Semiparametric Estimation of Markov Decision Processes: Continuous Control*

Sorawoot Srisuma†

London School of Economics and Political Science

January 2010

Abstract

We develop a new two-step estimator for a class of Markov decision processes with continuous control that is simple to implement and does not require a parametric specification on the distribution of the observables. Making use of the monotonicity assumption similar to Bajari, Benkard and Levin (2007), we estimate the continuation value functions nonparametrically in the first stage. In the second stage we minimize a minimum distance criterion that measures the divergence between the nonparametric conditional distribution function and a simulated semiparametric counterpart. We show under some regularity conditions that our minimum distance estimator is asymptotically normal and converges at the parametric rate. We estimate the continuation value function by kernel smoothing and derive its pointwise distribution theory. We propose to use a semiparametric bootstrap to estimate the standard error for inference since the asymptotic variance of the finite dimensional parameter will generally have a complicated form. Our estimation methodology also forms a basis for the estimation of dynamic models with both discrete and continuous controls, which can be used to estimate dynamic models of imperfect competition and other dynamic games.

Keywords: Markov Decision Models, Empirical Processes, Kernel Smoothing, Markovian Games, Semiparamatric Estimation

*I am most grateful to my advisor, Oliver Linton, for his encouragement and guidance. I also wish to thank Javier Hidalgo, Stefan Hoderlein, Tatiana Komarova, Martin Pesendorfer, Marcia Schafgans, Myung Seo and seminar participants at the LSE for valuable comments and suggestions. This research is partially supported by the ESRC.

†Department of Economics, London School of Economics, Houghton Street, London, WC2A 2AE, United Kingdom. E-mail address: `s.t.srisuma@lse.ac.uk`

1 Introduction

The estimation of dynamic programming models plays an important role in helping us understand the behavior of forward looking economic agents. In this paper, we develop a new estimator that is capable of estimating a class of Markovian 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. We show that our estimation methodology can be extended in several directions, including estimating models with discrete and continuous controls and estimating dynamic games.

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. The value functions and their corresponding continuation values generally have no closed form but are defined as solutions to some nonlinear functional equations. Much work in the dynamic estimation literature focuses on how to alleviate or avoid repeatedly solving such equations. We follow the strand of research that uses a two-step approach 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. This theme of estimation has been growing since the early work of Hotz and Miller (1993, hereafter HM) in their estimation of an optimal stopping problem. 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 known as 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". This is the main insight from the two-step estimation method proposed by HM that has led to many subsequent two-step estimation procedures in the literature. We note that, although HM proposes an estimation methodology for a single agent model with discrete controls with finite time horizon, their insight can be readily adapted to other frameworks. 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.

An alternative methodology, that also relies on Hotz and Miller's inversion theorem, is introduced in the paper by Hotz, Miller, Sanders and Smith (1994), who use Monte Carlo method to estimate the value functions instead of solving a linear equation in the first stage. Bajari, Benkard and Levin (2007, hereafter BBL) apply a closely related simulation idea 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 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 paper we focus on models with observable state variables that take finitely many values, our estimator can also accommodate continuous state variable. Recently, Bajari, Chernozhukov, Hong and Nekipelov (2009) and Srisuma and Linton (2009) have independently shown how to estimate such dynamic models with purely discrete choice and continuous state variable. The main technical difficulty that arises when the state variable is continuous 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, see Srisuma and Linton for a formal discussion. Bajari et al. (2009) provide some conditions for the per period payoff function to be nonparametrically identified and use the method of sieves in their nonparametric estimation. Srisuma and Linton, using kernel smoothing, generalize the methodology of Pesendorfer and Schmidt-Dengler (2008) and provide a set of weak conditions under time series framework to ensure root$-T$ consistent estimation of the finite dimensional parameters and also provide pointwise distribution theory for the conditional value functions and the continuation value functions.

We comment that there is comparatively less work on the development of estimation methodology with purely continuous control.[1] One well known exception to this is a 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 class of Markov decision models we are interested in do not fall into this class, for more details see Rust (1996). Our framework is more closely related to the study of

[1]The other paper that we are aware of that estimates purely continuous control problem in the I.O. literature is Berry and Pakes (2002), but it is based on quite a different set of assumptions.

dynamic auction and oligopoly models, which often allow for discrete choice as well (e.g. entry/exit decisions);[2] we refer to the surveys of Pakes (1994), and more recently, Ackerberg, Benkard, Berry and Pakes (2005). Although we focus on the estimation of continuous choice models, as illustrated by Arcidiacono and Miller (2008), our methodology forms a basis which can be used to estimate models with continuous and discrete decisions with general patterns of unobserved heterogeneity.

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 simulation 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 paper 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 3 provides the

[2] 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.

conditions to obtain the desired distribution theory. We discuss inference based on semiparametric bootstrap in Section 4. Section 5 reports a Monte Carlo study of our estimator and illustrates the affects of ignoring the model dynamics. Section 6 extends our methodology to estimate Markovian processes with discrete/continuous controls as well as a class of Markovian games and considers the estimation problem when the observable state space is uncountably infinite. Section 7 concludes. The proofs of all theorems can be found in the Appendix A. We collect the Figures and Tables at the end of the paper.

In this paper: 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 $g : A \times \Theta \to \mathbb{R}$, 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 θ; when $\mathcal{G}^J$ is a class of continuous $\mathbb{R}^J$ valued functions $(g_j(\cdot, \theta))$, then denote $\|g\|_{\mathcal{G}} = \max_{1 \le j \le J} \sup_{\theta \in \Theta} \|g_j(\cdot, \theta)\|_{\infty}$, where $\|g\|_{\infty} = \max_{1 \le j \le J} \sup_{a \in A} |g_j(a)|$, and, $\|g\|_{\mathcal{G}} = \|g\|_{\infty}$ when g does not depend on θ.

2 Markov Decision Processes

2.1 Basic Framework

We first briefly describe the Markov decision model that our methodology can estimate. The random variables in the model are the control and state variables, denoted by a and s respectively. The control variable, a, belongs to some convex set $A \subset \mathbb{R}^{L_A}$. The state variables, s, is an element in $\mathbb{R}^{L_X + 1}$. Time is indexed by t, the economic agent is forward looking in solving an infinite horizon intertemporal problem. At time period t, 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 $u(a_t, s_t)$ and agent's 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 some rate $\beta \in (0, 1)$. Formally the agent is represented by a triple of primitives (u, p, β), who 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* for any time τ

$$V(s_\tau) = \sup_{\mathcal{A}_\tau} E\left[\left.\sum_{t=\tau}^{\infty} \beta^{t-\tau} u(a_t, s_t)\right| s_\tau\right], \text{s.t. } a_t \in A(s_t) \text{ for all } t \ge \tau.$$

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

$$\alpha(s_t) = \arg\sup_{a \in A(s_t)} \{u(a, s_t) + \beta E[V(s_{t+1})|s_t, a_t = a]\} \quad \text{for all } t \ge 1. \tag{1}$$

Furthermore, the value function, V, is the unique solution to the *Bellman's equation*

$$V(s_t) = \sup_{a \in A(s_t)} \{u(a, s_t) + \beta E[V(s_{t+1}) | s_t, a_t = a]\}. \tag{2}$$

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; ε_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

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

In Section 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.

We end this subsection by providing some model assumptions that we assume throughout the paper.

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

Assumption M2: *(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 M3: *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* ε_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 M4: *(Monotone Choice) The per period payoff function* $u_\theta : A \times X \times \mathcal{E} \to \mathbb{R}$ *has increasing differences in* (a, ε) *for all* x *and* θ; u_θ *is specified upto some unknown parameters* $\theta \in \Theta \subset \mathbb{R}^L$.

The first two assumptions are familiar from the discrete control problems; M2 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 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 ε_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 ε_t can be weakened to the knowledge of the conditional distribution of ε_t given x_t upto some finite dimensional unknown parameters. However, unlike dynamic discrete choice models, the support of ε_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 M4, which we will discuss further below, it essentially ensures the policy function (1) is invertible on $\mathcal{E}$ for each state $x \in X$.

2.2 Value Functions

Before moving on to the estimation strategy, it will be useful to first discuss in details of our treatment regarding the value function. In dynamic structural estimation, it is often necessary to have a numerical representation for the continuation value function, under M2 this function can be written as $E[V_\theta(s_{t+1}) | x_t, a_t]$. As mentioned in the introduction, we aim to estimate the continuation value function rather than approximate it for each θ, cf. the methods discussed in Pakes (1994) and Rust (1994). Under some additional assumptions on the DGP, this conditional expectation is nonparametrically identified if we observe s_t and know $V_\theta(\cdot)$, the latter is defined in (2). Since we know neither, a standard approach is to consider the value function on the optimal path, which is defined as the solution to some linear equation; this is the approach taken by HM. By marginalizing out the unobservable states in the linear characterization of the aforementioned linear equation we have the conditional value function defined as a solution to a linear equation, called the *policy value equation*. The continuation value function can then be written as a linear transform of the solution to the policy value equation. More formally, M1 implies $a_t = \alpha_{\theta_0}(s_t)$, where $\alpha_{\theta_0} : X \times \mathcal{E} \to \mathbb{R}$ denotes the policy function defined in (1) that reflects the parameterization by the true structural parameter θ_0. On the optimal path, the value function is a stationary solution to the policy value equation when $\theta = \theta_0$, cf. (2)

$$V_\theta(s_t) = u_\theta(a_t, s_t) + \beta E[V_\theta(s_{t+1})|s_t]. \tag{3}$$

Note that the equation above is also well defined for any θ that is not equal to θ_0; then V_θ is interpreted as the value function for an economic agent whose underlying preference is θ but is using the policy function that is optimal with respect to θ_0. Marginalizing out the unobserved states in (3), under M2 we have the following characterization of the value functions

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

then, again by M2, that the continuation value function can be written as

$$E[V_\theta(s_{t+1})|x_t, a_t] = E[E[V_\theta(s_{t+1})|x_{t+1}]|x_t, a_t]. \tag{5}$$

In what follows, it will be convenient to write (4) succinctly as

$$m_\theta = r_\theta + \mathcal{L} m_\theta. \tag{6}$$

where for each $j, k = 1, \ldots, J$: $r_\theta(j)$ denotes $E[u_\theta(a_t, s_t)|x_t = j]$; $\mathcal{L}$ is a $J \times J$ stochastic matrix whose $(k, j)-$th entry represents $\beta \Pr[x_{t+1} = j|x_t = k]$; $m_\theta(j)$ denotes $E[V_\theta(s_t)|x_t = j]$. So we can define $\mathcal{R}_0 = \{(r_\theta(j)) = E[u_\theta(a_t, s_t)|x_t = j] \text{ for } j = 1, \ldots, J : \theta \in \Theta\} \subset \mathbb{R}^J$ to be a set of vectors of expected per period payoff for an agent whose true taste parameter is $\theta \in \Theta$, for all states in X, but behaves optimally according to θ_0. Note that $(I - \mathcal{L})$ is invertible by the dominant diagonal theorem, so the solution to (6) exists and is unique.[3] The conditional value function m_θ is therefore defined as the solution to (6), we denote such a subset of $\mathbb{R}^J$ by $\mathcal{M}_0 = \left\{m_\theta = (I - \mathcal{L})^{-1} r_\theta : r_\theta \in \mathcal{R}_0\right\}$. Finally, the continuation value function can be defined by the following linear transformation

$$g_\theta = \mathcal{H} m_\theta. \tag{7}$$

Here $\mathcal{H}$ is a conditional expectation operator that maps $\mathbb{R}^J$ to $\mathcal{G}^J$, where $\mathcal{G}^J$ denotes a Cartesian product of J normed space of functions $\mathcal{G}_j$ defined on A (to be defined more precisely later), so that for any $m \in \mathbb{R}^J$, j and $a \in A$, $\mathcal{H}m(j, a) = \sum_{k=1}^J m_k \Pr[x_{t+1} = k|x_t = j, a_t = a]$. We denote the set

[3] 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).

of continuation value functions of interest by $\mathcal{G}_0^J = \{\mathcal{H}m_\theta(j,\cdot) \text{ for } j=1,\ldots,J : m_\theta \in \mathcal{M}_0\}$, where $\mathcal{G}_0^J = \times_{j=1}^J \mathcal{G}_{0,j}$ with $\mathcal{G}_{0,j} = \{\mathcal{H}m_\theta(j,\cdot) : m_\theta \in \mathcal{M}_0\}$ and generally we have that $\mathcal{G}_{0,j} \subset \mathcal{G}_j$ for all j, so it follows that $\mathcal{G}_0^J \subset \mathcal{G}^J$. In this paper we denote a generic element of $\mathcal{G}_0^J$ that depends on θ by $g_0(\cdot,\theta) = (g_{0,j}(\cdot,\theta))$ where $g_{0,j}(\cdot,\theta) \in \mathcal{G}_{0,j}$ for any j; other vector of functions in $\mathcal{G}^J$ that depends on θ by $g(\cdot,\theta) = (g_j(\cdot,\theta))$ where $g_j(\cdot,\theta) \in \mathcal{G}_j$ for each j; all other functions of $\mathcal{G}^J$ that need not depend on θ by $g(\cdot)$. Since we will need to repeatedly work with the derivative of g w.r.t. a, it will be convenient to use $\partial_a g$ as a shorthand notation for $\frac{\partial}{\partial a} g(\cdot)$. Analogous to $\left(\mathcal{G}_j, \mathcal{G}_{0,j}, \mathcal{G}^J, \mathcal{G}_0^J\right)$, we denote $\left(\mathcal{G}_j^{(1)}, \mathcal{G}_{0,j}^{(1)}, \mathcal{G}^{(1)J}, \mathcal{G}_0^{(1)J}\right)$ to be sets of functions that $(\partial_a g(\cdot,\theta), \partial_a g_0(\cdot,\theta), (\partial_a g_j(\cdot,\theta)), (\partial_a g_{0,j}(\cdot,\theta)))$ belongs to. It is natural that our estimate of $g_0(\cdot,\theta)$ satisfies the empirical relations of (6) and (7). We are now ready to discuss the estimation methodology.

2.3 Estimation Methodology

Since we anticipate a cross-section time-series data, we first introduce an additional index i for different economic agent. To motivate the choice of our criterion function, we first consider the natural approach of generating some moments from conditional moment restrictions, see Pesendorfer and Schmidt-Dengler (2008) and Srisuma and Linton (2009) when the control is discrete. Analogous to the conditional choice probabilities, we can generate a class of estimators from the following conditional moment restrictions

$$E\left[\mathbf{1}\left[a_{it} \leq a\right] - F_{A|X}\left(a|j;\theta_0, \partial_a g_{0,j}(\cdot,\theta_0)\right) \middle| x_{it} = j\right] = 0, \text{ for } a \in A \text{ and } j = 1,\ldots,J, \tag{8}$$

where $F_{A|X}(a|j;\theta_0, \partial_a g_{0,j}(\cdot,\theta_0))$ is the conditional distribution function of a_{it} given $x_{it} = j$ that we can estimate, the dependence on $\partial_a g$ will become clear shortly. We see that the empirical counterpart of (8) is a random function over A. By allowing for profiling, one would expect that under some conditions optimal instruments should exist that will allow the corresponding estimator to achieve semiparametric efficiency bound, cf. Ai and Chen (2003). Efficiency aside, since we have a continuum of moment restrictions here, cf. Carrasco and Florens (2000), no general theory for semiparametric moment estimation with a continuum of moments is available at present. Further, we show below that $F_{A|X}(a|j;\theta,\partial_a g_j)$ can be written as an integral that we can approximate by Monte Carlo simulation, which introduces non-smoothness in the objective function. Another alternative to the moment based estimator is to maximize the conditional maximum likelihood function, however the maximum likelihood estimator (MLE) is more computationally demanding, we provide more discussion on MLE in later part of the paper.

Instead, we focus on another class of minimum distance estimators. Wolfowitz (1953) introduce the minimum distance method that since has developed into a general estimation technique that have 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

$$M_N\left(\theta,\widehat{g}\left(\cdot,\theta\right)\right)=\sum_{j=1}^{J}\int_A\left[\widetilde{F}_{A|X}\left(a|j;\theta,\partial_a\widehat{g}_j\left(\cdot,\theta\right)\right)-\widehat{F}_{A|X}\left(a|j\right)\right]^2\mu_j\left(da\right),\tag{9}$$

where for each $j=1,\ldots,J$: $\widehat{F}_{A|X}\left(\cdot|j\right),\widetilde{F}_{A|X}\left(\cdot|j;\theta,\partial_a g_j\right)$ and $\partial_a\widehat{g}_j\left(\cdot,\theta\right)$ are defined below in (10), (14) and (19) respectively; $F_{A|X}\left(\cdot|j\right)=E\left[\mathbf{1}\left[a_{it}\leq\cdot\right]|\,x_{it}=j\right]$ denotes the true conditional distribution function, which is equal to $F_{A|X}\left(\cdot|j;\theta_0,\partial_a g_{0,j}\left(\cdot,\theta_0\right)\right)$; μ_j is some user chosen sigma-finite measure on A. Clearly the property of $\widehat{\theta}$ will generally depend on the choice of $\left\{\mu_j\right\}_{j=1}^{J}$, similarly to the papers from minimum distance from independence literature mentioned in the introduction, we shall leave the issue of choosing $\left\{\mu_j\right\}_{j=1}^{J}$ for future work.

We now discuss how to feasibly compute the objective function in (9); for simplicity we assume $\dim\left(A\right)=1$. We remark that there are various numerical methods that can approximate the integral above arbitrarily well, for example see Judd (1998), hence we make no further comments regarding the numerical error that arises from the integral approximation. Under M1 - M4 and further assumptions on the DGP, to be made precise below, $F_{A|X}\left(a|j\right)$ will be nonparametrically identified for each a,j. $\widehat{F}_{A|X}\left(a|j\right)$ can be generally written as

$$\widehat{F}_{A|X}\left(a|j\right)=\frac{1}{NT}\sum_{i=1,t=1}^{N,T}w_{itN}\left(j\right)\mathbf{1}\left[a_{it}\leq a\right],\tag{10}$$

depending on the nature of X, the weighting function $w_{itN}\left(\cdot\right)$ will take different form. In this paper, when x_{it} is a discrete random variable we choose $w_{itN}\left(j\right)=\mathbf{1}\left[x_{it}=j\right]/\sum_{i=1,t=1}^{N,T}\mathbf{1}\left[x_{it}=j\right]$, which yields the frequency estimator, we consider the case when x_{it} is a continuous random variable in Section 6. We now describe the estimation of $F_{A|X}\left(a|j;\theta,\partial_a g_j\right)$. For any $\theta\in\Theta$, $j=1,\ldots,J$ and $g_j\in\mathcal{G}_j$, we define a function that is analogous to the discounted expected utility objective in (2)

$$\Xi_j\left(a,\varepsilon,\theta,g_j\right)=u_\theta\left(a,j,\varepsilon\right)+\beta g_j\left(a\right),\tag{11}$$

note that, under M2 and M3, we can rewrite (2) as

$$V_{\theta_0}\left(j,\varepsilon_{it}\right)=\sup_{a\in A}\Xi_j\left(a,\varepsilon_{it},\theta_0,g_{0,j}\right)\qquad\text{for }j=1,\ldots,J.$$

It will be convenient, at least for theoretically purposes, to assume that the optimal rule is characterized by the first order condition as displayed above. Taking first derivative of (11) w.r.t. a, we obtain the following map

$$\Xi_j^{(1)}\left(a,\varepsilon,\theta,\partial_a g_j\right)=\partial_a u_\theta\left(a,j,\varepsilon\right)+\beta\partial_a g_j\left(a\right).\tag{12}$$

Suppose that the maximizer to $\Xi_j(\cdot,\varepsilon,\theta,g_j)$, for any given (ε,θ,g_j), is characterized by the zero to (12); using the implicit function theorem in Banach space, we define

$$\alpha_j(\varepsilon,\theta,\partial_a g_j) = \arg\sup_{a\in A} \Xi_j(a,\varepsilon,\theta,g_j)\,.^4 \tag{13}$$

Therefore the policy profile (α_j), such that $\alpha_j : \mathcal{E}\times\Theta\times\mathcal{G}_j^{(1)} \to \mathbb{R}$, corresponds to the policy function (1) when $\partial_a g_j \in \mathcal{G}_{0,j}^{(1)}$ for all j. More generally, for any $\theta\in\Theta$, $j=1,\ldots,J$ and $\partial_a g_j \in \mathcal{G}_j^{(1)}$ we have

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

If we know the policy function α_j, we can approximate $F_{A|X}(a|j;\theta,\partial_a g_j)$ to an arbitrary degree of accuracy by Monte Carlo integration, since we assume the knowledge of Q_ε. In this paper, for simplicity we use

$$\widetilde{F}_{A|X}(a|j;\theta,\partial_a g_j) = \frac{1}{R}\sum_{r=1}^{R}\mathbf{1}\left[\alpha_j(\varepsilon_r,\theta,\partial_a g_j)\le a\right], \tag{14}$$

where $\{\varepsilon_r\}_{r=1}^R$ is a random sample from Q. We can also compute the policy profile $\alpha=(\alpha_j)$ to any degree of accuracy for any (j,ε). In particular the implication of M4, by Topkis' theorem, is that $\alpha_j(\varepsilon,\theta,\partial_a g_j)$ is non-decreasing in ε for all j, $\partial_a g_j$ and θ.[5,6] Since we are going to be working with a smooth utility function and ε has a convex support, we will assume that the strong form of monotone choice that will ensure that the policy function is strictly increasing in ε. For this reason, the convex support of a_{it} and ε_{it} must be either both bounded or unbounded to avoid internal inconsistency. For $j=1,\ldots,J$, we denote the inverse function of (13) by ρ_j, so that for all $\varepsilon\in\mathcal{E}$

$$\rho_j(\alpha_j(\varepsilon,\theta,\partial_a g_j),\theta,\partial_a g_j) = \varepsilon,$$

for any $(\theta,\partial_a g_j)$. But M4 permits ε to enter u_θ in a general way, which makes the estimation of r_θ more complicated since we do not observe ε_{it}. However, by Topkis' theorem, we can generate these terms nonparametrically. To see this, for any given pair (a,j), $F_{A|X}(a|j) = \Pr\left[\alpha_j(\varepsilon,\theta_0,\partial_a g_{0,j})\le a\right]$, by the invertibility of α_j, $\Pr\left[\alpha_j(\varepsilon_{it},\theta_0,\partial_a g_{0,j})\le a\right] = \Pr\left[\varepsilon_{it}\le\rho_j(a,\theta_0,\partial_a g_{0,j})\right] = Q_\varepsilon\left(\rho_j(a,\theta_0,\partial_a g_{0,j})\right)$, where Q_ε denotes the known distribution function of ε_{it}. So we can uniquely recover ε_{it}, which satisfies $\alpha_j(\varepsilon_{it},\theta_0,\partial_a g_{0,j}) = a_{it}$ when $x_{it}=j$ by the relation

$$\rho_j(a_{it},\theta_0,\partial_a g_{0,j}) = Q_\varepsilon^{-1}\left(F_{A|X}(a_{it}|x_{it})\right). \tag{15}$$

[4] 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).

[5] Topkis's theorem states that: if f is supermodular in (x,θ), and D is a lattice, then $x^*(\theta) = \arg\max_{x\in D} f(x,\theta)$ is nondecreasing in θ.

[6] The monotone choice condition is a common assumption used to analyze comparative statics in economic theory, for example see Athey (2002) and the reference therein.

In this paper, we use the following analogue to estimate ε_{it}

$$\widehat{\varepsilon}_{it} = Q_{\varepsilon}^{-1}\left(\widehat{F}_{A|X}\left(a_{it}|x_{it}\right)\right). \tag{16}$$

Since $F_{A|X}$ is nonparametrically identified, we can generate $\{\widehat{\varepsilon}_{it}\}_{i=1,t=1}^{N,T+1}$ nonparametrically.[7] We note that M4 is also essential in the forward simulation method of BBL (Section 3.2.2) where it is used in the reverse direction to simulate the optimal choice a_r given (j, ε_r), more specifically using our notation, $\alpha_j\left(\varepsilon_r, \theta_0, \partial_a g_{0,j}\right) = F_{A|X=j}^{-1}\left(Q_{\varepsilon}\left(\varepsilon_r\right)\right)$ where $F_{A|X=j}^{-1}$ denotes the inverse of the conditional distribution function $F_{A|X}\left(\cdot|j\right)$. Given $\{a_{it}, x_{it}, \widehat{\varepsilon}_{it}\}_{i=1,t=1}^{N,T+1}$, we can easily estimate r_θ for each θ

$$\widetilde{r}_\theta\left(j\right) = \frac{1}{NT}\sum_{i=1,t=1}^{N,T} w_{itN}\left(j\right) u_\theta\left(a_{it}, x_{it}, \widehat{\varepsilon}_{it}\right), \tag{17}$$

where, for simplicity, $w_{itN}\left(j\right)$ is the same as the one used in (10). Once we have an estimator for r_θ, the estimator of $g_0\left(\cdot, \theta\right)$ is easy to obtain. In particular, assuming further that $\widehat{p}_X\left(j\right) > 0$ for $j = 1, \ldots, J$, where $\widehat{p}_X\left(j\right)$ denotes the frequency estimator of $p_X\left(j\right) = \Pr\left[x_{it} = j\right]$, then the dominant diagonal theorem implies $\left(I - \widehat{\mathcal{L}}\right)^{-1}$ exists. So we can uniquely obtain

$$\widehat{m}_\theta = \left(I - \widehat{\mathcal{L}}\right)^{-1}\widetilde{r}_\theta. \tag{18}$$

For the estimator $\widehat{g}\left(\cdot, \theta\right)$, we can use various nonparametric estimators of a regression function, for simplicity we use the Nadaraya Watson estimator to approximate the operator $\mathcal{H}$, therefore

$$\widehat{g}_\theta = \widehat{\mathcal{H}}\widehat{m}_\theta, \tag{19}$$

such that, for any a, j, k

$$\widehat{g}_j\left(a, \theta\right) = \sum_{k=1}^{J}\widehat{m}_\theta\left(k\right)\frac{\widehat{p}_{X',X,A}\left(k, j, a\right)}{\widehat{p}_{X,A}\left(j, a\right)}, \tag{20}$$

$$\widehat{p}_{X',X,A}\left(k, j, a\right) = \frac{1}{NT}\sum_{i=1,t=1}^{N,T}\mathbf{1}\left[x_{it+1} = k, x_{it} = j\right]K_h\left(a_{it} - a\right), \tag{21}$$

$$\widehat{p}_{X,A}\left(j, a\right) = \frac{1}{NT}\sum_{i=1,t=1}^{N,T}\mathbf{1}\left[x_{it} = j\right]K_h\left(a_{it} - a\right), \tag{22}$$

[7] We can perform a simple test to check for the validity that ε_{it} has distribution Q by constructing the following Cramér von-Mises statistic,

$$\omega_N = NT\int_{\mathcal{E}}\left[\widehat{Q}_{\varepsilon,N}\left(e\right) - Q_{\varepsilon}\left(e\right)\right]^2 dQ_{\varepsilon}\left(e\right),$$

where $\widehat{Q}_{\varepsilon,N}\left(e\right) = \frac{1}{NT}\sum_{i=1,t=1}^{N,T}\mathbf{1}\left[\widehat{\varepsilon}_{it} \leq e\right]$. We can use the standard nonparametric bootstrap to approximate the asymptotic distribution under the null of ω_N.

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 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. By choosing a differentiable kernel K, (21) and (22) will also be differentiable in a, a simple estimator of $\partial_a g(\cdot, \theta)$ can be obtained by differentiating $\widehat{g}(\cdot, \theta)$. We note that other nonparametric estimators for $g_0(\cdot, \theta)$ may be preferred to the Nadaraya-Watson estimator. We choose the local constant estimator for simplicity, so when A has a bounded support we will need to trim out the boundary since the bias near the boundary will generally be of a higher order of magnitude than that of the interior. Although trimming at the boundary is simple to implement, as we show below, we can alternatively use other nonparametric estimators that do not require any boundary correction, e.g. the local linear regression that is design adaptive, see Fan (1992) for details.

2.4 Practical Aspects

Here we briefly discuss the practicality of computing (14); the arguments in this section is valid irrespective of whether X is finite or uncountably infinite. We split the discussion into two parts, the computation of $\widehat{g}(\cdot, \theta)$ and then $\alpha(\cdot, \theta, \widehat{g}(\cdot, \theta))$.[8]

The main computational burden in the first part lies in the estimation of the conditional value function. This involves solving a matrix equation to compute (or approximate) (18), see Pesendorfer and Schmidt-Dengler (2008) for the case that X is finite, and, Srisuma and Linton (2009) when X includes intervals. We note that the solving of the matrix equation will be a well-posed problem, even with a continuous state space (with fine enough grids), with probability approaching to 1 under some regularity conditions. Since we estimate $\mathcal{L}$ nonparametrically, independent of θ, the approximation of a potentially large matrix, $\left(I - \widehat{\mathcal{L}}\right)^{-1}$, only has to be computed once. It is even more straightforward to compute the sequence of nonparametrically generated residuals, defined in (16), this also only needs to be computed once. The estimation of r_θ on (or the grid approximating) X can then be obtained trivially as our nonparametric estimator of r_θ has a closed-form, see (17). Once we have the estimates for m_θ, it is straightforward to obtain the estimates of g_θ as defined in (20) - (22). A further computational gain is possible if the parameterization of θ in u is linear. This is a common feature given the linearity of the policy value equation, as noted by HM, Hotz et al. (1994) and BBL. In that case $\widetilde{r}_\theta(j)$ becomes $\sum_{i=1,t=1}^{N,T} w_{itN}(j)\, u(a_{it}, x_{it}, \widehat{\varepsilon}_{it})'\,\theta$ for each j, and we can write $\widetilde{r}_\theta = W\widehat{u}_\theta = W_{\widehat{u}}\theta$. Following the linearity of the inverse operator of $\left(I - \widehat{\mathcal{L}}\right)^{-1}$ we can then

[8]MATLAB programs that carry out the computations in this paper are available upon request. A programming suite to perform the estimation proposed in this paper will soon be made available from the website http://personal.lse.ac.uk/srisuma.

easily compute $\widehat{m}_\theta = \left(I - \widehat{\mathcal{L}}\right)^{-1} W_{\widehat{u}}\theta$ for any θ. When x_{it} contains a continuous state variable, more detailed discussion of the approximation, computation and solvability of the empirical version of (6), and whether to allow the size of the linear system to be independent or grow with sample size, can be found in Srisuma and Linton (2009).

Now we consider the prospect of approximating the policy function (13). One approach is through finding the zero of the empirical analogue of (12), i.e.

$$0 = \Xi_j^{(1)}\left(a, \varepsilon, \theta, \partial_a \widehat{g}_j(\cdot, \theta)\right)\Big|_{a=\alpha_j(\varepsilon,\theta,\partial_a\widehat{g}_j(\cdot,\theta))}.$$

Since we know the functional forms of u_θ and $\widehat{g}_j(\cdot, \theta)$, their derivatives can be explicitly derived, the above display can then be easily programmed in practice. Alternatively, since the dimension of A is generally small, it may be more convenient to approximate (13) directly from the approximation of (11) by grid-search. As explained in the previous paragraph, we can straightforwardly produce a vector of estimates of the continuation value function on a grid that approximates A. This approach is less demanding than it first appears since the empirical version of (11) always has an explicit form, which is easy to program and can be readily computed for each θ, this is especially true if θ enters u linearly as discussed in the previous paragraph. Another advantage of using the method of grid-search over the search for stationary points is that, at least in small sample, a particular stationary point may not necessary pick up the global maximizer.

3 Asymptotic Theory

Our minimum distance estimator falls in the class of a profiled semiparametric M-estimator with non-smooth objective function since (14) is discontinuous in θ. 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 (9) when $\{\mu_j\}_{j=1}^J$ is a sequence of non-random measures. Although the focus of our paper 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 (14), 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 $T : X \rightarrow Y$ 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 $D_T : X \rightarrow Y$ such that $T(x+f) - T(x) = D_T(x) f + o(\|f\|)$ with $\|f\| \rightarrow 0$ 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 θ is a finite dimensional Euclidean element, the first and second Fréchet derivatives coincide with the usual (partial-) derivatives. For Theorem G below, let θ_0 and g_0 denote the true finite and infinite dimensional parameters that lie in Θ and $\mathcal{G}$ respectively. Since we only need to focus on the local behavior around (θ_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 δ 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 Θ_δ, and the choice of δ for Θ_δ 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 $M_N\left(\widehat{\theta}, \widehat{g}\left(\cdot, \widehat{\theta}\right)\right) \leq \inf_{\theta\in\Theta} M_N(\theta, \widehat{g}(\cdot,\theta)) + o_p(N^{-1})$

G2 For all θ, $\widehat{g}(\cdot,\theta) \in \mathcal{G}_{\delta_N}$ *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* θ *at* θ_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* $\|S(\theta_0, g(\cdot,\theta_0)) - D_g S(\theta_0, g_0(\cdot,\theta_0))[g(\cdot,\theta_0) - g_0(\cdot,\theta_0)]\| \leq B_N \times \sup_{\theta\in\Theta} \|g(\cdot,\theta) - g_0(\cdot,\theta)\|_\infty^2$ *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

$$\begin{aligned} & \mathcal{D}_N\left(\theta, \widehat{g}(\cdot, \theta)\right) \qquad (23) \\ = \ & \sqrt{N}\frac{\left[M_N\left(\theta, \widehat{g}(\cdot, \theta)\right) - M_N\left(\theta_0, \widehat{g}(\cdot, \theta_0)\right) - \left(M\left(\theta, \widehat{g}(\cdot, \theta)\right) - M\left(\theta_0, \widehat{g}(\cdot, \theta_0)\right)\right) - \left(\theta - \theta_0\right)' C_N\right]}{\|\theta - \theta_0\|}. \end{aligned}$$

G6 For some finite positive definite matrices Ω_0 and Ω, we have the following weak convergence $\sqrt{N}C_N \Rightarrow N\left(0, \Omega_0\right)$ and $\sqrt{N}D_N = \sqrt{N}\left(C_N + D_g S\left(\theta_0, g_0\left(\cdot, \theta_0\right)\right)\left[\widehat{g} - g_0\right]\right) \Rightarrow \mathcal{N}\left(0, \Omega\right)$.

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 (23), will be the same as the case that $g_0(\cdot, \theta)$ is known.[9]

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 exogenous initial values;

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

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

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

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

(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;

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

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

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

(vii) The distribution function of ε_{it}, Q_ε, is Lipschitz continuous and twice continuously differentiable;

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

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

(xii) The simulation size R satisfies $N/R = o(1)$;

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) restricts $\dim(A)$ to 1 for the sake of simplicity. To allow for higher dimension of A, we will need to ensure that the policy functions (13) is invertible. $\dim(A)$ 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 paper 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 θ_0. We assume there does not exist any other $\theta \in \Theta \backslash \{\theta_0\}$ that can generate the same policy profile which θ_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 (8) 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 ε_{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 ε_{it} given $x_{it} = j$. Lastly, given θ, 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 θ_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_ε 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 M4. 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, ε and θ 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 (14) 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 (α_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 $\left(\mathcal{G}, \|\cdot\|_{\mathcal{G}}\right)$, 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* $N\left(\varepsilon, \mathcal{G}, \|\cdot\|_{\mathcal{G}}\right)$ and the *covering number with bracketing* $N_{[]}\left(\varepsilon, \mathcal{G}, \|\cdot\|_{\mathcal{G}}\right)$ respectively, see van der Vaart and Wellner (1996) for their definitions. We need the covering numbers of $\left(\mathcal{G}, \|\cdot\|_{\mathcal{G}}\right)$ 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 η, and

$$\|g\|_{\infty,\eta} = \max_{|\eta| \leq \underline{\eta}} \sup_{a} \left|\partial_a^{|\eta|} g(a)\right| + \max_{|\eta| = \underline{\eta}} \sup_{a \neq a'} \frac{\left|\partial_a^{|\eta|} g(a) - \partial_a^{|\eta|} g(a')\right|}{\|a - a'\|^{\eta - \underline{\eta}}}, \tag{24}$$

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 $\|\cdot\|_{\mathcal{G}}$.[10] 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, θ, 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 $\sup_{\theta, a, g_j \in \Theta \times A \times \mathcal{G}_j} \left\|D_g \rho_j(a, \theta, \partial_a g_j)\right\| < \infty$;

[10] 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.

(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 ρ_j w.r.t. $\partial_a g$ in the direction $[\partial_a \widehat{g}_j(\cdot, \theta_0) - \partial_a g_{0,j}(\cdot, \theta_0)]$ is asymptotically linear: in particular for any $a \in int(A)$

$$D_g \rho_j(a, \theta_0, \partial_a g_j(\cdot, \theta_0)) [\partial_a \widehat{g}_j(\cdot, \theta_0) - \partial_a g_{0,j}(\cdot, \theta_0)] = \frac{1}{NT} \sum_{i=1,t=1}^{N,T} \psi_{0,j}(a_{it}, x_{it}; a) + o_p\left(N^{-1/2}\right), \quad (25)$$

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 defined as in (13), 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.[11] Although these assumptions are hard to verify in practice, they are mostly mild conditions on the smoothness of ρ 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 $L^2(A, \Pi)$,[12] 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 ϱ_j, it is not difficult to show the validity of (25) 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 a valid linearization. However, since our (ρ_j) does not have a closed form it is not clear how one can obtain (ϱ_j). Once we obtain (25), 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

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

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

element in $l^{\infty}(A)$ we need $\psi_{N,j} \rightsquigarrow \psi_j$ as $N \to \infty$, where $\rightsquigarrow$ denotes weak convergence and ψ_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 1: *Under E1: For any $a \in int(A), \theta \in \Theta$ and $j = 1, \ldots, J$, if $\widehat{g}_j(\cdot, \theta)$ satisfies (19) then*

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

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_θ is known. In other words, the nonparametric estimation of m_θ, as well as the generation of the nonparametric residuals (16), 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 (9)), 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: *Under E1-E2: For $(\widehat{g}(\cdot,\theta_0))$ that satisfies* (19)*, if $\widehat{\theta}$ satisfies G1 with $M_N(\theta, g(\cdot,\theta))$ as defined in (9) then $\widehat{\theta} \xrightarrow{p} \theta_0$.*

THEOREM 3: *Under E1-E2: For $(\widehat{g}(\cdot,\theta_0))$ that satisfies* (19)*, if $\widehat{\theta}$ satisfies G1 with $M_N(\theta, g(\cdot,\theta))$ as defined in (9) then*

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

where

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

$$H_0 = 2\sum_{j=1}^{J}\int_A \left(D_\theta F_{A|X}\left(a|j;\theta_0,\partial_a g_{0,j}\left(\cdot,\theta_0\right)\right)\right)\left(D_\theta F_{A|X}\left(a|j;\theta_0,\partial_a g_{0,j}\left(\cdot,\theta_0\right)\right)\right)'\mu_j\left(da\right).$$

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

Theorem 4: *Under E1-E2: For any $a\in int\left(A\right)$ and $j=1,\ldots,J$, if $\widehat{g}_j\left(\cdot,\theta\right)$ satisfies (19) and $\widehat{\theta}$ satisfies G1 then*

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

where $B_{N,j}\left(a;m_{\theta_0}\right)$ 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 $\left(\widehat{g}_j\left(\cdot,\theta_0\right)\right)$ have the same first order asymptotic. This follows since $\widehat{g}$ (and g) is smooth in θ, and $\widehat{\theta}$ converges to θ_0 at a faster rate than $1/\sqrt{Nh}$. Note that, if we want to construct consistent confidence intervals for $g_{0,j}\left(a,\theta_0\right)$, we may use a different bandwidth in estimating $\widehat{g}$ to the one used in computing $\widehat{\theta}$.

4 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 paper.

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 paper 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 (β, u_θ, p), since we assume the complete knowledge of the discounting factor and the law of the unobserved error, the remaining primitives are θ 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, ε_{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 $\left(\widehat{\theta}^*, \widehat{g}^*(\cdot,\theta)\right)$ as done with the original data.

STEP 5:
Steps 2-4 is repeated $B-$times to obtain $B-$bootstrap estimates of $\left\{\widehat{\theta}^*_{(b)}, \widehat{g}^*_{(b)}(\cdot,\theta)\right\}_{b=1}^{B}$.

Then $\left\{\widehat{\theta}^*_{(b)}, \widehat{g}^*_{(b)}(\cdot,\theta)\right\}_{b=1}^{B}$ 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 paper.

5 Numerical Example

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 ε_t has a known distribution. $\overline{D}$ can be interpreted as the upper bound of the supply and (θ_1, θ_2) are the parameters representing the market elasticities. Unlike x_t, the evolution of the private shocks ε_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

$$\alpha(x_t, \varepsilon_t) = \left(\overline{D} + \theta_2 (x_t + \varepsilon_t) + c\theta_1 - \beta \frac{\lambda_{\theta,1} - \lambda_{\theta,2}}{\overline{a} - \underline{a}}\right) / 2\theta_1, \tag{26}$$

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

$$\alpha_s\left(x_t,\varepsilon_t\right) = \left(\overline{D}+\theta_2\left(x_t+\varepsilon_t\right)+c\theta_1\right)/2\theta_1. \tag{27}$$

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 θ_1) to be positive, i.e. the latter implies $\alpha_s\left(x,\varepsilon\right) > \alpha\left(x,\varepsilon\right)$ for any pair of $\left(x,\varepsilon\right)$. From (26) - (27), 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 $\varepsilon_t \sim Uni\left[-1,1\right]$. It can be shown that $\overline{a}-\underline{a}=1$ and

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

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 (26) 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 $\left(\theta_1,\theta_2\right)$ and assume the knowledge of $\left(\overline{D},c\right)$. In estimating the nonparametric estimator of $g_0\left(\cdot,\theta\right)$, 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 $\left\{h_\varsigma = 1.06s\left(NT\right)^{-\varsigma} : \varsigma = \frac{1}{5},\frac{1}{6},\frac{1}{7},\frac{1}{8},\frac{1}{9}\right\}$. We provide two estimators for each of $\left(\theta_1,\theta_2\right)$, 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 policy function (13) by using grid-search instead of computing the derivative of the continuation value. The measures (μ_1,μ_2) we use in constructing the minimum distance estimator in (9) simply put equal weights on all a and x.

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 paper. We report the bias, median of the bias, standard deviation and interquartile range (scaled by 1.349) for the estimators of θ_1 and θ_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.[13] However, the effect of trimming is unclear. In the case of the estimator of θ_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 θ_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_ς. 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 1 and 2 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 under static environment do not converge to

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

$(\theta_1, \theta_2) = (1, 0.5)$, 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 (27) 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.

6 Extensions

6.1 Discrete and Continuous Controls

In this subsection we outline how one can estimate dynamic models with discrete as well as 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 $E\left[u\left(a_t, d_t, x_t, \varepsilon_t, v_t^K\right)|x_t, v_t^K, d_t = k\right] + \beta E\left[V_\theta(s_{t+1})|x_t, v_t^K, d_t = k\right]$, sequentially, she then observes ε_t and chooses a that maximizes $u\left(a, d_t, x_t, \varepsilon_t, v_t^K\right) + \beta E\left[V_\theta(s_{t+1})|s_t, d_t, a_t = a\right]$. 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

$$\begin{aligned}
\delta\left(x_t, v_t^K\right) &= \arg\max_{1 \le k \le K} \left\{E\left[u\left(a_t, d_t, x_t, \varepsilon_t, v_t^K\right)|x_t, d_t = k\right] + \beta E\left[V(s_{t+1})|x_t, d_t = k\right]\right\}, \\
\alpha\left(x_t, \varepsilon_t, v_t^K, d_t\right) &= \sup_{a \in A}\left\{u\left(a_t, d_t, x_t, \varepsilon_t, v_t^K\right) + \beta E\left[V(s_{t+1})|x_t, a_t = a, d_t\right]\right\}.
\end{aligned}$$

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 β.*

Assumption DC2: *(Conditional Independence) The transitional distribution has the following factorization:* $p\left(x_{t+1}, \varepsilon_{t+1}, v_{t+1}^K | x_t, \varepsilon_t, v_t^K, a_t, d_t\right) = z\left(\varepsilon_{t+1}, v_{t+1}^K | x_{t+1}\right) p_{X'|X,A,D}\left(x_{t+1} | x_t, a_t, d_t\right)$.

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* ε_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* $u : A \times D \times X \times \mathcal{E} \times \mathcal{V}^K \to \mathbb{R}$ *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, $u_\theta^C : A \times D \times X \times \mathcal{E} \to \mathbb{R}$ *has increasing differences in* (a, ε) *for all* d, x *and* θ, *where* u_θ^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 ε_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 M4.

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. (2)

$$V_\theta\left(s_t\right) = u_\theta\left(a_t, d_t, x_t, \varepsilon_t, v_t^K\right) + \beta E\left[V_\theta\left(s_{t+1}\right) | s_t\right],$$

where, given the sequential framework, by DC1 - DC4 $d_t = \delta_\theta\left(x_t, v_t^K\right)$ and $a_t = \alpha_\theta\left(x_t, \varepsilon_t, d_t\right)$ such that

$$\begin{aligned}
\delta_\theta\left(x_t, v_t^K\right) &= \arg\max_{1 \leq k \leq K}\left\{E\left[u_\theta^C\left(a_t, d_t, x_t, \varepsilon_t\right) | x_t, d_t = k\right] + v_t(k) + \beta E\left[V\left(s_{t+1}\right) | x_t, d_t = k\right]\right\}, \\
\alpha_\theta\left(x_t, \varepsilon_t, d_t\right) &= \sup_{a \in A}\left\{u_\theta^C\left(a, d_t, x_t, \varepsilon_t\right) + \beta E\left[V\left(s_{t+1}\right) | x_t, a_t = a, d_t\right]\right\}.
\end{aligned}$$

Marginalizing out the unobserved states of the value function, under DC2, we obtain the following familiar characterization of the value functions

$$E\left[V_{\theta}\left(s_{t}\right)|x_{t}\right]=E\left[u_{\theta}\left(a_{t},d_{t},x_{t},\varepsilon_{t},v_{t}^{K}\right)|x_{t}\right]+\beta E\left[E\left[V_{\theta}\left(s_{t+1}\right)|x_{t+1}\right]|x_{t}\right]. \tag{28}$$

As seen previously, by DC2, that the continuation value function (onto the next time period) can be written as

$$E\left[V_{\theta}\left(s_{t+1}\right)|x_{t},a_{t},d_{t}\right]=E\left[E\left[V_{\theta}\left(s_{t+1}\right)|x_{t+1}\right]|x_{t},a_{t},d_{t}\right]. \tag{29}$$

To estimate θ_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 (6), 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 ε_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 θ. 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 (28), 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 (6) 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 (29) 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 (7) 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 θ_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 (9), as described in Section 2.

6.2 Markovian Games

The development of empirical methods to analyze of dynamic games has been growing in the empirical industrial organization literature, we refer to Ackerberg, Benkard, Berry and Pakes (2005) and Aguirregabiria and Mira (2008) for recent surveys, where the latter specialize on discrete action games. A class of Markovian games can be defined by considering a finite set of endogenously linked MDP, whose interactions are to be made precise below.

(Cf. Section 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, player i obtains information $s_{it} = (x_t, \varepsilon_{it}) \in X \times \mathcal{E}_i$, where x_t denotes public information and ε_{it} denotes the private information, and chooses an action $a_{it} \in A_i$ in order to maximize her discounted expected utility

$$V_{\theta,i}\left(s_{it};\boldsymbol{\alpha}_{-i}\right) = \sup_{a_i \in A_i} \left\{ \begin{array}{c} E\left[u_{\theta,i}\left(a_{it},\boldsymbol{\alpha}_{-i}\left(s_{-it}\right),s_{it}\right)|s_{it},a_{it}=a_i\right] \\ +\beta_i E\left[V_{\theta,i}\left(s_{it+1};\boldsymbol{\alpha}_{-i}\right)|s_{it},a_{it}=a_i\right] \end{array} \right\}, \tag{30}$$

where the present period utility is time separable and is denoted by $u_{\theta,i}\left(\mathbf{a}_t, s_{it}\right)$, with $\mathbf{a}_t = (a_{it}, \mathbf{a}_{-it})$ and $\mathbf{a}_{-it}$ denotes the actions of all other players except player i; $\boldsymbol{\alpha}_{-i} = (\alpha_j)_{j\neq i}$ denotes a profile of (pure) strategies of all other players apart from player i, where for each i a strategy can be represented by a map $\alpha_i : X \times \mathcal{E}_i \rightarrow A_i$. A strategy profile $\boldsymbol{\alpha} = (\alpha_i, \boldsymbol{\alpha}_{-i})$ constitutes to a stationary Markov perfect equilibrium if for each i, for all alternative Markov strategies α_i', α_i satisfies

$$\begin{aligned} E\left[V_{\theta,i}\left(s_{it};\boldsymbol{\alpha}_{-i}\right)|x_t;\alpha_i,\boldsymbol{\alpha}_{-i}\right] &\geq E\left[V_{\theta,i}\left(s_{it};\boldsymbol{\alpha}_{-i}\right)|x_t;\alpha_i',\boldsymbol{\alpha}_{-i}\right] \\ &= E\left[\left. \begin{array}{c} u_{\theta,i}\left(\alpha_i'\left(s_{it}\right),\boldsymbol{\alpha}_{-i}\left(s_{-it}\right),s_{it}\right) \\ +\beta_i E\left[V_{\theta,i}\left(s_{it+1};\boldsymbol{\alpha}_{-i}\right)|s_{it},\mathbf{a}_t=\left(\alpha_i'\left(s_{it}\right),\boldsymbol{\alpha}_{-i}\left(s_{-it}\right)\right)\right] \end{array} \right| x_t \right], \end{aligned} \tag{31}$$

where $E\left[V_{\theta,i}\left(s_{it};\boldsymbol{\alpha}_{-i}\right)|x_t;\alpha_i,\boldsymbol{\alpha}_{-i}\right]$ denotes the integration over the unobserved states assuming that strategies $(\alpha_i, \boldsymbol{\alpha}_{-i})$ are in play.[14] We now introduce a set of assumptions that are analogous to Conditions M1 - M4:

Assumption M1′: *For each market, the observed data* $\{a_{it}, x_{it}\}_{i=1,t=1}^{N,T+1}$ *are the controlled stochastic processes that satisfying (30) for all* i *with a unique equilibrium profile* α *that satisfies (31) with exogenously known* $\{\beta_i\}_{i=1}^{N}$.

Assumption M2′: *(Conditional Independence) The transitional distribution has the following factorization:* $p\left(x_{t+1}, \boldsymbol{\varepsilon}_{t+1}|x_t, \boldsymbol{\varepsilon}_t, \mathbf{a}_t\right) = q\left(\boldsymbol{\varepsilon}_{t+1}|x_{t+1}\right) p_{X'|X,\mathbf{A}}\left(x_{t+1}|x_t, \mathbf{a}_t\right)$.

Assumption M3′: *The support of* $\mathbf{s}_t = (x_t, \boldsymbol{\varepsilon}_t)$ *is* $X \times \mathcal{E}_1 \times \ldots \times \mathcal{E}_N$, *where* $X = \{1, \ldots, J\}$ *for some* $J < \infty$ *that denotes the observable state space. For all* i, $\mathcal{E}_i$ *is a (potentially strict) subset of* $\mathbb{R}$, *the distribution of* ε_{it}, *denoted by* Q_i, *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_i *on* $\mathcal{E}_i$.

[14] Doraszelski and Satterthwaite (2007) provide conditions for the existence of the equilibrium in a closely related model.

Assumption M4$'$: *(Monotone Choice) The per period payoff function* $u_{\theta,i}: A \times A^{N-1} \times X \times \mathcal{E}_i \rightarrow \mathbb{R}$ *has increasing differences in* (a_i, ε) *for all* a_{-i}, x *and* θ, *where* u_θ *is specified upto some unknown parameters* $\theta \in \Theta \subset \mathbb{R}^L$

Assumption M5$'$: *(Private Values)* $\boldsymbol{\varepsilon}_t$ *is also jointly independent across all players, i.e.* $q(\boldsymbol{\varepsilon}_t) = \prod_{i=1}^N q_i(\varepsilon_t)$.

From comparing the above assumptions with M1 - M4, with the exception of M5$'$, we observe that, for each player i, the controlled process $\{a_{it}, x_{it}\}_{i=1,t=1}^{N,T+1}$ only differs from the single agent case in that the per period payoff function and transition law are affected by other players' actions, and each player forms an expectation, in (30), using the beliefs she has over the distribution of other players' actions. In addition to the conditional independence and monotone choice assumptions, the private value assumption is a standard condition for the estimation of a dynamic game, see the surveys mentioned earlier.

The practical aspect of the estimation is very similar to the single agent case. First consider the estimation of the continuation value functions. For each player i, by marginalizing out the private information of all the other players of (30) in equilibrium (cf. (3)), using M2$'$, we have the generalized policy value equation

$$E[V_{\theta,i}(s_{it})|x_t] = E[u_{\theta,i}(a_{it}, \mathbf{a}_{-it}, s_{it})|x_t] + \beta_i E[E[V_{\theta,i}(s_{t+1})|x_{t+1}]|x_t]. \tag{32}$$

As seen previously, for each i, (32) can be expressed analogously to the matrix equation (6) in Pesendorfer and Schmidt-Dengler (2008).as

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

where the meaning of $(m_{\theta,i}, r_{\theta,i}, \mathcal{L}_i)$ is now obvious. As seen previously, $\{\mathcal{L}_i\}_{i=1}^N$ will be a sequence of contraction maps such that $\mathcal{L}_i = \beta_i \mathbf{L}$ for each i, where $\mathbf{L}$ is a $J \times J$ stochastic matrix that represents whose (k,j) −th entry represents $\Pr[x_{t+1} = j | x_t = k]$. M4$'$ implies that the optimal strategy (policy function) of each player will be strictly increasing in ε_{it} therefore Topkis' theorem can be applied to allow feasible estimation of $r_{\theta,i}$. The sequence of N linear equations can be estimated and solved independently. Hence we only need to approximate the operator $(I - \mathcal{L}_i)^{-1}$ once for each player. The characterization of the action specific value function is then completed by

$$g_{\theta,i} = \mathcal{H}_i m_{\theta,i},$$

where $\mathcal{H}_i : \mathbb{R}^J \rightarrow \mathcal{G}_i$ is a linear operator such that $\mathcal{H}_i m(j, a_i) = \sum_{k=1}^J m(k) \Pr[x_{t+1} = k | x_t = j, a_{it} = a_i]$ for any $m \in \mathbb{R}^J$, $a_i \in A_i$ and $j = 1, \ldots, J$. The estimator for $g_{\theta,i}$ can then be used to estimate

$F_{A_i|X}(a_i|x;\theta,\partial_a g_i)$, the only change here is how we compute the optimal strategy for each player. Analogously to (11), we define for each state j

$$\Xi_{i,j}(a_i,\varepsilon_i,\theta,g_{i,j}) = E_{-i}[u_{\theta,i}(a_i,\mathbf{a}_{-it},j,\varepsilon_i)] + \beta_i g_{i,j}(a_i),$$

since player i forms a belief (i.e. a distribution over $\mathbf{a}_{-it}$ given x_{it}), E_{-i} denotes the conditional expectation consistent with player i's belief. For each j, we approximate $E_{-i}[u_{\theta,i}(a_i,\mathbf{a}_{-it},j,\varepsilon_i)]$ by $\int u_{\theta,i}(a_i,\mathbf{a}_{-i},j,\varepsilon_i)\,d\widehat{F}_{A_{-i}|X}(\mathbf{a}_{-i}|j)$ where $\widehat{F}_{A_{-i}|X}(\mathbf{a}_{-i}|j)$ denotes any consistent nonparametric estimate of the conditional distribution function of all other players. For example, if we use the empirical analogue of $\widehat{F}_{A_{-i}|X}(\mathbf{a}_{-i}|j)$ then $\int u_{\theta,i}(a_i,\mathbf{a}_{-i},j,\varepsilon_i)\,d\widehat{F}_{A_{-i}|X}(\mathbf{a}_{-i}|j)$ is simply $\frac{1}{T}\sum_{t=1}^{T} u_{\theta,i}(a_i,\mathbf{a}_{-it},x_{it},\varepsilon_i)$ $\times\mathbf{1}[x_{it}=j]/\widehat{p}(j)$. Although we do not observe ε_{it}, it can be generated nonparametrically by (16) by replacing $\widehat{F}_{A|X}$ with $\widehat{F}_{A_i|X}$. Then for each j, $F_{A_i|X}(a_i|j;\theta,\partial_a g_{i,j})$ can be estimated using Monte Carlo method by

$$\widetilde{F}_{A_i|X}(a|j;\theta,\partial_a g_{i,j}) = \frac{1}{R}\sum_{r=1}^{R}\mathbf{1}[\alpha_{i,j}(\varepsilon_r,\theta,\partial_a g_{i,j}) \le a],$$

where $\alpha_{i,j}:\mathcal{E}_i\times\Theta\times\mathcal{G}_{i,j}^{(1)}\to\mathbb{R}$ is defined by

$$\alpha_{i,j}(\varepsilon_i,\theta,\partial_a g_{i,j}) = \arg\max_{a_i\in A_i}\widehat{\Xi}_{i,j}(a_i,\varepsilon_i,\theta,\partial_a g_{i,j}),$$

with $\widehat{\Xi}_{i,j}$ differing from $\Xi_{i,j}$ by replacing $E_{-i}[u_{\theta,i}(a_i,\mathbf{a}_{-i},x,\varepsilon)]$ by $\int u_{\theta,i}(a_i,\mathbf{a}_{-i},j,\varepsilon_i)\,d\widehat{F}_{A_{-i}|X}(\mathbf{a}_{-i}|j)$. A minimum distance estimator can then be constructed, with any nonparametric estimates $\widehat{F}_{A_i|X}$ and $\{\widehat{g}_{i,j}(\cdot,\theta)\}_{i=1,j=1}^{N,J}$

$$\widehat{\theta} = \arg\min_{\theta\in\Theta}\sum_{i=1,j=1}^{N,J}\int\left[\widetilde{F}_{A_i|X}(a_i|j;\theta,\partial_a\widehat{g}_{i,j}(\cdot,\theta)) - \widehat{F}_{A_i|X}(a|j)\right]^2 d\mu_{i,j}(a_i),$$

where $\mu_{i,j}$ is some $\sigma-$finite measure for each i and j. We note that the framework discussed in this section is to have fixed N, to appeal to the asymptotic results of Section 3 we need to assume T is also fixed but we have i.i.d. data across a number of markets, say W markets, so the asymptotic results will correspond to the limiting case when $W\to\infty$.

6.3 Continuous State Space X

It is easy to see that when X contains an interval the principles of the methodology described in Section 2 is still valid. The key difference lies in the estimation of $g_0(\cdot,\theta)$. As shown in Srisuma and Linton (2009), instead of solving a matrix equation in (6), one instead needs to solve an integral equation of type II that intuitively behaves just like a matrix equation. Their method is directly applicable when a_{it} is a continuous random variable since equation (6) is defined regardless of the nature of a_{it}. If X is some compact interval, we can choose $w_{itN}(x)$ in (10) can be $K_h(x_{it} - j) / \sum_{i=1,t=1}^{N,T} K_h(x_{it} - j)$, this will yield the local constant kernel estimator. m_θ and r_θ then become elements on some Banach space $(C(X), \|\cdot\|_X)$, and $\mathcal{L}$ now represents a linear operator that generalizes the discounted stochastic matrix mentioned previously. Srisuma and Linton provide weak conditions that ensure the approximation of the infinite dimensional, empirical analogue of (6) is a well-posed inverse problem and its solution has good convergence properties. In practice, the approximation of the integral equation is done on a finite grid, which can be represented by a matrix equation that is invertible with w.p.a. 1.

However, one must be aware of some theoretical differences. Clearly, in addition to the dimension of A, the number of continuously distributed state variables in x_{it} will contribute to the curse of dimensionality in the estimation of $g(\cdot, \theta)$. This is in contrast to the estimation of discrete choice models, where only the dimensionality of X can cause slower rate of convergence as A is finite. In terms of the asymptotic distribution of the estimator of the continuation value, we note that the asymptotic variance in Theorems 1 and 4 will remain the same, only that $p_{X,A}$ now denotes a joint density of continuously distributed random variables instead of a mixed continuous-discrete density. The reason behind this is the fact that the stochastic term in the estimation of $g(\cdot, \theta)$ is a higher dimension object than the estimator of m_θ, we also see this in Theorem 1 and 4 when X is finite, therefore the variance of the nonparametric estimator of conditional expectation operator $\mathcal{H}$ dominates. In contrast, the bias from the nonparametric estimation of $\left(\{\varepsilon_{it}\}_{i=1,t=1}^{N,T+1}, r_\theta, \mathcal{L}\right)$ will now have the same order of magnitude as the bias from estimating $\mathcal{H}$. Therefore there will be a change in the bias term in Theorem 1 and 4, it can be shown that these terms can be written explicitly as a linear transform of $(I - \mathcal{L})^{-1}$ and $\mathcal{H}$, the steps in the calculations of analogous results in estimating a discrete choice model, which is directly applicable, can be found Srisuma and Linton (2009). As for Theorem 2 and 3, we need to adjust E1 and E2 to ensure that our nonparametric estimators converge at an appropriate rate and full weak convergence of the appropriate terms to hold to maintain the consistency and asymptotic normality of our finite dimensional parameters. Essentially, we will need to impose more smoothness on various functions, see the comments to the assumptions E1 and E2 in Section 3.

7 Conclusion

In this paper we develop a new two-step estimator for a class of Markov decision processes with continuous control that. 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. Unlike the methodology of BBL, which is also capable of estimating the same class of models without having to solve for the equilibrium, we do not need to impose any distributional assumptions on the transition law of the observables. This additional flexibility is very 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 and compare our estimates to the ones which ignore the model dynamics. We also highlight how our estimation methodology with purely continuous control problem can be used to estimate more complex dynamic models, in particular we consider models which contains discrete as well as continuous control variables, dynamic games and accommodate for continuously distributed observable state variables.

By construction, the two-step estimators along the line of HM that we and others have developed consist of the estimation of the continuation value function which is then use in the second stage optimization. We note that the two steps are independent of one another. This is not uncommon in semiparametric estimation, see MINPIN estimators of Andrews (1994a). Hence there is a variety of criterion functions one can choose to define the finite dimensional structural parameters. These choices will lead to varying degree of the ease of use, robustness and efficiency consideration. In this paper we propose a minium distance criterion that is easy to compute and leads to estimators with good robustness property but not necessarily efficient. We are currently working on a semiparametric maximum likelihood version of the estimator, which frees us from the need to select $\{\mu_j\}_{j=1}^{J}$ arbitrarily and should be generally more efficient, however, this estimator is computationally more demanding than the proposed minimum distance estimator.

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 point identified through some a conditional moment restrictions. There are ongoing research on the nonparametric and semiparametric identification of Markov 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.

A Appendix

A.1 Proofs of Theorems G, 1-4

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 θ_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

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

The first equality follows from the definition of $\mathcal{D}_N$ in (23). For the inequality, we expand $M\left(\widehat{\theta},\widehat{g}\left(\cdot,\widehat{\theta}\right)\right)$ around θ_0, since $H\left(\theta,g\right)$ is continuous around $\left(\theta_0,g_0\right)$ and H_0 is positive definite by G3, there exists some $C_0>0$ such that, w.p.a. 1, $\left(\theta-\theta_0\right)'H\left(\theta_0,\widehat{g}\left(\cdot,\theta_0\right)\right)\left(\theta-\theta_0\right)+o_p\left(\left\|\theta-\theta_0\right\|^2\right)\geq C_0\left\|\theta-\theta_0\right\|^2$. 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

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

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, $J_N\left(\theta\right)=D_N\left(\theta-\theta_0\right)+\frac{1}{2}\left(\theta-\theta_0\right)'H_0\left(\theta-\theta_0\right)$, note that $J_N\left(\theta\right)$ is defined for each $\widehat{g}\left(\cdot,\theta\right)$ that satisfies the conditions of Theorem G2, implicit in D_N. $J_N\left(\theta\right)$ is a quadratic approximation of $M_N\left(\theta,\widehat{g}\left(\cdot,\theta\right)\right)-M_N\left(\theta_0,\widehat{g}\left(\cdot,\theta_0\right)\right)$, 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 $J_N\left(\theta\right)$ 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 Θ_{δ_N},

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

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

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

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

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

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: $M(\theta,g(\cdot,\theta)) = \sum_{j=1}^J \int_A E_j^2(\theta,\partial_a g_j(\cdot,\theta))\,d\mu_j$ where $E_j(\theta,\partial_a g_j(\cdot,\theta)) = F_{A|X=j}(\theta,\partial_a g_j(\cdot,\theta)) - F_{A|X=j}$, and, $F_{A|X=j}(\theta,\partial_a g_j(\cdot,\theta))$ and $F_{A|X=j}$ are functions defined on A that are the shorthand notations for $F_{A|X}(\cdot|j;\theta,\partial_a g_j(\cdot,\theta))$ and $F_{A|X}(\cdot|j)$ respectively; $M_N(\theta,g(\cdot,\theta)) = \sum_{j=1}^J \int_A E_{N,j}^2(\theta,\partial_a g_j(\cdot,\theta))\,d\mu_j$ where $E_{N,j}(\theta,\partial_a g_j(\cdot,\theta)) = \widetilde{F}_{A|X=j}(\theta,\partial_a g_j(\cdot,\theta)) - \widehat{F}_{A|X=j}$, and, $\widetilde{F}_{A|X=j}(\theta,\partial_a g_j(\cdot,\theta))$ and $\widehat{F}_{A|X=j}$ are functions defined on A that are the shorthand notations for $\widetilde{F}_{A|X}(\cdot|j;\theta,\partial_a g_j(\cdot,\theta))$ and $\widehat{F}_{A|X}(\cdot|j)$

respectively; $F_{0,j}$ is a function defined on A that is the shorthand notation for $F_0(\cdot|j)$. In addition, for $j = 1, \ldots, J$, we define the class of functions $\mathcal{F}_j = \left\{\mathbf{1}\left[\cdot \leq \rho_j(a, \theta, \partial_a g_j)\right] : a \in A, \theta \in \Theta, g_j \in \mathcal{G}\right\}$, and let $\nu_{R,j}$ denote the empirical process indexed by $(\theta, \partial_a g_j) \in \Theta \times \mathcal{G}_j^{(1)}$ to be a random element that takes value over A, i.e. $\nu_{R,j}(\theta, \partial_a g_j) = \frac{1}{\sqrt{R}} \sum_{r=1}^{R} \mathbf{1}\left[\varepsilon_r \leq \rho_j(\cdot, \theta, \partial_a g_j)\right] - Q_\varepsilon\left(\rho_j(\cdot, \theta, \partial_a g_j)\right)$. We will continue to use the multi-index notation to define higher order derivatives $\partial_a^{|\eta|}$ and $\partial_\theta^{|\eta|}$, of a and θ respectively for some natural number η, as seen in (24). We next present the some lemmas that will be useful in proofing Theorems 1-3.

Lemma 1. *Under E1* $\left\|\widehat{\mathcal{L}} - \mathcal{L}\right\| = O_p\left(N^{-1/2}\right)$.

Lemma 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 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 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 5. *Under E1: For all* $j = 1, \ldots, J, \max_{0 \leq l \leq 2, 1 \leq j \leq J} \sup_{\theta, a \in \Theta \times A_N} \left|\partial_a^{|l|} \widehat{g}_j(a, \theta) - \partial_a^{|l|} g_{0,j}(a, \theta)\right| = o_p(1)$.

Lemma 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 7. *Under E1 and E2: for all* $j = 1, \ldots, J$, $\mathcal{F}_j$ *is a Donsker class.*

Lemma 8 *Under E1 and E2: For any* $j = 1, \ldots, J$ *and some positive sequence* $\delta_N = o(1)$ *as* $N \to \infty$

$$\lim_{N \to \infty} \sup_{\substack{(a, \theta, \partial_a g_j) \in A \times \Theta \times \mathcal{G}_j^{(1)}, \\ \left\|\left(a' - a, \theta' - \theta, \partial_a g_j' - \partial_a g_j\right)\right\| < \delta_N}} \left| \begin{array}{c} \frac{1}{N} \sum_{i=1}^{N} \left(\mathbf{1}\left[\varepsilon_i \leq \rho_j\left(a', \theta', \partial_a g_j'\right)\right] - Q_\varepsilon\left(\rho_j\left(a', \theta', \partial_a g_j'\right)\right)\right) \\ -\frac{1}{N} \sum_{i=1}^{N} \left(\mathbf{1}\left[\varepsilon_i \leq \rho_j(a, \theta, \partial_a g_j)\right] - Q_\varepsilon\left(\rho_j(a, \theta, \partial_a g_j)\right)\right) \end{array} \right| = 0.$$

Lemma 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 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 1. This follows from Lemma 4. For the asymptotic distribution, we only have to calculate the variance of (40), the rest follows by standard CLT. Asymptotic independence will follow if we can show $\sqrt{Nh}\text{cov}(\widehat{g}_j(a,\theta), \widehat{g}_k(a',\theta)) = o(1)$ for any $k \neq j$ and $a' \neq a$, this is trivial to show.■

PROOF OF THEOREM 2. We first show that $M(\theta, g_0(\cdot,\theta))$ has a well separated minimum at θ_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 θ in the compact set Θ 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 θ given assumptions (vii) and (xiii), this ensures a well-separated minimum. By standard arguments, consistency will now follow if we can show

$$\sup_{\theta\in\Theta}\left|M_N(\theta,\widehat{g}(\cdot,\theta)) - M(\theta,g_0(\cdot,\theta))\right| = o_p(1). \tag{33}$$

By the triangle inequality, we have

$$\begin{aligned}
\left|M_N(\theta,\widehat{g}(\cdot,\theta)) - M(\theta,g_0(\cdot,\theta))\right| &\leq 4\sum_{j=1}^{J}\int\left|\widetilde{F}_{A|X=j}(\theta,\partial_a\widehat{g}_j(\cdot,\theta)) - F_{A|X=j}(\theta,\partial_a\widehat{g}_j(\cdot,\theta))\right|d\mu_j \\
&\quad +4\sum_{j=1}^{J}\int\left|F_{A|X=j}(\theta,\partial_a\widehat{g}_j(\cdot,\theta)) - F_{A|X=j}(\theta,\partial_a g_{0,j}(\cdot,\theta))\right|d\mu_j \\
&\quad +4\sum_{j=1}^{J}\int\left|\widehat{F}_{A|X=j} - F_{A|X=j}\right|d\mu_j \\
&= A_1 + A_2 + A_3.
\end{aligned}$$

For A_1, for each j and any $\eta > 0$ we have

$$\begin{aligned}
& \Pr\left[\sup_{\theta\in\Theta}\left|\widetilde{F}_{A|X}(a|x;\theta,\partial_a\widehat{g}(\cdot,\theta)) - F_{A|X}(a|x;\theta,\partial_a\widehat{g}(\cdot,\theta))\right| > \eta\right] \\
\leq\ & \Pr\left[\sup_{\theta,a\in\Theta\times A_N}\left|\frac{1}{R}\sum_{r=1}^{R}\mathbf{1}\left[\varepsilon_r \leq \rho_j(a,\theta,\partial_a\widehat{g}_j)\right] - Q_{\varepsilon}\left(\rho_j(a,\theta,\partial_a\widehat{g}_j)\right)\right| > \eta\right] \\
\leq\ & \Pr\left[\sup_{\theta,a,\partial_a g_j\in\Theta\times A_N\times\mathcal{G}_j^{(1)}}\left|\frac{1}{R}\sum_{r=1}^{R}\mathbf{1}\left[\varepsilon_r \leq \rho_j(a,\theta,\partial_a g_j)\right] - Q_{\varepsilon}\left(\rho_j(a,\theta,\partial_a g_j)\right)\right| > \eta\right] \\
& +\Pr\left[\partial_a\widehat{g}_j(\cdot,\theta) \notin \mathcal{G}_j^{(1)}\right].
\end{aligned}$$

From Lemma 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 $R\rightarrow\infty$ by assumption (xii). By Lemma 6, $\Pr\left[\partial_a\widehat{g}_j(\cdot,\theta) \notin \mathcal{G}_j^{(1)}\right] = o(1)$ hence by finiteness of μ_j it follows that $|A_1| = o_p(1)$ uniformly over Θ. For A_2, for each j we have

$$
\begin{aligned}
\left|F_{A|X=j}\left(\theta,\partial_a\widehat{g}_j(\cdot,\theta)\right)-F_{A|X=j}\left(\theta,\partial_a g_j(\cdot,\theta)\right)\right| &= \left|Q_\varepsilon\left(\rho_j\left(a,\theta,\partial_a\widehat{g}_j(\cdot,\theta)\right)\right)-Q_\varepsilon\left(\rho_j\left(a,\theta,\partial_a g_{0,j}(\cdot,\theta)\right)\right)\right| \\
&\leq C_0\left|\rho_j\left(a,\theta,\partial_a\widehat{g}_j(\cdot,\theta)\right)-\left(\rho_j\left(a,\theta,\partial_a g_{0,j}(\cdot,\theta)\right)\right)\right|,
\end{aligned}
$$

where the inequality follows from the mean value theorem (MVT) and the fact that the derivative of Q_ε is uniformly bounded. Given the smoothness assumption on $\left(\rho_j\right)$ in assumption (xiii), by MVT in Banach space $\sup_{a\in A_N}\left|\rho_j\left(a,\theta,\partial_a\widehat{g}_j(\cdot,\theta)\right)-\left(\rho_j\left(a,\theta,\partial_a g_{0,j}(\cdot,\theta)\right)\right)\right|\leq\sup_{\theta,a,\partial_a g_j\in\Theta\times A\times\mathcal{G}_j^{(1)}}\left\|D_{\partial_a g}\rho_j\left(a,\theta,\partial_a g_j\right)\right\|$ $\times\sup_{\theta,a\in\Theta\times A_N}\left|\partial_a\widehat{g}_j(a,\theta)-\partial_a g_{0,j}(a,\theta)\right|$. Since μ_j has zero measure on the boundary of A, by Lemma 5, $\int\left|F_{A|X=j}\left(\theta,\partial_a\widehat{g}_j(\cdot,\theta)\right)-F_{A|X=j}\left(\theta,\partial_a g_{0,j}(\cdot,\theta)\right)\right|d\mu_j\leq C_0\sup_{\theta,a\in\Theta\times A_N}\left|\partial_a\widehat{g}_j(a,\theta)-\partial_a g_{0,j}(a,\theta)\right|+2\mu_j\left(A\backslash A_N\right)=o_p(1)$. So we also have $\left|A_2\right|=o_p(1)$ uniformly over Θ. 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 $\widehat{F}_{A,X}(a,j)=\frac{1}{NT}\sum_{i=1,t=1}^{N,T}\mathbf{1}\left[a_{it}\leq a,x_{it}=j\right]$, 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 9, the class of functions $\left\{\mathbf{1}\left[\cdot\leq a,x_{it}=j\right]-F_{A,X}(\cdot,j):a\in A\right\}$ 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 θ, the uniform convergence in (33) holds and consistency follows.■

Proof of Theorem 3. To proof Theorem 3 we set out to show that our assumptions imply we satisfy all the conditions of Theorem G. We showed consistency in Theorem 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 $\sup_{\theta\in\Theta}\left\|\partial_a\widehat{g}_j(\cdot,\theta)-\partial_a g_{0,j}(\cdot,\theta)\right\|_\infty$ $=o_p\left(N^{-1/4}\right)$ for all $j=1,\ldots,J$. The former is implied by Lemma 6, from the proof of Lemma 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

$$
\begin{aligned}
&M_N\left(\theta,\widehat{g}(\cdot,\theta)\right)-M_N\left(\theta_0,\widehat{g}(\cdot,\theta_0)\right)-\left(M\left(\theta,\widehat{g}(\cdot,\theta)\right)-M\left(\theta_0,\widehat{g}(\cdot,\theta_0)\right)\right)-\left(\theta-\theta_0\right)'C_N \quad (34)\\
=\ &o_p\left(\left\|\theta-\theta_0\right\|^2+\frac{\left\|\theta-\theta_0\right\|}{\sqrt{N}}+\frac{1}{N}\right),
\end{aligned}
$$

holds uniformly for $\|\theta-\theta_0\|<\delta_N$. Then for any pair $(\theta, \partial_a \widehat{g}_j(\cdot,\theta))$ we can write

$$\begin{aligned}E_j^2(\theta, \partial_a \widehat{g}_j(\cdot,\theta)) - E_j^2(\theta_0, \partial_a \widehat{g}_j(\cdot,\theta_0)) &= \left(F_{A|X=j}(\theta, \partial_a \widehat{g}_j(\cdot,\theta)) - F_{A|X=j}(\theta_0, \partial_a \widehat{g}_j(\cdot,\theta_0))\right) \times \\ &\quad \left(F_{A|X=j}(\theta, \partial_a \widehat{g}_j(\cdot,\theta)) + F_{A|X=j}(\theta_0, \partial_a \widehat{g}_j(\cdot,\theta_0)) - 2F_{A|X=j}\right),\end{aligned}$$

and analogously

$$\begin{aligned}E_{N,j}^2(\theta, \partial_a \widehat{g}_j(\cdot,\theta)) - E_{N,j}^2(\theta_0, \partial_a \widehat{g}_j(\cdot,\theta_0)) &= \left(\widetilde{F}_{A|X=j}(\theta, \partial_a \widehat{g}_j(\cdot,\theta)) - \widetilde{F}_{A|X=j}(\theta_0, \partial_a \widehat{g}_j(\cdot,\theta_0))\right) \times \\ &\quad \left(\widetilde{F}_{A|X=j}(\theta, \partial_a \widehat{g}_j(\cdot,\theta)) + \widetilde{F}_{A|X=j}(\theta_0, \partial_a \widehat{g}_j(\cdot,\theta_0)) - 2\widehat{F}_{A|X=j}\right).\end{aligned}$$

Combing these we have

$$\begin{aligned}
& M_N(\theta, \widehat{g}(\cdot,\theta)) - M_N(\theta_0, \widehat{g}(\cdot,\theta_0)) \\
= & \sum_{j=1}^{J} \int \left[\begin{array}{c} \left[\begin{array}{c} R^{-1/2}\left(\nu_{R,j}(\theta, \partial_a \widehat{g}_j(\cdot,\theta)) - \nu_{R,j}(\theta_0, \partial_a \widehat{g}_j(\cdot,\theta_0))\right) \\ + \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] \\ \times \left[\begin{array}{c} R^{-1/2}\left(\nu_{R,j}(\theta, \partial_a \widehat{g}_j(\cdot,\theta)) + \nu_{R,j}(\theta_0, \partial_a \widehat{g}_j(\cdot,\theta_0))\right) \\ + \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)) - 2\widehat{F}_{A|X=j}\right) \end{array} \right] \end{array} \right] d\mu_j \\
= & \sum_{j=1}^{J} \int E_j^2(\theta, \partial_a \widehat{g}_j(\cdot,\theta)) - E_j^2(\theta_0, \partial_a \widehat{g}_j(\cdot,\theta_0))\, d\mu_j \\
& -2\sum_{j=1}^{J} \int \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)\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}(\theta, \partial_a \widehat{g}_j(\cdot,\theta)) - \nu_{R,j}(\theta_0, \partial_a \widehat{g}_j(\cdot,\theta_0))\right] \\ \times \left[\begin{array}{c} 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)) - 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}(\theta, \partial_a \widehat{g}_j(\cdot,\theta)) + \nu_{R,j}(\theta_0, \partial_a \widehat{g}_j(\cdot,\theta_0))\right] \\ \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_j \\
& +R^{-1}\sum_{j=1}^{J} \int \left[\begin{array}{c} \left[\nu_{R,j}(\theta, \partial_a \widehat{g}_j(\cdot,\theta)) - \nu_{R,j}(\theta_0, \partial_a \widehat{g}_j(\cdot,\theta_0))\right] \\ \times \left[\nu_{R,j}(\theta, \partial_a \widehat{g}_j(\cdot,\theta)) + \nu_{R,j}(\theta_0, \partial_a \widehat{g}_j(\cdot,\theta_0))\right] \end{array} \right] d\mu_j \\
= & M(\theta, \widehat{g}(\cdot,\theta)) - M(\theta_0, \widehat{g}(\cdot,\theta_0)) + 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 (34), 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

$$
\begin{aligned}
B_1 &= -2(\theta-\theta_0)' \sum_{j=1}^{J} \int D_\theta F_{A|X=j}\left(\overline{\theta}_j, \partial_a \widehat{g}_j\left(\cdot, \overline{\theta}_j\right)\right)\left(\widehat{F}_{A|X=j} - F_{A|X=j}\right) d\mu_j \\
&= -2(\theta-\theta_0)' \sum_{j=1}^{J} \int D_\theta F_{A|X=j}\left(\theta_0, \partial_a g_{0,j}(\cdot,\theta_0)\right)\left(\widehat{F}_{A|X=j} - F_{A|X=j}\right) d\mu_j \\
&\quad -2(\theta-\theta_0)' \sum_{j=1}^{J} \int \left[\begin{array}{c} \left[D_\theta F_{A|X=j}\left(\overline{\theta}_j, \partial_a \widehat{g}_j\left(\cdot, \overline{\theta}_j\right)\right) - D_\theta F_{A|X=j}\left(\theta_0, \partial_a g_{0,j}(\cdot,\theta_0)\right)\right] \\ \times \left[\widehat{F}_{A|X=j} - F_{A|X=j}\right] \end{array} \right] d\mu_j \\
&= B_{11} + B_{12},
\end{aligned}
$$

where for each $j, \overline{\theta}_j$ is some intermediate value between θ and θ_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 (34) and that $\sqrt{N} C_N$ converges to a normal random variable. By Lemma 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 9, it is easy to show that $D_\theta F_{A|X=j}\left(\theta_0, \partial_a g_{0,j}(\cdot,\theta_0)\right)\sqrt{N}(\widehat{F}_{A|X=j} - F_{A|X=j})$ also converges weakly to a tight Gaussian process. To see the latter, note that for any $\partial_a g_j(\cdot,\theta) \in \mathcal{G}_j^{(1)}$

$$
\begin{aligned}
& D_\theta F_{A|X}\left(a|j;\theta, \partial_a g_j(\cdot,\theta)\right) \\
= \; & q\left(\rho_j\left(a,\theta,\partial_a g_j(\cdot,\theta)\right)\right)\left(\partial_\theta \rho_j\left(a,\theta,\partial_a g_j(\cdot,\theta)\right) + D_{\partial_a g}\rho_j\left(a,\theta,\partial_a g_j(\cdot,\theta)\right)\left[\partial_\theta \partial_a g(\cdot,\theta)\right]\right),
\end{aligned}
$$

where, ∂_θ denotes the ordinary $L-$dimensional partial derivative, $\partial/\partial\theta$, w.r.t. in the argument θ. 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 $\mathbb{T}_j f = \int D_\theta F_{A|X=j}\left(\theta_0, \partial_a g_{0,j}(\cdot,\theta_0)\right) f d\mu$ for any $f \in l^\infty(A)$ then the map is linear and continuous, the boundedness follows from the observation that $\sup_{a\in A}\left\|D_\theta F_{A|X=j}\left(\theta_0, \partial_a g_{0,j}(\cdot,\theta_0)\right)\right\| < \infty$. 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

$$
\sqrt{N} C_N = \sum_{j=1}^{J} \int D_\theta F_{A|X=j}\left(\theta_0, \partial_a g_{0,j}(\cdot,\theta_0)\right)\sqrt{N}\left(\widehat{F}_{A|X=j} - F_{A|X=j}\right) d\mu_j, \tag{35}
$$

then $\sqrt{N} C_N$ also converges a Gaussian variable.

For B_{12}, for each j, by Cauchy Schwarz inequality we have

$$
\begin{aligned}
& \left|(\theta-\theta_0)'\int\left(D_\theta F_{A|X=j}\left(\overline{\theta}_j,\partial_a\widehat{g}_j\left(\cdot,\overline{\theta}_j\right)\right)-D_\theta F_{A|X=j}\left(\theta_0,\partial_a g_{0,j}\left(\cdot,\theta_0\right)\right)\right)\left(\widehat{F}_{A|X=j}-F_{A|X=j}\right)d\mu_j\right| \\
\leq & \left[(\theta-\theta_0)'\int\left[\begin{array}{c}\left(D_\theta F_{A|X=j}\left(\overline{\theta}_j,\partial_a\widehat{g}_j\left(\cdot,\overline{\theta}_j\right)\right)-D_\theta F_{A|X=j}\left(\theta_0,\partial_a g_{0,j}\left(\cdot,\theta_0\right)\right)\right) \\ \times\left(D_\theta F_{A|X=j}\left(\overline{\theta}_j,\partial_a\widehat{g}_j\left(\cdot,\overline{\theta}_j\right)\right)-D_\theta F_{A|X=j}\left(\theta_0,\partial_a g_{0,j}\left(\cdot,\theta_0\right)\right)\right)'\end{array}\right]d\mu_j\left(\theta-\theta_0\right)\right]^{1/2} \\
& \times\left[\int\left[\widehat{F}_{A|X=j}-F_{A|X=j}\right]^2 d\mu_j\right]^{1/2},
\end{aligned}
$$

where for each $j,\overline{\theta}_j$ is some intermediate value between θ_j and $\theta_{0,j}$ that corresponds to the MVT w.r.t. the $j-th$ summand. Let ∂_{θ_l} denotes the $l-$th element of ∂_θ then

$$
\begin{aligned}
& \left|D_{\theta_l}F_{A|X=j}\left(\overline{\theta}_j,\partial_a\widehat{g}_j\left(\cdot,\overline{\theta}_j\right)\right)-D_{\theta_l}F_{A|X=j}\left(\theta_0,\partial_a g_{0,j}\left(\cdot,\theta_0\right)\right)\right| \\
\leq & \left|\begin{array}{c} q\left(\rho_j\left(a,\overline{\theta}_j,\partial_a\widehat{g}_j\left(\cdot,\overline{\theta}_j\right)\right)\right)\partial_{\theta_l}\rho_j\left(a,\overline{\theta}_j,\partial_a\widehat{g}_j\left(\cdot,\overline{\theta}_j\right)\right) \\ -q\left(\rho_j\left(a,\theta_0,\partial_a g_{0,j}\left(\cdot,\theta_0\right)\right)\right)\partial_{\theta_l}\rho_j\left(a,\theta_0,\partial_a g_{0,j}\left(\cdot,\theta_0\right)\right)\end{array}\right| \\
& +\left|\begin{array}{c} q\left(\rho_j\left(a,\overline{\theta}_j,\partial_a\widehat{g}_j\left(\cdot,\overline{\theta}_j\right)\right)\right)D_{\partial_a g}\rho_j\left(a,\overline{\theta}_j,\partial_a\widehat{g}_j\left(\cdot,\overline{\theta}\right)\right)\left[\partial_{\theta_l}\partial_a\widehat{g}_j\left(\cdot,\overline{\theta}_j\right)\right] \\ -q\left(\rho_j\left(a,\theta_0,\partial_a g_{0,j}\left(\cdot,\theta_0\right)\right)\right)D_{\partial_a g}\rho_j\left(a,\theta_0,\partial_a g_{0,j}\left(\cdot,\theta_0\right)\right)\left[\partial_{\theta_l}\partial_a g_{0,j}\left(\cdot,\theta_0\right)\right]\end{array}\right|.
\end{aligned}
$$

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 5 and continuity in θ of $\partial_a g_{0,j}\left(\cdot,\theta\right)$. Then it will follow by DCT that

$$
\left[(\theta-\theta_0)'\int\left[\begin{array}{c}\left[D_\theta F_{A|X}\left(\overline{\theta}_j,\partial_a\widehat{g}_j\left(\cdot,\overline{\theta}_j\right)\right)-D_\theta F_{A|X}\left(\theta_0,\partial_a g_{0,j}\left(\cdot,\theta_0\right)\right)\right] \\ \times\left[D_\theta F_{A|X}\left(\overline{\theta}_j,\partial_a\widehat{g}_j\left(\cdot,\overline{\theta}_j\right)\right)-D_\theta F_{A|X}\left(\theta_0,\partial_a g_{0,j}\left(\cdot,\theta_0\right)\right)\right]'\end{array}\right]d\mu\left(\theta-\theta_0\right)\right]^{1/2}=o_p\left(\left\|\theta-\theta_0\right\|\right).
$$

From Lemma 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 $\left|B_{12}\right|=o_p\left(N^{-1/2}\left\|\theta-\theta_0\right\|\right)$.

For B_2 :

For each j

$$
\begin{aligned}
& F_{A|X=j}\left(\theta,\partial_a\widehat{g}_j\left(\cdot,\theta\right)\right)+F_{A|X=j}\left(\theta_0,\partial_a\widehat{g}_j\left(\cdot,\theta_0\right)\right)-2\widehat{F}_{A|X=j} \\
= & \left(F_{A|X=j}\left(\theta,\partial_a\widehat{g}_j\left(\cdot,\theta\right)\right)-F_{A|X=j}\left(\theta_0,\partial_a\widehat{g}_j\left(\cdot,\theta_0\right)\right)\right) \\
& +2\left(F_{A|X=j}\left(\theta_0,\partial_a\widehat{g}_j\left(\cdot,\theta_0\right)\right)-F_{A|X}\left(\theta_0,\partial_a g_{0,j}\left(\cdot,\theta_0\right)\right)\right) \\
& -2\left(\widehat{F}_{A|X=j}-F_{A|X=j}\right),
\end{aligned}
$$

then we can write B_2 as

$$
\begin{aligned}
B_2 &= R^{-1/2}\sum_{j=1}^{J}\int\left[\begin{array}{c}\left[\nu_{R,j}\left(\theta,\partial_a\widehat{g}_j(\cdot,\theta)\right)-\nu_{R,j}\left(\theta_0,\partial_a\widehat{g}_j(\cdot,\theta_0)\right)\right]\\ \times\left[F_{A|X=j}\left(\theta,\partial_a\widehat{g}_j(\cdot,\theta)\right)-F_{A|X=j}\left(\theta_0,\partial_a\widehat{g}_j(\cdot,\theta_0)\right)\right]\end{array}\right]d\mu_j \\
&\quad +2R^{-1/2}\sum_{j=1}^{J}\int\left[\begin{array}{c}\left[\nu_{R,j}\left(\theta,\partial_a\widehat{g}_j(\cdot,\theta)\right)-\nu_{R,j}\left(\theta_0,\partial_a\widehat{g}_j(\cdot,\theta_0)\right)\right]\\ \times\left[F_{A|X=j}\left(\theta_0,\partial_a\widehat{g}_j(\cdot,\theta_0)\right)-F_{A|X}\left(\theta_0,\partial_a g_{0,j}(\cdot,\theta_0)\right)\right]\end{array}\right]d\mu_j \\
&\quad -2R^{-1/2}\sum_{j=1}^{J}\int\left(\nu_{R,j}\left(\theta,\partial_a\widehat{g}_j(\cdot,\theta)\right)-\nu_{R,j}\left(\theta_0,\partial_a\widehat{g}_j(\cdot,\theta_0)\right)\right)\left(\widehat{F}_{A|X=j}-F_{A|X=j}\right)d\mu_j \\
&= B_{21}+B_{22}+B_{23}.
\end{aligned}
$$

We first show $\int\left[\nu_{R,j}\left(\theta,\partial_a\widehat{g}_j(\cdot,\theta)\right)-\nu_{R,j}\left(\theta_0,\partial_a\widehat{g}_j(\cdot,\theta_0)\right)\right]^2 d\mu_j = o_p(1)$ for any j. By Lemma 6 $\partial_a\widehat{g}_j \in \mathcal{G}_j^{(1)}$ w.p.a. 1, and by Lemma 8 it suffices to show that $\|\partial_a\widehat{g}_j(\cdot,\theta)-\partial_a\widehat{g}_j(\cdot,\theta_0)\| \xrightarrow{p} 0$ 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 $\|\partial_a\widehat{g}_j(\cdot,\theta)-\partial_a g_{0,j}(\cdot,\theta)\|+\|\partial_a\widehat{g}_j(\cdot,\theta_0)-\partial_a g_{0,j}(\cdot,\theta_0)\|+\|\partial_a g_{0,j}(\cdot,\theta)-\partial_a g_{0,j}(\cdot,\theta_0)\|$, and the fact that the first two terms of the majorant converge to zero by Lemma 5 and the last term converges to zero by the continuity of $\partial_a g_{0,j}(\cdot,\theta)$ in θ. For B_{21}

$$
\begin{aligned}
B_{21} &= 2R^{-1/2}\sum_{j=1}^{J}\int\left[\begin{array}{c}\left[\nu_{R,j}\left(\theta,\partial_a\widehat{g}_j(\cdot,\theta)\right)-\nu_{R,j}\left(\theta_0,\partial_a\widehat{g}_j(\cdot,\theta_0)\right)\right]\\ \times\left[F_{A|X=j}\left(\theta,\partial_a\widehat{g}_j(\cdot,\theta)\right)-F_{A|X=j}\left(\theta_0,\partial_a\widehat{g}_j(\cdot,\theta_0)\right)\right]\end{array}\right]d\mu \\
&= 2R^{-1/2}\sum_{j=1}^{J}\int\left[\begin{array}{c}\left[\nu_{R,j}\left(\theta,\partial_a\widehat{g}_j(\cdot,\theta)\right)-\nu_{R,j}\left(\theta_0,\partial_a\widehat{g}_j(\cdot,\theta_0)\right)\right]\\ \times D_\theta F_{A|X=j}\left(\overline{\theta}_j,\partial_a\widehat{g}_j\left(\cdot,\overline{\theta}_j\right)\right)'\end{array}\right]d\mu\left(\theta-\theta_0\right),
\end{aligned}
$$

by Cauchy Schwarz inequality

$$
\begin{aligned}
|B_{21}| &\leq o_p\left(R^{-1/2}\right)\times \\
&\quad \max_{1\leq j\leq J}\left[\left(\theta-\theta_0\right)'\int\left[D_\theta F_{A|X=j}\left(\overline{\theta}_j,\partial_a\widehat{g}_j\left(\cdot,\overline{\theta}_j\right)\right)D_\theta F_{A|X=j}\left(\overline{\theta}_j,\partial_a\widehat{g}_j\left(\cdot,\overline{\theta}_j\right)\right)'\right]d\mu\left(\theta-\theta_0\right)\right]^{1/2} \\
&= o_p\left(R^{-1/2}\right)O_p\left(\|\theta-\theta_0\|\right) \\
&= o_p\left(N^{-1/2}\|\theta-\theta_0\|\right),
\end{aligned}
$$

the first inequality follows from the stochastic equicontinuity condition of Lemma 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 10 and Lemma 9 respectively that $|B_{22}| = o\left(N^{-1}\right)$ and $|B_{23}| = o\left(N^{-1}\right)$.

For B_3 :

For each j

$$\begin{aligned}\nu_{R,j}\left(\theta,\partial_a\widehat{g}_j\left(\cdot,\theta\right)\right)+\nu_{R,j}\left(\theta_0,\partial_a\widehat{g}_j\left(\cdot,\theta_0\right)\right) &= 2\nu_R\left(\theta_0,\partial_a g_{0,j}\left(\cdot,\theta_0\right)\right)\\ &\quad+\left(\nu_{R,j}\left(\theta,\partial_a\widehat{g}_j\left(\cdot,\theta\right)\right)-\nu_R\left(\theta_0,\partial_a g_{0,j}\left(\cdot,\theta_0\right)\right)\right)\\ &\quad+\left(\nu_{R,j}\left(\theta_0,\partial_a\widehat{g}_j\left(\cdot,\theta_0\right)\right)-\nu_R\left(\theta_0,\partial_a g_{0,j}\left(\cdot,\theta_0\right)\right)\right),\end{aligned}$$

then we can write B_3 as

$$\begin{aligned}B_3 &= 2S^{-1/2}\sum_{j=1}^{J}\int\nu_{R,j}\left(\theta_0,\partial_a g_{0,j}\left(\cdot,\theta_0\right)\right)\left(F_{A|X=j}\left(\theta,\partial_a\widehat{g}_j\left(\cdot,\theta\right)\right)-F_{A|X=j}\left(\theta_0,\partial_a\widehat{g}_j\left(\cdot,\theta_0\right)\right)\right)d\mu_j\\ &\quad+R^{-1/2}\sum_{j=1}^{J}\int\left[\begin{array}{c}\left[\nu_{R,j}\left(\theta,\partial_a\widehat{g}_j\left(\cdot,\theta\right)\right)-\nu_R\left(\theta_0,\partial_a g_{0,j}\left(\cdot,\theta_0\right)\right)\right]\\ \times\left[F_{A|X=j}\left(\theta,\partial_a\widehat{g}_j\left(\cdot,\theta\right)\right)-F_{A|X=j}\left(\theta_0,\partial_a\widehat{g}_j\left(\cdot,\theta_0\right)\right)\right]\end{array}\right]d\mu_j\\ &\quad+R^{-1/2}\sum_{j=1}^{J}\int\left[\begin{array}{c}\left[\nu_{R,j}\left(\theta_0,\partial_a\widehat{g}_j\left(\cdot,\theta_0\right)\right)-\nu_R\left(\theta_0,\partial_a g_{0,j}\left(\cdot,\theta_0\right)\right)\right]\\ \times\left[F_{A|X=j}\left(\theta,\partial_a\widehat{g}_j\left(\cdot,\theta\right)\right)-F_{A|X=j}\left(\theta_0,\partial_a\widehat{g}_j\left(\cdot,\theta_0\right)\right)\right]\end{array}\right]d\mu_j\\ &= B_{31}+B_{32}+B_{33}.\end{aligned}$$

For each j: we have $\left[\int\left[F_{A|X=j}\left(\theta,\partial_a\widehat{g}_j\left(\cdot,\theta\right)\right)-F_{A|X=j}\left(\theta_0,\partial_a\widehat{g}_j\left(\cdot,\theta_0\right)\right)\right]^2d\mu_j\right]^{1/2}=O_p\left(\|\theta-\theta_0\|\right)$ by Cauchy Schwarz inequality; from Donsker theorem and CMT, $\left[\int\left[\nu_R\left(\theta_0,\partial_a g_{0,j}\left(\cdot,\theta_0\right)\right)\right]^2d\mu_j\right]^{1/2}=O_p\left(1\right)$. Then it follows that $|B_{31}|\leq o_p\left(N^{-1/2}\|\theta-\theta_0\|\right)$. By a similar argument, using Cauchy Schwarz inequality, continuity of $\partial_a g\left(\cdot,\theta\right)$ in θ, Lemma 5, 6 and 8, $|B_{32}|$ and $|B_{33}|$ 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

$$\begin{aligned}B_4 &= 2S^{-1}\sum_{j=1}^{J}\int\nu_{R,j}\left(\theta_0,\partial_a g_{0,j}\left(\cdot,\theta_0\right)\right)\left(\nu_{R,j}\left(\theta,\partial_a\widehat{g}_j\left(\cdot,\theta\right)\right)-\nu_{R,j}\left(\theta_0,\partial_a\widehat{g}_j\left(\cdot,\theta_0\right)\right)\right)d\mu_j\\ &\quad+R^{-1}\sum_{j=1}^{J}\int\left[\nu_{R,j}\left(\theta,\partial_a\widehat{g}_j\left(\cdot,\theta\right)\right)-\nu_{R,j}\left(\theta_0,\partial_a\widehat{g}_j\left(\cdot,\theta_0\right)\right)\right]^2d\mu_j\\ &\quad+R^{-1}\sum_{j=1}^{J}\int\left[\begin{array}{c}\left[\nu_{R,j}\left(\theta_0,\partial_a\widehat{g}_j\left(\cdot,\theta_0\right)\right)-\nu_R\left(\theta_0,\partial_a g_{0,j}\left(\cdot,\theta_0\right)\right)\right]\\ \times\left[\nu_{R,j}\left(\theta,\partial_a\widehat{g}_j\left(\cdot,\theta\right)\right)-\nu_{R,j}\left(\theta_0,\partial_a\widehat{g}_j\left(\cdot,\theta_0\right)\right)\right]\end{array}\right]d\mu_j\\ &= B_{41}+B_{42}+B_{43}.\end{aligned}$$

By repeatedly using Cauchy Schwarz inequality, continuity of $\partial_a g\left(\cdot,\theta\right)$ in θ, and Lemma 5,6 and 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 10.■

Proof of Theorem 4. From (19) 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 θ_0. It is easy to see that, for $j = 1, \ldots, J$

$$\begin{aligned}\left\|\widehat{g}_j\left(\cdot, \widehat{\theta}\right) - \widehat{g}_j\left(\cdot, \theta_0\right)\right\|_{\infty} &= O_p\left(\left\|\widehat{\theta} - \theta_0\right\|\right) \\ &= O_p\left(N^{-1/2}\right),\end{aligned}$$

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 $\sqrt{Nh}\left|\widehat{g}_j\left(a, \widehat{\theta}\right) - \widehat{g}_j\left(a, \theta_0\right)\right|$ $= o_p(1)$. 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 $\left\|\widehat{g}_j\left(\cdot, \widehat{\theta}\right) - \widehat{g}_j\left(\cdot, \theta_0\right)\right\|_{\infty} = O_p\left(N^{-1/2}\right)$.■

A.2 Proofs of Lemmas 1-10

These lemmas are used in the proofs of Theorem 1 - 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 1. We can write, for any $1 \leq k, j \leq J$

$$\widehat{p}_{X'|X}(k|j) - p_{X'|X}(k|j) = \frac{\widehat{p}_{X',X}(k,j) - p_{X',X}(k,j)}{p_X(j)} - \frac{\widehat{p}_{X'|X}(k|j)}{p_X(j)\,\widehat{p}_X(j)}\left(\widehat{p}_X(j) - p_X(j)\right).$$

Given the simple nature of our DGP, by standard CLT and LLN, we have $\widehat{p}_{X',X}(k,j) - p_{X',X}(k,j) = O_p\left(N^{-1/2}\right), \widehat{p}_X(j) - p_X(j) = O_p\left(N^{-1/2}\right)$ and $\widehat{p}_X(j)^{-1} = O_p(1)$, so it follows that $\widehat{p}_{X'|X}(k|j) - p_{X'|X}(k|j) = O_p\left(N^{-1/2}\right)$ 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 $\left(\left(\widehat{\mathcal{L}} - \mathcal{L}\right) m\right)_j = \beta \sum_{k=1}^{J}(\widehat{p}(k|j) - p(k|j)) m_i = O_p\left(N^{-1/2}\right)$ 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. For any $j = 1, \ldots, J$ and $\theta \in \Theta$, $\widetilde{r}_\theta(j)$ is defined in (17) with $w_{itN}(j) = \mathbf{1}[x_{it} = j] / \widehat{p}_X(j)$ and define $\widehat{r}_\theta(j) = \sum_{i=1,t=1}^{N,T} w_{itN}(j) u_\theta(a_{it}, x_{it}, \varepsilon_{it})$. Then we write

$$\widetilde{r}_\theta(j) - r_\theta(j) = (\widehat{r}_\theta(j) - r_\theta(j)) + (\widetilde{r}_\theta(j) - \widehat{r}_\theta(j)), \tag{36}$$

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

$$\begin{aligned}
\widehat{r}_\theta(j) - r_\theta(j) &= \frac{1}{\widehat{p}_X(j)} \frac{1}{NT} \sum_{i=1,t=1}^{N,T} \mathbf{1}[x_{it} = j] (u_\theta(a_{it}, x_{it}, \varepsilon_{it}) - r_\theta(j)) \\
&= \frac{1}{\widehat{p}_X(j)} \frac{1}{NT} \sum_{i=1,t=1}^{N,T} \upsilon_{\theta,it} \mathbf{1}[x_{it} = j],
\end{aligned}$$

where for each θ, $\upsilon_{\theta,it} = u_\theta(a_{it}, x_{it}, \varepsilon_{it}) - r_\theta(x_{it})$ is a zero mean random variable, note that $\mathbf{1}[x_{it} = j] \times (r_\theta(x_{it}) - r_\theta(j)) = 0$ for all i, j and t. Define $\Upsilon_{N,j}(\theta)$ as the sample average of i.i.d. random variables $\left\{ \sum_{t=1}^{T} \frac{1}{T} \upsilon_{\theta,it} \mathbf{1}[x_{it} = j] \right\}_{i=1}^{N}$, 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 θ by standard CLT. We want to obtain the uniform rate of convergence of $\Upsilon_{N,j}(\theta)$ over Θ. 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 θ. The pointwise rate can then be made uniform by Lipschitz continuity of $\upsilon_{\theta,it}$ (in θ) and compactness of Θ. More precisely, we first show that $\sup_{\theta \in \Theta} \mathrm{var}(\Upsilon_{N,j}(\theta)) = O(N^{-1})$. Since $\mathrm{var}(\Upsilon_{N,j}(\theta))$ is just a variance of $\sum_{t=1}^{T} \frac{1}{T} \upsilon_{\theta,it} \mathbf{1}[x_{it} = j]$ by divided by N, the numerator takes the following form

$$\begin{aligned}
\mathrm{var}\left(\frac{1}{T} \sum_{t=1}^{T} \upsilon_{\theta,it} \mathbf{1}[x_{it} = j] \right) &= \frac{1}{T} \sum_{t=1}^{T} \mathrm{var}(\upsilon_{\theta,it} \mathbf{1}[x_{it} = j]) \\
&\quad + \frac{2}{T} \sum_{s=1}^{T-1} \left(1 - \frac{s}{T}\right) Cov(\upsilon_{\theta,i0} \mathbf{1}[x_{i0} = j], \upsilon_{\theta,is} \mathbf{1}[x_{is} = j]), \\
&= Y_{\theta,1,j} + Y_{\theta,2,j}.
\end{aligned}$$

The covariance structure in $Y_{\theta,2,j}$ follows from the strict stationarity assumption, which also implies we can write $Y_{\theta,1,j} = E\left[\left[\upsilon_{\theta,it}^2 | x_{it}\right] \mathbf{1}[x_{it} = j]\right]$. Since $u_\theta(a, x, \varepsilon)$ is continuous in θ for all a, x and ε, it follows that $\sup_{\theta \in \Theta} Y_{\theta,1,j} < \infty$. For the covariance term, by Cauchy-Schwarz inequality, $Cov(\upsilon_{\theta,i0} \mathbf{1}[x_{i0} = j], \upsilon_{\theta,is} \mathbf{1}[x_{is} = j]) \leq E\left[\upsilon_{\theta,i0}^2 \mathbf{1}[x_{i0} = j]\right]^2 < \infty$, since $\sup_{\theta \in \Theta} \left|\upsilon_{\theta,i0,j}^2\right| < \infty$, 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 θ, 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

$$\Pr\left[\left|N\Upsilon_{N,j}(\theta)\right| > N\delta_N\right] \leq 2\exp\left\{-\frac{N^2\delta_N^2}{4Var\left(N\Upsilon_{N,j}(\theta)\right)+2CN\delta_N}\right\}. \tag{37}$$

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 Θ, 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 (37) 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 $\sup_{\theta}\left|\Upsilon_{N,j}(\theta)\right| = o_p\left(N^{-\lambda}\right)$. Then w.p.a. 1

$$\begin{aligned}
\sup_{\theta\in\Theta}\left|\widehat{r}_{\theta}(j)-r_{\theta}(j)\right| &\leq \frac{\max_{1\leq j\leq J}\sup_{\theta\in\Theta}\left|\Upsilon_{N,j}(\theta)\right|}{\min_{1\leq j\leq J}p_X(j)} \\
&= o_p\left(N^{-\lambda}\right).
\end{aligned}$$

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

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

where $\widehat{\varepsilon}_{it} = \chi\left(\widehat{F}_{A|X}\left(a_{it}|x_{it}\right)\right)$ with $\chi \equiv Q_{\varepsilon}^{-1}$. Using mean value expansion, $u_{\theta}\left(a_{it},x_{it},\widehat{\varepsilon}_{it}\right) - u_{\theta}\left(a_{it},x_{it},\varepsilon_{it}\right) = \frac{\partial}{\partial\varepsilon}u_{\theta}\left(a_{it},x_{it},\overline{\varepsilon}_{it}\right)\chi'\left(\overline{F}_{A|X}\left(a_{it}|x_{it}\right)\right)\left(\widehat{F}_{A|X}\left(a_{it}|x_{it}\right)-F_{A|X}\left(a_{it}|x_{it}\right)\right)$, where $\overline{\varepsilon}_{it}$ and

$\overline{F}_{A|X}(a_{it}|x_{it})$ are some intermediate points between $\widehat{\varepsilon}_{it}$ and ε_{it}, and, $\widehat{F}_{A|X}(a_{it}|x_{it})$ and $F_{A|X}(a_{it}|x_{it})$, respectively. Then it follows that

$$\begin{aligned}\widetilde{r}_\theta(j)-\widehat{r}_\theta(j) &= \frac{1}{NT}\sum_{i=1,t=1}^{N,T} w_{itN}(j)\left(u_\theta(a_{it},x_{it},\widehat{\varepsilon}_{it})-u_\theta(a_{it},x_{it},\varepsilon_{it})\right)\\ &= \frac{1}{NT}\sum_{i=1,t=1}^{N,T}\frac{\mathbf{1}[x_{it}=j]}{p_X(j)}\varkappa_\theta(a_{it},x_{it},\varepsilon_{it})\left(\widehat{F}_{A|X}(a_{it}|x_{it})-F_{A|X}(a_{it}|x_{it})\right)+O_p\left(N^{-1}\right),\end{aligned}$$

where $\varkappa_\theta(a_{it},x_{it},\varepsilon_{it})=\frac{\partial}{\partial\varepsilon}u_\theta(a_{it},x_{it},\varepsilon_{it})\chi'\left(F_{A|X}(a_{it}|x_{it})\right)$. In addition, the $O_p(N^{-1})-$term holds uniformly over θ and j, this follows from Markov inequality since $\frac{\partial^2}{\partial\varepsilon^2}u$ and χ'' are uniformly bounded over all of their arguments, $\max_{1\leq j\leq J}|\widehat{p}_X(j)-p_X(j)|=O_p\left(N^{-1/2}\right)$, and, $\max_{1\leq j\leq J}\sup_{a\in A}|\widehat{F}_{A|X}(a|j)-F_{A|X}(a|j)|=O_p\left(N^{-1/2}\right)$ by Lemma 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 $\zeta_\theta(\omega(x_{it};\cdot),x_{it};j)=\frac{1}{p_X(j)}\int\omega(x_{it};a_{js})\left[\int\varkappa_\theta(a_{js},x_{it},\varepsilon_{js})\mathbf{1}[x_{it}=j]\frac{f_{A,X,\mathcal{E}}(a_{js},x_{it},\varepsilon_{js})}{p(x_{it})}d\varepsilon_{js}\right]da_{js}$ and $f_{A,X,\mathcal{E}}$ denotes the joint continuous-discrete density of $(a_{it},x_{it},\varepsilon_{it})$. Note that ζ_θ is random with respect to ω_{it} and x_{it}, and $E[\omega(x_{it};\cdot)|x_{it}]=0$, so ζ_θ has zero mean. Given the boundedness and smoothness conditions on $\varkappa_\theta$, then $\frac{1}{NT}\sum_{i=1,t=1}^{N,T}\zeta_\theta(\omega(x_{it};\cdot),x_{it};j)$ 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 $\widetilde{r}_{\theta}(j)=r_{\theta}(j)+\widetilde{r}_{\theta}^{R}(j)$ with

$$\begin{aligned}
\widetilde{r}_{\theta}^{R}(j) &= \frac{1}{p_{X}(j)}\frac{1}{NT}\sum_{i=1,t=1}^{N,T}\mathbf{1}\left[x_{it}=j\right]\left(u_{\theta}\left(a_{it},x_{it},\varepsilon_{it}\right)-r_{\theta}(j)\right) \\
&\quad +\frac{1}{NT}\sum_{i=1,t=1}^{N,T}\zeta_{\theta}\left(\omega\left(x_{it};\cdot\right),x_{it};j\right)+o_{p}\left(N^{-\lambda}\right) \\
&= o_{p}\left(N^{-\lambda}\right),
\end{aligned}$$

where the smaller order term holds uniformly over j and θ.■

Proof of Lemma 3. Since $0<\|\mathcal{L}\|<1$ 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, $(I-\mathcal{L})^{-1}$ and $\left(I-\widehat{\mathcal{L}}\right)^{-1}$ are bounded linear operators since $\left\|(I-\mathcal{L})^{-1}\right\|\leq(1-\|\mathcal{L}\|)^{-1}<\infty$ 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 $\theta\in\Theta$, $\widehat{m}_{\theta}(j)$ is defined in (18), we write $\widehat{m}_{\theta}(j)=\left(I-\widehat{\mathcal{L}}\right)^{-1}\left(r_{\theta}(j)+\widetilde{r}_{\theta}^{R}(j)\right)$, given the results from Lemma 2, it follows that $\max_{1\leq j\leq J}\sup_{\theta\in\Theta}\left|\left(I-\widehat{\mathcal{L}}\right)^{-1}\widetilde{r}_{\theta}^{R}(j)\right|=o_{p}\left(N^{-\lambda}\right)$, since $\left\|\left(I-\widehat{\mathcal{L}}\right)^{-1}\right\|=O_{p}(1)$. For first term, we can write $\left(I-\widehat{\mathcal{L}}\right)^{-1}r_{\theta}(j)=m_{\theta}(j)+\widehat{m}_{\theta}^{A}(j)$ where $\widehat{m}_{\theta}^{A}(j)=\left(I-\widehat{\mathcal{L}}\right)^{-1}\left(\widehat{\mathcal{L}}-\mathcal{L}\right)m_{\theta}(j)$. Since we know $\left\|\left(I-\widehat{\mathcal{L}}\right)^{-1}\right\|=O_{p}(1)$ from earlier, from Lemma 1 $\left\|\widehat{\mathcal{L}}-\mathcal{L}\right\|=O_{p}\left(N^{-1/2}\right)$, and, $\max_{1\leq j\leq J}\sup_{\theta\in\Theta}\left|m_{\theta}(j)\right|=O(1)$ as $m_{\theta}(j)$ is a continuous function on a compact set Θ 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 4. The empirical analogue of (7) 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 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 $m = (m_1 \dots m_J)' \in \mathbb{R}^J, a$ and j

$$
\begin{aligned}
\left(\widehat{\mathcal{H}} - \mathcal{H}\right) m(j,a) &= \sum_{k=1}^{J} m_k \left(\frac{\widehat{p}_{X',X,A}(k,j,a)}{\widehat{p}_{X,A}(j,a)} - \frac{p_{X',X,A}(k,j,a)}{p_{,X,A}(j,a)} \right) \\
&= \sum_{k=1}^{J} m_k \left(\frac{\widehat{p}_{X',X,A}(k,j,a) - p_{X',X,A}(k,j,a)}{p_{X,A}(j,a)} \right) \\
&\quad - \sum_{k=1}^{J} m_k \left(\frac{\widehat{p}_{X',X,A}(k,j,a)}{\widehat{p}_{X,A}(j,a)\, p_{X,A}(j,a)} \left(\widehat{p}_{X,A}(j,a) - p_{X,A}(j,a)\right) \right),
\end{aligned} \tag{38}
$$

where

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

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 $\max_{1 \le j,k \le J} \sup_{a \in A_N} |I_{12}(k,j,a)| = O(h^4)$ 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 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

$$\text{var}\left(\sqrt{NTh}I_{11}(k,j,a)\right) = p_{X',X,A}(k,j,a)\,\kappa_2(K) + o(1),$$

where the display above for any j and k uniformly over A_N. Combining terms we have

$$\begin{aligned}
& \max_{1\leq j,k\leq J}\sup_{a\in A_N}\left|\sum_{k=1}^{J} m_k\left(\frac{\widehat{p}_{X',X,A}(k,j,a) - p_{X',X,A}(k,j,a)}{p_{X,A}(j,a)}\right)\right| \\
\leq\ & J\frac{\max_{1\leq j\leq J}|m_j|}{\min_{1\leq j\leq J}\inf_{a\in A_N}|p_{X,A}(j,a)|}\times\max_{1\leq j,k\leq J}\sup_{a\in A_N}|\widehat{p}_{X',X,A}(k,j,a) - p_{X',X,A}(k,j,a)| \\
=\ & O_p\left(h^4 + \frac{N^{\xi}}{\sqrt{Nh}}\right),
\end{aligned}$$

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

$$\begin{aligned}
I_{22}(j,a) &= \frac{1}{2}h^4\mu_4(K)\frac{\partial^4}{\partial a^2}p_{X,A}(j,a) + o\left(h^2\right), \\
\text{var}\left(\sqrt{NTh}I_{21}(k,j,a)\right) &= p_{X,A}(j,a)\,\kappa_2(K) + o(1),
\end{aligned}$$

then we have

$$\begin{aligned}
& \max_{1\leq j,k\leq J}\sup_{a\in A_N}\left|\sum_{k=1}^{J} m_k\left(\frac{\widehat{p}_{X',X,A}(k,j,a)}{\widehat{p}_{X,A}(j,a)\,p_{X,A}(j,a)}(\widehat{p}_{X,A}(j,a) - p_{X,A}(j,a))\right)\right| \\
\leq\ & J\frac{\max_{1\leq j\leq J}|m_j|}{\min_{1\leq j\leq J}\inf_{a\in A_N}|p_{X,A}(j,a)|^2}\times\max_{1\leq j\leq J}\sup_{a\in A_N}|\widehat{p}_{X,A}(j,a) - p_{X,A}(j,a)| \\
=\ & O_p\left(h^4 + \frac{N^{\xi}}{\sqrt{Nh}}\right).
\end{aligned}$$

So we can write for each j

$$\begin{aligned}
\left(\widehat{\mathcal{H}} - \mathcal{H}\right)m(j,a) &= \sum_{k=1}^{J} m_k\left(\frac{\widehat{p}_{X',X,A}(k,j,a) - p_{X',X,A}(k,j,a)}{p_{X,A}(j,a)}\right) \\
&\quad -\sum_{k=1}^{J} m_k\left(\frac{p_{X',X,A}(k,j,a)}{p^2_{X,A}(j,a)}(\widehat{p}_{X,A}(j,a) - p_{X,A}(j,a))\right) + W_{N,j}(a;m) \\
&= B_{N,j}(a;m) + V_{N,j}(a;m) + W_{N,j}(a;m),
\end{aligned}$$

where

$$B_{N,j}(a;m) = \frac{1}{2}h^4\mu_4(K)\sum_{k=1}^{J} m_k\left(\frac{\frac{\partial^4}{\partial a^4}p_{X',X,A}(k,j,a)}{p_{X,A}(j,a)} + \frac{p_{X',X,A}(k,j,a)\frac{\partial^4}{\partial a^4}p_{X,A}(j,a)}{p^2_{X,A}(j,a)}\right), \tag{39}$$

$$V_{N,j}(a;m) = \sum_{k=1}^{J} m_k\left(\begin{array}{c}\frac{1}{p_{X,A}(j,a)}\frac{1}{NT}\sum_{i=1,t=1}^{N,T}\left(\begin{array}{c}\mathbf{1}[x_{it+1}=k,x_{it}=j]K_h(a_{it}-a)\\ -E[\mathbf{1}[x_{it+1}=k,x_{it}=j]K_h(a_{it}-a)]\end{array}\right)\\ -\frac{p_{X',X,A}(k,j,a)}{p^2_{X,A}(j,a)}\frac{1}{NT}\sum_{i=1,t=1}^{N,T}\left(\begin{array}{c}\mathbf{1}[x_{it}=j]K_h(a_{it}-a)\\ -E[\mathbf{1}[x_{it}=j]K_h(a_{it}-a)]\end{array}\right)\end{array}\right), \tag{40}$$

$$W_{N,j}(a;m) = \sum_{k=1}^{J} m_k\left(\begin{array}{c}\frac{1}{p_{X,A}(j,a)}\left(\frac{\widehat{p}_{X',X,A}(k,j,a)}{\widehat{p}_{X,A}(j,a)} - \frac{p_{X',X,A}(k,j,a)}{p_{X,A}(j,a)}\right)\\ \times(\widehat{p}_{X,A}(j,a)-p_{X,A}(j,a))\end{array}\right). \tag{41}$$

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 $\max_{1\leq j\leq J}\sup_{a\in A_N} W_{N,j}(a;m) = o_p(B_{N,j}(a;m)+V_{N,j}(a;m))$ 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

$$\widehat{g}^A_j(a,\theta) = \widehat{g}^B_j(a,\theta) + \widehat{g}^S_j(a,\theta) + W_{N,j}(a;m_\theta),$$

where, from (39) - (40), $\widehat{g}^B_j(a,\theta) = B_{N,j}(a;m_\theta)$ and $\widehat{g}^S_j(a,\theta) = V_{N,j}(a;m_\theta)$. It also follows that these terms have the desired rate of convergence that holds uniformly over Θ as well since $\mathcal{H}$ is independent of θ and m_θ is a vector of $J-$real value functions that are continuous on Θ. Finally, we define $\widehat{g}^R_j(a,\theta)$ to be $W_{N,j}(a;m_\theta)+\widehat{\mathcal{H}}\widehat{m}^R_\theta(j,a)$. By the previous reasoning $W_{N,j}(a;m_\theta)$ already has the desired stochastic order so the proof of Lemma 4 will be complete if we can show, generally, that $\max_{1\leq j\leq J}\sup_{\theta,a\in\Theta\times A_N}\left|\widehat{\mathcal{H}}\widehat{m}^R_\theta(j,a)\right| = o_p\left(h^4+\frac{N^{\xi}}{\sqrt{Nh}}\right)$. 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 $\max_{1\leq j\leq J}\sup_{\theta,a\in\Theta\times A_N}\left|\widehat{\mathcal{H}}\widehat{m}^R_\theta(j,a)\right| = o_p\left(N^{-\lambda}\right)$.■

Proof of Lemma 5. When $l=0$, this follows from Lemma 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 (38) we obtain

$$\begin{aligned}
\frac{\partial}{\partial a}\left(\widehat{\mathcal{H}}-\mathcal{H}\right)m(j,a) &= \sum_{k=1}^{J} m_k \frac{\partial}{\partial a}\left(\frac{\widehat{p}_{X',X,A}(k,j,a)-p_{X',X,A}(k,j,a)}{p_{X,A}(j,a)}\right) \\
&\quad -\sum_{k=1}^{J} m_k \frac{\partial}{\partial a}\left(\frac{\widehat{p}_{X',X,A}(k,j,a)}{\widehat{p}_{X,A}(j,a)\,p_{X,A}(j,a)}\left(\widehat{p}_{X,A}(j,a)-p_{X,A}(j,a)\right)\right) \\
&= \sum_{k=1}^{J} m_k \left(\begin{array}{c} \frac{1}{p_{X,A}(j,a)}\frac{\partial}{\partial a}\left(\widehat{p}_{X',X,A}(k,j,a)-p_{X',X,A}(k,j,a)\right) \\ -\frac{\frac{\partial}{\partial a}p_{X,A}(j,a)}{p^2_{X,A}(j,a)}\left(\widehat{p}_{X',X,A}(k,j,a)-p_{X',X,A}(k,j,a)\right)\end{array}\right) \\
&\quad -\sum_{k=1}^{J} m_k \left(\begin{array}{c} \frac{\widehat{p}_{X',X,A}(k,j,a)}{\widehat{p}_{X,A}(j,a)p_{X,A}(j,a)}\frac{\partial}{\partial a}\left(\widehat{p}_{X,A}(j,a)-p_{X,A}(j,a)\right) \\ +\left(\frac{\partial}{\partial a}\frac{\widehat{p}_{X',X,A}(k,j,a)}{\widehat{p}_{X,A}(j,a)p_{X,A}(j,a)}\right)\left(\widehat{p}_{X,A}(j,a)-p_{X,A}(j,a)\right)\end{array}\right).
\end{aligned}$$

As seen in the proof of Lemma 4, it will be sufficient to show that $\max_{1\leq j,k\leq J}\sup_{a\in A_N}|\frac{\partial}{\partial a}\widehat{p}_{X',X,A}(k,j,a)$ $-\frac{\partial}{\partial a}p_{X',X,A}(k,j,a)| = o_p(1)$, and, $\max_{1\leq j,k\leq J}\sup_{a\in A_N}|\frac{\partial}{\partial a}\widehat{p}_{X,A}(j,a)-\frac{\partial}{\partial a}p_{X,A}(j,a)| = o_p(1)$ 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 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, it then follows that $\max_{1\leq j,k\leq J}\sup_{a\in A_N}\left|\frac{\partial}{\partial a}\widehat{p}_{X',X,A}(k,j,a)-\frac{\partial}{\partial a}p_{X',X,A}(k,j,a))\right|$ $=O_p(h^4+\frac{N^{\xi}}{\sqrt{Nh^3}})$. Similarly one can show $\max_{1\leq j,k\leq J}\sup_{a\in A_N}\left|\frac{\partial}{\partial a}\widehat{p}_{X,A}(j,a)-\frac{\partial}{\partial a}p_{X,A}(j,a)\right|=O_p(h^4$ $+\frac{N^{\xi}}{\sqrt{Nh^3}})$. It is easy to see that choosing $h=O\left(N^{-1/7}\right)$ will imply $\max_{1\leq j\leq J}\sup_{\theta,a\in\Theta\times A_N}|\frac{\partial}{\partial a}\widehat{g}_{\theta}(j,a)$ $-\frac{\partial}{\partial a}g_{\theta}(j,a)|=o_p(1)$.■

Proof of Lemma 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 5. Consider the case when $p=1$ and $l=0$, for all $1\leq j\leq J, 1\leq k\leq L$ and $\lambda<1/2$, the exact same arguments used in proofing Lemma 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 $\max_{1\leq j\leq J}\sup_{\theta,a\in\Theta\times A_N}\left|\frac{\partial}{\partial\theta_k}\widetilde{r}^R_{\theta}(j)\right|=o_p\left(N^{-\lambda}\right)$, and since $\mathcal{L}$ is independent of θ, the same arguments found in Lemma 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_θ, m_θ) 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 θ. Since $\mathcal{H}$ is independent of θ, the arguments used in Lemma 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 (39) - (41). 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 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

$$\lim_{N\to\infty}\left(E\left[\sup_{\left\|\left(a'-a,\theta'-\theta,\partial_a g_j'-\partial_a g_j\right)\right\|<\delta_N}\left|\mathbf{1}\left[\varepsilon_i\leq\rho_j\left(a',\theta',\partial_a g_j'\right)\right]-\mathbf{1}\left[\varepsilon_i\leq\rho_j(a,\theta,\partial_a g_j)\right]\right|^2\right]\right)^{1/2} = 0. \tag{42}$$

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

$$\begin{aligned}
\left|\rho_j\left(a',\theta',\partial_a g_j'\right)-\rho_j(a,\theta,\partial_a g_j)\right| &\leq C_0\left\{\left\|(a'-a,\theta'-\theta)\right\|+\left\|\partial_a g_j'-\partial_a g_j\right\|_{\mathcal{G}}\right\} \\
&\quad +o\left(\left\|(a'-a,\theta'-\theta)\right\|^2+\left\|\partial_a g_j'-\partial_a g_j\right\|_{\mathcal{G}}^2\right) \\
&\leq C_0\delta_N + o(\delta_N),
\end{aligned}$$

this follows from Taylor's theorem in Banach Space since ρ_j is twice Fréchet differentiable, see Chapter 4 of Zeidler (1986). Ignoring the smaller order term, this implies

$$\begin{aligned}
\rho_j(a,\theta,\partial_a g_j)-C_0\delta_N &\leq \rho_j\left(a',\theta',\partial_a g_j'\right)\leq\rho_j(a,\theta,\partial_a g_j)+C_0\delta_N, \\
\rho_j(a,\theta,\partial_a g_j)-C_0\delta_N &\leq \rho_j(a,\theta,\partial_a g_j)\leq\rho_j(a,\theta,\partial_a g_j)+C_0\delta_N.
\end{aligned}$$

Combining the inequalities above, it follows that $\sup_{\left\|\left(a'-a,\theta'-\theta,\partial_a g_j'-\partial_a g_j\right)\right\|<\delta_N} \left|\mathbf{1}\left[\varepsilon_i \leq \rho_j\left(a',\theta',\partial_a g_j'\right)\right]\right.$ $\left.-\mathbf{1}\left[\varepsilon_i \leq \rho_j\left(a,\theta,\partial_a g_j\right)\right]\right|$ 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_ε, (42) 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 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 (42), Lemma 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 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

$$\begin{aligned}\mathfrak{F}_{1,N}(a,j) &= \frac{1}{T\widehat{p}_X(j)}\times\frac{1}{\sqrt{N}}\sum_{i=1,t=1}^{N,T}\left(\mathbf{1}\left[a_{it}\leq a,x_{it}=j\right]-F_{A,X}(a,j)\right),\\ \mathfrak{F}_{2,N}(a,j) &= -\frac{\sqrt{T}F_{A|X}(a|j)}{\widehat{p}_X(j)}\times\sqrt{N}\left(\widehat{p}_X(j)-p_X(j)\right).\end{aligned}$$

Define $\mathcal{C}_a=\{y_a\in\mathbb{R}:y_a\leq a\}$, 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, $\frac{1}{\sqrt{NT}}\sum_{i=1,t=1}^{N,T}\left(\mathbf{1}\left[a_{it}\leq\cdot,x_{it}=j\right]-F_{A,X}(\cdot,x)\right)$ 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 $\{\frac{1}{\sqrt{NT}}\sum_{i=1,t=1}^{N,T}\left(\mathbf{1}\left[a_{it}\leq\cdot,x_{it}=j\right]-F_{A,X}(\cdot,j)\right)\}$. 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 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 9 with Slutsky theorem can be used to complete the proof.■

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 (2007): "Sequential Estimation of Dynamic Discrete Games," *Econometrica*, **75**, 1-53.

[4] Aguirregabiria, V., and P. Mira (2008): "Dynamic Discrete Choice Structural Models: A Survey," *Journal of Econometrics,* forthcoming.

[5] Ai. C., and X. Chen (2003): "Efficient Estimation of Models with Conditional Moment Restrictions Containing Unknown Functions," *Econometrica*, **71**, 1795-1844.

[6] Andrews, D.W.K. (1994a): "Asymptotics for Semiparametric Econometric Models via Stochastic Equicontinuity," *Econometrica*, **62**, 43-72.

[7] Andrews, D.W.K. (1994b): "Empirical Processes Methods in Econometrics," *Handbook of Econometrics*, vol. 4, eds. R.F. Engle and D. McFadden. North-Holland.

[8] Andrews, D.W.K. (1995): "Nonparametric Kernel Estimation for Semiparametric Econometric Models," *Econometric Theory,* **11**, 560-596.

[9] 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.

[10] Arak, T. and A. Zaizsev (1988): "Uniform Limit Theorems for Sums of Independent Random Variables," *Publ. Steklov Math. Institute* **1**.

[11] Arcidiacono, P. and R.A. Miller (2008): "CCP Estimation of Dynamic Discrete Choice Models with Unobserved Heterogeneity," *Working Paper.*

[12] 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.

[13] Athey, S.C. (2002): "Monotone Comparative Statics Under Uncertainty," *Quarterly Journal of Economics*, **117**, 187-223.

[14] Bajari, P., C. L. Benkard, and J. Levin (2007): "Estimating Dynamic Models of Imperfect Competition," *Econometrica*, **75**, 1331-1370.

[15] Bajari, P., V. Chernozhukov, H. Hong and D. Nekipelov (2009): "Semiparametric Estimation of a Dynamic Game of Incomplete Information," *Working paper.*

[16] Berry, S., and A. Pakes (2002): "Estimation from the Optimality Conditions for Dynamic Controls," *Mimeo.*

[17] Bose, A. (1988): "Edgeworth Correction by Bootstrap in Autoregressions," *Annals of Statistics*, **16**, 1709-1722.

[18] Bose, A. (1990): "Bootstrap in Moving Average Models," *Annals of the Institute of Statistical Mathematics*, **42**, 753-768.

[19] Bosq, D. (1996): *Nonparametric Statistics for Stochastic Processes*, Springer.

[20] Brown, D. J. and M. H. Wegkamp (2002): "Weighted Minimum Mean-Square Distance from Independence Estimation," *Econometrica,* **70**, 2035-2051.

[21] Carrasco, M., and J.P. Florens (2000): "Generalization Of Gmm To A Continuum Of Moment Conditions," *Econometric Theory*, **16**, 797-834.

[22] Chen, X., and X. Shen (1998): "Sieve Extremum Estimates for Weakly Dependent Data," *Econometrica*, **66**, 289-314.

[23] 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.

[24] Chen, X., and D. Pouzo (2008): "Estimation of Nonparametric Conditional Moment Models with Possibly Nonsmooth Moments," *Journal of Econometrics*, forthcoming.

[25] Domínguez, M.A., and I.N. Lobato (2004): "Consistent Estimation of Models defined by Conditional Moment Restrictions," *Econometrica*, **72**, 1601-1615.

[26] Doraszelski, U. and M. Satterthwaite (2007): "Computable Markov-Perfect Industry Dynamics: Existence, Purification and Multiplicity," *Mimeo*, Harvard University.

[27] 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.

[28] Efron, B. (1979): "Bootstrap Methods: Another look at the Jacknife," *Annals of Statistics*, **7**, 1-26.

[29] Efron, B., and R.J. Tibshirani (1993): *An Introduction to the Bootstrap,* Chapman & Hall, New York.

[30] Fan, J., (1992): "Design-adaptive nonparametric regression," *Journal of the American Statistical Association*, **87**, 998-1004.

[31] Hall, P. (1992): *The Bootstrap and Edgeworth Expansion,* Springer-Verlag, New York.

[32] Hansen, L.P. and K.J. Singleton (1982): "Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models," *Econometrica*, **50**, 1269-1286.

[33] Härdle, W., J. Horowitz and J.P. Kreiss (2003): "Bootstrap methods for time series," *International Statistical Review*, **71**, 435-459.

[34] Heckman, J. and S. Navarro (2007): "Dynamic Discrete Choice and Dynamic Treatment Effects," *Journal of Econometrics*, **136**, 341-396.

[35] Horowitz, J. (1998): "A Smoothed Maximum Score Estimator for the Binary Response Model," *Econometrica*, **60**, 1327-1352.

[36] Horowitz, J. (2001): "The Bootstrap," *Handbook of Econometrics*, vol. 5, eds. J. Heckman and E. Leamer. North-Holland.

[37] Hotz, V., and R.A. Miller (1993): "Conditional Choice Probabilities and the Estimation of Dynamic Models," *Review of Economic Studies*, **60**, 497-531.

[38] 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.

[39] Hu. Y. and M. Shum (2009): "Nonparametric Identification of Dynamic Models with Unobserved State Variables," *Working Paper*.

[40] Ichimura, H. and S. Lee (2006): "Characterization of the Asymptotic Distribution of Semiparametric M-Estimators," *CEMMAP Working Paper*.

[41] Jofre-Bonet, M. and M. Pesendorfer (2003): "Estimation of a Dynamic Auction Game," *Econometrica*, **71**, 1443-1489.

[42] Judd, K. (1998): *Numerical Methods in Economics,* MIT Press, Cambridge MA.

[43] Kasahara, H. and K. Shimotsu (2008a): "Pseudo-Likelihood Estimation and Bootstrap Inference for Structural Discrete Markov Decision Models," *Journal of Econometrics*, **146**, 92-106.

[44] Kasahara, H. and K. Shimotsu (2008b): "Nonparametric Identification of Finite Mixture Models of Dynamic Discrete Choices," *Econometrica*, **77**, 135-175.

[45] Komunjer, I. and A. Santos (2009): "Semiparametric Estimation of Nonseparable Models: A Minimum Distance from Independence Approach," *Working Paper*.

[46] Koul, H.L. (2002): *Weighted Empiricals in Dynamic Nonlinear Models*, Lecture Notes in Statistics, Springer, New York.

[47] Lee, A.J. (1990): *U-Statistics: Theory and Practice*, Marcel Dekker, New York.

[48] Linton, O.B., and E. Mammen (2005): "Estimating Semiparametric Arch(∞) Models by Kernel Smoothing Methods," *Econometrica*, **73**, 771-836.

[49] Magnac, T. and D. Thesmar (2002): "Identifying Dynamic Discrete Decision Processes," *Econometrica*, **70**, 801-816.

[50] Masry, E. (1996): "Multivariate Local Polynomial Regression for Time Series: Uniform Strong Consistency and Rates," *Journal of Time Series Analysis* **17**, 571–599.

[51] 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.

[52] 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.*

[53] Pesendorfer, M., and P. Schmidt-Dengler (2008): "Asymptotic Least Squares Estimator for Dynamic Games," *Reviews of Economics Studies*, **75**, 901-928.

[54] Pollard, D. (1985): "New Ways to Prove Central Limit Theorems," *Econometric Theory*, **1**, 295-314.

[55] Pollard, D. (1990): *Empirical Processes: Theory and Applications*, NSF-CBMS Regional Conference Series in Probability and Statistics.

[56] Powell, J.L., J.H. Stock and T.M. Stoker (1989): "Semiparametric Estimation of Index Coefficients," *Econometrica,* **57**, 1403-1430.

[57] Robinson, P.M. (1988): "Root-N-consistent semiparametric regression," *Econometrica,* **56**, 931-954.

[58] Rust, J. (1987): "Optimal Replacement of GMC bus engines: An Empirical Model of Harold Zurcher," *Econometrica*, 55, 999-1033.

[59] Rust, J. (1994): "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.*

[60] Rust, J. (1996): "Numerical Dynamic Programming in Economics," *Handbook of Computational Economics*, vol. 4, eds. H. Amman, D. Kendrick and J. Rust. North-Holland.

[61] Srisuma, S., and O. B. Linton (2009): "Semiparametric Estimation of Markov Decision Processes with Continuous State Space," *Working Paper.*

[62] Taussky, O. (1949): "A Recurring Theorem on Determinants," *American Mathematical Monthly*, **56**, 672–676.

[63] van der Vaart, A. W. and J. A. Wellner (1996): *Weak Convergence and Empirical Processes,* Springer-Verlag, New York.

[64] Wolfowitz, J. (1953): "Estimation by Minimum Distance Method," *Ann. Inst. Stat. Math.*, **5**, 9–23.

[65] Zeidler, E. (1986): *Nonlinear Functional Analysis and its Applications I: Fixed-Point Theorem,* Springer-Verlag, New York.

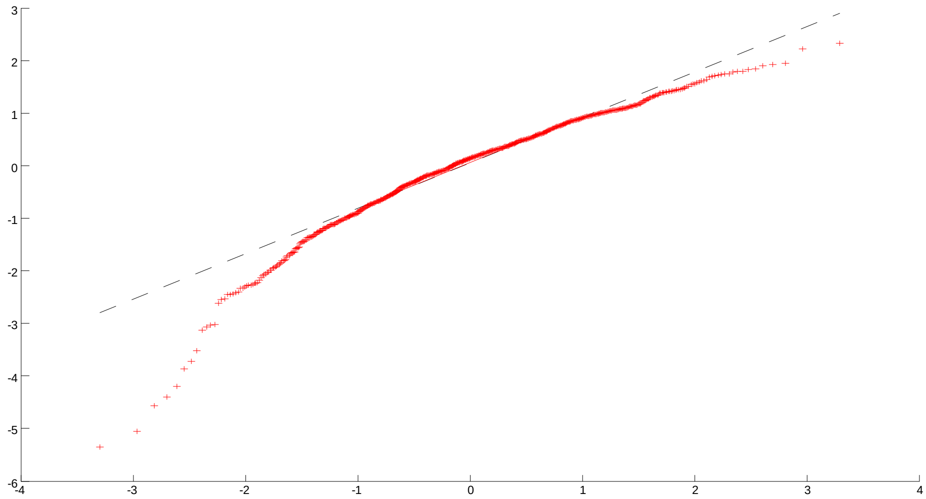

Figure 1: QQ Plot of sample (standardized) $\widehat{\theta}_1$ versus standard normal, $NT = 100$.

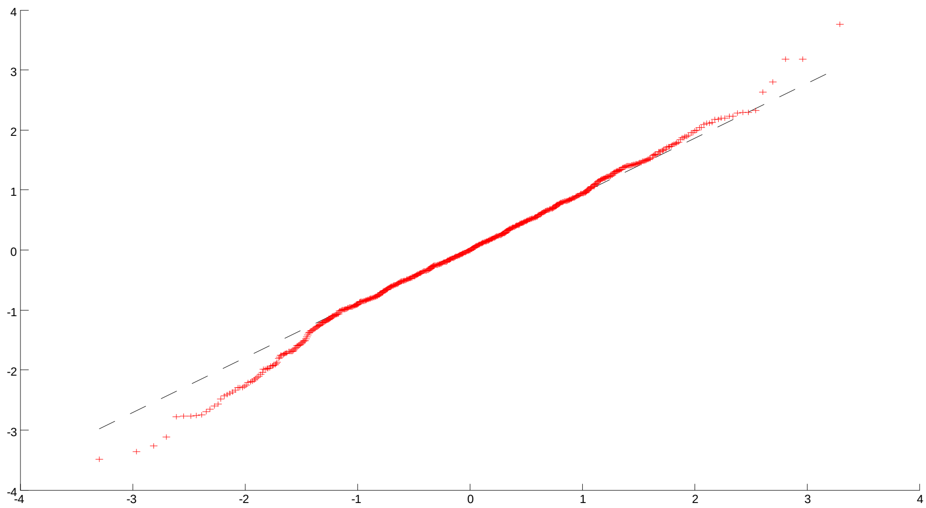

Figure 2: QQ Plot of sample (standardized) $\widehat{\theta}_1$ versus standard normal, $NT = 500$.

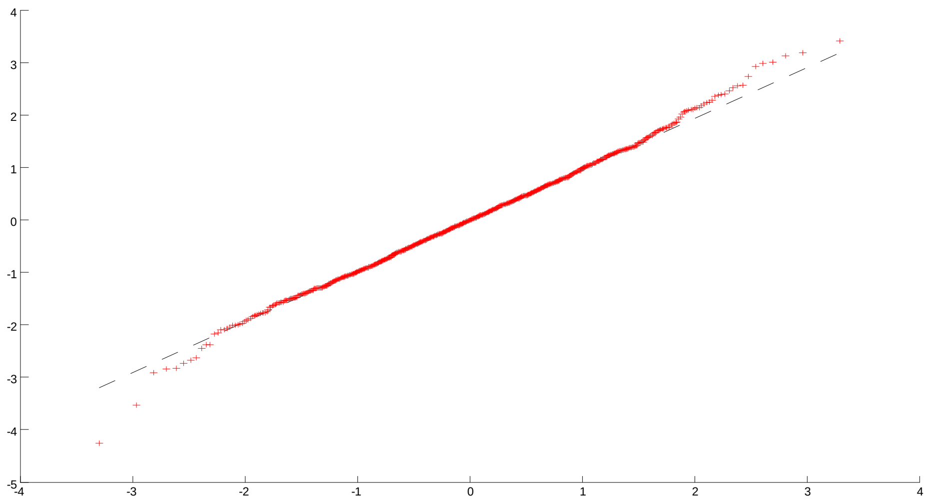

Figure 3: QQ Plot of sample (standardized) $\widehat{\theta}_1$ versus standard normal, $NT = 1000$.

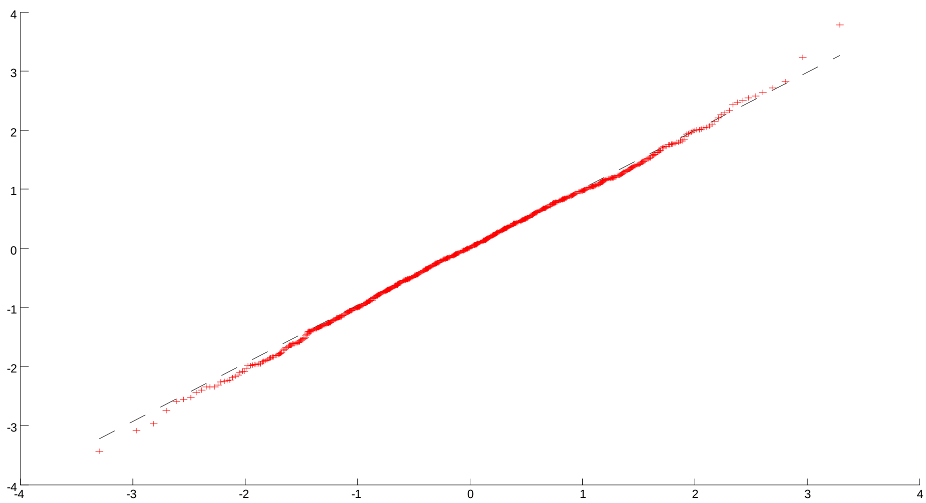

Figure 4: QQ Plot of sample (standardized) $\widehat{\theta}_1$ versus standard normal, $NT = 2500$.

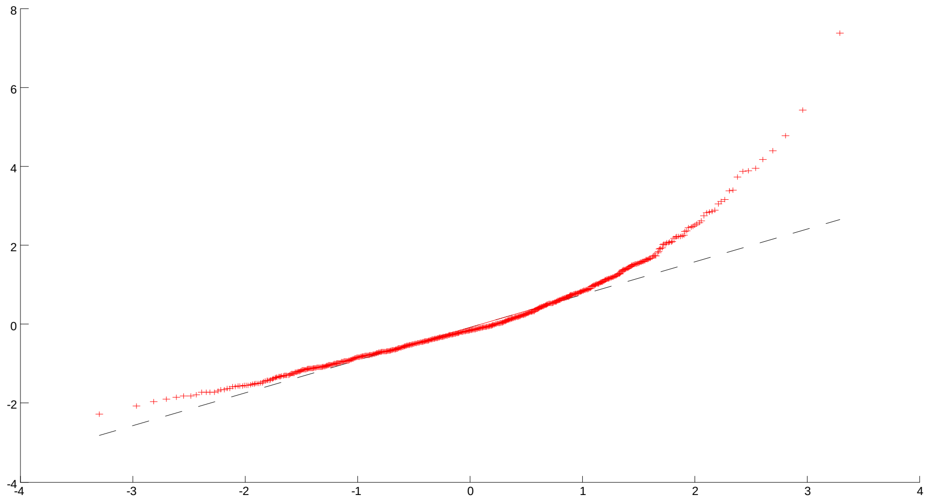

Figure 5: QQ Plot of sample (standardized) $\widehat{\theta}_2$ versus standard normal, $NT = 100$.

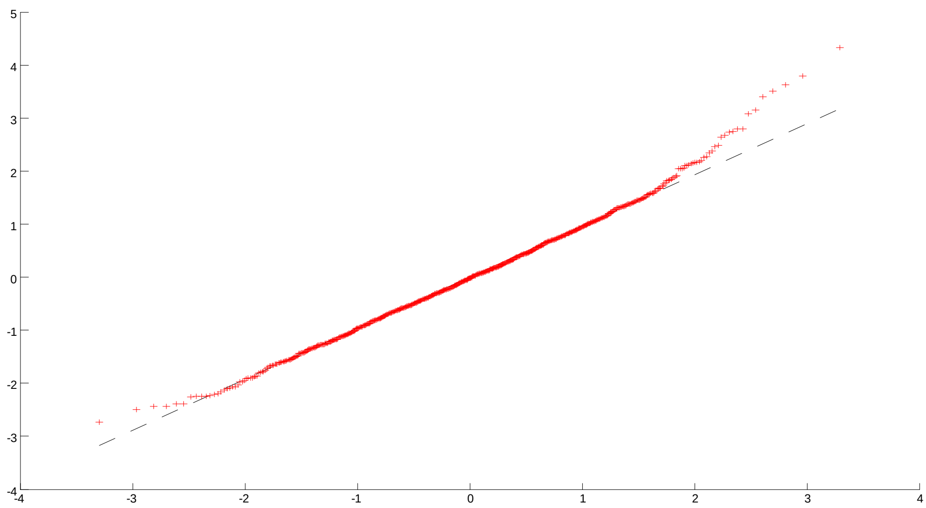

Figure 6: QQ Plot of sample (standardized) $\widehat{\theta}_2$ versus standard normal, $NT = 500$.

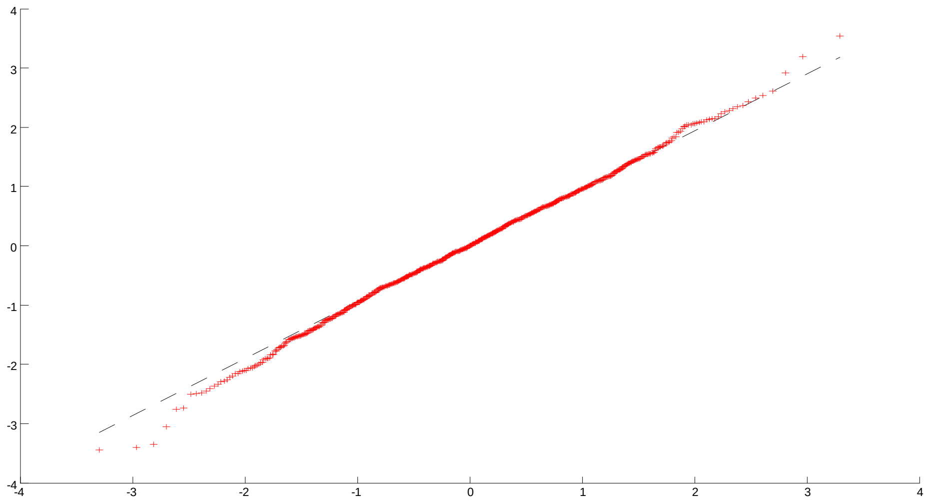

Figure 7: QQ Plot of sample (standardized) $\widehat{\theta}_2$ versus standard normal, $NT = 1000$.

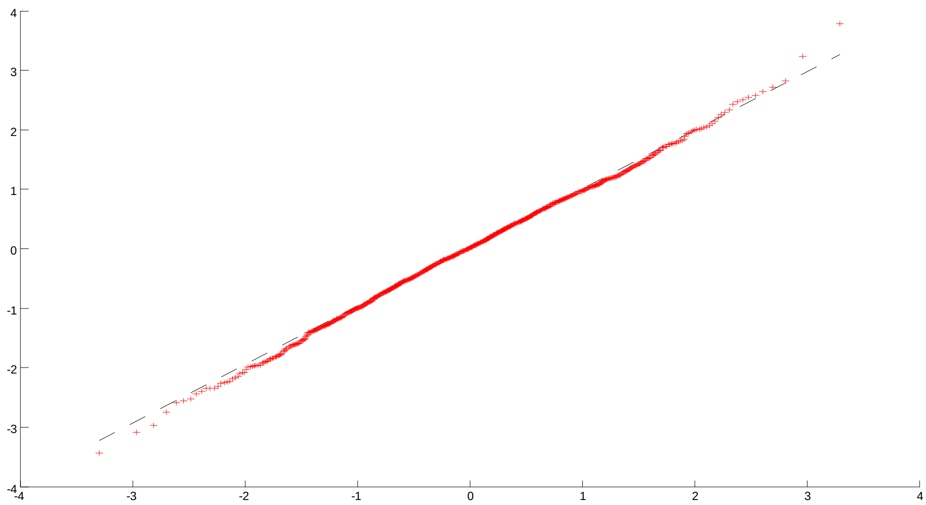

Figure 8: QQ Plot of sample (standardized) $\widehat{\theta}_2$ versus standard normal, $NT = 2500$.

		$\widehat{\theta}_1$				$\widehat{\theta}_1^{trim}$			
NT	ς	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 1: $h_\varsigma = 1.06s(NT)^{-\varsigma}$ is the bandwidth, for various choices of ς, used in the nonparametric 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	ς	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 2: $h_\varsigma = 1.06s(NT)^{-\varsigma}$ is the bandwidth, for various choices of ς, used in the nonparametric 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*.