THE REVIEW OF ECONOMIC STUDIES

Review of Economic Studies (2010) **77**, 273–304

0034-6527/09/00411011$02.00
doi: 10.1111/j.1467-937X.2009.00576.x

Pairwise-Difference Estimation of a Dynamic Optimization Model

HAN HONG
Stanford University

and

MATTHEW SHUM
California Institute of Technology

First version received December 2004; *final version accepted April* 2009 *(Eds.)*

We develop a new estimation methodology for dynamic optimization models with unobserved shocks and deterministic accumulation of the observed state variables. Investment models are an important example of such models. Our pairwise-difference approach exploits two common features of these models: (1) the monotonicity of the agent's decision (policy) function in the shocks, conditional on the observed state variables; and (2) the state-contingent nature of optimal decision making which implies that, conditional on the observed state variables, the variation in observed choices across agents must be due to randomness in the shocks across agents. We illustrate our procedure by estimating a dynamic trading model for the milk production quota market in Ontario, Canada.

1. INTRODUCTION

In this paper, we propose a new estimation methodology for a dynamic optimization model with preference and/or payoff shocks which are unobserved by the econometrician, but are observed by agents when they make their dynamic choices. The two-step estimator we propose relies on two common features of the dynamic optimization problem we consider. First, we exploit the monotonicity of the agent's decision (policy) function in the unobserved shocks, conditional on the observed state variables. Second, we exploit the state-contingent nature of optimal decision making which implies that, conditional on the observed state variables, the variation in observed choices across agents must be due to randomness in the shocks across agents.

This paper makes two contributions. First, the two-step pairwise-difference estimator we propose applies to the estimation of continuous–discrete choice dynamic models. To our knowledge, our approach represents the first application of pairwise-differencing methods, which have primarily been used in static cross-sectional and panel data contexts (cf. Honoré and Powell, 1994), to forward-looking structural dynamic optimization problems.

Second, our two-step estimation approach has a practical benefit in reducing the computational burden associated with estimating dynamic models due to the need for numeric dynamic programming. A number of model parameters can be estimated in the first step, which is computationally simple and does not involve numeric dynamic programming. Since the second step may require numeric dynamic programming or forward simulation in order to recover the value function, estimating a subset of the parameters in the first step significantly reduces the number of times that the value function must be computed in the second step, thereby lowering an important computational hurdle in estimating dynamic models.

Our approach is related to some recent work that exploits monotonicity assumptions to identify and estimate structural equations. Earlier, Olley and Pakes (1996) exploited such an assumption in order to invert out the unobservable shock and to derive a semi-parametric estimator for production functions with serially correlated unobservables. Matzkin (2003) exploited the quantile invariance implication of monotonicity to estimate non-parametrically functions which are non-linear in the error term. Bajari and Benkard (2005) also used this principle in their study of hedonic discrete-choice models of demand for differentiated products.

The model considered in this paper can be applied to any investment or consumption problem where the accumulation equation of the asset variable is deterministic and does not contain unobserved variables. The applications include the management of production quotas (which is the empirical illustration presented later in this paper), hiring/firing of employees by firms, and household consumption–savings problems. It does exclude cases where the evolution of the asset variable is stochastic (such as human capital investment) or consumption/investment cases when all variables in the asset accumulation equation are not observed.

This paper complements the existing literature on identification and estimation in discrete-choice dynamic optimization models (cf. Pakes and Simpson, 1989; Hotz and Miller, 1993; Taber, 2000; Magnac and Thesmar, 2002; Aguirregabiria, 2005). It is also related to recent literature on the identification and estimation of dynamic game models (e.g. Pesendorfer and Schmidt-Dengler, 2008; Aguirregabiria and Mira, 2007; Berry, Ostrovsky and Pakes, 2007; Bajari, Benkard and Levin, 2007). While we do not focus on dynamic games here, one contribution that we make is the consideration of situations where agents have both continuous action spaces and continuous state spaces.

The plan of the paper is as follows. In the next section, we present a single-agent dynamic optimization problem and state our model assumptions. We describe our two-step estimation approach in Section 3. In Section 4, we illustrate our methodology by estimating a dynamic model of trading behaviour in monthly exchanges operated by provincial regulatory agencies in Ontario, Canada, to allocate milk production quotas across milk farmers. We conclude in Section 5.

2. EMPIRICAL FRAMEWORK

Consider the following dynamic optimization problem of an agent i:

$$\max_{\{q_{it}\}_{t=0}^{\infty}} E\left[\sum_{t=0}^{\infty} \beta^t U(x_{it}, s_{it}, q_{it}; \theta) \mid \{q_{it}\}_{t=0}^{\infty}\right] \tag{1}$$

subject to the Markov transition probabilities for the state variables

$$F(x_{i,t+1}, s_{i,t+1} | x_{it}, s_{it}, q_{it}). \tag{2}$$

In this problem, x_{it} and s_{it} are the two state variables, with the distinction that x_{it} is observed by the econometrician, but s_{it} is not. The agent's choice variable is denoted by q_{it}. An example of such a model is an investment model where x_{it} can be interpreted as a stock and the control q_{it} as investment, or incremental additions to the stock which can be purchased at some fixed price. The unobserved variable s_{it} would be a time-varying idiosyncratic shock which affects agent i's period-t investment decisions. For convenience, we will sometimes refer to x_{it} as the "stock" and q_{it} as "investment" in this paper, in reference to this example.

$U(x_{it}, s_{it}, q_{it}; \theta)$ is a per-period utility function, parameterized by the parameter vector θ. The per-period utility depends on the current stock x_{it} and the idiosyncratic shock s_{it}, which is known to agent i before he makes his choice of q_{it}. We assume that the shock s_{it} is observed by the optimizing agent at the time she makes her period t decision, but not by the econometrician. This usage differs from a measurement error interpretation of a "shock", where a "shock" is often unobserved by both the econometrician as well as the optimizing agent when she makes her decision. The presence of the unobserved shock s_{it} induces, from the econometrician's point of view, randomness in the observed choices of the control q_{it}. As in Rust (1996), we also assume:

Assumption 1 (Conditional independence). *The Markov transition probabilities for the state variables can be factored as:*

$$F(x_{i,t+1}, s_{i,t+1}|x_{it}, s_{it}, q_{it}) = F(x_{i,t+1}|x_{it}, s_{it}, q_{it}) \cdot F_{\mathrm{s}}(s_{i,t+1}; \gamma). \tag{3}$$

Note that, without any restrictions, the following factorization holds:

$$F(x_{i,t+1}, s_{i,t+1}|x_{it}, s_{it}, q_{it}) = F(x_{i,t+1}|x_{it}, s_{it}, q_{it}) \cdot F_{\mathrm{s}}(s_{i,t+1}|x_{i,t+1}, x_{it}, s_{it}, q_{it}). \tag{4}$$

Hence, Assumption 1 consists to two restrictions. First, the law of motion for the observed state variable $x_{i,t+1}$ implies that $(x_{i,t+1}, s_{i,t+1})$ are independent, conditional on (x_{it}, s_{it}, q_{it}). Second, it implies that $F_s(s_{i,t+1}|x_{i,t+1}, x_{it}, s_{it}, q_{it}) = F(s_{i,t+1})$. The right-hand side is then assumed to come from a parametric family denoted by γ. While this rules out the important case of serial correlation in the unobserved shocks over time (arising perhaps from unobserved agent-specific fixed effects), it is a common assumption made in the literature on estimation of dynamic models. On the other hand, it is straightforward to extend the *i.i.d.* assumption to one where heterogeneity in the distribution of the shock s_{it} across agents and time is explicitly parameterized to depend on observed conditioning covariates.

Assumption 2 (Deterministic accumulation). *The stocks evolve in the following deterministic manner:*

$$x_{it+1} = x_{it} + q_{it}, \ \forall i, t. \tag{5}$$

This assumption is quite specific, but it arises naturally in in investment models, and also in our empirical illustration below. This assumption is important for the practical application of our estimator.[1]

1. As we discuss in Section 3.1.2 below, we could relax this assumption to allow x_{it} to evolve in a non-linear deterministic manner. However, the resulting estimator would involve an additional differencing step, making it less convenient and attractive.

Given these assumptions, and assuming stationarity, the agent's optimal policy function can be expressed as the maximizer of Bellman's equation: for each t,

$$\begin{aligned} q(x_{it}, s_{it}; \theta, \gamma) &= \text{argmax}_q \left\{ U(x_{it}, s_{it}, q; \theta) + \beta \mathbb{E}_{x_{it+1}, s_{it+1}|x_{it}, s_{it}, q} V(x_{it+1}, s_{it+1}; \theta, \gamma) \right\} \\ &= \text{argmax}_q \left\{ U(x_{it}, s_{it}, q; \theta) + \beta \mathbb{E}_{s_{it+1}|x_{it}, s_{it}, q} V(x_{it} + q, s_{it+1}; \theta, \gamma) \right\} \\ &= \text{argmax}_q \left\{ U(x_{it}, s_{it}, q; \theta) + \beta \int V(x_{it} + q, s_{it+1}; \theta, \gamma) F_s(ds_{it+1}; \gamma) \right\} \end{aligned} \quad (6)$$

where

$$V(x_{i,t+1}, s_{i,t+1}; \theta, \gamma) \equiv \max_{\{q_{i\tau}\}_\tau} \mathbb{E} \left[\sum_{\tau=t+1}^{\infty} \beta^{\tau-t-1} U(x_{i\tau}, s_{i\tau}, q_\tau; \theta) \mid \{q_{i\tau}\}_\tau, x_{i,t+1}, s_{i,t+1} \right]. \quad (7)$$

In equation (6), Assumption 2 is used to substitute $x_{it} + q_{it}$ for x_{it+1} in the second line, and Assumption 1 is used to get from the second to the third lines.

In what follows, we simplify notation by defining

$$\mathcal{V}(x_{it} + q_{it}; \theta, \gamma) \equiv \int V(x_{it} + q_{it}, s; \theta, \gamma) F_s(ds; \gamma), \quad (8)$$

the *ex ante* value function at time t, where the expectation is over $s_{i,t+1}$, the future realization of the shock.

2.1. *Monotonicity and quantile invariance*

We assume that the policy functions are monotonic in the unobserved state variable, conditional on a particular value for the observed state variable.

Assumption 3 (Monotonicity). *The policy functions $q(x_{it}, s_{it}; \theta, \gamma)$ are non-decreasing in s_{it}, conditional on x_{it}.*

Remark 1. Given Assumptions 1 and 2, a sufficient condition for Assumption 3 is that U is supermodular in (q, s), for all x.

Proof. (Remark 1): The optimal policy $q(x, s)$ is given by

$$\text{argmax}_q \overline{U}(x, s, q) \equiv \{U(x, s, q; \theta) + \beta \mathcal{V}(x + q; \theta, \gamma)\}. \quad (9)$$

In order for $q\,(s, x; \theta, \gamma)$ to be non-decreasing in s given x, we require $\overline{U}\,(x, s, q; \theta)$ to be supermodular in (q, s), for all x. This is equivalent to supermodularity of $U\,(x, s, q; \theta)$ in (q, s) given x, because the expected continuation value function $\mathcal{V}\,(x + q; \theta, \gamma)$ does not depend on s, from Assumption 1. ‖

An important implication of Assumption 3 is **quantile invariance**: conditional on x_{it}, the τ-th quantile of q_{it} conditional on x_{it} is $q\,(x_{it}, s_\tau; \theta, \gamma)$, where s_τ is the τ-th quantile of $F_s\,(\cdot)$. This implication of monotonicity was also exploited by Matzkin (2003) in her non-parametric estimation methodology for random functions that are non-additive in the error term.

The independence assumption that the distribution function F_s does not depend on x allows us to accommodate situations (such as atoms in $F(q|x)$) where we only have *weak* monotonicity of q in s, given x. This allows the investment decision to be a mixed discrete–continuous choice variable, with a point mass at zero (indicating no investment). This accommodates models of non-convex adjustment costs (cf. Eberly, 1994), and is appropriate for the empirical illustration we consider below.

3. ESTIMATION APPROACH

The parameters we wish to estimate are θ and γ, which are, respectively, the utility function and shock distribution parameters. To simplify notation, we assume that our data are a balanced panel: $\{q_{it}, x_{it}\}$, $i = 1, \ldots, N$, $t = 1, \ldots, T$. This is not critical, as our estimator also applies to cases where the number of cross-sectional observations differs across-time periods. Furthermore, the choice variable q can have both discrete and continuous components. In the discussion below, we assume for convenience that q has a mass point at zero, but also takes continuous non-zero values.

From the data, we can estimate the empirical distribution of q given x for each x. Denote each element of this family of distributions (indexed by x) by $\hat{F}(q|x)$. Therefore, $\hat{F}(q_{it}|x_{it})$ denotes the estimated conditional probability of $q(x_{it}, s_{it}) \leq q_{it}$, conditional on the observed state variable being equal to x_{it}.

Since the conditioning variable x is continuous, a kernel estimator can be employed to estimate these conditional cumulative distribution functions (CDFs):

$$\hat{F}(q|x) = \frac{\frac{1}{T}\frac{1}{N}\sum_{t=1}^{T}\sum_{i=1}^{N}\mathbf{1}(q_{it} \leq q)\, K\left(\frac{x-x_{it}}{h}\right)}{\frac{1}{T}\frac{1}{N}\sum_{t=1}^{T}\sum_{i=1}^{N} K\left(\frac{x-x_{it}}{h}\right)} \tag{10}$$

where $K(\cdot)$ is a kernel weighting function and h is a bandwidth sequence. In computing $\hat{F}(q|x)$, we employ all the observations, including those for which $q_{it} = 0$ (i.e., for which the agent remained at a corner solution and investment is zero).

We make the following assumptions on the kernel function:

Assumption 4. *1. $K(\cdot)$ is an r-th order kernel function, with $r \geq 2$: (i) $\int K(u)\, du = 1$; (ii) $\int u^{\xi} K(u)\, du = 0$ for $\xi = 1, \ldots, r-1$; and (iii) $\int u^{r} K(u)\, du < \infty$.*
2. As $N \to \infty$, the bandwidth satisfies (i) $h \to 0$; (ii) $\frac{Nh}{\log N} \to \infty$; and (iii) $\sqrt{N}h^{r} \to 0$.

Furthermore, we also require smoothness assumptions on the shock distribution and the per-period utility function:

Assumption 5. *(i) The shock distribution $F_s(s)$ has continuous derivatives up to order r that are uniformly bounded. The shock density $f_s(s)$ is bounded away from 0 on any compact set on its support. (ii) The function $U(x, s, q; \theta)$ has continuous partial derivatives in (x, s, q) of order $r + 1$ (where r is the order of the kernel from the previous step). The expectations of all derivatives with respect to x, s, q of order up to $r + 1$ exist. (iii) The density $f(x)$ of the observed state variable is uniformly bounded, continuous, and bounded away from 0 on any compact set on its support.*

Conditions 1.(iii) and 2.(iii) of Assumption 4 above are standard conditions for reducing the asymptotic bias in the kernel estimates. Assumption 5 ensures that the asymptotic bias of the limit pairwise-differencing estimating function (described below) can be approximated up to the r-th order of the bandwidth parameter (as in Powell, Stock and Stoker, 1989). Next, we describe our proposed two-step estimation approach.

3.1. *First step: pairwise-differencing of first-order conditions*

In the first step, we obtain estimates of γ, the parameters of the shock distribution, as well as a subset of the parameters θ in the utility function, by exploiting the first-order condition of the maximization problem in equation (6).[2] This step exploits the state-contingent nature of optimal decision making, which implies that, conditional on the observed state variables, the variation in observed choices across agents must be due to randomness in the unobserved state variables across agents.

First, the deterministic accumulation nature of the stock evolution process implies that the maximization problem for any agent i can be rewritten as

$$q\left(x_{it}, s_{it}; \theta, \gamma\right) = \operatorname{argmax}_q \left\{U\left(x_{it}, s_{it}, q; \theta\right) + \beta \mathcal{V}\left(x_{it} + q; \theta, \gamma\right)\right\}. \tag{11}$$

For any agent i who invests a non-zero amount $q_{it} \neq 0$, her choice of q_{it} satisfies the first-order condition

$$U_q\left(x_{it}, s_{it}, q_{it}; \theta\right) + \beta \mathcal{V}'\left(x_{it} + q_{it}; \theta, \gamma\right) = 0 \tag{12}$$

where $U_q\left(x_{it}, s_{it}, q_{it}; \theta\right)$ refers to the derivative of $U\left(x_{it}, s_{it}, q_{it}; \theta\right)$ with respect to its third argument. For any pair of agents i and j in period t such that $x_{it} + q_{it} = x_{jt} + q_{jt}$,

$$\mathcal{V}'\left(x_{it} + q_{it}; \theta, \gamma\right) = \mathcal{V}'\left(x_{jt} + q_{jt}; \theta, \gamma\right). \tag{13}$$

Hence we can condition on such pairs of agents in order to control for the unknown form of the expected value function.

Second, from the quantile invariance Assumption 3 and the assumption that s is distributed independently of x, we know that any individual i with a (q_{it}, x_{it}) pair must have received a shock s_{it} equal to $F_s^{-1}\left(\hat{F}(q_{it}|x_{it}); \gamma\right)$, the $\hat{F}(q_{it}|x_{it})$-th quantile of the shock distribution. This suggests that the cross-sectional variation in q given x for a collection of quantiles allows us to recover the corresponding quantiles of F_s, and hence estimate the γ parameters.

The considerations above suggest a pairwise-difference estimator for the first-stage parameters. Consider a pair of individuals i and j in period t with the same $x_{it} + q_{it} = x_{jt} + q_{jt}$. If we difference the first-order conditions for these two observations, we obtain

$$\left\{U_q\left(x_{it}, s\left(\hat{F}(q_{it}|x_{it}); \gamma\right), q_{it}; \theta\right) - U_q\left(x_{jt}, s\left(\hat{F}(q_{jt}|x_{jt}); \gamma\right), q_{jt}; \theta\right)\right\} = 0, \tag{14}$$

where $s(\tau; \gamma) \equiv F_s^{-1}(\tau; \gamma)$, the τ-th quantile of F_s.

Let θ_1 denote the subset of the parameters θ which enter equation (14). Precisely, θ_1 is the subset of the parameters θ which are not eliminated by either (i) taking the derivative of the utility function U with respect to its third argument; or (ii) taking the difference of the utility function derivative U_q between two individuals. The remaining parameters $\theta_2 \equiv \{\theta \setminus \theta_1\}$ will be estimated in the second step of our procedure.

Let $\psi \equiv (\theta_1, \gamma)$, the parameters estimated in the first step, and define I_{it} to be the indicator $\mathbf{1}(q_{it} \neq 0)$. Furthermore, we use $z_{it} \equiv (x_{it}, q_{it})$ to denote the data variables observed for agent i in period t. The pairwise-difference estimator of ψ takes the following form:

$$\min_{\theta_1, \gamma} \frac{1}{(NT)^2} \sum_{t=1}^{T} \sum_{i=1}^{N} \sum_{t'=1}^{T} \sum_{j=1}^{N} \left\{\frac{1}{h} K\left(\frac{(x_{it} + q_{it}) - \left(x_{jt'} + q_{jt'}\right)}{h}\right) \cdot I_{it} I_{jt'} \cdot \hat{m}\left(z_{it}, z_{jt'}; \psi\right)^2\right\} \tag{15}$$

2. Recently, Berry and Pakes (2000) also exploit the first-order condition to derive estimates of structural parameters for models of *multi-agent* dynamic games. However, their model is different from ours because unobserved state variables are not present in their model.

where $\hat{m}\left(z_{it}, z_{jt'}; \psi\right)$ denotes the differenced first-order condition:

$$\hat{m}\left(z_{it}, z_{jt'}, \psi\right) \equiv U_q\left(x_{it}, F_s^{-1}\left(\hat{F}\left(q_{it}|x_{it}\right); \gamma\right), q_{it}; \theta_1\right) - U_q\left(x_{jt'}, F_s^{-1}\left(\hat{F}\left(q_{jt'}|x_{jt'}\right); \gamma\right), q_{jt'}; \theta_1\right). \tag{16}$$

In computing the objective function (15) above, we only include observations with non-zero investment ($q \neq 0$) because only for these observations is the first-order condition (12) satisfied.[3]

The first-stage estimate $\hat{\gamma}$ of the parameters in the shock distribution function can be used to derive an estimate of the optimal policy function

$$\tilde{q}\left(x, s\right) \equiv \hat{F}_{q|x}^{-1}\left(F_s\left(s; \hat{\gamma}\right)\right), \ \forall s. \tag{17}$$

This estimate of the period t investment choice q_t at a given state (x, s) is just the $F_s(s; \hat{\gamma})$-th quantile of $\hat{F}(q|x)$, the empirical conditional distribution of q given x.

3.1.1. Asymptotic theory for first step. Ahn and Powell (1993) and Honoré and Powell (1994) pioneered the use of pairwise-differencing methods in econometrics and developed the techniques for deriving their asymptotic distributions. The objective function (15) resembles weighted least squares, where each pair of observations is weighted by a kernel function which takes on small values when certain features of the pair of observations are very far apart.

From equation (15), we can alternatively express the pairwise-difference estimate for ψ as that which solves the following sample score function:

$$W_{NT}(\hat{\psi}) \equiv \frac{1}{(NT)^2} \sum_i \sum_t \sum_j \sum_{t'} \hat{r}\left(z_{it}, z_{jt'}, \hat{\psi}\right) = 0 \tag{18}$$

where

$$\hat{r}\left(z_{it}, z_{jt'}, \hat{\psi}\right) \equiv \frac{1}{h} K\left(\frac{x_{it+1} - x_{jt'+1}}{h}\right) \hat{m}\left(z_{it}, z_{jt'}, \hat{\psi}\right) \frac{\partial}{\partial \psi}\left[\hat{m}\left(z_{it}, z_{jt'}, \hat{\psi}\right)\right] I_{it} I_{jt'}. \tag{19}$$

The limit objective function of the first-step estimator is

$$G_0(\psi) \equiv E_{x,q} E_{x',q'} \Big\{ \mathbf{1}\left(x + q = x' + q', \ q \neq 0, \ q' \neq 0\right) \cdot \left[U_q\left(x, F_s^{-1}\left(F(q|x); \gamma\right), q, \theta_1\right) - U_q\left(x', F_s^{-1}\left(F(q'|x'); \gamma\right), q', \theta_1\right)\right]^2 \Big\}. \tag{20}$$

Also define $m\left(z_{it}, z_{jt'}, \psi\right)$ and $r\left(z_{it}, z_{jt'}, \psi\right)$ analogous to $\hat{m}\left(z_{it}, z_{jt'}, \psi\right)$ and $\hat{r}\left(z_{it}, z_{jt'}, \psi\right)$ except that $\hat{F}\left(q_{it}|x_{it}\right)$ in equation (16) is replaced by the unknown true $F\left(q_{it}|x_{it}\right)$.

The regularity conditions required for the asymptotic results are collected in the following assumption:

Assumption 6. *Regularity conditions for first step:*

3. Even though we only use observations where $q_{it} \neq 0$, there is no selection issue here because, given the monotonicity assumption, we control for the selection by substituting in estimates of the random shocks (the s_t's) in the first-order conditions (which was a similar device used by Olley and Pakes (1996) in their productivity analysis). We thank a referee for pointing this out.

i. $\psi \in \Psi$, *a compact subset of* $\mathbf{R}^P$, *and the true value* $\psi^0 \in int(\Psi)$.
ii. $G_0(\psi)$ *is uniquely minimized at* ψ^0.
iii. $m\left(z_{it}, z_{jt'}; \psi\right)$ *is three times continuously differentiable in* $\psi \in \Psi$ *with probability 1.*
iv. $\sup_{\psi\in\Psi} |r(z_{it}, z_{jt'}; \psi)| < \overline{r}(z_{it}, z_{jt'})$ *for some function* $\overline{r}(\cdot)$ *with* $E\left[\overline{r}(z_{it}, z_{jt'})\right] < \infty$.
v. *Define* $\tilde{v}(z_{it}, \psi) \equiv E\left[r(z_{it}, z_{jt'}, \psi)|z_{it}\right]$, *and* $\lambda(\psi) \equiv \lim_{NT\to\infty} E\tilde{v}(z_{it}, \psi)$.
v.i. $\lambda(\psi^0) = 0$ *and is differentiable at* ψ^0, *with non-singular Jacobian matrix A.*
v.ii. *The expectation* $\sup_{N,T,\psi\in\Psi} E\left[||r(z_{it}, z_{jt'}, \psi)||^2\right]$ *exists and is finite.*

The conditions listed in Assumption 6 are standard identification, continuity, differentiability, and boundedness conditions on the limiting objective function. They are analogous to the conditions required for Theorem 2 in Honoré and Powell (1994).

The asymptotic normality of our first-step estimates of ψ is given in the following theorem, the full proof of which is in Appendix A.1.

Theorem 1. *Given Assumptions 1, 2, 3, 4, 5, and 6,*

$$\sqrt{NT}\left(\hat{\psi} - \psi^0\right) \xrightarrow{d} N\left(0, A^{-1}\Omega A^{-1}\right) \tag{21}$$

as $NT \to \infty$, *where A and* Ω *are defined in equations (A4) and (A14) in the Appendix.*

Note that, if we had a perfect estimate of the conditional distributions $F(q|x)$, the differenced first-order condition $\hat{m}(z_{it}, z_{jt'}, \psi^0)$ (defined in equation (16)) would be identically zero for all values of $z_{it}, z_{jt'}$ such that $x_{it} + q_{it} = x_{jt'} + q_{jt'}$. Hence, the sampling variation in the estimate of ψ will be determined completely from the sampling variation in the non-parametric estimates of the conditional distributions $F(q|x)$ using equation (10).

3.1.2. Remarks on first step. Next, we discuss several of the assumptions we made previously, and how they may be relaxed. First, our econometric framework is parametric, in the sense that both the utility function and shock distributions are assumed to be of known parametric form. In principle, the shock distribution F_s can be given a very flexible parametric form. In our empirical work below, we consider a flexible piecewise-linear specification for F_s. We let $s_k \equiv F_s^{-1}(\tau_k)$ denote the τ_k-th quantile of the shock distribution F_s, and let κ denote the total number of quantiles to be estimated (and the corresponding quantile values by $\tau_1 < \tau_2 < \cdots < \tau_\kappa$). For any fixed κ, we approximate the distribution of the shocks F_s via a piecewise linear function tied down at the origin as well as the κ points $\{s_k, \tau_k\}_{k=1}^{\kappa}$. That is, we approximate the inverse CDF of F_s as

$$\hat{F}_s^{-1}(\tau) \equiv \begin{cases} \tau\frac{s_1}{\tau_1} & \text{if } \tau \in [0, \tau_1] \\ s_{i-1} + (\tau - \tau_{i-1})\frac{s_i - s_{i-1}}{\tau_i - \tau_{i-1}} & \text{if } \tau \in (\tau_{i-1}, \tau_i],\ i = 2, \ldots, \kappa - 1. \\ s_{\kappa-1} + (\tau - \tau_{\kappa-1})\frac{s_\kappa - s_{\kappa-1}}{\tau_\kappa - \tau_{\kappa-1}} & \text{if } \tau \in (\tau_{\kappa-1}, 1]. \end{cases} \tag{22}$$

The parameters of this specification of the shock distribution which are to be estimated are $\gamma \equiv \{s_1, \ldots, s_\kappa\}$.

Second, the deterministic accumulation assumption 2 is crucial. Specifically, the linearity of x_{t+1} in q_t is critical to the applicability of the pairwise-differencing step. If, instead, the law of motion for x were non-linear, such as $x_{it+1} = l(x_{it}, q_{it})$ (and the non-linear functional form of l were known), the derivative dl/dq would also appear in the second term of the first-order condition (12), and we would also need to match on this quantity in the pairwise-differencing step, which reduces the attractiveness of our estimator. Similarly, if we wished to introduce additional observable (and possibly time-varying) characteristics z_{it} specific to individual i and period t, these would be additional variables which we need to match upon. Therefore, the simplicity of the law of motion is a restriction that could be relaxed at some cost, and the development of those cases would be an interesting extension.

Given the deterministic accumulation equation, we could re-parameterize the problem so that the per-period utility function is a function of x_t and x_{t+1} (rather than x_t and q_t), and we take next period's stock x_{t+1} as the choice variable in period t. In that case, in order for the monotonicity assumption 2 to be obtained, it would suffice that the per-period utility function be supermodular in s_t and x_{t+1}, which has the intuitive economic interpretation that the shocks increase the marginal utility of x_{t+1}.

The independence assumption 1 is important for the feasibility of the procedure. For example, if the distribution of the shock s_{t+1} were dependent on x_t (so that the conditional distribution $F(s_{t+1}|x_{t+1}, x_t, s_t) = F(s_{t+1}|x_t)$ varies depending on x_t), then the expected value function $\mathcal{V} = E_{s_{t+1}|x_t} V(x_{t+1}, s_{t+1})$ would also be a function of x_t, and the pairwise-differencing step would require matching individuals with the same $x_{t+1} = x_t + q_t$ as well as x_t. These individuals would also have the same q_t, leaving no degrees of freedom for the estimating equation (14). On the other hand, if $F(s_{t+1}|x_{t+1}, x_t, s_t) = F(s_{t+1}|x_{t+1})$, then for a parametric specification of $F(s_{t+1}|x_{t+1}; \gamma)$, it may still be possible to use equation (14) in the first step to estimate the first-step parameters ψ.

Finally, we note that because the shock s is unobserved, we could also follow Matzkin (2003) to assume that the shock is uniformly distributed on [0,1]. Since the shock s is distributed according to $F_s(\cdot; \gamma)$, we could define $\epsilon = F_s(s; \gamma)$ and re-parameterize the utility function so that

$$U(x, s, q; \theta) = U(x, F_s(\epsilon; \gamma), q; \theta) \equiv \tilde{U}(x, \epsilon, q; \theta, \gamma). \tag{23}$$

The monotonicity assumption 3 is a natural consequence of this re-parameterization: holding x fixed, q is monotonic in ϵ.[4] With this re-parameterization, $\epsilon = F(q|x)$, and

$$\tilde{U}_q(x_{it}, F(q_{it}|x_{it}), q_{it}) = \tilde{U}_q(x_{jt'}, F(q_{jt'}|x_{jt'}), q_{jt'}) \tag{24}$$

for $x_{it} + q_{it} = x_{jt'} + q_{jt'}$.

Discussion of identification. Before proceeding to the second step, we also present some discussion of identification. Consider how the parameters γ of the shock distribution are pinned down in the pairwise-differencing step when $\theta_1 = \{\}$ (so that there are no θ_1 parameters to estimate). In order to do pairwise-differencing, we need pairs of observations (it, jt') such that

$$0 = U_q\left(x_{it}, s(F(q_{it}|x_{it}); \gamma), q_{it}\right) - U_q\left(x_{jt'}, s(F(q_{jt'}|x_{jt'}); \gamma), q_{jt'}\right). \tag{25}$$

In order for equation (25) to be a non-trivial function, observations it and jt' must satisfy two conditions: (i) $x_{it} + q_{it} = x_{jt'} + q_{jt'}$; but (ii) $x_{it} \neq x_{jt'}$ (and hence $q_{it} \neq q_{jt'}$). The question of identification then relies crucially on the existence of such pairs of individuals, which in turn depends on the model.

For a specific example, consider the following dynamic firm investment problem where the individual-specific subscript i is omitted:

$$\max_{\{q_t\}} \sum_{t=0}^{\infty} \beta^t (p_t y_t - q_t s_t - \frac{1}{2} q_t^2) \tag{26}$$

4. We thank a referee for pointing this out.

subject to

$$p_t = 1 - y_t; \quad y_t = 2x_t; \quad x_{t+1} = x_t + q_t; \quad s_t \sim iid\ N(0, \sigma^2). \tag{27}$$

The interpretation here is that x_t is capital, which is transformed into final goods y_t via the production technology $F(x_t) = 2x_t$. The inverse demand curve for final goods is given by $p_t = 1 - y_t$. The shock s_t affects the linear component of investment costs. For this shock distribution, the CDF is given by $F_s(s) = \Phi(s/\sigma)$ and quantile function by $F_s^{-1}(\tau) = \sigma\Phi^{-1}(\tau)$ where $\Phi(\cdot)$ denotes the standard normal CDF function. The standard deviation σ is the only parameter to be estimated, in this example.

This is a linear-quadratic problem, and the optimal policy function (taking the discount rate $\beta = 0.95$) is given in the Appendix by

$$\begin{aligned} q_t =& c_1 - c_2 x_t - c_3 s_t, \\ & c_1 = 0.2235, \ c_2 = 0.8942, \ c_3 = 0.1058. \end{aligned} \tag{28}$$

In the context of this example, we can discuss identification in more detail. We observe the sequences (x_{it}, q_{it}) across many agents i, where for each x_{it}, the investment q_{it} is generated from the optimal policy function (28). From these observations, we can estimate the conditional CDFs $\hat{F}_{q|x}(\cdot|\cdot)$.

For a fixed value $\overline{C}$, we consider the locus of points (x, q) such that $x + q = \overline{C}$, or, for each x, the corresponding investment is $q = \overline{C} - x$. Along this locus, for every pair $x_{it} \neq x_{jt'}$, the corresponding pair of investments $q_{it} \neq q_{jt'}$. For every pair of distinct points (x_{it}, q_{it}) and $(x_{jt'}, q_{jt'})$ on this locus, equation (16) is

$$\sigma\left[\Phi^{-1}(\hat{F}_{q|x}(\overline{C} - x_{it}|x_{it})) - \Phi^{-1}(\hat{F}_{q|x}(\overline{C} - x_{jt'}|x_{jt'}))\right] + (x_{it} - x_{jt'}) = 0 \tag{29}$$

which serves as the estimating equation for σ.

Let σ_0 denote the true value of σ. Now, note that

$$\begin{aligned} F_{q|x}(\overline{C} - x_{it}|x_{it}) =& \ P\left(c_1 - c_2 x_{it} - c_3 s_{it} \leq \overline{C} - x_{it}|x_{it}\right) \\ =& \ P\left(s_{it} \geq \tfrac{1}{c_3}\left(c_1 + (1 - c_2)x_{it} - \overline{C}\right)\right) \\ =& \ \Phi\left(-\tfrac{1}{\sigma_0 c_3}\left(c_1 + (1 - c_2)x_{it} - \overline{C}\right)\right). \end{aligned} \tag{30}$$

Therefore equation (29) further simplifies into (using the relation $1 - c_2 = c_3$):

$$\left(x_{it} - x_{jt'}\right) = \frac{\sigma}{\sigma_0}\left(x_{it} - x_{jt'}\right), \tag{31}$$

which holds if and only if $\sigma = \sigma_0$. This clearly shows that σ is identified.

Identification in our framework differs from those in other papers in the literature. For example, Magnac and Thesmar (2002), Pesendorfer and Schmidt-Dengler (2008), and Aguirregabiria (2005) all consider dynamic discrete-choice models, and focus on the non-parametric identification of the utility functions using the observed choice probabilities. We focus on the case where the agents' choice variable (q_t) has a continuous component, and where the utility function takes an assumed parametric form. Because of these features, identification issues are different in our setting. The continuous component of agents' choices allows us to use pairwise-differencing methods to identify quantiles of the shock distribution, while the identification of the utility function is facilitated by parametric assumptions.

Alternative estimation methods. Our pairwise-difference-based estimation method is also different from Euler equation-based methods, which are often used to estimate dynamic optimization models. On the one hand, our approach accommodates dynamic optimization models in which agents' choices are both continuous and discrete (as in our empirical example below), for which conventional Euler equation methods are either not applicable or difficult. Furthermore, our approach also accommodates shocks that are observed by agents at the time they make their decisions but unobserved to the econometrician. Conventional Euler equation-based estimation methods generally have difficulties accommodating unobserved shocks because the estimating moment conditions are derived from the rational expectations implication that deviations between predicted and observed actions are orthogonal to any information available at time t, which includes all state variables which affect an agent's period t choice. Therefore, to form the sample analogues of these orthogonality conditions, the econometrician needs to know the value of all the state variables (including the shocks) at times t and $t+1$. Pakes (1994, pp. 188–189) provides a more thorough discussion.

On the other hand, our approach works best under Assumption 2, which represents a restriction on the law of motion for the state variable x_t which is not required for Euler equation methods. Hence, for models where this assumption does not hold, Euler equation methods may be a more attractive estimation option.

3.2. *Second step*

Not all model parameters can be identified from the first-step pairwise-differencing approach. In the second step, we use the first-order condition again to derive moment restrictions to estimate utility parameters in θ which were not in the subset θ_1 estimated in the first step, denoted as $\theta_2 \equiv \{\theta \setminus \theta_1\}$.

Given $\hat{\gamma}$ and $\hat{\theta}_1$, respectively, the shock distribution parameters and the subset of the utility function parameters which were estimated in the first step, define the first-order condition for observation (i, t) with non-zero investment level as follows:

$$\begin{aligned} &h_{it}\left(x_{it}, q_{it}; \hat{\psi}, \theta_2, \hat{F}_{q|x}(\cdot|\cdot)\right) \\ &\equiv U_q\left(x_{it}, s(\hat{F}(q_{it}|x_{it}); \hat{\gamma}), q_{it}; \hat{\theta}_1, \theta_2\right) + \beta\mathcal{V}'\left(x_{it}+q_{it}; \hat{\psi}, \theta_2, \hat{F}_{q|x}(\cdot|\cdot)\right). \end{aligned} \tag{32}$$

In what follows, we will use $\hat{F}_s$ as shorthand for $F_s(\cdot; \hat{\gamma})$.

Assume that we are able to compute the expected value function $\mathcal{V}(x_{it}; \psi, \theta_2)$ for every set of parameters ψ and θ_2 (we delay discussion of how this can be done until later). Because of the sampling error from estimating θ_1, γ and $F(q|x)$ in the first step, the first-order condition $h_{it}\left(x_{it}, q_{it}; \hat{\psi}, \theta_2, \hat{F}_{q|x}(\cdot|\cdot)\right)$ need not be identically zero, even at the true parameter vector θ_0. Therefore, we estimate θ_2 via a least squares procedure:

$$\hat{\theta}_2 = \text{argmin}_{\theta_2} \frac{1}{NT} \sum_{i=1}^{N} \sum_{t=1}^{T} I_{it} \cdot \left[h_{it}(x_{it}, q_{it}; \hat{\psi}, \theta_2, \hat{F}_{q|x}(\cdot|\cdot))\right]^2. \tag{33}$$

As in the first step, we can only include observations with non-zero investment ($q \neq 0$) in the objective function. Both steps of our estimation procedure are based on agents' first-order conditions, and thus only use the observations where $q_{it} \neq 0$. The observations with $q_{it} = 0$ are employed only in the construction of the conditional distributions $\hat{F}(q|x)$ [cf. equation (10)].[5]

3.2.1. Computing the expected value function by simulation. The expected value function $\mathcal{V}(\cdot; \psi, \theta_2)$ does not have a closed-form solution and needs to be evaluated numerically. Standard numerical dynamic programming methods for problems with both discrete and continuous controls, as described in Rust (1996) and Judd (1998), can be difficult since they involve solving for the optimal policy function $q(x, s)$ at every point (x, s) in the state space.

When the datasets available to the researcher are large (as in the dataset we consider later), an attractive alternative in the spirit of Hotz and Miller (1993) is available to avoid numerical computation of the dynamic programming problem, in which the value function is computed by a forward integration procedure. This procedure exploits the representation of the value function at time t as the expected discounted sum of future utilities [cf. equation (7)] which underlies numeric dynamic programming algorithms. Hotz and Miller (1993) recognize that, given enough data, and a particular parametric form of the per-period utility function $U(x, s, q; \theta)$, the expectation over future states in equation (7) can be represented as forward integration over the observed conditional probabilities $\hat{F}(q|x)$ (cf. equation (17) in Hotz and Miller, 1993).

Under the conditional independence assumption 1, this approach can be used in the case where agent i's control variable is continuous. More precisely, the agent's expected value function at a particular initial point x_1 is approximated as

$$\begin{aligned}\mathcal{V}\left(x_t; \hat{\psi}, \theta_2, \hat{F}_{q|x}(\cdot|\cdot)\right) = \int\int\cdots\int\Bigg\{\left[\sum_{z=t}^{T}\beta^{z-t}U\left(x_z, s_z, \hat{F}^{-1}_{q|x_z}(F_s(s_z)); \hat{\theta}_1, \theta_2\right)\right] \\ + \beta^T CV(x_{T+1})\Bigg\} dF\left(s_t; \hat{\gamma}\right) dF\left(s_{t+1}; \hat{\gamma}\right)\cdots dF\left(s_T; \hat{\gamma}\right).\end{aligned} \tag{34}$$

Here, $CV(x_{T+1})$ denotes the continuation value function, when the state after T periods is x_{T+1}. The sequence of stocks x_z is given by the initial condition x_t and $x_z = x_{z-1} + \hat{F}^{-1}_{q|x_{z-1}}(F_s(s_{z-1}))$ for $z = t+1, \ldots, T$.

More succinctly, let $\{\tau\} = \{\tau_t, \tau_{t+1}, \ldots, \tau_T\}$ denote a sequence of *i.i.d.* $U[0, 1]$ random variables. The expected value function can be written as

$$\begin{aligned}&\mathcal{V}\left(x_t; \hat{\psi}, \theta_2, \hat{F}_{q|x}(\cdot|\cdot)\right) \\ &= E_{\{\tau\}}\left\{\sum_{z=t}^{T}\beta^{z-t}U\left(x_z, s\left(\tau_z; \hat{\gamma}\right), \hat{F}^{-1}_{q|x_z}(\tau_z); \hat{\theta}_1, \theta_2\right) + \beta^T CV(x_{T+1})\right\}.\end{aligned} \tag{35}$$

In the above expression, given the starting value x_t, the subsequent sequence of stocks $x_{t+1}, x_{t+2}, \ldots$ is related to the uniform random draws τ's by the relation $x_z = x_{z-1} + \hat{F}^{-1}_{q|x_{z-1}}(\tau_{z-1})$. In our implementation below, we treat the continuation value function $CV(x_{T+1})$ as a nuisance parameter, and assume that it is approximated by a flexible finite-order polynomial in x_{T+1}.

In practice, the multidimensional integration involved in computing the expected value function [equations (34) or (35)] presents computational challenges, and so we simulate the

5. The square norm in (33) is chosen for convenience. Other norms, such as absolute deviation, may also be used. Furthermore, weighting schemes could be introduced to improve the efficiency of the estimation procedure. We have not considered these alternative possibilities.

expected value function by following Hotz *et al.* (1994). Let S denote the number of simulation draws. Using the parameters $\hat{\psi}$ and the conditional distributions $\hat{F}_{q|x}$ estimated from the first step, $\mathcal{V}\left(x_1; \hat{\psi}, \theta_2, \hat{F}_{q|x}(\cdot|\cdot)\right)$ [using equation (35)] can be simulated by

$$\begin{aligned}&\mathcal{V}^S\left(x_t; \hat{\psi}, \theta_2, \hat{F}_{q|x}(\cdot|\cdot)\right)\\ &= \frac{1}{S}\sum_{l=1}^{S}\left\{\left[\sum_{z=t}^{T}\beta^{z-t}U\left(x_z^l, s(\tau_z^l; \hat{\gamma}), \hat{F}^{-1}_{q|x_z^l}(\tau_z^l); \hat{\theta}_1, \theta_2\right)\right] + \beta^T CV(x_{T+1}^l)\right\}\end{aligned} \tag{36}$$

where

- $\tau_z^l, \; l = 1, \ldots, S, \; z = t, \ldots, T$ are *i.i.d.* $U[0, 1]$.
- $x_z^l = \begin{cases} x_1 & \text{for } z = t \\ x_{z-1}^l + q\left(x_{z-1}^l, s(\tau_z^l; \hat{\gamma})\right) & \text{for } z = t+1, \ldots, T+1. \end{cases}$

In order to implement the second-step estimator, we must also compute the derivative of the expected value function. This is approximated by a numeric finite-difference method:

$$\mathcal{V}^{S'}(x_{it}; \hat{\psi}, \theta_2, \hat{F}_{q|x}(\cdot|\cdot)) \approx \frac{\mathcal{V}^S(x_{it} + \Delta; \hat{\psi}, \theta_2, \hat{F}_{q|x}(\cdot|\cdot)) - \mathcal{V}^S(x; \hat{\psi}, \theta_2, \hat{F}_{q|x}(\cdot|\cdot))}{\Delta} \tag{37}$$

for Δ small. By plugging in equation (37) for $\mathcal{V}'\left(x_{it} + q_{it}; \hat{\psi}, \theta_2, \hat{F}_{q|x}(\cdot|\cdot)\right)$ into equation (32), we can estimate θ_2 by minimizing the objective function (33).

3.2.2. Asymptotic theory for second step. In this section, we present the limit distribution for the second-step estimator $\hat{\theta}_2$. In deriving the asymptotics, we ignore the approximation error in simulating the expected value function (as well as its derivative), and treat the expected function $\mathcal{V}(x_{it}; \hat{\psi}, \theta_2, \hat{F}_{q|x}(\cdot|\cdot))$ as a known function for all $(\hat{\psi}, \theta_2, \hat{F}_{q|x}(\cdot|\cdot))$.[6] The second-step estimator $\hat{\theta}_2$ solves the sample score function

$$J_{NT}(\hat{\theta}_2) \equiv \frac{1}{NT}\sum_{i=1}^{N}\sum_{t=1}^{T}\overline{h}\left(x_{it}, q_{it}; \hat{\psi}, \hat{\theta}_2, \hat{F}_{q|x}(\cdot|\cdot)\right) = 0 \tag{38}$$

where

$$\begin{aligned}&\overline{h}\left(x_{it}, q_{it}; \hat{\psi}, \theta_2, \hat{F}_{q|x}(\cdot|\cdot)\right)\\ &\equiv I_{it} * \left[U_q\left(x_{it}, s\left(\hat{F}(q_{it}|x_{it}); \hat{\gamma}\right), q_{it}; \hat{\theta}_1, \theta_2\right) + \beta\mathcal{V}^{S'}\left(x_{it} + q_{it}; \hat{\psi}, \theta_2, \hat{F}_{q|x}(\cdot)\right)\right] *\\ &\quad \frac{\partial}{\partial\theta_2}\left[U_q\left(x_{it}, s\left(\hat{F}(q_{it}|x_{it}); \hat{\gamma}\right), q_{it}; \hat{\theta}_1, \theta_2\right) + \beta\mathcal{V}^{S'}\left(x_{it} + q_{it}; \hat{\psi}, \theta_2, \hat{F}_{q|x}(\cdot)\right)\right].\end{aligned} \tag{39}$$

The notation $\hat{F}_{q|x}(\cdot)$ denotes the whole set of estimated conditional quota distributions, estimated as in equation (10). By including the entire conditional distribution $\hat{F}_{q|x}(\cdot|\cdot)$ as an argument, we recognize that the expected value function $\mathcal{V}(x_{it+1})$ [cf. equation (35)] depends

6. For the simulation-based approximation of the expected value function, we require that the number of simulation draws S increases quickly enough as $N \to \infty$ so that variation due to the simulation itself is small enough and does not affect the asymptotic variance. A sufficient condition for the asymptotic variance to be unaffected from simulation error is that $S/\sqrt{N} \to \infty$ (Gourieroux and Monfort, 1996).

on the entire set of functions $\hat{F}_{q|x}^{-1}(\cdot)$, and not just on any one of these functions evaluated at a particular quantile.

The asymptotics of the second step are complicated by dependence of the second-step objective function on the entire estimated function $F_{q|x}$. In characterizing the asymptotic distribution of $\hat{\theta}_2$, we follow Newey (1994). The regularity conditions required for deriving the asymptotic result of the second-step estimator are collected in the following assumption. Let F be the shorthand for $F_{q|x}(\cdot|\cdot)$, and $\hat{F}$ and F^0 denote the estimated and true values for this function, respectively. Also, let $||\cdot||$ denote a functional norm. Let $P_2 \equiv \dim(\theta_2)$, and define

$$H_0(\theta_2) \equiv E\mathbf{1}(q \neq 0)\left[h(x, q; \psi^0, \theta_2, F_{q|x}(\cdot|\cdot))\right]^2 \tag{40}$$

as the limit objective function of the second-step estimator.

Assumption 7. *1. $\theta_2 \in \Theta_2$, a compact subset of $\mathbf{R}^{P_2}$, and true value $\theta_2^0 \in int(\Theta_2)$.*

2. $H_0(\theta_2)$ is uniquely maximized at θ_2^0.

3. $h(x, q; \psi, \theta_2, F)$ is twice continuously differentiable in θ_2 and ψ with probability 1. Both the function and its derivatives are uniformly bounded by an integrable function.

4. The Hessian $\overline{A}$ of $H_0(\theta_2)$ with respect to θ_2 is non-singular at θ_2^0.

5. $\overline{h}(z, F) \equiv \overline{h}(x, q, \psi^0, \theta_2^0, F)$ is Fréchet-differentiable in F at F^0; that is, for all z, and for all F with $||F - F^0||$ small enough, there exists a linear operator $D(z, F)$ such that

$$|\overline{h}(z, F) - \overline{h}(z, F^0) - D(z, F - F_0)| \leq b(z)||F - F^0||^2.$$

Moreover, $\sqrt{NT} \cdot Eb(z)||\hat{F} - F^0||^2 = o_p(1)$.

6. Stochastic equicontinuity:

$$\frac{1}{\sqrt{NT}} \sum_i \sum_t \left[D(z_{it}, \hat{F} - F^0) - ED(z, \hat{F} - F^0)\right] = o_p(1). \tag{41}$$

7. There exists a function $\alpha(z)$ satisfying $E[a(z_i)] = 0$ and $E[||\alpha(z_i)|^2] < \infty$ such that

$$\sqrt{NT}\left[ED(z, \hat{F} - F^0) - \frac{1}{NT}\sum_i \sum_i \alpha(z_{it})\right] = o_p(1). \tag{42}$$

These assumptions are drawn from Newey (1994). Conditions 1–4 are standard conditions for consistency and asymptotic normality. Conditions 5–7 are useful for characterizing the effect of the estimated function $\hat{F}_{q|x}$ on the estimates of θ_2. Condition 5 assumes that the remainder term from a first-order functional Taylor expansion of $\overline{h}(z, F)$ around F^0 is asymptotically negligible. Condition 7 requires that the expectation of the Fréchet derivative $D(z, \hat{F} - F^0)$ at F^0 can be approximated by a sample average of $\alpha(z_{it})$, for some function $\alpha(\cdot)$. The form of this "influence function" $\alpha(\cdot)$ is discussed below.

Theorem 2. *Given Assumptions 1 to 7, the sample score function satisfies a central limit theorem:*

$$\sqrt{NT} J_{NT}(\theta_2^0) \xrightarrow{d} N(0, \overline{\Omega}). \tag{43}$$

In addition,

$$\sqrt{NT}\left(\hat{\theta}_2-\theta_2^0\right)\xrightarrow{d}N\left(0,\overline{A}^{-1}\overline{\Omega}\,\overline{A}^{-1}\right) \tag{44}$$

as $N\to\infty$, where $\overline{A}$ is the Hessian matrix defined in assumption 7.7, and

$$\overline{\Omega}=limVar\left(\frac{1}{\sqrt{NT}}\sum_{i,t}\alpha\left(z_{it}\right)+B\cdot\sqrt{NT}\left(\hat{\psi}-\psi_0\right)\right). \tag{45}$$

In the above, $B=\frac{\partial}{\partial\psi}E\overline{h}\left(x_{it},q_{it};\psi^0,\theta_2^0,F_{q|x}^0\left(\cdot|\cdot\right)\right)$.

A key element in the asymptotic distribution is the influence function $\alpha(z)$, which describes the effect of the estimation of $F_{q|x}$ on the estimates of θ_2. The form of $\alpha(z)$ is model-specific and follows Newey (1994). To summarize, under the above assumptions, if for any one-dimensional parametric subpath F_η of the function $F_{q|x}\left(\cdot|\cdot\right)$ estimated in the first stage, the "pathwise derivative" of $h(z,F)$ with respect to F can be written as

$$\frac{\partial}{\partial\eta}E\overline{h}\left(z_{it},F_\eta\right)=E\alpha\left(z_{it}\right)S_\eta\left(z_{it}\right), \tag{46}$$

where $S_\eta\left(z_{it}\right)$ is the score function of the joint density of the data $z_{it}=(x_{it},q_{it})$, then $\alpha(\cdot)$ is the desired influence function. In the proof of Theorem 2, we use equation (46) to calculate $\alpha(z)$ explicitly for a simpler two-period version of the model. We were not able to derive a general analytic functional form of $\alpha(z)$ for the more general model because of the recursive way the estimated $\hat{F}$ enters the expected value function [cf. equation (35)]. However, the influence function for the simpler model adequately illustrates the form that the influence function will take for the full model. Moreover, since in practice we use the bootstrap to estimate the standard errors for the two-step estimator, the characterization of the asymptotic distribution is mostly useful for justifying the use of the bootstrap, so that the analysis of the simpler model suffices.

At the true values of ψ, θ_2, and $F(q|x)$, the first-order condition (32) is identically zero for all values of (x_{it},q_{it}) which are optimally chosen. Hence, the second-step estimation procedure introduces no source of sampling variation beyond that which arises from the first-step estimation of ψ and the conditional distributions $F(q|x)$.

In principle, given the parametric assumptions on $F_s(\cdot;\gamma)$, the parameters θ and γ could be jointly estimated in the second step, without requiring the pairwise-differencing first step. However, by estimating θ_1 and γ in the first step, we reduce the number of parameters that must be estimated in the second step. Since the second step potentially involves numeric dynamic programming in order to recover the value function, reducing the dimensionality of the parameter space also reduces significantly the number of times that the value function must be computed, thereby reducing the computational burden. Such a "two-step" approach was also taken in Rust's (1987) dynamic discrete-choice model of bus engine replacement, in which the parameters describing the mileage Markov transition matrix were estimated in a first step to reduce the computational burden in the second step, which involved computationally intensive value function iteration (e.g. Rust *et al.* 2002).

4. EMPIRICAL ILLUSTRATION: MARKETS FOR MILK PRODUCTION QUOTA

As an illustration of our methodology, we estimate a dynamic trading model of the milk production quota market. In Ontario, Canada, milk production is controlled via production quotas which grant holders the right to produce a certain quantity of milk per year. Since 1980, in the province of Ontario these quotas have been traded among dairy farmers in monthly double auctions administered by the Dairy Farmers of Ontario (DFO) (cf. Biggs, 1990). This paper analyses data from the 11 auctions between September 1997 and July 1998. Our goal is to estimate the parameters of agents' utility functions, and the distribution of the unobserved state variables, using the two-step pairwise-differencing methodology described earlier.

Each quota exchange is a double-auction market. All producers who wish to sell quota submit offers to the exchange indicating that they have a certain volume of quota for sale and at a certain minimum price per unit. Producers who wish to buy quota submit bids to the exchange indicating that they would like to buy a certain volume of quota and that they are willing to pay a specific maximum price per unit. Units are traded at a market clearing price (MCP) at which the total quantity demanded (approximately) equals the total quantity supplied.

In order to fit the milk-quota trading market into our dynamic framework, we consider a dynamic, forward-looking model of the quota demand and supply process, in which each individual trader faces a dynamic optimization problem. Timing is as follows. At the beginning of month t, trader i owns x_{it} units of production quota. She experiences a shock s_{it} and must decide the amount of quota q_{it} to trade at any price p_t. Generally, the optimal amount is given by a function $q\left(x_{it}, s_{it}, p_t\right)$ which takes values in $(-\infty, \infty)$. For positive values of q, this can be interpreted as a demand function, and when negative it can be interpreted as a supply function. The amount actually transacted would be $q\left(x_{it}, s_{it}, p_t^*\right)$, where p_t^* denotes the realized market-clearing price for period t.

An important simplifying assumption that we make is that the market-clearing price p_t^* is taken as given and known by bidders when they are deciding how much quota q_{it} to buy. This assumption is consistent with the dynamic competitive equilibrium path of a continuum market, on which agents will have perfect foresight about the sequence of market-clearing prices, even though at the individual trader level there is uncertainty about the shocks received by other traders. As a result, equilibrium strategies in this market can be characterized as optimal policies of a non-stationary dynamic optimization problem solved by each trader individually. Because agents' quota decisions in period t will depend on p_t^*, the market-clearing price in period t, which we model as a deterministic time-varying covariate, the empirical model is non-stationary, which differs from the stationary problem used in the previous sections in describing our estimation procedure.[7]

Specifically, we model each trader i as choosing a sequence $\{q_{it}\}$ to maximize the expected discounted present value of its utility from its milk-quota trading operations:

$$\max_{\{q_{it}\}_t} \mathbb{E} \sum_{t=0}^{\infty} \beta^t U_t\left(x_{it}, s_{it}, q_{it}, p_t^*; \theta\right) \tag{47}$$

subject to

$$x_{it+1} = x_{it} + q_{it}; \quad s_{it} \sim F_s, i.i.d. \text{ over } t; \quad p_0^*, p_1^* \ldots \text{known}. \tag{48}$$

Note the t subscript on the per-period utility function, which emphasizes that the dynamic problem is non-stationary due to the presence of the market-clearing prices. The expectation

7. In principle, if we observed many more months of data, we could consider a stationary problem in which the evolution of the monthly market-clearing prices could be estimated directly from the data. Estimation would be more complicated, as we would also need to match on p_t (in addition to $x_t + q_t$) in the first stage, and then we also need to take draws of the price process in simulating the value function for the second stage. We do not undertake this extension in the empirical application because we only have 11 observations of the price process.

is over the sequences of s_{it} and x_{it} induced by trader i's chosen sequence $\{q_{it}\}$. Each trader i's optimal policy in period t is given by a period-specific function $q_t(x_{it}, s_{it})$, which satisfies Bellman's equation:

$$q_t(x_{it}, s_{it}) = \text{argmax}_q U\left(x_{it}, s_{it}, q, p_t^*; \theta\right) + \beta \mathcal{V}_{t+1}(x_{it} + q; \psi, \theta_2) \tag{49}$$

where

$$\mathcal{V}_{t+1}(x_{it} + q; \psi, \theta_2) \equiv \mathbb{E}_{s_{it+1}} V_{t+1}(x_{it} + q, s_{it+1}; \psi, \theta_2). \tag{50}$$

Accommodating non-stationarity in our estimation procedure requires several changes from the procedure presented in the first part of this paper. First, because agents' policy functions will be period-specific in a non-stationary problem, we estimate the conditional quota purchase distributions $F_{q|x,t}$ [cf. equation (10)] separately for each period t. Second, because the expected value functions are no longer time-homogeneous in a non-stationary problem, we can no longer match agents across periods in the pairwise-differencing step. As a result, the objective function used in this step is

$$\min_{\theta_1,\gamma} \frac{1}{(N)^2 T} \sum_{t=1}^{T} \left[\sum_{i=1}^{N} \sum_{j=1}^{N} \left\{ \frac{1}{h_1} K\left(\frac{(x_{it} + q_{it}) - \left(x_{jt} + q_{jt}\right)}{h_1} \right) \cdot I_{it} I_{jt} \cdot \left[U_q\left(x_{it}, s\left(\hat{F}(q_{it}|x_{it}); \gamma\right), q_{it}; \theta_1\right) - U_q\left(x_{jt}, s\left(\hat{F}\left(q_{jt}|x_{jt}\right); \gamma\right); \theta_1\right)\right]^2 \right\}\right] \tag{51}$$

which differs from the objective function for the stationary case [equation (9)] because we do not match across agents (i, j) in different periods.

Finally, given that we only observe 11 periods of data, we assume that agents solve a finite-horizon model with $T = 11$ but allow the continuation value of the problem (after the 11th month) to depend on x_{T+1}, the stock that a given trader has after the first 11 months. More specifically, for months $t = 1, \ldots, T$, we simulate the expected value function as

$$\mathcal{V}_t^S(x_t; \psi, \theta_2) = \frac{1}{S} \sum_{l=1}^{S} \left\{ \left[\sum_{z=t}^{T} \beta^{z-t} U\left(x_z^l, s(\tau_z^l; \gamma), \hat{F}^{-1}_{q|x_z^l, z}(\tau_z^l); \theta\right)\right] + \beta^{T+1-t} CV(x_{T+1}^l) \right\} \tag{52}$$

where

- $\tau_z^l,\ l = 1, \ldots, S,\ z = t, \ldots, T$ are *i.i.d.* $U[0, 1]$.
- $x_z^l = \begin{cases} x_t & \text{for } z = t \\ x_{z-1}^l + q\left(x_{z-1}^l, s(\tau_z^l; \gamma)\right) & \text{for } z = t+1, \ldots, T+1. \end{cases}$
- the continuation value function is a flexible (firth-order) polynomial in x_T:

$$CV(x_{T+1}) = \sum_{j=1}^{5} \eta_j \cdot x_{T+1}^j. \tag{53}$$

We estimate the polynomial coefficients $\eta_1, \ldots, \eta_5$ are jointly with θ_2 in the second step of our procedure.

4.1. *Data: summary statistics*

Summary statistics are presented in Table 1. The trading unit for quota is expressed in kilograms of butterfat, and 1 kg of quota purchased on the exchange allows a producer to ship 1 kg of

TABLE 1
Summary statistics for each quota exchange

(A) Year	(B) Month	(C) MCP*	(D) #non-** participants	(E) #participants:	(F) *of which* #sellers	(G) (%success)‡	(H) #buyers	(I) (%success)§	(J) #zero*** bids	(K) #non-zero† bids
1997	9	15999.00	2065	509	219	63.4%	290	59.7%	2377	197
1997	10	15250.00	2178	396	248	57.2%	148	84.5%	2445	129
1997	11	15025.00	2103	471	253	84.6%	218	82.6%	2497	77
1997	12	15510.00	2155	419	163	94.4%	256	46.5%	2428	146
1998	1	16150.00	2146	428	126	91.2%	302	39.1%	2379	195
1998	2	16360.00	1995	579	182	85.7%	397	53.9%	2365	209
1998	3	16501.00	2042	532	214	93.0%	318	75.8%	2482	92
1998	4	15499.00	2127	447	212	27.4%	235	94.0%	2406	168
1998	5	14500.00	1999	575	247	52.2%	328	98.5%	2451	123
1998	6	14500.25	1949	625	178	86.5%	447	72.5%	2427	147
1998	7	15025.00	2128	446	105	88.6%	341	44.0%	2371	203

*Canadian dollars per kilogram of butterfat per day.
**Computed as 2574−(E), where 2574=#producers who participated in at least one quota exchanges between 9/1997 and 7/1998.
***i.e. number of bidders who submitted zero bids, computed as (D)+(G)*(F)+(I)*(H).
†i.e. number of bidders who submitted non-zero bids, computed as 2574−(J).
‡% of sellers who sold in the exchange, i.e., who submitted bids at or below the MCP.
§% of buyers who bought in the exchange, i.e., who submitted ask prices at or above the MCP.

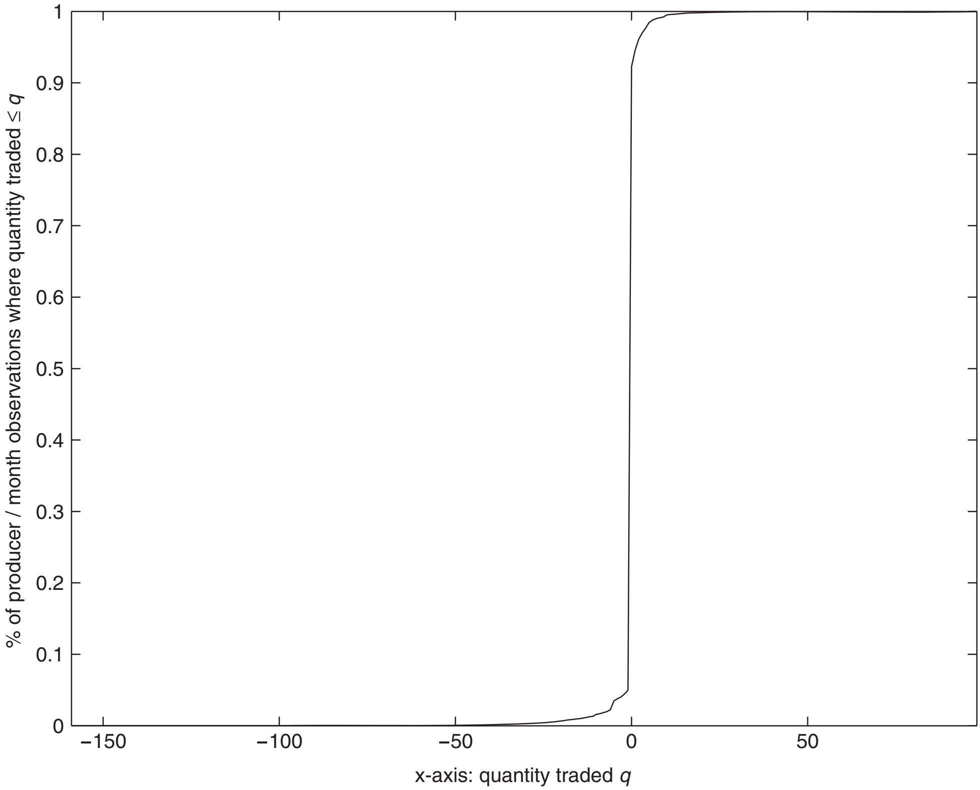

FIGURE 1
Empirical CDF of quantity traded per trader/month; x-axis: quantity traded q; y-axis: % of producer/month observations where quantity traded $\leq q$

butterfat per day, in perpetuity, for as long as the unit of quota is held.[8] Over the 11 exchanges, we observe the bids placed by 2574 distinct producers. For each trader, we have data on her total quota stock in September 1997 (the first month in our sample), as well as her purchases/sales of quota in each subsequent month, which we used to construct her total quota for each month.

Column E in Table 1 shows that a large number of sellers and buyers participate in each exchange, which suggests that there may not be much scope for strategic behaviour, which is not accommodated in our empirical model.

Across all auctions, column J shows that about 90% of the producers submit zero bids. In our empirical application, given the assumption that traders have perfect foresight about the market-clearing prices, a zero bid is attributed to two events: (i) non-participation in an auction (which, on average, is attributed to 2000 potential bidders in each auction); and (ii) submission of a non-zero bid, but not consummating a sale because it was either a sell price higher than the MCP, or a buy price lower than the MCP. In Figure 1, we present the empirical CDF of the

8. Prior to September 1997, a unit of quota conferred on its owner the right to produce milk containing 1 kg of butterfat *per year*. In September 1997, however, the trading unit for quota was redefined in kilograms of butterfat per day.

quantity traded per month, across all the monthly auctions. This shows clearly that over 90% of the observations are zero bids. Despite the large numbers of zero bids, however, columns G and I of Table 1 also indicate that each bidder's chance of getting their order filled (i.e., submitting selling bids below the MCP, or submitting buying bids above the MCP) is quite high across most of the exchanges.

Conditional on trading, there is a wide dispersion of trade amounts, ranging from about –150 to 100 units of quota. Given this large dispersion, we model a producer's choice of q, conditional on trade, as a continuous variable, even though trade is actually restricted to integer units.

4.2. *Utility function parameterization*

We assume an exponential constant absolute risk aversion (CARA) form for the utility function:

$$U(w_{it}) = -\exp\left(-rw_{it}\right), \tag{54}$$

and the following linear specification for trader i's period t payoff:

$$w_{it} = x_{it} \cdot s_{it} - p_t \cdot q_{it} - K \cdot \mathbf{1}\left(q_{it} \neq 0\right). \tag{55}$$

The per-period payoffs for each trader are as follows. Each period, trader i receives some profits $x_{it} \cdot s_{it}$ from producing and selling milk under its current stock of quota, but pays an amount $p_t \cdot q\left(x_{it}, s_{it}, p_t\right)$ to acquire additional quota. Furthermore, she incurs a fixed adjustment cost K which is associated with any non-zero transaction of quota (and the magnitude of which is not dependent on the amount of quota transacted): this would accommodate not only bidding costs but also general fixed costs associated with expanding/contracting the scale of milk production and is required to rationalize the large number of zero bids, as summarized in column J of Table 1.[9] Given this specification, s_{it} can be interpreted as stochastic production shocks which affect a trader's profits from his milk production.

In this parameterization, the only parameters identified in the first pairwise-differencing step are γ, the parameters of the shock distribution F_s. Too see this, note that U_q, the marginal utility, is equal to $-pr\exp[-r(xs - pq - K)]$ for our exponential specification. When we difference the marginal utilities for agents i and j, however, the pairwise-differencing estimating equation (14) becomes

$$-pe^{rK}\left[\exp(-r\left(x_i s_i(\gamma) - pq_i\right)) - \exp(-r\left(x_j s_j(\gamma) - pq_j\right))\right], \tag{56}$$

where $s_i\left(\gamma\right) \equiv s\left(F\left(q_i|x_i\right); \gamma\right)$. The constant proportion pe^{rK} does not have any sampling variation, and hence is not identified in the first-stage estimation using equation (14). Furthermore, from inspection of the above equation, we see that it holds if and only if

$$(x_i s_i(\gamma) - pq_i) = (x_j s_j(\gamma) - pq_j),\ \forall(i, j):\ x_i + q_i = x_j + q_j \tag{57}$$

which involves only the shock distribution parameters γ, and none of the utility parameters (r, K). Hence, condition (56) provides no information to pin down r and K, which must be estimated in the second step.

9. We may wish to allow the adjustment cost K to be a trader-specific fixed effect which varies across traders but is fixed across time. This could help explain the large number of $q_{it} = 0$ observations in the data. In principle, our estimation procedure can accommodate this, as we would amend the pairwise-differencing step to only match on $x_t + q_t$ using the across-time observations for each trader. While this is feasible in applications where we observe a long time series for each agent, it is not practical here, because we only observe 11 monthly observations for each trader.

Accordingly, in our empirical work, our first-stage estimator minimizes the following objective function, which is a least squares version of equation (57):

$$\min_{\theta_1,\gamma} \frac{1}{(N)^2 T} \sum_{t=1}^{T} \Bigg[\sum_{i=1}^{N} \sum_{j=1}^{N} \Bigg\{ \frac{1}{h_1} K \left(\frac{(x_{it}+q_{it}) - \left(x_{jt}+q_{jt}\right)}{h_1} \right) \cdot I_{it} I_{jt} \\ \left[(x_{it} s_{it}(\hat{\gamma}) - p_t q_{it}) - (x_{jt} s_{jt}(\hat{\gamma}) - p_t q_{jt}) \right]^2 \Bigg\} \Bigg] \tag{58}$$

with $\hat{s}_{it}(\gamma) \equiv s(\hat{F}_{t,q|x}(q_{it}|x_{it}); \gamma)$.

While we have derived the asymptotic covariance matrix for our estimator in Theorems 1 and 2 above, in practice it is fairly tedious and involved to compute. Therefore, in the empirical implementation, we obtained standard errors for our estimates using a bootstrap resampling procedure. The derivation of the asymptotic distribution in Theorems 1 and 2 serves to validate the use of bootstrap methods for our estimator.

For each specification, we used the bootstrap as follows: we resampled (with replacement) sequences from the dataset, and re-estimated the model for each resampled dataset. The reported bootstrap confidence intervals are therefore the empirical quantiles of the distribution of parameter estimates obtained in this fashion.

4.3. *Estimation results*

Log-normal shock distribution parameterization. First, we present results from a tightly parameterized model, assuming a log-normal specification for F_s, whereby $\log s \sim N\left(\mu, \sigma^2\right)$. The parameter estimates are shown in Tables 2 and 3.

These magnitudes imply that the mean (and median) shock is 6.928. Given the specification of the agents' payoffs [equation (55)], this can be interpreted as the monthly return from a unit of quota (in 1986 thousands of Canadian dollars). At a price of about $11,000 (again in 1986 CAD) per unit of quota, these magnitudes imply that the median producer would "recoup" her investment in less than 2 months: this seems quite an unrealistically small figure. The estimates of K and r indicate, respectively, very small adjustment costs (around 30 cents) and a very low level of risk aversion. In the top graph of Figure 2, we present our estimate of the implied period 1 (September 1997) policy function $q_1(x, s)$ for the log-normal distribution results. The policy function is estimated using equation (17).

Piecewise-linear shock distribution parameterization. Second, we present results using a more flexible piecewise-linear form for the shock distribution F_s, as described in equation (22) above. In the first step, we jointly estimated the 0.15, 0.25, 0.5, 0.75, and 0.85 quantiles for F_s. The estimated CDF is graphed in Figure 3. The median shock is estimated to be about 1.24, implying (using the same reasoning as in the previous paragraph) that the median trader recoups his investment in about nine months: this appears more realistic than the estimate obtained from the log-normal parameterization.[10]

In the bottom graph of Figure 2, we present our estimate of the implied period 1 (September 1997) policy function $q_1(x, s)$ for the F_s with linear interpolation estimated in the first step and plotted in Figure 3. The estimate of K implies that the magnitude of fixed adjustment costs are $119.70, which is much higher than the estimates obtained using the log-normal specification. The estimate of r, the coefficient of absolute risk aversion, remains very small (0.0072).

10. We also considered another specification allowing F_s to vary across periods. However, we found that the covariates had little effect and left the results virtually unchanged. Therefore, we do not report those results.

TABLE 2
Parameter estimates: log-normal specification for F_s $\log s \sim N\left(\mu, \sigma^2\right)$

	Estimate	Standard error*
K	0.0003	0.6750
r	0.0320	0.0101
μ	−0.6706	0.0772
σ	2.2830	0.1268

*Obtained via bootstrap resamples.

TABLE 3
Parameter estimates: flexible piecewise-linear specification for F_s

	Estimate	Standard errors*
Step 1 parameters		
$F_s^{-1}(0.15)$	0.0028	0.0064
$F_s^{-1}(0.25)$	0.6994	0.2761
$F_s^{-1}(0.50)$	1.2400	1.2014
$F_s^{-1}(0.75)$	1.3344	0.5365
$F_s^{-1}(0.85)$	1.6058	0.5010
Step 2 parameters**		
K	0.1197	0.0340***
r	0.0072	0.0023

Notes: Fifth-order polynomial approximation employed for terminal value (cf. end of Section 4).
*Standard deviation of parameter estimates obtained from 99 bootstrap resamples.
**Number of simulation draws used to evaluate expected value function: $L = 10$.
***These standard errors account for estimation error in the first-step estimates.

5. CONCLUSIONS

In this paper, we proposed a two-step pairwise-differencing procedure for structural estimation of a dynamic optimization model with unobserved state variables. Our estimator represents an innovative application of pairwise-difference methods, which have primarily been used in cross-sectional contexts (cf. Honoré and Powell, 1994) to structural dynamic optimization problems.

The most restrictive assumption made in this paper is that the unobserved state variables are independent across time. Accommodating serial correlation requires considering carefully

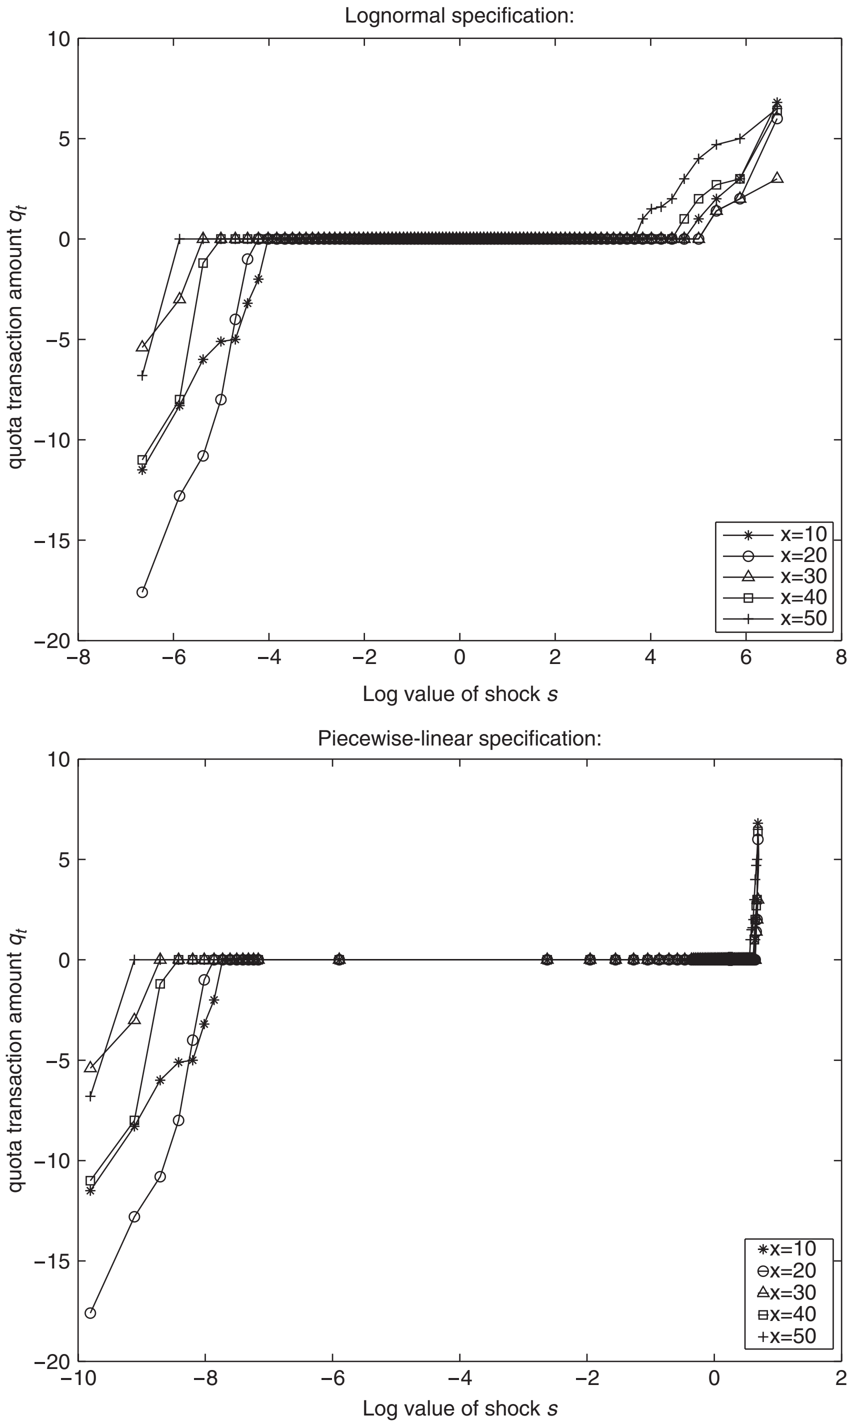

FIGURE 2

Estimated policy function for September 1997. Lognormal specification: piecewise-linear specification: x-axis: log value of shock s; y-axis: quota transaction amount q_t

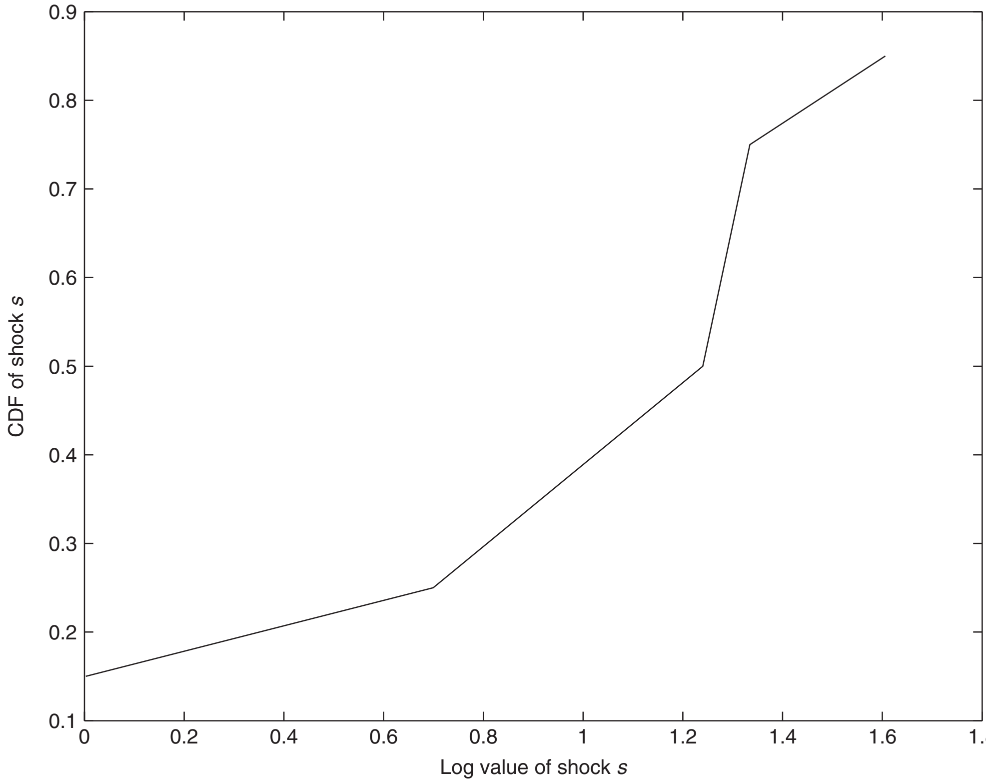

FIGURE 3
Estimated CDF of shock s. Estimated using equation (15). x-axis: log value of shock s; y-axis: CDF of shock s. Five quantiles were estimated: 0.15, 0.25, 0.5, 0.75, 0.85

the problem of initial conditions which, in turn, is very closely related to the issue of unobserved individual-specific heterogeneity (cf. Heckman, 1981). Unobserved heterogeneity, which we do not exploit, introduces substantial complexity in dynamic structural models and forms an active area of ongoing current research.

The estimation procedure only accommodates univariate unobserved state variables in agents' policy functions. This rules out multi-agent models in which the unobserved state variables of *all* the agents enter into each agent's policy function, as in the dynamic oligopoly model considered by Berry and Pakes (2000), where one firm's optimal investment is affected by the productivity state of every firm in the market, and all of these productivities are unobservable by the econometrician. It will be interesting to investigate in future work whether monotonicity and quantile invariance can be useful in these situations.

APPENDIX A. PROOFS

A.1. *Proof of Theorem 1*

For convenience, we sometimes use ψ in this proof to denote ψ^0, the true value. Also, let $m(z_{it}, z_{jt'}; \psi)$ and $r(z_{it}, z_{jt'}; \psi)$ denote, respectively, equations (16) and (19) evaluated at the actual (i.e. error-free) conditional distributions $F_{q|x}$.

Due to Assumption 5, the following approximation holds uniformly, up to $o_p\left(\frac{\sqrt{\log N}}{\sqrt{NTh}} + h^r\right)$:

$$\begin{aligned}
&\hat{F}(q|x) - F(s) \\
&\approx \frac{1}{f(x)}\left[\frac{1}{NTh}\sum_{l=1}^{N}\sum_{t=1}^{T}\mathbf{1}(q_{lt} < q)\,K\left(\frac{x_{lt}-x}{h}\right) - f(x)\,F(s)\right] \\
&\quad - \frac{F(s)}{f(x)}\left[\frac{1}{NTh}\sum_{l=1}^{N}\sum_{t=1}^{T}K\left(\frac{x_{lt}-x}{h}\right) - f(x)\right] \\
&= \frac{1}{f(x)}\left[\frac{1}{NTh}\sum_{l=1}^{N}\sum_{t=1}^{T}\mathbf{1}(q_{lt} < q)\,K\left(\frac{x_{lt}-x}{h}\right)\right] - \frac{F(s)}{f(x)}\left[\frac{1}{NTh}\sum_{l=1}^{N}\sum_{t=1}^{T}K\left(\frac{x_{lt}-x}{h}\right)\right].
\end{aligned} \tag{A1}$$

Together with other smoothness conditions in Assumption 5, uniform consistency of $\hat{F}(q|x)$ implies uniform convergence of the estimand (15) to the population limit $G_0(\psi)$, which in turn implies the consistency of $\hat{\psi}$ due to Assumption 6(ii).

To derive the asymptotic distribution, using a standard first-order Taylor expansion argument, we can approximate the estimator by

$$\sqrt{NT}\left(\hat{\psi} - \psi^0\right) = -A_{NT}^{-1}\left(1 + o_p(1)\right)\frac{1}{(NT)^{3/2}}\sum_{i=1}^{N}\sum_{t=1}^{T}\sum_{j=1}^{N}\sum_{t'=1}^{T}\hat{r}\left(z_{it}, z_{jt'}, \psi^0\right) \tag{A2}$$

where the Jacobian term is defined as (ψ^* is a set of intermediate values between ψ and $\hat{\psi}$):

$$A_{NT} \equiv \frac{1}{(NT)^2}\sum_{i,t,j,t'}\frac{1}{h}K\left(\frac{x_{it+1} - x_{jt'+1}}{h}\right)I_{it}I_{jt'}\frac{\partial}{\partial\psi}\left(\hat{m}\left(z_{it}, z_{jt'}, \psi^*\right)\frac{\partial}{\partial\psi}\left[\hat{m}\left(z_{it}, z_{jt'}, \psi^*\right)\right]\right). \tag{A3}$$

The Jacobian term A_{NT} can be approximated successively, each time up to $o_p(1)$, by replacing ψ^* with the true ψ^0, $\hat{m}(\cdot)$ with $m(\cdot)$, and the double summation with double expectations. As a consequence, $A_{NT} \xrightarrow{p} A$, where

$$A \equiv E_{z_{jt'}}E_{z_{it}}\left[I_{it}I_{jt'}\frac{\partial}{\partial\psi}\left[m\left(z_{it}, z_{jt'}, \psi\right)\right]\frac{\partial}{\partial\psi}\left[m\left(z_{it}, z_{jt'}, \psi\right)\right]\mathbf{1}_{x_{it+1} = x_{jt'+1}}\right] \tag{A4}$$

is the same matrix as stated in condition v.i of Assumption 6.

The form of A takes into account the fact that

$$E_{z_{jt'}}E_{z_{it}}\left[I_{it}I_{jt'}m\left(z_{it}, z_{jt'}, \psi\right)\frac{\partial^2}{\partial\psi\partial\psi'}m\left(z_{it}, z_{jt'}, \psi\right)\mathbf{1}_{x_{it+1} = x_{jt'+1}}\right] \equiv 0, \tag{A5}$$

which follows by assumption from the identity that $m\left(z_{it}, z_{jt'}, \psi\right) \equiv 0$ when $x_{it+1} = x_{jt'+1}$.

Next, we address the terms that appear behind the quadruple summation in (A2). Define

$$\begin{aligned}
\hat{w}\left(z_{it}, z_{jt'}\right) &\equiv \hat{m}\left(z_{it}, z_{jt'}, \psi^0\right)\frac{\partial}{\partial\psi}\left[\hat{m}\left(z_{it}, z_{jt'}, \psi^0\right)\right], \\
w\left(z_{it}, z_{jt'}\right) &\equiv m\left(z_{it}, z_{jt'}, \psi^0\right)\frac{\partial}{\partial\psi}\left[m\left(z_{it}, z_{jt'}, \psi^0\right)\right].
\end{aligned} \tag{A6}$$

Note that $\hat{w}\left(z_{it}, z_{jt'}\right)$ can be approximated up to $o_p(1)$ by the first-order linearization

$$\begin{aligned}
&w\left(z_{it}, z_{jt'}\right) + \frac{\partial w\left(z_{it}, z_{jt'}\right)}{\partial F_s\left(s_{it}\right)}\left(\hat{F}\left(q_{it}|x_{it}\right) - F_s\left(s_{it}\right)\right) + \frac{\partial w\left(z_{it}, z_{jt'}\right)}{\partial F_s\left(s_{jt'}\right)}\left(\hat{F}\left(q_{jt'}|x_{jt'}\right) - F_s\left(s_{jt'}\right)\right) \\
&\equiv \left(1 + o_p(1)\right)\left[w\left(z_{it}, z_{jt'}\right) + \frac{1}{NT}\sum_{l=1}^{N}\sum_{t''=1}^{T}v\left(z_{it}, z_{jt'}, z_{lt''}\right)\right],
\end{aligned} \tag{A7}$$

where, substituting in equation (A1) above:

$$\begin{aligned}
v\left(z_{it}, z_{jt'}, z_{lt''}\right) =& \frac{\partial w\left(z_{it}, z_{jt'}\right)}{\partial F_s\left(s_{it}\right)}\frac{1}{h}K\left(\frac{x_{lt''} - x_{it}}{h}\right)\frac{1}{f\left(x_{it}\right)}\left[\mathbf{1}\left(q_{lt''} < q_{it}\right) - F_s\left(s_{it}\right)\right] \\
&+ \frac{\partial w\left(z_{it}, z_{jt'}\right)}{\partial F_s\left(s_{jt'}\right)}\frac{1}{h}K\left(\frac{x_{lt''} - x_{jt'}}{h}\right)\frac{1}{f\left(x_{jt'}\right)}\left[\mathbf{1}_{q_{lt''} < q_{jt'}} - F_s\left(s_{jt'}\right)\right].
\end{aligned} \tag{A8}$$

Hence, we can approximate the linear term in equation (A2) by a U-statistic representation:

$$\frac{1}{(NT)^{3/2}}\sum_{i=1}^{N}\sum_{t=1}^{T}\sum_{j=1}^{N}\sum_{t'=1}^{T}\frac{1}{h}K\left(\frac{x_{it+1}-x_{jt'+1}}{h}\right)w\left(z_{it},z_{jt'}\right)I_{it}I_{jt'} + \frac{1}{(NT)^{5/2}}\sum_{i=1}^{N}\sum_{t=1}^{T}\sum_{j=1}^{N}\sum_{t'=1}^{T}\sum_{l=1}^{N}\sum_{t''=1}^{T}\frac{1}{h}K\left(\frac{x_{it+1}-x_{jt'+1}}{h}\right)v\left(z_{it},z_{jt'},z_{lt''}\right)I_{it}I_{jt'}. \quad \text{(A9)}$$

Given Assumption 4 in the main text on the kernel and bandwidth sequence, the bias terms in the non-parametric kernel estimation are asymptotically negligible and the conditions for Lemma 3.1 in Powell, Stock and Stoker (1989) hold. Hence, we can invoke the projection representation of (A9). For the first term in equation (A9), we have

$$\begin{aligned}
&\frac{1}{(NT)^{3/2}}\sum_{i=1}^{N}\sum_{t=1}^{T}\sum_{j=1}^{N}\sum_{t'=1}^{T}\frac{1}{h}K\left(\frac{x_{it+1}-x_{jt'+1}}{h}\right)I_{it}I_{jt'}w\left(z_{it},z_{jt'}\right)\\
&=\frac{1}{\sqrt{NT}}\sum_{i=1}^{N}\sum_{t=1}^{T}E\left(\frac{1}{h}K\left(\frac{x_{it+1}-x_{jt'+1}}{h}\right)I_{it}I_{jt'}w\left(z_{it},z_{jt'}\right)\bigg|z_i\right)\\
&\quad+\frac{1}{\sqrt{NT}}\sum_{j=1}^{N}\sum_{t'=1}^{T}E\left(\frac{1}{h}K\left(\frac{x_{it+1}-x_{jt'+1}}{h}\right)I_{it}I_{jt'}w\left(z_{it},z_{jt'}\right)\bigg|z_j\right)+o_p(1) \qquad \text{(A10)}\\
&=\frac{1}{\sqrt{NT}}\sum_{i=1}^{N}\sum_{t=1}^{T}E_{z_j}\left[I_{it}I_{jt'}w\left(z_{it},z_{jt'}\right)\mathbf{1}_{x_{jt'+1}=x_{it+1}}\right]f\left(x_{it+1}\right)\\
&\quad+\frac{1}{\sqrt{NT}}\sum_{j=1}^{N}\sum_{t'=1}^{T}E_{z_i}\left[I_{it}I_{jt'}w\left(z_{it},z_{jt'}\right)\mathbf{1}_{x_{it+1}=x_{jt'+1}}\right]f\left(x_{jt'+1}\right)+o_p(1)=o_p(1).
\end{aligned}$$

Both terms in the above display vanish asymptotically for the same reasoning that leads to (A5). This makes explicit the feature that the pairwise-differencing step introduces no additional variation to the parameter estimate $\hat{\psi}$. The non-parametric estimates of $F_{q|x}$ produce all the first-order variation, which is reflected in the non-negligible limit for the second term of equation (A9):

$$\begin{aligned}
&\frac{1}{(NT)^{5/2}}\sum_{i=1}^{N}\sum_{t=1}^{T}\sum_{j=1}^{N}\sum_{t'=1}^{T}\sum_{l=1}^{N}\sum_{t''=1}^{T}\frac{1}{h}K\left(\frac{x_{it+1}-x_{jt'+1}}{h}\right)I_{it}I_{jt'}v\left(z_{it},z_{jt'},z_{lt''}\right)\\
&=\frac{1}{\sqrt{NT}}\sum_{l=1}^{N}\sum_{t''=1}^{T}E\left(\frac{1}{h}K\left(\frac{x_{it+1}-x_{jt+1}}{h}\right)I_{it}I_{jt'}v\left(z_{it},z_{jt'},z_{lt''}\right)\bigg|z_l\right)+o_p(1) \qquad \text{(A11)}\\
&\equiv\frac{1}{\sqrt{NT}}\sum_{l=1}^{N}\sum_{t''=1}^{T}\tilde{v}_{t''}\left(z_l\right)+o_p(1).
\end{aligned}$$

The first equality follows from Assumption 4, which implies that the other two projection terms vanish. After tedious but straightforward calculations, we can write $\tilde{v}_{t''}(z_l)$ as

$$\begin{aligned}
&E_{z_i}\left[E_{z_j}\left(I_{it}I_{jt'}\frac{\partial w\left(z_{it},z_{jt'}\right)}{\partial F_s\left(s_{it}\right)}\mathbf{1}_{x_{jt'+1}=x_{it+1}}\right)f\left(x_{it+1}\right)\left(\mathbf{1}\left(q_{lt''}<q_{it}\right)-F_s\left(s_{it}\right)\right)\mathbf{1}_{x_{it}=x_{lt''}}\right]\\
&+E_{z_j}\left[E_{z_i}\left(I_{it}I_{jt'}\frac{\partial w\left(z_{it},z_{jt'}\right)}{\partial F_s\left(s_{jt}\right)}\mathbf{1}_{x_{it+1}=x_{jt'+1}}\right)f\left(x_{jt'+1}\right)\left(\mathbf{1}\left(q_{lt''}<q_{jt'}\right)-F_s\left(s_{jt'}\right)\right)\mathbf{1}_{x_{jt'}=x_{lt''}}\right].
\end{aligned} \quad \text{(A12)}$$

Therefore, we conclude that

$$\sqrt{NT}\left(\hat{\psi}-\psi\right)\xrightarrow{d}N\left(0,A^{-1}\Omega A^{-1}\right) \quad \text{(A13)}$$

where

$$\Omega=\frac{1}{T}E\left(\sum_{t=1}^{T}\tilde{v}_t\left(z_l\right)\right)\left(\sum_{t=1}^{T}\tilde{v}_t\left(z_l\right)\right)' \quad \text{(A14)}$$

and A is defined in (A4) above. The asymptotic variance can be consistently estimated using resampling methods or empirical analogues. ||

A.2. *Proof of Theorem 2*

Under the conditions of Assumption 7, it follows from a standard first-order Taylor expansion that $\sqrt{NT}\left(\hat{\theta}_2 - \theta_2^0\right)$ is asymptotically equivalent to

$$\overline{A}^{-1}\left(\frac{1}{\sqrt{NT}}\sum_i\sum_t \overline{h}\left(x_{it}, q_{it}; \hat{F}_{q|x}(\cdot|:), \hat{\psi}, \theta_2^0\right)\right), \tag{A15}$$

which is, in turn, asymptotically equivalent to

$$\overline{A}^{-1}\left(\frac{1}{\sqrt{NT}}\sum_{i,t}\left(\overline{h}\left(x_{it}, q_{it}; \hat{F}_{q|x}(\cdot|:), \psi_0, \theta_2^0\right) - \overline{h}\left(x_{it}, q_{it}; F_{q|x}(\cdot|:), \psi_0, \theta_2^0\right)\right) + B\sqrt{NT}\left(\hat{\psi} - \psi^0\right)\right). \tag{A16}$$

The above is asymptotically equivalent to

$$\overline{A}^{-1}\left(\sqrt{NT}E_{x,q}\left(\overline{h}\left(x, q; \hat{F}_{q|x}(\cdot|:), \psi_0, \theta_2^0\right) - \overline{h}\left(x, q; F_{q|x}(\cdot|:), \psi_0, \theta_2^0\right)\right) + B\sqrt{NT}\left(\hat{\psi} - \psi^0\right)\right). \tag{A17}$$

It then follows from the powerful results of Newey (1994) and the other regularity conditions in Assumption 7 that,

$$\begin{aligned}&\sqrt{NT}E_{x,q}\left(\overline{h}\left(x, q; \hat{F}_{q|x}(\cdot|:), \psi_0, \theta_2^0\right) - \overline{h}\left(x, q; F_{q|x}(\cdot|:), \psi_0, \theta_2^0\right)\right)\\ &= \frac{1}{\sqrt{NT}}\sum_i\sum_t \alpha(z_{it}) + o_p(1).\end{aligned} \tag{A18}$$

Hence, combining the above two equations (A17) and (A18), the asymptotic variance of $\sqrt{NT}(\hat{\theta}_2 - \theta_2^0)$ is equal to the expression given in the statement of Theorem 2. Interestingly, Newey (1994) shows that as long as sufficient regularity conditions are met, this asymptotic linear representation holds regardless of the non-parametric estimation method used in the first stage to construct $\hat{F}_{q|x}(\cdot|\cdot)$.

It remains to derive the expression for the influence function $\alpha(z)$. Recall that the second-stage estimation is based on equating to zero the first-order condition (with respect to current choice q_0) of a simulated version of the following objective function:

$$E_{x_0, q_0, \{\tau\}}\left[u\left(x_0, q_0, s\left(\hat{F}_{q|x}(q_0|x_0)\right), \psi, \theta_2\right) + \sum_t \beta^t u\left(x_t, q_t, s(\tau_t), \psi, \theta_2\right)\right] \tag{A19}$$

subject to $x_1 = x_0 + q_0$ and for all periods $t \geq 1$,

$$x_{t+1} = x_t + q_t; \quad q_t = \hat{F}_{q|x}^{-1}(\tau_t|x_t). \tag{A20}$$

Because of this recursive structure, the influence function $\alpha(z_i)$ also has a recursive structure, $\alpha(z_i) = \sum_t \alpha_t(z_i)$, and it is not possible to provide a complete analytical description. In the following, we will provide the derivation of the first two terms $\alpha_1(z_i)$ and $\alpha_2(z_i)$, corresponding to a two-period version of the model. Higher order terms can be derived using similar calculations, but are too tedious to describe analytically. We believe this derivation adequately illustrates the form that the influence function will take in the general model. In any case, since in practice we use bootstrap inference for the second-stage parameters, we only require the existence of the asymptotic influence function and do not require explicit knowledge of its functional form.

We introduce the shorthand notation $u(x, q, \tau) = u(x, q, s(\tau), \psi, \theta_2)$, and let u_1 and u_2 denote, respectively, the derivatives of $u(x, q, \tau)$ with respect to the first and second arguments. In the two-period model, the limiting first-order condition (corresponding to the $\overline{h}(\cdots)$ function above) satisfied by the parameters θ_2 is

$$0 = E\left[A(x_0, q_0)\left\{u_2\left(x_0, q_0, s\left(F_{q|x}(q_0|x_0)\right)\right) + \beta\frac{\partial}{\partial q_0}u(x_0 + q_0, F_{q|x}^{-1}(\tau_1|x_0 + q_0), s(\tau_1))\right\}\right] \tag{A21}$$

where [cf. equation (38)]

$$A(x_0, q_0) = \frac{\partial}{\partial \theta_2}\left[u_2\left(x_0, q_0, s\left(F_{q|x}(q_0|x_0)\right)\right) + \beta\frac{\partial}{\partial q_0}u(x_0 + q_0, F_{q|x}^{-1}(\tau_1|x_0 + q_0), s(\tau_1))\right]. \tag{A22}$$

Let η index a one-parameter family (or "sub-path") of parametric specifications for the conditional CDF $F(q|x)$. Members of this family are denoted $F_\eta(q|x)$, with $F_{\eta=0}(q|x) = F(q|x)$, the true value. (In the following, we omit the subscript $q|x$ from this distribution for convenience.) Following Newey (1994), the influence function for the two-period model satisfies the "pathwise derivative" (i.e. derivative with respect to η), evaluated at $\eta = 0$, of the first-order condition (A21). This influence function will be the sum of two terms, corresponding to the two terms within the curly brackets of equation (A21).[11]

The first term $\alpha_1(z)$ in the influence function satisfies the pathwise derivative of the first term of the first-order condition; that is:

$$\frac{\partial}{\partial\eta} E\left[A\left(x_0, q_0\right) u_2\left(x_0, q_0, s\left(F\left(q_0|x_0\right)\right)\right)\right] = E\alpha_1(z)S_\eta(z). \tag{A23}$$

To derive $\alpha_1(z)$, start with

$$\frac{\partial}{\partial\eta} E\left[A\left(x_0, q_0\right) u_2\left(x_0, q_0, s\left(F\left(q_0|x_0\right)\right)\right)\right] = E\left[\overline{u}_F\left(x_0, q_0\right) \cdot \frac{\partial}{\partial\eta} F_\eta\left(q_0|x_0\right)\right], \tag{A24}$$

where $\overline{u}_F\left(x_0, q_0\right) \equiv A\left(x_0, q_0\right) \frac{\partial}{\partial\tau} u_2\left(x_0, q_0, s\left(\tau\right)\right)|_{\tau=F(q_0|x_0)}$. The right-hand side can be in turn written as

$$\begin{aligned} E\left[\overline{u}_F\left(x, q\right) \cdot \frac{\partial}{\partial\eta} F_\eta\left(q|x\right)\right] &= \int\int\int \overline{u}_F\left(x, q\right) \mathbf{1}_{u\le q} \frac{\partial}{\partial\eta} f_\eta\left(u|x\right) du f\left(q, x\right) dq dx \\ &= \int\int \alpha_F\left(x, u\right) \frac{\partial}{\partial\eta} f_\eta\left(u|x\right) f\left(x\right) du dx \\ &= \int\int \alpha_F\left(x, u\right) \left[\frac{\partial}{\partial\eta} f_\eta\left(u, x\right) - f(u|x) \frac{\partial}{\partial\eta} f_\eta\left(x\right)\right] du dx \\ &= \int\int \alpha_F\left(x, u\right) \frac{\partial}{\partial\eta} f_\eta\left(u, x\right) du dx - \int E\left[\alpha_F\left(x, u\right)|x\right] \frac{\partial}{\partial\eta} f_\eta\left(x\right) dx \end{aligned} \tag{A25}$$

where $\alpha_F\left(x, u\right) = \int \overline{u}_F\left(x, q\right) \mathbf{1}_{u\le q} f\left(q|x\right) dq$. In the above display, the second line exchanges the order of differentiation, and the third line applies the chain rule. Hence, the first term in the influence function equals

$$\alpha_1\left(z\right) = \alpha_F\left(x, q\right) - E\left[\alpha_F\left(x, q\right)|x\right]. \tag{A26}$$

Similarly, the second term $\alpha_2(z)$ satisfies

$$\frac{\partial}{\partial\eta} E\left[A\left(x, q\right) \frac{\partial}{\partial q} u(x + q, F_\eta^{-1}(\tau_1|x + q), s\left(\tau_1\right))\right] = E\alpha_2(z)S_\eta(z). \tag{A27}$$

Note that

$$\begin{aligned} &\frac{\partial}{\partial q} u(x + q, F_\eta^{-1}(\tau_1|x + q), s(\tau_1)) \\ =& u_1(x + q, F_\eta^{-1}(\tau_1|x + q), s(\tau_1)) + u_2(x + q, F_\eta^{-1}(\tau_1|x + q), s(\tau_1)) \cdot \frac{\partial}{\partial q} F_\eta^{-1}(\tau_1|x + q). \end{aligned} \tag{A28}$$

Corresponding to the two terms on the right-hand side, we can further decompose $\alpha_2\left(z\right) = \alpha_{21}\left(z\right) + \alpha_{22}\left(z\right)$. For the first subcomponent (defining $q_1 = F_0^{-1}(\tau_1|x + q)$),

$$\begin{aligned} &\frac{\partial}{\partial\eta} E_{x,q,\tau_1}\left[A\left(x, q\right) u_1\left(x + q, F_\eta^{-1}\left(\tau_1|x + q\right), s(\tau_1)\right)\right] \\ =& - E_{x,q,\tau_1}\left[\frac{A\left(x, q\right) u_{12}\left(x + q, F_\eta^{-1}\left(\tau_1|x + q\right), s(\tau_1)\right)}{f\left(F^{-1}\left(\tau_1|x + q\right)|x + q\right)} \frac{\partial}{\partial\eta} F_\eta\left(q_1|x + q\right)\right]. \end{aligned} \tag{A29}$$

11. For completeness, $A(x_0, q_0)$ also depends on $F(q|x)$, and hence the pathwise derivative should also include terms involving the pathwise derivative of $A(x_0, q_0)$. However, because the term within curly brackets in equation (A21) is identically equal to zero for all x_0 and q_0, these additional terms will all evaluate to zero, so that we can treat $A(x_0, q_0)$ as known (and hence not affecting the asymptotic variance of θ_2) without loss of generality.

For $y = x + q$ and $h(y,u) = \int_0^1 E(A(x,q)|y) \frac{u_{12}\left(y, F_0^{-1}(\tau_1|y), \tau_1\right)}{f\left(F_0^{-1}(\tau_1|y)|y\right)} \mathbf{1}_{u \le F_0^{-1}(\tau_1|y)} d\tau_1$, the above can be written as

$$\begin{aligned} &\int\int h(y,u) \frac{\partial}{\partial \eta} f_\eta(u|y)\, du f(y)\, dy \\ =&\int\int h(y,u) \frac{\partial}{\partial \eta} f_\eta(u,y)\, du dy - \int\int E\left[h(y,u)|y\right] \frac{\partial}{\partial \eta} f_\eta(y)\, dy. \end{aligned} \tag{A30}$$

Therefore $\alpha_{21}(z) = h(q, x+q) - E\left[h(q, x+q)|x+q\right]$.

Finally, consider the second subcomponent $\alpha_{22}(z)$ for

$$\frac{\partial}{\partial \eta} E_{x,q,\tau_1}\left\{A(x,q)\left[u_2\left(x+q, F_\eta^{-1}(\tau_1|x+q), s(\tau_1)\right) \frac{\partial}{\partial q} F_\eta^{-1}(\tau_1|x+q)\right]\right\}. \tag{A31}$$

From the chain rule, the above pathwise derivative has two components, which correspond to two subcomponents of $\alpha_{22}(z)$, denoted $\alpha_{221}(z)$ and $\alpha_{222}(z)$. The first subcomponent $\alpha_{221}(z)$ comes from

$$\frac{\partial}{\partial \eta} E_{x,q,\tau_1}\left\{A(x,q)\left[u_2\left(x+q, F_\eta^{-1}(\tau_1|x+q), s(\tau_1)\right) \frac{\partial}{\partial q} F^{-1}(\tau_1|x+q)\right]\right\}. \tag{A32}$$

Its derivation is similar to that of $\alpha_{21}(z)$, except that the definition of $h(y,u)$ in $\alpha_{21}(z)$ is now replaced by

$$h(y,u) = \int E\left[A(x,q)|y\right] u_{22}\left(y, F_0^{-1}(\tau_1|y), \tau_1\right) \frac{\partial}{\partial q} F_0^{-1}(\tau_1|x+q) \frac{\mathbf{1}_{u \le F_0^{-1}(\tau_1|y)}}{-f\left(F_0^{-1}(\tau_1|y)|y\right)} d\tau_1. \tag{A33}$$

The last subcomponent is more complex and requires the use of integration by parts. Consider, for $y = x + q$,

$$\begin{aligned} &\frac{\partial}{\partial \eta} E_{y,\tau_1}\left[E\left[A(x,q)|y\right] u_2\left(y, F_0^{-1}(\tau_1|y), \tau_1\right) \frac{\partial}{\partial y} F_\eta^{-1}(\tau_1|y)\right] \\ =&E_{y,\tau_1} E\left[A(x,q)|y\right]\left[\frac{u_2\left(y, F_0^{-1}(\tau_1|y), \tau_1\right)}{-f\left(F_0^{-1}(\tau_1|y)\right)} \frac{\partial^2}{\partial \eta \partial y'} F_\eta\left(F_0^{-1}(\tau_1|y)|y'\right)\Big|_{y'=y}\right]. \end{aligned} \tag{A34}$$

If we define

$$\overline{h}(y,u) = \int_0^1 E\left[A(x,q)|y\right] \frac{u_2\left(y, F_0^{-1}(\tau_1|y), \tau_1\right)}{f\left(F_0^{-1}(\tau_1|y)\right)} \mathbf{1}_{u < F_0^{-1}(\tau_1|y)} d\tau_1. \tag{A35}$$

Then the pathwise derivative can be written as

$$\int\int \overline{h}(y,u) \frac{\partial^2}{\partial \eta \partial y} f_\eta(u|y)\, du f(y)\, dy. \tag{A36}$$

Through a sequence of integration by parts, this can be further written as

$$-\frac{\partial}{\partial \eta}\int\int \frac{\partial \overline{h}(y,u) f(y)}{\partial y} \frac{1}{f(y)} f_\eta(u|y) f(y)\, du dy. \tag{A37}$$

Careful inspection shows that the corresponding asymptotic influence function is given by

$$\alpha_{222}(z) = \tilde{h}(x+q, q) - E\left[\tilde{h}(x+q, q)|x+q\right], \tag{A38}$$

where $\tilde{h}(y,u) = -\frac{\partial \overline{h}(y,u) f(y)}{\partial y} \frac{1}{f(y)}$. This completes the derivation of the influence function for the two-period model (A21). ||

APPENDIX B. DERIVATION OF OPTIMAL POLICY FUNCTION IN EXAMPLE

Here we derive the optimal policy function (28) for the example problem we introduced in Section 3.1.2. The dynamic problem (26) falls into the framework of the "general" LQ problem in (Sargent, 1987, p. 51):

$$\max_{\{q_t\}} \sum_{t=0}^{\infty} \beta^t \left\{ (z_t', q_t) \begin{pmatrix} \overline{R} & \overline{W} \\ \overline{W}' & \overline{Q} \end{pmatrix} \begin{pmatrix} z_t \\ q_t \end{pmatrix} \right\} \tag{B1}$$

subject to

$$z_{t+1} = \overline{A} z_t + \overline{B} q_t + s_{t+1}$$

where

$$z_t = \begin{pmatrix} x_t \\ s_t \\ 1 \end{pmatrix} \tag{B2}$$

$$\overline{R} = \begin{pmatrix} 4 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix} \quad \overline{W} = \begin{pmatrix} 0 \\ -\frac{1}{2} \\ 0 \end{pmatrix} \quad \overline{Q} = -\frac{1}{2}$$

$$\overline{A} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad \overline{B} = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}. \tag{B3}$$

Following Sargent (1987), we re-parameterize the model. Define

$$\begin{aligned} R &= \overline{R} - \overline{W}\overline{Q}^{-1}\overline{W}' \\ Q &= \overline{Q} \\ A &= \overline{A} - \overline{B}\overline{Q}^{-1}\overline{W}' \\ B &= \overline{B}. \end{aligned} \tag{B4}$$

Then re-parameterized problem is

$$\max_{\{v_t\}} E \sum_{t=0}^{\infty} \beta^t \left\{ z_t' R z_t + v_t' Q v_t \right\} \tag{B5}$$

subject to

$$z_{t+1} + A z_t + B v_t + s_{t+1} \tag{B6}$$

where

$$v_t = \overline{Q}^{-1}\overline{W}' z_t + u_t. \tag{B7}$$

The optimal policy function for the re-parameterized problem is $v_t = -F z_t$ where

$$F_t = \beta \left(Q + \beta B' P B \right)^{-1} B' P A \tag{B8}$$

and P satisfies the Ricatti equation

$$P = R + \beta A' P A - \beta^2 A' P B \left(Q + \beta B' P B \right)^{-1} B' P A. \tag{B9}$$

Then, correspondingly, the optimal policy function for the original problem is

$$\begin{aligned} \overline{Q}^{-1}\overline{W}' z_t + u_t &= -F z_t \\ \Leftrightarrow u_t &= -(F + \overline{Q}^{-1}\overline{W}') z_t \end{aligned} \tag{B10}$$

which evaluates to the optimal policy function given in equation (28).

APPENDIX C. REMARKS ON EMPIRICAL ILLUSTRATION

In our empirical application, we make the assumption that the price p_t is taken as given and known by bidders when they are deciding how much quota q_t to buy. Here, we show that this assumption is consistent with a perfect foresight equilibrium in a dynamic competitive market composed of individually atomistic traders, similar to Jovanovic (1982) and Hopenhayn (1992). Prices for each period are determined by a market-clearing condition: given policies $q(x_{it}, s_{it}, p_t)$, $\forall i$,

$$p_t : \int \int q(x, s, p_t)\, \mathcal{J}_t(dx)\mathcal{H}_t(ds) = 0\,, \quad \forall t \tag{C1}$$

where $\mathcal{J}_t(\cdot)$ and $\mathcal{H}_t(\cdot)$ denote, respectively, the distribution of quota stocks and shocks in the cross-section of traders during period t. Given our *i.i.d.* assumption on the shock distribution, it is immediately clear that

$$\mathcal{H}_t(s) = F_s(s)\,, \ \forall t. \tag{C2}$$

Similarly, the cross-sectional distribution of stocks $\mathcal{J}_t(x)$ evolves according to

$$\mathcal{J}_t(x) = \int \int \mathbf{1}(z + q(z, s, p_{t-1}) \leq x)\, \mathcal{H}_{t-1}(ds)\mathcal{J}_{t-1}(dz). \tag{C3}$$

Given any initial stock distribution $\mathcal{J}_0$, the sequences $\{\mathcal{J}_t\}$ and $\{\mathcal{H}_t\}$ are both deterministic, and evolve according to (C2) and (C3). Therefore, by the market-clearing conditions (C1), the sequence $\{p_t\}_t$ is also deterministic. Hence, in competitive equilibrium in this market, all traders will have perfect foresight about the evolution of prices.

Acknowledgements. We thank Phil Cairns at the Dairy Farmers of Ontario for providing the data and patiently answering our questions. We thank J. Adda, V. Aguirregabiria, P. Bajari, L. Benkard, R. Blundell, B. Honoré, H. Ichimura, S. Khan, C. Meghir, E. Miravete, S. Ng, A. Pakes, M. Pesendorfer, G. Ridder, L. Roberts, F. Wolak, the editor and two referees, and seminar participants at various universities and conferences for helpful discussions. We are grateful to the NSF (SES-0079495, SES-0721015), the Sloan Foundation and SSHRC for financial support.

REFERENCES

AGUIRREGABIRIA, V. (2005), "Nonparametric Identification of Behavioral Responses to Counterfactual Policy Interventions in Dynamic Discrete Decision Processes", *Economics Letters,* **87**, 393–398.

AGUIRREGABIRIA, V. and MIRA, P. (2007), "Sequential Estimation of Dynamic Discrete Games", *Econometrica,* **75**, 1–53.

AHN, H. and POWELL, J. (1993), "Semiparametric Estimation of Censored Selection Models with a Nonparametric Selection Mechanism", *Journal of Econometrics,* **58**, 3–29.

BAJARI, P. and BENKARD, L. (2005), "Demand Estimation with Heterogeneous Consumers and Unobserved Product Characteristics: A Hedonic Approach", *Journal of Political Economy,* **113**, 1239–1276.

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

BERRY, S., OSTROVSKY, M. and PAKES, A. (2007), "Simple Estimators for the Parameters of Discrete Dynamic Games (With Entry Exit Examples)", *RAND Journal of Economics,* **38**, 373–399.

BERRY, S. and PAKES, A. (2000), "Estimation from the Optimality Conditions for Dynamic Controls" (Manuscript, Yale University).

BIGGS, E. (1990), "The Challenge of Achievement: The Ontario Milk Marketing Board's First 25 Years of Operation", The Ontario Milk Marketing Board.

EBERLY, J. (1994), "Adjustment of Consumers' Durables Stocks: Evidence from Automobile Purchases", *Journal of Political Economy,* **102**, 403–436.

GOURIEROUX, C. and MONFORT, A. (1996), *Simulation-Based Econometric Methods* (Oxford: Oxford University Press).

HECKMAN, J. (1981), "The Incidental Parameters Problem and the Problem of Initial Conditions in Estimating a Discrete Time-Discrete Data Stochastic Process", in C. Manski and D. McFadden (eds.) *Structural Analysis of Discrete Data* (Cambridge, MA: MIT Press).

HONORÉ, B. and POWELL, J. (1994), "Pairwise Difference Estimators of Censored and Truncated Regression Models", *Journal of Econometrics,* **64**, 241–278.

HOPENHAYN, H. (1992), "Entry, Exit, and Firm Dynamics in Long Run Equilibrium", *Econometrica,* **60**, 1127–1150.

HOTZ, J. and MILLER, R. (1993), "Conditional Choice Probabilities and the Estimation of Dynamic Models", *Review of Economic Studies,* **60**, 497–529.

HOTZ, J., MILLER, R., SANDERS, S. and SMITH, J. (1994), "A Simulation Estimator for Dynamic Models of Discrete Choice", *Review of Economic Studies,* **61**, 265–289.

JOVANOVIC, B. (1982), "Selection and the Evolution of Industry", *Econometrica,* **50**, 649–670.

JUDD, K. (1998), *Numerical Methods in Economics* (Cambridge, MA: MIT Press).

MAGNAC, T. and THESMAR, D. (2002), "Identifying Dynamic Discrete Decision Processes", *Econometrica,* **70**, 801–816.

MATZKIN, R. (2003), "Nonparametric Estimation of Nonadditive Random Functions", *Econometrica,* **71**, 1339–1376.

NEWEY, W. (1994), "The Asymptotic Variance of Semiparametric Estimators", *Econometrica,* **62**, 1349–1382.

OLLEY, S. and PAKES, A. (1996), "The Dynamics of Productivity in the Telecommunications Equipment Industry", *Econometrica,* **64**, 1263–1297.

PAKES, A. (1994), "Dynamic Structural Models, Problems and Prospects: Mixed Continuous Discrete Controls and Market Interactions", in C. Sims and J.-J. Laffont (eds.) *Advances in Econometrics: Proceedings of the 1990 Meetings of the Econometric Society*, Vol. 2 (Cambridge, UK: Cambridge University Press) 171–259.

PAKES, A. and SIMPSON, M. (1989), "Patent Renewal Data", *Brookings Papers on Economic Activity: Microeconomics,* **1**, 331–403.

PESENDORFER, M. and SCHMIDT-DENGLER, P. (2008), "Asymptotic Least Squares Estimators for Dynamic Games", *Review of Economic Studies,* **75**, 901–928.

POWELL, J., STOCK, J. and STOKER, T. (1989), "Semiparametric Estimation of Index Coefficients", *Econometrica,* **57**, 1403–1430.

RUST, J. (1987), "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher", *Econometrica,* **55**, 999–1033.

RUST, J. (1996), "Numerical Dynamic Programming in Economics", in D. K. H. Amman and J. Rust (eds.) *Handbook of Computational Economics* (New York: North Holland).

RUST, J., TRAUB, J. and WOZNIAKOWSKI, H. (2002), "Is there a Curse of Dimensionality for Contraction Fixed Points in the Worst Case?" *Econometrica,* **70**, 285–329.

SARGENT, T. (1987), *Dynamic Macroeconomic Theory* (Cambridge, MA: Harvard University Press).

TABER, C. (2000), "Semiparametric Identification and Heterogeneity in Discrete-Choice Dynamic Programming Models", *Journal of Econometrics,* **96**, 201–229.