

Structural Nested Mean Models for Modified Treatment Policies

Zach Shahn¹

¹CUNY School of Public Health, New York, NY, USA

October 8, 2025

Abstract

There is a growing literature on estimating effects of treatment strategies based on the natural treatment that would have been received in the absence of intervention, often dubbed ‘modified treatment policies’ (MTPs). MTPs are sometimes of interest because they are more realistic than interventions setting exposure to an ideal level for all members of a population. In the general time-varying setting, Richardson and Robins [2013] provided exchangeability conditions for nonparametric identification of MTP effects that could be deduced from Single World Intervention Graphs (SWIGs). Díaz et al. [2023] provided multiply robust estimators under these identification assumptions that allow for machine learning nuisance regressions. In this paper, we fill a remaining gap by extending Structural Nested Mean Models (SNMMs) [Robins, 1994a, 2004, Vansteelandt and Joffe, 2014] to MTP settings, which enables characterization of (time-varying) heterogeneity of MTP effects. We do this both under the exchangeability assumptions of Richardson and Robins [2013] and under parallel trends assumptions, which enables investigation of (time-varying heterogeneous) MTP effects in the presence of some unobserved confounding.

1 Introduction

There is a growing literature on estimating effects of treatment strategies based on the natural treatment that would have been received in the absence of intervention [Robins et al., 2004, Richardson and Robins, 2013, Young et al., 2014, Haneuse and Rotnitzky, 2013, Muñoz and Van Der Laan, 2012, Díaz et al., 2023, Sani et al., 2020]. Examples of such strategies include ‘exercise 20 minutes longer than you normally would’ or ‘discharge patients from the intensive care unit one day later than under usual care’. Strategies

depending on the natural value of treatment have been dubbed ‘modified treatment policies’ (MTPs), a term which has gained some traction and we hence adopt here. MTPs are sometimes of interest because they are more realistic than interventions setting exposure to an ideal level for all members of a population.

In the general time-varying setting, Richardson and Robins [2013] provided exchangeability conditions for nonparametric identification of MTP effects that could be deduced from Single World Intervention Graphs (SWIGs). Díaz et al. [2023] provided multiply robust estimators under these identification assumptions that allow for machine learning nuisance regressions. In this paper, we fill a remaining gap by extending Structural Nested Mean Models (SNMMs) [Robins, 1994a, 2004, Vansteelandt and Joffe, 2014] to MTP settings, which enables characterization of (time-varying) heterogeneity of MTP effects. We do this both under the exchangeability assumptions of Richardson and Robins [2013] and under parallel trends assumptions [Shahn et al., 2022], which enables investigation of (time-varying heterogeneous) MTP effects in the presence of some unobserved confounding.

The structure of the paper is as follows. In Section 2, we summarize notation and review definitions of MTPs and SNMMs. Once defined, we also provide some hypothetical examples of when SNMMs for MTP effects might be of interest. In Section 3, we provide identification results and Neyman orthogonal estimators under exchangeability assumptions. In Section 4, we do the same under parallel trends assumptions. In Section 5, we provide simulations showing that our estimators are unbiased and that the sandwich variance estimates provide nominal coverage. In Section 6, we present a real data analysis estimating the effects of shifting mobility from its natural level on subsequent county level Covid-19 incidence. In Section 7, we conclude and discuss some directions for future work.

2 Notation, MTPs, and SNMMs

Suppose we observe a cohort of N subjects indexed by $i \in \{1, \dots, N\}$. Assume that each subject is observed at regular intervals from baseline time 0 through end of follow-up time K , and there is no loss to follow-up. At each time point t , the data are collected on $O_t = (Z_t, Y_t, A_t)$ in that temporal order. A_t denotes the (possibly multidimensional with discrete and/or continuous components) treatment received at time t , Y_t denotes the outcome of interest at time t , and Z_t denotes a vector of covariates at time t excluding Y_t . Hence, Z_0 constitutes the vector of baseline covariates other than Y_0 . For arbitrary time varying variable X : we denote by $\bar{X}_t = (X_0, \dots, X_t)$ the history of X through time t ; we denote by $\underline{X}_t = (X_t, \dots, X_K)$ the future of X from time t through time K ; and whenever the negative index X_{-1}

appears it denotes the null value with probability 1. In Section 3, when we work under exchangeability assumptions, we define $\bar{L}_t$ to be $(\bar{Z}_t, \bar{Y}_t)$, i.e. the joint covariate and outcome history through time t , but we do not require that the outcome is necessarily measured at all time points. In Section 4, when we work under parallel trends assumptions, we define $\bar{L}_t$ to be $(\bar{Z}_t, \bar{Y}_{t-1})$, i.e. the joint covariate and outcome history through time t excluding the most recent outcome Y_t , and we assume that the outcome is measured at each time point. Hence, in the parallel trends section, L_0 is Z_0 . In all sections, we let $H_t = (\bar{L}_t, \bar{A}_{t-1})$ denote the relevant pre-treatment history. We denote random variables by capital letters and their realizations using lower case letters. We adopt the counterfactual framework for time-varying treatments [Robins, 1986] which posits that corresponding to each time-varying treatment regime $\bar{a}_t$, each subject has a counterfactual or potential outcome $Y_{t+1}(\bar{a}_t)$ that would have been observed had that subject received treatment regime $\bar{a}_t$.

Let $g = (g_1, \dots, g_K)$ denote an arbitrary MTP with each $g_t : (h_t, a_t) \rightarrow a_t^+$ a treatment rule setting the modified treatment (a_t^+) as a function of observed history through t . $A_t(g)$ is the treatment that would naturally occur at time t had strategy g been imposed through $t - 1$. $A_t^+(g)$ is the treatment received at time t under g , possibly a function of $A_t(g)$. $Y_t(g) \equiv Y_t(A_{t-1}^+(g))$ is the counterfactual outcome had treatment been assigned according to g . We will use the shorthand $Y_t(\bar{A}_m, \underline{g})$ to denote the counterfactual outcome under the observed regime through time m followed by regime g thereafter.

A general regime SNMM for MTP g models the contrasts

$$\gamma_{tk}^g(h_t, a_t) = E[Y_k(\bar{A}_t, \underline{g}) - Y_k(\bar{A}_{t-1}, \underline{g}) | H_t = h_t, A_t = a_t]. \quad (1)$$

(1) represents the conditional lasting effects among subjects with observed history (h_t, a_t) of receiving the observed treatment at time t then switching to MTP g thereafter, compared to switching to g at time t and continuing to follow it thereafter. Thus, (1) might be interpreted as the conditional effect of one final blip of the observational regime before switching to an MTP.

With knowledge of γ_{tk}^g , we can ‘blip down’ or strip away these effects from observed outcomes to obtain consistent estimates of conditional counterfactual outcomes under the MTP g .

Lemma 1. $E[Y_k - \sum_{j=t}^{k-1} \gamma_j(H_j, A_j) | H_t, A_t] = E[Y_k(\bar{A}_{t-1}, \underline{g}) | H_t, A_t]$

The proof of Lemma 1 is in the Appendix. It follows as a special case of Lemma 1 that given γ_{tk}^g , $E[Y_k(g)]$ is identified as $E[Y_k - \sum_{j=0}^{k-1} \gamma_{jk}(A_j, H_j)]$.

A parametric SNMM specifies a functional form

$$\gamma_{tk}^g(h_t, a_t) = \gamma_{tk}^g(h_t, a_t; \psi^*) \quad (2)$$

with $\gamma_{tk}^g(h_t, a_t; \psi)$ a known function of finite dimensional parameter ψ taking unknown true value ψ^* such that $\gamma_{tk}^g(h_t, a_t; \psi) = 0$ whenever $g(h_t, a_t) = a_t$ or $\psi = 0$.

As one substantive example, suppose A_t is a continuous measure of medication adherence (e.g. proportion of prescribed doses taken), and it is thought that efficacy starts significantly declining below some threshold δ . One might be interested in the effect of a partial adherence enforcing intervention $g(h_t, a_t) = \mathbf{1}\{a_t < \delta\}\delta + \mathbf{1}\{a_t \geq \delta\}a_t$ that sets treatment to δ if its natural value is below that threshold and leaves the natural value unchanged otherwise. In this setting, $\gamma_{tk}(h_t, a_t)$ is the lasting effect of adherence level a_t versus $g(h_t, a_t)$ in month t followed by regime g thereafter. If $a_t \geq \delta$, then the effect would be 0. For $a_t < \delta$, parametric model (2) could characterize how the lasting effects of a blip of non-adherence relative to regime g depend on the magnitude of nonadherence ($\delta - a_t$) and health history.

As another example, this time concerning an exposure more often studied under parallel trends assumptions, suppose A_t denotes state minimum wage in year t . Perhaps interest centers on the effect on poverty of increasing minimum wage by \$2 above its natural value if its natural value is less than \$10. (1) then defines the conditional effects on poverty of natural wages less than \$10 relative to their value had they been increased by \$2. (1) might, for example, reveal that \$2 wage hikes relative to observed wages would have been more impactful in states that actually had lower wages. This is a different estimand than the effect of treatment on the treated (i.e. the effect of wage hikes from the previous year in states that implemented them) that economists typically study with Difference in Differences.

3 Identification and Estimation Under Exchangeability Assumptions

3.1 Point Exposure Setting

In the point exposure setting where $K = 1$, we will drop time subscripts for simplicity. We will make the standard consistency assumption that

$$\textbf{Consistency: } Y(a) = Y \text{ whenever } A = a. \quad (3)$$

We also make a positivity assumption with respect to regime g

$$f_{A,L}(a, l) > 0 \text{ implies } f_{A,L}(g(a, l), l) > 0, \quad (4)$$

which states that for any treatment and covariate values that might be observed in the data, the corresponding g -modified treatment value with the same covariates might also be observed. Finally, Haneuse and Rotnitzky [2013] introduced the MTP exchangeability assumption

$$E[Y(a')|A = a, L = l] = E[Y(a')|A = a', L = l] \text{ for all } a' = g(a, l) \quad (5)$$

and showed that under assumptions (3), (4), and (5), conditional MTP effects are identified by

$$E[Y(g)|A = a, L = l] = E[Y|A = g(a, l), L = l]. \quad (6)$$

Identification of a SMM follows immediately:

$$\gamma^g(l, a) \equiv E[Y - Y(g(l, a))|A = a, L = l] = E[Y|A = a, L = l] - E[Y|A = g(a, l), L = l]. \quad (7)$$

Consider a parametric SMM $\gamma^g(l, a; \psi^*)$ as defined in (2). Then (7) implies that ψ^* satisfies

$$E[q(L, A)(Y - \gamma^g(L, A; \psi^*) - \mu(g(A, L), L))] = 0 \quad (8)$$

where $\mu(a, l) = E[Y|A = a, L = l]$. Thus, a consistent and asymptotically normal estimator of ψ^* can be obtained by solving the corresponding estimating equations

$$\mathbb{P}_N \left[\begin{array}{c} b(A, L)\{Y - \mu(A, L; \beta)\} \\ q(A, L)\{Y - \gamma^g(A, L; \psi) - \mu(g(A, L); \beta)\} \end{array} \right] = 0 \quad (9)$$

where $\mu(A, L; \beta)$ is a model for $\mu(A, L)$ smooth in an r -dimensional parameter β (equal to β^* under the true law) and $q(A, L)$ and $b(A, L)$ are conformable analyst selected index functions.

The estimator solving (9) is very sensitive to misspecification of μ , and asymptotic normality may not hold if, to avoid misspecification, μ is estimated by machine learning. Therefore, we introduce an estimator that has the Neyman-orthogonality property [Chernozhukov et al., 2018] crucial for enabling machine learning estimation of nuisance functions.

Let T_g act on functions h by $(T_g h)(A, L) = h(g(A, L), L)$, and define the L -conditional inner product weighted by the observed treatment density:

$$\langle u, v \rangle = E \left[\int u(L, a) v(L, a) f_{A|L}(a | L) da \right].$$

Let $T_g^\dagger$ be the adjoint to T_g such that

$$\langle q, T_g h \rangle_w = \langle T_g^\dagger q, h \rangle_w.$$

For a shift $g(a) = a + \delta$,

$$(T_g)^\dagger q(L, A) = q(L, g(A)) \frac{f_{A|L}(A - \delta | L)}{f_{A|L}(A | L)}.$$

In particular, for $\tilde{q} := T_g^\dagger q$ we have the identity

$$E[q(A, L) h(g(A, L), L) f_{A|L}(A|L)] = E[\tilde{q}(A, L) h(A, L) f_{A|L}(A|L)] \quad \forall h. \quad (10)$$

Thus, $T_g^\dagger$ is essentially a change of variables operator for g . If $g(\cdot, l)$ is bijective and differentiable, then

$$\tilde{q}(l, a) = q(l, g^{-1}(a, l)) \frac{\pi(g^{-1}(a, l) | l)}{\pi(a | l)} |\det J_a g^{-1}(a, l)|.$$

If $g(\cdot, l)$ is possibly many-to-one and A is discrete,

$$\tilde{q}(l, a) = \sum_{a': g(a', l)=a} q(l, a') \frac{\pi(a' | l)}{\pi(a | l)}.$$

Theorem 1. (*Neyman orthogonal estimator, point exposure*) Assume Consistency, Positivity, MTP Exchangeability, and regularity conditions. Consider the score

$$\phi(O; \psi, \eta) = \underbrace{(q - \tilde{q})(Y - \mu(A, L))}_{\text{augmentation}} + \underbrace{q\{\mu(A, L) - \mu(g(A, L), L) - \gamma_g(L, A; \psi)\}}_{\text{identifying}}, \quad \eta = (\mu, \pi). \quad (11)$$

Then:

1. **Identification.** At (ψ^*, η^*) , $E[\phi(O; \psi^*, \eta^*)] = 0$.

2. **Neyman orthogonality.** For any regular parametric submodel $t \mapsto \eta_t = (\mu_t, \pi_t)$ with $\eta_0 = \eta^*$,

$$\left. \frac{d}{dt} \right|_{t=0} E[\phi(O; \psi^*, \eta_t)] = 0.$$

3. **Asymptotic variance** Assume that $\hat{\eta}$ is learned on held-out folds and $\|\hat{\eta} - \eta^*\|_{L_2} = o_p(1)$. If $\dim(q) = \dim(\psi) = d$ and $\hat{\psi}$ solves $\frac{1}{n} \sum_{i=1}^n \hat{\phi}_i(\psi) = 0$, then for

$$G := E\left[\frac{\partial}{\partial \psi^\top} \phi(O; \psi^*, \eta^*)\right] \in R^{d \times d}$$

and

$$\Sigma := \text{Var}(\phi(O; \psi^*, \eta^*)) \in R^{d \times d},$$

$$\sqrt{n}(\hat{\psi} - \psi^*) \rightsquigarrow \mathcal{N}(0, V), \quad IF(O) = -G^{-1}\phi(O; \psi^*, \eta^*), \quad V = G^{-1}\Sigma(G^{-1})^\top.$$

Proof. See Appendix 8.1. □

Algorithm 1 computes the estimator from Theorem 1. Note that the theorem requires only that $T_g^\dagger$ satisfies the density-weighted adjoint identity (10). The familiar formulas for the settings where g is a continuous bijection or discrete are convenient instances but not required for the theorem to hold.

Algorithm 1 Cross-fitted Neyman–orthogonal estimator for a point-exposure MTP

Input: Data $\{(Y_i, A_i, H_i)\}_{i=1}^n$; basis $s(H)$; shift δ ; folds K

Output: $\hat{\psi}$ and IF-sandwich covariance $\hat{V}$

```

1: Partition indices into  $K$  folds  $(\mathcal{T}^{(k)}, \mathcal{I}^{(k)})$ 
2: Initialize  $M \leftarrow 0, b \leftarrow 0$ 
3: for  $k = 1, \dots, K$  do ▷ Nuisance fits on train fold
4:   Fit  $\hat{\mu}$  on  $\{(Y_j, A_j, H_j)\}_{j \in \mathcal{T}^{(k)}}$  for  $y \approx E[Y \mid A, H]$ 
5:   Fit  $\hat{m}$  on  $\{(A_j, H_j)\}_{j \in \mathcal{T}^{(k)}}$  for  $a \approx E[A \mid H]$ 
6:   Estimate  $\hat{\sigma}$  as SD of residuals  $A - \hat{m}(H)$  on  $\mathcal{T}^{(k)}$  (if using Normal ratio)
7:   for all  $i \in \mathcal{I}^{(k)}$  do
8:      $e_i \leftarrow Y_i - \hat{\mu}(A_i, H_i); \mu_i^g \leftarrow \hat{\mu}(A_i + \delta, H_i)$ 
9:      $s_i \leftarrow s(H_i); q_i \leftarrow s_i$ 
10:    if Normal working residual for  $A \mid H$  then
11:       $\hat{r}_i \leftarrow \frac{\phi(A_i - \delta; \hat{m}(H_i), \hat{\sigma})}{\phi(A_i; \hat{m}(H_i), \hat{\sigma})}$ 
12:    else
13:      Estimate  $\hat{r}_i \approx \frac{\hat{f}_{A|H}(A_i - \delta \mid H_i)}{\hat{f}_{A|H}(A_i \mid H_i)}$ 
14:       $\tilde{q}_i \leftarrow q_i \cdot \hat{r}_i$  ▷ Accumulate normal equations  $M\hat{\psi} = b$ 
15:       $M \leftarrow M + \delta s_i s_i^\top$ 
16:       $b \leftarrow b - [s_i \{\hat{\mu}(A_i, H_i) - \mu_i^g\} + (q_i - \tilde{q}_i) e_i]$ 
17: Solve and variance:
18:  $\hat{\psi} \leftarrow (M + \lambda I)^{-1} b$  ▷ optional small ridge  $\lambda \geq 0$ 
19:  $G \leftarrow (\delta/n) \sum_{i=1}^n s_i s_i^\top$ 
20:  $\phi_i \leftarrow (q_i - \tilde{q}_i) e_i + s_i \{\hat{\mu}(A_i, H_i) - \hat{\mu}(A_i + \delta, H_i) + \delta s_i^\top \hat{\psi}\}$ 
21:  $\text{IF}_i \leftarrow G^{-1} \phi_i; \hat{V} \leftarrow \text{Var}_n(\text{IF}_i)/n$ 

```

3.2 Time varying treatments

For didactic purposes, consider identification in the two time point case. We extend the results by induction to the general time varying case in the Appendix. We know that the one time step ahead effects $\gamma_{12}^g(a_1, h_1)$ and $\gamma_{01}^g(a_0, l_0)$ are identified from the point exposure case. So we will focus on identification of $\gamma_{02}^g(a_0, l_0)$.

$$\gamma_{02}^g(a_0, l_0) = E[Y_2(a_0, g_1) - Y_2(g_0, g_1) \mid L_0 = l_0, A_0 = a_0]$$

By iterated expectations and Lemma 1,

$$\begin{aligned}
& E[Y_2(a_0, g_1)|A_0 = a_0, L_0 = l_0] \\
&= E[E[Y_2(a_0, g_1)|A_0 = a_0, L_0 = l_0, L_1, A_1]|A_0 = a_0, L_0 = l_0] \\
&= E[E[Y_2 - \gamma_{12}^g(A_1, H_1)|A_0 = a_0, L_0 = l_0, L_1, A_1]|A_0 = a_0, L_0 = l_0] \\
&= E[Y_2 - \gamma_{12}^g(A_1, H_1)|A_0 = a_0, L_0 = l_0]
\end{aligned}$$

Thus, the first term of $\gamma_{02}^g(a_0, l_0)$ is identified. The second term is identified as

$$\begin{aligned}
& E[Y_2(g_0, g_1)|L_0 = l_0, A_0 = a_0] = \\
& \sum_{l_1, a_1} E[Y_2|L_0 = l_0, A_0 = g(l_0, a_0), L_1 = l_1, A_1 = g(l_0, g(l_0, a_0), a_1)] \\
& p(l_1|L_0 = l_0, A_0 = g(l_0, a_0))p(a_1|L_0 = l_0, L_1 = l_1, A_0 = g(l_0, a_0)) \\
& \equiv \mu_{02}(l_0, g(l_0, a_0))
\end{aligned}$$

by the extended g-formula under the assumptions of Section 5 of [Richardson and Robins, 2013]. Following Díaz et al. [2023], $\mu_{02}(l_0, a_0)$ can be represented in terms of iterative regressions by

$$\mu_{02}(l_0, a_0) = E[\mu_{12}(l_0, L_1, g(l_0, a_0), g(l_0, L_1, g(l_0, a_0), A_1))|L_0 = l_0, A_0 = a_0] \quad (12)$$

with $\mu_{12}(l_0, l_1, a_0, a_1) = E[Y_2|L_0 = l_0, L_1 = l_1, A_0 = a_0, A_1 = a_1]$. Then for parametric model $(\gamma_{01}^g(l_0, a_0), \gamma_{12}^g(h_1, a_1), \gamma_{02}^g(l_0, a_0)) = (\gamma_{01}^g(l_0, a_0; \psi^*), \gamma_{12}^g(h_1, a_1; \psi^*), \gamma_{02}^g(l_0, a_0; \psi^*))$, parameter ψ^* satisfies

$$E \left[\begin{array}{c} q_{01}(L_0, A_0)\{Y_1 - \gamma_{01}^g(L_0, A_0; \psi^*) - \mu_{01}(L_0, g(L_0, A_0))\} \\ q_{12}(H_1, A_1)\{Y_1 - \gamma_{12}^g(H_1, A_1; \psi^*) - \mu_{12}(H_1, g(H_1, A_1))\} \\ q_{02}(L_0, A_0)\{Y_1 - \gamma_{12}^g(H_1, A_1; \psi^*) - \gamma_{02}^g(L_0, A_0; \psi^*) - \mu_{02}(L_0, g(L_0, A_0))\} \end{array} \right] = 0, \quad (13)$$

which again leads to natural estimating equations. In particular, a consistent and asymptotically normal estimator of ψ^* can be obtained by solving

$$\mathbb{P}_N \begin{bmatrix} b_{01}(L_0, A_0)\{Y_1 - \mu_{01}(L_0, A_0; \beta)\} \\ b_{12}(\bar{L}_1, \bar{A}_1)\{Y_2 - \mu_{12}(\bar{L}_1, \bar{A}_1; \beta)\} \\ b_{02}(L_0, A_0)\{\mu_{12}(L_0, L_1, g(L_0, A_0), g(L_0, L_1, g(L_0, A_0), A_1)) - \mu_{02}(L_0, A_0; \beta)\} \\ q_{01}(L_0, A_0)\{Y_1 - \gamma_{01}^g(L_0, A_0; \psi) - \mu_{01}(L_0, g(L_0, A_0))\} \\ q_{12}(H_1, A_1)\{Y_1 - \gamma_{12}^g(H_1, A_1; \psi) - \mu_{12}(H_1, g(H_1, A_1))\} \\ q_{02}(L_0, A_0)\{Y_1 - \gamma_{12}^g(H_1, A_1; \psi) - \gamma_{02}^g(L_0, A_0; \psi) - \mu_{02}(L_0, g(L_0, A_0))\} \end{bmatrix} = 0 \quad (14)$$

where $\mu_{mk}(\bar{A}_m, \bar{L}_m; \beta)$ are models for $\mu_{mk}(\bar{A}_m, \bar{L}_m)$ smooth in an r -dimensional parameter β (equal to β^* under the true law) and $q_{mk}(\bar{A}_m, \bar{L}_m)$ and $b_{mk}(\bar{A}_m, \bar{L}_m)$ are conformable analyst selected index functions.

As in the point exposure setting, a Neyman orthogonal estimator would be desirable to enable machine learning estimation of nuisance functions. We now develop this estimator for the general time-varying setting. Define the one step ahead conditional mean

$$\mu_t(a, h) \equiv E[V_{t+1} | A_t = a, H_t = h]$$

where we set $V_T = Y$ and recursively set for $t = T - 1, \dots, 0$

$$V_t = \mu_t(g(H_t, A_t), H_t).$$

Now, for each time t , define the operator $T_{gt}h \equiv h(H_t, g(H_t, A_t))$. And let $T_{gt}^\dagger$ be its adjoint with respect to

$$\langle u, v \rangle_{w_t} \equiv E[u(H_t, A_t)v(H_t, A_t)w_t]$$

for $w_t = \pi(A_t | H_t)$. As in the point exposure setting, let $\tilde{q}_t \equiv T_{gt}^\dagger q_t$.

Theorem 2 (Neyman orthogonal estimator, time-varying). *Assume Sequential consistency, positivity, and MTP exchangeability (i.e. the assumptions of Section 5 of Richardson and Robins [2013]) hold. Consider the score*

$$\Phi(O; \psi, \eta) = \sum_{t=0}^{T-1} \phi_t(O; \psi, \eta),$$

with

$$\phi_t = (q_t - \tilde{q}_t) \{V_{t+1} - \mu_t(H_t, A_t)\} + q_t \left\{ \mu_t(H_t, A_t) - \mu_t(H_t, g_t(A_t, H_t)) - \gamma_t(H_t, A_t; \psi) \right\},$$

where $\eta = \{\mu_t, \pi_t, \tilde{q}_t : t = 0, \dots, T-1\}$.

1. **Identification** At the truth (ψ^*, η^*) , $\mathbb{E}[\Phi(O; \psi^*, \eta^*)] = 0$
2. **Neyman orthogonality** The score is Neyman-orthogonal: the pathwise derivative of $\mathbb{E}[\Phi(O; \psi^*, \eta)]$ with respect to each nuisance μ_t and π_t vanishes at η^* .
3. **Asymptotic variance** Suppose ψ is estimated by the root $\hat{\psi}$ of the cross-fitted sample moment $n^{-1} \sum_{i=1}^n \hat{\phi}_i(\psi) = 0$, with nuisances learned on held-out folds and consistent in L_2 . Let

$$G := E \left[\frac{\partial}{\partial \psi^\top} \phi(O; \psi^*, \eta^*) \right], \quad \Sigma := \text{Var}(\phi(O; \psi^*, \eta^*)).$$

Then under standard regularity conditions,

$$\sqrt{n}(\hat{\psi} - \psi^*) \xrightarrow{d} \mathcal{N}(0, V), \quad \text{with IF } IF(O) = -G^{-1}\phi(O; \psi^*, \eta^*),$$

and

$$V = G^{-1} \Sigma (G^{-1})^\top.$$

Proof. See Appendix 8.2

□

Algorithm 2 implements the estimator from Theorem 2.

Algorithm 2 Cross-fitted Neyman-orthogonal estimator for longitudinal MTPs

Input: Data $\{(Y_i, \{A_{t,i}, H_{t,i}\}_{t=0}^{T-1})\}_{i=1}^n$; bases $s_t(H_t)$; shifts $\{\delta_t\}$; folds K

Output: $\hat{\psi} = (\hat{\psi}_0^\top, \dots, \hat{\psi}_{T-1}^\top)^\top$ and IF-sandwich $\hat{V}$

```

1: Initialize  $V_{T,i} \leftarrow Y_i$  for all  $i$ ;  $M_t \leftarrow 0, b_t \leftarrow 0$  for  $t = 0, \dots, T-1$ 
2: Partition indices into  $K$  folds  $(\mathcal{T}^{(k)}, \mathcal{I}^{(k)})$ 
3: for  $t = T-1, \dots, 0$  do
4:   for  $k = 1, \dots, K$  do ▷ Nuisance fits on train fold at time  $t$ 
5:     Fit  $\hat{\mu}_t$  on  $\{(V_{t+1,j}, A_{t,j}, H_{t,j})\}_{j \in \mathcal{T}^{(k)}}$  for  $v \approx E[V_{t+1} \mid A_t, H_t]$ 
6:     Fit  $\hat{m}_t$  on  $\{(A_{t,j}, H_{t,j})\}_{j \in \mathcal{T}^{(k)}}$  for  $a \approx E[A_t \mid H_t]$ 
7:     Estimate  $\hat{\sigma}_t$  as SD of  $A_t - \hat{m}_t(H_t)$  on  $\mathcal{T}^{(k)}$  (if using Normal ratio)
8:     for all  $i \in \mathcal{I}^{(k)}$  do
9:        $e_{t,i} \leftarrow V_{t+1,i} - \hat{\mu}_t(A_{t,i}, H_{t,i}); \mu_{t,i}^g \leftarrow \hat{\mu}_t(A_{t,i} + \delta_t, H_{t,i})$ 
10:       $s_{t,i} \leftarrow s_t(H_{t,i}); q_{t,i} \leftarrow s_{t,i}$ 
11:      if Normal working residual then
12:         $\hat{r}_{t,i} \leftarrow \frac{\phi(A_{t,i} - \delta_t; \hat{m}_t(H_{t,i}), \hat{\sigma}_t)}{\phi(A_{t,i}; \hat{m}_t(H_{t,i}), \hat{\sigma}_t)}$ 
13:      else
14:        Estimate  $\hat{r}_{t,i} \approx \frac{\hat{f}_{A_t|H_t}(A_{t,i} - \delta_t \mid H_{t,i})}{\hat{f}_{A_t|H_t}(A_{t,i} \mid H_{t,i})}$ 
15:       $\tilde{q}_{t,i} \leftarrow q_{t,i} \cdot \hat{r}_{t,i}$  ▷ Accumulate time- $t$  normal equations
16:       $M_t \leftarrow M_t + \delta_t s_{t,i} s_{t,i}^\top$ 
17:       $b_t \leftarrow b_t - [s_{t,i} \{\hat{\mu}_t(A_{t,i}, H_{t,i}) - \mu_{t,i}^g\} + (q_{t,i} - \tilde{q}_{t,i}) e_{t,i}]$ 
18:       $V_{t,i} \leftarrow \mu_{t,i}^g$  ▷ backward recursion target
19: Solve and variance:
20: for  $t = 0, \dots, T-1$  do
21:    $\hat{\psi}_t \leftarrow (M_t + \lambda_t I)^{-1} b_t$ 
22:    $G \leftarrow \text{diag}((\delta_0/n) \sum_i s_{0,i} s_{0,i}^\top, \dots, (\delta_{T-1}/n) \sum_i s_{T-1,i} s_{T-1,i}^\top)$ 
23:   For each  $i$ :  $\phi_{t,i} \leftarrow (q_{t,i} - \tilde{q}_{t,i}) e_{t,i} + s_{t,i} \{\hat{\mu}_t(A_{t,i}, H_{t,i}) - \hat{\mu}_t(A_{t,i} + \delta_t, H_{t,i}) + \delta_t s_{t,i}^\top \hat{\psi}_t\}$ 
24:   Stack  $\phi_i \leftarrow (\phi_{0,i}^\top, \dots, \phi_{T-1,i}^\top)^\top$ ; IF  $i \leftarrow G^{-1} \phi_i$ ;  $\hat{V} \leftarrow \text{Var}_n(\text{IF}_i)/n$ 

```

4 Identification and Estimation Under Parallel Trends Assumptions

Recall that in this section $\bar{L}_t$ excludes the most recent outcome Y_t . In the parallel trends setting, we present identification results and the natural accompanying parametric outcome regression based estimators. We leave derivation of Neyman orthogonal estimators under parallel trends for future work.

4.1 Point Exposure Setting

Again, we suppress time subscripts for treatment and the baseline covariate to reduce notational clutter in the point exposure setting. We propose the MTP parallel trends assumption

$$E[Y_1(g(a, l)) - Y_0 | A = a, L = l] = E[Y_1(g(a, l)) - Y_0 | A = g(a, l), L = l] \text{ for all } (a, l) \text{ with } p(a, l) > 0. \quad (15)$$

In words, the conditional expected counterfactual trend in the outcome under the MTP does not depend on whether treatment takes its natural or modified value. Under (15), $\gamma^g(l, a)$ is identified as follows:

$$\begin{aligned} \gamma^g(l, a) &= E[Y_1 - Y_1(g(a, l)) | A = a, L = l] \\ &= E[Y_1 - Y_0 | A = a, L = l] - E[Y_1(g(a, l)) - Y_0 | A = a, L = l] \\ &= E[Y_1 - Y_0 | A = a, L = l] - E[Y_1(g(a, l)) - Y_0 | A = g(a, l), L = l] \\ &= E[Y_1 - Y_0 | A = a, L = l] - E[Y_1 - Y_0 | A = g(a, l), L = l] \end{aligned}$$

Thus, MTPs can be studied under parallel trends assumptions. Furthermore, ψ^* from (2) satisfies

$$E[q(L, A)(Y_1 - Y_0 - \gamma^g(L, A; \psi^*) - \mu^d(L, g(A, L)))] = 0, \quad (16)$$

with $\mu^d(l, a) = E[Y_1 - Y_0 | L = l, A = a]$, which again suggests a natural estimation procedure. A consistent and asymptotically normal estimator of ψ^* can be obtained by solving the corresponding estimating equations

$$\mathbb{P}_N \left[\begin{array}{c} b(A, L) \{Y_1 - Y_0 - \mu^d(L, A; \beta)\} \\ q(A, L) \{Y_1 - Y_0 - \gamma^g(A, L; \psi) - \mu^d(L, A; \beta)\} \end{array} \right] = 0 \quad (17)$$

where $\mu^d(L, A; \beta)$ is a model for $\mu^d(L, A)$ smooth in an r -dimensional parameter β (equal to β^* under the true law) and $q(A, L)$ and $b(A, L)$ are conformable analyst selected index functions.

4.2 Time-Varying Treatments

In the time-varying setting, we make the parallel trends assumptions

$$E[Y_k(g) - Y_{k-1}(g) | A_m = a_m, H_m = h_m] = E[Y_k(g) - Y_{k-1}(g) | A_m = g(h_m, a_m), H_m = h_m] \text{ for all } h_m, a_m, k > m. \quad (18)$$

In words, conditional expected future counterfactual trends under g from time m onwards do not depend on whether treatment at time m took its natural or modified value. For the two time point setting, these assumptions can be enumerated

$$\begin{aligned} E[Y_2(g) - Y_1(g(a_0, l_0)) | A_0 = a_0, L_0 = l_0] &= E[Y_2(g) - Y_1(g(l_0, a_0)) | A_0 = g(l_0, a_0), L_0 = l_0] \\ E[Y_1(g(l_0, a_0)) - Y_0 | A_0 = a_0, L_0 = l_0] &= E[Y_1(g(l_0, a_0)) - Y_0 | A_0 = g(l_0, a_0), L_0 = l_0] \\ E[Y_2(g) - Y_1 | A_1 = a_1, H_1 = h_1] &= E[Y_2(g) - Y_1 | A_1 = g(h_1, a_1), H_1 = h_1]. \end{aligned} \quad (19)$$

We know that the one time step ahead effects $\gamma_{12}^g(a_1, h_1)$ and $\gamma_{01}^g(a_0, l_0)$ are identified from the point exposure case. Identification of $\gamma_{02}^g(a_0, l_0)$ can be demonstrated as follows.

$$\begin{aligned} \gamma_{02}(a_0, l_0) &= E[Y_2(a_0, g) - Y_2(g) | A_0 = a_0, L_0 = l_0] \\ &= E[(Y_2(a_0, g) - Y_1(g)) - (Y_2(g) - Y_1(g)) | A_0 = a_0, L_0 = l_0] \end{aligned}$$

The first term of the conditional expectation on the right hand side is identified as

$$E[Y_2(a_0, g) - Y_1(g) | A_0 = a_0, L_0 = l_0] = E[Y_2 - \gamma_{12}(l_0, L_1, a_0, A_1) - (Y_1 - \gamma_{01}(l_0, a_0)) | A_0 = a_0, L_0 = l_0].$$

The second term is identified as

$$\begin{aligned} E[Y_2(g) - Y_1(g) | A_0 = a_0, L_0 = l_0] &= E[Y_2(g) - Y_1(g) | A_0 = g(a_0, l_0), L_0 = l_0] \\ &= \int_{l_1, a_1} E[Y_2(g) - Y_1 | A_0 = g(a_0, l_0), A_1 = a_1, L_0 = l_0, L_1 = l_1] p(l_1, a_1 | A_0 = g(a_0, l_0), L_0 = l_0) da_1 dl_1 \\ &= \int_{l_1, a_1} E[Y_2 - Y_1 | A_0 = g(a_0, l_0), A_1 = g(l_0, l_1, g(l_0, a_0), a_1), L_0 = l_0, L_1 = l_1] p(l_1, a_1 | A_0 = g(a_0, l_0), L_0 = l_0) da_1 dl_1 \\ &= E[\mu_{12}^d(l_0, L_1, g(a_0, l_0), g(l_0, L_1, g(l_0, a_0), A_1)) | A_0 = g(l_0, a_0), L_0 = l_0] \\ &\equiv \mu_{02}^d(g(l_0, a_0), l_0) \end{aligned}$$

where $\mu_{12}^d(l_0, l_1, a_0, a_1) = E[Y_2 - Y_1 | L_0 = l_0, L_1 = l_1, A_0 = a_0, A_1 = a_1]$. The derivation is simply repeated applications of iterated expectations, parallel trends, and consistency. It follows that for parametric model $(\gamma_{01}^g(l_0, a_0), \gamma_{12}^g(h_1, a_1), \gamma_{02}^g(l_0, a_0)) = (\gamma_{01}^g(l_0, a_0; \psi^*), \gamma_{12}^g(h_1, a_1; \psi^*), \gamma_{02}^g(l_0, a_0; \psi^*))$, parameter ψ^* satisfies

$$E \begin{bmatrix} q_{01}(L_0, A_0) \{Y_1 - Y_0 - \gamma_{01}^g(L_0, A_0; \psi^*) - \mu_{01}^d(L_0, g(L_0, A_0))\} \\ q_{12}(H_1, A_1) \{Y_2 - Y_1 - \gamma_{12}^g(H_1, A_1; \psi^*) - \mu_{12}^d(H_1, g(H_1, A_1))\} \\ q_{02}(L_0, A_0) \{Y_2 - Y_1 - \gamma_{12}^g(H_1, A_1; \psi^*) - \gamma_{02}^g(L_0, A_0; \psi^*) + \gamma_{01}^g(L_0, A_0; \psi^*) - \mu_{02}^d(L_0, g(L_0, A_0))\} \end{bmatrix} = 0, \quad (20)$$

which again leads to natural estimating equations. In particular, a consistent and asymptotically normal estimator of ψ^* can be obtained by solving

$$\mathbb{P}_N \begin{bmatrix} b_{01}(L_0, A_0) \{Y_1 - Y_0 - \mu_{01}^d(L_0, A_0; \beta)\} \\ b_{12}(\bar{L}_1, \bar{A}_1) \{Y_2 - Y_1 - \mu_{12}^d(\bar{L}_1, \bar{A}_1; \beta)\} \\ b_{02}(L_0, A_0) \{\mu_{12}^d(L_0, L_1, g(L_0, A_0), g(L_0, L_1, g(L_0, A_0), A_1)) - \mu_{02}^d(L_0, A_0; \beta)\} \\ q_{01}(L_0, A_0) \{Y_1 - Y_0 - \gamma_{01}^g(L_0, A_0; \psi) - \mu_{01}^d(L_0, g(L_0, A_0))\} \\ q_{12}(H_1, A_1) \{Y_2 - Y_1 - \gamma_{12}^g(H_1, A_1; \psi) - \mu_{12}^d(H_1, g(H_1, A_1))\} \\ q_{02}(L_0, A_0) \{Y_2 - Y_1 - \gamma_{12}^g(H_1, A_1; \psi) - \gamma_{02}^g(L_0, A_0; \psi) + \gamma_{01}^g(L_0, A_0; \psi) - \mu_{02}^d(L_0, g(L_0, A_0))\} \end{bmatrix} = 0 \quad (21)$$

where $\mu_{mk}^d(\bar{A}_m, \bar{L}_m; \beta)$ are models for $\mu_{mk}^d(\bar{A}_m, \bar{L}_m)$ smooth in an r -dimensional parameter β (equal to β^* under the true law) and $q_{mk}(\bar{A}_m, \bar{L}_m)$ and $b_{mk}(\bar{A}_m, \bar{L}_m)$ are conformable analyst selected index functions.

4.3 Neyman orthogonal estimators

To derive Neyman orthogonal estimators enabling machine learning estimation of nuisance functions, we proceed completely analogously to the (sequential) exchangeability case. In fact, we can simply substitute the difference $\Delta Y_t \equiv Y_t - Y_{t-1}$ for Y_t and the regressions $\mu_{mk}^d(H_m, A_m) = E[\Delta Y_k | H_m, A_m]$ for $\mu_{mk}(H_m, A_m)$ in the scores from Theorems 1 and 2. The proofs of those theorems then go through exactly as before. Here, we state the theorem for the general longitudinal parallel trends case.

Theorem 3 (Neyman–orthogonal estimator under parallel trends, longitudinal). *Set*

$$\mu_t^d(h_t, a_t) := \mathbb{E}[\Delta Y_t \mid H_t = h_t, A_t = a_t].$$

For each t , equip functions with $\langle u, v \rangle_t := \mathbb{E}[\int u(H_t, a_t)v(H_t, a_t)f_{A_t|H_t}(a_t \mid H_t) da_t]$. Let T_{g_t} act by $(T_{g_t}q_t)(h_t, a_t) = q_t(h_t, g_t(h_t, a_t))$ with adjoint $T_{g_t}^\dagger$ under $\langle \cdot, \cdot \rangle_t$. Choose $q_t(H_t, A_t) := s_t(H_t)$ and set $\tilde{q}_t := T_{g_t}^\dagger q_t$. Let $\Gamma_t(H_t; \psi_t)$ be linear in $s_t(H_t)$.

Assume for each t : $\mathbb{E}[\Delta Y_t(0) \mid H_t, A_t] = \mathbb{E}[\Delta Y_t(0) \mid H_t]$. Further assume sequential consistency, positivity, and regularity conditions.

Define the time- t score

$$\phi_t = (q_t - \tilde{q}_t)\{\Delta Y_t - \mu_t^d(H_t, A_t)\} + s_t(H_t)\{\mu_t^d(H_t, A_t) - \mu_t^d(H_t, g_t(H_t, A_t)) + \Gamma_t(H_t; \psi_t)\}.$$

Let $\phi = (\phi_0^\top, \dots, \phi_{T-1}^\top)^\top$ and $\psi = (\psi_0^\top, \dots, \psi_{T-1}^\top)^\top$.

Then:

- (i) **Identification.** For each t , there is a unique ψ_t^* s.t. $\mathbb{E}[s_t\{\mu_t^d(H_t, A_t) - \mu_t^d(H_t, g_t(H_t, A_t)) + \Gamma_t(H_t; \psi_t)\}] = 0$, and $\mathbb{E}[\phi_t(\psi_t^*, \eta_t^*)] = 0$. Hence $\mathbb{E}[\phi(\psi^*, \eta^*)] = 0$.
- (ii) **Neyman orthogonality.** For any regular paths $\eta_{t,\varepsilon} = (\mu_{t,\varepsilon}^d, f_{t,\varepsilon})$ through η_t^* , $\frac{d}{d\varepsilon}\mathbb{E}[\phi_t(\psi_t^*, \eta_{t,\varepsilon})]\big|_{\varepsilon=0} = 0$ for each t .
- (iii) **Asymptotic normality.** Suppose ψ is estimated by the root $\hat{\psi}$ of the cross-fitted sample moment $n^{-1}\sum_{i=1}^n \hat{\phi}_i(\psi) = 0$, with nuisances learned on held-out folds and consistent in L_2 . Let

$$G := E\left[\frac{\partial}{\partial \psi^\top}\phi(O; \psi^*, \eta^*)\right], \quad \Sigma := \text{Var}(\phi(O; \psi^*, \eta^*)).$$

Then under standard regularity conditions,

$$\sqrt{n}(\hat{\psi} - \psi^*) \xrightarrow{d} \mathcal{N}(0, V), \quad \text{with IF } IF(O) = -G^{-1}\phi(O; \psi^*, \eta^*),$$

and

$$V = G^{-1}\Sigma(G^{-1})^\top.$$

5 Simulations

5.1 Point exposure

We simulated data according to the data generating process:

$$\begin{aligned} L &\sim N(0, 1) \\ A|L &\sim N(m(L), 1) \text{ with } m(L) = \theta_0 + \theta_1 L + \theta_2 L^2 \\ Y &= \xi_0 + \xi_1 L + \xi_2 L^2 + (\beta_0 + \beta_1 L)A + \epsilon; \quad \epsilon \sim N(0, 1) \end{aligned}$$

We then sought to estimate the heterogeneous effects of the shift MTP $g(A, L) = A + \delta$. Are interested in the SNMM

$$\gamma(A, L) = E[Y|A, L] - E[Y|A + \delta, L] = -\delta(\beta_0 + \beta_1)L$$

Thus, we specify the correct parametric SNMM

$$\gamma(A, L; \psi) = -\delta s(L)^T \psi$$

with $s(L) = (1, L)$ and $\psi = (\beta_0, \beta_1)$. We then define

$$\tilde{q}(a, l) = q(a - \delta, l) \frac{p(a - \delta|l)}{p(a|l)}$$

and choose $q(A, L) = s(L)$. We set true parameter values to be:

$$\theta = (0.2, 0.8, -0.4); \quad \xi = (0.3, 0.5, 0.2); \quad \beta = (0.8, -0.6).$$

We then estimated the effect of a shift by $\delta = 0.5$ by solving the estimating equations of Theorem 1 with cross-fitting for data sets of sizes 400, 1000, and 3000. We used correctly specified parametric nuisance models. Table 1 shows that the estimator is unbiased and that sandwich confidence intervals had nominal coverage for each component of ψ .

5.2 Time-varying treatment

We consider a two-time-point longitudinal setting with baseline covariate L_0 , treatments $A_0, A_1 \in \mathbb{R}$, intermediate covariate L_1 , and outcome $Y \in \mathbb{R}$. The modified treatment policy (MTP) shifts each

Table 1: Monte Carlo performance — point-exposure continuous-shift orthogonal estimator (stable score; $R = 200$).

n	ψ_0					ψ_1				
	Bias	RMSE	EmpSD	Mean SE	Cov95	Bias	RMSE	EmpSD	Mean SE	Cov95
400	0.002	0.056	0.056	0.057	0.96	0.004	0.065	0.065	0.059	0.93
1000	0.001	0.037	0.037	0.035	0.93	0.002	0.037	0.037	0.035	0.93
3000	0.000	0.020	0.020	0.020	0.94	-0.002	0.020	0.019	0.020	0.96

Notes: Bias = mean of $\hat{\psi} - \psi^*$; EmpSD = empirical SD of $\hat{\psi}$ across replications; Mean SE = average sandwich SE; Cov95 = empirical coverage of 95% Wald CIs from sandwich SEs.

treatment by a constant

$$g_0(a_0) = a_0 + \delta_0, \quad g_1(a_1) = a_1 + \delta_1,$$

with fixed $\delta_0, \delta_1 > 0$. The blip functions relative to $g = (g_0, g_1)$ are parameterized by low-dimensional bases

$$s_0(H_0) = (1, L_0)^\top, \quad s_1(H_1) = (1, L_1)^\top,$$

and true parameters $\psi_0 = (\psi_{0,0}, \psi_{0,1})^\top$, $\psi_1 = (\psi_{1,0}, \psi_{1,1})^\top$, so that

$$\gamma_0^g(H_0, A_0) = -\delta_0 s_0(H_0)^\top \psi_0, \quad \gamma_1^g(H_1, A_1) = -\delta_1 s_1(H_1)^\top \psi_1.$$

Observed data-generating process (DGP). We calibrated our generating process such that we could derive the true parameter values for our blip model. Let n i.i.d. observations be generated as follows; all noises are mutually independent and independent of past history.

$$L_0 \sim \mathcal{N}(0, 1),$$

$$A_0 \mid L_0 \sim \mathcal{N}(0.4 L_0, \sigma_{A_0}^2),$$

$$L_1 \mid (L_0, A_0) := \rho_0 + \rho_1 L_0 + \rho_2 A_0 + \nu, \quad \nu \sim \mathcal{N}(0, \sigma_{L_1}^2),$$

$$A_1 \mid (L_0, L_1, A_0) \sim \mathcal{N}(\kappa_0 + \kappa_2 L_0, \sigma_{A_1}^2),$$

$$\mu_1(H_1, a_1) := b_1(L_0, L_1) + \{\psi_{1,0} + \psi_{1,1} L_1\} a_1,$$

$$Y \mid (H_1, A_1) := \mu_1(H_1, A_1) + \varepsilon, \quad \varepsilon \sim \mathcal{N}(0, \sigma_Y^2),$$

where $b_1(L_0, L_1) = \beta_{10} + \beta_{1L_1} L_1 + \beta_{1L_0} L_0$ is a baseline outcome component. With this construction,

$$\mu_1(H_1, a_1) - \mu_1(H_1, a_1 + \delta_1) = -\delta_1 \{\psi_{1,0} + \psi_{1,1} L_1\} = \gamma_1^g(H_1, A_1)$$

holds by design for any H_1 .

To ensure the time-0 blip identity also holds,

$$\mu_0(H_0, a_0) - \mu_0(H_0, a_0 + \delta_0) = -\delta_0 s_0(H_0)^\top \psi_0 = \gamma_0^g(H_0, A_0),$$

we calibrate ψ_0 so that the slope in a_0 of

$$\mathbb{E}[\mu_1(H_1, A_1 + \delta_1) \mid H_0, A_0 = a_0]$$

equals $s_0(H_0)^\top \psi_0$ for all L_0 . Under the simple A_1 law above (mean depending on L_0 but not on A_0 or L_1), a short calculation yields

$$\psi_{0,0} = \rho_2 \beta_{1L_1} + \rho_2 \psi_{1,1} (\kappa_0 + \delta_1), \quad \psi_{0,1} = \rho_2 \psi_{1,1} \kappa_2$$

so that $\mu_0(H_0, a_0) := \mathbb{E}[\mu_1(H_1, A_1 + \delta_1) \mid H_0, A_0 = a_0]$ is linear in a_0 with slope $s_0(H_0)^\top \psi_0$.

Default parameter values. In all experiments we use

$$\delta_0 = 0.4, \quad \delta_1 = 0.5, \quad \psi_1 = (0.5, 0.3), \quad (\rho_0, \rho_1, \rho_2) = (0.1, 0.6, 0.8),$$

$$(\kappa_0, \kappa_2) = (0.2, 0.35), \quad (\beta_{10}, \beta_{1L_1}, \beta_{1L_0}) = (0.25, 0.5, 0.2),$$

and standard deviations $(\sigma_{A_0}, \sigma_{L_1}, \sigma_{A_1}, \sigma_Y) = (1.0, 0.5, 1.0, 1.0)$. The calibration then implies the *true* time-0 blip coefficients

$$\psi_{0,0} = 0.8 \times 0.5 + 0.8 \times 0.3 \times (0.2 + 0.5) = 0.568, \quad \psi_{0,1} = 0.8 \times 0.3 \times 0.35 = 0.084.$$

Thus the overall target parameter vector is

$$\psi^\star = (\psi_{0,0}, \psi_{0,1}, \psi_{1,0}, \psi_{1,1})^\top = (0.568, 0.084, 0.5, 0.3)^\top.$$

Results Table 2 displays the bias, bootstrap coverage, and empirical standard errors over 500 data sets of size $n=1,000$. We observe very low bias and near nominal coverage.

Table 2: Monte Carlo (R=500) for longitudinal MTP estimator with continuous shifts; bootstrap 95% coverage reported.

n	$\psi_{0,0}$			$\psi_{0,1}$			$\psi_{1,0}$			$\psi_{1,1}$		
	Bias	RMSE	EmpSD	Bias	RMSE	EmpSD	Bias	RMSE	EmpSD	Bias	RMSE	EmpSD
1000	0.000	0.049	0.049	-0.002	0.039	0.039	0.000	0.033	0.033	0.002	0.026	0.026
<i>Boot 95% cov.</i>		<i>0.93</i>			<i>0.94</i>			<i>0.96</i>			<i>0.96</i>	

6 Real Data Application: Effect of Shifting Workplace Mobility on Covid Incidence

Data. We linked the Google Community Mobility Reports [Google LLC, 2022] (workplaces index; daily % point deviation from the Jan 3–Feb 6, 2020 baseline) with New York Times county-level COVID-19 case counts [The New York Times, 2025]. Units were counties. The exposure A is the 7-day average of the workplace mobility index ending at t_0 (units: percentage points change from baseline mobility in the January 3– February 6 period). We defined the calendar date t_0 to be the date with a mobility measurement closest to June 1, 2020, breaking ties arbitrarily. We excluded 74 counties without an eligible t_0 . Outcomes are future incident cases per 100,000 over the 7 days beginning 14 days after t_0 .

Intervention (MTP). We consider a constant shift policy $g(a) = a + \delta$ with $\delta = -5$ percentage points. This means that under this intervention each county’s percentage point change in mobility score from Jan3-Feb6 would be shifted down 5 percentage points from its observed value.

Adjustment covariates. We adjusted for: population, Rural–Urban Continuum (RUC) score (1=most urban, 9=most rural), percent Black, percent Hispanic, Republican vote share in 2016, poverty rate, unemployment rate, uninsurance rate, land area, population density, and per-capita income.

Estimand and blip specification. We model the blip as linear in the RUC score:

$$\Delta(H) \equiv \mathbb{E}\{Y(A + \delta) - Y(A) \mid H\} \approx \delta s(H)^\top \psi, \quad s(H) = (1, \text{RUC})^\top.$$

Here, ψ_1 is the average effect per 1 p.p. shift at RUC= 0 (intercept on the RUC scale) and ψ_2 captures linear effect modification by RUC. With $\delta = -5$, a positive ψ_2 means that a decreasing mobility is more protective in more rural areas. Treating RUC as a nominal variable when it is really ordinal is a slight

abuse. It surely pales in comparison to the omission of other covariates from the blip function. Keep in mind that this is an illustrative analysis.

Estimator. We used the point exposure Neyman-orthogonal estimator under exchangeability assumptions from Theorem 1. Nuisance functions were estimated with gradient-boosted trees (XGBoost): $\mu(a, h) \approx \mathbb{E}[Y \mid A = a, H = h]$ and $m(h) \approx \mathbb{E}[A \mid H = h]$, with a Normal working residual for $A \mid H$ to form the density-ratio pullback $p(a - \delta \mid h)/p(a \mid h)$. We used 5-fold cross-fitting. Standard errors come from the influence-function (IF) sandwich; Wald CIs are reported. Pointwise CIs for effects at specific RUC values use the delta method, i.e., $\widehat{\text{Var}}\{\delta s(H)^\top \hat{\psi}\} = \delta^2 s(H)^\top \widehat{\text{Var}}(\hat{\psi}) s(H)$.

Results. Estimated blip coefficients (per 1 p.p. shift):

$$\hat{\psi}_1 = 14.3 \quad (\text{SE} = 14.0; 95\% \text{ CI: } -13.2, 41.8), \quad \hat{\psi}_2 = -3.1 \quad (\text{SE} = 5.0; 95\% \text{ CI: } -13.0, 6.8).$$

With the policy shift $\delta = -5$ p.p., the implied change in cases per 100,000 at a given RUC is $\hat{\Delta}(\text{RUC}) = \delta (\hat{\psi}_1 + \hat{\psi}_2 \text{RUC})$. Delta-method estimates (and 95% CIs) at representative RUC values are:

$$\begin{aligned} \text{RUC} = 1 : \quad \hat{\Delta} &= -7.24, \text{ SE} = 38.08, \text{ CI } [-81.88, 67.40]; \\ \text{RUC} = 5 : \quad \hat{\Delta} &= -20.79, \text{ SE} = 22.46, \text{ CI } [-64.82, 23.24]; \\ \text{RUC} = 9 : \quad \hat{\Delta} &= -34.34, \text{ SE} = 41.65, \text{ CI } [-115.97, 47.29]. \end{aligned}$$

Interpretation. Point estimates suggest that a 5 p.p. reduction in workplace mobility index relative to what was observed would have lowered future COVID-19 cases per 100,000, with larger reductions in more rural counties (more negative at higher RUC). However, all 95% CIs include zero, and effects are imprecisely estimated in this purely illustrative analysis.

7 Discussion and Future Work

Estimating heterogeneous effects of MTPs with SNMMs can be both of scientific interest and important for planning realistic interventions. Suppose policy makers want to know the effect of a campaign to decrease opioid prescription dosing on some continuous quality of life utility measure. They believe the impact of the program might be approximated by the effect of an MTP: ‘prescribe a 20% lower dose of opioids than you normally would’. One might be interested in how the effect of this intervention varies

with the natural dose. Perhaps it turns out that reductions are more impactful when the natural dose is moderate, as opposed to low or high. Learning this might lead to hypotheses about the mechanism of opioid addiction that could be tested in further studies. The knowledge might also help to design targeted opioid reduction interventions.

Furthermore, economists who commonly apply DiD can now study MTPs under similar parallel trends assumptions. For example, they can estimate the effect of increasing minimum wage by 1 dollar more than it actually increased, and how this effect varies with the amount it actually increased. This is not an estimand that can be targeted by standard DiD.

There are many directions for future work. First, SNMMs can be estimated under instrumental variable assumptions [Robins, 1994b], and instrumental variable estimation of MTPs would also be of interest. Second, SNMM variants have been developed to estimate effects on survival [Picciotto et al., 2012] and binary [Wang et al., 2023] outcomes. Effects of MTPs on survival and binary outcomes are of course of interest as well. Third, sensitivity analysis for SNMMs [Robins, 2004, Robins et al., 2000] is well developed and would hopefully easily port over to the MTP setting, but that should be confirmed. Finally, marginal SNMMs that model heterogeneity as a function of a subset of covariates required for adjustment have been developed elsewhere, and it might be useful to transport them to the MTP setting, too.

8 Acknowledgments

I am extremely indebted to Jamie Robins for lots of things, but most relevant for this Acknowledgments section is that he suggested the general idea for this paper (‘SNMMs for MTPs’). However, all errors are my own.

References

- Thomas S Richardson and James M Robins. Single world intervention graphs (swigs): A unification of the counterfactual and graphical approaches to causality. *Center for the Statistics and the Social Sciences, University of Washington Series. Working Paper*, 128(30):2013, 2013.
- Iván Díaz, Nicholas Williams, Katherine L Hoffman, and Edward J Schenck. Nonparametric causal effects based on longitudinal modified treatment policies. *Journal of the American Statistical Association*, 118(542):846–857, 2023.

- James M. Robins. Correcting for non-compliance in randomized trials using structural nested mean models. *Communications in Statistics-Theory and Methods*, 23(8), 1994a.
- James M Robins. Optimal structural nested models for optimal sequential decisions. In *Proceedings of the Second Seattle Symposium in Biostatistics: analysis of correlated data*, pages 189–326. Springer, 2004.
- Stijn Vansteelandt and Marshall Joffe. Structural nested models and g-estimation: the partially realized promise. *Statistical Science*, 29(4):707–731, 2014.
- James M Robins, Miguel A Hernán, and UWE SiEBERT. Effects of multiple interventions. *Comparative quantification of health risks: global and regional burden of disease attributable to selected major risk factors*, 1:2191–2230, 2004.
- Jessica G Young, Miguel A Hernán, and James M Robins. Identification, estimation and approximation of risk under interventions that depend on the natural value of treatment using observational data. *Epidemiologic methods*, 3(1):1–19, 2014.
- Sebastian Haneuse and Andrea Rotnitzky. Estimation of the effect of interventions that modify the received treatment. *Statistics in medicine*, 32(30):5260–5277, 2013.
- Iván Díaz Muñoz and Mark Van Der Laan. Population intervention causal effects based on stochastic interventions. *Biometrics*, 68(2):541–549, 2012.
- Numair Sani, Jaron Lee, and Ilya Shpitser. Identification and estimation of causal effects defined by shift interventions. In *Conference on Uncertainty in Artificial Intelligence*, pages 949–958. PMLR, 2020.
- Zach Shahn, Oliver Dukes, Meghana Shamsunder, David Richardson, Eric Tchetgen Tchetgen, and James Robins. Structural nested mean models under parallel trends assumptions. *arXiv preprint arXiv:2204.10291*, 2022.
- 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.
- Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. Double/debiased machine learning for treatment and structural parameters, 2018.

- Google LLC. Covid-19 community mobility reports. <https://www.google.com/covid19/mobility/>, 2022. Historical data publicly available; updates ended Oct 15, 2022.
- The New York Times. Covid-19 data repository by the new york times. <https://github.com/nytimes/covid-19-data>, 2025. County-level daily cumulative cases/deaths; includes `us-counties.csv`.
- 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, 1994b.
- Sally Picciotto, Miguel A Hernán, John H Page, Jessica G Young, and James M Robins. Structural nested cumulative failure time models to estimate the effects of interventions. *Journal of the American Statistical Association*, 107(499):886–900, 2012.
- Linbo Wang, Xiang Meng, Thomas S Richardson, and James M Robins. Coherent modeling of longitudinal causal effects on binary outcomes. *Biometrics*, 79(2):775–787, 2023.
- James M Robins, Andrea Rotnitzky, and Daniel O Scharfstein. Sensitivity analysis for selection bias and unmeasured confounding in missing data and causal inference models. In *Statistical models in epidemiology, the environment, and clinical trials*, pages 1–94. Springer, 2000.

Appendix

8.1 Proof of Theorem 1

Proof. Step 1: Identification. By the SNMM restriction,

$$E[q\{\mu(A, L) - \mu(g(A, L), L) - \gamma_g(L, A; \psi^*)\}] = 0.$$

Moreover $E[Y - \mu(A, L) \mid A, L] = 0$, hence the augmentation term in (11) has mean zero. Summing yields $E[\phi(O; \psi^*, \eta^*)] = 0$.

Step 2: Orthogonality w.r.t. μ . Consider $\mu_t = \mu^* + t\delta\mu + o(t)$ with π fixed. Only $Y - \mu_t$ and the identifying bracket depend on t . Differentiate the expectation of (11) at $t = 0$:

$$\left. \frac{d}{dt} \right|_0 E[(q - \tilde{q})(Y - \mu_t)] = -E[(q - \tilde{q})\delta\mu] = -\langle q - \tilde{q}, \delta\mu \rangle_w.$$

For the identifying term,

$$\left. \frac{d}{dt} \right|_0 E[q\{\mu_t - T_g \mu_t - \gamma_g(\psi^*)\}] = E[q \delta \mu] - E[q T_g \delta \mu].$$

Rewrite under $\langle \cdot, \cdot \rangle_w$ and apply the definition of $\tilde{q}$:

$$E[q \delta \mu] - E[q T_g \delta \mu] = \langle q, \delta \mu \rangle_w - \langle q, T_g \delta \mu \rangle_w = \langle q, \delta \mu \rangle_w - \langle \tilde{q}, \delta \mu \rangle_w = \langle q - \tilde{q}, \delta \mu \rangle_w.$$

Summing the two derivatives: $-\langle q - \tilde{q}, \delta \mu \rangle_w + \langle q - \tilde{q}, \delta \mu \rangle_w = 0$.

Step 3: Orthogonality w.r.t. π . Let $\pi_t = \pi^* + t \delta \pi + o(t)$ and $w_t = \pi_t$. Then both w_t and $\tilde{q}_t := T_g^\dagger q$ vary with t . Write

$$F(t) := E[\phi(O; \psi^*, \mu^*, \pi_t)] = \underbrace{E[(q - \tilde{q}_t)(Y - \mu^*)]}_{=: A(t)} + \underbrace{E[q\{\mu^* - T_g \mu^* - \gamma_g(\psi^*)\}]}_{=: B \text{ (const. in } t)}.$$

Hence $F'(0) = A'(0)$. Differentiate $A(t)$ at 0:

$$A'(0) = -E[\dot{\tilde{q}}(Y - \mu^*) w^*] + E[(q - \tilde{q}^*)(Y - \mu^*) \dot{w}],$$

where dots denote t -derivatives at 0, and $w^* = \pi^*$. Now *condition on* (A, L) . Since $E[Y - \mu^*(A, L) | A, L] = 0$, both terms vanish:

$$E[\dot{\tilde{q}}(Y - \mu^*) w^*] = E[E[Y - \mu^* | A, L] \dot{\tilde{q}} w^*] = 0,$$

and similarly for the second term. Therefore $A'(0) = 0$, i.e. $F'(0) = 0$.

An equivalent route is to differentiate the identity $\langle q, T_g h \rangle_{w_t} = \langle \tilde{q}_t, h \rangle_{w_t}$ at $t = 0$ with $h(a, l) = Y - \mu^*(a, l)$. This yields

$$E[(q - \tilde{q}^*) T_g h \dot{w}] = E[\dot{\tilde{q}} h w^*],$$

which is exactly the cancellation needed in $A'(0)$ when h is replaced by $Y - \mu^*$. $\square$

8.2 Proof of Theorem 2

Proof. Step 1: Identification. By construction, for each t , $E[V_{t+1} - \mu_t(H_t, A_t) \mid H_t, A_t] = 0$, hence the augmentation term of ϕ_t has mean zero. The longitudinal SNMM restriction gives $E[q_t\{\mu_t(H_t, A_t) - \mu_t(H_t, g_t(A_t, H_t)) - \gamma_t(H_t, A_t; \psi^*)\}] = 0$. Summing over t yields $E[\Phi(O; \psi^*, \eta^*)] = 0$.

Step 2: Orthogonality w.r.t. μ_t . Fix t and perturb $\mu_{t,\varepsilon} = \mu_t^* + \varepsilon\delta\mu_t + o(\varepsilon)$, holding all $\{\mu_s, \pi_s\}_{s \neq t}$ and π_t fixed. Note that V_{t+1} is treated as fixed when differentiating ϕ_t in its own nuisances. Differentiate the expectation of ϕ_t at $\varepsilon = 0$:

$$\left. \frac{d}{d\varepsilon} \right|_0 E[(q_t - \tilde{q}_t)(V_{t+1} - \mu_{t,\varepsilon}(H_t, A_t))] = -\langle q_t - \tilde{q}_t, \delta\mu_t \rangle_{w_t},$$

and

$$\left. \frac{d}{d\varepsilon} \right|_0 E[q_t\{\mu_{t,\varepsilon} - T_{g,t}\mu_{t,\varepsilon} - \gamma_t(\psi^*)\}] = \langle q_t, \delta\mu_t \rangle_{w_t} - \langle q_t, T_{g,t}\delta\mu_t \rangle_{w_t} = \langle q_t - \tilde{q}_t, \delta\mu_t \rangle_{w_t},$$

by the adjoint identity for time t . The two derivatives cancel exactly. Because ϕ_t does not depend on μ_s for $s \neq t$ (given V_{t+1} fixed), we obtain $D_{\mu_t}E[\phi_t] = 0$.

Step 3: Orthogonality w.r.t. π_t . Perturb $\pi_{t,\varepsilon}$; then $w_{t,\varepsilon} = \pi_{t,\varepsilon}$ and $\tilde{q}_{t,\varepsilon} = T_{g,t}^{\dagger, w_{t,\varepsilon}} q_t$ vary with ε . Write

$$F_t(\varepsilon) := E[\phi_t(O; \psi^*, \mu_t^*, \pi_{t,\varepsilon})] = \underbrace{E[(q_t - \tilde{q}_{t,\varepsilon})(V_{t+1} - \mu_t^*)]}_{=: A_t(\varepsilon)} + \underbrace{E[q_t\{\mu_t^* - T_{g,t}\mu_t^* - \gamma_t(\psi^*)\}]}_{=: B_t \text{ const.}}.$$

Hence $F'_t(0) = A'_t(0)$. Differentiate at 0:

$$A'_t(0) = -E[\dot{\tilde{q}}_t (V_{t+1} - \mu_t^*) w_t^*] + E[(q_t - \tilde{q}_t^*) (V_{t+1} - \mu_t^*) \dot{w}_t].$$

Conditioning on (H_t, A_t) and using $E[V_{t+1} - \mu_t^*(H_t, A_t) \mid H_t, A_t] = 0$, both terms are zero, hence $A'_t(0) = 0$. Equivalently, differentiate the time- t adjoint identity $\langle q_t, T_{g,t}h \rangle_{w_{t,\varepsilon}} = \langle \tilde{q}_{t,\varepsilon}, h \rangle_{w_{t,\varepsilon}}$ at $\varepsilon = 0$ with $h(H_t, A_t) = V_{t+1} - \mu_t^*(H_t, A_t)$ to see the same cancellation.

Summing the time-local derivatives over t gives orthogonality for Φ . $\square$