Identification in Nonparametric Models for Dynamic Treatment Effects*

Sukjin Han
Department of Economics
University of Texas at Austin
sukjin.han@austin.utexas.edu

First Draft: August 12, 2017
This Draft: March 13, 2018

Abstract

This paper develops a nonparametric model that represents how sequences of outcomes and treatment choices are influenced by each other in a dynamic manner. In this setting, we are interested in identifying the average outcome of 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) as well as the optimal treatment regime, namely, the regime that maximizes the (weighted) sum of the average potential outcomes, possibly less the cost of 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. We show that the parameters of interest are identified under two-way exclusion restrictions in each period, i.e., with instruments excluded from the outcome-determining process and other exogenous variables excluded from the treatment selection process. We also consider partial identification in the case where the latter variables are not available.

JEL Numbers: C14, C35, C57
Keywords: Dynamic treatment effects, optimal treatment regimes, endogenous treatments, binary dependent variables, average treatment effects, instrumental variables.

1 Introduction

This paper develops a nonparametric model that represents how sequences of outcomes and treatment choices are influenced by each other in a dynamic manner. Often times, treatments are repeatedly chosen multiple times over a horizon. Examples of sequential treatments affecting outcomes are medical interventions affecting health outcomes, educational interventions affecting academic achievements, job training programs affecting employment, or online advertisements affecting consumers' preferences and purchase decisions. The relationship of

*The author is very grateful to Dan Ackerberg, Xiaohong Chen and Ed Vytlacil for thoughtful comments and discussions.

interest is dynamic in the sense that the current outcome is determined as a function of the past outcomes as well as the current and past treatments, and the current treatment as a function of the past outcomes as well as the past treatments. Such dynamic relationships are clearly present in the examples mentioned. A static model misrepresents the nature of the problem (e.g., state dependence, learning) and fails to capture important policy questions (e.g., optimal sequences of interventions).

In this setting, we are interested in identifying the causal effect of a sequence of endogenous treatments on a sequence of outcomes, or on a terminal outcome, e.g., the survival of a patient at a given period, college attendance, employment status, or a sale of a product, in the respective examples above. Specifically, we are interested in learning about the average of an outcome in each period had a particular treatment sequence been assigned *up to that period*, which defines the potential outcome in the dynamic setting. We are also interested in the average treatment effects (ATE's) defined based on the average potential outcome. For example, one may be interested in whether the success rate of a particular outcome 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 an optimal treatment regime, namely, a sequence of treatments that maximizes the (weighted) sum of the average potential outcomes, possibly less the cost of treatments. For example, a firm may be interested in the optimal timing of advertisements that maximizes the aggregate sales probabilities over time, or a sequence of educational programs may be aimed to maximize the college attendance rate.

Dynamic treatment effects have been extensively studied in the biostatistics literature for decades (Robins (1986, 1987), Murphy et al. (2001), Murphy (2003), among others). In this literature, the crucial condition for identification of the average potential outcome is a dynamic version of a random assignment assumption, called the *sequential randomization*. It 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 suitable in experimental studies with perfect compliance of subjects, but hard to justify in studies with partial compliance or observational studies as in the examples above. Another common feature in this literature is the use of a counterfactual framework with a sequence of treatments. Under this framework, however, it is hard to disentangle the dynamic mechanism and it is not straightforward to understand the definition of potential outcomes.

The main contribution of this paper is to relax the assumption of sequential randomization widely used in the literature by introducing a flexible choice-theoretic framework for a sequence of endogenous treatments. Towards this end, we consider a simple nonparametric structural model for a dynamic endogenous selection process and dynamic outcome formation. In this model, individuals are allowed not to fully comply with each period's assignment in experimental settings (e.g., clinical trials, field experiments) or allowed to make an endogenous choice in each period in observational settings. The joint distribution of the full history of unobservable variables in the outcome and treatment equations is left unspecified, allowing for abitrary forms of treatment endogeneity and serial correlation. Unlike in the counterfactual framework, the dynamic mechanism is clearly formulated in this structural model, which

[1] In the econometrics literature, Vikström et al. (2016) consider treatment effects on a process of transition to a destination state, and carefully analyze what the sequential randomization assumption can identify under the presence of dynamic selection.

in turn facilitates our identification analysis.

We show that the average potential outcome, or equivalently the average structural function (ASF) given the structural model we introduce, is identified under two-way exclusion restrictions (Vytlacil and Yildiz (2007), Shaikh and Vytlacil (2011)), i.e., with instruments excluded from the outcome-determining process and other exogenous variables excluded from the treatment selection process. Examples of the former would be randomized treatments or a sequence of policy shocks, examples of the latter would be factors that agents cannot anticipate when making treatment decisions but that determine the outcome. The identification of each period's ASF allows us to identify the ATE's as well as the optimal treatment regime. We show that the optimal regime is a natural extension of a static object commonly sought for in the literature, namely the sign of the ATE in a static environment. Analogous to a static setting, the knowledge about the optimal treatment regime may have important policy implications. For example, a social planner can at least hope to rule out specific sequences of treatments that are on average harmful.

In this paper, we also consider the cases where the two-way exclusion restriction are violated in the sense that only a standard exclusion restriction holds or where the variation of the exogenous variables are limited. In these cases, we can calculate the bounds on the ASF and the ATE.

Despite its importance, studies on the effects of dynamic endogenous treatments are limited in the literature. To best of our knowledge, Heckman and Navarro (2007), Cunha et al. (2010) and Heckman et al. (2016) are the only existing econometric research on the topic.[2] Building on Cunha et al. (2010) and Heckman and Navarro (2007), Heckman et al. (2016) consider dynamic treatment effects with a sequence of up-or-out treatment choices and a sequence of associated outcomes, and extends the literature on the marginal treatment effects. An interesting feature of the results in the paper is that, in ordered and unordered choice models, dynamic treatment effects are decomposed into direct effects and continuation values. The present paper complements to these paper in that it considers a different form of dynamics for the treatment choices and outcomes, different identifying assumptions, and it focuses 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 seminal work by Rust (1987) and more recently by, e.g., Blevins (2014) and Buchholz et al. (2016). This literature typically considers a single rational agent's optimal decision, whereas this paper considers multiple heterogenous agents with no assumptions on agents' rationality nor strong parametric assumptions. Most importantly, this paper's focus is on the identification of the effects of treatments formed as agents' decisions. The robust approach we take in this paper is similar to Heckman and Navarro (2007) and Heckman et al. (2016). But unlike these papers, we do not necessarily invoke infinite variation of exogenous variables while remaining flexible for economic and non-economic components of the model. Lastly, Torgovitsky (2016) extends the literature on dynamic discrete choice models (with no treatment) by considering a counterfactual framework without imposing parametric assumptions. In his framework, Y_{t-1} takes the role of a treatment for Y_t and the "treatment effect" captures the state dependence. In the present paper, we consider the effects of the treatment D_t on Y_t, and introduce

[2] As related work, Angrist and Imbens (1995)'s model for multiple treatments effects in can be applied to a dynamic setting.

a selection equation for D_t as an important component of the model.

In terms of notation, for a r.v. W_t, we write $\boldsymbol{W}^t \equiv (W_1, .., W_t)$ with its realization $\boldsymbol{w}^t$. We sometimes write $\boldsymbol{W} \equiv \boldsymbol{W}^T$ for convenience. For a vector $\boldsymbol{W}$ without the t-th element, we write $\boldsymbol{W}_{-t} \equiv (W_1, ..., W_{t-1}, W_{t+1}, ..., W_T)$ with realization $\boldsymbol{w}_{-t}$. All the boldface letters represent vectors in this paper. Lastly for r.v.'s Y and W, we sometimes abbreviate $\Pr[Y = y|W = w]$ to $\Pr[Y = y|w]$ or $P[y|w]$.

2 Robins's Framework

We first introduce Robins's 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 $t = 1, ..., T$ with fixed T, Y_t is the outcome at t with realization y_t and D_t is a binary treatment at t with realization d_t. The underlying data structure is panel data (cross-sectional index i suppressed throughout, unless necessary). We may be interested in the terminal outcome Y_T; e.g., the survival of a patient at a given period, college attendance, employment status, or online product sales. We call Y_t for $t < T$ a transition outcome; e.g., intermediate health status, test scores, intermediate employment status, or surveyed preferences (or click behaviors).

Consider a treatment regime $\boldsymbol{d} \equiv (d_1, ..., d_T) \in \mathcal{D} \subseteq \{0, 1\}^T$, which is defined as a predetermined hypothetical sequence of interventions over time, i.e., a sequence of each period's assignment decisions of whether to treat or not, or whether to assign treatment A or treatment B.[3] Then a potential outcome at t can be written as $Y_{t,\boldsymbol{d}}$. It can be understood as an outcome of an individual had a particular treatment sequence been assigned. Although the genesis of $Y_{t,\boldsymbol{d}}$ 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 $Y_{t,\boldsymbol{d}}$ becomes more transparent later with a structural model introduced in this paper.

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

$$(Y_{0,\boldsymbol{d}}, ..., Y_{T,\boldsymbol{d}}) \perp D_t | \boldsymbol{Y}^{t-1} = \boldsymbol{y}^{t-1}, \boldsymbol{D}^{t-1} = \boldsymbol{d}^{t-1} \tag{2.1}$$

for $t = 1, ..., T$ and all $\boldsymbol{y}^{t-1}$. This assumption asserts that, holding the history of outcomes and treatments fixed, the current treatment is fully randomized. In the next section, we relax this assumption and specify dynamic selection equations for a sequence of treatments that are allowed to be endogenous.

3 A Structural Model and Objects of Interest

We introduce the main framework of this paper. We consider a structural model and identifying assumptions that are plausible in observational studies and are economically interpretable, relative to the counterfactual framework in the previous section. The structural

[3] This is called a static regime in the biostatistics literature. A dynamic regime is a sequence of treatment assignments, each of which is a predetermined function of past outcomes; see e.g., Murphy et al. (2001). A static regime can be seen as being its special case where this function is constant.

model includes a sequence of selection processes for endogenous treatment choices.

Consider a *dynamic structural function* for the outcome Y_t: For $t = 1, ..., T$,

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

where $\mu_t(\cdot)$ is an unknown scalar-valued function, X_t is a set of exogenous variables, which we discuss the details later, U_t is the unobservable variable that may contain permanent and transitory components, i.e., $U_{it} = h_t(\alpha_i, \epsilon_{it})$, and $Y_0 = 0$ for convenience. Given this structural equation, we can express the potential outcome $Y_{t,\boldsymbol{d}}$ using a recursive structure:

$$\begin{aligned} Y_{t,\boldsymbol{d}} &= \mu_t(Y_{t-1,\boldsymbol{d}}, d_t, X_t, U_t), \\ &\vdots \\ Y_{2,\boldsymbol{d}} &= \mu_2(Y_{1,\boldsymbol{d}}, d_2, X_2, U_2), \\ Y_{1,\boldsymbol{d}} &= \mu_1(Y_{0,\boldsymbol{d}}, d_1, X_1, U_1), \end{aligned}$$

where $Y_{0,\boldsymbol{d}} = 0$ for convenience.[4] This recursive structure provides a useful interpretation of the potential outcome $Y_{t,\boldsymbol{d}}$ in a dynamic setting, and thus facilitates our identification analysis. Note that, conditional on $\boldsymbol{X}^t \equiv (X_1, ..., X_t)$, the remaining randomness in $Y_{t,\boldsymbol{d}}$ comes from $\boldsymbol{U}^t \equiv (U_1, ..., U_t)$. By an iterative argument, one can show that the potential outcome equals the observed outcome when the observed treatments are consistent with the regime: $Y_{t,\boldsymbol{d}} = Y_t$. Or equivalently,

$$Y_t = \sum_{\boldsymbol{d} \in \mathcal{D}} 1\{\boldsymbol{D} = \boldsymbol{d}\} Y_{t,\boldsymbol{d}}.$$

Given the structural model for Y_t, the following implies the sequential randomization assumption (2.1):

$$\boldsymbol{U}^T \perp D_t | \boldsymbol{Y}^{t-1} = \boldsymbol{y}^{t-1}, \boldsymbol{D}^{t-1} = \boldsymbol{d}^{t-1}$$

for $t = 1, ..., T$ and all $\boldsymbol{y}^{t-1}$ and $\boldsymbol{d}^{t-1}$, which help understand the assumption.

In this paper, we consider the *average potential terminal outcome* (equivalently, the ASF in the terminal period) conditional on $\boldsymbol{X} \equiv (X_1, ..., X_T)$ as a fundamental parameter of interest:

$$E[Y_{T,\boldsymbol{d}} | \boldsymbol{X} = \boldsymbol{x}]. \tag{3.1}$$

In general, we can consider the average potential outcome of any time period, i.e., $E[Y_{t,\boldsymbol{d}} | \boldsymbol{X}^t = \boldsymbol{x}^t]$ for any t, but we focus on $E[Y_{T,\boldsymbol{d}} | \boldsymbol{X} = \boldsymbol{x}]$ just for concreteness. The knowledge on the ASF is useful to recover other related parameters. First, we are interested in the conditional ATE:

$$E[Y_{T,\boldsymbol{d}} - Y_{T,\tilde{\boldsymbol{d}}} | \boldsymbol{X} = \boldsymbol{x}] \tag{3.2}$$

[4] To be more precise in terms of notation, we should write $Y_{t,\boldsymbol{d}} = Y_{t,\boldsymbol{d}^t}$, but we maintain the notation $Y_{t,\boldsymbol{d}}$ for simplicity.

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

As an additional parameter of interest, we consider an *optimal treatment regime*:

$$\boldsymbol{d}^*(\boldsymbol{x}) = \arg\max_{\boldsymbol{d}\in\mathcal{D}} E[Y_{T,\boldsymbol{d}}|\boldsymbol{X}=\boldsymbol{x}] \tag{3.3}$$

with $|\mathcal{D}| = 2^T$. That is, we are interested in a treatment regime that delivers the maximum expected outcome conditional on $\boldsymbol{X} = \boldsymbol{x}$.[5] Note that in a static model, the identification of $\boldsymbol{d}^*$ is equivalent to the identification of the sign of the static ATE, which is information typically sought for from a policy point of view. One can view $\boldsymbol{d}^*$ as a natural extension of this information to a dynamic setting, which is identified by establishing the signs of *all* the possible ATE's defined as (3.2). The optimal regime may serve as a guideline in developing future policies. Moreover, it may be too costly to find a customized treatment scheme for each individual and it may be a realistic goal from a social planner's point of view to find a scheme that maximizes the average benefit. Yet, the optimal regime is customized up to observed characteristics, as it is a function of covariates values $\boldsymbol{x}$. More ambitious than the identification of $\boldsymbol{d}^*(\boldsymbol{x})$ may be an optimal regime based on an objective function that delivers a cost-benefit analysis, granting than each d_t can be costly:

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

where

$$\Pi(\boldsymbol{d};\boldsymbol{x}) \equiv w_1 E[Y_{T,\boldsymbol{d}}|\boldsymbol{X}=\boldsymbol{x}] - w_0 \sum_{t=1}^{T} d_t$$

or

$$\Pi(\boldsymbol{d};\boldsymbol{x}) \equiv \sum_{t=1}^{T} w_t E[Y_{t,\boldsymbol{d}}|\boldsymbol{X}=\boldsymbol{x}] - w_0 \sum_{t=1}^{T} d_t$$

with (w_0, w_1) and $(w_0, \boldsymbol{w})$ being predetermined weights. The latter objective function concerns the weighted sum of the average potential outcomes throughout the entire period less the cost for treatments. Note that establishing the signs of ATE's will not identify $\boldsymbol{d}^\dagger$, and a stronger identification result, i.e., the point identification of $E[Y_{T,\boldsymbol{d}}|\boldsymbol{X}=\boldsymbol{x}]$'s for all $\boldsymbol{d}$, is required.

In order to identify the parameters of interest without assuming sequential randomization, we introduce a sequence of selection equations for binary endogenous treatments D_t's: For $t = 1, ..., T$

$$D_t = 1\{\pi_t(Y_{t-1}, D_{t-1}, Z_t) \geq V_t\},$$

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 $Y_0 = 0$ and $D_0 = 0$ for convenience. This dynamic selection process represents the agent's

[5] A dynamic version of an optimal treatment regime is considered in, e.g., Murphy (2003).

endogenous choices over time, e.g., due to learning or other optimal behaviors. The nonparametric threshold-crossing structure, however, posits a minimal notion of optimality for the agent. We take an agnostic and robust approach by avoiding strong assumptions of the standard dynamic economic models, e.g., forward looking behaviors and being able to compute a present value discounted flow of utilities, conditional independence assumptions which rule out persistence, and other parametric functional forms. When we are to maintain these assumptions, the selection model can be seen as a reduced-form approximation of a solution to a dynamic programming problem.

To summarize, the full model we consider in this paper is

$$Y_t = \mu_t(Y_{t-1}, D_t, X_t, U_t), \tag{3.4}$$
$$D_t = 1\{\pi_t(Y_{t-1}, D_{t-1}, Z_t) \geq V_t\}. \tag{3.5}$$

All other covariates common to (3.4) and (3.5) are suppressed for simplicity of exposition. Importantly in this model, the joint distribution of $\boldsymbol{U} \equiv (U_1, ..., U_T)$ and $\boldsymbol{V} \equiv (V_1, ..., V_T)$ is not specified, that is, U_t and $V_{t'}$ for any t, t' are allowed to be arbitrarily correlated to each other (allowing endogeneity) as well as within themselves across time (allowing serial correlation, e.g., via individual effects). Note that becasue of this, Y_t and D_t are not Markov processes unless conditional on both the observables and unobservables.

4 Identification Analysis

We first identify the ASF's, i.e., $E[Y_{t,\boldsymbol{d}}|\boldsymbol{X}^t]$ for each $\boldsymbol{d}$ and t, which will then identify the ATE's and the optimal regimes $\boldsymbol{d}^*$ and $\boldsymbol{d}^\dagger$. We impose assumptions on $(\boldsymbol{U}, \boldsymbol{V})$ as well as on $\boldsymbol{Z} \equiv (Z_1, ..., Z_T)$ and $\boldsymbol{X} \equiv (X_1, ..., X_T)$.

Assumption C. *The distribution of* $(\boldsymbol{U}, \boldsymbol{V})$ *has strictly positive density with respect to Lebesgue measure on* $\mathbb{R}^{2T}$.

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

Assumption C is a regularity condition to ensure smoothness of relevant conditional probabilities. Assumption SX imposes strict exogeneity. The variable Z_t is standard excluded instruments. Examples are sequential randomized treatments or a sequence of policy shocks. In addition to Z_t, we introduce exogenous variables X_t in the outcome equation (3.4), that is excluded from the selection equation (3.5). We assume that, at a given period t, there are outcome-determining factors that the agent cannot anticipate when making a treatment decision.

Assuming binary Y_t, consider the following model with weak separability imposed in the outcome equation as in the treatment equation:

$$Y_t = 1\{\mu_t(Y_{t-1}, D_t, X_t) \geq U_t\}, \tag{4.1}$$
$$D_t = 1\{\pi_t(Y_{t-1}, D_{t-1}, Z_t) \geq V_t\}. \tag{4.2}$$

Binary Y_t and weak separability may not be necessary (Vytlacil and Yildiz (2007), Han (2018)) for the results of this paper, but they simplify the exposition. Define the following

period-specific quantity readily identified from the data:

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

Lemma 4.1. *Suppose Assumptions C and SX hold. For each t and $(\boldsymbol{x}^{t-1}, \boldsymbol{z}^{t-1}, \boldsymbol{d}^{t-1}, y_{t-1})$, suppose z_t and $\tilde{z}_t$ satisfy*

$$\Pr[D_t = 1 | \boldsymbol{x}^{t-1}, \boldsymbol{z}^t, \boldsymbol{d}^{t-1}, y_{t-1}] - \Pr[D_t = 1 | \boldsymbol{x}^{t-1}, \tilde{z}_t, \boldsymbol{z}^{t-1}, \boldsymbol{d}^{t-1}, y_{t-1}] > 0. \tag{4.3}$$

Then for given $(x_t, \tilde{x}_t)$, the sign of $h_t(z_t, \tilde{z}_t, x_t, \tilde{x}_t; \boldsymbol{z}^{t-1}, \boldsymbol{x}^{t-1}, \boldsymbol{d}^{t-1}, y_{t-1})$ equals the sign of

$$\mu_t(y_{t-1}, 1, x_t) - \mu_t(y_{t-1}, 0, \tilde{x}_t).$$

The sign of $\mu_t(y_{t-1}, 1, x_t) - \mu_t(y_{t-1}, 0, \tilde{x}_t)$ with $x_t = \tilde{x}_t$ immediately identifies the sign of the *period-specific ATE* $E[Y_{t,1} - Y_{t,0} | Y_{t-1} = y_{t-1}, X_t = x_t]$ where $Y_{t,d_t} = \mu_t(Y_{t-1}, d_t, X_t, U_t)$ is the *period-specific potential outcome* at time t. The sign of $\mu_t(y_{t-1}, 1, x_t) - \mu_t(y_{t-1}, 0, \tilde{x}_t)$ itself is already useful for calculating bounds on the ASF's and thus on the ATE's defined in (3.2); we discuss the partial identification in Section 5.

Proof of Lemma 4.1: For the analysis of this paper which deals with a dynamic model, it is convenient to define the $\boldsymbol{U}$-set and $\boldsymbol{V}$-set, namely the sets of the histories of the unobservable variables that determine the current outcome and current treatment given the dynamic nature. To focus our attention on this dependence of the potential outcomes on the unobservables, let $Y_{t,\boldsymbol{d},\boldsymbol{x}} \equiv Y_{t,\boldsymbol{d}^t,\boldsymbol{x}^t} \equiv \mu_t(Y_{t-1,\boldsymbol{d},\boldsymbol{x}}, d_t, x_t, U_t)$ be the potential outcome given $(\boldsymbol{d}, \boldsymbol{x})$ and let $y_{t,\boldsymbol{d},\boldsymbol{x}} = \mu_t(y_{t-1,\boldsymbol{d},\boldsymbol{x}}, d_t, x_t, u_t)$ be its realization. For $t = 1, ..., T$, define a set of $\boldsymbol{U}^t$ as

$$\mathcal{U}_{\boldsymbol{d}^t}(\boldsymbol{x}^t; y_t) \equiv \{\boldsymbol{u}^t : y_t = 1\{\mu_t(y_{t-1,\boldsymbol{d}^{t-1},\boldsymbol{x}^{t-1}}, d_t, x_t) \geq u_t\}\}$$

with $y_0 = 0$. Also, for $t = 2, ..., T$, define a set of $\boldsymbol{V}^t$ as

$$\mathcal{V}^*_{\boldsymbol{d}^t}(\boldsymbol{u}^{t-1}, \boldsymbol{x}^{t-1}, \boldsymbol{z}^t) \equiv \{\boldsymbol{v}^t : d_s = 1\{v_s \leq \pi^*_s(\boldsymbol{u}^{s-1}, \boldsymbol{x}^{s-1}, \boldsymbol{d}^{s-1}, z_s)\} \text{ for } s = 2, ..., t\},$$

where $\pi^*_s(\boldsymbol{u}^{s-1}, \boldsymbol{x}^{s-1}, \boldsymbol{d}^{s-1}, z_s) \equiv \pi_s(y_{s-1,\boldsymbol{d}^{s-1},\boldsymbol{x}^{s-1}}, d_{s-1}, z_s)$, and define

$$\mathcal{V}^*_{d_1}(z_1) \equiv \{v_1 : d_1 = 1\{v_1 \leq \pi_1(0, 0, z_1)\}\}.$$

Consider

$$\begin{aligned}
&\Pr[D_t = 1 | \boldsymbol{x}^{t-1}, \boldsymbol{z}^t, \boldsymbol{d}^{t-1}, y_{t-1}] \\
= &\Pr[V_t \leq \pi_t(y_{t-1}, d_{t-1}, z_t) | \boldsymbol{x}^{t-1}, \boldsymbol{z}^t, \boldsymbol{V}^{t-1} \in \mathcal{V}^*_{\boldsymbol{d}^{t-1}}(\boldsymbol{U}^{t-2}), \boldsymbol{U}^{t-1} \in \mathcal{U}_{\boldsymbol{d}^{t-1}}] \\
= &\Pr[V_t \leq \pi_t(y_{t-1}, d_{t-1}, z_t) | \boldsymbol{V}^{t-1} \in \mathcal{V}^*_{\boldsymbol{d}^{t-1}}(\boldsymbol{U}^{t-2}), \boldsymbol{U}^{t-1} \in \mathcal{U}_{\boldsymbol{d}^{t-1}}],
\end{aligned}$$

where the last equality is by Assumption SX, and $\mathcal{V}^*_{\boldsymbol{d}^{t-1}}(\boldsymbol{U}^{t-2}) \equiv \mathcal{V}^*_{\boldsymbol{d}^{t-1}}(\boldsymbol{U}^{t-2}, \boldsymbol{x}^{t-2}, \boldsymbol{z}^{t-1})$ and $\mathcal{U}_{\boldsymbol{d}^{t-1}} \equiv \mathcal{U}_{\boldsymbol{d}^{t-1}}(\boldsymbol{x}^{t-1}; y_{t-1})$ as abbreviation. Note that the sets $\mathcal{V}^*_{\boldsymbol{d}^{t-1}}(\boldsymbol{U}^{t-2})$ and $\mathcal{U}_{\boldsymbol{d}^{t-1}}$ do not

change at the change in z_t. Then, under Assumption C,

$$\begin{aligned}0 &< \Pr[D_t = 1|\boldsymbol{x}^{t-1}, \boldsymbol{z}^t, \boldsymbol{d}^{t-1}, y_{t-1}] - \Pr[D_t = 1|\boldsymbol{x}^{t-1}, \tilde{z}_t, \boldsymbol{z}^{t-1}, \boldsymbol{d}^{t-1}, y_{t-1}] \\ &= \Pr[V_t \leq \pi_t(y_{t-1}, d_{t-1}, z_t)|\boldsymbol{V}^{t-1} \in \mathcal{V}^*_{\boldsymbol{d}^{t-1}}(\boldsymbol{U}^{t-2}), \boldsymbol{U}^{t-1} \in \mathcal{U}_{\boldsymbol{d}^{t-1}}] \\ &\quad - \Pr[V_t \leq \pi_t(y_{t-1}, d_{t-1}, \tilde{z}_t)|\boldsymbol{V}^{t-1} \in \mathcal{V}^*_{\boldsymbol{d}^{t-1}}(\boldsymbol{U}^{t-2}), \boldsymbol{U}^{t-1} \in \mathcal{U}_{\boldsymbol{d}^{t-1}}]\end{aligned}$$

implies $\pi_t(y_{t-1}, d_{t-1}, z_t) > \pi_t(y_{t-1}, d_{t-1}, z_t)$. Next, consider

$$\begin{aligned}&\Pr[Y_t = 1, D_t = 1|\boldsymbol{z}^t, \boldsymbol{x}^t, \boldsymbol{d}^{t-1}, y_{t-1}] \\ &= \Pr[U_t \leq \mu_t(y_{t-1}, 1, x_t), V_t \leq \pi_t(y_{t-1}, d_{t-1}, z_t)|\boldsymbol{V}^{t-1} \in \mathcal{V}^*_{\boldsymbol{d}^{t-1}}(\boldsymbol{U}^{t-2}), \boldsymbol{U}^{t-1} \in \mathcal{U}_{\boldsymbol{d}^{t-1}}]\end{aligned}$$

by Assumption SX. Again, note that $\mathcal{V}^*_{\boldsymbol{d}^{t-1}}(\boldsymbol{U}^{t-2})$ and $\mathcal{U}_{\boldsymbol{d}^{t-1}}$ do not change at the change in (z_t, x_t), which is key. Then, we have

$$\begin{aligned}&h_t(z_t, \tilde{z}_t, x_t, \tilde{x}_t; \boldsymbol{z}^{t-1}, \boldsymbol{x}^{t-1}, \boldsymbol{d}^{t-1}, y_{t-1}) \\ &= \Pr[U_t \leq \mu_t(y_{t-1}, 1, x_t), \pi_t(y_{t-1}, d_{t-1}, \tilde{z}_t) \leq V_t \leq \pi_t(y_{t-1}, d_{t-1}, z_t)|\mathcal{V}^*_{\boldsymbol{d}^{t-1}}, \mathcal{U}_{\boldsymbol{d}^{t-1}}] \\ &\quad - \Pr[U_t \leq \mu_t(y_{t-1}, 0, \tilde{x}_t), \pi_t(y_{t-1}, d_{t-1}, \tilde{z}_t) \leq V_t \leq \pi_t(y_{t-1}, d_{t-1}, z_t)|\mathcal{V}^*_{\boldsymbol{d}^{t-1}}, \mathcal{U}_{\boldsymbol{d}^{t-1}}],\end{aligned}$$

of which sign identifies the sign of $\mu_t(y_{t-1}, 1, x_t) - \mu_t(y_{t-1}, 0, \tilde{x}_t)$. For example, when this quantity is zero then $\mu_t(y_{t-1}, 1, x_t) - \mu_t(y_{t-1}, 0, \tilde{x}_t) = 0$. □

We make a further assumption on the variation of the exogenous variables $(\boldsymbol{Z}, \boldsymbol{X})$ for point identification of the ASF's. Define the following sets:

$$\begin{aligned}\mathcal{S}_t(y_{t-1}, d_t) &\equiv \left\{(x_t, \tilde{x}_t) : \mu_t(y_{t-1}, d_t, x_t) = \mu_t(y_{t-1}, \tilde{d}_t, \tilde{x}_t) \text{ for } \tilde{d}_t \neq d_t\right\}, \\ \mathcal{T}_t &\equiv \{(x_t, \tilde{x}_t) : \exists(z_t, \tilde{z}_t) \text{ with } (x_t, z_t), (\tilde{x}_t, z_t), (x_t, \tilde{z}_t), (\tilde{x}_t, \tilde{z}_t) \in \text{Supp}(X_t, Z_t)\}, \\ \mathcal{X}_t(y_{t-1}, d_t) &\equiv \{x_t : \exists\tilde{x}_t \text{ with } (x_t, \tilde{x}_t) \in \mathcal{S}_t(y_{t-1}, d_t) \cap \mathcal{T}_t\}, \\ \mathcal{X}_t(d_t) &\equiv \mathcal{X}_t(0, d_t) \cap \mathcal{X}_t(1, d_t).\end{aligned}$$

Assumption SP. *For each t and d_t,* $\Pr[X_t \in \mathcal{X}_t(d_t)] > 0$.

This assumption requires that X_t varies enough to achieve $\mu_t(y_{t-1}, d_t, x_t) = \mu_t(y_{t-1}, \tilde{d}_t, \tilde{x}_t)$ while holding Z_t to be z_t and $\tilde{z}_t$, respectively.[6] It is a dynamic version of the support assumption found in Vytlacil and Yildiz (2007). Note that even though this assumption seems to be written in terms of the unknown object $\mu_t(\cdot)$, it is testable because the sets above have an empirical analog as shown in Lemma 4.1. Now we are ready to state the main identification result of this paper.

Theorem 4.1. *Under Assumptions C, SX and SP, $E[Y_{T,\boldsymbol{d}}|\boldsymbol{x}]$ is identified for $\boldsymbol{d} \in \mathcal{D}$ and $x_t \in \mathcal{X}_t(d_t)$ for all t.*

Based on Theorem 4.1, we can identify the ATE's and the optimal treatment regimes $\boldsymbol{d}^*(\boldsymbol{x})$ and $\boldsymbol{d}^\dagger(\boldsymbol{x})$:

[6] Although Assumption SP requires sufficient rectangular variation in (X_t, Z_t), it clearly differs from the large variation assumptions in Heckman and Navarro (2007) and Heckman et al. (2016), which are employed for identification at infinity arguments to identify different objects of interest than ours.

Corollary 4.1. *Under Assumptions C, SX and SP, $E[Y_{T,\boldsymbol{d}} - Y_{T,\tilde{\boldsymbol{d}}}|\boldsymbol{x}]$ for $\boldsymbol{d}, \tilde{\boldsymbol{d}} \in \mathcal{D}$ and $x_t \in \mathcal{X}_t(d_t) \cap \mathcal{X}_t(\tilde{d}_t)$ for all t is identified, and $\boldsymbol{d}^*(\boldsymbol{x})$ and $\boldsymbol{d}^\dagger(\boldsymbol{x})$ are identified for $x_t \in \mathcal{X}_t(0) \cap \mathcal{X}_t(1)$ for all t.*

Proof of Theorem 4.1: To begin with, note that $E[Y_{T,\boldsymbol{d}}|\boldsymbol{x}] = E[Y_{T,\boldsymbol{d}}|\boldsymbol{x}, \boldsymbol{z}]$ by Assumption SX. As the first step of identifying $E[Y_{T,\boldsymbol{d}}|\boldsymbol{x}, \boldsymbol{z}]$ for given $\boldsymbol{d} = (d_1, ..., d_T)$, $\boldsymbol{x} = (x_1, ..., x_T)$ and $\boldsymbol{z} = (z_1, ..., z_T)$, we apply the result of Lemma 4.1. Fix $t = 2, ..., T$ and $y_{t-1} \in \{0, 1\}$. Suppose x_t' is such that $\mu_t(y_{t-1}, d_t, x_t) = \mu_t(y_{t-1}, d_t', x_t')$ with $d_t' \neq d_t$ by applying Lemma 4.1. The existence of x_t' is guaranteed by Assumption SP. Then, by the definition of the $\boldsymbol{U}$-set, $\boldsymbol{U} \in \mathcal{U}_{\boldsymbol{d}}(\boldsymbol{x}; y_T)$ is equivalent of $\boldsymbol{U} \in \mathcal{U}_{d_t', \boldsymbol{d}_{-t}}(x_t', \boldsymbol{x}_{-t}; y_T)$ *conditional on* $Y_{t-1,\boldsymbol{d}^{t-1},\boldsymbol{x}^{t-1}} = y_{t-1}$ for all $\boldsymbol{x}_{-t}$ and $\boldsymbol{d}_{-t}$.[7] Based on this result, we can equate the unobserved quantity $E[Y_{T,\boldsymbol{d}}|\boldsymbol{x}, \boldsymbol{z}, \boldsymbol{d}^{t-1}, d_t', y_{t-1}]$ with an observed quantity by matching the element d_t of the assigned treatments to be consistent with the element d_t' of the observed treatments:

$$
\begin{aligned}
&E[Y_{T,\boldsymbol{d}}|\boldsymbol{x}, \boldsymbol{z}, \boldsymbol{d}^{t-1}, d_t', y_{t-1}] \\
=&\Pr[\boldsymbol{U} \in \mathcal{U}_{\boldsymbol{d}}(\boldsymbol{x}; 1)|\boldsymbol{x}, \boldsymbol{z}, \boldsymbol{V}^t \in \mathcal{V}^*_{\boldsymbol{d}^{t-1}, d_t'}(\boldsymbol{U}^{t-1}), \boldsymbol{U}^{t-1} \in \mathcal{U}_{\boldsymbol{d}^{t-1}}] \\
=&\Pr[\boldsymbol{U} \in \mathcal{U}_{\boldsymbol{d}}(\boldsymbol{x}; 1)|\boldsymbol{V}^t \in \mathcal{V}^*_{\boldsymbol{d}^{t-1}, d_t'}(\boldsymbol{U}^{t-1}), \boldsymbol{U}^{t-1} \in \mathcal{U}_{\boldsymbol{d}^{t-1}}] \\
=&\Pr[\boldsymbol{U} \in \mathcal{U}_{d_t', \boldsymbol{d}_{-t}}(x_t', \boldsymbol{x}_{-t}; 1)|\boldsymbol{V}^t \in \mathcal{V}^*_{\boldsymbol{d}^{t-1}, d_t'}(\boldsymbol{U}^{t-1}), \boldsymbol{U}^{t-1} \in \mathcal{U}_{\boldsymbol{d}^{t-1}}] \\
=&\Pr[\boldsymbol{U} \in \mathcal{U}_{d_t', \boldsymbol{d}_{-t}}(x_t', \boldsymbol{x}_{-t}; 1)|x_t', \boldsymbol{x}_{-t}, \boldsymbol{z}, \boldsymbol{V}^t \in \mathcal{V}^*_{\boldsymbol{d}^{t-1}, d_t'}(\boldsymbol{U}^{t-1}), \boldsymbol{U}^{t-1} \in \mathcal{U}_{\boldsymbol{d}^{t-1}}] \\
=&E[Y_{T,(d_t', \boldsymbol{d}_{-t})}|x_t', \boldsymbol{x}_{-t}, \boldsymbol{z}, \boldsymbol{d}^{t-1}, d_t', y_{t-1}], \qquad (4.4)
\end{aligned}
$$

where the second and fourth equalities are by Assumption SX, and $\mathcal{V}^*_{\boldsymbol{d}^{t-1}, d_t'}(\boldsymbol{U}^{t-1}) \equiv \mathcal{V}^*_{\boldsymbol{d}^{t-1}, d_t'}(\boldsymbol{U}^{t-1}, \boldsymbol{x}^{t-1}, \boldsymbol{z}^t)$ and $\mathcal{U}_{\boldsymbol{d}^{t-1}} \equiv \mathcal{U}_{\boldsymbol{d}^{t-1}}(\boldsymbol{x}^{t-1}; y_{t-1})$ as abbreviation. We use this result in the next step.

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

$$
\begin{aligned}
E[Y_{T,\boldsymbol{d}}|\boldsymbol{x}, \boldsymbol{z}, \boldsymbol{d}^{t-1}] =& \Pr[D_t = d_t|\boldsymbol{x}, \boldsymbol{z}, \boldsymbol{d}^{t-1}] E[Y_{T,\boldsymbol{d}}|\boldsymbol{x}, \boldsymbol{z}, \boldsymbol{d}^{t-1}, d_t] \\
&+ \Pr[D_t = d_t'|\boldsymbol{x}, \boldsymbol{z}, \boldsymbol{d}^{t-1}] E[Y_{T,\boldsymbol{d}}|\boldsymbol{x}, \boldsymbol{z}, \boldsymbol{d}^{t-1}, d_t'].
\end{aligned}
$$

The only unobserved term on the r.h.s. can be shown to satisfy

$$
\begin{aligned}
E[Y_{T,\boldsymbol{d}}|\boldsymbol{x}, \boldsymbol{z}, \boldsymbol{d}^{t-1}, d_t'] =& \Pr[Y_{t-1} = 1|\boldsymbol{x}, \boldsymbol{z}, \boldsymbol{d}^{t-1}, d_t'] E[Y_{T,\boldsymbol{d}}|\boldsymbol{x}, \boldsymbol{z}, \boldsymbol{d}^{t-1}, d_t', Y_{t-1} = 1] \\
&+ \Pr[Y_{t-1} = 0|\boldsymbol{x}, \boldsymbol{z}, \boldsymbol{d}^{t-1}, d_t'] E[Y_{T,\boldsymbol{d}}|\boldsymbol{x}, \boldsymbol{z}, \boldsymbol{d}^{t-1}, d_t', Y_{t-1} = 0]. \qquad (4.5)
\end{aligned}
$$

But note that

$$
E[Y_{T,\boldsymbol{d}}|\boldsymbol{x}, \boldsymbol{z}, \boldsymbol{d}^{t-1}, d_t', y_{t-1}] = E[Y_{T,(d_t', \boldsymbol{d}_{-t})}|x_t', \boldsymbol{x}_{-t}, \boldsymbol{z}, \boldsymbol{d}^{t-1}, d_t', y_{t-1}] \qquad (4.6)
$$

by (4.4), which is assumed to be identified from the previous step. Therefore $E[Y_{T,\boldsymbol{d}}|\boldsymbol{x}, \boldsymbol{z}, \boldsymbol{d}^{t-1}]$

[7] The following analysis is significantly simplified when $\mu_t(y_{t-1}, d_t, x_t) = \mu_t(y_{t-1}, d_t', x_t')$ satisfies *for all* y_{t-1}. This situation, however, is hard to occur.

is identified. Lastly when $t = 1$,

$$E[Y_{T,\boldsymbol{d}}|\boldsymbol{x},\boldsymbol{z}] = \Pr[D_1 = d_1|\boldsymbol{x},\boldsymbol{z}]E[Y_{T,\boldsymbol{d}}|\boldsymbol{x},\boldsymbol{z},d_1] + \Pr[D_1 = d_1'|\boldsymbol{x},\boldsymbol{z}]E[Y_{T,\boldsymbol{d}}|\boldsymbol{x},\boldsymbol{z},d_1'].$$

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 4.1. Then

$$\begin{aligned}
E[Y_{T,\boldsymbol{d}}|\boldsymbol{x},\boldsymbol{z},d_1'] &= \Pr[\boldsymbol{U} \in \mathcal{U}_{\boldsymbol{d}}(\boldsymbol{x};1)|\boldsymbol{x},\boldsymbol{z},V_1 \in \mathcal{V}_{d_1'}^*(z_1)] \\
&= \Pr[\boldsymbol{U} \in \mathcal{U}_{\boldsymbol{d}}(\boldsymbol{x};1)|V_1 \in \mathcal{V}_{d_1'}^*(z_1)] \\
&= \Pr[\boldsymbol{U} \in \mathcal{U}_{d_1',\boldsymbol{d}_{-1}}(x_1',\boldsymbol{x}_{-1};1)|V_1 \in \mathcal{V}_{d_1'}^*(z_1)] \\
&= \Pr[\boldsymbol{U} \in \mathcal{U}_{d_1',\boldsymbol{d}_{-1}}(x_1',\boldsymbol{x}_{-1};1)|x_1',\boldsymbol{x}_{-1},\boldsymbol{z},V_1 \in \mathcal{V}_{d_1'}^*(z_1)] \\
&= E[Y_{T,(d_1',\boldsymbol{d}_{-1})}|x_1',\boldsymbol{x}_{-1},\boldsymbol{z},d_1'],
\end{aligned}$$

which is identified from the previous step for $t = 2$. Therefore $E[Y_{T,\boldsymbol{d}}|\boldsymbol{x},\boldsymbol{z}]$ is identified, which completes the proof of Theorem 4.1. □

Note that this proof provides a closed form expression for $E[Y_{T,\boldsymbol{d}}|\boldsymbol{x}]$ in an iterative manner, which can be immediately used for estimation. For concreteness, we provide an expression for $E[Y_{T,\boldsymbol{d}}|\boldsymbol{x}]$ when $T = 2$:

$$\begin{aligned}
E[Y_{2,\boldsymbol{d}}|\boldsymbol{x}] =& P[\boldsymbol{d}|\boldsymbol{x},\boldsymbol{z},d_1]E[Y_T|x,z,d] + P[d_1,d_2'|x,z]\mu_{Y_T,d_1,d_2'} \\
&+ P[d_1',d_2|x,z]E[Y_T|x_1',x_2,z,d_1',d_2] + P[d_1',d_2'|x,z]\mu_{Y_T,d_1',d_2'},
\end{aligned} \tag{4.7}$$

where

$$\begin{aligned}
\mu_{Y_T,d_1,d_2'} \equiv& P[y_1|x,z,d_1,d_2']E[Y_T|x_1,x_2',z,d_1,d_2',y_1] \\
&+ P[y_1'|x,z,d_1,d_2']E[Y_T|x_1,x_2'',z,d_1,d_2',y_1'], \\
\mu_{Y_T,d_1',d_2'} \equiv& P[y_1|x_1',x_2,z,d_1',d_2']E[Y_T|x_1',x_2',z,d_1',d_2',y_1] \\
&+ P[y_1'|x_1',x_2,z,d_1',d_2']E[Y_T|x_1',x_2'',z,d_1',d_2',y_1']
\end{aligned}$$

for (x_1', x_2', x_2'') such that $\mu_1(0,d_1,x_1) = \mu_1(0,d_1',x_1')$, $\mu_2(y_1,d_2,x_2) = \mu_2(y_1,d_2',x_2')$, and $\mu_2(y_1',d_2,x_2) = \mu_2(y_1',d_2',x_2'')$. Before closing this section, it is worth mentioning that, in estimating the identified parameters, one can improve efficiency by aggregating the conditional expectations (4.6) with respect to x_t' over the following set:

$$\lambda_t(x_t;\boldsymbol{z}^{t-1},\boldsymbol{x}^{t-1},\boldsymbol{d}^{t-1},y_{t-1}) \equiv \{\tilde{x}_t : h_t(z_t,\tilde{z}_t,x_t,\tilde{x}_t;\boldsymbol{z}^{t-1},\boldsymbol{x}^{t-1},\boldsymbol{d}^{t-1},y_{t-1}) = 0 \text{ for some } (z_t,\tilde{z}_t)\}.$$

5 Partial Identification

Suppose Assumption SP does not hold in that X_t does not exhibit rectangular variation, or there is no X_t that is excluded from the selection equation at time t. In this case, we partially identify the ASF's and ATE's. The partial identification of $\boldsymbol{d}^*(\boldsymbol{x})$ (or $\boldsymbol{d}^\dagger(\boldsymbol{x})$) may not yield informative bounds unless there are a sufficient number of ATE's whose bounds are informative about their signs.

We briefly illustrate the calculation of bounds on the ASF when the sufficient rectangular variation is not guaranteed.[8] For each unknown term $E[Y_{T,\boldsymbol{d}}|\boldsymbol{x},\boldsymbol{z},\boldsymbol{d}^{t-1},d_t']$ in the proof of Theorem 4.1, we can calculate its upper and lower bounds depending on the sign of $\mu_t(y_{t-1},1,x_t)-\mu_t(y_{t-1},0,\tilde{x}_t)$, which is identified in Lemma 4.1. Under the situation considered in this section, $\tilde{x}_t$ does *not* necessarily differ from x_t. To begin with a simple example, for given $t=2,...,T$, suppose $\mu_t(y_{t-1},d_t,x_t)-\mu_t(y_{t-1},d_t',x_t')\geq 0$ *for all* y_{t-1} where x_t' is allowed to equal x_t. Then, by the definition of the $\boldsymbol{U}$-set, it satisfies that $\mathcal{U}_{\boldsymbol{d}}(\boldsymbol{x};1)\supseteq\mathcal{U}_{d_t',\boldsymbol{d}_{-t}}(x_t',\boldsymbol{x}_{-t};1)$ regardless of the value of y_{t-1}, which yields a lower bound as

$$\begin{aligned}
E[Y_{T,\boldsymbol{d}}|\boldsymbol{x},\boldsymbol{z},\boldsymbol{d}^{t-1},d_t'] &= \Pr[\boldsymbol{U}\in\mathcal{U}_{\boldsymbol{d}}(\boldsymbol{x};1)|\boldsymbol{x},\boldsymbol{z},\boldsymbol{V}^t\in\mathcal{V}^*_{\boldsymbol{d}^{t-1},d_t'}(\boldsymbol{U}^{t-1})]\\
&= \Pr[\boldsymbol{U}\in\mathcal{U}_{\boldsymbol{d}}(\boldsymbol{x};1)|\boldsymbol{V}^t\in\mathcal{V}^*_{\boldsymbol{d}^{t-1},d_t'}(\boldsymbol{U}^{t-1})]\\
&\geq \Pr[\boldsymbol{U}\in\mathcal{U}_{d_t',\boldsymbol{d}_{-t}}(x_t',\boldsymbol{x}_{-t};1)|\boldsymbol{V}^t\in\mathcal{V}^*_{\boldsymbol{d}^{t-1},d_t'}(\boldsymbol{U}^{t-1})]\\
&= \Pr[\boldsymbol{U}\in\mathcal{U}_{d_t',\boldsymbol{d}_{-t}}(x_t',\boldsymbol{x}_{-t};1)|x_t',\boldsymbol{x}_{-t},\boldsymbol{z},\boldsymbol{V}^t\in\mathcal{V}^*_{\boldsymbol{d}^{t-1},d_t'}(\boldsymbol{U}^{t-1})]\\
&= E[Y_{T,(d_t',\boldsymbol{d}_{-t})}|x_t',\boldsymbol{x}_{-t},\boldsymbol{z},\boldsymbol{d}^{t-1},d_t'].
\end{aligned}$$

As a more realistic example, when the sign of $\mu_t(y_{t-1},d_t,x_t)-\mu_t(y_{t-1},d_t',x_t')$ is identified *for each* y_{t-1}, it is possible to calculate the bounds by (4.5). For instance, if $\mu_t(1,d_t,x_t)\geq\mu_t(1,d_t',x_t')$, then

$$E[Y_{T,\boldsymbol{d}}|\boldsymbol{x},\boldsymbol{z},\boldsymbol{d}^{t-1},d_t',Y_{t-1}=1]\geq E[Y_T|x_t',\boldsymbol{x}_{-t},\boldsymbol{z},\boldsymbol{d}^{t-1},d_t',Y_{t-1}=1]$$

since

$$\begin{aligned}
&E[Y_{T,\boldsymbol{d}}|\boldsymbol{x},\boldsymbol{z},\boldsymbol{d}^{t-1},d_t',Y_{t-1}=1]\\
=&\Pr[\boldsymbol{U}\in\mathcal{U}_{\boldsymbol{d}}(\boldsymbol{x};1)|\boldsymbol{x},\boldsymbol{z},\boldsymbol{V}^t\in\mathcal{V}^*_{\boldsymbol{d}^{t-1},d_t'}(\boldsymbol{U}^{t-1}),U_{t-1}\in\mathcal{U}_{\boldsymbol{d}^{t-1}}(\boldsymbol{x}^{t-1};1)]\\
=&\Pr[\boldsymbol{U}\in\mathcal{U}_{\boldsymbol{d}}(\boldsymbol{x};1)|\boldsymbol{V}^t\in\mathcal{V}^*_{\boldsymbol{d}^{t-1},d_t'}(\boldsymbol{U}^{t-1}),U_{t-1}\in\mathcal{U}_{\boldsymbol{d}^{t-1}}(\boldsymbol{x}^{t-1};1)]
\end{aligned}$$

and $\mathcal{U}_{\boldsymbol{d}}(\boldsymbol{x};1)\supseteq\mathcal{U}_{d_t',\boldsymbol{d}_{-t}}(x_t',\boldsymbol{x}_{-t};1)$ given $Y_{t-1,\boldsymbol{d}^{t-1},\boldsymbol{x}^{t-1}}=1$. Once the bounds on $E[Y_{T,\boldsymbol{d}}|\boldsymbol{x},\boldsymbol{z},\boldsymbol{d}^{t-1},d_t']$ are established, the bounds on $E[Y_{T,\boldsymbol{d}}|\boldsymbol{x}]=E[Y_{T,\boldsymbol{d}}|\boldsymbol{x},\boldsymbol{z}]$ can be calculated using the iterative scheme introduced in the proof of Theorem 4.1. Lastly, depending on how much we learn about the signs of the ATE's, we may be able to construct informative bounds on $\boldsymbol{d}^*(\boldsymbol{x})$, which will be expressed as strict subsets of $\mathcal{D}$.

References

Angrist, J. D., Imbens, G. W., 1995. Two-stage least squares estimation of average causal effects in models with variable treatment intensity. Journal of the American statistical Association 90 (430), 431–442. 2

Blevins, J. R., 2014. Nonparametric identification of dynamic decision processes with discrete and continuous choices. Quantitative Economics 5, 531–554. 1

[8] The case where X_t does not exist can be dealt in a similar way, which is omitted.

Buchholz, N., Shum, M., Xu, H., 2016. Semiparametric estimation of dynamic discrete choice models. 1

Cunha, F., Heckman, J. J., Schennach, S. M., 2010. Estimating the technology of cognitive and noncognitive skill formation. Econometrica 78 (3), 883–931. 1

Han, S., 2018. Multiple treatments with strategic interaction. 4

Heckman, J. J., Humphries, J. E., Veramendi, G., 2016. Dynamic treatment effects. Journal of econometrics 191 (2), 276–292. 1, 6

Heckman, J. J., Navarro, S., 2007. Dynamic discrete choice and dynamic treatment effects. Journal of Econometrics 136 (2), 341–396. 1, 6

Murphy, S. A., 2003. Optimal dynamic treatment regimes. Journal of the Royal Statistical Society: Series B (Statistical Methodology) 65 (2), 331–355. 1, 2, 5

Murphy, S. A., van der Laan, M. J., Robins, J. M., Group, C. P. P. R., 2001. Marginal mean models for dynamic regimes. Journal of the American Statistical Association 96 (456), 1410–1423. 1, 2, 3

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. Mathematical Modelling 7 (9-12), 1393–1512. 1, 2

Robins, J., 1987. A graphical approach to the identification and estimation of causal parameters in mortality studies with sustained exposure periods. Journal of chronic diseases 40, 139S–161S. 1, 2

Rust, J., 1987. Optimal replacement of gmc bus engines: An empirical model of harold zurcher. Econometrica: Journal of the Econometric Society, 999–1033. 1

Shaikh, A. M., Vytlacil, E. J., 2011. Partial identification in triangular systems of equations with binary dependent variables. Econometrica 79 (3), 949–955. 1

Torgovitsky, A., 2016. Nonparametric inference on state dependence with applications to employment dynamics. Tech. rep., University of Chicago. 1

Vikström, J., Ridder, G., Weidner, M., 2016. Bounds on treatment effects on transitions. 1

Vytlacil, E., Yildiz, N., 2007. Dummy endogenous variables in weakly separable models. Econometrica 75 (3), 757–779. 1, 4, 4