SEMIPARAMETRIC ESTIMATION OF DYNAMIC DISCRETE CHOICE MODELS

NICHOLAS BUCHHOLZ[*], MATTHEW SHUM[†], AND HAIQING XU[‡]

ABSTRACT. We consider the estimation of dynamic discrete choice models in a semiparametric setting, in which the per-period utility functions are known up to a finite number of parameters, but the distribution of utility shocks is left unspecified. This semiparametric setup differs from most of the existing identification and estimation literature for dynamic discrete choice models. To show identification we derive and exploit a new Bellman-like recursive representation for the unknown quantile function of the utility shocks. Our estimators are straightforward to compute, and resemble classic closed-form estimators from the literature on semiparametric regression and average derivative estimation. Monte Carlo simulations demonstrate that our estimator performs well in small samples. To highlight features of this estimator, we estimate a structural model of dynamic labor supply for New York City taxicab drivers.

Keywords: Semiparametric estimation, Dynamic discrete choice model, Average derivative estimation, Taxicab industry, Labor supply

JEL: C14, D91, C41, L91

[*]Department of Economics, University of Texas at Austin, nibu@utexas.edu.
[†]Division of Humanities and Social Sciences, California Institute of Technology, mshum@caltech.edu.
[‡]Department of Economics, University of Texas at Austin, h.xu@austin.utexas.edu.

1. INTRODUCTION

The dynamic discrete choice (DDC) framework, pioneered by Wolpin (1984), Pakes (1986), Rust (1987, 1994), has gradually become the workhorse model for modelling dynamic decision processes in structural econometrics. Such models, which can be considered an extension of McFadden's 1978; 1980 classic random utility model to a dynamic decision setting, have been used to model a variety of economic phenomenon ranging from labor and health economics to industrial organization, public finance, and political economy. More recently, the DDC framework has also been the starting point for the empirical dynamic games literature in industrial organization.

In this paper, we consider identification and estimation of a class of semiparametric DDC models in which the utility indices are parametrically specified but the shock distribution is left unspecified. Since the utility shocks are typically interpreted as idiosyncratic and unpredictable shocks to preferences, which cause agents' choices to vary over time even under largely unchanging economic environments, it is reasonable to leave their distribution unspecified. We study conditions under which the model structure (consisting of the finite-dimensional parameters in the utility indices, and the infinite-dimensional nonparametric shock distribution) is identified. Our identification argument is constructive, and we propose an estimator based upon it.

The semiparametric DDC framework that we focus on in this paper is novel relative to most of the existing literature on the identification and estimation of DDC models, which considers the case where the utility shocks are fully (parametrically) specified. This reflects an important result in Magnac and Thesmar (2002), who argue that in these models, the single-period utility indices for the choices are (nonparametrically) identifiable only when the distribution of the utility shocks is completely specified. Based on Magnac–Thesmar's "impossibility" result, many recent estimators for and applications of DDC models have considered a structure in which the single-period utility indices are left unspecified, but the utility shock distribution is fully specified (and usually logistic, leading to the convenient multinomial logit choice probabilities).

For identification, we derive a new recursive representation for the unknown quantile function of the utility shocks. Accordingly, we obtain a single-index representation for the conditional choice probabilities in the model, which permits us to estimate the model using classic estimators

from the existing semiparametric binary choice model literature. Specifically, we use Powell, Stock and Stoker's (1989,PSS) kernel-based average-derivative estimator; we show that, under additional mild conditions, our estimator has the same asymptotic properties as PSS's original estimator (which was applied to static discrete-choice models). Moreover, this estimator is computationally quite simple because it can be expressed in closed-form. Monte Carlo simulations demonstrate that our estimator performs well in small samples. To highlight features of this estimator, we estimate a structural model of dynamic labor supply for New York City taxicab drivers. We use our estimated model to develop insight into the labor-leisure tradeoff in light of the dynamic decision environment faced by drivers.

1.1. **Literature.** This paper builds upon several strands in the existing literature. The semiparametric binary choice literature (e.g. Manski (1975, 1985), Powell, Stock, and Stoker (1989), Ichimura and Lee (1991), Horowitz (1992), Klein and Spady (1993), and Lewbel (1998), among many others) is an important antecedent. There is an important substantive difference, however: because these papers focus on a static model, the shock distribution is treated as a nuisance element. As such, estimation of these shocks is not considered. In contrast, the shock distribution in a dynamic model must be estimated since it affects the beliefs that decision makers have regarding their future payoffs. Hence, the need to estimate both the utility parameters as well as the shock distribution represents an important point of divergence between our paper and the previous semiparametric discrete choice literature; nevertheless, as we will point out, the estimators we propose take a very similar form to the estimators in these papers.

   To our knowledge, only a handful of papers consider identification of dynamic models in which the error distribution is left unspecified. Norets and Tang (2014) focus on the discrete state case, and derive (joint) bounds on the error distribution and per-period utilities which are consistent with an observed vector of choice probabilities. We consider the case with continuous state variables, and discuss nonparametric identification and estimation. With a discrete state space, there can never be point identification when the error distribution has continuous support. When the state space is continuous, however, point identification is possible under some support conditions and a location–scale normalization on the error distribution, as we show.

   Chen (2014) considers the identification of dynamic models, and, as we do here, obtains estimators for the model parameters which resemble familiar estimators in the semiparametric

3

discrete choice literature. His approach exploits exclusion restrictions (that is, that a subset of the state variables affect only current utility, but not agents' beliefs about future utilities). In contrast, we do not use exclusion restrictions, but rather exploit the optimality conditions to derive a new recursive representation of the quantile function for the unobserved shocks which allows us to identify and estimate both the model parameters as well as the shock distribution. Blevins (2014) considers a very general class of dynamic models in which agents can make both discrete and continuous choices, and the shock distribution can depend on some of the state variables. Under exclusion restrictions, he shows the nonparametric identification of both the per-period utility functions as well as the error distribution.

Another important related paper is Srisuma and Linton (2012), who pioneered the use of tools for solving type 2 integral equations for estimating dynamic discrete-choice models. We show that, besides the Bellman equation, other structural relations in the dynamic model also take the form of type 2 integral equations. In particular, when viewed as a function of the choice probability, the (unknown) quantile function for the utility shocks can also be recursively characterized via a Bellman-type equation, and hence methods for solving for the value function in the "usual" Bellman equation (either value function iteration or "forward simulation") can also be applied in order to solve for this quantile function.

2. SINGLE AGENT DYNAMIC DISCRETE CHOICES MODEL

Following Rust (1987), we consider a single–agent infinite-horizon binary decision problem. At each time period t, the agent observes state variables $X_t \in \mathbb{R}^k$, and chooses a binary decision $Y_t \in \{0, 1\}$ to maximize her expected utility. The per–period utility is given by

$$u_t(Y_t, X_t, \epsilon_t) = \begin{cases} W_1(X_t)^\intercal \theta_1 + \epsilon_{1t}, & \text{if } Y_t = 1; \\ W_0(X_t)^\intercal \theta_0 + \epsilon_{0t}, & \text{if } Y_t = 0. \end{cases}$$

In the above $W_0(X_t) \in \mathbb{R}^{k_0}$ (resp. $W_1(X_t) \in \mathbb{R}^{k_1}$) denote known transformations of the state variables X_t which affect the per–period utility from choosing $Y_t = 0$ (resp. $Y_t = 1$), and $\epsilon_t \equiv (\epsilon_{0t}, \epsilon_{1t})^\intercal \in \mathbb{R}^2$ are the agent's action-specific payoff shocks; For $d \in \{0, 1\}$, $\theta_d \in \mathbb{R}^{k_d}$ are structural parameters of our interest. This specification of the per-period utility functions as single-indices of the transformed state variables $W_0(X)$ and $W_1(X)$ encompasses a majority

of the existing applications of dynamic discrete-choice models, and thus imposes little loss in generality. The utility of action 0 is not normalized to be zero for reasons discussed in Norets and Tang (2014).

Moreover, let $\beta \in (0,1)$ be the discount factor and $f_{X_{t+1},\epsilon_{t+1}|X_t,\epsilon_t,Y_t}$ be the Markov transition probability density function that depends on the state variable as well as the decision. For notational simplicity, we will use the shorthand W_d for $W_d(X)$ ($d = 0, 1$) and suppress the explicit dependence upon the state variables X.

The agent maximizes the expected discounted sum of the per-period payoffs:

$$\max_{\{y_t,y_{t+1},\ldots\}} \mathbb{E}\left\{\sum_{s=t}^{\infty} \beta^{s-t} u_s(y_s, X_s, \epsilon_s)|X_t,\epsilon_t\right\}, \quad \text{s.t.} \quad f_{X_{s+1},\epsilon_{s+1}|X_s,\epsilon_s,Y_s}.$$

We assume stationarity of the problem, which implies that the problem is invariant to the period t. Because of this, we can omit the t subscripts and use primes ($'$) to denote next period values. Let $V(X,\epsilon)$ be the value function given X and ϵ. By Bellman's equation, the value function can be written as

$$V(X,\epsilon) = \max_{y\in\{0,1\}} \left\{\mathbb{E}[u(y,X,\epsilon)|X,\epsilon] + \beta\mathbb{E}[V(X',\epsilon')|X,\epsilon,Y=y]\right\},$$

and then the agent's optimal decision is given by

$$Y = \operatorname{argmax}_{y\in\{0,1\}} \left\{\mathbb{E}[u(y,X,\epsilon)|X,\epsilon] + \beta\mathbb{E}[V(X',\epsilon')|X,\epsilon,Y=y]\right\}.$$

Unlike much of the existing literature, we do not assume the distribution of the utility shocks $(\epsilon_{0t}, \epsilon_{1t})$ to be known, but treat their distribution as a nuisance element for the estimation of θ. In a static setting, such flexibility may not be necessary, as a flexible specification of $u(X, Y)$ may be able to accommodate any observed pattern in the choice probabilities even when the distribution of utility shocks is parametric.[1] However, in a dynamic setting, the distribution of utility shocks also plays the role of agents' beliefs about the future evolution of state variables (i.e. they are a component in the transition probabilities $f_{X',\epsilon'|X,\epsilon,Y}$) and hence parametric assumptions on this distribution are not innocuous.

---

[1]McFadden and Train (2000) show such properties for the mixed logit specification of static multinomial choice models.

5

2.1. **Characterization of the value function.** In this subsection, we characterize the value function $V(X, \epsilon)$ and the expected value function given X, i.e., $V^e(X) \equiv \mathbb{E}[V(X, \epsilon)|X]$. Both value functions are useful to characterize the equilibrium in our dynamic model. Let F_A and $F_{A|B}$ denote the CDF and the conditional CDF for generic random variables A and B, respectively.

Assumption A (Conditional Independence Assumption). *The transition density satisfies the following condition: $F_{X',\epsilon'|X,\epsilon,Y} = F_{\epsilon'} \times F_{X'|X,Y}$. Moreover, $F_{\epsilon'} = F_\epsilon$.*

Assumption A is strong but fundamental in the current literature; our approach requires no additional exclusion restrictions beyond those typically assumed in the parametric dynamic discrete-choice literature (see e.g. Rust (1987) and Srisuma and Linton (2012)).[2]

Under assumption A, the value function can be written as

$$V(X, \epsilon) = \max \left\{ W_1^\mathsf{T}\theta_1 + \epsilon_1 + \beta\mathbb{E}[V(X', \epsilon')|X, Y = 1], \ W_0^\mathsf{T}\theta_0 + \epsilon_0 + \beta\mathbb{E}[V(X', \epsilon')|X, Y = 0] \right\}.$$

Let $\eta = \epsilon_0 - \epsilon_1$. Then the equilibrium decision maximizing the value function can be written as

$$Y = \mathbb{1}\{\eta \leq \eta^*(X)\},$$

where the cutoff $\eta^*(X)$ is defined as

$$\eta^*(X) \equiv W_1^\mathsf{T}\theta_1 - W_0^\mathsf{T}\theta_0 + \beta\left\{ \mathbb{E}[V(X', \epsilon')|X, Y = 1] - \mathbb{E}[V(X', \epsilon')|X, Y = 0] \right\}. \tag{1}$$

Moreover, let $u^e(X)$ be the expected per–period utility conditional on X, i.e.,

$$u^e(X) \equiv \mathbb{E}(\epsilon_0) + W_1^\mathsf{T}\theta_1 \cdot F_\eta\big(\eta^*(X)\big) + W_0^\mathsf{T}\theta_0 \cdot [1 - F_\eta\big(\eta^*(X)\big)] - \mathbb{E}\big\{\eta \cdot \mathbb{1}[\eta \leq \eta^*(X)]\big\}, \tag{2}$$

where F_η is the CDF of η. Thus, the Bellman equation can be rewritten as

$$V^e(X) = u^e(X) + \beta \cdot \mathbb{E}[V^e(X')|X]. \tag{3}$$

It is worth pointing out that eq. (3) is essentially a Fredholm Integral Equation of the second kind (FIE–2); See e.g. Zemyan (2012). Essentially, FIE–2 is a linear equation system in functional space,

---

[2]On the other hand, as in Blevins (2014), it is possible to relax the independence assumption to one where the state variables can be divided into two groups $X = (X_A, X_B)$ such that $\epsilon \perp X_B|X_A$ (ϵ is independent of X_B given X_A). It appears the identification and estimation procedure described in this paper follow through, with the additional conditioning on X_A at every step.

which is well–known to have a unique analytic solution under some sufficient and necessary conditions.

Assumption B. *For all $s \geq 1$, we have $\mathbb{E}\big(\|W_d^{[s]}\|\big|X\big) < \infty$, where $(^{[s]})$ denotes the next s period values.*

Assumption B holds when $W_d(\cdot)$ are bounded functions.

Given these assumptions, the next lemma applies the Fredholm theorem to obtain a solution of the expected value function to the Bellman equation. (Similar results are utilized in Srisuma and Linton (2012).)

Lemma 1. *Suppose assumptions A and B hold. Then, we have*

$$V^e(x) = u^e(x) + \beta \int_{\mathscr{S}_X} R^*(x', x; \beta) \cdot u^e(x')dx', \ \ \forall x \in \mathscr{S}_X, \tag{4}$$

where $R^*(x', x; \beta) = \sum_{s=1}^{\infty} \beta^{s-1} f_{X^{[s]}|X}(x'|x)$ is the resolvent kernel generated by the FIE eq. (3).

More succinctly, eq. (4) can be rewritten as

$$V^e(X) = u^e(X) + \sum_{s=1}^{\infty} \beta^s \cdot \mathbb{E}[u^e(X^{[s]})|X]. \tag{5}$$

In operator notation, eq. (5) denotes exactly the "forward integration" representation of the value function, which is familiar from many two-step procedures for estimating dynamic discrete choice models (see e.g. Hotz and Miller, 1993; Bajari, Benkard, and Levin, 2007; Hong and Shum, 2010). In the special case when the state variables X are finite and discrete-valued (taking $k < \infty$ values), the Bellman equation is a system of linear equations which can be solved for the value function (cf. Aguirregabiria and Mira, 2007; Pesendorfer and Schmidt-Dengler, 2008) and in that case, the resolvent kernel is just the inverse matrix $(I - \beta F_{X'|X})^{-1}$ where $F_{X'|X}$ denotes the $k \times k$ transition matrix for X.

2.2. **Equilibrium Condition.** To characterize the equilibrium, the key of our approach is to solve for the cutoff value η^* that depends on the state variables X (through the transformations $W_1(X)$ and $W_0(X)$). By using eq. (5), along with Lemma 1, eq. (1) becomes

$$\eta^*(X) = W_1^\mathsf{T}\theta_1 - W_0^\mathsf{T}\theta_0 + \sum_{s=1}^{\infty} \beta^s \left\{ \mathbb{E}[u^e(X^{[s]})|X, Y = 1] - \mathbb{E}[u^e(X^{[s]})|X, Y = 0] \right\}. \tag{6}$$

Moreover, let $\phi_d(X) \equiv (-1)^{d+1} W_d + \sum_{s=1}^{\infty} \beta^s \Big\{ \mathbb{E}\big[W_d^{[s]} \mathbb{1}_{Y^{[s]}=d} \big| X, Y = 1\big] - \mathbb{E}\big[W_d^{[s]} \mathbb{1}_{Y^{[s]}=d} \big| X, Y = 0\big] \Big\}$. Then, it follows from (2),

$$
\eta^*(X) = \phi^{\mathsf{T}}(X) \cdot \theta
$$

$$
- \sum_{s=1}^{\infty} \beta^s \Big\{ \mathbb{E}\big[\eta^{[s]} \mathbb{1}(\eta^{[s]} \leq \eta^*(X^{[s]})) \big| X, Y = 1\big] - \mathbb{E}\big[\eta^{[s]} \mathbb{1}(\eta^{[s]} \leq \eta^*(X^{[s]})) \big| X, Y = 0\big] \Big\}, \quad (7)
$$

where $\phi(X) = (\phi_0^{\mathsf{T}}(X), \phi_1^{\mathsf{T}}(X))^{\mathsf{T}}$ and $\theta = (\theta_0^{\mathsf{T}}, \theta_1^{\mathsf{T}})^{\mathsf{T}}$.

Eq. (7) characterizes the equilibrium decision rule in the single–agent infinite-horizon binary decision problem. Alternatively, we can rewrite it using a resolvent kernel:

$$
\eta^*(x) = \phi^{\mathsf{T}}(x) \cdot \theta - \beta \int_{\mathscr{S}_X} \mathbb{E}[\eta' \cdot \mathbb{1}(\eta' \leq \eta^*(x')] \cdot g(x', x; \beta) dx', \quad \forall x \in \mathscr{S}_X,
$$

where $g(x', x; \beta) = \sum_{s=1}^{\infty} \beta^{s-1} [f_{X^{[s]}|X,Y}(x'|x, 1) - f_{X^{[s]}|X,Y}(x'|x, 0)]$. Given the structural parameters θ_0, θ_1, F_η and $f_{X'|X,Y}$, in principal one can solve the threshold $\eta^*(\cdot)$.[3]

3. IDENTIFICATION

Next, we develop an identification strategy that does not involve solving the Markov equilibrium in the dynamic decision problem. To clarify ideas, we first provide identification of structural parameter $\theta \in \Theta \subseteq \mathbb{R}^{k_\theta}$ (where $k_\theta \equiv k_0 + k_1$) in a fully parametric model, i.e., assuming F_η is known. Then, we establish semiparametric identification of our model by a two–step approach: we first identify F_η up to the finite dimensional parameter θ. In the second step, we represent the agent's choice by a single–index representation. Therefore, the identification of θ follows the literature.

A key feature in our semiparametric identification is that we require (at least) one argument in the state variables X_t to have continuous variation, which is also the case in the semiparametric identification of the single–index binary response model in the static setting. See e.g. Manski (1975). Moreover, we show that the quantile function of F_η is identified on the support of the agent's choice probabilities under a location–scale normalization. This result also corresponds to the findings in static binary response models.

---

[3]However, if one were to use this equation to solve for $\eta^*(\cdot)$ via simulation or computation, note that $g(x', x; \beta)$ also contains $\eta^*(\cdot)$ implicitly through the transition density $f_{X^{[s]}|X,Y}(\cdot|\cdot, \cdot)$.

3.1. **Intermediate step: Parametric Identification.** As a building block towards the more general semiparametric results below, we first consider parametric identification of the model, assuming that the researcher know F_η, the distribution function for the utility shocks. (This is only an intermediate step; in the sections to follow we show how identification obtains even when F_η is not known.) Parametric identification in this setting (with continuous state variables) has also been established previously in Srisuma and Linton (2012). Our identification of θ is constructive, which has an single index structure and suggests an OLS-like (i.e. closed form) estimator. To begin with, we introduce the following assumption.

Assumption C. *Let η be continuously distributed with the full support $\mathbb{R}$.*

Assumption C is a weak condition widely used in semiparametric binary response models (see e.g. Horowitz, 2009). Under assumption C, F_η is strictly increasing on its support $\mathbb{R}$. Let Q be the quantile function of F_η, i.e., $Q = F_\eta^{-1}$.

Let $p(x) = \mathbb{P}(Y = 1 | X = x)$, which obtains directly from the data. Under assumption C, $0 < p(x) < 1$ for all $x \in \mathscr{S}_X$ and $\eta^*(x) = Q(p(x))$. Moreover, using the substitution $\tau \to Q(\tau)$, we have

$$\mathbb{E}[\eta \cdot \mathbb{1}(\eta \le Q(p)] = \int \tau \cdot \mathbb{1}(\tau \le Q(p)) dF_\eta(\tau) = \int_0^p Q(\tau) d\tau.$$

From the above discussion, it is straightforward that we obtain the following lemma.

Lemma 2. *Suppose assumptions A to C hold. Then we have*

$$Q(p(X)) + \sum_{s=1}^\infty \beta^s \left\{ \mathbb{E}\left[\int_0^{p(X^{[s]})} Q(\tau) d\tau \Big| X, Y = 1 \right] - \mathbb{E}\left[\int_0^{p(X^{[s]})} Q(\tau) d\tau \Big| X, Y = 0 \right] \right\}$$

$$= \phi^\mathsf{T}(X) \cdot \theta. \quad (8)$$

Eq. (8) is the key restriction for our identification and estimation analysis, where the number of restrictions equals to the size of the support $\mathscr{S}_X$.

When Q_η is given, then everything in (8) is known except for θ. If, in addition, the matrix $\mathbb{E}[\phi(X)\phi^\mathsf{T}(X)]$ is invertible, then θ can be estimated using nonlinear least-squares on eq. (8). This approach is related to Pesendorfer and Schmidt-Dengler (2008). The full rank of $\mathbb{E}[\phi(X)\phi^\mathsf{T}(X)]$ requires that if the transformed state variables $W_0(X)$ and $W_1(X)$ contain a common component $W_c(X)$, then $\mathbb{E}[W_c(X')|X, Y = 0] \ne \mathbb{E}[W_c(X')|X, Y = 1]$. Such a necessary condition rules out

the case that variables without any dynamic transition (e.g. the constant) are included in both transformed state variables.[4]

3.2. **Semiparametric Identification.** Without making any distributional assumptions on η, we now discuss the identification of θ as well as Q_η. In what follows, we will assume $W_0(X)$ or $W_1(X)$ has at least one continuous argument such that $p(X)$ is continuously distributed. It is well–known that the continuity of covariates is crucial for the semiparametric identification in the static binary response model. This is still the case in our dynamic binary decision model.

Intuitively, the number of restrictions imposed by (8) depends on the richness of the support $\mathscr{S}_X$. For identification of Q_η (up to θ), we only exploit variation in the choice probabilities $p(X)$. For notational simplicity, we assume the choice probability $p(X)$ is continuously distributed on a closed interval.[5]

Assumption D. *(i) Let $p(X)$ be continuously distributed; (ii) let the support of $p(X)$ be a closed interval, i.e., $[\underline{p}, \overline{p}] \subseteq [0, 1]$.*

In contrast, when $p(X)$ only has discrete variation (which typically arises when the state variables X themselves have only discrete variation), Norets and Tang (2014) show that the distribution of η is partially identified.

For each $p \in [\underline{p}, \overline{p}]$, let $z(p) = \mathbb{E}[\phi(X)|p(X) = p]$. We now take the conditional expectation given $p(X) = p$ on both sides of eq. (8). By the law of iterated expectation, we have

$$Q(p) + \sum_{s=1}^{\infty} \beta^s \left\{ \mathbb{E}\left[\int_0^{p(X^{[s]})} Q(\tau)d\tau \Big| p(X) = p, Y = 1 \right] - \mathbb{E}\left[\int_0^{p(X^{[s]})} Q(\tau)d\tau \Big| p(X) = p, Y = 0 \right] \right\}$$
$$= z(p)^{\mathsf{T}} \cdot \theta.$$

The above discussion is summarized by the following lemma.

---

[4]In this case, it is also required that the discount rate $\beta \neq 0$, otherwise $\phi_0(X) = W_0(X)$ and $\phi_1(X) = W_1(X)$, which clearly invalidates the rank condition due to the common term $W_c(X)$.

[5]This interval–support restriction can be relaxed at expositional expense. For instance, suppose $\mathscr{S}_{p(X)}$ is a non–degenerate compact subset of $[0, 1]$. All of our identification arguments below still hold by replacing the integral region $[\underline{p}, \overline{p}]$ with $\mathscr{S}_{p(X)}$.

Lemma 3. *Suppose assumptions A to D hold. Then we have*

$$Q(p) + \beta \int_{\underline{p}}^{\overline{p}} \int_{\underline{p}}^{p'} Q(\tau)d\tau \cdot \pi(p', p; \beta)dp' = z(p)^\mathsf{T} \cdot \theta, \quad \forall p \in [\underline{p}, \overline{p}], \tag{9}$$

where $\pi(p', p; \beta) = \sum_{s=1}^{\infty} \beta^{s-1}[f_{p(X^{[s]})|p(X),Y}(p'|p, 1) - f_{p(X^{[s]})|p(X),Y}(p'|p, 0)].$

By definition, $\pi(p', p; \beta)$ is the difference of the discounted aggregate densities of the future choice probabilities, conditional on the current choice probability and (exogenously given) action, which can be obtained directly from the data.

Note that eq. (9) is also an FIE–2. To see this, let $\Pi(p', p; \beta) \equiv \sum_{s=1}^{\infty} \beta^{s-1}[F_{p(X^{[s]})|p(X),Y}(p'|p, 1) - F_{p(X^{[s]})|p(X),Y}(p'|p, 0)]$. Then, the second term of eq. (9) can be rewritten as

$$\int_{\underline{p}}^{\overline{p}} \int_{\underline{p}}^{p'} Q(\tau)d\tau \cdot \pi(p', p; \beta)dp' = \int_0^1 Q(\tau) \cdot \int_{\underline{p}}^{\overline{p}} \mathbb{1}(\tau \le p') \cdot \pi(p', p; \beta)dp'd\tau$$

$$= -\int_0^1 Q(\tau) \cdot \left[\int_{\underline{p}}^{\overline{p}} \mathbb{1}(p' < \tau)\pi(p', p; \beta)dp' \right]d\tau$$

$$= -\int_{\underline{p}}^{\overline{p}} Q(\tau) \cdot \Pi(\tau, p; \beta)d\tau,$$

where the second step comes from the fact $\int_{\underline{p}}^{\overline{p}} \pi(p', p; \beta)dp' = 0$ and the last step is because $\Pi(p', p, \beta) = 0$ for all $p' \notin [\underline{p}, \overline{p}]$. Hence, we obtain the following FIE–2:

$$Q(p) - \beta \int_{\underline{p}}^{\overline{p}} Q(\tau) \cdot \Pi(\tau, p; \beta)d\tau = z(p)^\mathsf{T} \cdot \theta, \quad \forall p \in [\underline{p}, \overline{p}].$$

By solving this equation, we can identify $Q(\cdot)$ on $[\underline{p}, \overline{p}]$ up to the finite dimensional parameter θ.

Assumption E. *Let* $\beta^2 \cdot \int_{\underline{p}}^{\overline{p}} \int_{\underline{p}}^{\overline{p}} \Pi^2(p', p; \beta)dp'dp < 1.$

Assumption E is introduced for the uniqueness of the solution. This assumption is a high level condition but testable.

Lemma 4. *Suppose assumptions A to E hold. Then, Q_η is point identified on $[\underline{p}, \overline{p}]$ up to the finite dimensional parameter θ:*

$$Q(p) = \left\{ z(p) - \beta \int_{\underline{p}}^{\overline{p}} R(p', p; \beta) \cdot z(p')dp' \right\}^\mathsf{T} \cdot \theta, \quad \forall \, p \in [\underline{p}, \overline{p}] \tag{10}$$

11

where $R(p', p; \beta) = \sum_{s=1}^{\infty} (-\beta)^{s-1} K_s(p', p; \beta)$, *in which* $K_s(p', p; \beta) = \int_0^1 K_{s-1}(p', \tilde{p}; \beta) \cdot \Pi(\tilde{p}, p; \beta) d\tilde{p}$ *and* $K_1(p', p; \beta) = \Pi(p', p; \beta)$.

The solution (10) is proportional to θ, which is due to the linearity of the FIE system. Therefore, (10) can also be represented by a sequence of "basis" solutions. To see this, let $z_\ell(p)$ be the ℓ–th argument of $z(p)$. For $\ell = 1, \cdots, k_\theta$, let $b_\ell^*(\cdot)$ be the (unique) solution to the following equation

$$b_\ell(p) + \beta \int_{\underline{p}}^{\overline{p}} \int_{\underline{p}}^{p'} b_\ell(\tau) d\tau \cdot \pi(p', p; \beta) dp' = z_\ell(p). \tag{11}$$

By a similar argument to Lemma 4, we have

$$b_\ell^*(p) = z_\ell(p) - \beta \int_{\underline{p}}^{\overline{p}} R(p', p; \beta) \cdot z_\ell(p') dp', \quad \forall p \in [\underline{p}, \overline{p}]$$

as the unique solution to (11). Let $\mathcal{B}(\cdot) \equiv (b_1^*(\cdot), \cdots, b_{k_\theta}^*(\cdot))^\intercal$ be the sequence of solutions supported on $[\underline{p}, \overline{p}]$. Thus, the solution in eq. (10) can be written as

$$Q(p) = \mathcal{B}(p)^\intercal \cdot \theta, \quad \forall p \in [\underline{p}, \overline{p}]. \tag{12}$$

By Lemmas 1 to 4, we obtain a single–index representation of the semiparametric dynamic decision model, which is stated in the following theorem.

Theorem 1. *Suppose assumptions A to E hold. Then, the agent's dynamic decision can be represented by a static single–index model:*

$$\mathbb{P}(Y = 1|X) = F_\eta\big(m(X)^\intercal \cdot \theta\big)$$

where

$$m(X) = \phi(X) - \sum_{s=1}^{\infty} \beta^s \left\{ \mathbb{E}\left[\int_{\underline{p}}^{p(X^{[s]})} \mathcal{B}(\tau) d\tau \Big| X, Y = 1 \right] - \mathbb{E}\left[\int_{\underline{p}}^{p(X^{[s]})} \mathcal{B}(\tau) d\tau \Big| X, Y = 0 \right] \right\}.$$

Note that $\mathbb{P}(Y = 1|X) = F_\eta\big(Q(p(X))\big)$. Then, Theorem 1 obtains by combining eqs. (8) and (12). Given the identification of $\mathcal{B}(\cdot)$ on the support $[\underline{p}, \overline{p}]$, $m(\cdot)$ is then constructively identified on $\mathscr{S}_X$. Therefore, the identification of θ simply follows the single-index model literature, see e.g. Manski (1975, 1985).

It is worthing noting that any constant term in W_d remains a constant in the transformed linear–index $m(X)$. In other words, suppose, w.l.o.g., $W_{11} = 1$. Then the corresponding

argument in $m(X)$ also equals 1. To see this, note that the first argument in $\phi(X)$ is given by

$$1 + \sum_{s=1}^{\infty} \beta^s \left\{ \mathbb{E}\big[p(X^{[s]})|X, Y = 1\big] - \mathbb{E}\big[p(X^{[s]})|X, Y = 0\big] \right\},$$

which thereafter implies

$$z_1(p) = 1 + \sum_{s=1}^{\infty} \beta^s \left\{ \mathbb{E}\big[p(X^{[s]})|p(X) = p, Y = 1\big] - \mathbb{E}\big[p(X^{[s]})|p(X) = p, Y = 0\big] \right\}.$$

Using (11), one can verify that the solution is: $b_1^*(\cdot) = 1$. Then, we plug this solution into the the first element of $m(X)$, which gives us $m_1(X) = 1$. By a similar argument, a constant in W_0 implies the corresponding term in $m(X)$ equals -1.

By a similar argument as in the static binary response model literature, the index parameter θ is identified up to location and scale in the semiparametric setting. For notational simplicity, hereafter we assume the state vector X does not include a constant term in the semiparametric setting.[6] Moreover, we will introduce a scale normalization on θ which is also standard in the literature.

Assumption F. *We denote the first argument of $m(X)$ by $m_1(X)$ and the rest by $m_{-1}(X)$. Moreover, let $m_1(X)$ be continuously distributed on an interval given $m_{-1}(X)$ which is a vector of either discrete and/or continuous random variables. Let $f_{m_1(X)|m_{-1}(X)}$ be the conditional pdf. Moreover, the matrix $\mathbb{E}[m(X)m(X)^{\intercal}]$ is invertible.*

In Assumption F, the first half condition requires at least one argument in X_1 to be continuously distributed conditional on others; this rules out cases where, e.g. all the state variables are functions of a single variable X_1 (as in (Rust, 1987), where mileage and mileage-squared enter as state variables). The second half of Assumption F is a testable rank condition. Assumption F is a strong assumption, but almost indispensable in the semiparametric single index model literature; See Horowitz (2009).

Assumption G. *Let $\|\theta\| = 1$.*

---

[6]In our semiparametric setting, any constant term in the utility function will be absorbed by the error term since the distribution of the latter is left unspecified.

Assumption G is a scale normalization, which has also been used in PSS. Note that we implicitly normalize our location term by 0, since neither W_0 nor W_1 contains a constant term.

Theorem 2. *Suppose assumptions A to G hold. Then, the structural parameter θ is point identified.*

4. SEMIPARAMETRIC ESTIMATION

In this section, we describe and motivate the semiparametric estimation of our structural model. For expositional simplicity, we assume all variables in X are continuously distributed. A mixture of continuous and discrete regressors can be accommodated at the expense of notation. Let $\{(Y_t, X_t^\intercal)^\intercal : t = 1, \cdots, T\}$ be our sample of the Markov decision process. Our estimation procedure parallels the identification strategy, which takes multiple steps. Throughout, we use K and h to denote a Parzen–Rosenblatt kernel and a bandwidth, respectively.

First, we nonparametrically estimate the choice probabilities $p(\cdot)$ and the generated regressor $\phi(\cdot)$. In particular, let

$$\hat{p}(X_s) = \frac{\sum_{t=1}^{T} Y_t \times K_p\left(\frac{X_t - X_s}{h_p}\right)}{\sum_{t=1}^{T} K_p\left(\frac{X_t - X_s}{h_p}\right)}, \quad \forall s = 1, \cdots, T.$$

As is standard, we choose an optimal bandwidth, i.e., $h_p = 1.06 \times \hat{\sigma}(X) \times T^{-\frac{1}{2\iota + k}}$, where $\hat{\sigma}(X)$ is the sample standard deviation of X_t and ι ($\iota \geq 2$) is the order of the kernel function K_p. For example, if we choose K_p to be the pdf of the standard normal distribution, then $\iota = 2$. In addition, the support $[\underline{p}, \overline{p}]$ of $p(X)$ can be estimated by $[\min_{1 \leq s \leq T} \hat{p}(X_s), \max_{1 \leq s \leq T} \hat{p}(X_s)]$.

Moreover, recall that the transformed state variables $W_d(X)$ ($d = 0, 1$) are known. Then, for $s = 1, \cdots, S_T$, where $S_T = T - \ell_T$ for some integer ℓ_T satisfying $\ell_T \to +\infty$ and $S_T \to +\infty$ as $T \to +\infty$, let $\delta_{dt} = \sum_{s=1}^{\ell_T} \beta^s \cdot W_d(X_{t+s}) Y_{t+s}^d (1 - Y_{t+s})^{1-d}$. For $s = 1, \cdots, T$, let further

$$\hat{\phi}_d(X_s) = (-1)^{d+1} W_d(X_s) + \frac{\sum_{t=1}^{S_T} \delta_{dt} \cdot K_\phi\left(\frac{X_t - X_s}{h_\phi}\right) \mathbb{1}(Y_t = 1)}{\sum_{t=1}^{S_T} K_\phi\left(\frac{X_t - X_s}{h_\phi}\right) \mathbb{1}(Y_t = 1)} - \frac{\sum_{t=1}^{S_T} \delta_{dt} \cdot K_\phi\left(\frac{X_t - X_s}{h_\phi}\right) \mathbb{1}(Y_t = 0)}{\sum_{t=1}^{S_T} K_\phi\left(\frac{X_t - X_s}{h_\phi}\right) \mathbb{1}(Y_t = 0)}.$$

Similarly, we can choose h_ϕ in an optimal way. In above expression, the summation includes only the first S_T observations. This is because δ_{dt} is not well defined for all $t > S_T$. In practice, we choose ℓ_T in a way such that $\delta_{dt} - \sum_{s=1}^{+\infty} \beta^s W_d(X_{t+s}) Y_{t+s}^d (1 - Y_{t+s})^{1-d}$ is negligible relative

to the sampling error, which is feasible because the former converges to zero at an exponential rate.

In the second stage, we estimate $z(\cdot)$ and $\mathcal{B}(\cdot)$ on the support $[\underline{p}, \overline{p}]$. First, let

$$\hat{z}(p) = \frac{\sum_{t=1}^{T} \hat{\phi}(X_t) \cdot K_z\left(\frac{\hat{p}(X_t)-p}{h_z}\right)}{\sum_{t=1}^{T} K_z\left(\frac{\hat{p}(X_t)-p}{h_z}\right)}, \quad \forall\, p \in [\min_{1 \leq s \leq T} \hat{p}(X_s), \max_{1 \leq s \leq T} \hat{p}(X_s)].$$

According to Guerre, Perrigne, and Vuong (2000), h_z is chosen in a suboptimal way, i.e., $h_z = 1.06 \times \hat{\sigma}(p(X)) \times T^{-\frac{1}{2\iota+3}}$.

To estimate $b_\ell^*(\cdot)$ on the support $[\underline{p}, \overline{p}]$, we note that eq. (11) can be rewritten as

$$b_\ell(p) + \sum_{s=1}^{\infty} \beta^s \cdot \mathbb{E}\left[\int_{\underline{p}}^{p(X^{[s]})} b_\ell(\tau)d\tau \big| p(X) = p, Y = 1\right]$$

$$- \sum_{s=1}^{\infty} \beta^s \cdot \mathbb{E}\left[\int_{\underline{p}}^{p(X^{[s]})} b_\ell(\tau)d\tau \big| p(X) = p, Y = 0\right] = z_\ell(p).$$

This suggests an estimator $\hat{b}_\ell^*(\cdot)$ that solves

$$\hat{b}_\ell^*(p) + \frac{\sum_{t=1}^{S_T} \xi_t(\hat{b}_\ell^*) \times K_\xi\left(\frac{\hat{p}(X_t)-p}{h_\xi}\right) \times Y_t}{\sum_{t=1}^{S_T} K_\xi\left(\frac{\hat{p}(X_t)-p}{h_\xi}\right) \times Y_t} - \frac{\sum_{t=1}^{S_T} \xi_t(\hat{b}_\ell^*) \times K_\xi\left(\frac{\hat{p}(X_t)-p}{h_\xi}\right) \times (1-Y_t)}{\sum_{t=1}^{S_T} K_\xi\left(\frac{\hat{p}(X_t)-p}{h_\xi}\right) \times (1-Y_t)} = \hat{z}_\ell(p),$$

where $\xi_t(b_\ell) = \sum_{s=1}^{\ell_T} \beta^s \int_{\underline{p}}^{\hat{p}(X_{t+s})} b_\ell(\tau)d\tau$ for which the integration can be computed by numerical integration. Similarly, $h_z = 1.06 \times \hat{\sigma}(p(X)) \times T^{-\frac{1}{2\iota+3}}$ is chosen sub-optimally. A numerical solution of $\hat{b}_\ell^*$ can obtain using the iteration method: Let $\hat{b}_\ell^{[0]} = \hat{z}_\ell^\mathsf{T}(p)$. Then we set

$$\hat{b}_\ell^{[1]}(p) = \hat{z}_\ell^\mathsf{T}(p) - \left\{ \frac{\sum_{t=1}^{S_T} \xi_t(\hat{b}_\ell^{[0]}) \times K_\xi\left(\frac{\hat{p}(X_t)-p}{h_\xi}\right) \times Y_t}{\sum_{t=1}^{S_T} K_\xi\left(\frac{\hat{p}(X_t)-p}{h_\xi}\right) \times Y_t} - \frac{\sum_{t=1}^{S_T} \xi_t(\hat{b}_\ell^{[0]}) \times K_\xi\left(\frac{\hat{p}(X_t)-p}{h_\xi}\right) \times (1-Y_t)}{\sum_{t=1}^{S_T} K_\xi\left(\frac{\hat{p}(X_t)-p}{h_\xi}\right) \times (1-Y_t)} \right\}.$$

Repeat such an iteration until it converges. Then we obtain $\hat{b}_\ell^*(\cdot) = \hat{b}_\ell^{[\infty]}(\cdot)$ on $[\hat{\underline{p}}, \hat{\overline{p}}]$.

Next, we obtain the single–index variables $m(X_s)$ by: for $\ell = 1, \cdots, k_\theta$,

$$\hat{m}_\ell(X_s) = \hat{\phi}_\ell(X_s) - \left\{ \frac{\sum_{t=1}^{S_T} \xi_t(\hat{b}_\ell^*) \times K_m\left(\frac{X_t-X_s}{h_m}\right) \times Y_t}{\sum_{t=1}^{S_T} K_m\left(\frac{X_t-X_s}{h_m}\right) \times Y_t} - \frac{\sum_{t=1}^{S_T} \xi_t(\hat{b}_\ell^*) \times K_m\left(\frac{X_t-X_s}{h_m}\right) \times (1-Y_t)}{\sum_{t=1}^{S_T} K_m\left(\frac{X_t-X_s}{h_m}\right) \times (1-Y_t)} \right\}.$$

In particular, $h_m = 1.06 \times \hat{\sigma}(X) \times T^{-\frac{1}{2\iota+k}}$ is chosen optimally.

15

Finally, we apply PSS to estimate θ (up to scale).[7] Specifically, we define

$$\hat{\theta} = -\frac{2}{T(T-1)} \sum_{t=1}^{T} \sum_{s \neq t} \left[\frac{1}{h_\theta^{k_\theta+1}} \times \nabla K_\theta \left(\frac{\hat{m}(X_t) - \hat{m}(X_s)}{h_\theta} \right) \times Y_s \right]. \tag{13}$$

Following the standard kernel regression literature, we can show $\hat{\theta}$ is consistent given that $\sup_{x \in \mathscr{S}_X} |\hat{m}(x) - m(x)| = o_p(h_\theta)$, $h_\theta \to 0$ and $Th_\theta^{k_\theta+1} \to \infty$ as $T \to \infty$.

Similar to PSS, it is of particular interest to establish $\sqrt{T}$–consistency of $\hat{\theta}$. The argument follows closely to that in PSS. In particular, we need to choose a high order kernel K_θ and an under–smoothed bandwidth h_θ. However, it is more delicate in our setting because of the generated regressor $\hat{m}(X)$ contained in the kernel function of our estimator (13). Due to the first–stage estimation error, we must make the following additional assumptions on the convergence of $\hat{m}(X)$ to $m(X)$:

Assumption H. $h_\theta = T^{-\frac{1}{\gamma}}$ where $k_\theta + 2 < \gamma < k_\theta + 3 + \mathbb{1}(k_\theta$ is even).

Assumption I. *The support of the kernel function K_θ is a convex subset of $\mathbb{R}^{k_\theta}$ with nonempty interior, with the origin as an interior point. K_θ is a bounded differentiable function that obeys: $\int K_\theta(u)du = 1$, $K_\theta(u) = 0$ for all u belongs to the boundary of its support, $K_\theta(u) = K_\theta(-u)$ and*

$$\int u_1^{\ell_1} \cdots u_{k_\theta}^{\ell_{\rho'}} K_\theta(u)du = 0, \quad \text{for } \ell_1 + \cdots + \ell_{\rho'} < \frac{k_\theta + 3 + \mathbb{1}(k_\theta \text{ is even })}{2}, \text{ and}$$

$$\int u_1^{\ell_1} \cdots u_{k_\theta}^{\ell_{\rho'}} K_\theta(u)du \neq 0, \quad \text{for } \ell_1 + \cdots + \ell_{\rho'} = \frac{k_\theta + 3 + \mathbb{1}(k_\theta \text{ is even })}{2}.$$

where u_ℓ is the ℓ–th argument of u.

Assumption J. *(i) $\mathbb{E}\|\hat{m}(X) - m(X)\|^2 = o(T^{-\frac{1}{2}}h_\theta^3)$;*
(ii) $\mathbb{E}\|\mathbb{E}[\hat{m}(X)|X] - m(X)\| = o(T^{-\frac{1}{2}}h_\theta^2)$;
(iii) $\hat{m}(X_t) - \hat{m}_{t,-s} = o_p(T^{-\frac{1}{2}}h_\theta^2)$, where $\hat{m}_{t,-s}$ is the nonparametric estimator $\hat{m}(X_t)$, except for leaving the s–th observation out of the sample in its construction.

Assumptions H and I are introduced by PSS for the choice of bandwidth and kernel, respectively, to control the bias term in the estimation of θ.[8] The restriction on the bandwidth Assumption H

---

[7]One could also use alternative methods e.g. Klein and Spady (1993) and Ichimura (1993) to estimate θ.

[8]Note that we implicitly assume that Assumptions 1 – 3 in PSS hold, which impose smoothness conditions on $f_{m(X)}$ and $\mathbb{P}(Y_t = 1|m(X_t) = m)$ as well as other regularity conditions.

implies that h_θ is not an optimal bandwidth sequence (rather it is undersmoothed) such that the bias of estimating θ goes to zero faster than $\sqrt{T}$. Moreover, Assumption J encompasses high–level conditions that could be further established under primitive conditions. In particular, Assumption J(i) requires $\hat{m}(\cdot)$ converge to $m(\cdot)$ faster than $T^{-\frac{1}{4}}$. Similar conditions are also made in e.g. Ai and Chen (2003) for the regular convergence of finite–dimensional parameters in semiparametric models. Moreover, by Assumption J(ii), the bias term in the estimation of m uniformly converges to zero faster than $T^{-\frac{1}{2}}$. Hence, we need to use a higher order kernel in the estimation of $m(\cdot)$. Assumption J(iii) is not essential, which could be dropped if we exclude both t-th and s–th observations in the argument $\hat{m}(X_t) - \hat{m}(X_s)$ of the kernel function in (13).

Given these assumptions, we can show the following result (the proof is in the appendix):

Theorem 3. *Suppose assumptions H to J hold. Then, for some scalar $\lambda > 0$ specified below, $\sqrt{T}(\hat{\theta} - \lambda \cdot \theta)$ has a limiting multivariate normal distribution defined in Powell, Stock, and Stoker (1989, Theorem 3.1):*

$$\sqrt{T}(\hat{\theta} - \lambda \cdot \theta) \xrightarrow{d} N(0, \Sigma)$$

where $\Sigma \equiv 4\,\mathbb{E}(\zeta \cdot \zeta^\intercal) - 4\lambda^2 \times \theta \cdot \theta^\intercal$, $\zeta = f_m(m(X)) \cdot f_\eta(\eta^*(X)) \cdot \theta - \big[Y - F_\eta(\eta^*(X))\big] \cdot f'_m(m(X))$ and $\lambda = \mathbb{E}\big[f_m\big(m(X)\big) \times f_\eta(m(X)^\intercal \cdot \theta)\big]$.

In the above theorem, recall $\mathbb{P}(Y = 1|X) = F_\eta(\eta^*(X))$ and $\eta^*(X) = m(X)^\intercal \cdot \theta$ by Theorem 1. Our estimator $\hat{\theta}$ (as defined in Eq. (13)) has not imposed the scale restriction in Assumption G; thus $\lambda \in \mathbb{R}$ in the above theorem denotes the probability limit of $\|\hat{\theta}\|$; i.e., $\|\hat{\theta}\| = \lambda + O_p(T^{-1/2})$. Therefore, by rescaling our estimator $\hat{\theta}$ as $\hat{\theta}^* = \hat{\theta}/\lambda$, we obtain that

$$\sqrt{T}(\hat{\theta}^* - \theta) \xrightarrow{d} N(0, \Sigma/\lambda^2).$$

Given $\hat{\theta}^*$, a nonparametric estimator of $Q(\cdot)$ directly follows from Eq. (10). Namely, let

$$\hat{Q}(p) = \hat{z}^\intercal(p) \times \hat{\theta}^*, \quad \forall\, p \in \big[\min_{1 \leq s \leq T} \hat{p}(X_s), \max_{1 \leq s \leq T} \hat{p}(X_s)\big].$$

Because of the $\sqrt{T}$–consistency of $\hat{\theta}^*$, the estimator $\hat{Q}_p(p)$ is asymptotically equivalent to $\hat{z}^\intercal(p) \times \hat{\theta}^*$, which converges at a nonparametric rate.

4.1. **Monte Carlo.** The focus of our Monte Carlo is on the semiparametric estimation. In our experiments, let $u_t(0, X_t, \epsilon_t) = \theta_0 + \epsilon_{0t}$ and $u_t(1, X_t, \epsilon_t) = X_{1t}\theta_1 + X_{2t}\theta_2 + \epsilon_{1t}$, where X_{1t}, X_{2t}

TABLE 1. Monte Carlo Results

Sample Obs.	Parameter	True Value	Estimate	Std. Dev.	Bias
1000	θ_1	-1	-1.0182	0.3636	0.0182
	θ_2	-2	-1.9457	0.2158	-0.0543
2000	θ_1	-1	-1.0163	0.2913	0.0163
	θ_2	-2	-1.9618	0.1854	-0.0382
4000	θ_1	-1	-0.9985	0.2344	-0.0015
	θ_2	-2	-1.9836	0.1176	-0.0164

This table presents Monte Carlo results for different sample sizes. For each sample size, reported estimates, standard deviations and bias are computed as the mean across 150 simulation draws.

are random variables and $\theta_0, \theta_1, \theta_2 \in \mathbb{R}$. Moreover, we set the conditional distribution of X_{t+1} given X_t and Y_t as follows: for $k = 1, 2$

$$X_{k,t+1} = \begin{cases} X_{kt} + \nu_{kt}, & \text{if } Y_t = 0 \\ \nu_{kt} & \text{if } Y_t = 1 \end{cases},$$

where ν_{kt} conforms to $\ln \mathcal{N}(0, 1)$ and $\nu_{1t} \perp \nu_{2t}$. Moreover, let ϵ_{dt} be i.i.d. across $d = 0, 1$ and t, and conform to an extreme value distribution with the density function $f(e) = \exp(-e) \exp[-\exp(-e)]$. We set $\beta = 0.9$ and the parameter value as follows: $\theta_0 = -5$, $\theta_1 = -1$ and $\theta_2 = -2$.

Because we cannot estimate the constant θ_0 in the semiparametric framework, then we treat θ_0 as a nuisance parameter. Let $\theta = (\theta_1, \theta_2)^\mathsf{T}$. As a matter of fact, θ is only identified up to scale in the semiparametric setting. To compare the performance of the semiparametric estimators, we assume the scale of θ is known, i.e., $\|\theta\| = \sqrt{5}$, rather than imposing a different normalization, as assumption G. We present in Table 1 the bias and standard deviation of the semiparametric estimator.

5. EMPIRICAL APPLICATION: DYNAMIC LABOR SUPPLY FOR NYC TAXI DRIVERS

A growing literature has sprung up which seeks to estimate labor supply elasticities in markets where labor supply is continuously adjustable; several of these papers have studied the market for taxi rides, because taxi drivers choose their own hours. We next consider an application which aims to characterize the dynamic labor supply of taxi drivers. We first pose

and estimate a model of taxi driver's labor supply as a dynamic discrete choice over quitting for the day. Our model highlights the tradeoffs between working longer to earn extra income versus incurring increasing costs of effort. Estimating the structural model allows us to both analyze the labor-leisure tradeoff in a richer way compared to previous studies, and also to showcase some features of our semiparametric estimation procedure.

This work builds on a literature on labor supply in the taxi industry. Camerer, Babcock, Loewenstein, and Thaler (1997) found evidence of strong negative wage elasticities; they argued that negative elasticities reflected the presence of income-targeting on the part of drivers: for example, a labor supply policy of the form "I will work today until I earn $200." Farber (2005, 2008, 2014) consider static models of labor supply. The latter two papers integrate reference-dependent utility, which is the notion that agents' utility is not only a function of income but also reference-points or targets, where the marginal utility of income increases more quickly before the target is met than after it is met. Originally, Farber (2008) finds mixed evidence for the existence of reference-dependence, but Farber (2014) uses more comprehensive data and finds strong evidence that labor supply behavior is driven by the standard neoclassical prediction of upward sloping supply curves, as opposed to income-targeting and its associated negative elasticities. Crawford and Meng (2011) specify and estimate a dynamic model of labor supply incorporating reference-dependence in both income and hours-worked during a shift.

We estimate a dynamic structural model in which drivers solve a dynamic optimization problem to determine their hours worked, as a function of cumulative earned income and cumulative time spent working. Our model is based on the taxi labor supply model of Frechette, Lizzeri, and Salz (2015) [FLS], in which taxi drivers solve a dynamic competitive game by choosing the optimal starting times and length of time to stay on a shift.[9] As with FLS, our taxi drivers will decide how long to work by weighing the utility of earning revenue against the disutility of working longer. FLS utilizes the MPEC method to solve a dynamic entry game in an equilibrium framework, allowing the market to equilibrate via the waiting times experienced by passengers and taxis. While we do not consider these general equilibrium forces, we take advantage of our computationally light, semi-parametric estimation method to test for a variety

---

[9]The more taxis that are working, the less revenue is earned as a result of lower probabilities of finding a passenger.

of taxi driver stopping behaviors posited by previous authors. Thus our approach is partial equilibrium, as agents in our model take the waiting times and arrival of customers as given rather than endogenously determined in a dynamic equilibrium.

Both models can also be viewed as a stopping rule framework akin to the classic paper by Rust (1987). Rust models the decision to replace bus engines, which weighs routine maintenance costs against full engine replacement, the latter option preventing higher probabilities of catastrophic engine failure. In this setting, after each trip given, a taxi driver must weigh the opportunity for additional fares against a rising cost of fatigue in each day.[10] Similarly to Rust, we use our estimation method to recover parameters defining driver's cost functions.

5.1. **Behavioral model.**

5.1.1. *Period revenues and costs.* Taxi drivers are assumed to have costs of effort that are increasing in hours-worked each day. Each period is a ride. After each ride given, drivers face a discrete decision to continue searching for passengers or quit for the day. In this sense, their labor supply decision boils down to a comparison between the expected profit of searching for an additional unit of time versus the disutility of driving for that much more time. The period payoff function for driver i depends on the decision to quit ($y_{it} = 1$) or keep working ($y_{it} = 0$), and takes the following form:

$$u_i(s_{it}, h_{it}, y_{it}; \theta, X_t) = \begin{cases} \theta_u \cdot s_{it} + \varepsilon_i(1) & \text{if } y_{it} = 1 \\ \theta_{c,01} \cdot h_{it} + \theta_{c,02} \cdot h_{it}^2 + \varepsilon_i(0) & \text{if } y_{it} = 0 \end{cases}. \tag{14}$$

This dynamic labor supply model is an optimal stopping model, in which the taxi driver's dynamic problem ends once he decides to end his current shift. The terminal utility from ending the shift is given in the upper prong of the utility specification above. In this terminal utility, the term $\theta_u \cdot s_{it}$ captures the utility from earnings enjoyed by the taxi driver after ending his shift, and $\theta_{c,01} \cdot h_{it} + \theta_{c,02} \cdot h_{it}^2$ likewise captures the post-shift utility depending on the cumulative hours worked. When a driver continues to drive ($y_{it} = 0$), our specification assumes that he receives (dis-)utility from doing so, which depends on the cumulative hours worked so far in

---

[10]In other words, drivers experience increasingly large marginal utility of leisure as the remaining hours of leisure fall.

TABLE 2. Taxi Trip and Fare Summary Statistics

Trips/Shifts	Variable	Obs.	10%ile	Mean	90%ile	S.D.
Trip Statistics	Trip Revenue ($)	44,090	5.52	12.36	21.74	9.64
	Trip Time (min.)	44,090	3.67	11.50	22.0	8.33
Shift Statistics	Shift Revenue ($)	1857	163.58	293.37	417.09	138.61
	Shift Time (min.)	1857	300.37	529.53	689.20	287.34

Taxi trip and fare data come from New York Taxi and Limousine Commission (TLC) and refer to February 2012 data. The first set of statistics relates to individual taxi trips. The second set of statistics relate to cumulative earnings and time spent in individual driver shifts.

this shift. This lower prong of the utility function measures the cost of the effort exerted by the working driver, which may change as the cumulative hours h_{it} increases.

5.2. **Data.** In 2009, The Taxi and Limousine Commission of New York City (TLC) initiated the Taxi Passenger Enhancement Project, which mandated the use of upgraded metering and information technology in all New York medallion cabs. The technology includes the automated data collection of taxi trip and fare information. We use TLC trip data on all New York City medallion cab rides given in February, 2012. The sample analyzed here consists of 44,090 observations, or about 0.3% of the data. Each row in the data is information related to a single cab ride. Data include driver and medallion identifiers, the exact time and date of pickup and drop-offs, trip distance, and trip time for approximately 44,090 individual taxi rides. Table 2 provides summary statistics.

5.3. **Results.** The estimation results are presented in Table 3. For estimation, we scaled the cumulative time variable to be in units of five-minutes. We find that the terminal utility upon ending a shift grows with earnings, which is weighed against a negative effect of cumulative hours worked, the latter accumulating in each period of continued work. It is important to note that the relatively small coefficient on hours-worked is to be expected, since this utility accrues in every period that a driver continues working, while the utility benefit of earned income is only received once, when the driver stops working for the day. Indeed, we cannot infer the "relative importance" of hours vs. income directly from the estimates, since the utility of each accrues differently.

Given these parameter estimates, in Figure 1 we graph the implied quantile function for the difference in utility shocks $\eta \equiv \epsilon_1 - \epsilon_0$.[11] The density of $\hat{p}$ is plotted as well, which highlights a range over which choice probabilities are actually observed. Outside of this range, we are unable to identify the corresponding quantile function, and in the figure the blue dotted lines represent possible values of the quantile function outside the identified range. Using the density of $\hat{p}$ as a guide, we can recover the quantile function for the range of percentiles approximated by [0.05, 0.25]. A thin vertical dotted line depicts this range. The shocks take (even very large) positive values, with magnitudes in the hundreds; this may imply that there is a large fixed positive component to the terminal utility from quitting.

This feature that, as shown in Figure 1, our approach only yields an incomplete estimate of the error distribution, may be problematic for evaluating some counterfactual policies. For certain counterfactuals, knowledge of the entire distribution of the utility shocks is required, as this distribution feeds agents' beliefs about the future. In ongoing work, we are exploring ways for extrapolating this distribution beyond the range identified by our approach.

TABLE 3. Parameter Estimates

Parameter	Description	Estimate	Std. Error
θ_u	Earnings (upon quitting)	0.9907	0.0118
$\theta_{c,01}$	Cumul. hours (while working)	-0.1359	0.0759
$\theta_{c,02}$	Cumul. hours squared (while working)	-0.0004	0.0002

Note: Standard errors are computed by estimating the model 200 times and reporting the standard deviation of estimates.

While the empirical specification in Table 3 is simple, the behavioral implications of the dynamic model, which we illustrate in Figure 2, are quite rich. Theorem 1 shows that $m(X)'\theta$ corresponds to the difference in the choice-specific value functions for quitting and continuing, $m(X)'\theta = V_1(X) - V_0(X)$, at each value of the state variables $X = (hours\ worked, income)$. Hence, in Figure 2, we plot the estimated $m(X)'\theta$ function at different levels of hours-worked.

---

[11]This contrasts with much of the existing semiparametric estimation literature for discrete-choice models, in which the error distribution is treated as a completely nuisance component, and it is not straightforward to recover estimates of it even given estimates of the model parameters. Since we derive analytical expressions for the error distribution as part of our identification argument, we are able to estimate it once we have estimated the model parameters.

FIGURE 1. Estimated Quantile Function
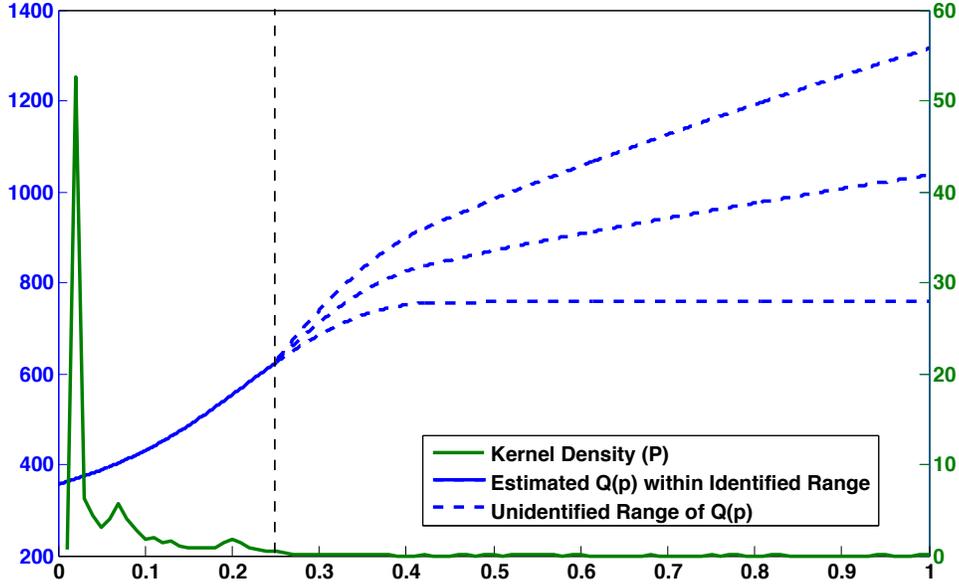
Clearly, the $m(X)'\theta$ curves are increasing in both income and hours-worked (except at hours-worked equal to ten). That is, for most values of the state variables X, the continuation benefits from quitting, $V_1(X)$, are increasing faster than the continuation benefits from continuing to drive, $V_0(X)$, as income rises and hours-worked increases. This implies that, holding hours-worked fixed, drivers are more likely to quit as their income increases; similarly, holding income fixed, drivers are more likely to quit as their hours-worked increases.

In Figure 2, we also plot, using stars, the actual average income in the raw data earned by drivers who quit at different values of hours-worked.[12] As we described above, much of the existing empirical literature on taxicab driver behavior has focused on testing whether drivers' wage elasticities are positive or negative, where positive elasticities are viewed as a corroboration of the classic model of labor-leisure choice, and negative elasticities are taken as evidence of a behavioral "income targeting" model. In both types of models, the wage rate is taken to be exogenous by the drivers and unchanging throughout the course of the day. Drivers then decide how many hours to work for each given wage rate.

---

[12]This illustrates the ranges of income for which the $m(X)'\theta$ functions would be estimated most precisely. Currently, we do not include standard error bands in Figure 2 as that would complicate the picture substantially.

FIGURE 2. Estimated Choice-specific Value Function Differences: $V_1(X) - V_0(X)$
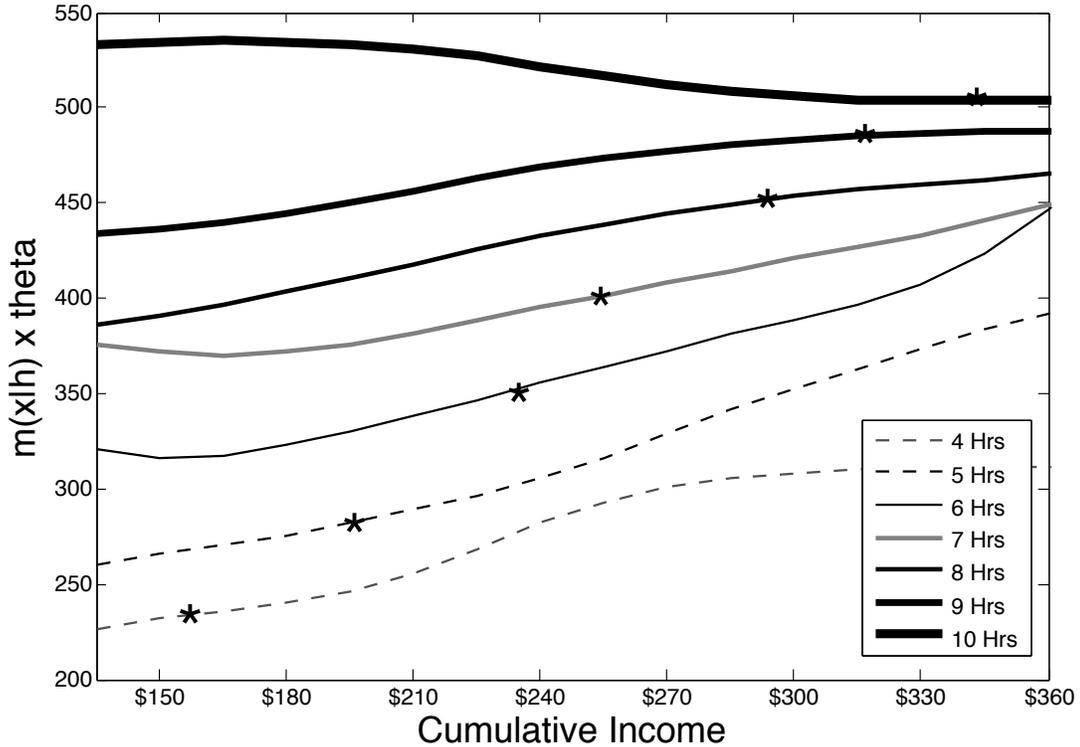
We report the estimates of the $m(X)'\theta$ function, as in Theorem 1, for $X = (hours\ worked, income)$. For fixed values of hours-worked, we graph $m(X)'\theta$ as a function of income. Stars (*) mark average income earned by drivers for a given hours-worked, as observed in the raw data.

In our model, however, income evolves stochastically; drivers' "wage rates" are random and vary across the shift. Accordingly, a driver reconsiders the decision to end the shift after each fare. Hence, implied wage elasticities are not straightforward to compute in our modelling framework. That caveat aside, the results in Figure 2 do have implications for wage elasticities if we associate higher cumulative income means with a higher wage rate (as is done in the existing literature). Under this interpretation, the graphs suggest that by holding hours fixed, the quitting probability is increasing in wage rate, which is consistent with negative wage elasticities. To illustrate this, consider a thought experiment where we have two drivers who have only driven four hours; the first driver already has income of \$300, while the second only has income of \$180.

The graphs suggest that the benefits from quitting are relatively larger for the first rather than the second driver, and hence that the first should quit more readily than the second.

On the other hand, if we just use reduced-form evidence from the raw data, the clear upward-sloping trend in the stars suggests the opposite result: that hours-worked is increasing in wage (positive income elasticities). This highlights the benefits of estimating a structural dynamic behavioral model. Indeed, our results show that, once dynamics are modelled, it is possible to obtain "nonstandard" (i.e., negative) wage elasticities from a model in which drivers' utility functions do not have any explicitly "nonstandard" features, such as reference dependence or loss aversion.

6. CONCLUSIONS

In this paper we consider the estimation of dynamic binary discrete choice models in a semiparametric setting, in which the per-period utility functions are parameterized as single-index functions, but the distribution of the utility shocks is left unspecified and treated as nuisance components of the model. This setup differs from most of the existing work on estimation and identification of dynamic discrete choice models. For identification, we derive a new recursive representation for the unknown quantile function of the utility shocks; our argument requires no additional exclusion restrictions beyond the conditional independence conditions assumed in the typical parametric dynamic-discrete choice literature (e.g. Rust (1987, 1994)). Accordingly, we obtain a single-index representation for the conditional choice probabilities in the model, which permits us to estimate the model using classic estimators from the existing semiparametric binary choice literature.

In particular, we use Powell, Stock and Stoker's (1989) kernel-based estimator to estimate the dynamic discrete choice model. We show that the estimator has the same asymptotic properties as PSS's original estimator (for static discrete-choice models), under mild conditions. Significantly, the computational procedure is quite simple, because the estimator for the parameters can be expressed in closed-form. Monte Carlo simulations show that the estimator works well even in moderately-sized samples. We provide an empirical application to estimate the dynamic labor supply problem for taxicab drivers in New York City.

25

More broadly, the analysis in this paper has opened possibilities for the use of classic closed-form estimators from the semiparametric literature, which were proposed for estimation of static models, to dynamic models. We will continue exploring these possibilities in future work.

References

AGUIRREGABIRIA, V., AND P. MIRA (2007): "Sequential estimation of dynamic discrete games," *Econometrica*, pp. 1–53.

AI, C., AND X. CHEN (2003): "Efficient estimation of models with conditional moment restrictions containing unknown functions," *Econometrica*, 71(6), 1795–1843.

BAJARI, P., L. BENKARD, AND J. LEVIN (2007): "Estimating Dynamic Models of Imperfect Competition," *Econometrica*, 75(5), 1331–1370.

BLEVINS, J. (2014): "Nonparametric identification of dynamic decision processes with discrete and continuous choices," *Quantitative Economics*, 5, 531–554.

CAMERER, C., L. BABCOCK, G. LOEWENSTEIN, AND R. THALER (1997): "Labor supply of New York City cabdrivers: One day at a time," *Quarterly Journal of Economics*, 112(2), 407–441.

CHEN, L.-Y. (2014): "Identification of Discrete Choice Dynamic Programming Models with Nonparametric Distribution of Unobservables," Working Paper, Academia Sinica.

CRAWFORD, V. P., AND J. MENG (2011): "New york city cab drivers' labor supply revisited: Reference-dependent preferences with rational expectations targets for hours and income," *American Economic Review*, 101(5), 1912–1932.

FARBER, H. S. (2005): "Is Tomorrow Another Day? The Labor Supply of New York City Cabdrivers," *Journal of Political Economy*, 113(1), 46–82.

——— (2008): "Reference-dependent preferences and labor supply: The case of New York City taxi drivers," *American Economic Review*, 98(3), 1069–1082.

——— (2014): "Why You Can't Find a Taxi in the Rain and Other Labor Supply Lessons from Cab Drivers," Discussion paper, National Bureau of Economic Research.

FRECHETTE, G., A. LIZZERI, AND T. SALZ (2015): "Frictions in a Competitive, Regulated Market: Evidence from Taxis," Working Paper, New York University.

GUERRE, E., I. PERRIGNE, AND Q. VUONG (2000): "Optimal Nonparametric Estimation of First-Price Auctions," *Econometrica*, 68, 525–74.

HONG, H., AND M. SHUM (2010): "Pairwise-difference estimation of a dynamic optimization model," *Review of Economic Studies*, 77(1), 273–304.

HOROWITZ, J. L. (1992): "A smoothed maximum score estimator for the binary response model," *Econometrica*, pp. 505–531.

——— (2009): *Semiparametric and nonparametric methods in econometrics*. Springer.

HOTZ, V. J., AND R. A. MILLER (1993): "Conditional choice probabilities and the estimation of dynamic models," *Review of Economic Studies*, 60(3), 497–529.

ICHIMURA, H. (1993): "Semiparametric least squares (SLS) and weighted SLS estimation of single-index models," *Journal of Econometrics*, 58(1), 71–120.

ICHIMURA, H., AND L.-F. LEE (1991): "Semiparametric least squares estimation of multiple index models: single equation estimation," in *Nonparametric and semiparametric methods in econometrics and statistics: Proceedings of the Fifth International Symposium in Economic Theory and Econometrics. Cambridge*, pp. 3–49.

KLEIN, R., AND R. SPADY (1993): "An efficient semiparametric estimator for binary response models," *Econometrica*, 61(2), 387–421.

LEWBEL, A. (1998): "Semiparametric latent variable model estimation with endogenous or mismeasured regressors," *Econometrica*, pp. 105–121.

MAGNAC, T., AND D. THESMAR (2002): "Identifying dynamic discrete decision processes," *Econometrica*, 70(2), 801–816.

MANSKI, C. (1975): "Maximum score estimation of the stochastic utility model of choice," *Journal of Econometrics*, 3(3), 205–228.

——— (1985): "Semiparametric analysis of discrete response:: Asymptotic properties of the maximum score estimator," *Journal of Econometrics*, 27(3), 313–333.

MCFADDEN, D. (1978): *Modelling the choice of residential location*. Institute of Transportation Studies, University of California.

——— (1980): "Econometric models for probabilistic choice among products," *Journal of Business*, pp. S13–S29.

MCFADDEN, D., AND K. TRAIN (2000): "Mixed MNL models for discrete response," *Journal of Applied Econometrics*, 15, 447–470.

NORETS, A., AND X. TANG (2014): "Semiparametric Inference in dynamic binary choice models," *Review of Economic Studies*, 81(3), 1229–1262.

PAKES, A. (1986): "Patents as Options: Some Estimates of the Value of Holding European Patent Stocks," *Econometrica*, 54(4), 755–84.

PESENDORFER, M., AND P. SCHMIDT-DENGLER (2008): "Asymptotic least squares estimators for dynamic games," *Review of Economic Studies*, 75(3), 901–928.

POWELL, J., J. STOCK, AND T. STOKER (1989): "Semiparametric estimation of index coefficients," *Econometrica*, pp. 1403–1430.

RUST, J. (1987): "Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher," *Econometrica*, pp. 999–1033.

——— (1994): "Structural estimation of Markov decision processes," *Handbook of econometrics*, 4(4), 3081–3143.

SRISUMA, S., AND O. LINTON (2012): "Semiparametric estimation of Markov decision processes with continuous state space," *Journal of Econometrics*, 166(2), 320–341.

WOLPIN, K. I. (1984): "An estimable dynamic stochastic model of fertility and child mortality," *Journal of Political Economy*, pp. 852–874.

ZEMYAN, S. M. (2012): *The classical theory of integral equations: a concise treatment*. Springer Science & Business Media.

A.1. **Proof of Lemma 1.**

Proof. First, note that the resolvent kernel R^* is well–defined. This is because $\beta^{s-1} f_{X^{[s]}|X}(x'|x) \to 0$ as $s \to +\infty$. Under assumption B, the solution $V^e(x)$ is also well defined.

Because it is straightforward to verify that the solution in the lemma solves eq. (3), Hence, it suffices to show the uniqueness of the solution. Eq. (3) can be rewritten as

$$V^e(x) = u^e(x) + \beta \cdot \int V^e(x') \cdot f_{X'|X}(x'|x)dx', \ \ \forall \, x \in \mathscr{S}_X,$$

which is an FIE–2. Then, we apply the method of Successive Approximation (see e.g. Zemyan, 2012). Specifically, let $V^*(\cdot)$ be an alternative solution to (3). Then, we have

$$V^*(x) = u^e(x) + \beta \int_{\mathscr{S}_X} V^*(x') \cdot f_{X'|X}(x'|x)dx'.$$

Let $\nu(x) = V^e(x) - V^*(x)$. Then $\nu(x)$ satisfies the following equation:

$$\nu(x) = \beta \int_{\mathscr{S}_X} \nu(x') \cdot f_{X'|X}(x'|x)dx'.$$

It suffices to show that $\nu(\cdot)$ has the unique solution: $\nu(x) = 0$. To see this, we substitute the left–hand side as an expression of ν into the integrand:

$$\nu(x) = \beta^2 \int_{\mathscr{S}_X} \int_{\mathscr{S}_X} \nu(\tilde{x}) \cdot f_{X'|X}(\tilde{x}|x')d\tilde{x} \cdot f_{X'|X}(x'|x)dx' = \beta^2 \int_{\mathscr{S}_X} \nu(x') \cdot f_{X^{[2]}|X}(x'|x)dx'.$$

Repeating this process, then we have: for all $t \geq 1$

$$\nu(x) = \beta^t \int_{\mathscr{S}_X} \nu(x') \cdot f_{X^{[t]}|X}(x'|x)dx'.$$

For the stationary Markov equilibrium, $f_{X^{[t]}|X}(x'|x)$ converges to $f_X(x')$ as $t \to \infty$. Hence, the right–hand side converges to zero as t goes to infinity. It follows that $\nu(x) = 0$ for all $x \in \mathscr{S}_X$. $\qquad \square$

A.2. **Proof of Lemma 4.**

Proof. The result follows the Theorem of Successive Approximation (see e.g. Zemyan, 2012). $\qquad \square$

A.3. **Proof of Theorem 3.** The estimator is defined in (13). For the consistency of $\hat{\theta}$, we need $h_\theta \to 0$, $Th_\theta^{k_\theta+1} \to \infty$ and $\mathbb{E}|\hat{m}(X) - m(X)| = o(h_\theta)$ as $T \to \infty$. Note that the last condition ensures the estimation error in $\hat{m}$ is negligible.

29

Let $\tilde{\theta}$ be the infeasible estimator

$$\tilde{\theta} = -\frac{2}{T(T-1)} \sum_{t=1}^{T} \sum_{s \neq t} \left[\frac{1}{h_\theta^{k_\theta+1}} \times \nabla K_\theta \left(\frac{m(X_t) - m(X_s)}{h_\theta} \right) \times Y_s \right].$$

The asymptotic analysis for $\tilde{\theta}$ was done in Powell, Stock, and Stoker (1989). They show that the variance term in $\tilde{\theta}$ has the order T^{-1} if $Th_\theta^{k_\theta+2} \to \infty$, while the bias term has the order h_θ^P. Therefore, if $T^{1/2}h_\theta^p \to 0$, then the bias term disappear faster than $T^{-1/2}$. The leading term left is the variance term – the $\tilde{\theta}$ converges at the rate $T^{-1/2}$. Our arguments piggybacks off of this argument, as we will show here that $T^{1/2}(\hat{\theta} - \theta)$ is identical to $T^{1/2}(\tilde{\theta} - \theta)$ by a negligible factor; that is, our estimator and the infeasible estimator have the same limiting distribution (corresponding to that derived in Powell, Stock, and Stoker (1989)).

By Taylor expansion, we have

$$\hat{\theta} = \tilde{\theta} - \frac{2}{T(T-1)} \sum_{t=1}^{T} \sum_{s \neq t} \left[\frac{1}{h_\theta^{k_\theta+2}} \nabla^2 K_\theta \left(\frac{m(X_t) - m(X_s)}{h_\theta} \right) \times Y_s \times \left(\hat{m}(X_t) - m(X_t) \right) \right]$$

$$+ \frac{2}{T(T-1)} \sum_{t=1}^{T} \sum_{s \neq t} \left[\frac{1}{h_\theta^{k_\theta+2}} \nabla^2 K_\theta \left(\frac{m(X_t) - m(X_s)}{h_\theta} \right) \times Y_s \times \left(\hat{m}(X_s) - m(X_s) \right) \right]$$

$$+ O_p(h_\theta^{-3} \cdot \mathbb{E}\|\hat{m}(X) - m(X)\|^2) \equiv \tilde{\theta} + \mathbb{A}_1 + \mathbb{A}_2 + \mathbb{B} \quad (15)$$

We will show that $\mathbb{A}_1 + \mathbb{A}_2 + \mathbb{B}_2$ are all $o_p(T^{-1/2})$ implying $T^{1/2}(\hat{\theta} - \tilde{\theta})$ is negligible. First, by Assumption J(i), we have

$$h_\theta^{-3} \times \mathbb{E}\|\hat{m}(X) - m(X)\|^2 = h_\theta^{-3} \times o_p(T^{-1/2}h_\theta^3) = o_p(T^{-1/2}). \quad (16)$$

Then, $\mathbb{B} = o_p(T^{-1/2})$.

Next we show $\mathbb{A}_1$ and $\mathbb{A}_2 = o_p(T^{-1/2})$. For simplicity, we only provide an argument for $\mathbb{A}_1$ (that for $\mathbb{A}_2$ is analogous).

Define

$$\tilde{\mathbb{A}}_1 \equiv -\frac{2}{T(T-1)} \sum_{t=1}^{T} \sum_{s \neq t} \left[\frac{1}{h_\theta^{k_\theta+2}} \nabla^2 K_\theta \left(\frac{m(X_t) - m(X_s)}{h_\theta} \right) Y_s \times \left[\mathbb{E}[\hat{m}(X_t)|X_t, X_s] - m(X_t) \right] \right]$$

Clearly $\mathbb{E}(\mathbb{A}_1) = \mathbb{E}(\tilde{\mathbb{A}}_1)$. Following Powell, Stock, and Stoker (1989), we now establish two properties:

$$(a): \tilde{\mathbb{A}}_1 = o_p(T^{-1/2});$$

$$(b): T \times \text{Var}(\mathbb{A}_1 - \tilde{\mathbb{A}}_1) \to 0,$$

which together imply $\mathbb{A}_1 = o_p(T^{-1/2})$.

For property (a), by Assumption J(iii),

$$\mathbb{E}[\hat{m}(X_t)|X_t, X_s] = \mathbb{E}[\hat{m}_{t,-s}|X_t, X_s] + o_p(T^{-1/2}h_\theta^2) = \mathbb{E}[\hat{m}(X_t)|X_t] + o_p(T^{-1/2}h_\theta^2).$$

Then, we have

$$
\begin{aligned}
&\tilde{\mathbb{A}}_1 \\
=\ & -\frac{2}{T(T-1)} \sum_{t=1}^{T} \sum_{s \neq t} \left[\frac{1}{h_\theta^{k_\theta+2}} \nabla^2 K_\theta \left(\frac{m(X_t) - m(X_s)}{h_\theta} \right) Y_s \times \left[\mathbb{E}[\hat{m}(X_t)|X_t] - m(X_t) \right] \right] + o_p(T^{-1/2}) \\
\equiv\ & \mathbb{C}_1 + o_p(T^{-1/2}).
\end{aligned}
$$

Because

$$
\begin{aligned}
\mathbb{E}|\mathbb{C}_1| \ &\leq\ 2\mathbb{E} \left| \frac{1}{h_\theta^{k_\theta+2}} \nabla^2 K_\theta \left(\frac{m(X_t) - m(X_s)}{h_\theta} \right) \times \left[\mathbb{E}[\hat{m}(X_t)|X_t] - m(X_t) \right] \right| \\
&\leq\ 2\overline{C} \times \frac{1}{h_\theta^2} \mathbb{E} \left\| \mathbb{E}[\hat{m}(X) - m(X)|X] \right\|
\end{aligned}
$$

for some positive $\overline{C} < \infty$. Hence, by Assumption J(ii), property (a) obtains.

For property (b), note that

$$
\mathbb{A}_1 - \tilde{\mathbb{A}}_1 \ \equiv\ -\frac{2}{T(T-1)} \sum_{t=1}^{T} \sum_{s \neq t} \phi_{T,s,t} \times \left[\hat{m}(X_t) - \mathbb{E}[\hat{m}(X_t)|X_t] \right] + o_p(T^{-1/2}) \equiv \mathbb{C}_2 + o_p(T^{-1/2})
$$

where $\phi_{T,s,t} = \frac{1}{h_\theta^{k_\theta+2}} \nabla^2 K_\theta \left(\frac{m(X_t)-m(X_s)}{h_\theta} \right) Y_s$.

Clearly,

$$
\begin{aligned}
\text{Var}(\mathbb{C}_2) =\ & \frac{4}{T^2(T-1)^2} \sum_{t=1}^{T} \sum_{s \neq t} \text{Var}\left(\phi_{T,s,t} \times \left[\hat{m}(X_t) - \mathbb{E}[\hat{m}(X_t)|X_t] \right] \right) \\
+\ & \frac{4}{T^2(T-1)^2} \sum_{t=1}^{T} \sum_{s \neq t} \sum_{s' \neq t, s} \text{Cov}\left(\phi_{T,s,t}\left[\hat{m}(X_t) - \mathbb{E}[\hat{m}(X_t)|X_t] \right], \phi_{T,s',t}\left[\hat{m}(X_t) - \mathbb{E}[\hat{m}(X_t)|X_t] \right] \right) \\
+\ & \frac{4}{T^2(T-1)^2} \sum_{t=1}^{T} \sum_{s \neq t} \sum_{t' \neq t, s} \sum_{s' \neq t, s, t'} \text{Cov}\left(\phi_{T,s,t}\left[\hat{m}(X_t) - \mathbb{E}[\hat{m}(X_t)|X_t] \right], \phi_{T,s',t'}\left[\hat{m}(X_{t'}) - \mathbb{E}[\hat{m}(X_{t'})|X_{t'}] \right] \right) \\
=\ & O(T^{-2}h_\theta^{-k_\theta-4}) \times \mathbb{E}\{ \hat{m}(X) - \mathbb{E}[\hat{m}(X)|X] \}^2 \\
+\ & \frac{4}{T} \text{Cov}\left(\phi_{T,2,1}\left[\hat{m}(X_1) - \mathbb{E}[\hat{m}(X_1)|X_1] \right], \phi_{T,3,1}\left[\hat{m}(X_1) - \mathbb{E}[\hat{m}(X_1)|X_1] \right] \right) \\
+\ & 4 \text{Cov}\left(\phi_{T,2,1}\left[\hat{m}(X_1) - \mathbb{E}[\hat{m}(X_1)|X_1] \right], \phi_{T,4,3}\left[\hat{m}(X_3) - \mathbb{E}[\hat{m}(X_3)|X_3] \right] \right).
\end{aligned}
$$

Note that

$$\text{Cov}\Big(\phi_{T,2,1}\big[\hat{m}(X_1) - \mathbb{E}[\hat{m}(X_1)|X_1]\big], \phi_{T,4,3}\big[\hat{m}(X_3) - \mathbb{E}[\hat{m}(X_3)|X_3]\big]\Big)$$

$$= \quad \mathbb{E}\left\{\phi_{T,2,1}\phi_{T,4,3}\big[\hat{m}(X_1) - \mathbb{E}[\hat{m}(X_1)|X_1]\big] \times \big[\hat{m}(X_3) - \mathbb{E}[\hat{m}(X_3)|X_3]\big]\right\}$$

$$- \quad \mathbb{E}\left\{\phi_{T,2,1}\big[\hat{m}(X_1) - \mathbb{E}[\hat{m}(X_1)|X_1]\big]\right\} \times \mathbb{E}\left\{\phi_{T,4,3}\big[\hat{m}(X_3) - \mathbb{E}[\hat{m}(X_3)|X_3]\big]\right\}.$$

By Assumption J(iii),

$$\mathbb{E}\left\{\phi_{T,2,1}\big[\hat{m}(X_1) - \mathbb{E}[\hat{m}(X_1)|X_1]\big]\right\}$$

$$= \mathbb{E}\left\{\phi_{T,2,1}\big[\hat{m}_{1,-2} - \mathbb{E}[\hat{m}_{1,-2}|X_1]\big]\right\} + O_p(h_\theta^{-2}) \times o_p(T^{-1/2}h_\theta^2) = o_p(T^{-1/2}).$$

Furthermore, by the law of iterated expectation (conditioning on the sigma algebra: $\mathscr{F}_2, \mathscr{F}_4, \mathscr{F}_{5,\cdots,n}$),

$$\mathbb{E}\left\{\phi_{T,2,1}\phi_{T,4,3}\big[\hat{m}(X_1) - \mathbb{E}[\hat{m}(X_1)|X_1]\big] \times \big[\hat{m}(X_3) - \mathbb{E}[\hat{m}(X_3)|X_3]\big]\right\}$$

$$= \quad O_p(h_\theta^{-4}) \times o_p(T^{-1/2}h_\theta^2) \times o_p(T^{-1/2}h_\theta^2)$$

$$= \quad o_p(T^{-1}),$$

where the term $o_p(T^{-1/2}h_\theta^2)$ is due to the differences $\hat{m}(X_1) - \hat{m}_{1,-3}$ and $\hat{m}(X_3) - \hat{m}_{3,-1}$. Therefore, the last term in $\text{Var}(\mathbb{C}_2)$ is $o_p(T^{-1})$.

Moreover, because

$$\frac{1}{T}\text{Cov}\Big(\phi_{T,2,1}\big[\hat{m}(X_1) - \mathbb{E}[\hat{m}(X_1)|X_1]\big], \phi_{T,3,1}\big[\hat{m}(X_1) - \mathbb{E}[\hat{m}(X_1)|X_1]\big]\Big)$$

$$= \quad \frac{1}{T}\mathbb{E}\left\{\phi_{T,2,1}\phi_{T,3,1}\big[\hat{m}(X_1) - \mathbb{E}[\hat{m}(X_1)|X_1]\big]^2\right\}$$

$$= \quad O(T^{-1}h_\theta^{-4}) \times \mathbb{E}\left\{\hat{m}(X) - \mathbb{E}[\hat{m}(X)|X]\right\}^2.$$

Then a sufficient condition for property (b) is

$$\mathbb{E}\left\{\hat{m}(X) - \mathbb{E}[\hat{m}(X)|X]\right\}^2 = o(h_\theta^4).$$

Note that this condition is implied by Assumption J(i).

Hence, we have shown that our estimator $\hat{\theta}$ and the infeasible estimator $\tilde{\theta}$ differ by an amount which is $o_p(T^{-1/2})$. Hence, the asymptotic properties for $\hat{\theta}$ are the same as those for the infeasible estimator $\tilde{\theta}$, which were previously established in Powell, Stock, and Stoker (1989).