cemmap
centre for microdata methods and practice

Identification and Estimation of Dynamic Structural Models with Unobserved Choices

Yingyao Hu
Yi Xin

The Institute for Fiscal Studies
Department of Economics,
UCL

cemmap working paper CWP35/19

E·S·R·C ECONOMIC & SOCIAL RESEARCH COUNCIL
An ESRC Research Centre

Identification and Estimation of Dynamic Structural Models with Unobserved Choices *

Yingyao Hu †
Johns Hopkins University
Yi Xin ‡
California Institute of Technology

June 12, 2019

Abstract

This paper develops identification and estimation methods for dynamic structural models when agents' actions are unobserved by econometricians. We provide conditions under which choice probabilities and latent state transition rules are nonparametrically identified with a continuous state variable in a single-agent dynamic discrete choice model. Our identification results extend to (1) models with serially correlated unobserved heterogeneity and continuous choices, (2) cases in which only discrete state variables are available, and (3) dynamic discrete games. We apply our method to study moral hazard problems in US gubernatorial elections. We find that the probabilities of shirking increase as the governors approach the end of their terms.

Keywords: dynamic discrete choice models, unobserved choice, moral hazard, gubernatorial elections

JEL Code: C10, C14, C18, C51, D72, D82

*This paper is a revised version of the second chapter of Xin's PhD thesis. The usual disclaimer applies.

†Department of Economics, Johns Hopkins University. Email: yhu@jhu.edu.

‡Division of the Humanities and Social Sciences, California Institute of Technology. Email: yixin@caltech.edu.

1 Introduction

In a revealed preference framework, choices made by agents reflect their underlying preferences, thus are the key ingredients to further economic analysis. In reality, however, agents' decisions may not be directly observed by researchers. In a principal-agent framework, moral hazard problems occur when the agent's actions impose an externality on the principal but cannot be directly observed. Models with hidden actions have been applied to many contexts in economics and political sciences. For example, in credit markets, borrowers may have incentives to invest in riskier projects, which increases the default probability, but the investment decisions may not be perfectly monitored by lenders; a politician knows the amount of time and effort he spends generating economic growth, but this action is unlikely to be observed by voters. In such contexts where actions are private information, it is almost impossible for researchers to observe agents' choices.[1] An important research question therefore arises: when actions are hidden, can we still uncover the decision-making process and infer the preferences of the agents from the data?

In this paper, we study identification and estimation of dynamic structural models when agents' choices are *not* observed by econometricians. In the existing literature on dynamic discrete choice models, researchers mainly focused on the cases in which choices are observables.[2] Examples include manager's engine replacement decisions in Rust (1987), parental contraceptive choices in Hotz and Miller (1993), occupational choices in Keane and Wolpin (1997), employees' retirement decisions in Rust and Phelan (1997), retail firms' inventory strategies in Aguirregabiria (1999), and water authorities' pricing behavior in Timmins (2002), etc. The identification results of dynamic structural models in previous works require the observation of choices and state variables for a random set of agents for a period of time (see Rust, 1994; Magnac and Thesmar, 2002; Aguirregabiria, 2010; Abbring, 2010; and Norets and Tang, 2014). For estimation of this class of models, agents' choices are needed to construct (pseudo) likelihood or to do first-stage nonparametric estimation of the conditional choice probabilities (CCP's) and the state transition probabilities (see Rust, 1987; Hotz and Miller, 1993; Hotz et al., 1994; Aguirregabiria and Mira, 2002).

Given that the existing approaches are not generally effective when agents' choices are not observed by econometricians, in this paper, we propose new identification and estimation methods for dynamic structural models with unobserved choice variables. We consider a

[1] Another reason that choices may be hard for researchers to obtain relates to data collecting issues. For example, in many survey datasets, some key decisions, such as agents' investments in human capital, health, and child development etc, are not reported (or inquired).

[2] See Aguirregabiria and Mira (2010) for a comprehensive survey on dynamic discrete choice structural models.

single-agent finite-horizon dynamic discrete choice model with a continuous state variable in the baseline analysis. We specify the state transition process through a nonparametric regression model with an additive error and assume that the unobserved choices may shift the distribution of the future state but are independent with the error term conditional on the current state. The key intuition of our identification results is as follows. In a finite-horizon dynamic structural model, agents' choice probabilities are usually time-varying. For example, when an executive in a firm is close to retirement, he/she may have less incentives to exert effort; the probability of shirking may exhibit an upward trending. However, the stationarity of the state transition process is typically considered as an innocuous assumption in the literature.[3] In the executive's example, this assumption means that conditional on him/her working hard, the distribution of the future state given a fixed current state will remain the same no matter whether the executive is close to retirement or not. In the data, the differences in the observed state transition process across periods are driven by the differences in choice probabilities. Therefore, by exploiting variations in moments of the observed future state distributions across periods, we identify the unobserved choice probabilities and the latent state transition process.

In this paper, we consider several extensions to our baseline model. First, we incorporate individual serially correlated unobserved heterogeneity into the dynamic discrete choice model when choices are unobserved. Existing papers by Aguirregabiria and Mira (2007), Houde and Imai (2006), Kasahara and Shimotsu (2009), and Hu and Shum (2012) have provided solutions to deal with unobserved heterogeneity. Following Hu and Shum (2012), we use joint distribution of the observed state variable at four consecutive periods to identify the transition of the observed state conditional on the unobserved heterogeneity, with which we can apply our method directly to deal with unobserved choices. Second, we discuss the identification for infinite-horizon models. In finite-horizon models, time essentially serves as an exclusion restriction. We show that as long as there is an excluded variable that only shifts choice probabilities but does not affect the latent state transition process, the baseline identification results remain valid. Third, we provide conditions under which unobserved choice probabilities and the latent state transition process are identified when only discrete state variables are available. Our results rely on the assumption that the transition process of two discrete state variables are independent conditional on the agent's choice. When this assumption holds, intuitively, the future states can be viewed as "measurements" of the unobserved choice. If two continuous state variables are available, it is straightforward to extend our results to allow for continuous choices.

Our identification results are not limited to single-agent dynamic models. We also show

[3] See the discussions on stationary Markovian policy function in Rust (1987).

in this paper that the proposed approach can be extended to dynamic discrete games of incomplete information. In a game setting, multiple players interact with each other and make decisions simultaneously. Their choices naturally depend on the actions and states of other players. In some cases, however, it is reasonable to assume that the state transition process for a player only depends on his own actions and state variables in the past.[4] When this assumption holds, state of other players can be treated as an excluded variable (i.e., it only affects the choice probabilities, but not the state transition process), hence our identification results for single-agent models can be applied to deal with unobserved choices in dynamic discrete games.

Following our identification strategies, we propose a sieve maximum likelihood estimation strategy for the nonparametric functions in the state transition process and the agent's utility primitives. We conduct Monte Carlo simulations to examine the finite sample performances of our estimator. We also apply our method to study moral hazard problems in US gubernatorial elections. Specifically, we estimate a dynamic discrete choice model for governors' effort-exerting decisions in the United States from 1950–2000. In our model, governors' choices are not directly observed by voters or econometricians, but have an impact on the state variable (log per capita spending). Our empirical analysis suggests that the probabilities of shirking increase as the governors approach the end of their terms; the shirking probability is 31 percent higher in the last period compared to that in the first year. We also find governors who were serving their second terms are more likely to exert effort, potentially because of the selection effect of elections.

This paper is, to the best of our knowledge, the first to incorporate unobserved choice variables into a general framework of dynamic discrete choice models. There are few empirical papers focused on models with unobserved choices. Misra and Nair (2011) investigate sales-forces' dynamic effort allocation, treating unobserved effort levels as time-specific fixed effects. Copeland and Monnet (2009) consider a dynamic model of effort decisions under non-linear incentive schemes; their identification results rely on the exogenous variations in the threshold in the firm's daily bonus plan. Gayle and Miller (2015) study models of managerial compensation and assume that some levels of revenue can only be achieved through high effort. Perrigne and Vuong (2011) focus on a false moral hazard model, in which effort, though unobserved, is a deterministic function of type that can be backed out one-to-one from observed prices.[5] Xin (2019) studies adverse selection and moral hazard problems in

[4]For example, in a dynamic oligopoly game where the state variable is the firm's capacity levels and the choice is the incremental changes to capacity, it is reasonable to assume that firm's future capacity levels only depend on its own decisions, not on other firms' choices. See Aguirregabiria, Mira, and Roman (2007), Ryan (2012), Collard-Wexler (2013), and Takahashi (2015) for more details on empirical models of oligopoly dynamics.

[5]The identification arguments in Perrigne and Vuong (2011) relates to the nonparametric identification

online credit markets, where borrowers' default and late payment performances are used as measurements of the unobserved effort choices. Our paper differs from these preceding papers in the sense that we impose general assumptions on the state transition process. Our identification strategies do not rely on multiple measurements of effort levels, exogenous variations in incentive schemes, or one-to-one mapping between effort levels and observables.[6] In this paper, we also provide identification results for dynamic discrete games with unobserved choices. The existing papers that develop estimation techniques for dynamic discrete games all require the observation of choices (see Jofre-Bonet and Pesendorfer, 2003; Aguirregabiria and Mira, 2007; Bajari, Benkard, and Levin, 2007; Pakes, Ostrovsky, and Berry, 2007; Pesendorfer and Schmidt-Dengler, 2008, etc).

This paper is also related to the empirical literature in political economy focused on understanding the impact of institutional design of election rules (e.g., term limits) on politicians' performances (see Besley and Case, 1995; Alt, Bueno de Mesquita, and Rose, 2011; Aruoba, Drazen, and Vlaicu, 2019; Sieg and Yoon, 2017). We study governors' dynamic effort-exerting decisions within a term and provide new empirical evidence on moral hazard problems in US gubernatorial elections.

The rest of the paper is organized as follows. We outline a standard dynamic discrete choice model in Section 2. Identification and estimation results for the baseline model are provided in Sections 3 and 4, respectively. Section 5 provides simulation results. We consider extensions to the baseline model in Section 6 and apply our methods to study moral hazard problems in gubernatorial elections in Section 7. Section 8 concludes.

2 A Basic Model

We first fix the notation for a standard single-agent dynamic discrete choice model with $t = 0, 1, \cdots, T < \infty$. Let s_t represent the observed state variable and y_t denote agent's choice. ε_t represents the state variable that is unobserved to econometricians, such as utility shocks. An agent's flow utility depends on the current state and the choice, i.e. $u_t(s_t, \varepsilon_t, y_t)$. The sum of discounted utility stream of the agent is therefore defined as

$$U(\boldsymbol{s}, \boldsymbol{\varepsilon}, \boldsymbol{y}) = \sum_{t=0}^{T} \beta^t u_t(s_t, \varepsilon_t, y_t), \tag{2.1}$$

results in the auction literature, see Guerre, Perrigne, and Vuong (2000).

[6] In the case where only discrete state variables are available, we need multiple measurements of effort levels to identify the model primitives. The details are provided in Section 6.3.

where $\boldsymbol{s} = (s_0, \cdots, s_T)$, $\boldsymbol{\varepsilon} = (\varepsilon_0, \cdots, \varepsilon_T)$, $\boldsymbol{y} = (y_0, \cdots, y_T)$, and β is the discount factor. The agent's problem is to choose an optimal decision rule $\delta = (\delta_0, \cdots, \delta_T)$ that maximizes the expected sum of the discounted utility, i.e.

$$\max_{\delta=(\delta_0,\cdots,\delta_T)} \mathrm{E}(U(\boldsymbol{s}, \boldsymbol{\varepsilon}, \boldsymbol{y})),$$

where expectation is with respect to the partially controlled stochastic process of $\{s_t, \varepsilon_t, y_t\}$ induced by the decision rule δ. We now introduce the first assumption to restrict attention to certain classes of models.

Assumption 1. *The dynamic process of* $\{s_t, \varepsilon_t, y_t\}$ *satisfies*

(i) First-order Markov: $f_{s_{t+1},\varepsilon_{t+1},y_{t+1}|s_t,\varepsilon_t,y_t,\Omega_{<t}} = f_{s_{t+1},\varepsilon_{t+1},y_{t+1}|s_t,\varepsilon_t,y_t}$, *where* $\Omega_{<t} \equiv \{s_{t-1}, \cdots, s_0, \varepsilon_{t-1}, \cdots, \varepsilon_0, y_{t-1}, \cdots, y_0\}$.

(ii) The distribution of s_{t+1} *given* $(s_t, \varepsilon_t, y_t)$ *only depends on* (s_t, y_t) *and is denoted by* $f_{s_{t+1}|s_t,y_t}$*; the distribution of* ε_{t+1} *given* $(s_{t+1}, s_t, \varepsilon_t, y_t)$ *only depends on* s_{t+1} *and is denoted by* $f_{\varepsilon_{t+1}|s_{t+1}}$.

(iii) State transition probabilities $f_{s_{t+1}|s_t,y_t}$ *are time-invariant.*

Assumption 1(i), which imposes the first-order Markov property on the transition process of $\{s_t, \varepsilon_t, y_t\}$, is commonly adopted in the dynamic discrete choice framework and may be easily relaxed to allow for higher-order Markov process. Following Rust (1987), Assumption 1(ii) highlights two types of conditional independence: (1) given the state s_t, ε's are independent over time; and (2) conditional on the current state s_t and choice y_t, the future state s_{t+1} is independent of the unobserved state ε_t. The relaxation of this assumption is discussed in a recent literature on identification and estimation of dynamic discrete choice models when the unobserved state variables are serially correlated (see Aguirregabiria and Mira, 2007; Houde and Imai, 2006; Kasahara and Shimotsu, 2009; Hu and Shum, 2012). In Section 6.1, we show that our identification results can be easily generalized for the model that incorporates serially correlated unobserved heterogeneity when at least five periods are available in the data. In order to highlight the identification intuition related to unobserved choice variables, we first focus on the case when Assumption 1(ii) is invoked. Assumption 1(iii) which guarantees the stationarity of the state transition process is usually invoked in infinite-horizon dynamic models (see Rust (1987)). When choice variables are available, this assumption can be directly tested using the data. The dynamic process of the state and choice variables (s_t, y_t) that satisfies Assumption 1 is illustrated in Figure 1.

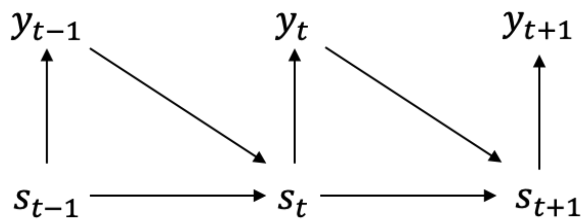

Figure 1: The Dynamic Process of (s_t, y_t)

Under Assumption 1, we represent the agent's optimization problem using the Bellman's equation as follows.

$$V_t(s_t, \varepsilon_t) = \max_y u(s_t, \varepsilon_t, y) + \beta\, \mathrm{E}[V_{t+1}(s_{t+1}, \varepsilon_{t+1})|s_t, y]. \tag{2.2}$$

The agent's decision rule is hence defined by

$$\delta_t(s_t, \varepsilon_t) = \arg\max_y \left\{ u(s_t, \varepsilon_t, y) + \beta\, \mathrm{E}[V_{t+1}(s_{t+1}, \varepsilon_{t+1})|s_t, y] \right\}. \tag{2.3}$$

At period t, the choice probability of alternative y_t conditional on the observed state s_t (also abbreviated as CCP) is defined in the following equation.

$$p_t(y_t|s_t) = \int \mathbf{1}\left\{y_t = \delta_t(s_t, \varepsilon)\right\} dF_{\varepsilon_t|s_t}(\varepsilon|s_t), \tag{2.4}$$

where $F_{\varepsilon_t|s_t}(\cdot|\cdot)$ denotes the cumulative density function of the unobserved state variable ε_t conditional on the current state s_t.

For the model described above, if the choice variable y_t is observed at each period, the two-step CCP method developed by Hotz and Miller (1993) can be easily adopted—in the first step the choice and state transition probabilities are nonparametrically identified and estimated. However, when y_t is not observed by econometricians, we cannot recover the decision rules and the state transition probabilities directly from the data in the first step. As a result, the existing methods fail to obtain sufficient ingredients for identifying and estimating structural primitives.

3 Identification

In this section, we provide new identification strategies to recover the unobserved choice probabilities $p_t(y_t|s_t)$ and latent state transition probabilities $f_{s_{t+1}|s_t,y_t}$ when only $\{s_t\}_{t=1}^T$

is observed. We focus on the case when s_t represents a continuous state variable in this section; the identification results of cases in which only discrete state variables are available are provided in Section 6.3. To highlight the feature that the choice variable is unobserved to econometricians, we use y_t^* to denote the unobserved choice variable hereafter.

When agents' choices are unobserved, neither conditional choice probabilities nor state transition rules can be directly recovered from the data. However, these two sets of unknowns are connected through the observed state transition process as shown in the following equation under Assumption 1(i)–(ii).

$$f_{s_{t+1}|s_t}(s'|s) = \sum_{y_t^*} f_{s_{t+1}|s_t,y_t^*}(s'|s, y_t^*)p_t(y_t^*|s), \tag{3.1}$$

where s' and s represent a realized value of s_{t+1} and s_t, respectively. In Equation (3.1), the probability density of the future state conditional on the current state is a mixture of the true latent state transition probabilities conditional on different alternatives; and the choice probabilities serve as the mixing weights. Under Assumption 1(iii), $f_{s_{t+1}|s_t,y_t^*}(s'|s, y_t^*)$ is stationary; while in finite-horizon models, $p_t(y_t^*|s)$ varies across different periods. The differences in $f_{s_{t+1}|s_t}(s'|s)$ across periods are therefore driven by the non-stationarity of the choice probabilities. In the rest of this section, we explore variations in moments of the observed state transition process to identify choice probabilities and latent state transition rules, for which the following assumption is invoked.

Assumption 2. $s_{t+1} = m(y_t^*, s_t) + \eta_t$, *where* $E(\eta_t|s_t) = 0$ *and* $\eta_t \perp y_t^*|s_t$.

Assumption 2 specifies the transition process of the continuous state variable s_t through a nonparametric regression model, where $m(\cdot, \cdot)$ is an unknown function and η_t represents the random shock realized in the transition process with conditional mean equal to zero. This assumption also requires that the regression error is independent of the unobserved choice conditional on the state variable. This ensures that the impact of the choices on the state transition process is only through the deterministic part but not through the error term. In other words, the choice variables only shifts the mean of the future state distribution. Combining Assumption 2 and Assumption 1(iii), we know that the conditional distribution of η_t is also stationary. That is, for $t, \tau \in \{0, 1, \cdots, T\}$, $f_{\eta_t|S_t}(\eta|s) = f_{\eta_\tau|S_\tau}(\eta|s)$, $\forall \eta, s$. By Assumption 2, the unknown function $m(\cdot, \cdot)$ and the conditional distribution of η_t jointly determine the state transition probabilities $f_{s_{t+1}|s_t,y_t^*}$, and thus are the key primitives in addition to the unobserved choice probabilities.

For illustration, we now consider the case in which the choice variable takes binary values, i.e. $y_t^* \in \{0, 1\}$. Identifying the function $m(y_t^*, s_t)$ is then equivalent to identifying two functions of s_t, i.e., $m(y_t^* = 0, s_t)$ and $m(y_t^* = 1, s_t)$. In the rest of this section, we consider the identification of model primitives for a fixed state s. Let $m_1 = m(1, s)$ and $m_0 = m(0, s)$. For the choice probability at period t, let $p_t = p_t(1|s)$. We define the first-, the second- and the third-order conditional moments of the observed state variable at $t + 1$ as follows.

$$
\begin{aligned}
\mu_{t+1} &= \mathrm{E}_{t+1}\left[s_{t+1}|s_t = s\right], \\
\nu_{t+1} &= \mathrm{E}_{t+1}\left[(s_{t+1} - \mu_{t+1})^2|s_t = s\right], \\
\xi_{t+1} &= \mathrm{E}_{t+1}\left[(s_{t+1} - \mu_{t+1})^3|s_t = s\right].
\end{aligned}
$$

Notice that all of these conditional moments can be directly estimated from the data, and are thus treated as known constants for identification purposes.

We rewrite the first-order conditional mean of the state variable at period $t + 1$ by replacing s_{t+1} with $m(y_t^*, s_t) + \eta_t$. Specifically,

$$
\mu_{t+1} = \sum_{y_t^*} p_t(y_t^*|s)\,\mathrm{E}_{t+1}[m(y_t^*, s) + \eta_t|s, y_t^*] = p_t m_1 + (1 - p_t) m_0, \tag{3.2}
$$

where the second equation holds because under Assumption 2, η_t and y_t^* are independent conditional on the state and $\mathrm{E}(\eta_t|s) = 0$. In Equation (3.2), μ_{t+1} is a weighted average of m_1 and m_0 with the choice probabilities $(p_t, 1 - p_t)$ serving as the mixing weights. Following similar arguments, we rewrite the second- and the third-order conditional moments of the state variable as follows.

$$
\begin{aligned}
\nu_{t+1} &= \sum_{y_t^*} p_t(y_t^*|s)\,\mathrm{E}_{t+1}\left[(m(y_t^*, s) + \eta_t - \mu_{t+1})^2|s, y_t^*\right] \\
&= \sum_{y_t^*} p_t(y_t^*|s)\left[(m(y_t^*, s) - \mu_{t+1})^2 + 2(m(y_t^*, s) - \mu_{t+1})\,\mathrm{E}[\eta_t|s] + \mathrm{E}[\eta_t^2|s]\right] \\
&= p_t(m_1 - \mu_{t+1})^2 + (1 - p_t)(m_0 - \mu_{t+1})^2 + \mathrm{E}\left[\eta_t^2|s\right]
\end{aligned} \tag{3.3}
$$

$$
\begin{aligned}
\xi_{t+1} &= \sum_{y_t^*} p_t(y_t^*|s)\,\mathrm{E}_{t+1}\left[(m(y_t^*,s)+\eta_t-\mu_{t+1})^3|s,y_t^*\right] \\
&= \sum_{y_t^*} p_t(y_t^*|s)\Bigg[(m(y_t^*,s)-\mu_{t+1})^3+\mathrm{E}[\eta_t^3|s]+3(m(y_t^*,s)-\mu_{t+1})^2\,\mathrm{E}[\eta_t|s] \\
&\qquad\qquad +3(m(y_t^*,s)-\mu_{t+1})\,\mathrm{E}[\eta_t^2|s]\Bigg] \\
&= p_t(m_1-\mu_{t+1})^3+(1-p_t)(m_0-\mu_{t+1})^3+E\left[\eta_t^3|s\right].
\end{aligned}
\tag{3.4}
$$

In Equations (3.3) and (3.4), $\mathrm{E}\left[\eta_t^2|s\right]$ and $\mathrm{E}\left[\eta_t^3|s\right]$ represent the second and third order moments of the error term respectively, but the values of these terms are not known.

To identify m_1, m_0, and the choice probabilities, we consider two periods t and τ along the dynamic process. Based on Equation (3.2), we have

$$
\begin{aligned}
\mu_{t+1} &= p_t m_1 + (1-p_t)m_0, \\
\mu_{\tau+1} &= p_\tau m_1 + (1-p_\tau)m_0.
\end{aligned}
$$

This system of two linear equations leads to the identification of p_t and p_τ for any given m_0 and m_1 as long as $m_0 \neq m_1$. Specifically,

$$
p_t = \frac{\mu_{t+1}-m_0}{m_1-m_0}, \quad p_\tau = \frac{\mu_{\tau+1}-m_0}{m_1-m_0}. \tag{3.5}
$$

Under Assumption 1(iii) and Assumption (2), we know that the conditional distribution of η_t is stationary, which implies that the higher order moments of the error term are time-invariant conditional on the same state s, i.e.,

$$
\mathrm{E}\left[\eta_t^2|s\right]=\mathrm{E}\left[\eta_\tau^2|s\right], \quad \mathrm{E}\left[\eta_t^3|s\right]=\mathrm{E}\left[\eta_\tau^3|s\right].
$$

By taking the difference of Equations (3.3) and (3.4) across the two periods t and τ, we get rid of the unknown moments of η_t and achieve the following two equations.

$$
\begin{aligned}
\nu_{t+1}-\nu_{\tau+1} =& p_t(m_1-\mu_{t+1})^2+(1-p_t)(m_0-\mu_{t+1})^2-p_\tau(m_1-\mu_{\tau+1})^2-(1-p_\tau)(m_0-\mu_{\tau+1})^2 \\
=&(p_t-p_\tau)(m_1+m_0)(m_1-m_0)-(\mu_{t+1}^2-\mu_{\tau+1}^2),
\end{aligned}
\tag{3.6}
$$

$$
\xi_{t+1}-\xi_{\tau+1} = p_t(m_1-\mu_{t+1})^3+(1-p_t)(m_0-\mu_{t+1})^3-p_\tau(m_1-\mu_{\tau+1})^3-(1-p_\tau)(m_0-\mu_{\tau+1})^3 \tag{3.7}
$$

Plugging the expressions of p_t and p_τ in Equation (3.5) into Equations (3.6) and (3.7), we obtain a system of equations for the unknown primitives m_1 and m_0. Specifically,

$$\begin{aligned}
&\nu_{t+1} - \nu_{\tau+1} = (\mu_{t+1} - \mu_{\tau+1})\Delta_1 - (\mu_{t+1}^2 - \mu_{\tau+1}^2), \\
&\xi_{t+1} - \xi_{\tau+1} = (\mu_{t+1}\Delta_1 - \Delta_2 - \mu_{t+1}^2)(\Delta_1 - 2\mu_{t+1}) - (\mu_{\tau+1}\Delta_1 - \Delta_2 - \mu_{\tau+1}^2)(\Delta_1 - 2\mu_{\tau+1}),
\end{aligned} \tag{3.8}$$

where $\Delta_1 = m_1 + m_0$ and $\Delta_2 = m_1 m_0$. It is easy to get analytical solutions for Δ_1 and Δ_2 from Equation (3.8).

$$\begin{aligned}
\Delta_1 &= \frac{\nu_{t+1} - \nu_{\tau+1} + (\mu_{t+1}^2 - \mu_{\tau+1}^2)}{\mu_{t+1} - \mu_{\tau+1}}, \\
\Delta_2 &= \frac{\xi_{t+1} - \xi_{\tau+1} - (\mu_{t+1}(\Delta_1 - \mu_{t+1})(\Delta_1 - 2\mu_{t+1}) - \mu_{\tau+1}(\Delta_1 - \mu_{\tau+1})(\Delta_1 - 2\mu_{\tau+1}))}{2(\mu_{t+1} - \mu_{\tau+1})}.
\end{aligned}$$

With Δ_1 and Δ_2 identified using the moments of the observed state transition process, m_0 and m_1 are the two roots of the equation $m^2 - \Delta_1 m + \Delta_2 = 0$ provided that $\Delta_1^2 - 4\Delta_2 > 0$. This condition can be directly tested from the data. To further decide the order of m_0 and m_1, we invoke the following assumption.

Assumption 3 (First Order Stochastic Dominance)**.** $F_{s_{t+1}|s_t,Y_t^*}(\cdot|s, y_t^* = 1)$ *first-order stochastically dominates* $F_{s_{t+1}|s_t,Y_t^*}(\cdot|s, y_t^* = 0)$.

Assumption 3 implies that $m_1 \geq m_0$ because $m_1 = \mathrm{E}(s_{t+1}|s, y_t^* = 1) \geq \mathrm{E}(s_{t+1}|s, y_t^* = 0) = m_0$. Intuitively, consider an example where s_t represents the outcome of the loan and y^* represents whether a borrower exerts effort to pay off the debt. In this case, it is reasonable to assume that when borrowers exert effort, the outcome distribution first-order stochastically dominates the one when borrowers exert no effort. Assumption 3 gives an example of how to decide the order of m_0 and m_1 from the state transition process; other assumptions arising from the model or consistent with the economic intuition would also work. Once m_0 and m_1 are recovered, it is straightforward to pin down the choice probabilities through Equation (3.5) provided that $m_0 \neq m_1$.

Last, we focus on the identification of the error term distribution. Given the additive structure of the state transition process and the independence of η_t and y_t^* conditional s_t, the observed state transition probability of $s_{t+1} = s'$ given $s_t = s$ can be written as a mixture of the conditional density of η_t evaluated at $s' - m_1$ and $s' - m_0$ with conditional choice probabilities serving as the mixing weights.

$$\begin{aligned}
f_{s_{t+1}|s_t}(s'|s) &= p_t f_{\eta_t|s_t}(s' - m_1|s) + (1 - p_t) f_{\eta_t|s_t}(s' - m_0|s), \\
f_{s_{\tau+1}|s_\tau}(s'|s) &= p_\tau f_{\eta_\tau|s_\tau}(s' - m_1|s) + (1 - p_\tau) f_{\eta_\tau|s_\tau}(s' - m_0|s).
\end{aligned} \tag{3.9}$$

Given the stationarity of η_t conditional on s_t,

$$f_{\eta_t|s_t}(s'-m_1|s) = f_{\eta_\tau|s_\tau}(s'-m_1|s), \quad f_{\eta_t|s_t}(s'-m_0|s) = f_{\eta_\tau|s_\tau}(s'-m_0|s).$$

Equation (3.9) identifies the conditional density function of η_t at $s'-m_1$ and $s'-m_0$ if p_t and p_τ are known and are not equal. Specifically,

$$\begin{aligned} f_{\eta_t|s_t}(s'-m_1|s) &= \frac{f_{s_{t+1}|s_t}(s'|s)(1-p_\tau) - f_{s_{\tau+1}|s_\tau}(s'|s)(1-p_t)}{p_t - p_\tau}, \\ f_{\eta_t|s_t}(s'-m_0|s) &= \frac{f_{s_{\tau+1}|s_\tau}(s'|s)p_t - f_{s_{t+1}|s_t}(s'|s)p_\tau}{p_t - p_\tau}. \end{aligned} \tag{3.10}$$

Notice that, the identification of p_t, p_τ, and the distribution of η_t requires that $m_0 \neq m_1$ and $p_t \neq p_\tau$. These two conditions are guaranteed if $\mu_{t+1} \neq \mu_{\tau+1}$, which is also empirically testable. We summarize the main identification results in the following theorem.

Theorem 1 (Identification)**.** *Suppose Assumptions 1–3 hold for the dynamic process of* $\{s_t, \varepsilon_t, y_t^*\}$*,* y_t^* *takes binary values, and* $\mu_{t+1} \neq \mu_{\tau+1}$*.* $f_{s_{t+1}|s_t}(\cdot|s)$ *and* $f_{s_{\tau+1}|s_\tau}(\cdot|s)$ *identify the latent state transition probabilities* $f_{s_{t+1}|s_t,y_t^*}(\cdot|s, y_t^*)$ *and the choice probabilities* $p_t(y_t^*|s)$ *and* $p_\tau(y_\tau^*|s)$ *for any* s *and* y_t^**.*

The economic intuition of Theorem 1 is as follows. In finite-horizon dynamic models, choice probabilities are non-stationary, while the stationarity of transition probabilities is often assumed. Conditional on the same state, the differences in observed future state distribution are driven by the differences in choice probabilities. Variations in moments (mean, variance, or higher order) of future state distributions create restrictions to identify the unobserved choice probabilities. We discuss several extensions of the main identification results in Section 6. In particular, we consider cases in which: (1) individual unobserved heterogeneity is allowed, (2) the model has infinite horizon, (3) only discrete state variables are available, and (4) there are multiple players making simultaneous decisions (i.e., dynamic discrete games).

4 Sieve Maximum Likelihood Estimation

Following our identification results in Section 3, we propose sieve maximum likelihood estimation (MLE) for the nonparametric function $m(\cdot,\cdot)$ and the distribution of the error term $f_{\eta_t|S_t}$ in the state transition process. Conditions under which the estimator is consistent are provided at the end of this section.

We first fix notations for our estimation strategies. Let $\theta^0 = (\alpha^0, m_0^0, m_1^0, f^0_{\eta_t|s_t})$ represent the vector of true parameter values of interest. $\alpha^0 \in \mathcal{A}$ is a vector of true utility parameters. $m_0^0 : \mathcal{S} \to \mathcal{S}$ and $m_1^0 : \mathcal{S} \to \mathcal{S}$ are two nonparametric functions in the state transition rules, where $\mathcal{S}$ denotes the state space. $f^0_{\eta_t|s_t} : \mathbb{R} \times \mathcal{S} \to \mathbb{R}^+$ is the probability density function of the error term conditional on the state variable. The sieve maximum likelihood estimator of θ^0 is denoted by $\hat{\theta}$.

We impose the following smoothness restrictions on m_0^0, m_1^0, and the density function $f^0_{\eta_t|S_t}$. To strengthen the definition of continuity, we introduce the notation for the space of Hölder continuous functions. If Ψ is an open set in $\mathbb{R}^n$, $\kappa \in \mathbb{N}$, and $\zeta \in (0, 1]$, then $\Gamma^{\kappa,\zeta}(\Psi)$ consists of all functions $m : \Psi \to \mathbb{R}$ with continuous partial derivatives in Ψ of order less than or equal to κ whose κ-th partial derivatives are locally uniformly Hölder continuous with exponent ζ in Ψ. Define a Hölder ball, which is a compact subset of $\Gamma^{\kappa,\zeta}(\Psi)$, as $\Gamma_c^{\kappa,\zeta}(\Psi) \equiv \left\{ m \in \Gamma^{\kappa,\zeta}(\Psi) \middle| \, \|m\|_{\Gamma^{\kappa,\zeta}(\Psi)} \leq c < \infty \right\}$ with respect to the norm

$$\|m\|_{\Gamma^{\kappa,\zeta}(\Psi)} \equiv \max_{|r|\leq\kappa} \sup_{\Psi} \|\partial^r m\|_e + \max_{|r|=\kappa} [\partial^r m]_{\zeta,\Psi}.$$

In the norm definition for the Hölder ball, $\|\cdot\|_e$ represents the Euclidean norm, and

$$[m]_{\zeta,\Psi} \equiv \sup_{x,x'\in\Psi, x\neq x'} \frac{\|m(x) - m(x')\|_e}{\|x - x'\|_e^{\zeta}}.$$

$\partial^r m$ represents the multi-index notation for partial derivatives with $r = (r_1, r_2, \cdots, r_{\dim(\Psi)})$ and $|r| = r_1 + r_2 + \cdots + r_{\dim(\Psi)}$. With the notations for the space of Hölder continuous functions, we define the functional space of m_0 and m_1 by $\mathcal{H} = \Gamma_c^{\kappa_1,\zeta_1}(\mathcal{S})$ with supremum norm $\|m\|_{\mathcal{H}} = \sup_{x\in\mathcal{S}} |m(x)|$. The space of the density function is

$$\mathcal{F} = \left\{ f_{\eta_t|s_t}(\cdot|\cdot) \in \Gamma_c^{\kappa_2,\zeta_2}(\mathbb{R}\times\mathcal{S}) : f_{\eta_t|s_t}(\cdot|s) > 0, \int_{\mathbb{R}} f_{\eta_t|s_t}(\eta|s) d\eta = 1, \mathrm{E}(\eta_t|s) = 0, \forall s \in \mathcal{S} \right\},$$

with norm defined by $\|f\|_{\mathcal{F}} = \sup_{x\in\mathbb{R}\times\mathcal{S}} \left| f(x)(1 + \|x\|_e^2)^{-\psi/2} \right|$, $\psi > 0$. Notice that the conditional mean of η_t for all density functions in $\mathcal{F}$ are equal to 0, which is consistent with Assumption 2. Let $\Theta = \mathcal{A} \times \mathcal{H} \times \mathcal{H} \times \mathcal{F}$ denote the space for all parameters of interest. Θ is an infinite-dimensional space as it contains nonparametric functions m_0, m_1, and $f_{\eta_t|s_t}$. The metric on Θ is defined by

$$d(\theta, \tilde{\theta}) = \|\alpha - \tilde{\alpha}\|_e + \|m_0 - \tilde{m}_0\|_{\mathcal{H}} + \|m_1 - \tilde{m}_1\|_{\mathcal{H}} + \left\| f_{\eta_t|s_t} - \tilde{f}_{\eta_t|s_t} \right\|_{\mathcal{F}}.$$

For $\theta = \{\alpha, m_0, m_1, f_{\eta_t|s_t}\} \in \Theta$, the log-likelihood evaluated at a single observation $D_i = \{s_{i,t}\}_{t=1}^T$ is derived in the following equation.

$$\begin{aligned} l(D_i;\theta) &= \sum_{t=1}^{T} \log\left(f_{s_{t+1}|s_t}(s_{i,t+1}|s_{i,t};\theta)\right) \\ &= \sum_{t=1}^{T} \log\left(f_{\eta_t|s_t}(s_{i,t+1}-m_1(s_{i,t}))p_{t,1}(s_{i,t};\theta) + f_{\eta_t|s_t}(s_{i,t+1}-m_0(s_{i,t}))p_{t,0}(s_{i,t};\theta)\right). \end{aligned} \tag{4.1}$$

In Equation (4.1), $p_{t,1}(s_{i,t};\theta)$ and $p_{t,0}(s_{i,t};\theta)$ are the choice probabilities for alternatives 1 and 0 conditional on state $s_{i,t}$ given the parameter value θ (including utility parameters and nonparametric functions m_0, m_1, and $f_{\eta_t|s_t}$ in the state transition rules). The population criterion function $Q : \Theta \to \mathbb{R}$ is hence defined by

$$Q(\theta) = \mathrm{E}(l(D_i;\theta)). \tag{4.2}$$

A sample counterpart of the objective function in Equation (4.2) is

$$\hat{Q}_n(\theta) = \frac{1}{n}\sum_{i=1}^{n} l(D_i;\theta). \tag{4.3}$$

In light of a finite sample, instead of searching parameters over an infinite-dimensional parameter space Θ, we use the sieve MLE to maximize the empirical criterion function over a sequence of approximating sieve spaces $\Theta_k = \mathcal{A} \times \mathcal{H}_{k_1} \times \mathcal{H}_{k_2} \times \mathcal{F}_{k_3}$, where

$$\begin{aligned} \mathcal{H}_{k_1} &= \left\{ m \in \mathcal{H} \,\middle|\, m : \mathcal{S} \to \mathbb{R}, m(s) = \sum_{q=1}^{k_1} \gamma_q h_q(s), \gamma_q \in \mathbb{R}, \forall q \right\} \\ \mathcal{H}_{k_2} &= \left\{ m \in \mathcal{H} \,\middle|\, m : \mathcal{S} \to \mathbb{R}, m(s) = \sum_{q=1}^{k_2} \gamma_q h_q(s), \gamma_q \in \mathbb{R}, \forall q \right\}, \\ \mathcal{F}_{k_3} &= \left\{ f \in \mathcal{F} \,\middle|\, f : \mathbb{R} \times \mathcal{S} \to \mathbb{R}^+, \sqrt{f(\eta|s)} = \boldsymbol{g}^{k_3}(\eta, s)^T \boldsymbol{\lambda}, \boldsymbol{\lambda} \in \mathbb{R}^{k_3} \right\}. \end{aligned}$$

In the definition of sieve spaces, $(h_1(\cdot), h_2(\cdot), h_3(\cdot), \cdots)$ represents a sequence of known basis functions, such as Hermite polynomials, power series, splines, etc. We use linear sieves to approximate square root of densities and $\boldsymbol{g}^k(\cdot,\cdot)$ is a $k \times 1$ vector of tensor product of spline basis functions on $\mathbb{R} \times \mathcal{S}$. Notice that it is standard to generate linear sieves of multivariate functions using tensor-product of linear sieves of univariate functions. With these settings, our sieve maximum likelihood estimator $\hat{\theta}_k$ is defined as

$$\hat{\theta}_k = \arg\sup_{\theta \in \Theta_k} \hat{Q}_n(\theta). \tag{4.4}$$

Chen (2007; Ch. 3) provides a general consistency theorem for sieve extremum estimators for various semi-/non-parametric models. Following Chen, Hu, and Lewbel (2008),Carroll, Chen, and Hu (2010), we provide lower level sufficient conditions tailored to our model for consistency of the sieve maximum likelihood estimator in Equation (4.4). [7]

Assumption 4 (Consistency)**.** *The following conditions are satisfied.*

(i) D_i *is i.i.d. across* i*;*

(ii) m_0 *and* $m_1 \in \mathcal{H}$ *with* $\kappa_1 + \zeta_1 > 1/2$*;* $f_{\eta|S} \in \mathcal{F}$ *with* $\kappa_2 + \zeta_2 > 1$*.*

(iii) $|Q(\theta^0)| < \infty$ *and* $Q(\theta)$ *is upper semicontinuous on* Θ *under the metric* $d(\cdot,\cdot)$*.*

(iv) There is a finite $\sigma > 0$ *and a random variable* $c(D_i)$ *with* $E(c(D_i)) < \infty$ *such that* $\sup_{\theta \in \Theta_k : d(\theta,\theta^0) \le \epsilon} |l(D_i;\theta) - l(D_i;\theta^0)| \le \epsilon^\sigma c(D_i)$*.*

(v) $k_1, k_2,$ *and* $k_3 \to \infty$*,* $k_1/n, k_2/n,$ *and* $k_3/n \to 0$*.*

Assumption 4 provides lower-level assumptions that imply the high-level conditions of Chen (2007; Ch. 3, Theorem 3.1). The following theorem for the consistency of our sieve maximum likelihood estimator is a direct application, therefore the proof is omitted.

Theorem 2 (Consistency)**.** *Suppose that all assumptions in Theorem 1 hold. If Assumption 4 is satisfied, then the sieve maximum likelihood estimator in Equation (4.4) is consistent with respect to the metric* $d(\cdot,\cdot)$*, i.e.,*

$$d(\hat{\theta}_k, \theta^0) = o_P(1).$$

Remark 1. *For general results on convergence rates, root-n asymptotic normality, and semi-parametric efficiency of sieve maximum likelihood estimators, see Shen and Wong (1994), Chen and Shen (1996), Shen (1997), Chen and Shen (1998), Ai and Chen (1999), Chen (2007; Theorem 3.2 and Theorem 4.3).*

[7] Chen, Hu, and Lewbel (2008) study identification and estimation of nonparametric regression model with discrete covariates measured with error. Carroll, Chen, and Hu (2010) consider general nonlinear errors-in-variables model using two samples.

5 Simulations

In this section, we present Monte Carlo simulation results when there is a continuous state variable. We assume that the utility function follows a very simple linear form

$$u(s_t, y_t^*) = \omega s_t - \rho y_t^*,$$

where $\omega = 0.8$ measures the marginal utility from higher values of the current state and $\rho = 0.3$ measures the marginal cost of exerting more effort. For this exercise, we consider a scenario in which choice variable only takes binary values $y_t^* \in \{0, 1\}$. The utility shock $\varepsilon_t(0)$ and $\varepsilon_t(1)$ independently follow the type I extreme value distribution and the discount factor is fixed at 0.95. We consider four data generating processes for the state transition process.

– DGP 1: $s_{t+1} = 0.8s_t + 0.5y_t^* + \eta_t$;

– DGP 2: $s_{t+1} = 0.8s_t + 0.5y_t^* + 0.3s_t \cdot y_t^* + \eta_t$.

– DGP 3: $s_{t+1} = 0.8s_t + 0.05s_t^2 + 0.5y_t^* + \eta_t$;

– DGP 4: $s_{t+1} = 0.2s_t + 0.1s_t^2 + 0.5y_t^* + \eta_t$;

In the first specification, $m_0(s_t) = 0.8s_t$ and $m_1(s_t) = 0.5 + 0.8s_t$, both taking a linear form and the marginal effects of the current state on the future state are the same given different choices. In the second specification, we add an interaction term between the state variable and the choice variable, so that the marginal effects of the current state vary across alternatives. Specifically, $m_0(s_t) = 0.8s_t$ and $m_1(s_t) = 0.5 + 1.1s_t$. For DGP's 3 and 4, we assume the transition rule is nonlinear in the current state s_t; while in the latter case, the nonlinearity is more important. For all specifications, we assume $\eta_t \sim N(0, 1)$ and choose $T = 10$; we run simulations for different sample sizes, $N = 100, 1000$, and 10000. The estimation results in this paper are based on 100 Monte Carlo replications.

For illustration of our identification intuition, we first plot the distribution of the state variable at different periods ($t = 1, 3, 5$, and 7) in Figure 2 under DGP 1. It is clear that the state distribution shifts to the right with a smaller variance as time goes by. The variations in the state distribution is mainly driven by the differences in choice probabilities across time periods. Figure 3 further confirms that the mean of the future state distribution conditional on $s_t = 0$ is decreasing over time. This observation suggests that the probability of agents exerting effort becomes lower as they approach the end of the game.

We summarize the estimation results for DGP's 1–4 with $N = 10000$ in Tables 1–4. The estimation results for different sample sizes are provided Tables 8–15 in Appendix A. In these

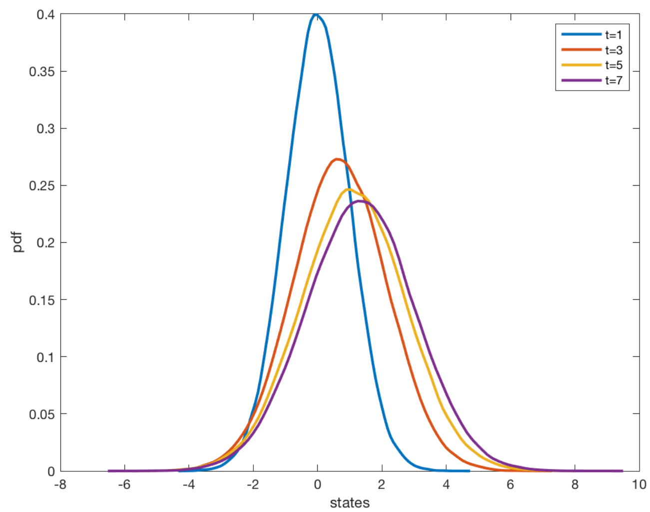

Figure 2: Distribution of the State Variable at Periods 1, 3, 5, and 7 under DGP1

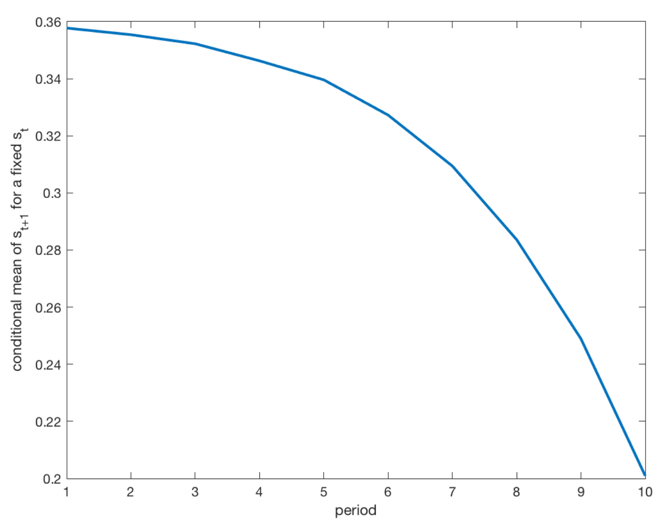

Figure 3: Mean of s_{t+1} Conditional on $s_t = 0$ under DGP 1

exercises, we use third-degree polynomials to approximate the nonparametric functions m_0 and m_1. Specifically,

$$m_0(s) \approx a_0 + a_1 s + a_2 s^2 + a_3 s^3,$$

$$m_1(s) \approx b_0 + b_1 s + b_2 s^2 + b_3 s^3.$$

For the square root of the density function f_{η_t}, we use fifth-degree polynomials. In Tables 1–4, we report Monte Carlo means, biases, standard deviations, mean absolute errors, and the root mean squared errors of the primitives of interest. Instead of showing the estimated coefficients for the η distribution, we report our estimates of μ_η and σ_η, which represent the mean and the standard deviation of the error distribution, respectively. The estimation results for the structural utility parameters are shown in the last two rows of each table. For all data generating processes, our Monte Carlo estimation results generally perform well; adding nonlinear effects of the current state in the transition process leads to slightly more imprecise estimates.

Table 1: Monte Carlo Simulation Results: DGP 1, N=1e4

	TRUE	MC Mean	MC Bias	MC Std	MAE	RMSE
$m_0 : a_0$	0.0000	0.0209	0.0209	0.0035	0.0209	0.0212
$m_0 : a_1$	0.8000	0.7626	-0.0374	0.0195	0.0375	0.0421
$m_0 : a_2$	0.0000	-0.0051	-0.0051	0.0004	0.0051	0.0051
$m_0 : a_3$	0.0000	0.0045	0.0045	0.0008	0.0045	0.0046
$m_1 : b_0$	0.5000	0.4951	-0.0049	0.0139	0.0114	0.0146
$m_1 : b_1$	0.8000	0.7682	-0.0318	0.0161	0.0319	0.0356
$m_1 : b_2$	0.0000	-0.0054	-0.0054	0.0006	0.0054	0.0055
$m_1 : b_3$	0.0000	0.0027	0.0027	0.0012	0.0027	0.0029
μ_η	0.0000	-0.0297	-0.0297	0.0000	0.0297	0.0297
σ_η	1.0000	1.0490	0.0490	0.0039	0.0490	0.0492
ω	0.8000	0.8656	0.0656	0.0135	0.0656	0.0670
ρ	0.3000	0.3301	0.0301	0.0225	0.0338	0.0375

To visualize our simulation results, we plot functions m_0 and m_1 using our estimates and the true parameter values in the data generating process in Figure 4. Our nonparametric estimates of m_0 and m_1 are generally close to the true parameter values. This is particularly the case when there is a linear effect of the current state in the transition process. For nonlinear cases, our estimates still predict the shape of the nonlinear function reasonably well. We also plot the predicted choice probabilities at each period using our estimates and compare those with the choice probabilities calculated using the simulated datasets. The results for the four data generating processes are shown in Figure 5. Except for DGP 3, choice probabilities at each period predicted using our estimates are very close to the ones

Table 2: Monte Carlo Simulation Results: DGP 2, N=1e4

	TRUE	MC Mean	MC Bias	MC Std	MAE	RMSE
$m_0 : a_0$	0.0000	0.0233	0.0233	0.0098	0.0233	0.0253
$m_0 : a_1$	0.8000	0.7744	-0.0256	0.0706	0.0518	0.0748
$m_0 : a_2$	0.0000	-0.0056	-0.0056	0.0026	0.0058	0.0062
$m_0 : a_3$	0.0000	0.0019	0.0019	0.0013	0.0019	0.0023
$m_1 : b_0$	0.5000	0.5515	0.0515	0.0748	0.0720	0.0905
$m_1 : b_1$	1.1000	1.0774	-0.0226	0.0433	0.0362	0.0487
$m_1 : b_2$	0.0000	-0.0066	-0.0066	0.0029	0.0067	0.0072
$m_1 : b_3$	0.0000	0.0010	0.0010	0.0009	0.0011	0.0014
μ_η	0.0000	-0.0297	-0.0297	0.0000	0.0297	0.0297
σ_η	1.0000	1.1096	0.1096	0.0462	0.1096	0.1189
ω	0.8000	0.8409	0.0409	0.0507	0.0542	0.0649
ρ	0.3000	0.3503	0.0503	0.0590	0.0619	0.0773

"observed" in the simulated datasets. These results support our identification and estimation strategies – even if we do not observe agents' choices in the dataset, we can still estimate the choice probabilities reasonably close to the first-step nonparametric estimates if choices were observed. The reason that predicted choice probabilities deviate from the estimates using the data in DGP3 is probably because there are some extremely large values of the state variable generated in the simulated dataset, which makes the nonparametric estimates of the observed state distribution very imprecise.

Table 3: Monte Carlo Simulation Results: DGP 3, N=1e4

	TRUE	MC Mean	MC Bias	MC Std	MAE	RMSE
$m_0 : a_0$	0.0000	0.0237	0.0237	0.0106	0.0239	0.0260
$m_0 : a_1$	0.8000	0.9372	0.1372	0.1165	0.1582	0.1796
$m_0 : a_2$	0.0500	0.0291	-0.0209	0.0145	0.0213	0.0254
$m_0 : a_3$	0.0000	0.0009	0.0009	0.0012	0.0009	0.0015
$m_1 : b_0$	0.5000	0.4837	-0.0163	0.0778	0.0558	0.0791
$m_1 : b_1$	0.8000	0.7797	-0.0203	0.0824	0.0596	0.0845
$m_1 : b_2$	0.0500	0.0279	-0.0221	0.0188	0.0234	0.0289
$m_1 : b_3$	0.0000	0.0025	0.0025	0.0021	0.0026	0.0033
μ_η	0.0000	-0.0297	-0.0297	0.0000	0.0297	0.0297
σ_η	1.0000	1.0701	0.0701	0.0195	0.0701	0.0727
ω	0.8000	0.8803	0.0803	0.0514	0.0844	0.0952
ρ	0.3000	0.3328	0.0328	0.0641	0.0581	0.0717

Table 4: Monte Carlo Simulation Results: DGP 4, N=1e4

	TRUE	MC Mean	MC Bias	MC Std	MAE	RMSE
$m_0 : a_0$	0.0000	0.0211	0.0211	0.0015	0.0211	0.0212
$m_0 : a_1$	0.2000	0.1514	-0.0486	0.0107	0.0486	0.0498
$m_0 : a_2$	0.1000	0.1120	0.0120	0.0051	0.0121	0.0131
$m_0 : a_3$	0.0000	0.0049	0.0049	0.0004	0.0049	0.0049
$m_1 : b_0$	0.5000	0.5133	0.0133	0.0128	0.0157	0.0184
$m_1 : b_1$	0.2000	0.2134	0.0134	0.0094	0.0138	0.0163
$m_1 : b_2$	0.1000	0.0708	-0.0292	0.0051	0.0292	0.0297
$m_1 : b_3$	0.0000	0.0049	0.0049	0.0004	0.0049	0.0050
μ_η	0.0000	-0.0297	-0.0297	0.0000	0.0297	0.0297
σ_η	1.0000	1.0476	0.0476	0.0032	0.0476	0.0477
ω	0.8000	0.8582	0.0582	0.0045	0.0582	0.0584
ρ	0.3000	0.3352	0.0352	0.0060	0.0352	0.0357

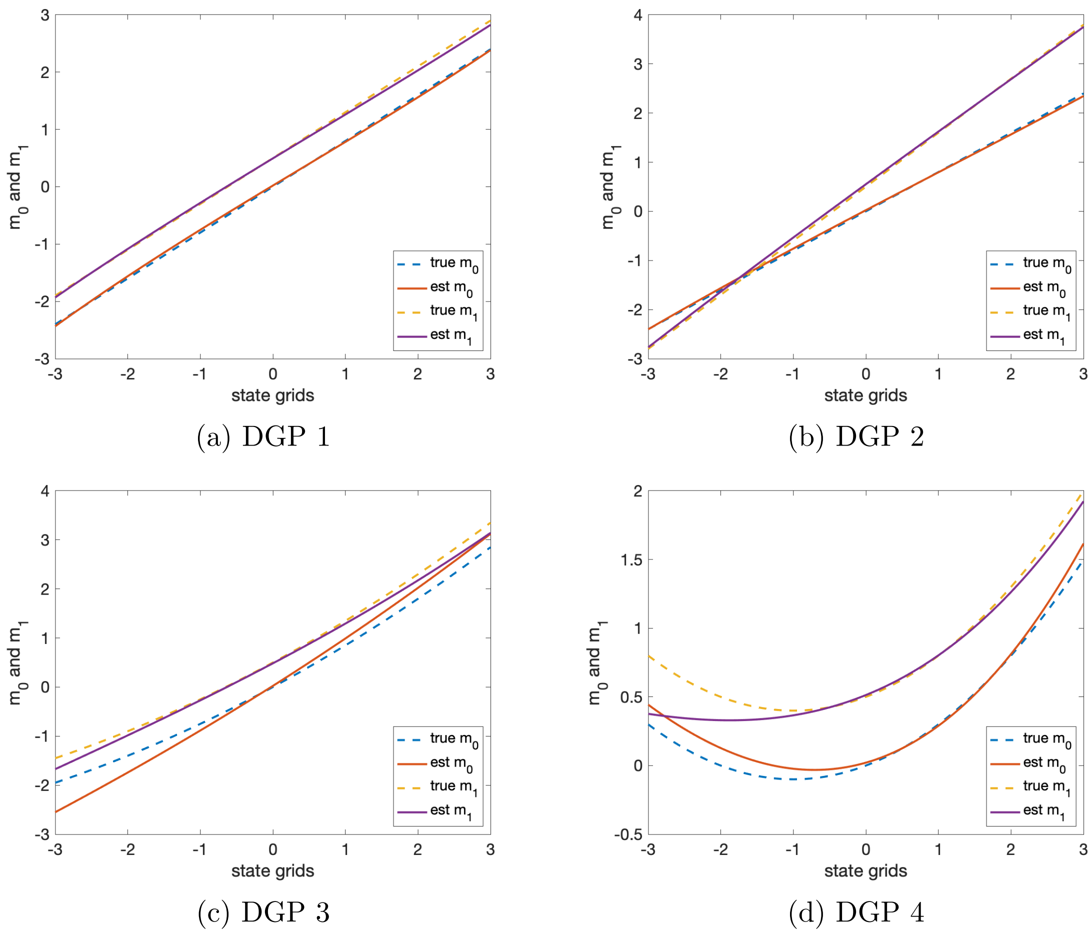

(a) DGP 1 (b) DGP 2

(c) DGP 3 (d) DGP 4

Figure 4: Plot m_0 and m_1 Using Estimates and the True Parameter Values

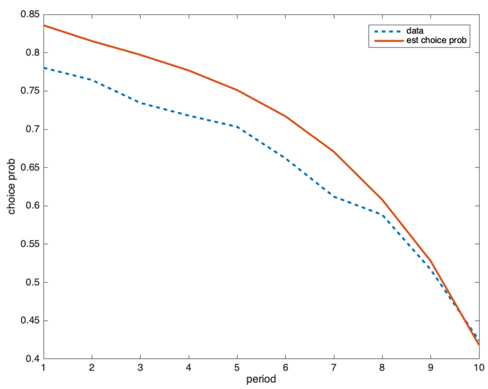

(a) DGP 1

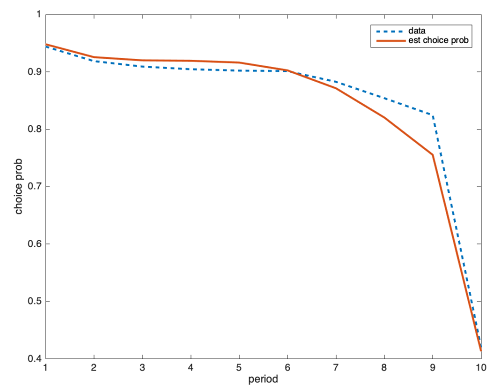

(b) DGP 2

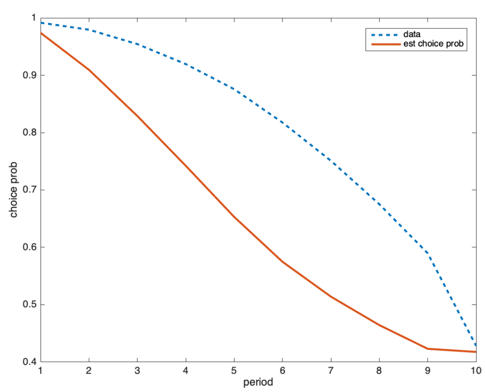

(c) DGP 3

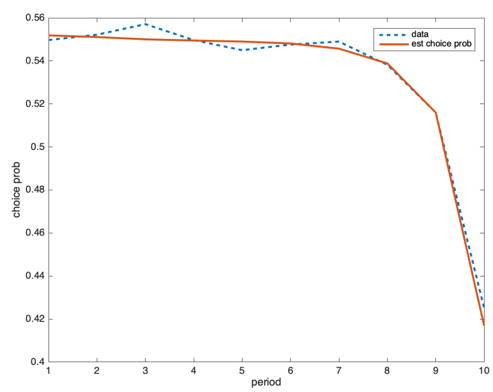

(d) DGP 4

Figure 5: Choice Probabilities: Model Predictions v.s. Data

6 Extensions

We focused on a single-agent finite-horizon dynamic discrete choice model with one continuous state variable to illustrate our main identification and estimation approaches. In the current section, we discuss extensions of our baseline identification results. In particular, we consider four scenarios: (1) when serially correlated unobserved heterogeneity is allowed, (2) when the model has infinite horizon, (3) when only discrete state variables are available in the data, and (4) when multiple players make simultaneous decisions in a game.

6.1 Serially Correlated Unobserved Heterogeneity

In this section, we consider a dynamic discrete choice model with serially correlated unobserved heterogeneity. Following our notations of the baseline model in Section 2, we use s_t to represent the observed state variable and y_t to denote the choice variable. Let (ε_t, x_t^*) represent the vector of unobserved state variables. We now impose assumptions on the dynamic process.

Assumption 5. *The dynamic process of* $\{s_t, \varepsilon_t, x_t^*, y_t\}$ *satisfies the following conditions.*

(i) First-order Markov: $f_{s_{t+1},\varepsilon_{t+1},x_{t+1}^*,y_{t+1}|s_t,x_t^*,\varepsilon_t,y_t,\Omega_{<t}} = f_{s_{t+1},\varepsilon_{t+1},x_{t+1}^*,y_{t+1}|s_t,\varepsilon_t,x_t^*,y_t}$, *where* $\Omega_{<t} \equiv \{s_{t-1}, \cdots, s_0, \varepsilon_{t-1}, \cdots, \varepsilon_0, x_{t-1}^*, \cdots, x_0^*, y_{t-1}, \cdots, y_0\}$.

(ii) The distribution of s_{t+1} *given* $(s_t, \varepsilon_t, x_t^*, y_t)$ *only depends on* (s_t, x_t^*, y_t) *and is denoted by* $f_{s_{t+1}|s_t,x_t^*,y_t}$*; the distribution of* ε_{t+1} *given* $(s_{t+1}, x_{t+1}^*, s_t, \varepsilon_t, x_t^*, y_t)$ *only depends on* (s_{t+1}, x_{t+1}^*) *and is denoted by* $f_{\varepsilon_{t+1}|s_{t+1},x_{t+1}^*}$*; the distribution of* x_{t+1}^* *given* $(s_{t+1}, s_t, \varepsilon_t, x_t^*, y_t)$ *only depends on* (s_{t+1}, x_t^*) *and is denoted by* $f_{x_{t+1}^*|s_{t+1},x_t^*}$.

(iii) State transition probabilities $f_{s_{t+1}|s_t,x_t^*,y_t}$ *are time invariant.*

In general, Assumption 5 is very similar to Assumption 1 invoked for the baseline model. The main difference between the two is that Assumption 5 imposes additional restrictions on the dynamic process related to the unobserved heterogeneity x_t^*. Specifically, Assumption 5(ii) allows that the transition of the observed state s_t depends on the unobserved heterogeneity in the last period; conditional on s_t and x_t^*, ε's are independent over time; and most importantly, the unobserved heterogeneity is serially correlated—the distribution of x_{t+1}^* depends on (s_{t+1}, x_t^*). Assumption 5 still holds if the unobserved heterogeneity is fixed over time, i.e., $x_{t+1}^* = x_t^*$.[8] The serial correlation of the unobserved heterogeneity invoked

[8] Aguirregabiria and Mira (2007), Houde and Imai (2006), and Kasahara and Shimotsu (2009) study the case with time-invariant discrete unobserved heterogeneity.

in Assumption 5(ii) is more general.[9] The dynamic process of the state (observed and unobserved) and choice variables (s_t, x_t^*, y_t) that satisfies Assumption 5 is illustrated in Figure 6. This graph indicates that now in the dynamic discrete choice model, agents' decisions depend on both the observed and unobserved state variables; the transition of the observed state variable also depends on the unobserved heterogeneity. The red dashed lines highlight the serial correlation of the unobserved heterogeneity.

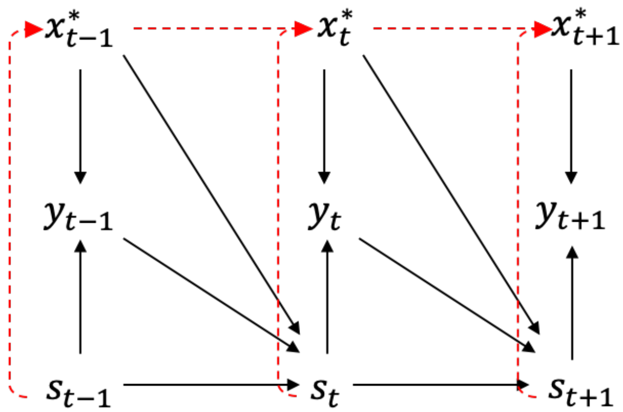

Figure 6: The Dynamic Process of (s_t, x_t^*, y_t)

When both unobserved choices and serially correlated unobserved heterogeneity are present, can we apply similar methodology developed in Section 3 to identify the primitives of interest, i.e., latent choice and state transition probabilities? Under Assumption 5 (i)–(ii), the transition probabilities of the observed state variable can be written as

$$f_{s_{t+1}|s_t,x_t^*}(s'|s,x^*) = \sum_{y_t^*} f_{s_{t+1}|s_t,x_t^*,y_t^*}(s'|s,x^*,y_t^*)p_t(y_t^*|s,x^*), \tag{6.1}$$

where $p_t(y_t^*|s,x^*)$ represents the choice probability of alternative y_t^* given the observed state variable $s_t = s$ and the unobserved heterogeneity $x_t^* = x^*$. Unlike Equation (3.1), both sides of the Equation (6.1) are consist of unobserved terms. On the left-hand side of this equation, the transition probability of the future state given the current state is not directly estimable from the data due to the existence of the unobserved heterogeneity x_t^*. It is clearly to see from Equation (6.1) that in order to apply our identification strategy developed in Section 3, the key is to first recover the transition process of the observed state conditional on the unobserved heterogeneity, i.e., $f_{s_{t+1}|s_t,x_t^*}$.

[9] Hu and Shum (2012) study identification of dynamic models with time-varying and continuous unobserved heterogeneity. Our assumption differs from the one made in their paper in terms of the time restriction. In our case, the unobserved heterogeneity x_t^* realizes after the state variable s_t.

In order to identify $f_{s_{t+1}|s_t,x_t^*}$, we consider the joint distribution of the observed state variable at four periods $(s_{t+2}, s_{t+1}, s_t, s_{t-1})$.

$$
\begin{aligned}
& f_{s_{t+2},s_{t+1},s_t,s_{t-1}} \\
&= \int_{x_{t+1}^*}\int_{x_t^*}\int_{x_{t-1}^*}\int_{y_{t+1}^*}\int_{y_t^*}\int_{y_{t-1}^*} f_{s_{t+2},y_{t+1}^*,x_{t+1}^*,s_{t+1},y_t^*,x_t^*,s_t,y_{t-1}^*,x_{t-1}^*,s_{t-1}} dF_{x_{t+1}^*}\cdots dF_{y_{t-1}^*} \\
&= \int_{x_{t+1}^*}\int_{x_t^*}\int_{x_{t-1}^*}\left(\int_{y_{t+1}^*} f_{s_{t+2}|s_{t+1},y_{t+1}^*,x_{t+1}^*}\times f_{y_{t+1}^*|s_{t+1},x_{t+1}^*} dF_{y_{t+1}^*}\right)\times f_{x_{t+1}^*|s_{t+1},x_t^*} \\
&\qquad \times\left(\int_{y_t^*} f_{s_{t+1}|s_t,y_t^*,x_t^*}\times f_{y_t^*|s_t,x_t^*} dF_{y_t^*}\right)\times f_{x_t^*|s_t,x_{t-1}^*} \\
&\qquad \times\left(\int_{y_{t-1}^*} f_{s_t|s_{t-1},y_{t-1}^*,x_{t-1}^*}\times f_{y_{t-1}^*|s_{t-1},x_{t-1}^*} dF_{y_{t-1}^*}\right)\times f_{x_{t-1}^*,s_{t-1}} dF_{x_{t+1}^*}\cdots dF_{x_{t-1}^*} \\
&= \int_{x_{t+1}^*}\int_{x_t^*}\int_{x_{t-1}^*} f_{s_{t+2}|s_{t+1},x_{t+1}^*}\times f_{x_{t+1}^*|s_{t+1},x_t^*}\times f_{s_{t+1}|s_t,x_t^*}\times f_{x_t^*|s_t,x_{t-1}^*}\times f_{s_t,x_{t-1}^*,s_{t-1}} dF_{x_{t+1}^*}\cdots dF_{x_{t-1}^*} \\
&= \int_{x_t^*}\left(\int_{x_{t+1}^*} f_{s_{t+2}|s_{t+1},x_{t+1}^*}\times f_{x_{t+1}^*|s_{t+1},x_t^*} dF_{x_{t+1}^*}\right)\times f_{s_{t+1}|s_t,x_t^*} \\
&\quad \times\left(\int_{x_{t-1}^*} f_{x_t^*|s_t,x_{t-1}^*}\times f_{s_t,x_{t-1}^*,s_{t-1}} dF_{x_{t-1}^*}\right) dF_{x_t^*} \\
&= \int_{x_t^*} f_{s_{t+2}|s_{t+1},x_t^*}\times f_{s_{t+1}|s_t,x_t^*}\times f_{x_t^*,s_t,s_{t-1}} dF_{x_t^*}
\end{aligned}
\tag{6.2}
$$

The second equality in Equation(6.2) holds under the first-order Markov property of the dynamic process and the conditional independence imposed in Assumption 5(i)–(ii). By integrating out the unobserved choice variables $(y_{t+1}^*, y_t^*, y_{t-1}^*)$, the third equality holds. We further integrate out the unobserved heterogeneity (x_{t+1}^*, x_{t-1}^*), which yields the last line of Equation (6.2). The key insight of the this equation is that the transition of the observed state variable reveals information of the underlying individual heterogeneity, hence can be considered as measurements of the unobserved heterogeneity; conditional on x_t^*, these measurements are independent.[10] Using the spectrum decomposition technique developed by Hu and Schennach (2008), $f_{s_{t+2}|s_{t+1},x_t^*}$, $f_{s_{t+1}|s_t,x_t^*}$, and $f_{x_t^*,s_t,s_{t-1}}$ are nonparametrically identified from the joint distribution of the observed state variable at four periods: $t+2, t+1, t$, and $t-1$.[11]

[10] Hu and Shum (2012) studied nonparametric identification of dynamic models with unobserved state variables. The main difference in this paper is that the choice variable is also unobserved. As a result, in Equation (6.2), we also have to integrate out choices.

[11] The assumptions that guarantee the validity and uniqueness of the spectrum decomposition are discussed in Hu and Shum (2012).

Given that $f_{s_{t+1}|s_t,x_t^*}$ is identified from the joint distribution of $(s_{t+2}, s_{t+1}, s_t, s_{t-1})$, the density function on the left-hand side of Equation (6.1) is identified and can be treated as known. Now in order to apply the identification results in Section 3, we need to find another period τ. Suppose $\tau = t + 1$. Then with the state variable at $t+3, t+2, t+1$, and t, we are able to identify $f_{s_{\tau+1}|s_\tau,x_\tau^*}$. The main takeaway from this is that the unobserved choice and state transition probabilities are identified when serially correlated heterogeneity is present if at least five periods of data are available.

Remark 2. *Identification of models with time-invariant unobserved heterogeneity, such as individual fixed effects, is a special case of our main identification results in Section 6.1 that allow for serially correlated unobserved heterogeneity. In addition, our results allow an individual's decision-making process to depend on his unobserved heterogeneity in a nonlinear way through the optimization process. In fact, it is easy to show that with constant individual unobserved heterogeneity (denoted by x^*), we can identify $f_{s_{t+1}|s_t,x^*}$ and $f_{s_t|s_{t-1},x^*}$ from the joint distribution of (s_{t+1}, s_t, s_{t-1}) using similar techniques as in Equation (6.2). This result indicates that three periods of observed state variables are sufficient to identify the unobserved choice and state transition probabilities conditional on individual fixed effects.*

6.2 Infinite Horizon

We focused on finite-horizon dynamic discrete choice models in the previous discussion. In the current section, we provide conditions under which the unobserved choice and state transition probabilities are identified in an infinite-horizon model.

In a finite-horizon model, the agent's choice probabilities vary over time. As a result, when the latent state transition rule is assumed to be stationary, variations in the moments of the future state distribution conditional on the same previous state can be attributed to the changes in choice probabilities across different periods. In other words, in a finite-horizon model, time serves as an exclusion restriction as it only affects the choice probabilities but not the latent state transition process. However, in an infinite-horizon model, agents' choice probabilities across different periods are the same conditional the same state variable. Consequently, time cannot be used as an excluded variable any more.

In an infinite-horizon model, we need to have an additional variable z_t that satisfies the following assumption serving as an exclusion restriction.

Assumption 6. *z_t enters agents' flow utility, i.e., $u(s_t, z_t, y_t, \varepsilon_t)$, but the transition rule of s_t does not depend on z_t.*

Assumption 6 ensures that the agent's choice probabilities vary with the values of z_t. The condition that the transition rule of s_t does not depend on z_t is an analogy to the stationarity assumption in the baseline model. To see this, for two distinct values of z_t, $\bar{z}$ and $\hat{z}$, we obtain the following two equations under Assumption 6.

$$\begin{aligned} f_{s_{t+1}|s_t,z_t}(s'|s,\bar{z}) &= \sum_{y_t^*} f_{s_{t+1}|s_t,y_t^*}(s'|s,y_t^*)p_t(y_t^*|s,\bar{z}), \\ f_{s_{t+1}|s_t,z_t}(s'|s,\hat{z}) &= \sum_{y_t^*} f_{s_{t+1}|s_t,y_t^*}(s'|s,y_t^*)p_t(y_t^*|s,\hat{z}). \end{aligned} \tag{6.3}$$

From Equation (6.3) it is straightforward to see that the variations in the moments of $f_{s_{t+1}|s_t,z_t}$ given different values of z_t are due to the differences in the choice probabilities. Similar identification arguments can be made as in Section 3, hence the details are omitted.

6.3 Discrete States

We discussed identification results with a continuous state variable in the baseline model. We now focus on a scenario where only discrete state variables are available. When there is only one discrete state variable, comparing future state distributions at two periods provides insufficient variations to identify the unobserved choice probabilities. This is because when the choice variable takes different values, not only the location of the future state distribution shifts, but also the shapes of the distribution changes. In this section, we consider a case where we have two discrete state variables $\{s_t, z_t\}$ that satisfy the following assumption.

Assumption 7 (Conditional Independence)**.** $f_{s_{t+1},z_{t+1}|s_t,z_t,y_t^*} = f_{s_{t+1}|s_t,y_t^*} f_{z_{t+1}|z_t,y_t^*}$.

Assumption 7 implies that the transition process of the two state variables are independent conditional on the choice variable. Specifically, s_t is excluded from the transition of z_t, and vice versa. But the choice probability depends on both state variables. We plot the dynamic process of (s_t, z_t, y_t^*) in Figure 7.

Under Assumption 7, the observed joint distribution of $\{s_{t+1}, z_{t+1}, s_t, z_t\}$ can be decomposed as follows.

$$f_{s_{t+1},z_{t+1},s_t,z_t}(s', z', s, z) = \sum_{y_t^*} f_{s_{t+1}|s_t,y_t^*}(s'|s,y_t^*) f_{z_{t+1}|z_t,y_t^*}(z'|z,y_t^*) f_{y_t^*,s_t,z_t}(y_t^*, s, z). \tag{6.4}$$

Let $j_s = 1, \cdots, J_s$, $j_z = 1, \cdots, J_z$, and $j_y = 1, \cdots, J_y$ index the categories of s_t, z_t and y_t^*, respectively. For simplicity, we consider the case where the cardinalities of s_t, z_t, and y_t are

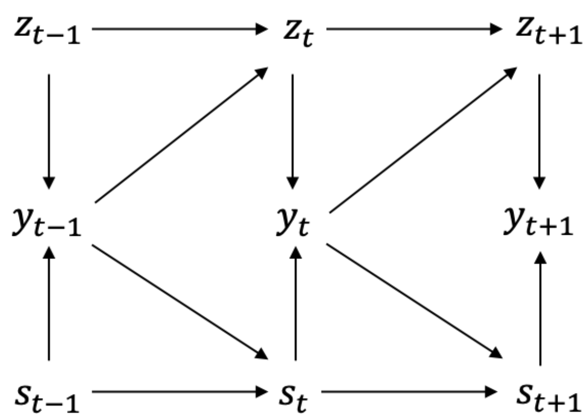

Figure 7: The Dynamic Process of (s_t, z_t, y_t)

equal, i.e., $J_s = J_z = J_y$. We define the following matrices for fixed (s, z):

$$
\begin{aligned}
M_{s_{t+1},z_{t+1},s,z} &= \left[f_{s_{t+1},z_{t+1},s_t,z_t}(s_{t+1}, z_{t+1}, s, z) \mid_{s_{t+1}=j_s, z_{t+1}=j_z} \right]_{j_s,j_z}, \\
M_{s_{t+1}|s,y_t^*} &= \left[f_{s_{t+1}|s_t,y_t^*}(s_{t+1}|s, y_t^*) \mid_{s_{t+1}=j_s, y_t^*=j_y} \right]_{j_s,j_y}, \\
M_{y_t^*,s,z} &= diag \left\{ \left[f_{y_t^*,s_t,z_t}(y_t^*, s, z) \mid_{y_t^*=j_y} \right]_{j_y=1,2,\cdots,J_y} \right\}, \\
M_{z_{t+1}|z,y_t^*} &= \left[f_{z_{t+1}|z_t,y_t^*}(z_{t+1}|z, y_t^*) \mid_{y_t^*=j_y, z_{t+1}=j_z} \right]_{j_y,j_z}.
\end{aligned}
$$

Equation (6.4) in matrix form is therefore

$$
M_{s_{t+1},z_{t+1},s,z} = M_{s_{t+1}|s,y_t^*} M_{y_t^*,s,z} M_{z_{t+1}|z,y_t^*}. \tag{6.5}
$$

We consider four combinations of observed states at t: $(\bar{s}, \bar{z})$, $(\hat{s}, \bar{z})$, $(\bar{s}, \hat{z})$, $(\hat{s}, \hat{z})$, and construct the following equations

$$
\begin{aligned}
M_t^s &= \left(M_{s_{t+1},z_{t+1},\bar{s},\bar{z}} \cdot M^{-1}_{s_{t+1},z_{t+1},\hat{s},\bar{z}} \right) \left(M_{s_{t+1},z_{t+1},\bar{s},\hat{z}} M^{-1}_{s_{t+1},z_{t+1},\hat{s},\hat{z}} \right)^{-1} \\
&= M_{s_{t+1}|\bar{s},y_t^*} \left(M_{y_t^*,\bar{s},\bar{z}} M^{-1}_{y_t^*,\hat{s},\bar{z}} M_{y_t^*,\hat{s},\hat{z}} M^{-1}_{y_t^*,\bar{s},\hat{z}} \right) M^{-1}_{s_{t+1}|\bar{s},y_t^*} \\
&\equiv M_{s_{t+1}|\bar{s},y_t^*} M_{y_t^*,\bar{s},\bar{z},\hat{s},\hat{z}} M^{-1}_{s_{t+1}|\bar{s},y_t^*},
\end{aligned} \tag{6.6}
$$

and

$$
\begin{aligned}
M_t^z &= \left(M^{-1}_{s_{t+1},z_{t+1},\bar{s},\bar{z}} \cdot M_{s_{t+1},z_{t+1},\bar{s},\hat{z}}\right)\left(M^{-1}_{s_{t+1},z_{t+1},\hat{s},\bar{z}} M_{s_{t+1},z_{t+1},\hat{s},\hat{z}}\right)^{-1} \\
&= M^{-1}_{z_{t+1}|\bar{z},y_t^*}\left(M^{-1}_{y_t^*,\bar{s},\bar{z}} M_{y_t^*,\bar{s},\hat{z}} M^{-1}_{y_t^*,\hat{s},\hat{z}} M_{y_t^*,\hat{s},\bar{z}}\right) M_{z_{t+1}|\bar{z},y_t^*} \\
&\equiv M^{-1}_{z_{t+1}|\bar{z},y_t^*} M_{y_t^*,\bar{z},\bar{s},\hat{z},\hat{s}} M_{z_{t+1}|\bar{z},y_t^*},
\end{aligned}
\tag{6.7}
$$

provided that the following assumption holds.

Assumption 8 (Invertibility)**.** *Matrices $M_{s_{t+1}|s,y_t^*}$, $M_{y_t^*,s,z}$, and $M_{z_{t+1}|z,y_t^*}$ are invertible for $(s,z) \in \{(\bar{s},\bar{z}),(\hat{s},\bar{z}),(\bar{s},\hat{z}),(\hat{s},\hat{z})\}$.*

To ensure the invertibility of $M_{s_{t+1}|s,y_t^*}$ and $M_{z_{t+1}|z,y_t^*}$, intuitively, we need the choice variable y_t^* to generate sufficient variations on the future state distributions of s_t and z_t. If for any combinations of (s,z), the choice probabilities of each alternative are nonzero, then the invertiblity of $M_{y_t^*,s,z}$ is guaranteed. With Assumption 8 satisfied, Equations (6.6) and (6.7) lead to eigenvalue-eigenvector decompositions of matrices M_t^s and M_t^z, respectively, although additional assumptions are required to guarantee the uniqueness of the decomposition. We provide one such example for the case where $y_t^* \in \{0,1\}$ and $s_{t+1} \in \{\bar{s},\hat{s}\}$, $\bar{s} < \hat{s}$.

Assumption 9 (Uniqueness)**.** *For any s, $f_{s_{t+1}|s_t,y_t^*}(\bar{s}|s, y_t^*=1) < f_{s_{t+1}|s_t,y_t^*}(\hat{s}|s, y_t^*=1)$ and $f_{s_{t+1}|s_t,y_t^*}(\bar{s}|s, y_t^*=0) > f_{s_{t+1}|s_t,y_t^*}(\hat{s}|s, y_t^*=0)$.*

Assumption 9 imposes restrictions on the state transition process given different choices. The economic intuition of this assumption can be illustrated using the executive's example. Suppose $y_t^* = 1$ represent the case where the executive exerts effort, and 0 otherwise; s_t represents the firm's revenue at period t. Assumption 9 implies that the distribution of the future revenue given exerting effort first order stochastically dominates the one given shirking. For more general assumptions, see the discussions in Hu (2008). Once matrices $M_{s_{t+1}|s,y_t^*}$, $M_{y_t^*,s,z}$, and $M_{z_{t+1}|z,y_t^*}$ are uniquely determined, the identification of all unknown densities in Equation (6.4) is achieved. This result is formally stated in the theorem below.

Theorem 3 (Identification)**.** *Suppose Assumptions 1 and 7–9 hold for the Markov process of $\{s_t, z_t, \varepsilon_t, y_t^*\}$. The joint distribution of $\{s_{t+1}, z_{t+1}, s_t, z_t\}$ identifies the state transition rules $f_{s_{t+1}|s_t,y_t^*}$ and $f_{z_{t+1}|z_t,y_t^*}$, and the choice probabilities $f_{y_t^*|s_t,z_t}$.*

The proof of Theorem 3 is a direct application of Hu (2008) on Equations (6.6) and (6.7), hence is omitted in this paper.

Remark 3. *When there is only one discrete state variable available in the data, we do not get point identification of the unobserved choice probabilities and the latent state transition probabilities. Following An, Hu, and Xiao (2018), we connect the unobserved choice probabilities and the latent state transition probabilities through (1) the observed state transition process, and (2) the agent's dynamic optimization problem. By constructing a sufficient number of nonlinear restrictions, we can locally identify the model primitives.*[12]

Remark 4. *When two continuous state variables are available, it is trivial to generalize our identification results in Section 6.3 to allow for continuous choice variables. Instead of using eigenvalue-eigenvector decompositions, spectrum decompositions proposed by Hu and Schennach (2008) can be easily applied.*

6.4 Dynamic Discrete Games

In the baseline model and extensions discussed in Sections 6.1–6.3, we all focus on single-agent dynamic discrete choice models. In this section, we show that our results can be extended to dynamic discrete games. We first provide a basic framework of dynamic discrete games of incomplete information and then provide identification strategies for conditional choice probabilities and state transition probabilities when players' choices are unobserved by econometricians.

Consider a game with I players, where $i = 1, 2, \cdots, I$ is the index of each individual. Players choose an action from the choice set $\mathcal{Y}$ simultaneously at each period $t = 1, 2, \cdots, \infty$. We use y_{it} to represent player i's action at t, so the action profile is denoted by $\boldsymbol{y}_t = (y_{1t}, y_{2t}, \cdots, y_{It}) \in \mathcal{Y}^I$. We use $s_{it} \in \mathcal{S}_i$ to denote the player's state variable that is publicly observed and $\varepsilon_{it} \in \mathcal{E}_i$ to denote utility shocks that are privately observed by player i (not by i's rivals or econometricians). Let $\boldsymbol{s}_t = (s_{1t}, s_{2t}, \cdots, s_{It}) \in \mathcal{S}$ and $\boldsymbol{\varepsilon}_t = (\varepsilon_{1t}, \varepsilon_{2t}, \cdots, \varepsilon_{It}) \in \mathcal{E}$ be the vector of observed states and private utility shocks at t, respectively. Define $\mathcal{S} = \times_{i=1}^{I} \mathcal{S}_i$ and $\mathcal{E} = \times_{i=1}^{I} \mathcal{E}_i$.

Unlike the single-agent case, a player's utility now depends on the action profile and state variables of all players and his own private information ε_{it}. We use $u(\boldsymbol{s}_t, \varepsilon_{it}, \boldsymbol{y}_t)$ to represent the player's per period flow utility. At each period t, all players choose their actions simultaneously to maximize their own expected sum of the discounted utility, i.e., $\mathrm{E}[\sum_{j=0}^{T-t} \beta^j u(\boldsymbol{s}_{t+j}, \varepsilon_{i,t+j}, \boldsymbol{y}_{t+j})]$, where the expectation is taken over other players' current and future actions, the future observed states, and i's private shocks in the future. We invoke

[12] We conduct Monte Carlo simulations for the scenario where only one discrete state variable is available. The basic setup remains the same as in Section 5. However, the state variable now only takes binary values: $s_t \in \{1, 2\}$. We present the simulation results in Table 16 in Appendix A.

the following assumption to restrict attention to certain classes of models.

Assumption 10. *The dynamic process of* $\{\boldsymbol{s}_t, \boldsymbol{\varepsilon}_t, \boldsymbol{y}_t\}$ *satisfies the following conditions.*

(i) First-order Markov: $f_{\boldsymbol{s}_{t+1},\boldsymbol{\varepsilon}_{t+1},\boldsymbol{y}_{t+1}|\boldsymbol{s}_t,\boldsymbol{\varepsilon}_t,\boldsymbol{y}_t,\boldsymbol{\Omega}_{<t}} = f_{\boldsymbol{s}_{t+1},\boldsymbol{\varepsilon}_{t+1},\boldsymbol{y}_{t+1}|\boldsymbol{s}_t,\boldsymbol{\varepsilon}_t,\boldsymbol{y}_t}$, *where* $\boldsymbol{\Omega}_{<t} \equiv \{\boldsymbol{s}_{t-1}, \cdots, \boldsymbol{s}_0, \boldsymbol{\varepsilon}_{t-1}, \cdots, \boldsymbol{\varepsilon}_0, \boldsymbol{y}_{t-1}, \cdots, \boldsymbol{y}_0\}$.

(ii) ε_{it} *are independently distributed over time and across players, and are drawn from a distribution* $F_i(\cdot|\boldsymbol{s}_t)$.

(iii) The distribution of $\boldsymbol{s}_{t+1}$ *given* $(\boldsymbol{s}_t, \boldsymbol{\varepsilon}_t, \boldsymbol{y}_t)$ *only depends on* $(\boldsymbol{s}_t, \boldsymbol{y}_t)$ *and is denoted by* $f_{\boldsymbol{s}_{t+1}|\boldsymbol{s}_t,\boldsymbol{y}_t}$.

Though typically invoked in the literature of dynamic discrete games of incomplete information, Assumption 10 imposes several restrictions on the model. First, it assumes that the distribution of observed state variables, utility shocks, and choices only depends on their values in the last period (i.e., they follow a first-order Markov process). Second, a conditional independence assumption that is very similar to the one imposed for single-agent models is invoked for private utility shocks. Assumption 10(ii) rules out the possibility that private shocks are serially correlated over time; in a game setting, allowing serial correlation could lead to complicated theoretical issues, including learning or strategic signaling behavior among players. Last, Assumption 10(iii) requires that the transition process of observed state variables does not depend on the private utility shocks in the previous periods

In the game described above, we consider pure strategy Markov Perfect Equilibrium (MPE) as our equilibrium concept, in which case players' actions only depend on the value of current states and utility shocks. In addition, we focus on stationary Markov strategies, so subscript t is dropped in the following definitions. We define a Markov strategy for player i as $a_i(\boldsymbol{s}_t, \varepsilon_{it})$ and i's belief that $\boldsymbol{y}_t$ is chosen at state $\boldsymbol{s}_t$ as $\sigma_i(\boldsymbol{y}_t|\boldsymbol{s}_t)$. Under Assumption 10(ii), the value function for player i given belief σ_i is

$$V_i(\boldsymbol{s}_t, \varepsilon_{it}; \sigma_i) = \max_{y \in \mathcal{Y}} \sum_{\boldsymbol{y}_{-i} \in \mathcal{Y}^{I-1}} \sigma_i(\boldsymbol{y}_{-i}|\boldsymbol{s}_t) \Big[u(\boldsymbol{s}_t, \varepsilon_{it}, (y, \boldsymbol{y}_{-i})) + \beta \, \mathrm{E}[V_i(\boldsymbol{s}_{t+1}, \varepsilon_{i,t+1}; \sigma_i)|\boldsymbol{s}_t, (y, \boldsymbol{y}_{-i})] \Big], \tag{6.8}$$

where $\boldsymbol{y}_{-i}$ represent the profile of actions for all other players except i. The optimal strategy of player i given state variable $\boldsymbol{s}_t$ and private utility shock ε_{it} under belief σ_i is therefore

$$a_i(\boldsymbol{s}_t, \varepsilon_{it}; \sigma_i) = \arg\max_{y \in \mathcal{Y}} V_i(\boldsymbol{s}_t, \varepsilon_{it}; \sigma_i). \tag{6.9}$$

After integrating out the player's private information, we can define i's choice probabilities given state variable $\boldsymbol{s}_t$ and belief σ_i as

$$p_i(y_{it}|\boldsymbol{s}_t;\sigma_i) = \int \mathbf{1}\{y_{it} = a_i(\boldsymbol{s}_t, \varepsilon_{it};\sigma_i)\} dF_i(\varepsilon_{it}|\boldsymbol{s}_t). \tag{6.10}$$

In a MPE, players' beliefs are consistent with their strategies, leading to a fixed point of a mapping in the space of conditional choice probabilities. Under certain regularity conditions, at least one Markov perfect equilibrium exists for dynamic discrete games of incomplete information, but multiplicity of equilibria may be possible.[13] In this paper, our goal is to analyze situations when players' actions are unobserved to econometricians, so we focus on the simplest case where the same equilibrium is played in the data.[14,15]

We define player i's equilibrium choice probabilities conditional on $\boldsymbol{s}_t$ as $p_i^*(y_{it}|\boldsymbol{s}_t)$. When agents' actions are observed to econometricians, following the two-step methods originally developed by Hotz and Miller (1993), we can directly estimate the conditional choice probabilities $p_i^*(y_{it}|\boldsymbol{s}_t)$ and state transition rules $f_{\boldsymbol{s}_{t+1}|\boldsymbol{s}_t,\boldsymbol{y}_t}$ from the data in the first step. Then different approaches have been developed in the literature to recover structural parameters of the game (see Jofre-Bonet and Pesendorfer, 2003; Aguirregabiria and Mira, 2007; Bajari, Benkard, and Levin, 2007; Pakes, Ostrovsky, and Berry, 2007; Pesendorfer and Schmidt-Dengler, 2008). However, when the actions are unobserved to researchers, the existing methods no longer work; in this paper, we invoke the following assumption to achieve the identification of structural parameters in dynamic discrete games with unobserved actions.

Assumption 11. *Conditional on the current values of player's own actions and states, the future states are independent across players, i.e.,*

$$f_{\boldsymbol{s}_{t+1}|\boldsymbol{s}_t,\boldsymbol{y}_t}(\boldsymbol{s}'|\boldsymbol{s},\boldsymbol{y}) = \prod_{i=1}^{I} f_{s_{i,t+1}|s_{it},y_{it}}(s_i'|s_i, y_i).$$

In general, Assumption 11 eliminates the "cross-effects": the transition process of the observed state variable only depends on player i's action and state in the last period, not on other players' actions or states. This assumption is motivated by the empirical setting

[13] Doraszelski and Satterthwaite (2010) provided conditions under which equilibrium exists. See the discussions in Bajari, Benkard, and Levin (2007) and Aguirregabiria and Mira (2010) for more details about multiple equilibria.

[14] Otsu, Pesendorfer, and Takahashi (2016) provide several statistical tests to examine whether the same (or unique) equilibrium is played when data from distinctive markets are pooled. Their method also requires the observation of players' choices to estimate CCPs and state transition probabilities in the first step.

[15] Luo, Xiao, and Xiao (2018) provides nonparametric identification results for dynamic discrete games of incomplete information when multiple equilibria and unobserved heterogeneity are present.

of dynamic oligopoly competition, where the state variable is the firm's capacity levels and the choice is the firm's incremental changes to capacity. In this case, it is natural to assume that the transition of the states only depends on the firm's own decisions, not on the other player's choices.[16] Under Assumption 11, we achieve the following equation for i's state transition process:

$$f_{s_{i,t+1}|\boldsymbol{s}_t}(s_i'|\boldsymbol{s}) = \sum_{y_{it}^* \in \mathcal{Y}} f_{s_{i,t+1}|s_{it},y_{it}^*}(s_i'|s_i, y_{it}^*) p_i^*(y_{it}^*|\boldsymbol{s}), \tag{6.11}$$

where y_{it}^* is used to represent player i's unobserved choice at period t. It is highlighted in Equation (6.11) that the transition process of $s_{i,t+1}$ does not depend on $\boldsymbol{s}_{-i,t}$; while in a game setting, all players interact with each other, so i's choices naturally depend on all other players' state variables. In dynamic games, $\boldsymbol{s}_{-i}$ can be used as an exclusion restriction. For two values of $\boldsymbol{s}_{-i}$, $\bar{\boldsymbol{s}}_{-i}$ and $\hat{\boldsymbol{s}}_{-i}$, we obtain the following two equations under Assumption 11.

$$\begin{aligned} f_{s_{i,t+1}|\boldsymbol{s}_t}(s_i'|s_i, \bar{\boldsymbol{s}}_{-i}) &= \sum_{y_{it}^* \in \mathcal{Y}} f_{s_{i,t+1}|s_{it},y_{it}^*}(s_i'|s_i, y_{it}^*) p_i^*(y_{it}^*|s_i, \bar{\boldsymbol{s}}_{-i}), \\ f_{s_{i,t+1}|\boldsymbol{s}_t}(s_i'|s_i, \hat{\boldsymbol{s}}_{-i}) &= \sum_{y_{it}^* \in \mathcal{Y}} f_{s_{i,t+1}|s_{it},y_{it}^*}(s_i'|s_i, y_{it}^*) p_i^*(y_{it}^*|s_i, \hat{\boldsymbol{s}}_{-i}), \end{aligned} \tag{6.12}$$

From Equation (6.12), it is clear to see that the variations in the moments of player i's state distribution conditional on other players' last-period states (i.e., $f_{s_{i,t+1}|\boldsymbol{s}_t}$) are due to the differences in the choice probabilities. Similar identification strategies as shown in Section 3 can be applied to identify the state transition probabilities and equilibrium choice probabilities for players $i = 1, 2, \cdots, I$. We therefore omit the details here.

Remark 5. *Our identification results do not require all state variables to satisfy Assumption 11. Depending on applications, we may have multiple dimensions of the state variable; as long as there exists one state variable whose transition process does not involve other players' actions or states, the equilibrium choice probabilities are identified. To identify the state transition process, we may relax Assumption 11 by allowing the transition of $s_{i,t+1}$ to depend on a subset of state variables from other players. When equilibrium choice probabilities are known, we can identify the state transition probabilities under certain rank conditions.*

[16] Ryan (2012) estimates a dynamic model of oligopoly to study the cost of environmental regulations on firms' entry, exit, and investment decisions. In this paper, it is assumed that the transition of the states (capacity) depend on firms' own current state variables and actions (i.e., entry, exit, or investment). In addition, the author assumes that the transition process is deterministic to reduce computational burden.

7 Empirical Application: Moral Hazard in US Gubernatorial Elections

In this section, we apply our methods to study moral hazard problems in US gubernatorial elections. As pointed out by the theoretical literature, the accountability of politicians are usually not observed by the voters—some incentives are necessary to motivate politicians to exert more effort. There is a strand of empirical literature in political economy focused on understanding the impact of institutional design of election rules (e.g., term limits) on politician's behavior, election outcomes, and voter's welfare.[17] A seminal paper by Besley and Case (1995) studies the effect of term limits on US governor's policy choices from 1950-1986; Alt, Bueno de Mesquita, and Rose (2011) extend the dataset and explore variations in gubernatorial term limits across states to separately identify the accountability and competence effects of elections. In two recent structural papers, Aruoba, Drazen, and Vlaicu (2019) develop and estimate a political agency model with asymmetric information between politicians and voters and they find significant incentive effects of reelections; Sieg and Yoon (2017) focus more on the adverse selection problem, treating the ideology of the politician as a source of unobserved heterogeneity instead of an effort-exerting decision.

In all of the papers mentioned above, governors are assumed to make one decision (exerting effort or shirking) for each term, which ignores the dynamics within a term and rules out the possibility of political business cycle (see Drazen (2000) for a comprehensive survey on this literature). The main goal of the empirical application in this paper is to estimate a dynamic structural model of politicians' within-term effort-exerting decisions to better understand the moral hazard problems in gubernatorial elections.

The dataset used for our empirical application comes from Alt, Bueno de Mesquita, and Rose (2011). This dataset contains all gubernatorial elections between 1950 and 2000 in the United States. During that period, different states may have adopted different term limits and the rules could also change over time.[18] We select governors serving their last terms for states that have four-year terms. The governors we select are essentially "lame ducks" who were not eligible for reelections.[19] For states that have adopted a limit of two consecutive terms, we only consider governors who were serving their second terms. In total, there are 142 governors in our sample. The summary statistics of whether the governor is a first-term lame duck, proportions of elderly people in the state, and whether the governor is a democratic politician are provided in the upper panel of Table 5. In our sample, about

[17] See Alt, Bueno de Mesquita, and Rose (2011) for a literature review.

[18] Detailed information about gubernatorial term limits can be found in the Book of the States.

[19] In Alt, Bueno de Mesquita, and Rose (2011), "lame ducks" refer to politicians who cannot run for reelection.

54% of the governors were serving their first terms and because of the term limits they were not eligible for reelections. The average proportion of elderly people is around 10% and 71% of the people we have are democratic governors.

Table 5: Summary Statistics

Variable	Mean	Std. Dev.	Min	Max	Obs
	observed characteristics				
first-term lame duck	0.5423	0.5000	0	1	142
proportions of elderly	0.1039	0.0235	0.0618	0.1848	142
democratic governor	0.7183	0.4514	0	1	142
	log of per capita spending				
year 0	6.6491	0.5675	5.4375	7.8136	142
year 1	6.6970	0.5599	5.4880	7.8375	142
year 2	6.7331	0.5399	5.5383	7.8445	142
year 3	6.7673	0.5276	5.5669	7.9448	142
year 4	6.8115	0.5108	5.6396	7.9362	142

In this application, we use log of per capita spending (reported in constant 1982 dollars) as the state variable. Let t be the index of years within a term. $t = 1$ refers to the year when a governor was elected (or reelected); $t = 0$ refers to the year before the term began. The summary statistics of the state variable for $t = 0, 1, \cdots, 4$ are provided in the lower panel of Table 5. We impose Assumption 2 on the transition process of the state variable, that is $s_{t+1} = m(s_t, y_t^*) + \eta_t$, where η_t is independent with the choice variable y_t^*. Although our identification results allow that the distribution of η_t depends on s_t, for this application we focus on the case in which η_t is also independent with s_t due to the small sample size. We assume the per period utility of a governor at t given the current state s_t and choice y_t^* has the following linear structure:

$$u(s_t, y_t^*) = \omega s_t - \rho y_t^*. \tag{7.1}$$

Let $y_t^* = 1$ if the governor exerts effort, and 0 otherwise. In Equation (7.1), ρ represents the marginal cost of exerting effort. In our estimation, we allow ρ to depend on individual observed characteristics, such as whether the governor is a first-term lame duck, proportions of elderly people in the state, and whether the governor is a democratic politician. Specifically, the following parametric form is considered in the estimation.

$$\rho = \rho_0 + \rho_1 \textit{First-Term} + \rho_2 \textit{Elderly-Prop} + \rho_3 \textit{Democratic}.$$

In addition to the deterministic part, the governor also receives a random utility shock ε_t, which is choice specific. Assume $(\varepsilon_t(0), \varepsilon_t(1))$ are drawn independently from the type

Table 6: Estimation Results

Panel (A) Estimates of m_0 and m_1

Parameters	Estimates	Std. Err.
$m_0 : a_0$	-3.4571	0.0794
$m_0 : a_1$	1.8077	0.0014
$m_0 : a_2$	0.0072	0.0053
$m_0 : a_3$	-0.0072	0.0006
$m_1 : b_0$	0.9173	0.2692
$m_1 : b_1$	0.8105	0.0041
$m_1 : b_2$	0.0077	0.0186
$m_1 : b_3$	0.0003	0.0019

Panel (B) Estimates of Utility Primitives

Parameters	Estimates	Std. Err.
ω	16.0678	7.2725
ρ_0	0.0000	4.0890
ρ_1	3.4327	1.1418
ρ_2	0.8428	31.9709
ρ_3	0.1832	0.7229
μ_η	-0.0453	0.1300
σ_η	0.0581	0.0008

I extreme value distribution. In summary, the parameters to be estimated in this model includes $\{\omega, \rho_0, \rho_1, \rho_2, \rho_3, m_0(\cdot), m_1(\cdot), f_\eta(\cdot)\}$, where the last three are unknown functions.

We estimate the model primitives following our sieve maximum likelihood estimation strategy developed in Section 4. The point estimates and the their standard errors are provided in Table 6.[20] From the estimation results of m_0 and m_1 in Panel (A) we can see that if governors exert effort, the distribution of the future state is on average better. The marginal utility governors get from the state variable is significantly positive; but exerting effort is costly. We find that the marginal cost of exerting effort for the first-term lame ducks are significantly higher compared to the second-term lame ducks. This suggests that governors who were reelected are potentially more competent, which is consistent with the selection effect of elections. We compute the probabilities of shirking for governors at each period using the estimated parameters. The results for the full sample and by each observed category are shown in Table 7. From this table we can see that the probabilities of shirking are increasing over time within a term. The probability of exerting no effort in the last period is 31% higher than that of the first period. This result is quite intuitive: governors have less incentives of exerting effort when they are approaching the end of the term. Overall, we observe a lower chance of exerting effort for first-term governors. The differences between democratic and republican politicians are not significant; having different proportions of elderly people also seems to have no significant impact on governors' shirking probabilities.

[20] Similar to our notations in Monte Carlo simulations in Section 5, parameters a_j and b_j for $j = 0, 1, 2, 3$ are coefficients in polynomials that approximate $m_0(\cdot)$ and $m_1(\cdot)$, respectively. Specifically, $m_0(s) \approx a_0 + a_1 s + a_2 s^2 + a_3 s^3$, and $m_1(s) \approx b_0 + b_1 s + b_2 s^2 + b_3 s^3$.

Table 7: Probabilities of Shirking at Each Period

	year 1	year 2	year 3	year 4
all sample	0.6031	0.6561	0.7130	0.7816
	By Category			
first-term lame duck	0.8224	0.9100	0.9624	0.9748
second-term lame duck	0.3432	0.3555	0.4176	0.5528
democratic governor	0.6313	0.6930	0.7542	0.8167
republican governor	0.5311	0.5623	0.6081	0.6921
lower percent of elderly	0.6933	0.7728	0.8343	0.8667
higher percent of elderly	0.4644	0.4771	0.5268	0.6509

8 Conclusion

In this paper, we provide new identification and estimation methods for dynamic structural models when agents' choices are unobserved by econometricians. We leverage on the variations in observed state transition process across different periods. In finite-horizon models, time serves as an exclusion restriction because it only affects the choice probabilities but not the state transition rules. We consider several extensions to our baseline model. First, we incorporate individual serially correlated heterogeneity into the dynamic discrete choice model. Second, we discuss the conditions under which infinite-horizon models with unobserved choices are also identified. Third, we consider the cases in which only discrete state variables are available in the data. Last, we extend the results to dynamic discrete games. We propose sieve maximum likelihood estimation strategy for nonparametric functions in the state transition rules and utility primitives. Monte Carlo simulations under various specifications confirm the validity of our proposed approaches.

We apply our method to study moral hazard problems in US gubernatorial elections. Our estimation results suggest that the probabilities of shirking for governors are generally increasing over time within a four-year term. The probability of exerting no effort in the last period is around 31% higher than that of the first period. These findings add new evidence to the empirical literature focused on understanding the impact of term limits on politician's behavior.

References

Abbring, Jaap H. 2010. "Identification of Dynamic Discrete Choice Models." *Annual Review of Economics,* 2(1): 367–394.

Ai, C., and Chen, X. 1999. "Efficient Sieve Minimum Distance Estimation of Semiparametric Conditional Moment Models." *Manuscript,* London School of Economics.

An, Yonghong, Yingyao Hu, and Ruli Xiao. 2018. "Dynamic Decisions under Subjective Expectations: A Structural Analysis." *Working Paper.*

Aguirregabiria, Victor. 1999. "The Dynamics of Markups and Inventories in Retailing Firms." *The Review of Economic Studies,* 66(2): 275–308.

Aguirregabiria, Victor, and Pedro Mira. 2002. "Swapping the Nested Fixed Point Algorithm: A Class of Estimators for Discrete Markov Decision Models." *Econometrica,* 70(4): 1519–1543.

Aguirregabiria, Victor, and Pedro Mira. 2007. "Sequential Estimation of Dynamic Discrete Games." *Econometrica,* 75(1): 1–53.

Aguirregabiria, Victor. 2010. "Another Look at the Identification of Dynamic Discrete Decision Processes: An Application to Retirement Behavior." *Journal of Business & Economic Statistics,* 28(2): 201–218.

Aguirregabiria, Victor, and Pedro Mira. 2010. "Dynamic Discrete Choice Structural Models: A Survey." *Journal of Econometrics,* 156(1): 38–67.

Aguirregabiria, Victor, Pedro Mira, and Hernan Roman. 2007. "An Estimable Dynamic Model of Entry, Exit and Growth in Oligopoly Retail Markets." *American Economic Review,* 97(2): 449–454.

Alt, James, Ethan Bueno de Mesquita, and Shanna Rose. 2011. "Disentangling Accountability and Competence in Elections: Evidence from US Term Limits." *The Journal of Politics,* 73(1): 171–186.

Aruoba, S. Boragan, Allan Drazen, and Razvan Vlaicu. 2019. "A Structural Model of Electoral Accountability." *International Economic Review,* Forthcoming.

Bajari, Patrick, C. Lanier Benkard, and Jonathan Levin. 2007. "Estimating Dynamic Models of Imperfect Competition." *Econometrica,* 75(5): 1331–1370.

Besley, Timothy, and Anne Case. 1995. "Does Electoral Accountability Affect Economic Policy Choices? Evidence from Gubernatorial Term Limits." *The Quarterly Journal of Economics,* 110(3): 769–798.

Carroll, Raymond J., Xiaohong Chen, and Yingyao Hu. 2010. "Identification and Estimation of Nonlinear Models Using Two Samples with Nonclassical Measurement Errors." *Journal of Nonparametric Statistics,* 22(4): 379–399.

Chen, Xiaohong. 2007. "Large Sample Sieve Estimation of Semi-nonparametric Models," *Handbook of Econometrics,* Vol. 6B, eds. J.J. Heckman and E. Leamer, Elsevier Science.

Chen, Xiaohong, Yingyao Hu, and Arthur Lewbel. 2008. "A Note on the Closed-Form Identification of Regression Models with a Mismeasured Binary Regressor." *Statistics & Probability Letters,* 78(12): 1473–1479.

Chen, X., Shen, X. 1996. "Asymptotic Properties of Sieve Extremum Estimates for Weakly Dependent Data with Applications." *Manuscript,* University of Chicago.

Chen, X., Shen, X. 1998. "Sieve Extremum Estimates for Weakly Dependent Data." *Econometrica,* 66: 289–314.

Collard-Wexler, Allan. 2013. "Demand Fluctuations in the Ready-Mix Concrete Industry." *Econometrica,* 81(3): 1003–1037.

Copeland, Adam, and Cyril Monnet. 2009. "The Welfare Effects of Incentive Schemes." *The Review of Economic Studies,* 76(1): 93–113.

Doraszelski, Ulrich, and Mark. Satterthwaite. 2010. "Computable Markov-Perfect Industry Dynamics." *RAND Journal of Economics,* 41(2): 215–243.

Drazen, Allan. 2000. "The Political Business Cycle after 25 Years." *NBER Macroeconomics Annual,* 15: 75–117.

Gayle, George-Levi, and Robert A. Miller. 2015. "Identifying and Testing Models of Managerial Compensation." *The Review of Economic Studies,* 82(3): 1074-1118.

Guerre, Emmanuel, Isabelle Perrigne, and Quang Vuong. 2000. "Optimal Nonparametric Estimation of First-Price Auctions." *Econometrica,* 68(3): 525–574.

Hotz, V. Joseph, and Robert A. Miller. 1993. "Conditional Choice Probabilities and the Estimation of Dynamic Models." *The Review of Economic Studies,* 60(3): 497–529.

Hotz, V. Joseph, Robert A. Miller, Seth Sanders, and Jeffrey Smith. 1994. "A Simulation Estimator for Dynamic Models of Discrete Choice." *The Review of Economic Studies,* 61(2): 265–289.

Houde, J. F., and S. Imai. 2006. "Identification and 2-step Estimation of DDC models with Unobserved Heterogeneity." *Working Paper,* Queen's University.

Hu, Yingyao. 2008. "Identification and Estimation of Nonlinear Models with Misclassification Error using Instrumental Variables: A General Solution." *Journal of Econometrics,* 144(1): 27–61.

Hu, Yingyao, and Susanne M. Schennach. 2008. "Instrumental Variable Treatment of Nonclassical Measurement Error Models." *Econometrica,* 76(1): 195–216.

Hu, Yingyao, and Matthew Shum. 2012. "Nonparametric Identification of Dynamic Models with Unobserved State Variables." *Journal of Econometrics,* 171(1): 32–44.

Jofre-Bonet, Mireia and Martin Pesendorfer. 2003. "Estimation of a Dynamic Auction Game." *Econometrica,* 71 (5): 1443–1489.

Kasahara, Hiroyuki, and Katsumi Shimotsu. 2009. "Nonparametric Identification of Finite Mixture Models of Dynamic Discrete Choices." *Econometrica,* 77(1): 135–175.

Keane, Michael P., and Kenneth I. Wolpin. 1997. "The Career Decisions of Young Men." *Journal of Political Economy,* 105(3): 473–522.

Luo, Yao, Ping Xiao, and Ruli Xiao. 2018. "Identification of Dynamic Games With Unobserved Heterogeneity and Multiple Equilibria: Global Fast Food Chains in China." *Working Paper.*

Magnac, Thierry, and David Thesmar. 2002. "Identifying Dynamic Discrete Decision Processes." *Econometrica,* 70(2): 801–816.

Misra, Sanjog, and Harikesh S. Nair. 2011. "A Structural Model of Sales-Force Compensation Dynamics: Estimation and Field Implementation." *Quantitative Marketing and Economics,* 9(3): 211-257.

Norets, Andri, and Xun Tang. 2014. "Semiparametric Inference in Dynamic Binary Choice Models." *Review of Economic Studies,* 81(3): 1229–1262.

Otsu, Taisuke, Martin Pesendorfer, and Yuya Takahashi. 2016. "Pooling Data Ccross Markets in Dynamic Markov Games." *Quantitative Economics,* 7(2):523–559.

Perrigne, Isabelle, and Quang Vuong. 2011. "Nonparametric Identification of a Contract Model with Adverse Selection and Moral Hazard." *Econometrica,* 79(5): 1499–1539.

Rust, John. 1994. "Structural Estimation of Markov Decision Processes." *Handbook of Econometrics*, Vol. 4, ed. by R. Engle and D. McFadden. Amsterdam: North Holland, pp. 3081–3143.

Rust, John. 1987. "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher." *Econometrica,* 55: 999–1033.

Rust, John, and Christopher Phelan. 1997. "How Social Security and Medicare Affect Retirement Behavior in a World of Incomplete Markets." *Econometrica,* 65: 781–831.

Ryan, Stephen. 2012. "The Costs of Environmental Regulation in a Concentrated Industry." *Econometrica,* 80(3): 1019–1061.

Shen, X. 1997. "On Methods of Sieves and Penalization." *The Annals of Statistics,* 25: 2555–2591.

Shen, X., and Wong, W. 1994. "Convergence Rate of Sieve Estimates." *Annals of Statistics,* 22: 580—615.

Sieg, Holger, and Chamna Yoon. 2017. "Estimating Dynamic Games of Electoral Competition to Evaluate Term Limits in US Gubernatorial Elections." *American Economic Review,* 107(7): 1824–57.

Takahashi, Yuya. 2015. "Estimating a War of Attrition: The Case of the US Movie Theater Industry." *American Economic Review,* 105(7): 2204–41.

Timmins, Christopher. 2002. "Measuring the Dynamic Efficiency Costs of Regulators' Preferences: Municipal Water Utilities in the Arid West." *Econometrica,* 70(2): 603–629.

Pakes, Ariel, Michael Ostrovsky, and Steve Berry. 2007. "Simple Estimators for the Parameters of Discrete Dynamic Games (with Entry/Exit Examples)." *RAND Journal of Economics,* 38(2): 373–399.

Pesendorfer, Martin, and Philipp Schmidt-Dengler. 2008. "Asymptotic Least Squares Estimators for Dynamic Games." *Review of Economic Studies,* 75(3): 901–928.

Xin, Yi. 2019. "Asymmetric Information, Reputation, and Welfare in Online Credit Markets." *Working Paper.*

A Additional Tables

Table 8: Monte Carlo Simulation Results: DGP 1, N=1000

	TRUE	MC Mean	MC Bias	MC Std	MAE	RMSE
$m_0 : a_0$	0.0000	0.0203	0.0203	0.0005	0.0203	0.0203
$m_0 : a_1$	0.8000	0.7941	-0.0059	0.0149	0.0109	0.0159
$m_0 : a_2$	0.0000	-0.0051	-0.0051	0.0002	0.0051	0.0051
$m_0 : a_3$	0.0000	0.0050	0.0050	0.0001	0.0050	0.0050
$m_1 : b_0$	0.5000	0.4626	-0.0374	0.0172	0.0383	0.0411
$m_1 : b_1$	0.8000	0.7489	-0.0511	0.0083	0.0511	0.0518
$m_1 : b_2$	0.0000	-0.0052	-0.0052	0.0001	0.0052	0.0052
$m_1 : b_3$	0.0000	0.0046	0.0046	0.0005	0.0046	0.0047
μ_η	0.0000	-0.0297	-0.0297	0.0000	0.0297	0.0297
σ_η	1.0000	1.0622	0.0622	0.0092	0.0622	0.0629
ω	0.8000	0.8508	0.0508	0.0041	0.0508	0.0510
ρ	0.3000	0.3476	0.0476	0.0062	0.0476	0.0480

Table 9: Monte Carlo Simulation Results: DGP 2, N=1000

	TRUE	MC Mean	MC Bias	MC Std	MAE	RMSE
$m_0 : a_0$	0.0000	0.0233	0.0233	0.0078	0.0235	0.0246
$m_0 : a_1$	0.8000	0.7590	-0.0410	0.0835	0.0706	0.0926
$m_0 : a_2$	0.0000	-0.0055	-0.0055	0.0021	0.0055	0.0059
$m_0 : a_3$	0.0000	0.0025	0.0025	0.0012	0.0025	0.0027
$m_1 : b_0$	0.5000	0.5423	0.0423	0.0899	0.0815	0.0990
$m_1 : b_1$	1.1000	1.0667	-0.0333	0.0373	0.0389	0.0498
$m_1 : b_2$	0.0000	-0.0062	-0.0062	0.0022	0.0062	0.0066
$m_1 : b_3$	0.0000	0.0012	0.0012	0.0009	0.0012	0.0015
μ_η	0.0000	-0.0297	-0.0297	0.0000	0.0297	0.0297
σ_η	1.0000	1.1272	0.1272	0.0417	0.1272	0.1338
ω	0.8000	0.8483	0.0483	0.0450	0.0577	0.0659
ρ	0.3000	0.3440	0.0440	0.0547	0.0560	0.0700

Table 10: Monte Carlo Simulation Results: DGP 3, N=1000

	TRUE	MC Mean	MC Bias	MC Std	MAE	RMSE
$m_0 : a_0$	0.0000	0.0221	0.0221	0.0100	0.0223	0.0242
$m_0 : a_1$	0.8000	0.8916	0.0916	0.1655	0.1713	0.1884
$m_0 : a_2$	0.0500	0.0242	-0.0258	0.0124	0.0261	0.0286
$m_0 : a_3$	0.0000	0.0017	0.0017	0.0016	0.0017	0.0024
$m_1 : b_0$	0.5000	0.5173	0.0173	0.0871	0.0676	0.0883
$m_1 : b_1$	0.8000	0.7788	-0.0212	0.0791	0.0630	0.0815
$m_1 : b_2$	0.0500	0.0261	-0.0239	0.0155	0.0245	0.0285
$m_1 : b_3$	0.0000	0.0029	0.0029	0.0019	0.0029	0.0035
μ_η	0.0000	-0.0297	-0.0297	0.0000	0.0297	0.0297
σ_η	1.0000	1.0800	0.0800	0.0268	0.0800	0.0843
ω	0.8000	0.8781	0.0781	0.0490	0.0837	0.0920
ρ	0.3000	0.3397	0.0397	0.0637	0.0560	0.0748

Table 11: Monte Carlo Simulation Results: DGP 4, N=1000

	TRUE	MC Mean	MC Bias	MC Std	MAE	RMSE
$m_0 : a_0$	0.0000	0.0207	0.0207	0.0008	0.0207	0.0207
$m_0 : a_1$	0.2000	0.1677	-0.0323	0.0118	0.0323	0.0344
$m_0 : a_2$	0.1000	0.1199	0.0199	0.0087	0.0200	0.0217
$m_0 : a_3$	0.0000	0.0049	0.0049	0.0002	0.0049	0.0049
$m_1 : b_0$	0.5000	0.4845	-0.0155	0.0238	0.0241	0.0283
$m_1 : b_1$	0.2000	0.2245	0.0245	0.0186	0.0283	0.0307
$m_1 : b_2$	0.1000	0.0697	-0.0303	0.0032	0.0303	0.0304
$m_1 : b_3$	0.0000	0.0050	0.0050	0.0002	0.0050	0.0050
μ_η	0.0000	-0.0297	-0.0297	0.0000	0.0297	0.0297
σ_η	1.0000	1.0500	0.0500	0.0071	0.0500	0.0505
ω	0.8000	0.8547	0.0547	0.0039	0.0547	0.0549
ρ	0.3000	0.3426	0.0426	0.0053	0.0426	0.0430

Table 12: Monte Carlo Simulation Results: DGP 1, N=100

	TRUE	MC Mean	MC Bias	MC Std	MAE	RMSE
$m_0 : a_0$	0.0000	0.0204	0.0204	0.0004	0.0204	0.0204
$m_0 : a_1$	0.8000	0.7839	-0.0161	0.0230	0.0199	0.0280
$m_0 : a_2$	0.0000	-0.0051	-0.0051	0.0001	0.0051	0.0051
$m_0 : a_3$	0.0000	0.0050	0.0050	0.0001	0.0050	0.0050
$m_1 : b_0$	0.5000	0.4501	-0.0499	0.0148	0.0503	0.0520
$m_1 : b_1$	0.8000	0.7462	-0.0538	0.0209	0.0539	0.0577
$m_1 : b_2$	0.0000	-0.0051	-0.0051	0.0002	0.0051	0.0051
$m_1 : b_3$	0.0000	0.0049	0.0049	0.0001	0.0049	0.0049
μ_η	0.0000	-0.0297	-0.0297	0.0000	0.0297	0.0297
σ_η	1.0000	1.0648	0.0648	0.0214	0.0648	0.0682
ω	0.8000	0.8508	0.0508	0.0026	0.0508	0.0509
ρ	0.3000	0.3482	0.0482	0.0024	0.0482	0.0482

Table 13: Monte Carlo Simulation Results: DGP 2, N=100

	TRUE	MC Mean	MC Bias	MC Std	MAE	RMSE
$m_0 : a_0$	0.0000	0.0205	0.0205	0.0046	0.0205	0.0210
$m_0 : a_1$	0.8000	0.7403	-0.0597	0.0768	0.0747	0.0969
$m_0 : a_2$	0.0000	-0.0051	-0.0051	0.0013	0.0052	0.0053
$m_0 : a_3$	0.0000	0.0042	0.0042	0.0009	0.0042	0.0043
$m_1 : b_0$	0.5000	0.4990	-0.0010	0.0767	0.0582	0.0763
$m_1 : b_1$	1.1000	0.9925	-0.1075	0.0387	0.1075	0.1142
$m_1 : b_2$	0.0000	-0.0056	-0.0056	0.0010	0.0056	0.0056
$m_1 : b_3$	0.0000	0.0035	0.0035	0.0013	0.0035	0.0038
μ_η	0.0000	-0.0297	-0.0297	0.0000	0.0297	0.0297
σ_η	1.0000	1.2163	0.2163	0.0673	0.2163	0.2264
ω	0.8000	0.8586	0.0586	0.0167	0.0586	0.0609
ρ	0.3000	0.3433	0.0433	0.0266	0.0455	0.0507

Table 14: Monte Carlo Simulation Results: DGP 3, N=100

	TRUE	MC Mean	MC Bias	MC Std	MAE	RMSE
$m_0 : a_0$	0.0000	0.0214	0.0214	0.0044	0.0214	0.0219
$m_0 : a_1$	0.8000	0.8117	0.0117	0.1352	0.1135	0.1350
$m_0 : a_2$	0.0500	0.0261	-0.0239	0.0065	0.0239	0.0247
$m_0 : a_3$	0.0000	0.0032	0.0032	0.0015	0.0032	0.0036
$m_1 : b_0$	0.5000	0.4999	-0.0001	0.0924	0.0754	0.0919
$m_1 : b_1$	0.8000	0.7105	-0.0895	0.1040	0.1089	0.1368
$m_1 : b_2$	0.0500	0.0429	-0.0071	0.0189	0.0150	0.0202
$m_1 : b_3$	0.0000	0.0039	0.0039	0.0014	0.0039	0.0042
μ_η	0.0000	-0.0297	-0.0297	0.0000	0.0297	0.0297
σ_η	1.0000	1.1349	0.1349	0.0694	0.1349	0.1515
ω	0.8000	0.8729	0.0729	0.0234	0.0729	0.0765
ρ	0.3000	0.3450	0.0450	0.0336	0.0478	0.0560

Table 15: Monte Carlo Simulation Results: DGP 4, N=100

	TRUE	MC Mean	MC Bias	MC Std	MAE	RMSE
$m_0 : a_0$	0.0000	0.0204	0.0204	0.0010	0.0204	0.0204
$m_0 : a_1$	0.2000	0.1738	-0.0262	0.0152	0.0262	0.0303
$m_0 : a_2$	0.1000	0.1274	0.0274	0.0118	0.0284	0.0298
$m_0 : a_3$	0.0000	0.0050	0.0050	0.0002	0.0050	0.0050
$m_1 : b_0$	0.5000	0.4626	-0.0374	0.0355	0.0450	0.0514
$m_1 : b_1$	0.2000	0.2313	0.0313	0.0348	0.0430	0.0466
$m_1 : b_2$	0.1000	0.0706	-0.0294	0.0026	0.0294	0.0295
$m_1 : b_3$	0.0000	0.0051	0.0051	0.0002	0.0051	0.0051
μ_η	0.0000	-0.0297	-0.0297	0.0000	0.0297	0.0297
σ_η	1.0000	1.0540	0.0540	0.0203	0.0541	0.0577
ω	0.8000	0.8524	0.0524	0.0058	0.0524	0.0528
ρ	0.3000	0.3470	0.0470	0.0059	0.0470	0.0474

Table 16: Monte Carlo Estimation Results: Discrete Case

Parameters	TRUE	MC Mean	MC Bias	MC Std	MAE	RMSE
ω	0.8000	0.7923	-0.0077	0.3654	0.2018	0.3653
ρ	0.3000	0.3346	0.0346	0.2264	0.1298	0.2290
$Pr(s_1 \vert s_1, y = 0)$	0.5000	0.5027	0.0027	0.0246	0.0183	0.0248
$Pr(s_2 \vert s_2, y = 0)$	0.2000	0.1918	-0.0082	0.0367	0.0237	0.0376
$Pr(s_1 \vert s_1, y = 1)$	0.3000	0.2924	-0.0076	0.0344	0.0232	0.0352
$Pr(s_2 \vert s_2, y = 1)$	0.6000	0.6189	0.0189	0.0553	0.0318	0.0584