

Hu, Yingyao; Xin, Yi

Working Paper

Identification and estimation of dynamic structural models with unobserved choices

cemmap working paper, No. CWP44/20

Provided in Cooperation with:

Institute for Fiscal Studies (IFS), London

Suggested Citation: Hu, Yingyao; Xin, Yi (2020) : Identification and estimation of dynamic structural models with unobserved choices, cemmap working paper, No. CWP44/20, Centre for Microdata Methods and Practice (cemmap), London, <https://doi.org/10.47004/wp.cem.2020.4420>

This Version is available at:

<https://hdl.handle.net/10419/241919>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

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 CWP44/20



Economic
and Social
Research Council

Identification and Estimation of Dynamic Structural Models with Unobserved Choices*

Yingyao Hu [†]
Johns Hopkins University

Yi Xin [‡]
California Institute of Technology

September 4, 2020

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 models with serially correlated unobserved heterogeneity, cases in which state variables are discrete or choices are partially unavailable, and dynamic discrete games. We propose a sieve maximum likelihood estimator for primitives in agents' utility functions and state transition rules. Monte Carlo simulation results support the validity of the proposed approach.

Keywords: dynamic discrete choice, unobserved choice, moral hazard, unobserved heterogeneity, dynamic discrete game, nonparametric identification

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

*We thank seminar and conference participants at Caltech, Bank of Canada, Tsinghua University, USC, UC Berkeley Simons Institute, and the Econometric Society North American Winter Meeting (2020). The usual disclaimer applies. The online appendix to this paper can be found [here](#).

[†]Department of Economics, Johns Hopkins University, 3400 N. Charles Street, Baltimore, MD 21218. Email: yhu@jhu.edu.

[‡]Division of the Humanities and Social Sciences, California Institute of Technology, 1200 East California Boulevard, Pasadena, CA 91125. 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 always be observed by researchers, and there are multiple reasons why this may occur. For example, in many panel survey datasets, some choice variables of interest (such as individuals' investment decisions on human capital, health, and child development, etc.) are not included as a result of the survey design. Data on consumer choices could be proprietary or highly regulated by the government due to privacy concerns. Moreover, there may be more than one dimension of choices (e.g., whether individuals search and the intensity of search) but not all are observed.¹ In some other scenarios, individuals may have inherent incentives *not* to disclose (or truthfully report) their choices. For instance, an executive officer or a politician may not be willing to reveal the actual amount of time and effort they spend promoting the growth of the company or the economy. In such contexts where a potential moral hazard problem exists, it is even more difficult for researchers to observe or collect data on agents' choices. An important research question therefore arises: when choices are unobserved (or at least not fully observed), can we still uncover the decision-making process and infer agents' preferences from the data?

In this paper, we provide novel identification results for dynamic structural models with unobserved choice variables, which have received little attention in the literature.² We focus on identifying the “first-step” objects, i.e., conditional choice probabilities (CCP's) and state transition rules, when agents' choices are not observed by econometricians.³ In the baseline analysis, we consider a single-agent finite-horizon dynamic discrete choice model with a continuous state variable. The state transition process is specified through a non-parametric regression model with an additive error. We 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 model, agents' choice probabilities are inherently *time-varying*.⁴ If the

¹We thank an anonymous referee for suggesting this.

² In the existing literature on dynamic discrete choice models, researchers mainly focused on the cases in which choices are observable. Classic examples include engine replacement decisions in [Rust \(1987\)](#), parental contraceptive choices in [Hotz and Miller \(1993\)](#), occupational choices in [Keane and Wolpin \(1997\)](#), employee retirement decisions in [Rust and Phelan \(1997\)](#), retail firm inventory strategies in [Aguirregabiria \(1999\)](#), and water authority pricing behavior in [Timmins \(2002\)](#), etc. See [Eckstein and Wolpin \(1989\)](#), [Rust \(1994\)](#), [Aguirregabiria and Mira \(2010\)](#), and [Arcidiacono and Ellickson \(2011\)](#) for comprehensive surveys on dynamic discrete choice structural models.

³Once these are identified, we can apply the results in [Magnac and Thesmar \(2002\)](#) and [Arcidiacono and Miller \(2020\)](#) to nonparametrically identify utility functions.

⁴For example, when an executive in a firm is close to retirement, s/he may have fewer incentives to exert effort; the probability of shirking may exhibit an upward trend.

state transition process is stationary, which is typically assumed in the literature, then the differences in the observed state transition process across periods are driven purely by the differences in choice probabilities.⁵ Therefore, exploiting variations in moments of observed future state distributions across periods helps identify the unobserved choice probabilities and the latent state transition process.

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 \(2009\)](#), [Arcidiacono and Miller \(2011\)](#), and [Hu and Shum \(2012\)](#) have provided solutions to deal with unobserved heterogeneity. Following [Hu and Shum \(2012\)](#), we use the joint distribution of the observed state variable at four consecutive periods to identify the transition of the observed state conditional on the unobserved heterogeneity, to which we can further apply our method to deal with unobserved choices. Second, we discuss identification for infinite-horizon models. In finite-horizon models, time essentially serves as an exclusion restriction. We show that as long as there exists an excluded variable that only shifts choice probabilities but does not affect the latent state transition process, the baseline identification results remain valid for infinite-horizon models. 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 processes 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.⁶ Fourthly, we show that similar identification arguments apply given partial unobservability of choices.

Our identification results are not limited to single-agent dynamic models. We show in this paper that the proposed approach can be extended to dynamic discrete games of incomplete information when choice data are not available.⁷ 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; however, the state transition process for a player may only depend on his own actions and state variables in the past.⁸ In that case, the state of other

⁵ For the identification of dynamic discrete choice models when the data generating processes are nonstationary and the panels are *short* (i.e., the time horizon of the agent extends beyond the length of the data), see [Arcidiacono and Miller \(2020\)](#).

⁶ If two continuous state variables are available, it’s possible to extend our results to continuous choices.

⁷ Existing papers that develop estimation techniques for dynamic discrete games generally require the observation of choices ([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.)

⁸ For example, in a dynamic oligopoly game where the state variable is the firm’s capacity levels and the choice is incremental changes to capacity, it is reasonable to assume that the firm’s future capacity levels only depend on its own decisions, not on other firms’ choices. See [Aguirregabiria, Mira, and Roman \(2007\)](#),

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 unobserved choices in dynamic discrete games.

Following the identification results, we propose a sieve maximum likelihood estimation strategy to jointly estimate primitives in agent’s utility functions and state transition rules. We conduct Monte Carlo simulations to examine the finite sample performances of our estimator under different data generating processes. Overall, our Monte Carlo simulations perform well. Compared to the estimation results assuming that the choices are observed by the econometricians, the results assuming choices are unobserved exhibit slightly larger finite sample biases. In addition to the simulations, for illustration purposes, we apply our method to a publicly available dataset containing all gubernatorial elections in the United States from 1950–2000. We estimate a dynamic discrete choice model for governors’ effort exerting decisions, which are not observed by econometricians. Our empirical analysis suggests that the probability of shirking increases as the governors approach the end of their terms.

Our paper is closely related to the literature on the identification of dynamic discrete choice models (Rust, 1994; Magnac and Thesmar, 2002; Abbring, 2010; Norets and Tang, 2014; Arcidiacono and Miller, 2020, etc.), which unexceptionally requires the observation of agents’ choices. Arcidiacono and Miller (2020) summarize the necessary and sufficient conditions for identifying a certain class of models, where the utility function is time-separable, the unobserved states are conditionally independent and additively separable, and the agents’ beliefs are rational.⁹ Several papers have explored using additional assumptions (e.g., parametric assumptions on utility functions, exclusion restrictions on which state variables affect payoffs, availability of terminating actions, etc.) to identify the discount factor, or counterfactual policies without normalizing per period payoffs (Aguirregabiria, 2010; Bajari et al., 2016; Abbring and Daljord, 2020, etc.). Our paper in general fits into this literature and we focus on relaxing the assumption of full data coverage (i.e., observations of state and choice variables for a random set of agents for a sufficient period of time). Relatedly, Arcidiacono and Miller (2020) study identification of dynamic discrete choice models when the panels are *short*, so that the choices and state transitions after the sample period are not observed. In this paper, we impose assumptions on the state transition process to achieve identification of model primitives when the data do not cover agents’ choices.

For the estimation of dynamic discrete choice models in the existing literature, agents’

Ryan (2012), Collard-Wexler (2013), and Takahashi (2015) for more details on empirical models of oligopoly dynamics.

⁹An, Hu, and Xiao (2018) studies identification of dynamic discrete choice models allowing subjective beliefs; Aguirregabiria and Magesan (2020) studies identification and estimation of dynamic discrete games allowing the players’ beliefs are not in equilibrium.

choices are usually needed to construct (pseudo) likelihood or to do first-stage nonparametric estimation of the conditional choice probabilities and the state transition probabilities (see Rust, 1987; Hotz and Miller, 1993; Hotz et al., 1994; Aguirregabiria and Mira, 2002, etc.). Our paper complements this literature by developing estimation strategies that do not rely on the availability of agents' choices.

There are few empirical papers focused on models with unobserved choices (Misra and Nair, 2011; Copeland and Monnet, 2009; Gayle and Miller, 2015; Gayle, Golan, and Miller, 2015; Perrigne and Vuong, 2011; Xin, 2020, etc.) In this paper, we incorporate unobserved choice variables into a general framework of dynamic discrete choice models. Different from the papers cited above, our identification strategies do not rely on multiple measurements of the unobserved choice, exogenous variations in incentive schemes, one-to-one mapping between the choice and observables, or variations in the support of observables across different choices. In addition, our framework allows the realization of distinct choices in the data generating process.¹⁰

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 Baseline Model

We first fix the notation for a standard single-agent dynamic discrete choice model with $t = 1, 2, \dots, T < \infty$. Let $s_t \in \mathcal{S}$ represent the observed state variable and $y_t \in \mathcal{Y} = \{1, 2, \dots, J\}$ denote the agent's choice. We use $\varepsilon_t = (\varepsilon_t(1), \varepsilon_t(2), \dots, \varepsilon_t(J)) \in \mathbb{R}^J$ to represent the state variable that is unobserved by 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 the discounted utility stream of the agent is therefore defined as

$$U(\mathbf{s}, \boldsymbol{\varepsilon}, \mathbf{y}) = \sum_{t=1}^T \beta^{t-1} u_t(s_t, \varepsilon_t, y_t), \quad (2.1)$$

where $\mathbf{s} = (s_1, \dots, s_T)$, $\boldsymbol{\varepsilon} = (\varepsilon_1, \dots, \varepsilon_T)$, $\mathbf{y} = (y_1, \dots, y_T)$, and β is the discount factor. The agent's problem is to choose an optimal decision rule $\delta = (\delta_1, \dots, \delta_T)$ that maximizes the

¹⁰Gayle, Golan, and Miller (2015) studies the executives' effort choice problems and focuses on the equilibrium in which all executives choose to work.

expected sum of the discounted utility, i.e.,

$$\max_{\delta=(\delta_1, \dots, \delta_T)} E(U(\mathbf{s}, \boldsymbol{\varepsilon}, \mathbf{y})),$$

where expectation is taken 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 the following conditions.*

- (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 $\Omega_{<t} \equiv \{s_{t-1}, \dots, s_1, \varepsilon_{t-1}, \dots, \varepsilon_1, y_{t-1}, \dots, y_1\}$.
- (ii) *Conditional independence:* 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) *Time-invariance:* 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 relaxed to allow for a 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 recent literature on identification and estimation of dynamic discrete choice models when the unobserved state variables are serially correlated.¹¹ In Section 6.1, we show that our identification results can be extended to models that incorporate serially correlated unobserved heterogeneity when at least five periods of data are available.¹² In order to highlight the identification intuition related to unobserved choice variables and for clarity of exposition, we first focus on the case when Assumption 1(ii) is invoked. Assumption 1(iii) guarantees that the state transition process does not change over time and is usually invoked in dynamic models.¹³ 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.

¹¹See Aguirregabiria and Mira (2007), Houde and Imai (2006), Kasahara and Shimotsu (2009), and Hu and Shum (2012) for examples.

¹²When the unobserved heterogeneity is fixed over time, we show that only four consecutive periods are needed.

¹³Bajari et al. (2016) impose the same assumption of time-homogeneous state transition for finite-horizon dynamic discrete choice models; see Rust (1987) for a discussion on stationary infinite-horizon models.




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

We further impose the additive separability assumption on the per-period utility function.

Assumption 2 (Additive Separability). $u_t(s_t, \varepsilon_t, y_t) = u_t^*(s_t, y_t) + \varepsilon_t(y_t)$.

Under Assumptions 1–2, we represent the agent’s optimization problem using the Bellman’s equation as follows.

$$V_t(s_t, \varepsilon_t) = \max_y u_t^*(s_t, y) + \varepsilon_t(y) + \beta \mathbb{E}[V_{t+1}(s_{t+1}, \varepsilon_{t+1}) | s_t, y]. \quad (2.2)$$

The agent’s decision rule is hence defined by

$$\delta_t(s_t, \varepsilon_t) = \arg \max_y \left\{ u_t^*(s_t, y) + \varepsilon_t(y) + \beta \mathbb{E}[V_{t+1}(s_{t+1}, \varepsilon_{t+1}) | s_t, y] \right\}. \quad (2.3)$$

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.

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

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, the primitives include mean utility functions $u_t^*(\cdot, \cdot)$, the distribution of the unobserved state variable $F_{\varepsilon_t | s_t}(\cdot | \cdot)$, the state transition rule $f_{s_{t+1} | s_t, y_t}$, and the discount factor β .¹⁴ This model is generally *not* identified even when the choice variable is observed. [Magnac and Thesmar \(2002, Proposition 2 and Corollary 3\)](#) provide formal results that utility functions in each alternative are nonparametrically identified from the conditional choice probabilities $p_t(y_t | s_t)$ and the state transition probabilities $f_{s_{t+1} | s_t, y_t}$ for $T = 2$, if the distribution of the unobserved state, the discount factor, and the utility in one reference alternative are known. [Arcidiacono and Miller \(2020, Theorem 3\)](#) characterize the identification of flow payoffs for a more general case nesting both the infinite and finite horizon models.

¹⁴ When choices are observed, structural transition probabilities, which can be directly recovered from the data, are usually not included in the list of parameters to be identified ([Magnac and Thesmar \(2002\)](#)).

When the choice variable y_t is not observed by econometricians, we cannot recover the decision rules nor the state transition probabilities directly from the data. These two objects are necessary for the identification results in [Magnac and Thesmar \(2002\)](#) and [Arcidiacono and Miller \(2020\)](#); moreover, they are the important first-step outputs for the two-step CCP method developed by [Hotz and Miller \(1993\)](#) for estimating dynamic discrete choice models. Overall, the existing methods fail to obtain sufficient ingredients for identifying and estimating structural primitives when agents' choices are not observed by econometricians.

In Section 3, we focus on identifying the conditional choice probabilities and the state transition rules when the choices made by agents are not observed in the data. Throughout the paper, we assume that (1) the researcher knows the set of alternatives from which the agent chooses and the horizon of the agents solving their problems, (2) the distribution of the unobserved state and the discount factor are known, and (3) payoff from one action is known for every state and time period. Once the conditional choice probabilities and the state transition rules are identified, it's straightforward to apply the Hotz-Miller inversion on the CCP's and nonparametrically identify the mean utility functions following [Magnac and Thesmar \(2002\)](#) and [Arcidiacono and Miller \(2020\)](#). In Section 4, we provide estimation strategies for jointly estimating primitives in utility functions and state transition rules.

3 Identification

In this section, we provide identification results for 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 $\{s_t\}_{t=1}^{T+1}$ (with $T \geq 2$) are observed for a random sample of agents.¹⁵ We first focus on the case in which s_t is a continuous state variable. To highlight the feature that the choice variable is unobserved by 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).

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

where s' and s represent realized values 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

¹⁵As described in Section 2, the agent solves the dynamic problem for $t = 1, 2, \dots, T$. Here we assume that after the final period T , we observe the realization of the state variable at $T + 1$. The observation of s_{T+1} is necessary for recovering the choice probabilities at T .

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 time-invariant; 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 exploit 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 3 (State Transition). $s_{t+1} = m(y_t^*, s_t) + \eta_t$, where $m(\cdot, \cdot)$ is continuously differentiable in s_t , $E(\eta_t|s_t) = 0$, $\eta_t \perp y_t^*|s_t$, and the conditional density function of the error term $f_{\eta_t|s_t}(\cdot|\cdot)$ is continuous in s_t .

Assumption 3 specifies the transition process of the continuous state variable s_t through a nonparametric regression model, where $m(\cdot, \cdot)$ is an unknown smooth function and η_t represents the random shock realized in the transition process with conditional mean equal to zero. Assumption 3 also requires that the regression error is independent of the unobserved choice conditional on the state variable. This conditional independence assumption ensures that the impact of the choices on the state transition process is only through the deterministic part but not through the error term—conditional on the current state, the choice made by the agent only shifts the mean of the future state distribution. Combining Assumption 3 and Assumption 1(iii), we know that the conditional distribution of η_t is stationary. That is, for $t, \tau \in \{1, \dots, T\}$, $f_{\eta_t|s_t}(\eta|s) = f_{\eta_\tau|s_\tau}(\eta|s)$, $\forall \eta, s$. Furthermore, we assume that the conditional density function of η_t is continuous in s_t . By Assumption 3, 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 to be identified in addition to the unobserved choice probabilities.

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 equivalent to identifying two functions of s_t , i.e., $m(y_t^* = 0, s_t)$ and $m(y_t^* = 1, s_t)$. We begin our analysis with a fixed state s . To simplify the notation, let $m_1 = m(1, s)$ and $m_0 = m(0, s)$. We use $p_t = p_t(y_t^* = 1|s)$ and $1 - p_t = p_t(y_t^* = 0|s)$ to denote the choice probabilities associated with choices 1 and 0, respectively, at period t . We define the first-, the second-, and the third-order conditional moments of the observed state variable at $t + 1$ as follows.

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

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

Note that μ_t , ν_t , and ξ_t are essentially weighted averages of moments of the future state conditional on the current state and the choice, where the choice probabilities $(p_t, 1 - p_t)$ serve as the mixing weights. Given that $s_{t+1} = m(y_t^*, s_t) + \eta_t$ and η_t and y_t^* are independent conditional on s_t under Assumption 3, μ_t , ν_t , and ξ_t can be represented as functions of m_1 , m_0 , p_t , and moments of η_t conditional on $s_t = s$. For example,

$$\mu_{t+1} = p_t m_1 + (1 - p_t) m_0 + E(\eta_t | s).$$

By Assumption 3, the conditional mean of η_t equals 0, i.e., $E(\eta_t | s) = 0$. Therefore, the choice probability

$$p_t = \frac{\mu_{t+1} - m_0}{m_1 - m_0}, \quad (3.2)$$

provided that $m_1 \neq m_0$.

Under Assumption 1(iii) and Assumption 3, the conditional distribution of η_t is stationary. This implies that the higher order moments of the error term η_t are time-invariant conditional on the same state s . Taking the difference of moments of the observed state variable (i.e., ν_{t+1} and ξ_{t+1}) across two periods t and τ eliminates the unknown moments of η_t and leads to a system of equations for m_1 and m_0 . We show that m_0 and m_1 are the two solutions to the equation of \bar{m} :

$$\bar{m}^2 - \Delta_1 \bar{m} + \Delta_2 = 0, \quad (3.3)$$

where

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

The existence of the solutions in Equation (3.3) requires $\Delta_1^2 - 4\Delta_2 \geq 0$. This condition is empirically testable as Δ_1 and Δ_2 are identified and can be directly estimated using moments of the observed state variable provided that $\mu_{t+1} \neq \mu_{\tau+1}$. The condition that $\mu_{t+1} \neq \mu_{\tau+1}$ implies $m_1 \neq m_0$ and $p_t \neq p_\tau$, and is also empirically testable from the data. If $m_1 = m_0$, the identification of m functions at state s is trivial as $m_1 = m_0 = \mu_{t+1} = \mu_{\tau+1}$, but the choice probabilities are not identified as shown in Equation (3.2). If $m_1 \neq m_0$ and $p_t = p_\tau$, we have no variations in observables from periods t and τ to identify the m functions and choice probabilities. We invoke the following assumption to restrict our attention to cases where $\mu_{t+1} \neq \mu_{\tau+1}$ and to pin down the order of m_0 and m_1 .

Assumption 4. *The following conditions are satisfied:*

(i) *There exist two periods t and τ such that $Pr(\tilde{\mathcal{S}}) = 0$ with*

$$\tilde{\mathcal{S}} := \{s \in \mathcal{S} : E_{t+1}[s_{t+1}|s_t = s] = E_{\tau+1}[s_{\tau+1}|s_\tau = s]\},$$

(ii) *$\frac{\partial m(1,s)}{\partial s} \neq \frac{\partial m(0,s)}{\partial s}$ holds for any $s \in \bar{\mathcal{S}} := \{s \in \tilde{\mathcal{S}} : m(1, s) = m(0, s)\}$,*

(iii) *There exists an $s_0 \in \mathcal{S}/\bar{\mathcal{S}}$ such that $m(1, s_0) > m(0, s_0)$.*

Assumption 4(i), imposed on observables, ensures that there are *at most* countable many values of the state variable at which m functions are not identified from Equations (3.3) using data from periods t and τ .¹⁶ The values of m functions for $s \in \tilde{\mathcal{S}}$ are, instead, identified using the continuity of $m(1, \cdot)$ and $m(0, \cdot)$ in s . Assumption 4(ii) rules out the possibility that the first-order derivatives of $m(1, \cdot)$ and $m(0, \cdot)$ are equal at the point they intersect (e.g., $m(1, \cdot)$ and $m(0, \cdot)$ are tangent to each other). Note that $\bar{\mathcal{S}}$ also has zero probability as it is a subset of $\tilde{\mathcal{S}}$. Under Assumption 4(i)–(ii), as long as there exists a state s at which we can order the two m functions, the smoothness condition (i.e., continuously differentiable) helps match $m(1, \cdot)$ and $m(0, \cdot)$ across all values of $s \in \mathcal{S}$.¹⁷ Assumption 4(iii) provides such an ordering condition as needed. Note that the assumption that $m(1, s_0) > m(0, s_0)$ is without loss of generality. Before linking the identified choice probabilities and conditional state transition probabilities to structural utility primitives, we can swamp the labels of the two choices.

Once m_1 and m_0 are identified (and if they are not equal), the conditional choice probabilities are identified from Equation (3.2). 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 $s' - m_0$; again, the choice probabilities $(p_t, 1 - p_t)$ serve as the mixing weights. With variations in choice probabilities across two periods (i.e., $p_t \neq p_\tau$) and the stationarity of η_t conditional on s_t , the conditional density function of η_t is also identified, i.e.,

$$\begin{aligned} f_{\eta_t|s_t}(s' - m_1|s) &= \frac{f_{s_{t+1}|s_t}(s'|s)(1 - p_\tau) - f_{s_{\tau+1}|s_\tau}(s'|s)(1 - p_t)}{p_t - p_\tau}, \\ f_{\eta_t|s_t}(s' - m_0|s) &= \frac{f_{s_{\tau+1}|s_\tau}(s'|s)p_t - f_{s_{t+1}|s_t}(s'|s)p_\tau}{p_t - p_\tau}. \end{aligned} \tag{3.4}$$

At the state s in which $p_t(\cdot|s) = p_\tau(\cdot|s)$, $f_{\eta_t|s_t}(\cdot|s)$ is instead identified by the continuity of the conditional density function of η_t in s_t as imposed in Assumption 3.¹⁸

¹⁶For $s \in \tilde{\mathcal{S}}$, the first-order conditional moments of the state variable are the same across two periods t and τ , i.e., $\mu_{t+1} = \mu_{\tau+1}$, so that the Equation (3.3) is not well defined.

¹⁷We thank an anonymous referee for suggesting this.

¹⁸Note that $\{s \in \mathcal{S} : p_t(\cdot|s) = p_\tau(\cdot|s)\}$ is a subset of $\tilde{\mathcal{S}}$, thus also has zero probability. Once $f_{\eta_t|s_t}(\cdot|s)$

We summarize the formal identification results in the following theorem.

Theorem 1 (Identification). *If Assumptions 1, 2, 3, and 4 hold for the dynamic process of $\{s_t, \varepsilon_t, y_t^*\}$ with $y_t^* \in \{0, 1\}$, then the observed conditional densities $f_{s_{t+1}|s_t}(\cdot|\cdot)$ and $f_{s_{\tau+1}|s_\tau}(\cdot|\cdot)$ for t and τ defined in Assumption 4 uniquely determine:*

(i) *the state transition $f_{s_{r+1}|s_r, y_r^*}$, including $m(1, s)$, $m(0, s)$, and $f_{\eta_r|s_r}(\cdot|s)$ for all $r \in \{1, 2, \dots, T\}$ and $s \in \mathcal{S}$;*

(ii) *choice probabilities $p_t(\cdot|s)$ and $p_\tau(\cdot|s)$ for $s \in \mathcal{S}/\bar{\mathcal{S}}$ with $Pr(\bar{\mathcal{S}}) = 0$.*

Proof. See Appendix A. □

Remark 1. *If additional smoothness assumptions are imposed on the utility functions, choice probabilities $p_t(\cdot|s)$ and $p_\tau(\cdot|s)$ are also identified for $s \in \bar{\mathcal{S}}$.*

Once the state transition rules are identified, choice probabilities $p_r(\cdot|s)$ are identified from $f_{s_{r+1}|s_r}(\cdot|s)$ for all $r \in \{1, 2, \dots, T\}$ and $s \in \mathcal{S}/\bar{\mathcal{S}}$ with $Pr(\bar{\mathcal{S}}) = 0$. In general, identifying the choice probabilities for all sample periods requires $\{s_t\}_{t=1}^{T+1}$ (with $T \geq 2$). If we are only interested in identifying choice probabilities for a certain period t , the observations of state variables for at least three consecutive time periods (including t and $t + 1$) are necessary.

The identification argument in our paper essentially follows a sequential approach.¹⁹ Theorem 1 provides identification results for the conditional choice probabilities and state transition rules when the agents' choice are not observed by econometricians up to *label swamping*. To associate choice probabilities and state transition rules with specific alternatives faced by the agent, various assumptions arising from the model or consistent with the economic intuition could be imposed. For example, consider a scenario where s_t represents the realized revenue of a loan and y_t^* represents the borrower's choice. Suppose we denote $y_t^* = 1$ if the borrower exerts effort to pay off the debt, and 0 otherwise; the agent's utility function has the following form $u^*(s_t, y_t^*) = \omega s_t - \rho \mathbf{1}\{y_t^* = 1\}$, where ρ represents the cost of exerting effort. It is reasonable to assume that when the borrower exerts effort, the revenue distribution first-order stochastically dominates the one when he exerts no effort, so that $m(1, s) \geq m(0, s)$ for all $s \in \mathcal{S}$. With this assumption, we can label the larger value of the m functions at state s as $m(1, s)$ and identify the probability that the borrower exerts effort accordingly.²⁰

is identified for all $s \in \mathcal{S}$ such that $p_t(\cdot|s) \neq p_\tau(\cdot|s)$, by the continuity of the density function, $f_{\eta_t|s_t}(\cdot|s)$ is identified for all $s \in \mathcal{S}$.

¹⁹ For more discussions on the sequential versus joint identification approaches in dynamic structural models, see Aguirregabiria and Mira (2019).

²⁰ If we swap the label and denote $y_t^* = 0$ if the borrower exerts effort to pay off the debt, the utility function can be written as $u^*(s_t, y_t^*) = \omega s_t - \rho \mathbf{1}\{y_t^* = 0\}$, and the larger value of the m functions is labeled as $m(0, s)$.

Once the conditional choice probabilities and state transition rules are recovered for each alternative, the identification of per-period utility functions in the second step follows immediately from [Magnac and Thesmar \(2002\)](#) and [Arcidiacono and Miller \(2020\)](#). The details are hence omitted in this paper. Note that if additional assumptions are imposed on the per-period utility functions, the data requirement of $\{s_t\}_{t=1}^{T+1}$ (with $T \geq 2$) may be relaxed. For example, if we assume that, in a finite horizon model, the last period utility is different from the utility in the previous periods (while all other per period utilities are the same), the observations of the state variables for at least $t = T - 1, T, T + 1$ are needed in order to identify utility functions. If we further restrict the flow utility to not depend on t , then choice probabilities at a certain period $t \in \{1, 2, \dots, T\}$ are needed, which require the observations of state variables in at least three consecutive time periods (including t and $t + 1$).

An important extension of the identification results in [Theorem 1](#) is to incorporate serially correlated unobserved heterogeneity into the model. Intuitively, when the unobserved heterogeneity is present, in order to apply our identification strategy, the key is to first recover the transition process of the observed state *conditional on* the unobserved heterogeneity. We show in [Section 6.1](#) that the state transition process given the unobserved heterogeneity is identified from the joint distribution of state variables at four consecutive periods. We discuss several other extensions in [Section 6](#), including infinite-horizon models ([Section 6.2](#)), cases in which state variables are discrete ([Section 6.3](#)), or choice data are partially unavailable ([Section 6.4](#)), and dynamic discrete games ([Section 6.5](#)).²¹

4 Estimation

In this section, we propose to use a sieve maximum likelihood approach to jointly estimate the utility primitives, the nonparametric function $m(\cdot, \cdot)$ and the distribution of the error term $f_{\eta_t|s_t}$ in the state transition process.²² For simplicity, we consider a case where the per-period utility functions are parametrized by a finite-dimensional parameter α .²³

For $\theta = \{\alpha, m_0, m_1, f_{\eta_t|s_t}\} \in \Theta$, the log-likelihood evaluated at a single observation

²¹For cases where the choice variable y_t^* takes multiple discrete values, the approach in [Chen, Hu, and Lewbel \(2009\)](#) may be used (under a generalized version of the current assumptions). We also show in [Section 6.3](#) that when two discrete state variables that are independent conditional on the choice variable are available, the number of alternatives is not limited to two.

²²It is also possible to estimate the model primitives using a two-step sequential approach. Our identification strategy is constructive and leads to the direct nonparametric estimation of the conditional choice probabilities and state transition rules in the first step. We can then use these two objects as the inputs for the second-step estimation of utility primitives following [Hotz and Miller \(1993\)](#).

²³The estimation strategy in this section can be extended to allow the nonparametric estimation of per-period utility functions.

$D_i = \{s_{it}\}_{t=1}^{T+1}$ is derived in the following equation.

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

In Equation (4.1), $p_t(1|s_{it}; \theta)$ and $p_t(0|s_{it}; \theta)$ are the choice probabilities for alternatives 1 and 0 conditional on state s_{it} 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). To evaluate the likelihood, choice probabilities are derived via the agent's optimization problem in Equation (2.3). The population criterion function $Q : \Theta \rightarrow \mathbb{R}$ is hence defined by

$$Q(\theta) = E(l(D_i; \theta)). \quad (4.2)$$

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

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

In light of a finite sample, instead of searching parameters over an infinite-dimensional parameter space Θ , we maximize the empirical criterion function over a sequence of approximating sieve spaces Θ_k . The sieve maximum likelihood estimator $\hat{\theta}_k$ is defined as

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

We discuss the details of constructing the sieve spaces and the asymptotic properties of the proposed estimator in Appendix B.

5 Simulations

In this section, we present Monte Carlo simulation results when there is a continuous state variable. Let $y_t^* = 1$ if the agent chooses to exert effort, and 0 otherwise. We assume that the mean utility function takes the following form when the agent exerts effort:

$$u^*(s_t, y_t^*) = 1 - \exp(-\omega s_t) - \rho y_t^*,$$

where $\omega = 0.8$, and $\rho = 0.3$ measures the marginal cost of exerting more effort. Assume that the utility level when no effort is exerted equals 0, i.e., $u^*(s_t, 0) = 0$. The utility shocks $\varepsilon_t(0)$ 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: $s_{t+1} = 0.8s_t + 0.5y_t^* + 0.2s_t \cdot y_t^* + \eta_t$.
- DGP 3: $s_{t+1} = 0.6s_t + 0.05s_t^2 + 0.5y_t^* + \eta_t$.
- 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 $m_1(s_t) = 0.5 + s_t$. 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 $T = 10$. We run simulations for different sample sizes, $N = 100, 1000$, and 5000. The estimation results shown 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 2a under DGP 1. It is clear that the state distribution shifts to the right with a smaller variance as time proceeds. The variations in the state distribution are driven by the differences in choice probabilities across time periods. Figure 2b 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 $N = 5000$ in Panel (A) of Tables 4–7 in appendix E. We also run the estimation assuming that the choices are observed by the econometricians and we provide the results in Panel (B) of Tables 4–7 in appendix E. The estimation results for sample sizes equal to 100 and 1000 are provided in Tables F.1–F.8 in online appendix F. In these exercises, we use third-degree polynomials to approximate the nonparametric functions m_0 and m_1 . Specifically,

$$m_0(s) \approx a_0 + a_1s + a_2s^2 + a_3s^3,$$

$$m_1(s) \approx b_0 + b_1s + b_2s^2 + b_3s^3.$$




Figure 2: Graphic Illustration of Variations in Observed States

Note: Figure 2a plots the distribution of the state variable at periods 1, 3, 5, 7 under DGP 1. Figure 2b plots the mean of s_{t+1} conditional on $s_t = 0$ under DGP 1 for $t = 1, 2, \dots, 10$.


For the square root of the density function f_{η_t} , we use fifth-degree polynomials. In Tables 4–7, 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 σ_η , which represents the standard deviation of the error distribution.²⁴ 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 simulations generally perform well; adding nonlinear effects of the current state to the transition process leads to slightly less precise estimates. Comparing between Panel (A) and Panel (B) in Tables 4–7, the estimation results when choices are unobserved exhibit slightly larger finite sample biases than the ones estimated assuming choices are observed. As the sample size increases, the differences between the two scenarios (i.e., when choices are unobserved versus when choices are observed) become smaller. To visualize the simulation results, we plot functions m_0 and m_1 using our estimates when choices are not observed and the true parameter values in the data generating process in Figure 3.²⁵ Our nonparametric estimates of m_0 and m_1 are generally close to the true parameter values, particularly 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 average probabilities of exerting effort at each period using our estimates and compare those with the ones observed in the simulated datasets. The results

²⁴We impose the restriction that the mean of the η distribution is zero in the estimation.


²⁵Figures 3 and 4 are based on estimation results with $N = 5000$.




(a) DGP 1



(b) DGP 2



(c) DGP 3




(d) DGP 4


Figure 3: Plot m_0 and m_1 Using Estimates and the True Parameter Values

for the four data generating processes are shown in Figure 4. Note that the probability of exerting effort depends on (1) which period the agent is at, and (2) the level of the state variable. As the agent is closer to the final period of the game, the probability of exerting effort conditional on the same state level decreases. However, as shown in Figure 2a, the distribution of the state shifts to the right as time proceeds. In other words, the average levels of state variables are improved as t increases. Because of these two counterforces, we may observe a non-monotone trend of average probabilities of exerting effort.


For all data generating processes, average probabilities of exerting effort at each period predicted using our estimates are very close to those “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.




(a) DGP 1



(b) DGP 2



(c) DGP 3



(d) DGP 4

Figure 4: Choice Probabilities: Model Predictions vs. Data

6 Extensions

We focus on a single-agent finite-horizon dynamic discrete choice model with one continuous state variable to illustrate our main identification and estimation approaches. In this section, we discuss extensions to the baseline identification results. In particular, we consider scenarios in which: (1) serially correlated unobserved heterogeneity is allowed, (2) the model has infinite horizon, (3) only discrete state variables are available, (4) choice data are partially unavailable, and (5) multiple players make simultaneous decisions in a game.

6.1 Serially Correlated Unobserved Heterogeneity

We now consider a model with serially correlated unobserved heterogeneity. Following our notations for 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 impose the following assumptions on the dynamic process.²⁶

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

- (i) *First-order Markov:* $f_{s_{t+1}, \varepsilon_{t+1}, x_{t+1}^*, y_{t+1} | s_t, \varepsilon_t, x_t^*, y_t, \Omega_{<t}} = f_{s_{t+1}, \varepsilon_{t+1}, x_{t+1}^*, y_{t+1} | s_t, \varepsilon_t, x_t^*, y_t}$, where $\Omega_{<t} \equiv \{s_{t-1}, \dots, s_1, \varepsilon_{t-1}, \dots, \varepsilon_1, x_{t-1}^*, \dots, x_1^*, y_{t-1}, \dots, y_1\}$.
- (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 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 more 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., $x_{t+1}^* = x_t^*$.²⁷ The serial correlation of the unobserved heterogeneity invoked in Assumption 5(ii) is more

²⁶To address issues related to initial conditions when serially correlated unobserved heterogeneity is included, we assume that the structural dynamic discrete choice model does not apply to pre-sample periods.

²⁷Aguirregabiria and Mira (2007), Houde and Imai (2006), and Kasahara and Shimotsu (2009) study cases with time-invariant discrete unobserved heterogeneity.

general.²⁸ The dynamic process of the state and choice variables (s_t, x_t^*, y_t) that satisfies Assumption 5 is illustrated in Figure 5. 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 5: The Dynamic Process of (s_t, x_t^*, y_t)

When both unobserved choices and serially correlated unobserved heterogeneity are present, can we apply a method similar to the one 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

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

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) 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 clear 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., $f_{s_{t+1}|s_t, x_t^*}$.

In order to identify $f_{s_{t+1}|s_t, x_t^*}$, we consider the joint distribution of the observed state

²⁸Hu 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 timing restriction. In our case, the unobserved heterogeneity x_t^* realizes after the state variable s_t .

variable at four consecutive periods $(s_{t+2}, s_{t+1}, s_t, s_{t-1})$:

$$f_{s_{t+2}, s_{t+1}, s_t, s_{t-1}} = \int_{x_t^*} f_{s_{t+2}|s_{t+1}, x_t^*} \times f_{s_{t+1}|s_t, x_t^*} \times f_{x_t^*, s_t, s_{t-1}} dF_{x_t^*}. \quad (6.2)$$

The derivation of Equation (6.2) is provided in Appendix C. The key insight of this equation is that correlation among observed state variables across different periods is induced by the underlying individual heterogeneity. We can treat $(s_{t+2}, s_{t+1}, s_t, s_{t-1})$ as measurements of x_t^* . 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 $t-1$.

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 here is that the identification of latent choice and state transition probabilities when serially correlated unobserved heterogeneity is present requires the availability of at least five periods of data, i.e., $\{s_t\}_{t=1}^{T+1}$ (with $T \geq 4$).

Remark 2. *Identification of models with time-invariant unobserved heterogeneity, such as individual fixed effects, is a special case of our results in Section 6.1 that allow for serially correlated unobserved heterogeneity.²⁹ If the unobserved heterogeneity is constant over time (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 $(s_{t+1}, s_t, s_{t-1}, s_{t-2})$ using similar techniques as in Equation (6.2). This result indicates that four periods of observed state variables are sufficient to identify the latent choice and state transition probabilities conditional on individual fixed effects.*

6.2 Infinite Horizon

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 time-homogeneous, variations in the moments of the future state distribution conditional on the same current state can be attributed to 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 on the same state variable. Consequently, time cannot be used as an excluded variable any more.

²⁹In 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 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 (Exclusion Restriction). 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 time-invariance 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), we can 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. Note that for stationary infinite horizon models, the observations of state variables for at least two consecutive periods are required.

6.3 Discrete States

We provide 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. When the choice variable takes different values, both the location and the shape of the future state distribution change. In this section, we consider a case where we have two discrete state variables $\{s_t, z_t\}$ for $t = 1, \dots, T+1$ 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 6.

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




Figure 6: The Dynamic Process of (s_t, z_t, y_t)

as follows.

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

Intuitively, the future states can be viewed as proxies of the unobserved choice. Following the results in the measurement literature (Hu (2008), Hu and Shum (2012)), Equation (6.4) leads to an eigenvalue-eigenvector decomposition of a matrix constructed using observed densities on its left-hand side. We provide the details for identifying 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 $f_{y_t^*|s_t, z_t}$ from the decomposition in Appendix D.

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.*

6.4 Partial Unobservability of Choices

Our framework is also applicable to cases where there are multiple dimensions of the choices made by the agents, but not all are observed in the data. For instance, researchers may know whether a consumer searches for a product or an apartment, but it is difficult to get information on the intensity of the search effort. It might be relatively easy to collect data on whether an employee goes to work on time, but it is rather difficult to observe the degree to which one works diligently.

To take partial unobservability of choices into account, we denote the vector of choices made by the agent at period t by $y_t = (y_t^o, y_t^*)$, where y_t^o represents the vector of choices observed by the econometrician, and y_t^* represents the ones not observed. Under Assumption 1(i)–(ii), the observed state transition process can be factorized in the following equation:

$$f_{s_{t+1}|s_t, y_t^o}(s'|s, y_t^o) = \sum_{y_t^*} f_{s_{t+1}|s_t, y_t^o, y_t^*}(s'|s, y_t^o, y_t^*) p_t(y_t^*|s, y_t^o), \quad (6.5)$$

where s' and s represent realized values of s_{t+1} and s_t , respectively. In Equation (6.5), $f_{s_{t+1}|s_t, y_t^o, y_t^*}$ denotes the conditional state transition rules and $p_t(y_t^*|s, y_t^o)$ denotes the probabilities of the unobserved choice conditional on the state and other dimensions of the choices observable to researchers. Both of the terms cannot be directly recovered from the data.

Equation (6.5) is a direct extension of Equation (3.1) in the baseline model. The only difference is that now the state transition probabilities are also conditional on the observable part of the choices. Imposing similar restrictions on the latent state transition rules helps to identify the unknown primitives on the right-hand side of Equation (6.5). The choice probabilities for y_t is therefore identified:

$$p_t(y_t|s_t) = p_t(y_t^*|s_t, y_t^o) \cdot p_t(y_t^o|s_t),$$

where $p_t(y_t^o|s_t)$ can be directly estimated from the data since both y_t^o and s_t are observables. Following our identification strategies given partial unavailability of choices, one may use the observed state transition rules (i.e., $f_{s_{t+1}|s_t, y_t^o}$) to test if other dimensions of the choices (e.g., the “intensity” margin of effort choices) are relevant for the empirical application before going all the way to the full structural estimation. Intuitively, if the moments of s_{t+1} conditional on s_t and y_t^o do not vary across periods (in a finite-horizon model), it suggests that unobserved dimensions of the choices might not play an important role in the agent’s problem.

6.5 Dynamic Discrete Games

In the baseline model and extensions discussed in Sections 6.1–6.4, we 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 describe a modeling framework of dynamic discrete games of incomplete information and then provide identification results for conditional choice probabilities and state transition rules when players’ choices are unobserved by econometricians.

Consider a game with I players, where $i = 1, 2, \dots, I$ is the index of each individual. Players choose an action from the choice set \mathcal{Y} simultaneously at each period $t =$

$1, 2, \dots, \infty$. We use y_{it} to represent player i 's action at t , so the action profile is denoted by $\mathbf{y}_t = (y_{1t}, y_{2t}, \dots, y_{It}) \in \mathcal{Y}^I$. 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 the utility shock that is privately observed by player i (not by i 's rivals or econometricians). Let $\mathbf{s}_t = (s_{1t}, s_{2t}, \dots, s_{It}) \in \times_{i=1}^I \mathcal{S}_i$ and $\boldsymbol{\varepsilon}_t = (\varepsilon_{1t}, \varepsilon_{2t}, \dots, \varepsilon_{It}) \in \times_{i=1}^I \mathcal{E}_i$ be the vector of observed states and private utility shocks at t , respectively.

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 $u(\mathbf{s}_t, \varepsilon_{it}, \mathbf{y}_t)$ 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., $E[\sum_{\tau=0}^{T-t} \beta^\tau u(\mathbf{s}_{t+\tau}, \varepsilon_{i,t+\tau}, \mathbf{y}_{t+\tau})]$, 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 the following assumption to restrict attention to certain classes of models.

Assumption 8. *The dynamic process of $\{\mathbf{s}_t, \boldsymbol{\varepsilon}_t, \mathbf{y}_t\}$ satisfies the following conditions.*

- (i) *First-order Markov:* $f_{\mathbf{s}_{t+1}, \boldsymbol{\varepsilon}_{t+1}, \mathbf{y}_{t+1} | \mathbf{s}_t, \boldsymbol{\varepsilon}_t, \mathbf{y}_t, \boldsymbol{\Omega}_{<t}} = f_{\mathbf{s}_{t+1}, \boldsymbol{\varepsilon}_{t+1}, \mathbf{y}_{t+1} | \mathbf{s}_t, \boldsymbol{\varepsilon}_t, \mathbf{y}_t}$,
where $\boldsymbol{\Omega}_{<t} \equiv \{\mathbf{s}_{t-1}, \dots, \mathbf{s}_1, \boldsymbol{\varepsilon}_{t-1}, \dots, \boldsymbol{\varepsilon}_1, \mathbf{y}_{t-1}, \dots, \mathbf{y}_1\}$.
- (ii) ε_{it} 's are independently distributed over time and across players, and are drawn from a distribution $F_i(\cdot | \mathbf{s}_t)$.
- (iii) The distribution of \mathbf{s}_{t+1} given $(\mathbf{s}_t, \boldsymbol{\varepsilon}_t, \mathbf{y}_t)$ only depends on $(\mathbf{s}_t, \mathbf{y}_t)$ and is denoted by $f_{\mathbf{s}_{t+1} | \mathbf{s}_t, \mathbf{y}_t}$.

Though typically invoked in the literature of dynamic discrete games of incomplete information, Assumption 8 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 8(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 8(iii) requires that the transition process of observed state variables does not depend on 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 $a_i(\mathbf{s}_t, \varepsilon_{it})$ and i 's belief that \mathbf{y}_t is chosen at state \mathbf{s}_t as $\sigma_i(\mathbf{y}_t|\mathbf{s}_t)$. Under Assumption 8, the value function for player i given belief σ_i is

$$V_i(\mathbf{s}_t, \varepsilon_{it}; \sigma_i) = \max_{y \in \mathcal{Y}} \sum_{\mathbf{y}_{-i} \in \mathcal{Y}^{I-1}} \sigma_i(\mathbf{y}_{-i}|\mathbf{s}_t) \left[u(\mathbf{s}_t, \varepsilon_{it}, (y, \mathbf{y}_{-i})) + \beta \mathbb{E}[V_i(\mathbf{s}_{t+1}, \varepsilon_{i,t+1}; \sigma_i) | \mathbf{s}_t, (y, \mathbf{y}_{-i})] \right], \quad (6.6)$$

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

$$a_i(\mathbf{s}_t, \varepsilon_{it}; \sigma_i) = \arg \max_{y \in \mathcal{Y}} V_i(\mathbf{s}_t, \varepsilon_{it}; \sigma_i). \quad (6.7)$$

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

$$p_i(y_{it}|\mathbf{s}_t; \sigma_i) = \int \mathbf{1}\{y_{it} = a_i(\mathbf{s}_t, \varepsilon_{it}; \sigma_i)\} dF_i(\varepsilon_{it}|\mathbf{s}_t). \quad (6.8)$$

In an 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.³⁰ In this paper, our goal is to analyze situations when players' actions are unobserved by econometricians, so we focus on the simplest case where the same equilibrium is played in the data.^{31,32}

We define player i 's equilibrium choice probabilities conditional on \mathbf{s}_t as $p_i^*(y_{it}|\mathbf{s}_t)$. When agents' actions are observed by econometricians, following the two-step methods originally developed by Hotz and Miller (1993), we can estimate the conditional choice probabilities $p_i^*(y_{it}|\mathbf{s}_t)$ and state transition rules $f_{\mathbf{s}_{t+1}|\mathbf{s}_t, \mathbf{y}_t}$ from the data in the first step. Different approaches have been developed in the literature to estimate structural parameters of the game along this direction.³³ However, when the actions are unobserved by researchers, the existing methods no longer work. In this paper, we invoke the following assumption to achieve the

³⁰Doraszelski and Satterthwaite (2010) provide 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.

³¹Otsu, 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.

³²Luo, Xiao, and Xiao (2018) and Aguirregabiria and Mira (2019) provide nonparametric identification results for dynamic discrete games of incomplete information when multiple equilibria and unobserved heterogeneity are present.

³³see Jofre-Bonet and Pesendorfer, 2003; Aguirregabiria and Mira, 2007; Bajari, Benkard, and Levin, 2007; Pakes, Ostrovsky, and Berry, 2007; Pesendorfer and Schmidt-Dengler, 2008).

identification of structural parameters in dynamic discrete games with unobserved actions.

Assumption 9. *Conditional on the current values of players' own actions and states, the future states are independent across players, i.e.,*

$$f_{s_{t+1}|s_t, y_t}(s'|s, y) = \prod_{i=1}^I f_{s_{i,t+1}|s_{it}, y_{it}}(s'_i|s_i, y_i),$$

where $s' = (s'_1, s'_2, \dots, s'_I)$, $s = (s_1, s_2, \dots, s_I)$, and $y = (y_1, y_2, \dots, y_I)$.

In general, Assumption 9 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 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 states only depends on the firm's own decisions, not on the other player's choices.³⁴ Under Assumption 9, we achieve the following equation for i 's state transition process:

$$f_{s_{i,t+1}|s_t}(s'_i|s) = \sum_{y_{it}^* \in \mathcal{Y}} f_{s_{i,t+1}|s_{it}, y_{it}^*}(s'_i|s_i, y_{it}^*) p_i^*(y_{it}^*|s), \quad (6.9)$$

where y_{it}^* is used to represent player i 's unobserved choice at period t . It is highlighted in Equation (6.9) that the transition process of $s_{i,t+1}$ does not depend on $s_{-i,t}$; 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, s_{-i} can be used as an excluded variable. For two values of s_{-i} , \bar{s}_{-i} and \hat{s}_{-i} , we obtain the following two equations under Assumption 9.

$$\begin{aligned} f_{s_{i,t+1}|s_t}(s'_i|s_i, \bar{s}_{-i}) &= \sum_{y_{it}^* \in \mathcal{Y}} f_{s_{i,t+1}|s_{it}, y_{it}^*}(s'_i|s_i, y_{it}^*) p_i^*(y_{it}^*|s_i, \bar{s}_{-i}), \\ f_{s_{i,t+1}|s_t}(s'_i|s_i, \hat{s}_{-i}) &= \sum_{y_{it}^* \in \mathcal{Y}} f_{s_{i,t+1}|s_{it}, y_{it}^*}(s'_i|s_i, y_{it}^*) p_i^*(y_{it}^*|s_i, \hat{s}_{-i}), \end{aligned} \quad (6.10)$$

From Equation (6.10), it is clear that the variations in the moments of player i 's state distribution conditional on other players' last-period states (i.e., $f_{s_{i,t+1}|s_t}$) 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 $i = 1, 2, \dots, I$. We therefore omit the details here.

³⁴Ryan (2012) estimates a dynamic model of oligopoly to study the cost of environmental regulations on firms' entry, exit, and investment decisions. In that 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.

Remark 4. *Our identification results do not require all state variables to satisfy Assumption 9. In cases with multiple dimensions of the state variable, as long as there exists one state variable of which the transition process does not involve other players’ actions or states, the equilibrium choice probabilities are identified.*

7 Empirical Illustration: Moral Hazard in U.S. Gubernatorial Elections

In this section, we apply our identification and estimation methods to a publicly available dataset containing all gubernatorial elections between 1950 and 2000 in the United States.³⁵ In the dataset, we observe the characteristics of the elected governors and a few policy outcome variables (e.g., log of per capita spending, unemployment rate, etc.) at the state level. However, politicians’ choices are not part of the data.

We estimate a dynamic structural model of politicians’ effort-exerting decisions, which are not observed by econometricians. This exercise is related to the 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 welfare (Besley and Case (1995), Alt, Bueno de Mesquita, and Rose (2011), Sieg and Yoon (2017), Aruoba, Drazen, and Vlaicu (2019), etc.). In this literature, political agency models are typically considered. For example, Alt, Bueno de Mesquita, and Rose (2011) considers an infinitely repeated game between a politician and a representative voter. The politician chooses a level of effort at each period; the voter observes the policy outcome but not the politician’s level of effort. 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.

Following the literature, in this application, we assume that the effort level exerted by the politician is not observed by the voters from a modeling perspective.³⁶ Different from papers cited above which assume that governors make one decision for each term, we focus on politicians’ within-term dynamic effort-exerting decisions. Our empirical results could potentially shed light on political business cycle (see Drazen (2000) for a comprehensive

³⁵Data source: <https://dataverse.harvard.edu/dataset.xhtml?persistentId=hdl:1902.1/14838>. For more discussions on this dataset, see Alt, Bueno de Mesquita, and Rose (2011).

³⁶It is also possible that the politician and the voters play a perfect information game, but the econometricians cannot observe the choices made by the politician. We thank an anonymous referee for pointing this out.

survey on this literature).

During the sample period, different states may have adopted different term limits and the rules could also change over time.³⁷ 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.³⁸ 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 1. In our sample, about 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 across states is around 10%, and 71% were Democratic governors.

Table 1: 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 $t = 0, 1, \dots, 4$ are provided in the lower panel of Table 1. We impose Assumption 3 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^* . Let $y_t^* = 1$ if the governor exerts effort, and 0 otherwise. 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

³⁷Detailed information about gubernatorial term limits can be found in the Book of the States.

³⁸In Alt, Bueno de Mesquita, and Rose (2011), “lame ducks” refer to politicians who cannot run for reelection. This includes two cases: (1) the state adopted a limit of one term, or (2) the state adopted a limit of two consecutive terms and the governor was in his second term.

a governor exerting effort at t given the current state s_t has the following linear structure:

$$u^*(s_t, y_t^* = 1) = \omega s_t - \rho y_t^*. \quad (7.1)$$

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 I extreme value distribution. We also assume that the utility level when no effort is exerted equals 0, i.e., $u(s_t, 0) = 0$. In summary, the parameters to be estimated in this model include $\{\omega, \rho_0, \rho_1, \rho_2, \rho_3, m_0(\cdot), m_1(\cdot), f_\eta(\cdot)\}$, where the last three are unknown functions.

We estimate the model primitives following the sieve maximum likelihood estimation strategy proposed in Section 4. The point estimates and their standard errors are provided in Table 2.³⁹ 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. In Figure 7, we plot m_0 and m_1 using model estimates in the range of the state variable observed in the data. 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 is higher compared to the second-term lame ducks, but the difference is not statistically significant. This finding 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 3. 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 13 percent higher than that of the first period. This result is quite intuitive: governors have fewer incentives to exert effort when they are approaching the end of the term. Overall, we observe a lower chance of exerting effort for first-term governors. Republican governors and those who have higher proportions of elderly people seem to have higher probabilities of exerting effort.

³⁹Similar 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, $m_0(s) \approx a_0 + a_1s + a_2s^2 + a_3s^3$, and $m_1(s) \approx b_0 + b_1s + b_2s^2 + b_3s^3$.

Table 2: Estimation Results

Panel (A) Estimates of m_0 and m_1			Panel (B) Estimates of Utility Primitives		
Parameters	Estimates	Std. Err.	Parameters	Estimates	Std. Err.
$m_0 : a_0$	1.0716	0.8072	ω	3.7953	1.3153
$m_0 : a_1$	0.8181	0.0259	ρ_0	21.4115	8.5479
$m_0 : a_2$	0.0029	0.0637	ρ_1	0.1816	1.0461
$m_0 : a_3$	-0.0005	0.0068	ρ_2	0.0003	17.4481
$m_1 : b_0$	1.1603	0.1724	ρ_3	0.4817	0.7990
$m_1 : b_1$	0.6802	0.0263	σ_η	0.0513	0.0015
$m_1 : b_2$	0.0283	0.0128			
$m_1 : b_3$	-0.0008	0.0012			




Figure 7: Plot m_0 and m_1 Using Model Estimates

Table 3: Probabilities of Shirking at Each Period

	Year 1	Year 2	Year 3	Year 4
All sample	0.0785	0.0769	0.0766	0.0884
By Category				
First-term lame duck	0.1346	0.1308	0.1307	0.1491
Second-term lame duck	0.0120	0.0130	0.0125	0.0164
Democratic governor	0.0997	0.0978	0.0977	0.1109
Republican governor	0.0245	0.0236	0.0228	0.0310
Lower percent of elderly	0.1233	0.1193	0.1181	0.1353
Higher percent of elderly	0.0098	0.0117	0.0130	0.0163

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 variations in the 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. Our identification results extend to infinite-horizon models, models with serially correlated unobserved heterogeneity, cases in which state variables are discrete or choices are partially unavailable, and dynamic discrete games. We propose a sieve maximum likelihood estimator for primitives in agents' utility functions and state transition rules. Monte Carlo simulations under various specifications demonstrate the good performance of the proposed approach. The identification and estimation results developed in this paper contribute to the body of our knowledge. Under mild assumptions on the state transition process, our methods can be applied to various empirical contexts in labor economics, industrial organization, political economy, and other related fields, when agents' choice data are not readily available.

A Proof of Theorem 1

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,

$$\mu_{t+1} = \sum_{y_t^* \in \{0,1\}} p_t(y_t^*|s) E_{t+1}[m(y_t^*, s) + \eta_t|s, y_t^*] = p_t m_1 + (1 - p_t)m_0, \quad (\text{A.1})$$

where the second equality holds because under Assumption 3, η_t and y_t^* are independent conditional on the state and $E(\eta_t|s) = 0$. In Equation (A.1), 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.

$$\begin{aligned} \nu_{t+1} &= \sum_{y_t^* \in \{0,1\}} p_t(y_t^*|s) E_{t+1}[(m(y_t^*, s) + \eta_t - \mu_{t+1})^2|s, y_t^*] \\ &= \sum_{y_t^* \in \{0,1\}} p_t(y_t^*|s) \left[(m(y_t^*, s) - \mu_{t+1})^2 + 2(m(y_t^*, s) - \mu_{t+1}) E[\eta_t|s] + E[\eta_t^2|s] \right] \\ &= p_t(m_1 - \mu_{t+1})^2 + (1 - p_t)(m_0 - \mu_{t+1})^2 + E[\eta_t^2|s] \end{aligned} \quad (\text{A.2})$$

$$\begin{aligned} \xi_{t+1} &= \sum_{y_t^* \in \{0,1\}} p_t(y_t^*|s) E_{t+1}[(m(y_t^*, s) + \eta_t - \mu_{t+1})^3|s, y_t^*] \\ &= \sum_{y_t^* \in \{0,1\}} p_t(y_t^*|s) \left[(m(y_t^*, s) - \mu_{t+1})^3 + E[\eta_t^3|s] + 3(m(y_t^*, s) - \mu_{t+1})^2 E[\eta_t|s] \right. \\ &\quad \left. + 3(m(y_t^*, s) - \mu_{t+1}) E[\eta_t^2|s] \right] \\ &= p_t(m_1 - \mu_{t+1})^3 + (1 - p_t)(m_0 - \mu_{t+1})^3 + E[\eta_t^3|s]. \end{aligned} \quad (\text{A.3})$$

In Equations (A.2) and (A.3), $E[\eta_t^2|s]$ and $E[\eta_t^3|s]$ are the second- and third-order conditional moments of the error η_t , 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 ($t \neq \tau$) along the dynamic process. Equation (A.1) identifies the choice probability for any given m_0 and m_1 as long as $m_0 \neq m_1$. Specifically, the choice probabilities at period t and τ are

$$p_t = \frac{\mu_{t+1} - m_0}{m_1 - m_0}, \quad p_\tau = \frac{\mu_{\tau+1} - m_0}{m_1 - m_0}. \quad (\text{A.4})$$

Under Assumption 1(iii) and Assumption 3, the conditional distribution of η_t is stationary. This implies that the higher order moments of the error term are time-invariant conditional on the same state s , i.e., $E[\eta_t^2|s] = E[\eta_\tau^2|s]$, and $E[\eta_t^3|s] = E[\eta_\tau^3|s]$. Taking the

difference of Equations (A.2) and (A.3) across the two periods t and τ , we eliminate the unknown moments of η_t and achieve the following two equations.

$$\begin{aligned} \nu_{t+1} - \nu_{\tau+1} &= p_t(m_1 - \mu_{t+1})^2 + (1 - p_t)(m_0 - \mu_{t+1})^2 - p_\tau(m_1 - \mu_{\tau+1})^2 - (1 - p_\tau)(m_0 - \mu_{\tau+1})^2 \quad (\text{A.5}) \\ &= (p_t - p_\tau)(m_1 + m_0)(m_1 - m_0) - (\mu_{t+1}^2 - \mu_{\tau+1}^2), \end{aligned}$$

$$\begin{aligned} \xi_{t+1} - \xi_{\tau+1} &= p_t(m_1 - \mu_{t+1})^3 + (1 - p_t)(m_0 - \mu_{t+1})^3 - p_\tau(m_1 - \mu_{\tau+1})^3 - (1 - p_\tau)(m_0 - \mu_{\tau+1})^3 \quad (\text{A.6}) \end{aligned}$$

We further plug the expressions of p_t and p_τ in Equation (A.4) into Equations (A.5) and (A.6), which leads to a system of equations for the unknown primitives m_1 and m_0 :

$$\nu_{t+1} - \nu_{\tau+1} = (\mu_{t+1} - \mu_{\tau+1})\Delta_1 - (\mu_{t+1}^2 - \mu_{\tau+1}^2), \quad (\text{A.7})$$

$$\xi_{t+1} - \xi_{\tau+1} = (\mu_{t+1}\Delta_1 - \Delta_2 - \mu_{t+1}^2)(\Delta_1 - 2\mu_{t+1}) - (\mu_{\tau+1}\Delta_1 - \Delta_2 - \mu_{\tau+1}^2)(\Delta_1 - 2\mu_{\tau+1}), \quad (\text{A.8})$$

where $\Delta_1 = m_1 + m_0$ and $\Delta_2 = m_1 m_0$. We obtain analytical solutions for Δ_1 and Δ_2 from Equations (A.7)–(A.8), provided that $\mu_{t+1} \neq \mu_{\tau+1}$.

$$\Delta_1 = \frac{\nu_{t+1} - \nu_{\tau+1} + (\mu_{t+1}^2 - \mu_{\tau+1}^2)}{\mu_{t+1} - \mu_{\tau+1}}, \quad (\text{A.9})$$

$$\Delta_2 = \frac{\xi_{t+1} - \xi_{\tau+1} - (\mu_{t+1}(\Delta_1 - \mu_{t+1})(\Delta_1 - 2\mu_{t+1}) - \mu_{\tau+1}(\Delta_1 - \mu_{\tau+1})(\Delta_1 - 2\mu_{\tau+1}))}{2(\mu_{t+1} - \mu_{\tau+1})}. \quad (\text{A.10})$$

We now focus on the case where $\mu_{t+1} \neq \mu_{\tau+1}$. With Δ_1 and Δ_2 identified using the moments of the observed state transition process as shown in Equations (A.9)–(A.10), m_0 and m_1 are the two distinctive solutions to the equation of \bar{m}

$$\bar{m}^2 - \Delta_1 \bar{m} + \Delta_2 = 0, \quad (\text{A.11})$$

when $\Delta_1^2 - 4\Delta_2 > 0$. Without loss of generality, we label the larger root of Equation (A.11) by m_1 and the smaller one by m_0 , i.e.,

$$m_1 = \frac{\Delta_1 + \sqrt{\Delta_1^2 - 4\Delta_2}}{2}, \quad m_0 = \frac{\Delta_1 - \sqrt{\Delta_1^2 - 4\Delta_2}}{2}.$$

Once m_1 and m_0 are identified (and they are not equal), the conditional choice probabilities p_t and p_τ are identified from Equation (A.4). 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 $s' - m_0$:

$$\begin{aligned} f_{s_{t+1}|s_t}(s'|s) &= p_t f_{\eta_t|s_t}(s' - m_1|s) + (1 - p_t) f_{\eta_t|s_t}(s' - m_0|s), \\ f_{s_{\tau+1}|s_\tau}(s'|s) &= p_\tau f_{\eta_\tau|s_\tau}(s' - m_1|s) + (1 - p_\tau) f_{\eta_\tau|s_\tau}(s' - m_0|s). \end{aligned} \quad (\text{A.12})$$

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 (A.12) identifies the conditional density function of η_t at $s' - m_1$ and $s' - m_0$ when $p_t \neq p_\tau$, i.e.,

$$\begin{aligned} f_{\eta_t|s_t}(s' - m_1|s) &= \frac{f_{s_{t+1}|s_t}(s'|s)(1 - p_\tau) - f_{s_{\tau+1}|s_\tau}(s'|s)(1 - p_t)}{p_t - p_\tau}, \\ f_{\eta_t|s_t}(s' - m_0|s) &= \frac{f_{s_{\tau+1}|s_\tau}(s'|s)p_t - f_{s_{t+1}|s_t}(s'|s)p_\tau}{p_t - p_\tau}. \end{aligned} \quad (\text{A.13})$$

So far, we have proved that for a state s at which $\mu_{t+1} \neq \mu_{\tau+1}$ and the order of the m functions is known, $m(0, s)$, $m(1, s)$, choice probabilities p_t and p_τ , and the conditional density function of the error term η_t are identified. Assumption 4(i) ensures that there are at most countable many values of the state variable at which $\mu_{t+1} = \mu_{\tau+1}$. Therefore, by the continuity of m functions, we also identify $m(1, s)$ and $m(0)$ for $s \in \tilde{\mathcal{S}}$. Assumption 4(iii) guarantees that there exists a state s_0 such that we can distinguish between $m(0, s_0)$, $m(1, s_0)$. Starting from s_0 , by the continuity of the first-order derivatives of the m functions imposed in Assumption 3 and the condition that $m(1, \cdot)$ and $m(0, \cdot)$ intersect at only countable many states with different derivatives imposed in Assumption 4(i)–(ii), we match $m(1, \cdot)$ and $m(0, \cdot)$ across all values of $s \in \mathcal{S}$.

Equation (A.13) identifies $f_{\eta_t|s_t}(\cdot|s)$ for any $s \in \mathcal{S}$ such that $p_t(\cdot|s) \neq p_\tau(\cdot|s)$. Given that $\{s \in \mathcal{S} : p_t(\cdot|s) = p_\tau(\cdot|s)\}$ has zero probability and the continuity of the density function, $f_{\eta_t|s_t}(\cdot|s)$ is identified for all $s \in \mathcal{S}$. The identification of m functions and the conditional density function of the error term η_t implies the identification of the state transition rules, i.e., $f_{s_{t+1}|s_t, y_t^*}(\cdot|\cdot, \cdot)$. Choice probabilities for periods t and τ are identified using Equation (A.4) for all $s \in \mathcal{S}/\bar{\mathcal{S}}$ with $Pr(\bar{\mathcal{S}}) = 0$. This completes the proof of Theorem 1. \square

B Sieve Maximum Likelihood Estimation

Let $\theta^0 = (\alpha^0, m_0^0, m_1^0, f_{\eta_t|s_t}^0)$ represent the vector of true parameter values of interest. For simplicity, consider a case where the per-period utility functions are parametrized by a finite-dimensional parameter $\alpha^0 \in \mathcal{A}$. $m_0^0 : \mathcal{S} \rightarrow \mathcal{S}$ and $m_1^0 : \mathcal{S} \rightarrow \mathcal{S}$ are two nonparametric functions in the state transition rules, where \mathcal{S} denotes the state space. $f_{\eta_t|s_t}^0 : \mathbb{R} \times \mathcal{S} \rightarrow \mathbb{R}^+$ is the probability density function of the error term conditional on the state variable.

We impose the following smoothness restrictions on m_0^0 , m_1^0 , and the density function $f_{\eta_t|s_t}^0$. 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 $m : \Psi \rightarrow \mathbb{R}$ 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 $\Gamma_c^{\kappa, \zeta}(\Psi) \equiv \left\{ m \in \Gamma^{\kappa, \zeta}(\Psi) \mid \|m\|_{\Gamma^{\kappa, \zeta}(\Psi)} \leq c < \infty \right\}$ 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 $r = (r_1, r_2, \dots, r_{\dim(\Psi)})$ and $|r| = r_1 + r_2 + \dots + r_{\dim(\Psi)}$. With 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

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

with norm defined by $\|f\|_{\mathcal{F}} = \sup_{x \in \mathbb{R} \times \mathcal{S}} |f(x)(1 + \|x\|_e^2)^{-\psi/2}|$, $\psi > 0$. Notice that the conditional mean of η_t for all density functions in \mathcal{F} are equal to 0, which is consistent with Assumption 3. $\Theta = \mathcal{A} \times \mathcal{H} \times \mathcal{H} \times \mathcal{F}$ denotes 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

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

In light of a finite sample, instead of searching parameters over an infinite-dimensional

parameter space Θ , we maximize the empirical criterion function in Equation (4.3) 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

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

In the definition of sieve spaces, $(h_1(\cdot), h_2(\cdot), h_3(\cdot), \dots)$ represents a sequence of known basis functions, such as Hermite polynomials, power series, splines, etc. We use linear sieves to approximate the square root of densities and $\mathbf{g}^{k_3}(\cdot, \cdot)$ is a $k_3 \times 1$ vector of the 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 a tensor product of linear sieves of univariate functions.

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) and Carroll, Chen, and Hu (2010), we provide sufficient conditions tailored to our model for consistency of the sieve maximum likelihood estimator in Equation (4.4).⁴⁰

Assumption 10 (Consistency). *The following conditions are satisfied: (i) all assumptions in Theorem 1 hold; (ii) D_i is i.i.d. across i ; (iii) m_0 and $m_1 \in \mathcal{H}$ with $\kappa_1 + \zeta_1 > 1/2$; $f_{\eta|S} \in \mathcal{F}$ with $\kappa_2 + \zeta_2 > 1$; (iv) $|Q(\theta^0)| < \infty$ and $Q(\theta)$ is upper semicontinuous on Θ under the metric $d(\cdot, \cdot)$; (v) the sieve spaces, Θ_k , are compact under $d(\cdot, \cdot)$; (vi) There is a finite $\sigma > 0$ and a random variable $c(D_i)$ with $E(c(D_i)) < \infty$ such that $\sup_{\theta \in \Theta_k : d(\theta, \theta^0) \leq \epsilon} |l(D_i; \theta) - l(D_i; \theta^0)| \leq \epsilon^\sigma c(D_i)$; (vii) k_1, k_2 , and $k_3 \rightarrow \infty$, $k_1/n, k_2/n$, and $k_3/n \rightarrow 0$.*

Assumption 10(i) ensures the identification of the model primitives. Assumption 10 overall provides a set of assumptions that imply the 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). *If Assumption 10 is satisfied, then the sieve maximum likelihood estimator in Equation (4.4) is consistent with respect to the metric $d(\cdot, \cdot)$, i.e.,*

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

⁴⁰Chen, Hu, and Lewbel (2008) study identification and estimation of a nonparametric regression model with discrete covariates measured with error. Carroll, Chen, and Hu (2010) consider a general nonlinear errors-in-variables model using two samples.

Remark 5. 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 \(1999\)](#), and [Chen \(2007; Theorem 3.2 and Theorem 4.3\)](#).

C Derivation of Equation (6.2)

We decompose the joint distribution of the observed state variable at four consecutive periods $(s_{t+2}, s_{t+1}, s_t, s_{t-1})$ as follows.

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

The second equality in Equation (C.1) 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 (C.1).

D Eigenvalue-Eigenvector Decomposition

In this section, we provide details for constructing eigenvalue-eigenvector decompositions based on Equation (6.4) in Section 6.3. Let $j_s = 1, \dots, J_s$, $j_z = 1, \dots, J_z$, and $j_y = 1, \dots, J_y$ index the categories of s_t , z_t and y_t^* , respectively. The agent's choice is not limited to taking binary values, i.e., $J_y \geq 2$. For simplicity, we consider the case where $J_s = J_z = J_y$.⁴¹ We define the following matrices for fixed (s, z) :

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

Equation (6.4) in matrix form is therefore

$$M_{s_{t+1}, z_{t+1}, s, z} = M_{s_{t+1}|s, y_t^*} M_{y_t^*, s, z} M_{z_{t+1}|z, y_t^*}. \quad (\text{D.1})$$

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 equation:

$$\begin{aligned} & \left(M_{s_{t+1}, z_{t+1}, \bar{s}, \bar{z}} \cdot M_{s_{t+1}, z_{t+1}, \hat{s}, \bar{z}}^{-1} \right) \left(M_{s_{t+1}, z_{t+1}, \bar{s}, \hat{z}} \cdot M_{s_{t+1}, z_{t+1}, \hat{s}, \hat{z}}^{-1} \right)^{-1} \\ &= M_{s_{t+1}|\bar{s}, y_t^*} \left(M_{y_t^*, \bar{s}, \bar{z}}^{-1} M_{y_t^*, \hat{s}, \bar{z}} M_{y_t^*, \bar{s}, \hat{z}}^{-1} M_{y_t^*, \hat{s}, \hat{z}} \right) M_{s_{t+1}|\bar{s}, y_t^*}^{-1} \\ &\equiv M D M^{-1}, \end{aligned} \quad (\text{D.2})$$

provided that the following assumption holds.

Assumption 11 (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 in 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 invertibility of $M_{y_t^*, s, z}$ is guaranteed. Under Assumption 11, Equation (D.2) leads to an eigenvalue-eigenvector decomposition of the observed matrix on its left-hand side, where M represents the matrix of eigenvectors and the diagonal elements in D are the

⁴¹When the number of possible values of the state variables is larger than the number of possible alternatives, we can regroup the state variables to obtain $J_s = J_z = J_y$.

corresponding eigenvalues. Additional assumptions are required to guarantee the uniqueness of the decomposition and to pin down the ordering of the eigenvectors.

Assumption 12 (Uniqueness). *Let $D(i)$ denote the i -th diagonal element in D . $D(i) \neq D(j)$ for any $i \neq j$.*

Assumption 13 (Ordering). *Suppose $y_t^* \in \{0, 1\}$. $E(s_{t+1}|\bar{s}, y_t^* = 1) > E(s_{t+1}|\bar{s}, y_t^* = 0)$.*

Assumption 12 rules out the possibility of duplicated eigenvalues. Assumption 13 imposes restrictions on the state transition process given different choices to determine which eigenvector corresponds to $y_t^* = 1$. The economic intuition of this assumption can be illustrated again using the borrower's example in Section 3. Suppose $y_t^* = 1$ represents the case where the borrower exerts effort to pay off his debts, and 0 otherwise; s_t represents the revenue at period t . Assumption 13 implies that the expected return given the borrower exerting effort is higher than the one when the borrower shirks. We summarize the identification results for the unknown densities in M in the following theorem.

Theorem 3 (Identification). *Suppose Assumptions 1, 7, 11, 12, and 13 hold for the Markov process of $\{s_t, z_t, \varepsilon_t, y_t^*\}$. The joint distribution of $\{s_{t+1}, z_{t+1}, s_t, z_t\}$ uniquely determines the state transition rules $f_{s_{t+1}|\bar{s}, y_t^*}$.*

Similar to Equation (D.2), we can derive the eigenvalue-eigenvector decompositions to identify matrices $M_{s_{t+1}|\hat{s}, y_t^*}$, $M_{z_{t+1}|\bar{z}, y_t^*}$, and $M_{z_{t+1}|\hat{z}, y_t^*}$, which essentially lead to the identification of the state transition rules $f_{s_{t+1}|s_t, y_t^*}$ and $f_{z_{t+1}|z_t, y_t^*}$. $M_{y_t^*, s, z}$ for any (s, z) are therefore identified from Equation (D.1); the diagonal elements of these matrices correspond to the unobserved choice probabilities $f_{y_t^*|s_t, z_t}$.

When two continuous state variables are available, we can generalize our identification results to allow for continuous choice variables. Instead of using eigenvalue-eigenvector decompositions, spectrum decompositions proposed by Hu and Schennach (2008) are applied.

E Monte Carlo Simulations: Tables

Table 4: Monte Carlo Simulation Results: DGP 1, N=5000

	TRUE	MC Mean	MC Bias	MC Std	MAE	RMSE
Panel (A) Unobserved Choices						
$m_0 : a_0$	0.0000	0.0208	0.0208	0.0020	0.0208	0.0209
$m_0 : a_1$	0.8000	0.7494	-0.0506	0.0103	0.0506	0.0516
$m_0 : a_2$	0.0000	-0.0052	-0.0052	0.0005	0.0052	0.0052
$m_0 : a_3$	0.0000	0.0051	0.0051	0.0004	0.0051	0.0051
$m_1 : b_0$	0.5000	0.5166	0.0166	0.0275	0.0256	0.0320
$m_1 : b_1$	0.8000	0.7834	-0.0166	0.0219	0.0226	0.0274
$m_1 : b_2$	0.0000	-0.0053	-0.0053	0.0005	0.0053	0.0054
$m_1 : b_3$	0.0000	0.0022	0.0022	0.0012	0.0022	0.0025
σ_η	1.0000	1.0452	0.0452	0.0047	0.0452	0.0454
ω	0.8000	0.8367	0.0367	0.0127	0.0367	0.0388
ρ	0.3000	0.3023	0.0023	0.0241	0.0185	0.0241
Panel (B) Observed Choices						
$m_0 : a_0$	0.0000	0.0203	0.0203	0.0006	0.0203	0.0203
$m_0 : a_1$	0.8000	0.7550	-0.0450	0.0063	0.0450	0.0454
$m_0 : a_2$	0.0000	-0.0050	-0.0050	0.0002	0.0050	0.0050
$m_0 : a_3$	0.0000	0.0050	0.0050	0.0002	0.0050	0.0050
$m_1 : b_0$	0.5000	0.5422	0.0422	0.0144	0.0422	0.0446
$m_1 : b_1$	0.8000	0.7527	-0.0473	0.0059	0.0473	0.0477
$m_1 : b_2$	0.0000	-0.0053	-0.0053	0.0003	0.0053	0.0053
$m_1 : b_3$	0.0000	0.0041	0.0041	0.0003	0.0041	0.0041
σ_η	1.0000	1.0503	0.0503	0.0037	0.0503	0.0504
ω	0.8000	0.8219	0.0219	0.0043	0.0219	0.0223
ρ	0.3000	0.3173	0.0173	0.0063	0.0173	0.0184

Table 5: Monte Carlo Simulation Results: DGP 2, N=5000

	TRUE	MC Mean	MC Bias	MC Std	MAE	RMSE
Panel (A) Unobserved Choices						
$m_0 : a_0$	0.0000	0.0220	0.0220	0.0126	0.0224	0.0254
$m_0 : a_1$	0.8000	0.7985	-0.0015	0.0721	0.0350	0.0718
$m_0 : a_2$	0.0000	-0.0023	-0.0023	0.0044	0.0039	0.0050
$m_0 : a_3$	0.0000	0.0005	0.0005	0.0015	0.0010	0.0016
$m_1 : b_0$	0.5000	0.4864	-0.0136	0.0594	0.0447	0.0606
$m_1 : b_1$	1.0000	1.0165	0.0165	0.0252	0.0245	0.0300
$m_1 : b_2$	0.0000	-0.0053	-0.0053	0.0032	0.0055	0.0062
$m_1 : b_3$	0.0000	0.0004	0.0004	0.0003	0.0004	0.0005
σ_η	1.0000	1.0520	0.0520	0.0063	0.0520	0.0524
ω	0.8000	0.8341	0.0341	0.0359	0.0393	0.0494
ρ	0.3000	0.2660	-0.0340	0.0645	0.0568	0.0726
Panel (B) Observed Choices						
$m_0 : a_0$	0.0000	0.0211	0.0211	0.0106	0.0212	0.0236
$m_0 : a_1$	0.8000	0.7181	-0.0819	0.0677	0.0921	0.1061
$m_0 : a_2$	0.0000	-0.0058	-0.0058	0.0030	0.0062	0.0065
$m_0 : a_3$	0.0000	0.0034	0.0034	0.0018	0.0034	0.0038
$m_1 : b_0$	0.5000	0.5211	0.0211	0.0962	0.0780	0.0980
$m_1 : b_1$	1.0000	0.9643	-0.0357	0.0369	0.0427	0.0513
$m_1 : b_2$	0.0000	-0.0051	-0.0051	0.0018	0.0051	0.0054
$m_1 : b_3$	0.0000	0.0019	0.0019	0.0013	0.0019	0.0023
σ_η	1.0000	1.0901	0.0901	0.0457	0.0901	0.1009
ω	0.8000	0.8245	0.0245	0.0349	0.0314	0.0425
ρ	0.3000	0.3195	0.0195	0.0678	0.0461	0.0702

Table 6: Monte Carlo Simulation Results: DGP 3, N=5000

	TRUE	MC Mean	MC Bias	MC Std	MAE	RMSE
Panel (A) Unobserved Choices						
$m_0 : a_0$	0.0000	0.0190	0.0190	0.0068	0.0194	0.0202
$m_0 : a_1$	0.6000	0.5680	-0.0320	0.0228	0.0347	0.0392
$m_0 : a_2$	0.0500	0.0351	-0.0149	0.0072	0.0152	0.0165
$m_0 : a_3$	0.0000	0.0043	0.0043	0.0014	0.0043	0.0045
$m_1 : b_0$	0.5000	0.5807	0.0807	0.0499	0.0866	0.0948
$m_1 : b_1$	0.6000	0.5643	-0.0357	0.0583	0.0584	0.0681
$m_1 : b_2$	0.0500	0.0375	-0.0125	0.0135	0.0147	0.0184
$m_1 : b_3$	0.0000	0.0036	0.0036	0.0011	0.0036	0.0037
σ_η	1.0000	1.0404	0.0404	0.0065	0.0404	0.0409
ω	0.8000	0.8398	0.0398	0.0270	0.0399	0.0480
ρ	0.3000	0.2983	-0.0017	0.0381	0.0255	0.0379
Panel (B) Observed Choices						
$m_0 : a_0$	0.0000	0.0193	0.0193	0.0035	0.0193	0.0196
$m_0 : a_1$	0.6000	0.5789	-0.0211	0.0107	0.0214	0.0236
$m_0 : a_2$	0.0500	0.0322	-0.0178	0.0032	0.0178	0.0181
$m_0 : a_3$	0.0000	0.0050	0.0050	0.0006	0.0050	0.0050
$m_1 : b_0$	0.5000	0.5601	0.0601	0.0271	0.0609	0.0659
$m_1 : b_1$	0.6000	0.5311	-0.0689	0.0283	0.0691	0.0744
$m_1 : b_2$	0.0500	0.0481	-0.0019	0.0068	0.0056	0.0070
$m_1 : b_3$	0.0000	0.0038	0.0038	0.0006	0.0038	0.0038
σ_η	1.0000	1.0507	0.0507	0.0042	0.0507	0.0509
ω	0.8000	0.8188	0.0188	0.0124	0.0198	0.0225
ρ	0.3000	0.3013	0.0013	0.0246	0.0183	0.0245

Table 7: Monte Carlo Simulation Results: DGP 4, N=5000

	TRUE	MC Mean	MC Bias	MC Std	MAE	RMSE
Panel (A) Unobserved Choices						
$m_0 : a_0$	0.0000	0.0205	0.0205	0.0015	0.0205	0.0205
$m_0 : a_1$	0.2000	0.1902	-0.0098	0.0069	0.0104	0.0120
$m_0 : a_2$	0.1000	0.1095	0.0095	0.0074	0.0107	0.0121
$m_0 : a_3$	0.0000	0.0050	0.0050	0.0003	0.0050	0.0050
$m_1 : b_0$	0.5000	0.4703	-0.0297	0.0159	0.0303	0.0336
$m_1 : b_1$	0.2000	0.2427	0.0427	0.0084	0.0427	0.0435
$m_1 : b_2$	0.1000	0.0667	-0.0333	0.0024	0.0333	0.0333
$m_1 : b_3$	0.0000	0.0050	0.0050	0.0003	0.0050	0.0050
σ_η	1.0000	1.0512	0.0512	0.0026	0.0512	0.0513
ω	0.8000	0.8268	0.0268	0.0025	0.0268	0.0270
ρ	0.3000	0.3150	0.0150	0.0069	0.0153	0.0165
Panel (B) Observed Choices						
$m_0 : a_0$	0.0000	0.0200	0.0200	0.0010	0.0200	0.0200
$m_0 : a_1$	0.2000	0.1823	-0.0177	0.0071	0.0179	0.0191
$m_0 : a_2$	0.1000	0.0925	-0.0075	0.0086	0.0104	0.0114
$m_0 : a_3$	0.0000	0.0049	0.0049	0.0002	0.0049	0.0049
$m_1 : b_0$	0.5000	0.4922	-0.0078	0.0143	0.0132	0.0162
$m_1 : b_1$	0.2000	0.2416	0.0416	0.0109	0.0416	0.0430
$m_1 : b_2$	0.1000	0.0723	-0.0277	0.0020	0.0277	0.0278
$m_1 : b_3$	0.0000	0.0052	0.0052	0.0003	0.0052	0.0052
σ_η	1.0000	1.0488	0.0488	0.0040	0.0488	0.0490
ω	0.8000	0.8227	0.0227	0.0042	0.0227	0.0231
ρ	0.3000	0.3096	0.0096	0.0064	0.0104	0.0115

References

- Abbring, Jaap H. 2010. "Identification of Dynamic Discrete Choice Models." *Annual Review of Economics*, 2(1): 367–394.
- Abbring, Jaap H., and Øystein Daljord. 2019. "Identifying the Discount Factor in Dynamic Discrete Choice Models." *Quantitative Economics*, 11(2): 471–501.
- 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. 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 Arvind Magesan. 2020. "Identification and Estimation of Dynamic Games when Players' Beliefs are Not In Equilibrium." *The Review of Economic Studies*, 87(2): 582–625.
- 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, and Pedro Mira. 2010. "Dynamic Discrete Choice Structural Models: A Survey." *Journal of Econometrics*, 156(1): 38–67.
- Aguirregabiria, Victor, and Pedro Mira. 2019. "Identification of Games of Incomplete Information with Multiple Equilibria and Unobserved Heterogeneity." *Quantitative Economics*, 10(4): 1659–1701.
- 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.
- Arcidiacono, Peter, and Paul B. Ellickson. 2011. “Practical Methods for Estimation of Dynamic Discrete Choice Models.” *Annual Review of Economics*, 3(1): 363–394.
- Arcidiacono, Peter, and Robert A. Miller. 2011. “Conditional Choice probability Estimation of Dynamic Discrete Choice Models with Unobserved Heterogeneity.” *Econometrica*, 79(6): 1823–1867.
- Arcidiacono, Peter, and Robert A. Miller. 2020. “Identifying Dynamic Discrete Choice Models off Short Panels.” *Journal of Econometrics*, 215(2): 473–485.
- Aruoba, S. Boragan, Allan Drazen, and Razvan Vlaicu. 2019. “A Structural Model of Electoral Accountability.” *International Economic Review*, Forthcoming.
- Bajari, Patrick, Chenghuan Sean Chu, Denis Nekipelov, and Minjung Park. 2016. “Identification and Semiparametric Estimation of a Finite Horizon Dynamic Discrete Choice Model with a Terminating Action.” *Quantitative Marketing and Economics*, 14(4): 271–323.
- 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, Xiaohong, Yingyao Hu, and Arthur Lewbel. 2009. “Nonparametric Identification and Estimation of Nonclassical Errors-in-variables Models without Additional Information.” *Statistica Sinica*, 19(3): 949–968.

- 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.
- Eckstein, Zvi, and Kenneth I. Wolpin. 1989. "The Specification and Estimation of Dynamic Stochastic Discrete Choice Models: A Survey." *The Journal of Human Resources*, 24(4): 562–598.
- Gayle, George-Levi, and Robert A. Miller. 2015. "Identifying and Testing Models of Managerial Compensation." *The Review of Economic Studies*, 82(3): 1074–1118.
- Gayle, George-Levi, Limor Golan, and Robert A. Miller. 2015. "Promotion, Turnover, and Compensation in the Executive Labor Market." *Econometrica*, 83(6): 2293–2369.
- 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 Non-classical 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 Across 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*.