

Pairwise-Difference Estimation of a Dynamic Optimization Model*

Han Hong
Dept. of Economics
Stanford University

Matthew Shum[†]
Dept. of Economics
Johns Hopkins University

This draft: July 6, 2007

Abstract

We develop a new estimation methodology for a dynamic optimization model with unobserved shocks. We propose a pairwise-difference approach which exploits two common features of the dynamic optimization problem we consider: (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.

Keywords: Pairwise-difference Estimators, Dynamic Optimization Models

JEL: C13, C50

*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. Honore, 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-0452143), the Sloan Foundation and SSHRC for financial support.

[†]Corresponding author: Matthew Shum, Dept. of Economics, Johns Hopkins University, 3400 N. Charles Street, Baltimore, MD 21218; mshum@jhu.edu (e-mail).

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 to the econometrician (but 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.

The two-step pairwise-difference estimator we propose represents a new approach to estimating 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 contexts (cf. [Honore and Powell \(1994\)](#)), to structural dynamic optimization problems. It 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 \(2005\)](#)).

Our approach is related to some recent work which exploits monotonicity assumptions to identify and estimate structural equations. Earlier, [Olley and Pakes \(1996\)](#) exploited such an assumption in order to invert out the unobserved shock to derive a semiparametric estimator for production functions with serially correlated unobservables. [Matzkin \(2003\)](#) exploited the quantile invariance implication of monotonicity to estimate nonparametrically functions which are nonlinear 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.

One advantage of our approach over alternative methods for estimating continuous choice dynamic optimization models, such as Euler Equation-based methods, is in accommodating shocks which are observed by agents at the time they make their decisions, but unobserved to the econometrician.¹ Our approach also accommodates dynamic optimization models

¹Conventional 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 analogs of these orthogonality conditions, the econometrician needs to know the value of all the state

in which agents' choices are both continuous and discrete, for which conventional Euler-Equation methods are either not applicable or difficult.

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 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 is also related to a recent literature on the identification and estimation of dynamic game models (e.g., [Magnac and Thesmar \(2002\)](#), [Pesendorfer and Schmidt-Dengler \(2003\)](#), [Aguirregabiria and Mira \(2007\)](#), [Berry, Ostrovsky, and Pakes \(2004\)](#), [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 behavior 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 :

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

subject to the Markov transition probabilities for the state variables

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

variables (including the shocks) at times t and $t + 1$. See [Pakes \(1994\)](#) (pp. 188–189) for a more thorough discussion.

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. q_{it} denotes the agent's choice variable. 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. 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.)

$U(\cdot \cdot \cdot ; \theta)$ 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.² 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} . We also assume:

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

$$F(x_{i,t+1}, s_{i,t+1} | x_{it}, s_{it}, q_{it}) = F(x_{i,t+1} | x_{it}, s_{it}, q_{it}) \cdot F_s(s_{i,t+1}; \gamma)$$

Specifically, the shocks s_{it} are assumed to be drawn i.i.d. across (i, t) from the marginal distribution $F_s(\cdot; \gamma)$, parameterized 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:*

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

It turns out that this assumption is more restrictive than required for our estimation procedure, but is convenient to make for explication purposes (and also natural for our empirical

²This usage differs from the macroeconomic literature, where a “shock” is often unobserved by both the econometrician as well as the optimizing agent when she makes her decision.

illustration below, and also in the investment example mentioned earlier). Later, we discuss how this assumption can be relaxed.

Given these assumptions, and assuming stationarity, the agent's optimal policy function can be expressed as the maximizer of Bellman's equation: for $t = 1, 2, 3, \dots$,

$$\begin{aligned} q(x_{it}, s_{it}; \theta, \gamma) &= \operatorname{argmax}_q \{U(x_{it}, s_{it}, q; \theta) + \beta \mathcal{E}_{x_{it+1}, s_{it+1} | x_{it}, s_{it}, q} V(x_{it+1}, s_{it+1}; \theta, \gamma)\} \\ &= \operatorname{argmax}_q \{U(x_{it}, s_{it}, q; \theta) + \beta \mathcal{E}_{s_{it+1} | x_{it}, s_{it}, q} V(x_{it} + q, s_{it+1}; \theta, \gamma)\} \end{aligned} \quad (3)$$

where:

$$V(x_{i,t+1}, s_{i,t+1}; \theta, \gamma) \equiv \max_{\{q_{i\tau}\}_\tau} \mathcal{E} \left[\sum_{\tau=t+1}^{\infty} \beta^{\tau-t-1} U(x_{i\tau}, s_{i\tau}, q_\tau; \theta) \mid \{q_{i\tau}\}_\tau, x_{i,t+1}, s_{i,t+1} \right]. \quad (4)$$

In what follows, we simplify notation by defining

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

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 nondecreasing in s_{it} , conditional on x_{it} .*

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

Proof: The optimal policy q is given by

$$\operatorname{argmax}_q \bar{U}(x, s, q) \equiv \{U(x, s, q; \theta) + \beta \mathcal{V}(x + q; \theta, \gamma)\}. \quad (5)$$

In order for $q(s, x; \theta, \gamma)$ to be non-decreasing in s given x , we require $\bar{U}(x, s, q; \theta)$ 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 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 nonparametric estimation methodology for non-additive (in the error term) random functions.

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 , respectively the utility function and shock distribution parameters. To simplify notation, we assume that our data are a balanced panel: $\{q_{it}, x_{it}\}$, $i = 1, \dots, N$, $t = 1, \dots, T$. This is not critical, as our estimator also applies to cases where the number of cross-sectional observations differs across time periods.

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 $q \leq q_{it}$, conditional on the observed state variable being equal to x_{it} .

Since the conditioning variable x is continuous, we employ a kernel estimator for these conditional CDFs:

$$\hat{F}(q|x) = \frac{\frac{1}{T} \frac{1}{N} \frac{1}{h} \sum_{t=1}^T \sum_{i=1}^N \mathbf{1}(q_{it} \leq q) K\left(\frac{x-x_{it}}{h}\right)}{\frac{1}{T} \frac{1}{N} \frac{1}{h} \sum_{t=1}^T \sum_{i=1}^N K\left(\frac{x-x_{it}}{h}\right)} \quad (6)$$

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 $q = 0$ (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 a r -th order kernel (with $r \geq 2$) function: (i) $\int K(u) du = 1$; (ii) $\int u^\xi K(u) du = 0$ for $\xi = 1, \dots, r-1$; and (iii) $\int u^r K(u) du < \infty$.

2. As $N \rightarrow \infty$, the bandwidth sequence (i) $h \rightarrow 0$; (ii) $\frac{Nh}{\log N} \rightarrow \infty$; and (iii) $\sqrt{N}h^r \rightarrow 0$.

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. (ii) The functions $U(x, s, q; \theta)$ and have 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 is uniformly bounded, continuous and bounded away from 0 on any compact set.

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 (as in Powell, Stock, and Stoker (1989)). Next, we describe our proposed two-step estimation approach.

3.1 First step: A 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 Eq. (3).³ 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 stock evolution process implies that the

³Recently, Berry and Pakes (2000) also exploit the first-order condition to derive estimates of structural parameters for models of *multi-agent* dynamic games. While we restrict our attention to single-agent problems, we focus on accommodating unobserved state variables, which are not present in the models considered by Berry and Pakes.

maximization problem for any agent i can be rewritten as

$$q(x_{it}, s_{it}; \theta, \gamma) = \operatorname{argmax}_q \{U(x_{it}, s_{it}, q; \theta) + \beta \mathcal{V}(x_{it} + q; \theta, \gamma)\}. \quad (7)$$

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

$$U_q(x_{it}, s_{it}, q_{it}; \theta) + \beta \mathcal{V}'(x_{it} + q_{it}; \theta, \gamma) = 0 \quad (8)$$

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

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

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 $F_s^{-1}(\hat{F}(q_{it}|x_{it}); \gamma)$, 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 immediately suggest a pairwise difference estimator for the 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 subtract the first-order conditions for these two observations, we obtain

$$\left\{ U_q \left(x_{it}, s \left(\hat{F}(q_{it}|x_{it}); \gamma \right), q_{it}; \theta \right) - U_q \left(x_{jt}, s \left(\hat{F}(q_{jt}|x_{jt}); \gamma \right), q_{jt}; \theta \right) \right\} = 0, \quad (9)$$

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

Let θ_1 denotes the subset of the parameters θ which enter Eq. (9). 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; (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:

$$\begin{aligned} \min_{\theta_1, \gamma} \frac{1}{(NT)^2} \sum_{t=1}^T \sum_{i=1}^N \sum_{t'=1}^T \sum_{j=1}^N \frac{1}{h} K \left(\frac{(x_{it} + q_{it}) - (x_{jt'} + q_{jt'})}{h} \right) \cdot I_{it} I_{jt'} \cdot \\ \left[U_q \left(x_{it}, s \left(\hat{F}(q_{it}|x_{it}); \gamma \right), q_{it}; \theta_1 \right) - U_q \left(x_{jt'}, s \left(\hat{F}(q_{jt'}|x_{jt'}); \gamma \right), q_{jt'}; \theta_1 \right) \right]^2 \quad (10) \\ \equiv \frac{1}{(NT)^2} \sum_{t=1}^T \sum_{i=1}^N \sum_{t'=1}^T \sum_{j=1}^N \left\{ \frac{1}{h} K \left(\frac{(x_{it} + q_{it}) - (x_{jt'} + q_{jt'})}{h} \right) \cdot I_{it} I_{jt'} \cdot \hat{m}(z_{it}, z_{jt'}; \psi)^2 \right\}. \end{aligned}$$

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

$$\begin{aligned} \hat{m}(z_{it}, z_{jt'}, \psi) \equiv U_q \left(x_{it}, F_s^{-1} \left(\hat{F}(q_{it}|x_{it}); \gamma \right), q_{it}; \theta_1 \right) \\ - U_q \left(x_{jt'}, F_s^{-1} \left(\hat{F}(q_{jt'}|x_{jt'}); \gamma \right), q_{jt'}; \theta_1 \right). \quad (11) \end{aligned}$$

The kernel function $K(\cdot)$ and bandwidth sequence $\{h\}$ obey Assumption 4 above. Moreover, in computing the objective function (10) above, we only include observations with non-zero investment ($q \neq 0$) because only for these observations is the first-order condition (8) satisfied.⁴

Given an estimate $\hat{\gamma}$ of the parameters in the shock distribution function, we can immediately derive an estimate of the optimal policy function

$$\tilde{q}(x, s) \equiv \hat{F}_{q|x}^{-1} (F_s(s; \hat{\gamma})), \quad \forall s. \quad (12)$$

Our 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 Honore and Powell (1994) pioneered the use of pairwise-differencing methods in econometrics. The objective function (10) resembles a weighted least squares objective, 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.

⁴ 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 in their earlier work). We thank a referee for pointing this out.

From Eq. (10), we can alternatively express the pairwise-difference estimate for ψ as that which solves the following sample score function:

$$W_{NT}(\hat{\psi}) \equiv \frac{1}{(NT)^2} \sum_i \sum_t \sum_j \sum_{t'} \hat{r}(z_{it}, z_{jt'}, \hat{\psi}) = 0 \quad (13)$$

where

$$\hat{r}(z_{it}, z_{jt'}, \hat{\psi}) \equiv \frac{1}{h} K\left(\frac{x_{it+1} - x_{jt'+1}}{h}\right) \hat{m}(z_{it}, z_{jt'}, \hat{\psi}) \frac{\partial}{\partial \psi} \left[\hat{m}(z_{it}, z_{jt'}, \hat{\psi}) \right] I_{it} I_{jt'}. \quad (14)$$

The limit objective function of the first-step estimator is

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

Also define $m(z_{it}, z_{jt'}, \psi)$ and $r(z_{it}, z_{jt'}, \psi)$ analogous to $\hat{m}(z_{it}, z_{jt'}, \hat{\psi})$ and $\hat{r}(z_{it}, z_{jt'}, \hat{\psi})$ except that $\hat{F}(q_{it}|x_{it})$ in (11) is replaced by the unknown true $F(q_{it}|x_{it})$.

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

Assumption 6 *Regularity conditions for first step:*

- i. $\psi \in \Psi$, a compact subset of \mathbf{R}^P , and true value $\psi^0 \in \text{int}(\Psi)$.
- ii. $G_0(\psi)$ is uniquely minimized at ψ^0 .
- iii. $r(z_{it}, z_{jt'}; \psi)$ is twice continuously differentiable in $\psi \in \Psi$ with probability 1.
- iv. $\sup_{\psi \in \Psi} |r(z_{it}, z_{jt'}; \psi)| < \bar{r}(z_{it}, z_{jt'})$ for some function $\bar{r}(\cdot)$ with $E[\bar{r}(z_{it}, z_{jt'})] < \infty$.
- v. Define $\tilde{v}(z_{it}, \psi) \equiv E[r(z_{it}, z_{jt'}, \psi)|z_{it}]$ and $\lambda(\psi) \equiv E\tilde{v}(z_{it}, \psi)$.
 - v.i. $\lambda(\psi^0) = 0$ and is differentiable at ψ^0 , with nonsingular Jacobian matrix A .
 - v.ii. The expectation $E[|r(z_{it}, z_{jt'}, \psi)|^2]$ exists and is finite.

The conditions listed 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 [Honore 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 the Appendix, Section A.1:

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

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

as $N \rightarrow \infty$, where A and Ω are defined in Eqs. (31) and (37) 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 Eq. (11)) would be identically zero for all values of $z_{it}, z_{jt'}$ such that $x_{it+1} = x_{jt'+1}$. Hence, the sampling variation in the estimate of ψ will be determined completely from the sampling variation in the nonparametric estimates of the conditional distributions $F(q|x)$ (using Eq. (6)).

3.1.2 Remarks on first step

Discussion of assumptions Next, we discuss several of the assumptions we made previously, and how they could potentially 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 , as follows:

Let $s_k \equiv F_s^{-1}(\tau_k)$ denote the τ_k -th quantile of the shock distribution F_s . Let κ denote the total number of quantiles to be estimated (and the corresponding quantile values by $\tau_1 < \tau_2 < \dots < \tau_\kappa$). For any fixed κ , we approximate the distribution of the shocks F_s via a piece-wise 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

$$\hat{F}_s^{-1}(\tau) \equiv \begin{cases} \tau \frac{s_1}{\tau_1} & \text{if } \tau \in [0, \tau_1] \\ s_{i-1} + (\tau - \tau_{i-1}) \frac{s_i - s_{i-1}}{\tau_i - \tau_{i-1}} & \text{if } \tau \in (\tau_{i-1}, \tau_i], \quad i = 2, \dots, \kappa - 1. \\ s_{\kappa-1} + (\tau - \tau_{\kappa-1}) \frac{s_\kappa - s_{\kappa-1}}{\tau_\kappa - \tau_{\kappa-1}} & \text{if } \tau \in (\tau_{\kappa-1}, 1]. \end{cases} \quad (15)$$

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

Second, 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 reparameterize the utility function so that

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

The monotonicity assumption 3 would be is a natural consequence of this reparametrization: holding x fixed, q is monotonic in ϵ .⁵ We do not do this for two reasons. First, an important difference with Matzkin’s (2003) paper is that she considers the case where U is nonparametric. In that case, the assumption that s is $U[0, 1]$ is a normalization in the sense that she would not be able to identify both U and F_s for some arbitrary unknown shock distribution. For us, U (and, indeed, also F_s) is parametric, so such a normalization is not necessary. Second, in the empirical application below, we are interested in interpreting the magnitude of the shocks, which would not be so straightforward if the shock were assumed to be uniformly distributed.

Third, the deterministic accumulation assumption (2) is simpler and more restrictive than necessary, but we have made it for expositional convenience, and because it is a natural one to make for the investment example (and later for our empirical application). More broadly, however, it suffices to match on x_{it+1} : in the pairwise-differencing step, the variation in (x_{it}, q_{it}) conditional on x_{it+1} (and $q_{it} \neq 0$) is the crucial variation which identifies the utility function parameters. Therefore, more complicated laws of motion for x can be accommodated, including nonlinear functions $x_{it+1} = l(x_{it}, q_{it}, \zeta)$ which include unknown parameters ζ .⁶ Moreover, we can also introduce additional observable (and potentially time-varying) characteristics z_{it} specific to individual i and period t —these would simply be additional variables which we need to match upon.⁷

The independence assumption (1), on the other hand, is crucial 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 $F(s_{t+1}|x_t)$ 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 , leading to a degenerate estimating equation (9).

⁵We thank a referee for pointing this out.

⁶In this case, since x_{it+1} , x_{it} , and q_{it} are observed, and the functional form of l is known, ζ can be estimated separately, apart from the rest of the model parameters. We thank a referee for pointing this out.

⁷Given the deterministic accumulation equation, we could reparameterize 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 obtain, 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} .

Discussion of identification Before proceeding to the second step, we also present some discussion of identification. Particularly, we want to consider how the parameters γ of the shock distribution are pinned down in the pairwise-differencing step. Since we focus on identifying F_s , we assume for convenience that $\theta_1 = \{\}$ (so that there are no θ_1 parameters to estimate). Consider how pairwise-differencing allows us to estimate γ . In order to do pairwise-differencing, we need (at least) two individuals (i, j) such that

$$0 = U_q(x_i, s(F(q_{it}|x_{it}); \gamma), q_i) - U_q(x_j, s(F(q_{jt}|x_{jt}); \gamma), q_j) \quad (16)$$

where $s(\tau; \gamma) \equiv F_s^{-1}(\tau; \gamma)$ denotes the τ -th quantile of $F_s(\cdot; \gamma)$. In order for Eq. (16) to not be a trivial function, individuals i and j 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 example, consider the linear-quadratic case, where the policy function will be linear in its arguments: $q = a + bx + cs$. Conditions (i) and (ii) from the previous paragraph require the existence of pairs (i, j) such that $x_i \neq x_j$ and also

$$\begin{aligned} x_i + q_i &= x_j + q_j \\ \Leftrightarrow x_i + a + bx_i + cs(\tau_i; \gamma) &= x_j + a + bx_j + cs(\tau_j; \gamma), \end{aligned}$$

where $\tau_i, \tau_j \in [0, 1]$ denote the quantiles of the shock distribution. Clearly the equation above admits multiple solutions of $(x_i, \tau_i, x_j, \tau_j)$ with $x_i \neq x_j$ and $\tau_i \neq \tau_j$. As long as this set of values contains enough observations where $q_i \equiv F^{-1}(\tau_i|x_i) \neq q_j \equiv F^{-1}(\tau_j|x_j)$, then γ can be recovered, by essentially running the nonlinear regression (16) using these observations.

Identification in our framework differs from other papers in the literature. For example, [Magnac and Thesmar \(2002\)](#), [Pesendorfer and Schmidt-Dengler \(2003\)](#), and [Aguirregabiria \(2005\)](#) all consider dynamic discrete choice models, and focus on the nonparametric 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 differences, identification issues are different (and simpler) 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.

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 parameters in θ which were not in the subset θ_1 estimated in the first step. Recall that $\theta_2 \equiv \{\theta \setminus \theta_1\}$ denotes the set of parameters which were not estimated in the first step.

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) where the investment level $q_{it} \neq 0$ as follows:

$$h_{it} \left(x_{it}, q_{it}; \hat{\psi}, \theta_2 \right) \equiv U_q \left(x_{it}, s(\hat{F}(q_{it}|x_{it}); \hat{\gamma}), q_{it}; \hat{\theta}_1, \theta_2 \right) + \beta \mathcal{V}' \left(x_{it} + q_{it}; \hat{\psi}, \theta_2 \right) = 0. \quad (17)$$

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). Due to sampling error from estimating θ_1 , γ and $F(q|x)$ in the first step, the first order condition $h_{it}(x_{it}, q_{it}; \hat{\psi}, \theta_2)$ need not be identically zero, even at the true parameter vector θ_0 . Therefore, we estimate θ_2 via a least squares procedure:⁸

$$\hat{\theta}_2 = \operatorname{argmin}_{\theta_2} \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T I_{it} \cdot \left[h_{it}(x_{it}, q_{it}; \hat{\psi}, \theta_2, \hat{F}_{q|x}(\cdot|\cdot)) \right]^2. \quad (18)$$

As in the first step, we can only include observations with non-zero investment ($q \neq 0$) in the objective function. Indeed, both step of our estimation procedure are based on agents' first-order conditions, and thus only use the observations where $q_{it} \neq 0$. In our estimation procedure, the observations with $q_{it} = 0$ are employed only in the construction of the conditional distributions $\hat{F}(q|x)$ (cf. Eq. (6)).

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 \(1998\)](#),

⁸The choice of a square norm is somewhat arbitrary; 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.

can be difficult since it involves 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 is available to avoid numerical computation of the dynamic programming problem. In this alternative, the value function is computed by a forward integration procedure in the spirit of [Hotz and Miller \(1993\)](#). This procedure exploits the representation of the value function at time t as the expected discounted sum of future utilities (cf. Eq. (4)) rather than the more familiar recursive representation via the Bellman equation (cf. Eq. (3)) 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 (4) can be represented as forward integration over the observed conditional probabilities $\hat{F}(q|x)$ (cf. Eq. (3.12) in [Hotz and Miller \(1993\)](#)).

Under the 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:

$$\begin{aligned} \mathcal{V}(x_1; \hat{\psi}, \theta_2) = & \int \int \cdots \int \left\{ \left[\sum_{z=1}^T \beta^{z-1} U(x_z, s_z, \hat{F}_{q|x_z}^{-1}(F_s(s_z)); \hat{\theta}_1, \theta_2) \right] \right. \\ & \left. + \beta^T CV(x_{T+1}) \right\} dF(s_1; \hat{\gamma}) dF(s_2; \hat{\gamma}) \cdots dF(s_T; \hat{\gamma}). \end{aligned} \quad (19)$$

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_1 and $x_z = x_{z-1} + \hat{F}_{q|x_{z-1}}^{-1}(F_s(s_{z-1}))$ for $z = 1, \dots, T$.

More succinctly, let $\{\tau\} = \{\tau_1, \tau_2, \dots, \tau_T\}$ denote a sequence of *i.i.d.* $U[0, 1]$ random variables. The expected value function can be written as

$$\mathcal{V}(x_1; \hat{\psi}, \theta_2) = E_{\{\tau\}} \left\{ \sum_{z=1}^T \beta^{z-1} U(x_z, s(\tau_z; \hat{\gamma}), \hat{F}_{q|x_z}^{-1}(\tau_z); \hat{\theta}_1, \theta_2) + \beta^T CV(x_{T+1}) \right\}. \quad (20)$$

In the above expression, given the starting value x_1 , the subsequent sequence of stocks x_2, x_3, \dots is related to the uniform random draws τ 's by the relation $x_z = x_{z-1} + \hat{F}_{q|x_{z-1}}^{-1}(\tau_{z-1})$. 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} . We will provide more details about this below.

In practice, the multidimensional integration involved to compute the expected value function (Eqs. (19) or (20)) presents computational challenges, and so we simulate the expected value function by following Hotz, Miller, Sanders, and Smith (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, $\mathcal{V}(x_1; \hat{\psi}, \theta_2)$ (using Eq. (20)) can be simulated by

$$\mathcal{V}^S(x_1; \hat{\psi}, \theta_2) = \frac{1}{S} \sum_{l=1}^S \left\{ \left[\sum_{z=1}^T \beta^{z-1} U(x_z^l, s(\tau_z^l; \gamma), \hat{F}_{q|x_z^l}^{-1}(\tau_z^l); \theta) \right] + \beta^T CV(x_{T+1}^l) \right\} \quad (21)$$

where

- $\tau_z^l, l = 1, \dots, S, z = 1, \dots, T$ are *i.i.d.* $U[0, 1]$.
- $x_z^l = \begin{cases} x_1 & \text{for } z = 1 \\ x_{z-1}^l + q(x_{z-1}^l, s(\tau_z^l; \gamma)) & \text{for } z = 2, \dots, T + 1. \end{cases}$

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

$$\mathcal{V}^{S'}(x_{it}; \hat{\psi}, \theta_2) \approx \frac{\mathcal{V}^S(x_{it} + \Delta; \hat{\psi}, \theta_2) - \mathcal{V}^S(x_{it}; \hat{\psi}, \theta_2)}{\Delta} \quad (22)$$

for Δ small. By plugging in Eq. (22) for $\mathcal{V}'(x_{it} + q_{it}; \hat{\psi}, \theta_2)$ in Eq. (17), we can estimate θ_2 by minimizing the objective function (18).

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 $\mathcal{V}(x_{it}; \hat{\psi}, \hat{\theta}_2)$ as a known function for all $(\hat{\psi}, \theta_2)$. For the simulation-based approximation of the expected value function, we require that the number of simulation draws S increases quickly enough as $N \rightarrow \infty$ so that variation due to the simulation itself is small enough and does not affect the asymptotic variance. From [Gourieroux and Monfort \(1996\)](#), a sufficient condition for the asymptotic variance to be unaffected from simulation error is that $S/\sqrt{N} \rightarrow \infty$.⁹

⁹For numeric dynamic programming methods, which usually are based on iterative function approximation algorithms, this generally requires that the accuracy of the function approximation (as measured in terms of the order of an approximating polynomial, or number of knot points in an approximating spline) increase as $N \rightarrow \infty$.

The second step estimator $\hat{\theta}_2$ solves the sample score function:

$$J_{NT}(\theta_2) \equiv \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \bar{h} \left(x_{it}, q_{it}; \hat{\psi}, \theta_2, \hat{F}_{q|x}(\cdot|\cdot) \right) I_{it} = 0. \quad (23)$$

where

$$\begin{aligned} & \bar{h} \left(x_{it}, q_{it}, \hat{F}_{q|x}(\cdot|\cdot), \hat{\psi}; \theta_2 \right) \\ \equiv & \left[U_q \left(x_{it}, F_s^{-1} \left(\hat{F}(q_{it}|x_{it}); \hat{\gamma} \right), q_{it}; \hat{\theta}_1, \theta_2 \right) + \beta \mathcal{V}^{S'} \left(x_{it} + q_{it}; \hat{\psi}, \theta_2, \hat{F}_{q|x}(\cdot) \right) \right] * \\ & \frac{\partial}{\partial \theta_2} \left[U_q \left(x_{it}, F_s^{-1} \left(\hat{F}(q_{it}|x_{it}); \hat{\gamma} \right), q_{it}; \hat{\theta}_1, \theta_2 \right) + \beta \mathcal{V}^{S'} \left(x_{it} + q_{it}; \hat{\psi}, \theta_2, \hat{F}_{q|x}(\cdot) \right) \right]. \end{aligned}$$

The notation $\hat{F}_{q|x}(\cdot)$ denotes the whole set of estimated conditional quota distributions, estimated as in Eq. (6). The inclusion of the entire conditional distribution $\hat{F}_{q|x}(\cdot|\cdot)$ as an argument in $\bar{h}(\dots)$ (in addition to $\hat{F}(q_{it}|x_{it})$) recognizes that the expected value function $\mathcal{V}(x_{it+1})$ (cf. Eq. (20)) depends on the entire set of functions $\hat{F}_{q|x}^{-1}(\cdot)$; $\forall x$, not just on any one of these functions evaluated at a particular quantile.

Let $P_2 \equiv \dim(\theta_2)$, and define

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

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

The regularity conditions required for deriving the asymptotic result of the second step estimator are collected in the following assumption.

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

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

3. $h(x, q; \psi, \theta_2, F_{q|x}(\cdot|\cdot))$ is twice continuously differentiable in θ_2 , ψ and $F_{q|x}(\cdot)$ with probability 1. The function and its derivatives are uniformly bounded by an integrable function.

4. The Jacobian \bar{A} of $\mu(\theta_2) \equiv E \bar{h} \left(x_{it}, q_{it}; \psi^0, \theta_2, F_{q|x}^0(\cdot|\cdot) \right) I_{it}$ is nonsingular at θ_2^0 .

The full proof of the following is in the appendix, section A.2.

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

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

In addition,

$$\sqrt{NT}(\hat{\theta}_2 - \theta_2^0) \xrightarrow{d} N(0, \bar{A}^{-1}\bar{\Omega}\bar{A}^{-1})$$

as $N \rightarrow \infty$, where \bar{A} and $\bar{\Omega}$ are defined in Eqs. (39) and (41) of the appendix.

At the true values of ψ^0 , θ_2^0 , and $F(q|x)$, the first-order condition (17) is identically zero for all values of (x_{it}, q_{it}) which are optimally chosen. Hence, the second-step estimation introduces no source of sampling variation beyond that which arises from the first-step estimation of ψ , and the estimation of the conditional distributions $F(q|x)$. This is made explicit in the proof of Theorem 2.

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 which 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, therefore 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 was estimated in a first-step to reduce the computational burden in the second-step, which involved value function iteration.

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 quota have been traded among dairy farmers in monthly double auctions administered by the Dairy Farmers of Ontario (DFO) (cf. Biggs (1990)). This paper analyzes data from the eleven 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/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 $q(x_{it}, s_{it}, p_t)$ which takes values in $(-\infty, \infty)$. For positive values of $q(\dots)$, 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 $q(x_{it}, s_{it}, p_t^*)$, 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_t to buy. In the appendix (section B), we show that 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 nonstationary dynamic optimization problem solved by each trader individually. The problem is nonstationary 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.¹⁰ The nonstationarity

¹⁰In 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 eleven observations of the price process.

of the dynamic problem is the main difference between the model used in our empirical application, and the stationary problem used in the previous sections in describing our estimation procedure.

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:

$$\max_{\{q_{it}\}_t} \mathcal{E}_{0|\{q_{it}\}} \sum_{t=0}^{\infty} \beta^t U_t(x_{it}, s_{it}, q_{it}, p_t^*; \theta) \quad (24)$$

subject to

$$x_{it+1} = x_{it} + q_{it}; \quad s_{it+1} \sim F_s; \quad p_0^*, p_1^* \dots \text{known.} \quad (25)$$

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 $\mathcal{E}_{0|\{q_{it}\}}$ is over the sequences of x_{it} and s_{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:

$$q_t(x_{it}, s_{it}) = \operatorname{argmax}_q U(x_{it}, s_{it}, q_{it}, p_t^*; \theta) + \beta \mathcal{V}_{t+1}(x_{it} + q_{it}; \psi, \theta_2) \quad (26)$$

where

$$\mathcal{V}_{t+1}(x_{it} + q_{it}; \psi, \theta_2) \equiv \mathcal{E}_{s_{it+1}} V(x_{it} + q_{it}, s_{it+1}).$$

Accommodating nonstationary 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 nonstationary problem, we estimate the conditional quota purchase distributions $F_t(q_t|x_t)$ (cf. Eq. (6)) separately for each period t . Second, because the expected value functions are no longer time-homogenous in a nonstationary 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

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

which differs from the objective function for the stationary case (Eq. (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 eleventh month) to depend on $x_{12} = x_{11} + q_{11}$, the stock that a given trader has after the first eleven months. More specifically, for months $t = 1, \dots, 11$, we simulate the expected value function as

$$\mathcal{V}_t^S(x_{t+1}; \hat{\psi}, \theta_2) = \frac{1}{S} \sum_{l=1}^S \left\{ \left[\sum_{z=t+1}^{11} \beta^{z-t} U(x_z^l, s(\tau_z^l; \gamma), \hat{F}_{q|x_z^l}^{-1}(\tau_z^l); \theta) \right] + \beta^{12-t} CV(x_{12}^l) \right\} \quad (27)$$

where

- τ_z^l , $l = 1, \dots, S$, $z = t + 1, \dots, T$ are *i.i.d.* $U[0, 1]$.
- $x_z^l = \begin{cases} x_t & \text{for } z = t + 1 \\ x_{z-1}^l + q(x_{z-1}^l, s(\tau_z^l; \gamma)) & \text{for } z = t + 2, \dots, 12. \end{cases}$
- the continuation value function is a flexible (fifth-order) polynomial in x_{12} :

$$CV(x_{12}) = \sum_{j=1}^5 \eta_j \cdot x_{12}^j.$$

We estimate the polynomial coefficients η_1, \dots, η_5 are jointly with θ_2 in the second step of our procedure.

4.1 Data: summary statistics

Summary statistics are presented in Table 3. The trading unit for quota is expressed in kilograms of butterfat, and one kilogram of quota purchased on the exchange allows a producer to ship one kilogram of butterfat per day, in perpetuity, for as long as the unit of quota is held.¹¹ Over the eleven exchanges, we observe the bids placed by 2,574 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.

¹¹Prior to September 1997, a unit of quota conferred on its owner the right to produce milk containing one kilogram of butterfat *per year*. In September 1997, however, the trading unit for quota was re-defined in kilograms of butterfat per day.

Column (E) in Table (3) shows that a large number of sellers and buyers participate in each exchange, which suggests that there may not be much scope for strategic behavior, which we have 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) nonparticipation 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 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 3 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. 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 CARA form for the utility function:

$$U(w_{it}) = -\exp(-rw_{it})$$

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

$$w_{it} = x_{it} \cdot s_{it} - p_t \cdot q_{it} - K \cdot \mathbf{1}(q_{it} \neq 0). \quad (28)$$

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 $p_t \cdot q(x_{it}, s_{it}, p_t)$ 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 3).¹² 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 . To 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 (9) becomes

$$-pre^{rK} (\exp(x_i s(F_s(q_i|x_i); \gamma) - pq_i) - \exp(x_j s(F_s(q_j|x_j); \gamma) - pq_j)).$$

The constant proportion pre^{rK} does not have any sampling variation, and hence is not identified in the first stage estimation using equation (9). The part that has sampling variation contains only the shock distribution parameters γ as the parameters estimated in the first step. Accordingly, the parameters which are identified in the second step are $\theta_2 \equiv (r, K)$.

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

For each specification, we used the bootstrap as follows: we re-sampled (with replacement) sequences from the dataset, and re-estimated the model for each re-sampled dataset. The reported bootstrap confidence intervals are therefore the empirical quantiles of the distribution of parameter estimates obtained in this fashion. We employed 50 bootstrap resamples in computing each set of standard errors.

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 $\log s \sim$

¹² 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.

$N(\mu, \sigma^2)$. The parameter estimates are shown in Table 1.

These magnitudes imply that the mean (and median) shock is 6.928. Given the specification of the agents' payoffs (Eq. (28)), this can be interpreted as the monthly return from a unit of quota (in 1986 Canadian \$'000). 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 two months ($=\frac{11,000}{6,928}$): 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 $\tilde{q}_1(x, s)$ for the log-normal distribution results. The policy function is estimated using Eq. (12) above.

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 Eq. (15) 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 ($=\frac{11,000}{1,240}$) months: this appears more realistic than the estimate obtained from the log-normal parameterization, reported above.¹³

In the bottom graph of Figure 2, we present our estimate of the implied period 1 (September 1997) policy function $\tilde{q}_1(x, s)$ 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).

5 Conclusions and Extensions

In this paper, we proposed a new two-step pairwise-differencing procedure for structural estimation of a dynamic optimization model with unobserved state variables. To our knowledge, our estimator represents the first application of pairwise-difference methods, which

¹³We 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.

have primarily been used in cross-sectional contexts (cf. [Honore 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. In accommodating serial correlation, we would have to consider carefully the problem of initial conditions which, in turn, is very closely related to the issue of unobserved individual-specific heterogeneity (cf. [Heckman \(1981\)](#)). In future work, we plan to explore extensions to our procedure to handle these issues.

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 to the econometrician. It will be interesting to investigate in future work whether monotonicity and quantile invariance can be useful in these situations.

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 P. MIRA (2007): "Sequential Estimation of Dynamic Discrete Games," Forthcoming, *Econometrica*.
- AHN, H., AND J. POWELL (1993): "Semiparametric Estimation of Censored Selection Models with a Nonparametric Selection Mechanism," *Journal of Econometrics*, 58, 3–29.
- AIT-SAHALIA, Y. (1996): "The Delta Method for Nonparametric Kernel Functionals," Department of Economics, Princeton University.
- BAJARI, P., AND L. BENKARD (2005): "Demand Estimation With Heterogeneous Consumers and Unobserved Product Characteristics: A Hedonic Approach," *Journal of Political Economy*, 113, 1239–1276.
- BAJARI, P., L. BENKARD, AND J. LEVIN (2007): "Estimating Dynamic Models of Imperfect Competition," Forthcoming, *Econometrica*.
- BERRY, S., M. OSTROVSKY, AND A. PAKES (2004): "Simple Estimators for the Parameters of Discrete Dynamic Games (with Entry Exit Examples)," NBER Working Paper, #10506.
- BERRY, S., AND A. PAKES (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 A. MONFORT (1996): *Simulation-Based Econometric Methods*. 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 *Structural Analysis of Discrete Data*, ed. by C. Manski, and D. McFadden. MIT Press.
- HONORE, B., AND J. POWELL (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 R. MILLER (1993): "Conditional Choice Probabilities and the Estimation of Dynamic Models," *Review of Economic Studies*, 60, 497–529.
- HOTZ, J., R. MILLER, S. SANDERS, AND J. SMITH (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*. MIT Press.
- MAGNAC, T., AND D. THESMAR (2002): "Identifying Dynamic Discrete Decision Processes," *Econometrica*, 70, 801–816.
- MATZKIN, R. (2003): "Nonparametric Estimation of Nonadditive Random Functions," *Econometrica*, 71, 1339–1376.
- OLLEY, S., AND A. PAKES (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 *Advances in Econometrics: Proceedings of the 1990 Meetings of the Econometric Society*, ed. by C. Sims, and J.-J. Laffont, vol. 2, pp. 171–259. Cambridge University Press.
- PAKES, A., AND M. SIMPSON (1989): "Patent Renewal Data," *Brookings Papers on Economic Activity: Microeconomics*, pp. 331–403.
- PESENDORFER, M., AND P. SCHMIDT-DENGLER (2003): "Identification and Estimation of Dynamic Games," NBER Working Paper, #9726.
- POWELL, J., J. STOCK, AND T. STOKER (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.

——— (1996): “Numerical Dynamic Programming in Economics,” in *Handbook of Computational Economics*, ed. by D. K. H. Amman, and J. Rust. North Holland.

TABER, C. (2000): “Semiparametric identification and heterogeneity in discrete-choice dynamic programming models,” *Journal of Econometrics*, 96, 201–229.

A Proofs

A.1 Proof of Theorem 1

For convenience, we 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, Eqs. (11) and (14) evaluated at the actual (ie. error-free) conditional distributions $F_{q|x}$.

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

$$\begin{aligned}
& \hat{F}(q|x) - F(s) \\
& \approx \frac{1}{f(x)} \left[\frac{1}{NT} \sum_{l=1}^N \sum_{t'=1}^T \mathbf{1}(q_{lt'} < q) K\left(\frac{x_{lt'} - x}{h}\right) - f(x) F(s) \right] \\
& \quad - \frac{F(s)}{f(x)} \left[\frac{1}{NT} \sum_{l=1}^N \sum_{t'=1}^T K\left(\frac{x_{lt'} - x}{h}\right) - f(x) \right] \\
& = \frac{1}{f(x)} \left[\frac{1}{NT} \sum_{l=1}^N \sum_{t'=1}^T \mathbf{1}(q_{lt'} < q) K\left(\frac{x_{lt'} - x}{h}\right) \right] - \frac{F(s)}{f(x)} \left[\frac{1}{NT} \sum_{l=1}^N \sum_{t'=1}^T K\left(\frac{x_{lt'} - x}{h}\right) \right].
\end{aligned} \tag{29}$$

Together with other smoothness conditions in assumption 5, uniform consistency of $\hat{F}(q|x)$ implies uniform convergence of the estimand (10) 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

$$\sqrt{NT}(\hat{\psi} - \psi^0) = -A_{NT}^{-1} (1 + o_p(1)) \frac{1}{(NT)^{3/2}} \sum_{i=1}^N \sum_{t=1}^T \sum_{j=1}^N \sum_{t'=1}^T \hat{r}(z_{it}, z_{jt'}, \psi^0) \tag{30}$$

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

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

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

$$A \equiv E_{z_{jt'}} E_{z_{it}} \left[I_{it} I_{jt'} \frac{\partial}{\partial \psi} [m(z_{it}, z_{jt'}, \psi)] \frac{\partial}{\partial \psi} [m(z_{it}, z_{jt'}, \psi)] \mathbf{1}(x_{it+1} = x_{jt'+1}) \right], \tag{31}$$

is the same matrix as stated in condition v.i of Assumption 6. The form of A takes into account the fact that

$$E_{z_{jt'}} E_{z_{it}} \left[I_{it} I_{jt'} m(z_{it}, z_{jt'}, \psi) \frac{\partial^2}{\partial \psi \partial \psi'} m(z_{it}, z_{jt'}, \psi) \mathbf{1}(x_{it+1} = x_{jt'+1}) \right] \equiv 0, \tag{32}$$

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

Next, we address the terms which appear behind the quadruple summation in (30). Define

$$\begin{aligned}\hat{w}(z_{it}, z_{jt'}) &\equiv \hat{m}(z_{it}, z_{jt'}, \psi^0) \frac{\partial}{\partial \psi} [\hat{m}(z_{it}, z_{jt'}, \psi^0)], \\ w(z_{it}, z_{jt'}) &\equiv m(z_{it}, z_{jt'}, \psi^0) \frac{\partial}{\partial \psi} [m(z_{it}, z_{jt'}, \psi^0)].\end{aligned}\tag{33}$$

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

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

where substituting in Eq. (29) above:

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

Hence, we can approximate the linear term in equation (30) by a U-statistic representation:

$$\begin{aligned}&\frac{1}{(NT)^{3/2}} \sum_{i=1}^N \sum_{t=1}^T \sum_{j=1}^N \sum_{t'=1}^T \frac{1}{h} K\left(\frac{x_{it+1} - x_{jt'+1}}{h}\right) w(z_{it}, z_{jt'}) I_{it} I_{jt'} \\ &+ \frac{1}{(NT)^{5/2}} \sum_{i=1}^N \sum_{t=1}^T \sum_{j=1}^N \sum_{t'=1}^T \sum_{l=1}^N \sum_{t''=1}^T \frac{1}{h} K\left(\frac{x_{it+1} - x_{jt'+1}}{h}\right) v(z_{it}, z_{jt'}, z_{lt''}) I_{it} I_{jt'}.\end{aligned}\tag{34}$$

Given our assumptions on the kernel and bandwidth sequence (Assumption 4 in the main text), the bias terms in the nonparametric 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 (34). For the first term in Eq. (34), we have

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

Both terms in the above display vanish asymptotically for the same reasoning that leads to (32). This makes explicit the feature that the pairwise-differencing step introduces no additional variation to the parameter estimate $\hat{\psi}$. The nonparametric 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 (34):

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

The first inequality follows from Assumption 4 which implies that the other two projection terms vanish. After tedious but straightforward calculations we can write $\tilde{v}(z_{lt''})$ as:

$$\begin{aligned} & E_{z_{it}} \left[E_{z_{jt'}} \left(I_{it} I_{jt'} \frac{\partial w(z_{it}, z_{jt'})}{\partial F_s(s_{it})} \Big|_{x_{jt'+1} = x_{it+1}} \right) f(x_{it+1}) (\mathbf{1}(q_{lt''} < q_{it}) - F_s(s_{it})) \Big|_{x_{it} = x_{lt''}} \right] \\ &+ E_{z_{jt'}} \left[E_{z_{it}} \left(I_{it} I_{jt'} \frac{\partial w(z_{it}, z_{jt'})}{\partial F_s(s_{jt'})} \Big|_{x_{it+1} = x_{jt'+1}} \right) f(x_{jt'+1}) (\mathbf{1}(q_{lt''} < q_{jt'}) - F_s(s_{jt'})) \Big|_{x_{jt'} = x_{lt''}} \right]. \end{aligned} \quad (35)$$

Therefore, we conclude that

$$\sqrt{NT} (\hat{\psi} - \psi) \xrightarrow{d} N(0, A^{-1} \Omega A^{-1}) \quad (36)$$

where

$$\Omega = \frac{1}{T} E \left(\sum_{t=1}^T \tilde{v}(z_{it}) \right) \left(\sum_{t=1}^T \tilde{v}(z_{it}) \right)' \quad (37)$$

and A is defined as in (31) above. The asymptotic variance can be consistently estimated using resampling methods or empirical analogs. \blacksquare

A.2 Proof of Theorem 2

In this proof, we abstract away from approximation error in computing the value function. If the value function is simulated, this requires that a sufficiently large number of simulation draws to compute the value function so that the variation due to the simulation itself is small enough and does not affect the asymptotic variance. However, the estimation error from the previous steps of estimating $\hat{\psi} = \{\hat{\theta}_1, \hat{\gamma}\}$ and $\hat{F}_{q|x}(\cdot|\cdot)$ will be reflected in the variance of the second-step estimator.

First we recall the following linear approximation, for $\tau = F(q|x)$:

$$\begin{aligned} & \hat{F}(q|x) - F(q|x) \\ &= (1 + o_p(1)) \frac{1}{NT} \sum_{l=1}^N \sum_{t=1}^T \frac{1}{f(x)} \left[\mathbf{1}(q_{lt} \leq F_{q|x}^{-1}(\tau)) - \tau \right] \frac{1}{h} K \left(\frac{x_{lt} - x}{h} \right) \\ &\equiv (1 + o_p(1)) \frac{1}{NT} \sum_{l=1}^N \sum_{t=1}^T G_h(q_{lt}, x_{lt}, \tau, x). \end{aligned} \quad (38)$$

Taking a Taylor expansion of (23) around $\hat{\theta}_2$, one obtains

$$\begin{aligned} 0 &= \frac{1}{NT} \sum_i \sum_t \bar{h} \left(x_{it}, q_{it}; \hat{F}_{q|x}(\cdot|\cdot), \hat{\psi}, \hat{\theta}_2 \right) I_{it} \\ &= \frac{1}{NT} \sum_i \sum_t \bar{h} \left(x_{it}, q_{it}; \hat{F}_{q|x}(\cdot|\cdot), \hat{\psi}, \theta_2^0 \right) I_{it} \\ &\quad + \left(\hat{\theta}_2 - \theta_2^0 \right) \frac{1}{NT} \sum_i \sum_t \frac{\partial}{\partial \theta_2} \bar{h} \left(x_{it}, q_{it}; \hat{F}_{q|x}(\cdot|\cdot), \hat{\psi}, \theta_2^0 \right) I_{it} + o_p(\hat{\theta}_2 - \theta_2^0). \end{aligned}$$

A standard law of large numbers applies to the Jacobian term in the above expression:

$$\begin{aligned} &\frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \frac{\partial}{\partial \theta_2} \bar{h} \left(x_{it}, q_{it}; \hat{F}_{q|x}(\cdot|\cdot), \hat{\psi}, \theta_2^0 \right) I_{it} \\ &\xrightarrow{p} \bar{A} \equiv E \frac{\partial}{\partial \theta_2} \bar{h} \left(x_{it}, q_{it}; F_{q|x}(\cdot|\cdot), \psi^0, \theta_2^0 \right) I_{it}. \end{aligned} \tag{39}$$

with \bar{A} the same matrix as specified in Assumption 7. Hence,

$$\sqrt{NT} \left(\hat{\theta}_2 - \theta_2^0 \right) = \bar{A}^{-1} \frac{1}{\sqrt{NT}} \sum_{i=1}^N \sum_{t=1}^T \bar{h} \left(x_{it}, q_{it}; \hat{F}_q(\cdot|\cdot), \hat{\psi}, \theta_2^0 \right) I_{it} + o_p(1). \tag{40}$$

The recursive use of the nonparametric estimates $\hat{F}_{q|x}(\cdot|\cdot)$ in the construction of the expected value function in Eq. (20) makes it tedious to derive explicit analytic expressions for the asymptotic linear representation of the nonlinear functional $\bar{h} \left(x_{it}, q_{it}, \hat{F}_{q|x}(\cdot|\cdot), \hat{\psi}, \theta_2 \right)$ as a function of $\hat{F}_{q|x}(\cdot|\cdot)$. Hence, in the following we will denote a linear function in $\hat{F}_{q|x}(\cdot|\cdot)$ by

$$g \left(x_{it}, q_{it}, \hat{F}_{q|x}(\cdot|\cdot) - F_{q|x}(\cdot|\cdot), \psi, \theta_2 \right)$$

without explicitly writing out its lengthy analytic formula.

Given this notation, the following asymptotically linear representation holds

$$\begin{aligned} &\frac{1}{\sqrt{NT}} \sum_{i=1}^N \sum_{t=1}^T \bar{h} \left(x_{it}, q_{it}, \hat{F}_{q|x}(\cdot|\cdot), \hat{\psi}, \theta_2 \right) I_{it} - \frac{1}{\sqrt{NT}} \sum_{i=1}^N \sum_{t=1}^T \bar{h} \left(x_{it}, q_{it}, F_{q|x}(\cdot|\cdot), \psi^0, \theta_2^0 \right) I_{it} \\ &= \frac{1}{\sqrt{NT}} \sum_{i=1}^N \sum_{t=1}^T g \left(x_{it}, q_{it}, \hat{F}_{q|x}(\cdot|\cdot) - F_{q|x}(\cdot|\cdot), \psi^0, \theta_2^0 \right) + B \sqrt{NT} \left(\hat{\psi} - \psi^0 \right) + o_p(1). \end{aligned}$$

In the above display,

$$B \equiv E \frac{\partial}{\partial \psi} \bar{h} \left(x_{it}, q_{it}, F_{q|x}(\cdot|\cdot), \psi^0, \theta_2^0 \right) I_{it}.$$

Using the linear approximation in (38), this can be further approximated as

$$\begin{aligned}
& \frac{1}{\sqrt{NT}} \sum_{l=1}^N \sum_{t=1}^T \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T g(x_{it}, q_{it}, G_h(q_{lt}, x_{lt}, \cdot, \cdot), \psi^0, \theta_2^0) + B\sqrt{NT} (\hat{\psi} - \psi^0) + o_p(1) \\
&= \frac{1}{\sqrt{NT}} \sum_{l=1}^N \sum_{t=1}^T E \left[g(x_{it}, q_{it}, G_h(q_{lt}, x_{lt}, \cdot, \cdot), \psi^0, \theta_2^0) \middle| q_{lt}, x_{lt} \right] + B\sqrt{NT} (\hat{\psi} - \psi^0) + o_p(1) \\
&= \frac{1}{\sqrt{NT}} \sum_{l=1}^N \sum_{t=1}^T E_{z_{it}} [g(x_{it}, q_{it}, G_h(q_{lt}, x_{lt}, \cdot, \cdot), \psi^0, \theta_2^0)] + B\sqrt{NT} (\hat{\psi} - \psi^0) + o_p(1)
\end{aligned}$$

In the above display G_h is defined as in equation (38).

Next, we use the modeling assumption that

$$\bar{h}(x_{it}, q_{it}, F_{q|x}(\cdot|\cdot), \psi^0, \theta_2^0) \equiv 0$$

in order to summarize the above analysis as

$$\begin{aligned}
& \frac{1}{\sqrt{NT}} \sum_{i=1}^N \sum_{t=1}^T \bar{h}(x_{it}, q_{it}, \hat{F}_{q|x}^{-1}(\cdot|\cdot), \hat{\psi}, \theta_2^0) I_{it} \\
&= \frac{1}{\sqrt{NT}} \sum_{l=1}^N \sum_{t=1}^T E_{z_{it}} [g(x_{it}, q_{it}, G_h(q_{lt}, x_{lt}, \cdot, \cdot), \psi^0, \theta_2^0)] + B\sqrt{NT} (\hat{\psi} - \psi^0) + o_p(1).
\end{aligned}$$

If we plug the linear representation of $\sqrt{NT}(\hat{\psi} - \psi^0)$ (equation (36) above) into the linear representation for $\sqrt{NT}(\hat{\theta}_2 - \theta_2^0)$ (equation (40) above), we conclude that

$$\sqrt{n}(\hat{\theta}_2 - \theta_2^0) = \bar{A}^{-1} \frac{1}{\sqrt{NT}} \sum_{l=1}^N \sum_{t=1}^T v^*(z_{lt''}, \psi^0, \theta_2^0),$$

where

$$v^*(z_{lt''}, \psi^0, \theta_2^0) = \bar{v}(z_{lt''}, \psi^0, \theta_2^0) + A^{-1} \tilde{v}(z_{lt''}, \psi^0).$$

In the above A and $\tilde{v}(z_{lt''}, \psi^0)$ are defined in Eqs. (31) and (35), and

$$\bar{v}(z_{lt''}, \psi^0, \theta_2^0) = \lim_{h \rightarrow 0} E_{z_{it}} [g(x_{it}, q_{it}, G_h(q_{lt}, x_{lt}, \cdot, \cdot), \psi^0, \theta_2^0)].$$

Therefore we conclude that

$$\sqrt{NT}(\hat{\theta}_2 - \theta_2^0) \xrightarrow{d} N(0, \bar{A}^{-1} \bar{\Omega} \bar{A}^{-1})$$

where

$$\bar{\Omega} = \frac{1}{T} E \left\{ \sum_{t=1}^T v^*(z_{lt''}, \psi^0, \theta_2^0) \sum_{t=1}^T v^*(z_{lt''}, \psi^0, \theta_2^0)' \right\}. \quad (41)$$

Moreover $\bar{\Omega}$ is the matrix specified in Assumption 7. ■

B 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 each period are determined by a market-clearing condition: given policies $q(x_{it}, s_{it}, p_t)$, $\forall i$,

$$p_t : \int \int q(x, s, p_t) \mathcal{J}_t(dx) \mathcal{H}_t(ds) = 0, \quad \forall t \quad (42)$$

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 immediate that

$$\mathcal{H}_t(s) = F_s(s), \quad \forall t. \quad (43)$$

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

$$\mathcal{J}_t(x) = \int \int \mathbf{1}(z + q(z, s, p_{t-1}) \leq x) \mathcal{H}_{t-1}(ds) \mathcal{J}_{t-1}(dz). \quad (44)$$

Given any initial stock distribution \mathcal{J}_0 , the sequences $\{\mathcal{J}_t\}$ and $\{\mathcal{H}_t\}$ are both deterministic, and evolve according to (43) and (44). Therefore, by the market clearing conditions (42), 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.

Table 1: Parameter estimates: log-normal specification for F_s
 $\log s \sim N(\mu, \sigma^2)$

	Estimate	Standard error ^a
K	0.0003	0.6750
r	0.0320	0.0101
μ	-0.6706	0.0772
σ	2.2830	0.1268

^a Obtained via bootstrap resamples.

Table 2: Parameter estimates: flexible piecewise-linear specification for F_s

	Estimate	Standard errors ^a
Step 1 parameters		
$F_s^{-1}(0.15)$	0.0028	0.0066
$F_s^{-1}(0.25)$	0.6994	0.2664
$F_s^{-1}(0.50)$	1.2400	0.7377
$F_s^{-1}(0.75)$	1.3344	0.4777
$F_s^{-1}(0.85)$	1.6058	0.4253
Step 2 parameters^b		
K	0.1197	0.0351 ^c
r	0.0072	0.0024

Fifth-order polynomial approximation employed for terminal value (cf. end of section 4).

^aStandard deviation of parameter estimates obtained from 49 bootstrap resamples.

^bNumber of simulation draws used to evaluate expected value function: $L = 10$.

^cThese standard errors account for estimation error in the first-step estimates.

Table 3: Summary statistics for each quota exchange

(A) Year	(B) Month	(C) MCP ^a	(D) #non- ^b participants	(E) #participants:	(F) <i>of which</i> #sellers	(G) (%success) ^e	(H) #buyers	(I) (%success) ^f	(J) #zero bids ^c	(K) #non-zero bids ^d
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

^aCanadian dollars per kilogram of butterfat per day.

^b Computed as $2574 - (E)$, where 2574 is the number of producers who participated in at least one of the monthly quota exchanges between September 1997 and July 1998.

^c ie. number of bidders who submitted zero bids, computed as $(D) + (G) * (F) + (I) * (H)$.

^d ie. number of bidders who submitted non-zero bids, computed as $2574 - (J)$.

^e % of sellers who sold in the exchange, i.e., who submitted bids at or below the MCP.

^f % 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 $\leq q$






Figure 2: Estimated Policy Functions
 Estimated policy function for September 1997.

Lognormal specification:

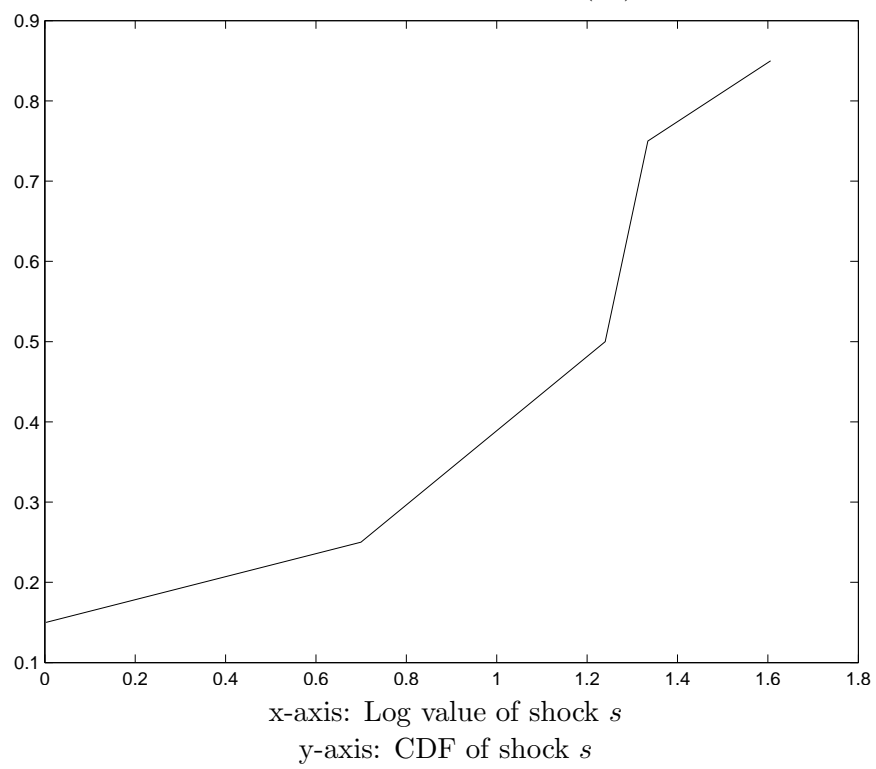


Piecewise-linear specification:



x-axis: Log value of shock s
 y-axis: quota transaction amount q_t

Figure 3: Estimated CDF of shock s
Estimated using equation (10).



Five quantiles were estimated: 0.15, 0.25, 0.5, 0.75, 0.85.