification in high dimensions: it requires modeling outcomes and *each* time-varying covariate as a function of all past covariates, outcomes, and treatment assignments. A second approach is to use inverse-probability weighting estimators for estimation and inference (Tchetgen and Shpitser, 2012; Vansteelandt et al., 2014). However, classical semi-parametric estimators are prone to instability in the estimated propensity score. There are two main reasons. First of all, the propensity score defines the joint probability of the *entire* treatment history and can be close to zero for moderately long treatment histories. Additionally, the propensity score can be misspecified in observational studies.[1]

Figure 1 presents an example. Here, the probability of being under treatment for two consecutive periods in an application from Acemoglu et al. (2019) shifts towards zero, making inverse-probability weighting estimators unstable in small sample.[2]

We overcome the problems above by proposing a parsimonious and easy-to-interpret model for potential outcomes. In the spirit of local projections (Jordà, 2005), we model the *potential* outcome as an (approximately) linear function of pre-

[1] An example for misspecification is when treatments depend on the decisions of forward-looking agents who maximize expected utility (Heckman and Navarro, 2007) with unknown utilities.

[2] Estimation is performed via logistic regression with a pooled regression with year, region fixed effects, and four lagged outcomes. The right panel also controls for the past treatment assignment.

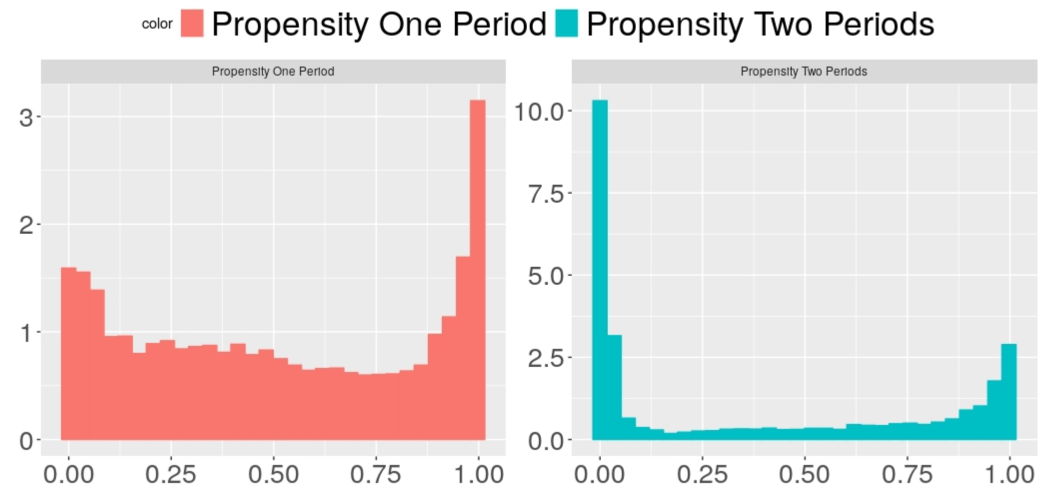

Figure 1: Data from Acemoglu et al. (2019). Estimated probability of treatment for one year (left-panel) and two consecutive years (right-panel).

vious potential outcomes and (high-dimensional) covariates.[3] Unlike the standard local projection framework, the model on potential outcomes allows researchers to be agnostic on the treatment assignment mechanism. In particular, assignments can depend on some unknown functions of arbitrary past information[4], and treatment effect have dynamics in outcomes and time-varying covariates.

Our method, entitled Dynamic Covariate Balancing (DCB), combines high dimensional estimation with covariate balancing. First, we study the identification of the model's parameters. Identification of treatment effects' paths consists of projecting conditional expectations recursively onto previous periods while controlling for past treatment history. We estimate each conditional expectation via penalized regression. We derive novel dynamic balancing conditions to circumvent the propensity score estimation and guarantee a negligible bias of the high-dimensional coefficients.

Balancing covariates is intuitive and common in practice: in cross-sectional studies, treatment and control units are comparable when the two groups have similar

[3]Potential outcomes and covariates are indexed by the past treatment history.

[4]The outcome model nests linear models with arbitrarily many auto-regressive components without, however, being required to estimate separate regressions for each time-varying covariate. It encompasses marginal structural models of Robins et al. (2000) under approximate linearity.

characteristics (Hainmueller, 2012; Imai and Ratkovic, 2014; Li et al., 2018; Ning et al., 2017). We generalize covariate balancing of Zubizarreta (2015) and Athey et al. (2018) to a dynamic setting. One important insight is to construct weights sequentially in time, where balancing weights in the current period depends on those estimated in the previous period. We show that our balancing procedure (i) allows for estimation and inference without knowledge of the propensity score; (ii) guarantees a vanishing bias of order faster than $n^{-1/2}$ and valid asymptotic inference; and (iii) solves a feasible quadratic program to find the weights with minimal variance and thus ensures robustness to poor overlap in a small sample. While identification and balancing differ from the cross-sectional studies due to the dynamic structure, theoretical derivations differ since they require to study the joint distribution of weighted combinations of residuals from local projections under arbitrary time dependence.

Our numerical studies show the advantage of the proposed method over state-of-the-art competitors. In our empirical application, we study the effect of negative advertisement on the election outcome and the effect of democracy on GDP growth.

The remainder of the paper is organized as follows. In Section 2, we discuss the framework and model in the presence of two periods. In Section 3 we discuss balancing with two periods. In Section 4 we extend to multiple periods and discuss theoretical guarantees. Numerical studies and the empirical application are included in Section 5 and Section 6 respectively. Section 7 concludes.

1.1 Related Literature

Dynamic treatments have been widely discussed in several independent strands of literature. Robins (1986), Robins et al. (2000), Hernán et al. (2001), Boruvka et al. (2018), Blackwell (2013), Bang and Robins (2005) and others discuss estimation and inference on dynamic treatments (for a review, Vansteelandt et al., 2014). Bojinov and Shephard (2019), Bojinov et al. (2020) study inverse-probability weighting estimators and characterize their properties from a design-based perspective. Doubly robust estimators (Robins et al., 1994) for dynamic treatment assignment have been studied by Jiang and Li (2015); Nie et al. (2021); Tchetgen and Shpitser (2012);

Zhang et al. (2013), Bodory et al. (2020). One drawback of the above references is that, while double-robust in low dimensions, in high-dimensions, they are sensitive to the misspecification of the propensity score and poor overlap in small sample.[5]

Our contribution to balancing with dynamic treatments is of independent interest. Differently from Zhou and Wodtke (2018), who extend entropy balancing of Hainmueller (2012), we do not estimate one model for each covariate given the past. Instead, we only estimate models for the end-line potential outcomes, which leads to computationally efficient estimators. DCB explicitly characterizes the high-dimensional model's bias in a dynamic setting to avoid overly conservative moment conditions, while Kallus and Santacatterina (2018) design balancing in the worst-case scenario. Different from Imai and Ratkovic (2015), the number of moment conditions here grows linearly with T and not exponentially. We do not require estimation of the propensity model's score function as in Yiu and Su (2018). Finally, Arkhangelsky and Imbens (2019) propose balancing weights assuming no carry-overs in treatment effects. None of the above references allows for high dimensional covariates.

Our problem more broadly connects to the literature on two-way fixed effects and Difference-in-Differences (Abraham and Sun, 2018; Athey and Imbens, 2022; Callaway and Sant'Anna, 2019; de Chaisemartin and d'Haultfoeuille, 2019; Goodman-Bacon, 2021; Imai and Kim, 2016). The above references mostly focus on staggered adoption, while allowing for time-invariant confounders. However, they prohibit dynamic selection into treatment based on past outcomes and time-varying covariates.[6] Here, we allow for dynamics in treatments assigned based on arbitrary past information and covariates to depend on the past assignments. Also, the above references either require correct specification of the propensity score, assume that there are no

[5]In particular, they require consistency of the propensity score function at an appropriate fast rate of convergence in the spirit of cross-sectional studies in high-dimensions (e.g., Farrell, 2015).

[6]The above references impose restrictions potential outcomes' momentconditionally on future assignments, imposing either strong exogeneity on future assignments or parallel trend assumptions. These often require conditional mean independence of potential outcomes with respect to the future treatment path, such as the one of "always being under control". However, in the presence of dynamic treatment assignments, past potential outcomes are predictive of future assignments (e.g., whether individuals do not receive the treatment may depend dynamically on their past outcomes). See also Ghanem et al. (2022) for a discussion.

high-dimensional covariates or both. Related methods also include discrete choice models and dynamic treatments with instruments (Heckman et al., 2016; Heckman and Navarro, 2007), which impose parametrizations on the propensity score.

Similarly, the literature on Synthetic Control (SC) methods (Abadie et al., 2010; Arkhangelsky et al., 2021; Doudchenko and Imbens, 2016) assumes staggered adoption with an exogenous treatment time, hence prohibiting dynamics in treatment assignments. Specifically, Ben-Michael et al. (2018, 2019) balance covariates as in Zubizarreta (2015), conditional on the treatment time. In their setting, staggered adoption motivates the construction of the same balancing weights for all post-treatment periods without allowing for dynamics in treatment assignments. Here, treatment assignments are time-varying and endogenously assigned based on past information.

In a few studies regarding high-dimensional panel data, researchers instead require correct specification of the propensity score (Belloni et al., 2016; Bodory et al., 2020; Chernozhukov et al., 2017, 2018; Lewis and Syrgkanis, 2020; Shi et al., 2018; Zhu, 2017), or impose homogeneous treatment effects (Kock and Tang, 2015; Krampe et al., 2020), differently from our study. An overview for panel data in econometrics can be found in Arellano and Bonhomme (2011), Abadie and Cattaneo (2018).

Differently from the time-series literature (Plagborg-Møller, 2019; Stock and Watson, 2018; White and Lu, 2010), this paper uses information from panel data and allows for arbitrary dependence of outcomes, covariates, and treatment assignments over time. Angrist and Kuersteiner (2011), and Angrist et al. (2018) study inference using inverse probability weights estimator without incorporating carryover effects in the weights.[7] This difference reflects different target estimands. In the context of local projections, Rambachan and Shephard (2019) discuss their causal interpretability for assignments independent of the past. Here, we derive identification results with serially correlated treatment assignments that depend on the past outcomes.

[7]Weights do not depend on the joint probability of treatment history.

2 Dynamics and potential projections

We first discuss the case of two time periods since it provides a simple illustration of the problem and our solution. Our focus is on ex-post evaluation, where treatment effects are evaluated after the entire history of interventions has been deployed.

In the presence of two periods only, we observe n i.i.d. copies of a random vector

$$\Big(X_{i,1}, D_{i,1}, Y_{i,1}, X_{i,2}, D_{i,2}, Y_{i,2}\Big) \sim_{i.i.d.} \mathcal{P}$$

where D_1 and D_2 are binary treatment assignments at time $t = 1, t = 2$, respectively. Here, $X_{i,1}$ and $X_{i,2}$ are covariates for unit i observed at time $t = 1$ and $t = 2$, respectively. We observe the outcome $Y_{i,t}$ right after $D_{i,t}$, but prior to $D_{i,t+1}$. That is, at time $t = 1$, we observe $\{X_{i,1}, D_{i,1}\}$. Outcome $Y_{i,1}$ is revealed after time $t = 1$ but before time $t = 2$. At time $t = 2$ we observe $\{X_{i,2}, D_{i,2}\}$ and finally, outcome $Y_{i,2}$ is revealed. An illustration is in Figure 2. Whenever we omit the index i, we refer to the vector of observations for all units.

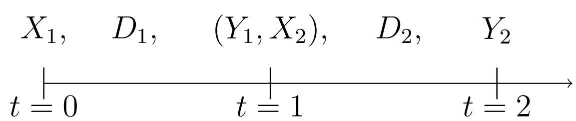

Figure 2: Illustration in two periods. At the baseline $t = 0$, we draw covariates X_1 first, and then D_1. At $t = 1$, we draw the outcome ad covariates, and then the assignment D_2. In the last period we draw the end-line outcome.

2.1 Potential outcomes and estimands

Potential outcomes and covariates are functions of the entire treatment history. Here,

$$\Big(Y_{i,2}(1,1), Y_{i,2}(1,0), Y_{i,2}(0,1), Y_{i,2}(0,0)\Big)$$

define the potential outcomes if individual i is under treatment for two consecutive periods, under treatment for the first but not the second period, the second but not the first, and none of the periods. We define compactly $Y_{i,1}(d_1, d_2)$ and $Y_{i,2}(d_1, d_2)$ the potential outcomes in period one and two, respectively, for unit i, under a treatment history that assigns treatment d_1 in the first period and d_2 in the second period. Throughout our discussion, we implicitly assume that SUTVA holds (Rubin, 1990).

Treatment histories may also affect future covariates. Therefore, we denote $X_{i,2}(d_1, d_2)$, the *potential covariates* for a treatment history (d_1, d_2), while $X_{i,1}$ denotes baseline covariates. The causal effect of interest is the long-run impact of two treatment histories (d_1, d_2),(d_1', d_2') on the potential outcomes conditional on baseline covariates. Let

$$\mu_2(d_1, d_2) = \frac{1}{n} \sum_{i=1}^{n} \mathbb{E}\Big[Y_{i,2}(d_1, d_2)\Big| X_{i,1}\Big]$$

denote the expectation of potential outcomes taken with respect to the conditional distribution of $Y_{i,2}(d_1, d_2)$ given $X_{i,1}$.[8] We construct

$$\text{ATE}(d_{1:2}, d_{1:2}') = \mu_2(d_1, d_2) - \mu_2(d_1', d_2'), \tag{1}$$

where $d_{1:2} = (d_1, d_2)$. A simple example is $\text{ATE}(\mathbf{1}, \mathbf{0})$, which denotes the effect of a policy when implemented on two consecutive periods against the effect of the policy when never implemented (Athey and Imbens, 2022).

The first condition we impose is the *no-anticipation*. This is defined below.

Assumption 1 (No Anticipation)**.** For $d_1 \in \{0, 1\}$, let (i) $Y_{i,1}(d_1, 1) = Y_{i,1}(d_1, 0)$, and (ii) $X_{i,2}(d_1, 1) = X_{i,2}(d_1, 0)$.

The no anticipation condition has two implications: (i) potential outcomes only depend on past but not future treatments; (ii) the treatment status at $t = 2$ has no contemporaneous effect on covariates. Observe that the no-anticipation allows for anticipatory effects governed by *expectation*, but it prohibits anticipatory effects

[8] Formally, $(Y_{i,2}(d_1, d_2), X_{i,1}) \sim \mathcal{D}(d_1, d_2)$.

based on the future treatment *realization*.[9] Also, observe that the no-anticipation is not imposed on the realized treatments, and it allows potential outcomes to be correlated with the future assignments (e.g., see Equation 2).

Example 2.1 (Observed outcomes)**.** Consider a dynamic model of the form (omitting time-varying covariates at time $t = 2$ for expositional convenience)

$$Y_{i,2} = g_2\Big(Y_{i,1}, X_{i,1}, D_{i,1}, D_{i,2}, \varepsilon_{i,2}\Big), \quad Y_{i,1} = g_1\Big(X_{i,1}, D_{i,1}, \varepsilon_{i,1}\Big),$$

for some arbitrary functions $g_1(\cdot), g_2(\cdot)$ and unobservables $(\varepsilon_{i,2}, \varepsilon_{i,1})$ which we assume are exogenous. Then we can write

$$Y_{i,2}(d_1, d_2) = g_2\Big(Y_{i,1}(d_1), X_{i,1}, d_1, d_2, \varepsilon_{i,2}\Big), \quad Y_{i,1}(d_1) = g_1\Big(X_{i,1}, d_1, \varepsilon_{i,1}\Big).$$

Since $g_1(\cdot)$ is not a function of d_2, Assumption 1 holds, for any (conditional) distribution of $(D_{i,1}, D_{i,2})$. □

With abuse of notation, in the rest of our discussion, we index potential outcomes and covariates by past treatment history only, letting Assumption 1 implicitly hold. We define $H_{i,2} = \Big[D_{i,1}, X_{i,1}, X_{i,2}, Y_{i,1}\Big]$, as the vector of past treatment assignments, covariates, and outcomes in the previous period. We refer to

$$H_{i,2}(d_1) = \Big[d_1, X_{i,1}, X_{i,2}(d_1), Y_{i,1}(d_1)\Big]$$

as the "potential history" under treatment status d_1 in the first period. In principle, $H_{i,2}$ can also contains interaction terms, omitted for the sake of brevity.

The second condition we impose is the sequential ignorability condition.

Assumption 2 (Sequential Ignorability)**.** Assume that for all $(d_1, d_2) \in \{0, 1\}^2$,

$$(A) \quad Y_{i,2}(d_1, d_2) \perp D_{i,2} \Big| D_{i,1}, X_{i,1}, X_{i,2}, Y_{i,1}$$
$$(B) \quad \Big(Y_{i,2}(d_1, d_2), H_{i,2}(d_1)\Big) \perp D_{i,1} \Big| X_{i,1},$$

[9] The no-anticipation assumption on potential outcomes has been previously discussed also in Bojinov and Shephard (2019), Imai et al. (2018), Boruvka et al. (2018) to cite some.

The Sequential Ignorability (Robins et al., 2000) is common in the literature on dynamic treatments. It states that treatment in the first period is randomized on baseline covariates only, while the treatment in the second period is randomized with respect to the observable characteristics in time $t = 2$.

Example 2.1 Cont'd We can equivalently write Assumption 2 as

$$D_{i,2} = f_2\Big(D_{i,1}, X_{i,1}, X_{i,2}, Y_{i,1}, \varepsilon_{D_{i,2}}\Big), \quad D_{i,1} = f_1\Big(X_{i,1}, \varepsilon_{D_{i,1}}\Big), \tag{2}$$

for some arbitrary (unknown) functions f_1, f_2, where the unobservables satisfy

$$\varepsilon_{D_{i,2}} \perp \varepsilon_{i,2}\Big|D_{1,i}, X_{i,1}, X_{i,2}, Y_{i,1}, \quad \varepsilon_{D_{i,1}} \perp (\varepsilon_{i,1}, \varepsilon_{i,2})\Big|X_{i,1}.$$

□

2.2 Potential projections

Next, we discuss the model for potential outcomes. Given baseline covariates $X_{i,1}$, for a treatment history (d_1, d_2), we denote

$$\theta_1(x_1, d_1, d_2) = \mathbb{E}\Big[Y_{i,2}(d_1, d_2)\Big|X_{i,1} = x_1\Big],$$
$$\theta_2(x_1, x_2, y_1, d_1, d_2) = \mathbb{E}\Big[Y_{i,2}(d_1, d_2)\Big|X_{i,1} = x_1, X_{i,2} = x_2, Y_{i,1} = y_1, D_{i,1} = d_1\Big],$$

respectively the conditional expectation of the potential outcome at the end-line period, given history at time $t = 1$ (base-line) and given the history at time $t = 2$.

In the same spirit of Jordà (2005) we model $\theta_1(\cdot), \theta_2(\cdot)$ linearly.

Assumption 3 (Model)**.** We assume that for some $\beta^{(1)}_{d_1,d_2} \in \mathbb{R}^{p_1}, \beta^{(2)}_{d_1,d_2} \in \mathbb{R}^{p_2}$

$$\theta_1(x_1, d_1, d_2) = x_1\beta^{(1)}_{d_1,d_2}, \quad \theta_2\Big(x_1, x_2, y_1, d_1, d_2\Big) = \Big[d_1, x_1, x_2, y_1\Big]\beta^{(2)}_{d_1,d_2}.$$

The above models can be seen as a *local projection* model on potential outcomes, with the end-line potential outcome depending linearly on information up to and from each period. An important feature of the proposed model is that we impose it directly on potential outcomes without imposing a model on treatment assignments. The coefficients $\beta^{(1)}_{d_1,d_2}, \beta^{(2)}_{d_1,d_2}$ are different and indexed by the treatment history, capturing the effects of (d_1, d_2) and heterogeneity (note that X_1 also contains an intercept).

Example 2.2 (Linear Model)**.** Let $X_{i,1}, X_{i,2}$ also contain an intercept. Consider the following set of conditional expectations

$$\mathbb{E}\Big[Y_{i,1}(d_1)\Big|X_{i,1}\Big] = X_{i,1}\alpha_{d_1}, \quad \mathbb{E}\Big[X_{i,2}(d_1)\Big|X_{i,1}\Big] = W_{d_1}X_{i,1}$$
$$\mathbb{E}\Big[Y_{i,2}(d_1, d_1)\Big|X_{i,1}, X_{i,2}, Y_{i,1}, D_{i,1} = d_1\Big] = \Big(X_{i,1}, X_{i,2}(d_1), Y_{i,1}(d_1)\Big)\beta^{(2)}_{d_1,d_2},$$

for some arbitrary parameters $\alpha_{d_1} \in \mathbb{R}^{p_1}$ and $\beta^{(2)}_{d_1,d_2} \in \mathbb{R}^{p_2}$. In the above display, W_{d_1}, V_{d_1} denote unknown matrices in $\mathbb{R}^{p_2 \times p_1}$. The model satisfies Assumption 3. □

Example 2.2 shows that the linearity condition imposed in Assumption 3 holds *exactly* whenever the potential outcomes follow a linear model and dependence between covariates is explained via an autoregressive structure.

Remark 1 (Linearity in high-dimensions as an approximation to the true model)**.** We observe that our results extend to the case where we relax Assumption 3 and assume only approximate linearity up to an order $\mathcal{O}(r_p)$, where r_p is an arbitrary sequence which depends on p with $r_p = o(n^{-1/2})$. Such a setting embeds empirical applications where many covariates (and their transformation) can approximate the conditional mean function as linear (see Belloni et al., 2014, for related discussions). We do not require valid model approximation on the propensity score. This is in contrast with high-dimensional settings which require valid model approximations for *both* the propensity score and conditional mean function (e.g., Farrell, 2015). □

As noted in Example 2.2, the local projection model has an important advantage (especially in high-dimensions): while valid under linearity of covariates and outcomes, it does not require specifying (and estimate) a structural model for each time-varying-covariate, which is cumbersome in high dimensions and prone to significant estimation error. Instead, the local projection model is parsimonious in the number of parameters. This motivates its large use in applications, dating back to Jordà (2005). Here, we revisit the model within a causal framework.

We conclude this discussion with the following identification result.

Lemma 2.1 (Identification of the potential outcome model)**.** *Let Assumption 1, 2, 3 hold. Then*[10]

$$\mathbb{E}\Big[Y_{i,2}\Big|H_{i,2}, D_{i,2} = d_2, D_{i,1} = d_1\Big] = \mathbb{E}\Big[Y_{i,2}(d_1, d_2)\Big|H_{i,2}, D_{i,1} = d_1\Big] = H_{i,2}(d_1)\beta^{(2)}_{d_1,d_2}$$
$$\mathbb{E}\Big[\mathbb{E}\Big[Y_{i,2}\Big|H_{i,2}, D_{i,2} = d_2, D_{i,1} = d_1\Big]\Big|X_{i,1}, D_{i,1} = d_1\Big] = \mathbb{E}\Big[Y_{i,2}(d_1, d_2)\Big|X_{i,1}\Big] = X_{i,1}\beta^{(1)}_{d_1,d_2}.$$

The proof is in the Appendix. Lemma 2.1 is new in the context of local projections. It connects to the literature in biostatistic on longitudinal data with marginal structural models (Bang and Robins, 2005; Robins et al., 2000). We identify coefficients that capture causal effects of treatment histories using information from recursive projections. As a result, we can first regress the observed outcome on the information in the second period. We then regress its (estimated) conditional expectation on information in the first period (see Algorithm 2). Note that the coefficients $\beta^{(1)}_{d_1,d_2}$ would not be consistently estimated by simple linear regressions of the observed outcomes on information in the first period (see Remark 2).

In the following section, we discuss as main contribution balancing conditions with dynamic treatments.

Remark 2 (Why a model on potential outcomes?)**.** The model on potential outcomes instead of observed outcomes more flexibly accomodates dependence with and between treatment assignments, motivating our choice. First, with *binary* (but not necessarily continuous) treatments, linearity of observed outcomes might be violated. Namely, suppose that covariates are time invariant and let (with $Y_{i,0} = 0$)

$$\begin{aligned} Y_{i,t} &= Y_{i,t-1}\alpha + D_{i,t}\beta + X_{i,1}\gamma + \varepsilon_{i,t} \\ &\Rightarrow \mathbb{E}\Big[Y_{i,2}|X_{i,1}, D_{i,1}\Big] = \alpha\beta D_{i,1} + \mathbb{E}\Big[\beta D_{i,2}|X_{i,1}, D_{i,1}\Big] + X_{i,1}(\gamma + \alpha\gamma). \end{aligned} \tag{3}$$

[10] Here, we are implicitely assuming that the event $(D_{i,2}, D_{i,2}) = (d_1, d_2)$ has non zero probability for the conditional expectation to be well defined. This holds by discreteness of the treatment assignments. An explicit condition is imposed in Assumption 4 (ii).

Observe that $\mathbb{E}\Big[D_{i,2}|X_{i,1}, D_{i,1}\Big]$ (and hence $\mathbb{E}\Big[Y_{i,2}|X_{i,1}, D_{i,1}\Big]$) is not a linear function of $X_{i,1}$ *unconditionally* on treatment assignments, for binary treatments. This issue does *not* arise if we impose the model directly on the *potential* outcomes, as we do in the proposed potential projections.[11] Second, note that by projecting the observed (instead of potential) outcome $Y_{i,2}$ onto $(D_{i,1}, X_{i,1})$, would also capture effects of $D_{i,1}$ mediated through $D_{i,2}$, whenever $D_{i,2}$ depends on $D_{i,1}$.[12] This differs from the model we propose on potential outcomes, whose parameters have an interpretation in terms of a fixed treatment path. □

3 Dynamic Covariate Balancing

In this section, we discuss the main algorithmic procedure. We start introducing an estimator based on doubly-robust scores.

3.1 Balancing histories for causal inference

Following previous literature on doubly-robust scores (Jiang and Li, 2015; Nie et al., 2021; Tchetgen and Shpitser, 2012; Zhang et al., 2013), we propose an estimator that exploits linearity while reweighting observations to guarantee balance. We adapt the estimator to the local projection model. Formally, we consider an estimator

$$\begin{aligned}\hat{\mu}_2(d_1, d_2) = &\sum_{i=1}^{n} \Big\{ \hat{\gamma}_{i,2}(d_1, d_2) Y_{i,2} - \Big(\hat{\gamma}_{i,2}(d_1, d_2) - \hat{\gamma}_{i,1}(d_1, d_2) \Big) H_{i,2} \hat{\beta}^{(2)}_{d_1,d_2} \Big\} \\ &- \sum_{i=1}^{n} \Big(\hat{\gamma}_{i,1}(d_1, d_2) - \frac{1}{n} \Big) X_{i,1} \hat{\beta}^{(1)}_{d_1,d_2},\end{aligned} \tag{4}$$

[11]Returning to the previous example, observe that $\mathbb{E}\Big[Y_{i,2}(d_1, d_2)|X_{i,1}\Big] = \alpha\beta d_1 + \beta d_2 + X_{i,1}(\gamma + \alpha\gamma)$, which is linear in $X_{i,1}$, hence satisfying Assumption 3.

[12]It is interesting to note that this difference also relates to the interpretation of impulse response functions (IRF). IRF capture the effect of a contemporaneous treatment also mediated through future assignments if treatments are serially correlated. This can be noted from Equation (3) where a local projection on $D_{i,1}$ would also capture its effect mediated through $D_{i,2}$ (and possibly past outcomes). Here, we are concerned with the effects of a *treatment history* such as $(d_1 = 1, d_2 = 0)$ in the spirit of the Neyman-Rubin potential outcome framework (Imbens and Rubin, 2015) as opposed to the effect of a treatment $d_1 = 1$ also mediated through future assignments. This motivates our model on potential outcomes directly.

where we discuss the choice of the parameters $\hat{\beta}^{(1)}_{d_1,d_2}, \hat{\beta}^{(2)}_{d_1,d_2}$ in Section 3.2.

A possible choice of the weights $\hat{\gamma}_1, \hat{\gamma}_2$ are inverse probability weights (IPW). As for multi-valued treatments (Imbens, 2000), these weights can be written as follows

$$w_{i,1}(d_1, d_2) = \frac{1\{D_{i,1} = d_1\}}{nP(D_{i,1} = d_1 | X_{i,1})}, \quad w_{i,2}(d_1, d_2) = \frac{w_{i,1}(d_1, d_2) 1\{D_{i,2} = d_2\}}{P(D_{i,2} = d_2 | Y_{i,1}, X_{i,1}, X_{i,2}, D_{i,1})}. \tag{5}$$

However, in high dimensions, IPW weights require the correct specification of the propensity score, which in practice may be unknown. Motivated by these considerations, we propose replacing IPW with more stable weights.

We start studying covariate balancing conditions induced by the local projection model. By denoting $\bar{X}_1$ the sample average of covariates X_1, we can write

$$\hat{\mu}_2(d_1, d_2) = \bar{X}_1 \beta^{(1)}_{d_1,d_2} + T_1 + T_2 + T_3, \tag{6}$$

where

$$T_1 = \Big(\hat{\gamma}_1(d_1, d_2)^\top X_1 - \bar{X}_1\Big)(\beta^{(1)}_{d_1,d_2} - \hat{\beta}^{(1)}_{d_1,d_2}) + \Big(\hat{\gamma}_2(d_1, d_2)^\top H_2 - \hat{\gamma}_1(d_1, d_2)^\top H_2\Big)(\beta^{(2)}_{d_1,d_2} - \hat{\beta}^{(2)}_{d_1,d_2}) \tag{7}$$

and

$$T_2 = \hat{\gamma}_2(d_1, d_2)^\top \Big[Y_2 - H_2 \beta^2_{d_1,d_2}\Big], \quad T_3 = \hat{\gamma}_1(d_1, d_2)^\top \Big[H_2 \beta^2_{d_1,d_2} - X_1 \beta^{(1)}_{d_1,d_2}\Big].$$

The covariate balancing conditions must control T_1, while the remaining two are centered around zero under regularity conditions.

Lemma 3.1 (Covariate balancing conditions)**.** *The following holds*

$$T_1 \leq \underbrace{\|\hat{\beta}^1_{d_1,d_2} - \beta^{(1)}_{d_1,d_2}\|_1 \Big|\Big|\bar{X}_1 - \hat{\gamma}_1(d_1, d_2)^\top X_1\Big|\Big|_\infty}_{(i)} + \underbrace{\|\hat{\beta}^2_{d_1,d_2} - \beta^{(2)}_{d_1,d_2}\|_1 \Big|\Big|\hat{\gamma}_2(d_1, d_2)^\top H_2 - \hat{\gamma}_1(d_1, d_2)^\top H_2\Big|\Big|_\infty}_{(ii)}.$$

Element (i) is equivalent to what is discussed in Athey et al. (2018) in one period setting. Element (ii) depends instead on the additional error induced by the presence of a second period. Therefore the above suggests controlling the following norms

$$\left|\left|\bar{X}_1 - \hat{\gamma}_1(d_1, d_2)^\top X_1\right|\right|_\infty, \quad \left|\left|\hat{\gamma}_2(d_1, d_2)^\top H_2 - \hat{\gamma}_1(d_1, d_2)^\top H_2\right|\right|_\infty. \tag{8}$$

By imposing that the first norm converges to zero, the weights in the first-period balance covariates in the first period only. The second condition requires that histories in the second period are balanced, *given* the weights in the previous period.

The remaining terms in (6) are mean zero under the following conditions.

Lemma 3.2 (Balancing error)**.** *Let assumptions 1 - 3 hold. Suppose that $\hat{\gamma}_1$ is measurable with respect to $\sigma(X_1, D_1)$ and $\hat{\gamma}_2$ is measurable with respect to $\sigma(X_1, X_2, Y_1, D_1, D_2)$. Suppose in addition that $\hat{\gamma}_{i,1}(d_1, d_2) = 0$ if $D_{i,1} \neq d_1$ and $\hat{\gamma}_{i,2}(d_1, d_2) = 0$ if $(D_{i,1}, D_{i,2}) \neq (d_1, d_2)$. Then*

$$\mathbb{E}\Big[T_2\Big|X_1, D_1, Y_1, X_2, D_2\Big] = 0, \quad \mathbb{E}\Big[T_3\Big|X_1, D_1\Big] = 0.$$

The proof is in the Appendix. Lemma 3.2 conveys a key insight: if we can guarantee that each component in Equation (8) is $o_p(1)$, under mild regularity assumptions, the estimator $\hat{\mu}$ is centered around the target estimand plus an estimation error which is asymptotically negligible. As a result, the estimation error of the (high-dimensional) coefficients does not affect the rate of convergence of the estimator.

Interestingly, we note that Lemma 3.2 imposes the following intuitive condition. The balancing weights in the first period are non-zero only for those units whose assignment *in the first period* coincide with the target assignment d_1, and this also holds in the second period with assignments (d_1, d_2). Moreover, we can only balance based on information observed before the realization of potential outcomes but not based on future information. A special case is IPW in Equation (5), for known propensity score. An illustrative example is provided in Figure 3.

Finally, note that under mild misspecification of linearity, i.e., for approximately linear model as in Remark 1, our results hold for $o_p(n^{-1/2})$ conditional expectations. The reader may refer to Section 4 for a discussion on estimation with long panels.

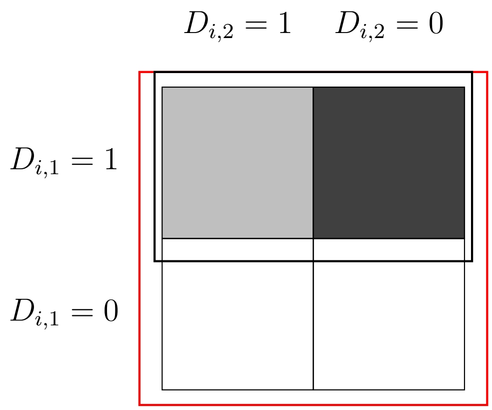

Figure 3: Illustrative description for balancing when estimating $\mathbb{E}[Y(1,1)]$, the average potential outcomes for those always under treatment. In the first period we balance covariates of those individuals with shaded areas (both light and dark gray) with covariates of all individuals in the region (red box). In the second period we balance covariates between the two shaded regions (black box).

3.2 Algorithm description

We can now introduce Algorithm 1. The algorithm works as follows. First, we construct weights in the first period that are nonzero only for those individuals with treatment at time $t = 1$ equal to the target treatment status d_1. We do the same for $\hat{\gamma}_{i,2}$ for the desired treatment history $(D_{i,1}, D_{i,2}) = (d_1, d_2)$. We then solve a quadratic program with linear constraints. In the first period, we balance covariates as in the one-period setting. In the second period, we balance present covariates with the same covariates, weighted by those weights obtained in the previous period. The weights sum to one, they are positive (to avoid aggressive extrapolation), and they do not assign the largest weight to few observations. We choose the weights to minimize their small sample variance to be robust to poor overlap in small samples.

Algorithm 2 summarizes the estimation of the regression coefficients.[13] The algorithm considers two separate model specifications which can be used. The first allows for all possible interactions of covariates and treatment assignments as in Assump-

[13]See Tran et al. (2019) for related procedures.

Algorithm 1 Dynamic covariate balancing (DCB): two periods

Require: Observations $(D_1, X_1, Y_1, D_2, X_2, Y_2)$, treatment history (d_1, d_2), finite parameters K, constraints $\delta_1(n, p_1), \delta_2(n, p_2)$.

1: Estimate $\beta^1_{d_{1:2}}, \beta^2_{d_{1:2}}$ as in Algorithm 2.

2: $\hat{\gamma}_{i,1} = 0$, if $D_{i,1} \neq d_1, \hat{\gamma}_{i,2} = 0$ if $(D_{i,1}, D_{i,2}) \neq (d_1, d_2)$

3: Estimate

$$\begin{aligned}
\hat{\gamma}_1 = \arg\min_{\gamma_1} ||\gamma_1||^2, \quad \text{s.t. } & \Big\|\bar{X}_1 - \frac{1}{n}\sum_{i=1}^n \gamma_{i,1} X_{i,1}\Big\|_\infty \leq \delta_1(n, p_1), \\
& 1^\top \gamma_1 = 1, \gamma_1 \geq 0, \|\gamma_1\|_\infty \leq \log(n) n^{-2/3}. \\
\hat{\gamma}_2 = \arg\min_{\gamma_2} ||\gamma_2||^2, \quad \text{s.t. } & \Big\|\frac{1}{n}\sum_{i=1}^n \hat{\gamma}_{i,1} H_{i,2} - \frac{1}{n}\sum_{i=1}^n \gamma_{i,2} H_{i,2}\Big\|_\infty \leq \delta_2(n, p_2), \\
& 1^\top \gamma_2 = 1, \gamma_2 \geq 0, \|\gamma_2\|_\infty \leq K \log(n) n^{-2/3}.
\end{aligned} \tag{9}$$

return $\hat{\mu}_2(d_1, d_2)$ as in Equation (4).

tion 3. The second is more parsimonious and assumes that treatment effects enter linearly in each equation, while it uses all the observations in the sample. The second specification can also contain linear interaction components, omitted for brevity. Note that the algorithm for the linear (second) specification builds predictions in the second period only for those units with $D_{i,1} = d_1$, and for all units in the first period. This is without loss of generality, since the remaining units receive a zero weight.

3.3 Existence and convergence rate

We conclude this introductory discussion by developing properties of the estimator. We first impose the following tail decay conditions.

Assumption 4. Let the following hold:

(i) $H^{(j)}_{i,2}$ is subgaussian given the past history for each entry $j \in \{1, \cdots, p_2\}$ and $X^{(j)}_{i,1}$ also is subgaussian for each $j \in \{1, \cdots, p_1\}$.

(ii) Assume that (i) $P(D_{i,1} = 1|X_{i,1}), P(D_{i,2} = 1|D_1, X_1, X_2, Y_1) \in (\delta, 1 - \delta), \delta \in (0, 1)$.

Algorithm 2 Coefficients estimation

Require: Observations, history $d_{1:2} = (d_1, d_2)$, `model` $\in$ {`full interactions`, `linear`}.

if `model` = `full interactions` **then**

  Estimate $\beta^2_{d_{1:2}}$ by regressing $Y_{i,2}$ onto $H_{i,1}$ for all $i : (D_{i,1:2} = d_{1:2})$;

  Estimate $\beta^1_{d_{1:2}}$ by regression $H_{i,1}\hat{\beta}^2_{d_{1:2}}$ onto $X_{i,1}$ for i that has $D_{i,1} = d_1$.

else

  Estimate $\beta^{(2)}$ by regressing $Y_{i,2}$ onto $(H_{i,1}, D_{i,2})$ for all i (without penalizing $(D_{i,1}, D_{i,2})$) and define $H_{i,2}\hat{\beta}_{d_1,d_2} = (H_{i,2}, d_2)\hat{\beta}^2$ for all $i : D_{i,1} = d_1$;

  Estimate $\beta^{(1)}$ by regressing $(H_{i,1}, d_2)\hat{\beta}^2$ onto $(X_{i,1}, D_{i,1})$ for all i (without penalizing $D_{i,1}$) and define $X_{i,1}\hat{\beta}^{(1)}_{d_1,d_2} = (X_{i,1}, d_1)\hat{\beta}^{(1)}$ for all i.

end if

The first condition states that histories are Sub-Gaussian. The second condition imposes overlap of the propensity score.

Theorem 3.3 (Existence of a feasible $\hat{\gamma}_t$)**.** *Let Assumptions 1 - 4 hold. Suppose that* $\delta_t(n, p_t) \geq c_0 \log^{3/2}(np_t)/n^{1/2}$*, for a finite constant* c_0*. Then, with probability* $\eta_n \to 1$*, for each* $t \in \{1, 2\}$*, for some* $N > 0$*,* $n > N$*, there exists a feasible* $\hat{\gamma}^*_t$*, solving the optimization in Algorithm 1, where*

$$\hat{\gamma}^*_{i,0} = 1/n, \quad \hat{\gamma}^*_{i,t} = \hat{\gamma}^*_{i,t-1} \frac{1\{D_{i,t} = d_t\}}{P(D_{i,t} = d_t|\mathcal{F}_{t-1})} \Big/ \sum_{i=1}^{n} \hat{\gamma}^*_{i,t-1} \frac{1\{D_{i,t} = d_t\}}{P(D_{i,t} = d_t|\mathcal{F}_{t-1})},$$

and $\mathcal{F}_0 = \sigma(X_1), \mathcal{F}_1 = \sigma(X_1, X_2, Y_1, D_1)$*.*

Theorem 3.3 has important practical implications. Inverse probability weights tend to be unstable in a small sample for moderately large periods. The algorithm thus finds weights that minimize the small sample variance, with the IPW weights being allowed to be *one* of the possible solutions. We formalize this below.

Corollary 1. *Under the conditions in Theorem 3.3, for some* $N > 0$*,* $n > N$*, with probability* $\eta_n \to 1$*,* $n||\hat{\gamma}_t||^2 \leq n||\hat{\gamma}^*_t||^2$*, for* $t \in \{1, 2\}$*.*

We now characterize the convergence rate.

Assumption 5. Let $\delta_t(n, p_t)$ is such that $\delta_t(n, p_t) \geq c_0 \log(np_t)/n^{1/4}$ for a finite constant c_0, $\|\hat{\beta}^t_{d_{1:2}} - \beta^t_{d_{1:2}}\|_1 \delta_t(n, p_t) = o_p(1/\sqrt{n})$, $t \in \{1, 2\}$, $\|\hat{\beta}^t_{d_{1:2}} - \beta^t_{d_{1:2}}\|_1 = o_p(n^{-1/4})$.

The above condition states that the estimation error of the linear regressor times the parameter $\delta_t(n, p_t) = o(1)$ is $\mathcal{O}(1/\sqrt{n})$. A simple example is an high-dimensional regression, where $\|\hat{\beta}^t_{d_{1:2}} - \beta^t_{d_{1:2}}\|_1 = O_p(\sqrt{\log(p_t)/n})$. Sufficient conditions for Assumption 5 are in Appendix B.1.2. Also, note that here the bound on covariates imbalance (controlled through $\delta_t(n, p_t)$) becomes less stringent as t increases, since p_t is increasing in the number of periods. However, since we might expect that p_t grows linearly in t, $\delta_t(n, p_t)$ grows logarithmically in t (for longer periods, researchers can condition on a larger set of covariates). See Section 4 for details.

Under Assumption 3, let

$$Y_{i,2}(d_1, d_2) = H_{i,2}(d_1)\beta^{(2)}_{d_1,d_2} + \varepsilon_{i,2}(d_1, d_2), \quad H_{i,2}(d_1)\beta^{(2)}_{d_1,d_2} = X_{i,1}(d_1)\beta^{(1)}_{d_1,d_2} + \nu_{i,1}(d_1),$$

where $\nu_{i,1}(d_1) = \mathbb{E}\Big[Y_{i,2}(d_1, d_2)|H_{i,2}(d_1)\Big] - \mathbb{E}\Big[Y_{i,2}(d_1, d_2)|X_{i,1}\Big]$ denotes the difference between the two local projections over two consecutive periods.

Assumption 6. Let the following hold:

(A) $\mathbb{E}[\varepsilon_2^4(d_1, d_2)|H_2], \mathbb{E}[\nu_1^4(d_1)|X_1] < C$ for a finite constant C almost surely;

(B) $\mathrm{Var}(\varepsilon_2(d_1, d_2)|H_{i,2}), \mathrm{Var}(\nu_1(d_1, d_2)|X_{i,1}) > u_{min} > 0$.

The above condition states that the residuals from projections in two consecutive time periods have non-zero variance and a bounded fourth moment.

Theorem 3.4. *Let Assumptions 1 - 6 hold. Then, whenever* $\log(n(p_1+p_2))/n^{1/4} \to 0$ *with* $n, p_1, p_2 \to \infty$,

$$\hat{\mu}_2(d_1, d_2) - \mu_2(d_1, d_2) = \mathcal{O}_P\Big(n^{-1/2}\Big).$$

The proofs are in the Appendix. Theorem 3.4 shows that the proposed estimator guarantees parametric convergence rate with high-dimensional covariates. Observe that the theorem does not require conditions on the convergence rate of the estimated propensity score, instead commonly encountered in the doubly-robust literature (Farrell, 2015). Inference is discussed with T periods in Section 4, where we construct confidence bands, using the square-root of the critical quantile of a chi-squared distribution with T degrees of freedom.

4 The general case: multiple time periods and inference

In this section we generalize our procedure to T time periods. Let $d_{1:T} = (d_1, \cdots, d_T)$. We define the estimand of interest as:

$$\text{ATE}(d_{1:T}, d'_{1:T}) = \mu_T(d_{1:T}) - \mu_T(d'_{1:T}), \quad \mu_T(d_{1:T}) = \frac{1}{n}\sum_{i=1}^{n} \mathbb{E}\Big[Y_T(d_{1:T})\Big|X_{i,1}\Big]. \tag{10}$$

This estimand denotes the difference in potential outcomes conditional on baseline covariates. We define $\mathcal{F}_t = \Big(D_1, \cdots, D_{t-1}, X_1, \cdots, X_t, Y_1, \cdots, Y_{t-1}\Big)$ the information at time t after excluding the treatment assignment D_t. We denote

$$H_{i,t} = \Big[D_{i,1}, \cdots, D_{i,t-1}, X_{i,1}, \cdots, X_{i,t}, Y_{i,1}, \cdots, Y_{i,t-1}\Big] \in \mathcal{H}_t \subseteq \mathbb{R}^{p_t} \tag{11}$$

the vector containing information from time one to time t, after excluding the treatment assigned in the present period D_t. Here $\mathcal{H}_t$ denotes the space of history at time t. Interaction components may also be considered in the above vector, and they are omitted for expositional convenience only. We let the potential history be

$$H_{i,t}(d_{1:(t-1)}) = \Big[d_{1:(t-1)}, X_{i,1:t}(d_{1:(t-1)}), Y_{i,1:(t-1)}(d_{1:(t-1)})\Big],$$

as a function of the treatment history. The following Assumption generalizes Assumptions 1-3 from the two-period setting: no-anticipation, sequential ignorability, and potential outcome models.

Assumption 7. For any $d_{1:T}, \in \{0,1\}^T$, and $t \leq T$,

(A) (No-anticipation) The potential history $H_{i,t}(d_{1:T})$ is constant in $d_{t:T}$;

(B) (Sequential ignorability) $\Big(Y_{i,T}(d_{1:T}), H_{i,t+1}(d_{1:(t+1)}), \cdots, H_{i,T-1}(d_{1:(T-1)})\Big) \perp D_{i,t}|\mathcal{F}_t$;

(C) (Potential projections) For some $\beta^{(t)}_{d_{1:T}} \in \mathbb{R}^{p_t}$,

$$\mathbb{E}\Big[Y_{i,T}(d_{1:T})|D_{i,1:(t-1)} = d_{1:(t-1)}, X_{i,1:t}, Y_{i,1:(t-1)}\Big] = H_{i,t}(d_{1:(t-1)})\beta^{(t)}_{d_{1:T}}.$$

Condition (A) imposes a non-anticipatory behavior of histories at each point in time, as commonly assumed in practice (Boruvka et al., 2018). With a slight abuse of notation, we implicitly impose (A), by referring to the potential history as $H_{i,t}(d_{1:(t-1)})$. Condition (B) states that treatment assignments are randomized based on the past only. Condition (C) states that the conditional expectation of the potential outcome at the end-line period is linear in the potential history, $H_{i,t}(d_{1:(t-1)})$. Identification follows similarly to Lemma 2.1 and omitted for brevity.

Once DCB weights are formed, we construct the estimator of $\mu_T(d_{1:T})$ as

$$\begin{aligned}\hat{\mu}_T(d_{1:T}) =& \sum_{i=1}^n \hat{\gamma}_{i,T}(d_{1:T})Y_{i,T} - \sum_{i=1}^n \sum_{t=2}^T \Big(\hat{\gamma}_{i,t}(d_{1:T}) - \hat{\gamma}_{i,t-1}(d_{1:T})\Big) H_{i,t}\hat{\beta}^{(t)}_{d_{1:T}} \\ & - \sum_{i=1}^n \Big(\hat{\gamma}_{i,1}(d_{1:T}) - \frac{1}{n}\Big) X_{i,1}\hat{\beta}^{(1)}_{d_{1:T}}.\end{aligned} \tag{12}$$

Lemma 4.1. *Suppose that $\hat{\gamma}_{i,T}(d_{1:T}) = 0$ if $D_{i,1:T} \neq d_{1:T}$. Then*

$$\begin{aligned}\hat{\mu}_T(d_{1:T}) - \mu_T(d_{1:T}) =& \underbrace{\sum_{t=1}^T \Big(\hat{\gamma}_t(d_{1:T})H_t - \hat{\gamma}_{t-1}(d_{1:T})H_t\Big)(\beta^{(t)}_{d_{1:T}} - \hat{\beta}^{(t)}_{d_{1:T}})}_{(I_1)} + \underbrace{\hat{\gamma}_T^\top(d_{1:T})\varepsilon_T}_{(I_2)} \\ & + \underbrace{\sum_{t=2}^T \hat{\gamma}_{t-1}(d_{1:T})\Big(H_t\beta^{(t)}_{d_{1:T}} - H_{t-1}\beta^{(t-1)}_{d_{1:T}}\Big)}_{(I_3)}\end{aligned} \tag{13}$$

where $\varepsilon_{i,t}(d_{1:T}) = Y_{i,T}(d_{1:T}) - H_{i,t}(d_{1:(t-1)})\beta^{(t)}_{d_{1:T}}$.

The proof is relegated to the Appendix. Lemma 4.1 decomposes the estimation error into three components. First, (I_1), depends on the estimation error of the coefficient and on balancing properties of the weights. (I_1) suggests imposing balancing conditions on

$$\Big|\Big|\hat{\gamma}_t(d_{1:T})H_t - \hat{\gamma}_{t-1}(d_{1:T})H_t\Big|\Big|_\infty$$

each period. The components characterizing the estimation error are $(I_2) = \hat{\gamma}_T(d_{1:T})^\top \varepsilon_T$, and (I_3). In the following lemma, we provide conditions such that (I_3) is mean zero.

Lemma 4.2. *Let Assumption 7 hold. Suppose that the sigma algebra $\sigma(\hat{\gamma}_t(d_{1:T})) \subseteq \sigma(\mathcal{F}_t, D_t)$. Suppose in addition that $\hat{\gamma}_{i,t}(d_{1:T}) = 0$ if $D_{i,1:t} \neq d_{1:t}$. Then*

$$\mathbb{E}\Big[\hat{\gamma}_{i,t-1}(d_{1:T})H_t\beta^{(t)}_{d_{1:T}} - \hat{\gamma}_{i,t-1}(d_{1:T})H_{t-1}\beta^{(t-1)}_{d_{1:T}}\Big|\mathcal{F}_{t-1}, D_{t-1}\Big] = 0.$$

The proof is presented in the Appendix. The above condition states that weights need to be estimated using observations that match the desired treatment path up at every t and are equal to zero on the other treatment paths. In the presence of long panels, we can relax the conditions that $D_{i,1:t} \neq d_{1:t}$ by studying carry-over effects over a limited time period h and imposing that the treatment path matches only over the last h periods. See Remark 4 for details.

Algorithm D.1 in the Appendix present the balancing algorithm for generic T periods (which follows similarly to two periods) and Algorithm D.3 shows how to choose tuning parameters adaptively. Coefficients are estimated recursively as discussed in the two periods setting (see Algorithm D.2 in the Appendix).

Remark 3 (Estimation error of the coefficients)**.** The estimation error $||\hat{\beta}^{(t)}_{d_{1:T}} - \beta^{(t)}_{d_{1:T}}||_1$ can scale either linearly or exponentially with T, depending on modeling assumptions. Whenever we let coefficients be different across entire different treatment histories, $||\hat{\beta}^{(t)}_{d_{1:T}} - \beta^{(t)}_{d_{1:T}}||_1$ would scale exponentially with T, since we would need to run different regressions over the subsample with treatment histories $D_{1:t} = d_{1:t}$. On the other hand, additional assumptions permit to estimate $\hat{\beta}^{(t)}_{d_{1:T}}$ using most or all in-sample information. A simple example, is to explicitely model the effect of the treatment history $d_{1:T}$ on the outcome, as in the linear model in Algorithm D.2. □

Remark 4 (Pooled regression and limited carry-overs)**.** In some application, we may be interested in a regression of the following form

$$Y_{i,t}(d_{1:t}) = \beta_0 + \beta_1 d_t + \beta_2 Y_{i,t-1}(d_{1:(t-1)}) + X_{i,t}(d_{1:(t-1)})\gamma + \tau_t + \varepsilon_{i,t},$$

where τ_t denotes fixed effects, and in the estimand

$$\mathbb{E}[Y_{i,t+h}(d_{1:t}, d_{t+1}, \cdots, d_{t+h})] - \mathbb{E}[Y_{i,t+h}(d_{1:t}, d'_{t+1}, \cdots, d'_{t+h})],$$

denoting the effect of changing treatment history in the past h periods. Estimation can be performed by considering each (i,t) as an observation for all $t > h$ and estimate its corresponding weight.[14] By considering the effect over a limited number of periods (h instead of T), construction of the weights requires that past treatment assignments coincide with the treatment history ($d_{(t+1):(t+h)}$) only over h periods. $\square$

4.1 Asymptotic properties

We now derive properties as long as $\log(\max_t p_t n)/n^{1/4} \to 0$ while $p_1, \cdots, p_T$ can potentially grow to infinity. For simplicity, we let $\max_t p_t = p_T$ (note that we expect p_t to grow linearly in t). We consider a finite-time horizon and $T < \infty$ regime.

We discuss the first regularity condition below, similarly to two periods.

Assumption 8 (Overlap and tails' conditions)**.** Assume that $P(D_{i,t} = d_t | \mathcal{F}_{t-1}, D_{t-1}) \in (\delta, 1-\delta), \delta \in (0,1)$ for each $t \in \{1, \cdots, T\}$. Assume also that $H_{i,t}^{(j)}, j \in \{1, \cdots, p_t\}$ is Sub-Gaussian given past history and similarly $X_{i,1}^{(j)}, j \in \{1, \cdots, p_1\}$.

The first condition is the overlap condition as in the case of two periods. The second condition is a tail restriction. In the following theorem, we characterize the existence of a solution to the optimization program.

[14] Note that here τ_t acts as an additional covariate for balancing. We obtain the corresponding variances after clustering residuals of the same individuals over different periods for a panel with finite T periods.

Theorem 4.3. *Let Assumptions 7, 8 hold. Consider* $\delta_t(n, p_t) \geq c_0 n^{-1/2} \log^{3/2}(p_t n)$ *for a finite constant* c_0*, and* $K_{2,t} = 2K_{2,t-1} b_t$ *for some constant* $b_t < \infty$*. Then, with probability* $\eta_n \to 1$*, for each* $t \in \{1, \cdots, T\}, T < \infty$*, for some* $N > 0$*,* $n > N$*, there exists a feasible* $\hat{\gamma}_t^*$*, solving the optimization in Algorithm D.1, where*

$$\hat{\gamma}_{i,0}^* = 1/n, \quad \hat{\gamma}_{i,t}^* = \hat{\gamma}_{i,t-1}^* \frac{1\{D_{i,t} = d_t\}}{P(D_{i,t} = d_t | \mathcal{F}_{t-1}, D_{t-1})} \Big/ \sum_{i=1}^n \hat{\gamma}_{i,t-1}^* \frac{1\{D_{i,t} = d_t\}}{P(D_{i,t} = d_t | \mathcal{F}_{t-1}, D_{t-1})}.$$

The above theorem shows existence of a feasible solution which encompasses stabilized inverse probability weights. Next, we characterize asymptotic properties of the estimator. We impose high-level assumptions on the coefficients similarly to the two-periods setting.

Assumption 9. Let the following hold: for every $t \in \{1, \cdots, T\}$, $d_{1:T} \in \{0, 1\}^T$,

(i) $\max_t \|\hat{\beta}_{d_{1:T}}^{(t)} - \beta_{d_{1:T}}^{(t)}\|_1 \delta_t(n, p_t) = o_p(1/\sqrt{n})$, $\delta_t(n, p_t) \geq c_{0,t} n^{-1/4} \log(2p_t n)$ for a finite constant $c_{0,t}$, $\max_t \|\hat{\beta}_{d_{1:T}}^{(t)} - \beta_{d_{1:T}}^{(t)}\|_1 = o_p(n^{-1/4})$;

(ii) For a finite constant C, $\mathbb{E}[\varepsilon_{i,T}^4 | H_T] < C$ almost surely, with $\varepsilon_{i,T} = Y_{i,T} - H_{i,T} \beta_{d_{1:T}}^{(t)}$; $\mathbb{E}[(H_{i,t} \beta_{d_{1:T}}^{(t)} - H_{i,t-1} \beta_{d_{1:T}}^{(t-1)})^4 | H_{i,t-1}] < C$;

(iii) $\mathrm{Var}(\varepsilon_{i,T} | H_{i,T}), \mathrm{Var}(H_{i,t} \beta_{d_{1:T}}^{(t)} - H_{i,t-1} \beta_{d_{1:T}}^{(t-1)} | H_{i,t-1}) > u_{min}$, almost surely, for some constant $u_{min} > 0$.

Assumption 9 imposes the consistency in estimation of the outcome models. Condition (i) is attained for many high-dimensional estimators, such as the lasso method, under regularity assumptions; see e.g., Bühlmann and Van De Geer (2011). A discussion is included in Example B.1 which is valid recursively for any finite T (see Appendix B.1.2). The remaining conditions impose moment assumptions similarly to the two periods setting. Similarly to what discussed in Section 3, the balancing constant $\delta_t(n, p_t)$ is less stringent for larger t, growing logarithmically in t as the dimension p_t grows linearly in t.

Theorem 4.4 (Asymptotic Inference)**.** *Let Assumptions 7 - 9 hold. Then, whenever* $\log(np_T)/n^{1/4} \rightarrow 0$*, as* $n, p_1, \cdots, p_T \rightarrow \infty$*,*

$$P\Big(\Big|\frac{\sqrt{n}\Big(\hat{\mu}(d_{1:T}) - \mu_T(d_{1:T})\Big)}{\hat{V}_T(d_{1:T})^{1/2}}\Big| > \sqrt{\chi_T(\alpha)}\Big) \leq \alpha, \quad \hat{\mu}_T(d_{1:T}) - \mu_T(d_{1:T}) = \mathcal{O}_P(n^{-1/2}), \tag{14}$$

where

$$\hat{V}_T(d_{1:T}) = n\sum_{i=1}^{n} \hat{\gamma}_{i,T}^2(d_{1:T})(Y_i - H_{i,T}\hat{\beta}^{(t)}_{d_{1:T}})^2 + \sum_{t=1}^{T-1} n \sum_{i=1}^{n} \hat{\gamma}_{i,t}^2(d_{1:t})(H_{i,t+1}\hat{\beta}^{t+1}_{d_{1:T}} - H_{i,t}\hat{\beta}^{t}_{d_{1:T}})^2$$

and $\chi_T(\alpha)$ *is* $(1-\alpha)$ *quantile of a chi-squared random variable with* T *degrees of freedom.*

The proofs of the above two theorems are contained in the Appendix. Note that the above theorem is also valid in low dimensional settings, where $p_1, \cdots, p_T$ are finite. We now discuss inference on the ATE.

Theorem 4.5 (Inference on ATE)**.** *Let the conditions in Theorem 4.4 hold. Let* $d_1 \neq d_1'$ *Then, whenever* $\log(np_T)/n^{1/4} \rightarrow 0$ *with* $n, p_1, \cdots, p_T \rightarrow \infty$*,*

$$P\Big(\Big|(\hat{V}_T(d_{1:T}) + \hat{V}_T(d'_{1:T}))^{-1/2}\sqrt{n}\Big(\hat{\mu}(d_{1:T}) - \hat{\mu}(d'_{1:T}) - \text{ATE}(d_{1:T}, d'_{1:T})\Big)\Big| > \sqrt{\chi_{2T}(\alpha)}\Big) \leq \alpha.$$

The proof is in the Appendix. The above theorem permits inference on the ATE. We conclude our discussion with a final remark.

Remark 5 (Tighter confidence bands under more restrictive conditions)**.** The confidence band depends on a chisquared random variable with T degrees of freedom. In Appendix C.2 we show that under additional conditions we can get

$$(\hat{V}_T(d_{1:T}) + \hat{V}_T(d'_{1:T}))^{-1/2}\sqrt{n}\Big(\hat{\mu}(d_{1:T}) - \hat{\mu}(d'_{1:T}) - \text{ATE}(d_{1:T}, d'_{1:T})\Big) \rightarrow_d \mathcal{N}(0,1)$$

and hence, tighter confidence bands. The assumptions needed is that $n||\hat{\gamma}_t||_2^2$ converge almost surely to a finite constant. This condition imposes restrictions on the degree of dependence of the optimal weights, and holds for a bernoulli design. □

5 Numerical Experiments

This section collects results from numerical experiments.[15] We estimate in two and three periods

$$\mathbb{E}\Big[Y_{i,T}(\mathbf{1}) - Y_{i,T}(\mathbf{0})\Big], \quad T \in \{2, 3\}.$$

We let the baseline covariates $X_{i,1}$ be drawn from as i.i.d. $\mathcal{N}(0, \Sigma)$ with $\Sigma^{(i,j)} = 0.5^{|i-j|}$. Covariates in the subsequent period are generated according to an auto-regressive model $\{X_{i,t}\}_j = 0.5\{X_{i,t-1}\}_j + \mathcal{N}(0, 1), j = 1, \cdots, p_t$. Treatments are drawn from a logistic model that depends on all previous treatments as well as previous covariates. Namely, $D_{i,t} \sim \text{Bern}\Big((1 + e^{\iota_{i,t}})^{-1}\Big)$ with

$$\iota_{i,t} = \eta \sum_{s=1}^{t} X_{i,s}\phi + \sum_{s=1}^{t-1} \delta_s (D_{i,s} - \bar{D}_s) + \xi_{i,t}, \quad \bar{D}_s = n^{-1} \sum_{i=1}^{n} D_{i,s} \tag{15}$$

and $\xi_{i,t} \sim \mathcal{N}(0, 1)$, for $t \in \{1, 2, 3\}$. Here, η, δ controls the association between covariates and treatment assignments. We consider values of $\eta \in \{0.1, 0.3, 0.5\}$, $\delta_1 = 0.5, \delta_2 = 0.25$. We let $\phi \propto 1/j$, with $\|\phi\|_2^2 = 1$, similarly to what discussed in Athey et al. (2018). Table 1 illustrates the behavior of the propensity score as a function of η. The larger the value of η, the weaker the overlap.

We generate the outcome according to the following equations:

$$Y_{i,t}(d_{1:t}) = \sum_{s=1}^{t} \Big(X_{i,s}\beta + \lambda_{s,t} Y_{i,s-1} + \tau d_s\Big) + \varepsilon_{i,t}(d_{1:t}), \quad t = 1, 2, 3,$$

where elements of $\varepsilon_{i,t}(d_{1:t})$ are i.i.d. $\mathcal{N}(0, 1)$ and $\lambda_{1,2} = 1, \lambda_{1,3}, \lambda_{2,3} = 0.5$. We consider three different settings: **Sparse** with $\beta^{(j)} \propto 1\{j \leq 10\}$, **Moderate** with moderately sparse $\beta^{(j)} \propto 1/j^2$ and the **Harmonic** setting with $\beta^{(j)} \propto 1/j$. We ensure $\|\beta\|_2 = 1$. Throughout our simulations we set $\tau = 1$. In Appendix E.1 we collect results in the presence of non-linear (misspecified) outcome models.

[15] Replication code is available on the website https://dviviano.github.io/projects/.

Table 1: Summary statistics of the distribution of the propensity score in two and three periods in a sparse setting with $\dim(X) = 300$.

	$\eta = 0.1$		$\eta = 0.3$		$\eta = 0.5$	
	T=2	T=3	T=2	T=3	T=2	T=3
Min	0.012	0.003	0.004	0.0002	0.001	0.00000
1st Quantile	0.126	0.049	0.105	0.031	0.079	0.018
Median	0.218	0.097	0.216	0.097	0.216	0.094
3rd Quantile	0.248	0.126	0.259	0.153	0.277	0.183
Max	0.352	0.175	0.377	0.226	0.429	0.286

5.1 Methods

We consider the following competing methodologies. **Augmented IPW**, with *known* propensity score and with *estimated* propensity score. The method replaces the balancing weights in Equation (4) with the (estimated or known) propensity score. Estimation of the propensity score is performed using a logistic regression (denoted as aIPWl) and a penalized logistic regression (denoted as aIPWh).[16] For both AIPW and IPW we consider stabilized inverse probability weights. We also compare to existing balancing procedures for dynamic treatments. Namely, we consider Marginal Structural Model (MSM) with balancing weights computed using the method in Yiu and Su (2018, 2020). The method consists of estimating *Covariate-Association Balancing weights* **CAEW (MSM)** as in Yiu and Su (2018, 2020), which consists in balancing covariates reweighted by marginal probabilities of treatments (estimated with a logistic regression), and use such weights to estimate marginal structural model of the outcome linear in past treatment assignments. We follow Section 3 in Yiu and Su (2020) for its implementation (here we do not also compare to Imai and Ratkovic 2015 for MSM since it is intractable in high-dimensions).[17] We also

[16]See for example Nie et al. (2021) and the recent work of Bodory et al. (2020) for a discussion on doubly-robust estimators.

[17]Estimation consists in projecting the outcome on the two or three past assignemnts, use the CAEW for reweighting. The reader can also refer to Blackwell (2013) for references on marginal

consider **"Dynamic" Double Lasso** that estimates the effect of each treatment assignment separately, after conditioning on the present covariate and past history for each period using the double lasso discussed in one period setting in Belloni et al. (2014).[18] **Naive Lasso** runs a regression controlling for covariates and treatment assignments only. Finally, **Sequential Estimation** estimates the conditional mean in each time period sequentially using the lasso method, and it predicts end-line potential outcomes as a function of the estimated potential outcomes in previous periods. For Dynamic Covariate Balancing, **DCB** choice of tuning parameters is data adaptive, and it uses a grid-search method discussed in Appendix D.[19] We estimate coefficients as in Algorithm 2 for DCB and (a)IPW, with a linear model in treatment assignments. Estimation of the penalty for the lasso methods is performed via cross-validation.

5.2 Results

We consider $\dim(\beta) = \dim(\phi) = 100$ and set the sample size to be $n = 400$. Under such design, the regression in the first period contains $p_1 = 101$ covariates, in the second period $p_2 = 203$ covariates, and in the third $p_3 = 305$ covariates.

In Table 2 we collect results for the average mean squared error in two and three periods. Throughout all simulations, the proposed method significantly outperforms any other competitor for $T = 3$, with one single exception for $T = 2$, good overlap and harmonic design. It also outperforms the case of known propensity score, consistently with our findings in Theorem 3.3. Improvements are particularly significant when (i) overlap deteriorates; (ii) the number of periods increases from two to three. This can also be observed in the panel at the bottom of Figure 4, where we report the decrease in MSE (in logarithmic scale) when using our procedure for $T = 3$. In Appendix E.1 we collect additional results with misspecified models.

structural models.

[18] See Lewis and Syrgkanis (2020) for related procedures.

[19] The grid-search procedure consists of finding the smallest feasible constraint-value through grid search, while choosing more stringent constants for those variables whose estimated coefficients are non-zero.

Table 2: Mean Squared Error (MSE) of Dynamic Covariate Balancing (DCB) across 200 repetitions with sample size 400 and 101 variables in time period 1. This implies that the number of variables in time period 2 and 3 are 203 and 304. Oracle Estimator is denoted with aIPW* whereas aIPWh(l) denote AIPW with high(low)-dimensional estimated propensity. CAEW (MSM) corresponds to the method in Yiu and Su (2020), D.Lasso is adaptation of Double Lasso (Belloni et al., 2014).

	$\eta = 0.1$			$\eta = 0.3$			$\eta = 0.5$		
	sparse	mod	harm	sparse	mod	harm	sparse	mod	harm
$T = 2$									
aIPW*	0.069	0.092	0.071	0.102	0.104	0.118	0.131	0.127	0.132
DCB	0.060	0.077	0.075	0.092	0.076	0.084	0.099	0.077	0.085
aIPWh	0.064	0.091	0.070	0.180	0.204	0.218	0.265	0.312	0.368
aIPWl	0.260	0.229	0.212	0.157	0.201	0.165	0.214	0.234	0.213
IPWh	2.37	1.78	2.80	10.19	6.49	11.72	15.25	8.09	16.67
Seq.Est.	0.932	1.333	0.692	1.388	1.787	1.152	1.759	1.795	1.664
Lasso	0.247	0.410	0.132	0.509	0.710	0.298	0.762	0.948	0.560
CAEW	0.432	0.444	0.517	1.934	1.274	1.974	3.376	2.168	4.423
Dyn.D.Lasso	0.124	0.118	0.256	0.208	0.147	0.430	0.218	0.153	0.554
$T = 3$									
aIPW*	0.226	0.296	0.261	0.403	0.251	0.339	0.472	0.496	0.562
DCB	0.155	0.208	0.199	0.257	0.217	0.329	0.294	0.267	0.455
aIPWh	0.201	0.273	0.280	0.595	0.747	0.835	0.999	1.328	1.607
aIPWl	0.823	0.625	0.829	0.623	0.704	0.638	1.078	1.396	1.234
IPWh	11.03	8.09	12.84	34.65	20.34	39.37	47.65	23.30	45.47
Seq.Est.	2.608	4.016	2.316	3.722	5.269	3.818	5.279	6.829	5.467
Lasso	0.409	0.492	0.514	0.559	0.732	0.507	1.290	1.315	1.174
CAEW	3.580	2.446	4.279	18.50	12.07	22.85	30.07	18.71	33.01
Dyn.D.Lasso	0.471	0.344	0.679	0.694	0.378	1.182	0.964	0.383	1.594

In the top panel of Figure 4 we report the length of the confidence interval and the point estimates. The length increases with the number of periods, and point estimates are more accurate for a non-harmonic (more sparse) setting.

Finally, we report finite sample coverage of the proposed method, DCB in Table 3 for estimating $\mu(1,1)$ and $\mu(1,1) - \mu(0,0)$ in the first two panel with $\eta = 0.5$ (see

Table 3: Conditional average Coverage Probability of Dynamic Covariate Balancing (DCB) over 200 repetitions, with $\eta = 0.5$ (*poor* overlap). Here, $n = 400$ and $p = 100$; implying that the number of variables at time 2 and time 3 are $2p$ and $3p$, respectively. Homoskedastic and heteroskedastic estimators of the variance are denoted with Ho and He, respectively. The first two panels use the square-root of the chi-squared critical quantiles as discussed in Theorems 4.4, 4.5 (see Table E.1 in the Appendix for the confidence intervals' length) and the last panel uses instead critical quantiles from the standard normal table (see Remark 5).

	$T=2$						$T=3$					
	Sparse		Moderate		Harmonic		Sparse		Moderate		Harmonic	
	Ho	He	Ho	He	Ho	He	Ho	He	Ho	He	Ho	He
$\mu(1,1)$: 95% Coverage Probability												
p=100	1.00	0.98	1.00	0.99	0.99	0.96	0.99	0.99	1.00	1.00	1.00	0.96
p=200	0.99	0.99	0.99	0.98	0.97	0.95	1.00	0.99	0.99	0.98	0.99	0.93
p=300	1.00	0.99	0.99	0.99	0.96	0.94	0.99	0.97	0.99	0.97	0.98	0.93
$\mu(1,1) - \mu(0,0)$: 95% Coverage Probability												
p=100	1.00	0.99	1.00	1.00	1.00	1.00	1.00	1.00	1.00	0.99	0.99	0.99
p=200	1.00	1.00	0.99	0.99	1.00	0.99	1.00	0.99	1.00	1.00	0.97	0.96
p=300	1.00	1.00	1.00	1.00	0.98	0.97	1.00	0.99	1.00	0.99	0.99	0.97
$\mu(1,1) - \mu(0,0)$: 95% Coverage Probability with *Gaussian* quantile												
p=100	0.96	0.94	0.98	0.96	0.98	0.94	0.97	0.90	0.98	0.92	0.90	0.79
p=200	0.97	0.94	0.98	0.92	0.91	0.85	0.98	0.92	0.98	0.91	0.75	0.64
p=300	0.99	0.96	0.99	0.95	0.89	0.84	0.92	0.85	0.94	0.86	0.73	0.61

Table E.1 in the Appendix for the confidence intervals' length).[20] The former is of interest when the effect under control is more precise and its variance is asymptotically neglegible compared to the estimated effect under treatment (e.g., many more individuals are not exposed to any treatment). The latter is of interest when both

[20]Results for $\eta \in \{0.1, 0.3\}$ present over-coverage of the chi-squared method, and correct or under-coverage (albeit less severe than $\eta = 0.5$) when considering a Gaussian critical quantile.

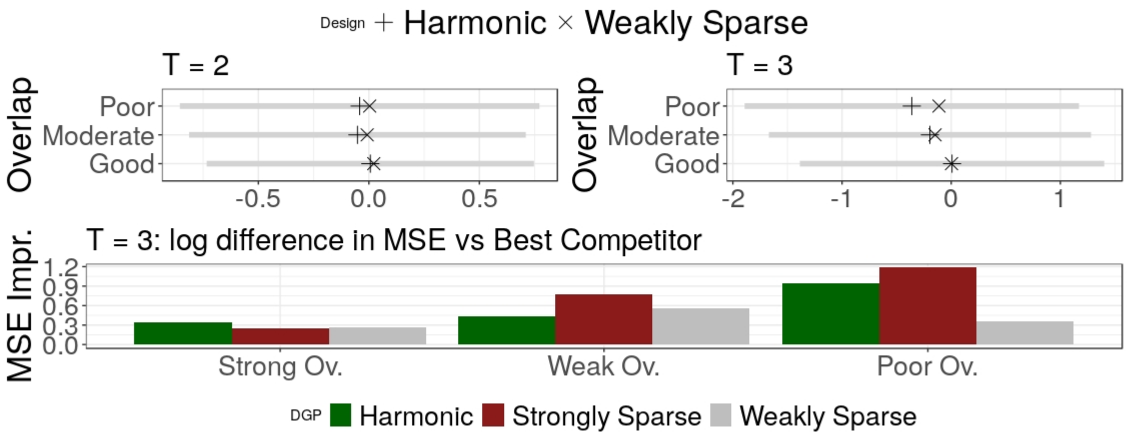

Figure 4: Top panels collect the point estimate (crosses), minus the true effect of the treatment, and confidence intervals of DCB for $p = 100$ across the three different designs. The bottom panel reports the decrease in MSE (in logarithmic scale) of the proposed method compared to the *best* competitor (excluding the one with known propensity score) for $T = 3$.

$\mu(1,1)$ and $\mu(0,0)$ are estimated from approximately a proportional sample. In the third panel, we report coverage when instead a Gaussian critical quantile (instead of the square root of a chi-squared quantile discussed in our theorems) is used. We observe that our procedure can lead to correct (over) coverage, while the Gaussian critical quantile leads to under-coverage in the presence of poor overlap and many variables, but correct coverage with fewer variables and two periods only.

Finally, we compare DCB and AIPW with high dimensional covariates with a longer time period. Namely, in Figure 5, we collect results for $T \in \{1, \cdots, 10\}$, where $T \leq 10$ is chosen for computational constraints. We generate data using a sparse model, $p = 100, n = 400$ over two-hundred replications. The outcome at time t depends on the contemporaneous treatment, covariates, and previous outcome at time $t-1$. To simulate a scenario where a strong correlation occurs between treatments over a long time period (similarly to applications in Figure 6), we generate

$$\mathbb{E}[D_{i,t}|D_{t-1}, X_t] = (1-\alpha)D_{i,t-1} + \alpha(1 + e^{\iota_{i,t}})^{-1},$$

where similarly to the propensity score model Equation (15), with $\iota_{i,t} = \frac{\eta}{\alpha} X_{i,t-1}\phi + \frac{\eta}{\alpha} X_{i,t}\phi + \frac{1}{2}(D_{i,t-1} - \bar{D}_{t-1}) + \xi_{i,t}, \xi_{i,t} \sim \mathcal{N}(0,1)$. Here η controls overlap together with α, where η/α has a similar role of the overlap constant in previous simulations.[21] In the figure we report results for $\alpha \in \{0.9, 0.7, 0.5\}$ (denoted as "High, Medium and Low correlation" respectively), and $\eta \in \{0.3, 0.5\}$. In Figure 5, we observe that for very strong time dependence between treatments (i.e., there are limited or no dynamics in assignments) the two methods are comparable. When instead, there are relatively more dynamics in treatment assignments the proposed method significantly improves in mean-squared error, with larger improvements in the presence of poorer overlap. In the Appendix, Figure E.1 we provide results also for very good overlap ($\eta = 0.1$), where the methods mostly provide comparable results on average. Such results illustrate the benefits of the method with dynamics in treatment assignments or poor overlap.

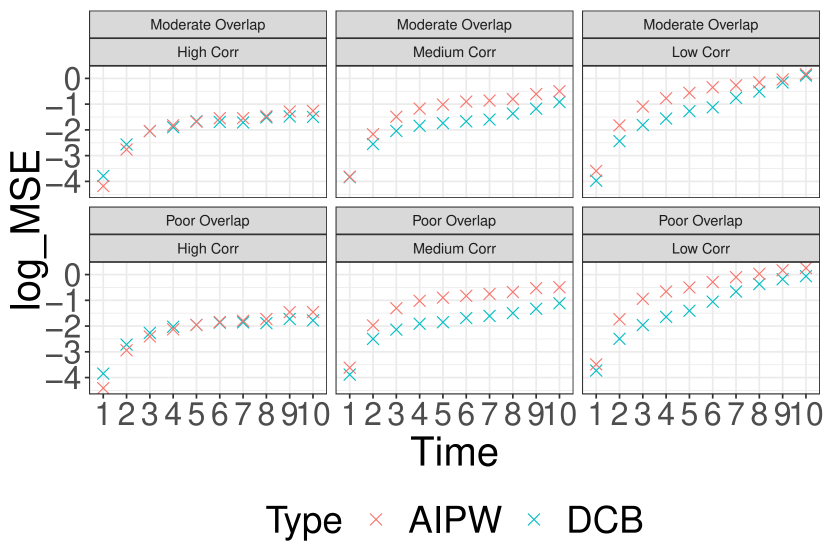

Figure 5: Mean-squared error in log-scale. Simulations for $T \leq 10, p = 100, n = 400$, two-hundred replications. Here high-correlation denotes strong serial depedence between treatment assignments with $\alpha = 0.9$, medium with $\alpha = 0.7$ and weak with $\alpha = 0.5$. $\eta \in \{0.3, 0.5\}$ for moderate and poor overlap, respectively.

[21]We take η/α as this plays approximately the same role of η in previous simulations from a simple linear approximation of $(1 + e^{\iota_{i,t}})^{-1}$ with respect to $\eta \approx 0$.

6 Empirical applications

6.1 The effect of negative advertisement on election outcome

Here, we study the effect of negative advertisements on the election outcome of democratic candidates. We use data from Blackwell (2013) who collects information on advertisement weeks before elections held in 2000, 2002, 2004, 2006.[22] There were 176 races during this period. We select a subsample of 148 races, removing the non-competitive races as in Blackwell (2013). Each race is associated with a different democratic candidate and a set of baseline and time-varying covariates. Negative advertisement is indicated by a binary variable as discussed in Blackwell (2013).[23]

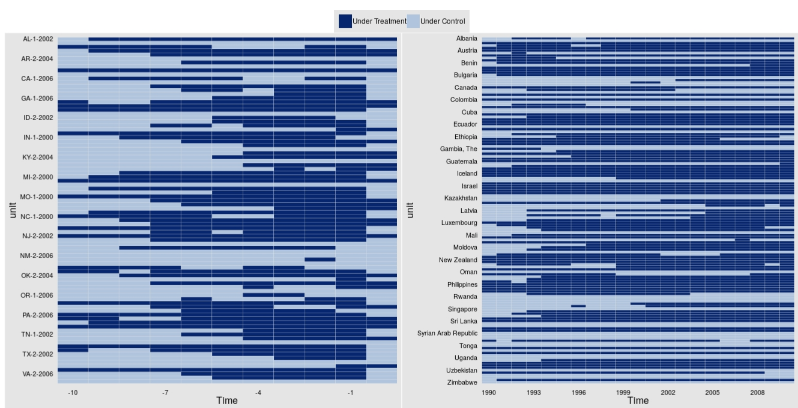

Figure 6: The figure illustrates the dynamics of treatment assignments for each application. The left-hand side is on the negative advertisement on the election outcome, and the right-hand side on the democratization on GDP.

As shown in Figure 6 (left-panel), each week, races may or may not "go negative" with treatment assignments exhibiting correlation in time. Hence, controlling for time-varying covariates and past assignments is crucial to avoid confounding. In a first model (Case 1), we control for the share of undecided voters in the previous week, whether the candidate is incumbent, the democratic polls, and whether the

[22] Data is available at https://dataverse.harvard.edu/dataset.xhtml?persistentId=hdl:1902.1/19801.

[23] This indicates whether at least ten percent of the negative advertisement.

democratic went negative in the previous week. Each of these variables (including treatment assignments) enters linearly in the regression. In Figure 7 we compare imbalance in covariates between the IPW weights estimated via logistic regression and the DCB weights for Case 1. We observe that imbalance is substantially smaller with the proposed weights, particularly for the share of undecided voters and the polls. The only exception is the left-bottom panel (second period's covariates for control individuals) where imbalance is approximately zero for both methods (the magnitude is 10^{-4} for this case). In Table 4 we collect results that demonstrate the negative effects of going negative for two consecutive periods. We also observe negative effects, albeit of smaller magnitude, when implementing a second specification (Case 2), which controls for a larger set of covariates.[24] When comparing to AIPW, we observe that DCB has a standard error twice as small as AIPW and larger point estimates in magnitude.

Table 4: The first row corresponds to Case 1, while the second row to Case 2. (A)-IPW refers to (Augmented)-inverse probability weights with stabilized weights. In parenthesis, the standard errors.

	ATE DCB	ATE AIPW
Case 1	-1.767	-0.57
	(0.70)	(1.45)
Case 2	-0.493	-0.22
	(0.764)	(1.33)

6.2 Effect of democracy on economic growth

Here, we revisit the study of Acemoglu et al. (2019) on the effects of democracy on economic growth.[25] The data consist of an extensive collection of countries observed

[24] These are campaign length, the baseline number of undecided voters, baseline share of democratic voters, assignments two periods before, the type of office, and year fixed effects.

[25] In the past, Mulligan et al. (2004) find no significant effect, while Giavazzi and Tabellini (2005) and Acemoglu et al. (2019) find positive and significant effects.

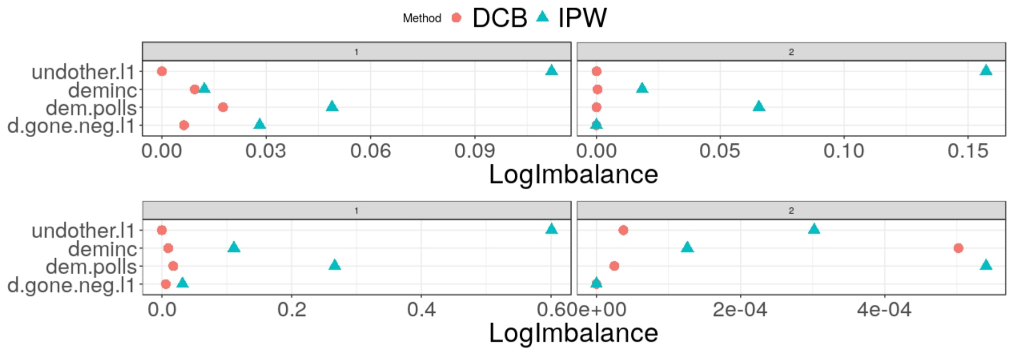

Figure 7: Effect of negative advertisement on election outcome: imbalance plot. Covariates are the share of undecided voters, whether the democratic candidate is incumbent, the democratic polls, and the treatment in the previous period. At the top, we report the imbalance on the treated and on the controls at the bottom. On the left panel, we illustrate the imbalance in the first period and on the right in the second period.

between 1960 and 2010.[26] We consider observations starting from 1989. After removing missing values, we run regressions with 141 countries. The outcome is the log-GDP in country i in period t as discussed in Acemoglu et al. (2019). Following Acemoglu et al. (2019) we capture democracy with a binary treatment based on international ranking. Studying the long-run impact of democracy has two challenges: (i) GDP growth depends on a long treatment history; (ii) unconfoundeness might hold only when conditioning on a large set of covariates and past outcomes.

For each country, we condition on lag outcomes in the past four years, following Acemoglu et al. (2019), past four treatment assignments which enter linearly in the regression (for example, for $t = 2010$ we condition on the outcomes in the four years before 2010, and similarly for every $t \geq 1989$). We consider a pooled regression (see Remark 4) and two alternative specifications. The first is parsimonious and include dummies for different regions and different intercepts for different periods.[27] A second one includes a larger set of covariates (in total 235 covariates). Coefficients are estimated with a penalized linear regression as described in Algorithm 2 (with

[26]Data available at https://www.journals.uchicago.edu/doi/suppl/10.1086/700936.
[27]Country specific fixed effects are not included.

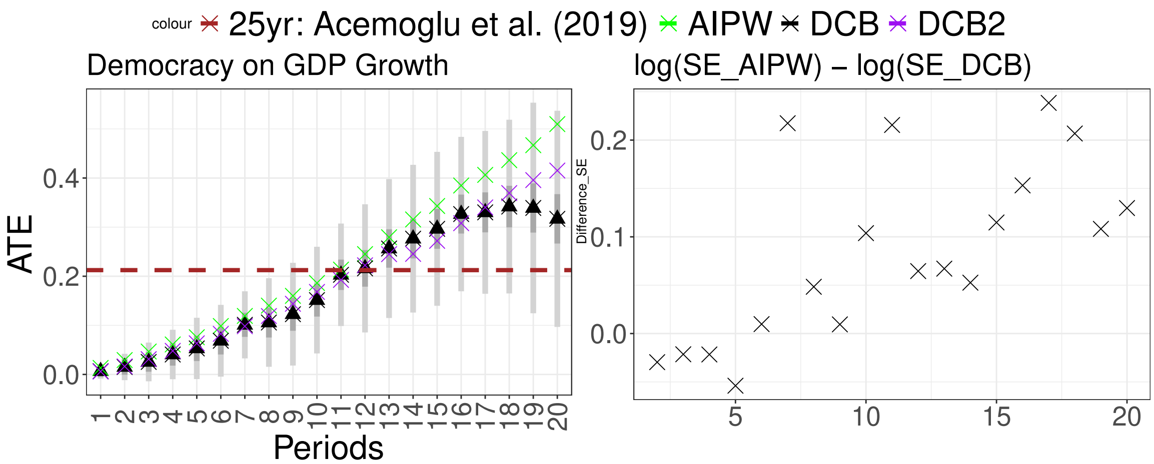

Figure 8: Left-hand side: pooled regression from $t \in \{1989, \cdots, 2010\}$ Gray region denotes the 90% confidence band for the least parsimonious model, with light-gray corresponding to the $\sqrt{\chi_{2T}(\alpha)}$ critical quantile, and darker area to the Gaussian critical quantile. DCB and DCB2 refer to two separate specification, with DCB corresponding to the more parsimonious one. The dotted line reports the effect after twenty five years of democracy discussed in Acemoglu et al. (2019). Right-hand side: log-difference in of standard error between AIPW and DCB for $t \in \{2, \cdots, 20\}$.

`model` = `linear`).[28] Tuning parameters are chosed as in the Appendix.

The estimand of interest is the t-long run effect of democracy.[29] It represents the effect of the past t consecutive years of democracy. In Figure 8 (left-panel) we collect our results, for endline outcomes pooled across 1989 to 2010. As t increases, the darker gray region (confidence intervals obtained with Gaussian quantiles) stays constant or slightly larger. The whiter region (chi-squared quantiles, which require weaker assumptions but can be conservative) becomes larger as T increases since the critical value also depends on the period length. Democracy has a statistically insignificant effect on the first years of GDP growth but a statistically significant positive impact on long-run GDP growth. The two specifications present similar results, showing the robustness of the results. Figure 8 illustrates the flexibility of

[28] In practice, we may also want to penalize only some of the coefficients and not others, which is also allowed in our framework.

[29] Formally $\mathbb{E}\Big[Y_{i,t}(\mathbf{1}_s, D_{i,(t-s):(-\infty)})\Big] - \mathbb{E}\Big[Y_{i,t}(\mathbf{0}_s, D_{i,(t-s):(-\infty)})\Big]$.

the method in capturing the effects of policies that are possibly non-linear in the exposure length. In the right-side panel we report the log-difference in standard errors between the proposed method and AIPW for $t \in \{2, \cdots, 20\}$. We note that for short periods, AIPW performs similarly or slightly outperforms (by less than 0.5 percentage points) DCB. For $t \geq 5$ improvements in standard errors of DCB upon AIPW are up to twenty percentage points, and with a positive trend in T.

7 Discussion

This paper discusses the problem of inference on dynamic treatments via covariate balancing. We allow for high-dimensional covariates, and we introduce novel balancing conditions that allow for the optimal $\sqrt{n}$-consistent estimation. The proposed method relies on computationally efficient estimators. Simulations and empirical applications illustrate its advantages over state-of-the-art methodologies.

Several questions remain open. First, the asymptotic properties crucially rely on cross-sectional independence while allowing for general dependence over time. A natural extension is where individuals exhibit dependence within clusters, which can be accommodated by our method with minor modifications. More broadly, future work should address more general extensions where cross-sectional *i.i.d.*-ness does not necessarily hold. Second, our asymptotic results assume a fixed period. This is an extension for future research, where the period is allowed to grow with the sample size. Third, our derivations impose a weak form of overlap when constructing balancing weights. A natural avenue for future research is whether conditions on overlap might be replaced by alternative (weaker) assumptions. One additional question, which cannot be directly handled in our framework, is how to handle dropouts during the study due to confounding variables.

Finally, the derivation of general balancing conditions which do not rely on a particular model specification remains an open research question.

References

Abadie, A. and M. D. Cattaneo (2018). Econometric methods for program evaluation. *Annual Review of Economics 10*, 465–503.

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

Abraham, S. and L. Sun (2018). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Available at SSRN 3158747*.

Acemoglu, D., S. Naidu, P. Restrepo, and J. A. Robinson (2019). Democracy does cause growth. *Journal of Political Economy 127*(1), 47–100.

Angrist, J. D., Ò. Jordà, and G. M. Kuersteiner (2018). Semiparametric estimates of monetary policy effects: string theory revisited. *Journal of Business & Economic Statistics 36*(3), 371–387.

Angrist, J. D. and G. M. Kuersteiner (2011). Causal effects of monetary shocks: Semiparametric conditional independence tests with a multinomial propensity score. *Review of Economics and Statistics 93*(3), 725–747.

Arellano, M. and S. Bonhomme (2011). Nonlinear panel data analysis. *Annu. Rev. Econ. 3*(1), 395–424.

Arkhangelsky, D., S. Athey, D. A. Hirshberg, G. W. Imbens, and S. Wager (2021). Synthetic difference-in-differences. *American Economic Review 111*(12), 4088–4118.

Arkhangelsky, D. and G. Imbens (2019). Double-robust identification for causal paneldata models. *arXiv preprint arXiv:1909.09412*.

Athey, S. and G. W. Imbens (2022). Design-based analysis in difference-in-differences settings with staggered adoption. *Journal of Econometrics 226*(1), 62–79.

Athey, S., G. W. Imbens, and S. Wager (2018). Approximate residual balancing: debiased inference of average treatment effects in high dimensions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology) 80*(4), 597–623.

Bang, H. and J. M. Robins (2005). Doubly robust estimation in missing data and causal inference models. *Biometrics 61*(4), 962–973.

Belloni, A., V. Chernozhukov, and C. Hansen (2014). Inference on treatment effects after selection among high-dimensional controls. *The Review of Economic Studies 81*(2), 608–650.

Belloni, A., V. Chernozhukov, C. Hansen, and D. Kozbur (2016). Inference in high-dimensional panel models with an application to gun control. *Journal of Business & Economic Statistics 34*(4), 590–605.

Ben-Michael, E., A. Feller, and J. Rothstein (2018). The augmented synthetic control method. *arXiv preprint arXiv:1811.04170*.

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

Blackwell, M. (2013). A framework for dynamic causal inference in political science. *American Journal of Political Science 57*(2), 504–520.

Bodory, H., M. Huber, and L. Lafférs (2020). Evaluating (weighted) dynamic treatment effects by double machine learning. *arXiv preprint arXiv:2012.00370*.

Bojinov, I., A. Rambachan, and N. Shephard (2020). Panel experiments and dynamic causal effects: A finite population perspective. *arXiv preprint arXiv:2003.09915*.

Bojinov, I. and N. Shephard (2019). Time series experiments and causal estimands: exact randomization tests and trading. *Journal of the American Statistical Association*, 1–36.

Boruvka, A., D. Almirall, K. Witkiewitz, and S. A. Murphy (2018). Assessing time-varying causal effect moderation in mobile health. *Journal of the American Statistical Association 113*(523), 1112–1121.

Bühlmann, P. and S. Van De Geer (2011). *Statistics for high-dimensional data: methods, theory and applications*. Springer Science & Business Media.

Callaway, B. and P. H. Sant'Anna (2019). Difference-in-differences with multiple time periods. *Available at SSRN 3148250*.

Chernozhukov, V., M. Goldman, V. Semenova, and M. Taddy (2017). Orthogonal machine learning for demand estimation: High dimensional causal inference in dynamic panels. *arXiv preprint arXiv:1712.09988*.

Chernozhukov, V., C. Hansen, Y. Liao, and Y. Zhu (2018). Inference for heterogeneous effects using low-rank estimations. *arXiv preprint arXiv:1812.08089*.

de Chaisemartin, C. and X. d'Haultfoeuille (2019). Two-way fixed effects estimators with heterogeneous treatment effects. Technical report, National Bureau of Economic Research.

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

Farrell, M. H. (2015). Robust inference on average treatment effects with possibly more covariates than observations. *Journal of Econometrics 189*(1), 1–23.

Finkelstein, A., S. Taubman, B. Wright, M. Bernstein, J. Gruber, J. P. Newhouse, H. Allen, K. Baicker, and O. H. S. Group (2012). The oregon health insurance experiment: evidence from the first year. *The Quarterly journal of economics 127*(3), 1057–1106.

Ghanem, D., P. H. Sant'Anna, and K. Wüthrich (2022). Selection and parallel trends. *arXiv preprint arXiv:2203.09001*.

Giavazzi, F. and G. Tabellini (2005). Economic and political liberalizations. *Journal of monetary economics 52*(7), 1297–1330.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*.

Hainmueller, J. (2012). Entropy balancing for causal effects: A multivariate reweighting method to produce balanced samples in observational studies. *Political Analysis 20*(1), 25–46.

Heckman, J. J., J. E. Humphries, and G. Veramendi (2016). Dynamic treatment effects. *Journal of econometrics 191*(2), 276–292.

Heckman, J. J. and S. Navarro (2007). Dynamic discrete choice and dynamic treatment effects. *Journal of Econometrics 136*(2), 341–396.

Hernán, M. A., B. Brumback, and J. M. Robins (2001). Marginal structural models to estimate the joint causal effect of nonrandomized treatments. *Journal of the American Statistical Association 96*(454), 440–448.

Imai, K. and I. S. Kim (2016). When should we use linear fixed effects regression models for causal inference with longitudinal data? *Unpublished working paper. Retrieved September 19*.

Imai, K., I. S. Kim, and E. Wang (2018). Matching methods for causal inference with time-series cross-section data.

Imai, K. and M. Ratkovic (2014). Covariate balancing propensity score. *Journal of the Royal Statistical Society: Series B (Statistical Methodology) 76*(1), 243–263.

Imai, K. and M. Ratkovic (2015). Robust estimation of inverse probability weights for marginal structural models. *Journal of the American Statistical Association 110*(511), 1013–1023.

Imbens, G. W. (2000). The role of the propensity score in estimating dose-response functions. *Biometrika 87*(3), 706–710.

Imbens, G. W. and D. B. Rubin (2015). *Causal inference in statistics, social, and biomedical sciences*. Cambridge University Press.

Jiang, N. and L. Li (2015). Doubly robust off-policy value evaluation for reinforcement learning. *arXiv preprint arXiv:1511.03722*.

Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American economic review 95*(1), 161–182.

Kallus, N. and M. Santacatterina (2018). Optimal balancing of time-dependent confounders for marginal structural models. *arXiv preprint arXiv:1806.01083*.

Kock, A. B. and H. Tang (2015). Inference in high-dimensional dynamic panel data models. *arXiv preprint arXiv:1501.00478*.

Krampe, J., E. Paparoditis, and C. Trenkler (2020). Impulse response analysis for sparse high-dimensional time series. *arXiv preprint arXiv:2007.15535*.

Lewis, G. and V. Syrgkanis (2020). Double/debiased machine learning for dynamic treatment effects via g-estimation. *arXiv preprint arXiv:2002.07285*.

Li, F., K. L. Morgan, and A. M. Zaslavsky (2018). Balancing covariates via propensity score weighting. *Journal of the American Statistical Association 113*(521), 390–400.

Mulligan, C. B., R. Gil, and X. Sala-i Martin (2004). Do democracies have different public policies than nondemocracies? *Journal of Economic Perspectives 18*(1), 51–74.

Negahban, S. N., P. Ravikumar, M. J. Wainwright, B. Yu, et al. (2012). A unified framework for high-dimensional analysis of m-estimators with decomposable regularizers. *Statistical science 27*(4), 538–557.

Nie, X., E. Brunskill, and S. Wager (2021). Learning when-to-treat policies. *Journal of the American Statistical Association 116*(533), 392–409.

Ning, Y., S. Peng, and K. Imai (2017). High dimensional propensity score estimation via covariate balancing.

Plagborg-Møller, M. (2019). Bayesian inference on structural impulse response functions. *Quantitative Economics 10*(1), 145–184.

Rambachan, A. and N. Shephard (2019). A nonparametric dynamic causal model for macroeconometrics. *Available at SSRN 3345325*.

Robins, J. (1986). A new approach to causal inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect. *Mathematical modelling 7*(9-12), 1393–1512.

Robins, J. M., M. A. Hernan, and B. Brumback (2000). Marginal structural models and causal inference in epidemiology.

Robins, J. M., A. Rotnitzky, and L. P. Zhao (1994). Estimation of regression coefficients when some regressors are not always observed. *Journal of the American statistical Association 89*(427), 846–866.

Rubin, D. B. (1990). Formal mode of statistical inference for causal effects. *Journal of statistical planning and inference 25*(3), 279–292.

Shi, C., A. Fan, R. Song, and W. Lu (2018). High-dimensional a-learning for optimal dynamic treatment regimes. *Annals of statistics 46*(3), 925–957.

Stock, J. H. and M. W. Watson (2018). Identification and estimation of dynamic causal effects in macroeconomics using external instruments. *The Economic Journal 128*(610), 917–948.

Tchetgen, E. J. T. and I. Shpitser (2012). Semiparametric theory for causal mediation analysis: efficiency bounds, multiple robustness, and sensitivity analysis. *Annals of statistics 40*(3), 1816.

Tran, L., C. Yiannoutsos, K. Wools-Kaloustian, A. Siika, M. Van Der Laan, and M. Petersen (2019). Double robust efficient estimators of longitudinal treatment effects: Comparative performance in simulations and a case study. *The international journal of biostatistics 15*(2).

Vansteelandt, S., M. Joffe, et al. (2014). Structural nested models and g-estimation: the partially realized promise. *Statistical Science 29*(4), 707–731.

White, H. and X. Lu (2010). Granger causality and dynamic structural systems. *Journal of Financial Econometrics 8*(2), 193–243.

Yiu, S. and L. Su (2018). Covariate association eliminating weights: a unified weighting framework for causal effect estimation. *Biometrika 105*(3), 709–722.

Yiu, S. and L. Su (2020). Joint calibrated estimation of inverse probability of treatment and censoring weights for marginal structural models. *Biometrics*.

Zhang, B., A. A. Tsiatis, E. B. Laber, and M. Davidian (2013). Robust estimation of optimal dynamic treatment regimes for sequential treatment decisions. *Biometrika 100*(3), 681–694.

Zhou, X. and G. T. Wodtke (2018). Residual balancing weights for marginal structural models: with application to analyses of time-varying treatments and causal mediation. *arXiv preprint arXiv:1807.10869*.

Zhu, Y. (2017). High-dimensional panel data with time heterogeneity: estimation and inference. *Available at SSRN 2665374*.

Zubizarreta, J. R. (2015). Stable weights that balance covariates for estimation with incomplete outcome data. *Journal of the American Statistical Association 110*(511), 910–922.

Appendix to "Dynamic covariate balancing: estimating treatment effects over time"

A Definitions

Throughout our discussion, we denote $y \lesssim x$ if the left-hand side is less or equal to the right-hand side up to a multiplicative constant term. We will refer to β^t as $\beta^t_{d_{1:T}}$ whenever clear from the context. Recall that when we omit the script i, we refer to the vector of all observations. We define

$$\nu_{i,t} = H_{i,t+1}\beta^{t+1}_{d_{1;T}} - H_{i,t}\beta^t_{d_{1:T}}, \quad \varepsilon_{i,T} = Y_{i,T}(d_{1:T}) - H_{i,T}\beta^T_{d_{1:T}}.$$

and $\hat{\nu}_{i,t}$ for estimated coefficients (omitting the argument $(d_{1:T})$ for notational convenience). In addition, we define $V_{i,t} = (Y_{i,1}, \cdots, Y_{i,t-1}, X_{i,1}, \cdots, X_{i,t})$ the vector of observations without including the treatment assignments.

B Lemmas

B.1 Lemmas in Section 3 and 4

B.1.1 Proof of Lemma 2.1

The first equation is a direct consequence of condition (A) in Assumption 2, and the linear model assumption. Consider the second equation. By condition (B) in Assumption 2, we have

$$\mathbb{E}\Big[Y_{i,2}(d_1, d_2)\Big|X_{i,1}\Big] = \mathbb{E}\Big[Y_{i,2}(d_1, d_2)\Big|X_{i,1}, D_{i,1} = d_1\Big].$$

Using the law of iterated expectations (since $X_{i,1} \subseteq H_{i,2}$)

$$\mathbb{E}\Big[Y_{i,2}(d_1, d_2)\Big|X_{i,1}, D_{i,1} = d_1\Big] = \mathbb{E}\Big[\mathbb{E}[Y_{i,2}(d_1, d_2)|H_{i,2}, D_{i,1} = d_1]|X_{i,1}, D_{i,1} = d_1\Big].$$

Using condition (A) in Assumption 2, we have

$$\mathbb{E}[Y_{i,2}(d_1, d_2)|H_{i,2}, D_{i,1} = d_1] = \mathbb{E}[Y_{i,2}(d_1, d_2)|H_{i,2}, D_{i,1} = d_1, D_{i,2} = d_2]$$

the proof completes as $\mathbb{E}[Y_{i,2}(d_1, d_2)|H_{i,2}, D_{i,1} = d_1, D_{i,2} = d_2] = \mathbb{E}[Y_{i,2}|H_{i,2}, D_{i,1} = d_1, D_{i,2} = d_2]$ as a consequence of condition (A) in Assumption 2.

B.1.2 Sufficient conditions for lasso

Lemma B.1 (Sufficient conditions for Lasso)**.** *Suppose that H_2, X_1 are uniformly bounded and $||\beta^2_{d_{1:2}}||_0, ||\beta^1_{d_{1:2}}||_0 \leq s, ||\beta^2_{d_{1:2}}||_\infty, ||\beta^1_{d_{1:2}}||_\infty < \infty$. Suppose that H_2, X_1 both satisfy the restricted eigenvalue assumption, and the column normalization condition (Negahban et al., 2012).*[30] *Suppose that $\hat{\beta}^1_{d_{1:2}}, \hat{\beta}^2_{d_{1:2}}$ are estimated with Lasso as in Algorithm 2 with a full interaction model and with penalty parameter $\lambda_n \asymp s\sqrt{\log(p)/n}$. Let Assumptions 1 - 4 hold. Let $\varepsilon_2(d_{1:2})|H_2$ be subgaussian almost surely and $\nu_1(d_1)|X_1$ be sub-gaussian almost surely. Then for each $t \in \{1, 2\}$,*

$$\Big|\Big|\hat{\beta}^t_{d_{1:2}} - \beta^t_{d_{1:2}}\Big|\Big|_1 = \mathcal{O}_p\Big(s^2\sqrt{\log(p)/n}\Big).$$

Therefore,

$$||\hat{\beta}^t_{d_{1:2}} - \beta^t_{d_{1:2}}||_1 \delta_t(n, p) = o_p(1/\sqrt{n}),$$

for $\delta_t(n, p) \asymp \log(np)/n^{1/4}$ and $s^2 \log^{3/2}(np)/n^{1/4} = \mathcal{O}(1)$.

□

The proof is discussed below and follows similarly to Negahban et al. (2012), with minor modifications. The above result provides a set of sufficient conditions such that Assumption 5 holds for a feasible choice of δ_t.

Proof. The result for

$$\Big|\Big|\hat{\beta}^2_{d_{1:2}} - \beta^2_{d_{1:2}}\Big|\Big|_1 = \mathcal{O}_p\Big(s\sqrt{\log(p)/n}\Big)$$

[30] Sufficient conditions that guarantee that the restricted eigenvalue assumption holds are discussed in (Negahban et al., 2012).

follows verbatim from Negahban et al. (2012) Corollary 2. For the result for $\hat{\beta}^1_{d_{1:2}}$ it suffices to notice, following the same argument from Negahban et al. (2012) (Corollary 2), that

$$\Big|\Big|\hat{\beta}^1_{d_{1:2}} - \beta^1_{d_{1:2}}\Big|\Big|_1 = O(s\lambda_n), \text{ for } \lambda_n \geq \Big|\Big|\frac{1}{n}X_1^\top \hat{\nu}_1\Big|\Big|_\infty,$$

since here we used the estimated outcome $H_2\hat{\beta}^2_{d_{1:T}}$ as the outcome of interest in our estimated regression instead of the true outcome.[31] The upper bound as a function of λ_n follows directly from Theorem 1 in Negahban et al. (2012).[32] We note that we can write

$$\begin{aligned}\Big|\Big|\frac{1}{n}X_1^\top \hat{\nu}_1\Big|\Big|_\infty &\leq \Big|\Big|\frac{1}{n}X_1^\top \nu_1\Big|\Big|_\infty + \Big|\Big|\frac{1}{n}X_1^\top (\nu_1 - \hat{\nu}_1)\Big|\Big|_\infty \\ &= \Big|\Big|\frac{1}{n}X_1^\top \nu_1\Big|\Big|_\infty + \Big|\Big|\frac{1}{n}X_1^\top H_2(\beta^2_{d_{1:2}} - \hat{\beta}^2_{d_{1:2}})\Big|\Big|_\infty \\ &\leq \Big|\Big|\frac{1}{n}X_1^\top \nu_1\Big|\Big|_\infty + ||X_1||_\infty ||H_2||_\infty ||\beta^2_{d_{1:2}} - \hat{\beta}^2_{d_{1:2}}||_1.\end{aligned}$$

We now study each component separately. By sub-gaussianity, since $\mathbb{E}[\nu_1|X_1] = 0$ by Assumption 3, we have for all $t > 0$, by Hoeffding inequality and the union bound,

$$P\Big(\Big|\Big|\frac{1}{n}X_1^\top \nu_1\Big|\Big|_\infty > t\Big|X_1\Big) \leq p\exp\Big(-M\frac{t^2 n}{s}\Big)$$

for a finite constant M. This result follows since $\nu_1 \leq ||\beta_1||_1 ||X_1^{(j)}||_\infty \leq Ms$. It implies that

$$\Big|\Big|\frac{1}{n}X_1^\top \nu_1\Big|\Big|_\infty = O_p(\sqrt{s\log(p)/n})$$

The second component instead is $O_p(s\sqrt{\log(p)/n})$ by the bound on $||\beta^2_{d_{1:2}} - \hat{\beta}^2_{d_{1:2}}||_1$. This complete the proof. Finally, observe also that the same argument follows recursively for any finite T, with the estimation error depending on T. □

[31]Formally, here to compute $\mathcal{R}^*(\nabla\mathcal{L}(\theta^*))$ in Negahban et al. (2012)'s notation we need to account for the loss function to depend on the estimated outcome.

[32]Note that Theorem 1 in Negahban et al. (2012) does not depend on the distribution of the data and is a deterministic statement which holds under strong convexity at the true regression parameter. For a linear model, strong convexity is satisfied under the restricted eigenvalue assumption which does not depend on the regression parameter.

B.1.3 Proof of Lemma 4.1

Since the Lemmas in Section 3 are a special case of those in Section 4 we directly prove the results for multiple periods.

Throughout the proof we omit the argument $d_{1:T}$ of $\hat{\gamma}_t(d_{1:T})$ for notational convenience. Recall that $\hat{\gamma}_{i,T} = 0$ if $D_{i,1:T} \neq d_{1:T}$. Therefore, by consistency of potential outcomes:

$$\hat{\gamma}_{i,T} Y_{i,T} = \hat{\gamma}_{i,T} Y_{i,T}(d_{1:T}) = \hat{\gamma}_{i,T}(H_{i,T}\beta^T_{d_{1:T}} + \varepsilon_{i,T}).$$

Then we can write

$$\begin{aligned}
&\sum_{i=1}^n \Big(\hat{\gamma}_{i,T} Y_{i,T} - \sum_{t=2}^T (\hat{\gamma}_{i,t} - \hat{\gamma}_{i,t-1}) H_{i,t}\hat{\beta}^t_{d_{1:T}} - (\hat{\gamma}_{i,1} - \frac{1}{n}) X_{i,1}\hat{\beta}^1_{d_{1:T}}\Big) \\
&= \sum_{i=1}^n \Big(\hat{\gamma}_{i,T} H_{i,T}\beta^T_{d_{1:T}} - \sum_{t=2}^T (\hat{\gamma}_{i,t} - \hat{\gamma}_{i,t-1}) H_{i,t}\hat{\beta}^t_{d_{1:T}} - (\hat{\gamma}_{i,1} - \frac{1}{n}) X_{i,1}\hat{\beta}^1_{d_{1:T}}\Big) + \hat{\gamma}_T^\top \varepsilon_T.
\end{aligned}$$

Consider first the term

$$\sum_{i=1}^n (\hat{\gamma}_{i,T} H_{i,T}\beta^T_{d_{1:T}} - (\hat{\gamma}_{i,T} - \hat{\gamma}_{i,T-1}) H_{i,T}\hat{\beta}^T_{d_{1:T}}) = (\hat{\gamma}_T H_T - \hat{\gamma}_{T-1} H_T)(\beta^T_{d_{1:T}} - \hat{\beta}^T_{d_{1:T}}) + \hat{\gamma}_{T-1} H_T \beta^T_{d_{1:T}}.$$

Notice now that for any $s > 1$,

$$\sum_{i=1}^n (\hat{\gamma}_{i,s} - \hat{\gamma}_{i,s-1}) H_{i,s}\hat{\beta}^s_{d_{1:T}} = (\hat{\gamma}_s H_s - \hat{\gamma}_{s-1} H_s)(\hat{\beta}^s_{d_{1:T}} - \beta^s_{d_{1:T}}) + \hat{\gamma}_s H_s \beta^s_{d_{1:s}} - \hat{\gamma}_{s-1} H_s \beta^s_{d_{1:s}}.$$

For $s = 1$ we have instead

$$\sum_{i=1}^n (\hat{\gamma}_{i,1} - \frac{1}{n}) X_{i,1}\hat{\beta}^1_{d_{1:T}} = (\hat{\gamma}_1 X_1 - \bar{X}_1)(\hat{\beta}^1_{d_{1:T}} - \beta^1_{d_{1:T}}) + \hat{\gamma}_1 X_1 \beta^1_{d_{1:s}} - \bar{X}_1 \beta^1_{d_{1:s}}.$$

Therefore, we can write

$$\sum_{i=1}^{n}\Big(\hat{\gamma}_{i,T}Y_{i,T} - \sum_{t=2}^{T}(\hat{\gamma}_{i,t} - \hat{\gamma}_{i,t-1})H_{i,t}\hat{\beta}^{t}_{d_{1:T}} - (\hat{\gamma}_{i,1} - \frac{1}{n})X_{i,1}\hat{\beta}^{1}_{d_{1:T}}\Big)$$

$$= (\hat{\gamma}_T H_T - \hat{\gamma}_{T-1}H_T)(\beta^{T}_{d_{1:T}} - \hat{\beta}^{T}_{d_{1:T}}) + \sum_{s=2}^{T-1}(\hat{\gamma}_s H_s - \hat{\gamma}_{s-1}H_s)(\beta^{s}_{d_{1:T}} - \hat{\beta}^{s}_{d_{1:T}}) + (\hat{\gamma}_1 X_1 - \bar{X}_1)(\beta^{1}_{d_{1:T}} - \hat{\beta}^{1}_{d_{1:T}})$$

$$+ \hat{\gamma}_T H_T \beta^{T}_{d_{1:T}} + \gamma_T^{\top}\varepsilon_T - \Big[\sum_{s=2}^{T-1}\hat{\gamma}_{s+1}H_s\beta^{s}_{d_{1:s}} - \hat{\gamma}_s H_s \beta^{s}_{d_{1:s}}\Big] - \hat{\gamma}_1 X_1 \beta^{1}_{d_{1:s}} + \bar{X}_1\beta^{1}_{d_{1:s}}.$$

The proof completes collecting the desired terms.

B.1.4 Proof of Lemma 4.2

Since $\hat{\gamma}_{i,t}(d_{1:T})$ is equal to zero if $D_{i,1:t} \neq d_{1:t}$ we can focus to the case where $D_{i,1:t} = d_{1:t}$. Since weights at time $t-1$ are measurable with respect to $\mathcal{F}_{t-1}, D_{t-1}$, we only need to show that

$$\mathbb{E}[\hat{\gamma}_{i,t-1}(d_{1:T})H_t\beta^{t}_{d_{1:T}}|\mathcal{F}_{t-1}, D_{-i,t-1}, D_{i,(1:(t-1))} = d_{1:(t-1)}] = \hat{\gamma}_{i,t}(d_{1:T})H_{t-1}\beta^{t-1}_{d_{1:T}}, \quad \text{(B.1)}$$

where, recall, weights $\hat{\gamma}_{i,t-1}(d_{1:T})$ are a measurable function of $\mathcal{F}_{t-1}, D_{t-1}$. On the event that $D_{i,(1:(t-1))} \neq d_{1:(t-1)}$ the expression is zero on both sides and the result trivially holds. Therefore, we can implicitely assume that $D_{i,(1:(t-1))} = d_{1:(t-1)}$ since otherwise the result trivially holds. Under Assumption 7 we can write

$$\begin{aligned}\mathbb{E}[\hat{\gamma}_{i,t-1}(d_{1:T})H_t\beta^{t}_{d_{1:T}}|\mathcal{F}_{t-1}, D_{t-1}] &= \mathbb{E}\Big[\hat{\gamma}_{i,t-1}(d_{1:T})\mathbb{E}[Y_{i,T}(d_{1:T})|\mathcal{F}_t, D_t]\Big|\mathcal{F}_{t-1}, D_{t-1}\Big] \\ &= \hat{\gamma}_{i,t-1}(d_{1:T})\mathbb{E}[Y_{i,T}(d_{1:T})|\mathcal{F}_{t-1}, D_{t-1}]\end{aligned} \quad \text{(B.2)}$$

by the tower property of the expectation and the definition of $\mathcal{F}_t$. Now notice that under Assumption 7, $\mathbb{E}[Y_{i,T}(d_{1:T})|\mathcal{F}_{t-1}, D_{t-1}] = \mathbb{E}[Y_{i,T}(d_{1:T})|\mathcal{F}_{t-1}]$. Therefore

$$\hat{\gamma}_{i,t-1}(d_{1:T})\mathbb{E}[Y_{i,T}(d_{1:T})|\mathcal{F}_{t-1}] = \hat{\gamma}_{i,t-1}(d_{1:T})H_{i,t-1}\beta^{t-1}_{d_{1:(t-1)}} \quad \text{(B.3)}$$

which follows since $\hat{\gamma}_{i,t-1}(d_{1:T}) = 0$ if $D_{1:t-1} \neq d_{1:t-1}$.

Corollary 2. *Lemma 3.2 holds.*

Proof. It follows directly choosing $t \in \{1,2\}$ from Lemma 4.2. $\square$

B.2 Additional auxiliary Lemmas

Lemma B.2. *(Existence of Feasible $\hat{\gamma}_1$) Suppose that $X_{i,1}^{(j)}$ is subgaussian for all $j \in \{1, \cdots, p_1\}$, $X_{i,1} \in \mathbb{R}^{p_1}$. Suppose that for $d_1 \in \{0,1\}$, $P(D_{i,1} = d_1|X_{i,1}) \in (\delta, 1-\delta)$. Then with probability $1 - 5/n$, for $\log(2np_1)/n \leq c_0$ for a constant $0 < c_0 < \infty$, where $\delta_1(n, p_1) \geq C\sqrt{2\log(2np_1)/n}$, for a constant $0 < C < \infty$, there exist a feasible $\hat{\gamma}_1$. In addition,*

$$\lim_{n\to\infty} P\Big(n||\hat{\gamma}_1||_2^2 \leq \mathbb{E}\Big[\frac{1}{P(D_{i,1} = d_1|X_{i,1})}\Big]\Big) = 1.$$

Proof of Lemma B.2. This proof follows in the same spirit of one-period setting (Athey et al., 2018). To prove existence of a feasible weight, we use a feasible guess. We prove the claim for a general $d_1 \in \{0,1\}$. Consider first

$$\hat{\gamma}_{i,1}^* = \frac{1\{D_{i,1} = d_1\}}{nP(D_{i,1} = d_1|X_{i,1})} \Big/ \Big(\frac{1}{n}\sum_{i=1}^n \frac{1\{D_{i,1} = d_1\}}{P(D_{i,1} = d_1|X_{i,1})}\Big). \tag{B.4}$$

For such weight to be well-defined, we need that the denominator is bounded away from zero. We now provide bounds on the denominator. Since $P(D_{i,1} = d_1|X_{i,1}) \in (\delta, 1-\delta)$ by Hoeffding inequality

$$P\Big(\Big|\frac{1}{n}\sum_{i=1}^n \frac{1\{D_{i,1} = d_1\}}{P(D_{i,1} = d_1|X_{i,1})} - 1\Big| > t\Big) \leq 2\exp\Big(-\frac{nt^2}{2a^2}\Big),$$

for a finite constant a. Therefore with probabiliy $1 - 1/n$,

$$\frac{1}{n}\sum_{i=1}^n \frac{1\{D_{i,1} = d_1\}}{P(D_{i,1} = d_1|X_{i,1})} > 1 - \sqrt{2a^2\log(2n)/n}. \tag{B.5}$$

Therefore for n large enough such that $\sqrt{2a^2\log(2n)/n} < 1 - \kappa$, weights are finite with high probability taking some arbitrary $\kappa \in (0,1)$. In addition, they sum up to one and they satisfy the requirement with probability $1 - 1/n$

$$\frac{1\{D_{i,1} = d_1\}}{nP(D_{i,1} = d_1 | X_{i,1})} \lesssim n^{-2/3} \Rightarrow \gamma^*_{i,1} \leq K_{2,1} n^{-2/3}$$

for a constant $K_{2,1}$, where the first inequality follows by the overlap assumption and the second by Equation (B.5). We are left to show that the first constraint is satisfied. First notice that under Assumption 7 $\mathbb{E}\Big[\frac{1}{n}\sum_{i=1}^n \frac{1\{D_{i,1}=d_1\}X_{i,1}^{(j)}}{P(D_{i,1}=1|X_{i,1})}|X_1\Big] = \bar{X}_1^{(j)}$. In addition, since $X_{i,1}$ is subgaussian, and $1/P(D_{i,1} = d_1|X_{i,1})$ is uniformly bounded, and the union bound

$$P\Big(\Big|\Big|\bar{X}_1 - \frac{1}{n}\sum_{i=1}^n \frac{1\{D_{i,1} = d_1\}}{P(D_{i,1} = 1|X_{i,1})} X_{i,1}\Big|\Big|_\infty > t\Big) \leq p_1 2 \exp\Big(-\frac{nt^2}{2a^2}\Big)$$

for a finite constant a^2. With trivial rearrangement, with probability $1 - 1/n$,

$$\Big|\Big|\bar{X}_1 - \frac{1}{n}\sum_{i=1}^n \frac{1\{D_{i,1} = d_1\}}{P(D_{i,1} = 1|X_{i,1})} X_{i,1}\Big|\Big|_\infty \leq a\sqrt{2\log(2np)/n} \tag{B.6}$$

Consider now the denominator. We have shown that the denominator concetrates around one at exponential rate, namely that with probability $1 - 1/n$,

$$\Big|\frac{1}{n}\sum_{i=1}^n \frac{1\{D_{i,1} = d_1\}}{P(D_{i,1} = d_1|X_{i,1})} - 1\Big| \leq 2a\sqrt{\log(2n)/n}. \tag{B.7}$$

Therefore, with probability $1-2/n$,

$$\Big|\Big|\bar{X}_1 - \frac{\frac{1}{n}\sum_{i=1}^n \frac{1\{D_{i,1}=d_1\}}{P(D_{i,1}=1|X_{i,1})}X_{i,1}}{\frac{1}{n}\sum_{i=1}^n \frac{1\{D_{i,1}=d_1\}}{P(D_{i,1}=1|X_{i,1})}}\Big|\Big|_\infty = \Big|\Big|\frac{\bar{X}_1\frac{1}{n}\sum_{i=1}^n \frac{1\{D_{i,1}=d_1\}}{P(D_{i,1}=1|X_{i,1})} - \frac{1}{n}\sum_{i=1}^n \frac{1\{D_{i,1}=d_1\}}{P(D_{i,1}=1|X_{i,1})}X_{i,1}}{\frac{1}{n}\sum_{i=1}^n \frac{1\{D_{i,1}=d_1\}}{P(D_{i,1}=1|X_{i,1})}}\Big|\Big|_\infty$$
$$= \Big|\Big|\frac{\bar{X}_1\frac{1}{n}\sum_{i=1}^n \frac{1\{D_{i,1}=d_1\}}{P(D_{i,1}=1|X_{i,1})} + \bar{X}_1 - \bar{X}_1 - \frac{1}{n}\sum_{i=1}^n \frac{1\{D_{i,1}=d_1\}}{P(D_{i,1}=1|X_{i,1})}X_{i,1}}{\frac{1}{n}\sum_{i=1}^n \frac{1\{D_{i,1}=d_1\}}{P(D_{i,1}=1|X_{i,1})}}\Big|\Big|_\infty$$
$$\le \Big|\Big|\frac{\bar{X}_1 - \frac{1}{n}\sum_{i=1}^n \frac{1\{D_{i,1}=d_1\}}{P(D_{i,1}=1|X_{i,1})}X_{i,1}}{\frac{1}{n}\sum_{i=1}^n \frac{1\{D_{i,1}=d_1\}}{P(D_{i,1}=1|X_{i,1})}}\Big|\Big|_\infty + \frac{2a\sqrt{\log(2n)/n}}{\frac{1}{n}\sum_{i=1}^n \frac{1\{D_{i,1}=d_1\}}{P(D_{i,1}=1|X_{i,1})}}$$
$$\le \frac{a\sqrt{2\log(2np)/n} + 2a\sqrt{\log(2n)/n}}{\frac{1}{n}\sum_{i=1}^n \frac{1\{D_{i,1}=d_1\}}{P(D_{i,1}=1|X_{i,1})}}, \tag{B.8}$$

where the first inequality follows by the triangular inequality and by concentration of the term $\frac{1}{n}\sum_{i=1}^n \frac{1\{D_{i,1}=d_1\}}{P(D_{i,1}=1|X_{i,1})}$ around one at exponential rate as in Equation (B.7). The second inequality follows by concentration of the numerator as in Equation (B.6). With probability $1-1/n$, the denominator is bounded away from zero. Therefore for a universal constant $C<\infty$,[33]

$$P\Big(\Big|\Big|\bar{X}_1 - \frac{\frac{1}{n}\sum_{i=1}^n \frac{1\{D_{i,1}=d_1\}}{P(D_{i,1}=1|X_{i,1})}X_{i,1}}{\frac{1}{n}\sum_{i=1}^n \frac{1\{D_{i,1}=d_1\}}{P(D_{i,1}=1|X_{i,1})}}\Big|\Big|_\infty \le Ca\sqrt{2\log(2np)/n}\Big) \ge 1-3/n. \tag{B.9}$$

We are left to provide bounds on $||\hat{\gamma}_1||_2^2$. For n large enough, with probability at least $1-5/n$, $||\hat{\gamma}_1||_2^2 \le ||\hat{\gamma}_1^*||_2^2$ since $\hat{\gamma}_1^*$ is a feasible solution. By overlap, the fourth moment of $1/P(D_{i,1}=d_1|X_{i,1})$ is bounded. By the strong law of large numbers and Slutsky theorem,

$$n||\hat{\gamma}_1^*||_2^2 = \sum_{i=1}^n \frac{1\{D_{i,1}=d_1\}}{nP(D_{i,1}=d_1|X_{i,1})^2}\Big/\Big(\sum_{i=1}^n \frac{1\{D_{i,1}=d_1\}}{nP(D_{i,1}=d_1|X_{i,1})}\Big)^2 \to_{as} \frac{\mathbb{E}[\frac{1\{D_{i,1}=d_1\}}{P(D_{i,1}=d_1|X_{i,1})^2}]}{\mathbb{E}[\frac{1\{D_{i,1}=d_1\}}{P(D_{i,1}=d_1|X_{i,1})}]^2} < \infty. \tag{B.10}$$

which completes the proof. □

[33] Here $3/n$ follows from the union bound.

Lemma B.3. *(Existence of a feasible $\hat{\gamma}_t$) Let*

$$Z_{i,t}(d_t) = \frac{1\{D_{i,t} = d_t\}}{P(D_{i,t} = d_t | Y_{i,1}, ..., Y_{i,t-1}, X_{i,1}, ..., X_{i,t-1}, D_{i,1}, ..., D_{i,t-1})}.$$

Assume that for $d_t \in \{0,1\}$. Assume that $H_{i,t}^{(j)} | H_{i,t-1}$ is sub-gaussian for all $j \in \{1, \cdots, p_t\}$ almost surely. Let Assumption 8 hold and let for a finite constant c_0,

$$\delta_t(n, p_t) \geq c_0 \frac{\log^{3/2}(p_t n)}{n^{1/2}}, \quad \text{and} \quad K_{2,t} = 2K_{2,t-1}\bar{c}, \quad \text{for some finite constant } \bar{c}.$$

Then with probability $\eta_n \to 1$, for some $N > 0$, $n \geq N$, there exists a feasible $\hat{\gamma}_t^*$ solving the optimization in Algorithm D.1, where

$$\hat{\gamma}_{i,t}^* = \hat{\gamma}_{i,t-1} Z_{i,t}(d_t) \Big/ \sum_{i=1}^n \hat{\gamma}_{i,t-1} Z_{i,t}(d_t)$$

In addition,

$$\lim_{n\to\infty} P\Big(n||\hat{\gamma}_t||_2^2 \leq C_t\Big) = 1 \tag{B.11}$$

for a constant $1 \leq C_t < \infty$ independent of (p_t, n).

Proof of Lemma B.3. The proof follows by induction. By Lemma B.2 we know that there exist a feasible $\hat{\gamma}_1$, with $\lim_{n\to\infty} P(n||\hat{\gamma}_1||_2^2 \leq C') = 1$. Suppose now that there exist feasible $\hat{\gamma}_1, ..., \hat{\gamma}_{t-1}$, such that

$$\lim_{n\to\infty} P(n||\hat{\gamma}_s||_2^2 \leq C_s) = 1 \tag{B.12}$$

for some finite constant C_s which only depends on s, and for all $s < t$. We want to show that the statement holds for $\hat{\gamma}_t$. We find γ_t^* that satisfies the constraint, with

$$\hat{\gamma}_{i,t}^* = \hat{\gamma}_{i,t-1} \frac{1\{D_{i,t} = d_t\}}{P(D_{i,t} = d_t | H_{i,t})} \Big/ \Big(\sum_{i=1}^n \hat{\gamma}_{i,t-1} \frac{1\{D_{i,t} = d_t\}}{P(D_{i,t} = d_t | H_{i,t})} \Big). \tag{B.13}$$

We break the proof into several steps.

Finite and Bounded Weights To show that such weights are finite, with high probability, we need to impose bounds on the numerator and the denominator. We want to bound for a universal constant $\bar{C} < \infty$,

$$P\Big(\Big\{\max_{i \in \{1,\dots,n\}} \hat{\gamma}_{i,t-1} \frac{1\{D_{i,t} = d_t\}}{P(D_{i,t} = d_t | H_{i,t})} > \bar{C} n^{-2/3} K_{2,t-1}\Big\} \bigcup \Big\{\sum_{i=1}^n \hat{\gamma}_{i,t} \frac{1\{D_{i,t} = d_t\}}{P(D_{i,t} = d_t | H_{i,t})} > \epsilon\Big\}\Big)$$
$$\leq \underbrace{P\Big(\max_{i \in \{1,\dots,n\}} \hat{\gamma}_{i,t-1} \frac{1\{D_{i,t} = d_t\}}{P(D_{i,t} = d_t | H_{i,t})} > \bar{C} n^{-2/3} K_{2,t-1}\Big)}_{(i)} + \underbrace{P\Big(\sum_{i=1}^n \hat{\gamma}_{i,t} \frac{1\{D_{i,t} = d_t\}}{P(D_{i,t} = d_t | H_{i,t})} > \epsilon\Big)}_{(ii)}.$$

We start by (i). Observe first that we can bound

$$\max_{i \in \{1,\dots,n\}} \hat{\gamma}_{i,t-1} \frac{1\{D_{i,t} = d_t\}}{P(D_{i,t} = d_t | H_{i,t})} \leq n^{-2/3} K_{2,t-1} \max_{i \in \{1,\dots,n\}} \frac{1\{D_{i,t} = d_t\}}{P(D_{i,t} = d_t | H_{i,t})} \leq K_{2,t-1} \bar{C} n^{-2/3}$$

for a finite constant $\bar{C}$. We now provide bounds on the denominator. Since $\sigma(H_{t-1}) \subseteq \sigma(H_t)$

$$\begin{aligned}\mathbb{E}\Big[\sum_{i=1}^n \hat{\gamma}_{i,t-1} \frac{1\{D_{i,t} = d_t\}}{P(D_{i,t} = d_t | H_{i,t})}\Big] &= \mathbb{E}\Big[\mathbb{E}\Big[\sum_{i=1}^n \hat{\gamma}_{i,t-1} \frac{1\{D_{i,t} = d_t\}}{P(D_{i,t} = d_t | H_{i,t})}\Big| H_{t-1}\Big]\Big] \\ &= \mathbb{E}\Big[\sum_{i=1}^n \hat{\gamma}_{i,t-1} \mathbb{E}\Big[\mathbb{E}\Big[\frac{1\{D_{i,t} = d_t\}}{P(D_{i,t} = d_t | H_{i,t})}\Big| H_t\Big]\Big| H_{t-1}\Big]\Big] = \sum_{i=1}^n \hat{\gamma}_{i,t-1} = 1.\end{aligned}$$

We show concentration of the denominator around its expectation to show that the denominator is bounded away from zero with high probability. Let C_{t-1} be the upper limit on $n||\hat{\gamma}_{t-1}||_2^2$, and let

$$c := 1\Big/C_{t-1} \quad \eta_{n,t} := P(||\hat{\gamma}_{t-1}||_2^2 \leq 1/(cn)), \tag{B.14}$$

for some constant c, which only depends on $t-1$ (the dependence with $t-1$ is suppressed for expositional convenience). Observe in addition that $\eta_{n,t} \to 1$ by the induction argument (see Equation (B.12)). We write for a finite constant a

$$
\begin{aligned}
&P\Big(\Big|\sum_{i=1}^n \hat{\gamma}_{i,t-1}\frac{1\{D_{i,t}=d_t\}}{P(D_{i,t}=d_t|H_{i,t})}-1\Big|>h\Big) \\
&\le P\Big(\Big|\sum_{i=1}^n \hat{\gamma}_{i,t-1}\frac{1\{D_{i,t}=d_t\}}{P(D_{i,t}=d_t|H_{i,t})}-1\Big|>h\Big|||\hat{\gamma}_{t-1}||_2^2\le 1/(cn)\Big)\eta_{n,t}+(1-\eta_{n,t}) \\
&\le 2\exp\Big(-\frac{ah^2}{2||\hat{\gamma}_{t-1}||_2^2}\Big|||\hat{\gamma}_{t-1}||_2^2\le 1/(cn)\Big)\eta_{n,t}+(1-\eta_{n,t}) \\
&\le 2\exp\Big(-\frac{ch^2an}{2}\Big)\eta_{n,t}+(1-\eta_{n,t}).
\end{aligned}
\tag{B.15}
$$

The third inequality follows from the fact that $\hat{\gamma}_{t-1}$ is measurable with respect to H_{t-1} and $\frac{1\{D_{i,t}=d_t\}}{P(D_{i,t}=d_t|H_{i,t})}$ is sub-gaussian conditional on $H_{i,t-1}$ (since uniformly bounded). Therefore with probability at least $1-\delta$,

$$
\Big|\sum_{i=1}^n \hat{\gamma}_{i,t-1}\frac{1\{D_{i,t}=d_t\}}{P(D_{i,t}=d_t|H_{i,t})}-1\Big|\le\sqrt{2\log(2\eta_{n,t}/(\delta+\eta_{n,t}-1))/(acn)}. \tag{B.16}
$$

By setting $\delta=\eta_{n,t}/n+(1-\eta_{n,t})$, with probability at least $1-\eta_{n,t}/n+(1-\eta_{n,t})$,

$$
\Big|\sum_{i=1}^n \hat{\gamma}_{i,t-1}\frac{1\{D_{i,t}=d_t\}}{P(D_{i,t}=d_t|H_{i,t})}-1\Big|\le\sqrt{2\log(2n)/acn},
$$

and hence the denominator is bounded away from zero for n large enough (recall that $\eta_{n,t}\to 1$).

First Constraint We now show that the proposed weights satisfy the first constraint in Algorithm D.1. The second trivially holds, while the third has been discussed in the first part of the proof. We write

$$\mathbb{E}\Big[\sum_{i=1}^n \hat\gamma_{i,t-1} H_{i,t}^{(j)} - \sum_{i=1}^n \hat\gamma_{i,t-1}\frac{1\{D_{i,t}=d_t\}}{P(D_{i,t}=d_t|H_{i,t})} H_{i,t}^{(j)}\Big]$$
$$= \mathbb{E}\Big[\mathbb{E}\Big[\sum_{i=1}^n \hat\gamma_{i,t-1} H_{i,t}^{(j)} - \sum_{i=1}^n \hat\gamma_{i,t-1}\frac{1\{D_{i,t}=d_t\}}{P(D_{i,t}=d_t|H_{i,t})} H_{i,t}^{(j)}\Big| H_t\Big]\Big] = 0.$$

We want to show concentration. First, we break the probability into two components:

$$P\Big(\Big|\Big|\sum_{i=1}^n \hat\gamma_{i,t-1} H_{i,t} - \sum_{i=1}^n \hat\gamma_{i,t-1}\frac{1\{D_{i,t}=d_t\}}{P(D_{i,t}=d_t|H_{i,t})} H_{i,t}\Big|\Big|_\infty > h\Big)$$
$$\leq \underbrace{P\Big(\Big|\Big|\sum_{i=1}^n \hat\gamma_{i,t-1} H_{i,t} - \sum_{i=1}^n \hat\gamma_{i,t-1}\frac{1\{D_{i,t}=d_t\}}{P(D_{i,t}=d_t|H_{i,t})} H_{i,t}\Big|\Big|_\infty > h\Big| ||\hat\gamma_{t-1}||_2^2 \leq 1/cn\Big)\eta_{n,t}}_{(I)} + \underbrace{(1-\eta_{n,t})}_{(II)},$$

where $\eta_{n,t} = P(||\hat\gamma_{t-1}||_2^2 \leq 1/cn)$ for some constant c. We study (I), whereas, by the induction argument $(II) \to 0$ (Equation (B.12)). For a constant $\bar c < \infty$, sub-gaussianity of $H_{i,t}|H_{t-1}$ and overlap, we can write for any $\lambda > 0$,

$$(I) \leq \sum_{j=1}^p \mathbb{E}\Big[\mathbb{E}\Big[\exp\Big(\lambda \bar c ||\hat\gamma_{t-1}||_2^2 - \lambda h\Big)|H_{t-1}, ||\hat\gamma_{t-1}||_2^2 \leq 1/cn\Big]\Big| ||\hat\gamma_{t-1}||_2^2 \leq 1/cn\Big]\eta_{n,t}. \tag{B.17}$$

Since $\hat\gamma_{t-1}$ is measurable with respect to H_{t-1}, we can write

$$\text{(B.17)} \leq \eta_n p_t \exp\Big(\lambda^2/(cn) - \lambda h\Big). \tag{B.18}$$

Choosing $\lambda = hcn/2$ we obtain that the above equation converges to zero as $\log(p_t)/n = o(1)$. After trivial rearrangement, with probability at least $1-(1-\eta_n)-1/n$ (recall that $\eta_n \to 1$) ,

$$\Big|\Big|\sum_{i=1}^n \hat\gamma_{i,t-1} H_{i,t} - \sum_{i=1}^n \hat\gamma_{i,t-1}\frac{1\{D_{i,t}=d_t\}}{P(D_{i,t}=d_t|H_{i,t})} H_{i,t}\Big|\Big|_\infty \lesssim \sqrt{\log(np_t)/n}. \tag{B.19}$$

As a result, we can write

$$
\begin{aligned}
&\Big|\Big|\sum_{i=1}^n \hat{\gamma}_{i,t-1} H_{i,t} - \frac{\sum_{i=1}^n \hat{\gamma}_{i,t-1} \frac{1\{D_{i,t}=d_t\}}{P(D_{i,t}=d_t|H_{i,t})} H_{i,t}}{\sum_{i=1}^n \hat{\gamma}_{i,t-1} \frac{1\{D_{i,t}=d_t\}}{P(D_{i,t}=d_t|H_{i,t})}}\Big|\Big|_\infty \\
&= \Big|\Big|\frac{\sum_{i=1}^n \hat{\gamma}_{i,t-1} H_{i,t} \sum_{i=1}^n \hat{\gamma}_{i,t-1} \frac{1\{D_{i,t}=d_t\}}{P(D_{i,t}=d_t|H_{i,t})} - \sum_{i=1}^n \hat{\gamma}_{i,t-1} \frac{1\{D_{i,t}=d_t\}}{P(D_{i,t}=d_t|H_{i,t})} H_{i,t}}{\sum_{i=1}^n \hat{\gamma}_{i,t-1} \frac{1\{D_{i,t}=d_t\}}{P(D_{i,t}=d_t|H_{i,t})}}\Big|\Big|_\infty \\
&\lesssim \underbrace{\Big|\Big|\frac{\sum_{i=1}^n \hat{\gamma}_{i,t-1} H_{i,t}\Big(1 - \sum_{i=1}^n \hat{\gamma}_{i,t-1} \frac{1\{D_{i,t}=d_t\}}{P(D_{i,t}=d_t|H_{i,t})}\Big)}{\sum_{i=1}^n \hat{\gamma}_{i,t-1} \frac{1\{D_{i,t}=d_t\}}{P(D_{i,t}=d_t|H_{i,t})}}\Big|\Big|_\infty}_{(i)} + \underbrace{\Big|\Big|\frac{\sum_{i=1}^n \hat{\gamma}_{i,t-1} H_{i,t}\Big(1 - \frac{1\{D_{i,t}=d_t\}}{P(D_{i,t}=d_t|H_{i,t})}\Big)}{\sum_{i=1}^n \hat{\gamma}_{i,t-1} \frac{1\{D_{i,t}=d_t\}}{P(D_{i,t}=d_t|H_{i,t})}}\Big|\Big|_\infty}_{(ii)}.
\end{aligned}
$$

Observe now that the denominators of the above expressions are bounded away from zero with high probability as discussed in Equation (B.16). The numerator of (ii) is bounded by Equation (B.19). We are left with the numerator of (i). Note first that

$$
\mathbb{E}\Big[\sum_{i=1}^n \hat{\gamma}_{i,t-1} \frac{1\{D_{i,t}=d_t\}}{P(D_{i,t}=d_t|H_{i,t})}\Big|H_{i,t}\Big] = 1.
$$

We can write

$$
\Big|\Big|\sum_{i=1}^n \hat{\gamma}_{i,t-1} H_{i,t}\Big(1 - \sum_{i=1}^n \hat{\gamma}_{i,t-1} \frac{1\{D_{i,t}=d_t\}}{P(D_{i,t}=d_t|H_{i,t})}\Big)\Big|\Big|_\infty \le \underbrace{\max_j \Big|\sum_{i=1}^n \hat{\gamma}_{i,t-1} H_{i,t}^{(j)}\Big|}_{(j)} \underbrace{\Big|1 - \sum_{i=1}^n \hat{\gamma}_{i,t-1} \frac{1\{D_{i,t}=d_t\}}{P(D_{i,t}=d_t|H_{i,t})}\Big|}_{(jj)}.
$$

Here (jj) is bounded as in Equation (B.16), with probability $1 - 1/n$ at a rate $\sqrt{\log(n)/n}$. The component (j) instead is bounded as

$$
(j) \le \max_{j,i} |H_{i,t}^{(j)}| \lesssim \log(p_t n)
$$

with probability $1 - 1/n$ using subgaussianity of $H_{i,t}^{(j)}$. As a result, all constraints are satisfied.

Finite Norm We now need to show that Equation (B.11) holds. With probability converging to one,

$$n||\hat{\gamma}_t||_2^2 \leq n||\hat{\gamma}_t^*||_2^2 = \sum_{i=1}^n n\hat{\gamma}_{i,t-1}^{*2} \frac{1\{D_{i,t} = d_t\}}{P(D_{i,t} = d_t|H_{i,t})^2} \Big/ \Big(\sum_{i=1}^n \hat{\gamma}_{i,t-1}^* \frac{1\{D_{i,t} = d_t\}}{P(D_{i,t} = d_t|H_{i,t})} \Big)^2.$$

The denominator converges in probability to one by Equation (B.16). The numerator can instead be bounded by $n||\gamma_{t-1}^*||^2$ up-to a finite multiplicative constant by Assumption 8. By the recursive argument $n||\gamma_t^*||^2 = O_p(1)$. □

Lemma B.4. *The weights solving the optimization problem in Algorithm D.1 are such that*

$$||\hat{\gamma}_t||_2^2 \geq 1/n.$$

Proof. Observe that for either algorithms, weights sum to one. The minimum under this constraint only is obtained at $\hat{\gamma}_{i,t} = 1/n$ for all i concluding the proof. □

C Proofs of the Main Theorems

Proof of Theorem 4.3

By Lemmas B.2 and B.3, Theorem 4.3 and Theorem 3.3 directly hold.

Proof of Theorem 4.4

Throughout the proof we will be omitting the script $d_{1:T}$ in the weights and coefficients whenever clear from the context. Note that Theorem 3.4 is a direct corollary of Theorem 4.4.

Weights do not diverge to infinity First notice that by Lemmas B.2, B.3, there exist a $\hat{\gamma}_t^*$ such that for N large enough, with probability converging to one, for some constant C, and $n > N$

$$n||\hat{\gamma}_t||_2^2 \leq n||\hat{\gamma}_t^*||_2^2 = O_p(1). \quad \text{(C.1)}$$

Similar reasoning also applies to $n\sum_{i=1}^n \gamma_{i,t}^2 \mathrm{Var}(\nu_{i,t}|H_{i,t})$ and $n\sum_{i=1}^n \gamma_{i,T}^2 \mathrm{Var}(\varepsilon_{i,T}|\mathcal{F}_T)$ since the conditional variances are uniformly bounded by the finite third moment condition.

Error Decomposition We denote $\bar{\sigma}^2$ the lower bound on the conditional variances and σ_{up}^2 a the upper bound on the variances. Recall $\nu_{i,t} = H_{i,t+1}\beta_{d_{1;T}}^{t+1} - H_{i,t}\beta_{d_{1:T}}^t$ and $\hat{\nu}_{i,t}$ for estimated coefficients, $\hat{\nu}_{i,t} = H_{i,t+1}\hat{\beta}_{d_{1;T}}^{t+1} - H_{i,t}\hat{\beta}_{d_{1:T}}^t$. First we write the expression as

$$\frac{\hat{\mu}(d_{1:T}) - \bar{X}_1\beta_{d_{1:T}}^1}{\sqrt{\hat{V}_T(d_{1:T})}} = \underbrace{\frac{\hat{\mu}(d_{1:T}) - \bar{X}_1\beta_{d_{1:T}}^1}{\sqrt{\sum_{i=1}^n \hat{\gamma}_{i,T}^2 \mathrm{Var}(\varepsilon_{i,T}|H_{i,T-1}) + \sum_{i=1}^n \sum_{t=1}^{T-1} \hat{\gamma}_{i,t}^2 \mathrm{Var}(\nu_{i,t}|H_{i,t})}}}_{(I)} \times$$
$$\times \underbrace{\frac{\sqrt{\sum_{i=1}^n \hat{\gamma}_{i,T}^2 \mathrm{Var}(\varepsilon_{i,T}|H_{i,T}) + \sum_{i=1}^n \sum_{t=1}^{T-1} \hat{\gamma}_{i,t}^2 \mathrm{Var}(\nu_{i,t}|H_{i,t})}}{\sqrt{\sum_{i=1}^n \hat{\gamma}_{i,T}^2 (Y_{i,T} - H_{i,T}\hat{\beta}_{d_{1:T}}^T)^2 + \sum_{t=1}^{T-1} \hat{\gamma}_{i,t}^2 \hat{\nu}_{i,t}^2}}}_{(II)}. \tag{C.2}$$

Term (I) We consider the term (I). By Lemma 4.1, we have

$$(I) = \underbrace{\frac{\sum_{t=1}^T (\beta^t - \hat{\beta}^t)^\top (\hat{\gamma}_t H_t - \hat{\gamma}_{t-1} H_t)}{\sqrt{\sum_{i=1}^n \hat{\gamma}_{i,T}^2 \mathrm{Var}(\varepsilon_{i,T}|H_{i,T-1}) + \sum_{i=1}^n \sum_{t=1}^{T-1} \hat{\gamma}_{i,t}^2 \mathrm{Var}(\nu_{i,t}|H_{i,t})}}}_{(j)}$$
$$+ \underbrace{\frac{\sum_{i=1}^n \hat{\gamma}_{i,T}\varepsilon_{i,T} + \sum_{t=1}^{T-1} \hat{\gamma}_{i,t}\nu_{i,t}}{\sqrt{\sum_{i=1}^n \hat{\gamma}_{i,T}^2 \mathrm{Var}(\varepsilon_{i,T}|H_{i,T-1}) + \sum_{i=1}^n \sum_{t=1}^{T-1} \hat{\gamma}_{i,t}^2 \mathrm{Var}(\nu_{i,t}|H_{i,t})}}}_{(jj)}.$$

We start from (j). Notice since $\sum_{i=1}^n \hat{\gamma}_{i,t} = 1$ and the variances are bounded from below (and Lemma B.4), it follows that

$$\sum_{i=1}^n \hat{\gamma}_{i,T}^2 \mathrm{Var}(\varepsilon_{i,T}|H_{i,T-1}) + \sum_{i=1}^n \sum_{t=1}^{T-1} \hat{\gamma}_{i,t}^2 \mathrm{Var}(\nu_{i,t}|H_{i,t}) \geq T\bar{\sigma}^2 \sum_{i=1}^n \frac{1}{n^2} = T\bar{\sigma}^2/n.$$

Therefore, since the denominator is bounded from below by $\bar{\sigma}\sqrt{T/n}$, and since, by Holder's inequality

$$\sum_{t=1}^{T}(\beta^t - \hat{\beta}^t)^\top(\hat{\gamma}_t H_t - \hat{\gamma}_{t-1}H_t) \lesssim T||\beta^t - \hat{\beta}^t||_1 \Big|\Big|\hat{\gamma}_t H_t - \hat{\gamma}_{t-1}H_t\Big|\Big|_\infty$$

we have

$$(j) \lesssim T \max_t \delta_t(n,p)||\beta^t - \hat{\beta}^t||_1 \rightarrow_p 0 \tag{C.3}$$

under Assumption 9 and the fact that T is fixed. We can now write

$$(I) = o_p(1) + \underbrace{\frac{\sum_{i=1}^n \hat{\gamma}_{i,T}\varepsilon_{i,T}}{\sqrt{\sum_{i=1}^n \hat{\gamma}_{i,T}^2 \text{Var}(\varepsilon_{i,T}|H_{i,T-1})}}}_{(i)} \times \underbrace{\sqrt{\frac{\sum_{i=1}^n \hat{\gamma}_{i,T}^2 \text{Var}(\varepsilon_{i,T}|H_{i,T-1})}{\sum_{i=1}^n \hat{\gamma}_{i,T}^2 \text{Var}(\varepsilon_{i,T}|H_{i,T-1}) + \sum_{i=1}^n\sum_{t=1}^{T-1}\hat{\gamma}_{i,t}^2 \text{Var}(\nu_{i,t}|H_{i,t})}}}_{(ii)}$$

$$+ \sum_{t=1}^{T-1} \underbrace{\frac{\sum_{i=1}^n \hat{\gamma}_{i,1}\nu_{i,t}}{\sqrt{\sum_i \text{Var}(\nu_{i,t}|H_{i,t})\hat{\gamma}_{i,t}^2}}}_{(iii)} \times \underbrace{\frac{\sqrt{\sum_i \text{Var}(\nu_{i,t}|H_{i,t})\hat{\gamma}_{i,t}^2}}{\sqrt{\sum_{i=1}^n \hat{\gamma}_{i,T}^2 \text{Var}(\varepsilon_{i,T}|H_{i,T-1}) + \sum_{i=1}^n\sum_{t=1}^{T-1}\hat{\gamma}_{i,t}^2 \text{Var}(\nu_{i,t}|H_{i,t})}}}_{(iv)}.$$

First, notice that $\sigma(\hat{\gamma}_T) \subseteq \sigma(D_T, \mathcal{F}_T)$, and by Assumption 7 $\varepsilon_T \perp D_T|\mathcal{F}_T$. Therefore,

$$\mathbb{E}[\hat{\gamma}_{i,T}\varepsilon_{i,T}|\mathcal{F}_T, D_T] = 0, \quad \bar{\sigma}^2||\hat{\gamma}_T||_2^2 \le \text{Var}\Big(\sum_{i=1}^n \hat{\gamma}_{i,T}\varepsilon_{i,T}|\mathcal{F}_T, D_T\Big) \le ||\hat{\gamma}_T||_2^2\sigma_\varepsilon^2,$$

where the first statement follows directly from 4.2 and the second statement holds for a finite constant σ_ε^2 by the third moment condition in Assumption 9. By the third moment conditions in Assumption 9 and independence of $\varepsilon_{i,T}$ of D_T given $\mathcal{F}_T$ in Assumption 7, for a constant $0 < C < \infty$,

$$\begin{aligned}\mathbb{E}\Big[\Big(\sum_{i=1}^n \hat{\gamma}_{i,T}\varepsilon_{i,T}\Big)^3\Big|\mathcal{F}_T, D_T\Big] &= \sum_{i=1}^n \hat{\gamma}_{i,T}^3\mathbb{E}[\varepsilon_{i,T}^3|\mathcal{F}_T] \\ &\le C\sum_{i=1}^n \hat{\gamma}_{i,T}^3 \le C||\hat{\gamma}_T||_2^2 \max_i |\hat{\gamma}_{i,T}| \lesssim \log(n)n^{-2/3}||\hat{\gamma}_T||_2^2.\end{aligned}$$

Thus,

$$\mathbb{E}\Big[\sum_{i=1}^{n}\hat{\gamma}_{i,T}^{3}\varepsilon_{i,T}^{3}\Big|\mathcal{F}_T, D_T\Big]\Big/\mathrm{Var}\Big(\sum_{i=1}^{n}\hat{\gamma}_{i,T}\varepsilon_{i,T}\Big|\mathcal{F}_T, D_T\Big)^{3/2} = O(\log(n)n^{-2/3}||\hat{\gamma}_T||_2^{-1}) = o(1).$$

By Liapunov theorem, we have

$$\frac{\sum_{i=1}^{n}\hat{\gamma}_{i,T}\varepsilon_{i,T}}{\sqrt{\sum_{i=1}^{n}\hat{\gamma}_{i,T}\mathrm{Var}(\varepsilon_{i,T}|\mathcal{F}_T)}}\Big|\sigma(\mathcal{F}_T, D_T) \rightarrow_d \mathcal{N}(0,\sigma^2).$$

Consider now (iii) for a generic time t. We study the behaviour of $\sum_{i=1}^{n}\hat{\gamma}_{i,t}\nu_{i,t}$ conditional on $\sigma(\mathcal{F}_t, D_t)$. Since $\sigma(\hat{\gamma}_t) \subseteq \sigma(\mathcal{F}_t, D_t)$, $\hat{\gamma}_t$ is deterministic given $\sigma(\mathcal{F}_t, D_t)$. By Lemma 4.2, $\mathbb{E}[\hat{\gamma}_{i,t}\nu_{i,t}|\mathcal{F}_t, D_t] = 0$. We now study the second moment. Notice that

$$\bar{\sigma}^2||\hat{\gamma}_t||_2^2 \leq \mathrm{Var}(\sum_{i=1}^{n}\hat{\gamma}_{i,t}\nu_{i,t}\big|\mathcal{F}_t, D_t) = \sum_{i=1}^{n}\hat{\gamma}_{i,t}^2\mathrm{Var}(\nu_{i,t}|\mathcal{F}_t, D_t) \leq \sum_{i=1}^{n}\hat{\gamma}_{i,t}^2\sigma_{ub}^2.$$

Finally, we consider the third moment. Under Assumption 9,

$$\mathbb{E}\Big[\sum_{i=1}^{n}\hat{\gamma}_{i,t}^{3}\nu_{i,t}^{3}\Big|X_1, D_1\Big] = \sum_{i=1}^{n}\hat{\gamma}_{i,t}^{3}\mathbb{E}[\nu_{i,t}^3|\mathcal{F}_t, D_t] \leq \sum_{i=1}^{n}\hat{\gamma}_{i,t}^{3}u_{max}^3 \lesssim \log(n)n^{-2/3}||\hat{\gamma}_t||_2^2.$$

Since $||\hat{\gamma}_t||_2 \geq 1/\sqrt{n}$ by Lemma B.4 and since $\mathrm{Var}(\nu_{i,t}|\mathcal{F}_t, D_t) > u_{min}$,

$$\mathbb{E}\Big[\sum_{i=1}^{n}\hat{\gamma}_{i,t}^{3}\nu_{i,t}^{3}\Big|\mathcal{F}_t, D_t\Big]\Big/\mathrm{Var}\Big(\sum_{i=1}^{n}\hat{\gamma}_{i,t}\nu_{i,t}\Big|\mathcal{F}_t, D_t\Big)^{3/2} = O(\log(n)n^{-2/3}||\hat{\gamma}_t||_2^{-1}) = o(1).$$

$$\Rightarrow \frac{\sum_{i=1}^{n}\hat{\gamma}_{i,t}\nu_{i,t}}{\sqrt{\sum_{i=1}^{n}\hat{\gamma}_{i,t}^2\mathrm{Var}(\nu_{i,t}|\mathcal{F}_t, D_t)}}\Big|\sigma(\mathcal{F}_t, D_t) \rightarrow_d \mathcal{N}(0,1).$$

Collecting our results it follows that

$$\frac{\sum_{i=1}^n \hat{\gamma}_{i,T}\varepsilon_{i,T}}{\sqrt{\sum_{i=1}^n \mathrm{Var}(\varepsilon_{i,T}|H_{i,T})\hat{\gamma}_{i,T}^2}} \quad \Bigg| \quad \sigma(\mathcal{F}_T, D_T) \quad \rightarrow_d \mathcal{N}(0,1)$$
$$\frac{\sum_{i=1}^n \hat{\gamma}_{i,t}\nu_{i,t}}{\sqrt{\sum_{i=1}^n \hat{\gamma}_{i,t}^2 \mathrm{Var}(\nu_{i,t}|\mathcal{F}_t, D_t)}} \quad \Bigg| \quad \sigma(\mathcal{F}_t, D_t) \quad \rightarrow_d \mathcal{N}(0,1), \quad \forall t \in \{1, ..., T-1\} \tag{C.4}$$

Notice now that $\sigma(\mathcal{F}_t, D_t)$ consistute a filtration and that

$$\mathbb{E}[\hat{\gamma}_{i,t}\varepsilon_{i,T}\hat{\gamma}_{i,t}\nu_{i,t}|\mathcal{F}_T, D_T] = \hat{\gamma}_{i,t}\nu_{i,t}\hat{\gamma}_{i,T}\mathbb{E}[\varepsilon_{i,T}|\mathcal{F}_T, D_T] = 0$$
$$\mathbb{E}[\hat{\gamma}_{i,t}\hat{\gamma}_{i,s}\nu_{i,s}\hat{\gamma}_{i,t}\nu_{i,t}|\mathcal{F}_{\max\{s,t\}}, D_{\max\{s,t\}}] = \hat{\gamma}_{i,t}\hat{\gamma}_{i,s}\nu_{i,\min\{t,s\}}\mathbb{E}[\nu_{i,\max\{s,t\}}|\mathcal{F}_{\max\{s,t\}}, D_{\max\{s,t\}}] = 0. \tag{C.5}$$

Since each component at time t converges conditionally on the filtration $\sigma(\mathcal{F}_t, D_t)$ and each component is measurable with respect to $\sigma(\mathcal{F}_{t+1}, D_{t+1})$, it follows the joint convergence result

$$\Big[Z_1, \cdots Z_T\Big]^\top \rightarrow_d \mathcal{N}\left(0, I\right),$$
$$Z_t = \frac{\sum_{i=1}^n \hat{\gamma}_{i,t}\nu_{i,t}}{\sqrt{\sum_{i=1}^n \hat{\gamma}_{i,t}^2 \mathrm{Var}(\nu_{i,t}|\mathcal{F}_t, D_t)}}, \quad t \in \{1, \cdots, T-1\}, \quad Z_T = \frac{\sum_{i=1}^n \hat{\gamma}_{i,T}\varepsilon_{i,T}}{\sqrt{\sum_{i=1}^n \mathrm{Var}(\varepsilon_{i,T}|H_{i,T})\hat{\gamma}_{i,T}^2}}.$$

We are left to consider the components (ii), (iv). Define

$$W_T = \sqrt{\frac{\sum_{i=1}^n \hat{\gamma}_{i,T}^2 \mathrm{Var}(\varepsilon_{i,T}|H_{i,T-1})}{\sum_{i=1}^n \hat{\gamma}_{i,T}^2 \mathrm{Var}(\varepsilon_{i,T}|H_{i,T-1}) + \sum_{i=1}^n \sum_{t=1}^{T-1} \hat{\gamma}_{i,t}^2 \mathrm{Var}(\nu_{i,t}|H_{i,t})}},$$
$$W_t = \frac{\sqrt{\sum_i \mathrm{Var}(\nu_{i,t}|H_{i,t})\hat{\gamma}_{i,t}^2}}{\sqrt{\sum_{i=1}^n \hat{\gamma}_{i,T}^2 \mathrm{Var}(\varepsilon_{i,T}|H_{i,T-1}) + \sum_{i=1}^n \sum_{t=1}^{T-1} \hat{\gamma}_{i,t}^2 \mathrm{Var}(\nu_{i,t}|H_{i,t})}}, \quad t \in \{1, \cdots, T-1\}.$$

Note that $||W||_2 = 1$. Note also that we can write the expression (I) as $\sum_{t=1}^T Z_t W_t$.

Therefore we write for any $t \geq 0$,

$$P\Big(\Big|\sum_{t=1}^{T} W_t Z_t\Big| > t\Big) \leq P\Big(||W||_2 \sqrt{\sum_{t=1}^{T} Z_t^2} > t\Big) = P\Big(\sum_{t=1}^{T} Z_t^2 > t^2\Big),$$

where the last equality follows from the fact that $||W||_2 = 1$. Note now that since Z_t are independent standard normal, $\sum_{t=1}^{T} Z_t^2$ is chisquared with T degrees of freedom. To complete the claim, we are only left to show that $(II) \rightarrow_p 1$ to then invoke Slutksy theorem.

Term (II) We can write

$$\begin{aligned}
|(II)^2 - 1| &= \Big|\frac{\sum_{i=1}^{n} \hat{\gamma}_{i,T}^2 (Y_{i,T} - H_{i,T}\hat{\beta}^T)^2 + \sum_{t=1}^{T-1} \hat{\gamma}_{i,t}^2 \hat{\nu}_{i,t}^2}{\sum_{i=1}^{n} \hat{\gamma}_{i,T}^2 \mathrm{Var}(\varepsilon_{i,T}|H_{i,T-1}) + \sum_{i=1}^{n}\sum_{t=1}^{T-1} \hat{\gamma}_{i,t}^2 \mathrm{Var}(\nu_{i,t}|H_{i,t})} - 1\Big| \\
&\lesssim \underbrace{\Big|\frac{n\sum_{i=1}^{n} \hat{\gamma}_{i,T}^2 \varepsilon_{i,T}^2 + n\sum_{t=1}^{T-1} \hat{\gamma}_{i,t}^2 \nu_{i,t}^2}{n\sum_{i=1}^{n} \hat{\gamma}_{i,T}^2 \mathrm{Var}(\varepsilon_{i,T}|H_{i,T-1}) + n\sum_{i=1}^{n}\sum_{t=1}^{T-1} \hat{\gamma}_{i,t}^2 \mathrm{Var}(\nu_{i,t}|H_{i,t})} - 1\Big|}_{(A)} \\
&\quad \underbrace{\Big|\frac{n\sum_{i=1}^{n} \hat{\gamma}_{i,T}^2 \Big[(Y_{i,T} - H_{i,T}\hat{\beta}^T)^2 - (Y_{i,T} - H_{i,T}\beta^T)^2\Big]}{n\sum_{i=1}^{n} \hat{\gamma}_{i,T}^2 \mathrm{Var}(\varepsilon_{i,T}|H_{i,T-1}) + n\sum_{i=1}^{n}\sum_{s=1}^{T} \hat{\gamma}_{i,s}^2 \mathrm{Var}(\nu_{i,s}|H_{i,s})}\Big|}_{(B)} \\
&\quad + \sum_{t=1}^{T-1} \underbrace{\Big|\frac{n\sum_{i=1}^{n} \hat{\gamma}_{i,t}^2 \Big[(H_{i,t+1}\beta^{t+1} - H_{i,t}\hat{\beta}^{t+1})^2 - (H_{i,t+1}\beta^{t+1} - H_{i,t}\beta^t)^2\Big]}{n\sum_{i=1}^{n} \hat{\gamma}_{i,T}^2 \mathrm{Var}(\varepsilon_{i,T}|H_{i,T-1}) + n\sum_{i=1}^{n}\sum_{s=1}^{T} \hat{\gamma}_{i,s}^2 \mathrm{Var}(\nu_{i,s}|H_{i,s})}\Big|}_{(C)}.
\end{aligned} \tag{C.6}$$

To show that (A) converges it suffices to note that the denominator is bounded from below by a finite positive constant by Lemmas B.2, B.3 and the fact that each variance component is bounded away from zero under Assumption 9. The conditional variance of each component in the numerator reads as follows (recall by the above lemmas that $n||\hat{\gamma}_t||^2 = O_p(1)$)

$$\text{Var}\Big(n\sum_{i=1}^n \hat{\gamma}_{i,T}^2 \varepsilon_{i,T}^2 \Big| H_T\Big) \leq n^2 \bar{C}||\hat{\gamma}_T||_4^4 \leq \log^2(n) n^2 \bar{C} n^{-4/3} ||\hat{\gamma}_T||_2^2 = O_p(1)\log^2(n) n n^{-4/3} = o_p(1),$$

$$\text{Var}\Big(n\sum_{i=1}^n \hat{\gamma}_{i,t}^2 \nu_{i,t}^2 \Big| H_t\Big) \leq \bar{C} n^2 ||\hat{\gamma}_T||_4^4 \leq n^2 \log^2(n) \bar{C} n^{-4/3} ||\hat{\gamma}_t||_2^2 = O_p(1)\log^2(n) n n^{-4/3} = o_p(1)$$

and hence (A) converges to zero by the continuous mapping theorem. For the term (B), the denominator is bounded from below away from zero as discussed for (A). The numerator is

$$n\sum_{i=1}^n \hat{\gamma}_{i,T}^2 \Big[(Y_{i,T} - H_{i,T}\hat{\beta}^T)^2 - (Y_{i,T} - H_{i,T}\beta^T)^2\Big] \leq n\sum_{i=1}^n \hat{\gamma}_{i,T}^2 \Big(H_{i,T}(\hat{\beta}^T - \beta^T)\Big)^2 \tag{C.7}$$

We can now write

$$n\sum_{i=1}^n \hat{\gamma}_{i,T}^2 \Big(H_{i,T}(\hat{\beta}^T - \beta^T)\Big)^2 \leq ||\hat{\beta}^T - \beta^T||_1^2 n||\hat{\gamma}_T||^2 || \max_i |H_{i,T}| ||_\infty^2.$$

Notice now that by sub-gaussianity, with probability $1-1/n$, we have $|| \max_i H_{i,T}||_\infty = O(\log(np))$.[34] Since $||\hat{\beta}^T - \beta^T||_1 = o_p(n^{-1/4}), n||\hat{\gamma}_T||^2 = O_p(1)$ and $\log(np)/n^{1/4} = o(1)$ the above expression is $o_p(1)$. Consider now

$$n\sum_{i=1}^n \hat{\gamma}_{i,t}^2 \Big[(H_{i,t+1}\beta^{t+1} - H_{i,t}\hat{\beta}^{t+1})^2 - (H_{i,t+1}\beta^t - H_{i,t}\beta^t)^2\Big] \leq n\sum_{i=1}^n \hat{\gamma}_{i,t}^2 \Big(H_{i,t}(\beta^t - \hat{\beta}^t)\Big)^2$$

which is $o_p(1)$ similarly to the term in Equation (C.7).

Rate of convergence is $n^{-1/2}$. To study the rate of convergences it suffices to show that (for fixed T)

$$n\Big[\sum_{i=1}^n \hat{\gamma}_{i,T}^2 \text{Var}(\varepsilon_{i,T}|H_{i,T-1}) + \sum_{i=1}^n \sum_{t=1}^{T-1} \hat{\gamma}_{i,t}^2 \text{Var}(\nu_{i,t}|H_{i,t})\Big] = O(1).$$

[34] To note this, we can write $P(\max_{i,j} |H_{i,T}^{(j)}| > t) \leq npP(|H_{i,T}^{(j)}| > t) \leq npe^{-t^2 v}$ for some finite constant v. Setting $npe^{-t^2 v} = 1/n$ the claim holds.

This follows directly from Lemma B.3, B.2 and the bounded conditional third moment assumption in Assumption 9.

C.1 Proof of Theorem 4.5

The proof of the corllary follows similarly to the proof of Theorem 4.4. In particular, note that we can write

$$\begin{aligned}
&\frac{\hat{\mu}(d_{1:T}) - \hat{\mu}(d'_{1:T}) - \bar{X}_1\beta^1_{d_{1:T}} + \bar{X}_1\beta^1_{d'_{1:T}}}{\sqrt{\sum_{d\in\{d_{1:T},,d'_{1:T}\}}\sum_{i=1}^n \hat{\gamma}^2_{i,T}(d)(Y_{i,T} - H_{i,T}\hat{\beta}^T_d)^2 + \sum_{t=1}^{T-1}\hat{\gamma}^2_{i,t}(d)\hat{\nu}^2_{i,t}(d)}} \\
&= \underbrace{\frac{\hat{\mu}(d_{1:T}) - \hat{\mu}(d'_{1:T}) - \bar{X}_1\beta^1_{d_{1:T}} + \bar{X}_1\beta^1_{d'_{1:T}}}{\sqrt{\sum_{d\in\{d_{1:T},,d'_{1:T}\}}\sum_{i=1}^n \hat{\gamma}^2_{i,T}(d)\mathrm{Var}(\varepsilon_{i,T}(d)|H_{i,T-1}) + \sum_{i=1}^n\sum_{t=1}^{T-1}\hat{\gamma}^2_{i,t}(d)\mathrm{Var}(\nu_{i,t}(d)|H_{i,t})}}}_{(I)} \times \\
&\times \underbrace{\frac{\sqrt{\sum_{d\in\{d_{1:T},,d'_{1:T}\}}\sum_{i=1}^n \hat{\gamma}^2_{i,T}(d)\mathrm{Var}(\varepsilon_{i,T}(d)|H_{i,T-1}) + \sum_{i=1}^n\sum_{t=1}^{T-1}\hat{\gamma}^2_{i,t}(d)\mathrm{Var}(\nu_{i,t}(d)|H_{i,t})}}{\sqrt{\sum_{d\in\{d_{1:T},,d'_{1:T}\}}\sum_{i=1}^n \hat{\gamma}^2_{i,T}(d)(Y_{i,T} - H_{i,T}\hat{\beta}^T_d)^2 + \sum_{t=1}^{T-1}\hat{\gamma}^2_{i,t}(d)\hat{\nu}^2_{i,t}(d)}}}_{(II)}.
\end{aligned} \tag{C.8}$$

The component (II) converges in probability to one as discussed in the proof of Theorem 4.4. The component (I) behaves similarly to the component (I) in Theorem 4.4 following verbatim the same argument with a single modification: here (I) can be written as $Z_t(d_{1:T})W_t + Z_t(d'_{1:T})W'_t$, where

$$\begin{aligned}
&\Big[Z_1(d_{1:T}), \cdots Z_T(d_{1:T}), Z_1(d'_{1:T}), \cdots Z_T(d'_{1:T})\Big]^\top \to_d \mathcal{N}(0, I), \\
&Z_t(d_{1:T}) = \frac{\sum_{i=1}^n \hat{\gamma}_{i,t}(d_{1:T})\nu_{i,t}(d_{1:T})}{\sqrt{\sum_{i=1}^n \hat{\gamma}(d_{1:T})^2_{i,t}\mathrm{Var}(\nu_{i,t}(d_{1:T})|\mathcal{F}_t, D_t)}}, \quad t \le T-1, \\
&Z_T(d_{1:T}) = \frac{\sum_{i=1}^n \hat{\gamma}_{i,T}(d_{1:T})\varepsilon_{i,T}}{\sqrt{\sum_{i=1}^n \mathrm{Var}(\varepsilon_{i,T}(d_{1:T})|H_{i,T})\hat{\gamma}^2_{i,T}}},
\end{aligned}$$

$$W_T = \sqrt{\frac{\sum_{i=1}^n \hat{\gamma}_{i,T}^2(d_{1:T})\text{Var}(\varepsilon_{i,T}|H_{i,T-1})}{\sum_{d\in\{d_{1:T},,d'_{1:T}\}}\sum_{i=1}^n \hat{\gamma}_{i,T}^2(d)\text{Var}(\varepsilon_{i,T}(d)|H_{i,T-1}) + \sum_{i=1}^n\sum_{t=1}^{T-1}\hat{\gamma}_{i,t}^2(d)\text{Var}(\nu_{i,t}(d)|H_{i,t})}},$$

$$W_t = \sqrt{\frac{\sum_i \text{Var}(\nu_{i,t}|H_{i,t})\hat{\gamma}_{i,t}(d_{1:t})^2}{\sum_{d\in\{d_{1:T},,d'_{1:T}\}}\sum_{i=1}^n \hat{\gamma}_{i,T}^2(d)\text{Var}(\varepsilon_{i,T}(d)|H_{i,T-1}) + \sum_{i=1}^n\sum_{t=1}^{T-1}\hat{\gamma}_{i,t}^2(d)\text{Var}(\nu_{i,t}(d)|H_{i,t})}},$$

and similarly W'_t corresponding to $d'_{1:t}$. Here, independence of $\Big[Z_1(d_{1:T}), \cdots Z_T(d_{1:T})\Big]$ of $\Big[Z_1(d'_{1:T}), \cdots Z_T(d'_{1:T})\Big]$ follows from the fact that $d_1 \neq d'_1$ and hence $\gamma_{i,t}(d_{1:T})\gamma_{i,s}(d'_{1:T}) = 0$ for all s, t conditional on X_1, D_1. The weights by construction satisfy $||(W, -W')||_2^2 = 1$. Therefore we write for any $t \geq 0$,

$$P\Big(\Big|\sum_{t=1}^T W_t Z_t(d_{1:T}) - \sum_{t=1}^T W'_t Z_t(d'_{1:T})\Big| > t\Big) \leq P\Big(||W||_2\sqrt{\sum_{d\in\{d_{1:T},d'_{1:T}\}}\sum_{t=1}^T Z_t^2(d_{1:T})} > t\Big)$$
$$= P\Big(\chi_{2T}^2 > t^2\Big),$$

with χ_{2T}^2 being a chi-squared random variable with $2T$ degrees of freedom.

C.2 Tighter asymptotic results

Theorem C.1 (Tighter confidence bands under more restrictive conditions)**.** *Suppose that the conditions in Theorem 4.5 hold. Suppose in addition that for all* $t \in \{1, \cdots, T-1\}$, $n\sum_{i=1}^n \hat{\gamma}_{i,t}^2\text{Var}(\nu_{i,t}|\mathcal{F}_{t-1}) \to_{as} c_t$, $n\sum_{i=1}^n \hat{\gamma}_{i,t}^2\text{Var}(\varepsilon_{i,T}|\mathcal{F}_{T-1}) \to_{as} C_T$ *for constants* $\{c_t\}_{t=1}^T$*. Then, whenever* $\log(np)/n^{1/4} \to 0$ *with* $n, p \to \infty$,

$$(\hat{V}_T(d_{1:T}) + \hat{V}_T(d'_{1:T}))^{-1/2}\sqrt{n}\Big(\hat{\mu}(d_{1:T}) - \hat{\mu}(d'_{1:T}) - \text{ATE}(d_{1:T}, d'_{1:T})\Big) \to_d \mathcal{N}(0,1). \quad \text{(C.9)}$$

Proof of Theorem C.1. The proof follows verbatim from the proof of Theorem 4.5, while here the components $W_t \to_{a.s.} c_t, W'_t \to_{a.s.} c'_t$ for constants c_t, c'_t. Note that by Lemma B.3, the asymptotic limits c_t must be finite since

$$n\sum_{i=1}^n \hat{\gamma}_{i,t}^2\text{Var}(\nu_{i,t}|\mathcal{F}_{t-1}) \leq \bar{u}n||\hat{\gamma}_t||^2 = O_p(1),$$

where $\bar{u}$ is a finite constant by Assumption 9 (ii). Following the same argument as in the proof of Theorem 4.5, we obtain that the left-hand side of Equation (C.9) converges to

$$\sum_{t=1}^{T} c_t Z_t - \sum_{t=1}^{T} c_t' Z_t', \quad (Z_1, \cdots, Z_T, Z_1', \cdots, Z_T') \sim \mathcal{N}(0, I).$$

The variance is therefore $\sum_{t=1}^{T} c_t^2 + \sum_{t=1}^{T} c_t^{2'} = 1$, since $||(W, -W)||^2 = 1$ as discussed in the proof of Theorem 4.5. □

D Additional Algorithms

Algorithm D.1 presents balancing with multiple periods. Algorithm D.2 presents estimation of the coefficients for multiple periods. Its extensions for a linear model on the treatment assignments (hence using all in-sample information) follows similarly to Algorithm 2. Algorithm D.3 presents the choice of the tuning parameters. The algorithm imposes stricter tuning on those covariates whose coefficients are non-zero. Whenever many coefficients (more than one-third) are non-zero, we impose a stricter balancing on those with the largest size.[35]

[35] E.g., 60 coefficients are prioritized for $T = 2, p = 100$ design, since the dimension is Tp.

Algorithm D.1 Dynamic covariate balancing (DCB): multiple time periods

Require: Observations $\{Y_{i,1}, X_{i,1}, D_{i,1}, \cdots, Y_{i,T}, X_{i,T}, D_{i,T}\}$, treatment history $(d_{1:T})$, finite parameters $\{K_{1,t}\}, K_{2,1}, K_{2,2}, \cdots, K_{2,T}$, constraints $\delta_1(n, p_1), \delta_2(n, p_2), \cdots, \delta_T(n, p_T)$.

Estimate $\beta^{(t)}_{d_{1:T}}$ as in Algorithm D.2 in Appendix D.
Let $\hat{\gamma}_{i,0} = 1/n$ and $t = 0$;
for each $t \leq T-1$ **do**
  $\hat{\gamma}_{i,t} = 0$, if $D_{i,1:t} \neq d_{1:t}$
  Estimate time t weights with

$$\hat{\gamma}_t = \arg\min_{\gamma_t} \sum_{i=1}^n \gamma_{i,t}^2, \quad \text{s.t. } \Big\| \frac{1}{n} \sum_{i=1}^n \hat{\gamma}_{i,t-1} H_{i,t} - \gamma_{i,t} H_{i,t} \Big\|_\infty \leq K_{1,t}\delta_t(n, p_t), \qquad \text{(D.1)}$$
$$1^\top \gamma_t = 1, \gamma_t \geq 0, \|\gamma_t\|_\infty \leq K_{2,t}\log(n) n^{-2/3}.$$

end for ▷ obtain T balancing vectors
  return Estimate of the average potential outcome as in Equation (12)

Algorithm D.2 Coefficients estimation with multiple periods

Require: Observations, history $(d_{1:2})$, `model` $\in$ {`full interactions, linear`}.

if `model` = `full interactions` **then**
  Estimate $\beta^T_{d_{1:T}}$ by regressing $Y_{i,T}$ onto $H_{i,T}$ for i with $D_{1:T} = d_{1:T}$.
  for $t \in \{T-1, \cdots, 1\}$ **do**
    Estimate $\beta^t_{d_{1:T}}$ by regressing $H_{i,t+1}\hat{\beta}^{t+1}_{d_{1:T}}$ onto $H_{i,t}$ for i that has the treatment history $(d_{1:t})$.
  end for
else
  Estimate β^T by regressing $Y_{i,T}$ onto $(H_{i,T}, D_{i,T})$ for all i (without penalizing $(D_{i,1:T})$) and define $H_{i,T}\hat{\beta}_{d_{1:T}} = (H_{i,T}, d_T)\hat{\beta}^T$ for all $i : D_{i,1:(T-1)} = d_{1:(T-1)}$;
  Repeat sequentially as in Algorithm 2
end if

Algorithm D.3 Tuning Parameters for DCB

Require: Observations $\{Y_{i,1}, X_{i,1}, D_{i,1}, ..., Y_{i,T}, X_{i,T}, D_{i,T}\}$, $\delta_t(n,p)$, treatment history $(d_{1:T})$, L_t, U_t, grid length G, number of grids R.

Estimate coefficients as in Algorithm D.2 and let $\hat{\gamma}_{i,0} = 1/n$;

Define R grids of length G, denoted as $\mathcal{G}_1, ..., \mathcal{G}_R$, equally between L_t an U_t.

Define

$$\mathcal{S}_1 = \{j : |\hat{\beta}^{t,(j)}| \neq 0\}, \quad \mathcal{S}_2 = \{j : |\hat{\beta}^{t,(j)}| = 0\}.$$

(Non-sparse regression): if $|\mathcal{S}_1|$ is too large (i.e., $> \dim(\hat{\beta}^t)/3$), select $\mathcal{S}_1$ the set of the $1/3^{rd}$ largest coefficients in absolute value and $\mathcal{S}_2 = \mathcal{S}_1^c$.

for each $s_1 \in 1 : G$ **do**

  for each $K_{1,t}^a \in \mathcal{G}_{s_1}$ **do**

    for each $K_{1,t}^b \in \mathcal{G}_{s_1}$ **do**

      Let $\hat{\gamma}_{i,t} = 0$, if $D_{i,1:t} \neq d_{1:t}$ and define $\hat{\gamma}_t := \text{argmin}_{\gamma_t} \sum_{i=1}^n \gamma_{i,t}^2$

$$\begin{aligned} \text{s.t. } & \Big|\frac{1}{n}\sum_{i=1}^n \hat{\gamma}_{i,t-1} H_{i,t}^{(j)} - \gamma_{i,t} H_{i,t}^{(j)}\Big| \leq K_{1,t}^a \delta_t(n,p), \quad \forall j : \hat{\beta}^{t,(j)} \in \mathcal{S}_1 \\ & \Big|\frac{1}{n}\sum_{i=1}^n \hat{\gamma}_{i,t-1} H_{i,t}^{(j)} - \gamma_{i,t} H_{i,t}^{(j)}\Big| \leq K_{1,t}^b \delta_t(n,p) \quad \forall j : \hat{\beta}^{t,(j)} \in \mathcal{S}_2 \\ & \sum_{i=1}^n \gamma_{i,t} = 1, \quad ||\gamma_t||_\infty \leq \log(n) n^{-2/3}, \gamma_{i,t} \geq 0. \end{aligned} \tag{D.2}$$

      Stop if: a feasible solution exists.

    end for

  end for

end for

  return $\hat{\mu}_T(d_{1:T})$

E More on simulations

Table E.1: Confidence intervals length for design in main text with chi-squared distribution.

	$t = 2, \eta = 0.3$	$t = 2, \eta = 0.5$	$t = 3, \eta = 0.3$	$t = 3, \eta = 0.5$
p = 50 - Sparse	1.640	1.806	3.107	3.370
p = 100 - Sparse	1.754	1.912	3.221	3.488
p = 200 - Sparse	1.691	1.829	3.052	3.474
p = 300 - Sparse	1.706	1.847	3.113	3.446
p = 50 - Moderate	1.565	1.686	3.160	3.335
p = 100 - Moderate	1.640	1.745	3.215	3.466
p = 200 - Moderate	1.559	1.658	3.085	3.323
p = 300 - Moderate	1.541	1.628	3.028	3.230
p = 50 - Harmonic	1.641	1.777	3.138	3.287
p = 100 - Harmonic	1.682	1.796	3.201	3.323
p = 200 - Harmonic	1.678	1.781	3.257	3.433
p = 300 - Harmonic	1.733	1.856	3.348	3.527

E.1 Simulations under misspecification

We simulate the outcome model over each period using non-linear dependence between the outcome, covariates, and past outcomes. The function that we choose for the dependence of the outcome with the past outcome and covariates follows similarly to Athey et al. (2018), where, differently, here, such dependence structure is applied not only to the first covariate only (while keeping a linear dependence with the remaining ones) but to all covariates, making the scenarios more challenging for the DCB method. Formally, the DGP is the following:

$$\begin{aligned} Y_2(d_1, d_2) = &\log(1 + \exp(-2 - 2X_1\beta_{d_1,d_2})) + \log(1 + \exp(-2 - 2X_2\beta_{d_1,d_2})) \\ &+ \log(1 + \exp(-2 - 2Y_1)) + d_1 + d_2 + \varepsilon_2, \end{aligned}$$

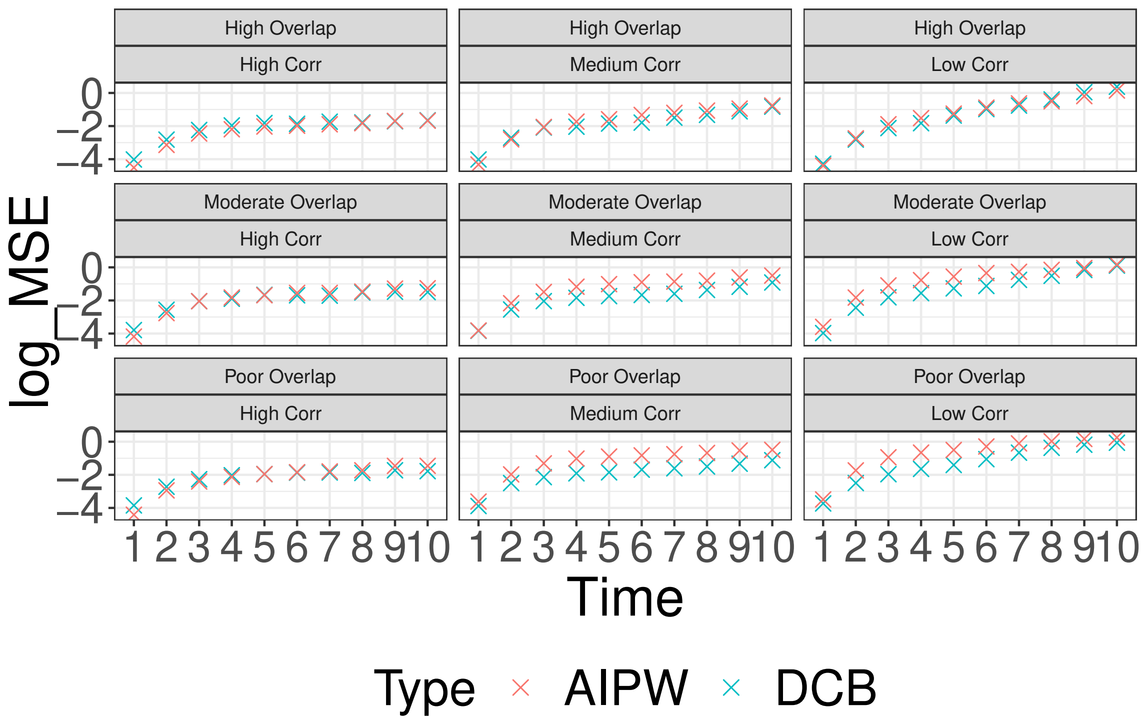

Figure E.1: Mean-squared error in log-scale. Simulations for $T \leq 10, p = 100, n = 400$, two-hundred replications. Here high-correlation denotes strong serial depedence between treatment assignments with $\alpha = 0.9$, medium with $\alpha = 0.7$ and weak with $\alpha = 0.5$. $\eta \in \{0.1, 0.3, 0.5\}$ for good, moderate and poor overlap, respectively.

and similarly for $Y_3(d_1, d_2, d_3)$, with also including covariates and outcomes in period $T = 2$. Coefficients β are obtained from the sparse model formulation discussed in the main text. Results are collected in Table E.3 for the MSE and for the bias and variance in he subsequent tables below. Interestingly, we observe that DCB performs relatively well under the misspecified model, even if our method does not use any information on the propensity score. We also note that our adaptation of the double lasso to dynamic setting performs comparable or better in the presence of two periods only or a sparse structure. However, as the number of periods increase or sparsity decreases Double Lasso's performance deteriorates.

Table E.2: MSE under misspecified model in a sparse setting.

	$T = 2$		$T = 3$	
	$\eta = 0.3$	$\eta = 0.5$	$\eta = 0.3$	$\eta = 0.5$
DCB	0.238	0.354	0.751	0.402
aIPW*	0.434	0.802	1.363	1.622
aIPWh	0.863	1.363	1.882	2.464
CAEW (MSM)	0.815	1.364	7.889	8.675
D. Lasso	0.121	0.142	0.689	0.503
Seq.Est	0.811	0.346	2.288	2.031

Table E.3: MSE under misspecified model in a moderately sparse setting.

	$T=2$		$T=3$	
	$\eta=0.3$	$\eta=0.5$	$\eta=0.3$	$\eta=0.5$
DCB	0.212	0.256	0.326	0.384
aIPW*	0.428	0.789	1.364	1.616
aIPWh	0.826	1.313	1.857	2.434
CAEW (MSM)	0.781	1.317	7.833	8.616
D. Lasso	0.115	0.133	0.675	0.494
Seq.Est	0.847	0.366	2.316	2.058

Table E.4: Bias for sparse setting under misspecified model.

	$t=2,\eta=0.3$	$t=2,\eta=0.5$	$t=3,\eta=0.3$	$t=3,\eta=0.5$
DCB	0.227	0.340	-0.467	-0.199
AIPW - Known Prop	0.146	0.288	-0.0003	0.318
AIPW - High Prop	0.852	1.119	1.245	1.459
AIPW - Low Prop	0.551	1.045	1.378	2.057
CAEW	0.760	1.086	2.718	2.872
Double Lasso	0.156	0.225	0.671	0.469
Seq.Est.	-0.793	-0.448	-1.391	-1.276

Table E.5: Variance for sparse setting under misspecified model.

	$t=2,\eta=0.3$	$t=2,\eta=0.5$	$t=3,\eta=0.3$	$t=3,\eta=0.5$
DCB	0.187	0.239	0.533	0.363
AIPW - Known Prop	0.413	0.719	1.364	1.521
Naive Lasso	0.273	0.259	1.058	1.194
AIPW - High Prop	0.138	0.111	0.333	0.336
AIPW - Low Prop	0.612	0.225	0.827	0.438
CAEW	0.237	0.184	0.500	0.425
Double Lasso	0.098	0.092	0.239	0.284
Seq.Est.	0.183	0.145	0.354	0.404

Table E.6: Bias for moderately sparse model under misspecification.

	$t = 2, \eta = 0.3$	$t = 2, \eta = 0.5$	$t = 3, \eta = 0.3$	$t = 3, \eta = 0.5$
This Paper	0.202	0.358	0.096	0.323
AIPW - Known Prop	0.123	0.266	-0.010	0.308
AIPW - High Prop	0.830	1.097	1.235	1.449
AIPW - Low Prop	0.529	1.023	1.367	2.047
CAEW	0.738	1.064	2.708	2.862
Double Lasso	0.134	0.202	0.661	0.459
Seq.Est.	-0.815	-0.470	-1.401	-1.286

Table E.7: Variance for moderately sparse model under misspecification.

	$t = 2, \eta = 0.3$	$t = 2, \eta = 0.5$	$t = 3, \eta = 0.3$	$t = 3, \eta = 0.5$
This Paper	0.171	0.129	0.317	0.280
AIPW - Known Prop	0.413	0.719	1.364	1.521
AIPW - High Prop	0.138	0.111	0.333	0.336
AIPW - Low Prop	0.612	0.225	0.827	0.438
CAEW	0.237	0.184	0.500	0.425
Double Lasso	0.098	0.092	0.239	0.284
Seq.Est.	0.183	0.145	0.354	0.404