



Identification in nonparametric models for dynamic treatment effects[☆]

Sukjin Han

Department of Economics, University of Bristol, United Kingdom



ARTICLE INFO

Article history:

Received 14 January 2019

Received in revised form 14 January 2019

Accepted 8 August 2019

Available online 27 November 2020

JEL classification:

C14, C32, C33, C36

Keywords:

Dynamic treatment effect

Endogenous treatment

Average treatment effect

Optimal treatment regime

Instrumental variable

ABSTRACT

This paper develops a nonparametric model that represents how sequences of outcomes and treatment choices influence one another in a dynamic manner. In this setting, we are interested in identifying the average outcome for individuals in each period, had a particular treatment sequence been assigned. The identification of this quantity allows us to identify the average treatment effects (ATE's) and the ATE's on transitions, as well as the optimal treatment regimes, namely, the regimes that maximize the (weighted) sum of the average potential outcomes, possibly less the cost of the treatments. The main contribution of this paper is to relax the sequential randomization assumption widely used in the biostatistics literature by introducing a flexible choice-theoretic framework for a sequence of endogenous treatments. This framework allows non-compliance of subjects in experimental studies or endogenous treatment decisions in observational settings. We show that the parameters of interest are identified under each period's exclusion restrictions, which are motivated by, e.g., a sequence of randomized treatment assignments or encouragements and a behavioral/information assumption on agents who receive treatments.

Crown Copyright © 2020 Published by Elsevier B.V. All rights reserved.

1. Introduction

This paper develops a nonparametric model that represents how sequences of outcomes and treatment choices influence one another in a dynamic manner. Often, treatments are chosen multiple times over a horizon, affecting a series of outcomes. Examples are medical interventions that affect health outcomes, educational interventions that affect academic achievements, or online advertisements that affect consumers' preferences or purchase decisions. Agents endogenously make decisions of receiving treatments, e.g., whether to comply with random assignments. The relationship of interest is dynamic in the sense that the current outcome is determined by the past outcomes as well as current treatment, and the current treatment is determined by the past outcomes and treatments. Such dynamic relationships are clearly present in the aforementioned examples. A static model misrepresents the nature of the problem (e.g., nonstationarity, state dependence, learning) and fails to capture important policy questions (e.g., optimal timing and schedule of interventions).

In this setting, we are interested in identifying the dynamic causal effect of a sequence of treatments on a sequence of outcomes or on a terminal outcome. We are interested in learning about the average of the outcome in each period, had a particular treatment sequence been assigned *up to that period*, which defines the potential outcome in this dynamic setting.

[☆] The author is grateful to Dan Akerberg, Tim Armstrong, Stephane Bonhomme, Xiaohong Chen, Jim Heckman, Pedro Sant'Anna, Ed Vytlačil, Shenshen Yang and Nese Yildiz for their helpful discussions. Also, comments from participants in seminars at UChicago, Rochester, USC, Penn State, Indiana, Boston College, NUS, SNU, and in the 2018 Asian Meeting and China Meeting of the Econometric Society, the 2018 International Panel Data Conference, the 2019 Texas Econometrics Camp are appreciated.

E-mail address: sukjin.han@gmail.com.

We are also interested in the average treatment effects (ATE's) and the transition-specific ATE's defined based on the average potential outcome, unconditional and conditional on the previous outcomes, respectively. For example, one may be interested in whether the success rate of a particular outcome (or the transition probability) is larger with a sequence of treatments assigned in relatively later periods rather than earlier, or with a sequence of alternating treatments rather than consistent treatments. Lastly, we are interested in the optimal treatment regimes, namely, sequences of treatments that maximize the (weighted) sum of the average potential outcomes, possibly less the cost of the treatments. For example, a firm may be interested in the optimal timing of advertisements that maximizes its aggregate sales probabilities over time, or a sequence of educational programs may be aimed to maximize the college attendance rate. Knowledge about the optimal treatment regime may be useful for policies, e.g., by at least excluding sequences that are suboptimal.

Dynamic treatment effects have been extensively studied in the biostatistics literature for decades under the counterfactual framework with a sequence of treatments (Robins (1986, 1987, 1997) and Murphy et al. (2001), Murphy (2003), among others). In this literature, the crucial condition used to identify the average potential outcome is a dynamic version of a random assignment assumption, called the *sequential randomization*. This condition assumes that the treatment is randomized in every period within those individuals who have the same history of outcomes and treatments.¹ This assumption is only suitable in experimental studies with the perfect compliance of subjects, which is often infeasible (Robins, 1994; Robins and Rotnitzky, 2004). When interventions continue for multiple periods as in the examples described above, non-compliance may become more prevalent than in one-time experiments, e.g., due to the cost of enforcement or the subjects' learning. In addition to partial compliance in experimental settings, sequential randomization is invalid in many observational contexts as well.

The main contribution of this paper is to relax the assumption of sequential randomization widely used in the literature by establishing a flexible choice-theoretic framework for a sequence of endogenous treatments. To this end, we consider a simple nonparametric structural model for a dynamic endogenous selection process and dynamic outcome formation. The heterogeneity in each period's potential outcome is given by recursively applying a switching-regression type of models with a sequential version of rank similarity. The joint distribution of the full history of unobservable variables in the outcome and treatment equations is still flexible, allowing for arbitrary forms of treatment endogeneity as well as serial correlation. Relative to the counterfactual framework, the dynamic mechanism is clearly formulated using this structural model, which in turn facilitates our identification analysis.

We show that the average potential outcome is identified under exclusion restrictions. That is, we first assume there exist (possibly binary) instruments excluded from the outcome-determining process. A leading example is a sequence of randomized treatment assignments or encouragements (Sexton and Hebel, 1984) from, e.g., clinical trials, field experiments, and A/B testings, and other examples include sequential policy shocks. We then show that exogenous variables excluded from the treatment-selection process are helpful (although not necessary) for identification especially with binary instruments. Examples include factors that agents cannot fully anticipate when making treatment or compliance choices but that determine the outcome. We show that such timing can be justified in this dynamic context, and some covariates in the outcome process may be valid candidates. Identification in nonseparable triangular models using these exclusion restrictions is pioneered by Vytlačil and Yildiz (2007) and subsequently appears in Shaikh and Vytlačil (2011) and Balat and Han (2020) among others, all in static settings. The dynamic structure introduced in this paper poses added challenges in using a similar strategy, since (i) the outcome and treatment structural functions depend on the vectors of lags, which in turn make each potential outcome a direct function of all the previous potential outcomes, (ii) the period specific knowledge analogous to that in Vytlačil and Yildiz (2007) does not directly recover any meaningful objects of interest in general, (iii) rank invariance substantially restricts heterogeneity in this dynamic setting, and (iv) the initial condition problem is present. In this paper, we address these challenges and show how to achieve identification. The proof is constructive and provides a closed form expression for the average potential outcome. The identification of each period's average potential outcome allows us to point identify the ATE's, transition-specific ATE's and the optimal treatment regimes.

This paper contributes to growing research on the identification of the effects of dynamic endogenous treatments that allows for treatment heterogeneity. Cunha et al. (2007) and Heckman and Navarro (2007) consider a semiparametric discrete-time duration model for the choice of the treatment timing and associated outcomes. Building on these works, Heckman et al. (2016) consider not only ordered choice models but also unordered choice models for up-or-out treatment choices.² As an important feature, these papers consider attrition based on the irreversible treatment decisions; see also Sasaki (2015). Similar to our approach, Heckman and Navarro (2007), Heckman et al. (2016) utilize exclusion restrictions. Unlike these papers, however, we do not necessarily invoke infinite supports of each period's instruments but instead require additional exogenous variables. Abraham and Sun (2018), Athey and Imbens (2018), Callaway and Sant'Anna (2020), and Imai et al. (2018) extend a difference-in-differences approach to dynamic settings without specifying fixed-effect panel data models. The first three papers consider the effects of treatment timing on the treated,

¹ This assumption is also called sequential conditional independence or sequential ignorability. In the econometrics literature, Vikström et al. (2018) consider treatment effects on a transition to a destination state, and carefully analyze what the sequential randomization assumption can identify in the presence of dynamic selection.

² As related works, the settings of Angrist and Imbens (1995), Jun et al. (2016), and Lee and Salanié (2018) for multiple (or multi-valued) treatment effects may be applied to a dynamic setting. Also, see Abbring and Heckman (2007) for a survey on dynamic treatment effects.

where the treatment process is irreversible as in the previous works. Unlike all the papers mentioned in this paragraph except Imai et al. (2018), we consider nonparametric dynamic models for treatment and outcome processes with a general form of evolution, where the processes can freely change states with an irreversible process being a special case. Moreover, we consider different identifying assumptions than those in the previous works and focus on the identification of the ATE's and related parameters.

This paper's structural approach is only relative to the counterfactual framework of Robins. A fully structural model of dynamic programming is considered in the seminal work by Rust (1987) and more recently by, e.g., Bleivins (2014) and Buchholz et al. (2019). This literature typically considers a single rational agent's optimal decision, whereas we consider a large group of heterogeneous agents with no assumptions on agents' rationality or strong parametric assumptions. Most importantly, our focus is on the identification of the effects of treatments formed as agents' decisions.

In the next section, we first introduce Robins' counterfactual outcome framework and discuss sequential randomization. Section 3 introduces the main structural model of this paper with parameters of interest, followed by a motivating example in Section 4. The main identifying conditions and identification results are present in Section 5, and an extension in Section 6. Section 7 briefly concludes. In the Appendix all the proofs are collected and estimation and inference are discussed.

In terms of notation, let $\mathbf{W}^t \equiv (W_1, \dots, W_t)$ denote a row vector that collects r.v.'s W_t across time up to t , and let \mathbf{w}^t be its realization. Note $\mathbf{W}^1 = \mathbf{W}_1$. We sometimes write $\mathbf{W} \equiv \mathbf{W}^T$ for convenience. For a vector \mathbf{W} without the t th element, we write $\mathbf{W}_{-t} \equiv (W_1, \dots, W_{t-1}, W_{t+1}, \dots, W_T)$ with realization \mathbf{w}_{-t} . Lastly, for r.v.'s Y and W , we sometimes abbreviate $\Pr[Y = y|W = w]$ and $\Pr[Y = y|W \in \mathcal{W}]$ to $\Pr[Y = y|w]$ (or $P[y|w]$) and $\Pr[Y = y|\mathcal{W}]$, respectively.

2. Robins' framework

We first introduce Robins' counterfactual framework and state the assumption of sequential randomization commonly used in the biostatistics literature (Robins, 1986, 1987; Murphy et al., 2001; Murphy, 2003). For a finite horizon $t = 1, \dots, T$ with fixed T , let Y_t be the outcome at t with realization y_t and let D_t be the binary treatment at t with realization d_t . The underlying data structure is panel data with a large number of cross-sectional observations over a short period of time (and the cross-sectional index i suppressed throughout, unless necessary). We call Y_T a *terminal outcome* and Y_t for $t \leq T - 1$ a *intermediate outcome*.³ Let \mathcal{Y} and $\mathcal{D} \subseteq \{0, 1\}^T$ be the supports of $\mathbf{Y} \equiv (Y_1, \dots, Y_T)$ and $\mathbf{D} \equiv (D_1, \dots, D_T)$, respectively. There can be other time-varying covariates present in this setup, but we omit them here.

Consider a treatment regime $\mathbf{d} \equiv (d_1, \dots, d_T) \in \mathcal{D}$, which is defined as a predetermined hypothetical sequence of interventions over time, i.e., a sequence of each period's decisions on whether to treat or not, or whether to choose treatment A or treatment B.⁴ Then, a potential outcome at t can be written as $Y_t(\mathbf{d})$. This can be understood as an outcome for an individual, had a particular treatment sequence been assigned. Although the genesis of $Y_t(\mathbf{d})$ can be very general under this counterfactual framework, the mechanism under which the sequence of treatments interacts with the sequence of outcomes is opaque. The definition of $Y_t(\mathbf{d})$ becomes more transparent later with the structural model introduced in this paper.

Given these definitions, we state the assumption of sequential randomization by Robins: For each $\mathbf{d} \in \mathcal{D}$,

$$(Y_1(\mathbf{d}), \dots, Y_T(\mathbf{d})) \perp D_t | \mathbf{Y}^{t-1}, \mathbf{D}^{t-1} \quad (2.1)$$

for $t = 1, \dots, T$. This assumption asserts that, holding the history of outcomes and treatments (and potentially other covariates) fixed, the current treatment is fully randomized. Under this assumption, the no-anticipation condition (discussed below), and the positivity condition,⁵ the biostatistics literature mentioned above identifies an average potential outcome. Sequential randomization, however, can be violated if agents partially comply with assignments and make decisions D_t based on time-varying or time-invariant factors, unobserved to the analyst. In the next section, we relax this assumption and specify dynamic selection equations for a sequence of treatments that are allowed to be endogenous, i.e., to be dependent on unobservable factors. We then introduce exclusion restrictions (with relevant support assumptions) and rank similarity to identify the average potential outcome and related treatment parameters. Apart from the sequential randomization assumption, we maintain the same preliminaries introduced in this section.

Remark 2.1 (Irreversibility). As a special case of our setting, the process of D_t may be irreversible in that the process only moves from an initial state to a destination state, i.e., the destination state is an absorbing state. The up-or-out treatment decision (or the treatment timing) can be an example where the treatment process satisfies $D_t = 1$ once $D_{t-1} = 1$ is reached, as in Heckman and Navarro (2007), Heckman et al. (2016), Abraham and Sun (2018) and Callaway and Sant'Anna (2020). Although it is not the main focus of this paper, the process of Y_t may as well be irreversible. This case, however, requires caution due to dynamic selection; see discussions later in this paper. The survival of patients ($Y_t = 0$) in discrete

³ The terminal period T may be an administrative end of follow-up time.

⁴ This is called a nondynamic regime in the biostatistics literature. A dynamic regime is a sequence of treatment assignments, each of which is a predetermined function of past outcomes. A nondynamic regime can be viewed as its special case, where this function is constant. See, e.g., Murphy et al. (2001) and Murphy (2003) for related discussions.

⁵ A treatment sequence relevant to defining the parameter has a positive probability of occurring.

time duration models can be an example where the transition of the outcome satisfies $Y_t = 1$ once $Y_{t-1} = 1$. In this case, it may be that D_t is missing when $Y_{t-1} = 1$, which can be dealt by conventionally assuming $D_t = 0$ if $Y_{t-1} = 1$. When processes are irreversible, the supports \mathcal{D} and \mathcal{Y} are strict subsets of $\{0, 1\}^T$.

Remark 2.2 (*Terminal Outcome of a Different Kind*). As in [Murphy et al. \(2001\)](#) and [Murphy \(2003\)](#), we may be interested in a terminal outcome that is of a different kind than that of the intermediate outcomes. For example, the terminal outcome can be college attendance, while the intermediate outcomes are secondary school performances. In this case, we replace Y_T with a random variable R_T to represent the terminal outcome, while maintaining Y_t for $t \leq T - 1$ to represent the intermediate outcomes. Analogously, $R_T(\mathbf{d})$ denotes the potential terminal outcome. Then, the analysis in this paper can be readily followed with the change of notation.⁶

3. A dynamic structural model and objects of interest

We now introduce the main framework of this paper. Consider a *dynamic structural function* for the outcomes, where Y_t depends on the entire history of outcomes (\mathbf{Y}^{t-1}) and the current treatment (D_t), and that has the form of switching regression models: For $t = 1, \dots, T$,

$$Y_t = \mu_t(\mathbf{Y}^{t-1}, D_t, X_t, U_t(D_t)),$$

where $\mu_t(\cdot)$ is an unknown scalar-valued function, X_t is a set of exogenous variables, which we discuss in detail later, and Y_0 is assumed to be exogenously determined, with $Y_0 = 0$ for convenience.⁷ Here, Y_t depends on (D_{t-s}, X_{t-s}) ($s \geq 1$) via Y_{t-s} . There can be other potentially endogenous covariates, which is suppressed in the model. The unobservable variable satisfies $U_t(D_t) = D_t U_t(1) + (1 - D_t) U_t(0)$, where $U_t(d_t)$ is the “rank variable” that captures the unobserved characteristics or rank, specific to treatment state d_t ([Chernozhukov and Hansen, 2005](#)). We allow $U_{it}(d_t)$ to contain a permanent component (i.e., individual effects) and a transitory component.⁸ Given this structural equation, we can express the potential outcome $Y_t(\mathbf{d})$ using a recursive structure:

$$\begin{aligned} Y_t(\mathbf{d}) &= Y_t(\mathbf{d}^t) = \mu_t(\mathbf{Y}^{t-1}(\mathbf{d}^{t-1}), d_t, X_t, U_t(d_t)), \\ &\quad \text{where } \mathbf{Y}^{t-1}(\mathbf{d}^{t-1}) \equiv (Y_1(d_1), Y_2(\mathbf{d}^2), \dots, Y_{t-1}(\mathbf{d}^{t-1})), \\ &\quad \vdots \\ Y_2(\mathbf{d}) &= Y_2(\mathbf{d}^2) = \mu_2(Y_1(d_1), d_2, X_2, U_2(d_2)), \\ Y_1(\mathbf{d}) &= Y_1(d_1) = \mu_1(Y_0, d_1, X_1, U_1(d_1)), \end{aligned}$$

where each potential outcome at time t is only a function of \mathbf{d}^t (not the full \mathbf{d}). This is related to the “no-anticipation” condition ([Abbring and Heckman, 2007](#)) or the “consistency” condition ([Robins, 2000](#)), which is implied from the structure of the model in our setting. The recursive structure provides us with a useful interpretation of the potential outcome $Y_t(\mathbf{d})$ in a dynamic setting, and thus facilitates our identification analysis. This structure allows us to exploit the variation in X_t in order to relax the sequential randomization assumption. Still, we allow rich channels in the evolution of potential outcomes, as $Y_t(\mathbf{d})$ is a function of *all* the past potential outcomes whose treatment indices are consistent with \mathbf{d} . Also, conditional on $\mathbf{X}^t \equiv (X_1, \dots, X_t)$, the heterogeneity in $Y_t(\mathbf{d})$ comes from the full vector $\mathbf{U}^t(\mathbf{d}^t) \equiv (U_1(d_1), \dots, U_t(d_t))$. By an iterative argument, we can readily show that the potential outcome is equal to the observed outcome when the observed treatments are consistent with the assigned regime: $Y_t(\mathbf{d}) = Y_t$ when $\mathbf{D} = \mathbf{d}$, or equivalently, $Y_t = \sum_{\mathbf{d} \in \mathcal{D}} 1\{\mathbf{D} = \mathbf{d}\} Y_t(\mathbf{d})$.

In this paper, we consider the *average potential terminal outcome*, conditional on $\mathbf{X} = \mathbf{x}$, as the fundamental parameter of interest:

$$E[Y_T(\mathbf{d}) | \mathbf{X} = \mathbf{x}]. \tag{3.1}$$

We also call this parameter the *average recursive structural function* (ARSF) in the terminal period, named after the recursive structure in the model for $Y_T(\mathbf{d})$. Generally, in defining this parameter and all others below, we can consider the potential outcome in any time period of interest, e.g., $E[Y_t(\mathbf{d}) | \mathbf{X}^t = \mathbf{x}^t]$ for any given t .⁹ We focus on the terminal potential outcome only for concreteness. The knowledge of the ARSF is useful in recovering other related parameters.

First, we are interested in the conditional ATE:

$$ATE(\mathbf{d}, \tilde{\mathbf{d}}; \mathbf{x}) \equiv E[Y_T(\mathbf{d}) - Y_T(\tilde{\mathbf{d}}) | \mathbf{X} = \mathbf{x}] \tag{3.2}$$

⁶ Extending this framework to incorporate the irreversibility of the outcome variables discussed in [Remark 2.1](#) is not straightforward. We leave this for future research.

⁷ This assumption of an exogenous initial outcome is *not* necessary but only introduced to simplify our analysis; see [Remark 5.1](#) for alternative assumptions.

⁸ In this case, it may make sense that the permanent component does not depend on each d_t , but that the transitory component does.

⁹ Related to this general object, the treatment effect is said to be dynamic, partly because the effect can vary depending upon the period of measurement even if the same set of treatments is assigned.

for two different regimes, \mathbf{d} and $\tilde{\mathbf{d}}$. For example, one may be interested in comparing more versus less consistent treatment sequences, or earlier versus later treatments.

Second, we consider the *optimal treatment regime*:

$$\mathbf{d}^*(\mathbf{x}) = \arg \max_{\mathbf{d} \in \mathcal{D}} E[Y_T(\mathbf{d})|\mathbf{X} = \mathbf{x}] \tag{3.3}$$

with $|\mathcal{D}| \leq 2^T$. That is, we are interested in a treatment regime that delivers the maximum expected potential outcome, conditional on $\mathbf{X} = \mathbf{x}$. Notice that, in a static model, the identification of \mathbf{d}^* is equivalent to the identification of the sign of the static ATE, which is the information typically sought from a policy point of view (Abrevaya et al., 2010; Machado et al., 2019). One can view \mathbf{d}^* as a natural extension of this information to a dynamic setting, which is identified by establishing the signs of *all* possible ATE's defined as in (3.2), or equivalently, by ordering all the possible ARSF's. The optimal regime may serve as a guideline in developing future policies. Moreover, it may be a realistic goal for a social planner to identify this kind of scheme that maximizes the average benefit, because it may be too costly to find a customized treatment scheme for every individual.¹⁰ Given $\mathbf{d}^*(\mathbf{x})$, we may be interested in $E[Y_T(\mathbf{d}^*(\mathbf{x}))]$ or the ATE for the effect of $\mathbf{d}^*(\mathbf{x})$ relative to another treatment sequence (e.g., the second best). More ambitious than the identification of $\mathbf{d}^*(\mathbf{x})$ may be recovering an optimal regime based on a cost–benefit analysis, granting that each d_t can be costly:

$$\mathbf{d}^\dagger(\mathbf{x}) = \arg \max_{\mathbf{d} \in \mathcal{D}} \Pi(\mathbf{d}; \mathbf{x}), \tag{3.4}$$

where

$$\Pi(\mathbf{d}; \mathbf{x}) \equiv wE[Y_T(\mathbf{d})|\mathbf{X} = \mathbf{x}] - \tilde{w} \sum_{t=1}^T d_t \quad \text{or} \quad \Pi(\mathbf{d}; \mathbf{x}) \equiv \sum_{t=1}^T w_t E[Y_t(\mathbf{d})|\mathbf{X} = \mathbf{x}] - \sum_{t=1}^T \tilde{w}_t d_t$$

with (w, \tilde{w}) and (w, \tilde{w}) being predetermined weights. The latter objective function concerns the weighted sum of the average potential outcomes throughout the entire period, less the cost of treatments. Note that establishing the signs of ATE's will not identify \mathbf{d}^\dagger , and a stronger identification result becomes important, i.e., the point identification of $E[Y_T(\mathbf{d})|\mathbf{X} = \mathbf{x}]$ for all \mathbf{d} (or $E[Y_t(\mathbf{d})|\mathbf{X} = \mathbf{x}]$ for all t and \mathbf{d}). Lastly, as an extension of the ATE, we are interested in the *transition-specific ATE*, which we introduce and discuss later in Section 6. The identification of this parameter can be paralleled by the analysis for the ARSF and ATE.

In order to facilitate identification of the parameters of interest without assuming sequential randomization, we introduce a sequence of selection equations for the binary endogenous treatments, where D_t depends on the entire history of outcomes and treatments (\mathbf{Y}^{t-1} and \mathbf{D}^{t-1}): For $t = 1, \dots, T$,

$$D_t = 1\{\pi_t(\mathbf{Y}^{t-1}, \mathbf{D}^{t-1}, \mathbf{Z}^t) \geq V_t\},$$

where $\pi_t(\cdot)$ is an unknown scalar-valued function, Z_t is the period-specific instruments, V_t is the unobservable variable that may contain permanent and transitory components, and D_0 is assumed to be exogenously given as $D_0 = 0$.¹¹ This dynamic selection process represents the agent's endogenous choices over time, e.g., as a result of learning or other optimal behaviors. Allowing the possibility of the agent's learning, the equation for D_t depends on the full history of the treatments and exogenous variables, as well as the full path of the outcomes. However, the nonparametric threshold-crossing structure posits a minimal notion of optimality for the agent. We take an agnostic approach by avoiding strong assumptions of the standard dynamic economic models pioneered by Rust (1987), such as forward looking behaviors and being able to compute a present value discounted flow of utilities. If we are to maintain the assumption of rational agents, the selection model can be viewed as a reduced-form approximation of a solution to a dynamic programming problem. Lastly, due to the dynamic structure, this selection equation does not necessarily imply the monotonicity assumption of Imbens and Angrist (1994) or vice versa.

To simplify the exposition, we consider binary Y_t and impose weak separability in the outcome equation as in the treatment equation. The binary outcome is *not* necessary for the result of this paper, and the analysis can be easily extended to the case of continuous or censored Y_t , maintaining weak separability; see Remark 3.3. Then, the full model can be summarized as

$$Y_t = 1\{\mu_t(\mathbf{Y}^{t-1}, D_t, X_t) \geq U_t(D_t)\}, \tag{3.5}$$

$$D_t = 1\{\pi_t(\mathbf{Y}^{t-1}, \mathbf{D}^{t-1}, \mathbf{Z}^t) \geq V_t\}. \tag{3.6}$$

In this model, the observable variables are $(\mathbf{Y}, \mathbf{D}, \mathbf{X}, \mathbf{Z})$. As discussed in Section 5, X_t may (or may not) be included in Z_t , while X_{t-1} is certainly allowed to be included in Z_t . All other (potentially endogenous) covariates are suppressed in the equations for simplicity of exposition. Importantly, in this model, the joint distribution of the unobservable variables $(\mathbf{U}(\mathbf{d}), \mathbf{V})$ for given \mathbf{d} is not specified, in that $U_t(d_t)$ and $V_{t'}$ for any t, t' are allowed to be arbitrarily correlated to each

¹⁰ We still view \mathbf{d}^* as a descriptive (rather than prescriptive) parameter in this paper.

¹¹ This is an alternative to simply assuming there is no treatment at $t = 0$. We maintain the current assumption to avoid additional definitions for $\pi_t(\cdot)$ and other relevant objects.

other (allowing endogeneity) as well as within themselves across time (allowing serial correlation, e.g., via time-invariant individual effects). Note that, because we allow an arbitrary form of persistence in the unobservables and the dependence of Y_t and D_t on the history, (Y_t, D_t) is *not* a Markov process even after conditioning on the observables. This is in contrast to the standard dynamic economic models, where conditional independence assumptions or Markovian unobservables are commonly introduced. By considering the nonparametric index functions that depend on t , we also avoid other strong assumptions on parametric functional forms or time homogeneity.

Remark 3.1 (Irreversibility—Continued). A process that satisfies $D_t = 1$ if $D_{t-1} = 1$ is consistent with having a structural function that satisfies $\pi_t(\mathbf{y}^{t-1}, \mathbf{d}^{t-1}, \mathbf{z}^t) = +\infty$ if $d_{t-1} = 1$. Similarly, processes that satisfy $Y_t = 1$ and $D_t = 0$ if $Y_{t-1} = 1$ are consistent with $\mu_t(\mathbf{y}^{t-1}, d_t, x_t) = +\infty$ and $\pi_t(\mathbf{y}^{t-1}, \mathbf{d}^{t-1}, \mathbf{z}^t) = -\infty$ if $y_{t-1} = 1$. This implies that $Y_t(\mathbf{d}^t) = 1$ for any \mathbf{d}^t if $Y_{t-1}(\mathbf{d}^{t-1}) = 1$. When Y_t is irreversible, the ARSF $E[Y_T(\mathbf{d})|\mathbf{X}]$ can be interpreted as (one minus) a potential survival rate. An important caveat is that, with irreversible Y_t , the ATE we define contains not only the treatment effect (the intensive margin) but also the effect on dynamic selection (the extensive margin), and the parameter may or may not be of interest depending on the application.

Remark 3.2 (Terminal Outcome of a Different Kind—Continued). When we replace Y_T with R_T to represent a terminal outcome of a different kind, we assume that the model (3.5) is only satisfied for $t \leq T - 1$ and introduce $R_T = 1\{\mu_T(\mathbf{Y}^{T-1}, D_T, X_T) \geq U_T(D_T)\}$ as the terminal structural function. The potential terminal outcome $R_T(\mathbf{d})$ can accordingly be expressed using the structural functions for $(Y_1, \dots, Y_{T-1}, R_T)$. The ARSF is written as $E[R_T(\mathbf{d})|\mathbf{X}]$, and the other parameters can be defined accordingly.

Remark 3.3 (Non-Binary Y_t). Even though we focus on binary Y_t in this paper, we can obtain similar identification results with continuous Y_t or limited dependent variable Y_t , by maintaining a general weak separability structure: $Y_t = m_t(\mu_t(\mathbf{Y}^{t-1}, D_t, X_t), U_t(D_t))$. As in the static settings of Vytlačil and Yildiz (2007) and Balat and Han (2020), we impose an assumption that guarantees certain monotonicity of each period's average structural function with respect to the index μ_t : For each t , $E[m_t(\mu_t, U_t(d_t))|\mathbf{V}^t, \mathbf{U}^{t-1}]$ is strictly monotonic in μ_t . Examples of the nonparametric model $m_t(\mu_t(\mathbf{y}^{t-1}, d_t, x_t), u_t)$ that satisfies this assumption are additively separable models or their transformation models, censored regression models, and threshold crossing models as in (3.5); see Vytlačil and Yildiz (2007) for more discussions. Also, see the second illustrative example in Appendix B in the Appendix where a linear model is considered.

4. Motivating example

A multi-period experiment with imperfect compliance is one motivating example of this paper's setup. Multi-period experiments are common in clinical trials, such as in the Fast Track Prevention Program (Conduct Problems Prevention Research Group, 1992), the Elderly Program randomized trial for the Systolic Hypertension (In the Elderly Program (SHEP) Cooperative Research Group, T.S.H., 1988), and the AIDS Clinical Trial Group;¹² also see the biostatistics literature referenced in the introduction for other examples. For instance, the Fast Track Prevention Program is a randomized trial to prevent conduct disorders and drug use in children at risk. Interventions are taken place at the end of each semester starting from first grade, by means of home visits and teacher consultations. In household visits, for example, it is reported that assignment deviation occurs for nearly 50% of the intervention children. Murphy et al. (2001) focus on the effect of treatment had there been no deviation, i.e., the intention-to-treat parameters. In this paper, we recover the average treatment effect parameters allowing for this type of imperfect compliance.

Based on these clinical trials, we consider the following stylized example for the structural model of this paper. A clinical research organization is interested in improving patients' symptoms (Y_t), and runs an experiment of randomly assigning treatments at each t (Z_t). Based on the assignment history (\mathbf{Z}^t), each patient decides whether or not to receive the treatment (D_t) by being a complier, defier, always-taker or never-taker. This information can be collected via a fidelity assessment as in the Fast Track Prevention Program. In making the compliance decision, the patient has a habit (\mathbf{D}^{t-1}) and takes into account her past symptoms (\mathbf{Y}^{t-1}). The current symptom (Y_t) is formed based on the past symptoms (\mathbf{Y}^{t-1}) and the current treatment take-up (D_t), and other symptom-influencing factors (X_t). As described in detail in the next section, it is useful to assume that patients cannot fully predict X_t when making the treatment decision D_t , although \mathbf{X}^{t-1} can be taken into account as X_{t-1} can be included in Z_t . For patients with potential respiratory diseases, temporal variation in air quality can be such a variable. In the Fast Track Prevention Program, the average quality of teachers or the average performance of non-risk peers randomly assigned every academic year can be a candidate. When it is not plausible to assume X_t being excluded from the treatment decision, one alternative is to assume the existence of two continuous exogenous variables commonly affecting Y_t and D_t . Based on the data $(\mathbf{Y}, \mathbf{D}, \mathbf{X}, \mathbf{Z})$ generated from this process, we want to learn about the ATE's and optimal treatment regimes.

¹² The AIDS Clinical Trials Group (<https://actgnetwork.org>) is one of the largest HIV clinical trials organizations in the world.

5. Main identification analysis

We first identify the ARSF's, i.e., $E[Y_t(\mathbf{d})|\mathbf{X}^t]$ for every \mathbf{d} and t , which will then be used to identify the ATE's and the optimal regimes \mathbf{d}^* and \mathbf{d}^\dagger . We maintain the following assumptions on (\mathbf{Z}, \mathbf{X}) and $(\mathbf{U}(\mathbf{d}), \mathbf{V})$ for every \mathbf{d} . These assumptions are written for the identification of $E[Y_T(\mathbf{d})|\mathbf{X}]$, and are sufficient but not necessary for the identification of $E[Y_t(\mathbf{d})|\mathbf{X}^t]$ for $t \leq T - 1$.

Assumption C. The distribution of \mathbf{V} has strictly positive density with respect to Lebesgue measure on \mathbb{R}^{2T} .

Assumption SX. (\mathbf{Z}, \mathbf{X}) and $(\mathbf{U}(\mathbf{d}), \mathbf{V})$ are independent.

Assumption C is a regularity condition to ensure the smoothness of relevant conditional probabilities. **Assumption SX** imposes strict exogeneity. It is implicit that the independence is conditional on the covariates suppressed in the model. Just as the treatments \mathbf{D} , these covariates may be correlated with the individual effects contained in $(\mathbf{U}(\mathbf{d}), \mathbf{V})$. The variable Z_t contains the standard excluded instruments, which is allowed to be binary. A leading example is a sequence of randomized treatment assignments. Other examples include sequential policy shocks. In addition to Z_t , we introduce exogenous variables X_t in the outcome Eq. (3.5), that may or may not be included in Z_t . When we assume X_t is not included in Z_t , i.e., X_t is excluded from the selection equation (3.6), we essentially make a behavioral/information assumption that there are outcome-determining factors that the agent cannot fully anticipate when making a treatment decision. Continuing with the stylized example in Section 4, when D_t is a compliance choice that a patient makes at the t th visit to the clinical facility, Y_{t-1} may be the symptom measured prior to the decision during the same visit. Then Y_t is the symptom measured upon the next visit, which may create enough time gap to prevent the patient from predicting X_t .¹³ It is worth noting that X_{t-1} is allowed to be included in Z_t , in which case \mathbf{X}^{t-1} directly affects the compliance decision D_t via \mathbf{Z}^t . Next, we introduce a sequential version of the rank similarity assumption (Chernozhukov and Hansen, 2005). Let $\mathbf{d}_{-t} \equiv (d_1, \dots, d_{t-1}, d_{t+1}, \dots, d_T)$ for given t .

Assumption RS. For each t and \mathbf{d}_{-t} , $\mathbf{U}(1, \mathbf{d}_{-t})$ and $\mathbf{U}(0, \mathbf{d}_{-t})$ are identically distributed conditional on \mathbf{V}^t .

When $T = 2$, for example, **Assumption RS** imposes that, for any given d_2 and d_1 ,

$$\begin{aligned} (U_1(1), U_2(d_2))|V_1 &\stackrel{d}{=} (U_1(0), U_2(d_2))|V_1, \\ (U_1(d_1), U_2(1))|(V_1, V_2) &\stackrel{d}{=} (U_1(d_1), U_2(0))|(V_1, V_2). \end{aligned}$$

Rank invariance (i.e., $\{\mathbf{U}(\mathbf{d})\}_d$ being equal to each other) is particularly restrictive in the multi-period context, because it requires that the same rank be realized across 2^T different treatment states.¹⁴ Significantly weaker than the rank invariance would be a joint rank similarity assumption that $\mathbf{U}(\mathbf{d})$'s are identically distributed across 2^T states, conditional on \mathbf{V} . This allows an individual to have different realized ranks across different \mathbf{d} 's. **Assumption RS**, which we call *sequential rank similarity*, relaxes this further by requiring that $\mathbf{U}(1, \mathbf{d}_{-t})$ and $\mathbf{U}(0, \mathbf{d}_{-t})$ are identically distributed, conditional on \mathbf{V}^t instead. That is, the assumption requires that, within individuals with the same history of the treatment unobservables, the joint distributions of the ranks are identical between just two states that differ by $d_t = 1$ and 0.¹⁵ Although it is weaker than rank invariance, **Assumption RS** imposes restrictions on the selection process and rules out, e.g., a dynamic version of Roy models.

Now, we are ready to derive a period-specific result. Let $\mathbf{O}_t = (Y_t, D_t, X_t, Z_t)$ be the vector of the observables. Define the following period-specific quantity directly identified from the data, i.e., from the distribution of $\mathbf{O} = (\mathbf{Y}, \mathbf{D}, \mathbf{X}, \mathbf{Z})$:

$$\begin{aligned} &h_t(z_t, \tilde{z}_t, z_t^\circ, \tilde{z}_t^\circ, x_t, \tilde{x}_t; \mathbf{o}^{t-1}) \\ &\equiv \Pr[Y_t = 1, D_t = 1|z_t, x_t, \mathbf{o}^{t-1}] + \Pr[Y_t = 1, D_t = 0|z_t^\circ, \tilde{x}_t, \mathbf{o}^{t-1}] \\ &\quad - \Pr[Y_t = 1, D_t = 1|\tilde{z}_t, x_t, \mathbf{o}^{t-1}] - \Pr[Y_t = 1, D_t = 0|\tilde{z}_t^\circ, \tilde{x}_t, \mathbf{o}^{t-1}] \end{aligned}$$

for $t \geq 1$, where $(\mathbf{Y}^0, \mathbf{D}^0, \mathbf{X}^0, \mathbf{Z}^0)$ is understood to mean that there is no conditioning. Let $P(z_t, \mathbf{o}^{t-1}) \equiv \Pr[D_t = 1|z_t, \mathbf{o}^{t-1}]$.

Lemma 5.1. Suppose **Assumptions C, SX and RS** hold. For each t and \mathbf{o}^{t-1} , suppose $(z_t, \tilde{z}_t, z_t^\circ, \tilde{z}_t^\circ)$ are such that

$$P(z_t, \mathbf{o}^{t-1}) = P(z_t^\circ, \mathbf{o}^{t-1}) \neq P(\tilde{z}_t, \mathbf{o}^{t-1}) = P(\tilde{z}_t^\circ, \mathbf{o}^{t-1}). \tag{5.1}$$

Then, for given (x_t, \tilde{x}_t) , the sign of $h_t(z_t, \tilde{z}_t, z_t^\circ, \tilde{z}_t^\circ, x_t, \tilde{x}_t; \mathbf{o}^{t-1})$ is equal to the sign of $\mu_t(\mathbf{y}^{t-1}, 1, x_t) - \mu_t(\mathbf{y}^{t-1}, 0, \tilde{x}_t)$.

¹³ In a static scenario, Balat and Han (2020) motivate this reverse exclusion restriction using the notion of externalities. In their setting where multiple treatments are strategically chosen (e.g., firms' entry decisions), factors that determine the outcome (e.g., pollution) are assumed not to appear in the firms' payoff functions.

¹⁴ In their static model, Vytlačil and Yildiz (2007) impose rank invariance, although, as they note, their analysis can be extended to a model with rank similarity.

¹⁵ In fact, we can further relax **Assumption RS** by allowing $U_t(d_t)$ to be a function of x_t from the outset; see Remark 5.2.

The requirement (5.1) is easily met as long as Z_t contains a relevant instrument and, e.g., by choosing $z_t = z_t^\circ$ and $\tilde{z}_t = \tilde{z}_t^\circ$. For the analysis of this paper which deals with a dynamic model, it is convenient to define the \mathbf{U} -set and \mathbf{V} -set, namely the sets of histories of the unobservable variables that determine the outcomes and treatments, respectively. To focus our attention on the dependence of the potential outcomes on the unobservables, we iteratively define the potential outcome given (\mathbf{d}, \mathbf{x}) as

$$Y_t(\mathbf{d}^t, \mathbf{x}^t) \equiv 1\{\mu_t(\mathbf{Y}^{t-1}(\mathbf{d}^{t-1}, \mathbf{x}^{t-1}), d_t, x_t) \geq U_t(d_t)\}$$

for $t \geq 2$, with $Y_1(d_1, x_1) = 1\{\mu_1(0, d_1, x_1) \geq U_1(d_1)\}$. Now, define the set of $\mathbf{U}^t(\mathbf{d}^t)$ as

$$\mathcal{U}^t(\mathbf{d}^t, \mathbf{y}^t) \equiv \mathcal{U}^t(\mathbf{d}^t, \mathbf{y}^t; \mathbf{x}^t) \equiv \{\mathbf{U}^t(\mathbf{d}^t) : y_s = Y_s(\mathbf{d}^s, \mathbf{x}^s) \text{ for all } s \leq t\}.$$

for $t \geq 1$. Then, $\mathbf{Y}^t = \mathbf{y}^t$ if and only if $\mathbf{U}^t(\mathbf{d}^t) \in \mathcal{U}^t(\mathbf{d}^t, \mathbf{y}^t; \mathbf{x}^t)$, conditional on $(\mathbf{D}^t, \mathbf{X}^t) = (\mathbf{d}^t, \mathbf{x}^t)$. The \mathbf{V} -set $\mathcal{V}^t(\mathbf{d}^t, \mathbf{u}^{t-1}) \equiv \mathcal{V}^t(\mathbf{d}^t, \mathbf{u}^{t-1}; \mathbf{z}^t, \mathbf{x}^{t-1})$ is similarly defined within the proof of Lemma 5.1 in Appendix C. Then, $\mathbf{D}^t = \mathbf{d}^t$ if and only if $\mathbf{V}^t \in \mathcal{V}^t(\mathbf{d}^t, \mathbf{U}^{t-1}(\mathbf{d}^{t-1}))$, conditional on $(\mathbf{Z}^t, \mathbf{X}^{t-1}) = (\mathbf{z}^t, \mathbf{x}^{t-1})$. Let $\pi_t \equiv \pi_t(\mathbf{y}^{t-1}, \mathbf{d}^{t-1}, \mathbf{z}^t) = \pi_t(\mathbf{y}^{t-1}, \mathbf{d}^{t-1}, \mathbf{z}^{t-1}, z_t^\circ)$ and $\tilde{\pi}_t \equiv \pi_t(\mathbf{y}^{t-1}, \mathbf{d}^{t-1}, \mathbf{z}^{t-1}, \tilde{z}_t) = \pi_t(\mathbf{y}^{t-1}, \mathbf{d}^{t-1}, \mathbf{z}^{t-1}, \tilde{z}_t^\circ)$ for abbreviation. Given the sets of unobservables, what we show in the proof of this lemma is that, under Assumptions C and SX,

$$\begin{aligned} & h_t(z_t, \tilde{z}_t, z_t^\circ, \tilde{z}_t^\circ, x_t, \tilde{x}_t; \mathbf{o}^{t-1}) \\ &= \Pr\{U_t(1) \leq \mu_t(\mathbf{y}^{t-1}, 1, x_t), \tilde{\pi}_t \leq V_t \leq \pi_t | \mathcal{V}^{t-1}(\mathbf{d}^{t-1}, \mathbf{U}^{t-2}(\mathbf{d}^{t-2})), \mathcal{U}^{t-1}(\mathbf{d}^{t-1}, \mathbf{y}^{t-1})\} \\ & \quad - \Pr\{U_t(0) \leq \mu_t(\mathbf{y}^{t-1}, 0, \tilde{x}_t), \tilde{\pi}_t \leq V_t \leq \pi_t | \mathcal{V}^{t-1}(\mathbf{d}^{t-1}, \mathbf{U}^{t-2}(\mathbf{d}^{t-2})), \mathcal{U}^{t-1}(\mathbf{d}^{t-1}, \mathbf{y}^{t-1})\}, \end{aligned}$$

the sign of which identifies the sign of $\mu_t(\mathbf{y}^{t-1}, 1, x_t) - \mu_t(\mathbf{y}^{t-1}, 0, \tilde{x}_t)$ by Assumption RS. For example, when this quantity is zero, then $\mu_t(\mathbf{y}^{t-1}, 1, x_t) - \mu_t(\mathbf{y}^{t-1}, 0, \tilde{x}_t) = 0$.

For the point identification of the ARSF's, the final assumption we introduce concerns the variation of the exogenous variables (\mathbf{Z}, \mathbf{X}) .¹⁶ Define the following sets:

$$\mathcal{S}_t(d_t, \mathbf{y}^{t-1}) \equiv \{(x_t, \tilde{x}_t) : \mu_t(\mathbf{y}^{t-1}, d_t, x_t) = \mu_t(\mathbf{y}^{t-1}, \tilde{d}_t, \tilde{x}_t) \text{ for } \tilde{d}_t \neq d_t\}, \tag{5.2}$$

$$\begin{aligned} \mathcal{T}_t(\mathbf{d}^{t-1}, \mathbf{y}^{t-1}; \mathbf{x}_{-t}, \mathbf{z}_{-t}) &\equiv \{(x_t, \tilde{x}_t) : \exists(z_t, \tilde{z}_t, z_t^\circ, \tilde{z}_t^\circ) \text{ such that (5.1) holds,} \\ & \quad (x_t, z_t), (\tilde{x}_t, z_t^\circ), (x_t, \tilde{z}_t), (\tilde{x}_t, \tilde{z}_t^\circ) \in \text{Supp}(X_t, Z_t | \mathbf{x}_{-t}, \mathbf{z}_{-t})\}, \end{aligned} \tag{5.3}$$

$$\mathcal{X}_t(\mathbf{d}^t, \mathbf{y}^{t-1}; \mathbf{x}_{-t}, \mathbf{z}_{-t}) \equiv \{x_t : \exists \tilde{x}_t \text{ with } (x_t, \tilde{x}_t) \in \mathcal{S}_t(d_t, \mathbf{y}^{t-1}) \cap \mathcal{T}_t(\mathbf{d}^{t-1}, \mathbf{y}^{t-1}; \mathbf{x}_{-t}, \mathbf{z}_{-t})\}, \tag{5.4}$$

$$\mathcal{X}_t(\mathbf{d}^t; \mathbf{x}_{-t}, \mathbf{z}_{-t}) \equiv \bigcap_{\mathbf{y}^{t-1}} \mathcal{X}_t(\mathbf{d}^t, \mathbf{y}^{t-1}; \mathbf{x}_{-t}, \mathbf{z}_{-t}), \tag{5.5}$$

where (5.2) is related to the sufficient variation of X_t and (5.3) is related to the rectangular variation of (X_t, Z_t) .

Assumption SP. For each t and \mathbf{d}^t , $\Pr\{X_t \in \mathcal{X}_t(\mathbf{d}^t; \mathbf{x}_{-t}, \mathbf{z}_{-t}) | \mathbf{x}_{-t}, \mathbf{z}_{-t}\} > 0$ almost everywhere.

This assumption requires that X_t varies sufficiently to achieve $\mu_t(\mathbf{y}^{t-1}, d_t, x_t) = \mu_t(\mathbf{y}^{t-1}, \tilde{d}_t, \tilde{x}_t)$, while holding Z_t to satisfy either $P(z_t, \mathbf{o}^{t-1}) = P(z_t^\circ, \mathbf{o}^{t-1})$ or $P(\tilde{z}_t, \mathbf{o}^{t-1}) = P(\tilde{z}_t^\circ, \mathbf{o}^{t-1})$, conditional on $(\mathbf{X}_{-t}, \mathbf{Z}_{-t})$. This is a dynamic version of the support assumption found in Vytlacil and Yildiz (2007). Although Assumption SP requires sufficient rectangular variation in (X_t, Z_t) , it clearly differs from the large variation assumptions in, e.g., Heckman and Navarro (2007) and Heckman et al. (2016). These papers employ identification-at-infinity arguments in each period that the support of explained variation (i.e., $\mu_t(\cdot)$ in our notation) is no smaller than the support of unobservables. On the other hand, Assumption SP only requires the existence of variation that equates $\mu_t(\cdot)$ for two different values of D_t . Apparently, this is trivially satisfied with the former assumption of large support.¹⁷ To discuss this comparison further, consider two cases where Assumption SP may hold: (a) X_t is not included in Z_t , although X_{t-1} can be included; or (b) X_t is an element in Z_t , i.e., $Z_t = (X_t, Z_{1t})$. Case (a) yields what we call two-way exclusion, i.e., Z_t is excluded from the outcome equation and X_t is excluded from the selection equation. Case (b) yields the standard exclusion restriction for Z_{1t} .

In Case (a), Assumption SP may hold with discrete Z_t provided X_t contains a continuous element. To see this, suppose X_t is scalar and $Z_t \in \{0, 1\}$. Suppose the index function is linear: $\mu_t(\mathbf{y}^{t-1}, d_t, x_t) = \alpha_t' \mathbf{y}^{t-1} + \beta_t d_t + \gamma_t x_t$. Then, we want x_t and \tilde{x}_t such that $\alpha_t' \mathbf{y}^{t-1} + \beta_t + \gamma_t x_t = \alpha_t' \mathbf{y}^{t-1} + \gamma_t \tilde{x}_t$, or equivalently, $\tilde{x}_t - x_t = \beta_t / \gamma_t$. Suppose $\text{Supp}(\mathbf{X}, \mathbf{Z}) = \text{Supp}(\mathbf{X}) \times \text{Supp}(\mathbf{Z})$ and $\text{Supp}(\mathbf{Z}) = \prod_{t=1}^T \text{Supp}(Z_t) = \{0, 1\}^T$, which can be motivated by the sequential randomization imposed on Z_t . Then, $\mathcal{X}_t(\mathbf{d}^t; \mathbf{x}_{-t}, \mathbf{z}_{-t}) = \mathcal{X}_t(\mathbf{x}_{-t})$ can be simply expressed as

$$\begin{aligned} & \{x_t : \exists \tilde{x}_t \text{ s.t. } h_t(z_t, \tilde{z}_t, x_t, \tilde{x}_t; \mathbf{o}^{t-1}) = 0 \text{ for } x_t, \tilde{x}_t \in \text{supp}(X_t | \mathbf{x}_{-t})\} \\ &= \{x_t : \exists \tilde{x}_t \text{ s.t. } \tilde{x}_t - x_t = \beta_t / \gamma_t \text{ for } x_t, \tilde{x}_t \in \text{supp}(X_t | \mathbf{x}_{-t})\}, \end{aligned}$$

¹⁶ Without this additional assumption, the sign of $\mu_t(\mathbf{y}^{t-1}, 1, x_t) - \mu_t(\mathbf{y}^{t-1}, 0, \tilde{x}_t)$ itself is already useful for calculating bounds on the ARSF's and thus on the ATE's, although we do not pursue it in this paper.

¹⁷ In our setting, it is possible that $\mathcal{X}_t(\mathbf{d}^t; \mathbf{x}_{-t}, \mathbf{z}_{-t})$ is nonempty even when Z_t is discrete, as long as X_t contains continuous elements with sufficient support (Vytlacil and Yildiz, 2007). In Heckman and Navarro (2007), Heckman et al. (2016) and the present paper, the support requirement is conditional on the exogenous variables in other periods; see also Cameron and Heckman (1998).

where $h_t(z_t, \bar{z}_t, x_t, \bar{x}_t; \mathbf{o}^{t-1}) \equiv h_t(z_t, \bar{z}_t, z_t^\circ, \bar{z}_t^\circ, x_t, \bar{x}_t; \mathbf{o}^{t-1})$ as $z_t = z_t^\circ$ and $\bar{z}_t = \bar{z}_t^\circ$ in this case. Let \underline{x}_t and \bar{x}_t be the smallest and largest values of x_t in $\text{Supp}(X_t|\mathbf{x}_{-t})$. When $T = 2$, **Assumption SP** is satisfied (i.e., $\mathcal{X}_t(\mathbf{d}^t; \mathbf{x}_{-t}, \mathbf{z}_{-t})$ is nonempty for $t = 1, 2$) if

$$\bar{x}_1 - \underline{x}_1 \geq \beta_1/\gamma_1 \text{ and } \bar{x}_2 - \underline{x}_2 \geq \beta_2/\gamma_2.$$

That is, as long as X_t is a relevant variable in explaining Y_t (so that γ_t is not small) for each t , the support condition can be substantially weaker than an infinite support assumption. In Case (b) where $Z_t = (X_t, Z_{1t})$, **Assumption SP** may hold either (i) with discrete Z_{1t} provided X_t contains two continuous elements, or (ii) with continuous Z_{1t} and X_t . The support requirement in this case is analogous to that in Case (a), which we discuss in **Appendix B**. In each of these cases, each period’s exclusion restrictions are of the same type as those in Vytlacil and Yildiz (2007), except for being conditioned on exogenous variables of other periods. Note that **Assumption SP** is testable because the sets defined above have empirical analogs, according to **Lemma 5.1**.

Let $\mathcal{X}_t(\mathbf{d}^t; \mathbf{x}_{-t}) \equiv \{x_t : x_t \in \mathcal{X}_t(\mathbf{d}^t; \mathbf{x}_{-t}, \mathbf{z}_{-t}) \text{ for some } \mathbf{z}_{-t} \in \text{Supp}(\mathbf{Z}_{-t}|\mathbf{x}_{-t})\}$ and $\mathcal{X}(\mathbf{d}) \equiv \{\mathbf{x} : x_t \in \mathcal{X}_t(\mathbf{d}^t; \mathbf{x}_{-t}) \text{ for some } (x_{t+1}, \dots, x_T) \in \text{Supp}(\mathbf{X}_{t+1:T}|\mathbf{x}_{-t})\}$, for $t \geq 1$, which sequentially collect $x_t \in \mathcal{X}_t(\mathbf{d}^t; \mathbf{x}_{-t}, \mathbf{z}_{-t})$ for all t . We are now ready to state the main identification result.

Theorem 5.1. Under **Assumptions C, SX, RS and SP**, $E[Y_T(\mathbf{d})|\mathbf{x}]$ is identified for $\mathbf{d} \in \mathcal{D}$ and $\mathbf{x} \in \mathcal{X}(\mathbf{d})$.

Based on **Theorem 5.1**, we can identify the ATE’s.¹⁸ Notably, **Assumptions RS and SP** allow us to identify the ATE rather than versions of the local ATE (Imbens and Angrist, 1994), without invoking an identification-at-infinity argument. Since the identification of all $E[Y_t(\mathbf{d})|\mathbf{x}^t]$ ’s can be shown analogously to **Theorem 5.1**, we can identify the optimal treatment regimes $\mathbf{d}^*(\mathbf{x})$ and $\mathbf{d}^\dagger(\mathbf{x})$ as well.

Corollary 5.1. Under **Assumptions C, SX, RS and SP**, $ATE(\mathbf{d}, \tilde{\mathbf{d}}; \mathbf{x})$ is identified for $\mathbf{d}, \tilde{\mathbf{d}} \in \mathcal{D}$ and $\mathbf{x} \in \mathcal{X}(\mathbf{d}) \cap \mathcal{X}(\tilde{\mathbf{d}})$, and $\mathbf{d}^*(\mathbf{x})$ and $\mathbf{d}^\dagger(\mathbf{x})$ are identified for $\mathbf{x} \in \bigcap_{\mathbf{d} \in \mathcal{D}} \mathcal{X}(\mathbf{d})$.

We sketch the identification analysis here; the full proof of **Theorem 5.1** is found in **Appendix C**. We consider the identification of $E[Y_T(\mathbf{d})|\mathbf{x}, \mathbf{z}]$, since $E[Y_T(\mathbf{d})|\mathbf{x}] = E[Y_T(\mathbf{d})|\mathbf{x}, \mathbf{z}]$ by **Assumption SX**.¹⁹ As the first step of identifying $E[Y_T(\mathbf{d})|\mathbf{x}, \mathbf{z}]$ for given $\mathbf{d} = (d_1, \dots, d_T)$, $\mathbf{x} = (x_1, \dots, x_T)$ and $\mathbf{z} = (z_1, \dots, z_T)$, we apply the result of **Lemma 5.1**. Fix $t \geq 2$ and $\mathbf{y}^{t-1} \in \{0, 1\}^{t-1}$. Suppose x'_t is such that $\mu_t(\mathbf{y}^{t-1}, d_t, x_t) = \mu_t(\mathbf{y}^{t-1}, d'_t, x'_t)$ with $d'_t \neq d_t$ by applying **Lemma 5.1**. The existence of x'_t is guaranteed by **Assumption SP**, as $x_t \in \mathcal{X}_t(\mathbf{d}^t, \mathbf{y}^{t-1}; \mathbf{x}_{-t}, \mathbf{z}_{-t}) \subset \mathcal{X}_t(\mathbf{d}^t; \mathbf{x}_{-t}, \mathbf{z}_{-t})$. The implication of $\mu_t(\mathbf{y}^{t-1}, d_t, x_t) = \mu_t(\mathbf{y}^{t-1}, d'_t, x'_t)$ for relevant \mathbf{U} -sets is as follows: Analogous to the \mathbf{U} -set defined earlier, define

$$\mathcal{U}^t(\mathbf{d}^t, y_t) \equiv \mathcal{U}^t(\mathbf{d}^t, y_t; \mathbf{x}^t) \equiv \{\mathbf{U}^t(\mathbf{d}^t) : y_t = Y_t(\mathbf{d}^t, \mathbf{x}^t)\}.$$

Then, by definition, $\mathbf{U} \in \mathcal{U}(\mathbf{d}, y_t; \mathbf{x})$ is equivalent to $\mathbf{U} \in \mathcal{U}^t(d'_t, \mathbf{d}_{-t}, y_t; x'_t, \mathbf{x}_{-t})$ conditional on $\mathbf{Y}^{t-1}(\mathbf{d}^{t-1}, \mathbf{x}^{t-1}) = \mathbf{y}^{t-1}$.²⁰ Based on this result, we equate the unobserved quantity $E[Y_T(\mathbf{d})|\mathbf{x}, \mathbf{z}, \mathbf{y}^{t-1}, \mathbf{d}^{t-1}, d'_t]$ with a quantity that partly matches the assigned treatment and the observed treatment as follows. First, we can show that

$$\begin{aligned} & E[Y_T(\mathbf{d})|\mathbf{x}, \mathbf{z}, \mathbf{y}^{t-1}, \mathbf{d}^{t-1}, d'_t] \\ &= \Pr \left[\mathbf{U}(\mathbf{d}) \in \mathcal{U}^T(\mathbf{d}, 1; \mathbf{x}) \mid \begin{array}{l} \mathbf{U}^{t-1}(\mathbf{d}^{t-1}) \in \mathcal{U}^{t-1}(\mathbf{d}^{t-1}, \mathbf{y}^{t-1}), \\ \mathbf{V}^t \in \mathcal{V}^t(\mathbf{d}^{t-1}, d'_t, \mathbf{U}^{t-1}(\mathbf{d}^{t-1})) \end{array} \right] \end{aligned}$$

for $t \geq 2$, by **Assumption SX**. Then, by **Assumption RS** and the discussion above, this quantity is shown to be equal to

$$\begin{aligned} & \Pr \left[\mathbf{U}(d'_t, \mathbf{d}_{-t}) \in \mathcal{U}^T(d'_t, \mathbf{d}_{-t}, 1; x'_t, \mathbf{x}_{-t}) \mid \begin{array}{l} \mathbf{U}^{t-1}(\mathbf{d}^{t-1}) \in \mathcal{U}^{t-1}(\mathbf{d}^{t-1}, \mathbf{y}^{t-1}), \\ \mathbf{V}^t \in \mathcal{V}^t(\mathbf{d}^{t-1}, d'_t, \mathbf{U}^{t-1}(\mathbf{d}^{t-1})) \end{array} \right] \\ &= E[Y_T(d'_t, \mathbf{d}_{-t})|x'_t, \mathbf{x}_{-t}, \mathbf{z}, \mathbf{y}^{t-1}, \mathbf{d}^{t-1}, d'_t], \end{aligned} \tag{5.6}$$

by **Assumption SX**. Note that this last quantity is still unobserved, since d_s for $s \geq t + 1$ are not realized treatments; e.g., when $T = 3$ and $t = 2$,

$$E[Y_3(\mathbf{d})|\mathbf{x}, \mathbf{z}, y_1, d_1, d'_2] = E[Y_3(d_1, d'_2, d_3)|x_1, x'_2, x_3, \mathbf{z}, y_1, d_1, d'_2].$$

The quantity, however, will be useful in the remaining proof where we use mathematical induction to recover $E[Y_T(\mathbf{d})|\mathbf{x}, \mathbf{z}]$; see **Appendix C** for this main part of the proof. Recall the abbreviations $\mathcal{V}^t(\mathbf{d}^{t-1}, d'_t, \mathbf{U}^{t-1}(\mathbf{d}^{t-1})) \equiv$

¹⁸ The identification of the unconditional ATE, $E[Y_T(\mathbf{d}) - Y_T(\tilde{\mathbf{d}})]$, requires an additional support condition that $\mathcal{X}(\mathbf{d}) \cap \mathcal{X}(\tilde{\mathbf{d}}) = \text{supp}(\mathbf{X})$.

¹⁹ When we are to identify the average potential outcome at t instead, the conditioning variables we use are the vectors of exogenous variables up to t , i.e., $E[Y_t(\mathbf{d}^t)|\mathbf{x}^t, \mathbf{z}^t]$. Then the entire proof can be easily modified based on this expression.

²⁰ The subsequent analysis is substantially simplified when $\mu_t(\mathbf{y}^{t-1}, d_t, x_t) = \mu_t(\mathbf{y}^{t-1}, d'_t, x'_t)$ is satisfied for all \mathbf{y}^{t-1} , but this situation is unlikely to occur. Therefore, it is important to condition on $\mathbf{Y}^{t-1}(\mathbf{d}^{t-1}, \mathbf{x}^{t-1})$ in the analysis.

$\mathcal{V}^t(\mathbf{d}^{t-1}, d'_t, \mathbf{U}^{t-1}(\mathbf{d}^{t-1}); \mathbf{z}^t, \mathbf{x}^{t-1})$ and $\mathcal{U}^{t-1}(\mathbf{d}^{t-1}, \mathbf{y}^{t-1}) \equiv \mathcal{U}^{t-1}(\mathbf{d}^{t-1}, \mathbf{y}^{t-1}; \mathbf{x}^{t-1})$. That is, in the derivation of (5.6), the key is to consider the average potential outcome for a group of individuals that is defined by the treatments at time t or earlier and the lagged outcome, for which x_t is excluded.

The proof of Theorem 5.1 is constructive in that it provides a closed-form expression for $E[Y_T(\mathbf{d})|\mathbf{x}]$ in an iterative manner, which can immediately be used for estimation. For concreteness, we provide an expression for $E[Y_T(\mathbf{d})|\mathbf{x}]$ when $T = 2$ and binary Z_t . Define

$$h_t^{d_t}(x_t; \mathbf{y}^{t-1}) \equiv h_t^{d_t}(x_t; \mathbf{z}^{t-1}, \mathbf{x}^{t-1}, \mathbf{d}^{t-1}, \mathbf{y}^{t-1}) \\ \equiv \Pr[Y_t = 1, D_t = d_t | Z_t = 1, \mathbf{z}^{t-1}, \mathbf{x}^t, \mathbf{d}^{t-1}, \mathbf{y}^{t-1}] - \Pr[Y_t = 1, D_t = d_t | Z_t = 0, \mathbf{z}^{t-1}, \mathbf{x}^t, \mathbf{d}^{t-1}, \mathbf{y}^{t-1}]$$

and

$$\lambda_t(x_t; \mathbf{y}^{t-1}) \equiv \{\tilde{x}_t : h_t^{d_t}(x_t; \mathbf{y}^{t-1}) + h_t^{d'_t}(\tilde{x}_t; \mathbf{y}^{t-1}) = 0\}$$

with $\lambda_1(x_1) \equiv \lambda_1(x_1; y_0)$. By Lemma 5.1, x'_t satisfies $\mu_t(\mathbf{y}^{t-1}, \mathbf{d}^t, \mathbf{x}^t) = \mu_t(\mathbf{y}^{t-1}, \mathbf{d}^{t-1}, d'_t, \mathbf{x}^{t-1}, x'_t)$ if and only if $x'_t \in \lambda_t(x_t; \mathbf{y}^{t-1})$. Then, our identification result suggests that

$$E[Y_2(\mathbf{d})|\mathbf{x}] = \int \{P[\mathbf{d}|\mathbf{x}, \mathbf{z}]E[Y_2|\mathbf{x}, \mathbf{z}, \mathbf{d}] + P[d_1, d'_2|\mathbf{x}, \mathbf{z}]m_{2,d_1,d'_2} \\ + P[d'_1, d_2|\mathbf{x}, \mathbf{z}]E[Y_2|\lambda_1(x_1), x_2, \mathbf{z}, d'_1, d_2] + P[d'_1, d'_2|\mathbf{x}, \mathbf{z}]m_{2,d'_1,d'_2}\}dF_{\mathbf{z}|\mathbf{x}}, \tag{5.7}$$

where

$$m_{2,d_1,d'_2} \equiv P[y_1|\mathbf{x}, \mathbf{z}, d_1, d'_2]E[Y_2|x_1, \lambda_2(x_2; y_1), \mathbf{z}, y_1, d_1, d'_2] \\ + P[y'_1|\mathbf{x}, \mathbf{z}, d_1, d'_2]E[Y_2|x_1, \lambda_2(x_2; y'_1), \mathbf{z}, y'_1, d_1, d'_2], \\ m_{2,d'_1,d'_2} \equiv P[y_1|\mathbf{x}, \mathbf{z}, d'_1, d'_2]E[Y_2|\lambda_1(x_1), \lambda_2(x_2; y_1), \mathbf{z}, y_1, d'_1, d'_2] \\ + P[y'_1|\mathbf{x}, \mathbf{z}, d'_1, d'_2]E[Y_2|\lambda_1(x_1), \lambda_2(x_2; y'_1), \mathbf{z}, y'_1, d'_1, d'_2].$$

The aggregation with respect to $\mathbf{Z} = \mathbf{z}$ conditional on $\mathbf{X} = \mathbf{x}$ is to improve efficiency. In Appendix A, we discuss further estimation and inference strategies for the parameter $E[Y_2(\mathbf{d})|\mathbf{x}]$.

Remark 5.1. The assumption that the initial condition Y_0 is exogenously determined is not necessary but imposed for convenience. Such an assumption appears in, e.g., Heckman and Navarro (2007). In an alternative setting where Y_0 is endogenously determined in the model, a similar identification analysis as in this section can be followed by modifying Assumption SX. We may consider two alternatives depending upon whether Y_0 is observable or not: (a) $(\mathbf{U}(\mathbf{d}), \mathbf{V})$ and (\mathbf{Z}, \mathbf{X}) are independent conditional on Y_0 ; or (b) $(\mathbf{U}(\mathbf{d}), \mathbf{V}, Y_0)$ and (\mathbf{Z}, \mathbf{X}) are independent. First, recall that each of these statements is “conditional on other covariates”. The assumption (a) can be imposed when Y_0 is observable, maybe because $t = 1$ is not the start of sample period. The assumption (b) can be imposed when Y_0 is unobservable, maybe because $t = 1$ is the start of sample period and the logical start of the process. The analysis in these alternative scenarios is omitted as it is a straightforward extension of the current one. In this analysis, there is no need to assume the distribution of initial conditions, unlike in the literature on dynamic models with random effects. Still, we recover certain treatment effects, unlike in the literature on nonseparable models with unobservable individual effects where, in general, partial effects are hard to recover. The trade-off is that we require variables that are independent of the individual effects, even though other covariates are allowed not to be.

Remark 5.2. In order to define the \mathbf{U} -set, recall that we use an alternative potential outcome $Y_t(\mathbf{d}^t, \mathbf{x}^t) = \mu_t(\mathbf{Y}^{t-1}(\mathbf{d}^{t-1}, \mathbf{x}^{t-1}), d_t, x_t, U_t(d_t))$. Motivated from this, we may consider a structural model that adds another dimension for heterogeneity by allowing $U_t(d_t)$ to be a function of x_t as well:

$$Y_t(\mathbf{d}^t, \mathbf{x}^t) = \mu_t(\mathbf{Y}^{t-1}(\mathbf{d}^{t-1}, \mathbf{x}^{t-1}), d_t, x_t, U_t(d_t, x_t)).$$

Given this extension, we can relax Assumption RS and impose that $\{\mathbf{U}(d_t, \mathbf{d}_{-t}, x_t, \mathbf{x}_{-t})\}_{d_t, x_t}$ are identically distributed conditional on \mathbf{V}^t . The current Assumption RS can be viewed as requiring rank invariance in terms of x_t , while it allows rank similarity in d_t .

6. Treatment effects on transitions

In fact, the identification strategy introduced in the previous section can tackle a more general problem. In this section, we extend the identification analysis of the ATE (Theorem 5.1 and Corollary 5.1) and show identification of the transition-specific ATE. Given the vector $\mathbf{Y}(\mathbf{d}) \equiv (Y_1(\mathbf{d}), \dots, Y_T(\mathbf{d}))$ of potential outcomes, let $\mathbf{Y}_-(\mathbf{d}) \equiv (Y_{t_1}(\mathbf{d}), \dots, Y_{t_L}(\mathbf{d})) \in \mathcal{Y}_- \subseteq \{0, 1\}^L$ be its $1 \times L$ subvector, where $t_1 < t_2 < \dots < t_L \leq T - 1$ and $L < T$. Then, the transition-specific ATE can be defined as $E[Y_T(\mathbf{d})|\mathbf{Y}_-(\mathbf{d}) = \mathbf{y}_-, \mathbf{X} = \mathbf{x}] - E[Y_T(\mathbf{d})|\mathbf{Y}_-(\mathbf{d}) = \mathbf{y}_-, \mathbf{X} = \mathbf{x}^*]$ for some sequences \mathbf{d} and \mathbf{d}^* .

Theorem 6.1. Under Assumptions C, SX, RS and SP, for each $\mathbf{y}_-, E[Y_T(\mathbf{d})|\mathbf{Y}_-(\mathbf{d}) = \mathbf{y}_-, \mathbf{X} = \mathbf{x}] - E[Y_T(\tilde{\mathbf{d}})|\mathbf{Y}_-(\tilde{\mathbf{d}}) = \mathbf{y}_-, \mathbf{X} = \mathbf{x}]$ is identified for $\mathbf{d}, \tilde{\mathbf{d}} \in \mathcal{D}$ and $\mathbf{x} \in \mathcal{X}(\mathbf{d}) \cap \mathcal{X}(\tilde{\mathbf{d}})$.

The proof of this theorem extends that of Theorem 5.1; see Appendix C.²¹ The transition-specific ATE defined in Theorem 6.1 concerns a transition from a state that is specified by the value of the vector of previous potential outcomes, $\mathbf{Y}_-(\mathbf{d})$. When $Y_T(\mathbf{d})$ is binary, $E[Y_T(\mathbf{d})|\mathbf{Y}_-(\mathbf{d}) = \mathbf{y}_-, \mathbf{X} = \mathbf{x}]$ can be viewed as a generalization of the transition probability. As a simple example, with $L = T - 1$, one may be interested in a transition to one state when all previous potential outcomes have stayed in the other state until $T - 1$. When $L = 1$ with $\mathbf{Y}_-(\mathbf{d}) = Y_{T-1}(\mathbf{d})$, the transition-specific ATE becomes $\Pr[Y_T(\mathbf{d}) = 1|Y_{T-1}(\mathbf{d}) = 0] - \Pr[Y_T(\tilde{\mathbf{d}}) = 1|Y_{T-1}(\tilde{\mathbf{d}}) = 0]$, which is a particular example of the treatment effect on the transition probability. The treatment effects on transitions have been studied by, e.g., Abbring and Van den Berg (2003), Heckman and Navarro (2007), Fredriksson and Johansson (2008) and Vikström et al. (2018).²² Torgovitsky (2019) extends the literature on dynamic binary response models (with no treatment) by considering a counterfactual framework without imposing parametric assumptions. In his framework, the lagged outcome plays the role of a treatment for the current outcome, and the “treatment effect” captures the state dependence. Here, on the other hand, we consider the causal effect of the treatments \mathbf{D} on the joint (or conditional) distribution of the potential outcomes $(\mathbf{Y}_-(\mathbf{d}), Y_T(\mathbf{d}))$.

Let $Y_t(d_t) \equiv \mu_t(\mathbf{Y}^{t-1}, \mathbf{D}^{t-1}, d_t, \mathbf{X}^t, U_t(d_t))$ be the period-specific potential outcome at time t . Since $Y_{t-1} = Y_{t-1}(\mathbf{D}^{t-1})$, the period-specific potential outcome can be expressed as $Y_t(d_t) = Y_t(\mathbf{D}^{t-1}, d_t)$ using the usual potential outcome. As a corollary of the result above, we also identify a related parameter that specifies the previous state by the observed outcome: $E[Y_T(1) - Y_T(0)|Y_{T-1} = y_{T-1}]$.

Corollary 6.1. Under Assumptions C, SX, RS and SP, for each y_{T-1} , $E[Y_T(1)|y_{T-1}, \mathbf{x}] - E[Y_T(0)|y_{T-1}, \mathbf{x}]$ is identified for $\mathbf{x} \in \mathcal{X}(\mathbf{d}) \cap \mathcal{X}(\tilde{\mathbf{d}})$.

The corollary is derived by observing that $Y_T(d_T) = Y_T(\mathbf{D}^{T-1}, d_T)$, and thus

$$E[Y_T(d_T)|y_{T-1}, \mathbf{x}] = \sum_{\mathbf{d}^{T-1} \in \mathcal{D}^{T-1}} \Pr[\mathbf{D}^{T-1} = \mathbf{d}^{T-1}|\mathbf{x}]E[Y_T(\mathbf{d}^{T-1}, d_T)|Y_{T-1}(\mathbf{d}^{T-1}) = y_{T-1}, \mathbf{D}^{T-1} = \mathbf{d}^{T-1}, \mathbf{x}],$$

where each $E[Y_T(\mathbf{d}^{T-1}, d_T)|Y_{T-1}(\mathbf{d}^{T-1}) = y_{T-1}, \mathbf{D}^{T-1} = \mathbf{d}^{T-1}, \mathbf{x}]$ is identified from the iteration at $t = T - 1$ in the proof of Theorem 6.1 by taking $\mathbf{Y}_-(\mathbf{d}) = Y_{T-1}(\mathbf{d}^{T-1})$.

7. Conclusions

In this paper, we consider identification in a nonparametric model for dynamic treatments and outcomes. We introduce a sequence of selection models, replacing the assumption of sequential randomization, which may be hard to justify under partial compliance or in observational settings. We consider treatment and outcome processes of general forms, and avoid making strong assumptions on distribution and functional forms, nor assumptions on rationality. We show that the treatment parameters and optimal treatment regimes are point identified under the exclusion restrictions and sequential rank similarity. We argue that the use of extra exogenous variables is a useful alternative tool for empirical researchers who seek identification of treatment effects in this type of nonseparable models with endogeneity. This source of variation may especially be easy to find and justify in a dynamic setting as in this paper. When the reverse exclusion restriction is violated, we show how to characterize bounds on these parameters.

Appendix A. Estimation and inference

The identification analysis is constructive and naturally suggests an estimation procedure by the sample analog principle. Here we illustrate that by considering (5.7), which can be alternatively expressed as

$$E[Y_2(\mathbf{d})|\mathbf{x}] = \int \{E[1_{\mathbf{d}}Y_2|\mathbf{x}, \mathbf{z}] + E[1_{d_1, d_2'}1_{y_1}Y_2|\lambda_1(x_2; y_1), \lambda_2(x_2; y_1), \mathbf{z}] + E[1_{d_1, d_2'}1_{y_1'}Y_2|\lambda_1(x_2; y_1'), \lambda_2(x_2; y_1'), \mathbf{z}] \\ + E[1_{d_1', d_2}Y_2|\lambda_1(x_1), x_2, \mathbf{z}] \\ + E[1_{d_1', d_2'}1_{y_1}|\mathbf{x}_1, \mathbf{z}]E[Y_2|\lambda_1(x_1), \lambda_2(x_2; y_1), \mathbf{z}, d_1', d_2', y_1] \\ + E[1_{d_1', d_2'}1_{y_1'}|\mathbf{x}_1, \mathbf{z}]E[Y_2|\lambda_1(x_1), \lambda_2(x_2; y_1'), \mathbf{z}, d_1', d_2', y_1']\}dF_{\mathbf{z}|\mathbf{x}},$$

²¹ As before, the parameters in Theorem 6.1 and Corollary 6.1 can be defined for any given period instead of the terminal period T . The identification analysis of such parameters is essentially the same, and thus omitted.

²² The definition of the treatment effect on the transition probability in this paper differs from those defined in the literature on duration models, e.g., that in Vikström et al. (2018). Since Vikström et al. (2018)'s main focus is on Y_t that is irreversible, they define a different treatment parameter that yields a specific interpretation under dynamic selection; see their paper for details. In addition, they assume sequential randomization and that treatments are assigned earlier than the transition of interest.

where $1_{\mathbf{d}} \equiv 1[\mathbf{D} = \mathbf{d}]$ for generic \mathbf{d} , the second term on the right hand side is by $E[1_{d_1, d'_2} 1_{y_1} | \mathbf{x}_1, \mathbf{z}] \times E[Y_2 | \mathbf{x}_1, \lambda_2(x_2; y_1), \mathbf{z}, d_1, d'_2, y_1] = E[1_{d_1, d'_2} 1_{y_1} Y_2 | \mathbf{x}_1, \lambda_2(x_2; y_1), \mathbf{z}]$ since $E[1_{d_1, d'_2} 1_{y_1} | \mathbf{x}_1, \mathbf{z}] = E[1_{d_1, d'_2} 1_{y_1} | \mathbf{x}_1, \mathbf{z}]$ (and similarly for the third term), and the fourth term is by $P[d'_1, d_2 | \mathbf{x}, \mathbf{z}] = P[d'_1, d_2 | \mathbf{z}]$. When $\lambda_1(\cdot)$ and $\lambda_2(\cdot; \cdot)$ are assumed to be known, the estimation of $g_{\mathbf{d}}(\mathbf{x}) \equiv E[Y_2(\mathbf{d}) | \mathbf{x}]$ and inference on its functionals can be dealt as a special case of the penalized sieve minimum distance (PSMD) estimation framework of [Chen and Pouzo \(2015\)](#) by constructing the following conditional moments:

$$\begin{aligned} \sum_{\mathbf{z} \in \{0,1\}^2} \Pr[\mathbf{Z}_i = \mathbf{z} | \mathbf{x}] \{ & g_{\mathbf{z}}^1(\mathbf{x}) + g_{\mathbf{z}}^2(x_1) + g_{\mathbf{z}}^3(x_1) \\ & + g_{\mathbf{z}}^4(x_2) + g_{\mathbf{z}}^5(x_1) \mu_{\mathbf{z}}^1 + g_{\mathbf{z}}^6(x_1) m_{\mathbf{z}}^2 \} - g_{\mathbf{d}}(\mathbf{x}) = 0, \\ & E[1_{\mathbf{d}} Y_2 | \mathbf{x}, \mathbf{z}] - g_{\mathbf{z}}^1(\mathbf{x}) = 0, \\ & E[1_{d_1, d'_2} 1_{y_1} Y_2 | \mathbf{x}_1, \lambda_2(x_2; y_1), \mathbf{z}] - g_{\mathbf{z}}^2(x_1) = 0, \\ & E[1_{d_1, d'_2} 1_{y_1} Y_2 | \mathbf{x}_1, \lambda_2(x_2; y'_1), \mathbf{z}] - g_{\mathbf{z}}^3(x_1) = 0, \\ & E[1_{d'_1, d_2} Y_2 | \lambda_1(x_1), x_2, \mathbf{z}] - g_{\mathbf{z}}^4(x_2) = 0, \\ & E[1_{d'_1, d'_2} 1_{y_1} | \mathbf{x}_1, \mathbf{z}] - g_{\mathbf{z}}^5(x_1) = 0, \\ & E[1_{d'_1, d'_2} 1_{y'_1} | \mathbf{x}_1, \mathbf{z}] - g_{\mathbf{z}}^6(x_1) = 0, \\ & E[Y_2 | \lambda_1(x_1), \lambda_2(x_2; y_1), \mathbf{z}, d'_1, d'_2, y_1] - m_{\mathbf{z}}^1 = 0, \\ & E[Y_2 | \lambda_1(x_1), \lambda_2(x_2; y'_1), \mathbf{z}, d'_1, d'_2, y'_1] - m_{\mathbf{z}}^2 = 0, \end{aligned}$$

where $g_{\mathbf{d}}(\mathbf{x})$, $\Pr[\mathbf{Z}_i = \mathbf{z} | \cdot]$ and $g_{\mathbf{z}}^j(\cdot)$ for $j \in \{1, \dots, 6\}$ are the nonparametric components, and $m_{\mathbf{z}}^1$ and $m_{\mathbf{z}}^2$ are the parametric components, for $\mathbf{z} \in \{0, 1\}^2$. It is worth noting that, since none of the nonparametric components has endogenous variables as its arguments, there is no ill-posed inverse problem. For a (possibly nonlinear) functional $\phi(\cdot)$ of $g_{\mathbf{d}}(\mathbf{x})$, the plug-in PSMD estimator $\phi(\hat{g}_{\mathbf{d}})$ is asymptotically normal with a consistent sieve variance estimator, which can be used to conduct inference. [Chen and Pouzo \(2015\)](#) also establish asymptotic theory for the sieve quasi likelihood ratio statistic, whose null distribution is tight no matter whether $\phi(g_{\mathbf{d}})$ is \sqrt{n} -estimable (as with a weighted derivative functional) or not (as with a point evaluation functional). Generalized weighted bootstrap can be used to calculate critical values for this test statistic. When the sets $\lambda_1(\cdot)$ and $\lambda_2(\cdot; \cdot)$ are estimated, estimation and inference become analogous to those with the matching estimator in [Heckman et al. \(1998\)](#).

Appendix B. Further Examples of Assumption SP

Continuing the discussion on the support condition in the main text, consider Case (b) where the two-way exclusion restriction does not hold: X_t is an element in Z_t , i.e., $Z_t = (X_t, Z_{1t})$. In this case, [Assumption SP](#) may hold either (i) with discrete Z_{1t} provided X_t contains two continuous elements, or (ii) with continuous Z_{1t} and X_t .²³ Suppose $\text{Supp}(\mathbf{Z}) = \text{Supp}(\mathbf{X}) \times \text{Supp}(\mathbf{Z}_1)$ and $\text{Supp}(\mathbf{Z}_1) = \prod_{t=1}^T \text{Supp}(Z_{1t}) = \{0, 1\}^T$.

First, consider Case (b-i): Z_{1t} is discrete and $X_t = (X_{1t}, X_{2t})$ where each element is continuously distributed. Let $R_t \equiv (Y_t, D_t, Z_t)$ with realization r_t . Suppose each element is scalar. Suppose $\mu_t(\mathbf{y}^{t-1}, d_t, x_t) = \alpha'_t \mathbf{y}^{t-1} + \beta_t d_t + \gamma_t x_t + \gamma_t^1 x_{1t} + \gamma_t^2 x_{2t}$ and $\pi_t(\mathbf{y}^{t-1}, \mathbf{d}^{t-1}, \mathbf{z}^t) = \delta'_t \mathbf{r}^{t-1} + \theta_t^1 x_{1t} + \theta_t^2 x_{2t} + \tau_t z_{1t}$. Then, [Assumption SP](#) is satisfied if

$$\begin{aligned} \bar{x}_{1t} - \underline{x}_{1t} &\geq \beta_t \theta_t^2 / \gamma_t^1 (\theta_t^1 \gamma_t^2 - \theta_t^2 \gamma_t^1), \\ \bar{x}_{2t} - \underline{x}_{2t} &\geq \beta_t \theta_t^1 / \gamma_t^2 (\theta_t^1 \gamma_t^2 - \theta_t^2 \gamma_t^1), \end{aligned}$$

where $\text{Supp}(X_t | \mathbf{X}_{-t} = \mathbf{x}_{-t}) = \text{Supp}(X_{1t} | \mathbf{X}_{-t} = \mathbf{x}_{-t}) \times \text{Supp}(X_{2t} | \mathbf{X}_{-t} = \mathbf{x}_{-t})$, and $(\bar{x}_{1t}, \underline{x}_{1t})$ and $(\bar{x}_{2t}, \underline{x}_{2t})$ are the largest and smallest values of x_t in $\text{Supp}(X_{1t} | \mathbf{X}_{-t} = \mathbf{x}_{-t})$ and $\text{Supp}(X_{2t} | \mathbf{X}_{-t} = \mathbf{x}_{-t})$, respectively.

Now, consider Case (b-ii): Z_{1t} and X_t are continuously distributed. Suppose each of them is scalar. Suppose $\mu_t(\mathbf{y}^{t-1}, d_t, x_t) = \alpha'_t \mathbf{y}^{t-1} + \beta_t d_t + \gamma_t x_t$ and $\pi_t(\mathbf{y}^{t-1}, \mathbf{d}^{t-1}, \mathbf{z}^t) = \delta'_t \mathbf{r}^{t-1} + \theta_t x_t + \tau_t z_{1t}$. Then, [Assumption SP](#) is satisfied if

$$\begin{aligned} \bar{x}_t - \underline{x}_t &\geq \beta_t / \gamma_t, \\ \bar{z}_{1t} - \underline{z}_{1t} &\geq \beta_t \theta_t / \gamma_t \tau_t, \end{aligned}$$

where $(\bar{x}_t, \underline{x}_t)$ and $(\bar{z}_{1t}, \underline{z}_{1t})$ are the largest and smallest values of x_t and z_{1t} in $\text{Supp}(X_t | \mathbf{X}_{-t} = \mathbf{x}_{-t})$ and $\text{Supp}(Z_{1t})$, respectively.

Now we turn to the second example where we assume a linear model for the outcome equation. This example allows continuous Y_t , which is related to the discussion in [Remark 3.3](#). Consider

$$Y_t = \delta'_t \mathbf{Y}^{t-1} + \beta_t D_t + \gamma_t X_t + U_t(D_t). \tag{B.1}$$

²³ One can find similar discussions in Appendix B of [Vytlačil and Yildiz \(2007\)](#).

With this linear model, the derivation for $h_t(z_t, \tilde{z}_t, z_t^\circ, \tilde{z}_t^\circ, x_t, \tilde{x}_t; \mathbf{o}^{t-1})$ is different from the one in Lemma 5.1 but the core features remain. Suppose $z_t = z_t^\circ$ and $\tilde{z}_t = \tilde{z}_t^\circ$ for simplicity. Also, recall $P_t = P(z_t, \mathbf{o}^{t-1}) \equiv \Pr[D_t = 1|z_t, \mathbf{o}^{t-1}]$, $\tilde{P}_t = P(\tilde{z}_t, \mathbf{o}^{t-1}) \equiv \Pr[D_t = 1|\tilde{z}_t, \mathbf{o}^{t-1}]$, $\pi_t \equiv \pi_t(\mathbf{y}^{t-1}, \mathbf{d}^{t-1}, \mathbf{z}^t)$, $\tilde{\pi}_t \equiv \pi_t(\mathbf{y}^{t-1}, \mathbf{d}^{t-1}, \mathbf{z}^{t-1}, \tilde{z}_t)$, $\nu^{t-1} \equiv \nu^{t-1}(\mathbf{d}^{t-1}, \mathbf{U}^{t-2}(\mathbf{d}^{t-2}))$ and $\mathcal{U}^{t-1} \equiv \mathcal{U}^{t-1}(\mathbf{d}^{t-1}, \mathbf{y}^{t-1})$. Then we can derive that

$$\begin{aligned} & h_t(z_t, \tilde{z}_t, x_t, \tilde{x}_t; \mathbf{o}^{t-1}) \\ & \equiv \Pr[Y_t = 1, D_t = 1|z_t, x_t, \mathbf{o}^{t-1}] + \Pr[Y_t = 1, D_t = 0|z_t, \tilde{x}_t, \mathbf{o}^{t-1}] \\ & \quad - \Pr[Y_t = 1, D_t = 1|\tilde{z}_t, x_t, \mathbf{o}^{t-1}] - \Pr[Y_t = 1, D_t = 0|\tilde{z}_t, \tilde{x}_t, \mathbf{o}^{t-1}] \\ & = P_t E[Y_t|D_t = 1, z_t, x_t, \mathbf{o}^{t-1}] + (1 - P_t) E[Y_t|D_t = 0, z_t, \tilde{x}_t, \mathbf{o}^{t-1}] \\ & \quad - \tilde{P}_t E[Y_t|D_t = 1, \tilde{z}_t, x_t, \mathbf{o}^{t-1}] - (1 - \tilde{P}_t) E[Y_t|D_t = 0, \tilde{z}_t, \tilde{x}_t, \mathbf{o}^{t-1}] \\ & = P_t \{ \delta_t' \mathbf{y}^{t-1} + \beta_t + \gamma_t x_t + E[U_t(1)|V_t \leq \pi_t, \nu^{t-1}, \mathcal{U}^{t-1}] \} \\ & \quad + (1 - P_t) \{ \delta_t' \mathbf{y}^{t-1} + \beta_t + \gamma_t \tilde{x}_t + E[U_t(0)|V_t > \pi_t, \nu^{t-1}, \mathcal{U}^{t-1}] \} \\ & \quad - \tilde{P}_t \{ \delta_t' \mathbf{y}^{t-1} + \beta_t + \gamma_t x_t + E[U_t(1)|V_t \leq \tilde{\pi}_t, \nu^{t-1}, \mathcal{U}^{t-1}] \} \\ & \quad - (1 - \tilde{P}_t) \{ \delta_t' \mathbf{y}^{t-1} + \beta_t + \gamma_t \tilde{x}_t + E[U_t(0)|V_t > \tilde{\pi}_t, \nu^{t-1}, \mathcal{U}^{t-1}] \}, \end{aligned}$$

where Assumption EX and the definitions of ν^{t-1} and \mathcal{U}^{t-1} are used in the last equality. Then, we can continue by deriving

$$\begin{aligned} & h_t(z_t, \tilde{z}_t, x_t, \tilde{x}_t; \mathbf{o}^{t-1}) \\ & = (P_t - \tilde{P}_t) (\delta_t' \mathbf{y}^{t-1} + \beta_t + \gamma_t x_t) + \int_{\tilde{\pi}_t}^{\pi_t} E[U_t(1)|V_t = v_t, \nu^{t-1}, \mathcal{U}^{t-1}] dv_t \\ & \quad - (P_t - \tilde{P}_t) (\delta_t' \mathbf{y}^{t-1} + \beta_t + \gamma_t \tilde{x}_t) - \int_{\tilde{\pi}_t}^{\pi_t} E[U_t(0)|V_t = v_t, \nu^{t-1}, \mathcal{U}^{t-1}] dv_t \\ & = (P_t - \tilde{P}_t) (\beta_t + \gamma_t (x_t - \tilde{x}_t)), \end{aligned}$$

where the two integrals are canceled out by Assumption RS. Therefore,

$$h_t(z_t, \tilde{z}_t, x_t, \tilde{x}_t; \mathbf{o}^{t-1}) = 0$$

if and only if $\beta_t + \gamma_t (\tilde{x}_t - x_t) = 0$ or

$$\tilde{x}_t - x_t = \beta_t / \gamma_t.$$

The remaining illustration is the same as before.

Appendix C. Proofs

C.1. Proof of Lemma 5.1

We first define the \mathbf{U} -set and \mathbf{V} -set. The \mathbf{U} -set is defined in the main text. Realizing the dependence of $Y_{s-1}(\mathbf{d}^{s-1}, \mathbf{x}^{s-1})$ on $(\mathbf{U}^{s-1}(\mathbf{d}^{s-1}), \mathbf{x}^{s-1}, \mathbf{d}^{s-1})$, let

$$\pi_s^*(\mathbf{U}^{s-1}(\mathbf{d}^{s-1}), \mathbf{x}^{s-1}, \mathbf{d}^{s-1}, \mathbf{z}^s) \equiv \pi_s(\mathbf{Y}^{s-1}(\mathbf{d}^{s-1}, \mathbf{x}^{s-1}), \mathbf{d}^{s-1}, \mathbf{z}^s),$$

and define the set of \mathbf{V}^t as

$$\nu^t(\mathbf{d}^t, \mathbf{u}^{t-1}) \equiv \nu^t(\mathbf{d}^t, \mathbf{u}^{t-1}; \mathbf{z}^t, \mathbf{x}^{t-1}) \equiv \{ \mathbf{V}^t : d_s = 1 \{ V_s \leq \pi_s^*(\mathbf{u}^{s-1}, \mathbf{x}^{s-1}, \mathbf{d}^{s-1}, \mathbf{z}^s) \} \text{ for all } s \leq t \}$$

for $t \geq 2$. Fix $t \geq 3$. Recall $\mathbf{o}_t = (z_t, x_t, d_t, y_t)$. Given (5.1), consider the case $P(z_t, \mathbf{o}^{t-1}) = P(z_t^\circ, \mathbf{o}^{t-1}) > P(\tilde{z}_t, \mathbf{o}^{t-1}) = P(\tilde{z}_t^\circ, \mathbf{o}^{t-1})$; the opposite case is symmetric. Using the definitions of the sets above, we have

$$\begin{aligned} & \Pr[D_t = 1|z_t, \mathbf{o}^{t-1}] \\ & = \Pr[V_t \leq \pi_t(\mathbf{y}^{t-1}, \mathbf{d}^{t-1}, \mathbf{z}^{t-1}, z_t) | \mathbf{z}^t, \mathbf{x}^{t-1}, \nu^{t-1}(\mathbf{d}^{t-1}, \mathbf{U}^{t-2}(\mathbf{d}^{t-2})), \mathcal{U}^{t-1}(\mathbf{d}^{t-1}, \mathbf{y}^{t-1})] \\ & = \Pr[V_t \leq \pi_t(\mathbf{y}^{t-1}, \mathbf{d}^{t-1}, \mathbf{z}^{t-1}, z_t) | \mathbf{z}^{t-1}, \mathbf{x}^{t-1}, \nu^{t-1}(\mathbf{d}^{t-1}, \mathbf{U}^{t-2}(\mathbf{d}^{t-2})), \mathcal{U}^{t-1}(\mathbf{d}^{t-1}, \mathbf{y}^{t-1})], \end{aligned}$$

where the last equality is by Assumption SX and because of the following: conditional on $(\mathbf{Z}^t, \mathbf{X}^{t-1}) = (\mathbf{z}^t, \mathbf{x}^{t-1})$, (i) $\mathbf{Y}^{t-1} = \mathbf{y}^{t-1}$ is equivalent to $\mathbf{U}^{t-1}(\mathbf{d}^{t-1}) \in \mathcal{U}^{t-1}(\mathbf{d}^{t-1}, \mathbf{y}^{t-1})$ conditional on $\mathbf{D}^{t-1} = \mathbf{d}^{t-1}$; (ii) $\mathbf{D}^{t-1} = \mathbf{d}^{t-1}$ is equivalent of $\mathbf{V}^{t-1} \in \nu^{t-1}(\mathbf{d}^{t-1}, \mathbf{U}^{t-2}(\mathbf{d}^{t-2}))$ conditional on $\mathbf{Y}^{t-1} = \mathbf{y}^{t-1}$. Note that the sets $\nu^{t-1}(\mathbf{d}^{t-1}, \mathbf{U}^{t-2}(\mathbf{d}^{t-2}))$ and $\mathcal{U}^{t-1}(\mathbf{d}^{t-1}, \mathbf{y}^{t-1})$ do not change with the change in z_t . Therefore, a parallel expression can be derived for $\Pr[D_t = 1|\tilde{z}_t, \mathbf{o}^{t-1}]$. Recall

$\pi_t \equiv \pi_t(\mathbf{y}^{t-1}, \mathbf{d}^{t-1}, \mathbf{z}^t) = \pi_t(\mathbf{y}^{t-1}, \mathbf{d}^{t-1}, \mathbf{z}^{t-1}, z_t^\circ)$ and $\tilde{\pi}_t \equiv \pi_t(\mathbf{y}^{t-1}, \mathbf{d}^{t-1}, \mathbf{z}^{t-1}, \tilde{z}_t) = \pi_t(\mathbf{y}^{t-1}, \mathbf{d}^{t-1}, \mathbf{z}^{t-1}, \tilde{z}_t^\circ)$. Then, under Assumption C,

$$\begin{aligned} & 0 < \Pr[D_t = 1 | z_t, \mathbf{o}^{t-1}] - \Pr[D_t = 1 | \tilde{z}_t, \mathbf{o}^{t-1}] \\ & = \Pr[V_t \leq \pi_t | \mathbf{z}^{t-1}, \mathbf{x}^{t-1}, \nu^{t-1}(\mathbf{d}^{t-1}, \mathbf{U}^{t-2}(\mathbf{d}^{t-2})), \mathcal{U}^{t-1}(\mathbf{d}^{t-1}, \mathbf{y}^{t-1})] \\ & \quad - \Pr[V_t \leq \tilde{\pi}_t | \mathbf{z}^{t-1}, \mathbf{x}^{t-1}, \nu^{t-1}(\mathbf{d}^{t-1}, \mathbf{U}^{t-2}(\mathbf{d}^{t-2})), \mathcal{U}^{t-1}(\mathbf{d}^{t-1}, \mathbf{y}^{t-1})], \end{aligned}$$

which implies $\pi_t > \tilde{\pi}_t$. The same argument applies by replacing (z_t, \tilde{z}_t) with $(z_t^\circ, \tilde{z}_t^\circ)$. Next, we have

$$\begin{aligned} & \Pr[Y_t = 1, D_t = 1 | z_t, x_t, \mathbf{o}^{t-1}] \\ & = \Pr[U_t(1) \leq \mu_t(\mathbf{y}^{t-1}, 1, x_t), V_t \leq \pi_t | \mathbf{z}^{t-1}, \mathbf{x}^{t-1}, \nu^{t-1}(\mathbf{d}^{t-1}, \mathbf{U}^{t-2}(\mathbf{d}^{t-2})), \mathcal{U}^{t-1}(\mathbf{d}^{t-1}, \mathbf{y}^{t-1})] \end{aligned}$$

by Assumption SX. Again, note that $\nu^{t-1}(\mathbf{d}^{t-1}, \mathbf{U}^{t-2}(\mathbf{d}^{t-2}))$ and $\mathcal{U}^{t-1}(\mathbf{d}^{t-1}, \mathbf{y}^{t-1})$ do not change with the change in (z_t, x_t) , which is key. Therefore, similar expressions can be derived for the other terms involved in h_t , and we have

$$\begin{aligned} & h_t(z_t, \tilde{z}_t, z_t^\circ, \tilde{z}_t^\circ, x_t, \tilde{x}_t; \mathbf{o}^{t-1}) \\ & = \Pr[U_t(1) \leq \mu_t(\mathbf{y}^{t-1}, 1, x_t), \tilde{\pi}_t \leq V_t \leq \pi_t | \nu^{t-1}(\mathbf{d}^{t-1}, \mathbf{U}^{t-2}(\mathbf{d}^{t-2})), \mathcal{U}^{t-1}(\mathbf{d}^{t-1}, \mathbf{y}^{t-1})] \\ & \quad - \Pr[U_t(0) \leq \mu_t(\mathbf{y}^{t-1}, 0, \tilde{x}_t), \tilde{\pi}_t \leq V_t \leq \pi_t | \nu^{t-1}(\mathbf{d}^{t-1}, \mathbf{U}^{t-2}(\mathbf{d}^{t-2})), \mathcal{U}^{t-1}(\mathbf{d}^{t-1}, \mathbf{y}^{t-1})], \end{aligned}$$

the sign of which identifies the sign of $\mu_t(\mathbf{y}^{t-1}, 1, x_t) - \mu_t(\mathbf{y}^{t-1}, 0, \tilde{x}_t)$ by Assumption RS. The case $t \leq 2$ can be shown analogously with $\nu^1(d_1) \equiv \nu^1(d_1; z_1) \equiv \{V_1 : d_1 = 1\{V_1 \leq \pi_1(0, 0, z_1)\}\}$. \square

C.2. Proof of Theorem 5.1

To begin with, note that $E[Y_T(\mathbf{d})|\mathbf{x}] = \Pr[Y_T(\mathbf{d}) = 1|\mathbf{x}] = \Pr[\mathbf{U}(\mathbf{d}) \in \mathcal{U}^T(\mathbf{d}, 1; \mathbf{x})|\mathbf{x}] = \Pr[\mathbf{U}(\mathbf{d}) \in \mathcal{U}^T(\mathbf{d}, 1; \mathbf{x})|\mathbf{x}, \mathbf{z}] = E[Y_T(\mathbf{d})|\mathbf{x}, \mathbf{z}]$ by Assumption SX. As the first step of identifying $E[Y_T(\mathbf{d})|\mathbf{x}, \mathbf{z}]$ for given $\mathbf{d} = (d_1, \dots, d_T)$, $\mathbf{x} = (x_1, \dots, x_T)$ and $\mathbf{z} = (z_1, \dots, z_T)$, we apply the result of Lemma 5.1. Fix $t \geq 2$ and $\mathbf{y}^{t-1} \in \{0, 1\}^{t-1}$. Suppose x'_t is such that $\mu_t(\mathbf{y}^{t-1}, d_t, x_t) = \mu_t(\mathbf{y}^{t-1}, d'_t, x'_t)$ with $d'_t \neq d_t$ by applying Lemma 5.1. The existence of x'_t is guaranteed by Assumption SP, as $x_t \in \mathcal{X}_t(\mathbf{d}^t, \mathbf{y}^{t-1}; \mathbf{x}_{-t}, \mathbf{z}_{-t}) \subset \mathcal{X}_t(\mathbf{d}^t; \mathbf{x}_{-t}, \mathbf{z}_{-t})$. Then, as discussed in the main text, $\mathbf{U} \in \mathcal{U}(\mathbf{d}, y_T; \mathbf{x})$ is equivalent to $\mathbf{U} \in \mathcal{U}^T(d'_t, \mathbf{d}_{-t}, y_T; x'_t, \mathbf{x}_{-t})$ conditional on $\mathbf{Y}^{t-1}(\mathbf{d}^{t-1}, \mathbf{x}^{t-1})$. Then, for $t \geq 2$,

$$\begin{aligned} & E[Y_T(\mathbf{d})|\mathbf{x}, \mathbf{z}, \mathbf{y}^{t-1}, \mathbf{d}^{t-1}, d'_t] \\ & = \Pr \left[\mathbf{U}(\mathbf{d}) \in \mathcal{U}^T(\mathbf{d}, 1; \mathbf{x}) \mid \begin{matrix} \mathbf{x}, \mathbf{z}, \\ \mathbf{U}^{t-1}(\mathbf{d}^{t-1}) \in \mathcal{U}^{t-1}(\mathbf{d}^{t-1}, \mathbf{y}^{t-1}), \\ \mathbf{V}^t \in \mathcal{V}^t(\mathbf{d}^{t-1}, d'_t, \mathbf{U}^{t-1}(\mathbf{d}^{t-1})) \end{matrix} \right] \\ & = \Pr \left[\mathbf{U}(\mathbf{d}) \in \mathcal{U}^T(\mathbf{d}, 1; \mathbf{x}) \mid \begin{matrix} \mathbf{x}_{-t}, \mathbf{z}, \\ \mathbf{U}^{t-1}(\mathbf{d}^{t-1}) \in \mathcal{U}^{t-1}(\mathbf{d}^{t-1}, \mathbf{y}^{t-1}), \\ \mathbf{V}^t \in \mathcal{V}^t(\mathbf{d}^{t-1}, d'_t, \mathbf{U}^{t-1}(\mathbf{d}^{t-1})) \end{matrix} \right], \end{aligned} \tag{C.1}$$

where the last equality follows from Assumption SX. Then, by Assumption RS and the discussion above, (C.1) is equal to

$$\begin{aligned} & \Pr \left[\mathbf{U}(d'_t, \mathbf{d}_{-t}) \in \mathcal{U}^T(d'_t, \mathbf{d}_{-t}, 1; x'_t, \mathbf{x}_{-t}) \mid \begin{matrix} \mathbf{x}_{-t}, \mathbf{z}, \\ \mathbf{U}^{t-1}(\mathbf{d}^{t-1}) \in \mathcal{U}^{t-1}(\mathbf{d}^{t-1}, \mathbf{y}^{t-1}), \\ \mathbf{V}^t \in \mathcal{V}^t(\mathbf{d}^{t-1}, d'_t, \mathbf{U}^{t-1}(\mathbf{d}^{t-1})) \end{matrix} \right] \\ & = \Pr \left[\mathbf{U}(d'_t, \mathbf{d}_{-t}) \in \mathcal{U}^T(d'_t, \mathbf{d}_{-t}, 1; x'_t, \mathbf{x}_{-t}) \mid \begin{matrix} x'_t, \mathbf{x}_{-t}, \mathbf{z}, \\ \mathbf{U}^{t-1}(\mathbf{d}^{t-1}) \in \mathcal{U}^{t-1}(\mathbf{d}^{t-1}, \mathbf{y}^{t-1}), \\ \mathbf{V}^t \in \mathcal{V}^t(\mathbf{d}^{t-1}, d'_t, \mathbf{U}^{t-1}(\mathbf{d}^{t-1})) \end{matrix} \right] \\ & = E[Y_T(d'_t, \mathbf{d}_{-t})|x'_t, \mathbf{x}_{-t}, \mathbf{z}^t, \mathbf{y}^{t-1}, \mathbf{d}^{t-1}, d'_t], \end{aligned} \tag{C.2}$$

where the first equality is by Assumption SX. We use the result (C.2) in the next step.

First, note that $E[Y_T(\mathbf{d})|\mathbf{x}, \mathbf{z}, \mathbf{y}^{T-1}, \mathbf{d}^T] = E[Y_T|\mathbf{x}, \mathbf{z}, \mathbf{y}^{T-1}, \mathbf{d}^T]$ is trivially identified for any generic values $(\mathbf{d}, \mathbf{x}, \mathbf{z}, \mathbf{y}^{T-1})$. We prove by means of mathematical induction. For given $2 \leq t \leq T - 1$, suppose $E[Y_T(\mathbf{d})|\mathbf{x}, \mathbf{z}^t, \mathbf{y}^{t-1}, \mathbf{d}^t]$ is identified for any generic values $(\mathbf{d}, \mathbf{x}, \mathbf{z}, \mathbf{y}^{t-1})$, and consider the identification of

$$\begin{aligned} & E[Y_T(\mathbf{d})|\mathbf{x}, \mathbf{z}, \mathbf{y}^{t-2}, \mathbf{d}^{t-1}] = \Pr[D_t = d_t | \mathbf{x}, \mathbf{z}, \mathbf{y}^{t-2}, \mathbf{d}^{t-1}] E[Y_T(\mathbf{d})|\mathbf{x}, \mathbf{z}, \mathbf{y}^{t-2}, \mathbf{d}^{t-1}, d_t] \\ & \quad + \Pr[D_t = d'_t | \mathbf{x}, \mathbf{z}, \mathbf{y}^{t-2}, \mathbf{d}^{t-1}] E[Y_T(\mathbf{d})|\mathbf{x}, \mathbf{z}, \mathbf{y}^{t-2}, \mathbf{d}^{t-1}, d'_t]. \end{aligned} \tag{C.3}$$

The first main term $E[Y_T(\mathbf{d})|\mathbf{x}, \mathbf{z}, \mathbf{y}^{t-2}, \mathbf{d}^{t-1}, d_t]$ in (C.3) is identified, by integrating over $y_{t-1} \in \{0, 1\}$ the quantity $E[Y_T(\mathbf{d})|\mathbf{x}, \mathbf{z}, \mathbf{y}^{t-1}, \mathbf{d}^t]$, which is assumed to be identified in the previous iteration. The remaining unknown term in (C.3) satisfies

$$\begin{aligned} & E[Y_T(\mathbf{d})|\mathbf{x}, \mathbf{z}, \mathbf{y}^{t-2}, \mathbf{d}^{t-1}, d_t'] \\ &= \Pr[Y_{t-1} = 1|\mathbf{x}, \mathbf{z}, \mathbf{y}^{t-2}, \mathbf{d}^{t-1}, d_t']E[Y_T(\mathbf{d})|\mathbf{x}, \mathbf{z}, (\mathbf{y}^{t-2}, 1), \mathbf{d}^{t-1}, d_t'] \\ &+ \Pr[Y_{t-1} = 0|\mathbf{x}, \mathbf{z}, \mathbf{y}^{t-2}, \mathbf{d}^{t-1}, d_t']E[Y_T(\mathbf{d})|\mathbf{x}, \mathbf{z}, (\mathbf{y}^{t-2}, 0), \mathbf{d}^{t-1}, d_t']. \end{aligned}$$

By applying (C.2) to the unknown terms in this expression, we have

$$E[Y_T(\mathbf{d})|\mathbf{x}, \mathbf{z}, \tilde{\mathbf{y}}^{t-1}, \mathbf{d}^{t-1}, d_t'] = E[Y_T(d_t', \mathbf{d}_{-t})|x_t', \mathbf{x}_{-t}, \mathbf{z}, \tilde{\mathbf{y}}^{t-1}, \mathbf{d}^{t-1}, d_t'] \tag{C.4}$$

for each $\tilde{\mathbf{y}}^{t-1}$, which is identified from the previous iteration. Therefore, $E[Y_T(\mathbf{d})|\mathbf{x}, \mathbf{z}, \mathbf{y}^{t-2}, \mathbf{d}^{t-1}]$ is identified. Note that when $t = 2$, \mathbf{Y}^0 is understood to mean there is no conditioning. Lastly, when $t = 1$,

$$E[Y_T(\mathbf{d})|\mathbf{x}, \mathbf{z}] = \Pr[D_1 = d_1|\mathbf{x}, \mathbf{z}]E[Y_T(\mathbf{d})|\mathbf{x}, \mathbf{z}, d_1] + \Pr[D_1 = d_1'|\mathbf{x}, \mathbf{z}]E[Y_T(\mathbf{d})|\mathbf{x}, \mathbf{z}, d_1'].$$

Noting that $Y_0 = 0$, suppose x_1' is such that $\mu_1(0, d_1, x_1) = \mu_1(0, d_1', x_1')$ with $d_1' \neq d_1$ by applying Lemma 5.1. Then,

$$\begin{aligned} E[Y_T(\mathbf{d})|\mathbf{x}, \mathbf{z}, d_1'] &= \Pr[\mathbf{U}(\mathbf{d}) \in \mathcal{U}^T(\mathbf{d}, 1; \mathbf{x})|\mathbf{x}_{-1}, \mathbf{z}, V_1 \in \mathcal{V}^1(d_1')] \\ &= \Pr[\mathbf{U}(d_1', \mathbf{d}_{-1}) \in \mathcal{U}^T(d_1', \mathbf{d}_{-1}, 1; x_1', \mathbf{x}_{-1})|\mathbf{x}_{-1}, \mathbf{z}, V_1 \in \mathcal{V}^1(d_1')] \\ &= E[Y_T(d_1', \mathbf{d}_{-1})|x_1', \mathbf{x}_{-1}, \mathbf{z}, d_1'], \end{aligned}$$

by Assumption SX, which is identified from the previous iteration for $t = 2$. Therefore, $E[Y_T(\mathbf{d})|\mathbf{x}, \mathbf{z}]$ is identified. \square

C.3. Proof of Theorem 6.1

We analyze the identification of $E[Y_T(\mathbf{d})|\mathbf{Y}_-(\mathbf{d}) = \mathbf{y}_-, \mathbf{x}]$. Since

$$E[Y_T(\mathbf{d})|\mathbf{Y}_-(\mathbf{d}) = \mathbf{y}_-, \mathbf{x}] = \Pr[Y_T(\mathbf{d}) = 1, \mathbf{Y}_-(\mathbf{d}) = \mathbf{y}_-|\mathbf{x}]/\Pr[\mathbf{Y}_-(\mathbf{d}) = \mathbf{y}_-|\mathbf{x}],$$

we focus on the identification of $\Pr[\mathbf{Y}(\mathbf{d}) = \mathbf{y}|\mathbf{x}] = \Pr[\mathbf{Y}(\mathbf{d}) = \mathbf{y}|\mathbf{x}, \mathbf{z}]$ so that each term in the fraction is identified by appropriately marginalizing $\Pr[\mathbf{Y}(\mathbf{d}) = \mathbf{y}|\mathbf{x}]$. The proof is parallel to that of Theorem 5.1. Generalizing the \mathbf{U} -sets introduced in Section 5, define the \mathbf{U} -sets as

$$\mathcal{U}(\mathbf{d}, \mathbf{y}) \equiv \mathcal{U}(\mathbf{d}, \mathbf{y}; \mathbf{x}) \equiv \{\mathbf{U}(\mathbf{d}) : y_s = Y_s(\mathbf{d}^s, \mathbf{x}^s) \text{ for all } s \leq T\}.$$

For $2 \leq t \leq T - 1$, we have

$$\begin{aligned} & \Pr[\mathbf{Y}(\mathbf{d}) = \mathbf{y}, \mathbf{D}^t = (\mathbf{d}^{t-1}, d_t')|\mathbf{x}, \mathbf{z}] \\ &= \Pr[\mathbf{U}(\mathbf{d}) \in \mathcal{U}^T(\mathbf{d}, \mathbf{y}; \mathbf{x}), \mathbf{V}^t \in \mathcal{V}^t(\mathbf{d}^{t-1}, d_t', \mathbf{U}^{t-1}(\mathbf{d}^{t-1}))] \\ &= \Pr \left[\begin{array}{l} \mathbf{U}(d_t', \mathbf{d}_{-t}) \in \mathcal{U}^T(d_t', \mathbf{d}_{-t}, \mathbf{y}; x_t', \mathbf{x}_{-t}), \\ \mathbf{V}^t \in \mathcal{V}^t(\mathbf{d}^{t-1}, d_t', \mathbf{U}^{t-1}(\mathbf{d}^{t-1})) \end{array} \right] \\ &= \Pr[\mathbf{Y}(d_t', \mathbf{d}_{-t}) = \mathbf{y}, \mathbf{D}^t = (\mathbf{d}^{t-1}, d_t')|x_t', \mathbf{x}_{-t}, \mathbf{z}], \end{aligned} \tag{C.5}$$

where the second equality uses x_t' such that $h_t^{d_t}(x_t; \mathbf{o}^{t-1}) + h_t^{d_t'}(x_t'; \mathbf{o}^{t-1}) = 0$ and thus $Y_t(\mathbf{d}^t, \mathbf{x}^t) = Y_t(\mathbf{d}^{t-1}, d_t', \mathbf{x}^{t-1}, x_t')$ by Lemma 5.1.

Now we use (C.5) to prove the lemma via mathematical induction. First, $\Pr[\mathbf{Y}(\mathbf{d}) = \mathbf{y}, \mathbf{D}^T = \mathbf{d}^T|\mathbf{x}, \mathbf{z}] = \Pr[\mathbf{Y} = \mathbf{y}, \mathbf{D}^T = \mathbf{d}^T|\mathbf{x}, \mathbf{z}]$ is trivially identified for any generic values $(\mathbf{d}, \mathbf{x}, \mathbf{z})$. For given $2 \leq t \leq T - 1$, suppose $\Pr[\mathbf{Y}(\mathbf{d}) = \mathbf{y}, \mathbf{D}^t = \mathbf{d}^t|\mathbf{x}, \mathbf{z}]$ is identified for any generic values $(\mathbf{d}, \mathbf{x}, \mathbf{z})$, and consider identification of

$$\begin{aligned} & \Pr[\mathbf{Y}(\mathbf{d}) = \mathbf{y}, \mathbf{D}^{t-1} = \mathbf{d}^{t-1}|\mathbf{x}, \mathbf{z}] \\ &= \Pr[\mathbf{Y}(\mathbf{d}) = \mathbf{y}, \mathbf{D}^t = (\mathbf{d}^{t-1}, d_t)|\mathbf{x}, \mathbf{z}] \\ &+ \Pr[\mathbf{Y}(\mathbf{d}) = \mathbf{y}, \mathbf{D}^t = (\mathbf{d}^{t-1}, d_t')|\mathbf{x}, \mathbf{z}]. \end{aligned} \tag{C.6}$$

The first term in the expression is identified from the previous iteration. The second unknown term in (C.6) satisfies

$$\begin{aligned} & \Pr[\mathbf{Y}(\mathbf{d}) = \mathbf{y}, \mathbf{D}^t = (\mathbf{d}^{t-1}, d_t')|\mathbf{x}, \mathbf{z}] \\ &= \Pr[\mathbf{Y}(d_t', \mathbf{d}_{-t}) = \mathbf{y}, \mathbf{D}^t = (\mathbf{d}^{t-1}, d_t')|x_t', \mathbf{x}_{-t}, \mathbf{z}], \end{aligned} \tag{C.7}$$

by (C.5), which is now identified from the previous iteration. Therefore, $\Pr[\mathbf{Y}(\mathbf{d}) = \mathbf{y}, \mathbf{D}^{t-1} = \mathbf{d}^{t-1} | \mathbf{x}, \mathbf{z}]$ is identified. Lastly, when $t = 1$,

$$\begin{aligned} \Pr[\mathbf{Y}(\mathbf{d}) = \mathbf{y} | \mathbf{x}, \mathbf{z}] &= \Pr[\mathbf{Y}(\mathbf{d}) = \mathbf{y}, D_1 = d_1 | \mathbf{x}, \mathbf{z}] \\ &+ \Pr[\mathbf{Y}(\mathbf{d}) = \mathbf{y}, D_1 = d'_1 | \mathbf{x}, \mathbf{z}]. \end{aligned}$$

The first term is identified from the iteration for $t = 2$. Noting that $Y_0 = 0$, suppose x'_1 is such that $\mu_1(0, d_1, x_1) = \mu_1(0, d'_1, x'_1)$ with $d'_1 \neq d_1$ by Lemma 5.1. Then, similarly to (C.5),

$$\begin{aligned} \Pr[\mathbf{Y}(\mathbf{d}) = \mathbf{y}, D_1 = d'_1 | \mathbf{x}, \mathbf{z}] \\ = \Pr[\mathbf{Y}(d'_1, \mathbf{d}_{-1}) = \mathbf{y}, D_1 = d'_1 | x'_1, \mathbf{x}_{-1}, \mathbf{z}], \end{aligned}$$

which is also identified from the previous iteration for $t = 2$. Therefore $\Pr[\mathbf{Y}(\mathbf{d}) = \mathbf{y} | \mathbf{x}, \mathbf{z}]$ is identified. \square

References

- Abbring, J.H., Van den Berg, G.J., 2003. The nonparametric identification of treatment effects in duration models. *Econometrica* 71 (5), 1491–1517.
- Abbring, J.H., Heckman, J.J., 2007. Econometric evaluation of social programs, part iii: Distributional treatment effects, dynamic treatment effects, dynamic discrete choice, and general equilibrium policy evaluation. *Handb. Econometrics* 6, 5145–5303.
- Abraham, S., Sun, L., 2018. Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects. MIT.
- Abrevaya, J., Hausman, J.A., Khan, S., 2010. Testing for causal effects in a generalized regression model with endogenous regressors. *Econometrica* 78 (6), 2043–2061.
- Angrist, J.D., Imbens, G.W., 1995. Two-stage least squares estimation of average causal effects in models with variable treatment intensity. *J. Amer. Statist. Assoc.* 90 (430), 431–442.
- Athey, S., Imbens, G.W., 2018. Design-based Analysis in Difference-in-Differences Settings with Staggered Adoption. National Bureau of Economic Research.
- Balat, J., Han, S., 2020. Multiple treatments with strategic substitutes. *J. Econometrics* (Conditionally Accepted).
- Blevins, J.R., 2014. Nonparametric identification of dynamic decision processes with discrete and continuous choices. *Quant. Econ.* 5, 531–554.
- Buchholz, N., Shum, M., Xu, H., 2019. Semiparametric estimation of dynamic discrete choice models. *J. Econometrics* (Accepted).
- Callaway, B., Sant'Anna, P.H., 2020. Difference-in-Differences with Multiple Time Periods. Temple University and Vanderbilt University.
- Cameron, S.V., Heckman, J.J., 1998. Life cycle schooling and dynamic selection bias: Models and evidence for five cohorts of American males. *J. Political Econ.* 106 (2), 262–333.
- Chen, X., Pouzo, D., 2015. Sieve wald and qlr inferences on semi/nonparametric conditional moment models. *Econometrica* 83 (3), 1013–1079.
- Chernozhukov, V., Hansen, C., 2005. An IV model of quantile treatment effects. *Econometrica* 73 (1), 245–261.
- Conduct Problems Prevention Research Group, 1992. A developmental and clinical model for the prevention of conduct disorder: The fast track program. *Dev. Psychopathol.* 4 (4), 509–527.
- Cunha, F., Heckman, J.J., Navarro, S., 2007. The identification and economic content of ordered choice models with stochastic thresholds. *Internat. Econom. Rev.* 48 (4), 1273–1309.
- Fredriksson, P., Johansson, P., 2008. Dynamic treatment assignment: the consequences for evaluations using observational data. *J. Bus. Econom. Statist.* 26 (4), 435–445.
- Heckman, J.J., Humphries, J.E., Veramendi, G., 2016. Dynamic treatment effects. *J. Econometrics* 191 (2), 276–292.
- Heckman, J.J., Ichimura, H., Todd, P., 1998. Matching as an econometric evaluation estimator. *Rev. Econom. Stud.* 65 (2), 261–294.
- Heckman, J.J., Navarro, S., 2007. Dynamic discrete choice and dynamic treatment effects. *J. Econometrics* 136 (2), 341–396.
- Imai, K., Kim, I.S., Wang, E., 2018. Matching methods for causal inference with time-series cross-section data. Princeton University.
- Imbens, G.W., Angrist, J.D., 1994. Identification and estimation of local average treatment effects. *Econometrica* 62 (2), 467–475.
- In the Elderly Program (SHEP) Cooperative Research Group, T.S.H., 1988. Rationale and design of a randomized clinical trial on prevention of stroke in isolated systolic hypertension. *J. Clin. Epidemiol.* 41 (12), 1197–1208.
- Jun, S.J., Pinkse, J., Xu, H., Yildiz, N., 2016. Multiple discrete endogenous variables in weakly-separable triangular models. *Econometrica* 4 (1), 7.
- Lee, S., Salanié, B., 2018. Identifying effects of multivalued treatments. *Econometrica* 86 (6), 1939–1963.
- Machado, C., Shaikh, A., Vytlačil, E., 2019. Instrumental variables and the sign of the average treatment effect. *J. Econometrics* 212 (2), 522–555.
- Murphy, S.A., 2003. Optimal dynamic treatment regimes. *J. R. Stat. Soc. Ser. B Stat. Methodol.* 65 (2), 331–355.
- Murphy, S.A., van der Laan, M.J., Robins, J.M., and C. P. P. R. Group, 2001. Marginal mean models for dynamic regimes. *J. Amer. Statist. Assoc.* 96 (456), 1410–1423.
- Robins, J., 1986. A new approach to causal inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect. *Math. Modelling* 7 (9–12), 1393–1512.
- Robins, J., 1987. A graphical approach to the identification and estimation of causal parameters in mortality studies with sustained exposure periods. *J. Chronic Dis.* 40, 1395–1615.
- Robins, J.M., 1994. Correcting for non-compliance in randomized trials using structural nested mean models. *Commun. Stat. - Theory Methods* 23 (8), 2379–2412.
- Robins, J.M., 1997. Causal inference from complex longitudinal data. In: *Latent Variable Modeling and Applications To Causality*. Springer, pp. 69–117.
- Robins, J.M., 2000. Marginal structural models versus structural nested models as tools for causal inference. In: *Statistical Models in Epidemiology, the Environment, and Clinical Trials*. Springer, pp. 95–133.
- Robins, J., Rotnitzky, A., 2004. Estimation of treatment effects in randomised trials with non-compliance and a dichotomous outcome using structural mean models. *Biometrika* 91 (4), 763–783.
- Rust, J., 1987. Optimal replacement of gmc bus engines: An empirical model of harold zurcher. *Econometrica* 999–1033.
- Sasaki, Y., 2015. Heterogeneity and selection in dynamic panel data. *J. Econometrics* 188 (1), 236–249.
- Sexton, M., Hebel, J.R., 1984. A clinical trial of change in maternal smoking and its effect on birth weight. *JAMA* 251 (7), 911–915.
- Shaikh, A.M., Vytlačil, E.J., 2011. Partial identification in triangular systems of equations with binary dependent variables. *Econometrica* 79 (3), 949–955.
- Torgovitsky, A., 2019. Nonparametric inference on state dependence in unemployment. *Econometrica* 87 (5), 1475–1505.
- Vikström, J., Ridder, G., Weidner, M., 2018. Bounds on treatment effects on transitions. *J. Econometrics* 205 (2), 448–469.
- Vytlačil, E., Yildiz, N., 2007. Dummy endogenous variables in weakly separable models. *Econometrica* 75 (3), 757–779.