

A CLOSED-FORM ESTIMATOR FOR DYNAMIC DISCRETE CHOICE MODELS: ASSESSING TAXICAB DRIVERS' DYNAMIC LABOR SUPPLY

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

Rewriting in progress: please excuse the débris!

ABSTRACT. We propose a new closed-form estimator for 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. Compared to other existing estimators for these models, our estimator requires no iterative nonlinear optimization, rendering issues of starting values or convergence criteria irrelevant. Using our approach, we estimate an optimal stopping model for taxicab drivers' labor supply decisions. Our results show that, once the inherent dynamic in taxicab drivers' work decisions are accounted for, 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.

Keywords: Dynamic discrete choice model, Closed form estimator, Optimal stopping, Taxicab industry, Labor supply, Negative wage elasticities, Semiparametric average derivative estimation

JEL: C14, D91, C41, L91

Date: Monday 30th October, 2017.

*We thank Hassan Afrouzi, Saroj Bhattarai, Kirill Pogorelskiy, Eduardo Souza-Rodrigues, Sorawoot (Tang) Srisuma, and seminar participants at Carnegie-Mellon, University of Arizona, UC-Irvine, University of Queensland, University of Melbourne, University of Washington (Seattle), UT Austin, Virginia, Wisconsin, University of Colorado at Boulder, University of Chicago, Shanghai University of Finance and Economics, Texas Metrics Camp, the URBAN conference (University of Copenhagen), and the Structural Dynamic Models Conference (University College London) for helpful discussions.

*Department of Economics, Princeton University, nbuchholz@princeton.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.

Despite its popularity, most of the estimators proposed for DDC models are time-consuming, typically involving iterative nonlinear optimization procedures which can be sensitive to starting values and convergence criteria. In this paper we introduce a new estimator for a large class of DDC models (including optimal stopping models) which circumvents these difficulties with the existing estimators. It provides *closed form* estimates for the model parameters – that is, our estimator is non-iterative, thus rendering starting values and convergence criteria irrelevant. In addition, our approach is also semiparametric, and does not require parametric assumptions to be made on the distribution of the error terms in the discrete choice model; this contrasts with many applications of dynamic discrete choice models, which assume that these errors follow the extreme-value distribution, leading to logit choice probabilities.

We apply our estimator to a model of optimal labor supply of taxicab drivers. The labor supply decisions of taxicab drivers has been an ongoing area of research ever since the seminal paper of Camerer, Babcock, Loewenstein, and Thaler (1997), who found evidence of negative income elasticities of labor supply– that is, working fewer hours under high wage rates. Such a finding is inconsistent with textbook neoclassical labor supply models but congruent with behavioral models of income targeting, in which agents with flexible hours set an income target and work until the target is reached. A large follow-up literature, using a variety of datasets, has both confirmed and disputed these findings.

In this paper, we use a new and comprehensive dataset of New York City taxi drivers (the world's largest taxicab market), and take a new approach to this question. We model taxicab drivers' labor supply decisions as emerging from an *optimal stopping* problem: in a stochastically evolving environment, drivers give rides and, after each fare, decide whether or not to continue

working or quit for the day. Their stopping rule is determined by both their cumulative income, as well as total amount of time worked, during the day.

Our results reconcile the previous literature to a certain extent. Estimates of drivers' optimal stopping rules show that, holding hours worked, drivers are more likely to quit at higher levels of cumulative income. The implied reduced-form depends on the specifics of and variability in the income process: on days with faster income accumulation, this may appear as negative income elasticities, while on days with slower income accumulation, this may look like positive income elasticities. More broadly, these findings suggest that once the inherent dynamic optimization aspect of taxicab drivers' labor supply decisions are accounted for, there is no need to add non-standard behavioral parameters to the model to explain their quitting behavior – it emerges as an outcome along the optimal dynamic decision-making path.

Our closed-form estimator for dynamic discrete choice models relies on a new recursive representation for the unknown quantile function of the utility shocks which we derive in this paper. This leads to a representation for the conditional choice probabilities which is linear in the utility function parameters, 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, which can be expressed in closed-form. 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). Monte Carlo simulations demonstrate that our estimator performs well in small samples.

In section X we do [...]

2. AN OPTIMAL STOPPING PROBLEM OF TAXICAB DRIVER LABOR SUPPLY

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. One main contribution in this paper is to 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.

Our dynamic modelling approach contrasts with much of the existing 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 first paper develops a static stopping rule model which explores similar forces to our model, showing that drivers stopping is most reliably predicted by hours instead of income. 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.¹ 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

¹The 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.² For the remainder of this section, we provide a utility specification for taxicab drivers in our model. We will return to the model in Section 5 below, after introducing and discussing our new closed form estimator for dynamic discrete choice models.

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:

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

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

²In other words, drivers experience increasingly large marginal utility of leisure as the remaining hours of leisure fall.

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.

3. SINGLE AGENT DYNAMIC DISCRETE CHOICE MODEL

In this section, we present our new closed-form estimator for dynamic discrete-choice models, which we will use for the estimation of the model parameters for the optimal stopping model from the previous section. For convenience, we will discuss our estimator in some degree of generality, instead of referring specifically to the taxicab labor supply model. Readers who wish to skip the econometric details may skip to Section X, where we discuss the results from the estimation of the taxicab driver optimal stopping model.

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

$$u_t(Y_t, X_t, \epsilon_t) = \begin{cases} W_1(X_t)^\top \theta_1 + \epsilon_{1t}, & \text{if } Y_t = 1; \\ W_0(X_t)^\top \theta_0 + \epsilon_{0t}, & \text{if } Y_t = 0. \end{cases} \quad (2)$$

In the above, $W_0(X_t) \in \mathbb{R}^{k_0}$ (resp. $W_1(X_t) \in \mathbb{R}^{k_1}$) denotes known transformations of the state variables X_t which affect the per-period utility from choosing $Y_t = 0$ (resp. $Y_t = 1$), and $\epsilon_t \equiv (\epsilon_{0t}, \epsilon_{1t})^\top \in \mathbb{R}^2$ are the agent's action-specific payoff shocks. The structural parameters which are of interest are $\theta_d \in \mathbb{R}^{k_d}$, for $d \in \{0, 1\}$. In what follows, let $W(X) \equiv \{W_0(X), W_1(X)\}$ denote the full set of transformed state variables at X . 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 when possible.

This specification of the per-period utility functions in eq. (2), as single-indices of the transformed state variables $W(X)$ 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 (which is assumed to be known for purposes of identification and estimation)³ 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.

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

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

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.⁴ 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.

To our knowledge, only a handful of papers consider estimation 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

³The assumption that β is known is commonplace in the applied DDC literature. See Magnac and Thesmar (2002) and Fang and Wang (2015), among others, for discussion on the identifiability of β .

⁴McFadden and Train (2000) show such properties for the mixed logit specification of static multinomial choice models.

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.

Aguirregabiria (2010) shows the joint nonparametric identification of utilities and the shock distribution in a class of finite-horizon dynamic binary choice models. His identification argument relies on the existence of a final period in the decision problem, and hence may not apply to infinite-horizon models as considered in this paper. 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 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). 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. Unlike these papers, we do not use exclusion restrictions for identification, 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.

The semiparametric static 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 to our work. 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.

3.1. 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, as it establishes that the shocks ϵ are fully independent of the observed state variables X . This rules out heteroskedasticity in the unobserved shocks, which is accommodated in other papers in the DDC literature (eg. Magnac and Thesmar (2002), Aguirregabiria (2010), among others). It is possible, as in Blevins (2014), 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), which allows for some degree of heteroskedasticity in ϵ . The identification and estimation procedure described in this paper follow through, with the additional conditioning on X_A at every step.

Under assumption A, the value function can be written as

$$V(X, \epsilon) = \max \left\{ W_1^\top \theta_1 + \epsilon_1 + \beta \mathbb{E}[V(X', \epsilon')|X, Y = 1], W_0^\top \theta_0 + \epsilon_0 + \beta \mathbb{E}[V(X', \epsilon')|X, Y = 0] \right\}.$$

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

$$\eta^*(X) \equiv W_1^\top \theta_1 - W_0^\top \theta_0 + \beta \left\{ \mathbb{E}[V(X', \epsilon')|X, Y = 1] - \mathbb{E}[V(X', \epsilon')|X, Y = 0] \right\}. \quad (3)$$

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

$$u^e(X) \equiv \mathbb{E}(\epsilon_0) + W_1^\top \theta_1 \cdot F_\eta(\eta^*(X)) + W_0^\top \theta_0 \cdot [1 - F_\eta(\eta^*(X))] - \mathbb{E}\{\eta \cdot \mathbb{1}[\eta \leq \eta^*(X)]\}, \quad (4)$$

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

$$V^e(X) = u^e(X) + \beta \cdot \mathbb{E}[V^e(X')|X]. \quad (5)$$

It is worth pointing out that eq. (5) 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, 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 $\mathbb{E}(\|W_d^{[s]}\| | X) < \infty$ a.s., where $(^{[s]})$ denotes the next s period values.

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

Srisuma and Linton (2012) pioneered the use of tools for solving type 2 integral equations for estimating dynamic discrete-choice models, and the following Lemma builds on their findings.

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

$$V^e(x) = u^e(x) + \beta \int_{\mathcal{S}_X} R^*(x', x; \beta) \cdot u^e(x') dx', \quad \forall x \in \mathcal{S}_X, \quad (6)$$

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. (5).

More succinctly, eq. (6) can be rewritten as

$$V^e(X) = u^e(X) + \sum_{s=1}^{\infty} \beta^s \cdot \mathbb{E}[u^e(X^{[s]}) | X]. \quad (7)$$

In operator notation, eq. (7) 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).⁵

3.2. Optimality Condition. To characterize the optimum, 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. (7), along with Lemma 1, eq. (3) becomes

$$\eta^*(X) = W_1^T \theta_1 - W_0^T \theta_0 + \sum_{s=1}^{\infty} \beta^s \left\{ \mathbb{E}[u^e(X^{[s]}) | X, Y = 1] - \mathbb{E}[u^e(X^{[s]}) | X, Y = 0] \right\}. \quad (8)$$

⁵In 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 .

Moreover, let $\phi_d(X) \equiv (-1)^{d+1}W_d + \sum_{s=1}^{\infty} \beta^s \left\{ \mathbb{E}[W_d^{[s]} \mathbb{1}_{Y^{[s]}=d} | X, Y = 1] - \mathbb{E}[W_d^{[s]} \mathbb{1}_{Y^{[s]}=d} | X, Y = 0] \right\}$. Then, it follows from (4),

$$\begin{aligned} \eta^*(X) &= \phi^\top(X) \cdot \theta \\ &= \sum_{s=1}^{\infty} \beta^s \left\{ \mathbb{E}[\eta^{[s]} \mathbb{1}(\eta^{[s]} \leq \eta^*(X^{[s]})) | X, Y = 1] - \mathbb{E}[\eta^{[s]} \mathbb{1}(\eta^{[s]} \leq \eta^*(X^{[s]})) | X, Y = 0] \right\}, \end{aligned} \quad (9)$$

where $\phi(X) = (\phi_0^\top(X), \phi_1^\top(X))^\top$ and $\theta = (\theta_0^\top, \theta_1^\top)^\top$. As a matter of fact, eq. (9) characterizes the optimal decision rule in the single-agent infinite-horizon binary decision problem.

4. CLOSED-FORM ESTIMATOR FOR MODEL PARAMETERS

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, $F_\eta(\cdot)$ is strictly increasing on its support \mathbb{R} . Let $Q(\cdot)$ be the quantile function of $F_\eta(\cdot)$, i.e., $Q(\cdot) = F_\eta^{-1}(\cdot)$.

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

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

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

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

$$Q(p(X)) + \sum_{s=1}^{\infty} \beta^s \left\{ \mathbb{E} \left[\int_0^{p(X^{[s]})} Q(\tau) d\tau | X, Y = 1 \right] - \mathbb{E} \left[\int_0^{p(X^{[s]})} Q(\tau) d\tau | X, Y = 0 \right] \right\} = \phi^\top(X) \cdot \theta. \quad (10)$$

When $Q(\cdot)$ is known, then everything in (10) is known except for θ . If, in addition, the matrix $\mathbb{E}[\phi(X)\phi^\top(X)]$ is invertible, then θ can be estimated using nonlinear least-squares on eq. (10). This approach is related to Pesendorfer and Schmidt-Dengler (2008).

However, in this paper, $Q(\cdot)$ is not specified, and we focus here on identifying θ even without knowledge of $Q(\cdot)$. For notational simplicity, we assume the choice probability $p(X)$ is continuously distributed on a closed interval.⁶

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

This assumption requires the state variables X to contain some continuous components. Letting X^D (resp. X^C) denote the discrete (resp. continuous) components of X , a more primitive statement of Assumption D would be that, for fixed values of the discrete components (say $X^D = x^d$), the support of $p(X^C, x^d)$ is a closed interval in $[0, 1]$. As is well-known, 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.⁷

For each $p \in [\underline{p}, \bar{p}]$, 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. (10). Applying the law of iterated expectation yields the following Lemma:

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

$$Q(p) + \beta \int_{\underline{p}}^{\bar{p}} \int_{\underline{p}}^{p'} Q(\tau) d\tau \cdot \pi(p', p; \beta) dp' = z(p)^\top \cdot \theta, \quad \forall p \in [\underline{p}, \bar{p}], \quad (11)$$

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. (11) is also an FIE-2. To see this, let $\Pi(\tau, p; \beta) \equiv \int_0^1 \mathbb{1}(p' \leq \tau) \pi(p', p; \beta) dp' = \sum_{s=1}^{\infty} \beta^{s-1} [F_{p(X^{[s]}|p(X), Y}(\tau|p, 1) - F_{p(X^{[s]}|p(X), Y}(\tau|p, 0)]$. Then, the second term of eq. (11) can

⁶This interval-support restriction can be relaxed at expositional expense. For instance, suppose $\mathcal{S}_{p(X)}$ is a non-degenerate compact subset of $[0, 1]$. All of our identification arguments below still hold by replacing the integral region $[\underline{p}, \bar{p}]$ with $\mathcal{S}_{p(X)}$.

⁷In 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 , even if it is continuous, is typically only identified at a set of isolated points.

be rewritten as

$$\begin{aligned} \int_{\underline{p}}^{\bar{p}} \int_0^{p'} Q(\tau) d\tau \cdot \pi(p', p; \beta) dp' &= \int_0^1 Q(\tau) \cdot \int_{\underline{p}}^{\bar{p}} \mathbb{1}(\tau \leq p') \cdot \pi(p', p; \beta) dp' d\tau \\ &= - \int_{\underline{p}}^{\bar{p}} Q(\tau) \cdot \Pi(\tau, p; \beta) d\tau, \end{aligned}$$

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

$$Q(p) - \beta \int_{\underline{p}}^{\bar{p}} Q(\tau) \cdot \Pi(\tau, p; \beta) d\tau = z(p)^\top \cdot \theta, \quad \forall p \in [\underline{p}, \bar{p}]. \quad (12)$$

The solution of this equation requires the following assumption:

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

Assumption E ensures that the mapping in eq. (12) is a contraction, so that the solution is unique. Note that this assumption is not a model restriction, but an identification condition, involving both structural primitives as well as variations of observed state variables. Though high-level, it is testable in principle as it depends only on data.

Lemma 4. Suppose assumptions A to E hold. Then, $Q(\cdot)$ is point identified on $[\underline{p}, \bar{p}]$ up to the finite dimensional parameter θ :

$$Q(p) = \left\{ z(p) - \beta \int_{\underline{p}}^{\bar{p}} R(p', p; \beta) \cdot z(p') dp' \right\}^\top \cdot \theta, \quad \forall p \in [\underline{p}, \bar{p}] \quad (13)$$

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 (13) is proportional to θ , which is due to the linearity of the FIE system. Therefore, (13) 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 $\ell = 1, \dots, k_\theta$, let $b_\ell^*(\cdot)$ be the (unique) solution to the following equation

$$b_\ell(p) + \beta \int_{\underline{p}}^{\bar{p}} \int_{\underline{p}}^{p'} b_\ell(\tau) d\tau \cdot \pi(p', p; \beta) dp' = z_\ell(p). \quad (14)$$

By a similar argument to Lemma 4, we have

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

as the unique solution to (14). Let $\mathcal{B}(\cdot) \equiv (b_1^*(\cdot), \dots, b_{k_\theta}^*(\cdot))^\top$ be the sequence of solutions supported on $[\underline{p}, \bar{p}]$. Thus, the solution in eq. (13) can be written as

$$Q(p) = \mathcal{B}(p)^\top \cdot \theta, \quad \forall p \in [\underline{p}, \bar{p}]. \quad (15)$$

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:*

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

where

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

or alternatively, $m(X) = \mathcal{B}(p(X))$.

Note that $\mathbb{P}(Y = 1|X) = F_\eta(Q(p(X)))$. Then, Theorem 1 obtains by combining eqs. (10) and (15). By a similar argument as in the static semiparametric binary choice literature ((e.g. Horowitz, 2009)), the index parameter θ is identified up to location and scale. For notational simplicity, hereafter we assume the state vector X does not include a constant term in the semiparametric setting.⁸ Moreover, we will introduce a scale normalization on θ which is also standard in the literature.

Assumption F. *Let $m(X)$ be continuously distributed with a joint probability density function, denoted by $f_m(\cdot)$. Let further the matrix $\mathbb{E}[m(X)m(X)^\top]$ be invertible.*

In Assumption F, the first half condition requires at least one argument of X to be continuously distributed, and the second half is a testable rank condition.

⁸Any constant term in the utility function will be absorbed by the error term since the distribution of the latter is left unspecified.

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

The scale normalization in Assumption G is commonplace.

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

Given Theorem 1, the proof of Theorem 2 follows e.g. Ichimura (1993).

For estimation, the structure of the DDC model as given in Theorem 1 is the same as a binary choice model with unknown distribution of the error term, thus making available the wide array of semiparametric estimators for this model which have been proposed in the econometrics literature. We utilize the Powell, Stock, and Stoker (1989) estimator, as it enjoys the important advantage of providing a closed-form (non-iterative) estimator for θ . The PSS estimator is defined as: Specifically, our closed-form estimator is:

$$\hat{\theta} = -\frac{2}{T(T-1)} \sum_{t=1}^T \sum_{s \neq t} \left[\frac{1}{h_\theta^{k_\theta+1}} \times \nabla K_\theta \left(\frac{\hat{m}(X_t) - \hat{m}(X_s)}{h_\theta} \right) \times Y_s \right]. \quad (16)$$

where $\hat{m}(X)$ denotes a nonparametric estimate of the $m(X)$ functions, and K_θ and h_θ denote, respectively, a kernel function and bandwidth. The definition of these objects, and full details of the estimation procedure, are provided in Appendix A. There, we also derive the asymptotic normality for this estimator of θ , which justifies the use of bootstrap in computing the standard errors in our empirical work below.

4.1. Monte Carlo. For the remainder of this section, we provide some Monte Carlo results on the performance of our estimator. 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} 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 $\theta_2 = -2$.

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. Estimation takes on average 6, 12, and 25 seconds respectively for each replication on a 4Ghz i7 computer.

Because we cannot estimate the constant θ_0 in the semiparametric framework, then we treat θ_0 as a nuisance parameter. Let $\theta = (\theta_1, \theta_2)^\top$. 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. ESTIMATES OF OPTIMAL STOPPING MODEL FOR NYC TAXI DRIVERS

In this section we return to the optimal stopping model of labor supply for taxicab drivers, introduced in Section 2 above. We first describe the data.

5.1. 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 10,000 observations, or about 0.1% 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

TABLE 2. Taxi Trip and Fare Summary Statistics

Trips/Shifts	Variable	Obs.	10%ile	Mean	90%ile	S.D.
Trip Statistics	Trip Revenue (\$)	10,000	5.88	12.47	21.80	8.86
	Trip Time (min.)	10,000	4.00	11.72	22.0	8.17
Shift Statistics	Shift Revenue (\$)	381	210.49	327.34	446.62	105.49
	Shift Time (min.)	381	368.00	559.39	730.17	204.90

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.

drop-offs, trip distance, and trip time for 10,000 individual taxi rides. Table 2 provides summary statistics.

It is worth noting that this data is essentially a complete record of all trips operated by licensed New York medallion taxis. While recent work such as Farber (2014) makes use of this data, the earlier research (including the work devoted to explicitly measuring labor supply elasticities) employ much smaller and less reliable taxi trip data. While the debate about model specification and behavioral biases may persist, we can be sure that data quality is no longer a concern.

5.2. Reduced-form results. Table 3 shows the results of elasticity regression of the form of Camerer, Babcock, Loewenstein, and Thaler (1997) and further analyzed (and critiqued) in Farber (2005). Each specification regresses $\log(hours)$ on $\log(wage)$, where “hours” refers to the cumulative time worked by a driver upon quitting for the day, and “wage” refers to the average hourly earnings achieved through the day. In these regressions, we derive a measure of labor supply elasticity as the parameter on $\log(wage)$. Specification (1) and (2) implement a simple OLS regression. As both of the above papers note, since wage is defined as cumulative revenue divided by cumulative hours worked, there will be a division bias by construction, as the variable *hours* appears in both the left- and right-hand sides of the regression. In the next two specifications, (3) and (4) we implement an instrumental variables regression (again mirroring the work of Camerer, Babcock, Loewenstein, and Thaler (1997)) which uses the 25th, 50th and 75th percentiles of average wages across all drivers in each day, along with day-of-week indicators, as instruments for driver wages. Specifications (2) and (4) also control for driver-specific fixed effects.

TABLE 3. Reduced Form Elasticity Regressions

	(1)	(2)	(3)	(4)
	OLS	OLS	IV	IV
Log Wage	-0.100** (0.011)	-0.098** (0.007)	0.499** (0.019)	0.486* (.247)
Weekday Dummy	-0.100** (0.002)	-0.099** (0.001)	-0.065** (0.002)	-0.058** (0.019)
Day shift	-0.066 (0.003)	-0.214 (0.006)	0.117** (0.004)	-0.049 (0.327)
Driver FE	x	✓	x	✓
<i>N</i>	623,482	623,482	623,482	623,482

Taxi trip and fare data come from New York Taxi and Limousine Commission (TLC) and refer to February 2012 data. Data record the final cumulative hours and average wage earned as of the last trip of each driver-shift. The IV specifications use the following instruments for wage: the 25th, 50th and 75th percentile across all driver wages each day, as well as a dummy for day-of-week. Standard Errors clustered at the driver-shift level.

We see from these regressions that, much like in the previous literature, the OLS specifications yield negative labor supply elasticities. By instrumenting for wages, however, the elasticities become positive, and more consistent with the standard models of labor supply.

5.3. Choice Probabilities. As Farber (2005) cautions, a conventional wage regression is somewhat inappropriate in settings where marginal wages are variable. The paper then formulates a hazard model of driver’s stopping decisions on the basis of cumulative income and hours. We can take advantage of the granularity of our data to show the influence of hours and earnings non-parametrically in the form of empirical choice probabilities. Table 4 provides a set of quitting probabilities by cumulative hours worked and cumulative earnings over a shift. This table reveals a broadly increasing pattern of increasing quit probabilities by both hour and income.⁹ These patterns are similar to those revealed by the hazard model estimates of Farber (2005).

5.4. Dynamic model results. The estimation results are presented in Table 5. 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

⁹The pattern comes with a caveat that the more extreme off-diagonal cells have relatively few observations despite the abundance of data, as those depicted in gray shading.

TABLE 4. Choice Probabilities by Cumulative Earnings and Hours

Cum. Hours Worked	Cumulative Income Earned									
	\$0	\$50	\$100	\$150	\$200	\$250	\$300	\$350	\$400	\$450
0	0.005	0.009	0.032	0.062	0.130	0.125	0.250	1.000	0.500	0.400
1	0.006	0.004	0.009	0.024	0.026	0.035	0.050	0.250	0.000	0.000
2	0.016	0.006	0.006	0.013	0.028	0.060	0.023	0.000	0.000	0.000
3	0.033	0.012	0.009	0.009	0.017	0.035	0.078	0.089	0.000	0.125
4	0.044	0.020	0.014	0.014	0.016	0.027	0.046	0.039	0.026	0.118
5	0.075	0.030	0.022	0.021	0.025	0.031	0.044	0.060	0.054	0.027
6	0.153	0.041	0.036	0.034	0.044	0.052	0.059	0.082	0.097	0.054
7	0.232	0.076	0.061	0.060	0.068	0.081	0.083	0.090	0.106	0.104
8	0.350	0.104	0.094	0.097	0.112	0.124	0.119	0.106	0.111	0.151
9	0.304	0.141	0.136	0.133	0.161	0.179	0.177	0.154	0.145	0.155
10	0.250	0.213	0.173	0.178	0.176	0.208	0.223	0.227	0.232	0.254
11	0.400	0.389	0.145	0.162	0.151	0.165	0.184	0.201	0.242	0.305

Data from TLC Data, February 2012. Each cell shows the fraction of time drivers in each category (of cumulative hours worked and income earned) quit for the day. Each category reflects values at or above the category label. For example, income category \$100 is read as “\$100-199.99” and hour category 1 is read as “1 hour 0 minutes - 1 hour 59 minutes”. Gray entries denote cells with fewer than 100 observations.

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.

Given these parameter estimates, in Figure 1 we graph the implied quantile function for the difference in utility shocks $\eta \equiv \epsilon_1 - \epsilon_0$.¹⁰ 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)

¹⁰This 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.

positive values, with magnitudes in the hundreds; this may imply that there is a large fixed positive component to the terminal utility from quitting.¹¹

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 5. Parameter Estimates

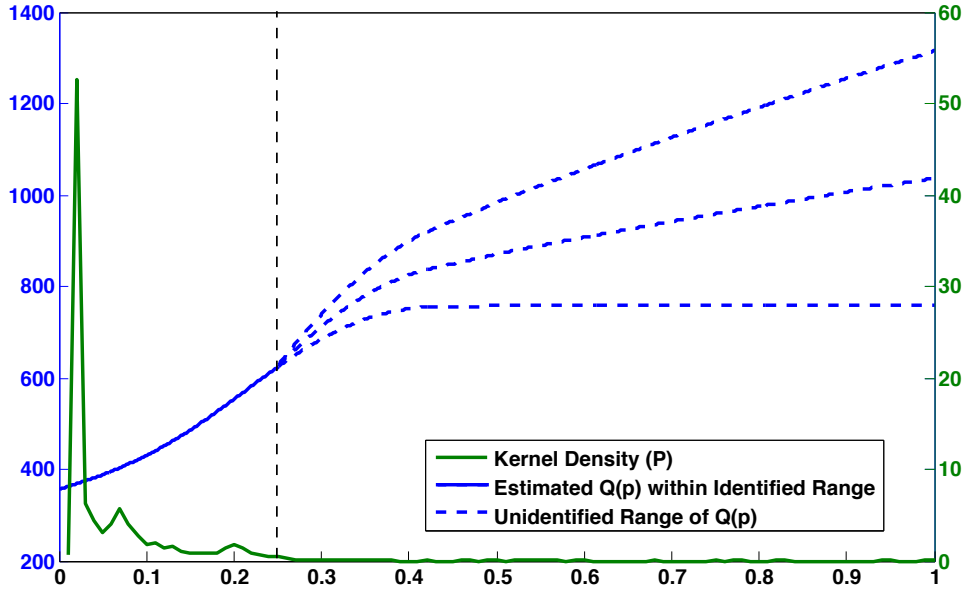
Parameter	Description	Estimate	Std. Error
θ_u	Earnings (upon quitting)	0.9907	0.0118
$\theta_{c,01}$	Cumul. hours (while working)	-0.1359	0.0759
$\theta_{c,02}$	Cumul. hours squared (while working)	-0.0004	0.0002

Note: Standard errors are computed by first sampling, with replacement, from each driver-shift (on average there are roughly 24 observations per driver-shift) to generate 200 resamples of approximately identical size to our original sample. We re-estimate the model for each resample and report the standard deviation of estimates.

While the empirical specification in Table 5 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 $X = (\text{hours worked}, \text{income})$. Hence, in Figure 2, we plot the estimated $m(X)'\theta$ function at different levels of hours-worked. 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.

¹¹Note that in estimating the quantile function, we have not fixed the scale and location for the utility shock difference η ; we have this flexibility because we imposed a scale normalization on the parameter vector β . In contrast, parametric estimation approaches for DDC models typically do not impose normalization on the parameters, but implicitly the researcher must set the scale and location for the utility shocks (a common assumption is zero mean and unit variance).


FIGURE 1. Estimated Quantile Function



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.

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 and, depending on how fast income accumulates during a shift, our optimal stopping rules can imply both negative or positive income elasticities. And indeed, our data shows substantial intra-day variability in the rate of income accumulation, which perhaps explains the instability of the wage coefficients in the reduced-form regressions, given in Table 3 above. At the same time, a point of emphasis here is that once taxi drivers’ quitting decision are (correctly, we argued) modeled in a dynamic optimal stopping framework, there is no need to add non-standard

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 $X = (\text{hours worked}, \text{income})$. 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.

behavioral parameters to the model to explain their quitting behavior – it emerges as an outcome along the optimal dynamic decision-making path.

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 apply this estimator to a new and comprehensive dataset of New York City taxi drivers (the largest single taxicab market in the United States). We take a new approach to a long-running question of drivers' wage elasticities by modelling taxicab drivers' labor supply decisions as emerging from a dynamic *optimal stopping* problem. Our results reconcile debates in the previous literature to a certain extent. Estimates of drivers' optimal stopping rules show that, holding hours worked, drivers are more likely to quit at higher levels of cumulative income. In reduced-form, such quitting rules can generate both "positive" and "negative" wage elasticities, depending on the specifics of the stochastic fare process. More broadly, these findings suggest that once the inherent dynamic optimization aspect of taxicab drivers' labor supply decisions are accounted for, there is no need to add non-standard behavioral parameters to the model to explain their quitting behavior.

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. (2010): "Another Look at the Identification of Dynamic Discrete Decision Processes: An Application to Retirement Behavior," *Journal of Business and Economic Statistics*, 28(2), 1795–1843.

- 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.
- FANG, H., AND Y. WANG (2015): "Estimating Dynamic Discrete Choice Models with Hyperbolic Discounting, with an Application to Mammography Decisions," *International Economic Review*, 56(2), 565–596.
- 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.

APPENDIX A. SEMIPARAMETRIC ESTIMATION: FULL DETAILS

In this section, we provide full details and asymptotic results for our closed-form estimator for θ as defined in Eq. (16) of the main text. 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 $\{(Y_t, X_t^\top)^\top : t = 1, \dots, T\}$ 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

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

As is standard, we choose an optimal bandwidth, i.e., $h_p = 1.06 \times \hat{\sigma}(X) \times T^{-\frac{1}{2\iota+8}}$, 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 $[p, \bar{p}]$ 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 $s = 1, \dots, S_T$, where $S_T = T - \ell_T$ for some integer ℓ_T satisfying $\ell_T \rightarrow +\infty$ and $S_T \rightarrow +\infty$ as $T \rightarrow +\infty$, 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 $s = 1, \dots, T$, 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 $[p, \bar{p}]$. First, let

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

According to Guerre, Perrigne, and Vuong (2000, Theorem 2), we choose an oversmoothing bandwidth h_z , since $p(X)$ is nonparametrically estimated. Specifically, $h_z = 1.06 \times \hat{\sigma}(p(X)) \times T^{-\frac{1}{2\iota+3}}$.

To estimate $b_\ell^*(\cdot)$ on the support $[\underline{p}, \bar{p}]$, we note that eq. (14) can be rewritten as

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

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, $h_z = 1.06 \times \hat{\sigma}(p(X)) \times T^{-\frac{1}{2\ell+3}}$ is chosen sub-optimally. A numerical solution of \hat{b}_ℓ^* can obtain using the iteration method: Let $\hat{b}_\ell^{[0]} = \hat{z}_\ell^\top(p)$. Then we set

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

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

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

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

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

Following the standard kernel regression literature, we can show our PSS-based estimator, $\hat{\theta}$ (defined in Eq. (16)) is consistent given that $\sup_{x \in \mathcal{S}_X} |\hat{m}(x) - m(x)| = o_p(h_\theta)$, $h_\theta \rightarrow 0$ and $Th^{k_\theta+1} \rightarrow \infty$ as $T \rightarrow \infty$.

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 (16). 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 $k_\theta + 2 < \gamma < k_\theta + 3 + \mathbb{1}(k_\theta \text{ is even})$.

Assumption I. The support of the kernel function K_θ is a convex subset of \mathbb{R}^{k_θ} 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 .¹² The restriction on the bandwidth Assumption H 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 $\hat{m}(\cdot)$ to converge to $m(\cdot)$ faster than $T^{-\frac{1}{4}}$. 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 (16). Assumption J is standard in the literature for the regular convergence of finite-dimensional parameters in semiparametric models (e.g. Ai and Chen, 2003), except for the polynomial terms of h_θ in the $o(\cdot)$ or $o_p(\cdot)$ which arises due to the average derivate estimator in the second stage.

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

¹²Note that we implicitly assume that Assumptions 1 – 3 in PSS hold, which impose smoothness conditions on f_m and $\mathbb{P}(Y_t = 1|m(X_t) = m)$ as well as other regularity conditions.

In the above theorem, recall $\mathbb{P}(Y = 1|X) = F_\eta(\eta^*(X))$ and $\eta^*(X) = m(X)^\top \cdot \theta$ by Theorem 1. Our estimator $\hat{\theta}$ (as defined in Eq. (16)) 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. (13). Namely, let

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

Because of the \sqrt{T} -consistency of $\hat{\theta}^*$, the estimator $\hat{Q}(p)$ is asymptotically equivalent to $\hat{z}^\top(p) \times \theta$, which converges at a nonparametric rate.¹³ Given the asymptotic normality established in this section, bootstrap inference is valid and we will use it for constructing standard errors in our empirical application below.

APPENDIX B. PROOFS

B.1. Proof of Lemma 1.

Proof. First, note that the resolvent kernel R^* is well-defined. This is because $\beta^{s-1} f_{X^{[s]}|X}(x'|x) \rightarrow 0$ as $s \rightarrow +\infty$. 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. (5), Hence, it suffices to show the uniqueness of the solution. Eq. (5) can be rewritten as

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

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 (5). Then, we have

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

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

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

¹³The asymptotic properties of $\hat{z}^\top(p)$ can be established by following Guerre, Perrigne, and Vuong (2000), who use nonparametrically estimated pseudo private values to construct a kernel estimator for the density function of bidders' private values in an independent private value auction model.

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:

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

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

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

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

B.2. Proof of Lemma 4.

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

B.3. Proof of Theorem 3. The estimator is defined in (16). For the consistency of $\hat{\theta}$, we need $h_\theta \rightarrow 0$, $Th_\theta^{k_\theta+1} \rightarrow \infty$ and $\mathbb{E}|\hat{m}(X) - m(X)| = o(h_\theta)$ as $T \rightarrow \infty$. Note that the last condition ensures the estimation error in \hat{m} is negligible.

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 $Th_\theta^{k_\theta+2} \rightarrow \infty$, while the bias term has the order h_θ^P . Therefore, if $T^{1/2}h_\theta^P \rightarrow 0$, 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

$$\begin{aligned} \hat{\theta} &= \tilde{\theta} - \frac{2}{T(T-1)} \sum_{t=1}^T \sum_{s \neq t} \left[\frac{1}{h_\theta^{k_\theta+2}} \nabla^2 K_\theta \left(\frac{m(X_t) - m(X_s)}{h_\theta} \right) \times Y_s \times (\hat{m}(X_t) - m(X_t)) \right] \\ &\quad + \frac{2}{T(T-1)} \sum_{t=1}^T \sum_{s \neq t} \left[\frac{1}{h_\theta^{k_\theta+2}} \nabla^2 K_\theta \left(\frac{m(X_t) - m(X_s)}{h_\theta} \right) \times Y_s \times (\hat{m}(X_s) - m(X_s)) \right] \\ &\quad + O_p(h_\theta^{-3} \cdot \mathbb{E}\|\hat{m}(X) - m(X)\|^2) \equiv \tilde{\theta} + \mathbb{A}_1 + \mathbb{A}_2 + \mathbb{B} \quad (17) \end{aligned}$$

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

$$h_\theta^{-3} \times \mathbb{E} \|\hat{m}(X) - m(X)\|^2 = h_\theta^{-3} \times o_p(T^{-1/2} h_\theta^3) = o_p(T^{-1/2}). \quad (18)$$

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

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

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

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

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

Because

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

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

For property (b), note that

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

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,

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

Note that

$$\begin{aligned} &\text{Cov} \left(\phi_{T,2,1} [\hat{m}(X_1) - \mathbb{E}[\hat{m}(X_1)|X_1]], \phi_{T,4,3} [\hat{m}(X_3) - \mathbb{E}[\hat{m}(X_3)|X_3]] \right) \\ &= \mathbb{E} \{ \phi_{T,2,1} \phi_{T,4,3} [\hat{m}(X_1) - \mathbb{E}[\hat{m}(X_1)|X_1]] \times [\hat{m}(X_3) - \mathbb{E}[\hat{m}(X_3)|X_3]] \} \\ &- \mathbb{E} \{ \phi_{T,2,1} [\hat{m}(X_1) - \mathbb{E}[\hat{m}(X_1)|X_1]] \} \times \mathbb{E} \{ \phi_{T,4,3} [\hat{m}(X_3) - \mathbb{E}[\hat{m}(X_3)|X_3]] \}. \end{aligned}$$

By Assumption J(iii),

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

Furthermore, by the law of iterated expectation (conditioning on the sigma algebra: $\mathcal{F}_2, \mathcal{F}_4, \mathcal{F}_5, \dots, n$),

$$\begin{aligned} &\mathbb{E} \{ \phi_{T,2,1} \phi_{T,4,3} [\hat{m}(X_1) - \mathbb{E}[\hat{m}(X_1)|X_1]] \times [\hat{m}(X_3) - \mathbb{E}[\hat{m}(X_3)|X_3]] \} \\ &= O_p(h_\theta^{-4}) \times o_p(T^{-1/2} h_\theta^2) \times o_p(T^{-1/2} h_\theta^2) \\ &= o_p(T^{-1}), \end{aligned}$$

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

$$\begin{aligned} & \frac{1}{T} \text{Cov} \left(\phi_{T,2,1} [\hat{m}(X_1) - \mathbb{E}[\hat{m}(X_1)|X_1]], \phi_{T,3,1} [\hat{m}(X_1) - \mathbb{E}[\hat{m}(X_1)|X_1]] \right) \\ &= \frac{1}{T} \mathbb{E} \left\{ \phi_{T,2,1} \phi_{T,3,1} [\hat{m}(X_1) - \mathbb{E}[\hat{m}(X_1)|X_1]]^2 \right\} \\ &= O(T^{-1}h_\theta^{-4}) \times \mathbb{E} \{ \hat{m}(X) - \mathbb{E}[\hat{m}(X)|X] \}^2. \end{aligned}$$

Then a sufficient condition for property (b) is

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

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).