# Essays on Semiparametric Estimation of Markov Decision Processes

Sorawoot Srisuma

Thesis submitted for the degree of

Doctor of Philosophy in Economics at the

London School of Economics and Political Science

## Declaration

I certify that this thesis I have presented for examination for the PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others.

Signature ..................................

Date 09/07/2010

# Abstract

Dynamic models of forward looking agents, whose goal is to maximize expected intertemporal payoffs, are useful modelling frameworks in economics. With an exception of a small class of dynamic decision processes, the estimation of the primitives in these models is computationally burdensome due to the presence of the value functions that has no closed form. We follow a popular two-step approach which estimates the functions of interest rather than use direct numerical approximation.

The first chapter, joint with Oliver Linton, considers a class of dynamic discrete choice models that contain observable continuously distributed state variables. Most papers on the estimation of dynamic discrete choice models assume that the observable state variables can only take finitely many values. We show that the extension to the infinite dimensional case leads to a well-posed inverse problem. We derive the distribution theory for the finite and the infinite dimensional parameters.

Dynamic models with continuous choice can sometimes avoid the numerical issues related to the value function through the use of Euler's equation. The second chapter considers models with continuous choice that do not necessarily belong to the Euler class but frequently arise in applied problems. In this chapter, a class of minimum distance estimators is proposed, their distribution theory along with the infinite dimensional parameters of the decision models are derived.

The third chapter demonstrates how the methodology developed for the discrete and continuous choice problems can be adapted to estimate a variety of other dynamic models.

The final chapter discusses an important problem, and provides an example, where some well-known estimation procedures in the literature may fail to consistently estimate an identified model. The estimation methodologies I propose in the preceding chapters may not suffer from the problems of this kind.

*To my parents and my wife*

## Acknowledgements

I am deeply indebted to my advisor Oliver Linton for his advice, encouragement and support throughout my PhD studies. In particular, I am very grateful for the opportunities Oliver gave me as his research assistant from an early stage. As an advisor, Oliver has been remarkably accessible and I have thoroughly enjoyed many conversations with him for his great sense of humour and insights on various topics ranging from econometrics and beyond.

I would like to express my sincere gratitude to Marcia Schafgans for her advice and support, since I was in the Mathematics Department at the LSE during my undergraduate degree, for convincing me to remain at the LSE and throughout my graduate studies as well as being her teaching assistant annually for various courses over the years. I also have benefited a great deal from various discussions with Martin Pesendorfer and, especially, Philipp Schmidt-Dengler who introduce me the field of Industrial Organization; particularly to the literature on estimation of structural dynamic models.

Finally I would like to acknowledge the financial support from the Economics Department at the LSE and the ESRC throughout my graduate studies.

*Sorawoot Srisuma*

*May 2010*

*London*

# Contents


**1 Dynamic Discrete Choice Problems with Continuous State** 8

1.1 Introduction . . . 8

1.2 Markov Decision Processes . . . 12

1.2.1 Definitions and Assumptions . . . 12

1.2.2 Policy Value Equation . . . 15

1.2.3 Linear Inverse Problems . . . 18

1.3 Estimation of Conditional Choice Probabilities . . . 20

1.4 Distribution Theory . . . 28

1.5 Simulation Study . . . 36

1.6 Markovian Games . . . 39

1.7 Conclusion . . . 44

1.8 Proofs of Theorems . . . 46

1.8.1 Outline of Asymptotic Approach . . . 46

1.8.2 Proofs of Theorems 1.1 - 1.2 . . . 53

1.8.3 Proofs of High Level Conditions A1.1 - A1.6 . . . 62

1.8.4 Proofs of Theorems 1.3 - 1.5 . . . 64

1.8.5 Additional Proofs . . . 77

1.9 Tables . . . 84

**2 Dynamic Models with Continuous Control** 88

2.1 Introduction . . . 88

2.2 Markov Decision Processes . . . 92

2.2.1 Basic Framework . . . 92

2.3 Estimation Methodology . . . 95

2.4 Practical Aspects . . . 102

2.5 Asymptotic Theory . . . . . . . . . . 102

2.6 Bootstrap Standard Errors . . . . . . . . . . 114

2.7 Simulation Study . . . . . . . . . . 116

2.8 Conclusion . . . . . . . . . . 121

2.9 Proofs of Theorems . . . . . . . . . . 123

2.9.1 Proof of Theorem G . . . . . . . . . . 123

2.9.2 Proofs of Theorems 2.1 - 2.4 . . . . . . . . . . 126

2.9.3 Proofs of Lemmas 2.1 - 2.10 . . . . . . . . . . 138

2.10 Tables and Figures . . . . . . . . . . 154

**3 Modelling and Estimating Other Dynamic Models** **160**

3.1 Introduction . . . . . . . . . . 160

3.2 Discrete-Continuous Control . . . . . . . . . . 160

3.3 Discrete and Continuous Controls . . . . . . . . . . 163

3.4 Ordered Discrete Response . . . . . . . . . . 167

3.5 Conclusion . . . . . . . . . . 169

**4 Consistent Estimation of an Identified Optimization Model** **170**

4.1 Introduction . . . . . . . . . . 170

4.2 The Model . . . . . . . . . . 171

4.3 Identification . . . . . . . . . . 173

4.4 BBL Methodology . . . . . . . . . . 175

4.5 An Example . . . . . . . . . . 177

4.6 Conclusion . . . . . . . . . . 181

# 1 Dynamic Discrete Choice Problems with Continuous State

## 1.1 Introduction

The inadequacy of static frameworks to model economic phenomena led to the development of recursive methods in economics. The mathematical theory underlying discrete time modelling is dynamic programming (DP) developed by Bellman (1957); for a review of its prevalence in modern economic theory, see Stokey and Lucas (1989). In this chapter we study the estimation of structural parameters and their functionals that underlie Markov decision processes (MDP) with discrete controls and time in the infinite horizon setting. The econometrics involved can be seen as an extension of the classical discrete choice analysis to a dynamic framework. Such models are popular in applied work, in particular in labor and industrial organization.

Discrete choice modelling has a long established history in the structural analysis of behavioral economics. McFadden (1974) pioneered the theory and methods of analyzing discrete choice in a static framework. The difficulty of estimating these discrete choice models parametrically arises in the form of computing multiple integrals. When estimating the finite dimensional parameters in a dynamic environment, these problems persist and are exacerbated by the need for researchers to estimate the (conditional) value functions defined recursively through the Bellman Equation. The treatment of these value functions determines the computational feasibility specific to the dynamic model. Our chapter contributes to the literature that deals with the computational complexity of this latter category.

The seminal paper of Rust (1988) proposed *additive separability* (AS) and *conditional independence* (CI) assumptions for the estimation of this type of dynamic models. These assumptions preserve the familiar structure of discrete choice problems of the static framework and have since served as the usual starting point for many applied and theoretical research that follow in the literature. In particular, Rust proposed a Nested Fixed Point (NFP) algorithm to estimate his parametric model by the maximum likelihood method. However, in practice, this method can post a considerable obstacle due to its requirement to repeatedly solve for the fixed point of some nonlinear map to obtain the value functions. Hotz and Miller (1993) avoided solving out for the value functions directly by showing the existence of an inversion map between the normalized value functions and the (conditional) choice probabilities. The value functions could be approximated, using nonparametric estimates of the choice probabilities, and used to estimate the structural parameters based on the method of moments. The nature of the inversion map is determined by the distribution of the unobserved state variables and will generally be nonlinear except for some special cases.

The semiparametric approach of Hotz and Miller significantly reduces the computational burden relative to the NFP algorithm. Their idea is central to several methodologies that followed, especially in the recent development of the estimation of dynamic games. A class of stationary infinite horizon Markovian games can be defined to include the MDP of interest as a special case. Various estimation procedures have been proposed to estimate the structural parameters of such empirical games. Pakes, Ostrovsky and Berry (2004), and Aguirregabiria and Mira (2007), considered two-step method of moments and pseudo maximum likelihood estimators respectively, which are included in the general class of minimum distance estimators defined by Pesendorfer and Schmidt-Dengler (2008). Bajari, Benkard and Levin (2007) generalizes the simulation-based estimators of Hotz et al. (1994) to the multiple agent setting. However, when the required transition density of the observed state variables is not specified parametrically, in both single and multiple agent settings, the aforementioned work assumed the observed state space is finite whenever the time horizon is infinite. As noted by

Aguirregabiria and Mira (2002,2007), we should be able to relax this requirement and allow for uncountable observable state space. The distinct attractive feature of the infinite horizon framework is that the value function is implicitly defined as a solution to a type II integral equation. This linear equation defines the value function through the smoothed (or integrated) Bellman equation (SBE) under the optimal decision rule, it is also known as the policy value equation (PVE). When the observable state space is finite, this linear equation is just a matrix equation whose statistical properties are well understood. The extension to allow for an uncountable state space is non-trivial. We also need to address the issue of the curse of dimensionality theoretically as well as in practice.

In this chapter, we propose a simple two-step semiparametric approach that falls in the general class of profiled semiparametric estimation discussed in Pakes and Olley (1995), and Chen, Linton and van Keilegom (2003). The criterion function will be based on some conditional moment restrictions that requires consistent estimators of the value functions. The additional difficulty here is due to the fact that the infinite dimensional parameter is defined through an integral equation. The study of the statistical properties of solutions to integral equations falls under the recent topic of research on inverse problem in econometrics, see Carrasco, Florens and Renault (2007) for a survey.[1] Type II integral equations are found, amongst others, in the study of additive models, see Mammen, Linton and Nielson (1995). We show that our problem is generally well-posed and utilize the approach similar to Linton and Mammen (2005) to estimate and provide the distribution theory for the infinite dimensional parameters of interest.

Our estimation strategy can be seen as a direct generalization of the unifying method of Pesendorfer and Schmidt-Dengler (2008), to estimate their Markovian games, that

[1] See also Carrasco's webpage on Inverse Problems in Econometrics at
`https://www.webdepot.umontreal.ca/Usagers/carrascm/MonDepotPublic/carrascm/inverse/index.html`

allows for continuous components in the observable state space. The main idea is to use the linearity of the policy value operator, where solving for the conditional value functions only requires the solving of some matrix equations when the state space is finite, to provide feasible estimator for the choice probabilities. To generalize this, we simply note that such matrix equations are linear integral equations in a finite dimensional space. We show that solving the analogous problem in an infinite dimensional space is also a well-posed problem for both population and empirical versions (at least for large sample size). In the first step, we flexibly estimate the integral equation using the method of kernel smoothing. The estimated PVE can be solved empirically, so we can provide estimates of the choice probabilities for any value of the structural parameter. The second stage involves minimizing some analogues of the moment restrictions over the parameter space based on feasible choice probabilities. The solving of the empirical integral equation in the first step requires us to approximate an inverse of a potentially large but invertible matrix but we only require to approximate this inverse once. We note that an independent work of Bajari, Chernozhukov, Hong and Nekipelov (2008) also proposes another estimation methodology that can estimate semiparametric Markovian games, based on the method of sieves, that allows for continuous observable state space. They focus on the case where the per period payoff utility function is linear in parameters and generate moment conditions based on the conditional value functions, some simple identification results are also provided in their paper. Therefore our methods are complementary in filling this gap in the literature. However, we feel that our estimation strategy, like its predecessor, is simpler and intuitive and by using the local approach of kernel smoothing, we can obtain the pointwise distribution theory of the infinite dimensional parameters that would otherwise be elusive with the series or splines expansion. Since the infinite dimensional parameters in MDP are the value functions, they may be of considerable interest themselves. Another advantage for the local estimator includes the optimality in the minimax sense for local linear estimators, see Fan (1993). In addition, we explicitly work under time series framework and provide the type of primitive conditions required for the validity of the methodology.

Since the main idea can be fully illustrated in the single agent setup, for most parts of the chapter, we initially consider the single agent setup and leave the discussion of the Markovian game estimation to the latter section. The chapter is organized as follows. Section 1.2 defines the MDP of interest and discusses SBE, PVE and the related linear inverse problem. Section 1.3 describes in detail the practical implementation of the procedure to obtain the feasible conditional choice probabilities. In Section 1.4, Primitive conditions and the consequent asymptotic distribution are provided, the semiparametric profiled likelihood estimator is illustrated as a special case. Section 1.5 discusses the extension to the dynamic games setting. Section 1.6 concludes.

## 1.2 Markov Decision Processes

We first define our time homogeneous MDP and introduce the main modelling assumptions and notation used throughout the chapter. We next outline the main issue of computational complexity for estimating MDP. The sources of the computational complexity are briefly reviewed and we introduce our estimator through the SBE and PVE that we view as an integral equation in 1.2.2 and discuss the inverse problem associated with solving such integral equations in 1.2.3.

### 1.2.1 Definitions and Assumptions

We index time by $t$, the agent is forward looking in solving the following infinite horizon intertemporal problem. The random variables in the model are the control and state variables, denoted by $a_t$ and $s_t$ respectively. The control variable, $a_t$, belongs to a finite set of alternatives $A = \{1, \ldots, K\}$. The state variables, $s_t$, is an element in

$\mathbb{R}^{J+K}$. At each period $t$, the agent observes $s_t$ and chooses an action $a_t$ in order to maximize her discounted expected utility. The present period utility is time separable and is represented by $u_\theta^0(a_t, s_t)$, for $\theta \in \Theta \subset \mathbb{R}^P$, and her action today directly affects the uncertain future states according to the (first order) Markovian transition density $p(ds_{t+1}|s_t, a_t)$. The next period utility is subjected to discounting at the rate $\beta \in (0,1)$. Formally, the agent is assumed to behave according to an optimal decision rule, $\mathcal{A}_\tau = \{\alpha_t(s_t)\}_{t=\tau}^\infty$, in solving the following sequential problem (SP) for any time $\tau$

$$V_\theta^0(s_\tau) = \sup_{\mathcal{A}_\tau} E\left[\left.\sum_{t=\tau}^{\infty} \beta^t u_\theta^0(a_t, s_t)\right| s_\tau\right]. \tag{1}$$

Under some regularity conditions, see Bertsekas and Shreve (1978) and Rust (1988), Blackwell's Theorem and its generalization ensure the following important properties. Firstly, there exists a deterministic and stationary Markovian optimal decision rule $\alpha_\theta^0(\cdot)$ so that $\alpha_\theta^0(s_t) = \alpha_\theta^0(s_{t+\tau})$ for any $s_t = s_{t+\tau}$ and any $t, \tau$, i.e.

$$\alpha_\theta^0(s_t) = \arg\max_{a \in A}\left\{u_\theta^0(a, s_t) + \beta E\left[V_\theta^0(s_{t+1})|s_t, a_t = a\right]\right\} \quad \text{for all } t \geq 1. \tag{2}$$

Secondly, the value function, $V_\theta^0$, is the unique solution to the Bellman's equation (BE)

$$V_\theta^0(s_t) = \max_{a \in A}\left\{u_\theta^0(a, s_t) + \beta E\left[V_\theta^0(s_{t+1})|s_t, a_t = a\right]\right\}. \tag{3}$$

In order to make this a more tractable econometric problem, the following set of conditions for the class of MDP of interest are imposed:

ASSUMPTION M1.1: *The observed data for each individual $\{a_t, x_t\}_{t=1}^{T+1}$ are the controlled stochastic processes satisfying (2) with exogenously known $\beta$.*

ASSUMPTION M1.2: *(Conditional Independence) The transitional distribution has* *the following factorization:* $p(x_{t+1}, \varepsilon_{t+1}|x_t, \varepsilon_t, a_t) = q(\varepsilon_{t+1}|x_{t+1}) f_{X'|X,A}(x_{t+1}|x_t, a_t)$ *for all* $t$.

ASSUMPTION M1.3: $s_t = (x_t, \varepsilon_t) \in X \times \mathbb{R}^K$, *where* $X = X^C \times X^D$ *is a compact subset of* $\mathbb{R}^J$. $X^C$ *includes intervals and* $X^D$ *is finite, they denote the observable state space with continuous and discrete components respectively.* $\varepsilon_t$ *is a vector of unobserved state variables, whose dimension is* $K$, *the cardinality of* $A$. *The distribution of* $\varepsilon_t$ *is also known and is absolutely continuous with respect to some Lebesgue measure with Radon Nikodym density* $q(\varepsilon_t|x_t)$ *with support* $\mathbb{R}^K$

ASSUMPTION M1.4: *(Additive Separability) The per period payoff function* $u_\theta^0 : A \times S \to \mathbb{R}$ *is specified upto some unknown parameters* $\theta \in \Theta \subset \mathbb{R}^P$*.and is additive separable w.r.t. unobservable state variables,* $u_\theta^0(a_t, x_t, \varepsilon_t) = u_\theta(a_t, x_t) + \varepsilon_{a_t,t}$.

Conditions M1.1 is a standard simplification to keep the model tractable. The knowledge of $\beta$ is important as it is generally not identified in MDP models. The popular infinite time framework yields us an elegant and simple linear equation to work with. We discuss the solving of such equations below; The fundamental assumption in the current literature is M1.2, the (CI) assumption of Rust (1988). The combination of M1.2 and M1.4 allows us to set our model in the familiar framework of static discrete choice modelling; The observable state space of M1.3 is usually assumed to be finite but we allow for it to include intervals. Compactness $X^C$ is not necessary, imposed here for the ease of exposition; The additive structure on the payoff function in M1.4 is also imposed by Rust (1994).

It is our goal to estimate the structural parameters as well as some functionals depending on them. Conditions M1.1 - M1.4 are crucial to the estimation methodology we propose. These conditions are standard in the literature. In particular, M1.3 is weaker than the usual finite $X$ assumption when no parametric assumption is assumed on $f_{X'|X,A}(x_{t+1}|x_t, a_t)$ in the infinite horizon framework. For departures of this framework see the discussion in the survey of Aguirregabiria and Mira (2008) and the references therein. Henceforth Conditions M1.1 - M1.4 will be assumed and later strengthened as appropriate.

### 1.2.2 Policy Value Equation

Similarly to the static discrete choice models, the estimation of our controlled process requires us to compute the choice probabilities. There are two numerical aspects that we need to consider in the evaluation of the choice probabilities. The first are the multiple integrals, that also arise in the static framework, where in practice many researchers avoid this issue via the use of conditional logit assumption of McFadden[2] (1974). The second is that we must compute the value function, directly or indirectly, as defined in (1) and (3) - this is unique to the dynamic setup. First we introduce what we also call, with an abuse of a terminology, the value function *defined* on the observed variables, $V_\theta$, which is a stationary solution to the *policy value equation* when for any $\theta$, cf. (3)

$$V_\theta(s_t) = u_\theta^0(a_t, s_t) + \beta E[V_\theta(s_{t+1})|s_t], \tag{4}$$

so that

$$V_\theta(s_t) = E\left[\sum_{\tau=t}^{\infty} \beta^{\tau-t} u_\theta^0(a_\tau, s_\tau) \middle| s_t\right].$$

In particular, we have must have $V_{\theta_0}(s_t) = V_{\theta_0}^0(s_t)$ and $a_t = \alpha_{\theta_0}^0(s_t)$. We stress that the equation above is also well defined for any $\theta$ that is not equal to $\theta_0$; then $V_\theta$ is interpreted as the value function for an economic agent whose underlying preference is $\theta$ but is using the policy function that is optimal with respect to $\theta_0$. To see precisely the difficulty we face, we first update our BE under the assumptions M1.1 - M1.4, M1.2

[2]Unlike in static models, we do not suffer from the undesirable I.I.A. when use i.i.d. extreme values errors of type I in the dynamic framework.

implies:

$$E\left[V_\theta\left(s_{t+1}\right)|s_t,a_t\right]=E\left[E\left[V_\theta\left(s_{t+1}\right)|x_{t+1}\right]|x_t,a_t\right].$$

Denote $u_\theta\left(a_t,x_t\right)+\beta E\left[E\left[V_\theta\left(s_{t+1}\right)|x_{t+1}\right]|x_t,a_t\right]$ by $v_\theta\left(a_t,x_t\right)$, we can define the model implied policy function $\alpha_\theta$ by

$$\alpha_\theta\left(x_t,\varepsilon_t\right)=a_t\Leftrightarrow v_\theta\left(a_t,x_t\right)+\varepsilon_{a_t,t}\geq v_\theta\left(a,x_t\right)+\varepsilon_{a,t}\text{ for }a\neq a_t. \tag{5}$$

Again, we must have $\alpha_{\theta_0}\left(s_t\right)=\alpha^0_{\theta_0}\left(s_t\right)$. The above is familiar from the static multinomial choice framework. In order to compute the choice probabilities we need to integrate across the domain of the unobservable states satisfying (5) as well as provide the values for $v_\theta$, for each $\theta$, denoting the conditional choice probabilities by $\{P\left(a_t|x_t,\theta\right)\}$

$$\begin{aligned}P\left(a_t|x_t,\theta\right) &= \Pr\left[v_\theta\left(a_t,x_t\right)+\varepsilon_{a_t,t}\geq v_\theta\left(a,x_t\right)+\varepsilon_{a,t}\text{ for }a\neq a_t|x_t\right] \qquad (6)\\ &= \int\mathbf{1}\left[\alpha_\theta\left(x_t,\varepsilon_t\right)=a_t\right]q\left(d\varepsilon_t|x_t\right).\end{aligned}$$

Generally, suppose that we know $v_\theta$, (6) will have no closed form and the task of performing multiple integrals numerically is often non-trivial.[3] Under M1.3, we can make distributional assumptions on the error terms, for example using the popular i.i.d. extreme value of type I - then we can avoid the multiple integrals as (6) has the well known multinomial logit form

$$P\left(a_t|x_t,\theta\right)=\frac{\exp\left(v_\theta\left(a_t,x_t\right)\right)}{\sum_{a\in A}\exp\left(v_\theta\left(a,x_t\right)\right)}. \tag{7}$$

Our estimation strategy accommodates for general form of distribution per M7. However, the problem we want to focus on is the fact that we generally do not know $v_\theta$. For

[3] See the discussion of Hajivassiliou and Ruud (1994) where they provided some form of escape via simulation methods.

$v_\theta(a_t, x_t)$ contains $E[E[V_\theta(s_{t+1})|x_{t+1}]|x_t, a_t]$, which we denote by $g_\theta(a_t, x_t)$, defined through some nonlinear functional equation that we need to solve for.

Under the rationality assumption of M1.1, we define the conditional value function by taking conditional expectation on (4) w.r.t. $x_t$,

$$
\begin{aligned}
E[V_\theta(s_t)|x_t] &= E\left[u_\theta^0(a_t, s_t)|x_t\right] + \beta E[E[V_\theta(s_{t+1})|s_t, a_t]|x_t] \qquad (8)\\
&= E\left[u_\theta^0(a_t, s_t)|x_t\right] + \beta E[E[V_\theta(s_{t+1})|x_{t+1}]|x_t].
\end{aligned}
$$

The latter equality follows from M1.2. Therefore we can express (8) generally as a linear integral equation of type II. Imposing the structural parameterization of M1.4 we can represent (8) by

$$m_\theta = r_\theta + \mathcal{L}m_\theta, \qquad (9)$$

where for any given $\theta$ and $x \in X$: $m_\theta(x)$ is the ex-ante (conditional) expected value function $E[V_\theta(s_t)|x_t = x]$. $r_\theta(x)$ is the ex-ante expected immediate payoff given state $x_t = x$, namely $E\left[u_\theta^0(a_t, s_t)|x_t = x\right]$. The integral operator $\mathcal{L}$ generates discounted expected next period values of its operands, e.g. $\mathcal{L}m_\theta(x) = \beta E[m_\theta(x_{t+1})|x_t = x]$. If we could solve for $m_\theta$ then we need another level of smoothing on $m_\theta$ to obtain the continuation value $v_\theta$ as defined in the previous subsection. In particular, we can define $g_\theta$ through the following linear transform

$$g_\theta = \mathcal{H}m_\theta, \qquad (10)$$

where for any given $\theta$ and $(x, a) \in X \times A$, $g_\theta(a, x)$ is the ex-ante expected value function $E[V_\theta(s_{t+1})|x_t = x, a_t = a]$ $(= E[m_\theta(x_{t+1})|x_t = x, a_t = a])$ and the integral operator $\mathcal{H}$ generates the expected next period values of its operands (cf. $\mathcal{L}$). The continuation value, net of unobserved shock, can be written in a linear functional notation

$$v_\theta = u_\theta + \beta \mathcal{H} m_\theta. \tag{11}$$

Before we can discuss the estimation of $v_\theta$, we need to address some issues regarding the solution of integral equations since $m_\theta$ is defined as a solution to the integral equation (9). It is natural to ask the fundamental question whether our problem is well-posed in the sense of Hadamard, namely, whether the solution of (9) exist and if so, whether it is unique and stable. The study of the solution to such integral equations falls in the general framework of linear inverse problems, and in what follows we show that our inverse problem is well-posed.

### 1.2.3 Linear Inverse Problems

The study of inverse problems is an old problem in applied mathematics. The type of inverse problems one commonly encounters in econometrics are integral equations. Carrasco et al. (2007) focused their discussion on ill-posed problems of integral equations of type I where recent works often needed regularizations in Hilbert Spaces to stabilize their solutions. Here we face an integral equation of type II, which is easier to handle, and in addition, the convenient structure of SBE allows us to easily show that the problem is well-posed in a familiar Banach Space. We now define the normed linear space and the operator of interest, and proof this claim. We shall simply state relevant results from the theory of integral equations. For definitions, proofs and further details on integral equations, readers are referred to Kress (1999) and the references therein.

From the Riesz Theory of operator equations of the second kind with compact operators on a normed space, say $A : X \to X$, we know that $I - A$ is injective if and only if it is surjective, and if it is bijective, then the inverse operator $(I - A)^{-1} : X \to X$ is bounded. We will be working on the Banach space $(B, \|\cdot\|)$, where $B = C(X)$ is a space of continuous functions defined on the compact subset of $\mathbb{R}^J$, equipped with the sup-norm, i.e. $\|\phi\| = \sup_{x \in X} |\phi(x)|$. $\mathcal{L}$ is a linear map, $\mathcal{L} : C(X) \to C(X)$ , such that, for any $\phi \in C(X)$ and $x \in X$,

$$\mathcal{L}\phi(x) = \beta \int_X \phi(x') f_{X'|X}(dx'|x),$$

where $f_{X|X}(dx_{t+1}|x_t)$ denotes the conditional density of $x_{t+1}$ given $x_t$.

In fact, the compactness of the operator is not required in this case since we know the existence, uniqueness and stability of the solution to (9) are assured as we can show $\mathcal{L}$ is a contraction. To see this, take any $\phi \in C(X)$ and $x \in X$,

$$|\mathcal{L}\phi(x)| \leq \beta \int_X |\phi(x')| f_{X'|X}(dx'|x) \leq \beta \sup_{x \in X} |\phi(x)|,$$

since the discounting factor $\beta \in (0, 1)$,

$$\|\mathcal{L}\phi\| \leq \beta \|\phi\| \Rightarrow \|\mathcal{L}\| \leq \beta < 1.$$

This implies that our inverse problem is well-posed. Further, the contraction property means we can represent the solution to (9) using the Neumann series,

$$\begin{aligned} m_\theta &= (I - \mathcal{L})^{-1} r_\theta && (12) \\ &= \lim_{T \longrightarrow \infty} \sum_{\tau=1}^{T} \mathcal{L}^\tau r_\theta. && (13) \end{aligned}$$

Therefore (13) provides one obvious way of approximating the solution to the integral equation which will converge geometrically fast to the true function. If $X$ is countable, then $\mathcal{L}^\tau$ would be represented by a $\tau$-step ahead transition matrix (scaled by $\beta^\tau$). Note that the operator for the (uncountable) infinite dimensional case share the analogous interpretation of $\tau$-step ahead transition operator with discounting.

Since our problem is well-posed, then it is reasonable to expect that with sufficiently good estimates of $(m_\theta, \mathcal{L}, \mathcal{H})$ our estimated integral equation is also well-posed and will lead to (uniform) consistent estimators for $(g_\theta, v_\theta)$. Our strategy is to use nonparametric methods to generate the empirical versions of (9) and (10), then use them to provide an approximate for $v_\theta$ necessary for computing the choice probabilities.

### 1.3 Estimation of Conditional Choice Probabilities

In this section we discuss the estimation of the nonparametric components necessary for the computation of the model implied choice probabilities. Our objective is to construct an estimator of $v_\theta$ as defined by (9), (10) and (11) from a time series $\{a_t, x_t\}_{t=1}^{T}$. A pure time series approach is assumed for notational simplicity, this can be trivially extended with $N$ independent realizations of the same controlled process. We proceed in two steps. First, we nonparametrically compute estimates of the kernels of $\mathcal{L}, \mathcal{H}$ and for each $\theta$, estimate $r_\theta$. Then obtain the estimate of $m_\theta$ by solving the empirical version of the integral equation (9) and estimate $g_\theta$ analogously from an empirical version of (10).

There are numerous choices available for empirically solving the integral equation in (9). We need to first decide on the nonparametric method. We will focus on the method of kernel smoothing due to its simplicity of use as well as its well established theoretical grounding. Our nonparametric estimation of the conditional expectations will be based on the Nadaraya-Watson estimator. The local constant estimator is chosen for its familiarity and simplicity of notation. However, since we will be working on bounded sets, it is necessary to address the boundary effects. The treatment of the boundary issues is straightforward, the precise trimming condition is described in Section 4. So we will assume to work on a smaller space $X_T \subset X$ where $X_T = \left(X_T^C, X^D\right)$

denotes a set where the support of the uncountable component is some strict compact subset of $X^C$ but increases to $X^C$ in $T$. When allowing for discrete components, we simply use the frequency approach, smoothing over the discrete components is also possible, see Aitchison and Aitken (1976). We will also need to make a decision on how to define and interpolate the solution to the empirical version of (9) in practice. We discuss two asymptotically equivalent options for this latter choice, whether the size of the empirical integral equation does or does not depend on the sample size, as one may have a preference given the relative size of the number of observations.

Estimation of $r_\theta, \mathcal{L}$ and $\mathcal{H}$:

We now define the nonparametric estimators, $(\widehat{r}_\theta, \widehat{\mathcal{L}}, \widehat{\mathcal{H}})$, of $(r_\theta, \mathcal{L}, \mathcal{H})$. Any generic density of a mixed continuous-discrete random vector $w_t = \left(w_t^c, w_t^d\right)$, $f_w : \mathbb{R}^{l^C} \times \mathbb{R}^{l^D} \to \mathbb{R}^+$ for some positive integers $l^C$ and $l^D$, is estimated as follows,

$$\widehat{f}_w\left(w^c, w^d\right) = \frac{1}{T}\sum_{t=1}^{T} K_h\left(w_t^c - w^c\right) \mathbf{1}\left[w_t^d = w^d\right],$$

where $K$ is some user chosen symmetric probability density function, $h$ is a positive bandwidth and for simplicity independent of $w^c$. $K_h(\cdot) = K(\cdot/h)/h$ and if $l^C > 1$ then $K_h\left(w_t^c - w^c\right) = \prod_{l=1}^{l^C} K_{h_l}\left(w_{t,l}^c - w_l^c\right)$, $\mathbf{1}[\cdot]$ denotes the indicator function, namely $\mathbf{1}[\mathcal{A}] = 1$ if event $\mathcal{A}$ occurs and takes value zero otherwise. Similar to the product kernel, the contribution from a multivariate discrete variable is represented by products of indicator functions. The conditional densities/probabilities are estimated using the ratio of the joint and marginal densities. The local constant estimator of any generic regression function, $E\left[z_t | w_t = w\right]$ is defined by,

$$\widehat{E}\left[z_t | w_t = w\right] = \frac{\frac{1}{T}\sum_{t=1}^{T} z_t K_h\left(w_t^c - w^c\right) \mathbf{1}\left[w_t^d = w^d\right]}{\widehat{f}_w(w)}. \tag{14}$$

(I) Estimation of $r_\theta$:

For any $x \in X_T$,

$$\begin{aligned} r_\theta(x) &= E\left[u_{0,\theta}(a_t, s_t) \mid x_t = x\right] \\ &= E\left[u_\theta(a_t, x_t) \mid x_t = x\right] + E\left[\varepsilon_{a_t} \mid x_t = x\right] \\ &= \rho_{1,\theta}(x) + \rho_{2,\theta}(x). \end{aligned}$$

The first term can be estimated by

$$\widehat{\rho}_{1,\theta}(x) = \sum_{a \in A} \widehat{P}(a|x)\, u_\theta(a, x), \tag{15}$$

or, alternatively, the Nadaraya-Watson estimator,

$$\widetilde{\rho}_{1,\theta}(x) = \widehat{E}\left[u_\theta(a_t, x_t) \mid x_t = x\right].$$

In (15), $\left\{\widehat{P}(a|x)\right\}_{a \in A}$ is a sequence kernel estimator of the conditional choice probabilities, and equivalently the Nadaraya-Watson estimator. Generally, by the inversion theorem of Hotz and Miller, it will be more convenient to use (15) since we have to compute $\left\{\widehat{P}(a|x)\right\}_{a \in A}$ in any case, as we shall see below.

The conditional mean of the unobserved states, $\rho_{2,\theta}$, is generally non-zero due to selectivity. By the inversion theorem of Hotz and Miller, we know $\rho_{2,\theta}$ can be expressed as a known smooth function of the choice probabilities. For example, the i.i.d. type I extreme value errors assumption will imply that

$$\rho_{2,\theta}(x) = \gamma + \sum_{a \in A} P(a|x) \log\left(P(a|x)\right), \tag{16}$$

where $\gamma$ is the Euler's constant. An estimator of $\rho_{2,\theta}$ can therefore be obtained by plugging in the local constant (linear) estimator of the choice probabilities.

Alternatively, we could use another consistent estimator, where the estimator of the conditional choice probability, $\left\{\widehat{P}\left(a|x\right)\right\}_{a\in A}$, can be estimated by the Nadaraya-Watson estimator of the regression of $\mathbf{1}\left[a_t = a\right]$ on $x_t = x$. This approach may be more convenient when sample size is relatively small, and we want to solve the empirical version of (9) by using purely nonparametric methods for interpolation, where we could use the local linear estimator to address the boundary effects.

(II) ESTIMATION OF $\mathcal{L}$ AND $\mathcal{H}$:

Suppose $X^D$ is empty. For the integral operators $\mathcal{L}$ and $\mathcal{H}$, if we would like to use the numerical integration to approximate the integral, we only need to provide the nonparametric estimators of their kernels, respectively, $\widehat{f}_{X'|X}\left(dx_{t+1}|x_t\right)$ and $\widehat{f}_{X'|X,A}\left(dx_{t+1}|x_t, a_t\right)$.

For any $\phi \in C\left(X_T\right)$, the empirical operators are defined as,

$$\begin{aligned} \widehat{\mathcal{L}}\phi\left(x\right) &= \int_{X_T} \phi\left(x'\right)\widehat{f}_{X'|X}\left(dx'|x\right), && (17) \\ \widehat{\mathcal{H}}\phi\left(a,x\right) &= \int_{X_T} \phi\left(x'\right)\widehat{f}_{X'|X,A}\left(dx'|x,a\right). && (18) \end{aligned}$$

So $\widehat{\mathcal{L}}$ and $\widehat{\mathcal{H}}$ are linear operators on the Banach space of continuous functions on $X_T$ with range $C\left(X_T\right)$ and $C\left(A\times X_T\right)$ respectively under sup-norm. Alternatively, we could use the Nadaraya-Watson estimator, defined in (14), to estimate the operators,

$$\begin{aligned} \widetilde{\mathcal{L}}\phi\left(x\right) &= \widehat{E}\left[\phi\left(x_{t+1}\right)|x_t = x\right], \\ \widetilde{\mathcal{H}}\phi\left(a,x\right) &= \widehat{E}\left[\phi\left(x_{t+1}\right)|x_t = x, a_t = a\right]. \end{aligned}$$

Note that, if $X$ is finite then the integrals in (17) and (18) will be defined with respect to discrete measures, then $(\widehat{\mathcal{L}},\widehat{\mathcal{H}})$ and $(\widetilde{\mathcal{L}},\widetilde{\mathcal{H}})$ can be equivalently represented by the same stochastic matrices.

ESTIMATION OF $m_\theta$,$g_\theta$ AND $v_\theta$:

We first describe the procedure used in Linton and Mammen (2005), by using $(\widehat{\mathcal{L}}, \widehat{\mathcal{H}})$, to solve the empirical integral equation. We define $\widehat{m}_\theta$ as any sequence of random functions defined on $X_T$ that approximately solves $\widehat{m}_\theta = \widehat{r}_\theta + \widehat{\mathcal{L}}\widehat{m}_\theta$. Formally, we shall assume that $\widehat{m}_\theta$ is any random sequence of functions that satisfy

$$\sup_{\theta\in\Theta,x\in X_T} \left|\left(I - \widehat{\mathcal{L}}\right)\widehat{m}_\theta(x) - \widehat{r}_\theta(x)\right| = o_p\left(T^{-1/2}\right). \tag{19}$$

In practice, we solve the integral equation on a finite grid points, which reduces it to a large linear system. Next we use $\widehat{m}_\theta$ to define $\widehat{g}_\theta$, specifically we define $\widehat{g}_\theta$ as any random sequence of functions that satisfy

$$\sup_{\theta\in\Theta,a\in A,x\in X_T} \left|\widehat{g}_\theta(a,x) - \widehat{\mathcal{H}}\widehat{m}_\theta(a,x)\right| = o_p\left(T^{-1/2}\right). \tag{20}$$

Once we obtain $\widehat{g}_\theta$, the estimator of $v_\theta$ is defined by

$$\sup_{\theta\in\Theta,a\in A,x\in X_T} \left|\widehat{v}_\theta(a,x) - u_\theta(a,x) - \beta\widehat{g}_\theta(a,x)\right| = o_p\left(T^{-1/2}\right).$$

For illustrational purposes, ignoring the trimming factors, we will assume that $X = [\underline{x}, \overline{x}] \subset \mathbb{R}$.

For any integrable function $\phi$ on $X$, define $J(\phi) = \int \phi(t)\,dt$. Given an ordered sequence of $n$ nodes $\{t_{j,n}\} \subset [a,b]$, and a corresponding sequence of weights $\{\omega_{j,n}\}$

such that $\sum_{j=1}^{n} \omega_{j,n} = b - a$, a valid integration rule would satisfy

$$\begin{aligned} \lim_{n\to\infty} J_n(\phi) &= J(\phi) \\ J_n(\phi) &= \sum_{j=1}^{n} \omega_{j,n} \phi(t_{j,n}), \end{aligned}$$

for example Simpson's rule and Gaussian quadrature both satisfy this property for smooth $\phi$. Therefore the empirical version of (9) can be approximated for any $x \in [a, b]$ by

$$\widehat{m}_\theta(x) = \widehat{r}_\theta(x) + \beta \sum_{j=1}^{n} \omega_{j,n} \widehat{f}_{X'|X}(t_{j,n}|x)\, \widehat{m}_\theta(t_{j,n}), \tag{21}$$

so the desired solution that approximately solves the empirical integral equation will satisfy the equation below at each node $\{t_{j,n}\}$,

$$\widehat{m}_\theta(t_{i,n}) = \widehat{r}_\theta(t_{i,n}) + \beta \sum_{j=1}^{n} \omega_{j,n} \widehat{f}_{X'|X}(t_{j,n}|t_{i,n})\, \widehat{m}_\theta(t_{j,n}). \tag{22}$$

This is equivalent to solving a system of $n$ equations with $n$ variables, the system of (22) can be written in a matrix notation as

$$\widehat{\mathbf{m}}_\theta = \widehat{\mathbf{r}}_\theta + \widehat{\mathbf{L}}\widehat{\mathbf{m}}_\theta, \tag{23}$$

where $\widehat{\mathbf{m}}_\theta = (\widehat{m}_\theta(t_{1,n}), \ldots, \widehat{m}_\theta(t_{n,n}))^\top, \widehat{\mathbf{r}}_\theta = (\widehat{r}_\theta(t_{1,n}), \ldots, \widehat{r}_\theta(t_{n,n}))^\top$, $I_n$ is an identity matrix of order $n$ and $\widehat{\mathbf{L}}$ is a square $n$ matrix such that $(\widehat{\mathbf{L}})_{ij} = \beta\omega_{j,n}\widehat{f}_{X'|X}(t_{j,n}|t_{i,n})$. Since $\widehat{f}_{X'|X}(\cdot|x)$ is a proper density for any $x$, with a sufficiently large $n$, $(I_n - \widehat{\mathbf{L}})$ is invertible by dominant diagonal theorem. So there is a unique solution to the system (23) for a given $\widehat{\mathbf{r}}_\theta$. In practice we have a variety of ways to solve for $\widehat{\mathbf{m}}_\theta$ with one obvious candidate being the successive approximation as mentioned in (13). Once we obtain $\widehat{\mathbf{m}}_\theta$, we can approximate $\widehat{m}_\theta(x)$ for any $x \in X$ by substituting $\widehat{\mathbf{m}}_\theta$ into the

RHS of (21). This is known as the Nyström interpolation. We need to approximate another integral to estimate $g_\theta$. This could be done using the conventional method of kernel regression as discussed in Section 3.1, or by appropriately selecting sequences of $r$ nodes $\{q_{j,r}\}$ and weights $\{\zeta_{j,n}\}$ so that

$$\widehat{g}_\theta(j,x)=\sum_{j=1}^{r}\zeta_{j,n}\widehat{f}_{X'|X,A}(q_{j,r}|x,j)\,\widehat{m}_\theta(q_{j,r}), \tag{24}$$

where the computation for the system of (24) is trivial. See Judd (1998) for a more extensive review of the methods and issues of approximating integrals and also the discussion of iterative approaches in Linton and Mammen (2003) for large grid sizes.

Alternatively, (again, ignoring the trimmed observations) we can form a matrix equation of size $T-1$,

$$\widetilde{\mathbf{m}}_\theta=\widetilde{\mathbf{r}}_\theta+\widetilde{\mathbf{L}}\widetilde{\mathbf{m}}_\theta,$$

to estimate (9) at the observed points with the $t$-th element defined below,

$$\widetilde{m}_\theta(x_t)=\widetilde{r}_\theta(x_t)+\beta\frac{\frac{1}{T-1}\sum_{t=1}^{T-1}\widetilde{m}_\theta(x_{t+1})K_h(x_t-x)}{\frac{1}{T-1}\sum_{\tau=1}^{T-1}K_h(x_\tau-x_t)}.$$

By the dominant diagonal theorem, the matrix equation above always has a unique solution for any $T\geq 2$. Once solved, the estimators of $\widetilde{m}_\theta$ can be interpolated by

$$\widetilde{m}_\theta(x)=\widetilde{r}_\theta(x)+\beta\widehat{E}\left[m(x_{t+1})\,|x_t=x\right],$$

for any $x\in X_T$. Similarly, $\widetilde{g}_\theta$ and $\widetilde{v}_\theta$ can be estimated nonparametrically without introducing any additional numerical error. Clearly, the more observation we have, the latter method will be more difficult as dimension of the matrix representing $\widetilde{\mathbf{L}}$ is large whilst the grid points for the former empirical equation is user-chosen.

Practical Discussion:

We reflect on the computational effort required of the proposed method. It will be helpful to have in mind the methodology of Pesendorfer and Schmidt-Dengler (2008) as our methods coincide when the $X$ is finite and there is only 1 player in the game (vice versa, extending from a single agent decision process to a dynamic game). For each $\theta$, the nonparametric estimates of $(r_\theta, \mathcal{L}, \mathcal{H})$ have closed form and are very easy to compute even with large dimensions, further, the empirical integral operators (or their approximations) only need to be computed once at the required nodes since they do not depend on $\theta$. Solving the empirical integral equation to obtain $\widehat{m}_\theta$, in (23), is the only potential complication that does not exist in a static problem. However, in this setup, this reduces to the need to invert a large matrix that approximates $(I - \mathcal{L})$ that only need to be done once at the beginning and stored for future computation with any other $\theta$. Estimators of $(m_\theta, g_\theta, v_\theta)$ are obtained trivially for any $\theta$, by simple matrix multiplication, once the empirical operator of $(I - \mathcal{L})^{-1}$ is obtained. We note further computational gain is possible if $u_\theta$ is linear in $\theta$. The reason for this is clear, from (15) and (16), linearity in $\theta$ implies $r_\theta = \sum_{l=1}^{L} \theta_l r_l$, utilizing the fact that the inverse of $(I - \mathcal{L})$ is a linear operator so we have $m_\theta = \sum_{l=1}^{L} \theta_l (I - \mathcal{L})^{-1} r_l$, where, once again, $\left(I - \widehat{\mathcal{L}}\right)^{-1} \widehat{r}_l$ only need to be computed once for each $l$. See Hotz, Miller, Sanders and Smith (1994) and Bajari, Benkard and Levin (2007) for related utilization of the repeated substitution concept.

However, it is important to note that, as we have decided on the kernel smoothing approach there is an issue of bandwidth selection which is important for small sample properties. Further, it is easy to see that the invertibility of the matrix $(I - \widehat{\mathbf{L}})$ and $(I - \widetilde{\mathbf{L}})$ are not dependent on the number of continuous and/or discrete components. Clearly, there are a lot of choices available regarding integral approximation and matrix inversion methods. It is beyond the scope of this chapter to analyze the finite sample performance of these various methodologies.

### 1.4 Distribution Theory

In this section we provide the type of primitives sufficient to obtain the distribution theory for both finite and infinite dimensional parameters of interest. A class of criterion functions can be generated from the following conditional moment restrictions,

$$E\left[\mathbf{1}\left[a_t = a\right] - P_{\theta_0}\left(a|x_t\right)|x_t\right] = 0. \tag{25}$$

We focus on a specific example of a profiled likelihood estimator.[4] In terms of the dimensionality of $X$, we restrict $X^C \subset \mathbb{R}$, the reason being this is the scenario that applied researchers may prefer to work with. This does not limit the usefulness of the primitives provided. For other estimation criterions, since two-step estimation problems of this type can be compartmentalized into nonparametric first stage and optimization in the second stage, the primitives below will be directly applicable. There might be other intrinsically continuous observable state variables that require discretizing but with increasing dimension in $X^C$, the practitioners will need to employ higher order kernels and/or undersmooth in order to obtain the parametric rate of convergence for the finite structural parameters, adaptation of the primitives are straightforward and will be discussed accordingly.

There are general large sample theory of profiled semiparametric estimators available that treat the estimators defined in our models. In particular, the work of Pakes and Olley (1995) and Chen, Linton and van Keilegom (2003) provide high level conditions for obtaining root$-T$ consistent estimators are directly applicable. The relevant large sample properties for the nonparametric first stage, under the time series

[4]This estimator can be derived from some conditional moment restrictions if the zero of the first order condition identifies the true parameter.

framework, for the pointwise results see the results of Roussas (1967,1969), Rosenblatt (1970,1971) and Robinson (1983). Roussas first provided central limit results for kernel estimates of Markov sequences, Rosenblatt established the asymptotic independence and Robinson generalized such results to the $\alpha$-mixing case. The uniform rates have been obtained for the class of polynomial estimators by Masry (1996), in particular, our method is closely related to the recent framework of Linton and Mammen (2005) who obtained the uniform rates and pointwise distribution theory for the solution of a linear integral equation of type II. We assume to possess a time series data $\{a_t, x_t\}_{t=1}^T$ generated from the MDP described in Section 1.2.

INFINITE DIMENSIONAL PARAMETERS:

We begin with some primitives. In addition to M1.1 - M1.4, the following sufficient conditions are weak enough to accommodate most of the existing empirical works in applied labor and industrial organization involving estimation of MDP.

We denote the strong mixing coefficient as

$$\alpha(k) = \sup_{t\in\mathbb{N}} \sup_{A\in\mathcal{F}_{t+k}^{\infty}, \mathcal{F}_{-\infty}^{t}} |\Pr(A\cap B) - \Pr(A)\Pr(B)| \quad \text{for } k \in \mathbb{Z},$$

where $\mathcal{F}_a^b$ denotes the sigma-algebra generated by $\{a_t, x_t\}_{t=a}^b$.

B1.1 *$X \times \Theta$ is a compact subset of $\mathbb{R}^J \times \mathbb{R}^L$ with $X^C = [\underline{x}, \overline{x}]$.*

B1.2 *The process $\{a_t, x_t\}_{t=1}^T$ is strictly stationary and strongly mixing, with a mixing coefficient $\alpha(k)$,such that for some $C \geq 0$ and some, possibly large $\chi > 0, \alpha(k) \leq Ck^{-\chi}$.*

B1.3 *The density of $x_t$ is absolutely continuous $f_{X^C,X^D}(dx_t, x_t^d)$ for each $x_t^d \in X^D$. The joint density of $(a_t, x_t)$ is bounded away from zero on $X^C$ and is twice* *continuously differentiable over $X^C$ for each $\left(x_t^d, a_t\right) \in X^D \times A$. The joint density of $\left(x_{t+1}, x_t, a_t\right)$ is twice continuously differentiable over $X^C \times X^C$ for each $\left(x_{t+1}^d, x_t^d, a_t\right) \in X^D \times X^D \times A$.*

B1.4 *The mean of the per period payoff function $u_\theta\left(a_t, x_t\right)$ is twice continuously differentiable on $X^C \times \Theta$ for each $\left(x_t^d, a_t\right) \in X^D \times A$.*

B1.5 *The kernel function is a symmetric probability density function with bounded support such that for some constant $C, |K(u) - K(v)| \leq C|u - v|$. Define $\mu_j(K) = \int u^j K(u)\,du$ and $\kappa_j(K) = \int K^j(u)\,du$.*

B1.6 *The bandwidth sequence $h_T$ satisfies $h_T = \gamma_0(T)\,T^{-1/5}$ and $\gamma_0(T)$ bounded away from zero and infinity.*

B1.7 *The triangular array of trimming factors $\{\mathfrak{c}_{t,T}\}$ is defined such that $\mathfrak{c}_{t,T} = \mathbf{1}\left[x_t^c \in X_T^C\right]$ where $X_T = \left[\underline{x} + c_T, \overline{x} - c_T\right]$ and $\{c_T\}$ is any positive sequence converging monotonically to zero such that $h_T < c_T$.*

B1.8 *The distribution of $\varepsilon_t$ is known to be distributed as i.i.d. extreme value of type I across $K$ alternatives, and is mean independent of $x_t$ and is i.i.d. across $t$.*

The compactness of the parameter space in B1.1 is standard. Compactness of the continuous component of the observable state space can be relaxed by using an increasing sequence of compact sets that cover the whole real line, see Linton and Mammen (2005) for the modelling in the tails of the distribution. The dimension of $X^C$ is assumed to be 1 for expositional simplicity, discussion on this is follows the theorems below. On the other hand, it is a trivial matter to add arbitrary (finite) number of discrete components to $X^D$.

Condition B1.2 is quite weak despite the value of $\chi$ can be large.

The assumptions of B1.3, B1.4 and B1.5 are standard in the kernel smoothing literature using second order kernel.

Here in B1.6 we use the bandwidth with the optimal MSE rate for a regular 1-dimensional nonparametric estimates.

The trimming factor in B1.7 provides the necessary treatment of the boundary effects. This would ensure all the uniform convergence results on the expanding compact subset $\{X_T\}$ whose limit is $X$. In practice we will want to minimize the trimming out of the data, we can choose $c_T$ close enough to $h_T$ to do this.

Condition B1.8 is not necessary for consistency and asymptotic normality for any of the parameters below, only M1.3 is required. In particular, B1.8 yields us the simple multinomial logit form, (7), that is often used in practice. For other distribution will result in the use of a more complicated inversion map, for example see Pesendorfer and Schmidt-Dengler (2003) for the Gaussian case.

Next we provide pointwise distribution theory for the nonparametric estimators obtained from the first stage, as described in Section 3, for any given set of values of the structural parameters. The bias and the variance terms are complicated, the explicit formulae can be found along with all proofs in the Appendix.

THEOREM 1.1. *Suppose B1.1 – B1.8 hold. Then for each $\theta \in \Theta$ and $x \in int(X)$, there exists deterministic functions $\eta_{m,\theta}$ and $\omega_{m,\theta}$ such that*

$$T^{2/5}\left(\widehat{m}_\theta(x) - m_\theta(x) - \frac{1}{2}\mu_2 h_T^2 \eta_{m,\theta}(x)\right) \Longrightarrow \mathcal{N}(0, \omega_{m,\theta}(x)),$$

*where $\widehat{m}_\theta(x)$ is defined as in (19) and*

$$\begin{aligned} \eta_{m,\theta}(x) &= (I - \mathcal{L})^{-1}(\eta_{r,\theta} + \eta_{\mathcal{L},\theta})(x), \\ \omega_{m,\theta}(x) &= \frac{\kappa_2}{f_X(x)}\left(\beta^2 \mathrm{var}(m_\theta(x_{t+1})|x_t = x) + \omega_{r,\theta}(x)\right). \end{aligned}$$

*Some components of the bias and variance are complicated, in particular the explicit form of $\eta_{r,\theta}$, $\eta_{\mathcal{L},\theta}$ and $\omega_{r,\theta}$ can be found in (45),(54) and (46) respectively. $\widehat{m}_\theta(x)$ and $\widehat{m}_\theta(x')$ are also asymptotically independent for any $x \neq x'$. Furthermore,*

$$\sup_{(x,\theta)\in X_T\times\Theta} |\widehat{m}_\theta(x) - m_\theta(x)| = o_p\left(T^{-1/4}\right).$$

The rate of convergence, $T^{-2/5}$, is the usual optimal rate (in the MSE sense) of a 1−dimensional nonparametric function. The above is obtained by using analogous arguments of Linton and Mammen (2005) after showing that the conditional density estimator that define the empirical integral operator converges uniformly (see Masry (1996)) over its domain. Similar to Theorem 1, we also obtain the following results for the estimator of $g_\theta$.

THEOREM 1.2. *Suppose $B1.1 - B1.8$ hold. Then for each $\theta \in \Theta, x \in int(X)$ and $a \in A$,*

$$T^{2/5}\left(\widehat{g}_\theta(a,x) - g_\theta(a,x) - \frac{1}{2}\mu_2 h_T^2 \eta_{g,\theta}(a,x)\right) \Longrightarrow \mathcal{N}\left(0, \omega_{g,\theta}(a,x)\right),$$

*where $\widehat{g}_\theta(a,x)$ is defined as in (20) and*

$$\begin{aligned}\eta_{g,\theta}(a,x) &= \mathcal{H}(I-\mathcal{L})^{-1}\left(\eta_{r,\theta} + \eta_{\mathcal{L},\theta}\right)(a,x) + \eta_{\mathcal{H},\theta}(a,x),\\ \omega_{g,\theta}(a,x) &= \frac{\kappa_2}{f_{X,A}(x,a)}\mathrm{var}\left(m_\theta(x_{t+1})\,|x_t = x, a_t = a\right).\end{aligned}$$

*The explicit form of $\eta_{r,\theta}$, $\eta_{\mathcal{L},\theta}$ and $\eta_{\mathcal{H},\theta}$ can be found in (45),(54) and (55) respectively. $\widehat{g}_\theta(a,x)$ and $\widehat{g}_\theta(a',x')$ are also asymptotically independent for any $x \neq x'$ and any $a$. Furthermore,*

$$\sup_{(x,a,\theta)\in X_T\times A\times\Theta} |\widehat{g}_\theta(a,x) - g_\theta(a,x)| = o_p\left(T^{-1/4}\right).$$

We end with a brief discussion of the change in primitives required to accommodate the case when the dimension of $X^C$ is higher than 1. Clearly, using the optimal (MSE) rates for $h_T$, $\dim\left(X^C\right)$ cannot exceed 3 with second order kernel if we were to have the uniform rate of convergence for our nonparametric estimates to be faster than $T^{-1/4}$ that is necessary for $\sqrt{T}-$consistency of the finite dimensional parameters. It is possible to overcome this by exploiting additional smoothness (if available) of our densities. This can be done by using higher order kernels to control the order of the bias, for details of their constructions and usages see Robinson (1988) and also Powell, Stock and Stoker (1989).

FINITE DIMENSIONAL PARAMETERS:

We first provide the notation for the objective functions:

$$Q_T(\theta) = \frac{1}{T}\sum_{t=1}^{T} q\left(a_t, x_t; \theta, g_\theta\right); \quad \widehat{Q}_T(\theta) = \frac{1}{T}\sum_{t=1}^{T} \mathfrak{c}_{t,T}\, q\left(a_t, x_t; \theta, \widehat{g}_\theta\right),$$

where $q$ denotes the log-likelihood function, $\widehat{Q}_T$ is the feasible criterion function, $Q_T$ is identical to $\widehat{Q}_T$ when the infinite dimensional component $\widehat{g}_\theta = g_\theta$. Here, $\{\mathfrak{c}_{t,T}\}$ is a triangular array of trimming factors, cf. B1.7. Define also the limiting objective function $Q(\theta) = \lim_{T\to\infty} EQ_T(\theta)$, which is assumed to exist. We define our estimator for the finite dimensional structural parameters, $\widehat{\theta}$, to be any sequence that satisfy the following inequality,

$$\widehat{Q}_T(\widehat{\theta}) \geq \sup_{\theta\in\Theta} \widehat{Q}_T(\theta) - o_p\left(T^{-1/2}\right).$$

In order to obtain consistency result and the parametric rate of convergence for $\widehat{\theta}$, we need to adjust some assumptions described in the previous subsection and add an identification assumption. Consider:

B1.6′ *The bandwidth sequence $h_T$ satisfies $h_T = \gamma_1(T)\, T^{-1/4}/\log T$ and $\gamma_1(T)$ bounded* *away from zero and infinity.*

B1.9 *The value $\theta_0 \in int(\Theta)$ is defined by, for any $\varepsilon > 0$*

$$\sup_{\|\theta-\theta_0\|\geq\varepsilon} Q(\theta_0) - Q(\theta) > 0.$$

The rate of undersmoothing (relative to B1.6) in Condition B1.6′ ensures that the bias from the nonparametric estimation disappears sufficiently quickly to obtain parametric rate of convergence for $\widehat{\theta}$. To accommodate for higher dimension of $X^C$, we generally cannot just proceed by undersmoothing but combining this with the use higher order kernels, again, see Robinson (1988) and also Powell, Stock and Stoker (1989).

Condition B1.9 assumes the identification of the parametric part. This is a high level assumption that might not be easy to verify due to the complication with the value function. In practice we will have to check for local maxima for robustness. We note that this is the only assumption concerning the criterion function, for other type of objective functions, obvious analogous identification conditions will be required.

The properties of $\widehat{\theta}$ can be obtained by application of the asymptotic theory for semiparametric profile estimators. This requires uniform expansion $\widehat{g}_\theta$ (and hence $\widehat{m}_\theta$) and their derivatives with respect to $\theta$.

THEOREM 1.3. *Suppose $B1.1 - B1.5, B1.6'$ and $B1.7 - B1.9$ hold. Then*

$$\sqrt{T}\left(\widehat{\theta} - \theta_0\right) \Longrightarrow \mathcal{N}\left(0, \mathcal{J}^{-1}\mathcal{I}\mathcal{J}^{-1}\right),$$

*where $\mathcal{I}$ is a complicated term representing the asymptotic variance of the leading terms in*

$\frac{1}{\sqrt{T}}\sum_{t=1}^{T}\frac{\partial q\left(a_t,x_t;\theta_0,\widehat{g}_{\theta_0}\right)}{\partial\theta}$ *and*

$$\mathcal{J}=E\left[\frac{\partial^2 q\left(a_t,x_t;\theta_0,g_{\theta_0}\right)}{\partial\theta\partial\theta'}\right].$$

The root-$T$ rate of convergence is common for such semiparametric estimators when the dimension of the continuous component of $X$ is not too large under some smoothness assumptions. We note that, unlike $\widehat{m}_\theta$ and $\widehat{g}_\theta$, the asymptotic variance of $\widehat{\theta}$ is rather complicated.

And finally we have,

THEOREM 1.4. *Suppose* $B1.1-B1.5, B1.6'$ and $B1.7-B1.9$ *hold. Then for any arbitrary estimator* $\widetilde{\theta}$ *such that* $||\widetilde{\theta}-\theta_0||=O_p(1)$ *and* $x\in int(X)$,

$$\sqrt{Th_T}\left(\widehat{m}_{\widetilde{\theta}}(x)-m_{\theta_0}(x)-\frac{1}{2}\mu_2h_T^2\eta_{m,\theta_0}(x)\right)\Longrightarrow\mathcal{N}\left(0,\omega_{m,\theta_0}(x)\right),$$

*where* $\widehat{m}_\theta,\eta_{m,\theta}$ *and* $\omega_{m,\theta}$ *are defined as those in Theorem 1.1 and,* $\widehat{m}_{\widetilde{\theta}}(x)$ *and* $\widehat{m}_{\widetilde{\theta}}(x')$ *are asymptotically independent for any* $x\neq x'$.

Similarly, for $g_\theta$ we have,

THEOREM 1.5. *Suppose* $B1.1-B1.5, B1.6'$ and $B1.7-B1.9$ *hold. Then for any arbitrary estimator* $\widetilde{\theta}$ *such that* $||\widetilde{\theta}-\theta_0||=O_p(1), x\in int(X)$ *and* $a\in A$,

$$\sqrt{Th_T}\left(\widehat{g}_{\widetilde{\theta}}(a,x)-g_{\theta_0}(a,x)-\frac{1}{2}\mu_2h_T^2\eta_{g,\theta_0}(a,x)\right)\Longrightarrow\mathcal{N}\left(0,\omega_{g,\theta_0}(a,x)\right),$$

*where* $\widehat{g}_\theta,\eta_{g,\theta}$ *and* $\omega_{g,\theta}$ *are defined as those in Theorem 2 and,* $\widehat{g}_{\widetilde{\theta}}(a,x)$ *and* $\widehat{g}_{\widetilde{\theta}}(a',x')$ *are asymptotically independent for any* $x\neq x'$ *and any* $a$.

Given the explicit forms of the bias and variance terms provided in the above

Theorems 1.1 - 1.2 and Theorems 1.4 - 1.5, inference can be conducted using large sample approximation based on obvious plug-in estimators. These expressions are useful as they provide insights into to the variation in the bias and variance of our estimators. However, for the estimator of $\widehat{\theta}$ in Theorem 1.3, due to their complicated form, bootstrap procedures would most likely be preferred in practice. Later in Chapter 2.4, we propose a bootstrap algorithm to estimate a class of dynamic models where the control variable is continuously distributed, the procedure can be readily adapted to estimate a discrete choice problem.

## 1.5 Simulation Study

In this section we illustrate some finite sample properties of our proposed estimator in a small scale Monte Carlo experiment. We replicate the setup of the bus engine replacement problem studied in Rust (1987).

THE RENEWAL PROBLEM:

Consider the decision problem of a manager who owns a bus that operates in each period. The observable state variable $x_t$ is the mileage reading from the machine's odometer. The manager's decision is to decide whether to replace the machine ($a_t = 1$) or maintain it ($a_t = 0$). In addition, the manager also takes into the account the net cost of goodwill gained/lost from her decision whether to take it out of service, which we denote by $(\varepsilon_{1t}, \varepsilon_{2t})$. The manager per period profit

$$u_{0,\theta}(a, x_t, \varepsilon_t) = \begin{cases} -\theta_1 + \varepsilon_{1t} & \text{if } a = 1, \\ -\theta_2 x_t + \varepsilon_{2t} & \text{otherwise.} \end{cases}$$

Therefore $\theta_1$ is the replacement cost of the machine engine while $\theta_2$ reflects the scaling factor for the cost of maintaining the machine, which depends on the machine's odometer reading. The evolution of the mileage $x_t$ is assumed to follow a regenerative random walk, in particular we assume

$$x_{t+1} = \begin{cases} \eta_t & \text{if } a_t = 1, \\ x_t + \eta_t & \text{otherwise,} \end{cases}$$

where $\eta_t \overset{i.i.d.}{\sim} \exp(1)$. This regenerative property is essential for maintaining the stationary structure we need to carry out our estimation procedure.

Given the framework above, we set $(\theta_1, \theta_2) = (5, 0.5)$ and use the fixed point iteration method of Rust (1987) to approximate the true continuation value function necessary to generate the data. Although the support of $x_t$ is the positive half-line we only restrict outselves to $[0, 20]$, and we base our approximation of functions on this interval by using 1000 equally-spaced grid points. We generate 1000 replications of such controlled Markov processes for various sizes of $T \in \{100, 500, 1000, 2500, 5000\}$.

IMPLEMENTATION:

We are interested in obtain estimates for the demand parameters $(\theta_1, \theta_2)$ when only $\{a_t, x_t\}_{t=1}^T$ are observed. In estimating the nonparametric estimator of $g_\theta$, we use a truncated Gaussian kernel with 3 different bandwidths $\{h_\varsigma = 1.06s\,(NT)^{-\varsigma} : \varsigma = \frac{1}{8}, \frac{1}{4}, \frac{3}{8}\}$, where $s$ is the standard deviation of observed $\{x_t\}_{t=1}^T$. Note that the rate of decay for $h_{\frac{1}{4}}$ is as specified in B1.6$'$ to ensure root$-T$ consistent estimation of $\theta$ based on using a second order kernel. For the conditional density of $x_{t+1}$ given $x_t$, we recall that $x_{t+1} - x_t = \eta_t$ so we base our estimation of the conditional density on the density of the random sample $\{\eta_t\}_{t=1}^T$. For estimating the nonparametric conditional choice probabilities we use the Nadaraya-Watson estimator for the regression function of $\mathbf{1}[a_t = 1]$ on $x_t$. And to deal with the boundary issue (from below), we employ the following simple boundary corrected kernel

$$K_h^b(x_t - x) = \begin{cases} \frac{K_h(x_t-x)}{\int_{v=-x/h}^{\infty} K(v)dv} & \text{for } x \in [0,h), \\ K_h(x_t - x) & \text{otherwise,} \end{cases}$$

where $K$ denotes the truncated Gaussian kernel. In addition, due to small number observations for large values of $x$, we set the values for the kernel estimators in the top 5−th percentile (of the observed data) to be equal to obtained from the 95−th percentile.

We also estimate the model by manually discretizing $\{x_t\}_{t=1}^{T}$. We do this by partitioning the support of $X$ by various grid points based on various number of grids ($d$). In particular, for different values of $d$:

- $d = 2, X = [0, 3] \cup [3, \infty)$
- $d = 3, X = [0, 2.5] \cup [2.5, 5] \cup [5, \infty)$
- $d = 4, X = [0, 2] \cup [2, 4] \cup [4, 6] \cup [6, \infty)$
- $d = 5, X = [0, 1.5] \cup [1.5, 3] \cup [3, 4.5] \cup [4.5, 6] \cup [6, \infty)$

The corresponding values for the discrete support is simply the mid-point values when the upper bound of the discretized is finite, otherwise it takes the following values $\{6, 7, 8, 8\}$ for $d = \{2, 3, 4, 5\}$ respectively.

Comments and Results:

The Tables can be found at the end of the chapter. We report the bias, median of the bias, standard deviation and interquartile range (scaled by 1.349) for the estimators of $\theta_1$ and $\theta_2$. For Tables 1 and 2, the rows are arranged according to the total sample size and bandwidths. We have the following general observations for both estimators: (i) the median of the bias is similar to the mean; (ii) the estimators converge to the true values as $N$ increases and their respective standard deviations are converging to zero; (iii) the standard deviation figures are similar to the corresponding scaled interquartile range.[5] However, using the bandwidth that decays at the rate specified in B1.6′ (Section 1.4) as the benchmark, undersmoothing seems to provide marginally better results in the MSE sense whilst the results for oversmoothing actually perform rather poorly as the bias term dominates the MSE, this is somewhat expected since the estimators are biased for the bandwidth with that rate.

We also report analogous summary statistics from using the manually discretized data. As expected, generally, larger support of the discretized state yields relative lower bias while increasing the variance. The estimates for $\theta_1$ and $\theta_2$ indicate that these estimators are typically inconsistent. Although an exception exists for the particular design with 4 discrete state appears to provide estimates that concentrates close to the true $\theta$ (only) for $\theta_2$.

## 1.6 Markovian Games

The development of empirical dynamic games is of recent interest especially in the empirical industrial organization literature. See Ackerberg et al. (2005) for an excellent survey. A class of Markovian games with discrete action and time can be defined by considering a finite set of endogenously linked MDP, whose interactions are to be made precise below. For some examples of the estimation techniques for such games see Aguirregabiria and Mira (2007), Bajari et al. (2008), Pakes et al. (2002), and Pesendorfer and Schmidt-Dengler (2008). Similarly to the single agent MDP, these papers, with the exception of Bajari et al., assume finite observable state space. In this section we discuss how we can use the methodology discussed in previous sections to estimate Markovian games.

[5] (iii) is a characteristic of a normal random variable.

(Cf. Section 1.2.1) For each period $t$ there are $N$ players, indexed by the ordered set $\{i\}$. Each player $i$ is forward looking in solving her intertemporal problem. At each period $t$,each player obtains some information $s_{i,t}$ and chooses an action $a_{i,t}$ in order to maximize her discounted expected utility. The present period utility is time separable and is represented by $u_{i,\theta}^{0}(a_t, s_t)$ for $\theta \in \Theta \subset \mathbb{R}^P$, where $a_t = (a_t, a_{-i,t})$, and $a_{-i,t}$ denotes the actions of all other players except player $i$, $s_t$ is defined analogously. The actions of all players today directly affect their uncertain future information according to the Markovian transition density $p(ds_{t+1}|s_t, a_t)$. The next period utility is subjected to discounting at the rate $\beta_i \in (0,1)$. Briefly, the stationary Markovian games of interest can be defined by assuming that the decision process of each player $i$ is characterized by the following amendments of Conditions M1.1 - M1.4:

M1.1′ *Player $i$ is represented by a triple $(u_i^0, p_i, \beta_i)$ and $\beta_i \in (0,1)$ and $T = \infty$, both are exogenously given and assumed to be known. The observed data $\{x_t, a_{it}\}_{t=1}^{T}$ is the controlled stochastic processes satisfying*

$$V_{i,\theta}^{0}(s_{it};\sigma_i) = \max_{a_i \in A}\left\{E\left[u_{i,\theta}^{0}(a_i, a_{-it}, s_{it})\,|s_{it}, a_i;\sigma_i\right] + \beta_i E\left[V_{i,\theta}^{0}(s_{it+1};\sigma_i)\,|s_{it}, a_i;\sigma_i\right]\right\}. \tag{26}$$

M1.2′ *(Conditional Independence) conditional independence of the state variables*

$$p(dx_{t+1}, d\varepsilon_{t+1}|x_t, \varepsilon_t, \mathbf{a}_t) = q(d\varepsilon_{t+1}|x_{t+1})\, f_{X'|X,\mathbf{A}}(dx_{t+1}|x_t, \mathbf{a})$$

M1.3′ *$s_i = (x, \varepsilon_i) \in X \times \mathbb{R}^K$ where $X = X^C \times X^D$ is a subset of $\mathbb{R}^J$. $X^C$ includes intervals and $X^D$ is finite, they denote the space of public information with continuous and discrete components respectively. $\varepsilon$ is a vector of private information, whose dimension is $K$ , the cardinality of $A$, the number of actions player $i$ can choose in each period.*[6] *The distribution of $\varepsilon_{it}$ is known and is absolutely continuous*

[6] Note that the choice of the notation for the state variables $(x, s)$ is consistent with that of Rust (1988) but the role of $x$ and $s$ is some times reveresed in some game papers.

*with respect to some Lebesgue measure with Radon Nikodym density* $q_i(d\varepsilon_{it}|x_t)$

M1.4′ *(Additive Separability) per period payoff function* $u^0_{i,\theta}: A^N \times S \to \mathbb{R}$ *is specified upto some unknown parameters* $\theta \in \Theta$ *and is additive separable w.r.t. unobservable state variables,* $u^0_{i,\theta}(a_{it}, a_{-it}, s_{it}) = u_{i,\theta}(a_{it}, a_{-it}, x_{it}) + \varepsilon_{ia_t,t}$.

M1.5′ *(Private Values)* $\varepsilon_t$ *is also jointly independent across all players, i.e.* $q(\varepsilon_t) = \prod_{i=1}^N q_i(\varepsilon_t)$.

The immediate observation on the Conditions M1.1′ - M1.4′ reveals that, for each player $i$, the controlled process $\{a_{it}, x_t\}_{t=1}^T$ only differs from the single agent case in that the per period payoff function and transition densities are affected by other players' actions, and each player forms an expectation, in (26), using the beliefs she has over the distribution of other players' actions. Although the private value assumption limits the applicability of our estimator, most of the existing literature on dynamic estimation of the same class of models make use of this assumption. We denote the distribution of beliefs of player $i$ over $a_{-i}$ by $\sigma_i$. We define the equilibrium concept through the notion of best response. The best response function is a map $\alpha^0_{i,\theta}: S \to A$ defined by

$$\alpha^0_{i,\theta}(s_{it};\sigma_i) = a_{it} \Leftrightarrow v^0_{i,\theta}(a_{it}, s_{it};\sigma_i) + \varepsilon_{ia_t,t} \geq v^0_{i,\theta}(a_i, s_{it};\sigma_i) + \varepsilon_{ia,t} \text{ for all } a_i \in A, \quad (27)$$

where $v^0_{i,\theta}$ computes the expected utility given any action and states for agent $i$ given her belief $\sigma_i$, namely the choice-specific continuation value, for any $(s_{it}, a_{it}, \sigma_i)$

$$v^0_{i,\theta}(a_{it}, s_{it};\sigma_i) = E\left[u_{i,\theta}(a_{it}, a_{-it}, x_{it})\,|s_{it}, a_{it};\sigma_i\right] + \beta_i E\left[V^0_{i,\theta}(s_{it+1};\sigma_i)\,|s_{it}, a_{it};\sigma_i\right]. \quad (28)$$

**Definition 1** *(Markov Perfect Equibrium) A collection* $(\mathbf{a}, \boldsymbol{\sigma}) = (a_1, \ldots, a_N, \sigma_1, \ldots, \sigma_N)$ *is a Markov Perfect Equibrium if*

(i) *for all $i$, $a_i$ is a best response to $a_{-i}$ given the beliefs $\sigma_i$ at any state $x_i$*

(ii) *all players use Markovian strategies*

(iii) *for all $i$, the beliefs $\sigma_i$ are consistent with the strategies $a_{-i}$*

See Maskin and Tirole (2001) for more details. As shown in Pesendorfer and Schmidt-Dengler (2008), we can also characterize our equilibrium concept through the choice probabilities conditioning on the public information.[7] These choice probabilities can be obtained from the definition of best response (27) by marginalizing out the private information of all the other players,

$$P_i^0\left(a_{it}|x_t;\theta,\sigma_i\right) = \int \mathbf{1}\left[\alpha_{i,\theta}^0\left(s_{it};\sigma_i\right) = a_{it}\right] q_i\left(d\varepsilon_{it}|x_t\right). \tag{29}$$

Collecting the choice probabilities in (29), for each player and across the action space, we have a map $\Psi : \mathcal{C} \to \mathcal{C}$, where $\mathcal{C}$ is a space of an $N \cdot K-$ vector of functions with each function mapping $X$ to $[0,1]$. This nonlinear operator maps the beliefs of the players into the choice probabilities of their actions determined as a consequence of their beliefs:

$$\mathbf{p} = \Psi_\theta\left(\sigma\right).$$

Assuming we observe equilibrium play $\mathbf{p}$ must be a fixed point of $\Psi_\theta$. This equilibrium condition, which is a set of conditional moment restrictions, leads to a class of minimum distance type criterion functions that we can use to estimate the structural parameters as discussed in Section 1.4. See Pesendorfer and Schmidt-Dengler for a thorough treatment of this idea when $X$ is finite. Dropping the dependence on $\sigma_i$, as seen from (27), (28) and (29), the computation of the choice probabilities will again depend on

[7] There are some fixed point theorems available for the case of infinite dimensional spaces, for example, Schauder fixed point theorem.

the model implied value functions $E\left[V_{i,\theta}\left(s_{it+1}\right)|s_{it},a_{it}\right]$, similar to (4), where

$$V_{i,\theta}\left(s_{it}\right)=E\left[u_{i,\theta}^{0}\left(\mathbf{a}_{t},s_{it}\right)|s_{it}\right]+\beta_{i}E\left[V_{i,\theta}\left(s_{it+1}\right)|s_{it}\right].$$

Therefore, fundamentally, the practical aspect of the estimation problem is essentially the same as in the single agent case. Our strategy will be the same as before. For each player $i$, by marginalizing out the private information of all the other players of (26) in the model implied equilibrium (cf. (8)), we have the generalized PVE

$$\begin{aligned} E\left[V_{i,\theta}\left(s_{it}\right)|x_{t}\right] &= E\left[E\left[u_{0,i,\theta}\left(\mathbf{a}_{t},s_{it}\right)|s_{it}\right]|x_{t}\right]+\beta_{i}E\left[E\left[V_{i,\theta}\left(s_{it+1}\right)|s_{it}\right]|x_{t}\right] \quad (30)\\ &= E\left[u_{0,i,\theta}\left(\mathbf{a}_{t},s_{it}\right)|x_{t}\right]+\beta_{i}E\left[E\left[V_{i,\theta}\left(s_{it+1}\right)|x_{t+1}\right]|x_{t}\right]. \end{aligned}$$

As seen previously, for each $i$, (8) can be expressed as,[8]

$$m_{i,\theta}=r_{i,\theta}+\mathcal{L}_{i}m_{i,\theta},$$

where $(m_{i,\theta},r_{i,\theta},\mathcal{L}_{i})$ have, by now, obvious meaning. By the same arguments used in Section 1.2.4, $\{\mathcal{L}_{i}\}_{i=1}^{N}$ will be a sequence of contraction maps. Therefore these integral equations can be estimated and solved independently. Hence we only need to approximate the operator $(I-\mathcal{L}_{i})^{-1}$ once for each player, where the only difference between the kernels of the integral operator of different players is their discounting factors. Note that there are some additional smoothings required for the per period payoff since, unlike the single agent case, players make decisions based on expected payoffs, see (28) and (30). For a direct comparison with the single agent case, under M1.2′ - M1.3′, we

[8]This is a direct generalization of the matrix equation ((6)) of Pesendorfer and Schmidt-Dengler (2008).

can write the expected choice-specific continuation value (28) as

$$v_{i,\theta}(a_{it}, x_t) = E\left[u_{i,\theta}(a_{it}, a_{-it}, x_t)|x_t, a_{it}\right] + \beta_i E\left[m_{i,\theta}(x_{t+1})|x_t, a_{it}\right].$$

This can be written in a linear functional notation (cf. (11))

$$v_{i,\theta} = \mathcal{K}_i u_{i,\theta} + \beta_i \mathcal{H}_i m_{i,\theta},$$

where $\mathcal{K}_i$ and $\mathcal{L}_i$ denote linear operators where the former is a conditional expectation operator of $a_{-it}$ given $(x_t, a_{it})$ and the latter is a conditional expectation on $x_{t+1}$ given $(x_t, a_{it})$.

Therefore we can use the model implied continuation value functions to construct the model implied best response and choice probabilities respectively

$$\begin{aligned}
\alpha_{i,\theta}(s_{it}) &= a_{it} \Leftrightarrow v_{i,\theta}(a_{it}, s_{it}) + \varepsilon_{ia_t,t} \geq v_{i,\theta}(a_i, s_{it}) + \varepsilon_{ia,t} \text{ for all } a_i \in A, \\
P_i(a_{it}|x_t; \theta, \sigma_i) &= \int \mathbf{1}\left[\alpha_{i,\theta}(s_{it}; \sigma_i) = a_{it}\right] q_i(d\varepsilon_{it}|x_t),
\end{aligned}$$

which can be used to define an analogous two-step semiparametric method as discussed in Section 1.3 and 1.4.

## 1.7 Conclusion

In this chapter, we provide a method to estimate a class of Markov decision processes that allows for continuous observable state space. The type of primitive conditions are provided the inference of the finite and infinite dimensional parameters in the model. Our estimation technique relies on the convenient well-posed linear inverse problem presented by the policy value equation. It inherits the computational simplicity of Pesendorfer and Schmidt-Dengler that is independent of the parameterization of the per period utility function. We also illustrate how this method can be extended naturally to the estimation of Markovian games in a similar setting to that of Bajari et al. Their identification results directly apply here.

There are some practical aspects of our estimators worth exploring. Although we performed some limited Monte Carlo experiment, it would be interesting to try to study the role of numerical error brought upon by approximating the integral in the case that we have large sample size compared to the purely nonparametric approximation. Second is to see how our estimator performs in practice relative to more extensive discretization schemes. Thirdly, some efficiency bounds should be obtainable in the special case of conditional logit assumption.

## 1.8 Proofs of Theorems

In this section, we provide a set of high level assumptions (A1.1 - A1.6) and their consequences (C1.1 - C1.4) of the nonparametric estimators described in Section 3. We outline the stochastic expansions required to obtain the asymptotic properties of $\widehat{m}_\theta$ and $\widehat{g}_\theta$. The high level assumptions are then proved under the primitives of M1.1 - M1.4 and B1.1 - B1.8. Consequences are simple and their proofs are omitted. In what follows, we refer frequently to Bosq (1998), Linton and Mammen (2005), Masry (1996) and Robinson (1983), so for brevity, we denote their references by [B], [LM], [M] and [R] respectively.

### 1.8.1 Outline of Asymptotic Approach

For notational simplicity we work on a Banach space, $(C(\mathcal{X}), \|\cdot\|)$, where $\mathcal{X} = \mathcal{X}^C \times X^D$ ,the continuous part of $\mathcal{X}$ is a compact set $[\underline{x}+\varepsilon, \overline{x}-\varepsilon]$ for some arbitrarily small $\varepsilon > 0$. We denote $B1.1'$, the analogous condition to B1.1 when we replace $X$ by $\mathcal{X}$. The approach taken here is similar to [LM], who worked on the $L^2$ Hilbert Space. The main difference between our problem and theirs is, after getting consistent estimates of (9), we require another level of smoothing (10) before plugging it into the criterion function. The first part here follows [LM].

ASSUMPTION A1.1. Suppose that for some sequence $\delta_T = o(1)$:

$$\sup_{x \in \mathcal{X}} \left| \left( \widehat{\mathcal{L}} - \mathcal{L} \right) m(x) \right| = o_p(\delta_T),$$

i.e., we have,

$$\left\| \left( \widehat{\mathcal{L}} - \mathcal{L} \right) m \right\| = o_p(\delta_T),$$

for any $m \in C(\mathcal{X})$.

CONSEQUENCE C1. Under A1.1:

$$\left\| \left( \left(I - \widehat{\mathcal{L}}\right)^{-1} - (I - \mathcal{L})^{-1} \right) m \right\| = o_p\left(\delta_T\right).$$

The rate of uniform approximation of the linear operator gets transferred to the inverse of $(I - \mathcal{L})$. This is summarized by C1.1 and is proven in [LM].

We supposed that $\widehat{r}_\theta(x) - r_\theta(x)$ can be decomposed into the following terms with some properties.

ASSUMPTION A1.2. For each $x \in$X:

$$\widehat{r}_\theta(x) - r_\theta(x) = \widehat{r}_\theta^B(x) + \widehat{r}_\theta^C(x) + \widehat{r}_\theta^D(x), \tag{31}$$

where $\widehat{r}_\theta^C, \widehat{r}_\theta^D$ and $\widehat{r}_\theta^E$ satisfy:

$$\begin{aligned}
\sup_{(x,\theta)\in\mathcal{X}\times\Theta} \left|\widehat{r}_\theta^B(x)\right| &= O_p\left(T^{-2/5}\right) \text{ with } \widehat{r}_\theta^B \text{ deterministic,} && (32)\\
\sup_{(x,\theta)\in\mathcal{X}\times\Theta} \left|\widehat{r}_\theta^C(x)\right| &= o_p\left(T^{-2/5+\xi}\right) \text{ for any } \xi > 0, && (33)\\
\sup_{(x,\theta)\in\mathcal{X}\times\Theta} \left|\mathcal{L}(I - \mathcal{L})^{-1}\widehat{r}_\theta^C(x)\right| &= o_p\left(T^{-2/5}\right), && (34)\\
\sup_{(x,\theta)\in\mathcal{X}\times\Theta} \left|\widehat{r}_\theta^D(x)\right| &= o_p\left(T^{-2/5}\right). && (35)
\end{aligned}$$

This is the standard bias+variance+remainder of local constant kernel estimates of the regression function under some smoothness assumptions. The intuition behind (34), as provided in [LM], is that the operator applies averaging to a local smoother and transforms it into a global average thereby reducing its variance. These terms are used to obtain the components of $\widehat{m}_\theta(x)$, for $j = B, C, D$, the terms $\widehat{m}_\theta^j(x)$ are solutions to the integral equations,

$$\widehat{m}_\theta^j = \widehat{r}_\theta^j + \widehat{\mathcal{L}}\widehat{m}_\theta^j \tag{36}$$

and $\widehat{m}_\theta^A$, from writing the solution $m_\theta + \widehat{m}_\theta^A$ to the integral equation

$$\left(m_\theta + \widehat{m}_\theta^A\right) = r_\theta + \widehat{\mathcal{L}}\left(m_\theta + \widehat{m}_\theta^A\right). \tag{37}$$

The existence and uniqueness of the solutions to (36) and (37) are assured, at least w.p.a. 1, under thet contraction property of the integral operator, so it follows from the linearity of $\left(I - \widehat{\mathcal{L}}\right)^{-1}$ that

$$\widehat{m}_\theta = m_\theta + \widehat{m}_\theta^A + \widehat{m}_\theta^B + \widehat{m}_\theta^C + \widehat{m}_\theta^D.$$

These components can be approximated by simpler terms. Define also $m_\theta^B$, as the solution to

$$m_\theta^B = \widehat{r}_\theta^B + \mathcal{L}m_\theta^B. \tag{38}$$

CONSEQUENCE C1.2. Under A1.1 - A1.2:

$$\sup_{(x,\theta)\in\mathcal{X}\times\Theta} \left|\widehat{m}_\theta^B(x) - m_\theta^B(x)\right| = o_p\left(T^{-2/5}\right), \tag{39}$$

$$\sup_{(x,\theta)\in\mathcal{X}\times\Theta} \left|\widehat{m}_\theta^C(x) - r_\theta^C(x)\right| = o_p\left(T^{-2/5}\right), \tag{40}$$

$$\sup_{(x,\theta)\in\mathcal{X}\times\Theta} \left|\widehat{m}_\theta^D(x)\right| = o_p\left(T^{-2/5}\right). \tag{41}$$

(39) and (41) follow immediately from (32), (35) and C1.1. (40) follows from (34), A1.1 and C1.1, as we can easily show that,

$$\left\|\widehat{\mathcal{L}}\left(I - \widehat{\mathcal{L}}\right)^{-1} - \mathcal{L}\left(I - \mathcal{L}\right)^{-1}\right\| = o_p\left(\delta_T\right).$$

We next, also, approximate $\widehat{m}_\theta^A$ by simpler terms, subtracting (9) from (37) yields

$$\widehat{m}_\theta^A = \left(\widehat{\mathcal{L}} - \mathcal{L}\right) m_\theta + \widehat{\mathcal{L}}\widehat{m}_\theta^A. \tag{42}$$

Assumption A1.3. For any $x \in \mathcal{X}$:

$$\left(\widehat{\mathcal{L}}-\mathcal{L}\right) m_\theta\left(x\right)=\widehat{r}_\theta^E\left(x\right)+\widehat{r}_\theta^F\left(x\right)+\widehat{r}_\theta^G\left(x\right),$$

where $\widehat{r}_\theta^E, \widehat{r}_\theta^F$ and $\widehat{r}_\theta^G$ satisfy:

$$\begin{aligned}
\sup_{(x,\theta)\in\mathcal{X}\times\Theta}\left|\widehat{r}_\theta^E\left(x\right)\right| &= O_p\left(T^{-2/5}\right) \text{ with } \widehat{r}_\theta^E \text{ deterministic},\\
\sup_{(x,\theta)\in\mathcal{X}\times\Theta}\left|\widehat{r}_\theta^F\left(x\right)\right| &= o_p\left(T^{-2/5+\xi}\right) \text{ for any } \xi>0,\\
\sup_{(x,\theta)\in\mathcal{X}\times\Theta}\left|\mathcal{L}\left(I-\mathcal{L}\right)^{-1}\widehat{r}_\theta^F\left(x\right)\right| &= o_p\left(T^{-2/5}\right),\\
\sup_{(x,\theta)\in\mathcal{X}\times\Theta}\left|\widehat{r}_\theta^G\left(x\right)\right| &= o_p\left(T^{-2/5}\right).
\end{aligned}$$

These terms are obtained by decomposing the conditional density estimates (cf. A1.2), then proceed as done previously, we define $\widehat{m}_\theta^j\left(x\right)$ for $j=E,F,G$ as the unique solutions to the estimated integral equation of (36), solving (42) we have,

$$\begin{aligned}
\widehat{m}_\theta^A &= \left(I-\widehat{\mathcal{L}}\right)^{-1}\left(\widehat{\mathcal{L}}-\mathcal{L}\right)m_\theta\\
&= \widehat{m}_\theta^E+\widehat{m}_\theta^F+\widehat{m}_\theta^G.
\end{aligned}$$

Such terms are asymptotically equivalent to more convenient terms (cf. C1.2), define also $m_\theta^E$ as the solution to the analogous integral equation of (38).

Consequence C1.3. Under A1.1 - A1.3:

$$\begin{aligned}
\sup_{(x,\theta)\in\mathcal{X}\times\Theta}\left|\widehat{m}_\theta^E\left(x\right)-m_\theta^E\left(x\right)\right| &= o_p\left(T^{-2/5}\right),\\
\sup_{(x,\theta)\in\mathcal{X}\times\Theta}\left|\widehat{m}_\theta^F\left(x\right)-\widehat{r}_\theta^F\left(x\right)\right| &= o_p\left(T^{-2/5}\right),\\
\sup_{(x,\theta)\in\mathcal{X}\times\Theta}\left|\widehat{m}_\theta^G\left(x\right)\right| &= o_p\left(T^{-2/5}\right).
\end{aligned}$$

C1.3 can be shown using the same reasonings used to obtain C1.2.

Combining these assumptions leads to Proposition 1 of [LM].

PROPOSITION 1.1. Suppose that [A1.1 - A1.3] holds for some estimators $\widehat{r}_\theta$ and $\widehat{\mathcal{L}}$. Define $\widehat{m}_\theta$ as any solution of $\widehat{m}_\theta = \widehat{r}_\theta + \widehat{\mathcal{L}}\widehat{m}_\theta$. Then the following expansion holds for $\widehat{m}_\theta$

$$\sup_{(x,\theta)\in\mathcal{X}\times\Theta} \left|\widehat{m}_\theta(x) - m_\theta(x) - m_\theta^B(x) - m_\theta^E(x) - \widehat{r}_\theta^C(x) - \widehat{r}_\theta^F(x)\right| = o_p\left(T^{-2/5}\right),$$

where all of the terms above have been defined previously.

The uniform expansion for the nonparametric estimators required in [LM] ends here. However, to obtain the uniform expansion of $\widehat{g}_\theta$ defined in (20), we need another level of smoothing. Note that the integral operator, $\mathcal{H}$, has a different range,

$$\mathcal{H} : C(\mathcal{X}) \to C(A \times \mathcal{X}),$$

where $C(A \times \mathcal{X})$ denotes a space of functions, say $g(a, x)$, which are continuous on X for each $a \in A$. So the relevant Banach Space is equipped with the sup-norm over $A \times \mathcal{X}$, which we also denote by $\|\cdot\|$ though this should not lead to any confusion. For notational simplicity, we first define,

$$\begin{aligned}
\overline{m}_\theta^B(x) &= m_\theta^B(x) + m_\theta^E(x), \\
\overline{m}_\theta^C(x) &= \widehat{r}_\theta^C(x) + \widehat{r}_\theta^F(x), \\
\overline{m}_\theta^D(x) &= \widehat{m}_\theta(x) - m_\theta(x) - \overline{m}_\theta^B(x) - \overline{m}_\theta^C(x).
\end{aligned}$$

We next define various components of the transformations (20), analogously to (36) and (38), for $j = B, C, D$ the terms $\widehat{g}_\theta^j$ are elements of the integral transform,

$$\begin{aligned}\widehat{g}_\theta^j &= \widehat{\mathcal{H}}\overline{m}_\theta^j,\\ g_\theta^j &= \mathcal{H}\overline{m}_\theta^j,\end{aligned}$$

and $\widehat{g}_\theta^A$ is defined by

$$\widehat{\mathcal{H}}m_\theta = g_\theta + \widehat{g}_\theta^A.$$

It follows from linearity of $\widehat{\mathcal{H}}$ that

$$\widehat{g}_\theta = g_\theta + \widehat{g}_\theta^A + \widehat{g}_\theta^B + \widehat{g}_\theta^C + \widehat{g}_\theta^D.$$

ASSUMPTION A1.4. Suppose that for some sequence $\delta_T$ as in A1.1:

$$\sup_{(a,x,\theta)\in A\times\mathcal{X}\times\Theta}\left|\left(\widehat{\mathcal{H}} - \mathcal{H}\right) m_\theta (a, x)\right| = o_p (\delta_T),$$

i.e., $\left\|\left(\widehat{\mathcal{H}} - \mathcal{H}\right) m\right\| = o_p (\delta_T)$ for any $m \in C (\mathcal{X})$.

A1.4 assumes the desirable properties of the conditional density estimators (cf. A1.1 and A1.3).

CONSEQUENCE C1.4. Under A1.1 - A1.4:

$$\begin{aligned}\sup_{(a,x,\theta)\in A\times\mathcal{X}\times\Theta}\left|\widehat{g}_\theta^B (a, x) - g_\theta^B (a, x)\right| &= o_p\left(T^{-2/5}\right),\\ \sup_{(a,x,\theta)\in A\times\mathcal{X}\times\Theta}\left|\widehat{g}_\theta^C (a, x) - g_\theta^C (a, x)\right| &= o_p\left(T^{-2/5}\right),\\ \sup_{(a,x,\theta)\in A\times\mathcal{X}\times\Theta}\left|\widehat{g}_\theta^D (a, x)\right| &= o_p\left(T^{-2/5}\right).\end{aligned}$$

This follows immediately from A1.5 and the properties of the elements defined in $\overline{m}_\theta^B (x)$.

Assumption A1.5. Suppose that:

$$\sup_{(a,x,\theta)\in A\times\mathcal{X}\times\Theta}\left|g_{\theta}^{C}(a,x)\right| = o_p\left(T^{-2/5}\right).$$

A1.5 follows since the operator $\mathcal{H}$ is a global smooth, hence it reduces the variance of $g_{\theta}^{C}$.

As with $\widehat{m}_{\theta}^{A}$ we can approximate $\widehat{g}_{\theta}^{A}$ by simpler terms.

Assumption A1.6. For any $m \in C(\mathcal{X})$ and for each $(a,x) \in A \times \mathcal{X}$:

$$\begin{aligned}\widehat{g}_{\theta}^{A}(a,x) &= \left(\widehat{\mathcal{H}}-\mathcal{H}\right)m_{\theta}(a,x)\\ &= \widehat{g}_{\theta}^{E}(a,x)+\widehat{g}_{\theta}^{F}(a,x)+\widehat{g}_{\theta}^{G}(a,x),\end{aligned}$$

where $\widehat{g}_{\theta}^{E}, \widehat{g}_{\theta}^{F}$ and $\widehat{g}_{\theta}^{G}$ satisfy:

$$\begin{aligned}\sup_{(a,x,\theta)\in A\times\mathcal{X}\times\Theta}\left|\widehat{g}_{\theta}^{E}(a,x)\right| &= O_p\left(T^{-2/5}\right) \text{ with } \widehat{g}_{\theta}^{E} \text{ deterministic,}\\ \sup_{(a,x,\theta)\in A\times\mathcal{X}\times\Theta}\left|\widehat{g}_{\theta}^{F}(a,x)\right| &= o_p\left(T^{-2/5+\xi}\right) \text{ for any } \xi>0,\\ \sup_{(a,x,\theta)\in A\times\mathcal{X}\times\Theta}\left|\widehat{g}_{\theta}^{G}(a,x)\right| &= o_p\left(T^{-2/5}\right).\end{aligned}$$

A1.6 follows from standard decomposition of the kernel conditional density estimator (cf. A1.3).

Proposition 1.2. Suppose that A1.1 - A1.6 holds for some estimators $\widehat{r}_{\theta}$, $\widehat{\mathcal{L}}$ and $\widehat{\mathcal{H}}$. Define $\widehat{m}_{\theta}$ as any solution of $\widehat{m}_{\theta} = \widehat{r}_{\theta} + \widehat{\mathcal{L}}\widehat{m}_{\theta}$ and $\widehat{g}_{\theta} = \widehat{\mathcal{H}}\widehat{m}_{\theta}$. Then the following expansion holds for $\widehat{g}_{\theta}$

$$\sup_{(a,x,\theta)\in A\times\mathcal{X}\times\Theta}\left|\widehat{g}_{\theta}(a,x)-g_{\theta}(a,x)-g_{\theta}^{B}(a,x)-\widehat{g}_{\theta}^{E}(a,x)-\widehat{g}_{\theta}^{F}(a,x)\right| = o_p\left(T^{-2/5}\right),$$

where all of the terms above have been defined previously, in particular $g_{\theta}^{B}$ and $\widehat{g}_{\theta}^{E}$ are non-stochastic and the leading variance terms is $\widehat{g}_{\theta}^{F}$. This can be rewritten in a similar notation to previously:

$$\begin{aligned}
\overline{g}_{\theta}^{B}(a,x) &= g_{\theta}^{B}(a,x)+\widehat{g}_{\theta}^{E}(a,x), \\
\overline{g}_{\theta}^{C}(a,x) &= \widehat{g}_{\theta}^{F}(a,x), \\
\overline{g}_{\theta}^{D}(a,x) &= \widehat{g}_{\theta}(a,x)-g_{\theta}(a,x)-\overline{g}_{\theta}^{B}(a,x)-\overline{g}_{\theta}^{C}(a,x).
\end{aligned}$$

### 1.8.2 Proofs of Theorems 1.1 - 1.2

We assume $B1.1'$ and B1.2 - B1.6 throughout this subsection. Set $\delta_T = T^{\xi-3/10}$, this rate is arbitrarily close to the rate of convergence of 1-dimensional nonparametric density estimates when $h_T$ decays at the rate specified by B1.6. For the ease of notation, we assume that $X^D$ is empty. The presence of discrete states do not affect any of the results below, we can simply replace any formula involving the density (and analogously for the conditional density) $f(dx_t)$ by $f\left(dx_t, x_t^d\right)$. We shall denote generic constants by $C_0$ that may take different values in different places. The uniform rate of convergence proof of various components utilize some exponential inequalities found in [B] as done in [LM], the details are deferred to Section 1.7.8.

Proof of Theorem 1.1. We proceed by providing the pointwise distribution theory for $\widehat{P}(a|x)$, for any $a \in A$ and $x \in int(X)$, and the functionals thereof. These are used to proof Theorem 1 and 2 and verify the high level conditions. $\widehat{P}(a|x)$ is the usual local constant regression estimator (or equivalently, the conditional probability estimator).

$$\widehat{P}(a|x)-P(a|x)=\frac{1}{T}\sum_{t=1}^{T}\left(\mathbf{1}\left[a_t=a\right]-P(a|x)\right)K_h\left(x_t-x\right)/\widehat{f}_X(x),$$

focusing on the numerator

$$\begin{aligned}
\frac{1}{T}\sum_{t=1}^{T}\left(\mathbf{1}\left[a_t=a\right]-P\left(a|x\right)\right)K_h\left(x_t-x\right) &= \frac{1}{T}\sum_{t=1}^{T}\left(P\left(a|x_t\right)-P\left(a|x\right)\right)K_h\left(x_t-x\right) \\
&\quad +\frac{1}{T}\sum_{t=1}^{T}e_{a,t}K_h\left(x_t-x\right) \\
&= A_{1,a,T}\left(x\right)+A_{2,a,T}\left(x\right),
\end{aligned}$$

where $e_{a,t}=\mathbf{1}\left[a_t=a\right]-P\left(a|x_t\right)$. The term $A_{1,a,T}\left(x\right)$ is dominated by the bias, by the usual change of variables and Taylor's expansion,

$$\begin{aligned}
E\left[A_{1,a,T}\left(x\right)\right] &= E\left[\left(P\left(a|x_t\right)-P\left(a|x\right)\right)K_h\left(x_t-x\right)\right] \\
&= \frac{1}{2}\mu_2 h_T^2\left(2\frac{\partial P\left(a|x\right)}{\partial x}\frac{\partial f_X\left(x\right)}{\partial x}+\frac{\partial^2 P\left(a|x\right)}{\partial x^2}f_X\left(x\right)\right)+o\left(h_T^2\right).
\end{aligned}$$

Recall that $E\left[e_{a,t}|x_t\right]=0$ for all $a$ and $t$. We next compute the variance of $A_{2,a,T}\left(x\right)$, this is dominated by the variances as covariance terms are of smaller order, e.g. see [M].

$$\begin{aligned}
\operatorname{var}\left(A_{2,a,T}\left(x\right)\right) &= \operatorname{var}\left(\frac{1}{T}\sum_{t=1}^{T}e_{a,t}K_h\left(x_t-x\right)\right) \\
&= \frac{1}{T}\operatorname{var}\left(e_{a,t}K_h\left(x_t-x\right)\right)+o\left(\frac{1}{Th_T}\right) \\
&= \frac{1}{T}E\left[\sigma_a^2\left(x_t\right)K_h\left(x_t-x\right)\right]+o\left(\frac{1}{Th_T}\right) \\
&= \frac{\kappa_2}{Th_T}\sigma_a^2\left(x\right)f_X\left(x\right)+o\left(\frac{1}{Th_T}\right),
\end{aligned}$$

note that

$$\begin{aligned}
\sigma_a^2\left(x\right) &= E\left[e_{a,t}^2|x_t=x\right] \\
&= \operatorname{var}\left(\mathbf{1}\left[a_t=a\right]|x_t=x\right) \\
&= P\left(a|x\right)\left(1-P\left(a|x\right)\right).
\end{aligned}$$

For the CLT, Lemma 7.1 of [R] can be used repeated throughout this section, using Bernstein blocking technique we obtain,

$$\sqrt{Th_T}\left(\widehat{P}(a|x) - P(a|x) - \frac{1}{2}\mu_2 h_T^2 \eta_{P_a}(x)\right) \Longrightarrow \mathcal{N}(0, \omega_{P_a}(x)),$$

where

$$\begin{aligned}
\eta_{P_a}(x) &= 2\frac{\frac{\partial P(a|x)}{\partial x}\frac{\partial f_X(x)}{\partial x}}{f_X(x)} + \frac{\partial P^2(a|x)}{\partial x^2}, \\
\omega_{P_a}(x) &= \kappa_2 \frac{\sigma_a^2(x)}{f_X(x)}.
\end{aligned} \tag{43}$$

For any $\theta \in \Theta$, recall from (15) and (16)

$$\widehat{r}_\theta(x) = \sum_{a \in A} \zeta_{x,a,\theta}\left(\widehat{P}(a|x)\right),$$

where,

$$\zeta_{x,a,\theta}(t) = t\left(u_\theta(a,x) + \log t\right) + \gamma,$$

by mean value theorem (MVT),

$$\begin{aligned}
& \zeta_{x,a,\theta}\left(\widehat{P}(a|x)\right) - \zeta_{x,a,\theta}\left(P(a|x) + \frac{1}{2}\mu_2 h_T^2 \eta_{P_a}(x)\right) \\
= \; & \zeta'_{x,a,\theta}(P(a|x))\left(\widehat{P}(a|x) - P(a|x) - \frac{1}{2}\mu_2 h_T^2 \eta_{P_a}(x)\right) + o_p(1),
\end{aligned}$$

and

$$\zeta'_{x,a,\theta}(t) = u_\theta(a,x) + \log t + 1. \tag{44}$$

By using MVT again, we can approximate $\zeta_{x,a,\theta}\left(P(a|x) + \frac{1}{2}\mu_2 h_T^2 \eta_{P_a}(x)\right)$ more conveniently as follows,

$$\zeta_{x,a,\theta}\left(P\left(a|x\right)+\frac{1}{2}\mu_2 h_T^2\eta_{P_a}\left(x\right)\right)=\zeta_{x,a,\theta}\left(P\left(a|x\right)\right)+\frac{1}{2}\mu_2 h_T^2\eta_{P_a}\left(x\right)\zeta'_{x,a,\theta}\left(P\left(a|x\right)\right)+o_p\left(h_T^2\right).$$

To obtain the asymptotic distribution for $\widehat{r}_\theta\left(x\right)$, we now provide the joint distribution of $\left\{\widehat{P}\left(a|x\right)\right\}$. It follows immediately, following [R], from Cramér-Wold device that

$$\sqrt{Th_T}\begin{pmatrix}\widehat{P}\left(1|x\right)-P\left(1|x\right)-\frac{1}{2}\mu_2 h_T^2\eta_{P_1}\left(x\right)\\ \vdots \\ \widehat{P}\left(K|x\right)-P\left(K|x\right)-\frac{1}{2}\mu_2 h_T^2\eta_{P_K}\left(x\right)\end{pmatrix}$$

$$\Longrightarrow\mathcal{N}\left(0,\frac{\kappa_2}{f_X\left(x\right)}\begin{pmatrix}\sigma_1^2\left(x\right) & \sigma_{2,1}^2\left(x\right) & \cdots & \sigma_{K,1}^2\left(x\right)\\ \sigma_{1,2}^2\left(x\right) & \ddots & & \vdots\\ \vdots & & \ddots & \sigma_{K,K-1}^2\left(x\right)\\ \sigma_{1,K}^2\left(x\right) & \cdots & \sigma_{K-1,K}^2\left(x\right) & \sigma_K^2\left(x\right)\end{pmatrix}\right),$$

where $\sigma_j^2\left(x\right)=P\left(j|x\right)\left(1-P\left(j|x\right)\right)$ and $\sigma_{j,k}^2\left(x\right)=\sigma_{k,j}^2\left(x\right)=-P\left(j|x\right)P\left(k|x\right)$ for $j,k\in A$. There are a couple of things to notice here, first there exist negative correlation between $\left\{\widehat{P}\left(j|x\right)\right\}$ across $A$, and the covariance matrix in the above display is rank deficient due to the constraint that $\sum_{j\in A}\widehat{P}\left(j|x\right)=1$ for any $x\in int\left(X\right)$. Using thet information from the display above, we have

$$\sqrt{Th_T}\left(\widehat{r}_\theta\left(x\right)-r_\theta\left(x\right)-\frac{1}{2}\mu_2 h_T^2\eta_{r,\theta}\right)\Longrightarrow\mathcal{N}\left(0,\omega_{\zeta_x,\theta}\left(x\right)\right),$$

where

$$\eta_{r,\theta}(x) = \sum_{j\in A}\eta_{P_j}(x)\,\zeta'_{x,j,\theta}(P(j|x)), \tag{45}$$

$$\omega_{r,\theta}(x) = \frac{\kappa_2}{f_X(x)}\left(\begin{array}{c}\sum_{j\in A}\left(\zeta'_{x,j,\theta}(P(j|x))\right)^2\sigma_j^2(x)\\ -2\sum_{j\neq k}\zeta'_{x,j,\theta}(P(j|x))\,\zeta'_{x,k,\theta}(P(k|x))\,P(j|x)\,P(k|x)\end{array}\right) \tag{46}$$

where $\left\{\eta_{P_j}\right\}_{j\in A}$ and $\left\{\zeta'_{x,j,\theta}\right\}_{j\in A}$ are defined in (45) and (44) respectively. Note we can relate components of the expansion of $\widehat{r}_\theta(x)$, in (31), to the terms above as follows,

$$r_\theta(x) = \sum_{j\in A}\zeta_{x,j,\theta}(P(j|x)), \tag{47}$$

$$\widehat{r}_\theta^B(x) = \frac{1}{2}\mu_2 h_T^2\eta_{r,\theta}(x), \tag{48}$$

$$\widehat{r}_\theta^C(x) = \sum_{j\in A}\frac{\zeta'_{x,j,\theta}(P(j|x))}{f_X(x)}\times\left(\frac{1}{T}\sum_{t=1}^{T}e_{j,t}K_h(x_t-x)\right). \tag{49}$$

We next provide the statistical properties for $\widehat{m}_\theta^A(x)$. First, $\left(\widehat{\mathcal{L}}-\mathcal{L}\right)m_\theta(x)$:

$$\begin{aligned}\left(\widehat{\mathcal{L}}-\mathcal{L}\right)m_\theta(x) &= \beta\int m_\theta(x')\left(\widehat{f}_{X'|X}(dx'|x)-f_{X'|X}(dx'|x)\right)\\ &= \frac{\beta}{f_X(x)}\int m_\theta(x')\left(\widehat{f}_{X,X}(dx',x)-f_{X,X}(dx',x)\right)\\ &\quad -\frac{\beta}{f_X(x)}\left(\widehat{f}_X(x)-f_X(x)\right)\int m_\theta(x')\,f_{X'|X}(dx'|x)+o_p\left(T^{-2/5}\right)\\ &= B_{1,\theta,T}(x)+B_{2,\theta,T}(x)+o_p\left(T^{-2/5}\right).\end{aligned}$$

To analyze $B_{1,\theta,T}(x)$, proceed with the usual decomposition of $\widehat{f}_{X',X}(x',x)-f_{X',X}(x',x)$ then integrate it over, note that the integral reduces the variance to that of a 1 dimensional nonparametric estimator, we have

$$B_{1,\theta,T}(x) = B_{1,\theta,T}^B(x)+B_{1,\theta,T}^C(x)+o_p\left(T^{-2/5}\right),$$

where

$$B_{1,\theta,T}^{B}(x) = \frac{1}{2}\mu_2 h_T^2 \beta \int \left( \frac{m_\theta(x')}{f_X(x)} \left( \frac{\partial^2 f_{X',X}(x',x)}{\partial x'^2} + \frac{\partial^2 f_{X',X}(x',x)}{\partial x^2} \right) \right) dx', \tag{50}$$

$$B_{1,\theta,T}^{C}(x) = \frac{\beta}{f_X(x)} \left( \frac{1}{T-1} \sum_{t=1}^{T-1} \int m_\theta(x') \left( \begin{array}{c} K_h(x_{t+1}-x') K_h(x_t - x) \\ -E\left[K_h(x_{t+1}-x') K_h(x_t-x)\right] \end{array} \right) \right) dx' \tag{51}$$

and it can be shown that

$$\sqrt{Th_T} B_{1,\theta,T}^{C}(x) \Longrightarrow \mathcal{N}\left(0, \frac{\beta^2}{f_X(x)} \kappa_2 \int (m_\theta(x'))^2 f_{X'|X}(dx'|x)\right).$$

For $B_{2,\theta,T}(x)$, this is just the kernel density estimator of $f_X(x)$ multiplied by a non-stochastic term,

$$B_{2,\theta,T}(x) = B_{2,\theta,T}^{B}(x) + B_{2,\theta,T}^{C}(x) + o_p\left(T^{-2/5}\right),$$

where

$$B_{2,\theta,T}^{B}(x) = -\frac{1}{2}\mu_2 h_T^2 \frac{\partial^2 f_X(x)}{\partial x^2} \left( \frac{\beta}{f_X(x)} \int m_\theta(x') f_{X'|X}(dx'|x) \right), \tag{52}$$

$$B_{2,\theta,T}^{C}(x) = -\left( \frac{\beta}{f_X(x)} \int m_\theta(x') f_{X'|X}(dx'|x) \right) \frac{1}{T} \sum_{t=1}^{T} \left( K_h(x_t - x) - E\left[K_h(x_t - x)\right] \right) \tag{53}$$

and it can be shown that

$$\sqrt{Th_T} B_{2,\theta,T}^{C}(x) \Longrightarrow \mathcal{N}\left(0, \kappa_2 f_X(x) \left( \frac{\beta}{f_X(x)} \int m_\theta(x') f_{X'|X}(dx'|x) \right)^2 \right).$$

Combining these we have,

$$\widehat{m}_\theta(x) = m_\theta(x) + \overline{m}_\theta^B(x) + \overline{m}_\theta^C(x) + o_p\left(T^{-2/5}\right),$$

where

$$\overline{m}_{\theta}^{B}(x) = (I-\mathcal{L})^{-1}\left(B_{1,\theta,T}^{B}+B_{2,\theta,T}^{B}+\widehat{r}_{\theta}^{B}\right)(x),$$

$$\overline{m}_{\theta}^{C}(x) = B_{1,\theta,T}^{C}(x)+B_{2,\theta,T}^{C}(x)+\widehat{r}_{\theta}^{C}(x).$$

Note also that

$$\sqrt{Th_T}\left(B_{1,\theta,T}^{C}(x)+B_{2,\theta,T}^{C}(x)\right) \Longrightarrow \mathcal{N}\left(0,\frac{\kappa_2\beta^2}{f_X(x)}\mathrm{var}\left(m_\theta(x_{t+1})\,|x_t=x\right)\right)$$

and

$$Cov\left(\sqrt{Th_T}\left(B_{1,\theta,T}^{C}(x)+B_{2,\theta,T}^{C}(x)\right),\sqrt{Th_T}\widehat{r}_{\theta}^{C}(x)\right) \rightarrow 0 \text{ as } n\rightarrow\infty.$$

This provides us with the pointwise theory for $\widehat{m}_\theta$ for any $x \in int(X)$ and $\theta\in\Theta$.

$$\sqrt{Th_T}\left(\widehat{m}_\theta(x)-m_\theta(x)-\frac{1}{2}\mu_2h_T^2\eta_{m,\theta}(x)\right)\Longrightarrow\mathcal{N}\left(0,\omega_{m,\theta}(x)\right),$$

where

$$\eta_{m,\theta}(x) = (I-\mathcal{L})^{-1}\left(\eta_{r,\theta}+\eta_{\mathcal{L},\theta}\right)(x),$$

$$\omega_{m,\theta}(x) = \frac{\kappa_2}{f_X(x)}\left(\beta^2\mathrm{var}\left(m_\theta(x_{t+1})\,|x_t=x\right)+\omega_{r,\theta}(x)\right),$$

where $\eta_{r,\theta}$ and $\omega_{r,\theta}$ are defined in (45) and (46), and

$$\eta_{\mathcal{L},\theta}(x)=\beta\begin{pmatrix}\frac{1}{f_X(x)}\int m_\theta(x')\left(\frac{\partial^2 f_{X',X}(x',x)}{\partial x'^2}+\frac{\partial^2 f_{X',X}(x',x)}{\partial x^2}\right)dx' \\ -\frac{\frac{\partial^2 f_X(x)}{\partial x^2}}{f_X(x)}\int m_\theta(x')\,f_{X'|X}(dx'|x)\end{pmatrix}, \tag{54}$$

$\eta_{r,\theta}$, $\omega_{r,\theta}$ are defined in (45) - (46). The proof of pairwise asymptotic independence across distinct $x$ is obvious. ■

PROOF OF THEOREM 1.2. From the decomposition from Theorem 1 we obtain the pointwise results for $\widehat{g}_\theta(a,x)$. Similarly to the decomposition of $\left(\widehat{\mathcal{L}}-\mathcal{L}\right)m_\theta(x)$, we have

$$\begin{aligned}\left(\widehat{\mathcal{H}}-\mathcal{H}\right)m_\theta(a,x) &= \int m_\theta(x')\left(\widehat{f}_{X'|X,A}(dx'|x,a)-f_{X'|X,A}(dx'|x,a)\right)\\ &= C_{1,\theta,T}(a,x)+C_{2,\theta,T}(a,x)+o_p\left(T^{-2/5}\right).\end{aligned}$$

The properties for $C_{1,\theta,T}$ and $C_{2,\theta,T}$ are closely related to that of $B_{1,\theta,T}$ and $B_{2,\theta,T}$.

$$C_{1,\theta,T}(a,x)=C^B_{1,\theta,T}(a,x)+C^C_{1,\theta,T}(a,x)+o_p\left(T^{-2/5}\right),$$

where

$$\begin{aligned}C^B_{1,\theta,T}(a,x) &= \frac{1}{2}\mu_2h_T^2\int\frac{m_\theta(x')}{f_{X,A}(x,a)}\left(\frac{\partial^2 f_{X',X,A}(x',x,a)}{\partial x'^2}+\frac{\partial^2 f_{X',X,A}(x',x,a)}{\partial x^2}\right)dx',\\ C^C_{1,\theta,T}(a,x) &= \frac{1}{f_{X,A}(x,a)}\int\left(\frac{1}{T-1}\sum_{t=1}^{T-1}m_\theta(x')\left(\begin{array}{c}K_h(x_{t+1}-x')K_h(x_t-x)\mathbf{1}[a_t=a]\\ -E\left[K_h(x_{t+1}-x')K_h(x_t-x)\mathbf{1}[a_t=a]\right]\end{array}\right)\right)\end{aligned}$$

and as in the case of $B^C_{1,\theta,T}$

$$\sqrt{Th_T}C^C_{1,\theta,T}(a,x)\Longrightarrow\mathcal{N}\left(0,\frac{\kappa_2}{f_{X,A}(x,a)}\int\left(m_\theta(x')\right)^2f_{X'|X,A}(dx'|x,a)\right).$$

Similarly for $C_{2,\theta,T}$,

$$\begin{aligned}C^B_{2,\theta,T}(a,x) &= -\frac{1}{2}\mu_2h_T^2\frac{\partial^2 f_{X,A}(x,a)}{\partial x^2}\left(\frac{1}{f_{X,A}(x,a)}\int m_\theta(x')f_{X'|X,A}(dx'|x,a)\right),\\ C^C_{2,\theta,T}(a,x) &= -\left(\frac{1}{f_{X,A}(x,a)}\int m_\theta(x')f_{X'|X,A}(dx'|x,a)\right)\frac{1}{T}\sum_{t=1}^{T}\left(\begin{array}{c}K_h(x_t-x)\mathbf{1}[a_t=a]\\ -E\left[K_h(x_t-x)\mathbf{1}[a_t=a]\right]\end{array}\right)\end{aligned}$$

and

$$\sqrt{Th_T}C^C_{2,\theta,T}(a,x) \Longrightarrow \mathcal{N}\left(0, \frac{\kappa_2}{f_{X,A}(x,a)}\int (m_\theta(x'))^2 f_{X'|X,A}(dx'|x,a)\right).$$

Combining these we have,

$$\widehat{g}_\theta(a,x) = g_\theta(a,x) + \overline{g}^B_\theta(a,x) + \overline{g}^C_\theta(a,x) + o_p\left(T^{-2/5}\right),$$

where

$$\begin{aligned}
\overline{g}^B_\theta(a,x) &= C^B_{1,\theta,T}(a,x) + C^B_{2,\theta,T}(a,x) + \mathcal{H}\widehat{r}^B_\theta(a,x), \\
\overline{g}^C_\theta(a,x) &= C^C_{1,\theta,T}(a,x) + C^C_{2,\theta,T}(a,x).
\end{aligned}$$

This provides us with the pointwise distribution theory for $\widehat{g}$ for any $x \in int(X), a \in A$ and $\theta \in \Theta$.

$$\sqrt{Th_T}\left(\widehat{g}(a,x) - g(a,x) - \frac{1}{2}\mu_2 h_T^2 \eta_{g,\theta}(a,x)\right) \Longrightarrow \mathcal{N}(0, \omega_{g,\theta}(a,x)),$$

where,

$$\begin{aligned}
\eta_{g,\theta}(a,x) &= \mathcal{H}(I-\mathcal{L})^{-1}\left(\eta_{r,\theta} + \eta_{\mathcal{L},\theta}\right)(a,x) + \eta_{\mathcal{H},\theta}(a,x), \\
\omega_{g,\theta}(a,x) &= \frac{\kappa_2}{f_{X,A}(a,x)}\text{var}(m_\theta(x_{t+1})|x_t = x, a_t = a),
\end{aligned}$$

$\eta_{r,\theta}$ and $\eta_{\mathcal{L},\theta}$ are as defined in the proof of Theorem 1, and

$$\begin{aligned}
\eta_{\mathcal{H},\theta}(a,x) &= \frac{1}{f_{X,A}(x,a)}\int m_\theta(x')\left(\frac{\partial^2 f_{X',X,A}(x',x,a)}{\partial^2 x'} + \frac{\partial^2 f_{X',X,A}(x',x,a)}{\partial^2 x}\right) \quad (55) \\
&\quad - \frac{\frac{\partial^2 f_{X,A}(x,a)}{\partial^2 x}}{f_{X,A}(x,a)}\int m_\theta(x') f_{X'|X,A}(dx'|x,a).
\end{aligned}$$

Pairwise asymptotic independence, across distinct $x$, completes the proof. ■

### 1.8.3 Proofs of High Level Conditions A1.1 - A1.6

PROOF OF A1.1. It suffices to show that

$$\begin{aligned}\sup_{(x',x)\in\mathcal{X}\times\mathcal{X}}\left|\widehat{f}_{X',X}(x',x)-f_{X',X}(x',x)\right| &= o_p(\delta_T),\\ \sup_{x\in\mathcal{X}}\left|\widehat{f}_{X}(x)-f_{X}(x)\right| &= o_p(\delta_T).\end{aligned}$$

These uniform rates are bounded by the rates for the bias squared and the rates of the centred process. The former is standard, and holds uniformly over $\mathcal{X}\times\mathcal{X}$ (and $\mathcal{X}$). See the Additional Proofs section below, where proof of A1.1 falls under Case 1. ■

PROOF OF A1.2. The components for the decomposition have been provided by (47) - (49). By uniform boundedness of $\eta_{P_a}$ and $\zeta_{x,a,\theta}$ over $A\times X\times\Theta$ and triangle inequalities, the order of the leading bias and remainder terms are as stated in (32) and (35) respectively. For the stochastic term, we can utilize the exponential inequality, see Case 2 of the Additional Proofs section. We next check (34). [LM] use eigen-expansion to construct the kernel of the new integral operator and showed that it had nice properties in their problem. We use the Neumann's series to construct our kernel, for any $\phi\in C(\mathcal{X})$

$$\mathcal{L}(I-\mathcal{L})^{-1}\phi=\sum_{j=1}^{\infty}\mathcal{L}^{j}\phi, \tag{56}$$

where $\mathcal{L}^j$ represents a linear operator of a $j$-step ahead predictor with discounting, this follows from Chapman-Kolmogorov equation for homogeneous Markov chains, for $\tau>1$

$$\begin{aligned}\mathcal{L}^{\tau}\phi(x) &= \beta^{\tau}\int\phi(x')\,f_{(\tau)}(dx'|x) \qquad (57)\\ f_{(\tau)}(x_{t+\tau}|x_t) &= \int f_{X'|X}(x_{t+\tau}|x_{t+\tau-1})\prod_{k=1}^{\tau-1}f_{X'|X}(dx_{t+\tau-k}|x_{t+\tau-k-1}),\end{aligned}$$

where $f_{(\tau)}(dx_{t+\tau}|x_t)$ denotes the conditional density of $\tau$-steps ahead. First, we note that $\mathcal{L}(I-\mathcal{L})^{-1}\phi \in C(\mathcal{X})$, this is always true since for any $\phi \in C(\mathcal{X})$ and $x \in \mathcal{X}$ since:

$$\begin{aligned}
\left|\mathcal{L}(I-\mathcal{L})^{-1}\phi(x)\right| &= \left|\sum_{\tau=1}^{\infty}\beta^{\tau}\int \phi(x') f_{(\tau)}(dx'|x)\right| \\
&\leq \sum_{\tau=1}^{\infty}\beta^{\tau}\int f_{(\tau)}(dx'|x)\,\|\phi\| \\
&\leq \frac{\beta}{1-\beta}\|\phi\| \\
&< \infty.
\end{aligned}$$

We denote the kernel of the integral transform (56) by the limit, $\varphi$, of the partial sum, $\varphi_J$,

$$\varphi_T(x', x) = \sum_{\tau=1}^{T}\beta^{\tau} f_{(\tau)}(x'|x), \tag{58}$$

where $\varphi$ is continuous on $\mathcal{X}\times\mathcal{X}$. This is easy to see since $f_{(\tau)}$ is continuous and is uniformly bounded for all $j$ by $\sup_{(x',x)\in\mathcal{X}\times\mathcal{X}}|f(x'|x)|$, by completeness, $\varphi_J$ converges to a continuous function (with Lipschitz constant no larger than $\frac{\beta}{1-\beta}\sup_{(x',x)\in\mathcal{X}\times\mathcal{X}}|f(x'|x)|$). To proof (34), for details see Case 3 of the Additional Proofs section, we apply exponential inequality to bound

$$\Pr\left(\frac{1}{T}\left|\sum_{t=1}^{T} e_{\theta,t}\nu(x_t, x)\right| > \delta_T\right), \tag{59}$$

for some positive sequence, $\delta_T = o\left(T^{-2/5}\right)$, where $\nu(x_t, x)$ is defined as

$$\begin{aligned}
\nu(x_t, x) &= \int \frac{K_h(x_t - x')}{f_X(x')}\varphi(dx', x) \\
&= \frac{\varphi(x_t, x)}{f_X(x_t)} + O\left(h_T^2\right),
\end{aligned} \tag{60}$$

and the latter equality holds uniformly on $\mathcal{X}$. ■

PROOF OF A1.3. Following the decomposition of $\widehat{f}(x'|x)$ we obtain the leading bias and variance terms are sum of (50) and (52), and, (51) and (53) respectively. The results rates of convergence follow similarly to the proof of A1.2. ■

PROOF OF A1.4. This is essentially the same as proof of A1.1. ■

PROOF OF A1.5. Notice that $\overline{m}_\theta^C$ consists of $\widehat{r}_\theta^C$ and $\widehat{r}_\theta^F$. We need to show,

$$
\begin{aligned}
\sup_{(a,x)\in A\times X} \left|\mathcal{H}\widehat{r}_\theta^C(a,x)\right| &= o_p\left(T^{-2/5}\right) \\
\sup_{(a,x)\in A\times X} \left|\mathcal{H}\widehat{r}_\theta^F(a,x)\right| &= o_p\left(T^{-2/5}\right).
\end{aligned}
$$

The proof follows from exponential inequalities, see the Additional Proofs section. ■

PROOF OF A6. This is essentially the same as proof of A1.3. ■

### 1.8.4 Proofs of Theorems 1.3 - 1.5

We begin with two lemmas for the uniform expansion of some partial derivatives of $\widehat{m}_\theta$ and $\widehat{g}_\theta$.

LEMMA 1.1: Under conditions $B1.1'$, B1.2 - B1.6 hold. Then the following expansion holds for $k = 0, 1, 2$ and $j = 1, \ldots, L$,

$$
\max_{1\leq j\leq L} \sup_{(x,\theta)\in\mathcal{X}\times\Theta} \left| \frac{\partial^k \widehat{m}_\theta(x)}{\partial\theta_j^k} - \frac{\partial^k m_\theta(x)}{\partial\theta_j^k} - \frac{\partial^k \overline{m}_\theta^B(x)}{\partial\theta_j^k} - \frac{\partial^k \overline{m}_\theta^C(x)}{\partial\theta_j^k} \right| = o_p\left(T^{-2/5}\right),
$$

where $\frac{\partial^k m_\theta}{\partial\theta^k}$ is defined as the solution to

$$
\frac{\partial^k m_\theta}{\partial\theta_j^k} = \frac{\partial^k r_\theta}{\partial\theta_j^k} + \mathcal{L}\frac{\partial^k m_\theta}{\partial\theta_j^k}, \tag{61}
$$

and $\frac{\partial^k \widehat{m}_\theta}{\partial\theta_j^k}$ defined as the solution to the analogous empirical integral equation. Standard definition for partial derivative applies for $\frac{\partial^k \overline{m}_\theta^b(x)}{\partial\theta_j^k}$ with $b = B, C$. Notice, when $k = 0$, this coincides with the terms previously defined in Proposition 1.1. Further,

$$\begin{aligned}
\max_{1\le j\le L}\sup_{(x,\theta)\in\mathcal{X}\times\Theta}\left|\frac{\partial^k\overline{m}_\theta^B(x)}{\partial\theta_j^k}\right| &= O_p\left(T^{-2/5}\right) \text{ with } \frac{\partial^k\overline{m}_\theta^B(x)}{\partial\theta^k} \text{ deterministic,}\\
\max_{1\le j\le L}\sup_{(x,\theta)\in\mathcal{X}\times\Theta}\left|\frac{\partial^k\overline{m}_\theta^C(x)}{\partial\theta_j^k}\right| &= o_p\left(T^{\xi-2/5}\right) \text{ for any } \xi>0.
\end{aligned}$$

Proof of Lemma 1.1. Comparing integral equations in (9) and (61), we notice that, these are just the integral equations with the same kernel but different intercepts. Since $\zeta_{x,j,\theta}$, $\zeta_{x,j,\theta}$ and $m_\theta$ are twice continuously differentiable in $\theta$ on $\Theta$ over $A\times X$, Dominated Convergence Theorem (DCT) can be utilized throughout, all arguments used to verify the definition of $\frac{\partial^k\overline{m}_\theta(x)}{\partial\theta_j^k}$ and their uniformity results analogous to A1.2 -A1.3 follow immediately. ∎

Lemma 1.2: Under conditions $B1.1'$, B1.2 - B1.6 hold. Then the following expansion holds for $k=0,1,2$ and $j=1,\ldots,L$,

$$\max_{1\le j\le L}\sup_{(x,a,\theta)\in A\times X\times\Theta}\left|\frac{\partial^k\widehat{g}_\theta(a,x)}{\partial\theta_j^k}-\frac{\partial^k g_\theta(a,x)}{\partial\theta_j^k}-\frac{\partial^k\overline{g}_\theta^B(a,x)}{\partial\theta_j^k}-\frac{\partial^k\overline{g}_\theta^C(a,x)}{\partial\theta_j^k}\right|=o_p\left(T^{-2/5}\right),$$

where all of the terms above are defined analogously to those found in Lemma 1.1 and, for $k=1,2$

$$\begin{aligned}
\max_{1\le j\le L}\sup_{(x,a,\theta)\in A\times X\times\Theta}\left|\frac{\partial^k\overline{g}_\theta^B(a,x)}{\partial\theta_j^k}\right| &= O_p\left(T^{-2/5}\right) \text{ with } \frac{\partial^k\overline{g}_\theta^B(a,x)}{\partial\theta_j^k} \text{ deterministic,}\\
\max_{1\le j\le L}\sup_{(x,a,\theta)\in A\times X\times\Theta}\left|\frac{\partial^k\overline{g}_\theta^C(a,x)}{\partial\theta_j^k}\right| &= o_p\left(T^{\xi-2/5}\right) \text{ for any } \xi>0.
\end{aligned}$$

Proof of Lemma 1.2: Same as the proof of Lemma 1.1. ∎

Proof of Theorem 1.3: We first proceed to show the consistency result of the estimator.

Consistency.

Consider any estimator $\theta_T$ of $\theta_0$ that asymptotically maximizes $\widehat{Q}_T(\theta)$:

$$Q_T(\theta_T) \geq \sup_{\theta\in\Theta} Q_T(\theta) - o_p(1).$$

Under B1.1 and B1.9, by standard arguments for example see McFadden and Newey (1994), consistency of such extremum estimators can be obtained if we have

$$\sup_{\theta\in\Theta} \left|\widehat{Q}_T(\theta) - Q(\theta)\right| = o_p(1). \tag{62}$$

By triangle inequality, (62) is implied by

$$\begin{aligned} \sup_{\theta\in\Theta} |Q_T(\theta) - Q(\theta)| &= o_p(1) && (63)\\ \sup_{\theta\in\Theta} \left|\widehat{Q}_T(\theta) - Q_T(\theta)\right| &= o_p(1). && (64) \end{aligned}$$

For (63), since $q : A \times X \times \Theta \to \mathbb{R}$ is continuous on the compact set $X \times \Theta$, for any $a \in A$, hence by Weierstrass Theorem

$$\max_{a\in A} \sup_{x\in X,\theta\in\Theta} |q(a, x; \theta, g_\theta)| < \infty. \tag{65}$$

This ensures that $E|q(a_t, x_t; \theta, v_\theta)| < \infty$, and by the LLN for ergodic and stationary processes we have

$$Q_T(\theta) \xrightarrow{p} Q(\theta) \quad \text{for each } \theta \in \Theta.$$

The convergence above can be made uniform since $Q_T$ is stochastic equicontinuous and $Q$ is uniformly continuous by DCT, with a majorant in (65). To proof (64) we partition $\widehat{Q}_T(\theta) - Q_T(\theta)$ into two components

$$\widehat{Q}_T(\theta) - Q_T(\theta) = \frac{1}{T}\sum_{t=1}^{T} \mathfrak{c}_{t,T}\left(q(a_t, x_t; \theta, \widehat{g}_\theta) - q(a_t, x_t; \theta, g_\theta)\right) + \frac{1}{T}\sum_{t=1}^{T} (1 - \mathfrak{c}_{t,T})\, q(a_t, x_t; \theta, \widehat{g}_\theta),$$

where the second term is $o_p(1)$. This follows since, we denote $1-\mathfrak{c}_{t,T}$ by $\mathfrak{d}_{t,T}$, ,

$$
\begin{aligned}
\left|\frac{1}{T}\sum_{t=1}^{T}\mathfrak{d}_{t,T}q\left(a_t,x_t;\theta,\widehat{g}_\theta\right)\right| &\leq \max_{a\in A}\sup_{x\in X,\theta\in\Theta}\left|q\left(a,x;\theta,g_\theta\right)\right|\frac{1}{T}\sum_{t=1}^{T}\mathfrak{d}_{t,T} \\
&= o_p(1).
\end{aligned}
$$

The first inequality holds w.p.a. 1 and the equality is the result of $\mathfrak{d}_{t,T}=o_p(\vartheta_T)$ for any rate $\vartheta_T\to\infty$. To proof (64), now it suffices to show,

$$
\max_{a\in A}\sup_{x\in X,\theta\in\Theta}\left|q\left(a,x;\theta,\widehat{g}_\theta\right)-q\left(a,x;\theta,g_\theta\right)\right|=o_p(1).
$$

Recall that

$$
\begin{aligned}
q\left(a,x;\theta,\widehat{g}_\theta\right)-q\left(a,x;\theta,g_0\right) &= \widehat{v}_\theta(a,x)-v_\theta(a,x)+\log\left(\frac{\sum_{\widetilde{a}\in A}\exp\left(v_\theta\left(\widetilde{a},x\right)\right)}{\sum_{\widetilde{a}\in A}\exp\left(\widehat{v}_\theta\left(\widetilde{a},x\right)\right)}\right), \\
v_\theta(a,x) &= u_\theta(a,x)+g_\theta(a,x), \\
\widehat{v}_\theta(a,x) &= u_\theta(a,x)+\widehat{g}_\theta(a,x).
\end{aligned}
$$

All the listed functions are in $C(\mathcal{X})$. We have shown earlier that for some $\delta_T=o(1)$

$$
\max_{a\in A}\sup_{x\in X,\theta\in\Theta}\left|\widehat{g}_\theta(a,x)-g_\theta(a,x)\right|=o_p(\delta_T),
$$

so we have uniform convergence for $\widehat{v}$ to $v$ at the same rate. We know for any continuously differentiable function $\phi$ (in this case, $\exp(\cdot)$ and $\log(\cdot)$), by MVT,

$$
\max_{a\in A}\sup_{x\in X,\theta\in\Theta}\left|\phi\left(\widehat{v}_\theta(a,x)\right)-\phi\left(v_\theta(a,x)\right)\right|=o_p(\delta_T).
$$

So we have

$$\sup_{x\in X,\theta\in\Theta}\left|\sum_{\widetilde{a}\in A}\exp\left(\widehat{v}_{\theta}\left(\widetilde{a},x\right)\right)-\sum_{\widetilde{a}\in A}\exp\left(v_{\theta}\left(\widetilde{a},x\right)\right)\right|=o_{p}\left(1\right),$$

and since we have, at least w.p.a. 1, $\exp\left(\widehat{v}_{\theta}\left(\widetilde{a},x\right)\right)$ and $\exp\left(v_{\theta}\left(\widetilde{a},x\right)\right)$ are positive a.s.

$$\left|\frac{\sum_{\widetilde{a}\in A}\exp\left(v_{\theta}\left(\widetilde{a},x\right)\right)}{\sum_{\widetilde{a}\in A}\exp\left(\widehat{v}_{\theta}\left(\widetilde{a},x\right)\right)}-1\right|=\left|\frac{1}{\sum_{\widetilde{a}\in A}\exp\left(\widehat{v}_{\theta}\left(\widetilde{a},x\right)\right)}\right|\left|\sum_{\widetilde{a}\in A}\exp\left(v_{\theta}\left(\widetilde{a},x\right)\right)-\sum_{\widetilde{a}\in A}\exp\left(\widehat{v}_{\theta}\left(\widetilde{a},x\right)\right)\right|,$$

and by Wierstrass Theorem, w.p.a. 1,

$$\min_{a\in A}\inf_{x\in X,\theta\in\Theta}\exp\left(\widehat{v}_{\theta}\left(a,x\right)\right)>0,$$

hence we have

$$\sup_{x\in X,\theta\in\Theta}\left|\frac{\sum_{\widetilde{a}\in A}\exp\left(v_{\theta}\left(\widetilde{a},x\right)\right)}{\sum_{\widetilde{a}\in A}\exp\left(\widehat{v}_{\theta}\left(\widetilde{a},x\right)\right)}-1\right|=o_{p}\left(1\right).$$

The proof of (64) is completed once we apply another mean value expansion, as done previously, to obtain

$$\sup_{x\in X,\theta\in\Theta}\left|\log\left(\frac{\sum_{\widetilde{a}\in A}\exp\left(v_{\theta}\left(\widetilde{a},x\right)\right)}{\sum_{\widetilde{a}\in A}\exp\left(\widehat{v}_{\theta}\left(\widetilde{a},x\right)\right)}\right)\right|=o_{p}\left(1\right).$$

Asymptotic Normality

Consider the first order condition

$$\frac{\partial\widehat{Q}_{T}\left(\widehat{\theta}\right)}{\partial\theta}=o_{p}\left(1\right),$$

from MVT we have

$$o_{p}\left(\sqrt{T}\right)=\sqrt{T}\frac{\partial\widehat{Q}_{T}\left(\theta_{0}\right)}{\partial\theta}+\frac{\partial^{2}\widehat{Q}_{T}\left(\overline{\theta}\right)}{\partial\theta^{2}}\sqrt{T}\left(\widehat{\theta}-\theta_{0}\right).$$

We show that for any sequence $\epsilon_T \to 0$ there exists some positive $C$ such that

$$\inf_{\|\theta-\theta_0\|<\epsilon_T} \lambda_{\min}\left(-\frac{\partial^2 \widehat{Q}_T(\theta)}{\partial\theta\partial\theta'}\right) > C + o_p(1) \tag{66}$$

$$\sqrt{T}\frac{\partial \widehat{Q}_T(\theta_0)}{\partial\theta} = O_p(1) \tag{67}$$

This implies

$$\sqrt{T}\left(\widehat{\theta}-\theta_0\right) = O_p(1).$$

To proof (66), we first show

$$\sup_{\|\theta-\theta_0\|<\epsilon_T}\left\|\frac{\partial^2 \widehat{Q}_T(\theta)}{\partial\theta\partial\theta'} - E\left[\frac{\partial^2 q(a_t,x_t;\theta,g_\theta)}{\partial\theta\partial\theta'}\right]\right\| = o_p(1). \tag{68}$$

Since the second derivative of $q : A \times X \times \Theta \to \mathbb{R}$ is continuous on the compact set $X \times \Theta$ and for each $a \in A$, standard arguments for uniform convergence implies that

$$\sup_{\|\theta-\theta_0\|<\epsilon_T}\left\|\frac{\partial^2 Q_T(\theta)}{\partial\theta\partial\theta'} - E\left[\frac{\partial^2 q(a_t,x_t;\theta,g_\theta)}{\partial\theta\partial\theta'}\right]\right\| = o_p(1).$$

By triangle inequality, (68) will hold if we can show,

$$\sup_{\|\theta-\theta_0\|<\epsilon_T}\left\|\frac{\partial^2 \widehat{Q}_T(\theta)}{\partial\theta\partial\theta'} - \frac{\partial^2 Q_T(\theta)}{\partial\theta\partial\theta'}\right\| = o_p(1).$$

This is similar to showing (64), as the above condition is implied by,

$$\max_{a\in A}\sup_{x\in X,\|\theta-\theta_0\|<\epsilon_T}\left\|\frac{\partial^2 q(a_t,x_t;\theta,\widehat{g}_\theta)}{\partial\theta\partial\theta'} - \frac{\partial^2 q(a_t,x_t;\theta,g_\theta)}{\partial\theta\partial\theta'}\right\| = o_p(1). \tag{69}$$

The expressions for the score of $q$ is,

$$\frac{\partial q(a_t,x_t;\theta,g_\theta)}{\partial\theta} = \frac{\partial v_\theta(a_t,x_t)}{\partial\theta} - \frac{\sum_{\widetilde{a}\in A}\left(\frac{\partial v_\theta(\widetilde{a},x_t)}{\partial\theta}\right)\exp\left(v_\theta(\widetilde{a},x_t)\right)}{\sum_{\widetilde{a}\in A}\exp\left(v_\theta(\widetilde{a},x_t)\right)}, \tag{70}$$

and for the Hessian

$$\begin{aligned}\frac{\partial^2 q(a_t,x_t;\theta,g_\theta)}{\partial\theta\partial\theta'} &= \frac{\partial^2 v_\theta(a_t,x_t)}{\partial\theta\partial\theta'} - \frac{\sum_{\tilde{a}\in A}\left(\frac{\partial^2 v_\theta(\tilde{a},x_t)}{\partial\theta\partial\theta'} + \frac{\partial v_\theta(\tilde{a},x_t)}{\partial\theta}\frac{\partial v_\theta(\tilde{a},x_t)}{\partial\theta'}\right)\exp(v_\theta(\tilde{a},x_t))}{\sum_{\tilde{a}\in A}\exp(v_\theta(\tilde{a},x_t))} \\ &+ \frac{\sum_{\tilde{a}\in A}\frac{\partial v_\theta(\tilde{a},x_t)}{\partial\theta}\frac{\partial v_\theta(\tilde{a},x_t)}{\partial\theta'}\exp(v_\theta(\tilde{a},x_t))^2}{\left(\sum_{\tilde{a}\in A}\exp(v_\theta(\tilde{a},x_t))\right)^2}.\end{aligned}$$

Proceed along the same line of arguments for proving (64), we show (69) holds by tedious but straightforward calculations. Essentially we need uniform convergence of the following partial derivatives,

$$\max_{a\in A,1\leq j\leq L}\sup_{x\in X,\theta\in\Theta}\left|\frac{\partial^k\hat{v}_\theta(a,x)}{\partial\theta_j^k} - \frac{\partial^k v_\theta(a,x)}{\partial\theta_j^k}\right| = o_p(1) \quad \text{for } k=0,1,2, \tag{71}$$

(69) follows from repeated mean value expansions as done in the proof of (64). The uniform convergence in (71) follows from Lemmas 1.1 and 1.2, this implies (66).

For (67),

$$\begin{aligned}\sqrt{T}\frac{\partial\hat{Q}_T(\theta_0)}{\partial\theta} &= \frac{1}{\sqrt{T}}\sum_{t=1}^T \mathfrak{c}_{t,T}\frac{\partial q(a_t,x_t;\theta_0,\hat{g}_{\theta_0})}{\partial\theta} \\ &= \frac{1}{\sqrt{T}}\sum_{t=1}^T\frac{\partial q(a_t,x_t;\theta_0,g_{\theta_0})}{\partial\theta} \\ &+\frac{1}{\sqrt{T}}\sum_{t=1}^T\mathfrak{c}_{t,T}\left(\frac{\partial q(a_t,x_t;\theta_0,\hat{g}_{\theta_0})}{\partial\theta} - \frac{\partial q(a_t,x_t;\theta_0,g_{\theta_0})}{\partial\theta}\right) \\ &\frac{1}{\sqrt{T}}\sum_{t=1}^T\mathfrak{d}_{t,T}\frac{\partial q(a_t,x_t;\theta_0,g_{\theta_0})}{\partial\theta} \\ &= D_{1,T}+D_{2,T}+D_{3,T},\end{aligned}$$

The term $D_{1,T}$ is asymptotically normal with mean zero and finite variance by the CLT for stationary and geometric mixing process,

$$\sqrt{T}D_{1,T} \Longrightarrow \mathcal{N}(0,\Lambda_1),$$

where

$$
\begin{aligned}
\Lambda_1 &= E\left[\frac{\partial q\left(a_t, x_t; \theta_0, g_{\theta_0}\right)}{\partial \theta}\frac{\partial q\left(a_t, x_t; \theta_0, g_{\theta_0}\right)'}{\partial \theta}\right] \\
&+ \lim_{T\to\infty}\frac{1}{T}\sum_{t=1}^{T}(T-t)\left(\begin{array}{c} E\left[\frac{\partial q\left(a_t,x_t;\theta_0,g_{\theta_0}\right)}{\partial\theta}\frac{\partial q\left(x_0,a_0;\theta_0,g_{\theta_0}\right)'}{\partial\theta}\right] \\ +E\left[\frac{\partial q\left(a_t,x_t;\theta_0,g_{\theta_0}\right)}{\partial\theta}\frac{\partial q\left(x_0,a_0;\theta_0,g_{\theta_0}\right)'}{\partial\theta}\right]' \end{array}\right).
\end{aligned}
$$

Note that $E\left[\frac{\partial q\left(a_t,x_t;\theta_0,g_{\theta_0}\right)}{\partial\theta}\right] = 0$ by definition of $\theta_0$. Next we show that $D_{2,T}$ also converges to a normal vector at the rate $1/\sqrt{T}$. Consider the $j$-th element of $D_{2,T}$, using the expression from the score function defined in (70),

$$
\begin{aligned}
\left(D_{2,T}\right)_j &= \frac{1}{\sqrt{T}}\sum_{t=1}^{T} c_{t,T}\left(\frac{\partial \widehat{v}_{\theta_0}\left(\widetilde{a}, x_t\right)}{\partial \theta_j} - \frac{\partial v_{\theta_0}\left(\widetilde{a}, x_t\right)}{\partial \theta_j}\right) \\
&- \frac{1}{\sqrt{T}}\sum_{t=1}^{T} c_{t,T}\left(\begin{array}{c} \left(\sum_{\widetilde{a}\in A}\frac{\partial \widehat{v}_{\theta_0}(\widetilde{a},x_t)}{\partial\theta_j}\exp\left(\widehat{v}_{\theta_0}\left(\widetilde{a},x_t\right)\right)\right) / \left(\sum_{\widetilde{a}\in A}\exp\left(\widehat{v}_{\theta_0}\left(\widetilde{a},x_t\right)\right)\right) \\ -\left(\sum_{\widetilde{a}\in A}\frac{\partial v_{\theta_0}(\widetilde{a},x_t)}{\partial\theta_j}\exp\left(v_{\theta_0}\left(\widetilde{a},x_t\right)\right)\right) / \left(\sum_{\widetilde{a}\in A}\exp\left(v_{\theta_0}\left(\widetilde{a},x_t\right)\right)\right) \end{array}\right),
\end{aligned}
$$

linearizing,

$$
\begin{aligned}
\left(D_{2,T}\right)_j &= \frac{1}{\sqrt{T}}\sum_{t=1}^{T} c_{t,T}\left(\frac{\partial \widehat{v}_{\theta_0}\left(a_t, x_t\right)}{\partial \theta_j} - \frac{\partial v_{\theta_0}\left(a_t, x_t\right)}{\partial \theta_j}\right) \\
&- \frac{1}{\sqrt{T}}\sum_{t=1}^{T}\sum_{\widetilde{a}\in A} c_{t,T}\psi_1\left(\widetilde{a}, x_t\right)\left(\frac{\partial \widehat{v}_{\theta_0}\left(\widetilde{a}, x_t\right)}{\partial \theta_j} - \frac{\partial v_{\theta_0}\left(\widetilde{a}, x_t\right)}{\partial \theta_j}\right) \\
&- \frac{1}{\sqrt{T}}\sum_{t=1}^{T}\sum_{\widetilde{a}\in A} c_{t,T}\psi_{2,j}\left(\widetilde{a}, x_t\right)\left(\widehat{v}_{\theta_0}\left(\widetilde{a}, x_t\right) - v_{\theta_0}\left(\widetilde{a}, x_t\right)\right) \\
&+ \frac{1}{\sqrt{T}}\sum_{t=1}^{T}\sum_{\widetilde{a}\in A} c_{t,T}\psi_{2,j}\left(\widetilde{a}, x_t\right)\left(\sum_{\widetilde{\widetilde{a}}\in A} P\left(\widetilde{\widetilde{a}}|x_t\right)\left(\widehat{v}_{\theta_0}\left(\widetilde{\widetilde{a}}, x_t\right) - v_{\theta_0}\left(\widetilde{\widetilde{a}}, x_t\right)\right)\right) + o_p(1) \\
&= \frac{1}{\sqrt{T}}\sum_{t=1}^{T}\left(E_{1,t,T}\right)_j + \frac{1}{\sqrt{T}}\sum_{t=1}^{T}\left(E_{2,t,T}\right)_j + \frac{1}{\sqrt{T}}\sum_{t=1}^{T}\left(E_{3,t,T}\right)_j + \frac{1}{\sqrt{T}}\sum_{t=1}^{T}\left(E_{4,t,T}\right)_j + o_p(1),
\end{aligned}
$$

where

$$\psi_1(\widetilde{a}, x_t) = P(\widetilde{a}|x_t), \tag{72}$$

$$\psi_{2,j}(\widetilde{a}, x_t) = P(\widetilde{a}|x_t)\frac{\partial v_{\theta_0}(\widetilde{a}, x_t)}{\partial \theta_j}, \tag{73}$$

and the remainder terms are of smaller order since our nonparametric estimates converge uniformly to the true at the rate faster than $T^{-1/4}$ on the trimming set, as proven in Theorem 1 and 2.

The asymptotic properties of these terms are tedious but simple to obtain. We utilize the projection results and law of large numbers for U-statistics, see Lee (1990). We also note that all of the relevant kernels for our statistics are uniformly bounded, along with the assumption [B1.1], this ensures the residuals from the projections can be ignored. Now we give some details for deriving the distribution of $\frac{1}{\sqrt{T}}\sum_{t=1}^{T}(E_{1,t,T})_j$. First we linearize $\frac{\partial^k \widehat{g}_{\theta_0}}{\partial \theta_j^k} - \frac{\partial^k g_{\theta_0}}{\partial \theta_j^k}$ for $k = 0, 1$,

$$\begin{aligned}\frac{\partial^k \widehat{g}_{\theta_0}}{\partial \theta_j^k} - \frac{\partial^k g_{\theta_0}}{\partial \theta_j^k} &= \widehat{\mathcal{H}}\frac{\partial^k \widehat{m}_{\theta_0}}{\partial \theta_j^k} - \mathcal{H}\frac{\partial^k m_{\theta_0}}{\partial \theta_j^k} \\ &\sim \left(\widehat{\mathcal{H}} - \mathcal{H}\right)\frac{\partial^k m_{\theta_0}}{\partial \theta_j^k} + \mathcal{H}(I - \mathcal{L})^{-1}\left(\widehat{\mathcal{L}} - \mathcal{L}\right)\frac{\partial^k m_{\theta_0}}{\partial \theta_j^k} \\ &\quad + \mathcal{H}(I - \mathcal{L})^{-1}\left(\frac{\partial^k \widehat{r}_{\theta_0}}{\partial \theta_j^k} - \frac{\partial^k r_{\theta_0}}{\partial \theta_j^k}\right),\end{aligned}$$

this expansion is valid, uniformly on the trimming set, inspite of the scaling of order $\sqrt{T}$. Consider the normalized sum of $\left(\widehat{\mathcal{H}} - \mathcal{H}\right)\frac{\partial^k m_{\theta_0}}{\partial \theta_j^k}$, with further linearization, see the decomposition $\widehat{\mathcal{L}} - \mathcal{L}$ and $\widehat{\mathcal{H}} - \mathcal{H}$ in the proof of [A1], we can obtain the following

U-statistics, scaled by $\sqrt{T}$, representation,

$$\frac{1}{\sqrt{T}}\sum_{t=1}^{T-1}\mathfrak{c}_{t,T}\left[\left(\widehat{\mathcal{H}}-\mathcal{H}\right)\frac{\partial m_{\theta_0}}{\partial\theta_j}(x_t,a_t)\right]$$

$$= \frac{1}{\sqrt{T}}\frac{1}{T-1}\sum_{t=1}^{T-1}\sum_{s\neq t}\mathfrak{c}_{t,T}\left(\begin{array}{c}\frac{\partial m_{\theta_0}(x_{s+1})}{\partial\theta_j}\frac{K_h(x_s-x_t)1[a_s=a_t]}{f_{X,A}(x_t,a_t)}-E\left[\left.\frac{\partial m_{\theta_0}(x_{t+1})}{\partial\theta_j}\right|x_t,a_t\right]\\ -\frac{K_h(x_s-x_t)1[a_s=a_t]-f_{X,A}(x_t,a_t)}{f_{X,A}(x_t,a_t)}E\left[\left.\frac{\partial m_{\theta_0}(x_{t+1})}{\partial\theta_j}\right|x_t,a_t\right]\end{array}\right)+o_p(1)$$

$$= \sqrt{T}\left(\begin{array}{c}T-1\\2\end{array}\right)^{-1}\sum_{t=1}^{T-1}\sum_{s>t}\frac{1}{2}\left(\begin{array}{c}\mathfrak{c}_{t,T}\frac{\partial m_{\theta_0}(x_{s+1})}{\partial\theta_j}\frac{K_h(x_s-x_t)1[a_s=a_t]}{f_{X,A}(x_t,a_t)}-\mathfrak{c}_{t,T}E\left[\left.\frac{\partial m_{\theta_0}(x_{t+1})}{\partial\theta_j}\right|x_t,a_t\right]\\ +\mathfrak{c}_{s,T}\frac{\partial m_{\theta_0}(x_{t+1})}{\partial\theta_j}\frac{K_h(x_t-x_s)1[a_t=a_s]}{f(x_s,a_s)}-\mathfrak{c}_{s,T}E\left[\left.\frac{\partial m_{\theta_0}(x_{t+1})}{\partial\theta_j}\right|x_s,a_s\right]\end{array}\right),$$

$$-\sqrt{T}\left(\begin{array}{c}T-1\\2\end{array}\right)^{-1}\sum_{t=1}^{T-1}\sum_{s>t}\frac{1}{2}\left(\begin{array}{c}\mathfrak{c}_{t,T}\frac{K_h(x_s-x_t)1[a_s=a_t]-f_{X,A}(x_t,a_t)}{f_{X,A}(x_t,a_t)}E\left[\left.\frac{\partial m_{\theta_0}(x_{t+1})}{\partial\theta_j}\right|x_t,a_t\right]\\ +\mathfrak{c}_{s,T}\frac{K_h(x_t-x_s)1[a_t=a_s]-f(x_s,a_s)}{f(x_s,a_s)}E\left[\left.\frac{\partial m_{\theta_0}(x_{t+1})}{\partial\theta_j}\right|x_s,a_s\right]\end{array}\right)+o_p$$

Hoeffding (H-)decomposition provides the following as leading term, disposing the trimming factor,

$$\frac{1}{\sqrt{T}}\sum_{t=1}^{T-1}\left(\frac{\partial m_{\theta_0}(x_{t+1})}{\partial\theta_j}-E\left[\left.\frac{\partial m_{\theta_0}(x_{t+1})}{\partial\theta_j}\right|x_t,a_t\right]\right). \tag{74}$$

To obtain the projection of the second term is more labor intensive. We first split it up into two parts,

$$\frac{1}{\sqrt{T}}\sum_{t=1}^{T-1}\mathfrak{c}_{t,T}\left[\mathcal{H}(I-\mathcal{L})^{-1}\left(\widehat{\mathcal{L}}-\mathcal{L}\right)\frac{\partial^k m_{\theta_0}}{\partial\theta_j^k}(x_t,a_t)\right]$$

$$= \frac{1}{\sqrt{T}}\sum_{t=1}^{T-1}\mathfrak{c}_{t,T}\left[\mathcal{H}\left(\widehat{\mathcal{L}}-\mathcal{L}\right)\frac{\partial^k m_{\theta_0}}{\partial\theta_j^k}(x_t,a_t)\right]+\frac{1}{\sqrt{T}}\sum_{t=1}^{T-1}\mathfrak{c}_{t,T}\left[\mathcal{H}\mathcal{L}(I-\mathcal{L})^{-1}\left(\widehat{\mathcal{L}}-\mathcal{L}\right)\frac{\partial^k m_{\theta_0}}{\partial\theta_j^k}(x_t,c$$

The summands of the first term takes the following form

$$\mathfrak{c}_{t,T}\beta\int\left(\begin{array}{c}\int\frac{\partial m_{\theta_0}(x'')}{\partial\theta_j}\frac{1}{f_X(x')}\left(\widehat{f}_{X'X}(dx'',x')-f_{X'X}(dx'',x')\right)\\ -\frac{\widehat{f}_X(x')-f_X(x')}{f_X(x')}E\left[\left.\frac{\partial m_{\theta_0}(x_{t+2})}{\partial\theta_j}\right|x_{t+1}=x'\right]\end{array}\right)f_{X'|X,A}\left(dx'|x_t,a_t\right),$$

with standard change of variable and usual symmetrization, this leads to the following kernel for the U-statistic,

$$\frac{\beta}{2}\left(\begin{array}{c} \mathfrak{c}_{t,T}\frac{\partial m_{\theta_0}(x_{s+1})}{\partial\theta_j}\frac{f'_{X|X,A}(x_s|x_t,a_t)}{f_X(x_s)} - \mathfrak{c}_{t,T}E\left[E\left[\frac{\partial m_{\theta_0}(x_{t+2})}{\partial\theta_j}\middle| x_{t+1}\right]\middle| x_t,a_t\right] \\ +\mathfrak{c}_{s,T}\frac{\partial m_{\theta_0}(x_{t+1})}{\partial\theta_j}\frac{f'_{X|X,A}(x_t|x_s,a_s)}{f_X(x_t)} - \mathfrak{c}_{s,T}E\left[E\left[\frac{\partial m_{\theta_0}(x_{s+2})}{\partial\theta_j}\middle| x_{s+1}\right]\middle| x_s,a_s\right] \end{array}\right)$$

$$-\frac{\beta}{2}\left(\begin{array}{c} \mathfrak{c}_{t,T}E\left[\frac{\partial m_{\theta_0}(x_{s+1})}{\partial\theta_j}\middle| x_s\right]\frac{f'_{X|X,A}(x_s|x_t,a_t)}{f_X(x_s)} - \mathfrak{c}_{t,T}E\left[E\left[\frac{\partial m_{\theta_0}(x_{t+2})}{\partial\theta_j}\middle| x_{t+1}\right]\middle| x_t,a_t\right] \\ +\mathfrak{c}_{s,T}E\left[\frac{\partial m_{\theta_0}(x_{t+1})}{\partial\theta_j}\middle| x_t\right]\frac{f'_{X|X,A}(x_t|x_s,a_s)}{f_X(x_t)} - \mathfrak{c}_{s,T}E\left[E\left[\frac{\partial m_{\theta_0}(x_{s+2})}{\partial\theta_j}\middle| x_{s+1}\right]\middle| x_s,a_s\right] \end{array}\right),$$

The leading term from H-decomposition leads to the following centered process

$$\frac{1}{\sqrt{T}}\sum_{t=1}^{T-1}\beta\left(\frac{\partial m_{\theta_0}(x_{t+1})}{\partial\theta_j} - E\left[\frac{\partial m_{\theta_0}(x_{t+1})}{\partial\theta_j}\middle| x_t\right]\right), \tag{75}$$

notice the conditional expectation term is a two-step ahead predictor, zero mean follows from stationarity assumption and the law of iterated expectation. As for the second part of the second term, using the Neumann series representation, see (56) and (57), the kernel of the relevant U-statistics is,

$$\frac{\beta}{2}\left(\begin{array}{c} \mathfrak{c}_{t,T}\frac{\partial m_{\theta_0}(x_{s+1})}{\partial\theta_j}\int\frac{\varphi(x_s|x')}{f_X(x_s)}f'_{X|X,A}(dx'|x_t,a_t) - \mathfrak{c}_{t,T}\sum_{\tau=1}^{\infty}\beta^{\tau}E\left[E\left[\frac{\partial m_{\theta_0}(x_{t+\tau+2})}{\partial\theta_j}\middle| x_{t+1}\right]\middle| x_t, a \\ +\mathfrak{c}_{s,T}\frac{\partial m_{\theta_0}(x_{t+1})}{\partial\theta_j}\int\frac{\varphi(x_t|x')}{f_X(x_t)}f'_{X|X,A}(dx'|x_s,a_s) - \mathfrak{c}_{s,T}\sum_{\tau=1}^{\infty}\beta^{\tau}E\left[E\left[\frac{\partial m_{\theta_0}(x_{s+\tau+2})}{\partial\theta_j}\middle| x_{s+1}\right]\middle| x_s, \end{array}\right.$$

$$-\frac{\beta}{2}\left(\begin{array}{c} \mathfrak{c}_{t,T}E\left[\frac{\partial m_{\theta_0}(x_{s+1})}{\partial\theta_j}\middle| x_s\right]\int\frac{\varphi(x_s|x')}{f_X(x_s)}f'_{X|X,A}(dx'|x_t,a_t) - \mathfrak{c}_{t,T}\sum_{\tau=1}^{\infty}\beta^{\tau}E\left[E\left[\frac{\partial m_{\theta_0}(x_{t+\tau+2})}{\partial\theta_j}\middle| x \right.\right. \\ +\mathfrak{c}_{s,T}E\left[\frac{\partial m_{\theta_0}(x_{t+1})}{\partial\theta_j}\middle| x_t\right]\int\frac{\varphi(x_t|x')}{f_X(x_t)}f'_{X|X,A}(dx'|x_s,a_s) - \mathfrak{c}_{s,T}\sum_{\tau=1}^{\infty}\beta^{\tau}E\left[E\left[\frac{\partial m_{\theta_0}(x_{s+\tau+2})}{\partial\theta_j}\middle| : \right.\right. \end{array}\right.$$

where $\varphi$ is defined as the limit of discounted sum of the conditional densities, see (58). The projection of the U-statistic with the above kernel yields,

$$\frac{1}{\sqrt{T}}\sum_{t=1}^{T-1}\left(\frac{\beta^2}{1-\beta}\right)\left(\frac{\partial m_{\theta_0}(x_{t+1})}{\partial\theta_j} - E\left[\frac{\partial m_{\theta_0}(x_{t+1})}{\partial\theta_j}\middle| x_t\right]\right). \tag{76}$$

The last term of $\frac{1}{\sqrt{T}}\sum_{t=1}^{T}\left(E_{1,t,T}\right)_j$ can be treated similarly, recall we have

$$\begin{aligned}
&\mathcal{H}\left(I-\mathcal{L}\right)^{-1}\left(\frac{\partial^k\widehat{r}_{\theta_0}}{\partial\theta_j^k}-\frac{\partial^k r_{\theta_0}}{\partial\theta_j^k}\right)\\
=\;&\mathcal{H}\left(\frac{\partial^k\widehat{r}_{\theta_0}}{\partial\theta_j^k}-\frac{\partial^k r_{\theta_0}}{\partial\theta_j^k}\right)+\mathcal{H}\mathcal{L}\left(I-\mathcal{L}\right)^{-1}\left(\frac{\partial^k\widehat{r}_{\theta_0}}{\partial\theta_j^k}-\frac{\partial^k r_{\theta_0}}{\partial\theta_j^k}\right).
\end{aligned}$$

Ignoring the bias term, that is negligible under assumptions B1.6 and B1.7,

$$\begin{aligned}
&\mathcal{H}\left(\frac{\partial^k\widehat{r}_{\theta_0}}{\partial\theta_j^k}-\frac{\partial^k r_{\theta_0}}{\partial\theta_j^k}\right)(a_t,x_t)\\
=\;&\frac{1}{T-1}\sum_{\widetilde{a}\in A}\sum_{s\neq t}\int\left(\frac{\partial^k\zeta'_{x,\widetilde{a},\theta_0}\left(P\left(\widetilde{a}|x'\right)\right)}{\partial\theta_j^k}\frac{e_{\widetilde{a},s}K_h\left(x_s-x'\right)}{f_X\left(x'\right)}\right)f'_{X|X,A}\left(dx'|x_t,a_t\right)+o_p\left(T^{-1/2}\right),\\
=\;&\frac{1}{T-1}\sum_{\widetilde{a}\in A}\sum_{s\neq t}f'_{X|X,A}\left(x_s|x_t,a_t\right)\frac{\partial^k\zeta'_{x_s,\widetilde{a},\theta_0}\left(P\left(\widetilde{a}|x\right)\right)}{\partial\theta_j^k}\frac{e_{\widetilde{a},s}}{f_X\left(x_s\right)}+o_p\left(T^{-1/2}\right).
\end{aligned}$$

Normalizing the projection of the corresponding U-statistics obtains

$$\frac{1}{\sqrt{T}}\sum_{t=1}^{T}\mathcal{H}\left(\frac{\partial^k\widehat{r}_{\theta_0}}{\partial\theta_j^k}-\frac{\partial^k r_{\theta_0}}{\partial\theta_j^k}\right)(a_t,x_t)=\frac{1}{\sqrt{T}}\sum_{\widetilde{a}\in A}\sum_{t=1}^{T}\frac{\partial^k\zeta'_{x_t,\widetilde{a},\theta_0}\left(P\left(\widetilde{a}|x_t\right)\right)}{\partial\theta_j^k}e_{\widetilde{a},t}+o_p\left(1\right).\tag{77}$$

The same can be done to the remaining term, in particular we obtain

$$\begin{aligned}
&\frac{1}{\sqrt{T}}\sum_{t=1}^{T-1}\mathcal{H}\mathcal{L}\left(I-\mathcal{L}\right)^{-1}\left(\frac{\partial^k\widehat{r}_{\theta_0}}{\partial\theta_j^k}-\frac{\partial^k r_{\theta_0}}{\partial\theta_j^k}\right)(a_t,x_t)\\
=\;&\frac{1}{\sqrt{T}}\sum_{\widetilde{a}\in A}\sum_{t=1}^{T}\frac{\beta}{1-\beta}\frac{\partial^k\zeta'_{x_t,\widetilde{a},\theta_0}\left(P\left(\widetilde{a}|x_t\right)\right)}{\partial\theta_j^k}e_{\widetilde{a},t}+o_p\left(1\right).
\end{aligned}\tag{78}$$

Collecting (74) - (78), for $k=1$,we obtain the leading terms of $\frac{1}{\sqrt{T}}\sum_{t=1}^{T}\left(E_{1,t,T}\right)_j$. For $\frac{1}{\sqrt{T}}\sum_{t=1}^{T}\left(E_{2,t,T}\right)_j$ and $\frac{1}{\sqrt{T}}\sum_{t=1}^{T}\left(E_{3,t,T}\right)_j$, we again use the projection technique of the U-statistics to obtain their leading terms. We gave a lot of details for the former case as remaining terms in $\left(D_{2,T}\right)$ can be treated in a similar fashion. In particular, it is simple to show that the projections of various relevant U-statistics, defined below with some elements $\varpi_k\in C\left(X\right),\varsigma_k\in C\left(A\times X\right)$ and $\widetilde{a}\in A$, have the following linear representation:

1. $\frac{1}{\sqrt{T}}\sum_{t=1}^{T}\varsigma_k(\widetilde{a},x_t)\left[\left(\widehat{\mathcal{H}}-\mathcal{H}\right)\varpi_k(\widetilde{a},x_t)\right]$
$= \frac{1}{\sqrt{T}}\sum_{t=1}^{T-1}\left(\frac{\varsigma_k(x_t,\widetilde{a})f_X(x_t)\mathbf{1}[a_t=\widetilde{a}]}{f(x_t,\widetilde{a})}\right)(\varpi_k(x_{t+1})-E[\varpi_k(x_{t+1})|\, x_t,a_t=\widetilde{a}])$
$+o_p(1)$.

2. $\frac{1}{\sqrt{T}}\sum_{t=1}^{T}\varsigma_k(\widetilde{a},x_t)\left[\mathcal{H}\left(\widehat{\mathcal{L}}-\mathcal{L}\right)\varpi_k(\widetilde{a},x_t)\right]$
$= \frac{1}{\sqrt{T}}\sum_{t=1}^{T-1}\beta\left(\frac{\int\varsigma_k(v,\widetilde{a})f'_{X|X,A}(x_t|v,\widetilde{a})f_X(dv)}{f_X(x_t)}\right)(\varpi_k(x_{t+1})-E[\varpi_k(x_{t+1})|\, x_t])$
$+o_p(1)$.

3. $\frac{1}{\sqrt{T}}\sum_{t=1}^{T}\varsigma_k(\widetilde{a},x_t)\left[\mathcal{H}\mathcal{L}(I-\mathcal{L})^{-1}\left(\widehat{\mathcal{L}}-\mathcal{L}\right)\varpi_k(\widetilde{a},x_t)\right]$
$= \frac{1}{\sqrt{T}}\sum_{t=1}^{T-1}\beta\left(\frac{\int\int\varsigma_k(v,\widetilde{a})\varphi(x_t|w)f'_{X|X,A}(dw|dv,\widetilde{a})f_X(dv)}{f_X(x_t)}\right)(\varpi_k(x_{t+1})-E[\varpi_k(x_{t+1})|\, x_t])$
$+o_p(1)$.

In correspondence of $(E_{k+1,t,T})_j$ for $k=1,2$, we have in mind

$$\begin{aligned}
\varsigma_1(\cdot) &= \psi_1(\widetilde{a},\cdot),\\
\varpi_1(\cdot) &= \frac{\partial m_{\theta_0}(\cdot)}{\partial\theta_j},\\
&\text{and}\\
\varsigma_2(\cdot) &= \psi_{2,j}(\widetilde{a},\cdot),\\
\varpi_2(\cdot) &= m_{\theta_0}(\cdot),
\end{aligned}$$

where $\psi_1$ and $\psi_{2,j}$ are defined in (72) - (73). Similarly, we also have

4. $\frac{1}{\sqrt{T}}\sum_{t=1}^{T}\varsigma_k(\widetilde{a},x_t)\mathcal{H}\left(\frac{\partial^k\widehat{r}_{\theta_0}}{\partial\theta_j^k}-\frac{\partial^k r_{\theta_0}}{\partial\theta_j^k}\right)(\widetilde{a},x_t)$
$= \frac{1}{\sqrt{T}}\sum_{a^*\in A}\sum_{t=1}^{T-1}\left(\int\varsigma_k(v,\widetilde{a})\,f'_{X|X,A}(x_t|v,\widetilde{a})\,f_X(dv)\right)\frac{\partial^k\zeta'_{x_t,\widetilde{a},\theta_0}(P(a^*|x_t))}{\partial\theta_j^k}\frac{e_{a^*,t}}{f_X(x_t)}$
$+o_p(1)$.

5. $\frac{1}{\sqrt{T}}\sum_{t=1}^{T}\varsigma_k(\widetilde{a},x_t)\left[\mathcal{H}\mathcal{L}(I-\mathcal{L})^{-1}\left(\frac{\partial^k\widehat{r}_{\theta_0}}{\partial\theta_j^k}-\frac{\partial^k r_{\theta_0}}{\partial\theta_j^k}\right)(\widetilde{a},x_t)\right]$
$= \frac{1}{\sqrt{T}}\sum_{a^*\in A}\sum_{t=1}^{T}\left[\int\int\varsigma_k(v,\widetilde{a})\,\varphi(x_t|w)\,f'_{X|X,A}(dw|v,\widetilde{a})\,f_X(dv)\right]\frac{\partial^k\zeta'_{x_t,\widetilde{a},\theta_0}(P(a^*|x_t))}{\partial\theta_j^k}\frac{e_{a^*,t}}{f_X(x_t)}$
$+o_p(1)$.

Notice that leading terms from all the projections above are mean zero processes. Collecting these terms, lots of covariance. Clearly $\frac{1}{\sqrt{T}}\sum_{t=1}^{T}(E_{k,t,T})_j = O_p(1)$ for $k = 1,2,3$ and $j = 1,\ldots,q$, this ensures the root-$T$ consistency $\widehat{\theta}$. The term $D_{3,T}$ is $o_p(1)$ since $\frac{\partial q(a_t,x_t;\theta_0,g_{\theta_0})}{\partial\theta}$ is uniformly bounded and $\mathfrak{d}_{t,T} = o_p\left(\sqrt{T}\right)$ for all $t$. In sum,

$$\sqrt{T}D_{2,T} \Rightarrow N(0,\Lambda_2),$$

$$\Lambda_2 = \lim_{T\to\infty} Var\left(\frac{1}{\sqrt{T}}\sum_{t=1}^{T}(E_{1,t,T}+E_{2,t,T}+E_{3,t,T})\right),$$

$$\sqrt{T}\left(\widehat{\theta}-\theta_0\right) \Rightarrow N\left(0,\mathcal{J}^{-1}\mathcal{I}\mathcal{J}^{-1}\right),$$

$$\begin{aligned}
\mathcal{I} &= \lim_{T\to\infty} Var(D_{1,T}+D_{2,T}), \\
\mathcal{J} &= E\left[\frac{\partial^2 q(a_t,x_t;\theta_0,g_{\theta_0})}{\partial\theta\partial\theta'}\right].
\end{aligned}$$

■

Proof of Theorem 1.4 and 1.5: Under the assumed smoothness assumptions, the results simply follow from MVT. ■

### 1.8.5 Additional Proofs

We now show various centered processes in the previous section converge uniformly at desired rates on a compact set X. We outline the main steps below and proof the results for relevant cases. The methodology here is similar to [LM] who employed the exponential inequality from [B] for various quantities similar to ours.

Consider some process $\mathfrak{l}_T(x) = \frac{1}{T}\sum l(x_t,x)$, where $l(x_t,x)$ has mean zero. For some positive sequence, $\delta_T$, converging monotonically to zero, we first show that $|\mathfrak{l}_T(x)| = o_p(\delta_T)$ pointwise on X, then we use the continuity property of $l(x_t,x)$ to show that this rate of convergence is preserved uniformly over X.

To obtain the pointwise rates, specializing Theorem 1.3 of [B], we have the following inequality.

$$\begin{aligned}
\Pr\left(\left|\mathrm{I}_T(x)\right| > \delta_T\right) &\leq 4\exp\left(-\frac{\delta_T^2 T^\beta}{8v^2\left(T^\beta\right)}\right) + 22\left(1+\frac{4b_T}{\delta_T}\right)^{1/2} T^\beta \alpha\left(\left\lfloor\frac{T^{1-\beta}}{2}\right\rfloor\right) \\
&\leq \exp\left(-G_{1,T}\right) + G_{2,T},
\end{aligned}$$

for some

$$\begin{aligned}
\beta &\in (0,1), \qquad (79)\\
b_T &= O\left(\sup_{(x',x)\in\mathcal{X}\times\mathcal{X}} l\left(x',x\right)\right), \\
v^2(\beta) &= \operatorname{var}\left(\frac{1}{\left\lfloor\frac{T^\beta}{2}+1\right\rfloor}\sum_{t=1}^{\left\lfloor\frac{T^\beta}{2}+1\right\rfloor} l\left(x_t,x\right)\right) + \frac{b_T\delta_T}{2}.
\end{aligned}$$

To have the first term converging to zero, at an exponential rate, we need $G_{1,T} \to \infty$. The main calculation here is the variance term in $v^2$. Following [M], we can generally show that the uniform order of such term comes from the variances and the covariances terms are of smaller order. We note that the bounds on these variances are independent on the trimming set. For our purposes, the natural choice of $\delta_T^2$ often reduces us to choosing $\beta$ to satisfy $b\delta_T = o\left(\delta_T^2 T^\beta\right)$. The rate of $G_{2,T}$ is easy to control since all of the quanitities involved increase (decrease) at a power rate, the mixing coefficient can be made to decay sufficiently fast so $G_{2,T} = O\left(T^{-\eta}\right)$ for some $\eta > 0$, hence $\Pr\left(\left|\mathrm{I}_T(x)\right| > \delta_T\right) = O\left(T^{-\eta}\right)$.

To obtain the uniform rates over $\mathcal{X}$, compactness implies there exist an increasing number, $Q_T$, of shrinking hyper-cubes $\{I_{Q_T}\}$ whose length of each side is $\{\epsilon_T\}$ with centres $\{x^{Q_T}\}$. These cubes cover $\mathcal{X}$, namely for some $C_0$ and $d$,

$$\epsilon_T^d Q_T \leq C_0 < \infty.$$

In particular, we will have $Q_T$ grow at a power rate in our applications. Then we have

$$\begin{aligned}\Pr\left(\sup_x |\mathsf{l}_T(x)| > \delta_T\right) &\leq \Pr\left(\max_{1\leq q\leq Q_T} |\mathsf{l}_T(x^q)| > \delta_T\right) + \Pr\left(\max_{1\leq q\leq Q_T}\sup_{x\in I_q} |\mathsf{l}_T(x) - \mathsf{l}_T(x^q)| > \delta_T\right) \\ &= G_{3,T} + G_{4,T},\end{aligned}$$

where $G_{3,T} = O\left(Q_T T^{-\eta}\right)$ by Bonferroni Inequality. Provided the rate of decay of the mixing coefficient, i.e. $\eta$, is sufficiently large relative to the rate $Q_T$ grows we shall have $Q_T = o(T^{\eta})$. For the second term, since the opposing behavior of $(\epsilon_T, Q_T)$ is independent of the mixing coefficient, $\max_{1\leq q\leq Q_T}\sup_{x\in I_q} |\mathsf{l}_T(x) - \mathsf{l}_T(x^q)| = o(\delta_T)$ can be shown using Liptschitz continuity when the hyper cubes shrink sufficiently fast.

Before we proceed with the specific cases we validate our treatment of the trimming factor. The pointwise rates are clearly unaffected by bias at the boundary so long $x \in int(X)$. The technique used to obtain uniformity also accommodates expanding space $\mathcal{X}$, so long we use the sequence $\{c_T\}$ to satisfy condition stated in [B1.9]. The uniform rate of convergence is also unaffected, when replace $\mathcal{X}$ with $X_T$, since the covering of an expanding of a compact subsets of a compact set can still grow (and shrink) at the same rate in each of the cases below. Therefore we could replace $\mathcal{X}$ everywhere by $X_T$.

Combining the results of uniform convergence of the zero mean processes and their biases, the uniform rates to various quantities in the previous section can now be established. We note that the treatment to allow for additional discrete observable states only requires trivial extension. We provide illustrate this for the first case of kernel density estimation, and for brevity, thenceforth assume that we only have purely continuous observable state variables.

Case 1

Here we deal with density estimators such as $\widehat{f}_X(x)$, $\widehat{f}_{X',X}(x',x)$ and $\widehat{f}_{X',X,A}(x',x,j)$ :

We first establish the pointwise rate of convergence of a de-meaned kernel density estimator.

$$\begin{aligned}
\mathsf{l}_T(x) &= \widehat{f}_X(x) - E\widehat{f}_X(x), \\
l(x_t, x) &= \prod_{l=0}^{d-1} K_h(x_{t-l} - x_{l+1}) - E\prod_{l=0}^{d-1} K_h(x_{t-l} - x_{l+1}).
\end{aligned}$$

The main elements for studying the rate of $G_{1,T}$ are

$$\begin{aligned}
\varpi &= \frac{1}{\sqrt{Th^d}}, \\
\delta_T &= T^{\xi}\varpi \text{ for some } \xi > 0, \\
b_T &= O\left(h^{-d}\right), \\
v^2\left(T^{\beta}\right) &= O\left(\varpi^2 T^{1-\beta} \vee T^{\xi}\varpi h^{-d}\right).
\end{aligned}$$

We obtain from simple algebra

$$G_{1,T} = O\left(\frac{T^{2\xi}T^{\beta}}{T^{1-\beta} + T^{\xi}T^{1/2}h^{-d/2}}\right).$$

As mentioned in the previous section, we have $d = 2$ and $h = O\left(T^{-1/5}\right)$. This means $\delta_T = T^{\xi-3/10}$, and if $\beta \in (7/10, 1)$ then we have $G_{1,T} \to \infty$ . Clearly, the same choice of $\beta$ will suffice for $d = 1$ as well.

To make this uniform on $X_T$, with product kernels and the Lipschitz continuity of $K$, we have for any $(x, x_q) \in I_q$,

$$|K_h(x_t - x) - K_h(x_t - x_q)| \le \frac{C_1}{h^3}\epsilon_T.$$

So it follows that

$$
\begin{aligned}
\delta_T^{-1} \max_{1 \leq q \leq Q_T} \sup_{x \in I_q} \left| \mathsf{l}_T(x) - \mathsf{l}_T(x_q) \right| &= O\left(\frac{\epsilon_T}{\delta_T h^3}\right) \\
&= O\left(\frac{T^{-\zeta/2}}{T^{\xi - 9/10}}\right).
\end{aligned}
$$

Define $Q_T = T^{\zeta}$, for some $\zeta > 0$, this requires $9/5 < \zeta < \eta$.

We can allow for additional discrete control variable and/or observable state variables. As an illustration, consider the density estimator of one continuous random variable and some discrete random variable, we have

$$
\begin{aligned}
\mathsf{l}_T(x) &= \widehat{f}_{X^C, X^D}(x_c, x_d) - E\widehat{f}_{X^C, X^D}(x_c, x_d), \\
l(x_t, x) &= K_h(x_{c,t} - x_c)\, \mathbf{1}(x_{d,t} = x_d) - EK_h(x_{c,t} - x_c)\, \mathbf{1}(x_{d,t} = x_d).
\end{aligned}
$$

Same rates as the purely continuous case apply. For the pointwise part, the variance is clearly of the same order. For the bounds on the uniform rates observe that,

$$
\left| K_h(x_{c,t} - x_c)\, \mathbf{1}(x_{d,t} = x_d) - K_h(x_{c,t} - x_c^q)\, \mathbf{1}(x_{d,t} = x_d) \right| \leq \left| K_h(x_{c,t} - x_c) - K_h(x_{c,t} - x_c^q) \right|.
$$

Same reasoning also applies for the kernel estimator of the density of the control and observable state variables.

CASE 2

Here we deal with $\widehat{r}^C(x)$ :

$$
l(x_t, x) = \frac{e_{\theta,t} K_h(x_t - x)}{f_X(x)}.
$$

Since $\{e_{\theta,T}\}$ is uniformly bounded (a.s.) it follows, as shown in Case 1, the choice $\beta \in (3/5, 1)$ will do to have $G_{1,T} \to \infty$.

To make this uniform on $X_T$, by boundedness of $\{e_{\theta,T}\}$, Lipschitz continuity of $K, f$ and their appropriate bounds, we have for any $(x, x_q) \in I_q$,

$$\left|K_h(x_t - x) - K_h(x_t - x_q)\right| \leq \frac{C_2}{h^2}\epsilon_T.$$

So it follows that

$$\begin{aligned}\delta_T^{-1} \max_{1\leq q\leq Q_T} \sup_{x\in I_q} \left|\mathfrak{l}_T(x) - \mathfrak{l}_T(x_q)\right| &= O\left(\frac{T^{-\zeta}}{T^{\xi-7/10}}\right) \\ &= o(1),\end{aligned}$$

for some $\zeta > 0$, this requires $7/10 < \zeta < \eta$.

Case 3

Here we deal with $\mathcal{L}(I - \mathcal{L})^{-1}\hat{r}_\theta^{\mathcal{C}}(x)$ :

$$l(x_t, x) = e_{\theta,t}\nu(x_t, x),$$

where the definition of $\nu$ is provided in (60). Using Billingsley's Inequality, it is straightforward to show that with the additional smoothing, the variance of $\mathfrak{l}_T$ is of parametric rate uniformly on $X_T$. Selecting $\beta \in (1/2, 1)$ will yield $G_{1,T} \to \infty$ for $\Pr\left(\left|\mathfrak{l}_T(x)\right| > T^{-2/5}\right) = o(1)$, for any $x \in X_T$.

To make this uniform on $X_T$, by boundedness of $\{e_{\theta,T}\}$ and Lipschitz continuity of $\varphi$, we have for any $(x, x_q) \in I_q$,

$$\left|e_{\theta,t}\nu(x_t, x) - e_{\theta,t}\nu(x_t, x_q)\right| \leq C_3\epsilon_T.$$

So it follows that

$$\delta_T^{-1} \max_{1 \leq q \leq Q_T} \sup_{x \in I_q} \left| \mathsf{I}_T(x) - \mathsf{I}_T(x_q) \right| = O\left( \frac{T^{-\zeta}}{T^{-2/5}} \right),$$

for some $\zeta > 0$, this requires $2/5 < \zeta < \eta$.

Case 4

Here we deal with $m^C_{1,\theta}(x)$ :

$$\mathsf{I}_T(x) = \frac{\beta}{f_X(x)} \int \left( \widehat{f}_{X',X}(x', x) - E\widehat{f}_{X',X}(x', x) \right) m_\theta(x')\, dx'.$$

As mentioned in the previous section, under our smoothness assumptions, we have uniformly on $X_T$,

$$\int \widehat{f}_{X',X}(x', x)\, m_\theta(x')\, dx' = \frac{1}{T-1} \sum_{t=1}^{T-1} K_h(X_t - x)\, m_\theta(X_{t+1}) + O\left(h^2\right).$$

The exact same choices found in Case 2 apply.

## 1.9 Tables

| | | $\widehat{\theta}_1$ | | | | |
|---|---|---|---|---|---|---|
| $T$ | $\varsigma$ | bias | mbias | std | iqr | mse |
| 100 | 1/8 | -0.1835 | 0.1103 | 1.4220 | 1.1469 | 2.0558 |
| | 1/4 | -0.2401 | 0.0520 | 1.4792 | 1.2459 | 2.2458 |
| | 3/8 | -0.1745 | 0.1242 | 1.3846 | 1.1022 | 1.9476 |
| 500 | 1/8 | 0.0857 | 0.1171 | 0.5254 | 0.4897 | 0.2834 |
| | 1/4 | 0.0176 | 0.0560 | 0.5418 | 0.5120 | 0.2939 |
| | 3/8 | 0.0245 | 0.0517 | 0.5676 | 0.5197 | 0.3228 |
| 1000 | 1/8 | 0.1725 | 0.1936 | 0.3602 | 0.3437 | 0.1595 |
| | 1/4 | 0.1011 | 0.1275 | 0.3664 | 0.3491 | 0.1445 |
| | 3/8 | 0.0970 | 0.1244 | 0.3753 | 0.3663 | 0.1503 |
| 2500 | 1/8 | 0.1580 | 0.1605 | 0.2232 | 0.2309 | 0.0748 |
| | 1/4 | 0.0884 | 0.0910 | 0.2277 | 0.2404 | 0.0597 |
| | 3/8 | 0.0793 | 0.0866 | 0.2316 | 0.2369 | 0.0599 |
| 5000 | 1/8 | 0.1438 | 0.1436 | 0.1647 | 0.1645 | 0.0478 |
| | 1/4 | 0.0764 | 0.0804 | 0.1684 | 0.1640 | 0.0342 |
| | 3/8 | 0.0668 | 0.0683 | 0.1714 | 0.1638 | 0.0338 |

Table 1: $h_\varsigma = 1.06s(NT)^{-\varsigma}$ is the bandwidth, for various choices of $\varsigma$, used in the nonparametric estimation, $s =$ denotes the standard deviation of $\{x_t\}_{t=1}^T$.

| $T$ | $\varsigma$ | bias | $\widehat{\theta}_2$ mbias | std | iqr | mse |
|---|---|---|---|---|---|---|
| 100 | 1/8 | -0.1495 | -0.1163 | 0.2357 | 0.2137 | 0.0779 |
| | 1/4 | -0.0958 | -0.0676 | 0.2211 | 0.1931 | 0.0580 |
| | 3/8 | -0.0425 | -0.0083 | 0.1953 | 0.1577 | 0.0399 |
| 500 | 1/8 | -0.0878 | -0.0840 | 0.0885 | 0.0867 | 0.0155 |
| | 1/4 | -0.0433 | -0.0396 | 0.0802 | 0.0763 | 0.0083 |
| | 3/8 | -0.0157 | -0.0124 | 0.0830 | 0.0728 | 0.0071 |
| 1000 | 1/8 | -0.0736 | -0.0728 | 0.0596 | 0.0592 | 0.0090 |
| | 1/4 | -0.0328 | -0.0316 | 0.0536 | 0.0548 | 0.0040 |
| | 3/8 | -0.0133 | -0.0138 | 0.0541 | 0.0538 | 0.0031 |
| 2500 | 1/8 | -0.0615 | -0.0605 | 0.0351 | 0.0358 | 0.0050 |
| | 1/4 | -0.0243 | -0.0240 | 0.0314 | 0.0321 | 0.0016 |
| | 3/8 | -0.0110 | -0.0111 | 0.0314 | 0.0315 | 0.0012 |
| 5000 | 1/8 | -0.0573 | -0.0570 | 0.0256 | 0.0252 | 0.0039 |
| | 1/4 | -0.0228 | -0.0227 | 0.0232 | 0.0236 | 0.0011 |
| | 3/8 | -0.0129 | -0.0129 | 0.0231 | 0.0234 | 0.0007 |

Table 2: $h_\varsigma = 1.06s(NT)^{-\varsigma}$ is the bandwidth, for various choices of $\varsigma$, used in the nonparametric estimation, $s =$ denotes the standard deviation of $\{x_t\}_{t=1}^T$.

| $T$ | $d$ | bias | $\widehat{\theta}_1$ mbias | std | iqr | mse |
|---|---|---|---|---|---|---|
| 100 | 2 | 1.7648 | 1.8530 | 0.6279 | 0.5871 | 3.5087 |
| | 3 | 1.6858 | 1.7495 | 1.0170 | 0.7800 | 3.8763 |
| | 4 | 2.0527 | 2.4576 | 1.0616 | 1.0683 | 5.3403 |
| | 5 | 1.8074 | 2.2551 | 1.2407 | 1.3758 | 4.8061 |
| 500 | 2 | 1.8606 | 1.8722 | 0.2469 | 0.2388 | 3.5227 |
| | 3 | 1.5970 | 1.6170 | 0.3036 | 0.2977 | 2.6424 |
| | 4 | 1.2602 | 1.2383 | 0.5112 | 0.3621 | 1.8494 |
| | 5 | 0.9338 | 0.8911 | 0.5825 | 0.4268 | 1.2113 |
| 1000 | 2 | 1.8878 | 1.8948 | 0.1811 | 0.1859 | 3.5966 |
| | 3 | 1.6311 | 1.6463 | 0.2069 | 0.2131 | 2.7035 |
| | 4 | 1.2393 | 1.2495 | 0.2690 | 0.2577 | 1.6083 |
| | 5 | 0.8906 | 0.9043 | 0.3093 | 0.3009 | 0.8889 |
| 2500 | 2 | 1.9075 | 1.9082 | 0.1103 | 0.1077 | 3.6509 |
| | 3 | 1.6532 | 1.6577 | 0.1280 | 0.1278 | 2.7494 |
| | 4 | 1.2528 | 1.2603 | 0.1589 | 0.1634 | 1.5948 |
| | 5 | 0.9177 | 0.9251 | 0.1773 | 0.1793 | 0.8736 |
| 5000 | 2 | 1.9045 | 1.9055 | 0.0807 | 0.0763 | 3.6335 |
| | 3 | 1.6495 | 1.6521 | 0.0924 | 0.0906 | 2.7295 |
| | 4 | 1.2513 | 1.2566 | 0.1161 | 0.1145 | 1.5793 |
| | 5 | 0.9124 | 0.9139 | 0.1323 | 0.1275 | 0.8500 |

Table 3: $d$ is the number of discretized support of $X$.

| $T$ | $\varsigma$ | bias | mbias | std | iqr | mse |
|---|---|---|---|---|---|---|
| | | | $\widehat{\theta}_2$ | | | |
| 100 | 2 | 0.0262 | 0.0364 | 0.1164 | 0.1050 | 0.0142 |
| | 3 | 0.0840 | 0.0883 | 0.1815 | 0.1405 | 0.0400 |
| | 4 | 0.1755 | 0.1915 | 0.1850 | 0.2097 | 0.0650 |
| | 5 | 0.1646 | 0.1843 | 0.1976 | 0.2394 | 0.0661 |
| 500 | 2 | 0.0451 | 0.0465 | 0.0472 | 0.0474 | 0.0043 |
| | 3 | 0.0486 | 0.0521 | 0.0571 | 0.0587 | 0.0056 |
| | 4 | -0.0031 | -0.0108 | 0.0959 | 0.0708 | 0.0092 |
| | 5 | -0.0327 | -0.0364 | 0.1059 | 0.0757 | 0.0123 |
| 1000 | 2 | 0.0490 | 0.0492 | 0.0333 | 0.0337 | 0.0035 |
| | 3 | 0.0525 | 0.0545 | 0.0380 | 0.0387 | 0.0042 |
| | 4 | -0.0105 | -0.0081 | 0.0491 | 0.0465 | 0.0025 |
| | 5 | -0.0443 | -0.0433 | 0.0541 | 0.0536 | 0.0049 |
| 2500 | 2 | 0.0533 | 0.0532 | 0.0210 | 0.0207 | 0.0033 |
| | 3 | 0.0573 | 0.0586 | 0.0237 | 0.0237 | 0.0038 |
| | 4 | -0.0077 | -0.0072 | 0.0285 | 0.0273 | 0.0009 |
| | 5 | -0.0395 | -0.0379 | 0.0309 | 0.0314 | 0.0025 |
| 5000 | 2 | 0.0522 | 0.0522 | 0.0152 | 0.0151 | 0.0030 |
| | 3 | 0.0562 | 0.0564 | 0.0167 | 0.0165 | 0.0034 |
| | 4 | -0.0087 | -0.0085 | 0.0210 | 0.0206 | 0.0005 |
| | 5 | -0.0408 | -0.0399 | 0.0229 | 0.0240 | 0.0022 |

Table 4: $d$ is the number of discretized support of $X$.

# 2 Dynamic Models with Continuous Control

## 2.1 Introduction

In this chapter, we develop a new estimator that is capable of estimating a class of Markovian decision processes with purely continuous control when one cannot utilize the Euler equation. Our estimation procedure is intuitive and it is also simple to implement since it does not solve the model equilibrium and, unlike the other existing estimator in the literature, we do not impose any parametric distributional assumption on the observables.

A well known obstacle in the estimation of many structural dynamic models in the empirical labor and industrial organization literature, regardless whether the controls are continuous, discrete or mixed, is the presence of the value functions. As seen in the previous chapter, the value functions and their corresponding continuation values generally have no closed form but are defined as solutions to some nonlinear functional equations. We show here how, analogous to the discrete choice framework, a two-step approach can be employed to estimate the value functions and continuation values in the first stage in order to reduce the burden of having to solve the model equilibrium. In particular, instead of solving out for the conditional value functions, one can use the linear characterization of the conditional value functions on the optimal path (i.e. the *policy value equation*) that is simple to estimate and solve. In a discrete choice setting, the policy value equation can be estimated nonparametrically by using Hotz and Miller's "inversion theorem". Of particular relevance to our methodology is the estimation of infinite horizon dynamic games with discrete actions, of Pesendorfer and Schmidt-Dengler (2008) who use Hotz and Miller's inversion theorem to estimate the conditional value function as a solution to some matrix equation in the first stage; the continuation value can then be estimated trivially and used to construct some least square criterion in the second stage.

We comment that there is comparatively less work on the development of estimation methodology with purely continuous control that have to deal with the presence of value functions.[9] This is in contrast to the well known subclass of a general Markov decision processes known as the Euler class, where one can bypass the issue of solving the Bellman's equation and use the Euler equation to generate some moment restrictions, for example see Hansen and Singleton (1982). However the general Markov decision models of significant economic interest do not fall into this class, for more details see Rust (1996). Our framework is more closely related to the study of dynamic auction and oligopoly models, which often allow for discrete choice as well (e.g. entry/exit decisions);[10] we refer to the surveys of Pakes (1994), and more recently, Ackerberg, Benkard, Berry and Pakes (2005).

Extending the idea of Hotz, Miller, Sanders and Smith (1994), Bajari, Benkard and Levin (2007), hereafter BBL, propose a closely related simulation estimator that is capable of estimating a large class of dynamic models that allows for continuous or discrete or mixed continuous-discrete controls. The "forward simulation" method of BBL uses the preliminary estimates of the policy function (optimal decision rule) and transition densities to simulate series of value functions for a given set of structural parameters; these simulated value functions are then used in constructing some minimum distance criterion based on the equilibrium conditions. The main assumption BBL use in estimating models that contain continuous control is that of *monotone choice*. We show that the monotone choice assumption can also be used to nonparametrically estimate the policy value equation, hence our methodology adopts HM's approach in

[9] The other two papers that we are aware of that estimates purely continuous control problem in the I.O. literature is Berry and Pakes (2002) and Hong and Shum (2009) but they are based on quite a different sets of assumptions.

[10] To our knowledge, Jofre-Bonet and Pesendorfer (2003) are the first to show that two-step estimation procedures can be used to estimate a dynamic game in their study of a repeated auction game.

the first stage estimation to estimate a continuous control problem. In addition, our estimator does not require any parametric specification of the transition law of the observables. This extra flexibility is of fundamental importance since the transition law is one of the model primitives that is required in the first stage estimation. In contrast, BBL explicitly require their preliminary estimator to converge at the parametric rate, this condition rules out the nonparametric estimation of the transition law on the observables whenever the control or the (observable) state variables are continuously distributed.

Although in this chapter we focus on models with observable state variables that take finitely many values, our estimator can also accommodate continuous state variable. As seen from the first chapter, we can easily allow the observable state variables to be continuously distributed. The main technical extension is that the policy value equation becomes an integral equation of type II, given the discounting factor, the solving of such equation is a well-posed inverse problem.

Our estimator originates from the large literature on minimum distance estimation, see the monograph by Koul (2002) for a review, where our criterion function measures the divergence between two estimators of the conditional distribution function. More specifically, we minimize some $L^2-$ distance between the nonparametric estimate of the conditional distribution function (implied by the data) to a simulated semiparametric counterpart (implied by the structural model). In finite samples, Monte Carlo integration causes our objective function to be discontinuous in the parameter, we use empirical process theory to ensure that our estimator converges to a normal random variable at the rate of $\sqrt{N}$ after an appropriate normalization. However, the asymptotic variance will generally be a complicated function(al) of various parameters; we discuss and propose the use of a semiparametric bootstrap method to estimate the standard errors. The analysis of the statistical properties of our estimator is similar to the work of Brown and Wegkamp (2002) on minimum distance from independence estimator, first introduced by Manski (1983). Brown and Wegkamp also show that nonparametric bootstrap can be used for inference in their problem. However, the estimator of Brown and Wegkamp does not depend on any preliminary estimator that converges slower than the rate of $\sqrt{N}$, so the treatment is essentially parametric. More recently, Komunjer and Santos (2009) consider the semiparametric problem of minimum distance estimators of nonseparable models under independence assumption. In this sense their work is more closely related to our estimator than that of Brown and Wegkamp. However, Komunjer and Santos use the method of sieves to simultaneously estimate their finite dimensional parameters and the infinite dimensional parameters in some sieve space and do not discuss estimation of the asymptotic variance. In our case, it is natural to take a two-step approach. The infinite dimensional parameter here is the continuation value function, which is defined as the regression of some unobervables to be estimated, and its structural relationship with the finite dimensional parameter is an essential feature in the methodology in this literature. We estimate the continuation value function using a simple Nadaraya-Watson estimator and provide its pointwise distribution theory.

The chapter proceeds as follows. The next section begins by describing the Markov decision model of interest for a single agent problem and provides a simple example that motivates our methodology, it then outlines the estimation strategy and discusses the computational aspect. Section 2.3 provides the conditions sufficient to obtain the desired distribution theory. We discuss inference based on semiparametric bootstrap in Section 2.4. Section 2.5 reports a Monte Carlo study of our estimator and illustrates the affects of ignoring the model dynamics. Section 2.6 concludes. The proofs of all theorems can be found in the Section 2.7. We collect the Figures and Tables at the end of the paper.

In this chapter: for any matrix $B = (b_{ij})$, define $\|B\|$ to be the Euclidean norm, namely $\sqrt{\lambda_{\max}(B'B)}$; when $\mathcal{G}$ is a class of real valued functions $g : A \times \Theta \rightarrow \mathbb{R}$, continuously defined on some compact Euclidean domain $A \times \Theta$, then denote $\|g\|_{\mathcal{G}} = \sup_{\theta\in\Theta} \|g(\cdot,\theta)\|_{\infty}$, where $\|g\|_{\infty} = \sup_{a\in A} |g(a)|$ is the usual supremum norm, and $\|g\|_{\mathcal{G}} = \|g\|_{\infty}$ when $g$ does not depend on $\theta$; when $\mathcal{G}^J$ is a class of continuous $\mathbb{R}^J$ valued functions $(g_j(\cdot,\theta))$, then denote $\|g\|_{\mathcal{G}} = \max_{1\leq j\leq J} \sup_{\theta\in\Theta} \|g_j(\cdot,\theta)\|_{\infty}$, where $\|g\|_{\infty} = \max_{1\leq j\leq J} \sup_{a\in A} |g_j(a)|$, and, $\|g\|_{\mathcal{G}} = \|g\|_{\infty}$ when $g$ does not depend on $\theta$.

## 2.2 Markov Decision Processes

### 2.2.1 Basic Framework

The random variables in the model are the control and state variables, denoted by $a_t$ and $s_t$ respectively. The support of control variable is a convex set $A \subset \mathbb{R}$. and the state space $S$ is a subset of $\mathbb{R}^{L+1}$. Time is indexed by $t$, the economic agent is forward looking in solving an infinite horizon intertemporal problem. In each period, the economic agent observes $s_t$ and chooses an action $a_t$ in order to maximize her discounted expected utility. The per period utility is time separable and is represented by a parametric function $u_\theta(a_t, s_t)$ for $\theta \in \Theta \subset \mathbb{R}^P$. The agent's action today affects the uncertain future states according to a Markovian transition law $p(ds_{t+1}|s_t, a_t)$. Next period's utility is subjected to discounting at a rate $\beta \in (0,1)$, which is assumed to be known. Formally the agent is represented by a triple of primitives $(u, p, \beta)$, who is assumed to behave according to an optimal decision rule, $\mathcal{A}_t = \{\alpha_\tau(s_\tau)\}_{\tau=t}^{\infty}$, in solving the following *sequential problem* for any time $\tau$

$$V_\theta^0(s_t) = \max_{\{\alpha(s_\tau)\}_{\tau=t}^{\infty}} E\left[\left.\sum_{\tau=t}^{\infty} \beta^{\tau-t} u_\theta(\alpha(s_\tau), s_\tau)\right| s_t\right], \text{s.t. } \alpha(s_t) \in A \text{ for all } t.$$

Under some regularity conditions, there exists a stationary Markovian optimal decision rule $\alpha_\theta(\cdot)$ so that

$$\alpha_\theta^0(s_t) = \arg\max_{a\in A}\left\{u_\theta(a, s_t) + \beta E\left[V_\theta^0(s_{t+1})\,|s_t, a_t = a\right]\right\}. \tag{80}$$

Furthermore, the value function, $V_\theta^0$, is the unique solution to the *Bellman's equation*

$$V_\theta^0(s_t) = \max_{a\in A}\left\{u_\theta(a, s_t) + \beta E\left[V_\theta^0(s_{t+1})\,|s_t, a_t = a\right]\right\}. \tag{81}$$

More details of related Markov decision models that are commonly used in economics can be found in Pakes (1994) and Rust (1994,1996). In order to avoid a degenerate model, we assume that the state variables $s_t = (x_t, \varepsilon_t)$ can be separated into two parts, which are observable and unobservable respectively to the econometrician, see Rust (1994) for various interpretations of the unobserved heterogeneity. We next provide an economic example that naturally fits in our dynamic decision making framework.

DYNAMIC PRICE SETTING EXAMPLE:

Consider a dynamic price setting problem for a firm. At the beginning of each period $t$, the firm faces a demand described by $D(a_t, x_t, \varepsilon_t)$ where: $a_t$ denotes the price that is assumed to belong to some subset of $\mathbb{R}$; $x_t$ is some measure of the consumer's satisfaction that affects the level of the demand for the immediate period that is publically observed; $\varepsilon_t$ is the firm's private demand shock. Within each period, the firm sets a price and earns the following immediate profit

$$u(a, x_t, \varepsilon_t) = D(a_t, x_t, \varepsilon_t)(a_t - c),$$

where $c$ denotes a constant marginal cost. The price setting decision in period $t$ affects the future sentiment of the demand of the consumers for the next period, $x_{t+1}$, that can be modelled by some Markov process. So the firm chooses price $a_t$ to maximize its discounted expected profit

$$a_t = \arg\max_{a \in A} \{u(a, x_t, \varepsilon_t) + \beta E[V(x_{t+1}, \varepsilon_{t+1}) | x_t, \varepsilon_t, a_t = a]\}$$

In Section 2.5, we focus on a specific example of this dynamic price setting decision problem and use a Monte Carlo experiment to illustrate the finite sample behavior of our estimator as well as the effects of ignoring the underlying dynamics in the DGP.

Unless stated otherwise, we assume the following set of assumptions, which are common in this literature, throughout the paper.

ASSUMPTION M2.1: *The observed data for each individual* $\{a_t, x_t\}_{t=1}^{T+1}$ *are the controlled stochastic processes satisfying (80) for some* $\theta_0 \in \Theta$ *with exogenously known* $\beta$.

ASSUMPTION M2.2: *(Conditional Independence) The transitional distribution has the following factorization:* $p(x_{t+1}, \varepsilon_{t+1} | x_t, \varepsilon_t, a_t) = q(\varepsilon_{t+1}) p_{X'|X,A}(x_{t+1} | x_t, a_t)$.

ASSUMPTION M2.3: *The support of* $s_t = (x_t, \varepsilon_t)$ *is* $X \times \mathcal{E}$, *where* $X = \{1, \ldots, J\}$ *for some* $J < \infty$ *that denotes the observable state space and* $\mathcal{E}$ *is a (potentially strict) subset of* $\mathbb{R}$. *The distribution of* $\varepsilon_t$, *denoted by* $Q$, *is known, it is also independent of* $x_t$ *and is absolutely continuous with respect to some Lebesgue measure with a positive Radon-Nikodym density* $q$ *on* $\mathcal{E}$.

ASSUMPTION M2.4: *(Monotone Choice) The per period payoff function* $u_\theta : A \times X \times \mathcal{E} \to \mathbb{R}$ *has increasing differences in* $(a, \varepsilon)$ *for all* $x$ *and* $\theta$.

The first two assumptions are familiar from the discrete control problems; M2.2 is introduced by Rust (1987). Finiteness of $X$ is imposed for the sake of simplicity, the generalization to more general compact set is discussed in Section 2.6. Notice that, unlike under the discrete choice setting, the estimation problem still requires an estimation of some infinite dimensional elements despite assuming that $X$ has finite elements since $A$ now includes an interval. The distribution of $\varepsilon_t$ is required to be known, this is a standard assumption in the estimation of structural dynamic programming models whether the control is continuous or discrete. The independence between $x_t$ and $\varepsilon_t$ can be weakened to the knowledge of the conditional distribution of $\varepsilon_t$ given $x_t$ upto some finite dimensional unknown parameters. However, unlike dynamic discrete choice models, the support of $\varepsilon_t$ need not be unbounded, since the unboundedness of $\mathcal{E}$ is required to utilize HM inversion theorem. In fact, as we shall see below, in many cases it is more natural to assume that $\mathcal{E}$ is a compact and convex subset of $\mathbb{R}$ when $A$ is also compact and convex. More important is the monotone choice assumption in M2.4, similar to Bajari et al. (2007) and Hong and Shum (2009), the monotonocity assumption is crucial in our methodology since $\varepsilon$ typically enters $u_\theta$ non-additively. However, this condition can be empirically motivated, in particular, the implication of M2.2 together with M2.4 is that policy function is increasing on $\mathcal{E}$. To see this, from (80) we have

$$\begin{aligned}
\alpha_\theta^0(s_t) &= \arg\max_{a\in A}\left\{u_\theta(a,s_t)+\beta E\left[V_\theta^0(s_{t+1})\,|s_t,a_t=a\right]\right\} \\
&= \arg\max_{a\in A}\left\{u_\theta(a,s_t)+\beta E\left[V_\theta^0(s_{t+1})\,|x_t,a_t=a\right]\right\},
\end{aligned}$$

since the function to be maximized on the RHS is supermodular in $(a,\varepsilon)$ for all $x$ and $\theta$, the claim follows from Topkis' theorem (see Topkis (1998)).

## 2.3 Estimation Methodology

Given a balanced panel data $\{a_{it},x_{it}\}$ of $N$ i.i.d. agents, our estimation strategy proceeds in two stages. First, we construct the model implied conditional distribution functions (CDFs). We then minimize the distance between CDFs obtained in the first stage with that from the data. The innovation that distinguishes most two-step estimators in this literature arises in the non-optimization first stage. Our first stage consists of three steps: (1) we estimate the model implied continuation value functions; (2) these functions are used to approximate the policy functions; (3) we simulate the CDFs from the policy functions. In more details:

STEP 1: VALUE FUNCTION

Based on the observed $\{a_{it}\}$, which corresponds to $\left\{\alpha^0_{\theta_0}(s_{it})\right\}$, for any $\theta \in \Theta$ we define a model implied value function, denoted by $W_\theta$, as a stationary solution to the following linear equation (cf. (81))

$$V_\theta(s_{it}) = u_\theta(a_{it}, s_{it}) + \beta E\left[V_\theta(s_{it+1})\,|s_{it}\right], \tag{82}$$

where $V_\theta$ can be written as

$$V_\theta(s_{it}) = E\left[\sum_{\tau=t}^{\infty} \beta^{\tau-t} u_\theta(a_{i\tau}, s_{i\tau}) \middle| s_{it}\right].$$

Therefore we must have $V_{\theta_0}(s_{it}) = V^0_{\theta_0}(s_{it})$. We can interpret $V_\theta$ as the value function for an economic agent whose underlying preference is $\theta$ but is using the policy function that is optimal with respect to $\theta_0$. Marginalizing out the unobserved states in (82), under M2.2, we have the following characterization of the model implied *conditional value functions*

$$E\left[V_\theta(s_{it})\,|x_{it}\right] = E\left[u_\theta(a_{it}, s_{it})\,|x_{it}\right] + \beta E\left[E\left[V_\theta(s_{it+1})\,|x_{it+1}\right]|x_{it}\right].$$

Since $|X| = J$, the equation above can be conveniently summarized by a matrix equation as

$$m_\theta = r_\theta + \mathcal{L}m_\theta. \tag{83}$$

where for each $j,k = 1,\ldots,J$: $r_\theta(j)$ denotes $E[u_\theta(a_{it},s_{it})|x_{it}=j]$; $\mathcal{L}$ is a $J\times J$ stochastic matrix whose $(k,j)-$th entry represents $\beta\Pr[x_{it+1}=j|x_{it}=k]$; $m_\theta(j)$ denotes $E[V_\theta(s_{it})|x_{it}=j]$. Note that $(I-\mathcal{L})$ is invertible by the dominant diagonal theorem, so the solution to (83) exists that uniquely defines the conditional value function.[11] To obtain the model implied *continuation value function*, again by M2.2, this function can be written as

$$E[V_\theta(s_{it+1})|x_{it},a_{it}] = E[E[V_\theta(s_{it+1})|x_{it+1}]|x_{it},a_{it}].$$

In a linear functional notation, the continuation value function can be defined by the following linear transformation

$$g_\theta = \mathcal{H}m_\theta. \tag{84}$$

Here $\mathcal{H}$ is a conditional expectation operator that maps $\mathbb{R}^J$ to a space of vector valued functions defined on $A$, so $g_\theta = (g_{0,j}(\cdot,\theta))_{j=1}^J$. In particular, for any $m\in\mathbb{R}^J$, we have $\mathcal{H}m(j,a) = \sum_{k=1}^J m_k \Pr[x_{it+1}=k|x_{it}=j,a_{it}=a]$ for $1\le j\le J$ and $a\in A$.

Given $\{a_{it},x_{it}\}$, to estimate $g_\theta$, we first estimate and solve (83) and transform the solution using the empirical counterpart of (84). First note that, we do not observe $\{\varepsilon_{it}\}$, which is not separable from $u_\theta$, using the monotonicity assumption we generate their nonparametric estimates from the following relation

$$\widehat{\varepsilon}_{it} = Q_\varepsilon^{-1}\left(\widehat{F}_{A|X}(a_{it}|x_{it})\right), \tag{85}$$

[11] A square matrix $P=(p_{ij})$ of size $n$ is said to be (strictly) diagonally dominant if $|p_{ii}| > \sum_{j\neq i}|p_{ij}|$ for all $i$. It is a standard result in linear algebra that a diagonally dominant matrix is non-singular, for example see Taussky (1977).

where $Q_\varepsilon$ is the known distribution function of $\varepsilon_{it}$; and $\widehat{F}_{A|X}(a|j) = \frac{1}{NT}\sum_{i=1,t=1}^{N,T} w_{itN}(j)\,\mathbf{1}[a_{it} \leq a]$ is an estimator for the true CDF of $a_{it}$ given $x_{it}$, where $w_{itN}(j) = \frac{\mathbf{1}[x_{it}=j]}{\widehat{p}_X(j)}$ with $\widehat{p}_X(j) = \frac{1}{NT}\sum_{i=1,t=1}^{N,T}\mathbf{1}[x_{it}=j]$.[12] Given $\{a_{it}, x_{it}, \widehat{\varepsilon}_{it}\}_{i=1,t=1}^{N,T+1}$, we can estimate $r_\theta$ by

$$\widetilde{r}_\theta(j) = \frac{1}{NT}\sum_{i=1,t=1}^{N,T} w_{itN}(j)\, u_\theta(a_{it}, x_{it}, \widehat{\varepsilon}_{it}). \tag{86}$$

Assuming further that $\widehat{p}_X(j) > 0$ for $j = 1, \ldots, J$, we simply use the frequency estimators to estimate each elements in $\mathcal{L}$. The dominant diagonal theorem implies $\left(I - \widehat{\mathcal{L}}\right)^{-1}$ exists, we can then uniquely estimate the conditional value function by the relation

$$\widehat{m}_\theta = \left(I - \widehat{\mathcal{L}}\right)^{-1}\widetilde{r}_\theta. \tag{87}$$

As seen above, there are also many nonparametric estimators available for the regression function, for simplicity we use the Nadaraya Watson estimator to approximate the operator $\mathcal{H}$, therefore

$$\widehat{g}_\theta = \widehat{\mathcal{H}}\widehat{m}_\theta, \tag{88}$$

such that, for any $1 \leq j, k \leq J$ and $a \in int(A)$

$$\begin{aligned}
\widehat{g}_j(a,\theta) &= \sum_{k=1}^{J}\widehat{m}_\theta(k)\frac{\widehat{p}_{X',X,A}(k,j,a)}{\widehat{p}_{X,A}(j,a)}, \qquad (89)\\
\widehat{p}_{X',X,A}(k,j,a) &= \frac{1}{NT}\sum_{i=1,t=1}^{N,T}\mathbf{1}[x_{it+1}=k, x_{it}=j]\,K_h(a_{it}-a),\\
\widehat{p}_{X,A}(j,a) &= \frac{1}{NT}\sum_{i=1,t=1}^{N,T}\mathbf{1}[x_{it}=j]\,K_h(a_{it}-a),
\end{aligned}$$

where $\widehat{p}_{X',X,A}$ denotes our choice of estimate for $p_{X',X,A}$, the mixed-continuous joint density of $(x_{it+1}, x_{it}, a_{it})$; $\widehat{p}_{X,A}$ and $p_{X,A}$ are defined similarly; $K_h(\cdot) = \frac{1}{h}K\left(\frac{\cdot}{h}\right)$ denotes

[12] BBL also uses the one to one correspondence between $a_{it}$ and $\varepsilon_{it}$ in their forward simulation method, where they draw $\{\varepsilon_b\}$ and generate the corresponding optimal choice from $\left\{\widehat{F}_{A|X}^{-1}(Q_\varepsilon(\varepsilon_b)|x_b)\right\}$, for any state $x_b$.

a user-chosen kernel and $h$ is the bandwidth that depends on the sample size but for the ease of notation we suppress this dependence. Regardless of the nature of the support of $A$, we may want to trim off the estimates near the boundaries or the tails of the distribution, this discussion is deferred until Section 3.

STEP 2: POLICY FUNCTION

For any $\theta \in \Theta$, $j = 1, \ldots, J$ and $g_j \in \mathcal{G}$, where $\mathcal{G}$ is a space of functions mapping $A$ to $\mathbb{R}$, we can define a generic objective function $\pi_j(\cdot, \cdot, \theta, g_j)$ where

$$\pi_j(a, \varepsilon, \theta, g_j) = u_\theta(a, j, \varepsilon) + \beta g_j(a),$$

so that $\theta$ indexes the current period payoff function and $g_j$ (that may also depend on $\theta$) summarizes the future expected payoff. Let $\partial_a$ denote the partial derivative $\frac{\partial}{\partial a}$, for each $j$, we can approximate the model implied *policy function*, denoted by $\alpha_j(\cdot, \theta, \partial_a g_j)$, to be the function that maximizes $\pi_j(a, \cdot, \theta, g_j)$ over $A$ by substituting $\widehat{g}_j(\cdot, \theta)$, from Step 1, in place of $g_j$.[13] Since $\dim(A) = 1$, the approximation can be done by direct grid-search or from finding the zero to the first derivative of $\pi_j$,

$$\partial_a \pi_j(a, \varepsilon, \theta, \partial_a g_j) = \partial_a u_\theta(a, j, \varepsilon) + \beta \partial_a g_j(a).$$

The model implied policy function is deliberately written to depend on the derivative of $g_j$. It will be convenient, at least for theoretical analysis, to assume that the optimal rule is characterized by the first order condition from differentiating $\pi_j$ w.r.t. $a$. This also has an important practical implication, in particular with regards to the bandwidth

[13]First note that, when $x_{it} = j$, $a_{it}$ must be equal to $\alpha_j(\varepsilon_{it}, \theta_0, g_{0,j}(\cdot, \theta_0))$, where $g_{0,j}(\cdot, \theta)$ denotes the true continuation value function as defined in (84). Also note that, under M2.1, M2.2 and M2.4, we can write (81) as

$$V_{\theta_0}^0(j, \varepsilon_{it}) = \max_{a \in A} \pi_j(a, \varepsilon_{it}, \theta_0, g_{0,j}(\cdot, \theta_0)) \quad \text{for } x = j = 1, \ldots, J.$$

choice, as it follows from the implicit function theorem in Banach space that the effective rate of convergence of the policy function will be determined by the rate of convergence of $\partial_a \widehat{g}_j(\cdot,\theta)$.[14] Note that the last statement is true irrespective of how we approximate $\{\alpha_j(\cdot,\theta,\partial_a g_j(\cdot,\theta))\}$.

STEP 3: DISTRIBUTION FUNCTION

Given a policy function, for any $\theta\in\Theta$, $j=1,\ldots,J$ and $g_j\in\mathcal{G}$, we can define the corresponding model implied *conditional distribution function* as follows

$$\begin{aligned}
F_{A|X}(a|j;\theta,\partial_a g_j) &= \Pr\left[\alpha_j(\varepsilon_{it},\theta,\partial_a g_j)\leq a\right] \\
&= E\left[\mathbf{1}\left[\alpha_j(\varepsilon_{it},\theta,\partial_a g_j)\leq a\right]\right].
\end{aligned}$$

In this notation, the true CDF that we write as $F_{A|X}(a|j)$ must be equaled to $F_{A|X}(a|j;\theta_0,\partial_a g_{0,j}(\cdot,$ The integral above can be approximated to any arbitrary degree of accuracy by Monte Carlo integration, since we assume the knowledge of $Q_\varepsilon$. In this paper, for simply we use

$$\widetilde{F}_{A|X}(a|j;\theta,\partial_a\widehat{g}_j)=\frac{1}{R}\sum_{r=1}^{R}\mathbf{1}\left[\alpha_j(\varepsilon_r,\theta,\partial_a\widehat{g}_j)\leq a\right], \tag{90}$$

where $\{\varepsilon_r\}_{r=1}^{R}$ is a random sample from $Q$. Note that the indicator function introduces some discontinuities in $\widetilde{F}_{A|X}$ (particularly with respect to $\theta$).

SECOND STAGE OPTIMIZATION

Similar to Hotz and Miller (1993), our estimator is derived from the following conditional moment restrictions

$$E\left[\mathbf{1}\left[a_{it}\leq a\right]-F_{A|X}(a|j;\theta,\partial_a g_{0,j}(\cdot,\theta))\middle|\,x_{it}=j\right]=0,\ \text{for } a\in A \text{ and } j=1,\ldots,J \text{ when } \theta=\theta_0. \tag{91}$$

[14]The implicit function theorem in Banach space is a well established result. The sufficient conditions for its validity generalizes the standard conditions used in Euclidean space, e.g. see Zeidler (1986).

The condition above represents a continuum of moment restrictions, cf. Carrasco and Florens (2000), however, no general theory for semiparametric moment estimation with a continuum of moments is available at present. Since we can equivalently write (91) as

$$F_{A|X}(a|j) = F_{A|X}(a|j;\theta_0, \partial_a g_{0,j}(\cdot,\theta_0)) \text{ for } a \in A \text{ and } j = 1,\ldots,J \text{ when } \theta = \theta_0,$$

we focus on a class of minimum distance estimators.[15] Wolfowitz (1953) introduce the minimum distance method that since has developed into a general estimation technique that has well known robustness and efficiency properties, see Koul (2002) for a review. In this paper, we define a class of estimators that minimize the following Cramér von-Mises type objective function that defines some $L^2-$distance between the CDF implied the model and that of the data

$$M_N(\theta, \widehat{g}(\cdot,\theta)) = \sum_{j=1}^{J} \int_A \left[\widetilde{F}_{A|X}(a|j;\theta, \partial_a \widehat{g}_j(\cdot,\theta)) - \widehat{F}_{A|X}(a|j)\right]^2 \mu_j(da). \tag{92}$$

Here $\{\mu_j\}$ is a sequence of user chosen sigma-finite (for now, assume non-random) measures on $A$. Clearly the property of $\widehat{\theta}$ will depend on the choice of $\{\mu_j\}$. In Section 3, we provide a discussion on how to select the measures to ensure consistent estimation under some regularity conditions, we leave the issue of choosing $\{\mu_j\}$ efficiently for future work.

[15] Another alternative to the moment based estimator is to maximize the conditional maximum likelihood function, however the maximum likelihood estimator (MLE) is much more computationally demanding. Although one can proceed with our minimum distance approach and perform Newton Ralphson type iterations to ensure we get the same first order asymptotic distribution as the conditional MLE.

## 2.4 Practical Aspects

First note that all elements we require to solve and transform the linear equations in (83) and (84) have explicit functional forms, so they are easy to program. In addition, similar to Hotz et al. (1994) and BBL, we can also take advantage of the linear structure the policy value equation. In particular, if the parameterization of $\theta$ in $u$ is linear, $\widetilde{r}_{\theta}(j)$ can be writen as $\frac{1}{NT}\sum_{i=1,t=1}^{N,T} w_{itN}(j)\, u(a_{it}, x_{it}, \widehat{\varepsilon}_{it})' \theta$ for each $j$, and we can write the vector $\widetilde{r}_{\theta} = W\widehat{u}_{\theta} = W_{\widehat{u}}\theta$ for some matrix $W_{\widehat{u}}$. To estimate the conditional value function $\widehat{m}_{\theta} = \left(I - \widehat{\mathcal{L}}\right)^{-1} W_{\widehat{u}}\theta$, since we estimate $\mathcal{L}$ nonparametrically, we only have to compute the matrix $\left(I - \widehat{\mathcal{L}}\right)^{-1} W_{\widehat{u}}$ once as it does not depend on $\theta$.

It is also straightforward to carry out our methodology in a parametric framework. One can choose to parameterize the transition law $p_{X'|X,A}(\theta_{tr})$, for some $\theta_{tr}$ that may have common elements with $\theta$. The continuation value function still satisfies (84) where the conditional expectation operator becomes $\mathcal{H}_{\theta_{tr}}$. For fix $\theta_{tr}$, we can estimate $\widehat{g}_{\theta}$ using the relation (88) by simply replacing $\frac{\widehat{p}_{X',X,A}}{\widehat{p}_{X,A}}$ in (89) by $p_{X'|X,A}(\theta_{tr})$. Although the conditional expectation operator $\mathcal{H}_{\theta_{tr}}$ depends on $\theta_{tr}$, it does not affect how we estimate $\widehat{m}_{\theta}$. Note also that all the subsequent stages of the methodology only assume we have $\widehat{g}_{\theta}$ and not how they are obtained, therefore the remaining steps in our procedure remains unchanged.

## 2.5 Asymptotic Theory

Our minimum distance estimator falls in the class of a profiled semiparametric M-estimator with non-smooth objective function since (90) is discontinuous in $\theta$. There are a few recent econometrics papers that treat general theories of semiparametric estimation that allows for non-smooth criterions; Chen, Linton and Van Keilegom (2003) provide some general theorems for a class of Z-estimators; Ichimura and Lee (2006) obtain the characterization of the asymptotic distribution of M-estimators; Chen and

Pouzo (2008) extend the results of Ai and Chen (2003), on conditional moments models, to the case with non-smooth residuals. The aforementioned papers put special emphasis on the criterion that is based on sample averages. However, minimum distance criterions generally do not fall into this category, for instance consider (92) when $\{\mu_j\}_{j=1}^{J}$ is a sequence of non-random measures. Although the focus of our chapter is not on the general theory of estimation, we find it convenient to proceed by providing a general asymptotic normality theorem for semiparametric M-estimators that naturally include minimum distance estimators as well as many others commonly used objective functions. We then provide a set of sufficient, more primitive, conditions specific to our problem. We note, as an alternative, the discontinuity in many criterion functions can be overcome by smoothing, e.g. see Horowitz (1998), and in some cases there may be statistical gains for doing so, e.g. a reduction in finite sample MSE. More specifically, we can overcome the discontinuity problem by smoothing over the indicators in (90), however, the use of unsmoothed empirical function is the most common approach we see in practice.

To analyze our estimator, it is necessary to introduce the notion of functional derivative in order to capture the effects from the nonparametric estimate. We denote the (partial-) Fréchet differential operators by $D_\theta, D_g, D_{\theta\theta}, D_{\theta g}$ and $D_{gg}$, where the indices denote the argument(s) used in differentiating and double indexing denotes second derivative. For any map $T : X \to Y$ and some Banach spaces $X$ and $Y$, we say that $T$ is Fréchet differentiable at $x$, that belongs to some open neighborhood of $X$, if and only if there exists a linear bounded map $D_T : X \to Y$ such that $T(x+f) - T(x) = D_T(x) f + o(\|f\|)$ with $\|f\| \to 0$ for all $f$ in some neighborhood of $x$; we denote the Fréchet differential at $x$ in a particular direction $f$ by $D_T(x)[f]$. Since $\theta$ is a finite dimensional Euclidean element, the first and second Fréchet derivatives coincide with the usual (partial-) derivatives. For Theorem G below, let $\theta_0$ and

$g_0$ denote the true finite and infinite dimensional parameters that lie in $\Theta$ and $\mathcal{G}$ respectively. Since we only need to focus on the local behavior around $(\theta_0, g_0)$, for any $\delta > 0$ we define $\Theta_\delta = \{\theta \in \Theta : \|\theta - \theta_0\| < \delta\}$ and $\mathcal{G}_\delta = \left\{g \in \mathcal{G} : \|g - g_0\|_{\mathcal{G}} < \delta\right\}$, here $\delta$ can also be replaced by some positive sequence $\delta_N = o(1)$. The pseudo-norm on $\mathcal{G}_\delta$ can be suitably modified to reflect the smaller parameter space $\Theta_\delta$, and the choice of $\delta$ for $\Theta_\delta$ and $\mathcal{G}_\delta$ can be distinct, but for notational simplicity we ignore this. Let $M(\theta, g(\cdot, \theta))$ denote the population objective function that is minimized at $\theta = \theta_0$, and $M_N(\theta, g(\cdot, \theta))$ denote the sample counterpart. Further, we denote $D_\theta M(\theta, g(\cdot, \theta))$ by $S(\theta, g(\cdot, \theta))$ and $D_{\theta\theta} M(\theta, g(\cdot, \theta))$ by $H(\theta, g(\cdot, \theta))$.

THEOREM G: *Suppose that $\widehat{\theta} \xrightarrow{p} \theta_0$, and for some positive sequence $\delta_N = o(1)$,*

*G1* $M_N\left(\widehat{\theta}, \widehat{g}\left(\cdot, \widehat{\theta}\right)\right) \leq \inf_{\theta \in \Theta} M_N(\theta, \widehat{g}(\cdot, \theta)) + o_p\left(N^{-1}\right)$

*G2 For all $\theta, \widehat{g}(\cdot, \theta) \in \mathcal{G}_{\delta_N}$ w.p.a. 1 and $\sup_{\theta \in \Theta} \|\widehat{g}(\cdot, \theta) - g_0(\cdot, \theta)\|_\infty = o_p\left(N^{-1/4}\right)$*

*G3 For some $\delta > 0$, $M(\theta, g)$ is twice continuously differentiable in $\theta$ at $\theta_0$ for all $g \in \mathcal{G}_\delta$. $H(\theta, g)$ is continuous in $g$ at $g_0$ for $\theta \in \Theta_\delta$. Further, $S(\theta_0, g_0(\cdot, \theta_0)) = 0$ and $H_0 = H(\theta_0, g_0(\cdot, \theta_0))$ is positive definite.*

*G4 For some $\delta > 0$, $S(\theta, g(\cdot, \theta))$ is (partial-) Fréchet differentiable with respect to $g$, for any $\theta \in \Theta_\delta$ and for all $g \in \mathcal{G}_\delta$. Further $\|S(\theta_0, g(\cdot, \theta_0)) - D_g S(\theta_0, g_0(\cdot, \theta_0)) [g(\cdot, \theta_0) - g_0(\cdot, \theta_0)]$ $B_N \times \sup_{\theta \in \Theta} \|g(\cdot, \theta) - g_0(\cdot, \theta)\|_\infty^2$ for some $B_N = O_p(1)$.*

*G5 (Stochastic Differentiability)*

$$\sup_{\|\theta - \theta_0\| < \delta_N} \left| \frac{\mathcal{D}_N(\theta, \widehat{g}(\cdot, \theta))}{1 + \sqrt{N}\|\theta - \theta_0\|} \right| = o_p(1),$$

*where there exist some sequence $C_N$, so that*

$$\begin{aligned} & \mathcal{D}_N(\theta, \widehat{g}(\cdot, \theta)) \qquad (93) \\ = \ & \sqrt{N} \frac{\left[M_N(\theta, \widehat{g}(\cdot, \theta)) - M_N(\theta_0, \widehat{g}(\cdot, \theta_0)) - (M(\theta, \widehat{g}(\cdot, \theta)) - M(\theta_0, \widehat{g}(\cdot, \theta_0))) - (\theta - \theta_0)' C_N\right]}{\|\theta - \theta_0\|}. \end{aligned}$$

*G6 For some finite positive definite matrices $\Omega_0$ and $\Omega$, we have the following weak convergence $\sqrt{N}C_N \Rightarrow N(0,\Omega_0)$ and $\sqrt{N}D_N = \sqrt{N}\left(C_N + D_g S(\theta_0, g_0(\cdot,\theta_0))[\widehat{g} - g_0]\right) \Rightarrow \mathcal{N}(0,\Omega)$.*

*Then*

$$\sqrt{N}\left(\widehat{\theta} - \theta_0\right) \Rightarrow N\left(0, H_0^{-1}\Omega H_0^{-1}\right).$$

COMMENTS ON THEOREM G:

Under the identification assumption and sufficient conditions for asymptotic normality, one can often show the consistency of the finite dimensional parameter in such models directly so we do not provide a separate theorem for it. Theorem G extends Theorem 7.1 in Newey and McFadden (1994) to a two-step semiparametric framework. G1 is the definition of the estimator. The way G1 - G4 accommodate for the preliminary nonparametric estimator is standard, cf. Chen et al. (2003), in fact, a weaker notion of functional derivative such as the Gâteaux derivative will also suffice here. G5 extends the stochastic differentiability condition of Pollard (1985) and Newey and McFadden (1994) to this more general case. We note that this is not the only way to impose the stochastic differentiability condition; we pose our equicontinuity condition in anticipation of a sequential stochastic expansion whilst Ichimura and Lee (2006) employ an expansion on both Euclidean and functional parameters simultaneously. Also, the first order properties of $C_N$, the stochastic derivative in (93), will be the same as the case that $g_0(\cdot,\theta)$ is known.[16]

ASSUMPTION E1:

*(i) $\{a_{it}, x_{it}\}_{i=1,t=1}^{N,T+1}$ is i.i.d. across $i$, within each $i$ $\{a_{it}, x_{it}\}_{t=1}^{T+1}$ is a strictly stationary realizations of the controlled Markov process for a fixed periods of $T+1$ with*

[16] An important special case of this theorem is when the preliminary function is independent of $\theta$. The formulation of the conditions for Theorem G remains valid since the profiling effects are implicit in the notation of $D_\theta$ and $D_{\theta\theta}$.

*exogenous initial values;*

*(ii) $A$ and $\mathcal{E}$ are compact and convex subsets of $\mathbb{R}$;*

*(iii) $\Theta$ is a compact subset of $\mathbb{R}^L$ then the following holds for all $j = 1, \ldots, J$*

$$\alpha_j\left(\cdot, \theta, \partial_a g_{0,j}\left(\cdot, \theta\right)\right) = \alpha_j\left(\cdot, \theta_0, \partial_a g_{0,j}\left(\cdot, \theta_0\right)\right) \qquad Q - a.e.$$

*if and only if $\theta = \theta_0$ where $\theta_0 \in int\left(\Theta\right)$;*

*(iv) For all $j = 1, \ldots, J, \mu_j$ is a finite measure on $A$ that dominates $Q$ and has zero measure on the boundary of $A$;*

*(v) For all $j = 1, \ldots, J$, the density $p_{X,A}\left(j, \cdot\right)$ is $5-$times continuously differentiable on $A$ and $\inf_{a \in A} p_{X,A}\left(j, a\right) > 0$;*

*(vi) For all $j, k = 1, \ldots, J$, the density $p_{X',X,A}\left(k, j, \cdot\right)$ is $5-$times continuously differentiable on $A$;*

*(vii) The distribution function of $\varepsilon_{it}$, $Q_\varepsilon$, is Lipschitz continuous and twice continuously differentiable;*

*(viii) For all $j = 1, \ldots, J$, $u_\theta\left(a, j, \varepsilon\right)$ is twice continuously differentiable in $\theta$ and $a$, once continuously differentiable in $\varepsilon$, these continuous derivatives exist for all $a, \varepsilon$ and $\theta$. In addition we assume $\frac{\partial^2}{\partial a \partial \varepsilon} u_\theta\left(a, j, \varepsilon\right) > 0$ and $\frac{\partial^4}{\partial a^2 \partial \varepsilon^2} u_\theta\left(a, j, \varepsilon\right)$ exists and is continuous for all $a, \varepsilon$ and $\theta$;*

*(ix) $K$ is a 4-th order even and continuously differentiable kernel function with support $[-1, 1]$, we denote $\int u^j K\left(u\right) du$ and $\int K^j\left(u\right) du$ by $\mu_j\left(K\right)$ and $\kappa_j\left(K\right)$ respectively;*

*(x) The bandwidth sequence $h_N$ satisfies $h_N = d_N N^{-\varsigma}$ for $1/8 < \varsigma < 1/6$, with $d_N$ is a sequence of real numbers that is bounded away from zero and infinity;*

*(xi) Trimming factor $\gamma_N = o\left(1\right)$ and $h_N = o\left(\gamma_N\right)$;*

*(xii) The simulation size $R$ satisfies $N/R = o\left(1\right)$;*

COMMENTS ON E1:

(i) assumes we have a large $N$ and small $T$ framework, common in microeconometric applications, and for simplicity we assume $T$ is the same for all $i$;

(ii) The dimension of $A$ determines the rate of convergence of the nonparametric estimate, if $\dim(A) > 1$ we can adjust our conditions in a straightforward way to ensure the root$-N$ consistency of finite dimensional parameters, e.g. see Robinson (1988) and Andrews (1995). Compactness of $A$ and $\mathcal{E}$ is also assumed for the sake of simplicity. We can use a well known trimming argument in nonparametric kernel literature if $A$ and $\mathcal{E}$ are both unbounded , see Robinson (1988); all of our theoretical results and techniques in this chapter hold on any compact subset of $A$ and $\mathcal{E}$, the compact support can then be made to increase without bounds at some appropriate rate;

(iii) is the main identification condition for $\theta_0$. We assume there does not exist any other $\theta \in \Theta \backslash \{\theta_0\}$ that can generate the same policy profile which $\theta_0$ generates when $(\partial_a g_{0,j}(\cdot,\theta))$ is known. It can be shown directly that the conditions we impose on the policy functions is equivalent to imposing that (91) holds if and only if $\theta = \theta_0$, which is the standard identification assumption in a parametric conditional moment model; in the case that $x_{it}$ and $\varepsilon_{it}$ are not independent we simply change $Q-a.e.$ to $Q_{\varepsilon|X_j}-a.e.$, where $Q_{\varepsilon|X_j}$ denotes the conditional distribution of $\varepsilon_{it}$ given $x_{it} = j$. Lastly, given $\theta$, under some primitive conditions on the DGP (contained in E1) $(\partial_a g_{0,j}(\cdot,\theta))$ will be nonparametrically identified hence we only have to consider the identification of $\theta_0$;

(iv) ensures that the identification condition of (iii) is not lost through the user chosen measures, cf. Domínguez and Lobato (2004). One simple choice of $\{\mu_j\}_{j=1}^J$ that satisfies this condition is a sequence of measures which are dominated by the Lebesgue measure on the interior of $A$ and has zero measure on the boundary. We can also allow the support of $a_{it}$ to depend on the conditioning state variable $x_{it}$ but common support is assumed for notational simplicity;

(v)-(vi) impose standard smoothness and boundedness restrictions on the underlying distribution of the observed random variables in the kernel estimation literature. They ensure we can carry out the usual expansion on our nonparametric estimators of $p_{X,A}$ and $p_{X',X,A}$ and their derivatives in anticipation of using a 4-th order kernel;

(vii) imposes standard smoothness on $Q_\varepsilon$ that is necessary for our statistical analysis;

(viii) imposes standard smoothness assumptions on the per period utility function, to be used in conjunction with earlier conditions, to obtain uniform rates of convergence for our nonparametric estimates. The cross partial derivative is the analytical equivalence of M2.4. We note that these conditions appear particularly straightforward, this is due to the fact that $u$ is a continuous function on a compact domain, so boundedness makes it simple to obtain uniform convergence results. On the other hand, had we allowed for unbounded $A$ and $\mathcal{E}$, then we will need some conditions to ensure the tail probability of $u_\theta(a_{it}, x_{it}, \varepsilon_{it})$ is sufficiently small. For example, one sufficient condition would be that all the functions mentioned belong in $L^2(P)$, and there exists a function $|u_\theta(a, x, \varepsilon)| \leq U(a, x, \varepsilon)$ for all $a, x, \varepsilon$ and $\theta$ such that $E[\exp\{CU(a_{it}, x_{it}, \varepsilon_{it})\}] < \infty$ for some $C > 0$. The latter is equivalent to the Cramér's condition, see Arak and Zaizsev (1988), that allows us to use Bernstein type inequalities for obtaining the uniform rate of convergence of the nonparametric estimates;

(ix) The use of a 4-th order kernel is necessary to ensure the asymptotic bias will disappear for certain range of bandwidths. The compact support assumption on the kernel is made to keep our proofs simple, other 4-th order kernel with unbounded support can also be used, e.g. if it satisfies the tail conditions of Robinson (1988);

(x) imposes the necessary condition on the rate of decay of the bandwidth corresponding to using a 4-th order kernel. The specified rate ensures the uniform convergence of the first two derivatives of a regular 1-dimension nonparametric density estimate, as well as, the uniform convergence of $\|\partial_a \widehat{g} - \partial_a g\|_{\mathcal{G}}$ at a rate faster than $N^{-1/4}$ and for the asymptotic bias (of order $\sqrt{N}h^4$) to converge to zero;

(xi) This is the rate that the trimming factor diminishes, it suffices to only trim out the region in a neighborhood the boundary where the order of the bias differs from other interior points;

(xii) The simulation size must increase at a faster rate than $N$ to ensure the simulation error from using (90) does not affect our first order asymptotic theory.

In relation to Theorem G, beyond the identification conditions (iii) - (iv), most of the conditions in Assumption E1 will ensure that G2 holds. We now must impose some additional smoothness conditions on $(\alpha_j)$ to satisfy the other conditions of Theorem G. In particular, in order to apply the results from empirical processes literature, we need to restrict the size of the class of functions that the continuation value functions belong to. For a general subset of some metric space $(\mathcal{G}, \|\cdot\|_{\mathcal{G}})$, two measures of the size, or level of complexity, of $\mathcal{G}$ that are commonly used in the empirical processes literature are the *covering number* $N(\varepsilon, \mathcal{G}, \|\cdot\|_{\mathcal{G}})$ and the *covering number with bracketing* $N_{[]}(\varepsilon, \mathcal{G}, \|\cdot\|_{\mathcal{G}})$ respectively, see van der Vaart and Wellner (1996) for their definitions. We need the covering numbers of $(\mathcal{G}, \|\cdot\|_{\mathcal{G}})$ to not increase too rapidly as $\varepsilon \to 0$ (to be made precise below) and this possible, for example, if the functions in $\mathcal{G}$ satisfy some smoothness conditions. We now define a class of real valued functions that is popular in nonparametric estimation, suppose $A \subset \mathbb{R}^{L_A}$, let $\underline{\eta}$ be the largest integer smaller than $\eta$, and

$$\|g\|_{\infty,\eta} = \max_{|\eta| \leq \underline{\eta}} \sup_{a} \left|\partial_a^{|\eta|} g(a)\right| + \max_{|\eta| = \underline{\eta}} \sup_{a \neq a'} \frac{\left|\partial_a^{|\eta|} g(a) - \partial_a^{|\eta|} g(a')\right|}{\|a - a'\|^{\eta - \underline{\eta}}}, \tag{94}$$

where $\partial_a^{|\eta|} = \partial^{|\eta|} / \partial a_1^{\eta_1} \ldots \partial a_{L_A}^{\eta_{L_A}}$ and $|\eta| = \sum_{l=1}^{L_A} \eta_l$, then $C_M^{\eta}(A)$ denotes the set of all continuous functions $g : A \to \mathbb{R}$ with $\|g\|_{\infty,\eta} \leq M < \infty$; let $l^{\infty}(A)$ denotes the class of bounded functions on $A$. If $\mathcal{G} = C_M^{\eta}(A)$, then by Corollary 2.7.3 of van der Vaart and Wellner $\log\left(N\left(\varepsilon, \mathcal{G}, \|\cdot\|_{\mathcal{G}}\right)\right) \leq const. \times \varepsilon^{-L_A/2}$. For our purposes, the precise condition for controlling the complexity of the class of functions is summarized by the following uniform entropy condition $\int_0^{\infty} \sqrt{\log N\left(\varepsilon, \mathcal{G}, \|\cdot\|_{\mathcal{G}}\right)} d\varepsilon < \infty$. So $\mathcal{G}$ satisfies the uniform entropy condition if $\eta > L_A/2$. Given the assumptions in E1 we can now be completely explicit regarding our space of functions and its norm. It is now clear that $\mathcal{G}_{0,j} \subset C_M^2(A) \subset l^{\infty}(A)$ for some $M > 0$ for each $j = 1, \ldots, J$ w.r.t. to the norm $\|\cdot\|_{\mathcal{G}}$ described in the introduction. Next, since we are required to define the notion of functional derivatives, it will be necessary to let our class of functions be an arbitrary open and convex set of functions that contains $\mathcal{G}_0$. So we define for all $j = 1, \ldots, J$, $\mathcal{G}_j = \{g(\cdot) \in C_M^2(A) : \sup_{a \in A} \|g(\cdot) - g_{0,j}(\cdot, \theta)\|_{\infty} < \delta$ for any $\theta \in \Theta\}$ for some $\delta > 0$, then it is also natural to also have $\mathcal{G}_j$ endowed with the norm $\|\cdot\|_{\mathcal{G}}$.[17] Finally, since we will be using results from empirical processes for a class of functions that are indexed by parameters in $A \times \Theta \times \mathcal{G}$, we define the norm for each element $(a, \theta, g)$ by $\|(a, \theta, g)\|_{\nu} = \|(a, \theta)\| + \|g\|_{\mathcal{G}}$.

ASSUMPTION E2:

*(xiii) For all $j = 1, \ldots, J$, the inverse of the policy function $\rho_j : A \times \Theta \times \mathcal{G}_j \to \mathbb{R}$ is twice Fréchet differentiable on $A \times \Theta \times \mathcal{G}_j$ and $\sup_{\theta, a, g_j \in \Theta \times A \times \mathcal{G}_j} \left\| D_g \rho_j(a, \theta, \partial_a g_j) \right\| < \infty$;*

*(xiv) For some $j = 1, \ldots, J$, the following $L \times L$ matrix*

$$\int_A [q(\rho_j(a, \theta_0, \partial_a g_j(\cdot, \theta_0)))]^2 D_\theta(\rho_j a, \theta_0, \partial_a g_{0,j}(\cdot, \theta_0)) D_\theta \rho_j(a, \theta_0, \partial_a g_{0,j}(\cdot, \theta_0))' \mu_j(da)$$

*is positive definite;*

*(xv) For all $j = 1, \ldots, J$, the Fréchet differential of $\rho_j$ w.r.t. $\partial_a g$ in the direction*

[17] Note that for any $g \in \mathcal{G}_j$ for any $j$, $\|g\|_{\mathcal{G}} \leq \delta + \max_{1 \leq j \leq J} \sup_{\theta, a \in A \times \Theta} |g_j(a, \theta)| < \infty$ holds by the triangle inequality.

$[\partial_a \widehat{g}_j(\cdot,\theta_0) - \partial_a g_{0,j}(\cdot,\theta_0)]$ *is asymptotically linear: in particular for any* $a \in int(A)$

$$D_g \rho_j(a,\theta_0,\partial_a g_j(\cdot,\theta_0))[\partial_a \widehat{g}_j(\cdot,\theta_0) - \partial_a g_{0,j}(\cdot,\theta_0)] = \frac{1}{NT}\sum_{i=1,t=1}^{N,T} \psi_{0,j}(a_{it},x_{it};a) + o_p\left(N^{-1/2}\right), \tag{95}$$

*with* $E\left[\psi_{0,j}(a_{it},x_{it};a)\right] = 0$ *and* $E\left[\psi_{0,j}^2(a_{it},x_{it};a)\right] < \infty$ *for all* $i,t$*; in addition, the display above holds uniformly on any compact subset* $A_N$ *of* $A$ *and* $\psi_{0,j}(a_{it},x_{it};\cdot) \in \Psi_{j,N}$ *where* $\Psi_{j,N}$ *is some class of functions on* $A_N$ *that is a Donsker class for all* $N$*.*

COMMENTS ON E2:

We first note that although it would appear more primitive to impose conditions on the policy function but the notation will be very cumbersome. Given the existence and smoothness of the inverse map we instead work with the inverse of the policy function, this is done without any loss of generality by using implicit, inverse and Taylor's theorems in Banach space.[18] Although these assumptions are hard to verify in practice, they are mostly mild conditions on the smoothness of $\rho$ that one would be quite comfortable in imposing if $\mathcal{G}$ belongs to a Euclidean space (at least for (xii) - (xiii)); in a similar spirit the same can be said regarding (xiv). For each $j$ and $a$, $D_g \rho_j(a,\theta_0,\partial_a g_j(\cdot,\theta_0))$ is a bounded linear functional and $[\partial_a \widehat{g}_j(\cdot,\theta_0) - \partial_a g_{0,j}(\cdot,\theta_0)]$ is a continuous and square integrable function in $L^2(A,\Pi)$,[19] by Riesz representation theorem there exists some $\varrho_j \in L^2(A,\Pi)$ such that $D_g \rho_j(a,\theta_0,\partial_a g_j(\cdot,\theta_0))[\partial_a \widehat{g}_j(\cdot,\theta_0) - \partial_a g_{0,j}(\cdot,\theta_0)] = \int \varrho_j(a';a)\,\partial_a \widehat{g}_j(a',\theta_0)$ $-\partial_a g_{0,j}(a',\theta_0)\,d\Pi(a')$. Given our assumptions, for a smooth $\varrho_j$, it is not difficult to show the validity of (95) since $\partial_a \widehat{g}_j(\cdot,\theta_0) - \partial_a g_{0,j}(\cdot,\theta_0)$ has an asymptotic linear form. This is not an uncommon approach when dealing with a general semiparametric estimator, see Newey (1994), Chen and Shen (1998), Ai and Chen (2003) and Chen et al. (2003), and in particular, Ichimura and Lee (2006) for the characterization of

[18]See Chapter 4 of Ziedler (1986) for these results.

[19]Here $L^2(A,\Pi)$ denotes a Banach space of measurable functions defined on $A$ that is square integrable w.r.t. some measure $\Pi$.

a valid linearization. However, since our $(\rho_j)$ does not have a closed form it is not clear how one can obtain $(\varrho_j)$. Once we obtain (95), standard CLT yields pointwise convergence in distribution (for each $a, x$) but this is still not enough for our minimum distance estimator since we will need a full weak convergence result, i.e. let $\psi_{N,j} = \frac{1}{\sqrt{NT}} \sum_{i=1,t=1}^{N,T} \psi_{0,j} (a_{it}, x_{it}; \cdot)$ be a random element in $l^{\infty}(A)$ we need $\psi_{N,j} \rightsquigarrow \psi_j$ as $N \to \infty$, where $\rightsquigarrow$ denotes weak convergence and $\psi_j$ is some tight Gaussian process that belongs to $l^{\infty}(A)$. The Donsker property can be satisfied for a large class of functions, see Van der Vaart and Wellner (1996). We note also that joint normality condition of G6 in Theorem G will also be easy to verify since we will end up working with sums of two Gaussian processes, each underlying asymptotic is driven by averages of zero mean functions of $\{(a_{it}, x_{it})\}_{i=1,t=1}^{N,T+1}$.

THEOREM 2.1: *Under E1: For any $a \in int(A), \theta \in \Theta$ and $j = 1, \ldots, J$, if $\widehat{g}_j(\cdot, \theta)$ satisfies (88) then*

$$\sqrt{Nh} \left(\widehat{g}_j(a,\theta) - g_{0,j}(a,\theta) - B_{N,j}(a; m_\theta)\right) \Rightarrow \mathcal{N}\left(0, \frac{\kappa_2(K)}{T p_{X,A}(j,a)} \text{var}\left(m_\theta(x_{it+1}) | x_{it} = j, a_{it} = a\right)\right),$$

*where*

$$B_{N,j}(a; m_\theta) = \frac{1}{4!} h^4 \mu_4(K) \sum_{k=1}^{J} m_\theta(k) \left( \frac{\frac{\partial^4}{\partial a^4} p_{X',X,A}(k,j,a)}{p_{X,A}(j,a)} + \frac{p_{X',X,A}(k,j,a) \frac{\partial^4}{\partial a^4} p_{X,A}(j,a)}{p_{X,A}^2(j,a)} \right),$$

*furthermore, $\widehat{g}_j(a,\theta)$ and $\widehat{g}_k(a',\theta)$ are asymptotically independent when $k \neq j$ or $a' \neq a$.*

We note that, for each $j$, the pointwise asymptotic property of $\widehat{g}_j(a,\theta)$ in Theorem 1 is identical to that of a Nadaraya-Watson estimator of $E[m_\theta(x_{it+1}) | x_{it} = j, a_{it} = a]$ when $m_\theta$ is known. In other words, the nonparametric estimation of $m_\theta$, as well as the generation of the nonparametric residuals (85), does not affect the first order asymptotic of $(\widehat{g}_j(\cdot,\theta))$. The reason behind this is due to the fact that $\left(\widetilde{r}_\theta, \widehat{m}_\theta, \widehat{\mathcal{L}}\right)$ converges uniformly (over $\Theta \times X$) in probability to $\left(\widetilde{r}_\theta, \widehat{m}_\theta, \widehat{\mathcal{L}}\right)$ at the rate close to $N^{-1/2}$, which is much faster than $1/Nh$.

In order to apply Theorem G, we now define the population and sample objective functions for our estimator. For any $\theta \in \Theta$ and $g(\cdot,\theta) \in \mathcal{G}$, we have defined $M_N(\theta, g(\cdot,\theta))$ earlier (see (92)), its population analogue is

$$M(\theta, g(\cdot,\theta)) = \sum_{j=1}^{J} \int_A \left[F_{A|X}(a|j;\theta,\partial_a g_j(\cdot,\theta)) - F_{A|X}(a|j)\right]^2 \mu_j(da).$$

THEOREM 2.2: *Under E1-E2: For $(\widehat{g}(\cdot,\theta_0))$ that satisfies* (88), *if $\widehat{\theta}$ satisfies G1 with $M_N(\theta, g(\cdot,\theta))$ as defined in (92) then $\widehat{\theta} \xrightarrow{p} \theta_0$.*

THEOREM 2.3: *Under E1-E2: For $(\widehat{g}(\cdot,\theta_0))$ that satisfies* (88), *if $\widehat{\theta}$ satisfies G1 with $M_N(\theta, g(\cdot,\theta))$ as defined in (92) then*

$$\sqrt{N}\left(\widehat{\theta} - \theta_0\right) \Rightarrow N\left(0, H_0^{-1}\Omega H_0^{-1}\right),$$

*where*

$$\begin{aligned}
\Omega &= \lim_{N\to\infty} \operatorname{var}\left(-2\sum_{j=1}^{J}\int \left[\sqrt{N}\left[\begin{array}{c} \left[D_\theta F_{A|X}(a|j;\theta_0,\partial_a g_{0,j}(\cdot,\theta_0))\right] \\ \times \\ \left(\widehat{F}_{A|X}(a|j) - F_{A|X}(a|j)\right) \\ -\left(D_g F_{A|X}(a|j;\theta_0,\partial_a g_{0,j}(\cdot,\theta_0))\left[\partial_a \widehat{g}_j(\cdot,\theta_0) - \partial_a g_{0,j}(\cdot,\theta_0)\right]\right. \end{array}\right.\right.\right. \\
H_0 &= 2\sum_{j=1}^{J}\int_A \left(D_\theta F_{A|X}(a|j;\theta_0,\partial_a g_{0,j}(\cdot,\theta_0))\right)\left(D_\theta F_{A|X}(a|j;\theta_0,\partial_a g_{0,j}(\cdot,\theta_0))\right)' \mu_j(da).
\end{aligned}$$

Next theorem provides the pointwise distribution theory of $\left(\widehat{g}_j\left(\cdot,\widehat{\theta}\right)\right)$ that can be used to estimate $(g_{0,j}(\cdot,\theta_0))$.

THEOREM 2.4: *Under E1-E2: For any $a \in int(A)$ and $j = 1, \ldots, J$ , if $\widehat{g}_j(\cdot, \theta)$ satisfies (88) and $\widehat{\theta}$ satisfies G1 then*

$$\sqrt{Nh}\left(\widehat{g}_j\left(a, \widehat{\theta}\right) - g_{0,j}(a, \theta_0) - B_{N,j}(a; m_{\theta_0})\right) \Rightarrow \mathcal{N}\left(0, \frac{\kappa_2(K)}{Tp_{X,A}(j,a)} \text{var}\left(m_{\theta_0}(x_{it+1}) | x_{it} = j, a_{it} = \right.\right.$$

*where $B_{N,j}(a; m_{\theta_0})$ has the same expression as in Theorem 1 when $\theta = \theta_0$. Furthermore, $\widehat{g}_j\left(a, \widehat{\theta}\right)$ and $\widehat{g}_k\left(a', \widehat{\theta}\right)$ are asymptotically independent when $k \neq j$ or $a' \neq a$.*

Theorem 4 implies that $\left(\widehat{g}_j\left(\cdot, \widehat{\theta}\right)\right)$ and $(\widehat{g}_j(\cdot, \theta_0))$ have the same first order asymptotic. This follows since $\widehat{g}$ (and $g$) is smooth in $\theta$, and $\widehat{\theta}$ converges to $\theta_0$ at a faster rate than $1/\sqrt{Nh}$. Note that, if we want to construct consistent confidence intervals for $g_{0,j}(a, \theta_0)$, we may use a different bandwidth in estimating $\widehat{g}$ to the one used in computing $\widehat{\theta}$.

## 2.6 Bootstrap Standard Errors

The asymptotic variance of the finite dimensional estimator in semiparametric models can have a complicated form that generally is a functional of the infinite dimensional parameters and their derivatives. Not only it is difficult to estimate such object, the estimate often works poorly in finite sample. In this section we propose to use semiparametric bootstrap to estimate the sampling distribution of the estimator described in this chapter.

The original bootstrap method was proposed by Efron (1979). The bootstrap is a general method that is very useful in statistics, for samples of its scope see the monographs by Hall (1992), Efron and Tibshirani (1993), as well as Horowitz (2001) for a survey that is specialized for an econometrics audience. In this chapter we concentrate on the use of bootstrap as a tool to estimate the standard error of $\widehat{\theta}$ defined in Theorem 3. Generally, bootstrap methods under i.i.d. framework are simpler to implement but are not appropriate for dependent data as it fails to capture the dependence structure of the underlying DGP. One well known exception to this rule is the case of the parametric bootstrap. Bose (1988, 1990) show that bootstrap approximation is valid and obtain higher order refinements for AR and MA processes. The main feature of an ARMA model is that the DGP is driven by the noise terms, since consistent estimators for the ARMA coefficients can be obtained under weak conditions, it is easy to construct bootstrap samples that mimic the dependence structure of the true DGP when the distribution of the noise terms is assumed.

The structural models we are interested in seem to possess enough structures suitable for a resampling scheme akin to that of the parametric bootstrap. Indeed, Kasahara and Shimotsu (2008a) has recently developed a bootstrap procedure for parametric discrete Markov decision models, where they use parametric bootstrap framework of Andrews (2002,2005) to obtain higher order refinements of their nested pseudo likelihood estimators. However, our problem is a semiparametric one. Recall that the primitives of the controlled Markov decision processes is the triple $(\beta, u_\theta, p)$, since we assume the complete knowledge of the discounting factor and the law of the unobserved error, the remaining primitives are $\theta$ and $p_{X'|X,A}$, both of which can be consistently estimated as shown in the previous sections. Therefore the semiparametric bootstrap seems to be a natural resampling method to use since we know the DGP for the controlled processes up to an estimation error. We now give the details to obtain the bootstrap samples.

Step 1:

Given the observations $\{a_{it}, x_{it}\}_{i=1,t=1}^{N,T+1}$ we obtain the estimators $\left(\widehat{\theta}, \widehat{g}\left(\cdot, \widehat{\theta}\right)\right)$ as described in Section 2.

Step 2:

We use $\{x_{i0}\}_{i=1}^{N}$ to construct the empirical distribution of the initial states, $F_N^{X_0}$ and draw (with replacement) $N$ bootstrap samples $\{x_{i1}^*\}_{i=1}^{N}$. These are to be used as the bootstrap initial states for each $i$ to construct $N$ series of length $T+1$.

STEP 3:

For each $i$, $\varepsilon_{it}^*$ is independently drawn from $Q$. Using the estimated policy profile $\left(\alpha_j\left(\cdot;\widehat{\theta},\partial\widehat{g}_j\left(\cdot,\widehat{\theta}\right)\right)\right)$, we compute for each $x_{it}^*=j$, $a_{it}^*=\alpha_j\left(\varepsilon_{it}^*;\widehat{\theta},\partial\widehat{g}_j\left(\cdot,\widehat{\theta}\right)\right)$. Also for each $x_{it}^*=j$ and $a_{it}^*$, $x_{it+1}^*$ is drawn from the nonparametric estimate of the transitional distribution $\widehat{p}_{X',X,A}\left(x_{it+1}^*,j,a_{it}^*\right)/\widehat{p}_{X,A}\left(j,a_{it}^*\right)$. Beginning with $t=0$, this process is continued successively to obtain $\{a_{it}^*,x_{it}^*\}_{i=1,t=1}^{N,T+1}$.

STEP 4:

Using $\{a_{it}^*,x_{it}^*\}_{i=1,t=1}^{N,T+1}$ to obtain the bootstrap estimates $\left(\widehat{\theta}^*,\widehat{g}^*\left(\cdot,\theta\right)\right)$ as done with the original data.

STEP 5:

Steps 2-4 is repeated $B-$times to obtain $B-$bootstrap estimates of $\left\{\widehat{\theta}_{(b)}^*,\widehat{g}_{(b)}^*\left(\cdot,\theta\right)\right\}_{b=1}^{B}$.

Then $\left\{\widehat{\theta}_{(b)}^*,\widehat{g}_{(b)}^*\left(\cdot,\theta\right)\right\}_{b=1}^{B}$ can be used as a basis to estimate the statistic of interest. One should be able to show that the method described above can be used to show the sampling distribution of $\sqrt{NT}\left(\widehat{\theta}-\theta_0\right)$ can be consistently estimated by $\sqrt{NT}\left(\widehat{\theta}^*-\widehat{\theta}\right)$, possibly with an additional bias correction term. The proof strategy analogous to the arguments of Arcones and Giné (1992), see also Brown and Wegkamp (2002), can be shown to accommodate a two-step semiparametric M-estimators considered in this chapter.

## 2.7 Simulation Study

In this section we illustrate some finite sample properties of our proposed estimator in a small scale Monte Carlo experiment. Since the generation of controlled Markov processes can be quite complicated, for simplicity, we consider a dynamic price setting problem for a representative firm described in Section 2 with the following specification.

DESIGN:

Each firm faces the following demand

$$D(a_t, x_t, \varepsilon_t) = \overline{D} - \theta_1 a_t + \theta_2 (x_t + \varepsilon_t).$$

such that $a_t$ belongs to some compact and convex set $A \subset \mathbb{R}$; $x_t$ takes value either 1 or $-1$, where 1 signifies an increase in demand towards the firm's product and vice versa; the firm's private shock in demand $\varepsilon_t$ has a known distribution. $\overline{D}$ can be interpreted as the upper bound of the supply and $(\theta_1, \theta_2)$ are the parameters representing the market elasticities. Unlike $x_t$, the evolution of the private shocks $\varepsilon_t$, are completely random and transitory. The distribution of the consumer satisfaction measure depends on the previous period's price set by the firm, which is summarized by $\Pr[x_{t+1} = -1 | x_t, a_t] = \frac{a_t - \underline{a}}{\overline{a} - \underline{a}}$, where $\underline{a}$ and $\overline{a}$ denote the minimum and maximum possible prices respectively. It is a simple exercise to show that the policy function can be characterized in terms of the conditional value function $E[V_\theta(x_{t+1}, \varepsilon_{t+1}) | x_t]$, in particular, the firm's optimal pricing strategy has the following explicit form

$$\alpha(x_t, \varepsilon_t) = \left(\overline{D} + \theta_2 (x_t + \varepsilon_t) + c\theta_1 - \beta \frac{\lambda_{\theta,1} - \lambda_{\theta,2}}{\overline{a} - \underline{a}}\right) / 2\theta_1, \tag{96}$$

where $\lambda_{\theta,1} = E[V_\theta(x_{t+1}, \varepsilon_{t+1}) | x_{t+1} = 1]$ and $\lambda_{\theta,2} = E[V_\theta(x_{t+1}, \varepsilon_{t+1}) | x_{t+1} = -1]$. It can be shown that $D(a_t, x_t, \varepsilon_t)(a_t - c)$ will be is supermodular in $(a_t, \varepsilon_t)$ if $(\theta_1, \theta_2)$ is positive, as expected from Topkis' theorem, the policy function above will then be strictly increasing in $\varepsilon_t$. If we ignore that the firm is forward looking, the optimal static profit can be characterized by the following pricing policy

$$\alpha_s(x_t, \varepsilon_t) = \left(\overline{D} + \theta_2(x_t + \varepsilon_t) + c\theta_1\right)/2\theta_1. \tag{97}$$

Intuitively, we expect firms which do not take into the account of the consumer's adverse response to high prices will overprice relative to their forward looking counterparts. This is confirmed by the expressions in the displays above since we expect $\lambda_{\theta,1} - \lambda_{\theta,2}$ (and $\theta_1$) to be positive, i.e. the latter implies $\alpha_s(x, \varepsilon) > \alpha(x, \varepsilon)$ for any pair of $(x, \varepsilon)$. From (96) - (97), identification issue aside, we also note that performing linear regression of $a_{it}$ on $x_{it}$ will yield estimable objects that are functions of the model primitives $\left(\overline{D}, \theta_1, \theta_2, c\right)$ that may have little economic interpretation.

In our design, we assign the following values to the parameters

$$\overline{D} = 3, \theta_1 = 1, \theta_2 = 1/2, c = 1,$$

and let $\varepsilon_t \sim Uni[-1, 1]$. It can be shown that $\overline{a} - \underline{a} = 1$ and

$$\begin{aligned}
\mathcal{L} &= \beta \begin{pmatrix} \Pr[x_{t+1} = 1 | x_t = 1] & \Pr[x_{t+1} = -1 | x_t = 1] \\ \Pr[x_{t+1} = 1 | x_t = -1] & \Pr[x_{t+1} = -1 | x_t = -1] \end{pmatrix} \\
&= \beta \begin{pmatrix} 0.25 & 0.75 \\ 0.75 & 0.25 \end{pmatrix}.
\end{aligned}$$

A numerical method that mirrors our estimation of the policy value equation in Section 2 can be used to show that $\lambda_{\theta,1} - \lambda_{\theta,2} = 1/1.45$. Combining these information, it is then straightforward to simulate the controlled Markov processes that are consistent with optimal pricing behavior in (96) that underlies the dynamic problem of interest. We generate 1000 replications of such controlled Markov processes with for various sizes of $N \in \{20, 100, 200, 500\}$ random samples of decision series over 5 time periods; this leads to five sets of experiments with the total sample size, $NT$, of $100, 500, 1000$ and 2500.

IMPLEMENTATION:

We are interested in obtain estimates for the demand parameters $(\theta_1, \theta_2)$ and assume the knowledge of $(\overline{D}, c)$. In estimating the nonparametric estimator of $g_0(\cdot, \theta)$, we construct a truncated 4-th order kernel based on the density of a standard normal random variable, see Rao (1983). For each replication, we experiment with 5 different bandwidths $\left\{h_\varsigma = 1.06s(NT)^{-\varsigma} : \varsigma = \frac{1}{5}, \frac{1}{6}, \frac{1}{7}, \frac{1}{8}, \frac{1}{9}\right\}$. We provide two estimators for each of $(\theta_1, \theta_2)$, one without trimming and another one that trims out calculations involving $\widehat{g}(\cdot, \theta)$ for $a$ that lies within a bandwidth neighborhood of the boundary. For the simulation of $F_{A|X}(a|j; \theta, \partial g_j)$, we take $R = N \log(N)$ random draws from $Q$. We approximate the model implied policy function by using grid-search instead of computing the derivative of the continuation value. The measures $(\mu_1, \mu_2)$ we use in constructing the minimum distance estimator in (92) simply put equal weights on all $a$ and $x$. It is much simpler to estimate the parameters when we assume the underlying model is static. Note that the policy function in such framework has a closed form as shown in (97), therefore we can simulate the model implied conditional distribution function directly from $\alpha_s$.

COMMENTS AND RESULTS:

The first observation is that our simulation design does not satisfy all of the conditions of E1. In particular, the support of price differs depending on the observable level of the popularity measure. This knowledge can be used in the estimation procedure without affecting any of our asymptotic results, as we commented in the previous sections, we assume common full support for each state for simplicity.

All of the Figures and Tables can be found at the end of the chapter. We report the bias, median of the bias, standard deviation and interquartile range (scaled by 1.349) for the estimators of $\theta_1$ and $\theta_2$. The rows are arranged according to the total sample size and bandwidths. We have the following general observations for both estimators regardless of bandwidth choice and trimming: (i) the median of the bias is similar to the mean; (ii) the estimators converge to the true values as $N$ increases and their respective standard deviations are converging to zero; (iii) the standard deviation figures are similar to the corresponding scaled interquartile range.[20] However, the effect of trimming is unclear. In the case of the estimator of $\theta_1$, the magnitude of the bias is significantly reduced by trimming that appear to far outweights the increase in variation in the MSE sense. On the contrary, trimming generally slightly increase the bias of the estimator for $\theta_2$. We check the distribution of our estimators by using QQ-plots. We only provide QQ-plots of the numerical results for the case of the trimmed estimator using $\varsigma = 1/7$ for the bandwidth $h_\varsigma$. Figures 1-4 plot the quantiles of $\left(\widehat{\theta}_1 - E\widehat{\theta}_1\right)/SE\left(\widehat{\theta}_1\right)$ with that of a standard normal for different sample sizes, where the dashed line denotes the 45 line and plots are marked by '+'; Figures 5-8 do the same for $\widehat{\theta}_2$. The distributional approximation supports our theory that $\widehat{\theta}$ behaves more like a normal random variable as $N$ increases. We find that the untrimmed estimators produce similar plots to their untrimmed counterparts across all bandwidths considered especially for the larger sample sizes, however, the quality of the QQ-plots varies across different bandwidth choices.

We also report analogous summary statistics for the structural estimation assuming the model is static, they can also be found in Table 5 and 6 in the rows labelled *static*. Note that the estimation of the static model does not involve the continuation value function so it does not depend on the bandwidth choice. It is clear that the estimators

[20](iii) is a characteristic of a normal random variable.

under static environment do not converge to $(\theta_1, \theta_2) = (1, 0.5)$, instead they converge to some values around $(1.26, 0.68)$ with very small standard errors. Since our minimum distance estimators reflect the model that best fit the observed data, the upwards bias of the elasticity parameters estimates is highly plausible. To see this, first recall from (97) that firms who do not take into the account of the future dynamics will overprice relative to the forward looking firms. The firms that only maximize their static profits will therefore, on average, need to expect the market elasticities to be more sensitive in order to generate more conservative pricing schemes consistent with the behaviors of their forward looking counterparts. Thus, in this example, ignoring the model dynamics leads to overestimating the elasticity parameters.

## 2.8 Conclusion

In this chapter we develop a new two-step estimator for a class of Markov decision processes with continuous control that forms a basis to estimate a larger class of structural dynamic models. Our criterion function has a simple interpretation and is also simple to construct; we minimize a minimum distance criterion that measures the divergence between two estimators of the conditional distribution function of the observables. In particular, we compare the conditional distribution functions, one implied purely by the data with another constructed from the structural model. We provide some primitive conditions to ensure our estimator is consistent for the structure parameter of interest when the model is identified. As an alternative estimator to BBL, which is designed to estimate the same class of models without having to solve for the equilibrium, in a parametric model we can simply use the empirical measure to construct our objective functions hence there is no additional decisions to be made by the practitioners (e.g. choosing classes of inequalities). We also explicitly work with the framework where we do not need to impose any distributional assumptions on the transition law of the observables. This additional flexibility is important since the transition law is a model primitive. We provide the distribution theory of both the finite dimensional parameters as well as the conditional value functions and propose to use semiparametric bootstrap to estimate the standard error for inference. We illustrate the performance of our estimator in a Monte Carlo experiment on a dynamic pricing problem, where we provide an intuition, based on our criterion function, for the direction of the biased of the estimator which ignore the model dynamics. We also demonstrate how the general approach we take to estimate dynamic models with continuous control, analogous to the discrete choice counterparts proposed by Hotz and Miller (1993) and Pesendorfer and Schmidt-Dengler (2008), can easily be adapted to estimate other class of interesting and practically relevant dynamic models.

There are also other important aspects of dynamic models we do not discuss in this paper. We end with a brief note of two issues that are particularly relevant to our framework. The first is regarding unobserved heterogeneity. The absence of unobserved heterogeneity has long been the main criticism against two-step approaches developed along the line of HM. Recently, finite mixtures have been used to add unobserved components in related two-step estimation methodologies, for example see Aguirregabiria and Mira (2007) and Arcidiacono and Miller (2008), Kasahara and Shimotsu (2008a,b). Finite mixture models can also be used with the estimator developed in this paper. Secondly, our paper focuses on estimation and assumes the model is identified. There are ongoing research on the nonparametric and semiparametric identification of dynamic decision models of single and multiple agents, for some samples, we refer interested readers to Aguirregabiria (2008), Bajari et al. (2009), Heckman and Navarro (2007) and Hu and Shum (2009) for examples.

## 2.9 Proofs of Theorems

### 2.9.1 Proof of Theorem G

The argument proceeds in a similar fashion to the case with no preliminary estimates of Newey and McFadden (1994, Theorem 7.1), see also Pollard (1985), by first showing that $\widehat{\theta}$ converge to $\theta_0$ at rate $N^{-1/2}$. By definition of the estimator, we have $M_N\left(\widehat{\theta},\widehat{g}\left(\cdot,\widehat{\theta}\right)\right)-M_N\left(\theta_0,\widehat{g}\left(\cdot,\theta_0\right)\right)\leq o_p\left(N^{-1}\right)$, and

$$
\begin{aligned}
& M_N\left(\widehat{\theta},\widehat{g}\left(\cdot,\widehat{\theta}\right)\right)-M_N\left(\theta_0,\widehat{g}\left(\cdot,\theta_0\right)\right) \\
= \; & M\left(\widehat{\theta},\widehat{g}\left(\cdot,\widehat{\theta}\right)\right)-M\left(\theta_0,\widehat{g}\right)+C_N'\left(\widehat{\theta}-\theta_0\right)+N^{-1/2}\left\|\widehat{\theta}-\theta_0\right\|\mathcal{D}_N\left(\widehat{\theta},\widehat{g}\left(\cdot,\widehat{\theta}\right)\right) \\
\geq \; & \left(C_N+S\left(\theta_0,\widehat{g}\left(\cdot,\widehat{\theta}\right)\right)\right)'\left(\widehat{\theta}-\theta_0\right)+C_0\left\|\widehat{\theta}-\theta_0\right\|^2(1+o_p(1))+N^{-1/2}\left\|\widehat{\theta}-\theta_0\right\|\mathcal{D}_N\left(\widehat{\theta},\widehat{g}\left(\cdot,\widehat{\theta}\right)\right) \\
= \; & O_p\left(N^{-1/2}\right)'\left(\widehat{\theta}-\theta_0\right)+C_0\left\|\widehat{\theta}-\theta_0\right\|^2+o_p\left(N^{-1/2}\left\|\widehat{\theta}-\theta_0\right\|+\left\|\widehat{\theta}-\theta_0\right\|^2\right).
\end{aligned}
$$

The first equality follows from the definition of $\mathcal{D}_N$ in (93). For the inequality, we expand $M\left(\widehat{\theta},\widehat{g}\left(\cdot,\widehat{\theta}\right)\right)$ around $\theta_0$, since $H(\theta,g)$ is continuous around $(\theta_0,g_0)$ and $H_0$ is positive definite by G3, there exists some $C_0>0$ such that, w.p.a. 1, $(\theta-\theta_0)'H\left(\theta_0,\widehat{g}\left(\cdot,\theta_0\right)\right)(\theta-\theta_0)+o_p\left(\|\theta-\theta_0\|^2\right)\geq C_0\|\theta-\theta_0\|^2$. Notice that $C_N+S\left(\theta_0,\widehat{g}\left(\cdot,\theta_0\right)\right)=O_p\left(N^{-1/2}\right)$, the first term follows from assumption G6 and the latter by G3 and G6 since

$$
\begin{aligned}
\left\|S\left(\theta_0,\widehat{g}\left(\cdot,\theta_0\right)\right)\right\| \leq \; & \left\|S\left(\theta_0,\widehat{g}\left(\cdot,\theta_0\right)\right)-D_gS\left(\theta_0,g_0\left(\cdot,\theta_0\right)\right)\left[\widehat{g}\left(\cdot,\theta_0\right)-g_0\left(\cdot,\theta_0\right)\right]\right\| \\
& +\left\|D_gS\left(\theta_0,g_0\left(\cdot,\theta_0\right)\right)\left[\widehat{g}\left(\cdot,\theta_0\right)-g_0\left(\cdot,\theta_0\right)\right]\right\| \\
\leq \; & o_p\left(N^{-1/2}\right)+O_p\left(N^{-1/2}\right) \\
= \; & O_p\left(N^{-1/2}\right).
\end{aligned}
$$

By completing the square

$$
\left(\left\|\widehat{\theta}-\theta_0\right\|+O_p\left(N^{-1/2}\right)\right)^2+o_p\left(N^{-1/2}\left\|\widehat{\theta}-\theta_0\right\|+\left\|\widehat{\theta}-\theta_0\right\|^2\right)\leq o_p\left(N^{-1}\right),
$$

thus $\left\|\widehat{\theta}-\theta_0\right\| = O_p\left(N^{-1/2}\right)$. To obtain the asymptotic distribution we define the following related criterion, $J_N(\theta) = D_N(\theta-\theta_0) + \frac{1}{2}(\theta-\theta_0)' H_0 (\theta-\theta_0)$, note that $J_N(\theta)$ is defined for each $\widehat{g}(\cdot,\theta)$ that satisfies the conditions of Theorem G2, implicit in $D_N$. $J_N(\theta)$ is a quadratic approximation of $M_N(\theta,\widehat{g}(\cdot,\theta)) - M_N(\theta_0,\widehat{g}(\cdot,\theta_0))$, whose unique minimizer is $\widetilde{\theta} = \theta_0 - H_0^{-1} D_N$, and $\sqrt{N}\left(\widetilde{\theta}-\theta_0\right) \Rightarrow \mathcal{N}\left(0, H_0^{-1}\Omega H_0^{-1}\right)$. Next, we show the approximation error of $J_N(\theta)$ from $M_N(\theta,\widehat{g}(\cdot,\theta)) - M_N(\theta_0,\widehat{g}(\cdot,\theta))$ is small. For any $\theta_N = \theta_0 + O_p\left(N^{-1/2}\right)$ in $\Theta_{\delta_N}$,

$$
\begin{aligned}
& M_N(\theta_N,\widehat{g}(\cdot,\theta_N)) - M_N(\theta_0,\widehat{g}(\cdot,\theta_0)) \\
= \;& M(\theta_N,\widehat{g}(\cdot,\theta_N)) - M(\theta_0,\widehat{g}(\cdot,\theta_0)) + C_N'(\theta_N-\theta_0) + \frac{\|\theta_N-\theta_0\|}{\sqrt{N}} \mathcal{D}_N(\theta_N,\widehat{g}(\cdot,\theta_N)) \\
= \;& (C_N + S(\theta_0,\widehat{g}(\cdot,\theta_0)))'(\theta_N-\theta_0) + \frac{1}{2}(\theta_N-\theta_0)' H\left(\overline{\theta},\widehat{g}\left(\cdot,\overline{\theta}\right)\right)(\theta_N-\theta_0) + \frac{\|\theta_N-\theta_0\|}{\sqrt{N}} \mathcal{D}_N(\theta_N \\
= \;& D_N'(\theta_N-\theta_0) + \frac{1}{2}(\theta_N-\theta_0)' H_0 (\theta_N-\theta_0) + o_p\left(\frac{\|\theta_N-\theta_0\|}{\sqrt{N}} + \|\theta_N-\theta_0\|^2\right) \\
= \;& J_N(\theta_N) + o_p\left(\frac{1}{N}\right).
\end{aligned}
$$

The equalities in the display follow straightforwardly from the definition of the $\mathcal{D}_N$, G3, G4 and G5. In particular, this implies that $M_N(\theta_N,\widehat{g}(\cdot,\theta_N)) - M_N(\theta_0,\widehat{g}(\cdot,\theta_0)) = J_N(\theta_N) + o_p\left(\frac{1}{N}\right)$ for $\theta_N = \widehat{\theta}$ and $\widetilde{\theta}$, hence we have

$$
\begin{aligned}
J_N\left(\widehat{\theta}\right) &= \left(J_N\left(\widehat{\theta}\right) - \left(M_N(\theta_N,\widehat{g}(\cdot,\theta_N)) - M_N(\theta_0,\widehat{g}(\cdot,\theta_0))\right)\right) \\
&\quad + \left(M_N(\theta_N,\widehat{g}(\cdot,\theta_N)) - M_N(\theta_0,\widehat{g}(\cdot,\theta_0))\right) \\
&\leq J_N\left(\widetilde{\theta}\right) + o_p\left(\frac{1}{N}\right),
\end{aligned}
$$

where the inequality follows from the relation derived from the previous display and G1. Since $J_N\left(\widetilde{\theta}\right) \le J_N\left(\widehat{\theta}\right)$,

$$\begin{aligned}
o_p\left(\frac{1}{N}\right) &= J_N\left(\widehat{\theta}\right) - J_N\left(\widetilde{\theta}\right) \\
&= D_N\left(\widehat{\theta}-\theta_0\right) + \frac{1}{2}\left(\widehat{\theta}-\theta_0\right)' H_0\left(\widehat{\theta}-\theta_0\right) - D_N\left(\widetilde{\theta}-\theta_0\right) - \frac{1}{2}\left(\widetilde{\theta}-\theta_0\right)' H_0\left(\widetilde{\theta}-\theta_0\right) \\
&= -\left(\widetilde{\theta}-\theta_0\right)' H_0\left(\widehat{\theta}-\theta_0\right) + \frac{1}{2}\left(\widehat{\theta}-\theta_0\right)' H_0\left(\widehat{\theta}-\theta_0\right) + \frac{1}{2}\left(\widetilde{\theta}-\theta_0\right)' H_0\left(\widetilde{\theta}-\theta_0\right) \\
&= \frac{1}{2}\left(\widehat{\theta}-\widetilde{\theta}\right)' H_0\left(\widehat{\theta}-\widetilde{\theta}\right),
\end{aligned}$$

this implies that $\left\|\widehat{\theta}-\widetilde{\theta}\right\|^2 = o_p\left(\frac{1}{N}\right)$. Since $\sqrt{N}\left(\widetilde{\theta}-\theta_0\right)$ has the desired asymptotic distribution, this completes the proof.■

For the proof of Theorems 2 and 3 we find it convenient to introduce the following notations: $M\left(\theta, g\left(\cdot,\theta\right)\right) = \sum_{j=1}^{J}\int_A E_j^2\left(\theta, \partial_a g_j\left(\cdot,\theta\right)\right) d\mu_j$ where $E_j\left(\theta, \partial_a g_j\left(\cdot,\theta\right)\right) = F_{A|X=j}\left(\theta, \partial_a g_j\left(\cdot,\theta\right)\right) - F_{A|X=j}$, and, $F_{A|X=j}\left(\theta, \partial_a g_j\left(\cdot,\theta\right)\right)$ and $F_{A|X=j}$ are functions defined on $A$ that are the shorthand notations for $F_{A|X}\left(\cdot|j;\theta,\partial_a g_j\left(\cdot,\theta\right)\right)$ and $F_{A|X}\left(\cdot|j\right)$ respectively; $M_N\left(\theta, g\left(\cdot,\theta\right)\right) = \sum_{j=1}^{J}\int_A E_{N,j}^2\left(\theta, \partial_a g_j\left(\cdot,\theta\right)\right) d\mu_j$ where $E_{N,j}\left(\theta, \partial_a g_j\left(\cdot,\theta\right)\right) = \widetilde{F}_{A|X=j}\left(\theta, \partial_a g_j\left(\cdot,\theta\right)\right) - \widehat{F}_{A|X=j}$, and, $\widetilde{F}_{A|X=j}\left(\theta, \partial_a g_j\left(\cdot,\theta\right)\right)$ and $\widehat{F}_{A|X=j}$ are functions defined on $A$ that are the shorthand notations for $\widetilde{F}_{A|X}\left(\cdot|j;\theta,\partial_a g_j\left(\cdot,\theta\right)\right)$ and $\widehat{F}_{A|X}\left(\cdot|j\right)$ respectively; $F_{0,j}$ is a function defined on $A$ that is the shorthand notation for $F_0\left(\cdot|j\right)$. In addition, for $j = 1,\ldots,J$, we define the class of functions $\mathcal{F}_j = \left\{\mathbf{1}\left[\cdot \le \rho_j\left(a,\theta,\partial_a g_j\right)\right] : a \in A, \theta \in \Theta\right.$ and let $\nu_{R,j}$ denote the empirical process indexed by $\left(\theta,\partial_a g_j\right) \in \Theta \times \mathcal{G}_j^{(1)}$ to be a random element that takes value over $A$, i.e. $\nu_{R,j}\left(\theta,\partial_a g_j\right) = \frac{1}{\sqrt{R}}\sum_{r=1}^{R}\mathbf{1}\left[\varepsilon_r \le \rho_j\left(\cdot,\theta,\partial_a g_j\right)\right] - Q_\varepsilon\left(\rho_j\left(\cdot,\theta,\partial_a g_j\right)\right)$. We will continue to use the multi-index notation to define higher order derivatives $\partial_a^{|\eta|}$ and $\partial_\theta^{|\eta|}$, of $a$ and $\theta$ respectively for some natural number $\eta$, as seen in (94).

### 2.9.2 Proofs of Theorems 2.1 - 2.4

We next present the following lemmas that will be useful in proofing Theorems 2.1-2.3.

LEMMA 2.1. *Under E1* $\left\|\widehat{\mathcal{L}}-\mathcal{L}\right\| = O_p\left(N^{-1/2}\right)$.

LEMMA 2.2. *Under E1: For any* $r_\theta \in \mathcal{R}_0$ *and* $j=1,\ldots,J$, $\widetilde{r}_\theta(j) = r_\theta(j) + \widetilde{r}_\theta^R(j)$ *such that* $\max_{1\leq j\leq J}\sup_{\theta\in\Theta}\left|\widetilde{r}_\theta^R(j)\right| = o_p\left(N^{-\lambda}\right)$ *for any* $\lambda < 1/2$.

LEMMA 2.3. *Under E1: For any* $m_\theta \in \mathcal{M}_0$ *and* $j=1,\ldots,J$, $\widehat{m}_\theta(j) = m_\theta(j) + \widehat{m}_\theta^R(j)$ *such that* $\max_{1\leq j\leq J}\sup_{\theta\in\Theta}\left|\widehat{m}_\theta^R(j)\right| = o_p\left(N^{-\lambda}\right)$ *for any* $\lambda < 1/2$.

LEMMA 2.4. *Under E1: For any* $\theta\in\Theta$, $j=1,\ldots,J$, *and* $a\in A$, $\widehat{g}_j(a,\theta) = g_j(a,\theta) + \widehat{g}_j^B(a,\theta) + \widehat{g}_j^S(a,\theta) + \widehat{g}_j^R(a,\theta)$ *such that*

$$\begin{aligned}
\max_{1\leq j\leq J}\sup_{\theta,a\in\Theta\times A_N}\left|\widehat{g}_j^B(a,\theta)\right| &= O_p\left(h^4\right),\\
\max_{1\leq j\leq J}\sup_{\theta,a\in\Theta\times A_N}\left|\widehat{g}_j^S(a,\theta)\right| &= o_p\left(\frac{N^\xi}{\sqrt{Nh}}\right),\\
\max_{1\leq j\leq J}\sup_{\theta,a\in\Theta\times A_N}\left|\widehat{g}_j^R(a,\theta)\right| &= o_p\left(h^4+\frac{N^\xi}{\sqrt{Nh}}\right).
\end{aligned}$$

LEMMA 2.5. *Under E1: For all* $j=1,\ldots,J$, $\max_{0\leq l\leq 2,1\leq j\leq J}\sup_{\theta,a\in\Theta\times A_N}\left|\partial_a^{|l|}\widehat{g}_j(a,\theta) - \partial_a^{|l|}g_{0,j}\right.$ $o_p(1)$.

LEMMA 2.6. *Under E1:* $\max_{0\leq l,p\leq 2,1\leq j\leq J}\sup_{\theta,a\in\Theta\times A_N}\left|\partial_a^{|l|}\partial_\theta^{|p|}\widehat{g}_j(a,\theta) - \partial_a^{|l|}\partial_\theta^{|p|}g_{0,j}(a,\theta)\right| =$ $o_p(1)$.

LEMMA 2.7. *Under E1 and E2: for all* $j=1,\ldots,J$, $\mathcal{F}_j$ *is a Donsker class.*

LEMMA 2.8 *Under E1 and E2: For any* $j=1,\ldots,J$ *and some positive sequence* $\delta_N = o(1)$ *as* $N\to\infty$

$$\lim_{N\to\infty}\sup_{\substack{(a,\theta,\partial_a g_j)\in A\times\Theta\times\mathcal{G}_j^{(1)},\\ \left\|\left(a'-a,\theta'-\theta,\partial_a g_j'-\partial_a g_j\right)\right\|<\delta_N}}\left|\begin{array}{c}\frac{1}{N}\sum_{i=1}^N\left(\mathbf{1}\left[\varepsilon_i\leq\rho_j\left(a',\theta',\partial_a g_j'\right)\right]-Q_\varepsilon\left(\rho_j\left(a',\theta',\partial_a g_j'\right)\right)\right)\\ -\frac{1}{N}\sum_{i=1}^N\left(\mathbf{1}\left[\varepsilon_i\leq\rho_j\left(a,\theta,\partial_a g_j\right)\right]-Q_\varepsilon\left(\rho_j\left(a,\theta,\partial_a g_j\right)\right)\right)\end{array}\right| = 0.$$

LEMMA 2.9 *Under E1: For any* $j = 1, \ldots, J$

$$\sqrt{N}\left(\widehat{F}_{A|X=j} - F_{A|X=j}\right) \rightsquigarrow \mathbb{F}_j,$$

*where* $\mathbb{F}_j$ *is a tight Gaussian process that belongs to* $l^{\infty}(A)$.

LEMMA 2.10: *Under E1 and E2: For any* $j = 1, \ldots, J$

$$\sqrt{N}(F_{A|X=j}(\theta_0, \partial_a \widehat{g}_j(\cdot, \theta_0)) - F_{A|X=j}(\theta_0, \partial_a g_{0,j}(\cdot, \theta_0))) \rightsquigarrow \mathbb{G}_j,$$

where $\mathbb{G}_j$ is a tight Gaussian process that belongs to $l^{\infty}(A)$.

PROOF OF THEOREM 2.1. This follows from Lemma 2.4. For the asymptotic distribution, we only have to calculate the variance of (105), the rest follows by standard CLT. Asymptotic independence will follow if we can show $\sqrt{Nh}\mathrm{cov}(\widehat{g}_j(a, \theta), \widehat{g}_k(a', \theta)) = o(1)$ for any $k \neq j$ and $a' \neq a$, this is trivial to show.■

PROOF OF THEOREM 2.2. We first show that $M(\theta, g_0(\cdot, \theta))$ has a well separated minimum at $\theta_0$. By assumption (ii) - (iii) and (vii) we have $M(\theta, g_0(\cdot, \theta)) \geq M(\theta_0, g_0(\cdot, \theta_0))$ for all $\theta$ in the compact set $\Theta$ with equality only holds for $\theta = \theta_0$. For each $a$ and $j$, we have $F_{A|X}(a|j; \theta, \partial_a g_j(\cdot, \theta)) = Q_\varepsilon\left(\rho_j(a, \theta, \partial_a g_0(\cdot, \theta))\right)$ which is continuous in $\theta$ given assumptions (vii) and (xiii), this ensures a well-separated minimum. By standard arguments, consistency will now follow if we can show

$$\sup_{\theta \in \Theta} |M_N(\theta, \widehat{g}(\cdot, \theta)) - M(\theta, g_0(\cdot, \theta))| = o_p(1). \tag{98}$$

By the triangle inequality, we have

$$
\begin{aligned}
|M_N(\theta,\widehat{g}(\cdot,\theta)) - M(\theta,g_0(\cdot,\theta))| &\le 4\sum_{j=1}^{J}\int \left|\widetilde{F}_{A|X=j}(\theta,\partial_a\widehat{g}_j(\cdot,\theta)) - F_{A|X=j}(\theta,\partial_a\widehat{g}_j(\cdot,\theta))\right| d\mu_j \\
&\quad +4\sum_{j=1}^{J}\int \left|F_{A|X=j}(\theta,\partial_a\widehat{g}_j(\cdot,\theta)) - F_{A|X=j}(\theta,\partial_a g_{0,j}(\cdot,\theta))\right| d\mu \\
&\quad +4\sum_{j=1}^{J}\int \left|\widehat{F}_{A|X=j} - F_{A|X=j}\right| d\mu_j \\
&= A_1 + A_2 + A_3.
\end{aligned}
$$

For $A_1$, for each $j$ and any $\eta > 0$ we have

$$
\begin{aligned}
& \Pr\left[\sup_{\theta\in\Theta}\left|\widetilde{F}_{A|X}(a|x;\theta,\partial_a\widehat{g}(\cdot,\theta)) - F_{A|X}(a|x;\theta,\partial_a\widehat{g}(\cdot,\theta))\right| > \eta\right] \\
\le\ & \Pr\left[\sup_{\theta,a\in\Theta\times A_N}\left|\frac{1}{R}\sum_{r=1}^{R}\mathbf{1}\left[\varepsilon_r \le \rho_j(a,\theta,\partial_a\widehat{g}_j)\right] - Q_\varepsilon\left(\rho_j(a,\theta,\partial_a\widehat{g}_j)\right)\right| > \eta\right] \\
\le\ & \Pr\left[\sup_{\theta,a,\partial_a g_j\in\Theta\times A_N\times\mathcal{G}_j^{(1)}}\left|\frac{1}{R}\sum_{r=1}^{R}\mathbf{1}\left[\varepsilon_r \le \rho_j(a,\theta,\partial_a g_j)\right] - Q_\varepsilon\left(\rho_j(a,\theta,\partial_a g_j)\right)\right| > \eta\right] \\
& +\Pr\left[\partial_a\widehat{g}_j(\cdot,\theta) \notin \mathcal{G}_j^{(1)}\right].
\end{aligned}
$$

From Lemma 2.7, $\mathcal{F}_j$ is $Q-$Glivenko-Cantelli by Slutsky's theorem, therefore the first term of the last inequality above converges to zero as $R \to \infty$ by assumption (xii). By Lemma 2.6, $\Pr\left[\partial_a\widehat{g}_j(\cdot,\theta) \notin \mathcal{G}_j^{(1)}\right] = o(1)$ hence by finiteness of $\mu_j$ it follows that $|A_1| = o_p(1)$ uniformly over $\Theta$. For $A_2$, for each $j$ we have

$$
\begin{aligned}
\left|F_{A|X=j}(\theta,\partial_a\widehat{g}_j(\cdot,\theta)) - F_{A|X=j}(\theta,\partial_a g_j(\cdot,\theta))\right| &= \left|Q_\varepsilon\left(\rho_j(a,\theta,\partial_a\widehat{g}_j(\cdot,\theta))\right) - Q_\varepsilon\left(\rho_j(a,\theta,\partial_a g_{0,j}(\right.\right. \\
&\le C_0\left|\rho_j(a,\theta,\partial_a\widehat{g}_j(\cdot,\theta)) - \left(\rho_j(a,\theta,\partial_a g_{0,j}(\cdot,\theta)\right)\right)
\end{aligned}
$$

where the inequality follows from the mean value theorem (MVT) and the fact that the derivative of $Q_\varepsilon$ is uniformly bounded. Given the smoothness assumption on $(\rho_j)$ in assumption (xiii), by MVT in Banach space $\sup_{a\in A_N}|\rho_j(a,\theta,\partial_a\widehat{g}_j(\cdot,\theta)) -$

$(\rho_j(a,\theta,\partial_a g_{0,j}(\cdot,\theta)))\,| \leq \sup_{\theta,a,\partial_a g_j\in\Theta\times A\times\mathcal{G}_j^{(1)}}\left\|D_{\partial_a g}\rho_j(a,\theta,\partial_a g_j)\right\| \times \sup_{\theta,a\in\Theta\times A_N}|\partial_a\widehat{g}_j(a,\theta)-$ $\partial_a g_{0,j}(a,\theta)\,|$. Since $\mu_j$ has zero measure on the boundary of $A$, by Lemma 2.5, $\int|F_{A|X=j}(\theta,\partial_a\widehat{g}_j(\cdot,\theta)) - F_{A|X=j}(\theta,\partial_a g_{0,j}(\cdot,\theta))\,|d\mu_j \leq C_0\sup_{\theta,a\in\Theta\times A_N}|\partial_a\widehat{g}_j(a,\theta)-\partial_a g_{0,j}(a,\theta)\,|+2\mu_j(A\backslash A_N) =$ $o_p(1)$. So we also have $|A_2| = o_p(1)$ uniformly over $\Theta$. Lastly for $A_3$, for each $j$ we write

$$\begin{aligned}\widehat{F}_{A|X}(a|j) - F_{A|X}(a|j) &= \frac{1}{p_X(j)}\left[\widehat{F}_{A,X}(a,j) - F_{A,X}(a,j)\right] \\ &\quad -\frac{\widehat{F}_{A|X}(a|j)}{p_X(j)}\left[\widehat{p}_X(j) - p_X(j)\right],\end{aligned}$$

where $\widehat{F}_{A,X}(a,j) = \frac{1}{NT}\sum_{i=1,t=1}^{N,T}\mathbf{1}[a_{it}\leq a, x_{it}=j]$, then w.p.a. 1

$$\begin{aligned}&\max_{1\leq j\leq J}\sup_{a\in A}\left|\widehat{F}_{A|X}(a|j) - F_{A|X}(a|j)\right| \\ \leq\ & \frac{1}{\min_{1\leq j\leq J}p_X(j)}\max_{1\leq j\leq J}\sup_{a\in A}\left|\widehat{F}_{A,X}(a,j) - F_{A,X}(a,j)\right| \\ &+\left|\frac{\max_{1\leq j\leq J}\left[\widehat{p}_X(j) - p_X(j)\right]}{\min_{1\leq j\leq J}p_X(j)}\right|.\end{aligned}$$

By Lemma 2.9, the class of functions $\{\mathbf{1}[\cdot\leq a, x_{it}=j] - F_{A,X}(\cdot,j) : a\in A\}$ is also a Glivenko-Cantelli class, so: the first term on the RHS of the inequality above converges in probability to zero; the second term also converges in probability to zero since $\widehat{p}_X(j) - p_X(j) = o_p(1)$ for each $x\in X$. Since $A_3$ is independent of $\theta$, the uniform convergence in (98) holds and consistency follows.■

PROOF OF THEOREM 2.3. To proof Theorem 2.3 we set out to show that our assumptions imply we satisfy all the conditions of Theorem G. We showed consistency in Theorem 2.2. G1 is the definition of the estimator. For G2, it suffices to show $\partial_a\widehat{g}_j(\cdot,\theta)\in\mathcal{G}_{\delta_N,j}$ w.p.a. 1 and $\sup_{\theta\in\Theta}\|\partial_a\widehat{g}_j(\cdot,\theta) - \partial_a g_{0,j}(\cdot,\theta)\|_\infty = o_p\left(N^{-1/4}\right)$ for all $j=1,\ldots,J$. The former is implied by Lemma 2.6, from the proof of Lemma 2.5, the latter holds if $h^4+\frac{N^\xi}{\sqrt{Nh^3}} = o\left(N^{-1/4}\right)$, this is certainly the case when $h$ is in the suggested range. G3 and G4 simply requires translating the smoothness we impose in E1 and E2 to satisfy these conditions. Now we show G5, in particular we need to show that

$$\begin{aligned}
& M_N\left(\theta, \widehat{g}\left(\cdot, \theta\right)\right) - M_N\left(\theta_0, \widehat{g}\left(\cdot, \theta_0\right)\right) - \left(M\left(\theta, \widehat{g}\left(\cdot, \theta\right)\right) - M\left(\theta_0, \widehat{g}\left(\cdot, \theta_0\right)\right)\right) - \left(\theta - \theta_0\right)'\Theta_N \qquad (9.9) \\
= \; & o_p\left(\left\|\theta - \theta_0\right\|^2 + \frac{\left\|\theta - \theta_0\right\|}{\sqrt{N}} + \frac{1}{N}\right),
\end{aligned}$$

holds uniformly for $\left\|\theta - \theta_0\right\| < \delta_N$. Then for any pair $\left(\theta, \partial_a \widehat{g}_j\left(\cdot, \theta\right)\right)$ we can write

$$\begin{aligned}
E_j^2\left(\theta, \partial_a \widehat{g}_j\left(\cdot, \theta\right)\right) - E_j^2\left(\theta_0, \partial_a \widehat{g}_j\left(\cdot, \theta_0\right)\right) \;=\; & \left(F_{A|X=j}\left(\theta, \partial_a \widehat{g}_j\left(\cdot, \theta\right)\right) - F_{A|X=j}\left(\theta_0, \partial_a \widehat{g}_j\left(\cdot, \theta_0\right)\right)\right) \times \\
& \left(F_{A|X=j}\left(\theta, \partial_a \widehat{g}_j\left(\cdot, \theta\right)\right) + F_{A|X=j}\left(\theta_0, \partial_a \widehat{g}_j\left(\cdot, \theta_0\right)\right) - 2F_A\right.
\end{aligned}$$

and analogously

$$\begin{aligned}
E_{N,j}^2\left(\theta, \partial_a \widehat{g}_j\left(\cdot, \theta\right)\right) - E_{N,j}^2\left(\theta_0, \partial_a \widehat{g}_j\left(\cdot, \theta_0\right)\right) \;=\; & \left(\widetilde{F}_{A|X=j}\left(\theta, \partial_a \widehat{g}_j\left(\cdot, \theta\right)\right) - \widetilde{F}_{A|X=j}\left(\theta_0, \partial_a \widehat{g}_j\left(\cdot, \theta_0\right)\right)\right) \times \\
& \left(\widetilde{F}_{A|X=j}\left(\theta, \partial_a \widehat{g}_j\left(\cdot, \theta\right)\right) + \widetilde{F}_{A|X=j}\left(\theta_0, \partial_a \widehat{g}_j\left(\cdot, \theta_0\right)\right) - \right.
\end{aligned}$$

Combing these we have

$$
\begin{aligned}
& M_N\left(\theta,\widehat{g}\left(\cdot,\theta\right)\right)-M_N\left(\theta_0,\widehat{g}\left(\cdot,\theta_0\right)\right) \\
= & \sum_{j=1}^{J}\int\left[\begin{array}{c}\left[\begin{array}{c}R^{-1/2}\left(\nu_{R,j}\left(\theta,\partial_a\widehat{g}_j\left(\cdot,\theta\right)\right)-\nu_{R,j}\left(\theta_0,\partial_a\widehat{g}_j\left(\cdot,\theta_0\right)\right)\right)\\+\left(F_{A|X=j}\left(\theta,\partial_a\widehat{g}_j\left(\cdot,\theta\right)\right)-F_{A|X=j}\left(\theta_0,\partial_a\widehat{g}_j\left(\cdot,\theta_0\right)\right)\right)\end{array}\right]\\ \times\left[\begin{array}{c}R^{-1/2}\left(\nu_{R,j}\left(\theta,\partial_a\widehat{g}_j\left(\cdot,\theta\right)\right)+\nu_{R,j}\left(\theta_0,\partial_a\widehat{g}_j\left(\cdot,\theta_0\right)\right)\right)\\+\left(F_{A|X=j}\left(\theta,\partial_a\widehat{g}_j\left(\cdot,\theta\right)\right)+F_{A|X=j}\left(\theta_0,\partial_a\widehat{g}_j\left(\cdot,\theta_0\right)\right)-2\widehat{F}_{A|X=j}\right)\end{array}\right]\end{array}\right]d\mu_j \\
= & \sum_{j=1}^{J}\int E_j^2\left(\theta,\partial_a\widehat{g}_j\left(\cdot,\theta\right)\right)-E_j^2\left(\theta_0,\partial_a\widehat{g}_j\left(\cdot,\theta_0\right)\right)d\mu_j \\
& -2\sum_{j=1}^{J}\int\left(F_{A|X=j}\left(\theta,\partial_a\widehat{g}_j\left(\cdot,\theta\right)\right)-F_{A|X=j}\left(\theta_0,\partial_a\widehat{g}_j\left(\cdot,\theta_0\right)\right)\right)\left(\widehat{F}_{A|X=j}-F_{A|X=j}\right)d\mu_j \\
& +R^{-1/2}\sum_{j=1}^{J}\int\left[\begin{array}{c}\left[\nu_{R,j}\left(\theta,\partial_a\widehat{g}_j\left(\cdot,\theta\right)\right)-\nu_{R,j}\left(\theta_0,\partial_a\widehat{g}_j\left(\cdot,\theta_0\right)\right)\right]\\ \times\left[\begin{array}{c}F_{A|X=j}\left(\theta,\partial_a\widehat{g}_j\left(\cdot,\theta\right)\right)\\+F_{A|X=j}\left(\theta_0,\partial_a\widehat{g}_j\left(\cdot,\theta_0\right)\right)-2\widehat{F}_{A|X=j}\end{array}\right]\end{array}\right]d\mu_j \\
& +R^{-1/2}\sum_{j=1}^{J}\int\left[\begin{array}{c}\left[\nu_{R,j}\left(\theta,\partial_a\widehat{g}_j\left(\cdot,\theta\right)\right)+\nu_{R,j}\left(\theta_0,\partial_a\widehat{g}_j\left(\cdot,\theta_0\right)\right)\right]\\ \times\left[F_{A|X=j}\left(\theta,\partial_a\widehat{g}_j\left(\cdot,\theta\right)\right)-F_{A|X=j}\left(\theta_0,\partial_a\widehat{g}_j\left(\cdot,\theta_0\right)\right)\right]\end{array}\right]d\mu_j \\
& +R^{-1}\sum_{j=1}^{J}\int\left[\begin{array}{c}\left[\nu_{R,j}\left(\theta,\partial_a\widehat{g}_j\left(\cdot,\theta\right)\right)-\nu_{R,j}\left(\theta_0,\partial_a\widehat{g}_j\left(\cdot,\theta_0\right)\right)\right]\\ \times\left[\nu_{R,j}\left(\theta,\partial_a\widehat{g}_j\left(\cdot,\theta\right)\right)+\nu_{R,j}\left(\theta_0,\partial_a\widehat{g}_j\left(\cdot,\theta_0\right)\right)\right]\end{array}\right]d\mu_j \\
= & M\left(\theta,\widehat{g}\left(\cdot,\theta\right)\right)-M\left(\theta_0,\widehat{g}\left(\cdot,\theta_0\right)\right)+B_1+B_2+B_3+B_4.
\end{aligned}
$$

We now show that, out of $\{B_i\}_{i=1}^4$, $B_1$ is the leading term that contains $C_N$ in (99), the rest are of smaller stochastic order. Since we are only interested in what happens as $\|\theta-\theta_0\|\to 0$, in what follows, the little '$o$' and big '$O$' terms will be implicitly assumed to hold with $\|\theta-\theta_0\|\to 0$ and $N\to\infty$.

For $B_1$ :

By mean value expansion

$$
\begin{aligned}
B_1 &= -2(\theta-\theta_0)'\sum_{j=1}^{J}\int D_\theta F_{A|X=j}\left(\overline{\theta}_j,\partial_a\widehat{g}_j\left(\cdot,\overline{\theta}_j\right)\right)\left(\widehat{F}_{A|X=j}-F_{A|X=j}\right)d\mu_j \\
&= -2(\theta-\theta_0)'\sum_{j=1}^{J}\int D_\theta F_{A|X=j}\left(\theta_0,\partial_a g_{0,j}\left(\cdot,\theta_0\right)\right)\left(\widehat{F}_{A|X=j}-F_{A|X=j}\right)d\mu_j \\
&\quad -2(\theta-\theta_0)'\sum_{j=1}^{J}\int\left[\begin{array}{c}\left[D_\theta F_{A|X=j}\left(\overline{\theta}_j,\partial_a\widehat{g}_j\left(\cdot,\overline{\theta}_j\right)\right)-D_\theta F_{A|X=j}\left(\theta_0,\partial_a g_{0,j}\left(\cdot,\theta_0\right)\right)\right] \\ \times\left[\widehat{F}_{A|X=j}-F_{A|X=j}\right]\end{array}\right]d\mu_j \\
&= B_{11}+B_{12},
\end{aligned}
$$

where for each $j,\overline{\theta}_j$ is some intermediate value between $\theta$ and $\theta_0$ that corresponds to the MVT w.r.t. the $j-th$ summand. We first show that $B_{11}$ is the leading term that is equal to $(\theta-\theta_0)'C_N$ in (99) and that $\sqrt{N}C_N$ converges to a normal random variable. By Lemma 2.9 $\sqrt{N}\left(\widehat{F}_{A|X=j}-F_{A|X=j}\right)\rightsquigarrow\mathbb{F}_j$ where $\mathbb{F}_j$ is a tight Gaussian process that belongs to $l^\infty(A)$ for all $j$, by Slutsky theorem and a similar argument used in the proof of Lemma 2.9, it is easy to show that $D_\theta F_{A|X=j}\left(\theta_0,\partial_a g_{0,j}\left(\cdot,\theta_0\right)\right)\sqrt{N}(\widehat{F}_{A|X=j}$ $-F_{A|X=j})$ also converges weakly to a tight Gaussian process. To see the latter, note that for any $\partial_a g_j(\cdot,\theta)\in\mathcal{G}_j^{(1)}$

$$
\begin{aligned}
& D_\theta F_{A|X}\left(a|j;\theta,\partial_a g_j\left(\cdot,\theta\right)\right) \\
&= q\left(\rho_j\left(a,\theta,\partial_a g_j\left(\cdot,\theta\right)\right)\right)\left(\partial_\theta\rho_j\left(a,\theta,\partial_a g_j\left(\cdot,\theta\right)\right)+D_{\partial_a g}\rho_j\left(a,\theta,\partial_a g_j\left(\cdot,\theta\right)\right)\left[\partial_\theta\partial_a g\left(\cdot,\theta\right)\right]\right),
\end{aligned}
$$

where, $\partial_\theta$ denotes the ordinary $L-$dimensional partial derivative, $\partial/\partial\theta$, w.r.t. in the argument $\theta$. This is continuous on $A$ for any $j$. Now, if we define a linear continuous map $\mathbb{T}_j:l^\infty(A)\rightarrow\mathbb{R}$ (w.r.t. sup-norm) so that $\mathbb{T}_jf=\int D_\theta F_{A|X=j}\left(\theta_0,\partial_a g_{0,j}\left(\cdot,\theta_0\right)\right)fd\mu$ for any $f\in l^\infty(A)$ then the map is linear and continuous, the boundedness follows from the observation that $\sup_{a\in A}\left\|D_\theta F_{A|X=j}\left(\theta_0,\partial_a g_{0,j}\left(\cdot,\theta_0\right)\right)\right\|<\infty$. Then, by continuous mapping theorem (CMT)

$$\int D_\theta F_{A|X=j}\left(\theta_0, \partial_a g_{0,j}(\cdot,\theta_0)\right)\sqrt{N}\left(\widehat{F}_{A|X=j} - F_{A|X=j}\right)d\mu_j \rightsquigarrow \int D_\theta F_{A|X=j}\left(\theta_0, \partial_a g_{0,j}(\cdot,\theta_0)\right)\mathbb{F}_j d\mu_j$$

Furthermore, the limit is also Gaussian since we know Gaussianity is preserved for any tight Gaussian process that is transformed by a linear continuous map, see Lemma 3.9.8 of VW. So we let

$$\sqrt{N}C_N = \sum_{j=1}^{J}\int D_\theta F_{A|X=j}\left(\theta_0, \partial_a g_{0,j}(\cdot,\theta_0)\right)\sqrt{N}\left(\widehat{F}_{A|X=j} - F_{A|X=j}\right)d\mu_j, \qquad (100)$$

then $\sqrt{N}C_N$ also converges a Gaussian variable.

For $B_{12}$, for each $j$, by Cauchy Schwarz inequality we have

$$\begin{aligned}
&\left|(\theta-\theta_0)'\int\left(D_\theta F_{A|X=j}\left(\overline{\theta}_j, \partial_a\widehat{g}_j(\cdot,\overline{\theta}_j)\right) - D_\theta F_{A|X=j}\left(\theta_0, \partial_a g_{0,j}(\cdot,\theta_0)\right)\right)\left(\widehat{F}_{A|X=j} - F_{A|X=j}\right)d\mu_j\right| \\
\leq &\left[(\theta-\theta_0)'\int\left[\begin{array}{c}\left(D_\theta F_{A|X=j}\left(\overline{\theta}_j, \partial_a\widehat{g}_j(\cdot,\overline{\theta}_j)\right) - D_\theta F_{A|X=j}\left(\theta_0, \partial_a g_{0,j}(\cdot,\theta_0)\right)\right) \\ \times\left(D_\theta F_{A|X=j}\left(\overline{\theta}_j, \partial_a\widehat{g}_j(\cdot,\overline{\theta}_j)\right) - D_\theta F_{A|X=j}\left(\theta_0, \partial_a g_{0,j}(\cdot,\theta_0)\right)\right)'\end{array}\right]d\mu_j(\theta-\theta_0)\right] \\
&\times\left[\int\left[\widehat{F}_{A|X=j} - F_{A|X=j}\right]^2 d\mu_j\right]^{1/2},
\end{aligned}$$

where for each $j, \overline{\theta}_j$ is some intermediate value between $\theta_j$ and $\theta_{0,j}$ that corresponds to the MVT w.r.t. the $j-th$ summand. Let $\partial_{\theta_l}$ denotes the $l-$th element of $\partial_\theta$ then

$$\begin{aligned}
&\left|D_{\theta_l}F_{A|X=j}\left(\overline{\theta}_j, \partial_a\widehat{g}_j(\cdot,\overline{\theta}_j)\right) - D_{\theta_l}F_{A|X=j}\left(\theta_0, \partial_a g_{0,j}(\cdot,\theta_0)\right)\right| \\
\leq &\left|\begin{array}{c} q\left(\rho_j\left(a,\overline{\theta}_j,\partial_a\widehat{g}_j(\cdot,\overline{\theta}_j)\right)\right)\partial_{\theta_l}\rho_j\left(a,\overline{\theta}_j,\partial_a\widehat{g}_j(\cdot,\overline{\theta}_j)\right) \\ -q\left(\rho_j\left(a,\theta_0,\partial_a g_{0,j}(\cdot,\theta_0)\right)\right)\partial_{\theta_l}\rho_j\left(a,\theta_0,\partial_a g_{0,j}(\cdot,\theta_0)\right)\end{array}\right| \\
&+\left|\begin{array}{c} q\left(\rho_j\left(a,\overline{\theta}_j,\partial_a\widehat{g}_j(\cdot,\overline{\theta}_j)\right)\right)D_{\partial_a g}\rho_j\left(a,\overline{\theta}_j,\partial_a\widehat{g}_j(\cdot,\overline{\theta})\right)\left[\partial_{\theta_l}\partial_a\widehat{g}_j(\cdot,\overline{\theta}_j)\right] \\ -q\left(\rho_j\left(a,\theta_0,\partial_a g_{0,j}(\cdot,\theta_0)\right)\right)D_{\partial_a g}\rho_j\left(a,\theta_0,\partial_a g_{0,j}(\cdot,\theta_0)\right)\left[\partial_{\theta_l}\partial_a g_{0,j}(\cdot,\theta_0)\right]\end{array}\right|.
\end{aligned}$$

First note that the terms on the RHS are uniformly bounded, it is easy to see that the terms on the RHS of the inequality are $o(1)$ as $\left\|\overline{\theta}_j - \theta_0\right\| \to 0$ since $\left(\overline{\theta}_j, \partial_a \widehat{g}_j\left(\cdot, \overline{\theta}_j\right)\right) \xrightarrow{p} \left(\theta_0, \partial_a g_{0,j}\left(\cdot, \theta_0\right)\right)$ by Lemma 2.5 and continuity in $\theta$ of $\partial_a g_{0,j}(\cdot, \theta)$. Then it will follow by DCT that

$$\left[ (\theta-\theta_0)' \int \left[ \begin{array}{c} \left[D_\theta F_{A|X}\left(\overline{\theta}_j, \partial_a \widehat{g}_j\left(\cdot, \overline{\theta}_j\right)\right) - D_\theta F_{A|X}\left(\theta_0, \partial_a g_{0,j}\left(\cdot, \theta_0\right)\right)\right] \\ \times \left[D_\theta F_{A|X}\left(\overline{\theta}_j, \partial_a \widehat{g}_j\left(\cdot, \overline{\theta}_j\right)\right) - D_\theta F_{A|X}\left(\theta_0, \partial_a g_{0,j}\left(\cdot, \theta_0\right)\right)\right]' \end{array} \right] d\mu \left(\theta - \theta_0\right) \right]^{1/2} = o_p\left(\|\right.$$

From Lemma 2.9 and CMT, $\left[\int \left[\widehat{F}_{A|X=j} - F_{A|X=j}\right]^2 d\mu\right]^{1/2} = O_p\left(N^{-1/2}\right)$. Since we have finite $j$ then $|B_{12}| = o_p\left(N^{-1/2}\left\|\theta - \theta_0\right\|\right)$.

For $B_2$ :

For each $j$

$$\begin{aligned} & F_{A|X=j}\left(\theta, \partial_a \widehat{g}_j(\cdot, \theta)\right) + F_{A|X=j}\left(\theta_0, \partial_a \widehat{g}_j\left(\cdot, \theta_0\right)\right) - 2\widehat{F}_{A|X=j} \\ = & \left(F_{A|X=j}\left(\theta, \partial_a \widehat{g}_j(\cdot, \theta)\right) - F_{A|X=j}\left(\theta_0, \partial_a \widehat{g}_j\left(\cdot, \theta_0\right)\right)\right) \\ & +2\left(F_{A|X=j}\left(\theta_0, \partial_a \widehat{g}_j\left(\cdot, \theta_0\right)\right) - F_{A|X}\left(\theta_0, \partial_a g_{0,j}\left(\cdot, \theta_0\right)\right)\right) \\ & -2\left(\widehat{F}_{A|X=j} - F_{A|X=j}\right), \end{aligned}$$

then we can write $B_2$ as

$$\begin{aligned} B_2 = & R^{-1/2}\sum_{j=1}^{J}\int \left[ \begin{array}{c} \left[\nu_{R,j}\left(\theta, \partial_a \widehat{g}_j(\cdot, \theta)\right) - \nu_{R,j}\left(\theta_0, \partial_a \widehat{g}_j\left(\cdot, \theta_0\right)\right)\right] \\ \times \left[F_{A|X=j}\left(\theta, \partial_a \widehat{g}_j(\cdot, \theta)\right) - F_{A|X=j}\left(\theta_0, \partial_a \widehat{g}_j\left(\cdot, \theta_0\right)\right)\right] \end{array} \right] d\mu_j \\ & +2R^{-1/2}\sum_{j=1}^{J}\int \left[ \begin{array}{c} \left[\nu_{R,j}\left(\theta, \partial_a \widehat{g}_j(\cdot, \theta)\right) - \nu_{R,j}\left(\theta_0, \partial_a \widehat{g}_j\left(\cdot, \theta_0\right)\right)\right] \\ \times \left[F_{A|X=j}\left(\theta_0, \partial_a \widehat{g}_j\left(\cdot, \theta_0\right)\right) - F_{A|X}\left(\theta_0, \partial_a g_{0,j}\left(\cdot, \theta_0\right)\right)\right] \end{array} \right] d\mu_j \\ & -2R^{-1/2}\sum_{j=1}^{J}\int \left(\nu_{R,j}\left(\theta, \partial_a \widehat{g}_j(\cdot, \theta)\right) - \nu_{R,j}\left(\theta_0, \partial_a \widehat{g}_j\left(\cdot, \theta_0\right)\right)\right)\left(\widehat{F}_{A|X=j} - F_{A|X=j}\right) d\mu_j \\ = & B_{21} + B_{22} + B_{23}. \end{aligned}$$

We first show $\int [\nu_{R,j}(\theta,\partial_a\widehat{g}_j(\cdot,\theta)) - \nu_{R,j}(\theta_0,\partial_a\widehat{g}_j(\cdot,\theta_0))]^2 d\mu_j = o_p(1)$ for any $j$. By Lemma 2.6 $\partial_a\widehat{g}_j \in \mathcal{G}_j^{(1)}$ w.p.a. 1, and by Lemma 2.8 it suffices to show that $\|\partial_a\widehat{g}_j(\cdot,\theta) - \partial_a\widehat{g}_j(\cdot,\theta_0)\|$ 0 as $\|\theta-\theta_0\|$. This follows from the triangle inequality since $\|\partial_a\widehat{g}_j(\cdot,\theta) - \partial_a\widehat{g}_j(\cdot,\theta_0)\|$ is bounded above by $\|\partial_a\widehat{g}_j(\cdot,\theta) - \partial_a g_{0,j}(\cdot,\theta)\| + \|\partial_a\widehat{g}_j(\cdot,\theta_0) - \partial_a g_{0,j}(\cdot,\theta_0)\| + \|\partial_a g_{0,j}(\cdot,\theta) - \partial_a g_{0,j}(\cdot,$ and the fact that the first two terms of the majorant converge to zero by Lemma 2.5 and the last term converges to zero by the continuity of $\partial_a g_{0,j}(\cdot,\theta)$ in $\theta$. For $B_{21}$

$$\begin{aligned}
B_{21} &= 2R^{-1/2}\sum_{j=1}^{J}\int\left[\begin{array}{c}[\nu_{R,j}(\theta,\partial_a\widehat{g}_j(\cdot,\theta)) - \nu_{R,j}(\theta_0,\partial_a\widehat{g}_j(\cdot,\theta_0))] \\ \times\left[F_{A|X=j}(\theta,\partial_a\widehat{g}_j(\cdot,\theta)) - F_{A|X=j}(\theta_0,\partial_a\widehat{g}_j(\cdot,\theta_0))\right]\end{array}\right]d\mu \\
&= 2R^{-1/2}\sum_{j=1}^{J}\int\left[\begin{array}{c}[\nu_{R,j}(\theta,\partial_a\widehat{g}_j(\cdot,\theta)) - \nu_{R,j}(\theta_0,\partial_a\widehat{g}_j(\cdot,\theta_0))] \\ \times D_\theta F_{A|X=j}\left(\overline{\theta}_j,\partial_a\widehat{g}_j\left(\cdot,\overline{\theta}_j\right)\right)'\end{array}\right]d\mu\,(\theta-\theta_0),
\end{aligned}$$

by Cauchy Schwarz inequality

$$\begin{aligned}
|B_{21}| &\leq o_p\left(R^{-1/2}\right)\times \\
&\quad \max_{1\leq j\leq J}\left[(\theta-\theta_0)'\int\left[D_\theta F_{A|X=j}\left(\overline{\theta}_j,\partial_a\widehat{g}_j\left(\cdot,\overline{\theta}_j\right)\right)D_\theta F_{A|X=j}\left(\overline{\theta}_j,\partial_a\widehat{g}_j\left(\cdot,\overline{\theta}_j\right)\right)'\right]d\mu\,(\theta-\theta_0)\right] \\
&= o_p\left(R^{-1/2}\right)O_p\left(\|\theta-\theta_0\|\right) \\
&= o_p\left(N^{-1/2}\|\theta-\theta_0\|\right),
\end{aligned}$$

the first inequality follows from the stochastic equicontinuity condition of Lemma 2.8, then it is easy to show the outer product term inside the integral is also bounded in probability and the last equality follows from $N = o(R)$. This same argument using Cauchy Schwarz inequality again be applied for $B_{22}$ and $B_{23}$, in particular, it follows from Lemma 2.10 and Lemma 2.9 respectively that $|B_{22}| = o\left(N^{-1}\right)$ and $|B_{23}| = o\left(N^{-1}\right)$.

For $B_3$ :

For each $j$

$$
\begin{aligned}
\nu_{R,j}\left(\theta, \partial_a \widehat{g}_j\left(\cdot, \theta\right)\right) + \nu_{R,j}\left(\theta_0, \partial_a \widehat{g}_j\left(\cdot, \theta_0\right)\right) \quad = \quad & 2\nu_R\left(\theta_0, \partial_a g_{0,j}\left(\cdot, \theta_0\right)\right) \\
& + \left(\nu_{R,j}\left(\theta, \partial_a \widehat{g}_j\left(\cdot, \theta\right)\right) - \nu_R\left(\theta_0, \partial_a g_{0,j}\left(\cdot, \theta_0\right)\right)\right) \\
& + \left(\nu_{R,j}\left(\theta_0, \partial_a \widehat{g}_j\left(\cdot, \theta_0\right)\right) - \nu_R\left(\theta_0, \partial_a g_{0,j}\left(\cdot, \theta_0\right)\right)\right),
\end{aligned}
$$

then we can write $B_3$ as

$$
\begin{aligned}
B_3 \quad = \quad & 2S^{-1/2} \sum_{j=1}^{J} \int \nu_{R,j}\left(\theta_0, \partial_a g_{0,j}\left(\cdot, \theta_0\right)\right) \left(F_{A|X=j}\left(\theta, \partial_a \widehat{g}_j\left(\cdot, \theta\right)\right) - F_{A|X=j}\left(\theta_0, \partial_a \widehat{g}_j\left(\cdot, \theta_0\right)\right)\right) d\mu_j \\
& + R^{-1/2} \sum_{j=1}^{J} \int \left[ \begin{array}{c} \left[\nu_{R,j}\left(\theta, \partial_a \widehat{g}_j\left(\cdot, \theta\right)\right) - \nu_R\left(\theta_0, \partial_a g_{0,j}\left(\cdot, \theta_0\right)\right)\right] \\ \times \left[F_{A|X=j}\left(\theta, \partial_a \widehat{g}_j\left(\cdot, \theta\right)\right) - F_{A|X=j}\left(\theta_0, \partial_a \widehat{g}_j\left(\cdot, \theta_0\right)\right)\right] \end{array} \right] d\mu_j \\
& + R^{-1/2} \sum_{j=1}^{J} \int \left[ \begin{array}{c} \left[\nu_{R,j}\left(\theta_0, \partial_a \widehat{g}_j\left(\cdot, \theta_0\right)\right) - \nu_R\left(\theta_0, \partial_a g_{0,j}\left(\cdot, \theta_0\right)\right)\right] \\ \times \left[F_{A|X=j}\left(\theta, \partial_a \widehat{g}_j\left(\cdot, \theta\right)\right) - F_{A|X=j}\left(\theta_0, \partial_a \widehat{g}_j\left(\cdot, \theta_0\right)\right)\right] \end{array} \right] d\mu_j \\
= \quad & B_{31} + B_{32} + B_{33}.
\end{aligned}
$$

For each $j$: we have $\left[\int \left[F_{A|X=j}\left(\theta, \partial_a \widehat{g}_j\left(\cdot, \theta\right)\right) - F_{A|X=j}\left(\theta_0, \partial_a \widehat{g}_j\left(\cdot, \theta_0\right)\right)\right]^2 d\mu_j\right]^{1/2} =$ $O_p\left(\|\theta - \theta_0\|\right)$ by Cauchy Schwarz inequality; from Donsker theorem and CMT, $\left[\int \left[\nu_R\left(\theta_0, \partial_a g_{0,j}\left(\cdot, \theta\right.\right.\right.\right.$ $O_p\left(1\right)$. Then it follows that $\left|B_{31}\right| \leq o_p\left(N^{-1/2}\|\theta - \theta_0\|\right)$. By a similar argument, using Cauchy Schwarz inequality, continuity of $\partial_a g\left(\cdot, \theta\right)$ in $\theta$, Lemmas 2.5, 2.6 and 2.8, $\left|B_{32}\right|$ and $\left|B_{33}\right|$ are also $o_p\left(N^{-1/2}\|\theta - \theta_0\|\right)$, in particular as we can use the triangle inequality to show $\left\|\left(\theta, \partial_a \widehat{g}_j\left(\cdot, \theta\right)\right) - \left(\theta_0, \partial_a g_{0,j}\left(\cdot, \theta_0\right)\right)\right\|_\nu$ and $\left\|\left(\theta_0, \partial_a \widehat{g}_j\left(\cdot, \theta_0\right)\right) - \left(\theta_0, \partial_a g_{0,j}\left(\cdot, \theta_0\right)\right)\right\|_\nu$ converge in probability to 0 as $\|\theta - \theta_0\| \to 0$ for all $j$.

For $B_4$ :

By the same argument above, we can re-express $B_4$

$$\begin{aligned}
B_4 &= 2S^{-1}\sum_{j=1}^{J}\int \nu_{R,j}\left(\theta_0,\partial_a g_{0,j}(\cdot,\theta_0)\right)\left(\nu_{R,j}\left(\theta,\partial_a\widehat{g}_j(\cdot,\theta)\right)-\nu_{R,j}\left(\theta_0,\partial_a\widehat{g}_j(\cdot,\theta_0)\right)\right)d\mu_j \\
&\quad +R^{-1}\sum_{j=1}^{J}\int\left[\nu_{R,j}\left(\theta,\partial_a\widehat{g}_j(\cdot,\theta)\right)-\nu_{R,j}\left(\theta_0,\partial_a\widehat{g}_j(\cdot,\theta_0)\right)\right]^2 d\mu_j \\
&\quad +R^{-1}\sum_{j=1}^{J}\int\left[\begin{array}{l}\left[\nu_{R,j}\left(\theta_0,\partial_a\widehat{g}_j(\cdot,\theta_0)\right)-\nu_R\left(\theta_0,\partial_a g_{0,j}(\cdot,\theta_0)\right)\right] \\ \times\left[\nu_{R,j}\left(\theta,\partial_a\widehat{g}_j(\cdot,\theta)\right)-\nu_{R,j}\left(\theta_0,\partial_a\widehat{g}_j(\cdot,\theta_0)\right)\right]\end{array}\right]d\mu_j \\
&= B_{41}+B_{42}+B_{43}.
\end{aligned}$$

By repeatedly using Cauchy Schwarz inequality, continuity of $\partial_a g(\cdot,\theta)$ in $\theta$, and Lemmas 2.5, 2.6 and 2.8, as seen in the analysis of $B_2$ and $B_3$, it follows easily that $|B_{4i}| = o_p\left(N^{-1}\right)$ for $i = 1,2,3$.

G6 then follows from Lemma 2.10.■

Proof of Theorem 2.4. From (88) we have

$$\begin{aligned}
\widehat{g}_{\widehat{\theta}}-\widehat{g}_{\theta_0} &= \widehat{\mathcal{H}}\left(I-\widehat{\mathcal{L}}\right)^{-1}\left(\widetilde{r}_{\widehat{\theta}}-\widetilde{r}_{\theta_0}\right) \\
&= \widehat{\mathcal{H}}\left(I-\widehat{\mathcal{L}}\right)^{-1}\left(\left(\widehat{\theta}-\theta_0\right)' D_\theta\widetilde{r}_{\overline{\theta}}\right)
\end{aligned}$$

where the expansion above follows from MVT and $\overline{\theta}$ denotes some intermediate value between $\widehat{\theta}$ and $\theta_0$. It is easy to see that, for $j = 1,\ldots,J$

$$\begin{aligned}
\left\|\widehat{g}_j\left(\cdot,\widehat{\theta}\right)-\widehat{g}_j(\cdot,\theta_0)\right\|_\infty &= O_p\left(\left\|\widehat{\theta}-\theta_0\right\|\right) \\
&= O_p\left(N^{-1/2}\right),
\end{aligned}$$

since $\left\|\widehat{\mathcal{H}}\left(I-\widehat{\mathcal{L}}\right)^{-1}\right\| = O_p(1)$, $\|\widetilde{r}_\theta\| = O_p(1)$ and $\sqrt{Nh} = o\left(N^{1/2}\right)$, then $\sqrt{Nh}\left|\widehat{g}_j\left(a,\widehat{\theta}\right)-\widehat{g}_j(a,\theta_0\right.$ $= o_p(1)$. It remains to show the asymptotic independence between any pair $\left(\widehat{g}_j\left(a,\widehat{\theta}\right),\widehat{g}_k\left(a',\widehat{\theta}\right)\right)$

for any $k \neq j$ and $a' \neq a$. Since

$$\begin{aligned}
& \operatorname{cov}\left(\widehat{g}_j\left(a,\widehat{\theta}\right),\widehat{g}_k\left(a',\widehat{\theta}\right)\right) \\
= & \operatorname{cov}\left(\widehat{g}_j\left(a,\theta_0\right),\widehat{g}_k\left(a',\theta_0\right)\right)+\operatorname{cov}\left(\widehat{g}_j\left(a,\theta_0\right),\widehat{g}_k\left(a',\widehat{\theta}\right)-\widehat{g}_k\left(a',\theta_0\right)\right) \\
& +\operatorname{cov}\left(\widehat{g}_k\left(a',\theta_0\right),\widehat{g}_j\left(a,\widehat{\theta}\right)-\widehat{g}_j\left(a,\theta_0\right)\right)+\operatorname{cov}\left(\widehat{g}_j\left(a,\widehat{\theta}\right)-\widehat{g}_j\left(a,\theta_0\right),\widehat{g}_k\left(a',\widehat{\theta}\right)-\widehat{g}_k\left(a',\theta_0\right)\right)
\end{aligned}$$

by Cauchy-Schwarz inequality, it suffices to show $\operatorname{var}\left(\sqrt{Nh}\left(\widehat{g}_k\left(a',\widehat{\theta}\right)-\widehat{g}_k\left(a',\theta_0\right)\right)\right)=o(1)$; this follows since $\left\|\widehat{g}_j\left(\cdot,\widehat{\theta}\right)-\widehat{g}_j\left(\cdot,\theta_0\right)\right\|_\infty=O_p\left(N^{-1/2}\right)$.■

### 2.9.3 Proofs of Lemmas 2.1 - 2.10

These lemmas are used in the proofs of Theorem 2.1 - 2.3. In what follows we let: $\xi>0$ be a number that is arbitrarily close to 0; $C_0$ denotes a positive constant that may take different values in various places; VW abbreviates van der Vaart and Wellner (1996).

PROOF OF LEMMA 2.1. We can write, for any $1\leq k,j\leq J$

$$\widehat{p}_{X'|X}\left(k|j\right)-p_{X'|X}\left(k|j\right)=\frac{\widehat{p}_{X',X}\left(k,j\right)-p_{X',X}\left(k,j\right)}{p_X\left(j\right)}-\frac{\widehat{p}_{X'|X}\left(k|j\right)}{p_X\left(j\right)\widehat{p}_X\left(j\right)}\left(\widehat{p}_X\left(j\right)-p_X\left(j\right)\right).$$

Given the simple nature of our DGP, by standard CLT and LLN, we have $\widehat{p}_{X',X}\left(k,j\right)-p_{X',X}\left(k,j\right)=O_p\left(N^{-1/2}\right),\widehat{p}_X\left(j\right)-p_X\left(j\right)=O_p\left(N^{-1/2}\right)$ and $\widehat{p}_X\left(j\right)^{-1}=O_p\left(1\right)$, so it follows that $\widehat{p}_{X'|X}\left(k|j\right)-p_{X'|X}\left(k|j\right)=O_p\left(N^{-1/2}\right)$ for any $k$ and $j$. Since $\mathcal{L}$ is a linear map on $\mathbb{R}^J$ to $\mathbb{R}^J$, for any vector $m\in\mathbb{R}^J$ we have $\left(\left(\widehat{\mathcal{L}}-\mathcal{L}\right)m\right)_j=\beta\sum_{k=1}^J(\widehat{p}\left(k|j\right)-p(k|j))m_i=O_p\left(N^{-1/2}\right)$ for all $j$ then it follows from the definition of an operator norm that $\left\|\widehat{\mathcal{L}}-\mathcal{L}\right\|=O_p\left(N^{-1/2}\right)$.■

PROOF OF LEMMA 2.2. For any $j=1,\ldots,J$ and $\theta\in\Theta$, $\widetilde{r}_\theta\left(j\right)$ is defined in (86) with $w_{itN}\left(j\right)=\mathbf{1}\left[x_{it}=j\right]/\widehat{p}_X\left(j\right)$ and define $\widehat{r}_\theta\left(j\right)=\sum_{i=1,t=1}^{N,T}w_{itN}\left(j\right)u_\theta(a_{it},x_{it},\varepsilon_{it})$.

Then we write

$$\widetilde{r}_{\theta}(j)-r_{\theta}(j)=\left(\widehat{r}_{\theta}(j)-r_{\theta}(j)\right)+\left(\widetilde{r}_{\theta}(j)-\widehat{r}_{\theta}(j)\right), \tag{101}$$

the first term is the usual term had we observed $\{\varepsilon_{it}\}$, the latter term arises due to the use of generated residuals. Treating them separately, for the first term

$$\begin{aligned}
\widehat{r}_{\theta}(j)-r_{\theta}(j) &= \frac{1}{\widehat{p}_X(j)}\frac{1}{NT}\sum_{i=1,t=1}^{N,T}\mathbf{1}\left[x_{it}=j\right]\left(u_{\theta}\left(a_{it},x_{it},\varepsilon_{it}\right)-r_{\theta}(j)\right)\\
&= \frac{1}{\widehat{p}_X(j)}\frac{1}{NT}\sum_{i=1,t=1}^{N,T}v_{\theta,it}\mathbf{1}\left[x_{it}=j\right],
\end{aligned}$$

where for each $\theta$, $v_{\theta,it}=u_{\theta}\left(a_{it},x_{it},\varepsilon_{it}\right)-r_{\theta}\left(x_{it}\right)$ is a zero mean random variable, note that $\mathbf{1}\left[x_{it}=j\right]\times\left(r_{\theta}\left(x_{it}\right)-r_{\theta}(j)\right)=0$ for all $i,j$ and $t$. Define $\Upsilon_{N,j}(\theta)$ as the sample average of i.i.d. random variables $\left\{\sum_{t=1}^{T}\frac{1}{T}v_{\theta,it}\mathbf{1}\left[x_{it}=j\right]\right\}_{i=1}^{N}$, given the assumptions on the DGP, in particular on the second moments, $\Upsilon_{N,j}(\theta)=O_p\left(N^{-1/2}\right)$ for any $\theta$ by standard CLT. We want to obtain the uniform rate of convergence of $\Upsilon_{N,j}(\theta)$ over $\Theta$. This can be achieved by using the arguments along the line of Masry (1996). We first obtain the uniform bound for the variance of $\Upsilon_{N,j}(\theta)$, some exponential inequality is then applied to get the rate of decay on the tail probability for any $\theta$. The pointwise rate can then be made uniform by Lipschitz continuity of $v_{\theta,it}$ (in $\theta$) and compactness of $\Theta$. More precisely, we first show that $\sup_{\theta\in\Theta}\text{var}(\Upsilon_{N,j}(\theta))=O\left(N^{-1}\right)$. Since $\text{var}(\Upsilon_{N,j}(\theta))$ is just a variance of $\sum_{t=1}^{T}\frac{1}{T}v_{\theta,it}\mathbf{1}\left[x_{it}=j\right]$ by divided by $N$, the numerator takes the following form

$$\begin{aligned}
\text{var}\left(\frac{1}{T}\sum_{t=1}^{T}v_{\theta,it}\mathbf{1}\left[x_{it}=j\right]\right) &= \frac{1}{T}\sum_{t=1}^{T}\text{var}\left(v_{\theta,it}\mathbf{1}\left[x_{it}=j\right]\right)\\
&\quad+\frac{2}{T}\sum_{s=1}^{T-1}\left(1-\frac{s}{T}\right)Cov\left(v_{\theta,i0}\mathbf{1}\left[x_{i0}=j\right],v_{\theta,is}\mathbf{1}\left[x_{is}=j\right]\right),\\
&= Y_{\theta,1,j}+Y_{\theta,2,j}.
\end{aligned}$$

The covariance structure in $Y_{\theta,2,j}$ follows from the strict stationarity assumption, which also implies we can write $Y_{\theta,1,j} = E\left[\left[v_{\theta,it}^2|x_{it}\right]\mathbf{1}\left[x_{it}=j\right]\right]$. Since $u_\theta\left(a,x,\varepsilon\right)$ is continuous in $\theta$ for all $a, x$ and $\varepsilon$, it follows that $\sup_{\theta\in\Theta} Y_{\theta,1,j} < \infty$. For the covariance term, by Cauchy-Schwarz inequality, $Cov\left(v_{\theta,i0}\mathbf{1}\left[x_{i0}=j\right], v_{\theta,is}\mathbf{1}\left[x_{is}=j\right]\right) \leq E\left[v_{\theta,i0}^2\mathbf{1}\left[x_{i0}=j\right]\right]^2 < \infty$, since $\sup_{\theta\in\Theta}\left|v_{\theta,i0,j}^2\right| < \infty$, it follows that $\sup_\theta Y_{\theta,2} < \infty$ for any finite $T$. Since $\Upsilon_{N,j}(\theta)$ is an average of $N-$i.i.d. sequence of random variables that, for each $\theta$, it satisfies the Cramér' conditions (since $u$ is uniformly bounded over all its arguments), then Bernstein's inequality, e.g. see Bosq (1998), can be used to obtain the following bound

$$\Pr\left[\left|N\Upsilon_{N,j}(\theta)\right| > N\delta_N\right] \leq 2\exp\left\{-\frac{N^2\delta_N^2}{4Var\left(N\Upsilon_{N,j}(\theta)\right)+2CN\delta_N}\right\}. \qquad (102)$$

Let $\delta_N = N^{(-1+\xi)/2}$, simple calculation of the display above yields $\Pr\left[\left|\Upsilon_{N,j}(\theta)\right| > \delta_N\right] = O\left(\exp\left(-N^\xi\right)\right)$. By compactness of $\Theta$, let $\{L_N\}_{N=1}^\infty$ be an increasing sequence of natural number, we can define a sequence $\{\theta_{iL_N}\}_{i=1}^{L_N}$ to be the centres of open balls, $\{\Theta_{iL_N}\}_{i=1}^{L_N}$, of radius $\{\epsilon_{L_N}\}_{i=1}^{L_N}$ such that $\Theta \subset \bigcup_{i=1}^{L_N}\Theta_{iL_N}$ and $L_N \times \epsilon_{L_N} = O(1)$, then it follows that

$$\begin{aligned}
\Pr\left[\sup_\theta\left|\Upsilon_{N,j}(\theta)\right| > \delta_N\right] &\leq \Pr\left[\max_{1\leq i\leq L_N}\left|\Upsilon_{N,j}(\theta_{iL_N})\right| > \delta_N\right] \\
&\quad +\Pr\left[\max_{1\leq i\leq N}\sup_{\theta\in\Theta_{iL_N}}\left|\Upsilon_{N,j}(\theta)-\Upsilon_{N,j}(\theta_{iL_N})\right| > \delta_N\right] \\
&\leq C_0 L_N\exp\left(-N^\xi\right)+\Pr\left[\epsilon_{L_N} > \delta_N\right] \\
&= o(1).
\end{aligned}$$

The second inequality from the display above follows from, Bonferroni inequality and (102) for the first term, and by Lipschitz continuity of $\Upsilon_{N,j}$ for the latter. Then the equality holds if we take $\epsilon_{L_N} = o(\delta_N)$ such that $L_N$ grows at some power rate. It then follows that that $\sup_\theta |\Upsilon_{N,j}(\theta)| = o_p(N^{-\lambda})$. Then w.p.a. 1

$$\begin{aligned}\sup_{\theta\in\Theta} |\widehat{r}_\theta(j) - r_\theta(j)| &\leq \frac{\max_{1\leq j\leq J}\sup_{\theta\in\Theta}|\Upsilon_{N,j}(\theta)|}{\min_{1\leq j\leq J} p_X(j)} \\ &= o_p\left(N^{-\lambda}\right).\end{aligned}$$

The procedure to obtain the uniform rate of convergence is shown above in detail to avoid repetition later since we will require to show many zero mean processes converge uniformly (either over the compact parameter space or the state space) to zero faster than some rates. The argument above can also be applied to nonparametric estimates, as well as some other appropriately (weakly) dependent zero mean process, see Linton and Mammen (2005), and especially Srisuma and Linton (2009) for such usages in closely related context. We comment here that, our paper along with the papers mentioned in the previous sentence, unlike Masry (1996), are not interested in sharp rate of uniform convergence so our proofs are comparatively more straightforward.

For the generated residuals, by definition

$$\widetilde{r}_\theta(j) - \widehat{r}_\theta(j) = \frac{1}{NT}\sum_{i=1,t=1}^{N,T} w_{itN}(j)\left(u_\theta(a_{it}, x_{it}, \widehat{\varepsilon}_{it}) - u_\theta(a_{it}, x_{it}, \varepsilon_{it})\right),$$

where $\widehat{\varepsilon}_{it} = \chi\left(\widehat{F}_{A|X}(a_{it}|x_{it})\right)$ with $\chi \equiv Q_\varepsilon^{-1}$. Using mean value expansion, $u_\theta(a_{it}, x_{it}, \widehat{\varepsilon}_{it}) - u_\theta(a_{it}, x_{it}, \varepsilon_{it}) = \frac{\partial}{\partial\varepsilon} u_\theta(a_{it}, x_{it}, \overline{\varepsilon}_{it})\, \chi'\left(\overline{F}_{A|X}(a_{it}|x_{it})\right)\left(\widehat{F}_{A|X}(a_{it}|x_{it}) - F_{A|X}(a_{it}|x_{it})\right)$, where $\overline{\varepsilon}_{it}$ and $\overline{F}_{A|X}(a_{it}|x_{it})$ are some intermediate points between $\widehat{\varepsilon}_{it}$ and $\varepsilon_{it}$, and, $\widehat{F}_{A|X}(a_{it}|x_{it})$ and $F_{A|X}(a_{it}|x_{it})$, respectively. Then it follows that

$$\begin{aligned}\widetilde{r}_\theta(j) - \widehat{r}_\theta(j) &= \frac{1}{NT}\sum_{i=1,t=1}^{N,T} w_{itN}(j)\left(u_\theta(a_{it}, x_{it}, \widehat{\varepsilon}_{it}) - u_\theta(a_{it}, x_{it}, \varepsilon_{it})\right) \\ &= \frac{1}{NT}\sum_{i=1,t=1}^{N,T} \frac{\mathbf{1}[x_{it} = j]}{p_X(j)} \varkappa_\theta(a_{it}, x_{it}, \varepsilon_{it})\left(\widehat{F}_{A|X}(a_{it}|x_{it}) - F_{A|X}(a_{it}|x_{it})\right) + O_p\left(N^{-}\right.\end{aligned}$$

where $\varkappa_\theta(a_{it}, x_{it}, \varepsilon_{it}) = \frac{\partial}{\partial \varepsilon} u_\theta(a_{it}, x_{it}, \varepsilon_{it}) \chi'(F_{A|X}(a_{it}|x_{it}))$. In addition, the $O_p(N^{-1})$ −term holds uniformly over $\theta$ and $j$, this follows from Markov inequality since $\frac{\partial^2}{\partial \varepsilon^2} u$ and $\chi''$ are uniformly bounded over all of their arguments, $\max_{1 \leq j \leq J} |\widehat{p}_X(j) - p_X(j)| = O_p(N^{-1/2})$, and, $\max_{1 \leq j \leq J} \sup_{a \in A} |\widehat{F}_{A|X}(a|j) - F_{A|X}(a|j)| = O_p(N^{-1/2})$ by Lemma 2.9. By a similar argument, using the leave one out estimator for $\widehat{F}_{A|X}$, the leading term for $\widetilde{r}_\theta(j) - \widehat{r}_\theta(j)$ can be simplified further to

$$\frac{1}{NT(NT-1)} \sum_{i=1,t=1}^{N,T} \sum_{j,s,(-it)}^{N,T} \varkappa_\theta(a_{it}, x_{it}, \varepsilon_{it}) \frac{\mathbf{1}[x_{it} = j]}{p_X(j)} \frac{\mathbf{1}[x_{js} = x_{it}]}{p(x_{it})} \left(\mathbf{1}[a_{js} \leq a_{it}] - F_{A|X}(a_{it}|x_{it})\right),$$

where $\sum_{j,s,(-it)}^{N,T}$ sums over the indices $j = 1, \ldots, N$ and $s = 1, \ldots, T$ but omits the $it^{th}$−summand. Subsequently, the term in the display above can be written as the following second order U-statistic

$$\begin{pmatrix} NT \\ 2 \end{pmatrix}^{-1} \frac{1}{2} \sum_{C((it),(js))} \begin{pmatrix} \varkappa_\theta(a_{it}, x_{it}, \varepsilon_{it}) \frac{\mathbf{1}[x_{it}=j]}{p_X(j)} \frac{\mathbf{1}[x_{js}=x_{it}]}{p(x_{it})} \left(\mathbf{1}[a_{js} \leq a_{it}] - F_{A|X}(a_{it}|x_{it})\right) \\ +\varkappa_\theta(a_{js}, x_{js}, \varepsilon_{js}) \frac{\mathbf{1}[x_{js}=j]}{p_X(j)} \frac{\mathbf{1}[x_{js}=x_{it}]}{p(x_{js})} \left(\mathbf{1}[a_{it} \leq a_{js}] - F_{A|X}(a_{js}|x_{js})\right) \end{pmatrix},$$

where $\sum_{C((it),(js))}$ sums over all distinct combinations of $C((it),(js))$. Note that $\mathbf{1}[a_{it} \leq a] = F_{A|X}(a|x_{it}) + \omega(x_{it}; a)$ where $E[\omega(x_{it}; a)|x_{it}] = 0$, so $\omega(x_{it}; \cdot)$ is a random element in $L^2(A)$. Then it is straightforward to obtain the leading term of the Hoeffding decomposition of our U-statistic, see Lee (1990), and, Powell, Stock and Stoker (1989), in particular we have for all $j$

$$\widetilde{r}_\theta(j) - \widehat{r}_\theta(j) = \frac{1}{NT} \sum_{i=1,t=1}^{N,T} \zeta_\theta(\omega(x_{it}; \cdot), x_{it}; j) + o_p\left(N^{-1/2}\right),$$

where $\zeta_\theta(\omega(x_{it}; \cdot), x_{it}; j) = \frac{1}{p_X(j)} \int \omega(x_{it}; a_{js}) \left[\int \varkappa_\theta(a_{js}, x_{it}, \varepsilon_{js}) \mathbf{1}[x_{it} = j] \frac{f_{A,X,\mathcal{E}}(a_{js}, x_{it}, \varepsilon_{js})}{p(x_{it})} d\varepsilon_{js}\right] da$ and $f_{A,X,\mathcal{E}}$ denotes the joint continuous-discrete density of $(a_{it}, x_{it}, \varepsilon_{it})$. Note that $\zeta_\theta$ is random with respect to $\omega_{it}$ and $x_{it}$, and $E[\omega(x_{it}; \cdot)|x_{it}] = 0$, so $\zeta_\theta$ has zero mean. Given the boundedness and smoothness conditions on $\varkappa_\theta$, then $\frac{1}{NT}\sum_{i=1,t=1}^{N,T}\zeta_\theta\left(\omega\left(x_{it};\cdot\right),x_{it};j\right)$ can be shown to converge uniformly in probability to zero faster than the rate $N^{-\lambda}$ as shown above. In sum, we have shown for $j=1,\ldots,J$ that $\widetilde{r}_\theta\left(j\right)=r_\theta\left(j\right)+\widetilde{r}_\theta^R\left(j\right)$ with

$$\begin{aligned}\widetilde{r}_\theta^R\left(j\right) &= \frac{1}{p_X\left(j\right)}\frac{1}{NT}\sum_{i=1,t=1}^{N,T}1\left[x_{it}=j\right]\left(u_\theta\left(a_{it},x_{it},\varepsilon_{it}\right)-r_\theta\left(j\right)\right)\\ &\quad+\frac{1}{NT}\sum_{i=1,t=1}^{N,T}\zeta_\theta\left(\omega\left(x_{it};\cdot\right),x_{it};j\right)+o_p\left(N^{-\lambda}\right)\\ &= o_p\left(N^{-\lambda}\right),\end{aligned}$$

where the smaller order term holds uniformly over $j$ and $\theta$.■

PROOF OF LEMMA 2.3. Since $0<\left\|\mathcal{L}\right\|<1$ and $0<\left\|\widehat{\mathcal{L}}\right\|<1$, the argument used in Linton and Mammen (2005) can be used to show

$$\left\|\left(I-\widehat{\mathcal{L}}\right)^{-1}-\left(I-\mathcal{L}\right)^{-1}\right\|=O_p\left(N^{-1/2}\right).$$

We note that, using the contraction property, $\left(I-\mathcal{L}\right)^{-1}$ and $\left(I-\widehat{\mathcal{L}}\right)^{-1}$ are bounded linear operators since $\left\|\left(I-\mathcal{L}\right)^{-1}\right\|\leq\left(1-\left\|\mathcal{L}\right\|\right)^{-1}<\infty$ and similarly $\left\|\left(I-\widehat{\mathcal{L}}\right)^{-1}\right\|\leq\left(1-\left\|\widehat{\mathcal{L}}\right\|\right)^{-1}<\infty$, this can be shown from the respective Neumann series representation of the inverses and by the basic properties of operator norms. We comment that these relations involving the empirical operator hold in finite sample since $X$ is finite, otherwise it will be true w.p.a. 1 by the same reasoning as used in Srisuma and Linton (2009). Then for each $x\in X$ and $\theta\in\Theta$, $\widehat{m}_\theta\left(j\right)$ is defined in (87), we write $\widehat{m}_\theta\left(j\right)=\left(I-\widehat{\mathcal{L}}\right)^{-1}\left(r_\theta\left(j\right)+\widetilde{r}_\theta^R\left(j\right)\right)$, given the results from Lemma 2.2, it follows that $\max_{1\leq j\leq J}\sup_{\theta\in\Theta}\left|\left(I-\widehat{\mathcal{L}}\right)^{-1}\widetilde{r}_\theta^R\left(j\right)\right|=o_p\left(N^{-\lambda}\right)$, since $\left\|\left(I-\widehat{\mathcal{L}}\right)^{-1}\right\|=O_p\left(1\right)$. For first term, we can write $\left(I-\widehat{\mathcal{L}}\right)^{-1}r_\theta\left(j\right)=m_\theta\left(j\right)+\widehat{m}_\theta^A\left(j\right)$ where $\widehat{m}_\theta^A\left(j\right)=\left(I-\widehat{\mathcal{L}}\right)^{-1}\left(\widehat{\mathcal{L}}-\mathcal{L}\right)m_\theta\left(j\right)$. Since we know $\left\|\left(I-\widehat{\mathcal{L}}\right)^{-1}\right\|=O_p\left(1\right)$ from earlier, from Lemma 2.1 $\left\|\widehat{\mathcal{L}}-\mathcal{L}\right\|=O_p\left(N^{-1/2}\right)$, and, $\max_{1\leq j\leq J}\sup_{\theta\in\Theta}\left|m_\theta\left(j\right)\right|=O\left(1\right)$ as $m_\theta\left(j\right)$

is a continuous function on a compact set $\Theta$ any $j$, this completes the proof with $\widehat{m}_{\theta}^{R}=\widehat{m}_{\theta}^{A}+\left(I-\widehat{\mathcal{L}}\right)^{-1}\widetilde{r}_{\theta}^{R}$.■

PROOF OF LEMMA 2.4. The empirical analogue of (84) is

$$\widehat{g}_{\theta}=\widehat{\mathcal{H}}\widehat{m}_{\theta},$$

where $\widehat{\mathcal{H}}$ is a linear operator that uses local constant approximation to estimate the conditional expectation operator $\mathcal{H}$. Then we proceed, similarly to the proof of Lemma 2.3, by writing $\widehat{g}_{j}(a,\theta)=g_{j}(a,\theta)+\widehat{g}_{j}^{A}(a,\theta)+\widehat{\mathcal{H}}\widehat{m}_{\theta}^{R}(j,a)$ where $\widehat{g}_{j}^{A}(a,\theta)=\left(\widehat{\mathcal{H}}-\mathcal{H}\right)m_{\theta}(j,a)$ for any $j$. The approach taken here is similar to that found in Srisuma and Linton (2009), we decompose $\widehat{g}_{j}^{A}(a,\theta)$ into variance+bias terms, note that the presence of discrete regressor only leads to a straightforward sample splitting in the local regression for each $x$. Since $A$ is a compact set, the bias term near the boundary for Nadaraya-Watson estimator has a slower rate of convergent there than in the interior, for this reason we will need to trim out values near the boundary of $A$. For the ease of notation we proceed by assuming that the support of $a_{it}$ is $A_{N}$, where $\{A_{n}\}_{n=1}^{N}$ is a sequence of increasing sets such that $\bigcup_{n=1}^{\infty}A_{n}=int(A)$, here the boundary of the set $A$ has zero measure w.r.t. any relevant measure to our problem so we can ignore the difference between $A$ and $int(A)$. In our case $A=[\underline{a},\overline{a}]$ then $A_{N}=[\underline{a}+\gamma_{N},\overline{a}-\gamma_{N}]$ such that $\gamma_{N}=o(1)$ and $h=o(\gamma_{N})$. So we only need the trimming factor to converge to zero (at any rate) slower than the bandwidth, the reason behind this is fact that, for large $N$, the boundary only effect exists within a neighborhood of a single bandwidth. Then for any $m = (m_1 \ldots m_J)' \in \mathbb{R}^J, a$ and $j$

$$\begin{aligned}
\left(\widehat{\mathcal{H}} - \mathcal{H}\right) m(j,a) &= \sum_{k=1}^{J} m_k \left( \frac{\widehat{p}_{X',X,A}(k,j,a)}{\widehat{p}_{X,A}(j,a)} - \frac{p_{X',X,A}(k,j,a)}{p_{,X,A}(j,a)} \right) \qquad (103) \\
&= \sum_{k=1}^{J} m_k \left( \frac{\widehat{p}_{X',X,A}(k,j,a) - p_{X',X,A}(k,j,a)}{p_{X,A}(j,a)} \right) \\
&\quad - \sum_{k=1}^{J} m_k \left( \frac{\widehat{p}_{X',X,A}(k,j,a)}{\widehat{p}_{X,A}(j,a)\, p_{X,A}(j,a)} \left(\widehat{p}_{X,A}(j,a) - p_{X,A}(j,a)\right) \right),
\end{aligned}$$

where

$$\begin{aligned}
\widehat{p}_{X',X,A}(k,j,a) &= \frac{1}{NT} \sum_{i=1,t=1}^{N,T} \mathbf{1}\left[x_{it+1} = k, x_{it} = j\right] K_h(a_{it} - a), \\
\widehat{p}_{X,A}(j,a) &= \frac{1}{NT} \sum_{i=1,t=1}^{N,T} \mathbf{1}\left[x_{it} = j\right] K_h(a_{it} - a).
\end{aligned}$$

For any $j, k$, then

$$\begin{aligned}
& \widehat{p}_{X',X,A}(k,j,a) - p_{X',X,A}(k,j,a) \\
= \;& \left(\widehat{p}_{X',X,A}(k,j,a) - E\left[\widehat{p}_{X',X,A}(k,j,a)\right]\right) + \left(E\left[\widehat{p}_{X',X,A}(k,j,a)\right] - p_{X',X,A}(k,j,a)\right) \\
= \;& I_{11}(k,j,a) + I_{12}(k,j,a),
\end{aligned}$$

where $I_{11}(k,j,a)$ has zero mean and $I_{12}(k,j,a)$ is nonstochastic for any $a \in A_N$. Under stationarity, by the standard change of variable and differentiability of $p_{X',X,A}(k,j,a)$ (w.r.t. $a$)

$$I_{12}(k,j,a) = \frac{1}{2} h^4 \mu_2(K) \frac{\partial^2}{\partial a^2} p_{X',X,A}(k,j,a) + o\left(h^2\right).$$

It then follows that $\max_{1 \leq j,k \leq J} \sup_{a \in A_N} |I_{12}(k,j,a)| = O\left(h^4\right)$ since $\frac{\partial^4}{\partial a^4} p_{X',X,A}(k,j,a)$ is a continuous function on $a$ for any $j$ and $k$. It is also straightforward to show by using the same arguments as in Lemma 2.2 that $\max_{1 \leq j,k \leq J} \sup_{a \in A_N} |I_{11}(k,j,a)| =$

$o_p\left(\frac{N^{\xi}}{\sqrt{Nh}}\right)$. In particular, this follows since

$$\mathrm{var}\left(\sqrt{NTh}I_{11}(k,j,a)\right) = p_{X',X,A}(k,j,a)\,\kappa_2(K) + o(1),$$

where the display above for any $j$ and $k$ uniformly over $A_N$. Combining terms we have

$$\begin{aligned}
&\max_{1\leq j,k\leq J}\sup_{a\in A_N}\left|\sum_{k=1}^{J} m_k\left(\frac{\widehat{p}_{X',X,A}(k,j,a) - p_{X',X,A}(k,j,a)}{p_{X,A}(j,a)}\right)\right| \\
\leq\ & J\frac{\max_{1\leq j\leq J}|m_j|}{\min_{1\leq j\leq J}\inf_{a\in A_N}|p_{X,A}(j,a)|}\times\max_{1\leq j,k\leq J}\sup_{a\in A_N}\left|\widehat{p}_{X',X,A}(k,j,a) - p_{X',X,A}(k,j,a)\right| \\
=\ & O_p\left(h^4 + \frac{N^{\xi}}{\sqrt{Nh}}\right),
\end{aligned}$$

where the inequality holds w.p.a. 1 since we know (to be shown next) $\widehat{p}_{X,A}$ converges to $p_{X,A}$ uniformly over $X\times A_N$. By the same type of argument as above, write for each $j$

$$\begin{aligned}
&\widehat{p}_{X,A}(j,a) - p_{X,A}(j,a) \\
=\ & (\widehat{p}_{X,A}(j,a) - E[\widehat{p}_{X,A}(j,a)]) + (E[\widehat{p}_{X,A}(j,a)] - p_{X,A}(j,a)) \\
=\ & I_{21}(j,a) + I_{22}(j,a),
\end{aligned}$$

then it is straightforward to show the followings hold uniformly over its arguments

$$\begin{aligned}
I_{22}(j,a) &= \frac{1}{2}h^4\mu_4(K)\frac{\partial^4}{\partial a^2}p_{X,A}(j,a) + o\left(h^2\right), \\
\mathrm{var}\left(\sqrt{NTh}I_{21}(k,j,a)\right) &= p_{X,A}(j,a)\,\kappa_2(K) + o(1),
\end{aligned}$$

then we have

$$
\begin{aligned}
& \max_{1\le j,k\le J}\sup_{a\in A_N}\left|\sum_{k=1}^{J} m_k\left(\frac{\widehat{p}_{X',X,A}(k,j,a)}{\widehat{p}_{X,A}(j,a)\,p_{X,A}(j,a)}\left(\widehat{p}_{X,A}(j,a)-p_{X,A}(j,a)\right)\right)\right| \\
\le\ & J\frac{\max_{1\le j\le J}|m_j|}{\min_{1\le j\le J}\inf_{a\in A_N}|p_{X,A}(j,a)|^2}\times\max_{1\le j\le J}\sup_{a\in A_N}\left|\widehat{p}_{X,A}(j,a)-p_{X,A}(j,a)\right| \\
=\ & O_p\left(h^4+\frac{N^{\xi}}{\sqrt{Nh}}\right).
\end{aligned}
$$

So we can write for each $j$

$$
\begin{aligned}
\left(\widehat{\mathcal{H}}-\mathcal{H}\right)m(j,a) &= \sum_{k=1}^{J} m_k\left(\frac{\widehat{p}_{X',X,A}(k,j,a)-p_{X',X,A}(k,j,a)}{p_{X,A}(j,a)}\right) \\
&\quad -\sum_{k=1}^{J} m_k\left(\frac{p_{X',X,A}(k,j,a)}{p^2_{X,A}(j,a)}\left(\widehat{p}_{X,A}(j,a)-p_{X,A}(j,a)\right)\right)+W_{N,j}(a;m) \\
&= B_{N,j}(a;m)+V_{N,j}(a;m)+W_{N,j}(a;m),
\end{aligned}
$$

where

$$
B_{N,j}(a;m) = \frac{1}{2}h^4\mu_4(K)\sum_{k=1}^{J} m_k\left(\frac{\frac{\partial^4}{\partial a^4}p_{X',X,A}(k,j,a)}{p_{X,A}(j,a)}+\frac{p_{X',X,A}(k,j,a)\frac{\partial^4}{\partial a^4}p_{X,A}(j,a)}{p^2_{X,A}(j,a)}\right), \tag{104}
$$

$$
V_{N,j}(a;m) = \sum_{k=1}^{J} m_k\left(\begin{array}{c}\frac{1}{p_{X,A}(j,a)}\frac{1}{NT}\sum_{i=1,t=1}^{N,T}\left(\begin{array}{c}\mathbf{1}\left[x_{it+1}=k,x_{it}=j\right]K_h(a_{it}-a)\\ -E\left[\mathbf{1}\left[x_{it+1}=k,x_{it}=j\right]K_h(a_{it}-a)\right]\end{array}\right)\\ -\frac{p_{X',X,A}(k,j,a)}{p^2_{X,A}(j,a)}\frac{1}{NT}\sum_{i=1,t=1}^{N,T}\left(\begin{array}{c}\mathbf{1}\left[x_{it}=j\right]K_h(a_{it}-a)\\ -E\left[\mathbf{1}\left[x_{it}=j\right]K_h(a_{it}-a)\right]\end{array}\right)\end{array}\right) \tag{105}
$$

$$
W_{N,j}(a;m) = \sum_{k=1}^{J} m_k\left(\begin{array}{c}\frac{1}{p_{X,A}(j,a)}\left(\frac{\widehat{p}_{X',X,A}(k,j,a)}{\widehat{p}_{X,A}(j,a)}-\frac{p_{X',X,A}(k,j,a)}{p_{X,A}(j,a)}\right)\\ \times\left(\widehat{p}_{X,A}(j,a)-p_{X,A}(j,a)\right)\end{array}\right). \tag{106}
$$

Note that $B_{N,j}$ is a deterministic term, $V_{N,j}$ is the zero mean process that will deliver CLT whilst, using the same arguments as above, it is straightforward to show that $\max_{1\le j\le J}\sup_{a\in A_N} W_{N,j}(a;m)=o_p\left(B_{N,j}(a;m)+V_{N,j}(a;m)\right)$ for any $m\in\mathbb{R}^J$. Then we can conclude $\left\|\widehat{\mathcal{H}}-\mathcal{H}\right\|=O_p\left(h^4+\frac{N^{\xi}}{\sqrt{Nh}}\right)$. Using the decomposition of $\widehat{\mathcal{H}}-\mathcal{H}$ above we have

$$\widehat{g}_j^A(a,\theta) = \widehat{g}_j^B(a,\theta) + \widehat{g}_j^S(a,\theta) + W_{N,j}(a;m_\theta),$$

where, from (104) - (105), $\widehat{g}_j^B(a,\theta) = B_{N,j}(a;m_\theta)$ and $\widehat{g}_j^S(a,\theta) = V_{N,j}(a;m_\theta)$. It also follows that these terms have the desired rate of convergence that holds uniformly over $\Theta$ as well since $\mathcal{H}$ is independent of $\theta$ and $m_\theta$ is a vector of $J-$real value functions that are continuous on $\Theta$. Finally, we define $\widehat{g}_j^R(a,\theta)$ to be $W_{N,j}(a;m_\theta) + \widehat{\mathcal{H}}\widehat{m}_\theta^R(j,a)$. By the previous reasoning $W_{N,j}(a;m_\theta)$ already has the desired stochastic order so the proof of Lemma 2.4 will be complete if we can show, generally, that $\max_{1\leq j\leq J}\sup_{\theta,a\in\Theta\times A_N}\left|\widehat{\mathcal{H}}\widehat{m}_\theta^R(j,a)\right| = o_p\left(h^4 + \frac{N^\xi}{\sqrt{Nh}}\right)$. This is indeed true, since we have already shown that $\left\|\widehat{\mathcal{H}} - \mathcal{H}\right\| = o_p\left(h^4 + \frac{N^\xi}{\sqrt{Nh}}\right)$ and given that $\mathcal{H}$ is a conditional expectation operator, this implies that $\|\mathcal{H}\| \leq 1$, it follows from triangle inequality and the definition of operator norm that $\max_{1\leq j\leq J}\sup_{\theta,a\in\Theta\times A_N}\left|\widehat{\mathcal{H}}\widehat{m}_\theta^R(j,a)\right| = o_p\left(N^{-\lambda}\right)$.■

PROOF OF LEMMA 2.5. When $l = 0$, this follows from Lemma 2.4 with $h = O\left(N^{-1/7}\right)$. Other values of $l$ can also be shown very similarly, only more tedious. Since $\dim(A) = 1$ then $\partial_a^{|l|} = \frac{\partial^l}{\partial a^l}$, when $l = 1$, taking a derivative w.r.t. $a$ on (103) we obtain

$$\begin{aligned}
\frac{\partial}{\partial a}\left(\widehat{\mathcal{H}} - \mathcal{H}\right)m(j,a) &= \sum_{k=1}^J m_k \frac{\partial}{\partial a}\left(\frac{\widehat{p}_{X',X,A}(k,j,a) - p_{X',X,A}(k,j,a)}{p_{X,A}(j,a)}\right) \\
&\quad - \sum_{k=1}^J m_k \frac{\partial}{\partial a}\left(\frac{\widehat{p}_{X',X,A}(k,j,a)}{\widehat{p}_{X,A}(j,a)\,p_{X,A}(j,a)}\left(\widehat{p}_{X,A}(j,a) - p_{X,A}(j,a)\right)\right) \\
&= \sum_{k=1}^J m_k \left(\begin{array}{c} \frac{1}{p_{X,A}(j,a)}\frac{\partial}{\partial a}\left(\widehat{p}_{X',X,A}(k,j,a) - p_{X',X,A}(k,j,a)\right) \\ -\frac{\frac{\partial}{\partial a}p_{X,A}(j,a)}{p_{X,A}^2(j,a)}\left(\widehat{p}_{X',X,A}(k,j,a) - p_{X',X,A}(k,j,a)\right) \end{array}\right) \\
&\quad - \sum_{k=1}^J m_k \left(\begin{array}{c} \frac{\widehat{p}_{X',X,A}(k,j,a)}{\widehat{p}_{X,A}(j,a)p_{X,A}(j,a)}\frac{\partial}{\partial a}\left(\widehat{p}_{X,A}(j,a) - p_{X,A}(j,a)\right) \\ +\left(\frac{\partial}{\partial a}\frac{\widehat{p}_{X',X,A}(k,j,a)}{\widehat{p}_{X,A}(j,a)p_{X,A}(j,a)}\right)\left(\widehat{p}_{X,A}(j,a) - p_{X,A}(j,a)\right) \end{array}\right).
\end{aligned}$$

As seen in the proof of Lemma 2.4, it will be sufficient to show that $\max_{1\leq j,k\leq J}\sup_{a\in A_N}|\frac{\partial}{\partial a}\widehat{p}_{X',X,A}($

$-\frac{\partial}{\partial a}p_{X',X,A}(k,j,a)\,| = o_p(1)$, and, $\max_{1\le j,k\le J}\sup_{a\in A_N}|\frac{\partial}{\partial a}\widehat{p}_{X,A}(j,a) - \frac{\partial}{\partial a}p_{X,A}(j,a)\,|$ $= o_p(1)$ since we assume that $\frac{\partial}{\partial a}p_{X',X,A}(k,j,a)$ and $\frac{\partial}{\partial a}p_{X,A}(j,a)$ are continuous functions on a compact set $A$ for any $j,k$. Proceeding as in the proof of Lemma 2.4, first note that for any $j,k$

$$\begin{aligned} E\left[\frac{\partial}{\partial a}\widehat{p}_{X',X,A}(k,j,a)\right] &= -\frac{1}{h}\int p_{X',X,A}(k,j,a+wh)\,dK(w) \\ &= \int \frac{\partial}{\partial a}p_{X',X,A}(k,j,a+wh)\,K(w)\,dw \\ &= \frac{\partial}{\partial a}p_{X',X,A}(k,j,a) + O\left(h^4\right). \end{aligned}$$

The first line in the display follows from a standard change of variable argument, then using integration by parts and Taylor's expansion, the last equality above holds uniformly over $A$. It is easy to verify that uniformly over $A$

$$\mathrm{var}\left(\sqrt{NTh^3}\frac{\partial}{\partial a}\widehat{p}_{X',X,A}(k,j,a)\right) = O(1).$$

As seen in Lemma 2.2, it then follows that $\max_{1\le j,k\le J}\sup_{a\in A_N}|\frac{\partial}{\partial a}\widehat{p}_{X',X,A}(k,j,a) - \frac{\partial}{\partial a}p_{X',X,A}(k,j,$ $= O_p(h^4+\frac{N^{\xi}}{\sqrt{Nh^3}})$. Similarly one can show $\max_{1\le j,k\le J}\sup_{a\in A_N}\left|\frac{\partial}{\partial a}\widehat{p}_{X,A}(j,a) - \frac{\partial}{\partial a}p_{X,A}(j,a)\right| =$ $O_p(h^4+\frac{N^{\xi}}{\sqrt{Nh^3}})$. It is easy to see that choosing $h = O\left(N^{-1/7}\right)$ will imply $\max_{1\le j\le J}\sup_{\theta,a\in\Theta\times A_N}|\frac{\partial}{\partial a}$ $-\frac{\partial}{\partial a}g_{\theta}(j,a)\,| = o_p(1)$.■

Proof of Lemma 2.6. Since $\mathcal{R}_0$ and $\mathcal{M}_0$ are $J-$dimensional subspace of twice continuously differentiable functions, DCT is applicable throughout. When $p = 0$ the result follows from Lemma 2.5. Consider the case when $p = 1$ and $l = 0$, for all $1\le j\le J, 1\le k\le L$ and $\lambda < 1/2$, the exact same arguments used in proofing Lemma 2.2 can then be used to show $\frac{\partial}{\partial\theta_k}\widetilde{r}_{\theta}(j) = \frac{\partial}{\partial\theta_k}r_{\theta}(j) + \frac{\partial}{\partial\theta_k}\widetilde{r}^R_{\theta}(j)$ with $\max_{1\le j\le J}\sup_{\theta,a\in\Theta\times A_N}\left|\frac{\partial}{\partial\theta_k}\widetilde{r}^R_{\theta}(j)\right| = o_p\left(N^{-\lambda}\right)$, and since $\mathcal{L}$ is independent of $\theta$, the same arguments found in Lemma 2.3 can be used to show $\frac{\partial}{\partial\theta_k}\widehat{m}_{\theta}(j) = \frac{\partial}{\partial\theta_k}m_{\theta}(j) +$

$\frac{\partial}{\partial\theta_k}\widehat{m}_\theta^R(j)$ with $\max_{1\leq j\leq J}\sup_{\theta,a\in\Theta\times A_N}\left|\frac{\partial}{\partial\theta_k}\widehat{m}_\theta^R(j)\right| = o_p\left(N^{-\lambda}\right)$. Apart from replacing $(r_\theta, m_\theta)$ everywhere by $\left(\frac{\partial}{\partial\theta_k}r_\theta, \frac{\partial}{\partial\theta_k}m_\theta\right)$, we note that it is here that we need $\frac{\partial^2}{\partial\varepsilon\partial\theta_k}u_\theta(a,j,\varepsilon)$ to be continuous on all $a, j$ and $\theta$. Since $\mathcal{H}$ is independent of $\theta$, the arguments used in Lemma 2.4 can be directly applied to show

$$\frac{\partial}{\partial\theta_k}\widehat{g}_j(a,\theta) = \frac{\partial}{\partial\theta_k}g_j(a,\theta) + \frac{\partial}{\partial\theta_k}\widehat{g}_j^B(a,\theta) + \frac{\partial}{\partial\theta_k}\widehat{g}_j^S(a,\theta) + \frac{\partial}{\partial\theta_k}\widehat{g}_j^R(a,\theta),$$

such that

$$\begin{aligned}
\max_{1\leq j\leq J}\sup_{\theta,a\in\Theta\times A_N}\left|\frac{\partial}{\partial\theta_k}\widehat{g}_j^B(a,\theta)\right| &= O_p\left(h^2\right), \\
\max_{1\leq j\leq J}\sup_{\theta,a\in\Theta\times A_N}\left|\frac{\partial}{\partial\theta_k}\widehat{g}_j^S(a,\theta)\right| &= o_p\left(\frac{N^\xi}{\sqrt{Nh}}\right), \\
\max_{1\leq j\leq J}\sup_{\theta,a\in\Theta\times A_N}\left|\frac{\partial}{\partial\theta_k}\widehat{g}_j^R(a,\theta)\right| &= o_p\left(h^2 + \frac{N^\xi}{\sqrt{Nh}}\right),
\end{aligned}$$

where $\frac{\partial}{\partial\theta_k}\widehat{g}_j^B(a,\theta) = B_{N,j}\left(a; \frac{\partial}{\partial\theta_k}m_\theta\right)$, $\frac{\partial}{\partial\theta_k}\widehat{g}_j^S(a,\theta) = V_{N,j}\left(a; \frac{\partial}{\partial\theta_k}m_\theta\right)$ and $\frac{\partial}{\partial\theta_k}\widehat{g}_j^R(a,\theta) = W_{N,j}\left(a; \frac{\partial}{\partial\theta_k}m_\theta\right) + \widehat{\mathcal{H}}\frac{\partial}{\partial\theta_k}\widehat{m}_\theta^R(j,a)$ and these terms are defined in (104) - (106). For $l = 2$ and $1 \leq k, d \leq L$, we simply replace $\frac{\partial}{\partial\theta_k}$ by $\frac{\partial^2}{\partial\theta_k\partial\theta_d}$ and the exact same reasoning used when $p = 1$ can be applied directly. All other cases of $0 \leq l, p \leq 2$ can be shown similarly.■

PROOF OF LEMMA 2.7. We first show that $\mathbf{1}\left[\cdot \leq \rho_j(a,\theta,\partial_a g_j)\right]$ is locally uniformly $L^2(Q)$ –continuous for all $j$ with respect to $a, \theta, \partial_a g_j$. More precisely, we need to show for a positive sequence $\delta_N = o(1)$ and any $(a,\theta,\partial_a g_j) \in A \times \Theta \times \mathcal{G}_j^{(1)}$ that

$$\lim_{N\to\infty}\left(E\left[\sup_{\left\|\left(a'-a,\theta'-\theta,\partial_a g_j'-\partial_a g_j\right)\right\|<\delta_N}\left|\mathbf{1}\left[\varepsilon_i \leq \rho_j\left(a',\theta',\partial_a g_j'\right)\right] - \mathbf{1}\left[\varepsilon_i \leq \rho_j(a,\theta,\partial_a g_j)\right]\right|^2\right]\right)^{1/2} = 0. \tag{107}$$

To do this, take any $\left\|\left(a'-a,\theta'-\theta,\partial_a g_j'-\partial_a g_j\right)\right\|<\delta_N$, then we have for all $j$

$$\begin{aligned}
\left|\rho_j\left(a',\theta',\partial_a g_j'\right)-\rho_j\left(a,\theta,\partial_a g_j\right)\right| &\leq C_0\left\{\left\|\left(a'-a,\theta'-\theta\right)\right\|+\left\|\partial_a g_j'-\partial_a g_j\right\|_{\mathcal{G}}\right\} \\
&\quad +o\left(\left\|\left(a'-a,\theta'-\theta\right)\right\|^2+\left\|\partial_a g_j'-\partial_a g_j\right\|_{\mathcal{G}}^2\right) \\
&\leq C_0\delta_N+o\left(\delta_N\right),
\end{aligned}$$

this follows from Taylor's theorem in Banach Space since $\rho_j$ is twice Fréchet differentiable, see Chapter 4 of Zeidler (1986). Ignoring the smaller order term, this implies

$$\begin{aligned}
\rho_j\left(a,\theta,\partial_a g_j\right)-C_0\delta_N &\leq \rho_j\left(a',\theta',\partial_a g_j'\right)\leq\rho_j\left(a,\theta,\partial_a g_j\right)+C_0\delta_N, \\
\rho_j\left(a,\theta,\partial_a g_j\right)-C_0\delta_N &\leq \rho_j\left(a,\theta,\partial_a g_j\right)\leq\rho_j\left(a,\theta,\partial_a g_j\right)+C_0\delta_N.
\end{aligned}$$

Combining the inequalities above, it follows that $\sup_{\left\|\left(a'-a,\theta'-\theta,\partial_a g_j'-\partial_a g_j\right)\right\|<\delta_N}|\mathbf{1}\left[\varepsilon_i\leq\rho_j\left(a',\theta',\partial_a g\right.\right.$ $-\mathbf{1}\left[\varepsilon_i\leq\rho_j\left(a,\theta,\partial_a g_j\right)\right]|$ is bounded above by $\mathbf{1}\left[\rho_j\left(a,\theta,\partial_a g_j\right)-C_0\delta_N<\varepsilon_i\leq\rho_j\left(a,\theta,\partial_a g_j\right)+C_0\delta_N\right.$ This majorant takes value 1 with probability $Q_\varepsilon\left(\rho_j\left(a,\theta,\partial_a g_j\right)+C_0\delta_N\right)-Q_\varepsilon\left(\rho_j\left(a,\theta,\partial_a g_j\right)-C_0\delta_N\right.$ and zero otherwise, then by Lipschitz continuity of $Q_\varepsilon$, (107) holds as required. Since $A\times\Theta$ is a compact Euclidean set it has a known covering number. For $\mathcal{G}_j^{(1)}$, since $\mathcal{G}_j\subset C^2(A)$ we have $\mathcal{G}_j^{(1)}\subset C^1(A)$; given that $\dim(A)=1$ we can apply Corollary 2.7.3 of VW to show that $\int_0^\infty\sqrt{\log N\left(\varepsilon,\mathcal{G}_j^{(1)},\|\cdot\|_{\mathcal{G}}\right)}d\varepsilon<\infty$, together with $L^2(Q)$ −continuity of $\mathbf{1}\left[\cdot\leq\rho_j\left(a,\theta,\partial_a g_j\right)\right]$, as shown in the proof of Theorem 3 (part (ii)) in Chen et al. (2003), $\mathcal{F}_j$ is $Q$−Donsker for each $j$.■

Proof of Lemma 2.8. For all $j$, $\mathcal{F}_j$ is $Q$−Donsker and is locally uniformly $L^2(Q)$ −continuous with respect to $a,\theta,\partial_a g_j$, as described in (107), Lemma 2.1 of Chen et al. (2003) implies that the stochastic equicontinuity also holds with respect to the parameters that index the functions in $\mathcal{F}_j$.■

PROOF OF LEMMA 2.9. For any $a$ and $j$ write

$$\sqrt{N}\left(\widehat{F}_{A|X}(a|j) - F_{A|X}(a|j)\right) = \mathfrak{F}_{1,N}(a,j) + \mathfrak{F}_{2,N}(a,j),$$

where

$$\begin{aligned}
\mathfrak{F}_{1,N}(a,j) &= \frac{1}{T\widehat{p}_X(j)} \times \frac{1}{\sqrt{N}} \sum_{i=1,t=1}^{N,T} \left(\mathbf{1}[a_{it} \le a, x_{it} = j] - F_{A,X}(a,j)\right), \\
\mathfrak{F}_{2,N}(a,j) &= -\frac{\sqrt{T} F_{A|X}(a|j)}{\widehat{p}_X(j)} \times \sqrt{N}\left(\widehat{p}_X(j) - p_X(j)\right).
\end{aligned}$$

Define $\mathcal{C}_a = \{y_a \in \mathbb{R} : y_a \le a\}$, then $\mathcal{C} = \bigcup_{a\in A} \mathcal{C}_a$ a classical VC-class of sets, for the definition VC-class of sets see Pollard (1990). Since $X$ is finite, it is also necessarily a VC-class of sets. Then for each $x$, $\frac{1}{\sqrt{NT}} \sum_{i=1,t=1}^{N,T} \left(\mathbf{1}[a_{it} \le \cdot, x_{it} = j] - F_{A,X}(\cdot, x)\right)$ converges weakly to some tight Gaussian process in $l^\infty(A)$ since $\mathcal{C} \times X$ is VC in $A \times X$, by Lemma 2.6.17 in VW, and VC-classes of functions is a Donsker class, see also Type I classes of Andrews (1994b). With an abuse of notation, for each $x$ let $\frac{1}{\widehat{p}_X(j)}$ $(\frac{1}{p_X(j)})$ also denote a random element that takes value in $l^\infty(A)$ such that the sample path of $\frac{1}{\widehat{p}_X(j)}$ $(\frac{1}{p_X(j)})$ is constant over $A$. By standard LLN $\frac{1}{\widehat{p}_X(j)} \xrightarrow{p} \frac{1}{p_X(j)}$ and it follows by Slutsky's theorem that $\mathfrak{F}_{1,N}(\cdot, x)$ converges weakly to a random element in $l^\infty(A)$. In particular, the limit of $\mathfrak{F}_{1,N}(\cdot, j)$ is also a tight Gaussian process. From the finite dimensional (fidi) weak convergence, Gaussianity is clearly preserved if we replace $\frac{1}{\widehat{p}_X(j)}$ by $\frac{1}{p_X(j)}$, but since $\widehat{p}_X(j) - p_X(j) = o_p(1)$ the remainder term from the expansion $\frac{1}{\widehat{p}_X(j)} - \frac{1}{p_X(j)}$ can be used to construct a random element that converges to zero in probability on $A$, so by an application of Slutsky's theorem Gaussianity is preserved. Tightness trivially follow since the multiplication of $\frac{1}{\widehat{p}_X(j)}$ does not affect the asymptotic tightness of $\{\frac{1}{\sqrt{NT}} \sum_{i=1,t=1}^{N,T} \left(\mathbf{1}[a_{it} \le \cdot, x_{it} = j] - F_{A,X}(\cdot, j)\right)\}$. Since the only random component of $\mathfrak{F}_{2,N}(\cdot, j)$ is from $\sqrt{NT}\left(\widehat{p}_X(j) - p_X(j)\right)$, which is a finite dimensional random variable, then a similar argument to the one used previously can trivially show that $\mathfrak{F}_{2,N}(\cdot, j)$ must also converge to a Gaussian process which is tight $l^{\infty}(A)$, where tightness follows from the (equi-)continuity of $F_{A|X}(a|j)$ on $A$. Therefore $\sqrt{N}\left(\widehat{F}_{A|X=j} - F_{A|X=j}\right)$ must converge to a tight Gaussian process in $l^{\infty}(A)$ for all $j$ since asymptotic tightness is closed under finite addition and, in this case, it is easy to see that Gaussianity is also closed under the sum.■

PROOF OF LEMMA 2.10. By MVT, for all $a$ and $j$

$$
\begin{aligned}
& F_{A|X}(a|j; \theta_0, \partial_a \widehat{g}(\cdot, \theta_0)) - F_{A|X}(a|j; \theta_0, \partial_a g_{0,j}(\cdot, \theta_0)) \\
= \; & q\left(\overline{\rho}_j(a, \theta_0, \partial_a g_{0,j}(\cdot, \theta_0))\right)\left(\rho_j(a, \theta_0, \partial_a \widehat{g}_j(\cdot, \theta_0)) - \rho_j(a, \theta_0, \partial_a g_{0,j}(\cdot, \theta_0))\right),
\end{aligned}
$$

where $\overline{\rho}_j(a, \theta_0, \partial_a g_{0,j}(\cdot, \theta_0))$ is some intermediate value between $\rho_j(a, \theta_0, \partial_a \widehat{g}_j(\cdot, \theta_0))$ and $\rho_j(a, \theta_0, \partial_a g_{0,j}(\cdot, \theta_0))$. Since $\rho_j(a, \theta_0, \partial_a g_j)$ is twice Fréchet continuously differentiable on $A$ at $\partial_a g_{0,j}(\cdot, \theta_0)$, using the linearization assumption, the argument analogous to Lemma 2.9 with Slutsky theorem can be used to complete the proof.■

## 2.10 Tables and Figures

| | | $\widehat{\theta}_1$ | | | | $\widehat{\theta}_1^{trim}$ | | | |
|---|---|---|---|---|---|---|---|---|---|
| $NT$ | $\varsigma$ | bias | mbias | std | iqr | bias | mbias | std | iqr |
| 100 | 1/5 | 0.0084 | 0.0309 | 0.1558 | 0.1350 | -0.0310 | 0.0004 | 0.1934 | 0.1683 |
| | 1/6 | 0.0278 | 0.0442 | 0.1359 | 0.1205 | -0.0150 | 0.0098 | 0.1741 | 0.1518 |
| | 1/7 | 0.0419 | 0.0541 | 0.1161 | 0.1035 | -0.0002 | 0.0214 | 0.1568 | 0.1357 |
| | 1/8 | 0.0536 | 0.0638 | 0.1092 | 0.0947 | 0.0109 | 0.0315 | 0.1375 | 0.1211 |
| | 1/9 | 0.0647 | 0.0743 | 0.0996 | 0.0874 | 0.0153 | 0.0373 | 0.1328 | 0.1143 |
| | *static* | 0.2620 | 0.2614 | 0.0187 | 0.0247 | 0.2620 | 0.2614 | 0.0187 | 0.0247 |
| 500 | 1/5 | 0.0193 | 0.0163 | 0.0618 | 0.0546 | -0.0038 | -0.0070 | 0.0739 | 0.0709 |
| | 1/6 | 0.0320 | 0.0291 | 0.0546 | 0.0476 | 0.0014 | 0.0031 | 0.0689 | 0.0609 |
| | 1/7 | 0.0422 | 0.0419 | 0.0497 | 0.0445 | 0.0063 | 0.0059 | 0.0635 | 0.0582 |
| | 1/8 | 0.0508 | 0.0512 | 0.0456 | 0.0396 | 0.0128 | 0.0145 | 0.0600 | 0.0564 |
| | 1/9 | 0.0597 | 0.0604 | 0.0414 | 0.0376 | 0.0195 | 0.0213 | 0.0573 | 0.0542 |
| | *static* | 0.2607 | 0.2606 | 0.0076 | 0.0108 | 0.2607 | 0.2606 | 0.0076 | 0.0108 |
| 1000 | 1/5 | 0.0150 | 0.0141 | 0.0428 | 0.0388 | -0.0045 | -0.0067 | 0.0506 | 0.0463 |
| | 1/6 | 0.0277 | 0.0264 | 0.0372 | 0.0343 | 0.0009 | 0.0006 | 0.0464 | 0.0429 |
| | 1/7 | 0.0375 | 0.0374 | 0.0344 | 0.0316 | 0.0041 | 0.0038 | 0.0437 | 0.0425 |
| | 1/8 | 0.0457 | 0.0468 | 0.0315 | 0.0294 | 0.0090 | 0.0085 | 0.0413 | 0.0413 |
| | 1/9 | 0.0536 | 0.0543 | 0.0291 | 0.0294 | 0.0143 | 0.0143 | 0.0398 | 0.0402 |
| | *static* | 0.2610 | 0.2608 | 0.0054 | 0.0073 | 0.2610 | 0.2608 | 0.0054 | 0.0073 |
| 2500 | 1/5 | 0.0119 | 0.0118 | 0.0258 | 0.0246 | -0.0023 | -0.0036 | 0.0305 | 0.0291 |
| | 1/6 | 0.0229 | 0.0235 | 0.0225 | 0.0221 | 0.0012 | 0.0017 | 0.0279 | 0.0269 |
| | 1/7 | 0.0320 | 0.0332 | 0.0206 | 0.0198 | 0.0032 | 0.0033 | 0.0270 | 0.0280 |
| | 1/8 | 0.0405 | 0.0411 | 0.0200 | 0.0198 | 0.0055 | 0.0062 | 0.0267 | 0.0266 |
| | 1/9 | 0.0482 | 0.0486 | 0.0193 | 0.0190 | 0.0089 | 0.0087 | 0.0263 | 0.0259 |
| | *static* | 0.2610 | 0.2609 | 0.0034 | 0.0045 | 0.2610 | 0.2609 | 0.0034 | 0.0045 |

Table 5: $h_\varsigma = 1.06s(NT)^{-\varsigma}$ is the bandwidth, for various choices of $\varsigma$, used in the non-parametric estimation, $s =$ denotes the standard deviation of $\{a_{it}\}_{i=1,t=1}^{N,T+1}$; the statistics from estimating the static model are reported under *static*.

| | | $\widehat{\theta}_2$ | | | | $\widehat{\theta}_2^{trim}$ | | | |
|---|---|---|---|---|---|---|---|---|---|
| $NT$ | $\varsigma$ | bias | mbias | std | iqr | bias | mbias | std | iqr |
| 100 | 1/5 | 0.0657 | 0.0299 | 0.2026 | 0.1532 | 0.1121 | 0.0477 | 0.2856 | 0.1810 |
| | 1/6 | 0.0632 | 0.0290 | 0.1843 | 0.1520 | 0.1014 | 0.0446 | 0.2529 | 0.1836 |
| | 1/7 | 0.0567 | 0.0299 | 0.1670 | 0.1388 | 0.0948 | 0.0371 | 0.2458 | 0.1805 |
| | 1/8 | 0.0513 | 0.0259 | 0.1535 | 0.1324 | 0.0871 | 0.0404 | 0.2201 | 0.1801 |
| | 1/9 | 0.0464 | 0.0225 | 0.1442 | 0.1275 | 0.0858 | 0.0347 | 0.2168 | 0.1664 |
| | *static* | 0.1303 | 0.1316 | 0.0326 | 0.0432 | 0.1303 | 0.1316 | 0.0326 | 0.0432 |
| 500 | 1/5 | 0.0383 | 0.0364 | 0.0820 | 0.0769 | 0.0513 | 0.0454 | 0.0996 | 0.0926 |
| | 1/6 | 0.0329 | 0.0304 | 0.0772 | 0.0728 | 0.0473 | 0.0398 | 0.0990 | 0.0920 |
| | 1/7 | 0.0330 | 0.0315 | 0.0742 | 0.0715 | 0.0472 | 0.0385 | 0.0964 | 0.0891 |
| | 1/8 | 0.0335 | 0.0321 | 0.0711 | 0.0705 | 0.0442 | 0.0330 | 0.0922 | 0.0849 |
| | 1/9 | 0.0346 | 0.0331 | 0.0660 | 0.0655 | 0.0430 | 0.0313 | 0.0891 | 0.0830 |
| | *static* | 0.1310 | 0.1314 | 0.0141 | 0.0195 | 0.1310 | 0.1314 | 0.0141 | 0.0195 |
| 1000 | 1/5 | 0.0267 | 0.0262 | 0.0590 | 0.0565 | 0.0346 | 0.0337 | 0.0669 | 0.0662 |
| | 1/6 | 0.0212 | 0.0212 | 0.0550 | 0.0529 | 0.0281 | 0.0261 | 0.0646 | 0.0619 |
| | 1/7 | 0.0214 | 0.0213 | 0.0519 | 0.0499 | 0.0277 | 0.0247 | 0.0616 | 0.0559 |
| | 1/8 | 0.0247 | 0.0247 | 0.0491 | 0.0461 | 0.0288 | 0.0250 | 0.0588 | 0.0562 |
| | 1/9 | 0.0263 | 0.0266 | 0.0458 | 0.0431 | 0.0296 | 0.0245 | 0.0560 | 0.0531 |
| | *static* | 0.1300 | 0.1302 | 0.0095 | 0.0137 | 0.1300 | 0.1302 | 0.0095 | 0.0137 |
| 2500 | 1/5 | 0.0202 | 0.0219 | 0.0369 | 0.0368 | 0.0259 | 0.0273 | 0.0401 | 0.0397 |
| | 1/6 | 0.0156 | 0.0160 | 0.0346 | 0.0345 | 0.0206 | 0.0210 | 0.0384 | 0.0386 |
| | 1/7 | 0.0154 | 0.0161 | 0.0335 | 0.0340 | 0.0186 | 0.0190 | 0.0381 | 0.0366 |
| | 1/8 | 0.0191 | 0.0206 | 0.0331 | 0.0337 | 0.0203 | 0.0213 | 0.0372 | 0.0366 |
| | 1/9 | 0.0237 | 0.0249 | 0.0322 | 0.0324 | 0.0231 | 0.0232 | 0.0365 | 0.0356 |
| | *static* | 0.1306 | 0.1305 | 0.0060 | 0.0079 | 0.1306 | 0.1305 | 0.0060 | 0.0079 |

Table 6: $h_\varsigma = 1.06s(NT)^{-\varsigma}$ is the bandwidth, for various choices of $\varsigma$, used in the non-parametric estimation, $s =$ denotes the standard deviation of $\{a_{it}\}_{i=1,t=1}^{N,T+1}$; the statistics from estimating the static model are reported under *static*.

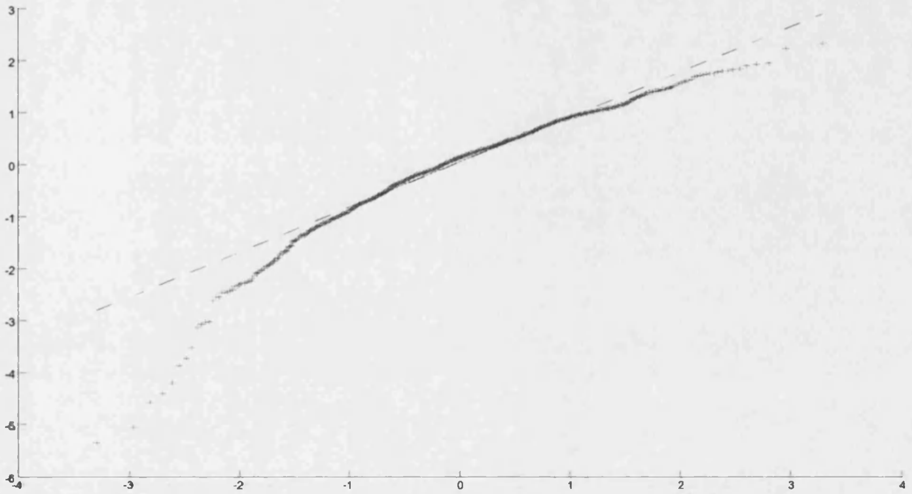


Figure 1: QQ Plot of sample (standardized) $\widehat{\theta}_1$ versus standard normal, $NT = 100$.

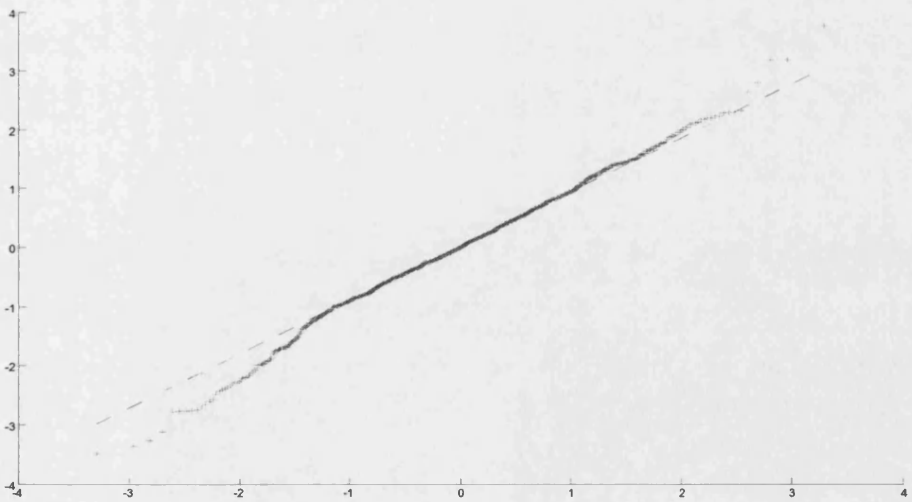


Figure 2: QQ Plot of sample (standardized) $\widehat{\theta}_1$ versus standard normal, $NT = 500$.

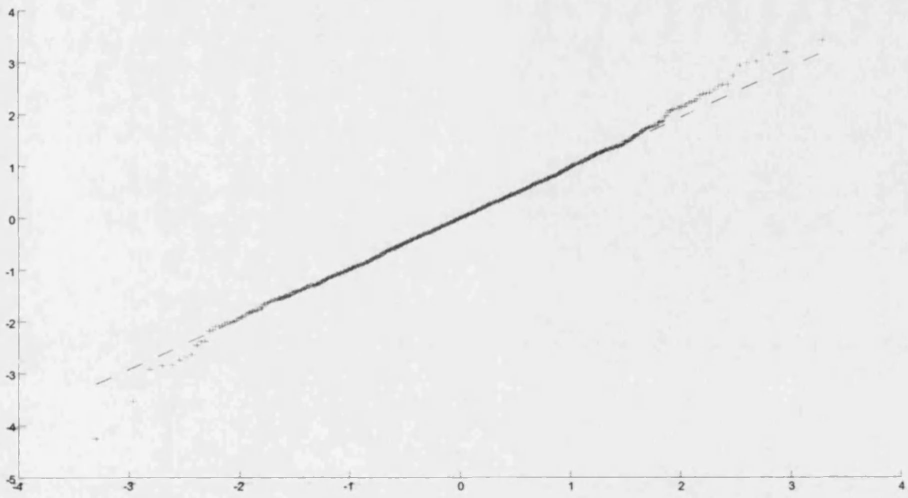


Figure 3: QQ Plot of sample (standardized) $\widehat{\theta}_1$ versus standard normal, $NT = 1000$.

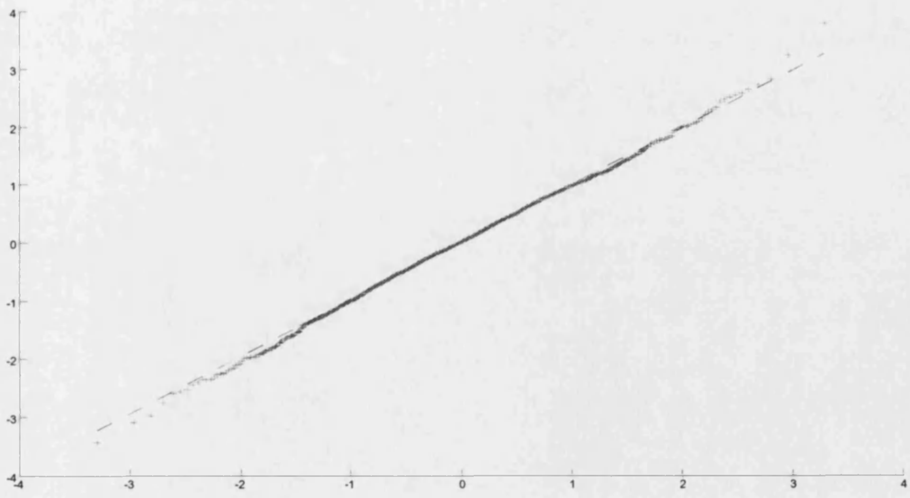


Figure 4: QQ Plot of sample (standardized) $\widehat{\theta}_1$ versus standard normal, $NT = 2500$.

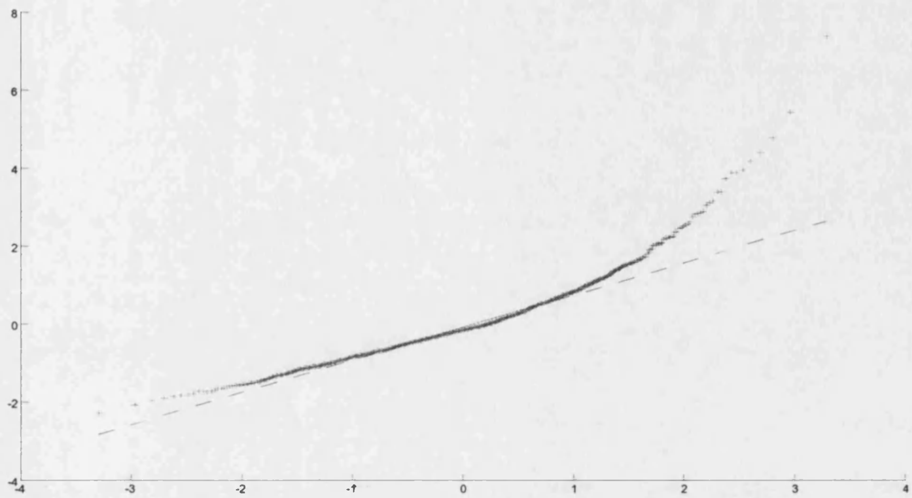


Figure 5: QQ Plot of sample (standardized) $\widehat{\theta}_2$ versus standard normal, $NT = 100$.

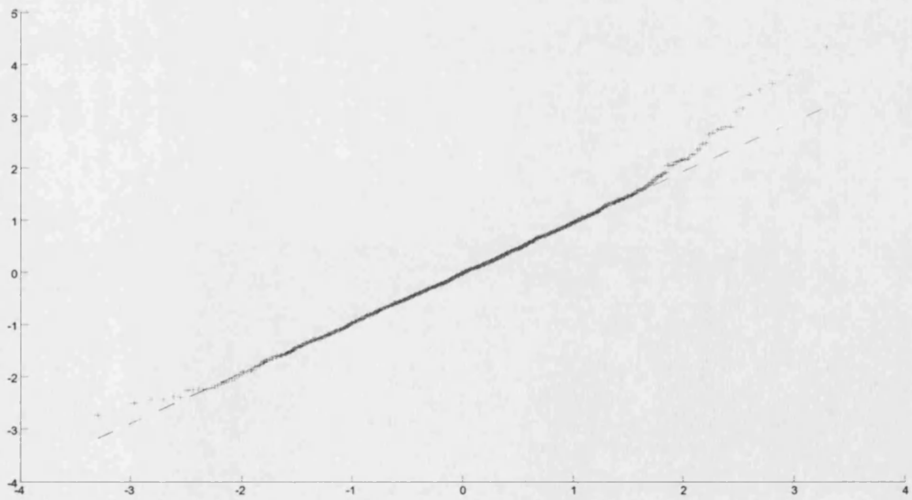


Figure 6: QQ Plot of sample (standardized) $\widehat{\theta}_2$ versus standard normal, $NT = 500$.

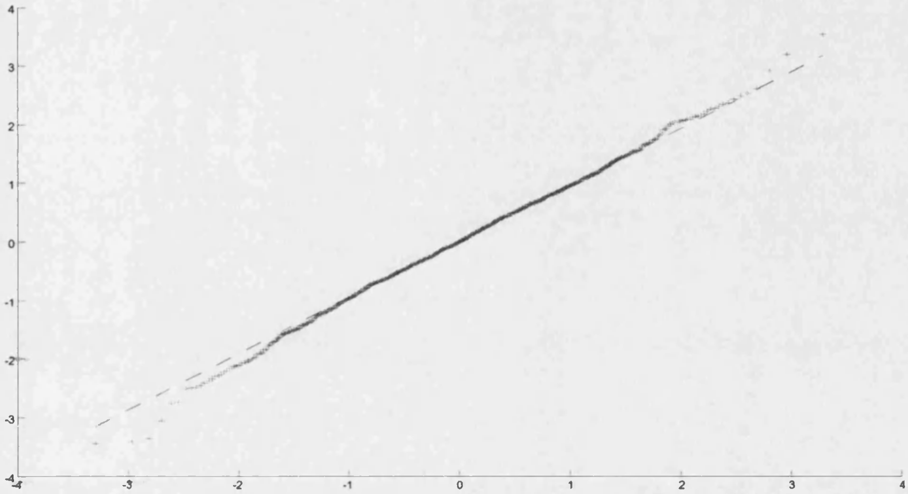


Figure 7: QQ Plot of sample (standardized) $\widehat{\theta}_2$ versus standard normal, $NT = 1000$.

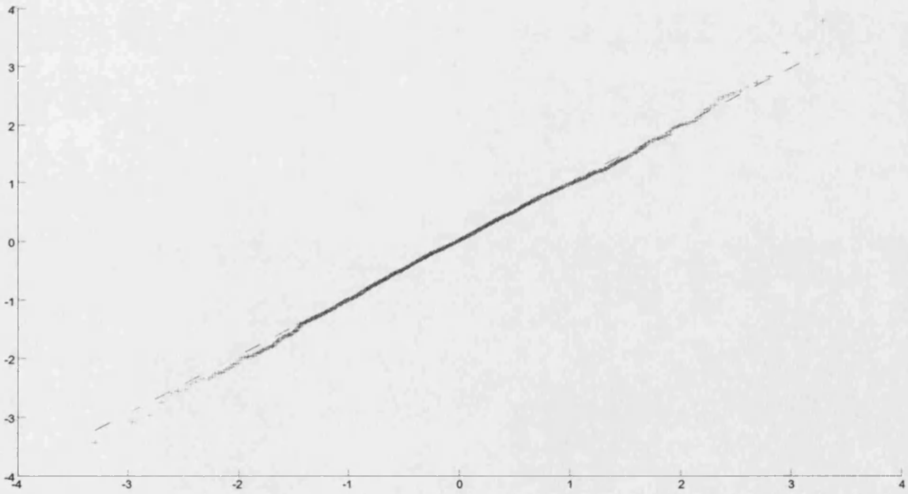


Figure 8: QQ Plot of sample (standardized) $\widehat{\theta}_2$ versus standard normal, $NT = 2500$.

# 3 Modelling and Estimating Other Dynamic Models

## 3.1 Introduction

As seen in the previous two chapters, the common insight in the two-step estimation procedures that we employ is the linear representation of the model implied value function which allow us to readily estimate the conditional value function and the continuation value function. In particular, when the observable state space is finite, the linear equation that defines the conditional value function is a matrix equation whilst Chapter 1 illustrates this can be generalized to the uncountably infinite dimensional framework by working with an integral equation. The first chapter (Section 1.5) also demonstrates how the estimation methodology for a single agent problem can be extended to estimate a class of Markovian games by taking averages across other players actions; this logic can straightforwardly be applied to the case of a continuous control as well.

Therefore in this chapter, we illustrate how to estimate the conditional value functions for other classes of dynamic programming models for a single agent problem, where we distinguish different models by the nature of the control variable(s). We only focus on the estimation of the conditional value function to avoid repetition since various objective functions can be constructed from the conditional moment restrictions implied by the policy function based on the conditional value function, see (25) and (91) for the discrete and continuous control cases respectively.

## 3.2 Discrete-Continuous Control

In many investment and pricing problems the distribution of the control variable has mass points as well as continuous component, for example firms may choose to not invest or prices are regulated to lie within certain bounds that is binding.

The framework of the decision problem here is similar to that found in Chapter 2. In particular, we assume the framework described in Section 2.2.1 as well as assumptions M2.1 - M2.3. For simplicity sake, let $A = [0, \overline{a}]$ for some $0 < \overline{a} \leq \infty$, we suppose that the control variable $a_t$ has a mixed distribution; it has a mass point at 0, with some probability $p_A(x_t) \in (0,1)$ for all possible values on the support of $x_t$, and has a density on $A \backslash \{0\}$. It is straightforward to allow for more than one mass points. We need to modify the monotone condition in M2.4 slightly to accommodate the mass point at zero.

ASSUMPTION M2.4$'$: *(Monotone Choice) The per period payoff function $u_\theta : A \times X \times \mathcal{E} \to \mathbb{R}$ has increasing differences in $(a, \varepsilon)$ for all $x$,$\theta$ and for $a \in A \backslash \{0\}$; $u_\theta$ is specified upto some unknown parameters $\theta \in \Theta \subset \mathbb{R}^L$.*

From assumptions M2.1 - M2.3, we can again obtain (82) and its conditional expectation, which we reproduce here for convenience:

$$
\begin{aligned}
V_\theta(s_t) &= u_\theta(a_t, x_t, \varepsilon_t) + \beta E[V_\theta(s_{t+1})|s_t], \\
E[V_\theta(s_t)|x_t] &= E[u_\theta(a_t, x_t, \varepsilon_t)|x_t] + \beta E[E[V_\theta(s_{t+1})|x_{t+1}]|x_t],
\end{aligned}
$$

where, as before

$$
\begin{aligned}
a_t &= \alpha^0_{\theta_0}(x_t, \varepsilon_t) \\
&= \max_{a \in A}\left\{u_\theta(a, x_t, \varepsilon_t) + \beta E\left[V^0_{\theta_0}(s_{t+1})|x_t, a_t = a\right]\right\}.
\end{aligned}
$$

We only need to show we can consistently estimate $E[V_\theta(s_t)|x_t]$, the solution to the matrix equation above. In order to do this, as seen in the last two chapters, we need to estimate two elements of the linear equation; the conditional expected payoff and the transition matrix, which we denote by $r_\theta$ and $\mathcal{L}$ respectively. The stochastic matrix $\mathcal{L}$

can be estimated in the exact same way as before. For $r_\theta$, first we write

$$\begin{aligned} E\left[u_\theta\left(a_t, x_t, \varepsilon_t\right)|x_t\right] &= \Pr\left[a_t = 0|x_t\right] E\left[u_\theta\left(0, x_t, \varepsilon_t\right)|x_t, a_t = 0\right] \\ &\quad + \Pr\left[a_t > 0|x_t\right] E\left[u_\theta\left(a_t, x_t, \varepsilon_t\right)|x_t, a_t > 0\right]. \end{aligned}$$

Clearly we do not have any problems providing estimates for $\Pr\left[a_t = 0|x_t\right]$ and $\Pr\left[a_t > 0|x_t\right]$, while we require M2.4$'$ to deal with the unobserved state variable $\varepsilon_t$ which enters $u_\theta$ non-additively. We can again rely on Topkis' Theorem, which ensures that the policy function is invertible on $A\backslash\{0\}$, to nonparametrically recover $\varepsilon_t$ by the relation $\widehat{\varepsilon}_t = Q_\varepsilon^{-1}\left(\widehat{F}_{A|X}\left(a_t|x_t\right)\right)$ for $a_t > 0$. The sequence $\left(a_t, x_t, \widehat{\varepsilon}_t\right)_{t=1}^T$ can then be used to estimate the regression function $E\left[u_\theta\left(a_t, x_t, \varepsilon_t\right)|x_t, a_t > 0\right]$. For the case when $a_t = 0$, although it is not possible to recover $\varepsilon_t$, by (weak) monotonicity of the policy function we know that

$$a_t = 0 \Leftrightarrow \varepsilon_t \leq \underline{\varepsilon}^c.$$

Using the equivalence condition above, the quantile invariance property between $(a_t, \varepsilon_t)$ implies that $\underline{\varepsilon}^c = Q_\varepsilon^{-1}\left(\Pr\left[a_t = 0|x_t\right]\right)$. We can then estimate the cutoff threshold by $Q_\varepsilon^{-1}\left(\widehat{\Pr\left[a_t = 0|x_t\right]}\right)$ and estimate $E\left[u_\theta\left(0, x_t, \varepsilon_t\right)|x_t, a_t = 0\right]$ by the empirical analogue of

$$E\left[u_\theta\left(0, x_t, \varepsilon_t\right)|x_t, a_t = 0\right] = \frac{\int_{\underline{\varepsilon}}^{\underline{\varepsilon}^c} u_\theta\left(0, x_t, \varepsilon\right) Q_\varepsilon\left(d\varepsilon\right)}{\Pr\left[a_t = 0|x_t\right]}.$$

Therefore the conditional value function is identified so we can proceed to estimate the continuation value and use it to construct some criterion function to estimate the parameter of interest $\theta_0$.

### 3.3 Discrete and Continuous Controls

The flexibility to estimate models with both discrete and continuous choices is very important, for example, the economic agents in the empirical study of oligopoly or dynamic auction models often endogenously choose whether to participate in the market before deciding on the price or investment decisions. The framework of the decision problem here is similar to Section 4 of Arcidiacono and Miller (2008). For each economic agent, the model now consists of the control variables $(a_t, d_t) \in A \times D$, where $A \subset \mathbb{R}$ and $D = \{1, \ldots, K\}$, and the state variables $s_t = \left(x_t, \varepsilon_t, v_t^K\right) \in X \times \mathcal{E} \times \mathcal{V}^K$, where $X = \{1, \ldots, J\}, \mathcal{E} \subset \mathbb{R}$ and $\mathcal{V}^K \subset \mathbb{R}^K$ so $v_t^K = (v_t(1), \ldots, v_t(K))$. The sequential decision problem can be stated as follows: at time $t$, the economic agent observes $\left(x_t, v_t^K\right)$ and choose an action $k \in \{1, \ldots, K\}$ to maximize $E\left[u\left(a_t, d_t, x_t, \varepsilon_t, v_t^K\right)|x_t, v_t^K, d_t = k\right] + \beta E\left[V(s_{t+1})|x_t, v_t^K, d_t = k\right]$, sequentially, she then observes $\varepsilon_t$ and chooses $a$ that maximizes $u\left(a, d_t, x_t, \varepsilon_t, v_t^K\right) + \beta E\left[V(s_{t+1})|s_t, d_t, a_t = a\right]$. The decisions made within and across period generally will affect the consequential state variables, we impose the conditions on the transition of the state variables within and across periods in the set of assumptions below. More formally, the decision problem (subject to the transition law) within each period $t$ leads to the following policy pair

$$\begin{aligned}
\delta\left(x_t, v_t^K\right) &= \arg\max_{1 \leq k \leq K} \left\{E\left[u\left(a_t, d_t, x_t, \varepsilon_t, v_t^K\right)|x_t, d_t = k\right] + \beta E\left[V(s_{t+1})|x_t, d_t = k\right]\right\}, \\
\alpha\left(x_t, \varepsilon_t, v_t^K, d_t\right) &= \sup_{a \in A} \left\{u\left(a_t, d_t, x_t, \varepsilon_t, v_t^K\right) + \beta E\left[V(s_{t+1})|x_t, a_t = a, d_t\right]\right\}.
\end{aligned}$$

We impose the following assumptions to ensure we can employ the estimation techniques that has been developed from purely discrete choice and continuous choice literature without much alteration.

ASSUMPTION DC1: *The observed data $\{a_t, d_t, x_t\}_{t=1}^T$ are the controlled stochastic processes described above with known $\beta$.*

Assumption DC2: *(Conditional Independence) The transitional distribution has the following factorization:* $p\left(x_{t+1}, \varepsilon_{t+1}, v_{t+1}^K | x_t, \varepsilon_t, v_t^K, a_t, d_t\right) = z\left(\varepsilon_{t+1}, v_{t+1}^K | x_{t+1}\right) \times p_{X'|X,A,D}\left(x_{t+1}|x_t, a_t, d_t\right)$.

Assumption DC3: *The support of* $s_t = \left(x_t, \varepsilon_t, v_t^K\right)$ *is* $X \times \mathcal{E} \times \mathcal{V}^K$, *where* $X = \{1, \ldots, J\}$ *for some* $J < \infty$ *that denotes the observable state space,* $\mathcal{E}$ *is a (potentially strict) subset of* $\mathbb{R}$ *and* $\mathcal{V}^K \subset \mathbb{R}^K$. *The distribution of* $v_t^K$ *is i.i.d. distributed across* $K-$*alternatives, denoted by* $W$, *is known, it is also independent of* $x_t$ *and is absolutely continuous with respect to some Lebesgue measure with a positive Radon-Nikodym densities* $w$. *The distribution of* $\varepsilon_t$, *denoted by* $Q$, *is known, it is also independent of* $x_t$ *and* $d_t$, *and it is absolutely continuous with respect to some Lebesgue measure with a positive Radon-Nikodym density* $q$ *on* $\mathcal{E}$.

Assumption DC4: *(Additive Separability) The per period payoff function* $u : A \times D \times X \times \mathcal{E} \times \mathcal{V}^K \rightarrow \mathbb{R}$ *can be written as* $u\left(a_t, d_t, x_t, \varepsilon_t, v_t^K\right) = u^C\left(a_t, d_t, x_t, \varepsilon_t\right) + v_t\left(d_t\right)$.

Assumption DC5: *(Monotone Choice) The per period payoff function, specific to discrete choice* $d_t$, $u_\theta^C : A \times D \times X \times \mathcal{E} \rightarrow \mathbb{R}$ *has increasing differences in* $(a, \varepsilon)$ *for all* $d, x$ *and* $\theta$, *where* $u_\theta^C$ *is specified upto some unknown parameters* $\theta \in \Theta \subset \mathbb{R}^L$.

Comments on DC1-DC5:

DC1 is standard. Similar to M2, DC2 implies that all the unobservable state variables are transitory shocks across time period. DC3 makes a simplifying assumption on the distribution of the unobservable state variables, for example, $v_t^K$ does not need to have random sampling across $K-$ alternatives, it is also straightforward to model the conditional distribution of $\varepsilon_t$ given $(x_t, d_t)$, and we do not need full independence of $\left(\varepsilon_t, v_t^K\right)$ and $x_t$ as commented in Section 2. DC4 imposes the additive separability of the choice specific unobserved shock, which is familiar from the discrete choice literature. DC5 ensures that the per period utility function for each discrete alternative satisfies the monotone choice assumption analogous to M2.4.

To illustrate how assumptions DC1 - DC5 put us on a familiar ground, consider the value function on the optimal path, which is a stationary solution to the following equation, cf. (4)

$$V_\theta(s_t) = u_\theta\left(a_t, d_t, x_t, \varepsilon_t, v_t^K\right) + \beta E\left[V_\theta(s_{t+1})\,|s_t\right],$$

where, given the sequential framework, by DC1 - DC4 $d_t = \delta_\theta\left(x_t, v_t^K\right)$ and $a_t = \alpha_\theta(x_t, \varepsilon_t, d_t)$ such that

$$\begin{aligned}
\delta_\theta\left(x_t, v_t^K\right) &= \arg\max_{1\leq k\leq K}\left\{E\left[u_\theta^C(a_t, d_t, x_t, \varepsilon_t)\,|x_t, d_t = k\right] + v_t(k) + \beta E\left[V_\theta(s_{t+1})\,|x_t, d_t = k\right]\right\}, \\
\alpha_\theta(x_t, \varepsilon_t, d_t) &= \sup_{a\in A}\left\{u_\theta^C(a, d_t, x_t, \varepsilon_t) + \beta E\left[V_\theta(s_{t+1})\,|x_t, a_t = a, d_t\right]\right\}.
\end{aligned}$$

Marginalizing out the unobserved states of the value function, under DC2, we obtain the following familiar characterization of the value functions

$$E\left[V_\theta(s_t)\,|x_t\right] = E\left[u_\theta\left(a_t, d_t, x_t, \varepsilon_t, v_t^K\right)|x_t\right] + \beta E\left[E\left[V_\theta(s_{t+1})\,|x_{t+1}\right]|x_t\right]. \qquad (108)$$

As seen previously, by DC2, that the continuation value function (onto the next time period) can be written as

$$E\left[V_\theta(s_{t+1})\,|x_t, a_t, d_t\right] = E\left[E\left[V_\theta(s_{t+1})\,|x_{t+1}\right]|x_t, a_t, d_t\right]. \qquad (109)$$

To estimate $\theta_0$, in the first step, we provide an estimate for the continuation value function. The main difference here lies in the estimation of the analogous equation to (83), where we need to nonparametrically estimate $E\left[u_\theta\left(a_t, d_t, x_t, \varepsilon_t, v_t^K\right)|x_t\right]$. Using

DC2 - DC4, we have

$$\begin{aligned} E\left[u_\theta\left(a_t,d_t,x_t,\varepsilon_t,v_t^K\right)|x_t\right] &= E\left[u_\theta^C\left(a_t,d_t,x_t,\varepsilon_t\right)|x_t\right]+E\left[v_t\left(d_t\right)|x_t\right], \\ &= \sum_{k=1}^{K}\Pr\left[d_t=k|x_t\right]E\left[u_\theta^C\left(a_t,d_t,x_t,\varepsilon_t\right)|x_t,d_t=k\right] \\ &\quad +\sum_{k=1}^{K}\Pr\left[d_t=k|x_t\right]E\left[v_t\left(d_t\right)|x_t,d_t=k\right]. \end{aligned}$$

The first term can be estimated nonparametrically using the method described in Section 2. In particular, under DC5, we can generate $\varepsilon_t$ by the relation $\widehat{\varepsilon}_t = Q_\varepsilon^{-1}\left(\widehat{F}_{A|X,D}\left(a_t|x_t,d_t\right)\right)$, where $\widehat{F}_{A|X,D}\left(a|j,k\right)$ is nonparametric estimator for $\Pr\left[a_t \leq a|x_t=j,d_t=k\right]$. Since the conditional choice probabilities are nonparametrically identified we can estimate the first term in the display above nonparametrically for any $\theta$. The second term is the selectivity term that arises from the discrete choice problem, which can be estimated nonparametrically by using Hotz and Miller's inversion theorem as in a purely discrete choice problem. Since $E\left[V_\theta\left(s_t\right)|x_t\right]$ is defined as the solution to (108), note that the transition probability in the linear equation is nonparametrically identified, we can estimate $E\left[V_\theta\left(s_t\right)|x_t\right]$ by solving a linear equation analogous to (83) once we have the estimate for $E\left[u_\theta\left(a_t,d_t,x_t,\varepsilon_t,v_t^K\right)|x_t\right]$. The continuation value in (109) can then be obtained by transforming $E\left[V_\theta\left(s_t\right)|x_t\right]$ by the a conditional expectation operator $E\left[\cdot|x_t,a_t,d_t\right]$, which differs from $\mathcal{H}$, see (84) for definition, precisely by increasing the conditioning variable to include $d_t$ in addition to $(x_t,a_t)$. The second step of the estimation procedure involves minimizing (maximizing) some criterion function to identify $\theta_0$. Obviously, one method is to construct a minimum distance criterion based on the conditional distribution function of $a_t$ given $(x_t,d_t)$, analogous to (92), as described in Section 2.2.

### 3.4 Ordered Discrete Response

The methodology to estimate a dynamic decision problem with ordered choice is the discrete counterpart of our main procedure described in Chapter 2. Practical application includes investment models where firms purchase or rent goods in discrete units, e.g. see Gowrisankaran et al. (2010).[21] Consider the following set of assumptions where the support of $a_t$ is an ordered set $\{a^1, \ldots, a^K\}$.

ASSUMPTION OC1: *The observed data for each individual* $\{a_t, x_t\}_{t=1}^{T+1}$ *are the controlled stochastic processes satisfying (81) with exogenously known* $\beta$*.*

ASSUMPTION OC2: *(Conditional Independence) The transitional distribution has the following factorization:* $p(x_{t+1}, \varepsilon_{t+1} | x_t, \varepsilon_t, a_t) = q(\varepsilon_{t+1} | x_{t+1}) p_{X'|X,A}(x_{t+1} | x_t, a_t)$ *for all* $t$*.*

ASSUMPTION OC3: *The support of* $s_t = (x_t, \varepsilon_t)$ *is* $X \times \mathcal{E}$*, where* $X = \{1, \ldots, J\}$ *for some* $J < \infty$ *that denotes the observable state space and* $\mathcal{E}$ *is a (potentially strict) subset of* $\mathbb{R}$*. The distribution of* $\varepsilon_t$*, denoted by* $Q$*, is known, it is also independent of* $x_t$ *and is absolutely continuous with respect to some Lebesgue measure with a positive Radon-Nikodym density* $q$ *on* $\mathcal{E}$*.*

ASSUMPTION OC4: *(Monotone Choice) The per period payoff function* $u_\theta : A \times X \times \mathcal{E} \to \mathbb{R}$ *has increasing differences (weakly w.r.t.* $a$*) in* $(a, \varepsilon)$ *for all* $x$ *and* $\theta$*;* $u_\theta$ *is specified upto some unknown parameters* $\theta \in \Theta \subset \mathbb{R}^L$*.*

COMMENTS ON OC1-OC4:

These conditions are essentially the analogue of M1 - M4 when $a_t$ is a discrete random variable.

[21] I thank Philipp Schmidt-Dengler for introducing to me a more general class of dynamic problems with the ordered discrete response component , which he and his co-authors in Gowrisankaran et al. (2010) are considering.

To see the intuition how one can identify the conditional value function, similar to (8), for any $\theta$ it is defined as the solution to

$$E\left[V_\theta\left(s_t\right)|x_t\right] = E\left[u_\theta\left(a_t, s_t\right)|x_t\right] + \beta E\left[E\left[V_\theta\left(s_{t+1}\right)|x_{t+1}\right]|x_t\right],$$

where

$$\begin{aligned}
a_t &= \alpha^0_{\theta_0}\left(x_t, \varepsilon_t, d_t\right) \\
&= \max_{1\leq k\leq K}\left\{u_{\theta_0}\left(a^k, x_t, \varepsilon_t\right) + \beta E\left[V_{\theta_0}\left(s_{t+1}\right)|x_t, a_t = a^k\right]\right\}.
\end{aligned}$$

Since the support of $a_t$ is finite we can write

$$E\left[u_\theta\left(a_t, s_t\right)|x_t\right] = \sum_{k=1}^{K}\Pr\left[a_t = a^k|x_t\right]E\left[u_\theta\left(a_t, s_t\right)|x_t, a_t = a_k\right],$$

where the potential issue again lies in the fact that we do not observe $\varepsilon_t$. But analogously to the case with continuous control, since the policy function is weakly monotone in $\varepsilon_t$, given the assumed distribution of $\varepsilon_t$ we can identify the conditional mean of the per-period payoff function by using the quantile invariance property between $(a_t, \varepsilon_t)$. In particular, for $k > 1$ let $\mathcal{I}_k = \left[Q_\varepsilon^{-1}\left(1 - F_{A|X}\left(a_k|x_t\right)\right), Q_\varepsilon^{-1}\left(1 - F_{A|X}\left(a_{k-1}|x_t\right)\right)\right]$, we have

$$\begin{aligned}
E\left[u_\theta\left(a_t, s_t\right)|x_t, a_t = a^k\right] &= E\left[u_\theta\left(a_t, s_t\right)|x_t, \varepsilon_t \in \mathcal{I}_k\right] \\
&= \frac{\int_{\mathcal{I}_k} u_\theta\left(a^k, x_t, \varepsilon\right)Q_\varepsilon\left(d\varepsilon\right)}{F_{A|X}\left(a^k|x_t\right) - F_{A|X}\left(a^{k-1}|x_t\right)},
\end{aligned}$$

and for $k = 1$, let $\mathcal{I}_1 = \left[Q_\varepsilon^{-1}\left(1 - F_{A|X}\left(a_1|x_t\right)\right), Q_\varepsilon^{-1}\left(0\right)\right]$

$$E\left[u_\theta\left(a_t, s_t\right)|x_t, a_t = a^1\right] = \frac{\int_{\mathcal{I}_1} u_\theta\left(a^1, x_t, \varepsilon\right)Q_\varepsilon\left(d\varepsilon\right)}{F_{A|X}\left(a^1|x_t\right)}.$$

Various nonparametric estimators discussed in the previous sections, for example the frequency estimator, can be used to estimate $\Pr\left[a_t = a^k | x_t\right]$ and $F_{A|X}\left(a^k | x_t\right)$ for $k = 1, \ldots, K$. Since we know $Q_\varepsilon$, we can estimate $\{\mathcal{I}_k\}$ and use them to estimate $\int_{\mathcal{I}_k} u_\theta\left(a^k, x_t, \varepsilon\right) Q_\varepsilon\left(d\varepsilon\right)$ for each $k$. All the remaining nonparametric estimators required to estimate the continuation value function are just the transitional probabilities of the observable state variable in the next period conditioning on this period's state (with or without the control). We can then approximate the model implied policy function as before.

## 3.5 Conclusion

In this chapter we show that the methodology proposed in the first two chapters has a much wider applicability than the given frameworks. We illustrate this by showing that various types of dynamic models in economics can be estimated through a familiar two-step approach, where differences between various models requires different modelling assumptions necessary to ensure we can identify the conditional value functions.

As seen from Chapter 1, although we only considered a single agent problem with observable state variable with finite support, it is straightforward to generalize the framework to allow for strategic interactions between players as well as observable state variables with continuous distributions.

# 4 Consistent Estimation of an Identified Optimization Model

## 4.1 Introduction

Bajari, Benkard and Levin (2007), henceforth BBL, propose a methodology to estimate a large class of structural dynamic models which has an extensive applicability. The motivation behind the construction of their estimator is conceptually appealing as it relies directly on the necessary conditions of an economic equilibrium. They suggest a forward simulation method which is not only easy to program, it also has a computationally attractive feature by making use of the linear structure of the problem. In addition, in the case that the model is not identified, they also propose a set estimator to estimate the partially identified model.

However, it is infeasible to approximate or provide an analytical expression for either the population or its empirical analogue of their criterion function even in a simple static optimization problem as they rely on an uncountably infinite set of inequality constraints, essentially indexed by functions. In practice only a strict subclass of such inequalities are considered. This may lead to the loss of identification since we are not using all of the relevant constraints imposed by the definition of an equilibrium. To the best of my knowledge, as suggested by BBL, all applications of their methodology only consider the class of alternative policies which are *translation shifts from the true*, for example see Sweeting (2007), Ryan (2009) and Santos (2009) amongst many others. We provide an example where the criterion functions constructed through this class of inequality constraints are not capable of consistently estimating an identified model. Although we do not provide specific examples for a partially identified model, we expect analogous findings to exist. In addition, most applications of BBL methodology only consider point estimation of the parameter of interest, which is attainable even if the objective function does not have a unique optimizer in the limit.

For a different reasoning, the swapping nested fixed point estimator of Aguirregabiria and Mira (2007) which has been used to estimate a class of Markovian games may also fail to consistently estimate an identified model. Pesendorfer and Schmidt-Dengler (2009) show that even when the observed data are generated from a single equilibrium of a game with multiple equilibria, the functional operator of Aguirregabiria and Mira may have multiple fixed points and if the equilibrium point is unstable their iterative method will generally lead to an inconsistent estimator.

Since the notations used in the study of many closely related Markov decision models vary in the literature, throughout this note we follow the notations used in BBL when possible.[22] For notational simplicity, we also focus on the decision problem of a single agent, an extension allowing for strategic interactions to a popular class of Markovian games considered is straightforward.

Next, we describe the class of Markov decision processes of interest in Section 2, then proceed to define the identification concept and summarizes the BBL methodology in Section 3 and 4. An analytical example that shows one can lose identification by using a BBL type of criterion is presented in Section 5. Section 6 concludes.

## 4.2 The Model

We now describe the dynamic models which are popular in Industrial Organization, amongst others fields, that rely on Rust (1987) conditional independence assumption. Although we do not restrict our attention to the discrete choice framework, much of this growing literature builds on the work of Rust (1987) and Hotz and Miller (1993) which fall under such setting. In particular, following Hotz and Miller, there has many subsequent two-step estimators proposed for dynamic discrete choice problems, for

[22] Of particular importance are the differences are the notations regarding the state variables and the policy function. Some authors use $(x_t, \varepsilon_t)$ to denote observed and unobserved state variables respectively. And more recently, motivated by applications to games, the policy (i.e. best response) function is often seen denoted by $\sigma$.

example see Hotz et al. (1994), Aguirregabiria and Mira (2002,2007), BBL, Pesendorfer and Schmidt-Dengler (2008), Bajari et al. (2009) and Srisuma and Linton (2009) amongst others. More recently, there are also general methods which are able estimate closely related dynamic models with continuous control, but otherwise rely on quite similar sets of modelling assumptions, see BBL, Hong and Shum (2010) and Srisuma (2010).

For each time period $t = 1, \ldots, \infty$, the economic agent makes a decision $a_t \in A$ given the state variables $(s_t, \nu_t)$, where $s_t \in S$ is observed by the econometrician whilst $\nu_t \in \mathcal{V}$ is known only to the agent. The agent problem is to choose a decision rule $\{\sigma_\tau\}_{\tau=t}^{\infty}$, where each $\sigma_t$ belongs to a set of functions $\Sigma$ which maps $S \times \mathcal{V}$ to $A$, to solve

$$\max_{\{\sigma_\tau\}_{\tau=t}^{\infty}} \mathbb{E}\left[\sum_{\tau=t}^{\infty} \beta^{\tau-t} \pi\left(\sigma_\tau, s_\tau, \nu_\tau\right) \middle| s_t, \nu_t\right], \tag{110}$$

where $\pi$ denotes the per period payoff function and $\beta \in (0,1)$ is the discounting rate. Under some regularity conditions the optimal time invariant policy function $\sigma$ exist, see Rust (1994) and the references therein for details. Assume further the conditional independence of Rust (1987), and that the agent has perfect expectations on the conditional laws of $s_{t+1}$ conditioning on $(a_t, s_t)$, denoted by $P\left(s_{t+1}|a_t, s_t\right)$, and of $\nu_t$ given $s_t$, denoted by $G\left(\nu_t|s_t\right)$. Then the observed actions $a_t$ is equal to $\sigma\left(s_t, \nu_t\right)$ such that

$$\sigma\left(s_t, \nu_t\right) = \arg\max_{a \in A}\left\{\pi\left(a, s_t, \nu_t\right) + \beta \int V\left(s'; \sigma\right) dP\left(s'|\sigma\left(s_t, \nu_t\right), s_t\right)\right\} \quad \text{for all } t \geq 1,$$

where the *ex-ante (or conditional) value function* $V$ satisfies

$$V\left(s_t; \sigma\right) = \mathbb{E}\left[\sum_{\tau=t}^{\infty} \beta^{\tau-t} \pi\left(\sigma\left(s_\tau, \nu_\tau\right), s_\tau, \nu_\tau\right) \middle| s_t\right]. \tag{111}$$

The integral $\int V\left(s'; \sigma\right) dP\left(s'|a_t, s_t\right)$, which is equal to $\mathbb{E}\left[V\left(s_{t+1}; \sigma\right)|a_t, s_t\right]$, is sometimes called *continuation value function.*

We note that we are dealing with a stationary problem. The time index $t$ is arbitrary and just denotes that various functions of interest are functions of random variables. Below, we also use $\tau$ to index future values when defining alternative policies to stress this point.

## 4.3 Identification

Thus far there has been relatively little identification results in Markov decision models of this kind. Rust (1994) showed that without any restrictions, these decision models are nonparametrically not identified. Magnac and Thesmar (2002) also show some negative results on the nonparametric identification of a class of single agent discrete choice problems. In the study of their discrete Markovian games, Pesendorfer and Schmidt-Dengler (2008) provide conditions to identify their parametric model. Bajari et al. (2009) extends the results of Pesendorfer and Schmidt-Dengler to the case when the support of the observable state variable can include intervals. Generally, it is fair to say that the identification problem for many of these parametric models are not well understood and identification is often assumed. Here we provide the formal definition of an identified model for the special case where the only unknown parameter in the model is the finite dimensional parameter indexing the payoff function.

For a single agent problem, the econometric model of such decision processes described in the previous section can be formally represented by the set of primitives $(\pi, \beta, G, P)$. Under some regularity assumptions, the data generating process $P$ is nonparametrically identified so we assume it to be known. In this literature, the value of $\beta$ is often assumed and $(\pi, G)$ are parametrically specified. The parametric form of $G$ is essential for the methodologies cited above. For policy purposes, the main objective is then to estimate the structural parameter indexing $\pi$, we denote this by

$\theta \in \Theta \subset \mathbb{R}^M$. For simplicity, in what follows, we assume the knowledge of $(\pi, \beta, G, P)$ upto some finite dimensional parameterization only on $\pi$. The identification problem of this parametric Markov decision processes can be stated in a familiar way using the following definitions.

DEFINITION 4.1: The *reduced-form* of an MDP model is the agent's optimal decision rule $\sigma \in \Sigma$.

DEFINITION 4.2: The *structure* of an MDP model is the map $\Lambda : \Theta \to \Sigma$ such that $\Lambda(\theta) = \sigma(s_t, \nu_t; \theta)$, where

$$\sigma(s_t, \nu_t; \theta) = \arg\max_{a \in A} \left\{ \pi(a, s_t, \nu_t; \theta) + \beta \int V\left(s'; \sigma(\cdot, \cdot; \theta); \theta\right) dP\left(s' | \sigma(s_t, \nu_t; \theta), s_t\right) \right\}.$$

We comment here that the notation for the conditional value function $V(s; \sigma; \theta)$ allows us to define the expected discounted value generated from a particular policy $\sigma$, that may but need not depend on $\theta$ (cf .(111)), so that

$$V(s_t; \sigma; \theta) = \mathbb{E}\left[ \sum_{\tau=t}^{\infty} \beta^{\tau-t} \pi\left(\sigma(s_\tau, \nu_\tau), s_\tau, \nu_\tau; \theta\right) \middle| s_t \right].$$

DEFINITION 4.3: The parameter points $\theta_1$ and $\theta_2$ are *observationally equivalent* if $\sigma(s_t, \nu_t; \theta_1) = \sigma(s_t, \nu_t; \theta_2)$ *a.s.*

DEFINITION 4.4: The model is *well specified* if the data is generated according to a decision rule $\sigma(s_t, \nu_t; \theta)$ for some $\theta \in \Theta$.

DEFINITION 4.5: Let $\Theta_0(\Lambda)$ be a set of observationally equivalence classes so that $\Theta_0(\Lambda)$ is a collection of sets $\dot{\theta}, \theta \in \Theta$ such that $\theta \in \dot{\theta}$ if and only if $\theta_1$ and $\theta$ are observationally equivalent. A well specified model is identified if and only if $\dot{\theta} = \{\theta_0\}$.

## 4.4 BBL Methodology

The estimator proposed by BBL is defined to satisfy the necessary conditions implied by the optimality conditions of an economic equilibrium. Suppose that the data is generated according to the true parameter $\theta_0 \in \Theta$. In what follows, it will be useful, to avoid potential confusion, to denote the underlying policy function $\sigma(\cdot,\cdot;\theta_0)$ by $\sigma_0(\cdot,\cdot)$.

By definition of the optimal policy, for any $\theta \in \Theta$ and any alternative Markov decision rule $\sigma' = \{\sigma'_\tau\}_{\tau=1}^{\infty}$, the solution to the sequential problem must satisfy (see (110))

$$\mathbb{E}\left[\sum_{\tau=t}^{\infty}\beta^{\tau-t}\pi\left(\sigma\left(s_\tau,\nu_\tau;\theta\right),s_\tau,\nu_\tau;\theta\right)\middle| s_t,\nu_t\right] \geq \mathbb{E}\left[\sum_{\tau=t}^{\infty}\beta^{\tau-t}\pi\left(\sigma'_\tau\left(s_\tau,\nu_\tau\right),s_\tau,\nu_\tau;\theta\right)\middle| s_t,\nu_t\right].$$

Integrating out the unobserved state variable from the expression above leads to the following inequality for the *model implied* conditional value function (cf. (111)), for all $(\sigma',\theta)$

$$V\left(s_t;\sigma\left(\cdot,\cdot;\theta\right);\theta\right) \geq V\left(s_t;\sigma';\theta\right). \tag{112}$$

The criterion function proposed by BBL is constructed based on the set of inequalities above, which are implied by the optimality conditions of an equilibrium. Let $x$ index the set of all *equilibrium conditions* (or also called *inequalities*) $\mathcal{X}$, we elaborate on this terminology used in BBL after introducing the function $g$ below. So $x$ denotes a particular pair $(s,\sigma')$. For an identified model, for any $\theta \in \Theta$, define $g(x;\theta)$ as follows

$$g\left(x;\theta\right) = V\left(s;\sigma_0;\theta\right) - V\left(s;\sigma';\theta\right).$$

Therefore $g(x;\theta)$ represents an equilibrium constraint in the sense that we must have $g(x;\theta_0) \geq 0$ for any $x$ in $\mathcal{X}$. In what follows, since we are analyzing the BBL procedure we continue to make use of their terminologies. Formally, the criterion function which

BBL proposed is represented by a pair $(\mathcal{X}, H)$, based on (112), defined as follows

$$Q(\theta) = \int (\min\{g(x;\theta), 0\})^2 \, dH(x),$$

where $H$ is a distribution over $\mathcal{X}$. Clearly the criterion function in the display must satisfy the condition $Q(\theta_0) = 0$.

However, $\mathcal{X}$ is an uncountably large set even if $S$ is finite. Although finiteness of $S$ is not an unusual assumption, allowing it to be uncountable is conceptually simple. It is more difficult to deal with the space of functions of alternative policies. To proceed, BBL suggest the practioner to consider a particular strict subset of $\mathcal{X}$. For some $\mathcal{X}_{\mathcal{E}} \subset \mathcal{X}$, we define the criterion function based on $\mathcal{X}_{\mathcal{E}}$ as

$$Q_{\mathcal{E}}(\theta) = \int (\min\{g(x_\epsilon;\theta), 0\})^2 \, dH_{\mathcal{E}}(x_\epsilon),$$

where we make the dependence on $\epsilon$ and $\mathcal{E}$ explicit. So a particular inequality $x_\epsilon$ denotes a certain optimality condition that belongs to $\mathcal{X}_{\mathcal{E}}$, a set indexed by $\epsilon$. To complete the construction of $Q_{\mathcal{E}}$, $H_{\mathcal{E}}$ denotes the underlying disribution of $x_\epsilon$. We next formally define the class of inequalities derived through *translating shift of the true policies* that we denote by $\mathcal{X}_{\mathcal{E}_1}$, this class of inequalities is informally introduced in BBL which consequently has been employed by most of their known applications. We note that the notations used to define $(\mathcal{X}_{\mathcal{E}_1}, H_{\mathcal{E}_1})$ reduce significantly for the example in the next section, where we analyze a much simpler setup, so the reader less familiar with BBL methodology may first wish to skip to that part for the intuition behind the setup. A typical element $x_\epsilon$ in $\mathcal{X}_{\mathcal{E}_1} = \{(s, \sigma') : s \in S \text{ and } \sigma' = \sigma_0 \oplus \epsilon \text{ for } \epsilon \in (\mathit{Support}(\Delta))^\infty \subset l^\infty\}$, where $l^\infty$ is the real sequence space and $(\mathit{Support}(\Delta))^\infty$ is the countably infinite Cartesian products of $\mathit{Support}(\Delta)$. Although $s$ requires no explanation, for each $\epsilon$ we define $\sigma_0 \oplus \epsilon$ to be a particular alternative policy defined through a particular infinite sequence of linear shifts $\{\epsilon_\tau\}_{\tau=1}^{\infty}$ so that

$$(\sigma_0 \oplus \epsilon)(s_\tau, \nu_\tau) = \sigma_0(s_\tau, \nu_\tau) + \epsilon_\tau \text{ for any } \tau = 1, \ldots, \infty.$$

So $(\sigma_0 \oplus \epsilon)$ is a sequence of functions representing a particular decision rule that differs from the true policy by a sequence of translations represented by $\{\epsilon_\tau\}_{\tau=1}^{\infty}$; more specifically, for each period $\tau$ it differs from the true policy by a translation $\epsilon_\tau$. The measure $H_{\mathcal{E}_1} = M_{\mathcal{E}_1}(S) \times N_{\mathcal{E}_1}(\Delta)$ is the product measure where $M_{\mathcal{E}_1}(S)$ denotes a uniformly distributed measure on $S$ (suppose $S$ is bounded) and $N_{\mathcal{E}_1}(\Delta)$ denotes the product measure for countably infinite product measurable space that generates the sequence of random samples $\{\epsilon_\tau\}_{\tau=1}^{\infty}$ from $\Delta$. Therefore the limiting criterion considered in BBL can be formally represented by the pair $(\mathcal{X}_{\mathcal{E}_1}, H_{\mathcal{E}_1})$.

The sample analogue of $Q_{\mathcal{E}_1}(\theta)$ can be constructed in practice using the forward simulation procedure outlined by BBL. Then one can construct a point or a set estimator based minimizing $Q_{\mathcal{E}_1}(\theta)$ depending on whether one thinks the model is identified or not.

## 4.5 An Example

The difficulty we face in trying to understand the type of criterion functions based on the optimality conditions implied by an equilibrium is that, it is generally unclear how one can mathematically show what happens to the value of $Q(\theta)$ for $\theta \in \Theta \backslash \{\theta_0\}$. Even if it may be true that $Q(\theta) = 0$ if and only if $\theta = \theta_0$, it is plausible that by considering a subset of all possible policies one may be able to find $\theta \neq \theta_0$ such that $Q_{\mathcal{E}}(\theta) = 0$ for some $\mathcal{E}$; clearly such criterion function cannot be used to consistently estimate $\theta_0$. The last statement is particularly relevant in practice since one cannot easily compare the inequality restrictions of such a large class of alternative policies, or even know what kind of alternative policies are sufficient to ensure we can consistently estimate a point identified model.

To illustrate this point we consider a much simpler optimization problem that also belongs to the class of models considered in BBL. In particular, we take a static optimization problem, which corresponds to the case that $\beta = 0$. Ignoring the presence of observable state variables, specify the payoff function to be

$$\pi(a, \nu; \theta) = -a^2 + 2\theta a \nu,$$

so $a$ and $\nu$ are values from the support of the control and state variables respectively. Let $G$ be some distribution for $\nu_t$ with zero mean.

It is easy to see that the (optimal) policy function, now reduces to $\sigma(\nu; \theta)$, is $\theta\nu$ for all $\theta, \nu$. Imposing $\Theta = \mathbb{R}^+$ ensures that the policy function will be increasing in the state variable, hence satisfying the *Monotone Choice* assumption that is essential to BBL's simulation method. Notice that if $\theta \neq \theta'$ then $\sigma(\nu_t; \theta) \neq \sigma(\nu_t; \theta')$ *a.s.*, therefore this model is identified so long that it is well specified. Then given a random sample of $\{a_t\}_{t=1}^T$ generated from $\sigma_0(\nu)$, for some $\theta_0 \in \Theta$, along with some standard regularity conditions one can construct a consistent estimator for $\theta_0$ by maximum likelihood or other minimum distance approach based on the moment condition below (the latter might be a preferred option for a more complicated dynamic model); as shown in Srisuma (2010), the uniqueness of the policy function with respect to $\theta$ implies

$$\mathbb{E}[\mathbf{1}[a_t \leq a] - F_A(a; \theta)] = 0 \text{ for all } a \in A \text{ if and only if } \theta = \theta_0, \qquad (113)$$

where $F_A(\cdot; \theta)$ is the distribution function of $\sigma(\nu_t; \theta)$.

However, we now show that the moment inequality approach of BBL may lead to set estimators that, in the limit, will only converge to a non-singleton set containing

$\theta_0$.

(I) TRANSLATION SHIFT:

We consider the criterion function constructed from $(\mathcal{X}_{\mathcal{E}_1}, H_{\mathcal{E}_1})$ described in the previous section. In this simpler setting , the class of inequalities $\mathcal{X}_{\mathcal{E}_1}$ is simply $\{\sigma' : \sigma' = \sigma_0 + \epsilon \text{ for } \epsilon \in \mathit{Support}(\Delta)\}$ and the distribution $H_{\mathcal{E}_1}$ is $\Delta$.[23] Note that an inequality $x_\epsilon$ in $\mathcal{X}_{\mathcal{E}_1}$ has a unique correspondence with each $\epsilon$ in the support of $\Delta$. For any $\epsilon$ and $\theta$, it is easy to see that

$$\pi\left(\sigma_0\left(\nu_t\right), \nu_t, \theta\right) - \pi\left(\sigma_0\left(\nu_t\right) + \epsilon, \nu_t, \theta\right) = \epsilon^2 + 2\epsilon\nu_t\left(\theta_0 - \theta\right).$$

To derive $g$, since $\nu_t$ has zero mean, integrating out $\nu_t$ in the display above leads to

$$g\left(x_\epsilon; \theta\right) = \epsilon^2 \geq 0.$$

Note that we can also write $Q_{\mathcal{E}_1}(\theta) = \int \left(\min\left\{\epsilon^2, 0\right\}\right)^2 dF_\Delta(\epsilon)$, where $F_\Delta$ denotes the distribution function which corresponds to $\Delta$. Clearly $Q_{\mathcal{E}_1}(\theta) = 0$ for all $\theta \in \Theta$ and any distribution $\Delta$ as long as $\epsilon$ is not degenerate at zero (which would then not represent an alternative policy); i.e., in this example, this class of alternative policies has no identifying power for $\theta_0$.

(II) MULTIPLICATIVE SCALE:

As another illustration, let's consider another class of policies, based on a multiplicative scale of the true policy. In particular let $\mathcal{X}_{\mathcal{E}_2} = \{\sigma' : \sigma' = \epsilon\sigma_0 \text{ for } \epsilon \in \mathit{Support}(\Delta)\}$ where $\Delta$ is the uniform distribution on $(0, 1)$, the non-negative support is chosen to

[23] As seen previously, it is conceptually straightforward to provide a precise formulation of $(\mathcal{X}_{\mathcal{E}_1}, H_{\mathcal{E}_1})$ to include the observable states and/or number of players for a game by using the direct product with respect to respective sets and measures.

ensure that the alternative policies are also monotone on $\mathcal{V}$. For any $x_\epsilon$ that belongs to $\mathcal{X}_{\mathcal{E}_2}$ and $\theta \in \Theta$ we have

$$g(x_\epsilon;\theta) = -\theta_0 (1-\epsilon)((1+\epsilon)\theta_0 - 2\theta)\int \nu^2 dG(\nu).$$

Therefore, when $\theta \geq \theta_0$ we have $g(x_\epsilon;\theta) \geq 0$, suppose further that $\mathbb{E}\nu_t^2 < \infty$, then any $\theta \in [\theta_0, \infty)$ will imply that $Q_{\mathcal{E}_2}(\theta) = 0$. On the other hand it is easy to show that any $\theta \in (0, \theta_0)$ will imply that $Q_{\mathcal{E}_2}(\theta) > 0$. The criterion function based on this class of alternative policies can only identify a set, although we can show that increasing the support of $\epsilon$ can reduce the identify set to a point, see below.

We can also try to find the class of alternative policies which will ensure that the criterion constructed from $(\mathcal{X}_{\mathcal{E}}, H_{\mathcal{E}})$ has a unique minimum. Recall that in this case, $\mathcal{X}$ corresponds to is the set of functions $\sigma \in \Sigma$ such that $\sigma : \mathcal{V} \rightarrow A \subset \mathbb{R}$. For any $x$ in $\mathcal{X}$, which is a function $\sigma(\cdot)$, it can be shown from simple algebra that for any $\theta \in \Theta$

$$g(x;\theta) = -(\theta-\theta_0)^2 \mathbb{E}\left[\nu_t^2\right] + \mathbb{E}\left[(\theta\nu_t - \sigma(\nu_t))^2\right].$$

In this case, the class of alternative policies which is a multiplicative scale from the true can ensure that we can construct criterion functions with a unique minimizer at $\theta_0$ so long that the support of $\epsilon$ is sufficiently large. To see this, note that the inequality we require is

$$g(x;\theta) < 0 \;\Leftrightarrow \mathbb{E}\left[(\theta\nu_t - \sigma(\nu_t))^2\right] < \mathbb{E}\left[\nu_t^2\right](\theta-\theta_0)^2.$$

For any $\theta = \theta_0 + \delta$, we see that by letting $\sigma(\nu_t) = (\theta - \eta)\nu_t$, for any $|\eta| < |\delta|$, will ensure the inequality above. This means that when $\Theta$ is a compact subset of $\mathbb{R}^+$ containing $\theta_0$, choosing $\mathcal{X}_{\mathcal{E}} = \{\sigma' : \sigma' = \epsilon\sigma_0 \text{ for } \epsilon \in \mathbb{R}\}$ and $H_{\mathcal{E}}$ to be any continuous distribution with full support on the real line (e.g. a standard normal) will be sufficient to ensure that

$Q_{\mathcal{E}}$ has a well separated minimum at $\theta_0$, which is implied by the high level assumptions of BBL (see their Assumption S2). In this case we would expect to be able to obtain consistent estimator for $\theta_0$ from minimizing $Q_{\mathcal{E}}(\theta)$.

## 4.6 Conclusion

We somewhat formalize the construct of the criterion functions introduced by BBL through the pair $(\mathcal{X}, H)$ of inequalities and its distribution. It is infeasible to work with $(\mathcal{X}, H)$ and we consider $(\mathcal{X}_{\mathcal{E}}, H_{\mathcal{E}})$, which represents the criterion functions that rely on a smaller set of inequalities. The example above reveals that applications of the moment inequality approach introduced by BBL may incorrectly lead one to believe that the model is only partially identified when it is actually point identified. In particular we show in a simple setup that when $\mathcal{X}_{\mathcal{E}}$ is the class of alternative policies which is a translation from the true policy has no identifying power of an identified model. We also show that the class of alternative policies which is a multiplicative scale of the true policy, which although is never used in practice, can succeed in constructing an appropriate criterion function to estimate the identified model when the range of the scaling factor is sufficiently large. We stress here that we are not implicating any relative potency between the two classes of inequalities in general. Also we would expect analogous results for a partially identified model where a particular criterion function $(\mathcal{X}_{\mathcal{E}}, H_{\mathcal{E}})$ may at best can consistently estimate a strictly larger bounds than the identified set.

The practical consequence is potentially serious since most applications of BBL use the point estimation method. In finite samples, when the Monte Carlo integration has yet to converge, various optimization techniques will produce point estimates that may not be informative at all. Fortunately, all known applications of BBL's methodology are used to estimate structural dynamic optimization problems that are more complicated than the simple example provided above, therefore alternative policies that introduce "noise" into a highly nonlinearly problem may give a better hope of preserving point identification in a point identified model. However, for general dynamic models of interest, the policy functions typically do not have closed form expression thereby preventing us from doing any kind of analytical analysis and the intuition that one may risk losing point identification by applying BBL type estimators cannot be ruled out.

The intuition behind this finding is somewhat similar to the identification issue studied in Domínguez and Lobato (2004), where they show one can lose identification in an identified conditional moment restriction model when one arbitrarily turns it into some unconditional moment restrictions with a finite number of instruments. Domínguez and Lobato also propose a minimum distance estimator that uses all of the available information to overcome this problem.

We end this note with two remarks on the risk of losing identification by using the criterion functions described above. First, this problem can potentially be alleviated by integrating over larger classes of policies. Second, such risk may be eliminated entirely by using alternative estimation methods based on different objective functions. For example, for the same class of structural dynamic models considered by BBL, if the model is identified the minimum distance estimator analogous to that found in Srisuma (2010), which relies on the generalization of (113), will be consistent under some regularity conditions.

# References


[1] Ackerberg, D., L. Benkard, S. Berry, and A. Pakes (2005): "Econometric Tools for Analyzing Market Outcome," *Handbook of Econometrics*, vol. 6, eds. J. Heckman and E. Leamer. North-Holland.

[2] Aguirregabiria, V. (2008): "Another Look at the Identification of Dynamic Discrete Decision Processes: An Application to Retirement Behavior," *Journal of Business and Economic Statistics*, forthcoming.

[3] Aguirregabiria, V., and P. Mira (2002): "Swapping Nested Fixed Point Algorithm: A Class of Estimators for Discrete Markov Decision Models," *Econometrica*, **70**, 1519-1543.

[4] Aguirregabiria, V., and P. Mira (2007): "Sequential Estimation of Dynamic Discrete Games," *Econometrica*, **75**, 1-53.

[5] Aguirregabiria, V., and P. Mira (2008): "Dynamic Discrete Choice Structural Models: A Survey," *Journal of Econometrics,* forthcoming.

[6] Aitchison, J., and C. G. G. Aitken (1976): "Multivariate Binary Discrimination by the Kernel Method," *Biometrika,* **63**, 413-420.

[7] Ai. C., and X. Chen (2003): "Efficient Estimation of Models with Conditional Moment Restrictions Containing Unknown Functions," *Econometrica*, **71**, 1795-1844.

[8] Andrews, D.W.K. (1994a): "Asymptotics for Semiparametric Econometric Models via Stochastic Equicontinuity," *Econometrica*, **62**, 43-72.

[9] Andrews, D.W.K. (1994b): "Empirical Processes Methods in Econometrics," *Handbook of Econometrics*, vol. 4, eds. R.F. Engle and D. McFadden. North-Holland.

[10] Andrews, D.W.K. (1995): "Nonparametric Kernel Estimation for Semiparametric Econometric Models," *Econometric Theory,* **11**, 560-596.

[11] Andrews, D.W.K. (2005): "Higher-order Improvements of the Parametric Bootstrap for Markov Processes," *Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg*, eds. D.W.K. Andrews and J.H. Stock. Cambridge Press.

[12] Arak, T. and A. Zaizsev (1988): "Uniform Limit Theorems for Sums of Independent Random Variables," *Publ. Steklov Math. Institute* **1**.

[13] Arcidiacono, P. and R.A. Miller (2008): "CCP Estimation of Dynamic Discrete Choice Models with Unobserved Heterogeneity," *Working Paper.*

[14] Arcones, M., and E. Giné, (1992): "On the Bootstrap of M-estimators and Other Statistical Functionals," *Exploring the Limits of Bootstrap*, eds. R. LePage and L. Billard. Wiley, New York.

[15] Athey, S.C. (2002): "Monotone Comparative Statics Under Uncertainty," *Quarterly Journal of Economics*, **117**, 187-223.

[16] Bajari, P., C. L. Benkard, and J. Levin (2007): "Estimating Dynamic Models of Imperfect Competition," *Econometrica*, **75**, 1331-1370.

[17] Bajari, P., V. Chernozhukov, H. Hong and D. Nekipelov (2009): "Semiparametric Estimation of a Dynamic Game of Incomplete Information," *Working paper.*

[18] Bellman, R. E. (1957): *Dynamic Programming,* Princeton University Press.

[19] Berry, S., and A. Pakes (2002): "Estimation from the Optimality Conditions for Dynamic Controls," *Mimeo.*

[20] Bose, A. (1988): "Edgeworth Correction by Bootstrap in Autoregressions," *Annals of Statistics*, **16**, 1709-1722.

[21] Bose, A. (1990): "Bootstrap in Moving Average Models," *Annals of the Institute of Statistical Mathematics*, **42**, 753-768.

[22] Bosq, D. (1996): *Nonparametric Statistics for Stochastic Processes*, Springer.

[23] Brown, D. J. and M. H. Wegkamp (2002): "Weighted Minimum Mean-Square Distance from Independence Estimation," *Econometrica,* **70**, 2035-2051.

[24] Carrasco, M., and J.P. Florens (2000): "Generalization Of Gmm To A Continuum Of Moment Conditions," *Econometric Theory*, **16**, 797-834.

[25] Carrasco, M., J.P. Florens and E. Renault (2007): "Linear Inverse Problems and Structural Econometrics Estimation Based on Spectral Decomposition and Regularization," *Handbook of Econometrics*, vol. 6, eds. J. Heckman and E. Leamer. North-Holland.

[26] Chen, X., and X. Shen (1998): "Sieve Extremum Estimates for Weakly Dependent Data," *Econometrica*, **66**, 289-314.

[27] Chen, X., O.B. Linton and I. van Keilegom (2003): "Estimation of Semiparametric Models when the Criterion Function is not Smooth," *Econometrica*, **71**, 1591-1608.

[28] Chen, X., and D. Pouzo (2008): "Estimation of Nonparametric Conditional Moment Models with Possibly Nonsmooth Moments," *Journal of Econometrics*, forthcoming.

[29] Domínguez, M.A., and I.N. Lobato (2004): "Consistent Estimation of Models defined by Conditional Moment Restrictions," *Econometrica*, **72**, 1601-1615.

[30] Doraszelski, U. and M. Satterthwaite (2007): "Computable Markov-Perfect Industry Dynamics: Existence, Purification and Multiplicity," *Mimeo*, Harvard University.

[31] Eckstein, Z. and K.I. Wolpin (1989): "The Specification and Estimation of Dynamic Stochastic Discrete Choice Models: A Survey," *Journal of Human Resources*, **24**, 562-598.

[32] Efron, B. (1979): "Bootstrap Methods: Another look at the Jacknife," *Annals of Statistics*, **7**, 1-26.

[33] Efron, B., and R.J. Tibshirani (1993): *An Introduction to the Bootstrap,* Chapman & Hall, New York.

[34] Fan, J., (1992): "Design-adaptive nonparametric regression," *Journal of the American Statistical Association*, **87**, 998-1004.

[35] Fan, J. (1993): "Local Linear Regression Smoothers and their Minimax Efficiencies," *Annal of Statistics,* **21**, 196-216.

[36] Hajivassiliou, V., and P. Ruud (1994): "Classical Estimation Methods for LDV Models Using Simulation," *Handbook of Econometrics*, vol. 4, eds. R.F. Engle and D.L. McFadden. North-Holland.

[37] Hall, P. (1992): *The Bootstrap and Edgeworth Expansion,* Springer-Verlag, New York.

[38] Hansen, L.P. and K.J. Singleton (1982): "Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models," *Econometrica*, **50**, 1269-1286.

[39] Härdle, W., and O.B. Linton (1994): "Applied Nonparametric Methods," *Handbook of Econometrics*, vol. 4, eds. R.F. Engle and D.L. McFadden. North-Holland.

[40] Härdle, W., J. Horowitz and J.P. Kreiss (2003): "Bootstrap methods for time series," *International Statistical Review*, **71**, 435-459.

[41] Heckman, J. and S. Navarro (2007): "Dynamic Discrete Choice and Dynamic Treatment Effects," *Journal of Econometrics*, **136**, 341-396.

[42] Hong, H. and M. Shum (2009): "Pairwise-Difference Estimation of a Dynamic Optimization Model," forthcoming, *Review of Economic Studies.*

[43] Horowitz, J. (1998): "A Smoothed Maximum Score Estimator for the Binary Response Model," *Econometrica*, **60**, 1327-1352.

[44] Horowitz, J. (2001): "The Bootstrap," *Handbook of Econometrics*, vol. 5, eds. J. Heckman and E. Leamer. North-Holland.

[45] Hotz, V., and R.A. Miller (1993): "Conditional Choice Probabilities and the Estimation of Dynamic Models," *Review of Economic Studies*, **60**, 497-531.

[46] Hotz, V., R.A. Miller, S. Sanders and J. Smith (1994): "A Simulation Estimator for Dynamic Models of Discrete Choice," *Review of Economic Studies*, **61**, 265-289.

[47] Hu. Y. and M. Shum (2009): "Nonparametric Identification of Dynamic Models with Unobserved State Variables," *Working Paper.*

[48] Ichimura, H. and S. Lee (2006): "Characterization of the Asymptotic Distribution of Semiparametric M-Estimators," *CEMMAP Working Paper.*

[49] Jofre-Bonet, M. and M. Pesendorfer (2003): "Estimation of a Dynamic Auction Game," *Econometrica*, **71**, 1443-1489.

[50] Judd, K. (1998): *Numerical Methods in Economics,* MIT Press, Cambridge MA.

[51] Kasahara, H. and K. Shimotsu (2008a): "Pseudo-Likelihood Estimation and Bootstrap Inference for Structural Discrete Markov Decision Models," *Journal of Econometrics*, **146**, 92-106.

[52] Kasahara, H. and K. Shimotsu (2008b): "Nonparametric Identification of Finite Mixture Models of Dynamic Discrete Choices," *Econometrica*, **77**, 135-175.

[53] Komunjer, I. and A. Santos (2009): "Semiparametric Estimation of Nonseparable Models: A Minimum Distance from Independence Approach," *Working Paper*.

[54] Koul, H.L. (2002): *Weighted Empiricals in Dynamic Nonlinear Models*, Lecture Notes in Statistics, Springer, New York.

[55] Kress, R. (1999): *Linear Integral Equations*, Springer.

[56] Lee, A.J. (1990): *U-Statistics: Theory and Practice*, Marcel Dekker, New York.

[57] Linton, O.B., and E. Mammen (2003): "Estimating Semiparametric Arch($\infty$) Models by Kernel Smoothing Methods," STICERD working paper EM/2003/453.

[58] Linton, O.B., and E. Mammen (2005): "Estimating Semiparametric Arch($\infty$) Models by Kernel Smoothing Methods," *Econometrica*, **73**, 771-836.

[59] Magnac, T. and D. Thesmar (2002): "Identifying Dynamic Discrete Decision Processes," *Econometrica*, **70**, 801-816.

[60] Masry, E. (1996): "Multivariate Local Polynomial Regression for Time Series: Uniform Strong Consistency and Rates," *Journal of Time Series Analysis* **17**, 571–599.

[61] McFadden, D.L. (1974): "Conditional Logit Analysis of Qualitative Choice Behavior," *Frontier in Econometrics*, eds P. Zarembka. New York.

[62] Newey, W.K. (1994): "The Asymptotic Variance of Semiparametric Estimators," *Econometrica,* **62**, 1349-1382.

[63] Newey, W.K. & D. McFadden, 1994. "Large Sample Estimation and Hypothesis Testing," *Handbook of Econometrics*, vol. 4, eds. R.F. Engle and D. McFadden. North-Holland.

[64] Pakes, A.(1994a): "Estimation of Dynamic Structural Models: Problems and Prospects Part II: Mixed Continuous-Discrete Control Models and Market Interactions," *Proceedings of the 6th World Congress of the Econometric Society, Cambridge University Press.*

[65] Pakes, A., and D. Pollard (1989): "Simulation and the Asymptotics of Optimization Estimators," *Econometrica,* **57**, 1027-1057.

[66] Pakes, A., and S. Olley (1995): "A Limit Theorem for a Smooth Class of Semiparametric Estimators," *Journal of Econometrics,* **65**, 295-332.

[67] Pakes, A., M. Ostrovsky, and S. Berry (2004): "Simple Estimators for the Parameters of Discrete Dynamic Games (with Entry/Exit Example)," *Working Paper.*

[68] Parzen, E. (1962): "On Estimation of a Probability Density Function and Mode," *The Annals of Mathematical Statistics,* **33**, 1065 - 1076.

[69] Pesendorfer, M., and P. Schmidt-Dengler (2003): "Identification and Estimation of Dynamic Games," *Working Paper.*

[70] Pesendorfer, M., and P. Schmidt-Dengler (2008): "Asymptotic Least Squares Estimator for Dynamic Games," *Reviews of Economics Studies*, **75**, 901-928.

[71] Pesendorfer, M., and P. Schmidt-Dengler (2009): "Sequential Estimation of Dynamic Discrete Games: A Comment," forthcoming, *Econometrica.*

[72] Pollard, D. (1985): "New Ways to Prove Central Limit Theorems," *Econometric Theory*, **1**, 295-314.

[73] Pollard, D. (1990): *Empirical Processes: Theory and Applications*, NSF-CBMS Regional Conference Series in Probability and Statistics.

[74] Powell, J.L., J.H. Stock and T.M. Stoker (1989): "Semiparametric Estimation of Index Coefficients," *Econometrica*, **57**, 1403-1430.

[75] Robinson, P.M. (1983): "Nonparametric Estimators for Time Series," *Journal of Time Series Analysis*, **4**, 185-207.

[76] Robinson, P.M. (1988): "Root-N-consistent semiparametric regression," *Econometrica*, **56**, 931-954.

[77] Rosenblatt, M. (1956): "Remarks on some Nonparametric Estimates of a Density Function," *Annals of Mathematical Statistics*, **27**, 832 - 837.

[78] Roussas, G. (1967): "Nonparametric Estimation in Markov Processes," *Annals of the Institute of Statistical Mathematics*, **21**, 73-87.

[79] Roussas, G. (1969): "Nonparametric Estimation of the Transition Distribution Function of a Markov Process," *Annals of Mathematical Statistics*, **40**, 1386-1400.

[80] Rust, J. (1987): "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," *Econometrica*, **55**, 999 - 1033.

[81] Rust, J. (1988): "Maximum Likelihood Estimation of Discrete Control Processes," *SIAM Journal on Control and Optimization*, **26**, 1006-1023.

[82] Rust, J. (1994a): "Estimation of Dynamic Structural Models: Problems and Prospects Part I: Discrete Decision Processes," *Proceedings of the 6th World Congress of the Econometric Society, Cambridge University Press.*

[83] Rust, J. (1994b): "Structural Estimation of Markov Decision Processes," *Handbook of Econometrics*, vol. 4, eds. R.F. Engle and D.L. McFadden. North-Holland.

[84] Rust, J. (1996): "Numerical Dynamic Programming in Economics," *Handbook of Computational Economics*, vol. 4, eds. H. Amman, D. Kendrick and J. Rust. North-Holland.

[85] Srisuma, S., and O. B. Linton (2009): "Semiparametric Estimation of Markov Decision Processes with Continuous State Space," *Working Paper.*

[86] Stokey, N.L., and R.E. Lucas (1989): *Recursive Methods in Economics Dynamics,* Harvard University Press, Cambridge MA.

[87] Taussky, O. (1949): "A Recurring Theorem on Determinants," *American Mathematical Monthly*, **56**, 672–676.

[88] Topkis, D. (1998): *Supermodularity and Complementarity*, Princeton University Press

[89] van der Vaart, A. W. and J. A. Wellner (1996): *Weak Convergence and Empirical Processes,* Springer-Verlag, New York.

[90] Wolfowitz, J. (1953): "Estimation by Minimum Distance Method," *Ann. Inst. Stat. Math.*, **5**, 9–23.

[91] Zeidler, E. (1986): *Nonlinear Functional Analysis and its Applications I: Fixed-Point Theorem,* Springer-Verlag, New York.