
Estimating the Long-Term Effects of Novel Treatments

Keith Battocchi*
Microsoft Research

Eleanor W. Dillon
Microsoft Research

Maggie Hei
Microsoft Research

Greg Lewis
Microsoft Research

Miruna Oprescu
Microsoft Research

Vasilis Syrgkanis
Microsoft Research

Abstract

Policy makers often need to estimate the long-term effects of newly-developed treatments, while only having historical data of older treatment options. We propose a surrogate-based approach using a long-term dataset where only past treatments were administered and a short-term dataset where novel treatments have been administered. Our approach generalizes previous surrogate-style methods, allowing for continuous treatments and serially-correlated treatment policies while maintaining consistency and root-n asymptotically normal estimates under a Markovian assumption on the data and the observational policy. Using a semi-synthetic dataset on customer incentives from a major corporation, we evaluate the performance of our method and discuss solutions to practical challenges when deploying our methodology.

1 Introduction

Businesses invent new ways of interacting with their customers. Marketing departments devise new campaigns. Pharmaceutical companies roll out trials of new drugs. In many cases, decision makers expect these novel treatments to affect outcomes like sales or health over several years, but they are eager to begin evaluating these innovations within months of deployment.

We propose an estimation methodology that leverages historical data, collected under earlier treatment regimes, to derive early estimates of the long-term effects of novel, newly-deployed treatments. Athey et al. (2020) develop a surrogate index as one solution to this problem. This method relies on an assumption that causal effects on long-term outcomes are channeled through a set of observable short-term proxies. This assumption allows one to use the historical long-term data set to learn a mapping from short-term signals to a projected long-term reward and subsequently estimate the causal effect of novel treatments on this index of surrogates.

In practice, treatment policies are often *dynamic*: treatments are assigned repeatedly, and their assignments depend on past treatments and short-term outcomes. Unfortunately, this pattern can easily break the assumptions needed for the surrogate index to work. For example, suppose that a firm offers multiple investments to a particular customer in the historical data and these investments are auto-correlated, i.e. if a customer receives an investment this month, then they will receive an investment with high probability in one of the subsequent months. These future investments can substantially increase the long-term outcome of interest, and this increase will be attributed to the short-term proxies. The surrogate index thus formed will tend to over-predict long-run outcomes, and so when it is used to measure the treatment effect of some new treatment in the short-run data, the estimated treatment effect will be bigger (in absolute magnitude) than the truth.

*Correspondence to: Vasilis Syrgkanis <vasy@microsoft.com>.

The main methodological innovation of this paper is to use the dynamic treatment effect analysis of Lewis & Syrgkanis (2020) on the historical data in order to create an unbiased, dynamically adjusted surrogate index. Our new index takes the interpretation of the projected long-term reward in the absence of any future treatments. This model produces unbiased causal effect estimates of the long-term effects of the novel treatments, even with auto-correlated treatments.

A second contribution of the paper is to generalize the final causal analysis step to allow for multiple continuous treatments rather than a single binary treatment, as in Athey et al. (2020). Our new estimator for this expanded surrogate approach allows for the construction of valid confidence intervals, even when using flexible machine learning models at both stages of the estimation to deal with high-dimensional data. In short, we show how one can combine three recently developed techniques, i) the surrogate index approach of Athey et al. (2020), ii) the double machine learning approach of Chernozhukov et al. (2018a) and iii) the dynamic treatment effect estimation approach of Lewis & Syrgkanis (2020), in a single data analysis pipeline to estimate treatment effects in the presence of dynamic treatment policies.

Our work lies in the broader field of estimating causal effects with machine learning and Neyman orthogonality (Neyman, 1979; Robinson, 1988; Ai & Chen, 2003; Chernozhukov et al., 2016; Chernozhukov et al., 2018b). Moreover, it relates to the work on machine learning estimation of treatment effects in the dynamic treatment regime (Nie et al., 2019; Thomas & Brunskill, 2016; Petersen et al., 2014; Kallus & Uehara, 2019b,a; Lewis & Syrgkanis, 2020; Bodory et al.; Singh et al., 2020) and on structural nested models in biostatistics (Robins, 1986; Robins et al., 1992; Robins, 1994; Robins & Ritov, 1997; Robins et al., 2000; Lok & DeGruttola, 2012; Vansteelandt et al., 2014; Vansteelandt & Sjolander, 2016). Finally, it relates to the surrogacy literature in causal inference (Prentice, 1989; Begg & Leung, 2000; Frangakis & Rubin, 2002; Freedman et al., 1992). In the following sections we build on insights in these works to design a new, streamlined method that resolves several difficult aspects of a common applied problem.

2 Problem and Methodology

We illustrate the difficulties of our problem and how our method resolves them with a running example of a firm making investments in its customers in order to increase subsequent purchases. In this section we describe the problem set up and data requirements and use a simple two-period model to build intuition for how our approach estimates long-term causal effects in this environment and why earlier approaches fail. In the following section we provide formal proofs of our methodology in the general, multi-period case.

2.1 Problem statement

A firm has a number of distinct treatments $T_1, T_2 \dots T_k$ it can offer to customers, such as discounts or access to special services. At each period t (say, a month), a vector of treatments $T_{i,t} = (T_{i,t,1}, \dots, T_{i,t,k}) \in \mathbb{R}^k$, is applied to each customer i . We also observe a vector of p characteristics $X_{i,t} = (X_{i,t,1}, \dots, X_{i,t,p}) \in \mathbb{R}^p$, some of which are constant within customer (e.g. industry) and some of which vary over time and customer (e.g. last month's revenue). We observe the outcome of interest $Y_{i,t} \in \mathbb{R}$ (e.g. monthly revenue).

Treatments may affect outcomes over many future periods, and this horizon could vary by treatment. We are therefore interested in identifying the average effect of each treatment at some period t , on the cumulative outcome in the subsequent M periods, i.e. $\bar{Y}_{i,t} = \sum_{\kappa=1}^M Y_{i,t+\kappa}$. Suppose direct inference is impossible because we have not yet observed the subsequent M periods following the introduction of a new treatment. Instead, we have access to a vector of d short-term proxies/surrogates, $S_{i,t} = (S_{i,t,1}, \dots, S_{i,t,d})$ that act as leading indicators for our target outcome. In this example, we would like surrogates that capture a customer's trajectory such as the next few months of revenue, intensity of product usage, new contract commitments, and participation in company events.

2.2 Dynamic Adjusted Surrogate Index

The key innovation of all surrogate-based solutions is to use the proxy variables to bridge two incomplete data sets to estimate $\bar{Y}_{i,t}$. A long-term *observational data set* need only include customer characteristics $X_{i,t}$, surrogates $S_{i,t+1}$, realized M -period outcomes $\bar{Y}_{i,t}$. A second short-term,

experimental data set, includes customer features, $X_{i,t}$, surrogates, $S_{i,t}$, and treatments $T_{i,t}$. This experimental data set is restricted to the period where all treatments of interest have been introduced, but only requires a few months of leading surrogates rather than M periods of leading outcomes. We propose the following estimation strategy:

1. Use the observational data set to estimate the expected long-term outcome net of any post-period t treatments using the recursive methodology of Lewis & Syrgkanis (2020). This adjusted outcome, \bar{Y}_{it}^{adj} , is the expected cumulative outcome starting in period t under the assumption that customer i receives no further treatments over the next M periods.
2. Still using the observational data set, construct a model that predicts this adjusted outcome from $S_{i,t}$ and $X_{i,t}$:

$$g_0^{adj}(S_{i,t}, X_{i,t}) := \mathbb{E}[\bar{Y}_{it}^{adj} | S_{i,t}, X_{i,t}]. \quad (\text{surr. model})$$

Any machine learning estimation algorithm can be used to construct \hat{g} .

3. Finally, calculate the adjusted surrogate index using the surrogates in the experimental data set:

$$I_{i,t} := \hat{g}(S_{i,t}, X_{i,t}). \quad (\text{predicted surrogate index})$$

This predicted index becomes the outcome in the final causal model.

This approach will be consistent for the treatment effects under three assumptions, which we illustrate as a causal graph in Figure 1.

1. **Conditional Independence:** There are no paths from current treatments to future outcomes except through current or future surrogates. Formally, $T_{it} \perp Y_{is} | X_{it}, S_{it}..S_{is}$ for all $s \geq t$.
2. **Stability:** Relationships between current and future surrogates and between surrogates and outcomes, marked in red in the figure, are constant over time and across samples.
3. **Unconfoundedness:** There are no variables beyond our observed set that directly affect both the treatment assignment and outcome.

Note that the first assumption allows relationships between current and future treatments and between current treatments and future surrogates, in contrast to the canonical surrogate model. The stability assumption does not extend to the relationships between treatments and surrogates, which allows for the introduction of novel treatments in the experimental data not found in the observational data. The model is able to account for steady growth in a natural way: as the surrogates grow, the predicted outcomes grow too.

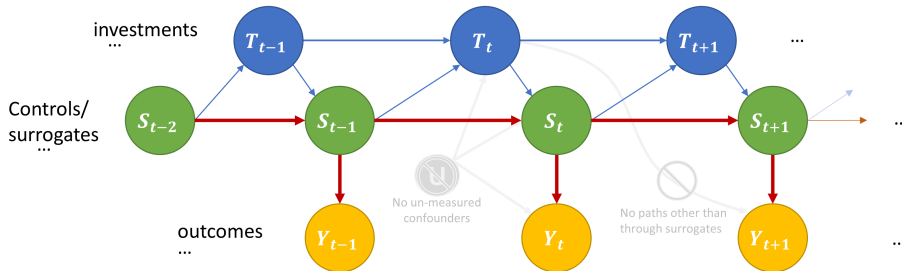


Figure 1: Causal graph representation of the main assumptions of our causal analysis.

2.3 A two-period example

For conciseness, we now collapse $X_{i,t}$ and $S_{i,t-1}$ so that the surrogates at each period and the controls in the next period are all denoted by $S_{i,t-1}$. In other words, all customer observable characteristics

in the next period are candidate surrogates and become controls for the period after next. This is without loss of generality. For the remainder of this section only, we further simplify by assuming a linear Markovian model of how the random variables evolve over time and focusing on a single scalar investment ($k = 1$). With these simplifications, the evolution of the random variables can be described in three equations:

$$\begin{aligned} S_{i,t} &= AT_{i,t} + BS_{i,t-1} + \epsilon_{i,t} && \text{(evolution of controls)} \\ Y_{i,t} &= \gamma' S_{i,t} + \zeta_{i,t} && \text{(outcome equation)} \\ T_{i,t+1} &= \kappa T_{i,t} + \lambda' S_{i,t} + \eta_{i,t} && \text{(treatment policy)} \end{aligned}$$

where A is an $p \times 1$ matrix, B is a $p \times p$ matrix, γ, λ are p -dimensional vectors and κ is a scalar. The terms $\epsilon_{i,t}, \zeta_{i,t}, \eta_{i,t}$ are exogenous mean-zero independent noise terms. For conciseness we drop the customer index i in the remaining equations.

Target effect estimand With two periods, the long term outcome $\bar{Y}_t = Y_{i,t} + Y_{i,t+1}$. The effect of treatment T_t on \bar{Y}_t can be derived from the structural model equations:

$$\bar{Y}_t = \gamma' S_t + \gamma' S_{t+1} + \text{noise} = \gamma' AT_{t+1} + \gamma' (B + I) AT_t + \gamma' (B + I) BS_{t-1} + \text{noise}^2$$

Thus we see that the effect of T_t on $Y_{t+1} + Y_t$, assuming future treatments are held constant, is:

$$\theta_0 := \gamma' (B + I) A. \quad (1)$$

Bias from a standard surrogate index Athey et al. (2020) propose training a surrogate index using the realized cumulative outcome \bar{Y}_t on S_t , essentially our proposed step 2 using raw instead of adjusted outcomes. Following that approach, we would estimate:

$$g_0(S_t) = \mathbb{E}[\bar{Y}_t | S_t] = \gamma' (I + B) S_t + \gamma' A \mathbb{E}[T_{t+1} | S_t] \quad (2)$$

Subsequently, if we estimate the causal effect of T_t on $g_0(S_t)$, then this effect would be:

$$\theta_* = \theta_0 + \gamma' A \frac{\mathbb{E}[\mathbb{E}[T_{t+1} | S_t] T_t]}{\mathbb{E}[T_t^2]} \quad (3)$$

The standard estimate contains a bias stemming from the fact that investment today can lead to higher investment tomorrow. As illustrated in equation above, this bias can stem from two channels, both of which are plausible in many use cases. In our example, a customer that is treated once may build a stronger relationship with the firm and receive further treatments (i.e. $\kappa \geq 0$ in the treatment policy). Alternatively, if the firm targets treatments to fast-growing customers, any treatment that successfully drives surrogate metrics higher will make the customer eligible for further treatment ($\lambda' A \geq 0$).

Dynamic adjustment With only two periods, our first step simply requires estimating the single-period causal effect of T_{t+1} on Y_{t+1} , controlling for S_t . This conditional expectation is equal to:

$$\mathbb{E}[Y_{t+1} | T_{t+1}, S_t] = \gamma' AT_{t+1} + \gamma' BS_t. \quad (4)$$

We can now create an *adjusted long-term outcome*, $\bar{Y}_t^{\text{adj}} := Y_t + Y_{t+1} - \alpha'_{t+1} T_{t+1}$, where α'_{t+1} is our estimate of $\gamma' A$ from this first step.

The *dynamically adjusted surrogate index* builds this adjustment into equation (2) above:

$$\begin{aligned} g_0^{\text{adj}}(S_t) &= \mathbb{E}[Y_t | S_t] + \mathbb{E}[Y_{t+1} - \gamma' AT_{t+1} | S_t] \\ &= \gamma' S_t + \mathbb{E}[\gamma' BS_t | S_{t+1}] \\ &= \gamma' (I + B) S_t \end{aligned} \quad (5)$$

This new index captures the projected M -period outcomes as if the customer was offered no future treatments.

When we estimate the effect of the treatment based on this adjusted surrogate index we recover:

$$\begin{aligned} \mathbb{E}[g_0^{\text{adj}}(S_t) | T_t, S_{t-1}] &= \gamma' (I + B) \mathbb{E}[S_t | T_t, S_{t-1}] \\ &= \gamma' (I + B) AT_t + \gamma' (I + B) BS_{t-1} \end{aligned} \quad (6)$$

With the dynamic adjustment, the coefficient in front of T_t that we recover is the true causal effect $\theta_0 = \gamma' (I + B) A$. Lewis & Syrgkanis (2020) show how this adjustment approach can be extended to many periods, via a recursive peeling process, and also to high-dimensional surrogates and controls via a dynamic double machine learning approach.

Estimating causal effects of new treatments With this adjusted surrogate index in hand, we can also estimate the long-term effect of any other treatment that was introduced more recently. In this example, the stationarity assumption of the proposed surrogate approaches requires that B , which governs how the surrogate evolves, and γ , which governs how surrogates translate to per-period outcomes, do not change between the observational and experimental data-sets. These two parameters govern how surrogates today relate to future outcomes in the absence of any treatment.

Under such a condition, if we introduce a new treatment T_t^{new} , which has a different effect A^{new} on the surrogates (and hence on the long-term outcome), then the effect of this treatment on the long-term outcome is $\theta_0^{\text{new}} = \gamma'(I + B)A^{\text{new}}$. This θ_0^{new} is exactly the outcome of estimating the causal effect of T_t^{new} on $g_0^{\text{adj}}(S_{t+1})$ controlling for S_t .

3 Formal Guarantees

In this section we present our main results. In addition to expanding to M periods, we introduce two substantive innovations over the basic strategy presented in the two period example above. First, we develop a generalization of the doubly robust estimation method in Athey et al. (2020) to the case of multiple continuous treatments. Second, we relax the linearity assumptions in the structural model and make use of orthogonal machine learning techniques Chernozhukov et al. (2018a) to allow for a rich set of potential confounders and valid analytic confidence intervals.

We now assume that we observe the full M -period time series, $(S_0, T_1, S_1, Y_1, T_2, S_2, Y_2, \dots, T_M, S_M, Y_M)$, for each sample in the observational population, denoted as o . In contrast, we only observe (S_0, T_1, S_1) for each sample from the experimental population, denoted as e .

Notation Throughout, we will denote with $\mathbb{E}_e[\cdot]$ the expectation conditional on the experimental population and $\mathbb{E}_o[\cdot]$ the expectation conditional on the observational population. Moreover, for any vector-valued function f that takes as input a random variable Z , we denote with:

$$\|f\|_{2,o} = \sqrt{\mathbb{E}_o[\|f(Z)\|_2^2]} \quad (7)$$

and analogously $\|f\|_{2,e}$. We denote with $\mathbb{E}_n[\cdot]$, the empirical expectation over all the samples, i.e. for any random variable Z , $\mathbb{E}_n[Z] := \frac{1}{n} \sum_i Z_i$, and with $\mathbb{E}_{e,n}$ and $\mathbb{E}_{o,n}$ the empirical expectation over the experimental and observational samples correspondingly, i.e. $\mathbb{E}_{e,n}[Z] = \frac{1}{n_e} \sum_{i \in e} Z_i$ and $\mathbb{E}_{o,n}[Z] = \frac{1}{n_o} \sum_{i \in o} Z_i$.

Our goal is to isolate the causal effect of treatments in period 1, T_1 , on the long term outcome $\bar{Y}_1 := \sum_{t=1}^M Y_t$. Formally, if we fix the future treatments that each sample receives and we change the treatment T_1 from some value t_0 to some other value t_1 , the treatment effect is:

$$\tau(t_1, t_0) := \mathbb{E}_e[\mathbb{E}[\bar{Y}_1 \mid do(T_1 = t_1), T_2, \dots, T_M] - \mathbb{E}[\bar{Y}_1 \mid do(T_1 = t_0), T_2, \dots, T_M]]. \quad (8)$$

We present our theoretical results in two steps. For exposition, we begin by assuming treatments happen only at period 1. This is the setting analyzed in Athey et al. (2020), albeit only for the case of a single binary treatment T_1 .³ We then show how this approach can be modified to incorporate a dynamic treatment policy in the observational and experimental sample.

3.1 Double/Debiased Correction without Dynamic Adjustment

To begin, we assume that $T_2, \dots, T_M = 0$. Throughout, we assume a *partially linear relationship* between the treatment and the long-term outcome:

$$\mathbb{E}_e[\bar{Y}_1 \mid T_1, S_0] = \theta_0^\top \phi(T_1, S_0) + b_0(S_0) \quad (\text{PLR})$$

³Athey et al. (2020) also allowed for estimation of average treatment effects, even in the case when there is arbitrary treatment effect heterogeneity. In this work, we assume that treatment effects are constant. A generalization to the case of arbitrary treatment effect heterogeneity is feasible, but would require the estimation of conditional covariance matrices, which would make the estimation algorithm more brittle and the exposition much more complex.

for some known feature map $\phi(\cdot, \cdot)$, but arbitrary function $b_0(\cdot)$.

Formally, the *invariance of the surrogate-outcome* relationship requires that the mean-relationship between the surrogates S_1 and the long-term outcome does not change between the observational and the experimental sample:

$$g_0(S_1) := \mathbb{E}_o[\bar{Y}_1 | S_1] = \mathbb{E}_e[\bar{Y}_1 | S_1]. \quad (\text{IR})$$

Under the PLR assumption and the causal graph governing our data, we have that:

$$\tau(t_1, t_0) = \theta_0^\top \mathbb{E}[\phi(t_1, S_0) - \phi(t_0, S_0)]. \quad (9)$$

We establish three consistent estimators for θ_0 that follow the same intuition as Athey et al. (2020), adapted to consider linear effects of continuous treatments rather than a single binary treatment.

Theorem 3.1 (Identification). *Denote the residual surrogate index and the residual treatment with:*

$$\tilde{g}_0(S_1) := g_0(S_1) - \mathbb{E}_o[g_0(S_1) | S_0] \quad (10)$$

$$\tilde{T}_1 := \phi(T_1, S_0) - \mathbb{E}_e[\phi(T_1, S_0) | S_0] \quad (11)$$

Then under assumptions (PLR) and (IR):

$$\theta_0 = \mathbb{E}_e[\tilde{T}_1 \tilde{T}_1^\top]^{-1} \mathbb{E}_e[\tilde{T}_1 \tilde{g}_0(S_1)] \quad (\text{surrogate index rep.})$$

$$\theta_0 = \mathbb{E}_e[\tilde{T}_1 \tilde{T}_1^\top]^{-1} \mathbb{E}_o \left[\frac{\Pr(e | S_1, S_0)}{\Pr(o | S_1, S_0)} \frac{\Pr(o)}{\Pr(e)} \mathbb{E}_e[\tilde{T}_1 | S_1, S_0] \bar{Y}_1 \right] \quad (\text{surrogate score rep.})$$

$$\theta_0 = \mathbb{E}_e[\tilde{T}_1 \tilde{T}_1^\top]^{-1} \left(\mathbb{E}_e[\tilde{T}_1 \tilde{g}_0(S_1)] + \mathbb{E}_o \left[\frac{\Pr(e | S_1, S_0)}{\Pr(o | S_1, S_0)} \frac{\Pr(o)}{\Pr(e)} \mathbb{E}_e[\tilde{T}_1 | S_1, S_0] (\bar{Y}_1 - g_0(S_1)) \right] \right) \quad (\text{orthogonal rep.})$$

The core estimation challenge that the surrogate approach resolves is that the treatments and outcome of interest are not observed in a single dataset. Intuitively, the first **surrogate index representation** approaches this challenge by using realized treatments from the experimental sample and, in place of realized outcomes, substitutes the expected outcome conditional on the surrogates, $g_0(S)$, which can be identified from the observational sample and then constructed in the experimental sample.

The second **surrogate score representation** reverses this substitution. The second term pairs an expectation of the treatment conditional on surrogates, $\mathbb{E}_e[\tilde{T}_1 | S_1, S_0]$, with the realized outcomes from the observational sample. This representation requires an added ratio of probabilities of appearing in each sample, $\Pr(e)$ and $\Pr(o)$, to adjust for sampling variation across the two datasets.

The third **orthogonal representation** blends the first two representations and satisfies *Neyman orthogonality*, which allows the construction of confidence intervals and double robustness.

The parameter identified in the equations in Theorem 3.1 is interpretable even if the partially linear assumption is violated. In this case, the equations are identifying the best linear projection of the variation in the long-term outcome that is not explained by the initial state. We formulate the estimation of θ_0 based on the third orthogonal representation as a Z-estimator based on a vector of moment equations that depends on a vector of nuisance functions f_0 , i.e.:

$$m(\theta_0; f_0) := \mathbb{E}[\psi(Z; \theta_0, f_0)] = 0 \quad (12)$$

and such that it satisfies the Neyman orthogonality condition:

$$D[m(\theta_0; f_0), f - f_0] := \frac{\partial}{\partial t} m(\theta_0; f_0 + t(f - f_0)) \Big|_{t=0} = 0$$

Subsequently, this will allow us to invoke the results in Chernozhukov et al. (2018b), to derive an asymptotic normal estimator, even when high-dimensional, regularized approaches are used to estimate the nuisance functions f_0 .

Theorem 3.2 (Orthogonal Moment Formulation). *Let $f = (g, q, p, h)$, denote a set of nuisance functions and define:*

$$\begin{aligned} q_0(S_1, S_0) &:= \frac{\Pr(e | S_1, S_0)}{\Pr(o | S_1, S_0)} \mathbb{E}_e[\tilde{T}_1 | S_1, S_0] \\ p_0(S_0) &:= \mathbb{E}_e[\phi(T_1, S_0) | S_0] \\ h_0(S_0) &:= \mathbb{E}_e[g(S_1) | S_0] \end{aligned}$$

Then θ_0 is the solution to the moment equation:

$$\begin{aligned} m(\theta_0; f_0) &:= m_e(\theta_0; f_0) + m_o(\theta_0; f_0) = 0 \\ m_e(\theta; f) &:= \mathbb{E} \left[1\{e\} (g(S_1) - h(S_0) - \theta^\top (\phi(T_1, S_0) - p(S_0))) \cdot (\phi(T_1, S_0) - p(S_0)) \right] \\ m_o(\theta_0; f) &:= \mathbb{E} [1\{o\} q(S_1, S_0)(\bar{Y}_1 - g(S_1))] \end{aligned}$$

Moreover, the moment m satisfies the Neyman orthogonality property with respect to f . Furthermore, it satisfies a stronger double robustness property with respect to g and q , i.e. if we denote with $\hat{\theta}$ the solution to $m(\theta; \hat{g}, \hat{q}, p_0, \hat{h}) = 0$, then:

$$\left\| \mathbb{E}_e[\tilde{T}_1 \tilde{T}_1^\top] (\hat{\theta} - \theta_0) \right\|_2 \leq \frac{\Pr(o)}{\Pr(e)} \|g_0 - \hat{g}\|_{2,o} \|q_0 - \hat{q}\|_{2,o}$$

We estimate θ_0 by method of moments, solving the empirical analogues of these moment conditions, after estimating the nuisance functions f . We show that this estimator is asymptotically normal under relatively mild rate conditions on the estimation of the nuisances in Appendix Theorem A.2 (Theorem 3.3, below, establishes the same result for a more general case).

3.2 Double/Debiased Correction with Dynamic Adjustment

We now consider the case of potentially serially correlated treatments in all periods. To achieve consistent estimates of the effects of period 1 treatments in this setting we need to again assume that all the relationships are linear, albeit potentially high-dimensional:

$$\begin{aligned} S_t &= AT_t + BS_{t-1} + \epsilon_t \\ Y_t &= CS_t + \eta_t \\ T_{t+1} &= DT_{t-1} + GS_{t-1} + \zeta_t \end{aligned} \tag{13}$$

where $\epsilon_t, \eta_t, \zeta_t$ are i.i.d. random shocks.

In this case, the main identification result in Lewis & Syrgkanis (2020), shows that:

$$\tau(t_1, t_0) = \sum_{t=1}^M \theta_{1,t}^\top (t_1 - t_0) = \theta_0^\top (t_1 - t_0) \tag{14}$$

where $\theta_0 := \sum_{t=1}^M \theta_{1,t}$ and $\theta_{1,t}$ are interpreted as the dynamic effect of treatment at period 1 on the outcome at period t , controlling for all future treatments. This quantity is equivalent to the effect assuming that all future treatments are zero, since the effects of each period treatment under this linear model are linearly separable. This property simplifies both the estimation and the interpretation of the effect quantity as many causal quantities of potential interest collapse in this setting to the same object.

The identification results in Lewis & Syrgkanis (2020) also imply that we can identify the quantity $\theta_{1,t}$ by estimating the effect of T_1 on the multi-period dynamically adjusted target outcome:

$$\bar{Y}_1^{\text{adj}} = \sum_{t=1}^M \left[Y_t - \sum_{\tau=2}^t \theta_{\tau,t}^\top T_\tau \right] = \theta_0^\top T_1 + \beta_0 S_0 + \mu_1 \tag{15}$$

for some exogenous mean-zero random shock variable μ_1 . This is the setup that we analyzed in the previous section, albeit with \bar{Y}_1 replaced by \bar{Y}_1^{adj} and with $\phi(T_1, S_0) = T_1$. The PLR assumption with respect to this dynamically adjusted target immediately holds by the equation above.

Furthermore, we need to assume that the invariance of the surrogate-outcome assumption holds, albeit now for the dynamically adjusted long-term outcome:

$$g_0^{\text{adj}}(S_1) := \mathbb{E}_o[\bar{Y}_1^{\text{adj}} | S_1] = \mathbb{E}_e[\bar{Y}_1^{\text{adj}} | S_1] \tag{dynIR}$$

This assumption is more permissive in practice than the standard invariance assumption as we no longer require that the dynamic treatment policy in the observational data be the same as in the

experimental data, but only that the adjusted outcomes retain the same relationship with the surrogates. For estimation, we can apply the same analysis as in the previous section. The differences are that 1) we now have this new target long-term outcome and 2) we must additionally account for the variance of the $\hat{\theta}_t$ terms that are being estimated by the dynamic adjustment algorithm. We articulate the full algorithm including these additions in Appendix B. We show below that the method of moments estimator with plug in nuisance estimates is asymptotically normal:

Theorem 3.3 (Surrogates with Dynamic Adjustment: Estimation and Asymptotic Normality). *Let:*

$$\begin{aligned} \psi_0(Z; \theta, f_0) &:= \underbrace{1\{e\} \left(\tilde{g}_0^{adj}(S_1) - \theta_0^\top \tilde{T}_1 \right) \tilde{T}_1 + 1\{o\} \left(\tilde{Y}_1^{adj} - g_0^{adj}(S_1) \right)}_{\psi_{0,e}(Z; \theta, f_0)} \underbrace{q_0(S_1, S_0)}_{\psi_{0,o}(Z; \theta, f_0)} \\ \forall 2 \leq \tau \leq t \leq M : \psi_{\tau,t}(Z; \theta, f_0) &:= 1\{o\} \left(\tilde{Y}_{t,\tau} - \sum_{\kappa=\tau+1}^t \theta_{\kappa,t}^\top \tilde{T}_{\kappa,\tau} - \theta_{\tau,t}^\top \tilde{T}_{\tau,\tau} \right) \tilde{T}_{\tau,\tau} \\ \forall 2 \leq t \leq M : \psi_t(Z; \theta, f_0) &= [\psi_{t,t}(Z; \theta, f_0) \quad \dots \quad \psi_{t,2}(Z; \theta, f_0)]^\top \end{aligned}$$

where we use the short-hand notation $\tilde{Y}_{t,\tau} = Y_t - \mathbb{E}[Y_t \mid S_{\tau-1}]$ and $\tilde{Y}_1^{adj} = \sum_{t=1}^M \left(Y_t - \sum_{\tau=2}^t \theta_{\tau,t}^\top T_\tau \right)$ and $\tilde{T}_{t,\tau} = T_t - \mathbb{E}[T_t \mid S_{\tau-1}]$. Let $h_{t,\tau}(S_{\tau-1}) = \mathbb{E}[Y_t \mid S_{\tau-1}]$ and $p_{t,\tau}(S_{\tau-1}) = \mathbb{E}[T_t \mid S_{\tau-1}]$. Let $f = \{g, q, p, h, \{h_{\tau,t}, p_{\tau,t}\}_{2 \leq \tau \leq t \leq M}\}$, denote the set of all nuisance functions and let f_0 denote their true value. Consider the estimator based on the empirical version of the orthogonal moment, with plug-in nuisance estimates \hat{f} trained on a separate sample, i.e. $\hat{\theta}$ solves the empirical moment equation:

$$m_n(\hat{\theta}; \hat{f}) = 0 \quad m_n(\theta; f) := \mathbb{E}_n[\psi(Z; \theta, f)] := \begin{bmatrix} \mathbb{E}_n[\psi_0(Z; \theta, f)] \\ \mathbb{E}_n[\psi_2(Z; \theta, f)] \\ \vdots \\ \mathbb{E}_n[\psi_M(Z; \theta, f)] \end{bmatrix} \quad (16)$$

If $\|\hat{r} - r_0\|_{2,o} = o_p(n^{-1/4})$, for $r \in \{\{h_{\tau,t}, p_{\tau,t}\}_{2 \leq \tau \leq t \leq M}\}$ and if $\|\hat{h} - h_0\|_e = o_p(n^{-1/4})$, $\|\hat{p} - p_0\|_e = o_p(n^{-1/4})$ and $\|g_0 - \hat{g}\|_{2,o} \cdot \|q_0 - \hat{q}\|_{2,o} = o_p(n^{-1/2})$, then:

$$\sqrt{n} (\hat{\theta}_0 - \theta_0) \rightarrow_d N(0, J_e^{-1} (\Sigma_e + \Sigma_{o,1} + \Sigma_{o,2}) J_e^{-1})$$

where $J_e, \Sigma_e, \Sigma_{o,1}, \Sigma_{o,2}$ are defined in the appendix.

4 Semi-Synthetic Experimental Evaluation

To evaluate the performance of our proposed estimation strategy we construct a semi-synthetic dataset that retains qualitative characteristics of data on real-world incentive investments in customers at a major corporation, while preserving confidentiality. The semi-synthetic data set preserves several common patterns that require thoughtful attention. The treatments, in this case incentive investments, are lumpy: in most periods most customers get no investments. Proxies, which include single period values of the outcome of interest, are highly auto-correlated over time. Treatments are also auto-correlated, and correlated with past values of proxies. The data also include time-invariant controls that affect proxies and treatments in all periods.

To build the data we estimate a series of moments from a real-world dataset: a full covariance matrix of all proxies, treatments, and controls in one period and a series of linear prediction models (lassoCV) of each proxy and treatment on a set of 6 lags of each treatment, 6 lags of each proxy, and time-invariant controls. Using these values, we draw new parameters from distributions matching the key characteristics of each family of parameters. Finally, we use these new parameters to simulate proxies, treatments, and controls by drawing a set of initial values from the covariance matrix and forward simulating to match intertemporal relationships from the transformed prediction models. Finally, we arbitrarily select one proxy to be the outcome of interest and construct the cumulative sum of this selected outcome over four or eight periods. For further details on the data generation process, see the appendix.

The true treatment effects in the synthetic data are known functions of parameters from the linear prediction models. We then compare these true causal effects to estimated effects using a variety of

approaches, gradually incorporating our proposed innovations. Figure 2 shows the distribution of the ℓ_2 estimation error ($\|\hat{\theta} - \theta_0\|_2$) for each approach over 100 simulated data sets. The top row plots the estimation error when estimating the effect on four periods of outcomes, increasing the sample size of each simulation from left to right, while the bottom row shows the same for the effect on eight periods of outcome.

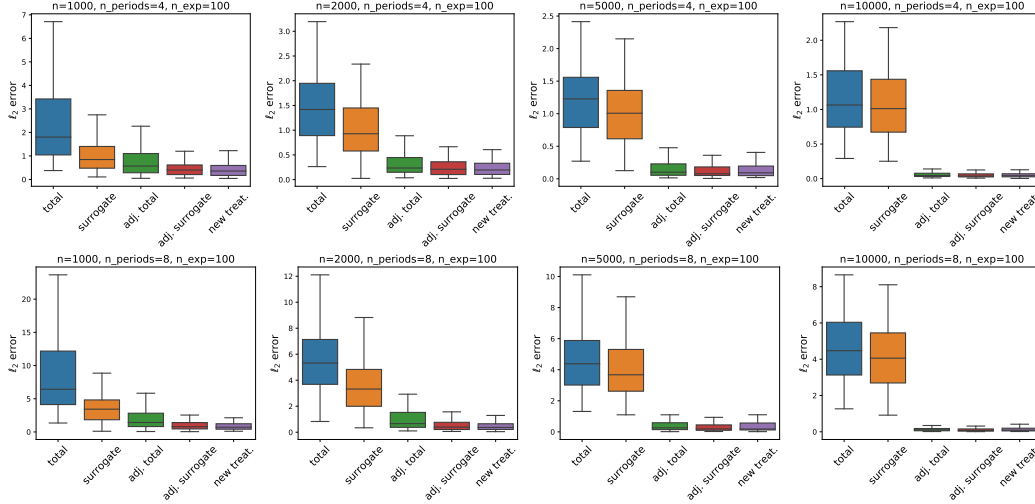


Figure 2: Experimental performance for $M = 4$ periods and $M = 8$ periods.

Because we construct a single, long synthetic data set for this experiment it is possible to estimate the treatment effects on realized long-term outcomes directly, unlike the typical use case for a surrogate approach. We begin by estimating the effect of each treatment at time t directly on the realized total outcomes, rather than the surrogate index, without dynamic adjustment. The blue, “total” bars in each panel of Figure 2 display substantial bias from this method, which does nothing to control for *future* treatments. Because current treatments are positively correlated with future treatments in the semi-synthetic data, this first approach consistently overestimates treatment effects.

We then estimate the same set of treatment effects using the surrogate approach described in Section 3.1, still without dynamic adjustment. The distribution of estimation errors from this approach is represented in the orange “surrogate” bars. The estimated treatment effects are still substantially larger than the true effects on average, but the surrogate model exhibits slightly less bias than the direct “total” approach. Intuitively, because the surrogate approach is only capturing the relationship between treatment and outcome that passes through the surrogates it picks up less of the bias resulting from future correlated treatments than the direct approach.

The third set of green “adj. total” bars plot errors when estimating treatment effects on adjusted realized outcomes using the method of Lewis & Syrgkanis (2020), the preferred approach when all treatments of interest and realized long-run outcomes are available in a single data set. This method removes the effects of future treatments from the long-run outcome in a first step and, as expected, exhibits significantly less bias than the first two methods, particularly for reasonably large samples.

The final two bars illustrate the success of the adjusted surrogate approach described in Section 3.2. As illustrated by the red “adj. surrogate” bars, our proposed approach is highly accurate in predicting long-term effects with a performance comparable to that of having access to the raw long-term outcome itself. The final purple “new treat.” bars show that the approach works equally well when considering the effect of a novel treatment that appears only in the experimental sample and was not part of the dynamic adjustment. Overall, this methodology overcomes a common data limitation when considering long-term effects of novel treatments and expands the surrogate approach to consider a common, and previously problematic, pattern of serially correlated treatments.

Acknowledgments and Disclosure of Funding

The authors have no sources of additional funding to disclose.

References

- Ai, C. and Chen, X. Efficient estimation of models with conditional moment restrictions containing unknown functions. *Econometrica*, 71(6):1795–1843, 2003.
- Athey, S., Chetty, R., Imbens, G., and Kang, H. Estimating treatment effects using multiple surrogates: The role of the surrogate score and the surrogate index, 2020.
- Begg, C. B. and Leung, D. H. On the use of surrogate end points in randomized trials. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 163(1):15–28, 2000.
- Bodory, H., Huber, M., and Laff ers, L. Evaluating (weighted) dynamic treatment effects by double machine learning.
- Chernozhukov, V., Escanciano, J. C., Ichimura, H., Newey, W. K., and Robins, J. M. Locally Robust Semiparametric Estimation. *arXiv e-prints*, art. arXiv:1608.00033, July 2016.
- Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., and Robins, J. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1):C1–C68, 01 2018a. ISSN 1368-4221. doi: 10.1111/ectj.12097. URL <https://doi.org/10.1111/ectj.12097>.
- Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., and Robins, J. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1):C1–C68, 2018b.
- Frangakis, C. E. and Rubin, D. B. Principal stratification in causal inference. *Biometrics*, 58(1): 21–29, 2002.
- Freedman, L. S., Graubard, B. I., and Schatzkin, A. Statistical validation of intermediate endpoints for chronic diseases. *Statistics in medicine*, 11(2):167–178, 1992.
- Kallus, N. and Uehara, M. Double reinforcement learning for efficient off-policy evaluation in markov decision processes, 2019a.
- Kallus, N. and Uehara, M. Efficiently breaking the curse of horizon in off-policy evaluation with double reinforcement learning, 2019b.
- Lewis, G. and Syrgkanis, V. Double/debiased machine learning for dynamic treatment effects, 2020.
- Lok, J. J. and DeGruttola, V. Impact of time to start treatment following infection with application to initiating haart in hiv-positive patients. *Biometrics*, 68(3):745–754, 2012.
- Neyman, J. $c(\alpha)$ tests and their use. *Sankhya*, pp. 1–21, 1979.
- Nie, X., Brunskill, E., and Wager, S. Learning when-to-treat policies, 2019.
- Petersen, M., Schwab, J., Gruber, S., Blaser, N., Schomaker, M., and van der Laan, M. Targeted maximum likelihood estimation for dynamic and static longitudinal marginal structural working models. *Journal of causal inference*, 2(2):147–185, 2014.
- Prentice, R. L. Surrogate endpoints in clinical trials: definition and operational criteria. *Statistics in medicine*, 8(4):431–440, 1989.
- Robins, J. 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.
- Robins, J. M. Correcting for non-compliance in randomized trials using structural nested mean models. *Communications in Statistics-Theory and methods*, 23(8):2379–2412, 1994.
- Robins, J. M. and Ritov, Y. Toward a curse of dimensionality appropriate (coda) asymptotic theory for semi-parametric models. *Statistics in medicine*, 16(3):285–319, 1997.

- Robins, J. M., Blevins, D., Ritter, G., and Wulfsohn, M. G-estimation of the effect of prophylaxis therapy for pneumocystis carinii pneumonia on the survival of aids patients. *Epidemiology*, pp. 319–336, 1992.
- Robins, J. M., Hernan, M. A., and Brumback, B. Marginal structural models and causal inference in epidemiology, 2000.
- Robinson, P. M. Root-n-consistent semiparametric regression. *Econometrica: Journal of the Econometric Society*, pp. 931–954, 1988.
- Singh, R., Xu, L., and Gretton, A. Kernel methods for policy evaluation: Treatment effects, mediation analysis, and off-policy planning. *arXiv preprint arXiv:2010.04855*, 2020.
- Thomas, P. S. and Brunskill, E. Data-efficient off-policy policy evaluation for reinforcement learning. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48*, ICML’16, pp. 2139–2148. JMLR.org, 2016.
- Vansteelandt, S. and Sjolander, A. Revisiting g-estimation of the effect of a time-varying exposure subject to time-varying confounding. *Epidemiologic Methods*, 5(1):37 – 56, 2016. doi: <https://doi.org/10.1515/em-2015-0005>. URL <https://www.degruyter.com/view/journals/em/5/1/article-p37.xml>.
- Vansteelandt, S., Joffe, M., et al. Structural nested models and g-estimation: the partially realized promise. *Statistical Science*, 29(4):707–731, 2014.

A Additional Theorems and Proofs from Section 3

A.1 Proof of Theorem 3.1

Surrogate index representation. By the causal graph assumption and the IR assumption we have:

$$\begin{aligned}\mathbb{E}_e[\bar{Y}_1 | T_1, S_0] &= \mathbb{E}_e[\mathbb{E}_e[\bar{Y}_1 | S_1, T_1, S_0] | T_1, S_0] = \mathbb{E}_e[\mathbb{E}_e[\bar{Y}_1 | S_1] | T_1, S_0] = \mathbb{E}_e[\mathbb{E}_o[\bar{Y}_1 | S_1] | T_1, S_0] \\ &= \mathbb{E}_e[g_0(S_1) | T_1, S_0]\end{aligned}$$

Thus also we have that $\mathbb{E}_e[\bar{Y}_1 | S_0] = \mathbb{E}_e[g_0(S_1) | S_0]$ and that:

$$\mathbb{E}_e[\bar{Y}_1 | T_1, S_0] - \mathbb{E}_e[\bar{Y}_1 | S_0] = \mathbb{E}_e[g_0(S_1) | T_1, S_0] - \mathbb{E}_e[g_0(S_1) | S_0] = \mathbb{E}_e[\tilde{g}_0(S_1) | T_1, S_0]$$

By the PLR assumption and the definition of $\tilde{T}_1 := \phi(T_1, S_0) - \mathbb{E}[\phi(T_1, S_0) | S_0]$, we have that:

$$\mathbb{E}_e[\bar{Y}_1 | T_1, S_0] - \mathbb{E}_e[\bar{Y}_1 | S_0] = \theta_0^\top \tilde{T}_1 \implies \mathbb{E}[\tilde{g}_0(S_1) - \theta_0^\top \tilde{T}_1 | T_1, S_0] = 0 \implies \mathbb{E}\left[\left(\tilde{g}_0(S_1) - \theta_0^\top \tilde{T}_1\right) \tilde{T}_1\right] = 0$$

Solving for θ_0 , yields the surrogate index representation:

$$\theta_0 := \mathbb{E}_e[\tilde{T}_1 \tilde{T}_1^\top]^{-1} \mathbb{E}_e[\tilde{g}_0(S_1) \tilde{T}_1] \quad (17)$$

Surrogate score representation. We start by the definition of the surrogate score representation and use the causal graph assumptions to derive:

$$\begin{aligned}I_{ss} &:= \mathbb{E}_e[\tilde{T}_1 \tilde{T}_1^\top]^{-1} \mathbb{E}_o \left[\frac{\Pr(e | S_1, S_0)}{\Pr(o | S_1, S_0)} \frac{\Pr(o)}{\Pr(e)} \bar{Y}_1 \mathbb{E}_e[\tilde{T}_1 | S_1, S_0] \right] \\ &= \mathbb{E}_e[\tilde{T}_1 \tilde{T}_1^\top]^{-1} \mathbb{E}_o \left[\frac{\Pr(S_1, S_0 | e)}{\Pr(S_1, S_0 | o)} \bar{Y}_1 \mathbb{E}_e[\tilde{T}_1 | S_1, S_0] \right] \quad (\text{Bayes-rule}) \\ &= \mathbb{E}_e[\tilde{T}_1 \tilde{T}_1^\top]^{-1} \mathbb{E}_o \left[\frac{\Pr(S_1, S_0 | e)}{\Pr(S_1, S_0 | o)} \mathbb{E}_o[\bar{Y}_1 | S_1, S_0] \mathbb{E}_e[\tilde{T}_1 | S_1, S_0] \right] \quad (\text{tower-law}) \\ &= \mathbb{E}_e[\tilde{T}_1 \tilde{T}_1^\top]^{-1} \mathbb{E}_e \left[\mathbb{E}_o[\bar{Y}_1 | S_1, S_0] \mathbb{E}_e[\tilde{T}_1 | S_1, S_0] \right] \quad (\text{change of measure}) \\ &= \mathbb{E}_e[\tilde{T}_1 \tilde{T}_1^\top]^{-1} \mathbb{E}_e \left[\mathbb{E}_e[\bar{Y}_1 | S_1, S_0] \mathbb{E}_e[\tilde{T}_1 | S_1, S_0] \right] \quad (\text{IR assumption}) \\ &= \mathbb{E}_e[\tilde{T}_1 \tilde{T}_1^\top]^{-1} \mathbb{E}_e \left[\mathbb{E}_e[\bar{Y}_1 \tilde{T}_1 | S_1, S_0] \right] \quad (\text{conditional independence}) \\ &= \mathbb{E}_e[\tilde{T}_1 \tilde{T}_1^\top]^{-1} \mathbb{E}_e \left[\bar{Y}_1 \tilde{T}_1 \right] \quad (\text{inverse tower-law}) \\ &= \mathbb{E}_e[\tilde{T}_1 \tilde{T}_1^\top]^{-1} \left(\mathbb{E}_e \left[\tilde{T}_1 \phi(T_1, S_0)^\top \theta_0 \right] + \mathbb{E}_e \left[b_0(S_0) \tilde{T}_1 \right] \right) \quad (\text{PLR assumption}) \\ &= \mathbb{E}_e[\tilde{T}_1 \tilde{T}_1^\top]^{-1} \mathbb{E}_e \left[\tilde{T}_1 \tilde{T}_1^\top \theta_0 \right] \quad (\text{mean-zero } \mathbb{E}[\tilde{T}_1 | S_0]) \\ &= \theta_0\end{aligned}$$

Orthogonal representation. This follows easily by the fact that the second term in the orthogonal representation is mean zero and the first term is equal to θ_0 by the surrogate-index argument.

A.2 Proof of Theorem 3.2

Orthogonal moment formulation. First we see that we can re-write the orthogonal representation from Theorem 3.1 as follows:

$$\theta_0 = \mathbb{E}_e[\tilde{T}_1 \tilde{T}_1^\top]^{-1} \left(\frac{\mathbb{E}[\tilde{g}_0(S_1) \tilde{T}_1 1\{e\}]}{\Pr(e)} + \frac{1}{\Pr(e)} \mathbb{E} \left[1\{o\} \frac{\Pr(e | S_1, S_0)}{\Pr(o | S_1, S_0)} (\bar{Y}_1 - g_0(S_1)) \mathbb{E}_e[\tilde{T}_1 | S_1, S_0] \right] \right) \quad (18)$$

$$= \mathbb{E}[\tilde{T}_1 \tilde{T}_1^\top 1\{e\}]^{-1} \left(\mathbb{E}[\tilde{g}_0(S_1) \tilde{T}_1 1\{e\}] + \mathbb{E} \left[1\{o\} \frac{\Pr(e | S_1, S_0)}{\Pr(o | S_1, S_0)} (\bar{Y}_1 - g_0(S_1)) \mathbb{E}_e[\tilde{T}_1 | S_1, S_0] \right] \right) \quad (19)$$

Equivalently, the solution to the orthogonal moment equation:

$$\mathbb{E} \left[1\{e\} (\tilde{g}_0(S_1) - \theta_0^\top \tilde{T}_1) \tilde{T}_1 + 1\{o\} \frac{\Pr(e \mid S_1, S_0)}{\Pr(o \mid S_1, S_0)} \mathbb{E}_e[\tilde{T}_1 \mid S_1, S_0] (\bar{Y}_1 - g_0(S_1)) \right] = 0 \quad (20)$$

By the definition of q_0, p_0 and h_0 the result follows that we can write θ_0 as the solution to $m(\theta_0; f_0) = 0$, where $f_0 = (q_0, p_0, h_0)$ and we remind that:

$$\begin{aligned} m(\theta; f) &:= m_e(\theta; f) + m_o(\theta; f) \\ m_e(\theta; f) &:= \mathbb{E} [1\{e\} (g(S_1) - h(S_0) - \theta^\top (\phi(T_1, S_0) - p(S_0))) (\phi(T_1, S_0) - p(S_0))] \\ m_o(\theta; f) &:= \mathbb{E} [1\{o\} q(S_1, S_0) (\bar{Y}_1 - g(S_1))] \end{aligned}$$

Orthogonality. Orthogonality with respect to p, h, q , follows since:

$$\begin{aligned} D_p[m(\theta_0; f_0), p - p_0] &:= \Pr(e) \left(\mathbb{E}_e[\tilde{T}_1 (p(S_0) - p_0(S_0))] - \mathbb{E}_e \left[\left(\tilde{g}_0(S_1) - \theta_0^\top \tilde{T}_1 \right) (p(S_0) - p_0(S_0)) \right] \right) = 0 \\ D_h[m(\theta_0; f_0), h - h_0] &:= \Pr(e) \mathbb{E}_e[-\tilde{T}_1 (h(S_0) - h_0(S_0))] = 0 \\ D_q[m(\theta_0; f_0), q - q_0] &:= \Pr(o) \mathbb{E}_o[(\bar{Y}_1 - g_0(S_1)) (q(S_1, S_0) - q_0(S_1, S_0))] \\ &= \Pr(o) \mathbb{E}_o[\mathbb{E}[\bar{Y}_1 - g_0(S_1) \mid S_1] (q(S_1, S_0) - q_0(S_1, S_0))] = 0 \end{aligned}$$

Orthogonality with respect to g is slightly more involved. First note that:

$$D_g[m(\theta_0; f_0), g - g_0] = \Pr(e) \mathbb{E}_e[\tilde{T}_1 (g(S_1) - g_0(S_1))] - \Pr(o) \mathbb{E}_o[q_0(S_1, S_0) (g(S_1) - g_0(S_1))]$$

Now using a sequence of derivations almost identical to the ones we invoked in the surrogate score representation in Theorem 3.1, we can show that:

Lemma A.1. For any scalar valued function $r(S_1, S_0)$,

$$\mathbb{E}_o[q_0(S_1, S_0) r(S_1, S_0)] = \frac{\Pr(e)}{\Pr(o)} \mathbb{E}_e[\tilde{T}_1 r(S_1, S_0)] \quad (21)$$

Proof.

$$\begin{aligned} \frac{\Pr(e)}{\Pr(o)} \mathbb{E}_o[q_0(S_1, S_0) r(S_1, S_0)] &:= \mathbb{E}_o \left[\frac{\Pr(e \mid S_1, S_0)}{\Pr(o \mid S_1, S_0)} \frac{\Pr(o)}{\Pr(e)} \mathbb{E}_e[\tilde{T}_1 \mid S_1, S_0] r(S_1, S_0) \right] \\ &\quad \text{(definition of } q_0) \\ &= \mathbb{E}_o \left[\frac{\Pr(S_1, S_0 \mid e)}{\Pr(S_1, S_0 \mid o)} \mathbb{E}_e[\tilde{T}_1 \mid S_1, S_0] r(S_1, S_0) \right] \quad \text{(Bayes-rule)} \\ &= \mathbb{E}_e \left[\mathbb{E}_e[\tilde{T}_1 \mid S_1, S_0] r(S_1, S_0) \right] \quad \text{(change of measure)} \\ &= \mathbb{E}_e \left[\mathbb{E}_e[\tilde{T}_1 r(S_1, S_0) \mid S_1, S_0] \right] \\ &= \mathbb{E}_e \left[\tilde{T}_1 r(S_1, S_0) \right] \quad \text{(inverse tower-law)} \end{aligned}$$

□

Applying Lemma A.1 for $r(S_1, S_0) = g(S_1) - g_0(S_1)$, we have:

$$D_g[m(\theta_0; f_0), g - g_0] = \Pr(e) \mathbb{E}_e[\tilde{T}_1 (g(S_1) - g_0(S_1))] - \Pr(o) \frac{\Pr(e)}{\Pr(o)} \mathbb{E}_e[\tilde{T}_1 (g(S_1) - g_0(S_1))] = 0$$

Double robustness. Observe that by re-writing the orthogonal moment equation, and applying the definition of $g_0(S_1)$ and Lemma A.1 we have:

$$\begin{aligned}
\mathbb{E}_e[\tilde{T}_1 \tilde{T}_1^\top] \hat{\theta} &= \mathbb{E}_e[\hat{g}(S_1) \tilde{T}_1] + \frac{\Pr(o)}{\Pr(e)} \mathbb{E}_o[(\bar{Y}_1 - \hat{g}(S_1)) \hat{q}(S_1, S_0)] \\
&= \mathbb{E}_e[(\hat{g}(S_1) - g_0(S_1)) \tilde{T}_1] + \frac{\Pr(o)}{\Pr(e)} \mathbb{E}_o[\mathbb{E}[\bar{Y}_1 - \hat{g}(S_1) \mid S_1] \hat{q}(S_1, S_0)] \\
&\quad \text{(tower-law and causal graph)} \\
&= \mathbb{E}_e[\hat{g}(S_1) \tilde{T}_1] + \frac{\Pr(o)}{\Pr(e)} \mathbb{E}_o[(g_0(S_1) - \hat{g}(S_1)) \hat{q}(S_1, S_0)] \quad \text{(definition of } g_0) \\
&= \frac{\Pr(o)}{\Pr(e)} \mathbb{E}_o[\hat{g}(S_1) q_0(S_1, S_0)] + \frac{\Pr(o)}{\Pr(e)} \mathbb{E}_o[(g_0(S_1) - \hat{g}(S_1)) \hat{q}(S_1, S_0)] \\
&\quad \text{(Lemma A.1)} \\
&= \frac{\Pr(o)}{\Pr(e)} \mathbb{E}_o[g_0(S_1) q_0(S_1, S_0)] + \frac{\Pr(o)}{\Pr(e)} \mathbb{E}_o[(g_0(S_1) - \hat{g}(S_1))(\hat{q}(S_1, S_0) - q_0(S_1, S_0))] \\
&= \mathbb{E}_e[\tilde{T}_1 g_0(S_1)] + \frac{\Pr(o)}{\Pr(e)} \mathbb{E}_o[(g_0(S_1) - \hat{g}(S_1))(\hat{q}(S_1, S_0) - q_0(S_1, S_0))] \\
&\quad \text{(Lemma A.1)} \\
&= \mathbb{E}_e[\tilde{T}_1 \tilde{T}_1^\top] \theta_0 + \mathbb{E}_o[(g_0(S_1) - \hat{g}(S_1))(\hat{q}(S_1, S_0) - q_0(S_1, S_0))]
\end{aligned}$$

Thus we get the desired property, i.e.:

$$\mathbb{E}_e[\tilde{T}_1 \tilde{T}_1^\top] (\hat{\theta} - \theta_0) = \frac{\Pr(o)}{\Pr(e)} \mathbb{E}_o[(g_0(S_1) - \hat{g}(S_1))(\hat{q}(S_1, S_0) - q_0(S_1, S_0))] \quad (22)$$

which implies that:

$$\left\| \mathbb{E}_e[\tilde{T}_1 \tilde{T}_1^\top] (\hat{\theta} - \theta_0) \right\|_2 \leq \frac{\Pr(o)}{\Pr(e)} \sqrt{\mathbb{E}_o[(g_0(S_1) - \hat{g}(S_1))^2]} \cdot \sqrt{\mathbb{E}_o[(\hat{q}(S_1, S_0) - q_0(S_1, S_0))^2]} \quad (23)$$

A.3 Estimation and Asymptotic Normality without Dynamic Adjustment

Theorem A.2 (Estimation and Asymptotic Normality). *Consider the estimator based on the empirical version of the orthogonal moment, with plug-in nuisance estimates $\hat{f} = (\hat{g}, \hat{q}, \hat{p}, \hat{h})$ trained on a separate sample, i.e. $\hat{\theta}$ solves the empirical moment equation:*

$$m_n(\hat{\theta}; \hat{f}) := m_{e,n}(\hat{\theta}; \hat{f}) + m_{o,n}(\hat{\theta}; \hat{f}) = 0$$

$$m_{e,n}(\theta; f) := \mathbb{E}_n \left[1\{e\} (g(S_1) - h(S_0) - \theta^\top (\phi(T_1, S_0) - p(S_0))) (\phi(T_1, S_0) - p(S_0)) \right]$$

$$m_{o,n}(\theta; f) := \mathbb{E}_n [1\{o\} q(S_1, S_0) (\bar{Y}_1 - g(S_1))]$$

If $\|\hat{h} - h_0\|_{2,e}, \|\hat{p} - p_0\|_{2,e} = o_p(n^{-1/4})$ and $\|g_0 - \hat{g}\|_{2,o} \cdot \|q_0 - \hat{q}\|_{2,o} = o_p(n^{-1/2})$, then:

$$\sqrt{n} (\hat{\theta} - \theta_0) \rightarrow_d N(0, J_e^{-1} (\Sigma_e + \Sigma_o) J_e^{-1})$$

where:

$$J_e := \mathbb{E}_e[\tilde{T}_1 \tilde{T}_1^\top]$$

$$\Sigma_e := \frac{1}{\Pr(e)} \mathbb{E}_e \left[(\tilde{g}_0(S_1) - \theta_0^\top \tilde{T}_1)^2 \tilde{T}_1 \tilde{T}_1^\top \right]$$

$$\Sigma_o := \frac{\Pr(o)}{\Pr(e)^2} \mathbb{E}_o [(\bar{Y}_1 - g_0(S_1))^2 q_0(S_1, S_0) q_0(S_1, S_0)^\top]$$

Proof. Employing the general theorem in Chernozhukov et al. (2018b) and since our estimator is the solution to the empirical analogue of a Neyman orthogonal moment equation:

$$m(\theta; f_0) = \mathbb{E}[\psi(Z; \theta, f_0)] = 0 \quad (24)$$

Thus as long as the nuisance functions f_0 are estimated on a separate sample and they satisfy an MSE rate of $o_p(n^{-1/4})$, then the final parameter estimate is asymptotically normal, i.e. $\sqrt{n}(\hat{\theta} - \theta_0) \rightarrow_d N(0, V)$ with $V = J^{-1}\Sigma J^{-\top}$, where $J = \nabla_{\theta} m(\theta_0; f_0)$ and $\Sigma = \mathbb{E}[\psi(Z; \theta_0, f_0) \psi(Z; \theta_0, f_0)^{\top}]$. We can also easily leverage the double robustness property in Theorem 3.2 to further relax the requirement on the nuisance functions g, q , so as to only require that the product of the two mean squared errors is of $o_p(n^{-1/2})$ as opposed to requiring the stronger property that each individual error is of $o_p(n^{-1/4})$.

Asymptotic variance characterization and estimation. We now analyze and simplify the variance term. First observe that in our setting:

$$J := \mathbb{E}[\tilde{T}_1 \tilde{T}_1^{\top} 1\{e\}] = \Pr(e) \mathbb{E}_e[\tilde{T}_1 \tilde{T}_1^{\top}] =: \Pr(e) J_e \quad (25)$$

Moreover, if we let:

$$\psi(Z; \theta, f) := 1\{e\} \psi_e(S_1, T_1, S_0; \theta, g, p, h) + 1\{o\} \psi_o(\bar{Y}_1, S_1, S_0; g, q) \quad (26)$$

$$:= 1\{e\} \psi_e(Z_e; \theta, g, p, h) + 1\{o\} \psi_o(Z_o; \theta, g, q) \quad (27)$$

$$\psi_e(Z_e; \theta, g, p, h) := (g(S_1) - h(S_0) - \theta_0^{\top} (\phi(T_1, S_0) - p(S_0))) (\phi(T_1, S_0) - p(S_0)) \quad (28)$$

$$\psi_o(Z_o; g, q) := q(S_1, S_0) (\bar{Y}_1 - g(S_1)) \quad (29)$$

Then we also have that:

$$\Sigma := \mathbb{E}[\psi(Z; \theta_0, f_0) \psi(Z; \theta_0, f_0)^{\top}] \quad (30)$$

$$= \Pr(e) \mathbb{E}_e[\psi_e(Z_e; \theta_0, g_0, p_0, h_0) \psi_e(Z_e; \theta_0, g_0, p_0, h_0)^{\top}] + \Pr(o) \mathbb{E}_o[\psi_o(Z_o; g_0, q_0) \psi_o(Z_o; g_0, q_0)^{\top}] \quad (31)$$

$$= \Pr(e)^2 \Sigma_e + \Pr(o)^2 \Sigma_o \quad (\text{definition of } \Sigma_e, \Sigma_o)$$

Thus the variance V can be re-written as desired:

$$V = J^{-1} \Sigma J^{-\top} = \frac{1}{\Pr(e)^2} J_e^{-1} \Sigma J_e^{-\top} = J_e^{-1} (\Sigma_e + \Sigma_o) J_e^{-\top} \quad (32)$$

□

A.4 Proof of Theorem 3.3

Definitions:

$$J_e := \mathbb{E}_e[\tilde{T}_1 \tilde{T}_1^{\top}]$$

$$\Sigma_e := \frac{1}{\Pr(e)} \mathbb{E}_e[\Psi_{o,e} \Psi_{o,e}^{\top}]$$

$$\Sigma_{o,1} := \frac{\Pr(o)}{\Pr(e)^2} \mathbb{E}_o[\Psi_{0,o} \Psi_{0,o}^{\top}] \quad \Sigma_{o,2} := \frac{\Pr(o)}{\Pr(e)^2} \sum_{t=2}^M B_t C_t^{-1} \cdot \left(\sum_{t'=2}^M \mathbb{E}_o[\Psi_t \Psi_{t'}^{\top}] C_{t'}^{-\top} B_{t'}^{\top} - \mathbb{E}_o[\Psi_t \Psi_{0,o}^{\top}] \right)$$

where we use the short-hand notation Ψ_t for the random variable $\psi_t(Z; \theta_0, f_0)$ and B_t, C_t are appropriately defined matrices.

Proof:

Employing the general theorem in Chernozhukov et al. (2018b) and since our estimator is the solution to the empirical analogue of a Neyman orthogonal moment equation:

$$m(\theta; f_0) = \mathbb{E}[\psi(Z; \theta, f_0)] = 0 \quad (33)$$

Neyman orthogonality of the moments that estimate the dynamic effects $\theta_{t,\tau}$, with respect to the nuisance functions $h_{t,\tau}, p_{t,\tau}$ is established in Lewis & Syrgkanis (2020), while orthogonality with respect to g, q, p, h can be established in a manner identical to the proof in Theorem 3.2.

Thus as long as the nuisance functions f_0 are estimated on a separate sample and they satisfy an MSE rate of $o_p(n^{-1/4})$, then the final parameter estimate is asymptotically normal, i.e. $\sqrt{n}(\hat{\theta} - \theta_0) \rightarrow_d N(0, V)$ with

$$V := [J^{-1} \Sigma J^{-\top}]_{1:k, 1:k} \quad J := \nabla_{\theta} m(\theta_0; f_0) \quad \Sigma := \mathbb{E}[\psi(Z; \theta_0, f_0) \psi(Z; \theta_0, f_0)^{\top}] \quad (34)$$

where:

$$\psi(Z; \theta_*, f_0) := [\psi_0(Z; \theta_*, f_0); \psi_2(Z; \theta_*, f_0); \psi_3(Z; \theta_*, f_0); \dots; \psi_M(Z; \theta_*, f_0)]^\top \quad (35)$$

and $[A]_{1:k, 1:k}$ is the $k \times k$ top left diagonal block of matrix A and k is the number of treatments and with θ_* we denote the true value of the whole vector of parameters $\theta_0, \{\theta_{t,\tau}\}_{2 \leq t \leq \tau \leq M}$. We can also easily leverage the double robustness property in Theorem 3.2 to further relax the requirement on the nuisance functions g, q , so as to only require that the product of the two mean squared errors is of $o_p(n^{-1/2})$ as opposed to requiring the stronger property that each error is of $o_p(n^{-1/4})$.

Asymptotic variance characterization and estimation. We further analyze and decompose the variance V above. In particular, observe that:

$$\begin{aligned} J &= - \begin{bmatrix} A \Pr(e) & B \Pr(o) \\ 0 & C \Pr(o) \end{bmatrix} \\ A &= \mathbb{E}_e[\tilde{T}_1 \tilde{T}_1^\top] \\ B &= [B_2 \quad \dots \quad B_M] \\ B_t &= [\mathbb{E}_o[q_0(S_1, S_0) T_t^\top] \quad \dots \quad \mathbb{E}_o[q_0(S_1, S_0) T_2^\top]] \end{aligned}$$

And:

$$C = \begin{bmatrix} C_2 & 0 & \dots & 0 \\ 0 & C_3 & \dots & 0 \\ \dots & \dots & \dots & \dots \\ 0 & 0 & \dots & C_M \end{bmatrix} \quad C_t = \begin{bmatrix} \mathbb{E}_o[\tilde{T}_{t,t} \tilde{T}_{t,t}^\top] & 0 & \dots & 0 \\ \mathbb{E}_o[\tilde{T}_{t,t-1} \tilde{T}_{t-1,t-1}^\top] & \mathbb{E}_o[\tilde{T}_{t-1,t-1} \tilde{T}_{t-1,t-1}^\top] & \dots & 0 \\ \dots & \dots & \dots & 0 \\ \mathbb{E}_o[\tilde{T}_{t,2} \tilde{T}_{2,2}^\top] & \mathbb{E}_o[\tilde{T}_{t-1,2} \tilde{T}_{2,2}^\top] & \dots & \mathbb{E}_o[\tilde{T}_{2,2} \tilde{T}_{2,2}^\top] \end{bmatrix}$$

Observe that we can write:

$$J = - \Pr(e) \begin{bmatrix} A & B \frac{\Pr(o)}{\Pr(e)} \\ 0 & C \frac{\Pr(o)}{\Pr(e)} \end{bmatrix} \quad J^{-1} = - \frac{1}{\Pr(e)} \begin{bmatrix} A^{-1} & -A^{-1} B C^{-1} \\ 0 & C^{-1} \frac{\Pr(e)}{\Pr(o)} \end{bmatrix} \quad (36)$$

Moreover, we can write:

$$\Sigma = \Pr(e) \begin{bmatrix} \mathbb{E}_e[\psi_{0,e}(Z; \theta_*, f_0) \psi_{0,e}(Z; \theta_*, f_0)^\top] & 0 \\ 0 & 0 \end{bmatrix} + \Pr(o) \mathbb{E}_o[\psi(Z; \theta_*, f_0) \psi(Z; \theta_*, f_0)^\top] \quad (37)$$

Leading to:

$$V = V_e + V_o \quad (38)$$

where:

$$[V_e]_{1:k, 1:k} = \frac{1}{\Pr(e)} A^{-1} \mathbb{E}_e[\psi_{0,e}(Z; \theta_*, f_0) \psi_{0,e}(Z; \theta_*, f_0)^\top] A^{-1} \quad (39)$$

and

$$V_o = \frac{\Pr(o)}{\Pr(e)^2} \begin{bmatrix} A & B \frac{\Pr(o)}{\Pr(e)} \\ 0 & C \frac{\Pr(o)}{\Pr(e)} \end{bmatrix}^{-1} \mathbb{E}_o[\psi(Z; \theta_*, f_0) \psi(Z; \theta_*, f_0)^\top] \begin{bmatrix} A & B \frac{\Pr(o)}{\Pr(e)} \\ 0 & C \frac{\Pr(o)}{\Pr(e)} \end{bmatrix}^{-\top} \quad (40)$$

$$= \frac{\Pr(o)}{\Pr(e)^2} \begin{bmatrix} A^{-1} & -A^{-1} B C^{-1} \\ 0 & C^{-1} \frac{\Pr(e)}{\Pr(o)} \end{bmatrix} \mathbb{E}_o[\psi(Z; \theta_*, f_0) \psi(Z; \theta_*, f_0)^\top] \begin{bmatrix} A^{-1} & 0 \\ -C^{-\top} B^\top A^{-\top} & C^{-1} \frac{\Pr(e)}{\Pr(o)} \end{bmatrix} \quad (41)$$

The V_e part of the variance is the variance if one ignores the uncertainty in the surrogate index model and simply estimates uncertainty as if the surrogate index was the target outcome. So this is the uncertainty in estimating the causal effect of the treatment on the surrogate index. The variance V_o is the influence of the uncertainty of estimating the surrogate index on the final treatment effect.

We can further expand and simplify the variance V_o in particular, we can write the top left $k \times k$ diagonal block in the following simplified form:

$$[V_o]_{1:k, 1:k} = \frac{\Pr(o)}{\Pr(e)^2} A^{-1} \left(\mathbb{E}_o[\Psi_{0,o} \Psi_{0,o}^\top] - \sum_{t=2}^M B_t C_t^{-1} \left(\mathbb{E}_o[\Psi_t \Psi_{0,o}^\top] - \sum_{t'=2}^M \mathbb{E}_o[\Psi_t \Psi_{t'}^\top] C_{t'}^{-\top} B_{t'}^\top \right) \right) A^{-1} \quad (42)$$

which nicely also decomposes into the part that we had without the dynamic effect estimation and the extra part that stems from the estimation of the dynamic effects.

B A Step-by-Step Algorithmic Description

In this section we describe in a step-by-step algorithmic manner the exact data analysis pipeline we describe in the main text.

We assume access to two data-sets:

1. a long-term historical data set that contains n_o samples of tuples $(X_{i,t}, T_{i,t}, S_{i,t}, Y_{i,t})$, that for each sample i consists of M periods
2. a short-term historical data set that contains n_e samples of triplets $(S_{i,0}, X_{i,1}, T_{i,1}^{\text{new}}, S_{i,1})$.

The treatments $T_{i,t}$ in the long-term dataset are any potential historical investments not necessarily overlapping with new treatments.

We then perform the following steps:

1. On the long-term data-set:
 - (a) Estimate dynamic effects, $\theta_{t,\tau}$ for each $2 \leq t \leq \tau \leq M$ of the dynamic effect of treatment at time t on outcome at time τ , via double/debiased dynamic treatment effect estimation method outlined in Lewis & Syrgkanis (2020), applied to each target period separately.
 - (b) For each sample i and for each $t = \{1, \dots, M\}$, calculate dynamically adjusted cumulative outcome:

$$Y_{i,t}^{\text{adj}} = Y_{i,t} - \sum_{\tau=2}^t \theta_{\tau,t}^\top T_{i,\tau} \quad (43)$$

- (c) For each sample i , calculate dynamically adjusted cumulative long-term outcome:

$$\bar{Y}_{i,1}^{\text{adj}} := \sum_{t=1}^M Y_{i,t}^{\text{adj}} \quad (44)$$

- (d) Using data $(X_{i,1}, S_{i,1}, \bar{Y}_{i,1}^{\text{adj}})$, for each sample i : estimate the model of the dynamically adjusted surrogate index \hat{g}^{adj} , by training a model to predict $\bar{Y}_{i,1}^{\text{adj}}$ from $S_{i,1}, X_{i,1}$. Equivalently, learn a model of

$$g_0^{\text{adj}}(S_{i,1}, X_{i,1}) := \mathbb{E}[\bar{Y}_{i,1}^{\text{adj}} \mid S_{i,1}, X_{i,1}]. \quad (45)$$

In our experiments, we estimated this model via Lasso regression. We used k -fold cross-validation to choose the weight on the Lasso penalty.

2. On the short-term data-set:

- (a) For each sample i , impute a surrogate index:

$$I_i = \hat{g}^{\text{adj}}(S_{i,1}, X_{i,1}) \quad (46)$$

- (b) Estimate the causal effect θ^{new} of $T_{i,t}^{\text{new}} \in \mathbb{R}^{k'}$ of the baseline investments on $I_{i,t}$, controlling for $X_{i,t}, S_{i,t-1}$, via Double Machine Learning Chernozhukov et al. (2018a):
 - i. Split the samples in half.
 - ii. On the first half, estimate a model \hat{p} that predicts $T_{i,t}^{\text{new}}$ from $X_{i,t}, S_{i,t-1}$, i.e. learns the conditional expectation $p_0(X_{i,t}, S_{i,t-1}) := \mathbb{E}[T_{i,t}^{\text{new}} \mid X_{i,t}, S_{i,t-1}]$. In our analysis we used Lasso regression for this prediction task, with a cross-validated penalty weight.
 - iii. On the first half, estimate a model \hat{q} that predicts $I_{i,t}$ from $X_{i,t}, S_{i,t-1}$, i.e. learns the conditional expectation $h_0(X_{i,t}, S_{i,t-1}) := \mathbb{E}[I_{i,t} \mid X_{i,t}, S_{i,t-1}]$. In our analysis we used Lasso regression for this prediction task, with a cross-validated penalty weight.

- iv. For every period and sample on the other half of customers, calculate surprise in treatment and surprise in surrogate index:

$$T_{i,t}^{\text{new,res}} = T_{i,t}^{\text{new}} - \hat{p}(X_{i,t}, S_{i,t-1}) \quad (47)$$

$$I_{i,t}^{\text{res}} = I_{i,t} - \hat{h}(X_{i,t}, S_{i,t-1}) \quad (48)$$

- v. Reverse the role of the two samples and repeat the last three steps, to calculate residuals for the other half of the samples.
 - vi. Using the treatment and surrogate index residuals from all samples, run a Linear Regression, regressing $I_{i,t}^{\text{res}}$ on $T_{i,t}^{\text{new,res}}$, to learn a vector of coefficients $\hat{\theta}^{\text{new}}$.
 - vii. Estimate the co-variance matrix $V \in \mathbb{R}^{k'} \times \mathbb{R}^{k'}$ of the estimate $\hat{\theta}^{\text{new}}$ using heteroskedasticity robust co-variance estimation for an ordinary least squares regression.
- (c) Construct 99% confidence intervals using asymptotic normal approximation of the distribution of estimate $\hat{\theta}_j^{\text{new}}$, with standard error $\sigma_j^2 := V_{j,j}/n_e$, i.e.:

$$\theta_{0,j} \in \hat{\theta}_j^{\text{new}} \pm 2.577 \cdot \sigma_j \quad (49)$$

Debiasing correction. To account for the uncertainty in the surrogate index and achieve asymptotically valid intervals on the true dynamic effect, we need to add a few further steps. In particular:

1. Split all samples in half (both in long-term and short-term dataset)
2. Train a classification model \hat{t} that predicts the population $\{o, e\}$, based on S_1, X_1, S_0 , i.e.:

$$t_0(S_1, X_1, S_0) = \Pr(e \mid S_1, X_1, S_0) \quad (50)$$

In our experiments we used a penalized logistic regression with a cross-validated penalty weight.

3. Train a regression model \hat{r} that predicts $\tilde{T}_{i,t}^{\text{new,res}}$ from S_1, X_1, S_0 , i.e.:

$$r_0(S_1, X_1, S_0) = \mathbb{E}[\tilde{T}_{i,t}^{\text{new,res}} \mid S_1, X_1, S_0] \quad (51)$$

In our experiments we used a lasso regression with a cross-validated penalty weight.

4. Evaluate these models on the other half to construct a surrogate score:

$$Q_{i,t} = \frac{\hat{t}(S_{i,1}, X_{i,1}, S_{i,0})}{1 - \hat{t}(S_{i,1}, X_{i,1}, S_{i,0})} \hat{r}(S_{i,1}, X_{i,1}, S_{i,0}) \quad (52)$$

5. Reverse the roles of the split-samples and repeat last three steps.
6. Add a de-biasing correction to the estimate $\hat{\theta}^{\text{new}}$:

$$\hat{\theta}^{\text{new,deb}} = \hat{\theta}^{\text{new}} + \frac{1}{n_o} \sum_{i \in o} \left(\bar{Y}_{i,1}^{\text{adj}} - \hat{g}^{\text{adj}}(S_{i,1}, X_{i,1}) \right) Q_{i,t} \quad (53)$$

7. Add variance correction to the computed variance V :

$$V^{\text{deb}} = \frac{n}{n_e} V + J_{e,n}^{-1} \Sigma_{o,n} J_{e,n}^{-1} \quad (54)$$

where $J_{e,n}$ and $\Sigma_{o,n}$ are empirical analogues of the quantities J_e and Σ_o of Theorem 3.3.

8. Construct 99% confidence intervals using asymptotic normal approximation of the distribution of estimate $\hat{\theta}_j^{\text{new}}$, with variance $\sigma_{j,\text{deb}}^2 := V_{j,j}^{\text{deb}}/n = V/n_e + J_{e,n}^{-1} \Sigma_{o,n} J_{e,n}^{-1}/n = \sigma_j^2 + J_{e,n}^{-1} \Sigma_{o,n} J_{e,n}^{-1}/n$, i.e.:

$$\theta_{0,j} \in \hat{\theta}_j^{\text{new,deb}} \pm 2.577 \cdot \sigma_{j,\text{deb}} \quad (55)$$

C Description of Data Generating Processes

C.1 Synthetic Data

For fully synthetic data we simply generate data based on the linear data generating process presented in Equation (13), with Gaussian exogenous shocks and randomly initialized parameter matrices.

C.2 Semi-Synthetic Data

As mentioned in Section 4, we generate the semi-synthetic data by leveraging the correlation matrix and some pre-trained models from a real world dataset. In this section, we describe in detail how we simulate the dataset.

The real world dataset contains approximately 10k customers. For each customer we collect a time series of their monthly investments, proxies, and revenue trajectory, along with a set of fixed customer characteristics. We extract meaningful information from the real data and then simulate new synthetic data as follows:

Generate data for initial period. From the real world dataset, we filter one month of data on some period t and derive the normalized covariance matrix across all variables. In order to not expose the real correlation among variables, we decompose this matrix and recreate new eigenvalues and eigenvectors that retain key characteristics. For eigenvalues, we keep the top 4 eigenvalues and fit a discontinuous linear regression on the true eigenvalues curve. For eigenvectors, we keep the corresponding 4 vectors and impute random remaining vectors. We then combine these new eigenvalues and eigenvectors to create a new covariance matrix. From this perturbed covariance matrix we draw a sample on a multi-variate gaussian distribution, which we use as the data for the initial period t_0 for each customer.

Learn intertemporal auto and cross-correlations. We first train linear models (e.g. LassoCV) on each proxy and investment to predict each of these outcomes using 6 lagged periods of each investment, 6 lagged periods of each proxy, and a set of time-invariant demographics. For both proxies and investments, we find a large amounts of auto-correlation and small cross-correlation effects, along with some effects from customer characteristics including customer size and level of engagement. Based on these insights, we then draw new coefficients for each model, with decaying trend on lagged variables in different scale, and randomly draw coefficients for demographics.

Simulate the residual distribution corresponding to each model. From the pre-trained models, we see how unobserved heterogeneity behaves on each investment and proxy. we find that majority of our residuals follow a normal distribution, however, we have some unpredictable large outliers on both sides. In order to capture those surprising behaviours, we fit a mixture of different models. For a residual distribution from a single model, we fit mixture gaussian model ($n_component=2$) on the 5% and 95% band of the residuals and for the two tails, we fit log normal distributions.

Build the panel dataset in a feed-forward manner. With the initial data set, proxy and investment parametric model coefficients and the residuals, we simulate the data following the equations below:

$$T_{i,t} = \sum_{j=1}^6 \kappa_j T_{i,t-j} + \sum_{j=1}^6 \alpha_j S_{i,t-j} + \lambda D_{i,t} + \eta_{i,t} \quad (56)$$

$$S_{i,t} = \theta T_{i,t} + \sum_{j=1}^6 \gamma_j S_{i,t-j} + \beta D_{i,t} + \epsilon_{i,t} \quad (57)$$

where T represents treatment, S represents surrogates(proxies), and D represents demographics. First of all, we use the investment model to predict the current period investment, and add random residuals drawing from our fitted residual mixture distribution. Then we use the proxy model to predict current period proxy from all the controls, and add residuals as well. Other than that, we also add the effect of current period investment to the proxy. Moving to the next period, all the predictive outcomes and controls in the current period will become controls for next period, we could then repeat the process mentioned above again to get the next period prediction. Figure 1 also shows how this forward-feeding iteration works. After repeating this process m times, we have a semi-synthetic panel data ready to run experiments.

D Coverage Results of Asymptotic Normal Based Intervals

D.1 Synthetic Data

$n_t =$	2000	5000	10000
$n_t = 2$	96	93.5	96
4	79	76	86
8	53.3	57.5	62

(a) Dynamic Adjusted Surrogate

$n_t =$	2000	5000	10000
$n_t = 2$	96.5	93	97
4	84.5	83.5	87
8	75	71.5	81

(b) De-biased Dynamic Adjusted Surrogate

Figure 3: Synthetic data, high-dimensional, lasso models. Coverage levels averaged across $n_{exp} = 100$ experiments and across the $k = 2$ treatments. n is number of samples and n_t number of periods of long-term outcome. Target coverage level is 99%.

$n_t =$	2000	5000	10000
$n_t = 2$	98.5	99	98.5
4	89	91	89.5
8	80	75	80

(a) Dynamic Adjusted Surrogate

$n_t =$	2000	5000	10000
$n_t = 2$	98.5	99	98.5
4	95.5	94	93
8	90.5	90.5	91

(b) De-biased Dynamic Adjusted Surrogate

Figure 4: Synthetic data, low-dimensional, linear regression models. Coverage levels averaged across $n_{exp} = 100$ experiments and across the $k = 2$ treatments. n is number of samples and n_t number of periods of long-term outcome. Target coverage level is 99%.

D.2 Semi-Synthetic Data

$n_t =$	2000	5000	10000
$n_t = 2$	53	42	34
4	49	42	36
8	42	26	

(a) Dynamic Adjusted Surrogate

$n_t =$	2000	5000	10000
$n_t = 2$	65	60	52
4	76	71	70
8	87	84	

(b) De-biased Dynamic Adjusted Surrogate

Figure 5: Semi-Synthetic data, lasso models. Coverage levels averaged across $n_{exp} = 100$ experiments and across the $k = 3$ treatments. n is number of samples and n_t number of periods of long-term outcome. Target coverage level is 99%.

E Bias from Auto-correlated Treatments

Our final exercise illustrates the increasing importance of the dynamic adjustment step as the auto-correlation between treatments increases. Using the fully synthetic data generation described in section ??, we create a series of simulated samples with increasingly strong auto-correlation in the treatment (D in equation 13). We then estimate the average treatment effect using the canonical unadjusted surrogate model and the dynamically adjusted surrogate model proposed in this paper, corresponding to the 2nd orange bar and the 4th red bar, respectively, in figure 7).

Figure 6 plots the mean MSE from the estimates of each model over 100 data draws at each degree of auto-correlation from $D = 0$ to $D = 0.8$. The auto-correlation in proxy variables, B in equation 13, is set to zero throughout this exercise to eliminate the indirect path for intertemporal correlation between treatments. As expected, the adjustment step matters less with no or weak correlation over time between treatments. In fact, the extra steps in the adjusted model lead to slightly higher estimation errors on average. However, as treatments become increasingly correlated, the bias in the unadjusted model becomes more severe.




Figure 6: Fully synthetic data with 10,000 samples in each draw and 8 periods. Figure plots the mean ℓ_2 error of recovered coefficients in each model across draws as the auto-correlation of the treatment, D in equation 13, increases from 0 to 0.8.

F MSE Results on Further Synthetic and Semi-Synthetic Data

We present results on the ℓ_2 error of the recovered coefficients, for synthetic and semi-synthetic data. The algorithms that whose performance we present is as follows:

1. total: estimating the effect with hypothetical access to the long-term outcome and no dynamic adjustment
2. surrogate: estimating the effect with hypothetical access to the long-term outcome and no dynamic adjustment, but first projecting to the surrogates and then estimating the effect on the surrogate index.
3. adj. total: estimating the effect with hypothetical access to the long-term outcome, applying dynamic adjustment
4. adj. surrogate: estimating the effect with hypothetical access to the long-term outcome, applying dynamic adjustment, after first projected the adjusted outcome on the surrogates and then estimating the effect on the dynamically adjusted surrogate index
5. new treat: our dynamically adjusted surrogate index algorithm that uses a separate long-term dataset to estimate the surrogate index and then applying it to the short term dataset.
6. deb new treat: the debiased (fully orthogonal) version of our dynamically adjusted surrogate index algorithm, based on orthogonal score.

E.1 Synthetic Data

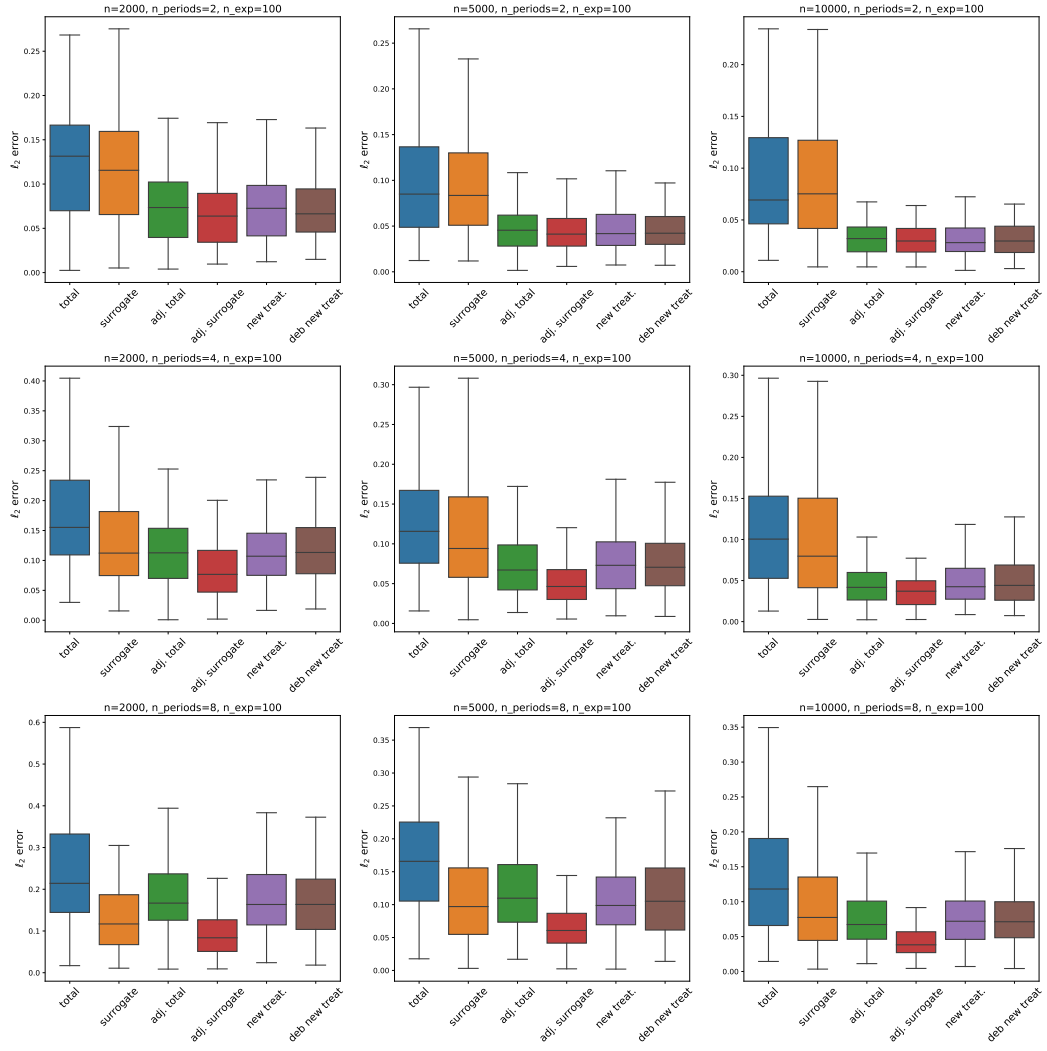


Figure 7: Synthetic data, high-dimensional, lasso models. n is number of samples and $n_{periods}$ number of periods of long-term outcome. ℓ_2 error of recovered coefficients.

F.2 Semi-Synthetic Data

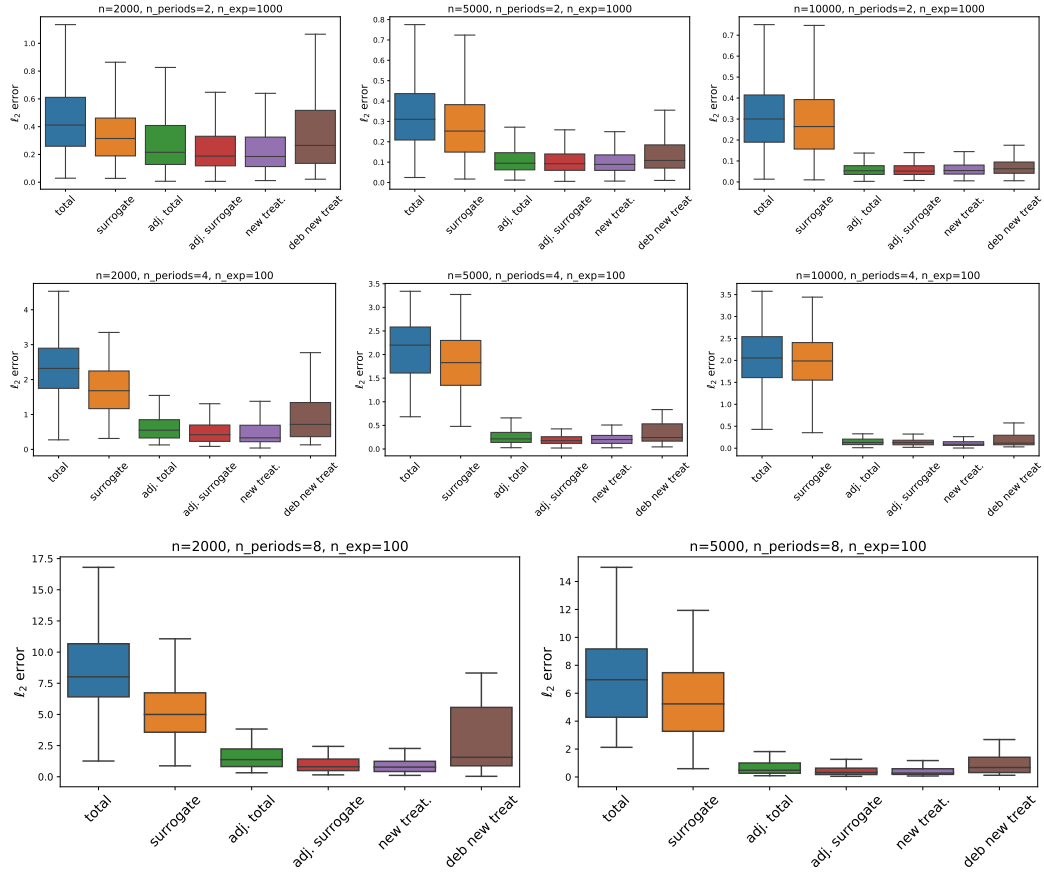


Figure 8: Semi-Synthetic data, high-dimensional, lasso models. n is number of samples and n_{periods} number of periods of long-term outcome. ℓ_2 error of recovered coefficients.