

DYNAMIC LABOR SUPPLY OF TAXICAB DRIVERS: A SEMIPARAMETRIC OPTIMAL STOPPING MODEL

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

Rewriting in progress: please excuse the débris!

ABSTRACT. We estimate an optimal stopping model for taxicab drivers' labor supply decisions, using a large sample of shifts for drivers of New York City taxicabs. Our results show that both "behavioral" and "neoclassical" wage responses are present in the data, with the behavioral income-targeting story explaining shorter shifts, and the standard neoclassical model explaining longer shifts. Hence these findings partially reconcile the divergent reduced-form results in the existing literature. A methodological contribution of this paper is to develop a new closed-form estimator for dynamic discrete choice models in a semiparametric setting, in which the distribution of utility shocks is left unspecified.

Keywords: Optimal stopping, Taxicab industry, Labor supply, Dynamic discrete choice model, Closed form estimator, Semiparametric estimation

JEL: C14, D91, C41, L91

Date: Friday 6th April, 2018.

^{*}We thank Hassan Afrouzi, Saroj Bhattarai, Ernan Haruvy, Kirill Pogorelskiy, Eduardo Souza-Rodrigues, Sorawoot (Tang) Srisuma, and seminar participants at Carnegie-Mellon, Arizona, UC-Irvine, Queensland, Melbourne, NYU, Rice, UCLA-Anderson, UT-Austin, UT-Dallas, Washington (Seattle), Virginia, Wisconsin, Colorado (Boulder), 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 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 both “behavioral” and “neoclassical” wage responses are present in the data, with the behavioral income-targeting story explaining shorter shifts, and the standard neoclassical model explaining longer shifts. Since these results are consistent with what in a static labor supply context would be called negative or positive wage elasticities, they may offer a partial reconciliation of the divergent reduced-form results in the existing literature. More broadly, these findings highlight how our (relatively simple) dynamic framework is rich enough to generate behavior which resembles both negative and positive wage elasticities from a static point of view; 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 behavior – it emerges as an outcome along the optimal dynamic decision-making path.¹

This paper also makes an important methodological contribution by introducing a new estimator for a large class of dynamic discrete choice (DDC) models (for which the optimal

¹For instance, Crawford and Meng (2011) and Farber (2014) build reference dependence explicitly into the utility specifications used in their analyses.

stopping model of taxicab drivers is one example) which is (i) *semiparametric* and (ii) can be computed in *closed form*. First, our semiparametric approach eschews parametric assumptions on the distribution of the error terms in the discrete choice model; in contrast, many applications of dynamic discrete choice models assume that these errors follow the extreme-value distribution, leading to logit choice probabilities. Second, our closed form estimator for the structural model parameters is non-iterative, which sidesteps computational pitfalls associated with most existing estimators for these models, which typically involve time-consuming iterative nonlinear optimization procedures which can be sensitive to starting values and convergence criteria.

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 apply Powell, Stock and Stoker's (1989,PSS) classic kernel-based estimator for *static* semiparametric binary choice models.

In section X we do [...]

2. LABOR SUPPLY FOR NEW YORK CITY TAXICAB DRIVERS

A growing literature has arisen aiming 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.

Our dynamic modeling 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. Thakral and Tô (2017) also take up the question of whether there are behavioral biases in drivers stopping decisions, showing that more recent income is a stronger determinant of quitting than income earned earlier in a shift. Our approach will somewhat reconcile this tension by treating work hours as a state variable in addition to just income, and by modeling drivers' stopping rules as stemming from a combination of cumulated hours and income.

We estimate a dynamic optimal stopping 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 (2016) [FLS], in which taxi drivers 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 estimate a dynamic optimal stopping model for taxicab drivers, which is new in the literature.

Taxi drivers are assumed to have costs of effort that are increasing in hours-worked each day. Each period is a fare. After each fare, 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.

Specifically, we consider a single-agent infinite-horizon binary decision problem. After each fare t , the agent observes state variables $X_t \in \mathcal{S}_x \subseteq \mathbb{R}^k$, experiences random utility shocks ϵ_t ,

and chooses a binary decision $Y_t \in \{0, 1\}$ to maximize her current and future expected payoffs. For the taxicab drivers' optimal stopping problem, the state variables X_t consist of (s_t, h_t) , the cumulative earnings and hours worked after t fares. The agent maximizes the expected discounted sum of the per-period utilities:

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

subject to $f_{X_{t+1}, \epsilon_{t+1} | X_t, \epsilon_t, Y_t}$ the Markov law of motion for the state variables (X, ϵ) . $\beta \in (0, 1)$ denotes the discount factor, and $u(y, X, \epsilon)$ denote the single-period payoff functions.

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:

$$u_i(s_{it}, h_{it}, y_{it}; \theta, X_t) = \begin{cases} u_{i1}(s_{it}) + \varepsilon_i(1) & \text{if } y_{it} = 1 \\ u_{i0}(h_{it}) + \varepsilon_i(0) & \text{if } y_{it} = 0 \end{cases} . \quad (2)$$

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 ($y_{it} = 1$). The terminal utility from ending the shift is given in the upper prong of the utility specification above. In this terminal utility, the term $u_{i1}(s_{it})$ captures the utility from earnings (s_{it}) enjoyed by the taxi driver after ending his shift. When a driver continues to drive ($y_{it} = 0$), as in the lower prong of the utility specification, he experiences (dis-)utility $u_{i0}(h_{it})$ which depends on h_{it} , the cumulative hours worked so far in this shift.

2.1. Data and reduced-form results. We start by introducing the dataset and presenting some reduced-form results. 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 trips, or about 0.1% of the data. Data include the exact time and date of pickup and drop-offs, trip distance, and trip time. Table 1 provides summary statistics.

This data set represents 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 there is continued debate about model specification and the presence of behavioral biases, the TLC data obviates most lingering worries about sample size and measurement error.²

Table 1 contains summary statistics for a random sample of all evening shifts in February 2012, as well as random subsamples broken down by weekday vs. weekend and rainy vs. non-rainy days. The top part of table shows that there is little heterogeneity across subsamples at the trip-level. The distribution and average of trip revenue and trip duration similar across weekdays/weekends, and across days with rain/no-rain.

In contrast, the bottom part of table shows interesting heterogeneity at the shift-level. We see that the difference between weekdays and weekends is substantially more prominent than between rainy and non-rainy days. Regardless of rain, the distributions of both shift revenue and shift duration on weekends are higher (ie. stochastically dominate) the distributions on weekdays. On average rainy (non-rainy) evenings, shift revenue is \$228.68 (\$239.23) on weekdays, increasing to \$331.15 (\$310.15) on weekends. Similarly, the average shift duration is 354 (417) mins. on weekdays, increasing to 549 (504) mins. on weekends. As we see from the top panel, this higher revenue and shift duration on weekends does not arise from longer trips or higher revenue per trip; rather, as the bottom of Table 1 shows, drivers typically make 25% more trips per shift on weekends (28 vs. 20 on rainy evenings, and 26 vs. 19 on non-rainy evenings). These big differences between weekday and weekend shifts may suggest that drivers' preferences and motivations for driving may be different across these settings. In our empirical work, we will estimate the model separately for the different subsamples.

Table 2 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

²Other recent research uses the New York TLC data set to study cab driver behavior. Buchholz (2017) studies how pricing regulations and location-based competition among taxi drivers leads to spatial equilibrium patterns of supply and demand. This paper focuses on the intensive-margin of how taxis choose where to search within the city, while taking the extensive margin (how many cabs supply labor at any one time) to be exogenous. Haggag, McManus, and Paci (2017) investigate whether drivers learn to accrue income faster over time, finding that tenure induces modest gains in drivers' shift earnings.

TABLE 1. Taxi Trip and Fare Summary Statistics

Category	Variable	Data Sample	Obs.	10%ile	Mean	90%ile	S.D.
Trip Statistics	Trip Revenue (\$)	Overall	10,007	6.01	12.97	22.42	9.29
		M-Th, No Rain	10,011	6.13	12.86	22.00	8.85
		M-Th, Rain	10,003	6.40	13.02	22.40	9.60
		F-Su, No Rain	10,003	6.01	12.50	20.90	8.64
		F-Su, Rain	10,025	6.10	12.91	21.86	9.19
	Trip Duration (min.)	Overall	10,007	3.98	11.42	21.32	7.72
		M-Th, No Rain	10,011	4.00	11.18	21.00	7.43
		M-Th, Rain	10,003	4.00	12.15	22.55	8.66
		F-Su, No Rain	10,003	4.00	11.52	21.00	7.52
		F-Su, Rain	10,025	4.00	11.61	21.35	7.66
Shift Statistics	Shift Revenue (\$)	Overall	402	132.80	274.64	407.11	107.34
		M-Th, No Rain	529	104.84	239.23	341.01	93.70
		M-Th, Rain	561	152.30	228.68	287.28	73.23
		F-Su, No Rain	385	151.80	310.15	438.63	115.09
		F-Su, Rain	284	230.48	331.15	425.72	80.67
	Shift Duration (min.)	Overall	402	237.00	463.97	651.53	173.18
		M-Th, No Rain	529	183.48	417.34	608.62	163.75
		M-Th, Rain	561	257.00	353.98	413.98	114.75
		F-Su, No Rain	385	274.00	503.80	669.83	180.66
		F-Su, Rain	284	412.08	549.05	666.00	118.81
	Trips per shift	Overall	-	8	21.9	34	9.6
		M-Th, No Rain	-	12	18.9	37	8.5
		M-Th, Rain	-	10	19.5	29	7.2
		F-Su, No Rain	-	12	25.8	37	9.8
		F-Su, Rain	-	16	28.0	37	8.5

Taxi trip and fare data come from New York Taxi and Limousine Commission (TLC) and refer to February 2012 data. Our sample extracts trip data from the “evening shift”, or shifts that begin after 4pm and end before 4am. Thus, Friday evening is treated as a weekend shift. 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. Samples are chosen by collecting enough driver shift series (at random) within the overall sample of all February 2012 data, so that at least 10,000 observations are collected. “Rain” means at least 1/10 of an inch of rain has been recorded in Central Park during the day in which the shift begins. The “overall” sample consists of a similar collection procedure, weighed so that the rain/weekday observations are proportional to their frequency of occurrence across the entire February 2012 period.

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

TABLE 2. Reduced Form Elasticity Regressions

<i>Dependent variable: Log shift duration</i>				
	(1)	(2)	(3)	(4)
	OLS	OLS	IV	IV
Log Wage	-0.106** (0.008)	-1.160** (0.007)	-0.485** (0.026)	-0.135** (.023)
Weekday Dummy	-0.121** (0.002)	-0.115** (0.001)	-0.102** (0.002)	-0.079** (0.002)
Rain > 1/10"	0.093** (0.002)	0.090** (0.002)	0.065** (0.002)	0.041** (0.002)
Day shift	-0.127** (0.002)	-0.355 (0.006)	-0.045** (0.004)	-0.265** (0.008)
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.

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. Specifications (3) and (4) utilize instrumental variables to adjust for endogeneity arising from “denominator bias” in the wages. Specifications (2) and (4) also control for driver-specific fixed effects.

We see from these regressions that, much like in the previous literature, the OLS specifications yield negative labor supply elasticities. Instrumenting for wages yields higher elasticities comparing (3) to (1), and less negative elasticities comparing (4) to (2). In all cases, negative elasticities conflict with the standard view that labor supply curves slope upwards, absent sufficiently strong income effects.

2.2. Nonparametric Choice Probabilities. As Farber (2005) cautions, conventional (static) wage regressions, as in Table 2, are somewhat inappropriate in settings like this where marginal wages are variable. And indeed, the differences between the OLS and IV regressions suggest that the data do not speak with one voice regarding the wage elasticities. As a further diagnostic

TABLE 3. 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.

step, we exploit the granularity of our data to show the influence of hours and earnings non-parametrically in the form of empirical choice probabilities. Table 3 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, although there are some interesting regions of quitting probabilities decreasing in income (eg. at \$300-\$400 and 8-9 hours of work).³ These patterns are similar to those revealed by the hazard model estimates of Farber (2005).

These nonparametric choice probabilities also highlight the inherent rich variation in hours worked, earnings, and quitting propensities, which will identify the parameters in our dynamic optimal stopping model. Next, we describe the novel semiparametric modeling framework and a new closed-form estimator which we propose for these models. Readers who wish to skip these methodological details may proceed directly to Section 4.

³The 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.

3. SEMIPARAMETRIC DYNAMIC DISCRETE CHOICE: A NEW CLOSED FORM ESTIMATOR

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. Since this estimator applies to a broad class of binary dynamic discrete choice models, we will discuss it in some degree of generality, instead of referring specifically to the taxicab labor supply model.

Consider the single-agent dynamic optimization problem (1) from the previous section. Following a majority of the existing applications of dynamic discrete-choice models, we assume that the single-period payoff functions are linear functions of parameters θ :

$$u_t(Y_t, X_t, \epsilon_t) = \begin{cases} W_1(X_t)^\top \theta_1 + \epsilon_{1t}, & \text{if } Y_t = 1; \\ W_0(X_t)^\top \theta_0 + \epsilon_{0t}, & \text{if } Y_t = 0. \end{cases} \quad (3)$$

$W_0(X_t) \in \mathbb{R}^{k_0}$ (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.⁴ 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 use the shorthand W_d for $W_d(X)$ ($d = 0, 1$) and suppress the explicit dependence upon the state variables X when possible. The structural parameters which are of interest are $\theta_d \in \mathbb{R}^{k_d}$, for $d \in \{0, 1\}$.⁵

A novel feature of our analysis, relative to much of the existing literature in dynamic discrete-choice models, is that 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 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.⁶

⁴The utility of action 0 is not normalized to be zero for reasons discussed in Norets and Tang (2014).

⁵The discount factor β is assumed to be known for purposes of estimation, which 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 .

⁶In a static setting, however, 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. McFadden and Train (2000) show such properties for the mixed logit model.

To our knowledge, only a handful of papers consider estimation of dynamic models in which the error distribution is left unspecified. Aguirregabiria (2010) shows the joint nonparametric identification of utilities and the shock distribution in a class of finite-horizon dynamic binary choice models, which may not apply to infinite-horizon models as considered in this paper. 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. In contrast, with continuous state variables, we show how identification and estimation of the error distribution is possible.

Chen (2016) 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 very general dynamic models in which agents can make both discrete and continuous choices, and shows, under exclusion restrictions, the nonparametric identification of both the per-period utility functions as well as the error distribution. Our approach eschews exclusion restrictions; rather, we exploit the optimality conditions to derive a novel recursive characterization 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.

3.1. The value function: two different characterizations. Let $V(X, \epsilon)$ be the value function given X and ϵ . Assuming stationarity, we drop the t subscripts and use primes ($'$) to denote next period values. The Bellman equation is

$$V(X, \epsilon) = \max_{y \in \{0,1\}} \{ [u(y, X, \epsilon) + \beta \mathbb{E}[V(X', \epsilon') | X, \epsilon, Y = y]] \}, \quad (4)$$

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

Assumption B (Conditional Independence). The law of motion satisfies: $F_{X', \epsilon' | X, \epsilon, Y} = F_{\epsilon'} \times F_{X' | X, Y}$. Moreover, $F_{\epsilon'} = F_{\epsilon}$.

Assumption A is a weak assumption which holds, for instance, when $W_d(\cdot)$ are bounded functions. Assumption B establishes that the shocks ϵ are fully independent of the observed

state variables X .⁷ Under this assumption, the value function can be written as

$$V(X, \epsilon) = \max \left\{ W_1^\top \theta_1 + \epsilon_1 + \beta \mathbb{E}[V(X', \epsilon') | X, Y = 1], W_0^\top \theta_0 + \epsilon_0 + \beta \mathbb{E}[V(X', \epsilon') | X, Y = 0] \right\}. \quad (5)$$

Let $\eta = \epsilon_0 - \epsilon_1$. Then the equilibrium decision maximizing the value function can be written as a cutoff rule $Y = \mathbb{1}\{\eta \leq \eta^*(X)\}$, where the cutoff $\eta^*(X)$ is defined as

$$\eta^*(X) \equiv W_1^\top \theta_1 - W_0^\top \theta_0 + \beta \left\{ \mathbb{E}[V(X', \epsilon') | X, Y = 1] - \mathbb{E}[V(X', \epsilon') | X, Y = 0] \right\}. \quad (6)$$

Taking expectations (over ϵ) in Eq. (5), we derive the “ex-ante” Bellman equation:

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

where $V^e(X) \equiv \mathbb{E}[V(X, \epsilon) | X]$ and

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

Mathematically, the ex-ante Bellman equation (7) is a *Fredholm Integral Equation of the second kind (FIE-2)*,⁸ a well-studied class of integral equations, for which solutions are well-known. 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, $\forall x \in \mathcal{S}_X$, we solve the FIE eq. (7) to obtain*

$$V^e(x) = u^e(X) + \sum_{s=1}^{\infty} \beta^s \cdot \mathbb{E}[u^e(X^{[s]}) | X] \quad (8)$$

That is, the ex-ante value function can equivalently be characterized by the Bellman equation (7), or as the discounted sum of current and future expected utilities (8). This is essentially an ex-ante version of Bellman’s principle of optimality. Moreover, these two equivalent representations of the $V^e(\dots)$ underlie the two prominent approaches for estimating *parametric* dynamic discrete-choice models; namely, the Rust (1987) nested-fixed point approach iterates over Eq. (7), while

⁷While this rules out heteroskedasticity in the unobserved shocks, it is possible, following Blevins (2014), to allow for some degree of heteroskedasticity by dividing the state variables into two groups $X = (X_A, X_B)$ such that $\epsilon \perp X_B | X_A$. The identification and estimation procedure described in this paper follow through, with the additional conditioning on X_A at every step.

⁸See 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.

the two-step approaches based on conditional choice probabilities (CCP's) use Eq. (8) to “forward simulate” the value function.⁹

Both the nested-fixed point and two-step CCP estimation approaches require the researcher to specify F_η , the distribution of the utility shocks. The key to our *semiparametric* identification approach is to show that the quantile function (the inverse of the CDF) of η also satisfies a Fredholm type-2 integral equation. This allows us to express the distribution of η in terms of quantities (such as the CCP's) estimable directly from the data, as we describe next.

3.2. A “Bellman-like” characterization of the quantile function of η . Let $p(x) = \mathbb{P}(Y = 1|X = x)$ denote the conditional choice probability given X .

Assumption C. η is continuously distributed with the full support \mathbb{R} .

Assumption D. $p(X)$ is continuously distributed with support on a closed interval, i.e., $[\underline{p}, \bar{p}] \subseteq [0, 1]$.

Assumption C is a weak assumption ensuring that the choice probabilities $p(X)$ are bounded away from 0 and 1 for all X .¹⁰ Assumption D requires the state variables X to contain some continuous components.¹¹ By Assumption C, $F_\eta(\cdot)$ is strictly increasing on its support \mathbb{R} . Hence we can unambiguously define the quantile function $Q(\cdot) = F_\eta^{-1}(\cdot)$.

Lemma 2. *Suppose assumptions A to D hold. For each $p \in [\underline{p}, \bar{p}]$, let $z(p) = \mathbb{E}[\phi(X)|p(X) = p]$. Then we have an FIE-2 for the quantile function $Q(\cdot)$:*

$$Q(p) = z(p)^\top \cdot \theta + \beta \int_{\underline{p}}^{\bar{p}} Q(\tau) \cdot \Pi(\tau, p; \beta) d\tau, \quad \forall p \in [\underline{p}, \bar{p}] \quad (9)$$

where $\Pi(\tau, p; \beta) \equiv \sum_{s=1}^{\infty} \beta^{s-1} [F_{p(X^{[s]}|p(X), Y}(\tau|p, 1) - F_{p(X^{[s]}|p(X), Y}(\tau|p, 0)]$.

The proof of this result, given in the Appendix, follows from algebraic manipulations of Eqs. (6) and (8). To provide intuition for Lemma 2, recall that the state variables X evolve

⁹(See e.g. Hotz and Miller, 1993; Bajari, Benkard, and Levin, 2007; Aguirregabiria and Mira, 2007; Pesendorfer and Schmidt-Dengler, 2008; Hong and Shum, 2010).

¹⁰It is implied in parametric discrete choice models in which η follows a logistic distribution, for instance.

¹¹Letting 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 $[0, 1]$. 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 , even if it is continuous, is typically only identified at a set of isolated points.

as a first-order Markov process, along the optimal dynamic decision path. As a result, the CCP's $p = P(Y = 1|X)$ likewise evolve as a first-order Markov process. Roughly speaking, then, $\Pi(p', p; \beta)$ represents the difference in the Markovian laws of motions for these choice probabilities between the two alternative choices $Y = 1$ and $Y = 0$. Since the CCP's $p(Y|X)$ can be estimated directly from the data, so can the law of motion $\Pi(p', p; \beta)$. Then Eq. (9) expresses, in recursive "Bellman-like" form, the quantile function $Q(p)$ at the current choice probability p as the sum of a current period term $z(p)^\top \cdot \theta$ involving p , and a discounted term involving the expectation of the quantile function $Q(p')$ at the next period's stochastic CCP.

By analogy with the solution to the value function in Eq. (8) as the discounted sum of future utilities, we expect that by solving Eq. (9), we can express $Q(p)$, the quantile function at the current value of p , as the discounted sum of current-period terms similar to $z(\cdot)'\theta$, evaluated at a sequence of current and future values of CCP's: $\{p, p', p'', p''', \dots\}$. Since these terms are all linear in θ , we expect then that $Q(p)$ itself will be linear in θ , as Lemma 3 confirms:

Lemma 3. *Suppose assumptions A to D hold, and $\beta^2 \cdot \int_{\underline{p}}^{\bar{p}} \int_{\underline{p}}^{\bar{p}} \Pi^2(p', p; \beta) dp' dp < 1$ (*). Then, on $[\underline{p}, \bar{p}]$, $Q(\cdot)$ is a linear function of the finite dimensional parameter θ :*

$$Q(p) = \left\{ z(p) + \beta \int_{\underline{p}}^{\bar{p}} R(p', p; \beta) \cdot z(p') dp' \right\}^\top \cdot \theta \equiv \mathcal{B}(p)^\top \theta, \quad \forall p \in [\underline{p}, \bar{p}] \quad (10)$$

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 condition (*) in Lemma 3 ensures that the mapping in eq. (9) is a contraction, so that the solution is unique. While this is a high-level condition on the law of motion of the (endogenous) choice probabilities structural primitives, it is testable in principle as this law of motion can be estimated from the data.

3.3. The Closed-Form Estimator. The linearity of $Q(p)$ in θ , as shown in Lemma 3, is critical for deriving a closed-form estimator for θ . Recall that $\mathbb{P}(Y = 1|X) = F_\eta(Q(p(X)))$. This leads to the following linear (in θ) index specification for the choice probabilities in the dynamic

discrete-choice model:¹²

$$\mathbb{P}(Y = 1|X) = F_{\eta}(m(X)^{\top} \cdot \theta), \quad (11)$$

where

$$m(X) = \phi(X) - \sum_{s=1}^{\infty} \beta^s \left\{ \mathbb{E} \left[\int_{\underline{p}}^{p(X^{[s]})} \mathcal{B}(\tau) d\tau | X, Y = 1 \right] - \mathbb{E} \left[\int_{\underline{p}}^{p(X^{[s]})} \mathcal{B}(\tau) d\tau | X, Y = 0 \right] \right\}. \quad (12)$$

The structure of the DDC model as given in Eq. (11) is identical to a *static* 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. (See, among many others, Manski (1975, 1985), Powell, Stock, and Stoker (1989), Ichimura and Lee (1991), Horowitz (1992), Klein and Spady (1993), and Lewbel (1998).) From these papers (see (Horowitz, 2009) for a summary), we know that θ can be estimated up to location and scale, which is ensured by the next two assumptions.¹³

Assumption E. $m(X)$ is continuously distributed with a joint probability density function, denoted by $f_m(\cdot)$. The matrix $\mathbb{E}[m(X)m(X)^{\top}]$ is invertible.

Assumption F. $\|\theta\| = 1$.

The first half of Assumption E requires at least one argument of X to be continuously distributed, and the second half is a testable rank condition. Assumption F imposes a scale normalization on θ .

While a number of semiparametric estimators are available for the binary choice model in Eq. (11), we utilize the Powell, Stock, and Stoker (1989) estimator, as it provides the important advantage of being a closed-form (non-iterative) estimator for θ . PSS show that θ can be estimated by $-2\mathbb{E}Y \cdot \nabla f_m(m)$; that is, as a weighted gradient of the density function of $m(X)$, the index functions in Eq. (11). Following PSS, we approximate the density function $f_m(\cdot)$

¹²This comes from plugging Eq. (9) into Eq. (17) in the Appendix.

¹³For notational simplicity, hereafter we assume the state vector X does not include a constant term in the 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.

nonparametrically using kernel functions, resulting in the estimator:

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

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.

Full details of the estimation procedure, are provided in Appendix A. Despite the mathematical complexity of the index functions $m(X)$ (cf. Eq. (12)), the computational curse of dimensionality associated with our estimator is no more than that associated with the usual CCP-based two-step estimators (as in Hotz and Miller (1993)). Specifically, the potential high-dimensionality of the state variables X does not affect the computation of the quantile function $Q(\cdot)$ at all, because X enters this function only through the choice probabilities $p(X)$ which are scalar quantities (cf. Eq. (10)).

In Appendix A, 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. A Monte Carlo exercise of our estimator, which is also in Appendix A, demonstrates that our estimator works well even in moderately-sized samples.

4. ESTIMATES OF OPTIMAL STOPPING MODEL FOR NYC TAXI DRIVERS

In this section we discuss the empirical results for the dynamic optimal stopping model of taxicab driver labor supply, which we presented in Section 2. The estimates of the structural payoff parameters are reported in Table 4. We considered four specifications, which differ in the nonlinear terms which we include in the model. Specification (1) allows a quadratic term in hours to enter the (dois)utility of driving. Specification (2) includes a quadratic term in earnings to enter the utility from quitting. Specification (3) adds quadratic terms on both hours and earnings. Finally, specification (4) replaces the quadratic earnings of specification (2) with a piecewise linear specification to estimate non-linear payoffs in income.

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

TABLE 4. Parameter Estimates

Parameter	Description	(1)	(2)	(3)	(4)	Logit
θ_u	Earnings (upon quitting)	0.9907 (0.0118)	0.7560 (0.0463)	0.9549 (0.0276)	0.8818 (0.0658)	0.8450 (0.4579)
$\theta_{u,03}$	Earnings * ($earnings \geq 300$)		-0.2946 (0.0670)		-0.2474 (0.0648)	-0.0229 (0.2279)
$\theta_{u,04}$	Earnings * ($earnings \geq 400$)		-0.2240 (0.0619)		-0.0264 (0.0697)	0.2509 (0.1984)
$\theta_{c,01}$	Cumul. hours (while working)	-0.1359 (0.0759)	-0.5399 (0.1031)	-0.1252 (0.0778)	-0.1821 (0.0980)	-0.1981 (0.4149)
$\theta_{c,02}$	Cumul. hours squared (while working)	-0.0004 (0.0002)				
$\theta_{c,03}$	Cum. hours * ($hours \geq 7$)			-0.3673 (0.0631)	-0.2923 (0.0851)	-0.4267 (0.2836)
$\theta_{c,04}$	Cum. hours * ($hours \geq 9$)			-0.0953 (0.0439)	-0.0951 (0.0517)	-0.0365 (0.1982)

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. Standard errors are shown in parentheses.

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. In specifications (3) and (4), which allows the contribution of earnings and hours to utility to change at large values, we see that the hours worked has an increasingly detrimental effect on utility while working, which is consistent with a increase marginal disutility from driving. On the other hand, the positive coefficient on earnings becomes smaller in magnitude at higher levels of earnings, implying a decreasing marginal utility from income. Below we will see how these features play out in drivers' implied quitting rules.

The final column in Table 4 presents estimates of specification (4), assuming a logit distribution for F_η and estimated using the Rust (1987) Nested Fixed Point procedure. For ease of comparison with the semiparametric estimates, we normalized the coefficient vector for the Logit specification to have a unitary norm, and report the normalized values. For most of the parameters, the logit and semiparametric estimates coincide. There are big differences, though, in the estimates of $\theta_{u,03}$ and $\theta_{u,04}$, with the logit estimates implying an increasing marginal utility

from earnings; however, the standard errors for these logit estimates are quite large compared to the semiparametric estimates, which cautions against drawing strong conclusions from them.¹⁴

TABLE 5. Parameter Estimates

Spec.	Parameter	Description	Weekday		Weekend	
			No Rain	Rain	No Rain	Rain
	θ_u	Earnings (upon quitting)	0.6673	0.8416	0.7867	0.8998
	$\theta_{u,03}$	Earnings $*(earnings \geq 300)$	-0.4715	-0.2724	-0.1781	-0.0099
	$\theta_{u,04}$	Earnings $*(earnings \geq 400)$	-0.0742	0.4583	0.4356	-0.1411
(4)	$\theta_{c,02}$	Cumul. hours (while working)	-0.3730	-0.0260	-0.2192	-0.4058
	$\theta_{c,03}$	Cum. Hours $*(hours \geq 7)$	-0.4219	-0.0812	-0.3170	-0.0077
	$\theta_{c,04}$	Cum. Hours $*(hours \geq 9)$	-0.0990	-0.0105	-0.1051	-0.0753


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. Standard errors are shown in parentheses.

In Table 5, we present the results estimated on the four subsamples broken down by weekday/weekend and rain/no-rain, as shown in Table 1. At face value, the parameter estimates are quantitatively quite similar across the subsamples, for each of the specifications. However, because of the normalization $\|\theta\| = 1$ required for the semiparametric estimation procedure, one cannot compare the magnitudes of the parameter estimates across specifications. There are nevertheless a few conclusions we can draw. Across all specifications, and all subsamples, the coefficient on earnings, θ_u , is largest in magnitude. However, for specifications (2) and (4), we see that for the “weekday/rain” and “weekend/no-rain” subsamples, the marginal utility from earnings is estimated to *increase* at higher levels of earnings.

While the empirical specifications in Table 4 are simple, the behavioral implications of the dynamic model, which we illustrate in Figure 1, are quite rich. Eqs. (11,12) above show that $m(X)'\theta$

¹⁴Furthermore, we also found that the logit estimates were quite sensitive to starting values in the optimization procedure.


FIGURE 1. Estimated Choice-specific Value Function Differences: $V_1(X) - V_0(X)$



We report the estimates of the $m(X)' \theta$ function defined in Eqs. (11,12) 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.

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 1, we plot the estimated $m(X)' \theta$ function at different levels of hours-worked, where estimates come from Specification 4. At smaller values for hours worked (the bottom three lines in Figure 1), the curves are upwardly-sloping in cumulative income. That is, holding hours fixed, a driver is more likely to quit at higher than lower income, implying a *negative* wage elasticity, which is consistent with the behavioral “income targeting” hypothesis. At higher-values of hours worked, however, the curves begin to slope downward, implying a positive elasticity. Figure 2 captures similar patterns by shows iso-contour lines for the difference in choice-specific value functions (panel I) and conditional choice probabilities (panel II) over a grid of income and hours. Again we see that by holding one dimension fixed and varying the other, we can reproduce both positive and negative “wage elasticities”. While the model is

FIGURE 2. Value Function Differences and Choice Probabilities





This figure depicts estimated results, showing iso-contours for the difference in choice-specific value functions (panel I) and conditional choice probabilities (panel II) over a grid of income and hours.

agnostic about the link between wage and hours per-se, the *mostly* horizontal gradients to these contours nevertheless suggest that hours are more predictive of drivers' quitting probabilities than incomes. Where the gradients are horizontal, incomes play a larger role.

In summary, then, we find that both "behavioral" and "neoclassical" wage responses are present in the data, with the behavioral income-targeting story explaining shorter shifts, and the standard neoclassical wage response explaining the longer shifts. At face value, these results are consistent with what in a static labor supply context would be called negative or positive wage elasticities, and hence they may offer a partial reconciliation of the divergent reduced-form results in the existing literature. Nevertheless, a point of emphasis here is that once taxi drivers' quitting decision are (correctly, we argued) modeled in a dynamic optimal stopping framework, the notion of wage makes less sense than in a static model, in which the wage rate is taken to be exogenous by the drivers and unchanging throughout the course of the day. Hence, implied wage elasticities are not natural to compute in our modeling framework, but the results here highlight how our (relatively simple) dynamic framework is rich enough to generate behavior which resembles both negative and positive wage elasticities from a static point of view.

Finally, in Figure 3 we graph the implied quantile function for the difference in utility shocks $\eta \equiv \epsilon_1 - \epsilon_0$, using our estimates from Specification 4. The density of \hat{p} is plotted as well, which highlights a range over which choice probabilities are actually observed. Outside of this range,

FIGURE 3. Estimated Quantile Function and Calibrated Parametric Approximations



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 $[0.01, 0.30]$. 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.¹⁵

This feature that, as shown in Figure 3, 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. One possibility may be to use our estimated quantile function, along the restricted range for which we obtain identification, to calibrate a distribution function for η along its full unrestricted range using distributions from parametric families.

This is illustrated in the bottom panel of Figure 3, where we plot the best-fitting quantile functions from four popular parametric families: Normal, Log-Normal, Gamma, and Extreme-Value. Apparently, all four alternative fit the estimated nonparametric quantile function reasonably within the identified range of $[0,0.3]$, but diverge appreciably outside of this range. It is an open question how robust counterfactuals might be to different specifications of the utility shock distribution at these large quantiles. It is interesting, however, that none of the fitted parametric quantile functions are able to match the "convex" shape of the nonparametric quantile function along its identified range.¹⁶

5. CONCLUSIONS

We develop a new closed-form estimator for semiparametric dynamic discrete-choice models. 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

¹⁵In 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).

¹⁶This may imply that truncated versions of the parametric distribution may fit our estimated quantile function better.

TABLE 6. Counterfactual Choice Probabilities For Quitting

		<i>Benchmark scenario ($\lambda = 1$)</i>									
Cum. Mins		Cumulative Income Earned									
Worked		\$20	\$49	\$76	\$104	\$133	\$163	\$196	\$232	\$277	\$352
20		0.0027	0.0042	0.0074	0.0129	0.0232	0.0412	0.0789	0.1597	0.5239	0.7433
65		0.0041	0.0062	0.0106	0.0175	0.0301	0.0511	0.0927	0.1779	0.5453	0.7565
110		0.0061	0.0089	0.0144	0.0229	0.0375	0.0610	0.1056	0.1938	0.5624	0.7669
155		0.0090	0.0128	0.0197	0.0298	0.0464	0.0719	0.1187	0.2088	0.5775	0.7760
204		0.0137	0.0187	0.0273	0.0391	0.0574	0.0843	0.1319	0.2219	0.5879	0.7815
255		0.0222	0.0289	0.0400	0.0542	0.0749	0.1038	0.1526	0.2428	0.6058	0.7919
310		0.0379	0.0473	0.0616	0.0788	0.1020	0.1322	0.1802	0.2677	0.6232	0.8014
370		0.0733	0.0874	0.1070	0.1287	0.1554	0.1873	0.2334	0.3131	0.6487	0.8135
445		0.1961	0.2236	0.2580	0.2927	0.3310	0.3716	0.4219	0.4969	0.7779	0.8877
565		0.1386	0.1621	0.1931	0.2254	0.2624	0.3031	0.3560	0.4374	0.7429	0.8688

		<i>Higher competition scenario ($\lambda = 0.5$)</i>									
Cum. Mins		Cumulative Income Earned									
Worked		\$20	\$49	\$76	\$104	\$133	\$163	\$196	\$232	\$277	\$352
20		0.0081	0.0123	0.0204	0.0325	0.0521	0.0814	0.1327	0.2265	0.5943	0.7867
65		0.0116	0.0170	0.0271	0.0413	0.0635	0.0952	0.1485	0.2434	0.6091	0.7951
110		0.0162	0.0230	0.0352	0.0516	0.0760	0.1096	0.1642	0.2592	0.6216	0.8020
155		0.0231	0.0318	0.0466	0.0657	0.0929	0.1288	0.1850	0.2802	0.6384	0.8115
204		0.0325	0.0433	0.0610	0.0828	0.1123	0.1500	0.2068	0.3012	0.6539	0.8201
255		0.0460	0.0595	0.0804	0.1051	0.1372	0.1766	0.2339	0.3271	0.6726	0.8304
310		0.0652	0.0816	0.1057	0.1330	0.1671	0.2075	0.2644	0.3554	0.6919	0.8411
370		0.0939	0.1132	0.1398	0.1689	0.2037	0.2437	0.2983	0.3847	0.7099	0.8506
445		0.1721	0.1974	0.2295	0.2624	0.2994	0.3392	0.3898	0.4670	0.7600	0.8779
565		0.1386	0.1621	0.1931	0.2254	0.2624	0.3031	0.3560	0.4374	0.7429	0.8688

		<i>Higher fares scenario ($\lambda = 1.5$)</i>									
Cum. Mins		Cumulative Income Earned									
Worked		\$20	\$49	\$76	\$104	\$133	\$163	\$196	\$232	\$277	\$352
20		0.0012	0.0018	0.0032	0.0056	0.0107	0.0209	0.0463	0.1115	0.4581	0.7000
65		0.0020	0.0029	0.0049	0.0081	0.0145	0.0267	0.0555	0.1261	0.4824	0.7172
110		0.0032	0.0046	0.0072	0.0112	0.0188	0.0324	0.0630	0.1350	0.4906	0.7209
155		0.0054	0.0074	0.0111	0.0165	0.0258	0.0416	0.0752	0.1513	0.5121	0.7349
204		0.0094	0.0124	0.0175	0.0245	0.0358	0.0537	0.0893	0.1673	0.5293	0.7456
255		0.0169	0.0214	0.0287	0.0380	0.0517	0.0716	0.1083	0.1861	0.5444	0.7538
310		0.0326	0.0399	0.0507	0.0636	0.0808	0.1035	0.1412	0.2170	0.5673	0.7661
370		0.0709	0.0836	0.1009	0.1198	0.1428	0.1699	0.2088	0.2784	0.6055	0.7836
445		0.2099	0.2392	0.2758	0.3124	0.3524	0.3944	0.4459	0.5207	0.7930	0.8961
565		0.1387	0.1622	0.1931	0.2254	0.2624	0.3031	0.3559	0.4373	0.7429	0.8687

Data from TLC Data, February 2012. Each cell shows estimates of optimal choice probabilities for model parameters in Specification (4), and error distribution G_η given by the Logistic distribution as shown in the bottom panel of Figure 3.

TABLE 7. Counterfactual Shift Durations and Revenues

λ	Revenue Mean	Duration Mean	Revenue Median	Duration Median	Revenue StDev	Duration StDev
0.50	125.74	40.63	125.95	38.11	45.60	16.28
0.75	136.48	47.70	140.71	48.51	54.26	17.48
1.00	178.65	60.67	187.86	59.40	59.62	19.08
1.25	207.37	67.05	223.00	71.14	58.60	19.87
1.50	226.56	71.82	223.00	71.14	55.58	19.16

question of drivers' wage elasticities by modeling 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.

Methodologically, 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.
- BUCHHOLZ, N. (2017): "Spatial Equilibrium, Search Frictions and Efficient Regulation in the Taxi Industry," Princeton University, working paper.
- 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. (2016): "Identification of Discrete Choice Dynamic Programming Models with Nonparametric Distribution of Unobservables," *Econometric Theory*.
- 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 (2016): "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.
- HAGGAG, K., B. MCMANUS, AND G. PACI (2017): "Learning by driving: Productivity improvements by New York City taxi drivers," *American Economic Journal: Applied Economics*, 9(1), 70–95.
- 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., 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., 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.
- 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.
- SRISUMA, S., AND O. LINTON (2012): "Semiparametric estimation of Markov decision processes with continuous state space," *Journal of Econometrics*, 166(2), 320–341.
- THAKRAL, N., AND L. T. TÔ (2017): "Daily Labor Supply and Adaptive Reference Points," Discussion paper, Working Paper.
- 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. (13) 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, $h_z = 1.06 \times \hat{\sigma}(p(X)) \times T^{-\frac{1}{2\tau+3}}$.

Let $(b_1^*(\cdot), \dots, b_{k_\theta}^*(\cdot))$ denote the components of $\mathcal{B}(p)$ in Eq. (10). For $\ell = 1, \dots, k_\theta$, $b_\ell^*(\cdot)$ satisfies

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

be the sequence of solutions supported on $[\underline{p}, \bar{p}]$. 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, $h_z = 1.06 \times \hat{\sigma}(p(X)) \times T^{-\frac{1}{2i+3}}$ 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, $h_m = 1.06 \times \hat{\sigma}(X) \times T^{-\frac{1}{2i+k}}$ is chosen optimally.

Following the standard kernel regression literature, we can show our PSS-based estimator, $\hat{\theta}$ (defined in Eq. (13)) 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 (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 G. $h_\theta = T^{-\frac{1}{\gamma}}$ where $k_\theta + 2 < \gamma < k_\theta + 3 + \mathbb{1}(k_\theta \text{ is even})$.

Assumption H. 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 I. (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 G and H are introduced by PSS for the choice of bandwidth and kernel, respectively, to control the bias term in the estimation of θ .¹⁷ The restriction on the bandwidth Assumption G 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 I encompasses high-level conditions that could be further established under primitive conditions. In particular, Assumption I(i) requires $\hat{m}(\cdot)$ to converge to $m(\cdot)$ faster than $T^{-\frac{1}{4}}$. By Assumption I(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 I(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). Assumption I 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:

¹⁷Note 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.

Theorem 1. *Suppose assumptions G to I 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 $\lambda = \mathbb{E}[f_m(m(X)) \times f_\eta(m(X)^\top \cdot \theta)]$.

In the above theorem, recall $\mathbb{P}(Y = 1|X) = F_\eta(\eta^*(X))$ and $\eta^*(X) = m(X)^\top \cdot \theta$ by ???. Our estimator $\hat{\theta}$ (as defined in Eq. (13)) has not imposed the scale restriction in Assumption F; 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

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

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.¹⁸ 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.

A.1. Monte Carlo. These Monte Carlo experiments illustrate the finite-sample 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 $\theta_2 = -2$.

¹⁸The 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.

TABLE 8. 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.

θ_0 is not identified in the semiparametric framework, so we treat it as a nuisance parameter. Let $\theta = (\theta_1, \theta_2)^\top$. Since θ is only identified up to scale in the semiparametric setting (cf. Assumption F), for comparing the performance of the semiparametric estimators, we treat the scale of θ as known, i.e., $\|\theta\| = \sqrt{5}$, rather than imposing a different normalization, as assumption F. We present in Table 8 the bias and standard deviation of the semiparametric estimator.

APPENDIX B. PROOFS

B.1. Proof of Lemma 1.

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

Because it is straightforward to verify that the solution in the lemma solves eq. (7), Hence, it suffices to show the uniqueness of the solution. Eq. (7) 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 (7). 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:

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

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 2.

Proof. By using eq. (8), along with Lemma 1, eq. (6) becomes

$$\eta^*(X) = W_1^\top \theta_1 - W_0^\top \theta_0 + \sum_{s=1}^{\infty} \beta^s \left\{ \mathbb{E}[u^e(X^{[s]})|X, Y = 1] - \mathbb{E}[u^e(X^{[s]})|X, Y = 0] \right\} \quad (15)$$

$$= \phi^\top(X) \cdot \theta - \sum_{s=1}^{\infty} \beta^s \left\{ \mathbb{E}[\eta^{[s]} \mathbb{1}(\eta^{[s]} \leq \eta^*(X^{[s]}))|X, Y = 1] - \mathbb{E}[\eta^{[s]} \mathbb{1}(\eta^{[s]} \leq \eta^*(X^{[s]}))|X, Y = 0] \right\} \quad (16)$$

where $\phi(X) = (\phi_0^\top(X), \phi_1^\top(X))^\top$ and $\theta = (\theta_0^\top, \theta_1^\top)^\top$ and $\phi_d(X) \equiv (-1)^{d+1} W_d + \sum_{s=1}^{\infty} \beta^s \left\{ \mathbb{E}[W_d^{[s]} \mathbb{1}_{Y^{[s]}=d}|X, Y = 1] - \mathbb{E}[W_d^{[s]} \mathbb{1}_{Y^{[s]}=d}|X, Y = 0] \right\}$.

Moreover, using the substitution $\tau \rightarrow Q(\tau)$, we have

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

Given the above equation, we can rewrite the expression for the cutoff value $\eta^*(X)$, from Eq. (16), as

$$Q(p(X)) = \phi^\top(X) \cdot \theta - \sum_{s=1}^{\infty} \beta^s \left\{ \mathbb{E} \left[\int_0^{p(X^{[s]})} Q(\tau) d\tau | X, Y = 1 \right] - \mathbb{E} \left[\int_0^{p(X^{[s]})} Q(\tau) d\tau | X, Y = 0 \right] \right\} \quad (17)$$

Eq. (17) is almost a Fredholm type-2 integral equation, except that the function on the LHS, $Q(p(X))$ is actually a composed function of X , and not a function of τ , the quantile. But it is easy to transform it into a bona-fide FIE-2. assumption D implies that the choice probability $p(X)$ is continuously distributed on a

closed interval, over which we can integrate.¹⁹ Accordingly, we take the conditional expectation given $p(X) = p$ on both sides of eq. (17) to obtain

$$Q(p) - \beta \int_{\underline{p}}^{\bar{p}} \int_{\underline{p}}^{p'} Q(\tau) d\tau \cdot \pi(p', p; \beta) dp' = z(p)^\top \cdot \theta, \quad \forall p \in [\underline{p}, \bar{p}]. \quad (18)$$

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

Eq. (18) is already an FIE-2. To see this, rewrite the second term 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 Eq. (9). \square

B.3. Proof of Lemma 3.

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

B.4. Proof of Theorem 1. The estimator is defined in (13). 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)).

¹⁹This 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 $\mathcal{S}_{p(X)}$.

By Taylor expansion, we have

$$\begin{aligned}\hat{\theta} = \tilde{\theta} &- \frac{2}{T(T-1)} \sum_{t=1}^T \sum_{s \neq t} \left[\frac{1}{h_\theta^{k_\theta+2}} \nabla^2 K_\theta \left(\frac{m(X_t) - m(X_s)}{h_\theta} \right) \times Y_s \times (\hat{m}(X_t) - m(X_t)) \right] \\ &+ \frac{2}{T(T-1)} \sum_{t=1}^T \sum_{s \neq t} \left[\frac{1}{h_\theta^{k_\theta+2}} \nabla^2 K_\theta \left(\frac{m(X_t) - m(X_s)}{h_\theta} \right) \times Y_s \times (\hat{m}(X_s) - m(X_s)) \right] \\ &+ O_p(h_\theta^{-3} \cdot \mathbb{E} \|\hat{m}(X) - m(X)\|^2) \equiv \tilde{\theta} + \mathbb{A}_1 + \mathbb{A}_2 + \mathbb{B} \quad (19)\end{aligned}$$

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 I(i), we have

$$h_\theta^{-3} \times \mathbb{E} \|\hat{m}(X) - m(X)\|^2 = h_\theta^{-3} \times o_p(T^{-1/2} h_\theta^3) = o_p(T^{-1/2}). \quad (20)$$

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 I(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 I(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

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

By Assumption I(iii),

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

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

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

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 I(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).