
Double/Debiased Machine Learning for Dynamic Treatment Effects

Greg Lewis
Microsoft Research
glewis@microsoft.com

Vasilis Syrgkanis
Microsoft Research
vasy@microsoft.com

Abstract

We consider the estimation of treatment effects in settings when multiple treatments are assigned over time and treatments can have a causal effect on future outcomes. We propose an extension of the double/debiased machine learning framework to estimate the dynamic effects of treatments and apply it to a concrete linear Markovian high-dimensional state space model and to general structural nested mean models. Our method allows the use of arbitrary machine learning methods to control for the high dimensional state, subject to a mean square error guarantee, while still allowing parametric estimation and construction of confidence intervals for the dynamic treatment effect parameters of interest. Our method is based on a sequential regression peeling process, which we show can be equivalently interpreted as a Neyman orthogonal moment estimator. This allows us to show root- n asymptotic normality of the estimated causal effects.

1 Introduction

Improving outcomes often requires multiple treatments: patients may need a course of drugs to manage or cure a disease, soil may need multiple additives to improve fertility, companies may need multiple marketing efforts to close the sale. To make data-driven decisions, policy-makers need precise estimates of what will happen when a new policy is pursued. Because of its importance, this topic has been studied by many communities and under multiple regimes and formulations; examples include the field of reinforcement learning [34], longitudinal analysis in biostatistics [2], the dynamic treatment regime in causal inference and adaptive clinical trials [19].

This paper offers a new method for estimating and making inferences about the counterfactual effect of a new treatment policy. Our method is designed to work with observational data, in an environment with multiple treatments (either discrete or continuous), and a high-dimensional state. Valid causal inference is necessary to correctly attribute changes in outcomes to the different treatments applied. But it is more challenging than in a static context, since there are multiple causal pathways from treatments to subsequent outcomes (e.g. directly, or by changing future states, or by affecting intermediate outcomes, or by influencing future treatments).

Our work bridges many distinct literatures. The first is the econometrics literature on semi-parametric inference [22, 33, 21, 1, 7, 6]. We extend this literature, which has typically focused on static treatment regimes, to consider a semi-parametric Markovian model with flexible high-dimensional state and more broadly dynamic treatment regimes. We propose an estimation algorithm that estimates the dynamic effects of interest, from observational data, at parametric root- n rates. We prove asymptotic normality of the estimates, even when the state is high-dimensional and machine learning algorithms are used to control for indirect effects of the state, as opposed to effects stemming from the treatments. Our formulation can be viewed as a dynamic version of Robinson’s classic partial linear model in semi-parametric inference [33], where the controls evolve over time and are affected by prior

treatments. Our estimation algorithm builds upon and generalizes the recent work of [6, 8] to estimate not only contemporaneous effects, but also dynamic effects over time. In particular, we propose a sequential residualization approach, where the effects at every period are estimated in a Neyman orthogonal manner and then peeled-off from the outcome, so as to define a new “calibrated” outcome, which will be used to estimate the effect of the treatment in the previous period.

In doing this, we build on results in the semi-parametric inference literature in bio-statistics [3, 37, 36] on estimating causal effects in structural nested models (see [39, 29, 16, 4] for recent overviews). In particular, our identification strategy for dynamic treatment effects is a variant of the well-studied g -formula and g -estimation for structural nested mean models (SNMMs) [27, 30, 28, 32, 31, 29, 20, 38]. The works cited above have provided doubly robust estimators for this setting. One practical challenge faced by these estimation approaches is that the nuisance functions required are hard to estimate when there are many treatments or treatments are continuous.

Our approach addresses these challenges by providing a Neyman orthogonal g -estimation algorithm for linear structural nested mean models [28, 29], that allows for both continuous and discrete treatments in each time period. We use a two stage algorithm, where in the first stage a sequence of regression and classification models are fitted and in the second stage a simple linear system of equations is solved. This approach also allows for an easy sample-splitting/cross-fitting approach, which allows the use of arbitrary machine learning approaches in the first stage.

While Neyman orthogonality is a weaker condition than double robustness, it is sufficient to achieve robustness to bias introduced by using machine learning to estimate the nuisance parameters. This approach thus permits machine learning for the nuisance model estimation, which is practically important in high-dimensional state space/control variable settings. We view this as the main contribution of this paper to the biostatistics literature. Other contributions are 1) we provide formal guarantees for the case of a single time-series dataset, as opposed to cross-sectional time-series; and 2) we extend to linear structural nested models with heterogeneous coefficients, thus allowing for infinite dimensional target functions. Using the recent framework of orthogonal statistical learning [12], these can also be estimated via machine learning techniques. The resulting algorithm is an extension of the popular in practice RLearner algorithm [24] for heterogeneous treatment effects, but in the dynamic treatment regime.

Our work is also closely related to the work on doubly robust estimation in longitudinal data analysis and in offline policy evaluation in reinforcement learning [23, 35, 26, 18, 17] from the causal machine learning community. However, one crucial point of departure from all of these works is that we formulate minimal parametric assumptions that allow us to avoid non-parametric rates or high-variance estimation. Typical fully non-parametric approaches in off-policy evaluation in dynamic settings, requires the repeated estimation of inverse propensity weights at each step, which leads to a dependence on quantities that can be very ill-posed in practice, such as the product of the ratios of states under the observational and the target policy. Moreover, these approaches are typically restricted to discrete states and actions. Our goal is to capture settings where the treatment can also be continuous (investment level, price discount level, drug dosage), and the state contains many continuous variables and is potentially high dimensional. Inverse propensity weighting approaches, though more robust in terms of the assumptions that they make on the model, can be quite prohibitive in these settings even in moderately large data sets.

Our work is also somewhat related to the online debiasing literature [11, 10, 42, 14], since we perform inference from adaptively collected data. But our work is different in two main ways. First, the inferential target in this paper is more general, since that literature restricts itself to inference on contemporaneous effects (i.e. effect of treatment on same period outcome) or on policy values in the absence of a dynamic endogenous state (i.e. assuming no dynamic effects through a Markovian state variable). Further, our results extend to the more general structural nested mean model setting. The second important difference is that, in our main setting, we assume multiple independent small chains of samples (each from a separate treated unit). This sidesteps one of the key technical difficulties in the literature above, which is performing inference from a single chain of serially correlated and adaptively collected samples. With multiple short independent chains, the main technical difficulty tackled by these works, namely the potential non-convergence of the covariance matrices that are used in the effect estimate, vanishes. In that respect, that literature is closer to our results on a single adaptive chain (see Section 6). However, for our results in this setting, we make a homoskedastic noise assumption in the treatment choice policy (i.e. that the noise term in the treatment policy enters additively and is drawn i.i.d.), which allows us to easily show convergence of the associated

covariance matrices of our Neyman orthogonal estimator and hence again sidestep a key technical difficulty. It is an interesting avenue for future research if this i.i.d. additive noise assumption can be dropped from our results on a single adaptive chain by utilizing ideas from this literature (e.g. adaptive re-weighted variants of our Neyman orthogonal estimator; as for instance described in [11]).

2 Model and Preliminaries

We state our main results in a stylized partially linear state space Markov decision process and in Section 7 we show that our results generalize to more complex Markovian models, known in the biostatistics and causal inference literature as Structural Nested Mean Models (SNMMs). For simplicity of exposition we focus on this simplified version, as it captures all the complexities needed to highlight our main contributions. We consider a sequence $\{X_t, T_t, Y_t\}_{t=1}^m$, where $X_t \in \mathbb{R}^p$ is the state at time t , $T_t \in \mathbb{R}^d$ is the action or treatment at time t and $Y_t \in \mathbb{R}$ is an observed outcome of interest at time t . We assume that these variables are related via a linear Markovian process:

$$\begin{aligned} \forall t \in \{1, \dots, m\} : \quad & X_t = A \cdot T_{t-1} + B \cdot X_{t-1} + \eta_t \\ & T_t = p(T_{t-1}, X_t) + \zeta_t \\ & Y_t = \theta'_0 T_t + \mu' X_t + \epsilon_t \end{aligned} \tag{1}$$

where η_t, ζ_t and ϵ_t are exogenous mean-zero random shocks, independent of all contemporaneous and lagged treatments and states, that for simplicity we assume are each drawn i.i.d. across time with $T_0 = X_0 = 0$. The function p could be any propensity function.

Our goal is to estimate the effect of a change in the treatment policy on the final outcome Y_m . More concretely, suppose that were to make an intervention and set each of the treatments to some sequence of values: $\{\tau_1, \dots, \tau_m\}$, then what would be the expected difference in the final outcome Y_m as compared to some baseline policy? For simplicity and without loss of generality, we will consider the baseline policy to be setting all treatments to zero. We will denote this expected difference as: $V(\tau_1, \dots, \tau_m)$. Equivalently, we can express the quantity we are interested in do-calculus: if we denote with $R(\tau) = \mathbb{E}[Y_m \mid do(T_1 = \tau_1, \dots, T_m = \tau_m)]$, then $V(\tau) = R(\tau) - R(0)$. In Appendix A, we will also analyze the estimation of the effect of adaptive counterfactual treatment policies, where the treatment at each step can be a function of the state.

3 Identification via Dynamic Effects

Our first observation is that we can decompose the quantity $V(\tau)$ into the estimation of the *dynamic treatment effects*: if we were to make an intervention and increase the treatment at period $m - \kappa$ by 1 unit, then what is the change θ_κ in the outcome Y_m , for $\kappa \in \{0, \dots, m\}$; assuming that we set all subsequent treatments to zero (or equivalently to some constant value). This quantity is the effect in the final outcome, that does not go through the changes in the subsequent treatments, due to our observational Markovian treatment policy, but only the part of the effect that goes through changes in the state space X_t , that is not part of our decision process. This effect can also be expressed in terms of the constants in our Markov process as:

$$\forall \kappa \in \{1, \dots, m - 1\} : \theta_\kappa = \mu' B^{\kappa-1} A$$

Lemma 1. *The counterfactual value function $V : \mathbb{R}^{d \cdot m} \rightarrow \mathbb{R}$, can be expressed in terms of the dynamic treatment effects as: $V(\tau_1, \dots, \tau_m) = \sum_{\kappa=0}^{m-1} \theta'_\kappa \tau_{m-\kappa}$*

Thus to estimate the function V , it suffices to estimate the dynamic treatment effects: $\theta_0, \dots, \theta_m$. We first start by showing that the parameters $\theta_0, \dots, \theta_m$ are identifiable from the observational data. Identification is not immediately obvious. Write the final outcome as a linear function of all the treatments and the initial state: $Y_m = \sum_{\kappa=0}^{m-1} \theta'_\kappa T_{m-\kappa} + \mu' \delta_m + \epsilon_m$. The shock δ_m contains components that are heavily correlated with the treatments, since the shocks at period t affect the state at period $t + 1$, which in turn affect the observed treatment at period $t + 1$. In other words, if we view the problem as a simultaneous treatment problem, where the final outcome is the outcome, then we essentially have a problem of unmeasured confounding (implicitly because we ignored the intermediate states).

However, we show that these dynamic effects can be identified from the observational data via a sequential peeling process, which as we show in the next section, leads to an estimation strategy that achieves parametric rates. This is one of the main contributions of the paper.

Theorem 2. *The dynamic treatment effects satisfy the following set of conditional moment restrictions: $\forall q \in \{0, \dots, m\}$*

$$\mathbb{E}[\bar{Y}_{m,m-q} - \theta_q T_{m-q} - \mu' B^q X_{m-q} \mid T_{m-q}, X_{m-q}] = 0$$

where: $\bar{Y}_{m,m-q} = Y_m - \sum_{\kappa=0}^{q-1} \theta_\kappa T_{m-\kappa}$. Moreover, if we let $\tilde{T}_{m-q} := T_{m-q} - \mathbb{E}[T_{m-q} \mid X_{m-q}]$ and we have that the covariance matrix $J := \mathbb{E}[\tilde{T}_{m-q} \tilde{T}_{m-q}^\top]$ is invertible, then these conditional moment restrictions uniquely identify θ_q .

Proof. For any $q \in \{0, \dots, m-1\}$, by repeatedly unrolling the state X_m , $q-1$ times we have that:

$$\begin{aligned} Y_m &= \theta'_0 T_m + \sum_{\kappa=0}^q \mu' B^{\kappa-1} A T_{m-\kappa} + \mu' B^q X_{m-q} + \sum_{\kappa=0}^{q-1} \mu' B^{\kappa-1} \eta_{m-\kappa} + \epsilon_m \\ &= \psi'_m T_m + \sum_{\kappa=0}^q \psi'_{m-\kappa} T_{m-\kappa} + \mu' B^q X_{m-q} + \sum_{\kappa=0}^q \mu' B^{\kappa-1} \eta_{m-\kappa} + \epsilon_m \\ &= \psi'_{m-q} T_{m-q} + \sum_{j=m-q+1}^m \psi'_j T_j + \mu' B^q X_{m-q} + \sum_{\kappa=0}^q \mu' B^{\kappa-1} \eta_{m-\kappa} + \epsilon_m \end{aligned}$$

Thus by re-arranging the equation, we have:

$$\bar{Y}_{m,m-q} = \psi'_{m-q} T_{m-q} + \mu' B^q X_{m-q} + \sum_{\kappa=0}^q \mu' B^{\kappa-1} \eta_{m-\kappa} + \epsilon_m$$

Since for all $t \geq m-q$, η_t, ϵ_t are subsequent random shocks to the variables T_{m-q}, X_{m-q} , we have that they are mean-zero conditional on T_{m-q}, X_{m-q} . Thus:

$$\mathbb{E}[\bar{Y}_{m,m-q} - \psi'_{m-q} T_{m-q} - \mu' B^q X_{m-q} \mid T_{m-q}, X_{m-q}] = 0$$

Thus we have concluded that for any $t \in \{1, \dots, m\}$:

$$\mathbb{E}[\bar{Y}_{m,t} - \psi'_t T_t - \mu' B^{m-q} X_t \mid T_t, X_t] = 0$$

which concludes the proof of the first part of the theorem.

For the second part, consider any q and let $t = m-q$. Note that by taking an expectation of the above condition for with respect to T_t conditional on X_t , we get that it implies

$$\mathbb{E}[\bar{Y}_{m,m-q} \mid X_{m-q}] = \psi'_{m-q} T_{m-q} - \mu' B^q X_{m-q} \mid T_{m-q}, X_{m-q}]$$

□

4 Dynamic DML Estimation

We now address the estimation problem. We assume that we are given access to n i.i.d. samples from the Markovian process, i.e. we are given n independent time-series, and we denote sample i , with $\{X_t^i, T_t^i, Y_t^i\}$. Our goal is to develop an estimator of the function V or equivalently of the parameter vector $\theta = (\theta_0, \dots, \theta_m)$. We will consider the case of a high-dimensional state space, i.e. $p \gg n$, but low dimensional treatment space and a low dimensional number of periods m , i.e. $d, m \ll n$ is a constant independent of n . We want to estimate the parameters θ at \sqrt{n} -rates and in a way that our estimator is asymptotically normal, so that we can construct asymptotically valid confidence intervals around our dynamic treatment effects and our estimate of the function V . The latter is a non-trivial task due to the high-dimensionality of the state space. For instance, the latter would be statistically impossible if we were to take the direct route of estimating the whole Markov process (i.e. the high-dimensional quantities A, B, μ): if these quantities have a number of non-zero coefficients that grows with n at any polynomial rate, then known results on sparse linear regression, preclude their estimation at root- n rates (see e.g. [40]). However, we are not really interested in these low-level parameters of the dynamic process, but solely on the low dimensional parameter vector θ . We will treat this problem as a semi-parametric inference problem and develop a Neyman orthogonal estimator for the parameter vector [22, 33, 1, 7, 6].

In particular, we consider a sequential version of the double machine learning algorithm proposed in [6]. In the case of a single time-period, i.e. $m = 0$, then [6], recommends the following estimator for θ_0 : using half of your data, fit a model $\hat{q}_0(X_0)$ of $\mathbb{E}[Y_0 | X_0]$, i.e. that predicts the outcome Y_0 from the controls X_0 and a model $\hat{p}_0(X_0)$ for $\mathbb{E}[T_0 | X_0]$. Then estimate θ_0 on the other half of the data, based on the estimating equation:

$$m(\theta; \hat{p}_0, \hat{q}_0) = \mathbb{E} \left[(\tilde{Y}_0 - \theta_0 \tilde{T}_0) \tilde{T}_0 \right] = 0$$

where $\tilde{Y}_0 = Y_0 - \hat{q}_0(X_0)$ and $\tilde{T}_0 = T_0 - \hat{p}_0(X_0)$ are the residual outcome and treatment.

We propose a sequential version of this process that we call *Dynamic DML for sequential double/debiased machine learning*. Intuitively our algorithm proceeds as follows:

1. We can construct an estimate $\hat{\theta}_0$ of θ_0 in a robust (Neyman orthogonal) manner, by applying the approach of [6] on the final step of the process, i.e. on time step T_m ; this will estimate all the contemporaneous effects of the treatments,
2. Subsequently we can remove the effect of the observed final step treatment from the observed final step outcome, i.e. by re-defining the random variable $Y_{m,m-1}^i = Y_m^i - \hat{\theta}_0 T_m^i$; doing this we have removed any effects on Y_m^i , caused by the final treatment T_m^i .
3. We can then estimate the one-step dynamic effect θ_1 , by performing the residual-on-residual estimation approach with target outcome the ‘‘calibrated’’ outcome $Y_{m,m-1}^i$, treatment T_{m-1}^i and controls X_{m-1}^i . Theorem 2 tells us that the required *conditional exogeneity* moment required to apply the residualization process is valid for these random variables. We can continue in a similar manner, by removing the estimated effect of T_{m-1}^i from $Y_{m,m-1}^i$ and repeating the above process.

We provide a formal statement of the Dynamic DML process in Algorithm 1, which also describes more formally the sample splitting and cross-fitting approach that we follow in order to estimate the nuisance models p and q required for calculating the estimated residuals.

Algorithm 1 Dynamic DML

Randomly split the n samples in S, S'

for each $\kappa \in \{0, \dots, m\}$ **do**

Regress Y_m on $X_{m-\kappa}$ using S to learn estimate \hat{q}_κ of model $q_\kappa(x) = \mathbb{E}[Y_m | X_{m-\kappa} = x]$ and calculate residuals $\tilde{Y}_{m,m-\kappa}^i = Y_m^i - \hat{q}_\kappa(X_{m-\kappa}^i)$ on other half; vice versa use S' to learn model and evaluate on S .

for each $\tau \in \{0, \dots, \kappa\}$ **do**

Regress $T_{m-\tau}$ on $X_{m-\kappa}$ using S to learn estimate $\hat{p}_{\tau,\kappa}$ of model $p_{\tau,\kappa}(x) = \mathbb{E}[T_{m-\tau} | X_{m-\kappa} = x]$ on the first half and calculate residuals $\tilde{T}_{m-\tau,m-\kappa}^i = T_{m-\tau}^i - \hat{p}_{\tau,\kappa}(X_{m-\kappa}^i)$ on other half, and vice versa.

end for

end for

Using all the data $S \cup S'$

for $\kappa = 0$ **to** m **do**

Regress $\tilde{Y}_{m,\kappa} = \tilde{Y}_{m,m-\kappa} - \sum_{\tau < \kappa} \hat{\theta}_\tau \tilde{T}_{m-\tau,m-\kappa}$ on $\tilde{T}_{m-\kappa,m-\kappa}$, i.e. find a solution $\hat{\theta}_\kappa$ to:

$$\frac{1}{n} \sum_{i=1}^n \left(\tilde{Y}_{m,\kappa}^i - \hat{\theta}_\kappa \tilde{T}_{m-\kappa,m-\kappa}^i \right) \tilde{T}_{m-\kappa,m-\kappa}^i = 0$$

end for

5 Estimation Rates and Normality

Our main theorems are to show that subject to the first stage models of the conditional expectations achieving a small (but relatively slow) estimation error, then the recovered parameters are root- n -consistent and asymptotically normal. Our asymptotic normality proof relies on showing that one can

re-interpret our *Dynamic DML* estimator as a Z -estimator¹ based on a set of moments that satisfy the property of Neyman orthogonality².

Let p, q denote the vector of all nuisance functions and p^*, q^* their true values. Moreover, let θ denote the vector of dynamic effect parameters, with θ_0 its true value. We provide both finite sample rates mean squared error rates and asymptotic normality of our estimates (proofs in Appendix B.3 and B.4). Moreover, in the appendix we discuss how the required guarantees for the nuisance functions can be easily satisfied using estimation algorithms such as the Lasso and under sparsity conditions.

Theorem 3 (Finite Sample). *Suppose that $\mathbb{E}[\zeta_t \zeta_t'] \succeq 2\lambda I$ and that each coordinate \hat{h} of each nuisance model $\{q_\kappa, p_{\tau, \kappa}\}_{\tau, \kappa}$ satisfy that w.p. $1 - \delta$:*

$$\|\hat{h} - h^*\|_2 = \sqrt{\mathbb{E}_X[(\hat{h}(X) - h(X))^2]} \leq \epsilon_{n, \delta}$$

where expectation is with respect to the corresponding input of each model. Moreover, suppose that.³

$$\left\| \sum_{\tau=0}^m \mathbb{E} \left[\tilde{T}_{m-\tau, m-\kappa}^* \zeta_t' \right] \right\|_\infty \leq c_m$$

Then w.p. $1 - 2\delta$:

$$\max_{t \in [m]} \|\hat{\theta}_t - \theta_{0, t}\|_2 \leq O \left(d C_m \left(\sqrt{\frac{\log(dm/\delta)}{n}} + \epsilon_{n, \delta}^2 \right) \right)$$

where $C_m := \max \left\{ 1, \left(\frac{c_m}{2\lambda - O(m d \epsilon_{n, \delta})} \right)^{m+1} \right\}$. If each coordinate \hat{h} of each nuisance model satisfies: $\mathbb{E}_{\hat{h}} [\|\hat{h} - h^*\|_2^4]^{(1/4)} \leq \epsilon_n$, then:

$$\sqrt{\mathbb{E} \left[\max_{t \in [m]} \|\hat{\theta}_t - \theta_{0, t}\|_2^2 \right]} \leq O \left(d C_m \left(\sqrt{\frac{\log(dm)}{n}} + \epsilon_n^2 \right) \right).$$

Theorem 4 (Asymptotic Normality). *Suppose that the nuisance models $\{q_\kappa, p_{\tau, \kappa}\}_{\tau, \kappa}$ satisfy that:*

$$\forall \kappa, \tau \leq \kappa : \|q_\kappa - q_\kappa^*\|_2, \|p_{\tau, \kappa} - p_{\tau, \kappa}^*\|_2 \leq o(n^{-1/4})$$

Moreover, suppose that all random variables are bounded and that $\mathbb{E}[\zeta_t \zeta_t'] \succeq \lambda I$. Then:

$$\sqrt{n}(\hat{\theta} - \theta_0) \rightarrow N(0, V)$$

where $V = M^{-1} \Sigma (M^{-1})'$ and M is a $m \times m$ block lower triangular matrix consisting of blocks of size $d \times d$, whose block entries are of the form:

$$\forall \tau \leq \kappa : M_{\kappa, \tau} = \mathbb{E}[(T_{m-\tau} - \mathbb{E}[T_{m-\tau} | X_{m-\kappa}]) \zeta_{m-\kappa}']$$

and Σ is a $m \times m$ block diagonal matrix whose diagonal block entries take the form: $\Sigma_{tt} = \mathbb{E}[\epsilon_{m-t}^2 \zeta_{m-t} \zeta_{m-t}']$.

Concrete Rates for Lasso Nuisance Estimates. Suppose that the observational policy p is also linear, i.e. $p(X) = AX$. Then all the models q_κ and $p_{\tau, \kappa}$ are high-dimensional linear functions of their input arguments, i.e. $q_\kappa(x) = \phi' x$ and $p_{\tau, \kappa, j}(x) = \pi'_{\tau, \kappa, j} x$. If these linear functions satisfy a sparsity constraint then under standard regularity assumptions we can guarantee if we use the Lasso regression to estimate each of these functions that w.p. $1 - \delta$, the estimation error of all nuisance models is $O \left(s \sqrt{\frac{\log(p/\delta)}{n}} \right)$, where s is an upper bound on the number of non-zero coefficients.

One sufficient regularity condition is that the expected co-variance matrix of every period's state has full rank, i.e. $\mathbb{E}[X_t X_t'] \succeq \lambda I$. Thus the requirements of the main theorems of this section would be satisfied as long as the sparsity grows as $s = o(n^{1/4})$. These conditions are for instance satisfied if s coordinates of the high-dimensional state have any effect on the final outcome (i.e. are outcome-relevant) and if s coordinates of the high-dimensional state enter the observational policy.

¹A Z -estimator is a solution to an empirical analogue of a vector of moment equations $\mathbb{E}[m(W; \theta, \nu)] = 0$, where W are all random variables, m is a vector valued function and $\theta \in \mathbb{R}^d$ is a vector parameter of interest and $\nu \in \mathcal{V}$ is a potentially infinite dimensional nuisance parameter. The true parameters satisfy $\mathbb{E}[m(W; \theta_0, \nu_0)] = 0$.

²A vector of moments satisfies Neyman orthogonality if $\forall \nu \in \mathcal{V} : \frac{\partial}{\partial t} \mathbb{E}[m(W; \theta_0, \nu_0 + t(\nu - \nu_0))] |_{t=0} = 0$

³Where by absolute value we denote coordinate-wise

6 Dependent Single Time-Series Samples

Thus far we have assumed that we are working with n independent time series, each of duration m . Though this is applicable to many settings where we have panel data with many units over time, in some other settings it is unreasonable to assume that we have many units over time, but rather that we have the same unit over a long period. In this case, we would want to do asymptotics as the number of periods grows. Our goal is still to estimate the dynamic treatment effects (i.e. the effect of a treatment at period t on an outcome in period $t + \kappa$, for $\kappa \in \{0, \dots, m\}$) for some fixed look-ahead horizon m .

These quantities can allow us to evaluate the effect of counterfactual treatment policies on the discounted sum of the outcomes, i.e. $\sum_{t=0}^{\infty} \gamma^t Y_t$ for $\gamma < 1$. We can write the counterfactual value function for any non-adaptive policy as: $V(\tau) = \sum_{t=0}^{\infty} \gamma^t \sum_{q \leq t} \theta_{t-q} \tau_q$. Assuming outcomes are bounded, the effect $\sum_{q \leq t} \theta_{t-q} \tau_q$ on any period t can be at most some constant. Thus taking m to be roughly $\log_{1/\gamma}(n)$, suffices to achieve a good approximation of the effect function $V(\pi)$, since the rewards vanish after that many periods, i.e. if we let: $V_m(\tau) = \sum_{t=0}^m \gamma^t \sum_{q \leq t} \theta_{t-q} \tau_q$, then observe that: $\|V_m(\tau) - V(\tau)\| \leq O(\gamma^m)$. Thus after $m = \log_{1/\gamma}(n)$, we have that the approximation error is smaller than $1/\sqrt{n}$. Thus it suffices to learn the dynamic treatment effect parameters for a small number of steps. To account for this logarithmic growth, we will make the dependence on m explicit in our theorems below.

For any m , we will estimate these parameters by splitting the time-series into sequential $B = n/m$ blocks of size m . Then we will treat each of these blocks roughly as independent observations and we denote the resulting estimate as $\hat{\theta}$. The main challenge in our proofs is dealing with the fact that these blocks are not independent but serially correlated. However, we can still apply techniques, such as martingale Bernstein concentration inequalities and martingale Central Limit Theorems to achieve the desired estimation rates.

The other important change that we need to make is in the way that we fit our nuisance estimates. To avoid using future samples to train models that will be used in prior samples (which would ruin the martingale structure), we instead propose a progressive nuisance estimation fitting approach, where at every period, all prior blocks are used to train the nuisance models and then they are evaluated on the next block. We present the normality theorem and defer the finite sample MSE result to the appendix.

Theorem 5. *Let \mathcal{F}_b denote the filtration up until block b . Suppose that $\mathbb{E}[\zeta_t \zeta_t' | \mathcal{F}_b] \succeq 2\lambda I$ and that each coordinate \hat{h} of each nuisance model $\{q_\kappa, p_{\tau, \kappa}\}_{\tau, \kappa}$ satisfies: $\forall b \geq B/2$*

$$\mathbb{E}_{\hat{h}}[\|\hat{h} - h^*\|_{b,2}^4]^{(1/4)} \leq o(B^{-1/4})$$

Moreover, suppose that: $\left\| \sum_{\tau=0}^m \mathbb{E} \left[\tilde{T}_{m-\tau, m-\kappa}^* \zeta_t' | \mathcal{F}_b \right] \right\|_{\infty} \leq c_m$ and that m satisfies that $m C_m = o(\sqrt{B})$, where $C_m := \max \left\{ 1, \left(\frac{c_m}{2\lambda - O(m d \epsilon_n, \delta)} \right)^{m+1} \right\} m$. Then:

$$\sqrt{B} \Sigma^{-1/2} M(\hat{\theta} - \theta_0) \rightarrow N(0, I)$$

where M, Σ are the same as in Theorem 4.

7 Generalization to Structural Nested Mean Models (SNMMs)

We now generalize our results by drawing a connection between our algorithm and g -estimation of structural nested models in biostatistics [27] and show how our DynamicDML algorithm can be applied almost verbatim with a small change of variables and estimate dynamic effects in a much more general setting. Consider an arbitrary time-series process $\{X_t, T_t\}_{t=1}^m$, with $X_t \in \mathcal{X}_t$ and $T_t \in \mathcal{T}_t$. Let Y denote some final outcome of interest. For any time t , let $\bar{X}_t = \{X_1, \dots, X_t\}$ and $\bar{T}_t = \{T_1, \dots, T_t\}$, denote the sequence of the variables up until time t and similarly, let $\underline{X}_t = \{X_t, \dots, X_m\}$ and $\underline{T}_t = \{T_t, \dots, T_m\}$. We will also denote with $\bar{x}_t, \bar{\tau}_t, \underline{x}_t, \underline{\tau}_t$, corresponding realizations of the latter random sequences. Let $\pi = (\pi_1, \dots, \pi_m)$ denote any dynamic policy, such that for each t , π_t maps a history $\bar{x}_t, \bar{\tau}_{t-1}$ into a next period action τ_t . For any such dynamic policy, let $Y^{(\pi)}$ denote the counterfactual outcome under policy d . For any static policy $\tau \in \times_{t=1}^m \mathcal{T}_t$, we will overload notation and let $Y^{(\tau)}$ denote the counterfactual outcome under this static treatment policy.

Moreover, for any two policies (static or dynamic) we will be denoting with $(\bar{\pi}'_t, \bar{\pi}_{t+1})$, the policy that follows π' up until time t and then continues with policy π . We let $0 \in \mathcal{T}_t$ denote a baseline policy value, which could be appropriately instantiated based on the context.

We assume that the data generating process satisfies the *sequential conditional exogeneity condition*:

$$\{Y^{(\tau)}, \tau \in \times_{t=1}^m \mathcal{T}_t\} \perp\!\!\!\perp T_t \mid \bar{T}_{t-1}, \bar{X}_t$$

Identification of mean counterfactual outcomes $\mathbb{E}[Y^{(\pi)}]$ for a target policy of interest π can be expressed in terms of the following conditional expectation functions:

$$\gamma_t(\bar{x}_t, \bar{\tau}_t) = \mathbb{E} \left[Y^{(\bar{\tau}_t, \bar{\pi}_{t+1})} - Y^{(\bar{\tau}_{t-1}, 0, \bar{\pi}_{t+1})} \mid \bar{T}_t = \bar{\tau}_t, \bar{X}_t = \bar{x}_t \right]$$

which corresponds to the mean change in outcome if we go to all units which received treatment τ_t up until time t and had observed state history \bar{x}_t and we remove their last treatment, while we subsequently always continue with the target policy π . These functions are known as the *blip* functions [4, 29]. Theorem 3.1 of [29] shows that via a telescoping sum argument and invoking the sequential randomization condition that:

$$\mathbb{E}[Y^{(\bar{\tau}_{t-1}, \bar{\pi}_t)} \mid \bar{X}_t, \bar{T}_t = \bar{\tau}_t] = \mathbb{E} \left[Y + \sum_{j=t}^m \rho_j(\bar{X}_j, \bar{T}_j) \mid \bar{X}_t, \bar{T}_t = \bar{\tau}_t \right]$$

where $\rho_j(\bar{X}_j, \bar{T}_j) := \gamma_j(\bar{X}_j, (\bar{T}_{j-1}, \pi(\bar{X}_j, \bar{T}_{j-1}))) - \gamma_j(\bar{X}_j, \bar{T}_j)$ and hence also: $\mathbb{E}[Y^{(\pi)}] = \mathbb{E} \left[Y + \sum_{t=1}^m \rho_t(\bar{X}_t, \bar{T}_t) \right]$. Importantly, the conditioning set contains the observed t periods treatment. Intuitively, each term ρ_j , removes from the outcome the *blip effect* of the observed action T_j and adds the *blip effect* of the target action $\pi(\bar{X}_j, \bar{T}_{j-1})$. Thus for any parameterization of the blip functions $\gamma_t(\bar{x}_t, \bar{\tau}_t; \psi_t)$, if we let $H_t(\psi) := Y + \sum_{j=t}^m \rho_j(\bar{X}_j, \bar{T}_j)$, then the true parameter vector must satisfy the moment restrictions:

$$\forall t \in [m], \forall f : \mathbb{E}[H_t(\psi) (f(\bar{X}_t, \bar{T}_t) - \mathbb{E}[f(\bar{X}_t, \bar{T}_t) \mid \bar{X}_t, \bar{T}_{t-1}])] = 0$$

If the blip functions admit a linear parametric form:

$$\gamma_t(\bar{x}_t, \bar{\tau}_t; \psi_t) := \psi_t' \phi(\bar{x}_t, \bar{\tau}_t)$$

for some known feature vector ϕ , satisfying $\phi(\bar{x}_t, (\bar{\tau}_{t-1}, 0)) = 0$, then we can consider the subset of the moment restrictions of the form:

$$\forall t \in [m] : \mathbb{E}[H_t(\psi) (\phi(\bar{X}_t, \bar{T}_t) - \mathbb{E}[\phi(\bar{X}_t, \bar{T}_t) \mid \bar{X}_t, \bar{T}_{t-1}])] = 0$$

Moreover, we can also subtract from $H_t(\psi)$, its conditional expectation $\mathbb{E}[H_t(\psi) \mid \bar{X}_t, \bar{T}_{t-1}]$ while maintaining the moment condition:

$$\forall t \in [m] : \mathbb{E}[(H_t(\psi) - \mathbb{E}[H_t(\psi) \mid \bar{X}_t, \bar{T}_{t-1}]) (\phi(\bar{X}_t, \bar{T}_t) - \mathbb{E}[\phi(\bar{X}_t, \bar{T}_t) \mid \bar{X}_t, \bar{T}_{t-1}])] = 0$$

This is the doubly robust moment condition proposed by [28, 29], where it is shown that an estimator of ψ based on this moment is correct if either the estimate of $\mathbb{E}[H_t(\psi) \mid \bar{X}_t, \bar{T}_{t-1}]$ or the estimate of $\mathbb{E}[\phi(\bar{X}_t, \bar{T}_t) \mid \bar{X}_t, \bar{T}_{t-1}]$ is correct. However, estimating $\mathbb{E}[H_t(\psi) \mid \bar{X}_t, \bar{T}_{t-1}]$, itself requires seemingly knowledge of ψ , which is cumbersome and can be im-practical (e.g. by constructing poor preliminary estimates of ψ). [16] (see Technical Point 21.5) note that for a binary treatment and when the target policy is the all-zero policy, then $\mathbb{E}[H_t(\psi) \mid \bar{X}_t, \bar{T}_t]$ can be estimated by regressing the outcome of the population that received zero subsequent treatment on the history. However, such a population can be quite small in practice and can have severe co-variate imbalances compared to the overall population. Moreover, this approach only applies to the case of a binary treatment and a static target policy.

Here, we can achieve a Neyman orthogonal moment for identifying ψ , which is sufficient for robustness to biases stemming from machine learning models used to train the nuisance components, while avoiding the cumbersome part of estimating the nuisance $\mathbb{E}[H(\psi) \mid \bar{X}_t, \bar{T}_{t-1}]$. Moreover, our approach leads to a strongly convex loss for the parameters at each step of the recursive process, which is beneficial for finite sample guarantees and subsequently for generalizing it to linear models with parameter heterogeneity with respect to exogenous covariates.

In particular, let $Q_j := \phi(\bar{X}_j, \bar{T}_j) - \phi(\bar{X}_j, (\bar{T}_{j-1}, \pi(\bar{X}_j, \bar{T}_{j-1})))$ and for any $t \leq j$, let $\tilde{T}_{j,t} := Q_j - \mathbb{E}[Q_j \mid \bar{X}_t, \bar{T}_{t-1}]$ and $\tilde{Y}_{m,t} := Y - \mathbb{E}[Y \mid \bar{X}_t, \bar{T}_{t-1}]$. Then, note that:

$$H_t(\psi) - \mathbb{E}[H_t(\psi) \mid \bar{X}_t, \bar{T}_{t-1}] = \tilde{Y}_{m,t} - \sum_{j=t}^m \psi_j' \tilde{T}_{j,t}$$

Moreover, note that since the second term in Q_t only depends on the conditioning set \bar{X}_t, \bar{T}_{t-1} , then $\tilde{T}_{t,t} = \phi(\bar{X}_t, \bar{T}_t) - \mathbb{E}[\phi(\bar{X}_t, \bar{T}_t) \mid \bar{X}_t, \bar{T}_{t-1}]$. Thus we conclude that the true parameters must be satisfying the moment restrictions:

$$\forall t \in [m] : \mathbb{E} \left[\left(\tilde{Y}_t - \sum_{j=t+1}^m \psi'_j \tilde{T}_{j,t} - \psi'_t \tilde{T}_{t,t} \right) \tilde{T}_{t,t} \right] = 0$$

These are exactly of the same form as the moment restrictions considered in the definition of the Dynamic DML algorithm and its analysis. Thus the results we presented so far directly extend to the estimation of structural nested mean models, with a linear parameterization of the blip functions, simply by using the different definition of the residual variables \tilde{Y}_t and $\tilde{T}_{t,t}$ and letting $\psi_t = \theta_{m-t}$ in the definition of Algorithm 1 and in Theorems 3 and 4. Moreover, note that the nuisance models that are trained in the first stage have a larger conditioning set which includes all past history of states and treatments \bar{X}_t, \bar{T}_{t-1} . If we made further restrictions that there was a "funnel state" S_t at each period that summarizes the history and such that any dependence of the future to the past is going through that funnel state, then conditioning only on that state would have been sufficient. This is what we essentially did in the linear Markovian model. Moreover, the linear Markovian model with a static policy τ , is a special case where the blip functions take the simple form: $\gamma_t(\bar{x}_t, \bar{\tau}_t) = \theta'_{m-t} \tau_t$ and are target policy independent.

Thus our Dynamic DML algorithm extends to the estimation of the structural parameters in a structural nested mean model for any target dynamic policy d and any user defined baseline policy $\bar{0}$ (referred to as a double regime structural nested mean model). It allows for the estimation of the nuisance functions with arbitrary machine learning algorithms, subject to a relatively slow mean-squared-error condition and reduces estimation to simple regression and classification oracles in the first stage, with only a simple linear system of equations in the second phase, which can also be solved in linear time in a recursive manner.

Dynamic effect heterogeneity Finally, we note that our moment condition that identifies each parameter ψ_t is the derivative of a square loss and can be written as the solution to the following square loss minimization problem: $\min_{\psi_t} \mathbb{E} \left[\left(\tilde{Y}_t - \sum_{j=t+1}^m \psi'_j \tilde{T}_{j,t} - \psi'_t \tilde{T}_{t,t} \right)^2 \right]$, fixing the solution ψ_j for any $j > t$, from previous iterations. This allows us to generalize the Dynamic DML to the case where we allow non-parametric heterogeneity in the parameters ψ_t , with respect to an exogenous fixed covariate vector of each sample, denoted as X_0 , i.e. $\gamma_t(\bar{x}_t, \bar{\tau}_t) = \psi_t(x_0)' \phi(\bar{x}_t, \bar{\tau}_t)$, for a known feature map ϕ and unknown heterogeneous parameters ψ . Thus we can essentially generalize the g -estimation approach to SNMMs to allow for infinite dimensional parameters of the blip functions, as long as the input to these infinite dimensional parameters is fixed and not changing endogenously by the treatments (e.g. fixed characteristics of a unit). This can be achieved by simply minimizing recursively the square loss:

$$\min_{\psi_t(\cdot)} \mathbb{E} \left[\left(\tilde{Y}_t - \sum_{j=t+1}^m \psi_j(X_0)' \tilde{T}_{j,t} - \psi_t(X_0)' \tilde{T}_{t,t} \right)^2 \right]$$

over arbitrary function spaces $\psi_t(\cdot)$. Using techniques from the recently introduced orthogonal statistical learning framework [12], we show in the appendix that this estimation method provides mean-squared-error guarantees on the recovered heterogeneous parameters ψ_t , that are robust to errors in the nuisance functions. This heterogeneous extension can also be viewed as an analogue of the R Learner meta-learner algorithm [24], generalized to the dynamic treatment regime setting.

8 Experimental Results

We consider data drawn from the DGP presented in Equation (1), with a linear observational policy: $p(T_{t-1}, X_t) := C T_{t-1} + D \cdot X_t$, with $X_0, T_0 = 0$ and $\epsilon_t, \zeta_t, \eta_t$ standard normal r.v.'s. (recall that d is the number of treatments and p the number of state variables). We consider the instance where: $A_{ij} = .5$, for all $i \in [p]$, $j \in [d]$, $B = .5 I_p$, $C = .2 I_d$, $D[:, 1 : 2] = .4$, $D[:, 3 : p] = 0$, $\mu[1 : 2] = .8$. We consider settings where the effect is constant, i.e. $\theta_0 \in \mathbb{R}^d$ or heterogeneous, where $\theta_0(X) = \theta_0 + \langle \beta_0, X[S] \rangle$ for some known low dimensional subset S of the states. We compare the dynamic DML to several benchmarks. The results are presented in Figures 1, comparing the estimates of the dynamic DML algorithm to a number of other alternatives on a single instance.

They fall into two categories. In the “static” set of approaches, each of the contemporaneous and lag effects is estimated one at a time, either by direct regression or (static) DML. So for example, to estimate the one period lag effect θ_1 , we would regress Y_t on T_{t-1} , with controls X . We consider direct regression with no controls (“no-ctrls”), direct and DML with controls from the initial period (i.e X_{t-1} , “init-ctrls” and “init-ctrls-dml”) and direct and DML with controls from the same period as the outcome (i.e X_t , “fin-ctrls” and “fin-ctrls-dml”). As an alternative to all of these, we try a “direct” dynamic approach, where initially we estimate θ_0 using a lasso regression of Y_t with all the controls and past treatments, and return the coefficient on T_t , and then “peel” off the estimated effect as in the main text before running another Lasso regression of $Y_t - \theta_0 T_t$ on T_{t-1} to get the first lag effect etc. So this approach incorporates the peeling effect but doesn’t do any orthogonalization. The point estimate for θ_0 , θ_1 and θ_2 are depicted in the three panels of Figure 1 and the error bars correspond to the constructed confidence interval. For all three, the dynamic DML is relatively close to the truth and the confidence interval contains the truth. The remaining approaches are not, although for the contemporaneous effect the approaches with final period controls have similar performance - it is really in the lagged effects that the differences become most apparent. Subsequently we run multiple experiments to evaluate the performance of DynamicDML. In each setting, we run 1000 Monte Carlo experiments, where each experiment draws $N = 500$ samples from the above DGP and then estimated the effects and lag effects based on our Dynamic DML algorithm. Figure 2 considers the case of two treatments - and shows that the algorithm performs well in terms of giving reasonable coverage guarantees - for a nominal 95% coverage, actual coverage varies from 91% to 94.5%. The right panel shows that the average estimates are close to the truth. In Figure 5 in Appendix E we repeat the experiments with $N = 2000$, and actual coverage is now tightly in the range 94% to 95%, and the average estimates remain relatively unbiased. Finally, in the appendix we explore a scenario with heterogeneous effects and find qualitatively similar performance.

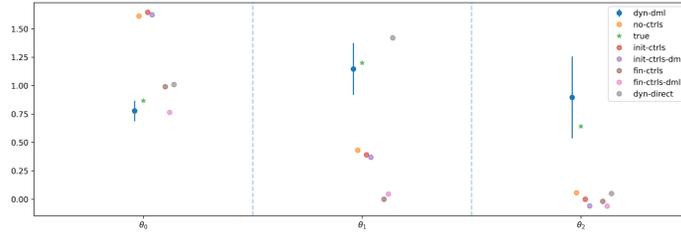


Figure 1: Comparison of DynamicDML (with confidence intervals) with benchmarks on a single instance. $n = 400$, $n_t = 2$, $n_x = 100$, $s = 10$, $\sigma(\epsilon_t) = .5$, $\sigma(\zeta_t) = .5$, $\sigma(\eta_t) = .5$, $C = 0$

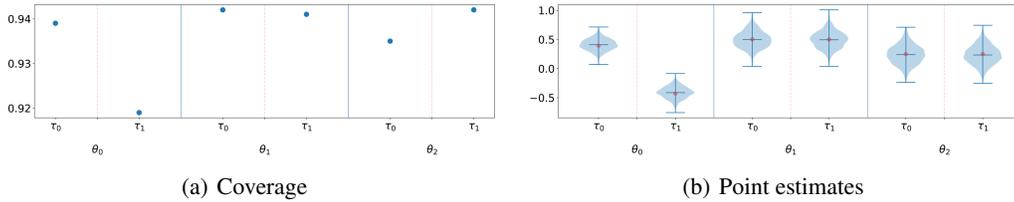


Figure 2: $n = 500$, $n_t = 2$, $n_x = 450$, $s = 2$, $\sigma(\epsilon_t) = 1$, $\sigma(\zeta_t) = .5$, $\sigma(\eta_t) = 1$

References

- [1] Chunrong Ai and Xiaohong Chen. Efficient estimation of models with conditional moment restrictions containing unknown functions. *Econometrica*, 71(6):1795–1843, 2003.
- [2] Heejung Bang and James M. Robins. Doubly robust estimation in missing data and causal inference models. *Biometrics*, 61(4):962–973, 2005.
- [3] Peter J Bickel, Chris AJ Klaassen, Peter J Bickel, Ya’acov Ritov, J Klaassen, Jon A Wellner, and YA’Acov Ritov. *Efficient and adaptive estimation for semiparametric models*, volume 4. Johns Hopkins University Press Baltimore, 1993.
- [4] Bibhas Chakraborty and Erica E. M. Moodie. *G-computation: Parametric Estimation of Optimal DTRs*, pages 101–112. Springer New York, New York, NY, 2013.
- [5] V. Chernozhukov, J. C. Escanciano, H. Ichimura, W. K. Newey, and J. M. Robins. Locally Robust Semiparametric Estimation. *ArXiv e-prints*, July 2016.
- [6] 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.
- [7] Victor Chernozhukov, Juan Carlos Escanciano, Hidehiko Ichimura, Whitney K. Newey, and James M. Robins. Locally Robust Semiparametric Estimation. *arXiv e-prints*, page arXiv:1608.00033, July 2016.
- [8] Victor Chernozhukov, Matt Goldman, Vira Semenova, and Matt Taddy. Orthogonal machine learning for demand estimation: High dimensional causal inference in dynamic panels. *arXiv preprint arXiv:1712.09988*, 2017.
- [9] Victor Chernozhukov, Denis Nekipelov, Vira Semenova, and Vasilis Syrgkanis. Plug-in regularized estimation of high-dimensional parameters in nonlinear semiparametric models. *arXiv preprint arXiv:1806.04823*, 2018.
- [10] Yash Deshpande, Adel Javanmard, and Mohammad Mehrabi. Online debiasing for adaptively collected high-dimensional data with applications to time series analysis. *arXiv preprint arXiv:1911.01040*, 2019.
- [11] Yash Deshpande, Lester Mackey, Vasilis Syrgkanis, and Matt Taddy. Accurate inference for adaptive linear models. In *International Conference on Machine Learning*, pages 1194–1203. PMLR, 2018.
- [12] Dylan J Foster and Vasilis Syrgkanis. Orthogonal statistical learning. *arXiv preprint arXiv:1901.09036*, 2019.
- [13] David A. Freedman. On tail probabilities for martingales. *Ann. Probab.*, 3(1):100–118, 02 1975.
- [14] Vitor Hadad, David A. Hirshberg, Ruohan Zhan, Stefan Wager, and Susan Athey. Confidence intervals for policy evaluation in adaptive experiments. *Proceedings of the National Academy of Sciences*, 118(15), 2021.
- [15] P. Hall and C.C. Heyde. 3 - the central limit theorem. In P. Hall and C.C. Heyde, editors, *Martingale Limit Theory and its Application*, Probability and Mathematical Statistics: A Series of Monographs and Textbooks, pages 51 – 96. Academic Press, 1980.
- [16] Miguel A Hernán and James M Robins. Causal inference: What if, 2010.
- [17] Nathan Kallus and Masatoshi Uehara. Double reinforcement learning for efficient off-policy evaluation in markov decision processes, 2019.
- [18] Nathan Kallus and Masatoshi Uehara. Efficiently breaking the curse of horizon in off-policy evaluation with double reinforcement learning, 2019.

- [19] Huitan Lei, Inbal Nahum-Shani, K Lynch, David Oslin, and Susan A Murphy. A " smart" design for building individualized treatment sequences. *Annual review of clinical psychology*, 8:21–48, 2012.
- [20] Judith J Lok and Victor DeGruttola. Impact of time to start treatment following infection with application to initiating haart in hiv-positive patients. *Biometrics*, 68(3):745–754, 2012.
- [21] Whitney K Newey. Semiparametric efficiency bounds. *Journal of applied econometrics*, 5(2):99–135, 1990.
- [22] Jerzy Neyman. $c(\alpha)$ tests and their use. *Sankhya*, pages 1–21, 1979.
- [23] Xinkun Nie, Emma Brunskill, and Stefan Wager. Learning when-to-treat policies, 2019.
- [24] Xinkun Nie and Stefan Wager. Quasi-oracle estimation of heterogeneous treatment effects. *arXiv preprint arXiv:1712.04912*, 2017.
- [25] Thomas Peel, Sandrine Anthoine, and Liva Ralaivola. Empirical bernstein inequality for martingales: Application to online learning. 2013.
- [26] Maya Petersen, Joshua Schwab, Susan Gruber, Nello Blaser, Michael Schomaker, and Mark van der Laan. Targeted maximum likelihood estimation for dynamic and static longitudinal marginal structural working models. *Journal of causal inference*, 2(2):147–185, 2014.
- [27] James 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.
- [28] 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.
- [29] James M. Robins. *Optimal Structural Nested Models for Optimal Sequential Decisions*, pages 189–326. Springer New York, New York, NY, 2004.
- [30] James M Robins, Donald Blevins, Grant Ritter, and Michael Wulfsohn. G-estimation of the effect of prophylaxis therapy for pneumocystis carinii pneumonia on the survival of aids patients. *Epidemiology*, pages 319–336, 1992.
- [31] James M Robins, Miguel Angel Hernan, and Babette Brumback. Marginal structural models and causal inference in epidemiology, 2000.
- [32] James M Robins and Ya’acov Ritov. Toward a curse of dimensionality appropriate (coda) asymptotic theory for semi-parametric models. *Statistics in medicine*, 16(3):285–319, 1997.
- [33] Peter M Robinson. Root-n-consistent semiparametric regression. *Econometrica: Journal of the Econometric Society*, pages 931–954, 1988.
- [34] Richard S. Sutton and Andrew G. Barto. *Introduction to Reinforcement Learning*. MIT Press, Cambridge, MA, USA, 1st edition, 1998.
- [35] Philip S. Thomas and Emma Brunskill. 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*, pages 2139–2148. JMLR.org, 2016.
- [36] Anastasios Tsiatis. *Semiparametric theory and missing data*. Springer Science & Business Media, 2007.
- [37] Mark J Van der Laan, MJ Laan, and James M Robins. *Unified methods for censored longitudinal data and causality*. Springer Science & Business Media, 2003.
- [38] Stijn Vansteelandt, Marshall Joffe, et al. Structural nested models and g-estimation: the partially realized promise. *Statistical Science*, 29(4):707–731, 2014.
- [39] Stijn Vansteelandt and Arvid Sjolander. Revisiting g-estimation of the effect of a time-varying exposure subject to time-varying confounding. *Epidemiologic Methods*, 5(1):37 – 56, 2016.

- [40] JM Wainwright. High-dimensional statistics: A non-asymptotic viewpoint. *preparation*. University of California, Berkeley, 2015.
- [41] Martin J. Wainwright. *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2019.
- [42] Kelly W. Zhang, Lucas Janson, and Susan A. Murphy. Inference for batched bandits. *CoRR*, abs/2002.03217, 2020.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
 - (b) Did you describe the limitations of your work? [Yes]
 - (c) Did you discuss any potential negative societal impacts of your work? [Yes]
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [Yes]
 - (b) Did you include complete proofs of all theoretical results? [Yes]
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes]
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [N/A]
 - (b) Did you mention the license of the assets? [N/A]
 - (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

Contents

A	Technical Exposition of Extension to Structural Nested Models (Section 7)	15
A.1	Heterogeneous Dynamic Effects	19
A.2	High Dimensional Sparse Linear Blip Functions	21
A.3	Proof of Lemma 6	22
A.4	Proof of Lemma 7	23
A.5	Proof of Lemma 9	23
A.6	Proof of Theorem 10	26
A.7	Proof of Theorem 11	28
B	Proofs of Results in Main Text	31
B.1	Proof of Lemma 1	31
B.2	Proof of Theorem 2	32
B.3	Proof of Theorem 3	32
B.4	Proof of Theorem 4	34
B.5	Proof of Theorem 5	35
C	Omitted Content from Section 5	36
D	Omitted Content from Section 6	36
E	Further Experimental Results	39
E.1	Constant Treatment Effects	39
E.2	Heterogeneous Treatment Effects	40
E.3	Counterfactual Policy Values	41
A	Technical Exposition of Extension to Structural Nested Models (Section 7)	

We present a more formal treatment of the extension of our main algorithm to g -estimation of structural nested models in biostatistics [27]. Consider an arbitrary time-series process $\{X_t, T_t\}_{t=1}^m$, with $X_t \in \mathcal{X}_t$ and $T_t \in \mathcal{T}_t$. Let Y denote some final outcome of interest. For any time t , let $\bar{X}_t = \{X_1, \dots, X_t\}$ and $\bar{T}_t = \{T_1, \dots, T_t\}$, denote the sequence of the variables up until time t and similarly, let $\underline{X}_t = \{X_t, \dots, X_m\}$ and $\underline{T}_t = \{T_t, \dots, T_m\}$. We will also denote with $\bar{x}_t, \bar{\tau}_t, \underline{x}_t, \underline{\tau}_t$, corresponding realizations of the latter random sequences. Let $d = (d_1, \dots, d_m)$ denote any dynamic policy, such that for each t , d_t maps a history $\bar{x}_t, \bar{\tau}_{t-1}$ into a next period action τ_t . For any such dynamic policy, let $Y^{(d)}$ denote the counterfactual outcome under policy d . For any static policy $\tau \in \times_{t=1}^m \mathcal{T}_t$, we will overload notation and let $Y^{(\tau)}$ denote the counterfactual outcome under this static treatment policy. Moreover, for any two policies (static or dynamic) we will be denoting with $(\bar{d}'_t, \underline{d}_{t+1})$, the policy that follows d' up until time t and then continues with policy d . We let $0 \in \mathcal{T}_t$ denote a baseline policy value, which could be appropriately instantiated based on the context.

We assume that the data generating process satisfies the following sequential conditional randomization condition:

ASSUMPTION 1 (*Sequential Conditional Exogeneity*). *The data generating process satisfies the following conditional independence conditions:*

$$\forall t \in [m] : \{Y^{(\tau)}, \tau \in \times_{t=1}^m \mathcal{T}_t\} \perp\!\!\!\perp T_t \mid \bar{T}_{t-1}, \bar{X}_t$$

Identification of mean counterfactual outcomes $\mathbb{E}[Y^{(d)}]$ for a target policy of interest d can be expressed in terms of the following conditional expectation functions:

$$\gamma_t(\bar{x}_t, \bar{\tau}_t) = \mathbb{E} \left[Y^{(\bar{\tau}_t, d_{t+1})} - Y^{(\bar{\tau}_{t-1}, 0, d_{t+1})} \mid \bar{T}_t = \bar{\tau}_t, \bar{X}_t = \bar{x}_t \right]$$

which corresponds to the mean change in outcome if we go to all units which received treatment $\bar{\tau}_t$ up until time t and had observed state history \bar{x}_t and we remove their last treatment, while we subsequently always continue with the target policy d . These functions are known as the *blip* functions [4, 29] and can be shown to be non-parametrically identifiable, assuming sequential conditional exogeneity and a sequential analogue of the positivity (aka overlap) assumption [29].

Theorem 3.1 of [29] shows that via a telescoping sum argument and invoking the sequential randomization condition the following lemma, which we adapt here to our notation and provide a proof for completeness:

Lemma 6. *For any dynamic policy d and under the sequential conditional exogeneity assumption, the following identity holds about the counterfactual outcomes:*

$$\mathbb{E}[Y^{(\bar{\tau}_{t-1}, d_t)} \mid \bar{X}_t, \bar{T}_t = \bar{\tau}_t] = \mathbb{E} \left[Y + \sum_{j=t}^m \rho_j(\bar{X}_j, \bar{T}_j) \mid \bar{X}_t, \bar{T}_t = \bar{\tau}_t \right] \quad (2)$$

where:

$$\rho_j(\bar{X}_j, \bar{T}_j) := \gamma_j(\bar{X}_j, (\bar{T}_{j-1}, d(\bar{X}_j, \bar{T}_{j-1}))) - \gamma_j(\bar{X}_j, \bar{T}_j).$$

Hence also:

$$\mathbb{E}[Y^{(d)}] = \mathbb{E} \left[Y + \sum_{t=1}^m \rho_t(\bar{X}_t, \bar{T}_t) \right]$$

Importantly, the conditioning set in Equation (2) contains the observed t periods treatment. Intuitively, each term ρ_j , removes from the outcome the *blip effect* of the observed action T_j and adds the *blip effect* of the target action $d(\bar{X}_j, \bar{T}_{j-1})$. Lemma 6, together with conditional sequential exogeneity also implies that the following set of moment restrictions must be satisfied (the following is an adaptation of Theorem 3.2 of [29] to our notation and we include its proof for completeness).

Lemma 7. *For any parameterization of the blip functions $\gamma_t(\bar{x}_t, \bar{\tau}_t; \psi_t)$, if we let the random variable $H_t(\psi) := Y + \sum_{j=t}^m \rho_j(\bar{X}_j, \bar{T}_j; \psi_j)$, then the true parameter vector ψ^* must satisfy the moment restrictions:*

$$\forall t \in [m], \forall f \in \mathcal{F} : \mathbb{E} [H_t(\psi^*) (f(\bar{X}_t, \bar{T}_t) - \mathbb{E}[f(\bar{X}_t, \bar{T}_t) \mid \bar{X}_t, \bar{T}_{t-1}])] = 0$$

where \mathcal{F} contains all functions mapping histories $\bar{x}_t, \bar{\tau}_t$ to \mathbb{R} .

To achieve parametric estimation rates for the quantities of interest, we will need to further make a semi-parametric assumption, i.e. that the *blip* functions take a low-dimensional parametric form:

ASSUMPTION 2 (Linear Blip Functions). *The blip functions admit a linear parametric form:*

$$\gamma_t(\bar{x}_t, \bar{\tau}_t; \psi_t) := \psi_t' \phi_t(\bar{x}_t, \bar{\tau}_t)$$

for some known r -dimensional feature vector maps ϕ_t , satisfying $\phi_t(\bar{x}_t, (\bar{\tau}_{t-1}, 0)) = 0$ and such that for some true ψ_t^* , $\gamma_t(\cdot, \cdot; \psi_t^*) = \gamma_t(\cdot, \cdot)$.

Then we can identify ψ^* by finding a parameter vector ψ that satisfies the subset of the moment restrictions of the form:

$$\forall t \in [m] : \mathbb{E}[H_t(\psi) (\phi_t(\bar{X}_t, \bar{T}_t) - \mathbb{E}[\phi_t(\bar{X}_t, \bar{T}_t) \mid \bar{X}_t, \bar{T}_{t-1}])] = 0$$

Moreover, we can also subtract from $H_t(\psi)$, the conditional expectation $\mathbb{E}[H_t(\psi^*) \mid \bar{X}_t, \bar{T}_{t-1}]$ while maintaining the moment condition:

$$\forall t \in [m] : \mathbb{E}[(H_t(\psi) - \mathbb{E}[H_t(\psi^*) \mid \bar{X}_t, \bar{T}_{t-1}]) (\phi_t(\bar{X}_t, \bar{T}_t) - \mathbb{E}[\phi_t(\bar{X}_t, \bar{T}_t) \mid \bar{X}_t, \bar{T}_{t-1}])] = 0$$

This is the doubly robust moment condition proposed by [28, 29], where it is shown that an estimator of ψ based on this moment is correct if either the estimate of $\mathbb{E}[H_t(\psi^*) \mid \bar{X}_t, \bar{T}_{t-1}]$ or the estimate of $\mathbb{E}[\phi_t(\bar{X}_t, \bar{T}_t) \mid \bar{X}_t, \bar{T}_{t-1}]$ is correct. However, estimating $\mathbb{E}[H_t(\psi^*) \mid \bar{X}_t, \bar{T}_{t-1}]$, itself requires

seemingly knowledge of ψ^* , which is cumbersome and can be impractical (e.g. by constructing poor preliminary estimates of ψ^*). This issue has also been discussed as one of the main points of not using the doubly robust correction in practice in g -estimation of structural nested models (see e.g. the discussion at the end of Section 6.1 of [38]). For a binary treatment and when the target policy is the all-zero policy [16] (see Technical Point 21.5) note that $\mathbb{E}[H_t(\psi^*) | \bar{X}_t, \bar{T}_t]$ can be estimated by regressing the outcome of the population that received zero subsequent treatment on the history. However, such a population can be quite small in practice and can have severe co-variate imbalances compared to the overall population. Moreover, this approach only applies to the case of a binary treatment and a static target policy.

Our approach In this work, we show that we can achieve a Neyman orthogonal moment for identifying ψ^* , which is sufficient for robustness to biases stemming from machine learning models used to train the nuisance components, while avoiding the cumbersome part of estimating the nuisance $\mathbb{E}[H(\psi^*) | \bar{X}_t, \bar{T}_{t-1}]$. Moreover, our approach leads to a strongly convex loss for the parameters at each step of the recursive process, which is beneficial for finite sample guarantees and subsequently for generalizing it to linear models with parameter heterogeneity with respect to exogenous co-variates.

In particular, instead of subtracting $\mathbb{E}[H_t(\psi^*) | \bar{X}_t, \bar{T}_{t-1}]$, we subtract $\mathbb{E}[H_t(\psi) | \bar{X}_t, \bar{T}_{t-1}]$, i.e. the function that we subtract is dynamically dependent on the estimate of ψ . Even though this moment is not doubly robust, we will show that it remains *locally robust* [5], aka *Neyman orthogonal* [6]. To define our moment restrictions and connect them to the main results of the paper we first define several convenient random variables: for any $j \in [m]$, let

$$Q_j := \phi_j(\bar{X}_j, \bar{T}_j) - \phi_j(\bar{X}_j, (\bar{T}_{j-1}, d(\bar{X}_j, \bar{T}_{j-1})))$$

and for any $1 \leq t \leq j$, let

$$\begin{aligned} \tilde{T}_{j,t} &:= Q_j - \mathbb{E}[Q_j | \bar{X}_t, \bar{T}_{t-1}] \\ \tilde{Y}_t &:= Y - \mathbb{E}[Y | \bar{X}_t, \bar{T}_{t-1}] \end{aligned}$$

Then, note that for any ψ :

$$H_t(\psi) - \mathbb{E}[H_t(\psi) | \bar{X}_t, \bar{T}_{t-1}] = \tilde{Y}_t - \sum_{j=t}^m \psi'_j \tilde{T}_{j,t}$$

Moreover, note that since the second term in Q_t only depends on the conditioning set \bar{X}_t, \bar{T}_{t-1} , then:

$$\tilde{T}_{t,t} = \phi_t(\bar{X}_t, \bar{T}_t) - \mathbb{E}[\phi_t(\bar{X}_t, \bar{T}_t) | \bar{X}_t, \bar{T}_{t-1}].$$

Thus we conclude that the true parameter ψ^* must be satisfying the moment restrictions:

$$\forall t \in [m] : \mathbb{E} \left[\left(\tilde{Y}_t - \sum_{j=t+1}^m \psi'_j \tilde{T}_{j,t} - \psi'_t \tilde{T}_{t,t} \right) \tilde{T}_{t,t} \right] = 0$$

These are exactly of the same form as the moment restrictions considered in the definition of the Dynamic DML algorithm and its analysis. Thus the results we presented so far directly extend to the estimation of structural nested mean models, with a linear parameterization of the blip functions, simply by using the different definition of the residual variables \tilde{Y}_t and $\tilde{T}_{t,t}$ and letting $\psi_t = \theta_{m-t}$ in the definition of Algorithm 1 and in Theorems 3 and 4.

Moreover, note that the nuisance models that are trained in the first stage have a larger conditioning set which includes all past history of states and treatments \bar{X}_t, \bar{T}_{t-1} . If we made further restrictions that there was a "funnel state" S_t at each period that summarizes the history and such that any dependence of the future to the past is going through that funnel state, then conditioning only on that state would have been sufficient. This is what we essentially did in the linear Markovian model. Moreover, the linear Markovian model with a static policy τ , is a special case where the blip functions take the simple form: $\gamma_t(\bar{x}_t, \bar{\tau}_t) = \theta'_{m-t} \tau_t$ and are target policy independent. When the blip functions are target policy independent, then any target policy can be used to estimate the structural parameters ψ^* . In our main development, we essentially used the baseline zero policy as a target policy to estimate the structural parameters.

Thus our Dynamic DML algorithm extends to the estimation of the structural parameters in a structural nested mean model for any target dynamic policy d and any user defined baseline policy $\bar{0}$ (referred to as a $(\pi, \bar{0})$ -double regime structural nested mean model). It allows for the estimation of the nuisance functions with arbitrary machine learning algorithms, subject to a relatively slow mean-squared-error condition and reduces estimation to simple regression and classification oracles in the first stage, with only a simple linear system of equations in the second phase, which can also be solved in linear time in a recursive manner.

For completeness, we present the generalization of Algorithm 1 to SNMMs in Algorithm 2 and we re-state the main Theorems in this broader context.

Algorithm 2 Dynamic DML for Structural Nested Mean Models (SNMMs)

Input: A data set consisting of n samples of m -length paths: $\{(X_1^i, T_1^i, \dots, X_m^i, T_m^i, Y^i)\}_{i=1}^n$

Input: A target dynamic policy d

Input: r -dimensional feature vector maps $\{\phi_t\}_{t=1}^m$, which parameterize the blip functions

Randomly split the n samples in S, S'

for each $t \in \{1, \dots, m\}$ **do**

Regress Y on \bar{X}_t, \bar{T}_{t-1} using S to learn estimate \hat{q}_t of model

$$q_t(\bar{x}_t, \bar{\tau}_{t-1}) := \mathbb{E}[Y \mid \bar{X}_t = \bar{x}_t, \bar{T}_{t-1} = \bar{\tau}_{t-1}]$$

Calculate residuals for each sample $i \in S'$: $\tilde{Y}_t^i := Y^i - \hat{q}_t(\bar{X}_t^i, \bar{T}_{t-1}^i)$

Vice versa use S' to learn model \hat{q}_t and calculate residuals on S .

for each $j \in \{t, \dots, m\}$ **do**

Let $Q_j := \phi_j(\bar{X}_j, \bar{T}_j) - \phi_j(\bar{X}_j, (\bar{T}_{j-1}, d(\bar{X}_j, \bar{T}_{j-1}))) 1\{j > t\}$

Regress Q_j on \bar{X}_t, \bar{T}_{t-1} using S to learn estimate $\hat{p}_{j,t}$ of model

$$p_{j,t}(\bar{x}_t, \bar{\tau}_{t-1}) := \mathbb{E}[Q_j \mid \bar{X}_t = \bar{x}_t, \bar{T}_{t-1} = \bar{\tau}_{t-1}]$$

Calculate residuals for each sample $i \in S'$: $\tilde{T}_{j,t}^i := Q_j^i - p_{j,t}(\bar{X}_t^i, \bar{T}_{t-1}^i)$

Vice versa use S' to learn model $\hat{p}_{j,t}$ and calculate residuals on S .

end for

end for

Using all the data $S \cup S'$

for $t = m$ **down to** 1 **do**

Regress $\tilde{Y}_t - \sum_{j=t+1}^m \hat{\psi}_j' \tilde{T}_{j,t}$ on $\tilde{T}_{t,t}$ with ordinary least squares:

$$\hat{\psi}_t := \arg \min_{\psi_t \in \Psi_t} \frac{1}{n} \sum_{i=1}^n \left(\tilde{Y}_t^i - \sum_{j=t+1}^m \hat{\psi}_j' \tilde{T}_{j,t}^i - \psi_t' \tilde{T}_{t,t}^i \right)^2$$

end for

Return: Structural parameter estimate $\hat{\psi} = (\hat{\psi}_1, \dots, \hat{\psi}_m)$ of $\psi^* = (\psi_1^*, \dots, \psi_m^*)$

We present here the generalized version of the asymptotic normality theorem for SNMMs and we defer the finite sample guarantee to the next section where we will analyze a more general setting of heterogeneous dynamic effects and present a finite sample bound. The proof is identical to that of Theorem 4 and so we omit it.

Theorem 8 (Asymptotic Normality). *Consider a Structural Nested Mean Model and a target dynamic policy d , such that Assumptions 1 and Assumption 2 are satisfied, for a constant feature map dimension r . Suppose that the nuisance models $\{q_\kappa, p_{\tau, \kappa}\}_{\tau, \kappa}$ satisfy that:*

$$\forall 1 \leq t \leq j \leq m : \|\hat{q}_t - q_t\|_2, \|\hat{p}_{j,t} - p_{j,t}\|_2 \leq o_p(n^{-1/4})$$

Moreover, suppose that all random variables are bounded and that

$$\mathbb{E}[Cov(Q_t, Q_t \mid \bar{X}_t, \bar{T}_{t-1})] \succeq 2\lambda I \quad (\text{average positivity})$$

for $\lambda > 0$. Then the estimate $\hat{\psi}$ produced by Algorithm 2 satisfies w.p. $1 - 2\delta$:

$$\sqrt{n}(\hat{\psi} - \psi^*) \rightarrow N(0, V)$$

where $V = M^{-1}\Sigma(M^{-1})'$ and M is a $m \times m$ block upper triangular matrix, consisting of blocks of size $r \times r$, whose block entries are of the form:

$$\forall t \leq j : M_{t,j} = Cov(Q_j, Q_t | \bar{X}_t, \bar{T}_{t-1})$$

and Σ is a $m \times m$ block diagonal matrix, consisting of blocks of size $r \times r$, whose diagonal block entries take the form:

$$\Sigma_{tt} = \mathbb{E}[\epsilon_t^2 Cov(Q_t, Q_t | \bar{X}_t, \bar{T}_{t-1})].$$

where $\epsilon_t = Y - \mathbb{E}[Y | \bar{X}_t, \bar{T}_{t-1}] - \sum_{j=t}^m (\psi_j^*)' (Q_j - \mathbb{E}[Q_j | \bar{X}_t, \bar{T}_{t-1}])$.

A.1 Heterogeneous Dynamic Effects

We note that our moment condition that identifies each parameter ψ_t is the derivative of a square loss and can be written as the solution to the following square loss minimization problem:

$$\min_{\psi_t} \mathbb{E} \left[\left(\tilde{Y}_t - \sum_{j=t+1}^m \psi_j' \tilde{T}_{j,t} - \psi_t' \tilde{T}_{t,t} \right)^2 \right], \text{ fixing the solution } \psi_j \text{ for any } j > t, \text{ from previous iterations.}$$

This allows us to generalize the Dynamic DML to the case where we allow non-parametric heterogeneity in the parameters ψ_t , with respect to an exogenous fixed covariate vector of each sample, denoted as X_0 , i.e. $\gamma_t(\bar{x}_t, \bar{\tau}_t) = \psi_t(x_0)' \phi(\bar{x}_t, \bar{\tau}_t)$, for a known feature map ϕ and unknown heterogeneous parameters ψ . Thus we can essentially generalize the g -estimation approach to SNMMs to allow for infinite or high dimensional parameters of the blip functions, as long as the input to these infinite dimensional parameters is fixed and not changing endogenously by the treatments (e.g. fixed characteristics of a unit). This can be achieved by simply minimizing recursively the square loss:

$$\hat{\psi}_t = \arg \min_{\psi_t(\cdot) \in \Psi_t} \mathcal{L}_{D,t}(\psi_t; \hat{\psi}_{t+1}, \hat{h}) := \mathbb{E} \left[\left(\tilde{Y}_t - \sum_{j=t+1}^m \hat{\psi}_j(X_0)' \tilde{T}_{j,t} - \psi_t(X_0)' \tilde{T}_{t,t} \right)^2 \right]$$

over arbitrary function spaces Ψ_t or by using any other machine learning techniques that achieve small excess risk with respect to the latter square loss problem (e.g. regularized least squares, early stopping, etc). In the latter, we denoted with \hat{h} an estimate of the vector of all nuisance functions (e.g. estimated in the first stage of the Dynamic DML Algorithm) and $\hat{\psi}_{t+1} = (\hat{\psi}_{t+1}, \dots, \hat{\psi}_m)$ the estimates of the target structural parameters constructed in previous iterations of the recursion.

Algorithm 3 Dynamic RLearner for Structural Nested Mean Models (SNMMs)

Input: A data set of n samples of m -length paths: $\{(X_0^i, X_1^i, T_1^i, \dots, X_m^i, T_m^i, Y^i)\}_{i=1}^n$

Input: A target dynamic policy d

Input: r -dimensional feature vector maps $\{\phi_t\}_{t=1}^m$, which parameterize the blip functions

Input: Function spaces for the heterogeneous structural parameters $\{\Psi_t\}_{t=1}^m$

Execute first stage identically to Algorithm 2

Using all the data $S \cup S'$

for $t = m$ **down to** 1 **do**

Estimate heterogeneous structural function $\hat{\psi}_t$ by minimizing the following square loss:

$$\hat{\psi}_t := \arg \min_{\psi_t \in \Psi_t} \frac{1}{n} \sum_{i=1}^n \left(\tilde{Y}_t^i - \sum_{j=t+1}^m \hat{\psi}_j(X_0^i)' \tilde{T}_{j,t}^i - \psi_t(X_0^i)' \tilde{T}_{t,t}^i \right)^2$$

end for

Return: Heterogeneous structural function estimates $\hat{\psi} = (\hat{\psi}_1, \dots, \hat{\psi}_m)$ of $\psi^* = (\psi_1^*, \dots, \psi_m^*)$

Using techniques from the recently introduced orthogonal statistical learning framework [12, 9], we show in the appendix that this estimation method provides mean-squared-error guarantees on the recovered heterogeneous parameters $\hat{\psi}_t$, that are robust to errors in the nuisance functions. This heterogeneous extension can also be viewed as an analogue of the RLearner meta-learner algorithm [24], generalized to the dynamic treatment regime setting.

To state our main results we will introduce some norm notation. For any vector valued function ψ , taking as input a random variable X and having output in \mathbb{R}^r , we will denote with:

$$\|\psi\|_{u,v} := \mathbb{E}[\|\psi(X)\|_u^v]^{1/v} = \mathbb{E} \left[\left(\sum_{j=1}^r \psi_j(X)^u \right)^{u/v} \right]^{1/v}$$

for any $u, v > 1$. If ψ is a parameter vector, then we will overload notation and let $\|\psi\|_{u,v} = \|\psi\|_u = \left(\sum_{j=1}^r \psi_j^u \right)^{1/u}$. If u or v equals ∞ , then this would designate the sup norm, e.g. $\|\psi\|_{\infty,v} := \mathbb{E}[\max_{j \in r} \psi_j(X)^v]^{1/v}$ and $\|\psi\|_{u,\infty} = \sup_{x \in \mathcal{X}} \|\psi(x)\|_u$. For any u, v , we will denote with \bar{u}, \bar{v} the parameters that correspond to the dual norm, i.e. $1/u + 1/\bar{u} = 1$ and similarly for \bar{v} . For any two functions f, g , taking as input random variables X, Y we will use the shorthand notation: $\|f \circ g\|_{u,v} = \mathbb{E}[\|f(X)\|_u^v \cdot \|g(Y)\|_v^u]^{1/v}$. For an $n \times m$ matrix A , we will denote with:

$$\|A\|_{u,v} := \left(\sum_{i=1}^n \|A_{\cdot,i}\|_v^u \right)^{1/u} = \left(\sum_{i=1}^n \left(\sum_{j=1}^m |A_{i,j}|^{v/u} \right)^u \right)^{1/u}$$

Lemma 9 (Algorithm-agnostic analysis). *Consider any estimation algorithm that produces a vector of estimates $\hat{\psi} = (\psi_1, \dots, \psi_m)$ with small plug-in excess risks, i.e.,*

$$\mathcal{L}_{D,t}(\hat{\psi}_t; \hat{\psi}_{t+1}, \hat{h}) - \mathcal{L}_{D,t}(\psi_t^*; \hat{\psi}_{t+1}, \hat{h}) \leq \epsilon(\hat{\psi}_t, \hat{h}).$$

Let $C_{t,j} := \text{Cov}(Q_t, Q_j \mid \bar{X}_t, \bar{T}_{t-1})$ and suppose that for all $t \in [m]$ and for all $x_0 \in \mathcal{X}_0$:

$$\mathbb{E}[C_{t,t} \mid X_0 = x_0] \succeq \lambda I \quad (\text{positivity})$$

and, either for all $t \in [m]$, Ψ_t is a convex set or for all $t \in [m]$: $\psi_t^* \in \Psi_t$, then

$$\begin{aligned} \frac{\lambda}{2} \|\hat{\psi}_t - \psi_t^*\|_{2,2}^2 - \frac{\sigma}{4} \|\hat{\psi}_t - \psi_t^*\|_{u,v}^2 &\leq \epsilon(\hat{\psi}_t, \hat{h}) + \sum_{j=t+1}^m \frac{c_{t,j}}{\sigma} \|\hat{\psi}_j - \psi_j^*\|_{u,\bar{v}}^2 \\ &\quad + \frac{4M}{\sigma} \sum_{j=t}^m \|(\hat{p}_{t,t} - p_{t,t}) \circ (\hat{p}_{j,t} - p_{j,t})\|_{2,2}^2 \\ &\quad + \frac{1}{\sigma} \|(\hat{p}_{t,t} - p_{t,t}) \circ (\hat{q}_t - q_t)\|_{\bar{u},\bar{v}}^2 \end{aligned}$$

with $c_{t,j} := \sup_{x_0 \in \mathcal{X}_0} \mathbb{E}[\|C_{t,j}\|_{\bar{u},\bar{u}}^{\bar{v}} \mid X_0 = x_0]^{2/\bar{v}}$ and $M := \max_{t \in [m], \psi_t \in \Psi_t} \|\psi_t\|_{u,\infty}$.

Note that for the case when $u, v = 2$ and setting $\sigma = \lambda$, then we get the following recursive bound on the MSE of the heterogeneous structural parameters:

$$\begin{aligned} \frac{\lambda}{4} \|\hat{\psi}_t - \psi_t^*\|_{2,2}^2 &\leq \epsilon(\hat{\psi}_t, \hat{h}) + \sum_{j=t+1}^m \frac{c_{t,j}}{\lambda} \|\hat{\psi}_j - \psi_j^*\|_{2,2}^2 \\ &\quad + \frac{4M}{\lambda} \sum_{j=t}^m \|(\hat{p}_{t,t} - p_{t,t}) \circ (\hat{p}_{j,t} - p_{j,t})\|_{2,2}^2 + \frac{1}{\lambda} \|(\hat{p}_{t,t} - p_{t,t}) \circ (\hat{q}_t - q_t)\|_{2,2}^2 \end{aligned}$$

where $c_{t,j} := \sup_{x_0 \in \mathcal{X}_0} \mathbb{E}[\|C_{t,j}\|_F^2 \mid X_0 = x_0]$ and $M := \sup_{t \in [m], x_0 \in \mathcal{X}_0, \psi_t \in \Psi_t} \|\psi_t(x_0)\|_2$.

In the next theorem, we apply this conclusion to the case when the estimate $\hat{\psi}$ is produced via running empirical risk minimization in the second stage of the Heterogeneous Dynamic DML algorithm, as described in Algorithm 3. To state the theorem we will use the notion of the critical radius, which is a measure of statistical complexity of a function space. For any function space \mathcal{F} , with functions having range in $[-1, 1]$, we consider the localized Rademacher complexity as:

$$R_n(\mathcal{F}; \delta) = \mathbb{E}_{\epsilon_{1:n}, X_{1:n}} \left[\sup_{f \in \mathcal{F}: \|f\|_2 \leq \delta} \frac{1}{n} \sum_{i=1}^n \epsilon_i f(X_i) \right]$$

We denote as the critical radius $\delta_n > 0$ of \mathcal{F} , any solution to the inequality:

$$R_n(\mathcal{F}; \delta) \leq \delta^2$$

For each function space Ψ_t , we denote with $\Psi_{t,i}$ the marginal function space corresponding to the i -th coordinate output of the functions in Ψ_t . Moreover, we denote with $\Psi_{t,i} - \psi_{t,i}^* = \{\psi_{t,i} - \psi_{t,i}^* : \psi_{t,i} \in \Psi_{t,i}\}$. Finally, we define the star hull of a function space as: $\text{star}(\mathcal{F}) := \{\tau f : f \in \mathcal{F}, \tau \in [0, 1]\}$. The critical radius is a well-established concept in modern statistical learning theory and has been characterized for many function spaces. Moreover, for many function spaces it yields minimax optimal statistical learning rates. See [41] for an overview.

Theorem 10 (MSE Rate for Dynamic RLearner). *Suppose that all random variables and functions are bounded and let δ_n be an upper bound on the critical radius of the star hull of all the function spaces $\{\Psi_{t,i} - \psi_{t,i}^*\}_{t \in [m], i \in [r]}$ and that $\delta_n = \Omega\left(\sqrt{\frac{r \log \log(n)}{n}}\right)$. Suppose that the quantities $m, \lambda, c_{t,j}, M$ (as defined in Lemma 9, for $u = v = 2$) are constants independent of n and that:*

$$\begin{aligned} \max_{1 \leq t \leq j \leq m} \mathbb{E}_{\hat{p}_{t,t}, \hat{p}_{j,t}} [\|(\hat{p}_{t,t} - p_{t,t}) \circ (\hat{p}_{j,t} - p_{j,t})\|_{2,2}^2] &= O(r^2 \delta_n^2) \\ \max_{1 \leq t \leq m} \mathbb{E}_{\hat{p}_{t,t}, \hat{q}_t} [\|(\hat{p}_{t,t} - p_{t,t}) \circ (\hat{q}_t - q_t)\|_{2,2}^2] &= O(r^2 \delta_n^2) \end{aligned}$$

Then the output of Algorithm 3 satisfies:

$$\max_{t \in [m]} \mathbb{E}[\|\hat{\psi}_t - \psi_t^*\|_{2,2}^2] = O(r^2 \delta_n^2)$$

Analogous results hold with high probability and exponential tail, if we make such exponential tail assumptions also on the guarantees provided by the nuisance functions. We omit them for succinctness.

Note that most conditions on the nuisance functions have a doubly robust flavor, i.e. we need the product of two different nuisance function errors to be small. Thus we need that either one or the other is modeled and estimated accurately. The only exception is the functions $p_{t,t}$, which also need to satisfy that: $\|\hat{p}_{t,t} - p_{t,t}\|_4^4 = o_p(\delta_n/2)$. Thus one step ahead treatment propensivities, need to be more accurately estimated than the remainder of the nuisance functions.

A.2 High Dimensional Sparse Linear Blip Functions

The results we have discussed so far assume that the linear feature map that parameterizes the blip functions is low dimensional, i.e. $r \ll n$. Observe for instance that Theorem 10 is vacuous when $r = \Omega(n)$. In this section we examine the case where $r \gg n$ and provide guarantees under sparsity conditions on the true structural parameters. For simplicity, we will not consider non-parametric heterogeneity of the sparse coefficients with respect to some initial state X_0 , i.e. we consider blip functions of the form: $\gamma_t(\bar{X}_t, \bar{T}_t; \psi_t) = \psi_t' \phi(\bar{X}_t, \bar{T}_t)$, with $\psi_t \in \mathbb{R}^r$ and $r \gg n$. However, we note that here the high-dimensionality of the feature map already offers a lot of modelling flexibility and one could encode heterogeneity of the blip effect through the feature map.

Apart from the explicit dependence on r , when the feature map is high dimensional then the $\ell_{2,2}$ norm of the errors of the nuisance functions in Theorem 10 can accumulate across their r -dimensional components. Instead, we would ideally only require a bound on the maximum error across the r dimensions of each nuisance function. Then an exponential tail bound on the MSE of each coordinate would imply a bound on the maximum that scales only logarithmically with r . To achieve this we can invoke Lemma 9 with $u = 1, v = \infty$. Then we get the following recursive bound on the MSE of the heterogeneous structural parameters:

$$\begin{aligned} \frac{\lambda}{2} \|\hat{\psi}_t - \psi_t^*\|_2^2 - \frac{\sigma}{4} \|\hat{\psi}_t - \psi_t^*\|_1^2 &\leq \epsilon(\hat{\psi}_t, \hat{h}) + \sum_{j=t+1}^m \frac{c_{t,j}}{\sigma} \|\hat{\psi}_j - \psi_j^*\|_1^2 \\ &\quad + \frac{4M}{\sigma} \sum_{j=t}^m \|(\hat{p}_{t,t} - p_{t,t}) \circ (\hat{p}_{j,t} - p_{j,t})\|_{\infty,1}^2 \\ &\quad + \frac{1}{\sigma} \|(\hat{p}_{t,t} - p_{t,t}) \circ (\hat{q}_t - q_t)\|_{\infty,1}^2 \end{aligned}$$

where $c_{t,j} := \mathbb{E}[\|C_{t,j}\|_{\infty}]$ and $M := \sup_{t \in [m], \psi_t \in \Psi_t} \|\psi_t\|_1$. However, we see that we incur a dependency on the ℓ_1 norm of the error $\hat{\psi}_t - \psi_t^*$. Thus we need to be able to relate the ℓ_1 with the ℓ_2 norm of the error of our estimate. This is a restricted cone property on our estimate. We will thus invoke sparsity assumptions on the true parameter ψ_t^* and augment the empirical risk minimization step with an ℓ_1 penalty on ψ_t . This enforces the estimate to be primarily supported on the s relevant dimensions. Within such a restricted cone the ℓ_1 and the ℓ_2 norm are equivalent with a constant that depends only on the sparsity level and not the ambient dimension.

Furthermore, the explicit dependence in Theorem 10 also stems from the fact that we invoked a multi-dimensional contraction inequality, across the r output dimensions of ψ_t . However, observe that the loss function depends on each ψ_t only through single indices of the form $\psi_t' \tilde{T}_{j,t}$. Thus we could instead control the critical radius of these m index spaces at each iteration of the recursion, which would avoid the explicit dependence on r . Together these insights yield the following theorem.

Algorithm 4 Sparse Linear Dynamic DML for Structural Nested Mean Models (SNMMs)

Input: A data set of n samples of m -length paths: $\{(X_0^i, X_1^i, T_1^i, \dots, X_m^i, T_m^i, Y^i)\}_{i=1}^n$

Input: A target dynamic policy d

Input: r -dimensional feature vector maps $\{\phi_t\}_{t=1}^m$, which parameterize the blip functions

Execute first stage identically to Algorithm 2

Using all the data $S \cup S'$

for $t = m$ **down to** 1 **do**

Estimate heterogeneous structural function $\hat{\psi}_t$ by minimizing the ℓ_1 -penalized square loss:

$$\hat{\psi}_t := \arg \min_{\psi_t \in \mathbb{R}^r: \|\psi_t\|_1 \leq M} \frac{1}{n} \sum_{i=1}^n \left(\tilde{Y}_t - \sum_{j=t+1}^m \hat{\psi}_j' \tilde{T}_{j,t} - \psi_t' \tilde{T}_{t,t} \right)^2 + \kappa_t \|\psi_t\|_1$$

end for

Return: Structural function estimates $\hat{\psi} = (\hat{\psi}_1, \dots, \hat{\psi}_m)$ of $\psi^* = (\psi_1^*, \dots, \psi_m^*)$

Theorem 11 (ℓ_2 -Error Rate for Sparse Linear Dynamic DML). *Suppose that ψ_t^* have only s non-zero coefficients and that:*

$$\mathbb{E}[\text{Cov}(Q_t, Q_t \mid \bar{X}_t, \bar{T}_{t-1})] \succeq \lambda I \quad (\text{average positivity})$$

Let c_0, c_1, c_3, c_4 be sufficiently large universal constants. Then there exists a sequence of regularization levels (see proof for construction) $\kappa_1, \dots, \kappa_t$, such that for $n \geq c_2 \frac{s^2 \log(mr/\zeta)}{\lambda}$, the estimate output by Algorithm 4, satisfies w.p. $1 - c_0 m \zeta$:

$$\forall t \in [m]: \|\hat{\psi}_t - \psi_t^*\|_2 \leq \|\hat{\psi}_t - \psi_t^*\|_1 \leq c_1 \left(\frac{s}{\lambda} \left(\sqrt{\frac{\log(mr/\zeta)}{n}} + \epsilon_n \right) \frac{c_n^{m-t+1} - 1}{c_n - 1} \right)$$

where:

$$\epsilon_n := \max_{m \geq j \geq t \geq 1, i, i' \in [r]} \max \{ \|\hat{p}_{t,i} - p_{t,i}\|_2 \|\hat{q}_{t,i'} - q_{t,i'}\|_2, M m \|\hat{p}_{t,i} - p_{t,i}\|_2 \|\hat{p}_{j,t,i'} - p_{j,t,i'}\|_2 \}$$

$$c_n := c_4 \frac{s}{\lambda} \max_{t \in [m]} \sum_{j=t+1}^m \left(\|\mathbb{E}[\text{Cov}(Q_t, Q_j \mid \bar{X}_t, \bar{T}_{t-1})]\|_\infty + \sqrt{\frac{\log(mr/\zeta)}{n}} \right)$$

A.3 Proof of Lemma 6

Proof. First we note that by invoking the sequential conditional exogeneity condition, for any dynamic policy d and for any treatment sequence $\bar{\tau}_m$, we have:

$$\begin{aligned} \gamma(\bar{x}_j, (\bar{\tau}_{j-1}, d(\bar{x}_j, \bar{\tau}_{j-1}))) &= \mathbb{E}[Y^{(\bar{\tau}_{j-1}, d_j)} - Y^{(\bar{\tau}_{j-1}, 0, d_{j+1})} \mid \bar{X}_j = \bar{x}_j, \bar{T}_{j-1} = \bar{\tau}_{j-1}, T_j = d(\bar{X}_j, \bar{T}_{j-1})] \\ &= \mathbb{E}[Y^{(\bar{\tau}_{j-1}, d_j)} - Y^{(\bar{\tau}_{j-1}, 0, d_{j+1})} \mid \bar{X}_j = \bar{x}_j, \bar{T}_{j-1} = \bar{\tau}_{j-1}, T_j = \tau_j] \\ &= \mathbb{E}[Y^{(\bar{\tau}_{j-1}, d_j)} - Y^{(\bar{\tau}_{j-1}, 0, d_{j+1})} \mid \bar{X}_j = \bar{x}_j, \bar{T}_j = \bar{\tau}_j] \end{aligned}$$

Using the latter fact and invoking the definition of $\rho_j(\bar{x}_j, \bar{\tau}_j)$, we can write:

$$\begin{aligned} \rho_j(\bar{x}_j, \bar{\tau}_j) &= \gamma(\bar{x}_j, (\bar{\tau}_{j-1}, d(\bar{x}_j, \bar{\tau}_{j-1}))) - \gamma(\bar{x}_j, \bar{\tau}_j) \\ &= \mathbb{E} \left[Y^{(\bar{\tau}_{j-1}, d_j)} - Y^{(\bar{\tau}_{j-1}, 0, d_{j+1})} - \left(Y^{(\bar{\tau}_j, d_{j+1})} - Y^{(\bar{\tau}_{j-1}, 0, d_{j+1})} \right) \mid \bar{X}_j = \bar{x}_j, \bar{T}_j = \bar{\tau}_j \right] \\ &= - \mathbb{E} \left[Y^{(\bar{\tau}_j, d_{j+1})} - Y^{(\bar{\tau}_{j-1}, d_j)} \mid \bar{X}_j = \bar{x}_j, \bar{T}_j = \bar{\tau}_j \right] \end{aligned} \quad (3)$$

Second, note that by the definition of the observed Y , we have that:

$$Y \equiv Y^{(\bar{T}_m)}$$

and that for any treatment sequence $\bar{\tau}_m$, via a telescoping sum argument, we can write:

$$Y^{(\tau)} - Y^{(\bar{\tau}_{t-1}, d_t)} = \sum_{j=t}^m Y^{(\bar{\tau}_j, d_{j+1})} - Y^{(\bar{\tau}_{j-1}, d_j)}$$

Thus, applying linearity of expectation, the tower law of expectations and Equation (3), we have:

$$\begin{aligned}
\mathbb{E} \left[Y - Y^{(\bar{\tau}_{t-1}, \bar{d}_t)} \mid \bar{X}_t, \bar{T}_t = \bar{\tau}_t \right] &= \mathbb{E} \left[Y^{(T)} - Y^{(\bar{T}_{t-1}, \bar{d}_t)} \mid \bar{X}_t, \bar{T}_t = \bar{\tau}_t \right] \\
&= \mathbb{E} \left[\sum_{j=t}^m Y^{(\bar{T}_j, \bar{d}_{j+1})} - Y^{(\bar{T}_{j-1}, \bar{d}_j)} \mid \bar{X}_t, \bar{T}_t = \bar{\tau}_t \right] \\
&= \sum_{j=t}^m \mathbb{E} \left[Y^{(\bar{T}_j, \bar{d}_{j+1})} - Y^{(\bar{T}_{j-1}, \bar{d}_j)} \mid \bar{X}_t, \bar{T}_t = \bar{\tau}_t \right] \\
&= \sum_{j=t}^m \mathbb{E} \left[\mathbb{E} \left[Y^{(\bar{T}_j, \bar{d}_{j+1})} - Y^{(\bar{T}_{j-1}, \bar{d}_j)} \mid \bar{X}_j, \bar{T}_j \right] \mid \bar{X}_t, \bar{T}_t = \bar{\tau}_t \right] \\
&= \sum_{j=t}^m -\mathbb{E} \left[\rho(\bar{X}_j, \bar{T}_j) \mid \bar{X}_t, \bar{T}_t = \bar{\tau}_t \right] \\
&= -\mathbb{E} \left[\sum_{j=t}^m \rho(\bar{X}_j, \bar{T}_j) \mid \bar{X}_t, \bar{T}_t = \bar{\tau}_t \right]
\end{aligned}$$

By re-arranging we conclude the desired property. \square

A.4 Proof of Lemma 7

Proof. For simplicity of notation, for any $f \in \mathcal{F}$, let:

$$\bar{f}(\bar{X}_t, \bar{T}_t) = f(\bar{X}_t, \bar{T}_t) - \mathbb{E}[f(\bar{X}_t, \bar{T}_t) \mid \bar{X}_t, \bar{T}_{t-1}]$$

and observe that by the definition of \bar{f} , we crucially have that for any $f \in \mathcal{F}$:

$$\mathbb{E}[\bar{f}(\bar{X}_t, \bar{T}_t) \mid \bar{X}_t, \bar{T}_{t-1}] = 0.$$

By Lemma 6, we have that at the true ψ^* :

$$\begin{aligned}
\mathbb{E} \left[H_t(\psi^*) \bar{f}(\bar{X}_t, \bar{T}_t) \right] &= \mathbb{E} \left[\mathbb{E} \left[H_t(\psi^*) \mid \bar{X}_t, \bar{T}_t \right] \bar{f}(\bar{X}_t, \bar{T}_t) \right] && \text{(tower law)} \\
&= \mathbb{E} \left[\mathbb{E} \left[Y^{(\bar{T}_{t-1}, \bar{d}_t)} \mid \bar{X}_t, \bar{T}_t \right] \bar{f}(\bar{X}_t, \bar{T}_t) \right] && \text{(Equation (2))} \\
&= \mathbb{E} \left[\mathbb{E} \left[Y^{(\bar{T}_{t-1}, \bar{d}_t)} \mid \bar{X}_t, \bar{T}_{t-1} \right] \bar{f}(\bar{X}_t, \bar{T}_t) \right] && \text{(sequential exogeneity)} \\
&= \mathbb{E} \left[\mathbb{E} \left[Y^{(\bar{T}_{t-1}, \bar{d}_t)} \mid \bar{X}_t, \bar{T}_{t-1} \right] \mathbb{E} \left[\bar{f}(\bar{X}_t, \bar{T}_t) \mid \bar{X}_t, \bar{T}_{t-1} \right] \right] && \text{(tower law)} \\
&= 0
\end{aligned}$$

\square

A.5 Proof of Lemma 9

Lemma 12. For any random vectors I, X, Y and any two vector-valued functions f, g , we have that:

$$\mathbb{E}[(X - f(I))(Y - g(I))' \mid I] = \text{Cov}(X, Y \mid I) + \mathbb{E}[(\mathbb{E}[X \mid I] - f(I))(\mathbb{E}[Y \mid I] - g(I))' \mid I]$$

Proof. Let:

$$\begin{aligned}
A(f, g) &:= \mathbb{E}[(X - f(I))(Y - g(I))'] \\
B(f, g) &:= \mathbb{E}[(\mathbb{E}[X \mid I] - f(I))(\mathbb{E}[Y \mid I] - g(I))']
\end{aligned}$$

By simple numeric manipulations we have:

$$\begin{aligned}
A(f, g) &= \mathbb{E}[(X - \mathbb{E}[X \mid I])(Y - g(I))' \mid I] + \mathbb{E}[(\mathbb{E}[X \mid I] - f(I))(Y - g(I))' \mid I] \\
&= \mathbb{E}[(X - \mathbb{E}[X \mid I])(Y - g(I))' \mid I] + \mathbb{E}[(\mathbb{E}[X \mid I] - f(I))(Y - \mathbb{E}[Y \mid I])' \mid I] + B(f, g) \\
&= \mathbb{E}[(X - \mathbb{E}[X \mid I])(Y - g(I))' \mid I] + B(f, g) \\
&= \mathbb{E}[(X - \mathbb{E}[X \mid I])(Y - \mathbb{E}[Y \mid I])' \mid I] + \mathbb{E}[(X - \mathbb{E}[X \mid I])(\mathbb{E}[Y \mid I] - g(I))' \mid I] + B(f, g) \\
&= \mathbb{E}[(X - \mathbb{E}[X \mid I])(Y - \mathbb{E}[Y \mid I])' \mid I] + B(f, g) \\
&= \text{Cov}(X, Y \mid I) + B(f, g)
\end{aligned}$$

□

Proof of Lemma 9. First we define for any functional $L(f)$ the Frechet derivative as:

$$D_f L(f)[\nu] = \frac{\partial}{\partial t} L(f + t\nu) |_{t=0}$$

Similarly, we can define higher order derivatives, denoted as $D_{g,f} L(f, g)[\mu, \nu]$ and $D_{h,g,f} L(f, g, h)[\kappa, \mu, \nu]$.

For simplicity, in the remainder of the proof we will use the short-hand notation $\mathcal{L}_{D,t} \equiv \mathcal{L}_D$, since the time t will be evident by the inputs to the loss. Moreover, we'll use the shorthand notation $I_t := (\bar{X}_t, \bar{T}_{t-1})$. Moreover, we will be using the shorthand notation $\tilde{T}_t \equiv \tilde{T}_{t,t}$ and $p_{t,t} \equiv p_t$, $\hat{p}_{t,t} \equiv \hat{p}_t$.

First, we argue that the plug-in loss function is functionally λ -strongly convex. Note that for any $\psi_t \in \Psi_t$:

$$\begin{aligned} D_{\psi_t, \psi_t} \mathcal{L}_D(\psi_t; \underline{\psi}_{t+1}, \hat{h})[\hat{\psi}_t - \psi_t^*, \hat{\psi}_t - \psi_t^*] &= \mathbb{E} \left[\left(\hat{\psi}_t(X_0) - \psi_t^*(X_0) \right)' \tilde{T}_t \tilde{T}_t' \left(\hat{\psi}_t(X_0) - \psi_t^*(X_0) \right) \right] \\ &= \mathbb{E} \left[\left(\hat{\psi}_t(X_0) - \psi_t^*(X_0) \right)' \text{Cov}(Q_t, Q_t | I_t) \left(\hat{\psi}_t(X_0) - \psi_t^*(X_0) \right) \right] \\ &\quad + \mathbb{E} \left[\left(\left(\hat{\psi}_t(X_0) - \psi_t^*(X_0) \right)' (p_t(I_t) - \hat{p}_t(I_t)) \right)^2 \right] \\ &\geq \mathbb{E} \left[\left(\hat{\psi}_t(X_0) - \psi_t^*(X_0) \right)' \mathbb{E}[\text{Cov}(Q_t, Q_t | I_t) | X_0] \left(\hat{\psi}_t(X_0) - \psi_t^*(X_0) \right) \right] \\ &\geq \lambda \|\hat{\psi}_t - \psi_t^*\|_{2,2}^2 \end{aligned}$$

where in the second identity we invoked Lemma 12 and in the last inequality we used our positivity assumption. By the λ -strong convexity of \mathcal{L}_D , we have that

$$\mathcal{L}_D(\hat{\psi}; \underline{\psi}_{t+1}, \hat{h}) \geq \mathcal{L}_D(\psi_t^*; \underline{\psi}_{t+1}, \hat{h}) + D_{\psi_t} \mathcal{L}_D(\psi_t^*; \underline{\psi}_{t+1}, \hat{h})[\hat{\psi}_t - \psi_t^*] + \frac{\lambda}{2} \|\hat{\psi}_t - \psi_t^*\|_{2,2}^2.$$

Furthermore, our excess risk assumption and the optimality of f_0 give us

$$\begin{aligned} \frac{\lambda}{2} \|\hat{\psi}_t - \psi_t^*\|_{2,2}^2 &\leq \underbrace{\mathcal{L}_D(\hat{\psi}_t; \underline{\psi}_{t+1}, \hat{h}) - \mathcal{L}_D(\psi_t^*; \underline{\psi}_{t+1}, \hat{h})}_{\text{excess risk of } \hat{f}} - D_{\psi_t} \mathcal{L}_D(\psi_t^*; \underline{\psi}_{t+1}, \hat{h})[\hat{\psi}_t - \psi_t^*] \\ &\stackrel{(a)}{\leq} \underbrace{\epsilon(\hat{\psi}_t, \hat{h}) - D_{\psi_t} \mathcal{L}_D(\psi_t^*; \underline{\psi}_{t+1}, h_0)[\hat{\psi}_t - \psi_t^*]}_{\geq 0 \text{ by optimality of } \psi_t^*} \\ &\quad + \underbrace{D_{\psi_t} (\mathcal{L}_D(\psi_t^*; \underline{\psi}_{t+1}, h_0) - \mathcal{L}_D(\psi_t^*; \underline{\psi}_{t+1}, \hat{h}))[\hat{\psi}_t - \psi_t^*]}_{=: A} \\ &\quad + \underbrace{D_{\psi_t} (\mathcal{L}_D(\psi_t^*; \underline{\psi}_{t+1}, \hat{h}) - \mathcal{L}_D(\psi_t^*; \underline{\psi}_{t+1}, \hat{h}))[\hat{\psi}_t - \psi_t^*]}_{=: B}. \end{aligned}$$

Note that:

$$D_{\psi_t} \mathcal{L}_D(\psi_t; \underline{\psi}_{t+1}, h)[\hat{\psi}_t - \psi_t^*] = \mathbb{E} \left[\left(\psi_t^*(X_0) - \hat{\psi}_t(X_0) \right)' \tilde{T}_t \left(\tilde{Y}_t - \sum_{j=t+1}^m \psi_j(X_0)' \tilde{T}_{j,t} - \psi_t(X_0)' \tilde{T}_t \right) \right]$$

Then by repeated applying Holder's inequality for the dual norms associated with u, \bar{u} and v, \bar{v} , we have:

$$\begin{aligned}
|B| &= \sum_{j=t+1}^m \left| \mathbb{E} \left[\left(\psi_t^*(X_0) - \hat{\psi}_t(X_0) \right)' \tilde{T}_t \tilde{T}_{j,t}' \left(\psi_j^*(X_0) - \hat{\psi}_j(X_0) \right) \right] \right| \\
&= \sum_{j=t+1}^m \left| \mathbb{E} \left[\left(\psi_t^*(X_0) - \hat{\psi}_t(X_0) \right)' \text{Cov}(Q_t, Q_j \mid I_t) \left(\psi_j^*(X_0) - \hat{\psi}_j(X_0) \right) \right] \right| \\
&\quad + \left| \mathbb{E} \left[\left(\psi_t^*(X_0) - \hat{\psi}_t(X_0) \right)' (\hat{p}_t(I_t) - p_t(I_t)) (\hat{p}_{j,t}(I_t) - p_{j,t}(I_t))' \left(\psi_j^*(X_0) - \hat{\psi}_j(X_0) \right) \right] \right| \\
&\leq \sum_{j=t+1}^m \|\hat{\psi}_t - \psi_t^*\|_{u,v} \mathbb{E} \left[\|C_{t,j}\|_{\bar{u},\bar{u}} \|\psi_j^*(X_0) - \hat{\psi}_j(X_0)\|_{\bar{u}} \right]^{1/\bar{v}} \\
&\quad + 2M \|\hat{\psi}_t - \psi_t^*\|_{u,v} \|(\hat{p}_t - p_t) \circ (\hat{p}_{j,t} - p_{j,t})\|_{\bar{u},\bar{v}} \\
&\leq \sum_{j=t+1}^m \sup_{x_0} \mathbb{E} \left[\|C_{t,j}\|_{\bar{u},\bar{u}} \mid X_0 = x_0 \right]^{1/\bar{v}} \|\hat{\psi}_t - \psi_t^*\|_{u,v} \|\hat{\psi}_j - \psi_j^*\|_{u,\bar{v}} \\
&\quad + 2M \|\hat{\psi}_t - \psi_t^*\|_{u,v} \|(\hat{p}_t - p_t) \circ (\hat{p}_{j,t} - p_{j,t})\|_{\bar{u},\bar{v}} \\
&= \sum_{j=t+1}^m \sqrt{c_{t,j}} \|\hat{\psi}_t - \psi_t^*\|_{u,v} \|\hat{\psi}_j - \psi_j^*\|_{u,\bar{v}} + 2M \|\hat{\psi}_t - \psi_t^*\|_{u,v} \|(\hat{p}_t - p_t) \circ (\hat{p}_{j,t} - p_{j,t})\|_{\bar{u},\bar{v}}
\end{aligned}$$

Finally, we analyze the quantity A . First it is straightforward to verify that the loss satisfies the Neyman orthogonality property, i.e. for any h :

$$D_{h,\psi_t} \mathcal{L}_D(\psi_t^*; \underline{\psi}_{t+1}^*, h_0)[h - h_0, \hat{\psi}_t - \psi_t^*] = 0$$

By a second order Taylor's theorem with integral remainder,

$$A = - \int_0^1 \underbrace{D_{h,h,\psi_t} \mathcal{L}_D(\psi_t^*; \underline{\psi}_{t+1}^*, h_\tau)}_{G(h_\tau, \hat{h}, \hat{\psi}_t)} [h - h_0, \hat{h} - h_0, \hat{\psi}_t - \psi_t^*] d\tau$$

where $h_\tau = h_0 + \tau(\hat{h} - h_0)$. Moreover:

$$\begin{aligned}
G(h_\tau, \hat{h}, \hat{\psi}_t) &= \mathbb{E} \left[\left(\psi_t^*(X_0) - \hat{\psi}_t(X_0) \right)' (\hat{p}_t(I_t) - p_t(I_t)) (\hat{q}_t(I_t) - q_t(I_t)) \right] \\
&\quad - \mathbb{E} \left[\left(\psi_t^*(X_0) - \hat{\psi}_t(X_0) \right)' (\hat{p}_t(I_t) - p_t(I_t)) (\hat{p}_{j,t}(I_t) - p_{j,t}(I_t))' \psi_j^*(X_0) \right] \\
&\quad - 2 \mathbb{E} \left[\left(\psi_t^*(X_0) - \hat{\psi}_t(X_0) \right)' (\hat{p}_t(I_t) - p_t(I_t)) (\hat{p}_{t,t}(I_t) - p_{t,t}(I_t))' \psi_t^*(X_0) \right]
\end{aligned}$$

Thus by repeatedly applying Holder's inequality, we have that:

$$|A| \leq \|\hat{\psi}_t - \psi_t^*\|_{u,v} \|(\hat{p}_t - p_t) \circ (\hat{q}_t - q_t)\|_{\bar{u},\bar{v}} + 3M \|\hat{\psi}_t - \psi_t^*\|_{u,v} \sum_{j=t}^m \|(\hat{p}_t - p_t) \circ (\hat{p}_{j,t} - p_{j,t})\|_{\bar{u},\bar{v}}$$

By an AM-GM inequality, for all $a, b \geq 0$ and $\sigma > 0$: $a \cdot b \leq \frac{1}{2}(\frac{2}{\sigma}a^2 + \frac{\sigma}{2}b^2)$. Combining all the above inequalities and applying the AM-GM inequality for each term that contains $\|\hat{\psi}_t - \psi_t^*\|_{u,v}$ and re-arranging, yields that:

$$\begin{aligned}
\frac{\lambda}{2} \|\hat{\psi}_t - \psi_t^*\|_{2,2}^2 - \frac{\sigma}{4} \|\hat{\psi}_t - \psi_t^*\|_{u,v}^2 &\leq \epsilon(\hat{\psi}_t, \hat{h}) + \sum_{j=t+1}^m \frac{c_{t,j}}{\lambda} \|\hat{\psi}_j - \psi_j^*\|_{u,\bar{v}}^2 \\
&\quad + \frac{4M}{\lambda} \sum_{j=t}^m \|(\hat{p}_t - p_t) \circ (\hat{p}_{j,t} - p_{j,t})\|_{\bar{u},\bar{v}}^2 + \frac{1}{\lambda} \|(\hat{p}_t - p_t) \circ (\hat{q}_t - q_t)\|_{\bar{u},\bar{v}}^2
\end{aligned}$$

□

A.6 Proof of Theorem 10

First we prove a preliminary auxiliary lemma.

Lemma 13. Consider any sequence of non-negative numbers a_1, \dots, a_m satisfying the inequality:

$$a_t \leq \mu_t + \sum_{j=t+1}^m \kappa_j a_j$$

with $\mu_t, \kappa_t \geq 0$. Let $c := \max_{t \in [m]} \sum_{j=t+1}^m \kappa_j$ and $\mu := \max_{t \in [m]} \mu_t$. Then it must also hold that:

$$a_t \leq \mu \frac{c^{m-t+1} - 1}{c - 1}$$

Proof. By the assumed inequality we have that:

$$a_t \leq \mu_t + \max_{j>t} a_j \sum_{j=t+1}^m \kappa_j \leq \mu + c \max_{j>t} a_j$$

We prove by induction that $a_t \leq \mu \sum_{\tau=0}^{m-t} c^\tau$. By the above inequality, the property holds for $t = m$. Assuming it holds for $j > t$, we have that:

$$a_t \leq \mu + c \mu \sum_{\tau=0}^{m-t-1} c^\tau = \mu + \mu \sum_{\tau=0}^{m-t-1} c^{\tau+1} = \mu \sum_{\tau=0}^{m-t} c^\tau$$

which completes the inductive proof. Finally, observe that $\sum_{\tau=0}^{m-t} c^\tau = \frac{c^{m-t+1} - 1}{c - 1}$. \square

Proof of Theorem 10. Let $\mathcal{L}_{S,t}, \mathcal{L}_{S',t}$ denote the empirical loss over the samples in set S and S' correspondingly and let $\hat{h}_S, \hat{h}_{S'}$ the nuisance functions trained on sample S and S' correspondingly.

Denote the random variable $Z = \{X_0, X_1, T_1, \dots, X_m, T_m, Y\}$ and consider the loss function:

$$u(Z, \underline{\psi}_t(X_0); \hat{h}) := \left(Y - \hat{q}(I_t) - \sum_{j=t}^m \psi_j(X_0)' (Q_j - \hat{p}_{j,t}(I_t)) - \psi_t(X_0)' (\phi(\bar{X}_t, \bar{T}_t) - \hat{p}_{t,t}(I_t)) \right)^2$$

and the centered loss function:

$$\ell(Z, \underline{\psi}_t(X_0); \hat{h}) := u(Z, \underline{\psi}_t(X_0); \hat{h}) - u(Z, \psi_t^*, \underline{\psi}_{t+1}(X_0); \hat{h})$$

When all random variables are bounded, then ℓ is L -Lipschitz, for some constant L , with respect to $\underline{\psi}_t(X_0)$, i.e.:

$$|\ell(Z; \underline{\psi}_t(X_0); \hat{h}) - \ell(Z; \underline{\psi}'_t(X_0); \hat{h})| \leq L \|\underline{\psi}_t(X_0) - \underline{\psi}'_t(X_0)\|_2$$

Moreover, note that:

$$\mathcal{L}_{S,t}(\psi_t; \underline{\psi}_{t+1}, \hat{h}) - \mathcal{L}_{S,t}(\psi_t^*; \underline{\psi}_{t+1}, \hat{h}) = \frac{1}{|S|} \sum_{i \in S} \ell(Z^i, \underline{\psi}_t(X_0^i); \hat{h})$$

and that $\ell(Z, \underline{\psi}_t^*(X_0); \hat{h}) = 0$. For succinctness, for any given \hat{h} , let:

$$\mathcal{L}_{S,t}^*(\underline{\psi}_t, \hat{h}) := \mathcal{L}_{S,t}(\psi_t; \underline{\psi}_{t+1}, \hat{h}) - \mathcal{L}_{S,t}(\psi_t^*; \underline{\psi}_{t+1}, \hat{h}) = \frac{1}{|S|} \sum_{i \in S} \ell(Z^i, \underline{\psi}_t(X_0^i); \hat{h})$$

$$\mathcal{L}_{D,t}^*(\underline{\psi}_t, \hat{h}) := \mathcal{L}_{D,t}(\psi_t; \underline{\psi}_{t+1}, \hat{h}) - \mathcal{L}_{D,t}(\psi_t^*; \underline{\psi}_{t+1}, \hat{h}) = \mathbb{E} \left[\ell(Z; \underline{\psi}_t(X_0); \hat{h}) \right]$$

Suppose that δ_n upper bounds the critical radius of the function classes $\{\Psi_{t,i}\}_{t \in [m], i \in [r]}$ and $\delta_n = \Omega\left(\frac{r \log(\log(n))}{n}\right)$ and let $\delta_{n,\zeta} = c_0 \sqrt{\frac{\log(c_1 m/\zeta)}{n}}$, for some appropriately defined universal constants c_0, c_1 . Then by Lemma 11 of [12], w.p. $1 - \zeta$, for all $t \in [m]$:

$$\left| \mathcal{L}_{S,t}^*(\hat{\psi}_t, \hat{h}_{S'}) - \mathcal{L}_{D,t}^*(\hat{\psi}_t, \hat{h}_{S'}) \right| \leq O\left(r m \delta_{n/2,\zeta} \|\hat{\psi}_t - \psi_t^*\|_{2,2} + r m \delta_{n/2,\zeta}^2\right)$$

Analogously we can exchange the roles of S and S' and define corresponding losses and properties.

Moreover, we have that by the definition of the final stage of Dynamic DML:

$$\frac{1}{2} \left(\mathcal{L}_{S,t}^*(\psi_t, \hat{h}_{S'}) + \mathcal{L}_{S',t}^*(\psi_t, \hat{h}_S) \right) \leq 0$$

Thus we have that w.p. $1 - 2\zeta$:

$$\frac{1}{2} \sum_{O \in \{S, S'\}} \mathcal{L}_{D,t}^*(\hat{\psi}_t) \leq O\left(r m \delta_{n/2,\zeta} \|\hat{\psi}_t - \psi_t^*\|_{2,2} + r m \delta_{n/2,\zeta}^2\right) \quad (4)$$

Moreover, note that instantiating Lemma 9 with nuisance functions \hat{h}_S and $\hat{h}_{S'}$ correspondingly, we have that in the notation of that lemma:

$$\mathcal{L}_{D,t}^*(\hat{\psi}_t, \hat{h}_S) := \epsilon(\hat{\psi}_t, \hat{h}_S)$$

Applying Theorem 9 for each $\hat{h}_S, \hat{h}_{S'}$ and averaging the final inequality we get:

$$\frac{\lambda}{4} \|\hat{\psi} - \psi^*\|_{2,2}^2 \leq \frac{1}{2} \sum_{O \in \{S, S'\}} \left(\epsilon(\hat{\psi}_t, \hat{h}_O) + \rho_t(\hat{h}_O) + \sum_{j=t+1}^m \frac{c_{t,j}}{\lambda} \|\hat{\psi}_j - \psi_j^*\|_{2,2}^2 \right).$$

where:

$$\rho_t(\hat{h}) := \frac{1}{\lambda} \|(\hat{p}_{t,t} - p_{t,t}) \circ (\hat{q}_t - q_t)\|_{2,2}^2 + \frac{4M}{\lambda} \sum_{j=t}^m \|(\hat{p}_{t,t} - p_{t,t}) \circ (\hat{p}_{j,t} - p_{j,t})\|_{2,2}^2$$

Plugging in the bound from Equation (4), we have that for some universal constant c_0 :

$$\frac{\lambda}{4} \|\hat{\psi} - \psi^*\|_{2,2}^2 \leq c_0 r m \delta_{n/2,\zeta} \|\hat{\psi}_t - \psi_t^*\|_{2,2} + c_0 r m \delta_{n/2,\zeta}^2 + \sum_{j=t+1}^m \frac{c_{t,j}}{\lambda} \|\hat{\psi}_j - \psi_j^*\|_{2,2}^2 + \frac{1}{2} \sum_{O \in \{S, S'\}} \rho_t(\hat{h}_O).$$

By an AM-GM inequality, for all $a, b \geq 0$: $a \cdot b \leq \frac{1}{2}(\frac{2}{\sigma} a^2 + \frac{\sigma}{2} b^2)$. Thus for any $\epsilon > 0$, we can achieve:

$$\frac{\lambda - \epsilon}{4} \|\hat{\psi}_t - \psi_t^*\|_{2,2}^2 \leq c_0 \frac{1 + \epsilon}{\epsilon} r^2 \delta_{n/2,\zeta}^2 + \sum_{j=t+1}^m \left(\frac{c_{t,j}}{\lambda} + \frac{\epsilon}{4} \right) \|\hat{\psi}_j - \psi_j^*\|_{2,2}^2 + \frac{1}{2} \sum_{O \in \{S, S'\}} \rho_t(\hat{h}_O).$$

Invoking Lemma 13, we thus get that:

$$\|\hat{\psi}_t - \psi_t^*\|_{2,2}^2 \leq \mu_\epsilon \frac{c_\epsilon^{m-t+1} - 1}{c_\epsilon - 1}$$

where:

$$c_\epsilon := \max_{t \in [m]} \frac{4}{\lambda - \epsilon} \sum_{j=t+1}^m \left(\frac{c_{t,j}}{\lambda} + \frac{\epsilon}{4} \right)$$

$$\mu_\epsilon := \frac{4c_0}{\lambda - \epsilon} \frac{1 + \epsilon}{\epsilon} r^2 \delta_{n/2,\zeta}^2 + \frac{2}{\lambda - \epsilon} \max_{t \in [m]} \sum_{O \in \{S, S'\}} \rho_t(\hat{h}_O).$$

Let $c := \max_{t \in [m]} \frac{4}{\lambda} \sum_{j=t+1}^m \frac{c_{t,j}}{\lambda}$. Choosing $\epsilon \leq \frac{c\lambda}{m^2}$ and $\epsilon \leq \frac{\lambda}{2m}$, we have:

$$c_\epsilon \leq \frac{4}{\lambda(1 - 1/2m)} \sum_{j=t+1}^m \frac{c_{t,j}}{\lambda} + \frac{1}{\lambda(1 - 1/2m)} m \epsilon \leq \frac{1 + 1/m}{1 - 1/2m} c$$

$$\mu_\epsilon \leq \frac{8c_0}{\lambda} \left(1 + \frac{2m^2}{\lambda} \max\left\{ \frac{m}{c}, 2 \right\} \right) r^2 \delta_{n/2,\zeta}^2 + \frac{8}{\lambda} \max_{t \in [m]} \sum_{O \in \{S, S'\}} \rho_t(\hat{h}_O)$$

Moreover:

$$c_\epsilon^m \leq \frac{(1 + 1/m)^m}{(1 - 1/2m)^m} c^m \leq 2ec^m$$

Thus we have that:

$$\|\hat{\psi}_t - \psi_t^*\|_{2,2}^2 \leq \frac{8}{\lambda} \frac{2ec^m - 1}{2ec - 1} \left(c_0 \left(1 + \frac{m}{\lambda} \max \left\{ \frac{m}{c}, 2 \right\} \right) r^2 \delta_{n/2, \zeta}^2 + \max_{\substack{t \in [m] \\ O \in \{S, S'\}}} r_t(\hat{h}_O) \right)$$

Taking expectation of this inequality and integrating the tail bound, we have that:

$$\mathbb{E}_{\hat{\psi}} \left[\|\hat{\psi}_t - \psi_t^*\|_{2,2}^2 \right] \leq \frac{8}{\lambda} \frac{2ec^m - 1}{2ec - 1} \left(c_0 \left(1 + \frac{m}{\lambda} \max \left\{ \frac{m}{c}, 2 \right\} \right) r^2 \delta_{n/2}^2 + 2m \max_{\substack{t \in [m] \\ O \in \{S, S'\}}} \mathbb{E}_{\hat{h}_O} [r_t(\hat{h}_O)] \right)$$

If $m, c_{t,j}, \lambda, M$ are constants and

$$\begin{aligned} \max_{1 \leq t \leq j \leq m} \mathbb{E}_{\hat{p}_{t,t}, \hat{p}_{j,t}} \left[\left\| (\hat{p}_{t,t} - p_{t,t}) \circ (\hat{p}_{j,t} - p_{j,t}) \right\|_{2,2}^2 \right] &= O(r^2 \delta_{n/2}^2) \\ \max_{1 \leq t \leq m} \mathbb{E}_{\hat{p}_{t,t}, \hat{q}_t} \left[\left\| (\hat{p}_{t,t} - p_{t,t}) \circ (\hat{q}_t - q_t) \right\|_{2,2}^2 \right] &= O(r^2 \delta_{n/2}^2) \end{aligned}$$

then we have that:

$$\max_{t \in [m]} \mathbb{E}_{\hat{\psi}} \left[\|\hat{\psi}_t - \psi_t^*\|_{2,2}^2 \right] \leq O \left(r^2 \delta_{n/2}^2 \right)$$

□

A.7 Proof of Theorem 11

Lemma 14 (Restricted Cone Property). *Assume that all random variables and functions are bounded. Suppose we set κ_t such that it satisfies:*

$$\frac{\kappa_t}{2} \geq \delta_{n,\zeta} + \sum_{j=t+1}^m (c_{t,j} + \delta_{n,\zeta}) \|\nu_j\|_1 + \epsilon_n$$

where:

$$\begin{aligned} \delta_{n,\zeta} &:= c_0 \sqrt{\frac{\log(c_1 m r / \zeta)}{n}} \\ \epsilon_n &:= \max_{m \geq j \geq t \geq 1, i, i' \in [r]} \max \{ \|\hat{p}_{t,i} - p_{t,i}\|_2 \|\hat{q}_{t,i'} - q_{t,i'}\|_2, 4M m \|\hat{p}_{t,i} - p_{t,i}\|_2 \|\hat{p}_{j,t,i'} - p_{j,t,i'}\|_2 \} \end{aligned}$$

where c_0, c_1 are universal constants and $c_{t,j} = \|\mathbb{E}[\text{Cov}(Q_t, Q_j \mid \bar{X}_t, \bar{T}_{t-1})]\|_\infty$. Let $\nu_t := \hat{\psi}_t - \psi_t^*$, let T be the support of ψ_t^* and T^c its complement. Let $\nu_t^{(T)}$ denote the sub-vector support on T and $\mathcal{C}(T; 3) := \{\nu_t \in \mathbb{R}^r : \|\nu_t^{(T^c)}\|_1 \leq 3\|\nu_t^{(T)}\|_1\}$. Then w.p. $1 - \zeta$: $\nu_t \in \mathcal{C}(T; 3)$.

Proof of Lemma 14. Let $\mathcal{L}_{n,t}(\psi_t; \underline{\psi}_{t+1}, \hat{h})$ denote the un-penalized empirical loss using the estimated first stage nuisance functions, where by $\hat{h} = (\hat{h}_S, \hat{h}_{S'})$ we denote the union of both nuisance functions estimated on each side of the split. Moreover, let $\bar{\mathcal{L}}_{D,t}(\psi_t; \underline{\psi}_{t+1}, \hat{h}) = \frac{1}{2} \sum_{O \in \{S, S'\}} \mathcal{L}_{D,t}(\psi_t; \underline{\psi}_{t+1}, \hat{h}_O)$. By optimality of $\hat{\psi}_t$ for the penalized empirical square loss:

$$\mathcal{L}_{n,t}(\hat{\psi}_t; \underline{\hat{\psi}}_{t+1}, \hat{h}) - \mathcal{L}_{n,t}(\psi_t^*; \underline{\psi}_{t+1}, \hat{h}) \leq \kappa_t \left(\|\psi_t^*\|_1 - \|\hat{\psi}_t\|_1 \right)$$

Let $\nu_t = \hat{\psi}_t - \psi_t^*$ and T denote the support of ψ_t^* (i.e. the set of non-zero coordinates). Denote with $\nu^{(T)}$ the sub-vector on the set T and T^c the complement of T . By additivity of the ℓ_1 norm and the triangle inequality:

$$\begin{aligned}\|\hat{\psi}_t\|_1 &= \|\psi_t^* + \nu_t^{(T)}\|_1 + \|\nu_t^{(T^c)}\|_1 \geq \|\psi_t^*\|_1 - \|\nu_t^{(T)}\|_1 + \|\nu_t^{(T^c)}\|_1 \\ \|\psi_t^*\|_1 - \|\hat{\psi}_t\|_1 &\leq \|\nu_t^{(T)}\|_1 - \|\nu_t^{(T^c)}\|_1\end{aligned}$$

Moreover, by convexity of the empirical square loss, we have:

$$\begin{aligned}\mathcal{L}_{n,t}(\hat{\psi}_t; \hat{h}) - \mathcal{L}_{n,t}(\psi_t^*; \hat{h}) &\geq D_{\psi_t} \mathcal{L}_{n,t}(\psi_t^*; \hat{\psi}_{t+1}, \hat{h})[\nu_t] \\ &= D_{\psi_t} \left(\mathcal{L}_{n,t}(\psi_t^*; \hat{\psi}_{t+1}, \hat{h}) - \bar{\mathcal{L}}_{D,t}(\psi_t^*; \hat{\psi}_{t+1}, \hat{h}) \right) [\nu_t] \\ &\quad + D_{\psi_t} \mathcal{L}_{D,t}(\psi_t^*; \hat{\psi}_{t+1}, \hat{h})[\nu_t] \\ &\geq \underbrace{D_{\psi_t} \left(\mathcal{L}_{n,t}(\psi_t^*; \hat{\psi}_{t+1}, \hat{h}) - \bar{\mathcal{L}}_{D,t}(\psi_t^*; \hat{\psi}_{t+1}, \hat{h}) \right)}_F [\nu_t] \\ &\quad + \underbrace{D_{\psi_t} \left(\bar{\mathcal{L}}_{D,t}(\psi_t^*; \hat{\psi}_{t+1}, \hat{h}) - \mathcal{L}_{D,t}(\psi_t^*; \hat{\psi}_{t+1}, h_0) \right)}_G [\nu_t]\end{aligned}$$

where in the last step we used the optimality of ψ_t^* for the oracle population loss. Observe that the second quantity G is equal to the terms $A + B$ from the proof of Lemma 9. Thus from very similar arguments with those in that proof (with $u = 1$ and $v = \infty$), albeit with slight modifications that take advantage of the fact that $v = \infty$,⁴ we have that:

$$\begin{aligned}|G| &\leq \|\nu_t\|_1 \left(\sum_{j=t+1}^m c_{t,j} \|\nu_j\|_1 + \|\|(\hat{p}_t - p_t)(\hat{q}_t - q_t)'\|_1\|_\infty + 4M \sum_{j=t}^m \|\|(\hat{p}_t - p_t)(\hat{p}_{j,t} - p_{j,t})'\|_1\|_\infty \right) \\ &\leq \|\nu_t\|_1 \left(\sum_{j=t+1}^m c_{t,j} \|\nu_j\|_1 + \underbrace{\max_{j \geq t, i, i' \in [r]} \{ \|\hat{p}_{t,i} - p_{t,i}\|_2 \|\hat{q}_{t,i'} - q_{t,i'}\|_2, 4M m \|\hat{p}_{t,i} - p_{t,i}\|_2 \|\hat{p}_{j,t,i'} - p_{j,t,i'}\|_2 \}}_{\epsilon_{n,t}} \right)\end{aligned}$$

where $c_{t,j} := \|\mathbb{E}[C_{t,j}]\|_\infty$ and $M := \max_{t \in [m], \psi_t \in \Psi_t} \|\psi_t\|_1$. Moreover, note that we can write:

$$\begin{aligned}|F| &\leq |D_{\psi_t} \left(\mathcal{L}_{n,t}(\psi_t^*; \hat{\psi}_{t+1}, \hat{h}) - \bar{\mathcal{L}}_{D,t}(\psi_t^*; \hat{\psi}_{t+1}, \hat{h}) \right) [\nu_t]| \\ &\quad + |D_{\psi_t} \left(\mathcal{L}_{n,t}(\psi_t^*; \hat{\psi}_{t+1}, \hat{h}) - \mathcal{L}_{n,t}(\psi_t^*; \hat{\psi}_{t+1}, \hat{h}) - (\bar{\mathcal{L}}_{D,t}(\psi_t^*; \hat{\psi}_{t+1}, \hat{h}) - \bar{\mathcal{L}}_{D,t}(\psi_t^*; \hat{\psi}_{t+1}, \hat{h})) \right) [\nu_t]| \\ &\leq \|\nabla_{\psi_t} \left(\mathcal{L}_{n,t}(\psi_t^*; \hat{\psi}_{t+1}, \hat{h}) - \bar{\mathcal{L}}_{D,t}(\psi_t^*; \hat{\psi}_{t+1}, \hat{h}) \right)\|_\infty \|\nu_t\|_1 \\ &\quad + \|\nabla_{\psi_t} \left(\mathcal{L}_{n,t}(\psi_t^*; \hat{\psi}_{t+1}, \hat{h}) - \mathcal{L}_{n,t}(\psi_t^*; \hat{\psi}_{t+1}, \hat{h}) - (\bar{\mathcal{L}}_{D,t}(\psi_t^*; \hat{\psi}_{t+1}, \hat{h}) - \bar{\mathcal{L}}_{D,t}(\psi_t^*; \hat{\psi}_{t+1}, \hat{h})) \right)\|_\infty \|\nu_t\|_1\end{aligned}$$

The first term can be written as the maximum among $2r$ centered empirical processes of n bounded i.i.d. random variables. Thus we have that for some universal constants c_0, c_1 , w.p. $1 - \zeta$:

$$\|\nabla_{\psi_t} \left(\mathcal{L}_{n,t}(\psi_t^*; \hat{\psi}_{t+1}, \hat{h}) - \bar{\mathcal{L}}_{D,t}(\psi_t^*; \hat{\psi}_{t+1}, \hat{h}) \right)\|_\infty \leq c_0 \sqrt{\frac{\log(c_1 r / \zeta)}{n}}$$

For the second term note that for each half-split, we have that:

$$\begin{aligned}\nabla_{\psi_t} \left(\mathcal{L}_{n,t}(\psi_t^*; \hat{\psi}_{t+1}, \hat{h}) - \mathcal{L}_{n,t}(\psi_t^*; \hat{\psi}_{t+1}, \hat{h}) \right) &:= \sum_{j=t+1}^m \mathbb{E}_n \left[\tilde{T}_{t,t} \tilde{T}'_{j,t} \right] \nu_j \\ \nabla_{\psi_t} \left(\bar{\mathcal{L}}_{D,t}(\psi_t^*; \hat{\psi}_{t+1}, \hat{h}) - \bar{\mathcal{L}}_{D,t}(\psi_t^*; \hat{\psi}_{t+1}, \hat{h}) \right) &:= \sum_{j=t+1}^m \mathbb{E} \left[\tilde{T}_{t,t} \tilde{T}'_{j,t} \right] \nu_j\end{aligned}$$

⁴In particular, for any ν_t, ν_j and random variables X, Y , instead of applying Holder's inequality as: $|\mathbb{E}[\nu_t' X Y' \nu_j]| \leq \mathbb{E}[\|\nu_t\|_1 \|X\|_\infty \|Y\|_\infty \|\nu_j\|_1]$, which would have yielded a dependence on $\|X \circ Y\|_{\infty, 1}$, we apply: $|\mathbb{E}[\nu_t' X Y' \nu_j]| = |\nu_t' \mathbb{E}[X Y'] \nu_j| \leq \|\nu_t\|_1 \|\mathbb{E}[X Y']\|_\infty \|\nu_j\|_1$, which instead yields a dependence on $\|X \cdot Y'\|_1$.

Thus we can upper bound that term by:

$$\sum_{j=t+1}^m \left\| \mathbb{E}_n [\tilde{T}_{i,t} \tilde{T}'_{j,t}] - \mathbb{E} [\tilde{T}_{i,t} \tilde{T}'_{j,t}] \right\|_{\infty} \|\nu_j\|_1$$

Thus we need to control the maximum of $2mr^2$ centered empirical processes. Hence, we have that this term is upper bounded by $c_2 \sqrt{\frac{\log(c_3 m r / \zeta)}{n}}$, for some universal constants c_2, c_3 . Thus for some universal constants c_0, c_1 , we have that, w.p. $1 - \zeta$:

$$|F| \leq \|\nu_t\|_1 \underbrace{c_0 \sqrt{\frac{\log(c_1 m r / \zeta)}{n}}}_{\delta_{n,\zeta}} \left(1 + \sum_{j=t+1}^m \|\nu_j\|_1 \right)$$

Thus we have that:

$$F + G \geq -\|\nu_t\|_1 \left(\delta_{n,\zeta} + \sum_{j=t+1}^m (c_{t,j} + \delta_{n,\zeta}) \|\nu_j\|_1 + \epsilon_{n,t} \right)$$

Assuming that $\frac{\kappa_t}{2} \geq \delta_{n,\zeta} + \sum_{j=t+1}^m (c_{t,j} + \delta_{n,\zeta}) \|\nu_j\|_1 + \epsilon_{n,t}$, we thus have that:

$$\kappa_t \left(\|\nu_t^{(T)}\|_1 - \|\nu_t^{(T^c)}\|_1 \right) \geq F + G \geq -\frac{\kappa_t}{2} \|\nu_t\|_1 \implies \|\nu_t^{(T^c)}\|_1 \leq 3\|\nu_t^{(T)}\|_1$$

Hence $\nu_t \in \mathcal{C}(T; 3)$. □

Proof of Theorem 11. First note that for any $\underline{\psi}_t$ and for any ν_t :

$$\begin{aligned} \nu_t' \nabla_{\psi_t, \psi_t} \mathcal{L}_{n,t}(\psi_t; \underline{\psi}_{t+1}, \hat{h}) \nu_t &= \nu_t' \mathbb{E}_n [\tilde{T}_{t,t} \tilde{T}'_{t,t}] \nu_t \\ &= \nu_t' \mathbb{E} [\tilde{T}_{t,t} \tilde{T}'_{t,t}] \nu_t + \nu_t' \left(\mathbb{E} [\tilde{T}_{t,t} \tilde{T}'_{t,t}] - \mathbb{E}_n [\tilde{T}_{t,t} \tilde{T}'_{t,t}] \right) \nu_t \\ &\geq \nu_t' \mathbb{E} [\text{Cov}(Q_t, Q_t \mid \bar{X}_t, \bar{T}_{t-1})] \nu_t - \left\| \mathbb{E} [\tilde{T}_{t,t} \tilde{T}'_{t,t}] - \mathbb{E}_n [\tilde{T}_{t,t} \tilde{T}'_{t,t}] \right\|_{\infty} \|\nu_t\|_1^2 \\ &\quad \text{(by Lemma 12)} \\ &\geq \lambda \|\nu_t\|_2^2 - \left\| \mathbb{E} [\tilde{T}_{t,t} \tilde{T}'_{t,t}] - \mathbb{E}_n [\tilde{T}_{t,t} \tilde{T}'_{t,t}] \right\|_{\infty} \|\nu_t\|_1^2 \end{aligned}$$

With probability $1 - \zeta$, for all t , the latter is at least:

$$\lambda \|\nu_t\|_2^2 - \delta_{n,\zeta} \|\nu_t\|_1^2$$

Since κ_t is chosen such that $\nu_t \in \mathcal{C}(T; 3)$, we have that:

$$\|\nu_t\|_1 \leq 4\|\nu_t^{(T)}\|_1 \leq 4\sqrt{s} \|\nu_t^{(T)}\|_2 \leq 4\sqrt{s} \|\nu_t\|_2.$$

Thus we can lower bound the latter by:

$$\underbrace{(\lambda - 16s\delta_{n,\zeta})}_{\lambda_n} \|\nu_t\|_2^2$$

Thus the loss function $\mathcal{L}_{n,t}(\psi_t; \underline{\psi}_{t+1}, \hat{h})$ is λ_n -strongly convex. Combining with the optimality of $\hat{\psi}_t$ and the lower bound on κ_t , we have:

$$\begin{aligned} \kappa_t \left(\|\nu_t^{(T)}\|_1 - \|\nu_t^{(T^c)}\|_1 \right) &\geq \mathcal{L}_{n,t}(\hat{\psi}_t; \hat{\psi}_{t+1}, \hat{h}) - \mathcal{L}_{n,t}(\psi_t^*; \hat{\psi}_{t+1}, \hat{h}) \\ &\geq D_{\psi_t} \mathcal{L}_{n,t}(\psi_t^*; \hat{\psi}_{t+1}, \hat{h})[\nu_t] + \frac{1}{2} D_{\psi_t, \psi_t} \mathcal{L}_{n,t}(\psi_t^*; \hat{\psi}_{t+1}, \hat{h})[\nu_t, \nu_t] \\ &\geq -\frac{\kappa_t}{2} \|\nu_t\|_1 + \frac{\lambda_n}{2} \|\nu_t\|_2^2 \end{aligned}$$

Re-arranging yields:

$$\frac{\lambda_n}{2} \|\nu_t\|_2^2 \leq \frac{3\kappa_t}{2} \|\nu_t^{(T)}\|_1 \leq \frac{3\kappa_t}{2} \sqrt{s} \|\nu_t\|_2$$

Thus we conclude that w.p. $1 - \zeta$: $\|\nu_t\|_2 \leq \kappa_t \frac{3\sqrt{s}}{\lambda_n}$ and $\|\nu_t\|_1 \leq \kappa_t \frac{12s}{\lambda_n}$.

We also note that we need to chose κ_t such that:

$$\kappa_t \geq 2(\delta_{n,\zeta} + \epsilon_{n,t}) + \sum_{j=t+1}^m 2(c_{t,j} + \delta_{n,\zeta})\|\nu_j\|_1$$

If we let $\mu_n = \max_{t \in [m]} 2(\delta_{n,\zeta} + \epsilon_{n,t})$ and $c_n = \max_{t \in [m]} \sum_{j=t+1}^m 2(c_{t,j} + \delta_{n,\zeta})$, then it suffices to choose:

$$\kappa_t \geq \mu_n + c_n \max_{j>t} \|\nu_j\|_1$$

Finally, let $\beta_n = \frac{12s}{\lambda_n}$. Then under the above condition we have that $\|\nu_t\|_1 \leq \beta_n \kappa_t$.

We will show that for some sequence of $\kappa_1, \dots, \kappa_m$, we have: $\|\nu_j\|_1 \leq \beta_n \mu_n \sum_{j'>j} \beta_n^{m-j} c_n^{m-j}$. This holds for $j = m$ if we choose $\kappa_m := \mu_n$. If we assume this property holds for any $j > t$, then we have that for time t , if we choose:

$$\kappa_t := \mu_n + c_n \beta_n \mu_n \sum_{j>t+1} \beta_n^{m-j} c_n^{m-j} = \mu_n + \mu_n \sum_{j>t+1} \beta_n^{m-j+1} c_n^{m-j+1} = \mu_n \sum_{j>t} \beta_n^{m-j} c_n^{m-j}$$

then it satisfies the desired property that $\kappa_t \geq \mu_n + c_n \max_{j>t} \|\nu_j\|_1$ and therefore also:

$$\|\nu_t\|_1 \leq \beta_n \mu_n \sum_{j>t} \beta_n^{m-j} c_n^{m-j}$$

completing the inductive step. Thus we conclude that for all $t \in [m]$:

$$\|\nu_t\|_2 \leq \|\nu_t\|_1 \leq \beta_n \mu_n \sum_{j>t} \beta_n^{m-j} c_n^{m-j} = \beta_n \mu_n \frac{(\beta_n c_n)^{m-t+1} - 1}{\beta_n c_n - 1}$$

□

B Proofs of Results in Main Text

B.1 Proof of Lemma 1

Proof. Observe that by repeatedly expanding the state space in terms of the previous state and action pairs we can write that under any non-state dependent sequence of interventions τ_1, \dots, τ_m , the counterfactual value of the state at time m , $X_m^{(\tau)}$ can be written as:

$$X_m^{(\tau)} = \sum_{\kappa=1}^{m-1} B^{\kappa-1} A \cdot \tau_{m-\kappa} + \delta_m$$

where δ_t is a linear function of the mean-zero exogenous random shocks at all periods. Thus under intervention τ , the final outcome is distributed according to the random variable:

$$\begin{aligned} Y_m^{(\tau)} &= \theta'_0 \tau_m + \sum_{\kappa=1}^{m-1} \mu' B^{\kappa-1} A \cdot \tau_{m-\kappa} + \mu' \delta_m + \epsilon_t \\ &= \psi'_m \tau_m + \sum_{\kappa=1}^{m-1} \psi'_{m-\kappa} \tau_{m-\kappa} + \mu' \delta_m + \epsilon_t \end{aligned}$$

Since $V(\tau)$ is the expected difference of this counterfactual final outcome and the counterfactual final outcome when $\tau = 0$, and since δ_t is mean-zero conditional on any exogenous sequence of interventions that do not depend on the state or prior actions, we get the desired result. □

B.2 Proof of Theorem 2

Proof. For any $q \in \{0, \dots, m-1\}$, by repeatedly unrolling the state $X_m, q-1$ times we have that:

$$\begin{aligned} Y_m &= \theta'_0 T_m + \sum_{\kappa=0}^q \mu' B^{\kappa-1} A T_{m-\kappa} + \mu' B^q X_{m-q} + \sum_{\kappa=0}^{q-1} \mu' B^{\kappa-1} \eta_{m-\kappa} + \epsilon_m \\ &= \psi'_m T_m + \sum_{\kappa=0}^q \psi'_{m-\kappa} T_{m-\kappa} + \mu' B^q X_{m-q} + \sum_{\kappa=0}^q \mu' B^{\kappa-1} \eta_{m-\kappa} + \epsilon_m \\ &= \psi'_{m-q} T_{m-q} + \sum_{j=m-q+1}^m \psi'_j T_j + \mu' B^q X_{m-q} + \sum_{\kappa=0}^q \mu' B^{\kappa-1} \eta_{m-\kappa} + \epsilon_m \end{aligned}$$

Thus by re-arranging the equation, we have:

$$\bar{Y}_{m,m-q} = \psi'_{m-q} T_{m-q} + \mu' B^q X_{m-q} + \sum_{\kappa=0}^q \mu' B^{\kappa-1} \eta_{m-\kappa} + \epsilon_m$$

Since for all $t \geq m-q$, η_t, ϵ_t are subsequent random shocks to the variables T_{m-q}, X_{m-q} , we have that they are mean-zero conditional on T_{m-q}, X_{m-q} . Thus:

$$\mathbb{E}[\bar{Y}_{m,m-q} - \psi'_{m-q} T_{m-q} - \mu' B^q X_{m-q} \mid T_{m-q}, X_{m-q}] = 0$$

Thus we have concluded that for any $t \in \{1, \dots, m\}$:

$$\mathbb{E}[\bar{Y}_{m,t} - \psi'_t T_t - \mu' B^{m-q} X_t \mid T_t, X_t] = 0$$

which concludes the proof of the first part of the theorem.

For the second part, consider any q and let $t = m-q$. Note that by taking an expectation of the above condition for with respect to T_t conditional on X_t , we get that it implies

$$\mathbb{E}[\bar{Y}_{m,m-q} \mid X_{m-q}] = \psi'_{m-q} T_{m-q} - \mu' B^q X_{m-q} \mid T_{m-q}, X_{m-q}]$$

□

B.3 Proof of Theorem 3

Proof. Let Z denote the random variables associated with a single sample series. Observe that if we let:

$$\ell_\kappa(Z; \theta, p, q) = \left(\bar{Y}_{m,\kappa}^i - \theta_\kappa \tilde{T}_{m-\kappa, m-\kappa}^i \right)^2$$

Then we have that: $\nabla_{\theta_\kappa} \ell_\kappa(Z; \theta, p, q) = \psi_\kappa(Z; \theta, p, q)$ (where ψ, m are as defined in the proof of Theorem 4).

Moreover, observe that by consistency of the nuisance functions we have that for any $\nu = \theta_\kappa - \theta'_\kappa$:

$$\begin{aligned} \frac{\nu' \nabla_{\theta_\kappa \theta_\kappa} \mathbb{E}[\ell_\kappa(Z; \theta, \hat{p}, \hat{q})] \nu}{\|\nu\|_2^2} &\geq \frac{\nu' \nabla_{\theta_\kappa \theta_\kappa} \mathbb{E}[\ell_\kappa(Z; \theta, p^*, q^*)] \nu}{\nu} - O\left(\epsilon_{n,\delta} \frac{\|\nu\|_1^2}{\|\nu\|_2^2}\right) \\ &= \frac{\nu' \mathbb{E}[\zeta_{m-\kappa} \zeta'_{m-\kappa}] \nu}{\nu} + O\left(\epsilon_{n,\delta}^2 \frac{\|\nu\|_1^2}{\|\nu\|_2^2}\right) \geq 2\lambda - O(m d \epsilon_{n,\delta}) \end{aligned}$$

where the first inequality follows from smoothness of the second derivatives of ℓ_κ with respect to the nuisances and the mean-squared-error consistency of the nuisance functions. Thus for some $\lambda_n \geq \lambda - O(m d \epsilon_{n,\delta})$, the function $\mathbb{E}[\ell_\kappa(Z; \theta, \hat{p}, \hat{q})]$ is λ_n -strongly convex in θ_κ . By strong convexity we have:

$$\mathbb{E}[\ell_\kappa(Z; \theta, \hat{p}, \hat{q}) - \ell_\kappa(Z; \theta_{-k}, \theta_k^*, \hat{p}, \hat{q})] \geq m_\kappa(\theta_{-k}, \theta_k^*; p, q)' (\theta_\kappa - \theta_k^*) + \lambda_n \|\theta_\kappa - \theta_k^*\|_2^2$$

Moreover, by convexity we have:

$$\mathbb{E}[\ell_\kappa(Z; \theta_{-k}, \theta_k^*, p, q) - \ell_\kappa(Z; \theta, p, q)] \geq m_\kappa(\theta; p, q)' (\theta_k^* - \theta_\kappa)$$

Combining the two yields:

$$\begin{aligned}\lambda_n \|\theta_\kappa - \theta_\kappa^*\|_2^2 &\leq (m_\kappa(\theta; p, q) + m_\kappa(\theta_{-\kappa}, \theta_\kappa^*; p, q))'(\theta_\kappa^* - \theta_\kappa) \\ &\leq \|m_\kappa(\theta; p, q) + m_\kappa(\theta_{-\kappa}, \theta_\kappa^*; p, q)\|_2 \|\theta_\kappa^* - \theta_\kappa\|_2\end{aligned}$$

Thus we have:

$$\|\theta_\kappa - \theta_\kappa^*\| \leq \frac{1}{\lambda_n} \|m_\kappa(\theta; p, q) + m_\kappa(\theta_{-\kappa}, \theta_\kappa^*; p, q)\|$$

Observe that by the definition of our estimate $\hat{\theta}$, we have that:

$$\mathbb{E}_n[\psi_\kappa(Z; \hat{\theta}, \hat{p}, \hat{q})] = 0$$

Moreover, on each side of the fold we have that for a fixed θ :

$$\mathbb{E}[\psi_\kappa(Z; \theta, \hat{p}, \hat{q})] = m(\theta; \hat{p}, \hat{q})$$

Thus by the lipschitzness of the moment ψ with respect to θ and standard empirical process theory, w.p. $1 - \delta$:

$$\sup_{\theta \in [-H, H]^d} \|\mathbb{E}_n[\psi_\kappa(Z; \theta, \hat{p}, \hat{q})] - m_\kappa(\theta; \hat{p}, \hat{q})\|_2 \leq O\left(H d \sqrt{\frac{\log(d/\delta)}{n}}\right)$$

Thus we have that w.p. $1 - \delta$:

$$\|m_\kappa(\theta; \hat{p}, \hat{q})\|_2 \leq O\left(d \sqrt{\frac{\log(d/\delta)}{n}}\right)$$

Moreover, by the identifying equation assumption we also have that:

$$m_\kappa(\theta_0; p^*, q^*) = 0$$

Thus it suffices to understand the deviation of $m_\kappa(\theta_{-\kappa}, \theta_\kappa^*; \hat{p}, \hat{q})$ from the above. For this we will perform a second order Taylor expansion around all arguments, θ, p, q . We observe that m_κ is independent of θ_t for $t \geq \kappa$ and also by Neyman orthogonality, the first order term with respect to the nuisance models will be zero. Thus we have:

$$\begin{aligned}m_\kappa(\theta_{-\kappa}, \theta_\kappa^*; \hat{p}, \hat{q}) - m_\kappa(\theta_0; p^*, q^*) &\leq \nabla_{\theta_{<\kappa}} m_\kappa(\theta_0; p^*, q^*)'(\theta_{<\kappa} - \theta_{<\kappa}^*) + O(\|\hat{p} - p^*\|_2^2 + \|\hat{q} - q^*\|_2^2) \\ &\leq c_m \sup_{t < \kappa} \|\theta_t - \theta_t^*\|_2 + \epsilon_{n,\delta}^2\end{aligned}$$

where we used the fact that:

$$\nabla_{\theta_\tau} m_\kappa(\theta_0; p^*, q^*) = \mathbb{E}\left[\tilde{T}_{m-\tau, m-\kappa}^* \zeta_t'\right]$$

and by our assumption $\left\|\sum_{\tau=0}^m \|\mathbb{E}\left[\tilde{T}_{m-\tau, m-\kappa}^* \zeta_t'\right]\right\|_\infty \leq c_m$, and subsequently we applied an $\ell_1 - \ell_\infty$ Cauchy-Schwarz inequality coordinate-wise for each coordinate of the treatment. Thus overall we have, for some absolute constant C :

$$\|\theta_\kappa - \theta_\kappa^*\|_2 \leq \frac{1}{\lambda_n} \left(C \left(H d \sqrt{\frac{\log(d/\delta)}{n}} + \epsilon_{n,\delta}^2 \right) + c_m \sup_{t < \kappa} \|\theta_t - \theta_t^*\|_2 \right)$$

Let $\mu_{n,\delta} = \frac{C}{\lambda_n} \left(H d \sqrt{\frac{\log(d/\delta)}{n}} + \epsilon_{n,\delta}^2 \right)$. Then observe that:

$$\|\theta_\kappa - \theta_\kappa^*\|_2 \leq \mu_{n,\delta} + (c_m/\lambda_n) \sup_{t < \kappa} \|\theta_t - \theta_t^*\|_2$$

Then we have that: $\|\theta_0 - \theta_0^*\|_2 \leq \mu_{n,\delta}$ and by induction, if we assume that $\|\theta_t - \theta_t^*\|_2 \leq \mu_{n,\delta} \sum_{\tau=0}^t (c_m/\lambda_n)^\tau$ for all $t < \kappa$, then:

$$\|\theta_\kappa - \theta_\kappa^*\|_2 \leq \mu_{n,\delta} + (c_m/\lambda_n) \mu_{n,\delta} \sum_{\tau=0}^{\kappa-1} (c_m/\lambda_n)^\tau = \mu_{n,\delta} \sum_{t=0}^{\kappa} (c_m/\lambda_n)^\tau$$

Thus overall we have:

$$\begin{aligned}
\forall \kappa \in \{0, \dots, m\} : \|\theta_\kappa - \theta_\kappa^*\|_2 &\leq \mu_{n,\delta} \sum_{t=0}^m (c_m/\lambda_n)^\tau \leq \mu_{n,\delta} \frac{(c_m/\lambda_n)^{m+1} - 1}{c_m/\lambda_n - 1} \\
&= O\left(\mu_{n,\delta} \max\{1, (c_m/\lambda_n)^{m+1}\}\right) \\
&= O\left(\mu_{n,\delta} \max\left\{1, \left(\frac{c_m}{2\lambda - O(md\epsilon_{n,\delta})}\right)^{m+1}\right\}\right)
\end{aligned}$$

Which concludes the proof of the theorem. \square

B.4 Proof of Theorem 4

Proof. Observe that our estimator can be equivalently viewed as the solution to an empirical version of the following vector of moment conditions:

$$\forall \kappa \in \{0, \dots, m\} : m_\kappa(\theta; p^*, q^*) = \mathbb{E}[\psi_\kappa(Z; \theta, p^*, q^*)] = 0$$

where Z is a random vector denoting the variables of a single sample series and:

$$\psi_\kappa(Z; \theta, p^*, q^*) = \left(\bar{Y}_{m,\kappa}^* - \sum_{\tau=0}^{\kappa} \theta_\tau \tilde{T}_{m-\tau, m-\kappa}^* \right) \tilde{T}_{m-\kappa, m-\kappa}^*$$

and

$$\tilde{Y}_{m,\kappa}^* := Y_m - q_\kappa^*(X_{m-\kappa}) \quad \tilde{T}_{m-\tau, m-\kappa}^* := T_{m-\tau} - p_{\tau,\kappa}^*(X_{m-\kappa})$$

If we denote with \hat{p} and \hat{q} the estimates of the nuisance models, then our estimate $\hat{\theta}$ is equivalent to the solution of the system of equations:

$$\mathbb{E}_n \left[\psi(Z; \hat{\theta}, \hat{p}, \hat{q}) \right] = \frac{1}{n} \sum_{i=1}^n \psi(Z^i; \hat{\theta}, \hat{p}, \hat{q}) = 0$$

We now show that the moment vector ψ satisfies the property of Neyman orthogonality with respect to all the nuisance models p, q . For this it suffices to show that for all $\kappa \in \{0, \dots, m\}$:

$$\mathbb{E}[\nabla_{\pi_{\tau,\kappa}, \gamma_\kappa} \psi_\kappa(Z; \theta_0, p^*, q^*) \mid X_{m-\kappa}] = 0$$

where $\nabla_{\pi_{\tau,\kappa}, \gamma_\kappa}$ is the derivative of ψ with respect to the output of the nuisance models $p_{\tau,\kappa}^*, q_\kappa^*$ evaluated at $X_{m-\kappa}$.

For any κ , the derivative with respect to $\pi_{\tau,\kappa}$ for $\tau < \kappa$ is equal to:

$$\mathbb{E}[\tilde{T}_{m-\kappa, m-\kappa}^* \mid X_{m-\kappa}] = 0$$

while the derivative with respect to $\pi_{\kappa,\kappa}$ is:

$$-\mathbb{E}[\bar{Y}_{m,\kappa}^* - \sum_{\tau=0}^{\kappa} \theta_\tau \tilde{T}_{m-\tau, m-\kappa}^* \mid X_{m-\kappa}] + \mathbb{E}[\tilde{T}_{m-\kappa, m-\kappa}^* \mid X_{m-\kappa}] = 0$$

Similarly, the derivative with respect to γ_κ is:

$$-\mathbb{E}[\tilde{T}_{m-\kappa, m-\kappa}^* \mid X_{m-\kappa}] = 0$$

Moreover, the Jacobian M of the moment vector m at the true values θ_0, p^*, q^* is a block lower triangular matrix whose block values are of the form:

$$\forall 0 \leq \tau \leq \kappa \leq m : M_{\tau,\kappa} = \mathbb{E}[\tilde{T}_{m-\tau, m-\kappa}^* \zeta_{m-\kappa}^{\prime}]$$

Thus its diagonal block values take the form $M_{\kappa,\kappa} = \mathbb{E}[\zeta_{m-\kappa} \zeta_{m-\kappa}^{\prime}] \succeq \lambda I$. Hence, the minimum eigenvalue of M is at least λ .

Thus our setting and our estimator falls exactly into the framework of orthogonal moment estimation of [6] and we can directly apply Theorem 3.1 of [6] to get the result (the exact form of the matrix Σ follows from the observation that $\psi_\kappa(Z; \theta_0, p^*, q^*) = \epsilon_{m-\kappa} \zeta_{m-\kappa}$ and the fact that $\Sigma = \mathbb{E}[\psi(Z; \theta_0, p^*, q^*) \psi(Z; \theta_0, p^*, q^*)^{\prime}]$). \square

B.5 Proof of Theorem 5

Consider blocks of size m and let Z_b denote the set of all random variables in block b , starting from the first state $X_{b,0}$ of that block and ending in the final outcome $Y_{b,m}$. Moreover, let \mathcal{F}_b denote the filtration before each block b (i.e. all past variables before b) and let:

$$m_{b,\kappa}(\theta; p, q) = \mathbb{E}[\psi(Z_b; \theta, p, q) \mid \mathcal{F}_b]$$

Let

$$m_\kappa(\theta; p, q) = \frac{1}{B} \sum_b m_{b,\kappa}(\theta; p, q)$$

$$\Psi_{B,\kappa}(\theta; p, q) = \frac{1}{B} \sum_b \psi_\kappa(Z_b; \theta, p, q)$$

By a second order Taylor expansion of $m(\theta; p, q)$ around θ_0, p_0, q_0 , and due to orthogonality of the moment with respect to p, q , the linearity of the moment with respect to θ , and the fact that $m(\theta_0; p_0, q_0) = 0$ we have:

$$m(\theta; p, q) = \nabla_\theta m(\theta_0; p_0, q_0) (\theta - \theta_0) + \xi$$

where ξ , satisfies that:

$$\|\xi\| = m d O(\|p - p_0\|_\infty^2, \|q - q_0\|_\infty^2)$$

If we let: $M = \nabla_\theta m(\theta_0; p_0, q_0)$, we have that:

$$M(\theta - \theta_0) = m(\theta; p, q) + \rho$$

where ρ also satisfies:

$$\|\rho\| = m d O(\|p - p_0\|_\infty^2, \|q - q_0\|_\infty^2)$$

Moreover, observe that by the definition of our estimator, we have:

$$\Psi_B(\hat{\theta}; \hat{p}, \hat{q}) = 0$$

Thus we conclude that:

$$M(\hat{\theta} - \theta_0) = \left(m(\hat{\theta}; \hat{p}, \hat{q}) - \Psi_B(\hat{\theta}; \hat{p}, \hat{q}) \right) + \rho$$

Moreover, observe that $m(\theta_0; p^*, q^*) = 0$. Thus we can further expand the right hand side as:

$$\begin{aligned} m(\hat{\theta}; \hat{p}, \hat{q}) - \Psi_B(\hat{\theta}; \hat{p}, \hat{q}) &= -\Psi_B(\theta_0; p^*, q^*) + m(\hat{\theta}; \hat{p}, \hat{q}) - \Psi_B(\hat{\theta}; \hat{p}, \hat{q}) - (m(\theta_0; p^*, q^*) - \Psi_B(\theta_0; p^*, q^*)) \\ &= \underbrace{-\Psi_B(\theta_0; p^*, q^*)}_{\text{Asymptotic normal term}} + \underbrace{m(\hat{\theta}; \hat{p}, \hat{q}) - m(\theta_0; p^*, q^*) - (\Psi_B(\hat{\theta}; \hat{p}, \hat{q}) - \Psi_B(\theta_0; p^*, q^*))}_{A=\text{Stochastic equicontinuous term}} \end{aligned}$$

First we note that the term A decays to zero faster than a root- B rate. By a martingale Bernstein inequality (see e.g. [13] or [25]), we have that for any θ :

$$\|A\|_\infty \leq O\left(\sup_{\kappa \in \{0, \dots, m-1\}, i \in \{1, \dots, d\}} \sqrt{\frac{\sup_b \mathbb{E}[(\psi_{\kappa,i}(Z_b; \theta, \hat{p}, \hat{q}) - \psi_{\kappa,i}(Z_b; \theta_0, \hat{p}, \hat{q}))^2 \mid \mathcal{F}_b] \log(m d / \delta)}{B}} + \frac{\log(m d / \delta)}{B} \right)$$

Moreover, by the Lipschitzness of ψ , we have that:

$$\begin{aligned} \sqrt{\mathbb{E}[(\psi_{\kappa,i}(Z_b; \theta, \hat{p}, \hat{q}) - \psi_{\kappa,i}(Z_b; \theta_0, \hat{p}, \hat{q}))^2 \mid \mathcal{F}_b]} &\leq \\ m d O\left(\sqrt{\mathbb{E}[\|\hat{\theta} - \theta_0\|_\infty^2]}, \sqrt{\mathbb{E}[\|\hat{p} - p^*\|_{b,\infty}^2]}, \sqrt{\mathbb{E}[\|\hat{q} - q^*\|_{b,\infty}^2]} \right) \end{aligned}$$

By our estimation rate theorem for $\hat{\theta}$ and our assumption on the estimation rate of \hat{p} and \hat{q} we have that:

$$\forall \kappa \in \{0, \dots, m-1\} : \sqrt{\mathbb{E}[\|\hat{\theta} - \theta_0\|_{\infty}^2]} \leq O\left(d C_m \left(\sqrt{\frac{\log(dB)}{B}} + \epsilon_B^2\right)\right)$$

and $\sqrt{\mathbb{E}[\|\hat{p} - p^*\|_{b,\infty}^2]}, \sqrt{\mathbb{E}[\|\hat{q} - q^*\|_{b,\infty}^2]} \leq \epsilon_B = o(B^{-1/4})$. Since by assumption $m C_m = o(\sqrt{B})$, we have:

$$\sqrt{B} m d O\left(\sqrt{\mathbb{E}[\|\hat{\theta} - \theta_0\|_{\infty}^2]}, \sqrt{\mathbb{E}[\|\hat{p} - p^*\|_{b,\infty}^2]}, \sqrt{\mathbb{E}[\|\hat{q} - q^*\|_{b,\infty}^2]}\right) = o_p(1)$$

Thus:

$$\sqrt{B}\|A\| = o_p(1)$$

Moreover, for the exact same reason:

$$\sqrt{B}\|\rho\| = o_p(1)$$

Thus it suffices to show that:

$$-\sqrt{B} \Psi_B(\theta_0; p^*, q^*) \rightarrow_d N(0, \Sigma)$$

Observe that under the true values of θ_0, p^*, q^* , the moment $\psi_{b,\kappa}(Z_b; \theta_0, p^*, q^*) = \epsilon_{b,\kappa} \zeta_{b,\kappa}$ is only a function of the residual random noises, $\zeta_{b,\kappa}, \epsilon_{b,\kappa}$. Since these variables are independent and identically distributed across all blocks⁵, we can conclude that the conditional variance:

$$\text{Var}(\psi_{\kappa}(Z_b; \theta_0, p^*, q^*) | \mathcal{F}_b) = \mathbb{E}[\epsilon_{b,\kappa}^2 \zeta_{b,\kappa} \zeta'_{b,\kappa} | \mathcal{F}_b] = \Sigma_{\kappa\kappa} \succ 2\sigma^2 \lambda I$$

where σ^2 is the variance of the shocks $\epsilon_{b,\kappa}$ and 2λ is the minimum eigenvalue of the co-variance of $\zeta_{b,\kappa}$. Thus Σ is a positive definite symmetric matrix with minimum eigenvalue independent of m . Thus $\text{Var}(\psi_{\kappa}(Z_b; \theta_0, p^*, q^*) | \mathcal{F}_b)$ is a constant positive definite matrix, independent of b and \mathcal{F}_b . We can thus apply a Martingale Central Limit Theorem (see e.g. [15]) to get the desired statement:

$$\sqrt{B}\Sigma^{-1/2}M(\hat{\theta} - \theta_0) \rightarrow_d N(0, I)$$

C Omitted Content from Section 5

Concrete Rates for Lasso Nuisance Estimates. Suppose that the observational policy p is also linear, i.e. $p(X) = AX$. Then all the models q_{κ} and $p_{\tau,\kappa}$ are high-dimensional linear functions of their input arguments, i.e. $q_{\kappa}(x) = \phi'x$ and $p_{\tau,\kappa,j}(x) = \pi'_{\tau,\kappa,j}x$. If these linear functions satisfy a sparsity constraint then under standard regularity assumptions we can guarantee if we use the Lasso regression to estimate each of these functions that w.p. $1 - \delta$, the estimation error of all nuisance models is $O\left(s\sqrt{\frac{\log(p/\delta)}{n}}\right)$, where s is an upper bound on the number of non-zero coefficients. One sufficient regularity condition is that the expected co-variance matrix of every period's state has full rank, i.e. $\mathbb{E}[X_t X_t'] \succeq \lambda I$. Thus the requirements of our main theorems of this section would be satisfied as long as the sparsity grows as $s = o(n^{1/4})$. These sparsity conditions are for instance satisfied if for instance, only s coordinates of the high-dimensional state have any effect on the final outcome (i.e. are outcome-relevant) and if only s coordinates of the high-dimensional state enter the observational policy.

D Omitted Content from Section 6

We present here a more formal description of the block dynamic DML algorithm and also provide a finite sample mean squared error theorem that complements our asymptotic normality result from the main part of the paper.

⁵Our theorem would easily extend if these variables form a martingale with a non-zero variance at each step

Algorithm 5 Block Dynamic DML with Progressive Nuisance Estimation

Partition the data into $B = n/m$ blocks of m periods
 Denote with X_t^b, T_t^b, Y_t^b the state, action, outcome pairs in the t -th period of block b .
for each block $b \in \{B/2, \dots, B\}$ **do**
 for each $\kappa \in \{0, \dots, m\}$ **do**
 Regress $Y_m^{b'}$ on $X_{m-\kappa}^{b'}$ using all blocks $b' < b$ to learn model $q_\kappa(x) = \mathbb{E}[Y_m \mid X_{m-\kappa} = x]$
 and calculate the residual $\tilde{Y}_{m,m-\kappa}^b = Y_m^b - q_\kappa(X_{m-\kappa}^b)$.
 for each $\tau \in \{0, \dots, \kappa\}$ **do**
 Regress each $T_{m-\tau}^{b'}$ on $X_{m-\kappa}^{b'}$ using blocks $b' < b$ to learn model $p_{\tau,\kappa}$ and calculate
 residual $\tilde{T}_{m-\tau,m-\kappa}^b = T_{m-\tau}^b - p_{\tau,\kappa}(X_{m-\kappa}^b)$ on the other half, and vice versa.
 end for
 end for
 Using all the blocks $b \in \{B/2, \dots, B\}$
for $\kappa = 0$ **to** m **do**
 Regress $\tilde{Y}_{m,\kappa} = \tilde{Y}_{m,m-\kappa} - \sum_{\tau < \kappa} \hat{\theta}'_\tau \tilde{T}_{m-\tau,m-\kappa}$ on $\tilde{T}_{m-\kappa,m-\kappa}$, i.e. find a solution $\hat{\theta}_\kappa$ to the
 estimating equation:

$$\frac{2}{B} \sum_{b=1}^{B/2} \left(\bar{Y}_{m,\kappa}^b - \theta_\kappa \tilde{T}_{m-\kappa,m-\kappa}^b \right) \tilde{T}_{m-\kappa,m-\kappa}^b = 0$$

end for

Theorem 15. Let \mathcal{F}_b denote the filtration up until block b . Suppose that $\mathbb{E}[\zeta_t \zeta_t' \mid \mathcal{F}_b] \succeq 2\lambda I$ and that each coordinate \hat{h} of each nuisance model $\{q_\kappa, p_{\tau,\kappa}\}_{\tau,\kappa}$ satisfy that w.p. $1 - \delta, \forall b \geq B/2$:

$$\|\hat{h} - h^*\|_{b,2} = \sqrt{\mathbb{E}_X[\hat{h}(X_b) - h(X_b)]^2 \mid \mathcal{F}_b} \leq \epsilon_{B,\delta}$$

where expectation is with respect to the corresponding input of each model from block b . Moreover, suppose that: $\left\| \sum_{\tau=0}^m \mathbb{E} \left[\tilde{T}_{m-\tau,m-\kappa}^* \zeta_t' \mid \mathcal{F}_b \right] \right\|_\infty \leq c_m$. Then w.p. $1 - 2\delta$:

$$\|\hat{\theta} - \theta_0\|_\infty \leq O \left(d C_m \left(\sqrt{\frac{\log(dmB/\delta)}{B}} + \epsilon_{B,\delta}^2 \right) \right)$$

where $C_m := \max \left\{ 1, \left(\frac{c_m}{2\lambda - O(m d \epsilon_{n,\delta})} \right)^{m+1} \right\} m$. If each coordinate \hat{h} of each nuisance model satisfies:

$$\forall b \geq B/2 : \mathbb{E}_{\hat{h}}[\|\hat{h} - h^*\|_{b,2}^4]^{(1/4)} \leq \epsilon_B$$

Then:

$$\sqrt{\mathbb{E}[\|\hat{\theta} - \theta_0\|_\infty^2]} \leq O \left(d C_m \left(\sqrt{\frac{\log(dmB)}{B}} + \epsilon_B^2 \right) \right).$$

Proof. Consider blocks of size m and let Z_b denote the set of all random variables in block b , starting from the first state $X_{b,0}$ of that block and ending in the final outcome $Y_{b,m}$. Moreover, let \mathcal{F}_b denote the filtration before each block b (i.e. all past variables before b) and let:

$$m_{b,\kappa}(\theta; p, q) = \mathbb{E}[\psi(Z_b; \theta, p, q) \mid \mathcal{F}_b]$$

Since by assumption $\mathbb{E}[\zeta_t \zeta_t' \mid \mathcal{F}_b] \succeq 2\lambda I$. Then:

$$\mathbb{E}[\ell_\kappa(Z_b; \theta, p, q) - \ell_\kappa(Z_b; \theta_{-\kappa}, \theta_\kappa^*, p, q) \mid \mathcal{F}_b] \geq m_{b,\kappa}(\theta_{-\kappa}, \theta_\kappa^*; p, q)'(\theta_\kappa - \theta_\kappa^*) + \lambda \|\theta_\kappa - \theta_\kappa^*\|^2$$

$$\mathbb{E}[\ell_\kappa(Z; \theta_{-\kappa}, \theta_\kappa^*, p, q) - \ell_\kappa(Z; \theta, p, q) \mid \mathcal{F}_b] \geq m_{b,\kappa}(\theta; p, q)'(\theta_\kappa^* - \theta_\kappa)$$

Combining the two yields:

$$\lambda \|\theta_\kappa - \theta_\kappa^*\|^2 \leq (m_{b,\kappa}(\theta; p, q) + m_{b,\kappa}(\theta_{-\kappa}, \theta_\kappa^*; p, q))'(\theta_\kappa^* - \theta_\kappa)$$

Averaging the latter inequality over the B blocks, applying Cauchy-Swartz inequality and dividing over by $\|\theta_\kappa - \theta_\kappa^*\|$:

$$\|\theta_\kappa - \theta_\kappa^*\| \leq \frac{1}{\lambda} \left\| \frac{1}{B} \sum_b (m_{b,\kappa}(\theta; p, q) + m_{b,\kappa}(\theta_{-\kappa}, \theta_\kappa^*; p, q)) \right\|$$

Observe that by the definition of our estimate $\hat{\theta}$, we have that:

$$\frac{1}{B} \sum_b \psi_\kappa(Z_b; \hat{\theta}, \hat{p}, \hat{q}) = 0$$

Moreover, conditioning on the first fold, we have that for any fixed θ :

$$\mathbb{E}[\psi_\kappa(Z_b; \theta, \hat{p}, \hat{q}) \mid \mathcal{F}_b] = m_{b,\kappa}(\theta; \hat{p}, \hat{q})$$

Thus by the lipschitzness of the moment ψ with respect to θ and martingale Azuma inequality and standard covering arguments, we have:

$$\sup_{\theta \in [-H, H]^d} \left\| \frac{1}{B} \sum_b (\psi_\kappa(Z_b; \theta, \hat{p}, \hat{q})) - m_{b,\kappa}(\theta; \hat{p}, \hat{q}) \right\| \leq O\left(\sqrt{\frac{d \log(dB/\delta)}{B}}\right)$$

Thus we have that w.p. $1 - \delta$:

$$\|m_{b,\kappa}(\hat{\theta}; \hat{p}, \hat{q})\| \leq O\left(\sqrt{\frac{d \log(dB/\delta)}{B}}\right)$$

Moreover, by the identifying equation assumption we also have that:

$$\forall b : m_{b,\kappa}(\theta_0; p^*, q^*) = 0$$

Thus it suffices to understand the deviation of $m_\kappa(\theta_{-\kappa}, \theta_\kappa^*; \hat{p}, \hat{q})$ from the above. For this we will perform a second order Taylor expansion around all arguments, θ, p, q . We observe that m_κ is independent of θ_t for $t \geq \kappa$ and also by Neyman orthogonality, the first order term with respect to the nuisance models will be zero. The remainder of the proof is identical to the final inductive part of the proof of Theorem 3. \square

E Further Experimental Results

E.1 Constant Treatment Effects

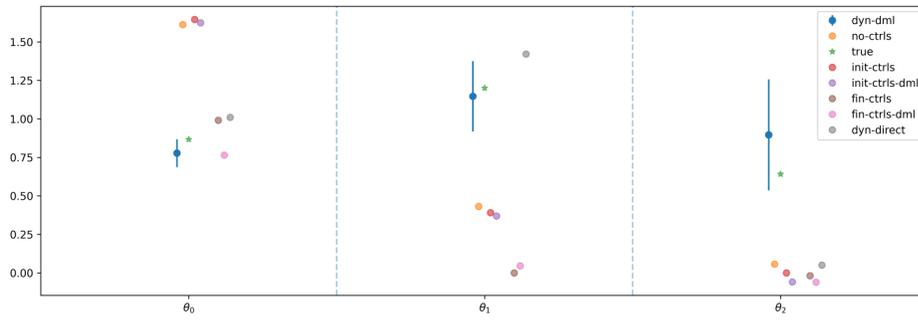
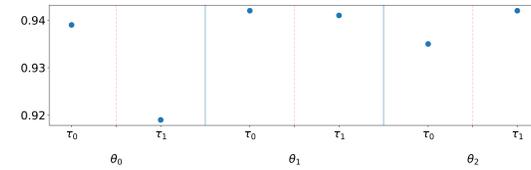
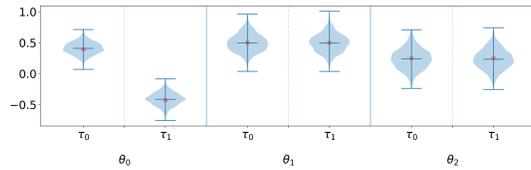


Figure 3: Comparison of DynamicDML (with confidence intervals) with benchmarks on a single instance. $n = 400$, $n_t = 2$, $n_x = 100$, $s = 10$, $\sigma(\epsilon_t) = .5$, $\sigma(\zeta_t) = .5$, $\sigma(\eta_t) = .5$, $C = 0$

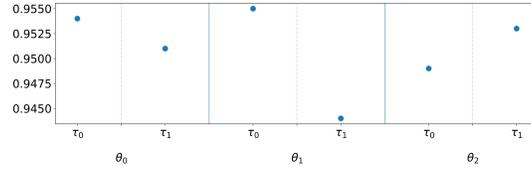


(a) Coverage

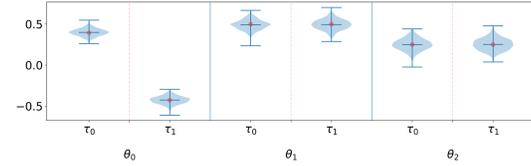


(b) Point estimates

Figure 4: $n = 500$, $n_t = 2$, $n_x = 450$, $s = 2$, $\sigma(\epsilon_t) = 1$, $\sigma(\zeta_t) = .5$, $\sigma(\eta_t) = 1$



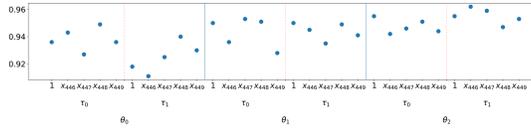
(a) Coverage



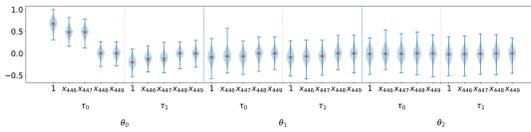
(b) Point estimates

Figure 5: $n = 2000$, $n_t = 2$, $n_x = 450$, $s = 2$, $\sigma(\epsilon_t) = 1$, $\sigma(\zeta_t) = .5$, $\sigma(\eta_t) = 1$

E.2 Heterogeneous Treatment Effects

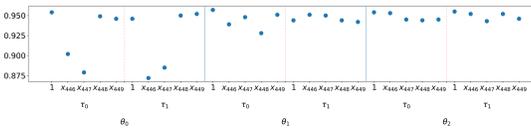


(a) Coverage

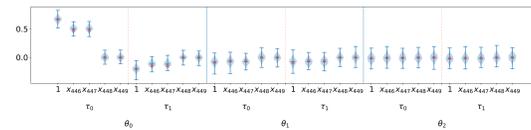


(b) Point estimates

Figure 6: Heterogeneous effects: $n = 500$, $n_t = 2$, $n_x = 450$, $s = 2$, $\sigma(\epsilon_t) = 1$, $\sigma(\zeta_t) = .5$, $\sigma(\eta_t) = 1$



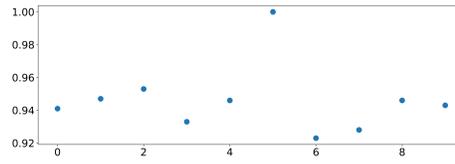
(a) Coverage



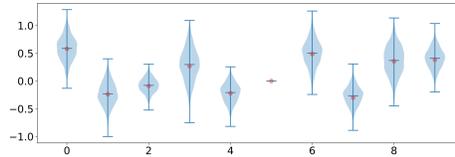
(b) Point estimates

Figure 7: Heterogeneous effects: $n = 2000$, $n_t = 2$, $n_x = 450$, $s = 2$, $\sigma(\epsilon_t) = 1$, $\sigma(\zeta_t) = .5$, $\sigma(\eta_t) = 1$

E.3 Counterfactual Policy Values

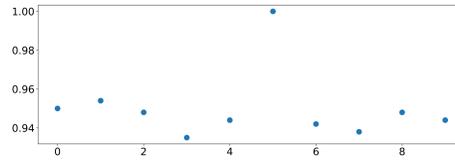


(a) Coverage

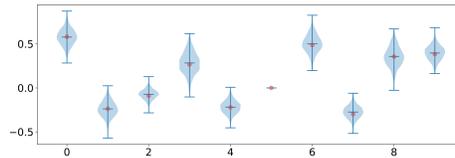


(b) Point estimates

Figure 8: Counterfactual Policy Values for 10 randomly chosen binary policies based on exogenous features: $n = 500$, $n_t = 2$, $n_x = 450$, $s = 2$, $\sigma(\epsilon_t) = 1$, $\sigma(\zeta_t) = .5$, $\sigma(\eta_t) = 1$



(a) Coverage



(b) Point estimates

Figure 9: Counterfactual Policy Values for 10 randomly chosen binary policies based on exogenous features: $n = 2000$, $n_t = 2$, $n_x = 450$, $s = 2$, $\sigma(\epsilon_t) = 1$, $\sigma(\zeta_t) = .5$, $\sigma(\eta_t) = 1$