JOINT REGRESSION ADJUSTMENT FOR SEQUENTIAL RANDOMIZED EXPERIMENTS WITH PRETRAINED PREDICTIONS

Anonymous authors
Paper under double-blind review

ABSTRACT

In longitudinal randomized experiments on internet platforms, researchers often need to assess the effects of sequential interventions on long-term outcomes. In such experiments, there are two types of potentially useful information for predicting outcomes: within-experiment covariates that evolve with past treatments, and fixed pretrained predictors trained using offline data available before the experiment. However, the former are affected by past treatments, while the latter may not transfer accurately to the current experiment. These difficulties limit the usability of such information. We propose a joint regression adjustment method for estimating average potential outcomes under fixed treatment paths and causal contrasts between these paths. It learns sequential outcome predictions through cross-fitting using time varying covariates within the experiment, then jointly adjusts for these predictions and fixed pretrained predictions. Under general conditions, we establish asymptotic normality and consistent variance estimation for valid Wald inference. Its asymptotic variance is no larger than that of Horvitz–Thompson, Hájek, and AIPW estimators. We characterize the incremental variance reduction from the pretrained predictions and give a necessary and sufficient condition for attaining the semiparametric efficiency bound, which can hold even when the internal predictions are misspecified. We demonstrate the advantages of our method through synthetic experiments and real-data experiments on a Chinese game platform.

1 INTRODUCTION

In online platforms, digital health, and other multistage experiments, researchers often need to evaluate the effects of different treatment paths on long term outcomes (Shi et al., 2023; Battalio et al., 2021). We consider sequential randomized experiments, where the mean potential outcome under a prespecified treatment path can be identified through inverse probability weighting (Robins et al., 2000; Robins & Hernán, 2009). However, standard IPW estimators such as Horvitz–Thompson and Hájek estimators use the treatment assignment probabilities to construct weights but do not use outcome models to reduce variance. They can therefore have large variances.

Online platforms often provide two sources of predictive information. One is time varying covariates that are affected by past treatments and may predict future outcomes. The other is fixed external predictors trained on large external datasets before analysis. The time varying covariates can be used to learn sequential outcome predictions within the current experiment, but their accuracy depends on the experiment's sample size and the choice of learner. External predictors can have strong predictive power in external data, but may not transfer accurately to the current experiment and can be arbitrarily biased. Thus, directly using these predictions does not guarantee improved estimation precision. A method is needed to combine the two sources of predictions without increasing asymptotic variance, even when either predictor is misspecified.

Existing methods address different parts of this problem. Longitudinal AIPW and LTMLE use sequential outcome regressions and attain the semiparametric efficiency bound when these regressions are correctly specified (Bang & Robins, 2005; van der Laan & Gruber, 2012). Recent work models the evolution of time varying covariates through transition kernels (Chihara et al., 2026). However, when the outcome regressions are misspecified, these methods can have an even larger asymptotic variance

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

than the IPW estimators (Rubin & van der Laan, 2008; Cohen & Fogarty, 2024; Kallus & Uehara, 2019). Another line of work uses pretreatment covariates or external predictions for regression adjustment (Tsiatis et al., 2008; Lin, 2013; Guo et al., 2021; Schuler et al., 2022; De Bartolomeis et al., 2025). However, these methods only focus on single-stage randomized experiments. Kallus & Uehara (2019) develop intrinsically efficient estimators for off-policy evaluation, which can be viewed as estimating average potential outcomes. They use a shared parameterized action-value model learned through temporal-difference learning. However, longitudinal causal inference generally involves stage-specific recursive outcome regressions, because the prediction target changes across treatment stages with the remaining treatment path. Their asymptotic analysis of learned outcome models relies on parameterized models and empirical process arguments for Donsker classes, limiting its applicability to flexible machine learning predictions. They also do not provide the central limit theorem or variance estimators, leaving no explicit procedure for constructing confidence intervals. Moreover, we may also be interested in causal effects. Separately estimating the two path means does not generally minimize the variance of the causal effect estimator, which also depends on their covariance. We study estimation and statistical inference for both average potential outcomes and causal contrasts between treatment paths in sequential randomized experiments, combining flexible machine learning predictions with fixed external predictors and providing valid inference and asymptotic variance guarantees even when both are misspecified.

In this work, we propose joint regression adjustment for longitudinal randomized experiments. For a single treatment path, we first learn internal sequential outcome predictions recursively backward in time through cross fitting. We then use the mean zero weight increments induced by sequential randomization to jointly adjust for these predictions and fixed external predictions. Under the stated conditions, even with misspecified predictions, we establish asymptotic normality and consistency of the residual and HC2 variance estimators, ensuring asymptotically valid Wald coverage. We further give an exact asymptotic variance decomposition, showing a variance no larger than those of HT, Hájek, and AIPW and characterizing the incremental information provided by pretrained predictions beyond internal predictions. For causal effects between two treatment paths, separately minimizing the variances of the two path mean estimators does not generally minimize the variance of their difference. We perform joint regression adjustment directly for the causal effect and establish the same inference guarantees and asymptotic variance advantages. Finally, synthetic experiments and experiments using real data from a Chinese gaming platform demonstrate precision gains over existing methods, including AIPW and LTMLE that use pretrained predictions as additional features.

The remainder of this paper is organized as follows. Section 2 reviews related work, and Section 3 introduces the problem setup. Section 4 presents joint regression adjustment for treatment path means and establishes its inferential and efficiency properties. Section 5 extends the framework to causal contrasts between treatment paths. Section 6 reports synthetic and real-data experiments, and Section 7 concludes the paper. Proofs of the theoretical results are provided in Appendix A.

2 Related Work

2.1 Estimation and Inference for Longitudinal Treatment Effects

Classical methods in longitudinal causal inference establish identification and estimation frameworks for potential outcomes under treatment paths (Robins, 1986; 1994; Robins et al., 2000). AIPW and LTMLE use sequential outcome regressions for estimation and inference (Bang & Robins, 2005; van der Laan & Gruber, 2012). Recent work models covariate evolution through transition models to improve estimation precision using histories affected by prior treatments (Chihara et al., 2026). Machine learning methods have also been developed to estimate longitudinal counterfactual outcomes and treatment effects (Bica et al., 2020; Melnychuk et al., 2022; Frauen et al., 2025). Building on orthogonal estimation and cross fitting in double machine learning (Chernozhukov et al., 2018), related work establishes estimation and inference methods for dynamic treatment effects (Lewis & Syrgkanis, 2021; Bodory et al., 2022). These methods attain semiparametric efficiency under suitable nuisance convergence conditions. Under outcome model misspecification, they may have larger asymptotic variance than unadjusted estimators (Rubin & van der Laan, 2008; Cohen & Fogarty, 2024; Kallus & Uehara, 2019). Our joint regression adjustment provides valid asymptotic inference and asymptotic variance no larger than that of HT, Hájek, and AIPW, even when both internal and external predictions are misspecified.

2.2 Regression Adjustment in Randomized Experiments

Regression adjustment uses covariates to improve estimation precision in randomized experiments (Tsiatis et al., 2008; Lin, 2013; Deng et al., 2013). Methods for high dimensional covariates and machine learning further expand the range of prediction models used for adjustment (Wager et al., 2016; Bloniarz et al., 2016; Guo et al., 2021; Lei & Ding, 2021). Another line of work uses prognostic scores or pretrained predictions to incorporate external information (Schuler et al., 2022; De Bartolomeis et al., 2025). Related work on prediction powered inference combines machine learning predictions with observed outcomes to improve inferential precision (Angelopoulos et al., 2023; Zrnic & Candès, 2024; Li & Ignatiadis, 2025; Wang et al., 2025). The above adjustment methods for randomized experiments only consider a single treatment assignment. We instead study sequential experiments in which treatment is randomized at each stage and the time varying covariates are affected by prior treatments.

To our knowledge, we develop the first joint regression adjustment framework for estimating treatment path means and contrasts in sequential randomized experiments that combines flexible sequential predictions with fixed external predictors, with valid asymptotic inference and asymptotic variance no larger than that of HT, Hájek, and AIPW even under prediction misspecification. In synthetic and real-data experiments, our joint regression adjustment is competitive with methods that directly include pretrained predictions as additional features in outcome regression.

3 Problem Setup

We consider the standard setting for longitudinal causal inference (Robins & Hernán, 2009).We consider n i.i.d. samples from a distribution P. The observed data for a single unit can be written as $O = (L_1, A_1, \ldots, L_T, A_T, Y)$, where T is fixed. At stage t, the covariate L_t is observed before the binary treatment $A_t \in \{0, 1\}$. The outcome Y is observed after A_T. Define $\bar{A}_t = (A_1, \ldots, A_t)$ and the pretreatment history $H_t = (L_1, A_1, \ldots, L_{t-1}, A_{t-1}, L_t)$. Let $f = (f_1, \ldots, f_T)$ be a fixed pretrained predictor. At stage t, $f_t(H_t)$ predicts Y from the pretreatment history H_t. The predictor can be arbitrarily misspecified.

For a prespecified static path $\bar{a} = (a_1, \ldots, a_T) \in \{0, 1\}^T$, define $\bar{a}_t = (a_1, \ldots, a_t)$ and let $Y(\bar{a})$ be the potential outcome. Its mean is $\mu(\bar{a}) = \mathbb{E}\{Y(\bar{a})\}$. For two prespecified static paths $\bar{a}$ and $\bar{b}$, the causal effect is $\tau(\bar{a}, \bar{b}) = \mu(\bar{a}) - \mu(\bar{b}) = \mathbb{E}\{Y(\bar{a}) - Y(\bar{b})\}$. The following standard assumptions ensure identification of these target parameters (Robins & Hernán, 2009).

Assumption 1. *The following conditions hold for every treatment path $\bar{a}$.* (i) *Each unit's potential outcome depends only on its own treatment path, and $Y = Y(\bar{A}_T)$.* (ii) *Sequential randomization holds: $A_t \perp Y(\bar{a}) \mid H_t$ for $t = 1, \ldots, T$.* (iii) *The propensity score $e_t(a \mid H_t) = \mathbb{P}(A_t = a \mid H_t)$ is known, and there exists $\epsilon > 0$ such that $e_t(a \mid H_t) \geq \epsilon$ for every $a \in \{0, 1\}$ and t.*

Define the products of inverse propensity scores by $W_0^{\bar{a}} = 1, W_t^{\bar{a}} = \prod_{s=1}^{t} \frac{\mathbf{1}(A_s = a_s)}{e_s(a_s \mid H_s)}, t = 1, \ldots, T$. Under Assumption 1, the standard inverse-probability weighting identification result (Robins et al., 2000) gives $\mu(\bar{a}) = \mathbb{E}\{W_T^{\bar{a}} Y\}$ and $\tau(\bar{a}, \bar{b}) = \mathbb{E}\{(W_T^{\bar{a}} - W_T^{\bar{b}})Y\}$.

Standard estimators include the HT, Hájek, and AIPW estimators (Robins et al., 2000; Bang & Robins, 2005; Robins & Hernán, 2009). Our goal is to combine sequential outcome predictions learned from the experiment with the fixed pretrained predictor f to improve precision while retaining valid inference.

Our construction utilizes the mean zero weight increments induced by sequential randomization. For any square integrable function $g_t(H_t)$,

$$\mathbb{E}\left\{\left(W_t^{\bar{a}} - W_{t-1}^{\bar{a}}\right) g_t(H_t)\right\} = \mathbb{E}\left[g_t(H_t)\mathbb{E}\left(W_t^{\bar{a}} - W_{t-1}^{\bar{a}} \mid H_t\right)\right] = \mathbb{E}\left[g_t(H_t) W_{t-1}^{\bar{a}} \left\{\frac{\mathbb{P}(A_t = a_t \mid H_t)}{e_t(a_t \mid H_t)} - 1\right\}\right] = 0.$$

Hence, for any collection of square integrable functions $g = (g_1, \ldots, g_T)$ and any coefficients $\beta = (\beta_1, \ldots, \beta_T) \in \mathbb{R}^T$,

$$\mathbb{E}\left\{W_T^{\bar{a}} Y - \sum_{t=1}^{T} \beta_t \left(W_t^{\bar{a}} - W_{t-1}^{\bar{a}}\right) g_t(H_t)\right\} = \mu(\bar{a}).$$

Thus, we can use any functions of H_t in regression adjustment without changing the identification target. Section 4 jointly uses the fixed predictor f and cross fitted outcome predictions as such functions.

4 Joint Regression Adjustment for Average Potential Outcomes

We first present the method for the mean potential outcome $\mu(\bar{a})$ under a prespecified treatment path.

4.1 Constructing the Estimator

Overview: Figure 1 summarizes our approach. For each prespecified treatment path $\bar{a}$, the construction has three steps. **Step A: sequential outcome prediction.** Motivated by the efficient influence function, we learn cross fitted sequential outcome predictions $\widehat{Q}_t(a_t, H_t)$ from the current experiment. We estimate these functions recursively and evaluate each prediction out of fold. **Step B: joint regression adjustment.** At every stage, we combine the fixed pretrained prediction $f_t(H_t)$ and the cross fitted sequential outcome prediction $\widehat{Q}_t(a_t, H_t)$ in a joint regression adjustment. The fitted intercept from this regression is our point estimator $\widehat{\mu}(\bar{a})$. **Step C: confidence intervals.** We estimate the asymptotic variance of $\widehat{\mu}(\bar{a})$ using the squared fitted residuals and construct a Wald confidence interval.

We describe the three steps in detail below.

Step A: sequential outcome prediction. We learn the sequential outcome regressions from the current experiment. On histories with positive probability under the target path, define, for $a \in \{0, 1\}$,

$$Q_T(a, H_T) = \mathbb{E}(Y \mid H_T, A_T = a), \quad Q_t(a, H_t) = \mathbb{E}\{Q_{t+1}(a_{t+1}, H_{t+1}) \mid H_t, A_t = a\}, \quad t = T-1, \ldots, 1.$$

Thus, $Q_t(a_t, H_t)$ is the conditional expectation of Y given H_t, treatment a_t at stage t, and subsequent adherence to the target path. As shown in (van der Laan & Gruber, 2012), the efficient influence function (EIF) for $\mu(\bar{a})$ is

$$\begin{aligned}\phi^{\bar{a}}_{\mathrm{EIF}}(O) &= Q_1(a_1, H_1) - \mu(\bar{a}) + \sum_{t=1}^{T-1} W^{\bar{a}}_t \{Q_{t+1}(a_{t+1}, H_{t+1}) - Q_t(a_t, H_t)\} + W^{\bar{a}}_T \{Y - Q_T(a_T, H_T)\} \\ &= W^{\bar{a}}_T Y - \mu(\bar{a}) - \sum_{t=1}^{T} \{W^{\bar{a}}_t - W^{\bar{a}}_{t-1}\} Q_t(a_t, H_t).\end{aligned}$$

Its variance $V^{\bar{a}}_{\mathrm{EIF}} = \mathbb{E}\{[\phi^{\bar{a}}_{\mathrm{EIF}}(O)]^2\}$ is the semiparametric efficiency bound for $\mu(\bar{a})$. Every regular asymptotically linear estimator has asymptotic variance no smaller than $V^{\bar{a}}_{\mathrm{EIF}}$ (Tsiatis, 2006). This form motivates the sequential outcome predictions. When $Q_t(a_t, H_t)$ is used in the joint regression adjustment, the adjusted estimation has the form of the efficient influence function.

In practice, $Q_t(a, H_t)$ is unknown. We estimate these functions through cross fitting (Chernozhukov et al., 2018). Partition the n samples into K folds $I_1, \ldots, I_K$. For each fold k, let $\widehat{Q}^{(-k)}_t$ denote the sequential outcome prediction function trained using only the samples in I^c_k. We follow the recursive regression approach of Frauen et al. (2023); Oprescu et al. (2025); Shirakawa et al. (2024). First, fit $\widehat{Q}^{(-k)}_T$ using Y as the response and (A_T, H_T) as the inputs. Then, for $t = T-1, \ldots, 1$, fit $\widehat{Q}^{(-k)}_t$ using $\widehat{Q}^{(-k)}_{t+1}(a_{t+1}, H_{t+1})$ as the response and (A_t, H_t) as the inputs. All recursive fitting

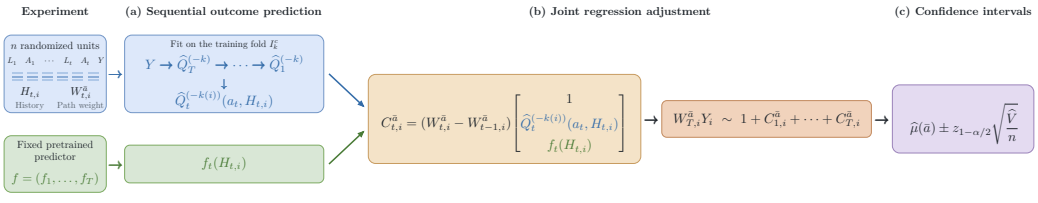

Figure 1: Joint regression adjustment with pretrained and sequential outcome predictions.

is performed within I_k^c. Let $k(i)$ denote the fold containing sample i. Its prediction at stage t is $\widehat{Q}_t^{(-k(i))}(a_t, H_{t,i})$.

Step B: joint regression adjustment. We combine the fixed pretrained predictor $f = (f_1, \ldots, f_T)$ and the sequential outcome predictions $\widehat{Q} = (\widehat{Q}_1, \ldots, \widehat{Q}_T)$ in regression adjustment. As in regression adjustment for randomized experiments (Lin, 2013; Lei & Ding, 2021), the regression serves as a linear working model and does not assume that the true outcome regressions are linear. For sample i at stage t, define

$$C_{t,i}^{\bar{a}}(\widehat{Q}, f) = \left(W_{t,i}^{\bar{a}} - W_{t-1,i}^{\bar{a}}\right) \begin{pmatrix} 1 \\ \widehat{Q}_t^{(-k(i))}(a_t, H_{t,i}) \\ f_t(H_{t,i}) \end{pmatrix}, \quad C_i^{\bar{a}}(\widehat{Q}, f) = \left\{C_{1,i}^{\bar{a}}(\widehat{Q}, f)^\mathsf{T}, \ldots, C_{T,i}^{\bar{a}}(\widehat{Q}, f)^\mathsf{T}\right\}^\mathsf{T}.$$

At each stage t, the first component adjusts for random treatment imbalance, and the other two use predictions based on time varying covariates to reduce variance. Let $\theta = (\theta_1^\mathsf{T}, \ldots, \theta_T^\mathsf{T})^\mathsf{T} \in \mathbb{R}^{3T}$, where $\theta_t \in \mathbb{R}^3$. Define

$$\{\widehat{\mu}(\bar{a}), \widehat{\theta}\} = \underset{m \in \mathbb{R},\, \theta \in \mathbb{R}^{3T}}{\arg\min} \frac{1}{n} \sum_{i=1}^{n} \left\{W_{T,i}^{\bar{a}} Y_i - m - \theta^\mathsf{T} C_i^{\bar{a}}(\widehat{Q}, f)\right\}^2. \tag{4.1}$$

The fitted intercept $\widehat{\mu}(\bar{a})$ estimates $\mu(\bar{a})$.

Step C: confidence intervals. A standard variance estimator is $\widehat{V} = \frac{1}{n}\sum_{i=1}^n \widehat{\varepsilon}_i^2$, and the $1-\alpha$ Wald interval is $\widehat{\mu}(\bar{a}) \pm z_{1-\alpha/2}\sqrt{\widehat{V}/n}$. For better finite-sample properties, we recommend using the HC2 variance estimator $\widehat{V}_{\mathrm{HC2}}$ (MacKinnon & White, 1985).

4.2 Why Is Joint Regression Adjustment Desirable?

Under the conditions below, joint regression adjustment has theoretical guarantees in two respects. It supports valid inference with flexible sequential outcome predictions and has an asymptotic variance no larger than that of the HT, Hájek, and cross fitted AIPW estimators.

Why is the above task challenging? First, both the fixed pretrained predictor f and the sequential outcome predictions $\widehat{Q}_t$ can be misspecified. We seek valid inference and an asymptotic variance no larger than those of the HT, Hájek, and cross fitted AIPW estimators even when the predictions are inaccurate. Second, the internal sequential predictions are learned recursively from the current experiment, so the prediction functions used in regression adjustment are themselves random. We must show that this additional randomness does not affect the estimator's first order asymptotic distribution, allowing valid confidence intervals without an additional correction for learning the predictions. We establish this result using cross fitting (Chernozhukov et al., 2018; Newey & Robins, 2018) and the conditionally mean zero weight increments induced by sequential randomization.

Valid inference.

Assumption 2. *The number of folds K is fixed, and the random partition is independent of the samples. For every k, $|I_k|/n \xrightarrow{p} \pi_k$ for some $\pi_k \in (0,1)$.*

Assumption 3. *For every t, there exists a deterministic function $Q_t^\circ(H_t)$ with $\|Q_t^\circ\|_{P,2} < \infty$ such that $\max_{1\le k\le K} \left\|\widehat{Q}_t^{(-k)}(a_t, H_t) - Q_t^\circ(H_t)\right\|_{P,2} \xrightarrow{p} 0$, where $\|g\|_{P,2} = \{\mathbb{E}[g(H_t)^2]\}^{1/2}$.*

For a fixed path $\bar{a}$, $Q_t^\circ(H_t)$ is the L_2 limit of $\widehat{Q}_t^{(-k)}(a_t, H_t)$. It may differ from the true sequential outcome regression $Q_t(a_t, H_t)$, and this condition imposes no explicit convergence rate.

Assumption 4. *$\mathbb{E}(Y^2) < \infty$ and $\mathbb{E}\{f_t(H_t)^2\} < \infty$ for every $t = 1, \ldots, T$. Define $C_t^{\bar{a}}(Q^\circ, f) = (W_t^{\bar{a}} - W_{t-1}^{\bar{a}})\{1, Q_t^\circ(H_t), f_t(H_t)\}^\mathsf{T}$. For every $t = 1, \ldots, T$, $\mathbb{E}\left\{C_t^{\bar{a}}(Q^\circ, f) C_t^{\bar{a}}(Q^\circ, f)^\mathsf{T}\right\} \succ 0$.*

The second moment conditions ensure that the adjustment variables and the population residual have finite variances. Positive definiteness ensures that the population regression adjustment coefficient is unique.

Let $\widehat{\mu}^\circ$ denote the OLS intercept obtained from Equation (4.1) after replacing every $\widehat{Q}_t^{(-k)}(a_t, H_t)$ with $Q_t^\circ(H_t)$. Define $C^{\bar{a}}(Q^\circ, f) = \{C_1^{\bar{a}}(Q^\circ, f)^\mathsf{T}, \ldots, C_T^{\bar{a}}(Q^\circ, f)^\mathsf{T}\}^\mathsf{T}$. The population minimizer

and residual variance are

$$\theta^\star = \arg\min_{\theta\in\mathbb{R}^{3T}} \mathbb{E}\left[\left\{W_T^{\bar{a}}Y - \mu(\bar{a}) - \theta^\mathsf{T} C^{\bar{a}}(Q^\circ, f)\right\}^2\right], \qquad V = \mathbb{E}\left[\left\{W_T^{\bar{a}}Y - \mu(\bar{a}) - (\theta^\star)^\mathsf{T} C^{\bar{a}}(Q^\circ, f)\right\}^2\right].$$

Theorem 1 (Valid inference)**.** *Under Assumptions 1, 2, 3, and 4,*

$$\sqrt{n}\{\widehat{\mu}(\bar{a}) - \mu(\bar{a})\} \xrightarrow{d} \mathcal{N}(0, V), \qquad \widehat{V} \xrightarrow{p} V.$$

If $V > 0$*, then* $\mathbb{P}\left\{\mu(\bar{a}) \in \left[\widehat{\mu}(\bar{a}) - z_{1-\alpha/2}\sqrt{\widehat{V}/n}, \widehat{\mu}(\bar{a}) + z_{1-\alpha/2}\sqrt{\widehat{V}/n}\right]\right\} \longrightarrow 1-\alpha.$

The first result shows that estimation error in $\widehat{Q}_t$ does not affect the limiting distribution of $\widehat{\mu}(\bar{a})$. The proof is given in Appendix A.1.

Corollary 1 (HC2 inference)**.** *Under the conditions of Theorem 1 with* $V > 0$*,*

$$\widehat{V}_{\mathrm{HC2}} \xrightarrow{p} V, \quad \mathbb{P}\left\{\mu(\bar{a}) \in \left[\widehat{\mu}(\bar{a}) - z_{1-\alpha/2}\sqrt{\widehat{V}_{\mathrm{HC2}}/n}, \widehat{\mu}(\bar{a}) + z_{1-\alpha/2}\sqrt{\widehat{V}_{\mathrm{HC2}}/n}\right]\right\} \longrightarrow 1-\alpha.$$

No larger asymptotic variance.

We first characterize the gap between V and the semiparametric efficiency bound and then compare joint regression adjustment with standard estimators. For square integrable functions $g(H_t)$ and $\widetilde{g}(H_t)$, define $\langle g, \widetilde{g}\rangle_t = \mathbb{E}\left[\mathbb{E}\left\{\left(W_t^{\bar{a}} - W_{t-1}^{\bar{a}}\right)^2 \mid H_t\right\} g(H_t)\widetilde{g}(H_t)\right], \|g\|_t^2 = \langle g, g\rangle_t$.

Theorem 2 (Asymptotic variance decomposition)**.** *Under the assumptions of Theorem 1,*

$$V = V_{\mathrm{EIF}}^{\bar{a}} + \sum_{t=1}^{T} \inf_{\beta\in\mathbb{R}^3} \left\|Q_t(a_t, H_t) - \beta_0 - \beta_1 Q_t^\circ(H_t) - \beta_2 f_t(H_t)\right\|_t^2.$$

Thus, $V - V_{\mathrm{EIF}}^{\bar{a}}$ is the sum of the weighted residual variances from linearly approximating $Q_t(a_t, H_t)$ using 1, $Q_t^\circ(H_t)$, and $f_t(H_t)$.

Let V_Q denote the residual variance obtained by setting the coefficient of $f_t(H_t)$ to zero at every stage. The asymptotic variances of the HT and Hájek estimators are $V_{\mathrm{HT}} = \mathbb{E}\left[\{W_T^{\bar{a}}Y - \mu(\bar{a})\}^2\right], V_{\mathrm{Haj}} = \mathbb{E}\left[\left\{W_T^{\bar{a}}\big(Y - \mu(\bar{a})\big)\right\}^2\right]$. The asymptotic variance of the cross fitted AIPW estimator is $V_{\mathrm{AIPW}} = \mathbb{E}\left[\left\{W_T^{\bar{a}}Y - \mu(\bar{a}) - \sum_{t=1}^{T}\left(W_t^{\bar{a}} - W_{t-1}^{\bar{a}}\right) Q_t^\circ(H_t)\right\}^2\right]$.

Corollary 2 (No larger asymptotic variance)**.** *Under the assumptions of Theorem 2,*

$$V_{\mathrm{EIF}}^{\bar{a}} \le V \le \min\left\{V_{\mathrm{HT}}, V_{\mathrm{Haj}}, V_{\mathrm{AIPW}}, V_Q\right\}.$$

Thus, adding the fixed predictor f cannot increase the asymptotic variance relative to regression adjustment using Q° alone. The asymptotic variance is also no larger than those of the HT, Hájek, and cross fitted AIPW estimators.

When $Q_t^\circ(H_t)$ differs from the true sequential outcome regression, $f_t(H_t)$ can provide additional variance reduction. Let $\Pi_t g$ denote the weighted least squares projection of g onto $\mathcal{S}_t = \{\beta_0 + \beta_1 Q_t^\circ(H_t) : \beta_0, \beta_1 \in \mathbb{R}\}$. Define $Q_t^\perp = Q_t(a_t, H_t) - \Pi_t Q_t(a_t, H_t)$ and $f_t^\perp = f_t(H_t) - \Pi_t f_t(H_t)$.

Corollary 3 (Incremental value of pretrained predictions)**.** *Under the assumptions of Theorem 2,*

$$V_Q - V = \sum_{t=1}^{T} \frac{\langle Q_t^\perp, f_t^\perp\rangle_t^2}{\left\|f_t^\perp\right\|_t^2}.$$

The stage t term is the reduction in asymptotic variance from adding $f_t(H_t)$ after using 1 and $Q_t^\circ(H_t)$.

The decomposition also gives the condition for semiparametric efficiency.

Corollary 4 (Semiparametric efficiency)**.** *Under the assumptions of Theorem 2,* $V = V_{\mathrm{EIF}}^{\bar{a}}$ *if and only if, for every t, there exists* $\beta_t = (\beta_{t0}, \beta_{t1}, \beta_{t2})^\mathsf{T} \in \mathbb{R}^3$ *such that*

$$\left\|Q_t(a_t, H_t) - \beta_{t0} - \beta_{t1} Q_t^\circ(H_t) - \beta_{t2} f_t(H_t)\right\|_t = 0.$$

Thus, even when Q_t° differs from the true sequential outcome regression, joint adjustment can attain the semiparametric efficiency bound if a linear combination of Q_t°, the external prediction, and a constant represents the true sequential outcome regression at each stage.

5 JOINT REGRESSION ADJUSTMENT FOR CAUSAL EFFECTS BETWEEN TREATMENT PATHS

Section 4 constructs $\widehat{\mu}(\bar{a})$ for a single treatment path. We extend the method to $\tau(\bar{a}, \bar{b}) = \mu(\bar{a}) - \mu(\bar{b})$. Estimating the two average potential outcomes separately and taking their difference gives a valid estimator. However, the two estimators use the same experimental sample, so the asymptotic variance of their difference also depends on their covariance:

$$\mathrm{Var}\{\widehat{\mu}(\bar{a}) - \widehat{\mu}(\bar{b})\} = \mathrm{Var}\{\widehat{\mu}(\bar{a})\} + \mathrm{Var}\{\widehat{\mu}(\bar{b})\} - 2\,\mathrm{Cov}\{\widehat{\mu}(\bar{a}), \widehat{\mu}(\bar{b})\}.$$

Therefore, selecting the regression adjustment coefficients separately for the two paths does not generally minimize the asymptotic variance of the path contrast. We instead apply regression adjustment directly to $\tau(\bar{a}, \bar{b})$.

Step A: sequential outcome prediction. We learn sequential outcome predictions for the two treatment paths. For sample i, the predictions at stage t are $\widehat{Q}_t^{\bar{a},(-k(i))}(a_t, H_{t,i})$ and $\widehat{Q}_t^{\bar{b},(-k(i))}(b_t, H_{t,i})$, where $t = 1, \ldots, T$.

Step B: joint regression adjustment. Using the control vectors defined in Section 4, let $C^{\tau}_{i,\mathrm{full}}(\widehat{Q}, f) = \{C_i^{\bar{a}}(\widehat{Q}^{\bar{a}}, f)^{\mathsf{T}}, -C_i^{\bar{b}}(\widehat{Q}^{\bar{b}}, f)^{\mathsf{T}}\}^{\mathsf{T}}$. As detailed in the Appendix, we remove completely linearly dependent components from $C^{\tau}_{i,\mathrm{full}}(\widehat{Q}, f)$ and denote the resulting vector by $C_i^{\tau}(\widehat{Q}, f)$. Define

$$\{\widehat{\tau}(\bar{a}, \bar{b}), \widehat{\theta}_{\tau}\} = \underset{m \in \mathbb{R},\, \theta}{\arg\min} \frac{1}{n} \sum_{i=1}^{n} \left[\{W_{T,i}^{\bar{a}} - W_{T,i}^{\bar{b}}\} Y_i - m - \theta^{\mathsf{T}} C_i^{\tau}(\widehat{Q}, f) \right]^2.$$

Step C: confidence intervals. Let $\widehat{\varepsilon}_{\tau,i} = \{W_{T,i}^{\bar{a}} - W_{T,i}^{\bar{b}}\} Y_i - \widehat{\tau}(\bar{a}, \bar{b}) - \widehat{\theta}_{\tau}^{\mathsf{T}} C_i^{\tau}(\widehat{Q}, f)$. Define $\widehat{V}_{\tau} = n^{-1} \sum_{i=1}^{n} \widehat{\varepsilon}_{\tau,i}^2$. The Wald interval is $\widehat{\tau}(\bar{a}, \bar{b}) \pm z_{1-\alpha/2} \sqrt{\widehat{V}_{\tau}/n}$. As in Section 4, we can also use the HC2 variance estimator $\widehat{V}_{\mathrm{HC2},\tau}$.

Assumption 5 (Path contrast conditions)**.** *For every t, there exist deterministic functions $Q_t^{\bar{a},\circ}(H_t)$ and $Q_t^{\bar{b},\circ}(H_t)$ with finite $L_2(P)$ norms such that* $\max_{1 \le k \le K} \left\| \widehat{Q}_t^{\bar{a},(-k)}(a_t, H_t) - Q_t^{\bar{a},\circ}(H_t) \right\|_{P,2} \xrightarrow{p} 0, \max_{1 \le k \le K} \left\| \widehat{Q}_t^{\bar{b},(-k)}(b_t, H_t) - Q_t^{\bar{b},\circ}(H_t) \right\|_{P,2} \xrightarrow{p} 0$. *Moreover,* $\mathbb{E}(Y^2) < \infty$ *and* $\mathbb{E}\{f_t(H_t)^2\} < \infty$. *Let $Q^{\circ} = (Q^{\bar{a},\circ}, Q^{\bar{b},\circ})$, and define $C^{\tau}(Q^{\circ}, f)$ by setting $\widehat{Q} = Q^{\circ}$ in $C_i^{\tau}(\widehat{Q}, f)$. We assume that* $\mathbb{E}\{C^{\tau}(Q^{\circ}, f) C^{\tau}(Q^{\circ}, f)^{\mathsf{T}}\} \succ 0$.

Define the population residual variance by $V_{\tau} = \inf_{\theta} \mathbb{E}\left[\left\{ (W_T^{\bar{a}} - W_T^{\bar{b}}) Y - \tau(\bar{a}, \bar{b}) - \theta^{\mathsf{T}} C^{\tau}(Q^{\circ}, f) \right\}^2 \right]$.

Let $V_{\mathrm{HT},\tau}$, $V_{\mathrm{Haj},\tau}$, and $V_{\mathrm{AIPW},\tau}$ denote the asymptotic variances of the HT, Hájek, and cross fitted AIPW path contrast estimators. Let $V_{Q,\tau}$ denote the asymptotic variance of regression adjustment for $\tau(\bar{a}, \bar{b})$ obtained by omitting f. Let $V_{\mathrm{EIF},\tau}$ denote the semiparametric efficiency bound, given by $V_{\mathrm{EIF},\tau} = \mathrm{Var}\{\phi_{\mathrm{EIF}}^{\bar{a}}(O) - \phi_{\mathrm{EIF}}^{\bar{b}}(O)\}$.

Theorem 3 (Valid inference and no larger asymptotic variance)**.** *Under Assumptions 1, 2, and 5,*

$$\sqrt{n}\{\widehat{\tau}(\bar{a}, \bar{b}) - \tau(\bar{a}, \bar{b})\} \xrightarrow{d} \mathcal{N}(0, V_{\tau}),$$
$$\widehat{V}_{\tau} \xrightarrow{p} V_{\tau}, \qquad \widehat{V}_{\mathrm{HC2},\tau} \xrightarrow{p} V_{\tau},$$
$$V_{\mathrm{EIF},\tau} \le V_{\tau} \le \min\{V_{\mathrm{HT},\tau}, V_{\mathrm{Haj},\tau}, V_{\mathrm{AIPW},\tau}, V_{Q,\tau}\}.$$

6 EXPERIMENTS

Synthetic experiments. We use data based on the Lorenz–96 model (Karimi & Paul, 2010), following the simulation framework of Chihara et al. (2026). This framework generates nonlinear longitudinal data in which treatments affect subsequent covariates and outcomes. We use sequential randomization for treatment assignment in the current experiment and train a fixed predictor on external observational data with different treatment assignment probabilities.

We consider experiments with $T = 6$ treatment stages and $n = 5000$ observations, evaluated over 400 Monte Carlo replications. The four estimands are $\mu_1 = \mu(101010)$, $\mu_2 = \mu(111000)$, $\tau_1 = \mu(111111) - \mu(000000)$, and $\tau_2 = \mu(100000) - \mu(000001)$. We compare our regression adjustment estimators with and without the external predictor f with HT, Hájek, cross fitted AIPW (Bang & Robins, 2005; Chernozhukov et al., 2018), cross fitted LTMLE (van der Laan & Gruber, 2012), and DynamicRA (Chihara et al., 2026). We also compare with AIPW and LTMLE that use f as an additional feature in the outcome regressions. Our estimators and the AIPW and LTMLE estimators use LightGBM with two folds for cross fitting of the sequential outcome regressions. We use HC2 variance estimation to construct 95% confidence intervals for our methods. The full data generating mechanism, learner specifications, implementation details, and additional experiments are provided in Appendix C.

Table 1: Results with $n = 5000$, $T = 6$, and 400 replications. MSE $\times 10^3$.

	μ_1			μ_2			τ_1			τ_2		
Method	MSE	Cov.	Length	MSE	Cov.	Length	MSE	Cov.	Length	MSE	Cov.	Length
Without f												
HT	263.65	95.25	1.95	265.40	93.75	1.93	523.84	94.25	2.80	516.47	95.50	2.80
Hájek	20.39	95.50	0.56	19.20	96.00	0.56	40.05	95.00	0.79	46.36	94.25	0.79
CF-AIPW	1.87	94.75	0.17	1.81	95.00	0.17	3.05	97.00	0.23	3.66	94.50	0.23
CF-LTMLE	1.89	95.00	0.17	1.83	95.25	0.17	2.97	96.75	0.23	3.69	94.50	0.23
DynamicRA	6.70	96.25	0.33	7.16	94.50	0.33	12.98	96.00	0.47	14.55	94.25	0.46
Regression adjustment	1.84	95.75	0.17	1.81	95.00	0.17	2.80	97.00	0.23	3.49	93.75	0.22
With f												
CF-AIPW	1.09	96.00	0.13	1.00	96.50	0.13	1.48	97.00	0.16	1.79	94.75	0.16
CF-LTMLE	1.10	96.00	0.13	1.02	95.50	0.13	1.49	96.00	0.16	1.78	95.00	0.16
Joint regression adjustment	**0.98**	97.25	**0.13**	**0.93**	96.75	**0.13**	**1.33**	96.00	**0.16**	**1.58**	96.00	**0.16**

Results. (i) Our estimators achieve coverage close to the nominal 95% level across the four estimands. (ii) Regression adjustment without f substantially reduces MSE and mean confidence interval length compared with HT, Hájek, and DynamicRA, and is competitive with CF-AIPW and CF-LTMLE. Adding f further reduces MSE by 46.9%–54.8% and shortens the mean confidence intervals. (iii) Joint regression adjustment achieves the lowest MSE across all four estimands, reducing MSE by 6.7%–11.9% compared with CF-AIPW and CF-LTMLE that use f as an additional feature, with shorter or comparable mean confidence intervals.

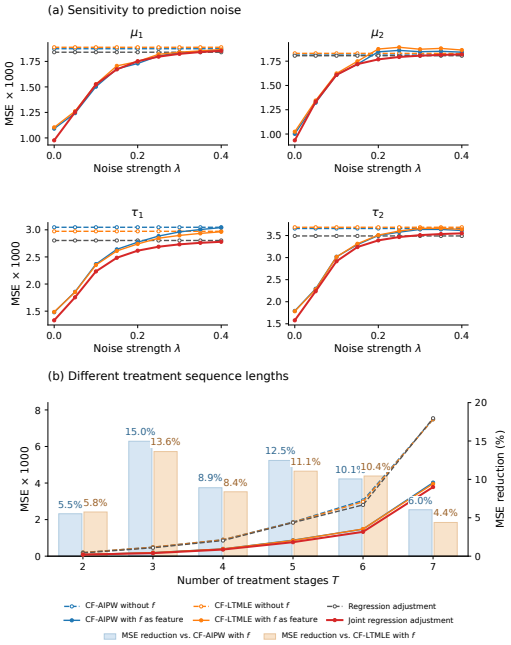

Figure 2: (a) Sensitivity to prediction noise. (b) Different treatment sequence lengths.

Sensitivity to the pretrained predictor. Keeping the experimental data and the pretrained predictor fixed, we perturb the predictions as $f_\lambda = f(1 + \lambda Z)$, where Z is independent noise and $\lambda \in [0, 0.4]$. Figure 2(a) shows that, as noise increases, the precision gains from the pretrained predictor diminish, and the performance of joint regression adjustment approaches that of regression adjustment without f. Across increasing values of λ, joint regression adjustment broadly preserves the estimation precision obtained without f, consistent with our theoretical guarantee of no larger asymptotic variance. Compared with CF-AIPW and CF-LTMLE that use f_λ as an additional feature, joint regression adjustment is competitive.

Different treatment sequence lengths. We vary the treatment sequence length from $T = 2$ to $T = 7$ and estimate $\tau_T = \mu(1^T) - \mu(0^T)$, using $n = 5000$ observations and 400 Monte Carlo replications for each setting. Figure 2(b) shows that MSE increases with T for all methods. Joint

regression adjustment achieves the lowest MSE among the six methods shown at every sequence length considered, reducing MSE by 5.5%–15.0% relative to CF-AIPW and 4.4%–13.6% relative to CF-LTMLE when both use f as an additional feature.

Real-world data. To evaluate estimation precision and confidence interval coverage, we conduct repeated A/A randomization experiments on longitudinal user records from a large online gaming platform. Both treatment labels correspond to the same product experience, so the true treatment effect is zero. This allows us to directly evaluate estimation error and coverage while preserving real user behavior (Guo et al., 2021; Deng et al., 2017).

We consider $T = 5$ stages and conduct 1000 replications. Each replication samples $n = 5000$ users without replacement from a pool of 16,123 users and independently assigns treatment labels with probability $1/2$ at each stage. At each stage, we use 20 time varying covariates available before treatment assignment, covering player feedback, activity, game mode preferences, gameplay performance, and progression. The outcome Y is the cumulative number of game sessions during the five days following the experimental window, and the true treatment effect is $\tau = \mathbb{E}[Y(11111) - Y(00000)] = 0$. The internal sequential outcome regressions use five-fold cross-fitting and LightGBM. The external predictor is a Transformer regression model trained on behavioral sequences from 816,881 users outside the evaluation pool. It predicts the cumulative five-day outcome using pretreatment histories and remains fixed across replications.

We compare HT, Hájek, CF-AIPW, CF-LTMLE, and our regression adjustment estimators with and without the external predictor. We also compare with CF-AIPW and CF-LTMLE that use the external predictions as additional features in the outcome regressions. We use HC2 variance estimation to construct 95% confidence intervals for our two regression adjustment estimators. Table 2 reports results.

Table 2: Results for the real A/A experiment with 1000 replications.

Method	Empirical variance	Coverage (%)	Mean CI length
Without f			
HT	217.95	95.4	57.85
Hájek	110.33	94.9	40.91
CF-AIPW	78.22	95.6	34.69
CF-LTMLE	78.19	94.7	34.53
Regression adjustment	**77.88**	94.8	**34.29**
With f			
CF-AIPW	74.06	95.4	33.61
CF-LTMLE	74.22	94.8	33.44
Joint regression adjustment	**69.54**	95.3	**32.73**

Results. (i) All methods achieve coverage close to the nominal 95% level. (ii) Regression adjustment without f reduces empirical variance by 29.4% compared with Hájek and performs comparably to CF-AIPW and CF-LTMLE. Adding f further reduces empirical variance by 10.7%. (iii) Joint regression adjustment achieves the lowest empirical variance and shortest mean confidence intervals among the methods compared. It reduces empirical variance by 6.1%–6.3% compared with CF-AIPW and CF-LTMLE that use f as an additional feature.

7 Conclusion

We propose joint regression adjustment for sequential randomized experiments, combining internal sequential outcome predictions with fixed pretrained predictions. We establish valid asymptotic inference and variance guarantees under prediction misspecification and characterize the additional variance reduction from external predictions. Simulations and an A/A experiment on a gaming platform demonstrate gains in estimation precision.

Our theory considers a fixed number of stages and assumes no interference between individuals. As the number of stages grows, fewer observations follow a given treatment path, which can increase the variance of the estimator. Interference between individuals may also arise in online platforms, where one user's treatment can affect other users' outcomes. Extending our method to these settings is a direction for future work.

AI USE STATEMENT

Language models assisted with early drafting of parts of this paper, code implementation, language polishing of parts of the proofs, and reference searches. The research ideas and all proofs were developed by the authors. References were verified and selected by the authors, and all AI-assisted content was reviewed and verified by the authors.

ETHICS STATEMENT

We adhere to the ICLR Code of Ethics. This study uses synthetic data and non-public user records provided by an online gaming platform.

REPRODUCIBILITY STATEMENT

We provide code for the synthetic experiments in the supplementary materials. The main text describes the research assumptions and estimation procedure, and the appendix provides complete proofs, experimental details, and additional experiments.

REFERENCES

Anastasios N. Angelopoulos, Stephen Bates, Clara Fannjiang, Michael I. Jordan, and Tijana Zrnic. Prediction-powered inference. *Science*, 382(6671):669–674, 2023.

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

Samuel L. Battalio, David E. Conroy, Walter Dempsey, Peng Liao, Marianne Menictas, Susan Murphy, Inbal Nahum-Shani, Tianchen Qian, Santosh Kumar, and Bonnie Spring. Sense2Stop: A micro-randomized trial using wearable sensors to optimize a just-in-time-adaptive stress management intervention for smoking relapse prevention. *Contemporary Clinical Trials*, 109:106534, 2021.

Ioana Bica, Ahmed M. Alaa, James Jordon, and Mihaela van der Schaar. Estimating counterfactual treatment outcomes over time through adversarially balanced representations. In *International Conference on Learning Representations*, 2020.

Adam Bloniarz, Hanzhong Liu, Cun-Hui Zhang, Jasjeet S. Sekhon, and Bin Yu. Lasso adjustments of treatment effect estimates in randomized experiments. *Proceedings of the National Academy of Sciences*, 113(27):7383–7390, 2016.

Hugo Bodory, Martin Huber, and Lukáš Lafférs. Evaluating (weighted) dynamic treatment effects by double machine learning. *The Econometrics Journal*, 25(3):628–648, 2022.

Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1):C1–C68, 2018.

Naoki Chihara, Tatsushi Oka, Yasuko Matsubara, Yasushi Sakurai, and Shota Yasui. Modeling covariate transition for efficient estimation of longitudinal treatment effects in randomized experiments. In *The 43rd International Conference on Machine Learning*, 2026.

Peter L. Cohen and Colin B. Fogarty. No-harm calibration for generalized Oaxaca–Blinder estimators. *Biometrika*, 111(1):331–338, 2024.

Piersilvio De Bartolomeis, Javier Abad, Guanbo Wang, Konstantin Donhauser, Raymond M. Duch, Fanny Yang, and Issa J. Dahabreh. Efficient randomized experiments using foundation models. In *Advances in Neural Information Processing Systems*, volume 38, pp. 159470–159509, 2025.

Alex Deng, Ya Xu, Ron Kohavi, and Toby Walker. Improving the sensitivity of online controlled experiments by utilizing pre-experiment data. In *Proceedings of the Sixth ACM International Conference on Web Search and Data Mining*, pp. 123–132, 2013.

Alex Deng, Jiannan Lu, and Jonathan Litz. Trustworthy analysis of online A/B tests: Pitfalls, challenges and solutions. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*, pp. 641–649, 2017.

Dennis Frauen, Tobias Hatt, Valentyn Melnychuk, and Stefan Feuerriegel. Estimating average causal effects from patient trajectories. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37 (6):7586–7594, 2023.

Dennis Frauen, Konstantin Hess, and Stefan Feuerriegel. Model-agnostic meta-learners for estimating heterogeneous treatment effects over time. In *The Thirteenth International Conference on Learning Representations*, 2025.

Yongyi Guo, Dominic Coey, Mikael Konutgan, Wenting Li, Chris Schoener, and Matt Goldman. Machine learning for variance reduction in online experiments. In *Advances in Neural Information Processing Systems*, volume 34, 2021.

Nathan Kallus and Masatoshi Uehara. Intrinsically efficient, stable, and bounded off-policy evaluation for reinforcement learning. In *Advances in Neural Information Processing Systems*, volume 32, 2019.

A. Karimi and M. R. Paul. Extensive chaos in the Lorenz-96 model. *Chaos*, 20(4):043105, 2010.

Lihua Lei and Peng Ding. Regression adjustment in completely randomized experiments with a diverging number of covariates. *Biometrika*, 108(4):815–828, 2021.

Greg Lewis and Vasilis Syrgkanis. Double/debiased machine learning for dynamic treatment effects. In *Advances in Neural Information Processing Systems*, volume 34, 2021.

Sida Li and Nikolaos Ignatiadis. Prediction-powered adaptive shrinkage estimation. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 34836–34875, 2025.

Winston Lin. Agnostic notes on regression adjustments to experimental data: Reexamining Freedman's critique. *The Annals of Applied Statistics*, 7(1):295–318, 2013.

James G. MacKinnon and Halbert White. Some heteroskedasticity-consistent covariance matrix estimators with improved finite sample properties. *Journal of Econometrics*, 29(3):305–325, 1985.

Valentyn Melnychuk, Dennis Frauen, and Stefan Feuerriegel. Causal transformer for estimating counterfactual outcomes. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 15293–15329, 2022.

Whitney K. Newey and James M. Robins. Cross-fitting and fast remainder rates for semiparametric estimation. *arXiv preprint arXiv:1801.09138*, 2018.

Miruna Oprescu, David Park, Xihaier Luo, Shinjae Yoo, and Nathan Kallus. GST-UNet: A neural framework for spatiotemporal causal inference with time-varying confounding. In *Advances in Neural Information Processing Systems*, volume 38, Main Conference, pp. 17295–17322, 2025.

James M. Robins. 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, 1986.

James M. Robins. Correcting for non-compliance in randomized trials using structural nested mean models. *Communications in Statistics—Theory and Methods*, 23(8):2379–2412, 1994.

James M. Robins and Miguel A. Hernán. Estimation of the causal effects of time-varying exposures. In Garrett Fitzmaurice, Marie Davidian, Geert Verbeke, and Geert Molenberghs (eds.), *Longitudinal Data Analysis*, chapter 23. Chapman & Hall/CRC, 2009.

James M. Robins, Miguel A. Hernán, and Babette Brumback. Marginal structural models and causal inference in epidemiology. *Epidemiology*, 11(5):550–560, 2000.

Daniel B. Rubin and Mark J. van der Laan. Empirical efficiency maximization: Improved locally efficient covariate adjustment in randomized experiments and survival analysis. *The International Journal of Biostatistics*, 4:Article 5, 2008.

Alejandro Schuler, David Walsh, Diana Hall, Jon Walsh, and Charles Fisher. Increasing the efficiency of randomized trial estimates via linear adjustment for a prognostic score. *The International Journal of Biostatistics*, 18(2):329–356, 2022.

Chengchun Shi, Xiaoyu Wang, Shikai Luo, Hongtu Zhu, Jieping Ye, and Rui Song. Dynamic causal effects evaluation in A/B testing with a reinforcement learning framework. *Journal of the American Statistical Association*, 118(543):2059–2071, 2023.

Toru Shirakawa, Yi Li, Yulun Wu, Sky Qiu, Yuxuan Li, Mingduo Zhao, Hiroyasu Iso, and Mark J. van der Laan. Longitudinal targeted minimum loss-based estimation with temporal-difference heterogeneous transformer. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 45097–45113, 2024.

Anastasios A. Tsiatis. *Semiparametric Theory and Missing Data*. Springer, 2006.

Anastasios A. Tsiatis, Marie Davidian, Min Zhang, and Xiaomin Lu. Covariate adjustment for two-sample treatment comparisons in randomized clinical trials: A principled yet flexible approach. *Statistics in Medicine*, 27(23):4658–4677, 2008.

Mark J. van der Laan and Susan Gruber. Targeted minimum loss based estimation of causal effects of multiple time point interventions. *The International Journal of Biostatistics*, 8(1), 2012.

Stefan Wager, Wenfei Du, Jonathan Taylor, and Robert J. Tibshirani. High-dimensional regression adjustments in randomized experiments. *Proceedings of the National Academy of Sciences*, 113 (45):12673–12678, 2016.

Yuxin Wang, Maresa Schröder, Dennis Frauen, Jonas Schweisthal, Konstantin Hess, and Stefan Feuerriegel. Constructing confidence intervals for average treatment effects from multiple datasets. In *The Thirteenth International Conference on Learning Representations*, 2025.

Tijana Zrnic and Emmanuel J. Candès. Cross-prediction-powered inference. *Proceedings of the National Academy of Sciences*, 121(15):e2322083121, 2024.

A PROOFS

A.1 PROOF OF THEOREM 1

Proof. Step 1: OLS coefficient consistency and asymptotic normality. Throughout this proof, we suppress the dependence on $\bar{a}$ and write $W_t = W_t^{\bar{a}}$, $C_t(\cdot) = C_t^{\bar{a}}(\cdot)$, $C_i(\cdot) = C_i^{\bar{a}}(\cdot)$, $\mu = \mu(\bar{a})$, and $\widehat{\mu} = \widehat{\mu}(\bar{a})$. For sample i, let $C_i^\circ = C_i(Q^\circ, f)$ and $\widehat{C}_i = C_i(\widehat{Q}, f)$. Let $\widehat{\theta}^\circ$ denote the OLS slope obtained from Equation (4.1) after replacing every $\widehat{Q}_t^{(-k(i))}(a_t, H_{t,i})$ with $Q_t^\circ(H_{t,i})$.

For every t,

$$\mathbb{E}\{C_t(Q^\circ, f)\} = \mathbb{E}\left[\begin{pmatrix} 1 \\ Q_t^\circ(H_t) \\ f_t(H_t) \end{pmatrix} \mathbb{E}\{W_t - W_{t-1} \mid H_t\}\right] = 0.$$

Therefore, $\mathbb{E}(C_i^\circ) = 0$. Since $\mathbb{E}(W_T Y) = \mu$, the population OLS intercept based on C_i° is μ.

For $s < t$, $C_s(Q^\circ, f)$ and $\{1, Q_t^\circ(H_t), f_t(H_t)\}^\mathsf{T}$ are measurable with respect to H_t. Therefore,

$$\mathbb{E}\{C_s(Q^\circ, f) C_t(Q^\circ, f)^\mathsf{T}\} = \mathbb{E}\left[C_s(Q^\circ, f)\,(1 \quad Q_t^\circ(H_t) \quad f_t(H_t))\,\mathbb{E}\{W_t - W_{t-1} \mid H_t\}\right] = 0.$$

Let $\Sigma^\circ = \mathbb{E}\{C_i^\circ (C_i^\circ)^\mathsf{T}\}$. The preceding equality shows that the off-diagonal stage-specific blocks of Σ° are zero. Assumption 4 makes each diagonal block positive definite, and hence $\Sigma^\circ \succ 0$.

Let $n_k = |I_k|$ and $\mathcal{F}_k = \sigma\{I_1, \ldots, I_K, O_j : j \notin I_k\}$. Conditional on $\mathcal{F}_k$, $\widehat{Q}_t^{(-k)}$ is fixed and the samples in I_k are independent. For every $i \in I_k$ and t,

$$C_{t,i}(\widehat{Q}, f) - C_{t,i}(Q^\circ, f) = (W_{t,i} - W_{t-1,i}) \begin{pmatrix} 0 \\ \widehat{Q}_t^{(-k)}(a_t, H_{t,i}) - Q_t^\circ(H_{t,i}) \\ 0 \end{pmatrix}.$$

Since $\mathbb{E}(W_t - W_{t-1} \mid H_t, \mathcal{F}_k) = \mathbb{E}(W_t - W_{t-1} \mid H_t) = 0$,

$$\mathbb{E}\left[(W_{t,i} - W_{t-1,i})\{\widehat{Q}_t^{(-k)}(a_t, H_{t,i}) - Q_t^\circ(H_{t,i})\} \middle| H_{t,i}, \mathcal{F}_k\right] = 0.$$

By positivity and fixed T, there is a constant $B < \infty$ such that $|W_t - W_{t-1}| \le B$ for every t. Let

$$S_{t,k,n} = \frac{1}{\sqrt{n_k}} \sum_{i \in I_k} (W_{t,i} - W_{t-1,i})\{\widehat{Q}_t^{(-k)}(a_t, H_{t,i}) - Q_t^\circ(H_{t,i})\}, \quad R_{t,k,n} = B^2 \left\| \widehat{Q}_t^{(-k)}(a_t, H_t) - Q_t^\circ(H_t) \right\|_{P,2}^2.$$

Conditional independence within I_k and the displayed conditional mean-zero equality give

$$\mathbb{E}(S_{t,k,n}^2 \mid \mathcal{F}_k) \le R_{t,k,n} = o_\mathbb{P}(1).$$

For arbitrary $\delta, \eta > 0$, conditional Chebyshev's inequality gives

$$\begin{aligned} \mathbb{P}(|S_{t,k,n}| > \delta) &= \mathbb{E}\left[\mathbb{P}(|S_{t,k,n}| > \delta \mid \mathcal{F}_k)\mathbf{1}(R_{t,k,n} \le \eta)\right] + \mathbb{E}\left[\mathbb{P}(|S_{t,k,n}| > \delta \mid \mathcal{F}_k)\mathbf{1}(R_{t,k,n} > \eta)\right] \\ &\le \frac{\eta}{\delta^2} + \mathbb{P}(R_{t,k,n} > \eta). \end{aligned}$$

Since $R_{t,k,n} = o_\mathbb{P}(1)$, $\limsup_{n\to\infty} \mathbb{P}(|S_{t,k,n}| > \delta) \le \frac{\eta}{\delta^2}$.

Letting $\eta \downarrow 0$ gives $S_{t,k,n} = o_\mathbb{P}(1)$. Since T and K are fixed and $n_k/n \xrightarrow{p} \pi_k$,

$$\frac{1}{\sqrt{n}} \sum_{i=1}^n (\widehat{C}_i - C_i^\circ) = \sum_{k=1}^K \sqrt{\frac{n_k}{n}} \frac{1}{\sqrt{n_k}} \sum_{i \in I_k} (\widehat{C}_i - C_i^\circ) = o_\mathbb{P}(1). \tag{A.1}$$

For every k, let

$$\Lambda_{k,n} = \frac{1}{n_k} \sum_{i \in I_k} \left\| \widehat{C}_i - C_i^\circ \right\|^2.$$

By the expression for $C_{t,i}(\widehat{Q}, f) - C_{t,i}(Q^\circ, f)$ and the definition of $R_{t,k,n}$,

$$\mathbb{E}(\Lambda_{k,n} \mid \mathcal{F}_k) \le \sum_{t=1}^T R_{t,k,n} = o_\mathbb{P}(1).$$

For arbitrary $\delta, \eta > 0$, conditional Markov's inequality gives

$$\mathbb{P}(\Lambda_{k,n} > \delta) = \mathbb{E}\left[\mathbb{P}(\Lambda_{k,n} > \delta \mid \mathcal{F}_k)\mathbf{1}\left\{\sum_{t=1}^{T} R_{t,k,n} \le \eta\right\}\right] + \mathbb{E}\left[\mathbb{P}(\Lambda_{k,n} > \delta \mid \mathcal{F}_k)\mathbf{1}\left\{\sum_{t=1}^{T} R_{t,k,n} > \eta\right\}\right]$$
$$\le \frac{\eta}{\delta} + \mathbb{P}\left\{\sum_{t=1}^{T} R_{t,k,n} > \eta\right\}.$$

Since T is fixed and $R_{t,k,n} = o_{\mathbb{P}}(1)$, $\limsup_{n\to\infty} \mathbb{P}(\Lambda_{k,n} > \delta) \le \frac{\eta}{\delta}$.

Letting $\eta \downarrow 0$ gives $\Lambda_{k,n} = o_{\mathbb{P}}(1)$. Since K is fixed and $n_k/n \xrightarrow{p} \pi_k$,

$$\frac{1}{n}\sum_{i=1}^{n} \left\|\widehat{C}_i - C_i^{\circ}\right\|^2 = \sum_{k=1}^{K} \frac{n_k}{n}\Lambda_{k,n} = o_{\mathbb{P}}(1). \tag{A.2}$$

Let $\overline{\widehat{C}} = \frac{1}{n}\sum_{i=1}^{n} \widehat{C}_i$, $\overline{C}^{\circ} = \frac{1}{n}\sum_{i=1}^{n} C_i^{\circ}$, and $\overline{W_T Y} = \frac{1}{n}\sum_{i=1}^{n} W_{T,i} Y_i$. Define

$$\widehat{\Sigma} = \frac{1}{n}\sum_{i=1}^{n}(\widehat{C}_i - \overline{\widehat{C}})(\widehat{C}_i - \overline{\widehat{C}})^{\mathsf{T}}, \qquad \widehat{\Sigma}^{\circ} = \frac{1}{n}\sum_{i=1}^{n}(C_i^{\circ} - \overline{C}^{\circ})(C_i^{\circ} - \overline{C}^{\circ})^{\mathsf{T}}.$$

The OLS normal equations give

$$\theta^{\star} = (\Sigma^{\circ})^{-1}\mathbb{E}\{C_i^{\circ}(W_T Y - \mu)\}, \qquad \widehat{\theta} = \widehat{\Sigma}^{-1}\frac{1}{n}\sum_{i=1}^{n}(\widehat{C}_i - \overline{\widehat{C}})(W_{T,i}Y_i - \overline{W_T Y}),$$

$$\widehat{\theta}^{\circ} = (\widehat{\Sigma}^{\circ})^{-1}\frac{1}{n}\sum_{i=1}^{n}(C_i^{\circ} - \overline{C}^{\circ})(W_{T,i}Y_i - \overline{W_T Y}).$$

For every i,

$$\widehat{C}_i\widehat{C}_i^{\mathsf{T}} - C_i^{\circ}(C_i^{\circ})^{\mathsf{T}} = (\widehat{C}_i - C_i^{\circ})(C_i^{\circ})^{\mathsf{T}} + C_i^{\circ}(\widehat{C}_i - C_i^{\circ})^{\mathsf{T}} + (\widehat{C}_i - C_i^{\circ})(\widehat{C}_i - C_i^{\circ})^{\mathsf{T}}.$$

Therefore, by Cauchy–Schwarz, Equation (A.2), and the law of large numbers,

$$\left\|\frac{1}{n}\sum_{i=1}^{n}\{\widehat{C}_i\widehat{C}_i^{\mathsf{T}} - C_i^{\circ}(C_i^{\circ})^{\mathsf{T}}\}\right\|$$
$$\le 2\underbrace{\left\{\frac{1}{n}\sum_{i=1}^{n}\left\|\widehat{C}_i - C_i^{\circ}\right\|^2\right\}^{1/2}}_{o_{\mathbb{P}}(1)}\underbrace{\left\{\frac{1}{n}\sum_{i=1}^{n}\|C_i^{\circ}\|^2\right\}^{1/2}}_{O_{\mathbb{P}}(1)} + \underbrace{\frac{1}{n}\sum_{i=1}^{n}\left\|\widehat{C}_i - C_i^{\circ}\right\|^2}_{o_{\mathbb{P}}(1)} = o_{\mathbb{P}}(1).$$

By Equation (A.1), $\sqrt{n}(\overline{\widehat{C}} - \overline{C}^{\circ}) = o_{\mathbb{P}}(1)$. Since $\mathbb{E}(C_i^{\circ}) = 0$ and C_i° has a finite second moment, $\overline{C}^{\circ} = O_{\mathbb{P}}(n^{-1/2})$. Therefore,

$$\left\|\overline{\widehat{C}}(\overline{\widehat{C}})^{\mathsf{T}} - \overline{C}^{\circ}(\overline{C}^{\circ})^{\mathsf{T}}\right\| \le 2\underbrace{\left\|\overline{\widehat{C}} - \overline{C}^{\circ}\right\|}_{o_{\mathbb{P}}(n^{-1/2})}\underbrace{\left\|\overline{C}^{\circ}\right\|}_{O_{\mathbb{P}}(n^{-1/2})} + \underbrace{\left\|\overline{\widehat{C}} - \overline{C}^{\circ}\right\|^2}_{o_{\mathbb{P}}(n^{-1})} = o_{\mathbb{P}}(1).$$

By the definitions of $\widehat{\Sigma}$ and $\widehat{\Sigma}^{\circ}$,

$$\widehat{\Sigma} - \widehat{\Sigma}^{\circ} = \underbrace{\frac{1}{n}\sum_{i=1}^{n}\{\widehat{C}_i\widehat{C}_i^{\mathsf{T}} - C_i^{\circ}(C_i^{\circ})^{\mathsf{T}}\}}_{o_{\mathbb{P}}(1)} - \underbrace{\left\{\overline{\widehat{C}}(\overline{\widehat{C}})^{\mathsf{T}} - \overline{C}^{\circ}(\overline{C}^{\circ})^{\mathsf{T}}\right\}}_{o_{\mathbb{P}}(1)} = o_{\mathbb{P}}(1).$$

Since $\sum_{i=1}^{n}(W_{T,i}Y_i - \overline{W_T Y}) = 0$,

$$\frac{1}{n}\sum_{i=1}^{n}(\widehat{C}_i - \overline{\widehat{C}})(W_{T,i}Y_i - \overline{W_T Y}) - \frac{1}{n}\sum_{i=1}^{n}(C_i^{\circ} - \overline{C}^{\circ})(W_{T,i}Y_i - \overline{W_T Y})$$
$$= \frac{1}{n}\sum_{i=1}^{n}(\widehat{C}_i - C_i^{\circ})(W_{T,i}Y_i - \overline{W_T Y}).$$

By Cauchy–Schwarz,

$$\left\| \frac{1}{n}\sum_{i=1}^{n}(\widehat{C}_i - C_i^\circ)(W_{T,i}Y_i - \overline{W_T Y}) \right\| \leq \underbrace{\left\{ \frac{1}{n}\sum_{i=1}^{n} \left\| \widehat{C}_i - C_i^\circ \right\|^2 \right\}^{1/2}}_{o_{\mathbb{P}}(1)} \underbrace{\left\{ \frac{1}{n}\sum_{i=1}^{n}(W_{T,i}Y_i - \overline{W_T Y})^2 \right\}^{1/2}}_{O_{\mathbb{P}}(1)} = o_{\mathbb{P}}(1).$$

The law of large numbers gives

$$\widehat{\Sigma}^\circ \xrightarrow{p} \Sigma^\circ, \qquad \frac{1}{n}\sum_{i=1}^{n}(C_i^\circ - \overline{C}^\circ)(W_{T,i}Y_i - \overline{W_T Y}) \xrightarrow{p} \mathbb{E}\{C_i^\circ(W_T Y - \mu)\}.$$

Therefore,

$$\widehat{\Sigma} - \Sigma^\circ = (\widehat{\Sigma} - \widehat{\Sigma}^\circ) + (\widehat{\Sigma}^\circ - \Sigma^\circ) \xrightarrow{p} 0,$$

and

$$\frac{1}{n}\sum_{i=1}^{n}(\widehat{C}_i - \overline{\widehat{C}})(W_{T,i}Y_i - \overline{W_T Y}) \xrightarrow{p} \mathbb{E}\{C_i^\circ(W_T Y - \mu)\}.$$

Since $\Sigma^\circ \succ 0$, $\widehat{\Sigma}$ and $\widehat{\Sigma}^\circ$ are nonsingular with probability tending to one. The OLS normal equations give

$$\widehat{\theta} \xrightarrow{p} \theta^\star, \qquad \widehat{\theta}^\circ \xrightarrow{p} \theta^\star, \qquad \widehat{\theta} - \widehat{\theta}^\circ = o_{\mathbb{P}}(1).$$

The OLS intercept identities are

$$\widehat{\mu} = \overline{W_T Y} - \widehat{\theta}^{\mathsf{T}}\overline{\widehat{C}}, \qquad \widehat{\mu}^\circ = \overline{W_T Y} - (\widehat{\theta}^\circ)^{\mathsf{T}}\overline{C}^\circ.$$

Therefore,

$$\sqrt{n}(\widehat{\mu} - \widehat{\mu}^\circ) = -(\widehat{\theta} - \widehat{\theta}^\circ)^{\mathsf{T}}\sqrt{n}\,\overline{\widehat{C}} - (\widehat{\theta}^\circ)^{\mathsf{T}}\sqrt{n}(\overline{\widehat{C}} - \overline{C}^\circ).$$

Since $\mathbb{E}(C_i^\circ) = 0$ and C_i° has a finite second moment, $\sqrt{n}\,\overline{C}^\circ = O_{\mathbb{P}}(1)$. By Equation (A.1), $\sqrt{n}(\overline{\widehat{C}} - \overline{C}^\circ) = o_{\mathbb{P}}(1)$, and hence $\sqrt{n}\,\overline{\widehat{C}} = O_{\mathbb{P}}(1)$. Since $\widehat{\theta} - \widehat{\theta}^\circ = o_{\mathbb{P}}(1)$ and $\widehat{\theta}^\circ = O_{\mathbb{P}}(1)$,

$$\sqrt{n}(\widehat{\mu} - \widehat{\mu}^\circ) = -\underbrace{(\widehat{\theta} - \widehat{\theta}^\circ)^{\mathsf{T}}}_{o_{\mathbb{P}}(1)}\underbrace{\sqrt{n}\,\overline{\widehat{C}}}_{O_{\mathbb{P}}(1)} - \underbrace{(\widehat{\theta}^\circ)^{\mathsf{T}}}_{O_{\mathbb{P}}(1)}\underbrace{\sqrt{n}(\overline{\widehat{C}} - \overline{C}^\circ)}_{o_{\mathbb{P}}(1)} = o_{\mathbb{P}}(1).$$

Define $\varepsilon_i^\star = W_{T,i}Y_i - \mu - (\theta^\star)^{\mathsf{T}}C_i^\circ$. Since $\mathbb{E}(C_i^\circ) = 0$, $\mathbb{E}(\varepsilon_i^\star) = 0$, and, by the definition of V, $\mathbb{E}\{(\varepsilon_i^\star)^2\} = V$. The OLS intercept identity gives

$$\sqrt{n}(\widehat{\mu}^\circ - \mu) = \frac{1}{\sqrt{n}}\sum_{i=1}^{n}\varepsilon_i^\star - (\widehat{\theta}^\circ - \theta^\star)^{\mathsf{T}}\sqrt{n}\,\overline{C}^\circ.$$

The preceding result gives $\widehat{\theta}^\circ - \theta^\star = o_{\mathbb{P}}(1)$, and $\sqrt{n}\,\overline{C}^\circ = O_{\mathbb{P}}(1)$, so the second term is $o_{\mathbb{P}}(1)$. Since the $\varepsilon_i^\star$ are i.i.d. with mean zero and variance V, the central limit theorem and Slutsky's theorem give

$$\sqrt{n}(\widehat{\mu} - \mu) = \sqrt{n}(\widehat{\mu} - \widehat{\mu}^\circ) + \sqrt{n}(\widehat{\mu}^\circ - \mu) = \frac{1}{\sqrt{n}}\sum_{i=1}^{n}\varepsilon_i^\star + o_{\mathbb{P}}(1) \xrightarrow{d} \mathcal{N}(0, V).$$

Step 2: Variance estimator consistency.

For $\varepsilon_i^\star$ defined above,

$$\widehat{\varepsilon}_i - \varepsilon_i^\star = -(\widehat{\mu} - \mu) - (\widehat{\theta} - \theta^\star)^{\mathsf{T}}\widehat{C}_i - (\theta^\star)^{\mathsf{T}}(\widehat{C}_i - C_i^\circ).$$

Hence,

$$\frac{1}{n}\sum_{i=1}^{n}(\widehat{\varepsilon}_i - \varepsilon_i^\star)^2 \leq 3(\widehat{\mu} - \mu)^2 + 3\left\|\widehat{\theta} - \theta^\star\right\|^2 \frac{1}{n}\sum_{i=1}^{n}\left\|\widehat{C}_i\right\|^2 + 3\left\|\theta^\star\right\|^2 \frac{1}{n}\sum_{i=1}^{n}\left\|\widehat{C}_i - C_i^\circ\right\|^2.$$

Since $\left\|\widehat{C}_i\right\|^2 \le 2\left\|C_i^\circ\right\|^2 + 2\left\|\widehat{C}_i - C_i^\circ\right\|^2$, the law of large numbers and Equation (A.2) give

$$\frac{1}{n}\sum_{i=1}^n \left\|\widehat{C}_i\right\|^2 \le \frac{2}{n}\sum_{i=1}^n \left\|C_i^\circ\right\|^2 + \frac{2}{n}\sum_{i=1}^n \left\|\widehat{C}_i - C_i^\circ\right\|^2 = O_{\mathbb{P}}(1).$$

Together with $\widehat{\mu} \xrightarrow{p} \mu$, $\widehat{\theta} \xrightarrow{p} \theta^\star$, and Equation (A.2), this gives $\frac{1}{n}\sum_{i=1}^n (\widehat{\varepsilon}_i - \varepsilon_i^\star)^2 = o_{\mathbb{P}}(1)$. The law of large numbers gives $\frac{1}{n}\sum_{i=1}^n (\varepsilon_i^\star)^2 \xrightarrow{p} V$. Since $\widehat{V} = \frac{1}{n}\sum_{i=1}^n \widehat{\varepsilon}_i^2$, the reverse triangle inequality gives

$$\left|\widehat{V}^{1/2} - \left\{\frac{1}{n}\sum_{i=1}^n (\varepsilon_i^\star)^2\right\}^{1/2}\right| = \left|\left\{\frac{1}{n}\sum_{i=1}^n \widehat{\varepsilon}_i^2\right\}^{1/2} - \left\{\frac{1}{n}\sum_{i=1}^n (\varepsilon_i^\star)^2\right\}^{1/2}\right| \le \left\{\frac{1}{n}\sum_{i=1}^n (\widehat{\varepsilon}_i - \varepsilon_i^\star)^2\right\}^{1/2} = o_{\mathbb{P}}(1).$$

Therefore, $\widehat{V} \xrightarrow{p} V$.

Step 3: Wald interval coverage.

If $V > 0$, Steps 1 and 2 and Slutsky's theorem give $\frac{\sqrt{n}(\widehat{\mu}-\mu)}{\sqrt{\widehat{V}}} \xrightarrow{d} \mathcal{N}(0,1)$. Therefore,

$$\mathbb{P}\left\{\mu \in \left[\widehat{\mu} - z_{1-\alpha/2}\sqrt{\widehat{V}/n}, \widehat{\mu} + z_{1-\alpha/2}\sqrt{\widehat{V}/n}\right]\right\} = \mathbb{P}\left\{\left|\frac{\sqrt{n}(\widehat{\mu}-\mu)}{\sqrt{\widehat{V}}}\right| \le z_{1-\alpha/2}\right\} \longrightarrow 1-\alpha.$$

□

A.2 HC2 VARIANCE ESTIMATION AND PROOF OF COROLLARY 1

We retain the notation of Appendix A.1 and suppress the dependence on $\bar{a}$. For sample i, write $\widehat{C}_i = C_i(\widehat{Q}, f)$ and $C_i^\circ = C_i(Q^\circ, f)$. Let

$$\overline{\widehat{C}} = \frac{1}{n}\sum_{i=1}^n \widehat{C}_i, \qquad \widehat{\Sigma} = \frac{1}{n}\sum_{i=1}^n (\widehat{C}_i - \overline{\widehat{C}})(\widehat{C}_i - \overline{\widehat{C}})^\mathsf{T}.$$

For the regression in Equation (4.1), block matrix inversion gives

$$h_{ii} = \frac{1}{n} + \frac{1}{n}(\widehat{C}_i - \overline{\widehat{C}})^\mathsf{T}\widehat{\Sigma}^{-1}(\widehat{C}_i - \overline{\widehat{C}}), q_i = 1 - \overline{\widehat{C}}^\mathsf{T}\widehat{\Sigma}^{-1}(\widehat{C}_i - \overline{\widehat{C}}),$$

where h_{ii} is the leverage of observation i, and q_i/n is the weight on $W_{T,i}Y_i$ in the OLS intercept. The HC2 variance estimate for the intercept is $\widehat{V}_{\mathrm{HC2}}/n$, where

$$\widehat{V}_{\mathrm{HC2}} = \frac{1}{n}\sum_{i=1}^n \frac{q_i^2 \widehat{\varepsilon}_i^2}{1-h_{ii}}. \tag{A.3}$$

Proof. The proof of Theorem 1 establishes

$$\frac{1}{n}\sum_{i=1}^n \left\|\widehat{C}_i - C_i^\circ\right\|^2 = o_{\mathbb{P}}(1), \qquad \sqrt{n}\,\overline{\widehat{C}} = O_{\mathbb{P}}(1), \qquad \widehat{\Sigma} \xrightarrow{p} \Sigma^\circ \succ 0, \qquad \widehat{V} \xrightarrow{p} V.$$

Since the C_i° are identically distributed with $\mathbb{E}\left\|C_i^\circ\right\|^2 < \infty$, for every $c > 0$,

$$\mathbb{P}\left(\max_{1\le i\le n} \left\|C_i^\circ\right\|^2 > nc\right) \le n\mathbb{P}(\left\|C_1^\circ\right\|^2 > nc) \le \frac{1}{c}\mathbb{E}\left[\left\|C_1^\circ\right\|^2 \mathbf{1}\{\left\|C_1^\circ\right\|^2 > nc\}\right] \longrightarrow 0.$$

By the triangle inequality and Equation (A.2),

$$\frac{\max_i \left\|\widehat{C}_i - \overline{\widehat{C}}\right\|}{\sqrt{n}} \le \frac{\max_i \left\|C_i^\circ\right\|}{\sqrt{n}} + \left\{\frac{1}{n}\sum_{i=1}^n \left\|\widehat{C}_i - C_i^\circ\right\|^2\right\}^{1/2} + \frac{\left\|\overline{\widehat{C}}\right\|}{\sqrt{n}} = o_{\mathbb{P}}(1).$$

Since $\widehat{\Sigma} \xrightarrow{p} \Sigma^\circ \succ 0$, $\widehat{\Sigma}$ is nonsingular with probability tending to one and $\left\|\widehat{\Sigma}^{-1}\right\| = O_{\mathbb{P}}(1)$. The definitions of h_{ii} and q_i give

$$h_{ii} = \frac{1}{n} + \frac{1}{n}(\widehat{C}_i - \overline{\widehat{C}})^{\mathsf{T}}\widehat{\Sigma}^{-1}(\widehat{C}_i - \overline{\widehat{C}}),$$
$$q_i = 1 - \overline{\widehat{C}}^{\mathsf{T}}\widehat{\Sigma}^{-1}(\widehat{C}_i - \overline{\widehat{C}}).$$

Hence, $\max_i \left|q_i^2/(1-h_{ii}) - 1\right| = o_{\mathbb{P}}(1)$. Using Equation (A.3) and $\widehat{V} = \frac{1}{n}\sum_{i=1}^n \widehat{\varepsilon}_i^2 = O_{\mathbb{P}}(1)$, we obtain

$$\left|\widehat{V}_{\mathrm{HC2}} - \widehat{V}\right| \le \max_i \left|\frac{q_i^2}{1-h_{ii}} - 1\right| \frac{1}{n}\sum_{i=1}^n \widehat{\varepsilon}_i^2 = o_{\mathbb{P}}(1).$$

Together with $\widehat{V} \xrightarrow{p} V$, this proves $\widehat{V}_{\mathrm{HC2}} \xrightarrow{p} V$. Since $V > 0$, Theorem 1 and Slutsky's theorem give $\frac{\sqrt{n}(\widehat{\mu}-\mu)}{\sqrt{\widehat{V}_{\mathrm{HC2}}}} \xrightarrow{d} \mathcal{N}(0,1)$. □

A.3 Proof of Theorem 2

Throughout this proof, we suppress the superscript $\bar{a}$ and write $W_t = W_t^{\bar{a}}$, $\mu = \mu(\bar{a})$, $Q_t(H_t) = Q_t(a_t, H_t)$, $\phi_{\mathrm{EIF}}(O) = \phi_{\mathrm{EIF}}^{\bar{a}}(O)$, and $V_{\mathrm{EIF}} = V_{\mathrm{EIF}}^{\bar{a}}$. Define $\omega_t(H_t) = \mathbb{E}\big[\{W_t - W_{t-1}\}^2 \mid H_t\big]$, then $\omega_t(H_t) = W_{t-1}^2 \frac{1-e_t(a_t|H_t)}{e_t(a_t|H_t)}$. Sequential randomization and the recursive definition of Q_t give $\mathbb{E}(W_T Y \mid H_t, A_t) = W_t Q_t(H_t)$ and $\mathbb{E}(W_T Y \mid H_t) = W_{t-1} Q_t(H_t)$. Consequently,

$$\mathbb{E}\{\omega_t(H_t) Q_t(H_t)^2\} = \mathbb{E}\left[\frac{1-e_t(a_t \mid H_t)}{e_t(a_t \mid H_t)}\{\mathbb{E}(W_T Y \mid H_t)\}^2\right] \le \epsilon^{-1}\mathbb{E}\{(W_T Y)^2\} < \infty.$$

The last inequality follows from positivity and $\mathbb{E}(Y^2) < \infty$. Positivity and fixed T also imply that ω_t is bounded. Hence 1, Q_t, Q_t°, and f_t have finite weighted second moments.

Lemma 1 (Control covariance)**.** *For every history function g_t with $\|g_t\|_t < \infty$,*

$$\mathbb{E}[W_T Y\{W_t - W_{t-1}\} g_t(H_t)] = \langle Q_t, g_t\rangle_t.$$

Proof. Conditioning on H_t, and using $W_t = W_{t-1}\mathbf{1}(A_t = a_t)/e_t(a_t \mid H_t)$, gives

$$\begin{aligned}&\mathbb{E}[W_T Y\{W_t - W_{t-1}\} g_t(H_t) \mid H_t] \\ &\quad = Q_t(H_t) g_t(H_t) \mathbb{E}[W_t\{W_t - W_{t-1}\} \mid H_t] = \omega_t(H_t) Q_t(H_t) g_t(H_t).\end{aligned}$$

Taking expectations proves the result. □

Lemma 2 (Stagewise orthogonality)**.** *For $s < t$ and history functions g_s and $\widetilde{g}_t$ with finite weighted second moments,*

$$\mathbb{E}[\{W_s - W_{s-1}\} g_s(H_s)\{W_t - W_{t-1}\}\widetilde{g}_t(H_t)] = 0.$$

Moreover,

$$\mathbb{E}\big[\{W_t - W_{t-1}\}^2 g_t(H_t)\widetilde{g}_t(H_t)\big] = \langle g_t, \widetilde{g}_t\rangle_t.$$

Proof. For $s < t$, $\{W_s - W_{s-1}\} g_s(H_s)\widetilde{g}_t(H_t)$ is H_t-measurable. Thus,

$$\begin{aligned}&\mathbb{E}[\{W_s - W_{s-1}\} g_s(H_s)\{W_t - W_{t-1}\}\widetilde{g}_t(H_t)] \\ &\quad = \mathbb{E}[\{W_s - W_{s-1}\} g_s(H_s)\widetilde{g}_t(H_t)\mathbb{E}\{W_t - W_{t-1} \mid H_t\}] = 0.\end{aligned}$$

The second identity follows from the definition of ω_t. □

Proof of Theorem 2. For every history function g_t with $\|g_t\|_t < \infty$, Lemmas 1 and 2, together with $\mathbb{E}[\{W_t - W_{t-1}\} g_t(H_t)] = 0$, give

$$\mathbb{E}[\phi_{\mathrm{EIF}}(O)\{W_t - W_{t-1}\} g_t(H_t)] = \langle Q_t, g_t\rangle_t - \langle Q_t, g_t\rangle_t = 0.$$

Fix $\beta_t = (\beta_{t0}, \beta_{t1}, \beta_{t2})^\mathsf{T} \in \mathbb{R}^3$ for $t = 1, \ldots, T$. Then

$$W_T Y - \mu - \sum_{t=1}^{T} \{W_t - W_{t-1}\}\{\beta_{t0} + \beta_{t1} Q_t^\circ(H_t) + \beta_{t2} f_t(H_t)\}$$
$$= \phi_{\mathrm{EIF}}(O) + \sum_{t=1}^{T} \{W_t - W_{t-1}\}\{Q_t(H_t) - \beta_{t0} - \beta_{t1} Q_t^\circ(H_t) - \beta_{t2} f_t(H_t)\}.$$

The preceding orthogonality identities give

$$\mathbb{E}\left[\left\{W_T Y - \mu - \sum_{t=1}^{T} \{W_t - W_{t-1}\}\{\beta_{t0} + \beta_{t1} Q_t^\circ(H_t) + \beta_{t2} f_t(H_t)\}\right\}^2\right]$$
$$= V_{\mathrm{EIF}} + \sum_{t=1}^{T} \left\|Q_t - \beta_{t0} - \beta_{t1} Q_t^\circ - \beta_{t2} f_t\right\|_t^2.$$

By the definition of V, minimizing the left-hand side over $\beta_1, \ldots, \beta_T$ minimizes the terms on the right-hand side separately, which proves the result. □

A.4 Proof of Corollary 2

Proof. Throughout this proof, we suppress the superscript $\bar{a}$ and write $W_t = W_t^{\bar{a}}$, $\mu = \mu(\bar{a})$, and $V_{\mathrm{EIF}} = V_{\mathrm{EIF}}^{\bar{a}}$. The lower bound $V_{\mathrm{EIF}} \le V$ follows from Theorem 2. Restricting $\beta_{t2} = 0$ for every t gives $V \le V_Q$. Taking $\beta_t = (0, 1, 0)^\mathsf{T}$ for every t gives the cross-fitted AIPW residual, while $\beta_t = (0, 0, 0)^\mathsf{T}$ gives the HT residual. Taking $\beta_t = (\mu, 0, 0)^\mathsf{T}$ for every t gives

$$W_T Y - \mu - \sum_{t=1}^{T} \mu\{W_t - W_{t-1}\} = W_T Y - \mu W_T = W_T(Y - \mu),$$

which is the Hájek residual. Therefore, $V \le V_{\mathrm{AIPW}}$, $V \le V_{\mathrm{HT}}$, and $V \le V_{\mathrm{Haj}}$. Combining these inequalities proves the result. □

A.5 Proof of Corollary 3

Proof. Throughout this proof, we suppress the superscript $\bar{a}$ and write $Q_t(H_t) = Q_t(a_t, H_t)$ and $V_{\mathrm{EIF}} = V_{\mathrm{EIF}}^{\bar{a}}$. By the definition of V_Q, $V_Q = V_{\mathrm{EIF}} + \sum_{t=1}^{T} \left\|Q_t^\perp\right\|_t^2$. By the definition of Π_t, $\inf_{\beta \in \mathbb{R}^3} \left\|Q_t - \beta_0 - \beta_1 Q_t^\circ - \beta_2 f_t\right\|_t^2 = \inf_{\beta_2 \in \mathbb{R}} \left\|Q_t^\perp - \beta_2 f_t^\perp\right\|_t^2$. Assumption 4 implies $\left\|f_t^\perp\right\|_t > 0$. The quadratic on the right is minimized at $\beta_2 = \langle Q_t^\perp, f_t^\perp \rangle_t / \left\|f_t^\perp\right\|_t^2$, with minimum $\left\|Q_t^\perp\right\|_t^2 - \langle Q_t^\perp, f_t^\perp \rangle_t^2 / \left\|f_t^\perp\right\|_t^2$. Substituting this expression gives

$$V_Q - V = \sum_{t=1}^{T} \frac{\langle Q_t^\perp, f_t^\perp \rangle_t^2}{\left\|f_t^\perp\right\|_t^2}.$$

□

A.6 Proof of Corollary 4

Proof. Throughout this proof, we suppress the superscript $\bar{a}$ and write $Q_t(H_t) = Q_t(a_t, H_t)$ and $V_{\mathrm{EIF}} = V_{\mathrm{EIF}}^{\bar{a}}$. Theorem 2 gives $V - V_{\mathrm{EIF}} = \sum_{t=1}^{T} \inf_{\beta \in \mathbb{R}^3} \left\|Q_t - \beta_0 - \beta_1 Q_t^\circ - \beta_2 f_t\right\|_t^2$. Each term is nonnegative. Assumption 4 ensures that the weighted least squares problem at every stage has a minimizer. Therefore, $V = V_{\mathrm{EIF}}$ if and only if, for every t, there exists $\beta_t = (\beta_{t0}, \beta_{t1}, \beta_{t2})^\mathsf{T} \in \mathbb{R}^3$ such that $\left\|Q_t - \beta_{t0} - \beta_{t1} Q_t^\circ - \beta_{t2} f_t\right\|_t = 0$. If $Q_t^\circ = Q_t$ almost surely for every t, this condition holds with $\beta_t = (0, 1, 0)^\mathsf{T}$. □

A.7 Proof of Theorem 3

Proof. Let $\tau = \tau(\bar{a}, \bar{b})$, and let $Z_i^\tau = (W_{T,i}^{\bar{a}} - W_{T,i}^{\bar{b}})Y_i$, $C_i^{\tau,\circ} = C_i^\tau(Q^\circ, f)$, and $\widehat{C}_i^\tau = C_i^\tau(\widehat{Q}, f)$. Let $\theta_\tau^\star$ be the unique minimizer defining V_τ, and let $\widehat{\tau}^\circ$ be the OLS intercept obtained after replacing $\widehat{Q}$ with Q°. Based on IPW identification, $\mathbb{E}(Z_i^\tau) = \tau$. After removing completely linearly dependent components, each component of $C_i^{\tau,\circ}$ contains either $W_t^{\bar{a}} - W_{t-1}^{\bar{a}}$ or $W_t^{\bar{b}} - W_{t-1}^{\bar{b}}$. Hence, $\mathbb{E}(C_i^{\tau,\circ}) = 0$. Assumption 5 gives $\mathbb{E}\{C_i^{\tau,\circ} C_i^{\tau,\circ\mathsf{T}}\} \succ 0$, and its moment conditions give finite second moments of Z_i^τ and $C_i^{\tau,\circ}$.

For each fold k, let $\mathcal{F}_k = \sigma(I_1, \ldots, I_K, O_j : j \notin I_k)$. Conditional on $\mathcal{F}_k$, the fitted functions for both paths are fixed. For $i \in I_k$, $\widehat{C}_i^\tau - C_i^{\tau,\circ}$ is a fixed combination of $\{W_t^{\bar{a}} - W_{t-1}^{\bar{a}}\}\{\widehat{Q}_t^{\bar{a},(-k)}(a_t, H_t) - Q_t^{\bar{a},\circ}(H_t)\}$ and $-\{W_t^{\bar{b}} - W_{t-1}^{\bar{b}}\}\{\widehat{Q}_t^{\bar{b},(-k)}(b_t, H_t) - Q_t^{\bar{b},\circ}(H_t)\}$. The conditional Chebyshev and Markov arguments in the proof of Theorem 1 show that

$$\frac{1}{\sqrt{n}}\sum_{i=1}^n (\widehat{C}_i^\tau - C_i^{\tau,\circ}) = o_{\mathbb{P}}(1), \qquad \frac{1}{n}\sum_{i=1}^n \left\|\widehat{C}_i^\tau - C_i^{\tau,\circ}\right\|^2 = o_{\mathbb{P}}(1).$$

Applying the OLS argument in the proof of Theorem 1 to Z_i^τ, τ, and C_i^τ gives $\sqrt{n}(\widehat{\tau} - \widehat{\tau}^\circ) = o_{\mathbb{P}}(1)$ and $\widehat{V}_\tau \xrightarrow{p} V_\tau$. It also gives

$$\sqrt{n}(\widehat{\tau} - \tau) = \frac{1}{\sqrt{n}}\sum_{i=1}^n \left[Z_i^\tau - \tau - (\theta_\tau^\star)^\mathsf{T} C_i^{\tau,\circ}\right] + o_{\mathbb{P}}(1).$$

The summands in the leading term are i.i.d. with mean zero and variance V_τ. The central limit theorem and Slutsky's theorem give the first conclusion of the theorem.The moment and positive definiteness conditions above, together with the two perturbation bounds, allow us to apply the HC2 variance consistency argument in Appendix A.2 to the path contrast regression. Replacing $\widehat{C}_i$, C_i°, and $\widehat{\varepsilon}_i$ in that argument with $\widehat{C}_i^\tau$, $C_i^{\tau,\circ}$, and $\widehat{\varepsilon}_{\tau,i}$ gives $\widehat{V}_{\mathrm{HC2},\tau} - \widehat{V}_\tau = o_{\mathbb{P}}(1)$, $\widehat{V}_{\mathrm{HC2},\tau} \xrightarrow{p} V_\tau$.

By definition, V_τ is the smallest residual variance over the coefficients in the regression adjustment for $\tau(\bar{a}, \bar{b})$. Regression adjustment without f, the cross fitted AIPW residual, and the HT residual belong to this space. Taking the constant adjustment coefficients for the two paths to be $\mu(\bar{a})$ and $\mu(\bar{b})$ gives the Hájek residual $W_T^{\bar{a}}\{Y - \mu(\bar{a})\} - W_T^{\bar{b}}\{Y - \mu(\bar{b})\}$. Therefore, $V_\tau \le V_{Q,\tau}$, $V_\tau \le V_{\mathrm{AIPW},\tau}$, $V_\tau \le V_{\mathrm{HT},\tau}$, and $V_\tau \le V_{\mathrm{Haj},\tau}$. Let $\phi_{\mathrm{EIF},\tau} = \phi_{\mathrm{EIF}}^{\bar{a}} - \phi_{\mathrm{EIF}}^{\bar{b}}$. By the recursive definition of the sequential outcome regressions, for any $\bar{c}, \bar{d} \in \{\bar{a}, \bar{b}\}$, $\mathbb{E}\{\phi_{\mathrm{EIF}}^{\bar{c}}(O) \mid H_t, A_t\} = \mathbb{E}\{\phi_{\mathrm{EIF}}^{\bar{c}}(O) \mid H_t\}$. Together with $\mathbb{E}(W_t^{\bar{d}} - W_{t-1}^{\bar{d}} \mid H_t) = 0$, this gives $\mathbb{E}[\phi_{\mathrm{EIF}}^{\bar{c}}(O)(W_t^{\bar{d}} - W_{t-1}^{\bar{d}})g_t(H_t)] = 0$ whenever $(W_t^{\bar{d}} - W_{t-1}^{\bar{d}})g_t(H_t)$ has a finite second moment. Define $R_\tau(O) = (W_T^{\bar{a}} - W_T^{\bar{b}})Y - \tau - (\theta_\tau^\star)^\mathsf{T} C^\tau(Q^\circ, f) - \phi_{\mathrm{EIF},\tau}(O)$. By the efficient influence function representation, R_τ is a sum of weight increments multiplied by history functions for the two paths. Hence, $\mathbb{E}\{\phi_{\mathrm{EIF},\tau}(O)R_\tau(O)\} = 0$, and

$$V_\tau = \mathbb{E}\left[\{\phi_{\mathrm{EIF},\tau}(O) + R_\tau(O)\}^2\right] = V_{\mathrm{EIF},\tau} + \mathbb{E}\{R_\tau(O)^2\} \ge V_{\mathrm{EIF},\tau}.$$

□

B Reduction Rules for Treatment Path Contrasts

The full adjustment vector in Section 5 is

$$C_{\mathrm{full}}^\tau(Q, f) = \left(C^{\bar{a}}(Q^{\bar{a}}, f)^\mathsf{T}, -C^{\bar{b}}(Q^{\bar{b}}, f)^\mathsf{T}\right)^\mathsf{T}.$$

Write $\Delta W_t^{\bar{r}} = W_t^{\bar{r}} - W_{t-1}^{\bar{r}}$ for $\bar{r} \in \{\bar{a}, \bar{b}\}$. For brevity, we suppress the arguments of the prediction functions. The full vector at stage t is

$$C_{t,\mathrm{full}}^\tau(Q, f) = \left(\Delta W_t^{\bar{a}},\ \Delta W_t^{\bar{a}} Q_t^{\bar{a}},\ \Delta W_t^{\bar{a}} f_t,\ -\Delta W_t^{\bar{b}},\ -\Delta W_t^{\bar{b}} Q_t^{\bar{b}},\ -\Delta W_t^{\bar{b}} f_t\right)^\mathsf{T}.$$

Before joint regression adjustment, we remove components that can be expressed exactly as linear combinations of the others with constant coefficients. The reduction is determined by the target paths, the known randomization design, and the prespecified prediction construction. We denote the

retained stage-specific vector by $C_t^\tau(Q, f)$ and concatenate these vectors to form $C^\tau(Q, f)$. The specific reduction rules are as follows:

For $\bar{a} \neq \bar{b}$, let $d = \min\{t : a_t \neq b_t\}$ be the first stage at which the paths differ.

Before the first divergence, $t < d$. Since $a_s = b_s$ for all $s \leq t$,

$$W_t^{\bar{a}} = W_t^{\bar{b}}, \qquad W_{t-1}^{\bar{a}} = W_{t-1}^{\bar{b}}, \qquad \Delta W_t^{\bar{a}} = \Delta W_t^{\bar{b}}.$$

The two weight increment components differ only in sign, as do the two external prediction adjustment components using the same f_t. Retaining one component from each pair gives

$$C_t^\tau(Q, f) = \big(\Delta W_t^{\bar{a}},\, \Delta W_t^{\bar{a}} Q_t^{\bar{a}},\, \Delta W_t^{\bar{a}} f_t,\, -\Delta W_t^{\bar{b}} Q_t^{\bar{b}}\big)^\mathsf{T}.$$

At the first divergence, $t = d$. For binary treatment,

$$-\Delta W_d^{\bar{b}} = \frac{e_d(a_d \mid H_d)}{e_d(b_d \mid H_d)} \Delta W_d^{\bar{a}}.$$

If the assignment probabilities do not depend on history, this ratio is constant, and

$$C_d^\tau(Q, f) = \big(\Delta W_d^{\bar{a}},\, \Delta W_d^{\bar{a}} Q_d^{\bar{a}},\, \Delta W_d^{\bar{a}} f_d,\, -\Delta W_d^{\bar{b}} Q_d^{\bar{b}}\big)^\mathsf{T}.$$

If the assignment probabilities depend on history and there are no other linear relations with constant coefficients, then $C_d^\tau(Q, f) = C_{d,\mathrm{full}}^\tau(Q, f)$.

After the first divergence, $t > d$. We have

$$\begin{aligned}
\Delta W_t^{\bar{a}} \Delta W_t^{\bar{b}} &= W_{t-1}^{\bar{a}} W_{t-1}^{\bar{b}} \left\{ \frac{\mathbf{1}(A_t = a_t)}{e_t(a_t \mid H_t)} - 1 \right\} \left\{ \frac{\mathbf{1}(A_t = b_t)}{e_t(b_t \mid H_t)} - 1 \right\} \\
&= \frac{\mathbf{1}(\bar{A}_{t-1} = \bar{a}_{t-1}) \mathbf{1}(\bar{A}_{t-1} = \bar{b}_{t-1})}{\prod_{s=1}^{t-1} e_s(a_s \mid H_s) e_s(b_s \mid H_s)} \left\{ \frac{\mathbf{1}(A_t = a_t)}{e_t(a_t \mid H_t)} - 1 \right\} \left\{ \frac{\mathbf{1}(A_t = b_t)}{e_t(b_t \mid H_t)} - 1 \right\} \\
&= 0.
\end{aligned}$$

We retain the components for both paths, so $C_t^\tau(Q, f) = C_{t,\mathrm{full}}^\tau(Q, f)$. If there are additional known relations, such as a constant f_t, we further reduce the vector using the same rule.

Adjustment space. For the prespecified prediction construction, there exists a fixed matrix A with full column rank such that $C_{\mathrm{full}}^\tau(Q, f) = AC^\tau(Q, f)$. Thus, the full and reduced vectors span the same linear adjustment space, and the reduction leaves the minimum population residual variance unchanged. The positive definiteness condition in Assumption 5 is imposed on the reduced vector.

A/A experiments. In an A/A experiment with equal-probability sequential randomization, the sequential outcome regressions under different treatment paths coincide at the population level. At and before the first divergence ($t \leq d$), we use only the sequential outcome predictions for one fixed path and retain

$$C_t^\tau(Q, f) = \big(\Delta W_t^{\bar{a}},\, \Delta W_t^{\bar{a}} Q_t^{\bar{a}},\, \Delta W_t^{\bar{a}} f_t\big)^\mathsf{T}.$$

After the first divergence, we follow the rules above.

C Lorenz–96 Simulation Design and Additional Experiments

C.1 Data generating process and experimental design

Our data generating process is based on the rotated Lorenz–96 model (Karimi & Paul, 2010) used by Chihara et al. (2026). We replace their treatment assignment with independent randomization at each stage. We also introduce a separate external environment for training the pretrained predictor f.

Randomized experiment. Let $X_{i,t} \in \mathbb{R}^d$ denote the state of individual i at stage t, where $d = 10$, let $Y_{i,t}$ denote the intermediate outcome, and let $A_{i,t} \in \{0, 1\}$ denote treatment. The pretreatment history is

$$H_{i,t} = (X_{i,1:t}, Y_{i,1:t-1}, A_{i,1:t-1}),$$

with dimensions $10, 22, 34, 46, 58, 70$ for $t = 1, \ldots, 6$. The final outcome is $Y_i = Y_{i,T}$. In the notation of the main text, we can set $L_{i,1} = X_{i,1}$ and $L_{i,t} = (X_{i,t}, Y_{i,t-1})$ for $t \geq 2$.

The initial states are sampled as $X_{i,1} \sim N(4\mathbf{1}_d, 0.8^2 I_d)$, $Y_{i,0} = 0$, and $A_{i,0} = 0$. Define

$$h(x) = \exp\left\{-\frac{(x_1 - 4)^2}{5}\right\} + \sin\left(\frac{x_2 x_3}{3}\right), \delta_t = 0.2 - 0.6\exp(-t/4) + 0.4\exp(-3t/5).$$

Outcomes follow

$$Y_{i,t} = h(X_{i,t}) + 0.8\log\{1 + \exp(Y_{i,t-1})\} + \delta_t A_{i,t} + \nu_{i,t}, \qquad \nu_{i,t} \overset{\text{iid}}{\sim} N(0, 0.2^2).$$

The randomized experiment independently assigns $A_{i,t} \sim \text{Bernoulli}(1/2)$ at every stage. At each stage, treatment is assigned before the outcome is generated, and the state is then updated for the next stage when $t < T$.

The state transition follows a rotated Lorenz–96 system. We generate a Gaussian $d \times d$ matrix using seed 1901, apply a QR decomposition to obtain a fixed orthogonal matrix R, and set $Z_{i,t}(0) = RX_{i,t}$. With cyclic coordinate indices, define

$$\psi_j(z; y, a) = (z_{j+1} - z_{j-2})z_{j-1} - z_j + 4 + 0.5y + 0.2a, \qquad j = 1, \ldots, d.$$

We solve

$$\frac{\mathrm{d}Z_{i,t,j}(u)}{\mathrm{d}u} = \psi_j\{Z_{i,t}(u); Y_{i,t}, A_{i,t}\}$$

over $u \in [0, 0.01]$ using one fourth-order Runge–Kutta step, and then set

$$X_{i,t+1} = R^{\mathsf{T}} Z_{i,t}(0.01) + 0.2\sqrt{0.01}\,\xi_{i,t}, \qquad \xi_{i,t} \overset{\text{iid}}{\sim} N(0, I_d).$$

The initial states, outcome noise, and state noise are mutually independent, and the noise terms are independent across individuals and stages. The same matrix R is used for the randomized experiment, external samples, interventional truth samples, and all Monte Carlo replications. This mechanism allows treatment to affect subsequent covariates and the final outcome through intermediate outcomes and state transitions.

Pretrained predictor. The external data use the same initial state distribution, outcome equation, and state transition equation as the randomized experiment, but treatment assignment depends on the observed history: $A_{i,t} \mid H_{i,t} \sim \text{Bernoulli}\{\pi_{\text{ext},t}(H_{i,t})\}$, where

$$\pi_{\text{ext},t}(H_{i,t}) = \text{clip}_{[0.03,\,0.97]}\left[\text{expit}\left\{2\big(h(X_{i,t}) - 0.1\big) + 0.8\tanh(Y_{i,t-1}/2) + 0.6\mathbf{1}\{t > 1\}(2A_{i,t-1} - 1)\right\}\right].$$

The external data are independent of the experimental sample and are used only to train f, which remains fixed across replications. Training details are provided in Appendix C.2, and additional sensitivity analyses in Appendix C.3.

Monte Carlo replications and ground truth. The main experiment uses $n = 5000$ and $T = 6$ over 400 Monte Carlo replications, with experimental samples generated using seeds 9100000–9100399. All methods use the same experimental sample within each replication. We consider two average potential outcomes, $\mu_1 = \mu(101010)$ and $\mu_2 = \mu(111000)$, and two causal contrasts, $\tau_1 = \mu(111111) - \mu(000000)$ and $\tau_2 = \mu(100000) - \mu(000001)$.

Ground truth is approximated using 500,000 simulated individuals independent of the experimental samples. For each treatment path, we fix treatments to the specified values at all stages and generate data using the outcome and state transition equations above. Table 3 reports the Monte Carlo ground truth for the four estimands.

Table 3: **Monte Carlo values for the four estimands.**

Estimand	Monte Carlo ground truth
μ_1	4.263857
μ_2	4.220931
τ_1	0.118167
τ_2	−0.079149

C.2 Implementation and evaluation

Sequential outcome predictions. CF-AIPW and CF-LTMLE without external prediction features and our two regression adjustment estimators use the same initial sequential outcome predictions, obtained using LightGBM with two folds for cross fitting. Within each training sample, the regressions proceed backward from the final outcome Y. At stage $t < T$, the regression response is the prediction from the next stage evaluated at the target treatment a_{t+1}. Each regression pools both treatment arms and uses the history and current treatment as inputs. Predictions for each individual are obtained from models trained without that individual's fold. For CF-AIPW and CF-LTMLE with external prediction features, we additionally include $f_t(H_t)$ in the regression inputs and use the same cross fitting procedure.

Pretrained predictor. For each external environment, we generate $50{,}000$ independent individuals and train a multilayer perceptron at each stage to predict the observed final outcome Y directly from the pretreatment history H_t. The predictor does not take a future treatment path as input, so the same $f_t(H_t)$ can be used for regression adjustment under different treatment paths. All predictors are trained before the experimental analysis and remain fixed across all Monte Carlo replications.

Table 4 reports the specifications of LightGBM and the external multilayer perceptrons.

Table 4: **Learner specifications for the synthetic experiments.**

Learner	Parameter	Setting
LightGBM	Loss	Squared error
	Boosting trees	120
	Learning rate	0.04
	Number of leaves	15
	Maximum depth	4
	Minimum observations per leaf	30
	Row subsampling	0.8
	Subsampling frequency	1
	Column subsampling	0.8
	L_1 penalty	0
	L_2 penalty	1.0
External MLP	Hidden layer widths	$256, 256, 128, 64$
	Hidden activation	GELU
	Output layer	Linear
	Loss	Squared error
	Optimizer	AdamW
	Learning rate	0.001
	Weight decay	0.0001
	Batch size	4096
	Validation fraction	10%
	Maximum epochs	30
	Early stopping patience	4

Compared estimators. The main experiment compares nine estimators, all using the known treatment assignment probabilities.

- **HT** estimates each potential outcome mean by averaging $W_T^{\bar{a}}Y$. For a path contrast, it averages $(W_T^{\bar{a}} - W_T^{\bar{b}})Y$. Standard errors use the empirical variance of these individual scores.
- **Hájek** estimates potential outcome means using normalized inverse probability weights. Contrasts are obtained by subtracting the two path mean estimates.
- **CF-AIPW** uses the sequential augmented inverse probability weighting construction of Bang & Robins (2005), with internal predictions obtained through cross fitting (Chernozhukov et al., 2018).
- **CF-LTMLE** applies backward targeting to pooled out of fold predictions using weighted logistic intercept updates, without refitting the initial outcome models (van der Laan & Gruber, 2012). We scale outcomes to $[0, 1]$ using the sample minimum and maximum, clip scaled initial predictions to $[10^{-6}, 1 - 10^{-6}]$, and transform the targeted predictions back to the original scale. Standard errors use the estimated influence functions, with their paired differences used for path contrasts.

- **CF-AIPW with f as a feature** uses sequential outcome regressions with $f_t(H_t)$ included as an additional input, followed by the same AIPW construction.
- **CF-LTMLE with f as a feature** uses sequential outcome regressions with $f_t(H_t)$ included as an additional input, followed by the same targeting procedure.
- **DynamicRA** uses the official implementation of Chihara et al. (2026), with gradient boosting regression for the outcome model and a multilayer perceptron for the transition model. The method uses five folds and 128 Monte Carlo samples for each forward integration. Standard errors and confidence intervals use the moment implementation provided in the official code.
- **Regression adjustment** uses the construction in the main text with internal sequential outcome predictions alone.
- **Joint regression adjustment** jointly incorporates the internal predictions and fixed pretrained predictions $f_t(H_t)$. Both regression adjustment estimators use HC2 variance estimation to construct 95% confidence intervals.

C.3 Sensitivity to the pretrained predictor

Prediction noise. We report the full numerical results for Figure 2(a). We use independent standard normal noise $Z_{i,t}$ and share the same noise realizations across noise levels, methods, and treatment paths within each replication. For CF-AIPW and CF-LTMLE with external predictions as features, we refit the sequential outcome regressions at each noise level. Table 5 reports MSE, coverage, and mean confidence interval length over 400 replications.

Table 5: Sensitivity to prediction noise. MSE is multiplied by 10^3. Cov. and Length denote coverage (%) and mean confidence interval length.

		μ_1			μ_2			τ_1			τ_2		
λ	Method	MSE	Cov.	Length	MSE	Cov.	Length	MSE	Cov.	Length	MSE	Cov.	Length
Without f													
–	CF-AIPW	1.87	94.75	0.17	1.81	95.00	0.17	3.05	97.00	0.23	3.66	94.50	0.23
–	CF-LTMLE	1.89	95.00	0.17	1.83	95.25	0.17	2.97	96.75	0.23	3.69	94.50	0.23
–	Regression adjustment	1.84	95.75	0.17	1.81	95.00	0.17	2.80	97.00	0.23	3.49	93.75	0.22
With f_λ													
0.00	CF-AIPW	1.09	96.00	0.13	1.00	96.50	0.13	1.48	97.00	0.16	1.79	94.75	0.16
	CF-LTMLE	1.10	96.00	0.13	1.02	95.50	0.13	1.49	96.00	0.16	1.78	95.00	0.16
	Joint regression adjustment	0.98	97.25	0.13	0.93	96.75	0.13	1.33	96.00	0.16	1.58	96.00	0.16
0.05	CF-AIPW	1.24	97.25	0.15	1.32	96.75	0.15	1.86	96.75	0.19	2.29	93.75	0.19
	CF-LTMLE	1.26	96.50	0.15	1.34	96.00	0.15	1.85	97.25	0.19	2.28	93.25	0.19
	Joint regression adjustment	1.25	95.25	0.15	1.33	95.75	0.15	1.76	97.00	0.18	2.24	95.00	0.18
0.10	CF-AIPW	1.50	96.50	0.16	1.61	96.00	0.16	2.37	96.75	0.21	3.02	94.00	0.21
	CF-LTMLE	1.53	96.25	0.16	1.62	95.50	0.16	2.35	96.25	0.21	3.02	94.25	0.21
	Joint regression adjustment	1.52	96.00	0.16	1.61	95.25	0.16	2.23	96.25	0.20	2.92	94.00	0.20
0.15	CF-AIPW	1.68	96.25	0.17	1.72	95.50	0.17	2.64	96.75	0.22	3.29	94.50	0.22
	CF-LTMLE	1.71	95.50	0.17	1.75	94.50	0.16	2.61	96.25	0.22	3.31	93.50	0.22
	Joint regression adjustment	1.67	95.50	0.16	1.72	95.00	0.16	2.48	96.50	0.21	3.24	93.25	0.21
0.20	CF-AIPW	1.73	96.00	0.17	1.85	94.75	0.17	2.77	97.00	0.23	3.50	94.75	0.23
	CF-LTMLE	1.75	95.25	0.17	1.87	94.00	0.17	2.74	97.00	0.22	3.51	94.75	0.22
	Joint regression adjustment	1.75	95.25	0.17	1.77	94.75	0.17	2.62	96.25	0.22	3.39	93.50	0.22
0.25	CF-AIPW	1.80	96.00	0.17	1.86	95.00	0.17	2.88	97.25	0.23	3.58	95.50	0.23
	CF-LTMLE	1.82	94.75	0.17	1.89	94.75	0.17	2.85	96.75	0.23	3.60	95.25	0.23
	Joint regression adjustment	1.80	95.00	0.17	1.79	94.75	0.17	2.69	96.25	0.22	3.47	93.75	0.22
0.30	CF-AIPW	1.84	95.25	0.17	1.85	94.75	0.17	2.96	97.00	0.23	3.64	95.25	0.23
	CF-LTMLE	1.84	94.75	0.17	1.87	94.00	0.17	2.90	96.50	0.23	3.66	94.50	0.23
	Joint regression adjustment	1.82	95.00	0.17	1.81	94.75	0.17	2.73	96.50	0.22	3.51	93.25	0.22
0.35	CF-AIPW	1.84	95.75	0.17	1.85	94.50	0.17	3.01	97.25	0.23	3.64	95.00	0.23
	CF-LTMLE	1.85	95.25	0.17	1.88	94.00	0.17	2.93	97.25	0.23	3.66	95.75	0.23
	Joint regression adjustment	1.84	95.00	0.17	1.81	94.50	0.17	2.76	96.25	0.22	3.53	93.00	0.22
0.40	CF-AIPW	1.86	95.00	0.17	1.84	95.00	0.17	3.04	96.75	0.23	3.61	94.50	0.23
	CF-LTMLE	1.87	95.00	0.17	1.86	94.00	0.17	2.96	96.75	0.23	3.64	94.50	0.23
	Joint regression adjustment	1.85	95.00	0.17	1.82	94.50	0.17	2.78	96.00	0.22	3.55	93.00	0.22

Random pretrained predictor. We randomly permute the final outcomes across individuals in the external data and retrain the pretrained predictor, denoted by random f. We compare the methods

under this setting to examine how uninformative external predictions affect estimation precision and confidence interval coverage. The trained predictor remains fixed across 400 Monte Carlo replications, and all other settings follow the main experiment. Table 6 reports MSE, coverage, and mean confidence interval length. With random f, joint regression adjustment performs similarly to regression adjustment without external predictions, supporting robustness of our method to uninformative external predictions.

Table 6: Results with a randomly trained pretrained predictor.

	μ_1			μ_2			τ_1			τ_2		
Method	MSE	Cov.	Length	MSE	Cov.	Length	MSE	Cov.	Length	MSE	Cov.	Length
Without f												
CF-AIPW	1.87	94.75	0.17	1.81	95.00	0.17	3.05	97.00	0.23	3.66	94.50	0.23
CF-LTMLE	1.89	95.00	0.17	1.83	95.25	0.17	2.97	96.75	0.23	3.69	94.50	0.23
Regression adjustment	1.84	95.75	0.17	1.81	95.00	0.17	2.80	97.00	0.23	3.49	93.75	0.22
With random f												
CF-AIPW	1.88	94.50	0.17	1.82	94.75	0.17	3.00	97.75	0.23	3.60	94.25	0.23
CF-LTMLE	1.90	94.50	0.17	1.85	94.75	0.17	2.93	97.25	0.23	3.63	94.25	0.23
Joint regression adjustment	1.87	95.75	0.17	1.79	95.00	0.17	2.77	97.25	0.23	3.53	93.50	0.22

Sensitivity to the external treatment assignment mechanism. For each individual, we generate a latent variable $U_i \sim N(0,1)$ that remains constant across stages. Outcomes in both the experimental and external data follow

$$Y_{i,t} = h(X_{i,t}) + 0.8\log\{1+\exp(Y_{i,t-1})\} + \delta_t A_{i,t} + \gamma U_i + \nu_{i,t}, \qquad \nu_{i,t} \overset{\text{iid}}{\sim} N(0, 0.2^2).$$

The treatment assignment probability in the external data is

$$\pi_{\text{ext},t}(H_{i,t}, U_i) = \text{clip}_{[0.03,\,0.97]}\Big[\text{expit}\Big\{2\big(h(X_{i,t}) - 0.1\big) + 0.8\tanh(Y_{i,t-1}/2) + 0.6\mathbf{1}\{t>1\}(2A_{i,t-1}-1) + \kappa U_i\Big\}\Big].$$

Here, γ controls the effect of the latent variable on outcomes, and κ controls its effect on external treatment assignment. The latent variable is not used as an input to any prediction model, and treatments in the experiment remain independently randomized with probability $1/2$. We consider $\gamma \in \{0, 0.2, 0.4\}$ and $\kappa \in \{0, 0.5, 1, 2, 4, 8\}$. For each fixed γ, we keep the experimental data unchanged, vary κ, and retrain the predictor on the corresponding external data. Tables 7, 8, and 9 report $\text{MSE} \times 10^3$, coverage, and mean confidence interval length over 400 Monte Carlo replications for the three values of γ. All other settings follow the main experiment.

Table 7: Sensitivity to the external treatment assignment mechanism with $\gamma = 0$.

		μ_1			μ_2			τ_1			τ_2		
κ	Method	MSE	Cov.	Length	MSE	Cov.	Length	MSE	Cov.	Length	MSE	Cov.	Length
Without f													
–	CF-AIPW	1.87	94.75	0.17	1.81	95.00	0.17	3.05	97.00	0.23	3.66	94.50	0.23
–	CF-LTMLE	1.89	95.00	0.17	1.83	95.25	0.17	2.97	96.75	0.23	3.69	94.50	0.23
–	Regression adjustment	1.84	95.75	0.17	1.81	95.00	0.17	2.80	97.00	0.23	3.49	93.75	0.22
With f													
0	CF-AIPW	1.09	96.00	0.13	1.00	96.50	0.13	1.48	97.00	0.16	1.79	94.75	0.16
	CF-LTMLE	1.10	96.00	0.13	1.02	95.50	0.13	1.49	96.00	0.16	1.78	95.00	0.16
	Joint regression adjustment	0.98	97.25	0.13	0.93	96.75	0.13	1.33	96.00	0.16	1.58	96.00	0.16
0.5	CF-AIPW	1.10	96.00	0.13	1.00	96.75	0.13	1.49	97.25	0.16	1.83	94.75	0.16
	CF-LTMLE	1.11	95.75	0.13	1.01	96.00	0.13	1.48	96.75	0.16	1.82	94.75	0.16
	Joint regression adjustment	0.99	96.75	0.13	0.95	96.75	0.13	1.32	96.25	0.16	1.59	96.25	0.16
1	CF-AIPW	1.09	96.75	0.13	0.98	96.75	0.13	1.49	97.00	0.16	1.80	95.00	0.16
	CF-LTMLE	1.10	96.75	0.13	1.00	96.25	0.13	1.50	96.25	0.16	1.79	95.25	0.16
	Joint regression adjustment	0.98	97.50	0.13	0.93	96.75	0.13	1.29	96.25	0.16	1.59	95.75	0.16
2	CF-AIPW	1.10	96.00	0.13	0.99	96.50	0.13	1.48	97.00	0.16	1.81	94.75	0.16
	CF-LTMLE	1.11	96.25	0.13	1.00	95.75	0.13	1.49	96.25	0.16	1.79	94.50	0.16
	Joint regression adjustment	0.97	97.00	0.13	0.94	97.00	0.13	1.29	96.25	0.16	1.58	95.25	0.16
4	CF-AIPW	1.10	96.25	0.13	0.98	97.25	0.13	1.44	97.50	0.16	1.80	95.50	0.16
	CF-LTMLE	1.10	96.00	0.13	0.99	96.25	0.13	1.46	97.00	0.16	1.77	95.50	0.16
	Joint regression adjustment	0.97	96.50	0.13	0.93	97.25	0.13	1.29	96.25	0.16	1.56	95.25	0.16

Table 7 continued

		μ_1			μ_2			τ_1			τ_2		
κ	Method	MSE	Cov.	Length	MSE	Cov.	Length	MSE	Cov.	Length	MSE	Cov.	Length
8	CF-AIPW	1.09	96.25	0.13	0.99	97.00	0.13	1.49	97.00	0.16	1.81	94.75	0.16
	CF-LTMLE	1.10	96.25	0.13	1.01	96.50	0.13	1.50	96.75	0.16	1.78	94.75	0.16
	Joint regression adjustment	0.97	96.25	0.13	0.94	97.25	0.13	1.28	96.50	0.16	1.56	95.50	0.16

Table 8: Sensitivity to the external treatment assignment mechanism with $\gamma = 0.2$.

		μ_1			μ_2			τ_1			τ_2		
κ	Method	MSE	Cov.	Length	MSE	Cov.	Length	MSE	Cov.	Length	MSE	Cov.	Length
Without f													
–	CF-AIPW	3.27	95.50	0.22	3.29	95.25	0.22	5.75	95.75	0.30	6.74	92.25	0.31
–	CF-LTMLE	3.31	95.00	0.22	3.24	94.00	0.22	5.76	96.00	0.30	6.76	91.75	0.30
–	Regression adjustment	3.25	95.75	0.22	3.34	94.00	0.22	5.49	96.00	0.30	6.61	92.75	0.30
With f													
0	CF-AIPW	1.78	96.00	0.16	1.65	94.75	0.16	2.73	96.00	0.21	2.99	94.25	0.21
	CF-LTMLE	1.81	95.50	0.16	1.64	94.75	0.16	2.78	95.25	0.21	2.99	93.75	0.21
	Joint regression adjustment	1.64	96.50	0.16	1.51	95.00	0.16	2.59	95.50	0.21	2.72	93.25	0.21
0.5	CF-AIPW	1.76	96.00	0.16	1.60	96.00	0.16	2.78	95.50	0.21	3.00	94.75	0.22
	CF-LTMLE	1.78	95.75	0.16	1.59	95.50	0.16	2.83	95.50	0.21	3.05	93.50	0.21
	Joint regression adjustment	1.62	96.00	0.16	1.52	95.25	0.16	2.55	95.50	0.21	2.70	94.25	0.21
1	CF-AIPW	1.79	96.00	0.17	1.63	95.75	0.17	2.82	96.00	0.22	3.09	94.50	0.22
	CF-LTMLE	1.80	95.50	0.16	1.60	95.25	0.16	2.78	96.00	0.21	3.12	94.00	0.21
	Joint regression adjustment	1.65	95.75	0.16	1.52	96.25	0.16	2.51	96.00	0.21	2.73	94.50	0.21
2	CF-AIPW	1.82	95.25	0.17	1.76	95.25	0.17	2.98	95.75	0.22	3.18	94.50	0.22
	CF-LTMLE	1.82	95.75	0.17	1.70	94.75	0.16	2.82	96.00	0.21	3.15	93.25	0.22
	Joint regression adjustment	1.68	95.75	0.16	1.59	95.75	0.16	2.54	96.25	0.21	2.77	95.00	0.21
4	CF-AIPW	1.89	95.75	0.17	1.81	95.50	0.17	3.06	96.25	0.22	3.23	94.25	0.23
	CF-LTMLE	1.87	95.75	0.17	1.73	94.75	0.17	2.89	96.00	0.22	3.19	93.75	0.22
	Joint regression adjustment	1.77	94.50	0.17	1.59	95.50	0.16	2.64	96.25	0.21	2.87	95.50	0.21
8	CF-AIPW	1.89	95.75	0.17	1.81	95.75	0.17	3.14	96.00	0.23	3.29	95.00	0.23
	CF-LTMLE	1.88	95.75	0.17	1.73	96.00	0.17	2.95	96.00	0.22	3.27	93.75	0.22
	Joint regression adjustment	1.76	96.00	0.17	1.59	96.50	0.16	2.70	96.25	0.21	2.84	95.25	0.21

Table 9: Sensitivity to the external treatment assignment mechanism with $\gamma = 0.4$.

		μ_1			μ_2			τ_1			τ_2		
κ	Method	MSE	Cov.	Length	MSE	Cov.	Length	MSE	Cov.	Length	MSE	Cov.	Length
Without f													
–	CF-AIPW	5.20	95.75	0.28	5.46	94.00	0.28	10.41	95.50	0.39	10.72	94.00	0.39
–	CF-LTMLE	5.20	95.50	0.28	5.42	94.50	0.28	10.48	94.50	0.39	10.65	93.75	0.39
–	Regression adjustment	5.08	95.75	0.28	5.46	93.00	0.28	10.06	94.75	0.39	10.62	93.25	0.39
With f													
0	CF-AIPW	2.51	95.00	0.19	2.38	94.75	0.19	4.30	95.00	0.26	4.37	94.75	0.26
	CF-LTMLE	2.52	94.75	0.19	2.37	95.50	0.19	4.37	94.25	0.26	4.35	94.75	0.26
	Joint regression adjustment	2.23	94.75	0.19	2.19	94.75	0.19	3.84	96.25	0.25	3.93	94.25	0.25
0.5	CF-AIPW	2.52	93.75	0.19	2.40	94.25	0.19	4.31	96.50	0.26	4.35	94.25	0.26
	CF-LTMLE	2.52	94.50	0.19	2.39	94.25	0.19	4.37	95.00	0.26	4.38	94.75	0.26
	Joint regression adjustment	2.22	95.25	0.19	2.19	94.75	0.19	3.70	96.50	0.25	3.82	94.00	0.25
1	CF-AIPW	2.48	93.50	0.19	2.37	94.50	0.19	4.32	95.75	0.26	4.32	94.75	0.26
	CF-LTMLE	2.47	94.50	0.19	2.35	94.50	0.19	4.33	94.75	0.26	4.37	94.25	0.26
	Joint regression adjustment	2.22	95.25	0.19	2.17	95.50	0.19	3.64	97.25	0.25	3.86	94.25	0.25
2	CF-AIPW	2.50	93.75	0.19	2.46	94.50	0.20	4.27	96.50	0.26	4.39	95.50	0.27
	CF-LTMLE	2.49	94.00	0.19	2.40	94.75	0.19	4.17	96.25	0.26	4.42	94.00	0.26
	Joint regression adjustment	2.24	95.00	0.19	2.20	94.75	0.19	3.59	97.25	0.25	3.94	94.25	0.25
4	CF-AIPW	2.54	95.00	0.20	2.50	94.50	0.20	4.58	95.75	0.27	4.42	95.00	0.27
	CF-LTMLE	2.50	94.50	0.19	2.43	94.50	0.19	4.37	96.00	0.26	4.39	94.00	0.26
	Joint regression adjustment	2.30	95.25	0.19	2.24	94.75	0.19	3.69	97.25	0.25	3.96	94.00	0.25
8	CF-AIPW	2.61	93.75	0.20	2.56	94.50	0.20	4.67	96.75	0.27	4.49	95.50	0.27
	CF-LTMLE	2.57	94.50	0.20	2.46	94.75	0.20	4.45	96.00	0.26	4.47	94.00	0.26
	Joint regression adjustment	2.33	95.00	0.19	2.22	94.50	0.19	3.82	96.25	0.25	3.98	94.50	0.25

C.4 SENSITIVITY TO THE INTERNAL LEARNER

We replace LightGBM with a multilayer perceptron (MLP) to examine sensitivity to the internal learner. The MLP has two hidden layers with widths 64 and 32, ReLU activation, and the Adam optimizer. The learning rate is 0.001, the L_2 regularization parameter is 10^{-4}, and the batch size is 64. Training runs for at most 200 epochs, with 15% of the training data used for validation and early stopping after 15 epochs without improvement exceeding 10^{-4}. We retain two folds for cross fitting and use the MLP for the sequential outcome regressions in CF-AIPW, CF-LTMLE, and our regression adjustment estimators, including the regressions with external predictions as additional features. We use $n = 5000$, $T = 6$, and $\gamma = \kappa = 0$. The external predictor and all other settings follow the main experiment. Table 10 reports results over 400 Monte Carlo replications. Both regression adjustment estimators use HC2 variance estimation.

Table 10: **Sensitivity to the internal learner.** MSE ($\times 10^3$), coverage (%), and mean confidence interval length using MLP sequential outcome regressions.

	μ_1			μ_2			τ_1			τ_2		
Method	MSE	Cov.	Len.	MSE	Cov.	Len.	MSE	Cov.	Len.	MSE	Cov.	Len.
Without f												
HT	263.65	95.25	1.95	265.40	93.75	1.93	523.84	94.25	2.80	516.47	95.50	2.80
Hájek	20.39	95.50	0.56	19.20	96.00	0.56	40.05	95.00	0.79	46.36	94.25	0.79
CF-AIPW	1.78	96.00	0.17	1.44	97.25	0.17	2.67	96.75	0.22	2.93	95.25	0.22
CF-LTMLE	1.80	94.75	0.17	1.49	96.00	0.17	2.67	95.75	0.22	3.00	95.00	0.22
Regression adjustment	1.77	95.50	0.17	1.53	96.75	0.17	2.76	96.25	0.22	3.06	94.75	0.22
With f												
CF-AIPW	1.67	93.50	0.15	1.35	95.50	0.16	2.10	97.25	0.20	2.31	95.50	0.20
CF-LTMLE	1.72	92.25	0.15	1.41	95.25	0.15	2.10	96.75	0.20	2.32	95.75	0.20
Joint regression adjustment	**0.97**	95.50	**0.13**	**0.94**	97.75	**0.13**	**1.31**	95.75	**0.16**	**1.55**	95.75	**0.16**

C.5 CONFIDENCE INTERVAL PERFORMANCE ACROSS SAMPLE SIZES

We compare 95% confidence intervals constructed using the standard and HC2 variance estimators at $n = 1000, 3000, 5000$. Both variance estimators use the same point estimates, and all other settings follow the main experiment. Table 11 reports coverage and mean confidence interval length over 400 Monte Carlo replications.

Table 11: **Confidence interval performance across sample sizes.** Coverage (%) and mean confidence interval length.

n	Method	Variance	μ_1		μ_2		τ_1		τ_2	
		estimator	Cov.	Len.	Cov.	Len.	Cov.	Len.	Cov.	Len.
1000	Regression adjustment	Standard	88.25	0.42	89.00	0.42	89.00	0.57	91.75	0.58
		HC2	94.25	0.49	94.75	0.49	93.50	0.67	95.25	0.66
	Joint regression adjustment	Standard	87.25	0.26	86.25	0.26	88.25	0.31	87.75	0.31
		HC2	94.25	0.32	93.25	0.32	97.00	0.38	95.50	0.39
3000	Regression adjustment	Standard	92.75	0.22	92.75	0.22	92.75	0.29	95.00	0.29
		HC2	93.50	0.23	95.25	0.23	94.25	0.30	95.75	0.30
	Joint regression adjustment	Standard	94.50	0.16	94.50	0.16	94.00	0.19	92.25	0.19
		HC2	95.50	0.17	96.00	0.17	95.75	0.21	93.50	0.21
5000	Regression adjustment	Standard	94.25	0.17	94.00	0.17	96.50	0.22	93.25	0.22
		HC2	95.75	0.17	95.00	0.17	97.00	0.23	93.75	0.22
	Joint regression adjustment	Standard	96.00	0.13	96.75	0.13	95.25	0.15	95.00	0.15
		HC2	97.25	0.13	96.75	0.13	96.00	0.16	96.00	0.16

Results. In small samples, HC2 substantially reduces the undercoverage of confidence intervals based on the standard variance estimator. At larger sample sizes, both methods achieve coverage close to the nominal level, with HC2 being slightly conservative for some estimands. We recommend using the HC2 variance estimator to construct confidence intervals.

C.6 DIFFERENT TREATMENT SEQUENCE LENGTHS

We fix $n = 5000$ and vary the number of treatment stages from $T = 2$ to $T = 7$. For each T, the estimand is $\tau_T = \mu(\mathbf{1}_T) - \mu(\mathbf{0}_T)$, where $\mathbf{1}_T$ and $\mathbf{0}_T$ denote the paths with treatment and without treatment at all T stages. We train an external predictor for each time horizon and keep it fixed across replications. The internal sequential outcome regressions use LightGBM with two folds for cross fitting. All other settings follow the main experiment. Table 12 reports MSE, coverage, and mean confidence interval length over 400 Monte Carlo replications. Figure 2(b) displays the MSE and relative MSE reductions for $T = 2, \ldots, 7$. Both regression adjustment estimators use HC2 variance estimation.

Table 12: **Different treatment sequence lengths.** MSE ($\times 10^3$), coverage (%), and mean confidence interval length for estimating τ_T.

	$T=2$			$T=3$			$T=4$		
Method	MSE	Cov.	Len.	MSE	Cov.	Len.	MSE	Cov.	Len.
Without f									
CF-AIPW	0.20	95.00	0.06	0.49	94.75	0.09	0.90	96.75	0.12
CF-LTMLE	0.20	95.25	0.06	0.49	94.75	0.09	0.90	96.75	0.12
Regression adjustment	0.18	95.50	0.05	0.46	94.50	0.08	0.85	96.50	0.12
With f									
CF-AIPW	0.09	93.75	0.04	0.19	95.50	0.05	0.40	94.25	0.08
CF-LTMLE	0.09	93.50	0.04	0.19	95.75	0.05	0.39	94.00	0.08
Joint regression adjustment	0.09	93.50	0.04	0.16	96.25	0.05	0.36	94.50	0.08
	$T=5$			$T=6$			$T=7$		
Method	MSE	Cov.	Len.	MSE	Cov.	Len.	MSE	Cov.	Len.
Without f									
CF-AIPW	1.87	95.00	0.17	3.05	97.00	0.23	7.48	94.75	0.32
CF-LTMLE	1.85	94.75	0.17	2.97	96.75	0.23	7.49	93.75	0.31
Regression adjustment	1.83	94.50	0.16	2.80	97.00	0.23	7.54	92.50	0.32
With f									
CF-AIPW	0.87	95.25	0.11	1.48	97.00	0.16	4.02	93.75	0.23
CF-LTMLE	0.86	94.75	0.11	1.49	96.00	0.16	3.95	93.25	0.23
Joint regression adjustment	0.76	95.25	0.11	1.33	96.00	0.16	3.78	93.50	0.23

Results. MSE and mean confidence interval length increase with the number of treatment stages for all methods. Joint regression adjustment has lower MSE than regression adjustment using internal predictions alone across all time horizons considered.