

CCP Estimation of Dynamic Discrete Choice Models with Unobserved Heterogeneity*

Peter Arcidiacono Robert A. Miller
Duke University Carnegie Mellon University

February 20, 2008

Abstract

We adapt the Expectation-Maximization (EM) algorithm to incorporate unobserved heterogeneity into conditional choice probability (CCP) estimators of dynamic discrete choice problems. The unobserved heterogeneity can be time-invariant, fully transitory, or follow a Markov chain. By exploiting finite dependence, we extend the class of dynamic optimization problems where CCP estimators provide a computationally cheap alternative to full solution methods. We also develop CCP estimators for mixed discrete/continuous problems with unobserved heterogeneity. Further, when the unobservables affect both dynamic discrete choices and some other outcome, we show that the probability distribution of the unobserved heterogeneity can be estimated in a first stage, while simultaneously accounting for dynamic selection. The probabilities of being in each of the unobserved states from the first stage are then taken as given and used as weights in the second stage estimation of the dynamic discrete choice parameters. Monte Carlo results for the three experimental designs we develop confirm that our algorithms perform quite well, both in terms of computational time and in the precision of the parameter estimates.

Keywords: dynamic discrete choice, unobserved heterogeneity

*We thank Esteban Aucejo, Lanier Benkard, Jason Blevins, Paul Ellickson, George-Levi Gayle as well as seminar participants at Duke University, Stanford University, University College London, UC Berkeley, University of Pennsylvania, University of Texas, IZA, and the NASM of the Econometric Society for valuable comments. Josh Kinsler and Andrew Beauchamp provided excellent research assistance. Financial support was provided for by NSF grants SES-0721059 and SES-0721098.

1 Introduction

Standard methods for solving dynamic discrete choice models involve calculating the value function either through backwards recursion (finite-time) or through the use of a fixed point algorithm (infinite-time).¹ Conditional choice probability (CCP) estimators, originally proposed by Hotz and Miller (1993), provide an alternative to these computationally-intensive procedures by exploiting the mappings from the value functions to the probabilities of making particular decisions. CCP estimators are much easier to compute than Maximum Likelihood (ML) estimators based on obtaining the full solution and have experienced a resurgence in the literature on dynamic games.² The computational gains associated with CCP estimation give researchers considerable latitude to explore different functional forms for their models.

Nevertheless, there are at least two reasons why researchers have been reticent to employ CCP estimators in practice.³ First, many believe that CCP estimators cannot be easily adapted to handle unobserved heterogeneity.⁴ Second, the mapping between conditional choice probabilities and value functions is simple only in specialized cases, and seems to rely heavily on the Type I extreme value distribution to be operational.⁵

This paper extends the application of CCP estimators to handle rich classes of probability distributions for unobservables. We develop estimators for dynamic structural models where there is time dependent unobserved heterogeneity and relax restrictive functional form assumptions about its within period probability distribution. In our framework, the unobserved state variables follow

¹The full solution or nested fixed point approach for discrete dynamic models was developed in Miller (1984), Pakes (1986), Rust (1987) and Wolpin(1984), and further refined by Keane and Wolpin (1994, 1997).

²Aguirregabiria and Mira (2008) have recently surveyed the literature on estimating dynamic models of discrete choice. For applications of CCP estimators to dynamic games in particular, see Aguirregabiria and Mira (2007), Bajari, Benkard, and Levin (2007), Pakes, Ostrovsky, and Berry (2004), and Pesendorfer and Schmidt-Dengler (2003).

³A third reason is that to perform policy experiments it is often necessary to solve the full model. While this is true, using CCP estimators would only involve solving the full model once for each policy simulation as opposed to multiple times in a maximization algorithm.

⁴Several studies based on CCP estimation have included fixed effects estimated from another part of the econometric framework. For example see Altug and Miller (1998), Gayle and Miller (2006) and Gayle and Golan (2007). As discussed in the text below, our approach is more closely related Aguirregaberia and Mira (2007), who similarly use finite mixture distributions in estimation.

⁵Bajari, Benkard and Levin (2007) provide an alternative method for relaxing restrictive functional form assumptions on the distributions of the unobserved disturbances to current utility. Building off the approach of Hotz et al. (1994), they estimate reduced form policy functions in order to forward simulate the future component of the dynamic discrete choice problem.

from a finite mixture distribution. The framework can readily be adapted to cases where the unobserved state variables are time-invariant, such as is standard in the dynamic discrete choice literature,⁶ as well as to cases where the unobserved states transition over time and, in the limit, are time independent. In this way we provide a unified approach to rectifying the two limitations commonly attributed to CCP estimators.

Our estimators adapt the EM algorithm, and in particular its application to sequential likelihood estimation developed in Arcidiacono and Jones (2003), to CCP estimation techniques. We construct several related algorithms for obtaining these estimators, derive their asymptotic properties, and investigate the small sample properties via three Monte Carlo studies. We show how to implement the estimator on a wide variety of dynamic optimization problems and games of incomplete information with discrete and continuous choices. To accomplish this, we generalize the concept of finite dependence developed in Altug and Miller (1998) to models where finite dependence is defined in terms of probability distributions rather than exact matches.

Our baseline algorithm iterates on three steps. First, given an initial guess on the parameter values and on the conditional choice probabilities (CCP's) where the conditioning is also on the unobserved state, we calculate the conditional probability of being in each of the unobserved states. We next follow the maximization step of the EM algorithm where the likelihood is calculated as though the unobserved state is observed and the conditional probabilities of being in each of the unobserved states are used as weights in the maximization. Finally, the CCP's for each state (both observed and unobserved) are updated using the new parameter estimates, recognizing the correlated structure of the unobservables when appropriate. The updated CCP's can come from the likelihoods themselves, or can be formed from an empirical likelihood as a weighted average of discrete choice decisions observed in the data, where the weights are the conditional probabilities of being in each of the unobserved states.

Our algorithm can be modified to situations where the data not only include records of discrete choices, but also outcomes on continuous choices, such as costs, sales, profits, and so forth that

⁶Aguirregabiria and Mira (2007) and Buchinsky, Hahn and Hotz (2005) both incorporate a time-invariant effect drawn from a finite mixture within their CCP estimation framework. Aguirregabiria and Mira (2007), in an algorithm later extended by Kasahara and Shimotsu (2007b), show how to incorporate unobserved characteristics of markets in dynamic games, where the unobserved heterogeneity is a time-invariant effect in the utility or payoff function. Our analysis also demonstrates how to incorporate unobserved heterogeneity into both the utility functions and the transition functions, and thereby account for the role of unobserved heterogeneity in dynamic selection. Buchinsky et al (2005) use the tools of cluster analysis, seeking conditions on the model structure that allow them to identify the unobserved type of each agent as the number of time periods per observation grows.

are also affected by the unobserved state variables. With observations on such outcomes, and the empirical distribution of the dynamic discrete choice decisions, we show how to estimate the distribution of unobserved heterogeneity in a first stage. The estimated probabilities of being in particular unobserved states obtained from the first stage are then used as weights when estimating the second stage parameters, namely those parameters entering the dynamic discrete choice problem that are not part of the first stage outcome equation. We show how the first stage of this modified algorithm can be paired with estimators proposed by Hotz et al (1994) and Bajari et al (2007) in the second stage. Our analysis complements their work by extending their applicability to unobserved time dependent heterogeneity.

We illustrate the small sample properties of our estimator using a set of Monte Carlo experiments designed to highlight the wide variety of problems that can be estimated with the algorithm. The first is a finite horizon model of teen drug use and schooling decisions where individuals learn about their preferences for drugs through experimentation. Here we illustrate both ways of updating the CCP's, using either the likelihoods themselves or the conditional probabilities of being in each of the unobserved states as weights. The second is a dynamic entry/exit example with unobserved heterogeneity in the demand levels for particular markets which in turn affects the values of entry and exit. The unobserved states are allowed to transition over time and the example explicitly incorporates dynamic selection. We estimate the model both by updating the CCP's with the model and by estimating the unobserved heterogeneity in a first stage. Our final Monte Carlo illustrates the performance of our methods in mixed discrete/continuous settings in the presence of unobserved heterogeneity. In particular, we focus on firms making discrete decisions about whether to run their plants and then, conditional on running, continuous decisions as to how much to produce. For all three sets of Monte Carlos, the estimators perform quite well both in terms of the precision of the estimates as well as the speed at which the estimates are obtained.

The techniques developed here are being used to estimate structural models in environmental economics, labor economics, and industrial organization. Bishop (2007) applies the reformulation of the value functions to the migration model of Kennan and Walker (2006) to accommodate state spaces that are computationally intractable using standard techniques. Joensen (2007) incorporates unobserved heterogeneity into a CCP estimator of educational attainment and work decisions. Finally, Finger (2007) estimates a dynamic game using our two-stage estimator to obtain estimates of the unobserved heterogeneity parameters in a first stage.

The rest of the paper proceeds as follows. Section 2 sets up the basic framework for our analysis.

Section 3 shows that, for many cases, the differences in conditional valuation functions only depend upon a small number of conditional choice probabilities. Section 4 extends the basic framework as well as applying the results of section 3 to the case when continuous choices are also present. Section 5 shows how to incorporate unobserved heterogeneity—including unobserved heterogeneity that transitions over time—into the classes of problems discussed in the preceding sections. Section 5 also shows how the parameters governing the unobserved heterogeneity can often be estimated in a first stage. Section 6 presents the asymptotics. Section 7 reports a series of Monte Carlos conducted to illustrate both the small sample properties of the algorithms as well as the broad classes of models that can be estimated using these techniques. Section 8 concludes. All proofs are in the appendix.

2 A Framework for Analyzing Discrete Choice

Consider a dynamic programming problem in which an individual makes a series of discrete choices d_t over his lifetime $t \in \{1, \dots, T\}$ for some $T \leq \infty$. The choice set has the same cardinality K at each date t , so we define d_t by the multiple indicator function $d_t = (d_{1t}, \dots, d_{Kt})$ where $d_{kt} \in \{0, 1\}$ for each $k \in \{1, \dots, K\}$ and

$$\sum_{k=1}^K d_{kt} = 1$$

A vector of characteristics (z_t, ε_t) fully describes the individual at each time t , where $\varepsilon_t \equiv (\varepsilon_{1t}, \dots, \varepsilon_{Kt})$ is independently and identically distributed over time with continuous support and distribution function $G(\varepsilon_{1t}, \dots, \varepsilon_{Kt})$, and the vector z_t evolves as a Markov process, depending stochastically on the choices of the individual. The probability of z_{t+1} conditional on being in z_t and making choice k at time t is given by $f_k(z_{t+1} | z_t)$ with the cumulative distribution function given by $F_k(z_{t+1} | z_t)$. At the beginning of each period t the individual observes $(z_t, \varepsilon_{1t}, \dots, \varepsilon_{Kt})$. The individual then makes a discrete choice d_t to sequentially maximize the expected discounted sum of utilities

$$E \left\{ \sum_{t=1}^T \sum_{k=1}^K \beta^{t-1} d_{kt} [u_k(z_t) + \varepsilon_{kt}] \right\}$$

where $u_k(z_t) + \varepsilon_{kt}$ denotes the current utility of an individual with characteristics z_t from choosing $d_{kt} = 1$. The discount factor is denoted by $\beta \in (0, 1)$, and the state z_t is updated at the end of each period.

Let $d_t^o \equiv (d_{1t}^o, \dots, d_{Kt}^o)$ denote the optimal decision rule given the current values of the state variables. Let $V(z_t)$ be the expected value of lifetime utility at date t as a function of the current state z_t but integrating over ε_t :

$$V(z_t) = E \left\{ \sum_{\tau=t}^T \sum_{k=1}^K \beta^{\tau-t} d_{k\tau}^o [u_k(z_\tau) + \varepsilon_{k\tau}] | z_t \right\}$$

The conditional valuation functions are given by current period utility for a particular choice net of ε_t plus the expected value of future utility. The expectation is taken with respect to next period's state variables conditional on the current state variables z_t and the choice $j \in \{1, \dots, K\}$:

$$v_j(z_t) = u_j(z_t) + \beta \sum_{z_{t+1}} V(z_{t+1}) f_j(z_{t+1} | z_t)$$

The inversion theorem of Hotz and Miller (1993) for dynamic discrete choice models implies there is a mapping from the conditional choice probabilities, defined by

$$p_j(z_t) = \int d_{jt}(z_t, \varepsilon_t) dG(\varepsilon_{1t}, \dots, \varepsilon_{Kt})$$

to differences in the conditional valuation functions which we now denote as

$$\psi_j^k[p(z_t)] = v_k(z_t) - v_j(z_t)$$

The inversion theorem can then be used to formulate the expected contribution of ε_t conditional on the choice. The expected contribution of the ε_{kt} disturbance to current utility, conditional on the state z_t , is found by integrating over the region in which the j^{th} action is taken, so appealing to the representation theorem

$$\begin{aligned} \int [d_{jt}(z_t, \varepsilon_t) \varepsilon_{jt}] dG(\varepsilon_t) &= \int 1 \{ \varepsilon_{jt} - \varepsilon_{kt} \geq v_k(z_t) - v_j(z_t) \text{ for all } k \in \{1, \dots, K\} \} \varepsilon_{jt} dG(\varepsilon_t) \\ &= \int 1 \{ \varepsilon_{jt} - \varepsilon_{kt} \geq \psi_j^k[p(z_t)] \text{ for all } k \in \{1, \dots, K\} \} \varepsilon_{jt} dG(\varepsilon_t) \\ &\equiv w_j[\psi[p(z_t)]] \end{aligned}$$

where $\psi[p(z_t)] \equiv (\psi_1^1[p(z_t)], \dots, \psi_1^K[p(z_t)])$. It now follows that the conditional valuation functions can be expressed as the sum of future discounted utility flows for each of the choices, weighted by the probabilities of each of these choices being optimal given the information set and then integrated over the state transitions. These discounted utility flows for each of the choices include the expected contribution of ε_t conditional on each of the choices being optimal. Hence, we can express $v_j(z_t)$ as:

$$v_j(z_t) = u_j(z_t) + E \left\{ \sum_{\tau=t+1}^T \sum_{k=1}^K \beta^{\tau-t} p_k(z_\tau) (u_k(z_\tau) + w_k[\psi[p(z_\tau)]) \mid d_{jt} = 1, z_t) \right\}$$

Two issues then remain for estimating dynamic discrete choice models using conditional choice probabilities. First, the mappings between the conditional probabilities and the expected ε_t contributions need to be explicit and we discuss a class of such models in the next subsection. Second, for a broad class of models the representation theorem itself can be used to avoid calculating conditional choice probabilities, flow utility terms, and transitions on the states across the T periods.

Indeed, as we show in section 3, it is often the case that only one-period-ahead transitions and choice probabilities are needed to fully capture the future utility terms.

2.1 Example 1: Generalized Extreme Value Distributions

We now illustrate how to map conditional choice probabilities into the expected contribution of ε_t as expressed through each $w_k[\psi[p(z_t)]]$. Suppose ε_t is drawn from the distribution function

$$G(\varepsilon_{1t}, \varepsilon_{2t}, \dots, \varepsilon_{Kt}) \equiv \exp[-\mathcal{H}(e^{-\varepsilon_{1t}}, e^{-\varepsilon_{2t}}, \dots, e^{-\varepsilon_{Kt}})]$$

where $\mathcal{H}(Y_1, Y_2, \dots, Y_K)$ satisfies the properties outlined for the generalized extreme value distribution in McFadden (1978).⁷ We first establish that essentially no computational cost is incurred from computing $w_k(\psi[p(z_t)])$ when the assumption of generalized extreme values holds and the mapping $\psi[p(z_t)]$ is known. In particular, Lemma 1 shows there is a log linear mapping relating the expected value of the disturbance to the specification of $\mathcal{H}(Y_1, Y_2, \dots, Y_K)$.

Lemma 1 *If ε_t is distributed generalized extreme value, then*

$$w_k(\psi[p(z_t)]) = \gamma + \log \mathcal{H}(e^{\psi_k^1[p(z_t)]}, e^{\psi_k^2[p(z_t)]}, \dots, e^{\psi_k^K[p(z_t)]})$$

The lemma demonstrates that the difficulty in mapping conditional choice probabilities into the expected contribution of ε_t comes from obtaining the inverse $\psi[p(z_t)]$, and not from mapping ψ into $w_k(\psi)$.⁸ The former mapping does, however, have a closed form in the nested logit case. Suppose there are R clusters and K_r alternatives within each cluster. Each period the person makes a choice by setting $d_{krt} = 1$ for some $r \in \{1, \dots, R\}$ and $k \in \{1, \dots, K_r\}$. We denote by p_{krt} the probability of making choice k in cluster at time t when the state is z_t , and define p_{rt} as the choice probability associated with the r^{th} cluster. That is

$$p_{rt} = \sum_{k=1}^{K_r} p_{krt}$$

⁷The properties are that $\mathcal{H}(Y_1, Y_2, \dots, Y_K)$ is a nonnegative real valued function of $(Y_1, Y_2, \dots, Y_K) \in \mathbb{R}_+^J$, homogeneous of degree one, with $\lim \mathcal{H}(Y_1, Y_2, \dots, Y_K) \rightarrow \infty$ as $Y_k \rightarrow \infty$ for all $k \in \{1, \dots, K\}$, and for any distinct (i_1, i_2, \dots, i_r) , the cross derivative $\partial \mathcal{H}(Y_1, Y_2, \dots, Y_K) / \partial Y_{i_1}, Y_{i_2}, \dots, Y_{i_r}$ is nonnegative for r odd and nonpositive for r even.

⁸The expression given in Lemma 1 can also be used to derive welfare effects outside of the conditional choice probability case. The differences in the v 's can be substituted back in for ψ giving the expected ε as a function of the parameters of the model. Hence, rather than attempting to draw errors from complicated GEV distributions in order to simulate welfare changes, the expected errors conditional on the choice can be calculated directly. As shown in Cardell (1997), even simulating draws from a nested logit distribution is difficult.

The distribution function of the disturbances, $G(\varepsilon_t) \equiv G(\varepsilon_{11t}, \varepsilon_{12t}, \dots, \varepsilon_{21t}, \dots, \varepsilon_{RK_Rt})$, is defined through $\mathcal{H}(Y) \equiv \mathcal{H}(Y_{11}, Y_{12}, \dots, Y_{21}, \dots, Y_{RK_R})$ by

$$\mathcal{H}(Y) = \sum_{r=1}^R \left[\sum_{k=1}^{K_r} Y_{kr}^{\delta_r} \right]^{1/\delta_r}$$

Bearing in mind that $\psi[p(z_t)]$ and $(w_1(\psi), \dots, w_K(\psi))$ typically enter linearly in CCP estimation, Lemma 2 below demonstrates that applying a CCP estimator to discrete choice dynamic models with a nested logit structure does not pose substantial computational challenges over and above the multinomial logit structure. Yet relaxing the multinomial logit assumption adds significantly to the flexibility of the estimator by introducing parameters that define the distribution of unobserved heterogeneity, in essentially the same way as in the static literature on random utility models.

Lemma 2 *The differences in the conditional valuation functions in the nested logit framework can be expressed as*

$$v_{krt} - v_{jst} = \frac{1}{\delta_r} \log(p_{krt}) - \frac{1}{\delta_s} \log(p_{jst}) + \left(1 - \frac{1}{\delta_r}\right) \log(p_{rt}) - \left(1 - \frac{1}{\delta_s}\right) \log(p_{st})$$

and the expected value of the disturbance conditional on an optimal choice can be written

$$E[\varepsilon_{jst} | d_{jst} = 1] = \gamma - \frac{1}{\delta_s} \log(p_{jst}) - \left(1 - \frac{1}{\delta_s}\right) \log(p_{st}) + \log \left\{ \sum_{r=1}^R p_{rt}^{1-1/\delta_r} \left[\sum_{j=1}^{K_r} p_{jrt}^{\delta_s/\delta_r} \right]^{1/\delta_s} \right\}$$

It is straightforward to generalize this framework to hierarchical clusters beyond two levels, and also to models where δ_r depends on the state z . Conversely, when all clusters are symmetric to the extent that $\delta = \delta_r = \delta_s$, the differences in conditional valuation functions simplify to

$$v_{krt} - v_{jst} = \frac{1}{\delta} [\log(p_{krt}) - \log(p_{jst})] + \left(1 - \frac{1}{\delta}\right) [\log(p_{rt}) - \log(p_{st})]$$

while the expected value of the disturbance conditional on making the k^{th} choice in cluster s becomes

$$E[\varepsilon_{jst} | d_{jst} = 1] = \gamma - \frac{1}{\delta} \log(p_{jst}) - \left(1 - \frac{1}{\delta}\right) \log(p_{st})$$

Specializing further, the multinomial logit is obtained by setting $\delta = 1$.

3 Finite Dependence

While Section 2 explored the mapping between CCP's and expected error contributions, in this section we exploit the Hotz-Miller inversion theorem directly to avoid calculating T period ahead conditional choice probabilities, flow utility terms, and transitions on the state variables. We show

that when a problem exhibits finite time dependence, a term we define below, the number of future conditional choice probabilities needed may shrink dramatically. This result relies upon two features of dynamic discrete choice problems. First, estimation relies upon differences in conditional valuation functions not the conditional valuation functions themselves. Second, the future utility terms can always be expressed as the conditional valuation function for one of the choices plus a term that only depends upon the differences in the conditional valuation functions. This latter term can then be expressed as a function of the CCP's. Hence, a sequence of normalizations on the future utility terms with respect to particular choices may lead to a cancellation of future utility terms after a particular point in time once we difference across the two alternatives. The rest of this section defines the class of models covered by finite dependence as well as showing how many future conditional choice probabilities are needed in estimation. We show that finite dependence covers a broad class of models in labor economics and industrial organization including but not limited to models with a terminal state or renewal.⁹

We begin by generalizing the concept of finite dependence developed in Altug and Miller (1998) to accommodate models where the outcome of choices on the state variables is endogenously random, as follows:

Definition 1 *Denote by $\lambda(j, z_t) \equiv \{\lambda_t(j, z_t), \dots, \lambda_{t+\rho}(j, z_t)\}$ a stochastic process of choices defined for at least ρ periods, starting at period t where the state at period t is z_t , the initial choice in the sequence is j , and the choice at period $\tau \in \{t, \dots, t + \rho\}$ is conditional on the current state z_τ (stochastically determined by realizations of the choice process). Also let $\kappa_\tau(z|j, z_t)$ denote the probability of state $z \in Z$ occurring at date τ , given the process $\lambda(j, z_t)$ and conditional only on z_t and $d_{jt} = 1$. A pair of choices, $j \in \{1, 2, \dots, J\}$ and $j' \in \{1, 2, \dots, J\}$, exhibits ρ period dependence for a state z_t , if there exists a process $\lambda(j, z_t)$ with the property that $\kappa_{t+\rho}(z|j, z_t) = \kappa_{t+\rho}(z|j', z_t)$ for all z_t and $t \in \{1, 2, \dots, T\}$.*

The basis for finite dependence comes from expanding the conditional valuation function $v_j(z_t)$ associated with choice j at time t one period into the future. For ease of notation, denote $\lambda_\tau(j) = \lambda_\tau(j, z_t)$. For the choice $\lambda_{t+1}(j)$ the Hotz-Miller inversion theorem implies $v_j(z_t)$ can be expressed

⁹Following Hotz and Miller (1993), a state is called terminal, and a choice which directly leads to it are called terminating, if there are no further decisions to be made in the dynamic program or game. In a renewal model, the initial state that can be reached from every other state via some decision sequence.

as:

$$v_j(z_t) = u_j(z_t) + \beta \sum_{z_{t+1}} \left\{ v_{\lambda_{t+1}(j)}(z_{t+1}) + \sum_{k=1}^K p_k(z_{t+1}) \left(\psi_{\lambda_{t+1}(j)}^k(p(z_{t+1})) + w_k[\psi[p(z_{t+1})]] \right) \right\} f_j(z_{t+1}|z_t) \quad (1)$$

Forming an equivalent expression for $v_{j'}(z_t)$, suppose the expected value of $v_{\lambda(j)}(z_{t+1})$ under the distribution $f_j(z_{t+1}|z_t)$ equals the expected value of $v_{\lambda_{t+1}(j')}(z_{t+1})$ under the distribution $f_{j'}(z_{t+1}|z_t)$

$$\sum_{z_{t+1}} v_{\lambda_{t+1}(j)}(z_{t+1}) f_j(z_{t+1}|z_t) = \sum_{z_{t+1}} v_{\lambda_{t+1}(j')}(z_{t+1}) f_{j'}(z_{t+1}|z_t)$$

The difference $v_j(z_t) - v_{j'}(z_t)$ could then be expressed in terms of this period's utilities and terms depending on next periods conditional choice probabilities $p(z_{t+1})$, plus the transition probabilities alone. Intuitively, aside from the two t period disturbances ε_{jt} and $\varepsilon_{j't}$, taking action j versus j' in period t would not matter if they are followed by actions $\lambda(j)$ and $\lambda(j')$ respectively, and also compensated for nonoptimal behavior by terms that are functions solely of the one-period-ahead conditional choice probabilities. Proposition (1), which follows directly from an induction argument, provides sufficient conditions for finite dependence to hold.

Proposition 1 *Differences in conditional valuation functions can be expressed in terms of future conditional choice probabilities up to ρ periods ahead if ρ -period finite dependence holds across all dates $t \in \{1, 2, \dots, T\}$, states $z_t \in Z$ and initial choices d_t . In that case there exists a choice process $\lambda(j, z_t)$ defined for all $j \in \{1, 2, \dots, K\}$, $\tau \in \{1, 2, \dots, T\}$ and $z_t \in Z$ such that:*

$$\begin{aligned} v_j(z_t) - v_{j'}(z_t) &= u_j(z_t) - u_{j'}(z_t) \\ &+ \sum_{\tau=t+1}^{t+\rho} \sum_{k=1}^K \sum_{z_{t+1}} \beta^{\tau-t} p_k(z_\tau) \left\{ \psi_{\lambda_\tau(j)}^k[p(z_\tau)] + u_k(z_\tau) + w_k[\psi[p(z_\tau)]] \right\} \kappa_\tau(z_\tau|j, z_t) \\ &- \sum_{\tau=t+1}^{t+\rho} \sum_{k=1}^K \sum_{z_{t+1}} \beta^{\tau-t} p_k(z_\tau) \left\{ \psi_{\lambda_\tau(j')}^k[p(z_\tau)] + u_k(z_\tau) + w_k[\psi[p(z_\tau)]] \right\} \kappa_\tau(z_\tau|j', z_t) \end{aligned}$$

We illustrate the finite dependence property with some examples that highlight the broad class of models that satisfy the finite dependence assumption, starting with renewal problems where only one-period-ahead CCP's are necessary to calculate the expected future utility differences.¹⁰

¹⁰The finite dependence property is also illustrated in the migration model of Bishop (2007), in which individuals choose among over fifty locations where to live. With state variables transitioning across locations, the finite dependence assumption allows Bishop to effectively reduce the dynamic discrete problem to a three period decision.

3.1 Example 2: Renewal

In renewal problems, such as Miller's (1984) job matching model or Rust's (1987) machine maintenance problem, the agent has an option to nullify all previous history by taking a renewal action, namely starting a new job in the job matching model, or replacing the bus engine in the maintenance problem. Formally, the first choice, say, is a renewal action if and only if $f_1(z_{t+1}|z_t) = f_1(z_{t+1})$ for all $z \in Z$. Renewal problems satisfy the finite dependence assumption, because for any two choices j and j' made in period t , the state at the beginning of period $t + 2$ will be identical if the renewal action is taken in period $t + 1$. Denoting the renewal action by the first choice

$$v_1(z_t) \equiv u_1(z_t) + \beta \sum_{z_{t+1}} V(z_{t+1}) f_1(z_{t+1}) \equiv u_1(z_t) + \beta V^*$$

Models with terminal states also have this property.

Suppose the disturbance associated with the renewal action (such as engine replacement), is independent of the disturbances associated with the other choices (such as different types of repair and servicing combined with different types of usage), which might be correlated with each other in any way the generalized extreme value distribution permits. When $G(\varepsilon_t) \equiv \exp[-\mathcal{H}(e^{-\varepsilon_1 t}, e^{-\varepsilon_2 t}, \dots, e^{-\varepsilon_K t})]$ is generalized extreme value, this is equivalent to saying

$$\mathcal{H}(Y_1, \dots, Y_K) \equiv \bar{\mathcal{H}}(Y_2, \dots, Y_K) + Y_1$$

where $\bar{G}(\varepsilon_t) \equiv \exp[-\bar{\mathcal{H}}(e^{-\varepsilon_2 t}, \dots, e^{-\varepsilon_K t})]$ is any generalized extreme value distribution of dimension $K - 1$. In this case, Lemma 3 establishes that the likelihood of any decision depends only on current flow utilities, the one-period-ahead probabilities of transitioning to each of the states, and the one-period-ahead probabilities of the renewal action.¹¹

Lemma 3 *If $\mathcal{H}(Y_1, \dots, Y_K) \equiv \bar{\mathcal{H}}(Y_2, \dots, Y_K) + Y_1$ in the generalized extreme value model and the first choice is a renewal action then*

$$v_j(z_t) = u_j(z_t) + \beta \left(\sum_{z_{t+1}} [u_1(z_{t+1}) - \log p_1(z_{t+1})] f_j(z_{t+1}|z_t) dz_{t+1} + \gamma + \beta V^* \right) \quad (2)$$

¹¹When z_t contains observed variables only, estimation proceeds as in the static problem. Note that in estimation we work with differences in conditional valuation functions. Since the last term in (2) is the same across all choices, the last term cancels out. The second to last term can be calculated outside the model by estimating the transitions on the state variables, for example by using a cell estimator to obtain an estimate of the probability of the renewal action. The first-stage estimate of the second term is then just subtracted off the flow utility in estimation. Note that this method applies whether the model is stationary or not, whether or not it has a finite or infinite horizon, and accommodates a rich pattern of correlations between nonrenewal choices.

Since the likelihood of any choice only depends upon differences in the conditional valuation functions, the constant $(\gamma + \beta V^*)$ cancels out.

3.2 Example 3: Dynamic Entry and Exit

Several empirical studies investigate the dynamics of entry and exit decisions.¹² To further illustrate finite dependence and demonstrate its applicability to this topic, we develop a prototype model of an infinite horizon dynamic entry/exit game, estimated in our second Monte Carlo study of \mathcal{N} distinct markets. Suppose a typical market is served by at most two firms, with up to one firm entering each market every period. Potential entrants choose whether to enter the market or not, and incumbents choose whether to exit or not. Choices by the incumbent and a potential entrant are made simultaneously. If an incumbent exits, it disappears forever, and firms only have one opportunity to enter.

The systematic component of the realized profit flow of a firm in period t , denoted by $u(E_t, M_t, z_t)$, depends on whether the firm is an entrant, $E_t = 1$, or an incumbent, $E_t = 0$, whether the firm operates as a monopolist, $M_t = 1$, or a duopolist, $M_t = 0$, and the state of demand, $z_t \in \{0, 1\}$. The state of demand transitions over time according to the Markov process $f(z_{t+1}|z_t)$. Finally, an independent and identically distributed Type I extreme value shock affects both the profits associated with participating or not participating in the market. These profit shocks are unobserved to rival firms and the firm's future profit shocks are independent over time and unknown to the firm.

The state variables determining the firm's expected value from entering or remaining in the industry depend upon whether the firm is an entrant $E_t = 1$ or an incumbent $E_t = 0$; whether there is an incumbent rival, which we denote by $R_t = 0$, or not (by setting $R_t = 1$); and the state of demand z_t . Let $p_0(E_t, R_t, z_t)$ denote the probability of not entering or exiting, and similarly let $p_1(E_t, R_t, z_t)$ denote the probability of remaining in or entering the market. In a symmetric equilibrium $p_0(E_t, 0, z_t)$ is the probability that a potentially entering rival stays out when facing competition from the firm as an incumbent, and $p_0(0, R_t, z_t)$ is the probability that an incumbent rival exits. We can then express the expected value from entering as the sum of the disturbance ε_{1t} plus:

$$v_1(E_t, R_t, z_t) \equiv E_t R_t \left\{ u(1, 1, z_t) + \beta \sum_{z_{t+1}=0}^1 V(0, 1, z_{t+1}) f(z_{t+1}|z_t) \right\} \quad (3)$$

$$+ (1 - E_t R_t) \sum_{k=0}^1 p_k(E_t, R_t, z_t) \left\{ u(E_t, 1 - k, z_t) + \beta \sum_{z_{t+1}=0}^1 V(0, 1 - k, z_{t+1}) f(z_{t+1}|z_t) \right\}$$

¹²See, for example, Beresteanu and Ellickson (2006), Collard-Wexler (2006), Dunne et al. (2006), and Ryan (2006).

where $V(0, R_{t+1}, z_{t+1})$ is the expected value of an incumbent firm at the beginning of period $t + 1$ conditional on R_{t+1} and z_{t+1} . The first expression on the right side of (3) reflects the fact that when $E_t R_t = 1$, the firm enjoys monopoly rents of $u(1, 1, z_t)$ for at least one period if it enters. Otherwise the rent is shared by the duopoly with probability $p_1(E_t, R_t, z_t)$, as indicated in the second expression. Since this framework has a terminating state, the previous example establishes that the conditional valuation function for entering/remaining can be expressed as:

$$v_1(E_t, R_t, z_t) = E_t R_t \left\{ u(1, 1, z_t) - \beta \sum_{z_{t+1}=0}^1 \log[p_0(0, 1, z_{t+1})] f(z_{t+1}|z_t) \right\} + \beta \gamma \quad (4)$$

$$+ (1 - E_t R_t) \sum_{k=0}^1 p_k(E_t, R_t, z_t) \left\{ u(E_t, 1 - k, z_t) - \beta \sum_{z_{t+1}=0}^1 \log[p_0(0, 1 - k, z_{t+1})] f(z_{t+1}|z_t) \right\}$$

where the value of exiting has been normalized to zero. Similar to the renewal case, everything except for flow profit terms can be calculated outside of the model where the calculations only involve one-period-ahead transition probabilities on the states as well as current and one-period-ahead probabilities of rival and own actions.

3.3 Example 4: Female Labor Supply

We now consider a case when more than one-period-ahead conditional choice probabilities are needed in estimation. In particular, we consider female labor supply where experience on the job increases human capital in an uncertain way, thus extending previous work on human capital accumulation on the job by Altug and Miller (1998), Gayle and Miller (2006) and Gayle and Golan (2007), where it is measured as an observed deterministic variable. Each period a woman chooses whether to work by setting $d_t = 1$, versus stay at home by setting $d_t = 0$. Earnings at work depend upon her human capital, denoted by h_t , and participation in the previous period d_{t-1} . Human capital h_t increases stochastically by $z \in \{1, 2, \dots, Z\}$, where $f(z)$ is the probability of drawing z . At the beginning of period t the woman receives utility of $u_j(h_t, d_{t-1})$ from setting $d_t = j \in \{0, 1\}$ plus a choice specific disturbance term denoted by ε_{jt} that is distributed Type 1 extreme value. Her goal is to maximize expected lifetime utility, the expected discounted sum of current utilities, by sequentially choosing whether to work or not each period until T . To show there is two period dependence in this model, we note that if the woman participates in period t and then does not participate in periods $t + 1$ and $t + 2$, her state variables in period $t + 3$ have the same probability distribution as if she does not participate in period t but participates in period $t + 1$ instead and then finally does not participate at $t + 2$. Applying Proposition 1, we obtain the difference in the conditional valuation functions directly:

Lemma 4 *The difference in conditional valuation functions between working and not working are given by:*

$$\begin{aligned}
& [v_1(h_t, d_{t-1}) - u_1(h_t, d_{t-1})] - [v_0(h_t, d_{t-1}) - u_0(h_t, d_{t-1})] \\
&= \sum_{z=1}^Z \{ \beta [u_0(h_t + z, 1) - \log p_0(h_t + z, 1)] + \beta^2 [u_0(h_t + z, 0) - \log p_0(h_t + z, 0)] \} f(z) \\
&\quad - \beta [u_1(h_t, 0) - \log p_1(h_t, 0)] - \sum_{z=1}^Z \{ \beta^2 [u_0(h_t + z, 1) - \log p_0(h_t + z, 1)] \} f(z)
\end{aligned} \tag{5}$$

Here the future utility terms are expressed as a function of the one-period-ahead flow utilities, the two-period ahead transitions on the state variables, and the two-period-ahead conditional choice probabilities.

4 Continuous Choices

Our framework is readily extended to incorporate continuous choices as follows. We now suppose that in addition to the discrete choices $d_t = (d_{1t}, \dots, d_{Kt})$, an individual also makes a sequence of continuous choices c_t over his lifetime $t \in \{1, \dots, T\}$. At each time t , the individual is now described by a vector of characteristics (z_t, ε_t) , where $\varepsilon_t \equiv (\varepsilon_{0t}, \dots, \varepsilon_{Kt})$ is independently and identically distributed over time with continuous support and distribution function $G_0(\varepsilon_{0t})G(\varepsilon_{1t}, \dots, \varepsilon_{Kt})$, and z_t is defined as before. Conditional on discrete choice $k \in \{1, \dots, K\}$ and continuous choice c , the transition probability from z_t to z_{t+1} is denoted by $f_{ck}(z_{t+1} | c_t, z_t)$. At the beginning of each period t the individual observes $(z_t, \varepsilon_{1t}, \dots, \varepsilon_{Kt})$, and makes a discrete choice d_t . The individual then observes ε_{0t} and chooses c_t . Both the discrete and choices are chosen to sequentially maximize the expected discounted sum of utilities

$$E \left\{ \sum_{t=1}^T \sum_{k=1}^K \beta^{t-1} d_{kt} [U_k(c_t, z_t, \varepsilon_{0t}) + \varepsilon_{kt}] \right\}$$

where $U_k(c, z_t, \varepsilon_{0t}) + \varepsilon_{kt}$ denotes the current utility an individual with characteristics (z_t, ε_t) receives from choosing (c, k) . We write $c_{kt}^o \equiv c_k(z_t, \varepsilon_{0t})$ for the optimal continuous choice the person would make conditional on discrete choice $k \in \{1, \dots, K\}$ after observing ε_{0t} .¹³

¹³The two most closely related papers to ours that incorporate both continuous and discrete choices are Altug and Miller (1998) and Bajari et al (2007). There are important differences between the three approaches, but one similarity is that we follow Bajari et al (2007) by including an independently distributed disturbance term, or private shock, and exploiting a monotonicity assumption relating that shock ε_{0t} to the continuous choice. They explicitly treat the case where there is a single continuous choice variable, but also note the difficulties in extending their approach to models where there is more than one continuous choice. In Altug and Miller (1998) choices may be discrete or continuous,

Substituting c_{kt}^o into current utility $U_k(c_{kt}^o, z_t, \varepsilon_{0t})$ and transition $f_{ck}(z_{t+1} | c_{kt}^o, z_t)$, then integrating over ε_0 yields the expected payoff of setting $d_{kt} = 1$ given z_t net of ε_{kt}

$$u_k(z_t) = \int U_k[c_k(z_t, \varepsilon_{0t}), z_t, \varepsilon_{0t}] dG(\varepsilon_{0t})$$

along with the state transition

$$f_k(z_{t+1} | z_t) \equiv \int f_{ck}(z_{t+1} | c_k(z_t, \varepsilon_{0t}), z_t) dG_0(\varepsilon_{0t})$$

for each $k \in \{1, \dots, K\}$. In this section we reinterpret $u_k(z_t)$ and $f_k(z_{t+1} | z_t)$ as reduced forms for $U_k(c_{kt}^o, z_t, \varepsilon_{0t})$ and $f_{ck}(z_{t+1} | c_{kt}^o, z_t)$ respectively, derived endogenously from the primitives and the optimal continuous choice rule. Data on (z_t, c_t, d_t) provide information linking the reduced form to the structural primitives. By exploiting these connections and adapting the methods we develop for estimating the reduced form $u_k(z_t)$ and $f_k(z_{t+1} | z_t)$, we can extend our estimation techniques to a mixture of discrete and continuous variables and thus estimate the primitives $U_k(c_t, z_t, \varepsilon_{0t})$, $f_{ck}(z_{t+1} | c_t, z_t)$ and $G_0(\varepsilon_{0t})$.

4.1 Two representations of the reduced form

More specifically, we exploit two representations derived below. They rely on the identity that, given the state and discrete choice $d_{kt} = 1$, the probability distribution for ε_0 induces a distribution on to $c(z_t, k, \varepsilon_{0t})$ defined by

$$\Pr\{c_t \leq c | k, z_t\} = \int 1\{c(z_t, k, \varepsilon_{0t}) \leq c\} dG_0(\varepsilon_{0t}) \equiv H_k(c_t | z_t)$$

Both representations assume monotonicity conditions relating the optimal continuous choice c_t^o to the value of the unobservable ε_{0t} .

The first representation holds when $c(z_t, k, \varepsilon_{0t})$ is strictly monotone (increasing) in ε_{0t} . Under this assumption the cumulative distribution functions $G_0(\varepsilon)$ and $H_k(c | z)$ are related through the optimal decision rule $c_k(z_t, \varepsilon_{0t})$ by the equations

$$G_0(\varepsilon) = \Pr[\varepsilon_0 \leq \varepsilon] = \Pr[c_k(z_t, \varepsilon_{0t}) \leq c_k(z_t, \varepsilon)] = H_k(c_k(z_t, \varepsilon) | z_t)$$

for all state and choice coordinate pairs (z, k) . It now follows that

$$\varepsilon_{0t} = G_0^{-1}[H_k(c_{kt}^o | z_t)]$$

and all decisions in period t , whether discrete or continuous, are made simultaneously. However they do not include a variable corresponding to ε_{0t} , so the policy function for the continuous choice c is a mapping from the discrete choice k and the state z alone. This facilitates their use of Euler equations to form orthogonality conditions in estimation, the continuous choice variable is a mapping of (z, k) .

Hence the reduced form utility and reduced form transition can be expressed as

$$u_k(z_t) = \int U_k [c_{kt}^o, z_t, G_0^{-1} [H_k(c_{kt}^o | z_t)]] dH_k(c_{kt}^o | z_t)$$

and

$$f_k(z_{t+1} | z_t) = \int f_{ck}(z_{t+1} | c_{kt}^o, z_t) dH_k(c_{kt}^o | z_t)$$

respectively. Given a parametric form for $G_0(\varepsilon)$, the induced dynamic discrete choice model can be estimated using the approach described in the other sections in this paper.

The second representation of $u_k(z_t)$ holds when c_{kt}^o satisfies a first order condition of the form

$$U_{1k}(c_{kt}^o, z_t, \varepsilon_{0t}) + \sum_{z_{t+1}} \beta V(z_{t+1}) \frac{\partial f_{ck}(z_{t+1} | c_{kt}^o, z_t)}{\partial c} = 0$$

and the marginal utility of consumption $U_{1k}(c_{kt}^o, z_t, \varepsilon_{0t}) \equiv \partial U_k(c_{kt}^o, z_t, \varepsilon_{0t}) / \partial c$ is strictly monotone in ε_0 for all (k, c, z) . The latter assumption implies $U_{1k}(c_{kt}^o, z_t, \varepsilon_{0t})$ has a partial inverse in ε_{0t} , denoted $\lambda(u, k, c, z)$, meaning that for all (ε, k, c, z)

$$\varepsilon_{0t} = \lambda_k [U_{1k}(c_{kt}^o, z_t, \varepsilon_{0t}), c_{kt}^o, z_t]$$

In that case the monotonicity assumption implies

$$\varepsilon_{0t} = -\lambda_k \left(\sum_{z_{t+1}} \beta V(z_{t+1}) \frac{\partial f_{ck}(z_{t+1} | c_{kt}^o, z_t)}{\partial c}, c_{kt}^o, z_t \right)$$

and hence $u_k(z_t)$ can be expressed as

$$u_k(z_t) \equiv \int U_k \left[c_{kt}^o, z_t, -\lambda_k \left(\sum_{z_{t+1}} \beta V(z_{t+1}) \frac{\partial f_{ck}(z_{t+1} | c_{kt}^o, z_t)}{\partial c}, c_{kt}^o, z_t \right) \right] dH_k(c_{kt}^o | z_t)$$

Given finite dependence of length ρ , we may express $V(z_{t+1})$ using its finite dependent representation, and thus ignore all the utility terms following period $t + \rho + 1$ in $V(z_{t+1})$. They are independent of z_{t+1} and therefore have no effect on the integrand since

$$\sum_{z_{t+1}} \frac{\partial f_{ck}(z_{t+1} | c_{kt}^o, z_t)}{\partial c} = 0$$

Given a parametric form for $U_k(c, z, \varepsilon_0)$ we can determine $\lambda_k(u, c, z)$ up to a parameterization and estimate the parameters from the induced discrete choice model together with orthogonality conditions constructed from the first order condition.

The monotonicity condition used in the first representation applies to the policy function for the continuous variable, so whether it is satisfied or not is partly determined by the definition of the probability transition which depends on the continuous choice. The monotonicity condition in the

second representation relies on regularity conditions that support an optimal interior solution, to be exploited in estimation, but does not impose any additional restrictions on the way continuous choices affect the transition probability. Another advantage of using the second representation is that it is not necessary to specify $G_0(\varepsilon)$ parametrically in order to estimate the other primitives of the model.

4.2 Example 5: Plant Production

At the beginning of each period t the owner manager of a manufacturing plant chooses between operating his plant by setting $d_{2t} = 1$, or temporarily idling it by setting $d_{1t} = 1$. For each discrete choice $k \in \{1, 2\}$ we model the costs of setting $d_{kt} = 1$ as $\alpha_k + \varepsilon_{kt}$, where α_k is the systematic component and ε_{kt} is a random variable, identically and independently distributed Type 1 extreme value. Three factors determine the net revenue generated from operating the plant and setting $d_{2t} = 1$: the condition of the plant $z_{2t} \in \{1, \dots, Z_2\}$, where higher levels of z_2 indicate that the plant is in worse condition; the variable input the manager assigns to determine the scale of the production function, which is a continuous choice variable denoted by $c_t \in (0, \infty)$; and two demand shocks. One of the shocks, denoted by ε_{0t} , is distributed $N(0, \sigma^2)$ and is independent across time. The other, denoted by z_{1t} , evolves stochastically but does not depend upon the choice. We interpret z_{1t} as a long run trend in demand (for example high or low) and ε_{0t} as indicating changes in demand elasticity and the attractiveness of different market segments. Given the condition of the plant z_{2t} , and the state of demand $(\varepsilon_{0t}, z_{1t})$, net revenue from operating the plant in period t and choosing c_t is a quadratic in the logarithm of c_t . The coefficient on the linear term is $(\varepsilon_{0t} + \alpha_3 z_{1t})$, the coefficient on the quadratic term is $\alpha_4 z_{2t}$, and $\alpha_3 > 0 > \alpha_4$. Increasing inputs c_t raises the probability that the machinery is in bad condition B next period $t + 1$, according to the formula $\gamma_0 / (\gamma_0 + c_t^{\gamma_1})$, where $\gamma_0, \gamma_1 > 0$.

In terms of our previous notation, $z_t \equiv (z_{1t}, z_{2t})$ and the systematic component to the utility from idling the plant is

$$U_1(c_t, z_t, \varepsilon_{0t}) = u_1(z_t) = \alpha_1$$

When the plant runs, utility is given by:

$$U_2(c_t, z_t, \varepsilon_{0t}) = (\varepsilon_{0t} + \alpha_3 z_{1t}) \ln c_t + \alpha_4 z_{2t} (\ln c_t)^2 + \alpha_2$$

The first reduced form of current utility from operating this plant in this example is therefore

$$u_2(z_t) = \int \left\{ \left(\Phi^{-1} \left[\frac{H_2(c_t | z_t)}{\sigma} \right] + \alpha_3 z_{1t} \right) \ln c_t + \alpha_4 z_{2t} (\ln c_t)^2 \right\} dH_2(c_t) + \alpha_2$$

where $H_{c_2}(c_t | z_t)$ is the distribution for c_t when the plant runs, and $\Phi(\cdot)$ is the standard normal distribution function.

To derive the second representation, it is straightforward to check that an interior solution is optimal and the conditional value functions are bounded. Consequently the optimal input choice for operating the plant must satisfy the first order and second order conditions for an optimum, and in this case the former can be expressed as

$$\varepsilon_{0t} + \alpha_3 z_{1t} + 2\alpha_4 z_{2t} (\ln c_t) = \left(\sum_{z_{t+1}} [V(z_{1t+1}, z_{2t}) - V(z_{1t+1}, z_{2t} + 1)] f(z_{1t+1} | z_{1t}) \right) \gamma_0 \gamma_1 c_t^{\gamma_1} (\gamma_0 + c_t^{\gamma_1})^{-2} \quad (6)$$

Given the Type I extreme value distributions for the costs of idling or running the plant, we know that $V(\cdot)$ can be expressed as $v_1(\cdot) - \ln(p_1(\cdot)) + \gamma$ where γ is Euler's constant. But, because the choice to idle is a renewal action for z_2 , $v_1(z_{1t+1}, z_{2t}) = v_1(z_{1t+1}, z_{2t} + 1)$. Hence, we can write equation (6) as:

$$\varepsilon_{0t} + \alpha_3 z_{1t} + 2\alpha_4 z_{2t} (\ln c_t) = \left(\sum_{z_{t+1}} [\ln(p_1(z_{1t+1}, z_{2t} + 1)) - \ln(p_1(z_{1t+1}, z_{2t}))] f(z_{1t+1} | z_{1t}) \right) \gamma_0 \gamma_1 c_t^{\gamma_1} (\gamma_0 + c_t^{\gamma_1})^{-2} \quad (7)$$

Substituting for ε_{0t} in $U_2(c_t, z_t, \varepsilon_{0t})$ and integrating over c_t implies that the alternative representation of current utility conditional on operating the plant is

$$u_2(z_t) = \alpha_2 + \int \left\{ \left(\gamma_0 \gamma_1 c_t^{\gamma_1} (\gamma_0 + c_t^{\gamma_1})^{-2} \ln c_t \sum_{z_{t+1}} [\ln(p_1(z_{1t+1}, z_{2t} + 1)) - \ln(p_1(z_{1t+1}, z_{2t}))] f(z_{1t+1} | z_{1t}) \right) - \alpha_4 z_{2t} (\ln c_t)^2 \right\} dH_2(c_t | z_t) \quad (8)$$

Totally differentiating the first order condition with respect to ε_{0t} and c_t , and appealing to the second order condition, it immediately follows that the second monotonicity condition is satisfied in this example, so the consumption policy function is strictly monotone increasing in ε_{0t} , thus establishing that both representations apply to one of the discrete choices. Finally we note that although the monotonicity conditions only apply to one discrete choice, this is sufficient for estimation purposes in this example, as we later demonstrate in our Monte Carlo application.

5 The Algorithm

This section develops algorithms for estimating dynamic optimization problems and games of incomplete information where there is unobserved heterogeneity that evolves over time as a stochastic process. We consider a panel data set of \mathcal{N} individuals. We observe \mathcal{T} choices for each individual $n \in \{1, \dots, \mathcal{N}\}$, along with a sub-vector of their state variables. Observations are independent across individuals. We partition the state variables z_{nt} into those observed by the econometrician, $x_{nt} \in \{x_1, \dots, x_X\}$, and those that are not observed, $s_{nt} \in \{1, \dots, S\}$. The n^{th} individual's unobserved state at time t , s_{nt} , may affect both the utility function and the transition functions on the observed variables and may also evolve over time. The initial probability of being assigned to unobserved state s is π_s . Unobserved states follow a Markov process with π_{jk} dictating the probability of transitioning from state j to state k . When unobserved heterogeneity is permanent, $\pi_{jk} = 0$ for $j \neq k$, and we write $\pi_{jj} = \pi_j$. When the unobserved states are completely transitory and there is no serial dependence, the elements of any given column in the transition matrix have the same value, and we write $\pi_{jk} = \pi_k$. We denote by π the $(S + 1) \times S$ matrix of initial and transitional probabilities for the unobserved states. The structural parameters that define the utility outcomes for the problem are denoted by $\theta \in \Theta$ and the set of CCP's, denoted by p , are treated as nuisance parameters in the estimation.

5.1 Data on discrete choices

Let $\mathcal{L}(d_{nt} | x_{nt}, s; \theta, \pi, p)$ be the likelihood of observing individual n make choice d_{nt} at time t , conditional on being in state (x_{nt}, s) , given structural parameter θ and CCP's p . Forming their product over the \mathcal{T} periods we obtain the likelihood of any given path of choices and $(d_{n1}, \dots, d_{n\mathcal{T}})$, conditional on the $(x_{n1}, \dots, x_{n\mathcal{T}})$ sequence and the unobserved state variables $(s(1), \dots, s(\mathcal{T}))$. Integrating the product over the initial unobserved state with probabilities π_j and the subsequent transitions π_{jk} then yields the likelihood of observing the choices d_n conditional on x_n given (θ, π, p) :

$$\begin{aligned} L(d_n | x_n, \theta, \pi, p) &\equiv \sum_{s(1)}^S \sum_{s(2)}^S \dots \sum_{s(\mathcal{T})}^S \pi_{s(1)} \mathcal{L}(d_{n1} | x_{n1}, s(1); \theta, \pi, p) \\ &\quad \times \prod_{t=2}^{\mathcal{T}} \pi_{s(t-1), s(t)} \mathcal{L}(d_{nt} | x_{nt}, s(t); \theta, \pi, p) \end{aligned}$$

Therefore the log likelihood for the sample is:

$$\sum_{n=1}^{\mathcal{N}} \log L(d_n | x_n, \theta, \pi, p) \tag{9}$$

When unobserved heterogeneity is permanent, the log likelihood for the sample reduces to:

$$\sum_{n=1}^{\mathcal{N}} \log \left(\sum_{s=1}^S \prod_{t=1}^T \pi_s \mathcal{L}_{nst} \right)$$

When the mixing distribution has no state dependence, the log likelihood for the sample reduces to:

$$\sum_{n=1}^{\mathcal{N}} \log \left(\sum_{s=1}^S \prod_{t=1}^T \pi_s \mathcal{L}_{nst} \right) = \sum_{n=1}^{\mathcal{N}} \sum_{t=1}^T \log \left(\sum_{s=1}^S \pi_s \mathcal{L}_{nst} \right)$$

Directly maximizing the log likelihood for such problems can be computationally infeasible. An alternative to maximizing (9) directly is to iteratively maximize the expected log likelihood function as follows.¹⁴ Given estimates of $\pi^{(m)}$, the initial probabilities of being in each of the unobserved states and later transitions, and $p^{(m-1)}$, estimates of the CCP's obtained from the previous iteration, the m^{th} iteration maximizes:

$$\sum_{n=1}^{\mathcal{N}} \sum_{s=1}^S \sum_{t=1}^T q_{nst}^{(m)} \log \mathcal{L} \left(d_{nt} \mid x_{nt}, s; \theta, \pi^{(m)}, p^{(m-1)} \right) \quad (10)$$

with respect to θ to obtain $\theta^{(m)}$. Here, $q_{nst}^{(m)} = q_{st}(d_n, x_n, \theta^{(m-1)}, \pi^{(m-1)}, p^{(m-1)})$, and is formally defined below as the probability that individual n is in state s at time t given parameter values (θ, π, p) , and conditional on the all the data about n . The information from the data is then $(d_n, x_n) \equiv (d_{n1}, x_{n1}, \dots, d_{nT}, x_{nT})$.

To define $q_{st}(d_n, x_n, \theta, \pi, p)$, let $L_{st}(d_n \mid x_n, \theta, \pi, p)$ denote the joint probability of state s occurring at date t for the n^{th} individual and observing the choice sequence d_n , conditional on the exogenous variables x_n , when the parameters take value (θ, π, p) . Abbreviating $\mathcal{L}(d_{nt} \mid x_{nt}, s; \theta, \pi, p)$ by \mathcal{L}_{nst} , we define $L_{st}(d_n \mid x_n, \theta, \pi, p)$ by:

$$\begin{aligned} & L_{st}(d_n \mid x_n, \theta, \pi, p) \\ &= \sum_{s(1)}^S \dots \sum_{s(t-1)}^S \sum_{s(t+1)}^S \dots \sum_{s(T)}^S \left(\prod_{r=2, r \neq t, r \neq t+1}^T \pi_{s(r-1), s(r)} \mathcal{L}_{n, s(r), r} \right) (\pi_{s(1)} \mathcal{L}_{n, s(1), 1} \pi_{s(t-1), s} \mathcal{L}_{nst} \pi_{s, s(t+1)} \mathcal{L}_{n, s(t+1), t+1}) \end{aligned}$$

where the summations of $s(1)$ and so on are over $s \in \{1, \dots, S\}$. When unobserved heterogeneity is permanent, $L_{st}(d_n \mid x_n, \theta, \pi, p)$ simplifies to

$$L_{st}(d_n \mid x_n, \theta, \pi, p) = \left(\prod_{r=2}^T \mathcal{L}_{nsr} \right) (\pi_s \mathcal{L}_{ns1})$$

for all t . Summing over all states $s \in S$ at any time t returns the likelihood of observing the choices d_n conditional on x_n given (θ, π, p) :

$$L(d_n \mid x_n, \theta, \pi, p) = \sum_{s=1}^S L_{st}(d_n \mid x_n, \theta, \pi, p)$$

¹⁴For applications of the EM algorithm in time series models with regime-switching, see Hamilton (1990).

Therefore the probability that individual n is in state s at time t given the parameter values (θ, π, p) conditional on all the data for n is:

$$q_{st}(d_n, x_n, \theta, \pi, p) \equiv \frac{L_{st}(d_n | x_n, \theta, \pi, p)}{L(d_n | x_n, \theta, \pi, p)} \quad (11)$$

Note that the denominator is the same across all time periods and all states. When the transitions are independent, the n^{th} individual's previous and future history is not informative about the current state, and in this case $q_{st}(d_n, x_n, \theta, \pi, p)$ reduces to

$$q_{st}(d_n, x_n, \theta, \pi, p) = \frac{\pi_s \mathcal{L}_{nst}}{\sum_{s'=1}^S \pi_{s'} \mathcal{L}_{ns't}}$$

To make the algorithm operational we must explain how to update π , the probabilities for the initial unobserved states and their transitions, θ , the other structural parameters, and p , the CCP's. The updating formula for the transitions is based on the identities:

$$\pi_{jk} \equiv \Pr\{k | j\} = \frac{\Pr\{k, j\}}{\Pr\{j\}} = \frac{E_n \{E[s_{nkt} | d_n, x_n, s_{njt-1}] E[s_{njt-1} | d_n, x_n]\}}{E_n \{E[s_{njt} | d_n, x_n]\}} \equiv \frac{E_n [q_{nkt|j} q_{njt}]}{E_n [q_{njt}]}$$

where the n subscript on an expectations operator indicates that the integration is over the whole sample population, s_{nkt} is an indicator for whether individual n is in state k at time t and $q_{nts|j} \equiv E[s_{ntk} | d_n, x_n, s_{n,t-1,j}]$ denotes the probability of individual n being type k at time t conditional on the data and also on being in unobserved state j at time $t-1$. This conditional probability is defined by the expression:

$$q_{nkt|j} = \frac{\pi_{jk} \mathcal{L}_{nkt} \left(\sum_{s(t+1)}^S \cdots \sum_{s(T)}^S \prod_{r=t+1}^T \pi_{s(r-1),s(r)} \mathcal{L}_{n,s(r),r} \right)}{\sum_{s'=1}^S \pi_{js(t)} \mathcal{L}_{ns(t)t} \left(\sum_{s(t+1)}^S \cdots \sum_{s(T)}^S \prod_{r=t+1}^T \pi_{s(r-1),s(r)} \mathcal{L}_{n,s(r),r} \right)}$$

Averaging $q_{nkt|j} q_{njt}$ over the sample to approximate the joint probability $E_n [q_{nkt|j} q_{njt}]$, and averaging q_{njt} over it to estimate $E_n [q_{njt}]$, we update π_{jk} using:

$$\pi_{jk}^{(m+1)} = \frac{\sum_{n=1}^N \sum_{t=2}^T q_{nkt|j}^{(m)} q_{njt}^{(m)}}{\sum_{n=1}^N \sum_{t=2}^T q_{njt}^{(m)}} \quad (12)$$

Setting $t = 1$ yields the conditional probability of the n^{th} individual being in unobserved state s in the first time period. We update the probabilities for the initial states by averaging the conditional probabilities obtained from the previous iteration over the sample population:

$$\pi_s^{(m+1)} = \frac{1}{N} \sum_{n=1}^N q_{ns1}^{(m)} \quad (13)$$

In a Markov stationary environment, the unconditional probabilities reproduce themselves each period. In that special case we can average over all the periods in the sample in the update formula

for π to obtain

$$\pi_s^{(m+1)} = \frac{1}{NT} \sum_{t=1}^T \sum_{n=1}^N q_{nst}^{(m)}$$

The other component to update is the vector of conditional choice probabilities. In contrast to models where unobserved heterogeneity is absent, initial consistent estimates of p cannot be cheaply computed prior to structural estimation, but must be iteratively updated along with (θ, π) . One way of updating the CCP's is to substitute in the likelihood evaluated at the previous iteration. Let $l_k(x_{nt}, s; \theta, \pi, p)$ denote the conditional likelihood of observing choice $k \in \{1, \dots, K\}$ for the state (x, s) when the parameters are (θ, π, p) , which implies

$$\mathcal{L}(d_{nt} | x_{nt}, s; \theta, \pi, p) = \sum_{k=1}^K d_{ntk} l_k(x_{nt}, s; \theta, \pi, p)$$

One updating rule for p is:

$$p_{kxs}^{(m+1)} = l_k(x, s; \theta^{(m+1)}, \pi^{(m+1)}, p^{(m)}) \quad (14)$$

Another way of updating p comes from exploiting the identities

$$\Pr\{d_{nkt} | x, s\} \Pr\{s | x\} = \Pr\{d_{nkt}, s | x\} \equiv E[d_{nkt}(s_{nt} = s) | x] = E[d_{nkt} E\{s_{nt} = s | d_n, x_n\} | x]$$

where the last equality follows from the law of iterated expectations and the fact that d_n includes d_{nkt} as a component. From its definition

$$q_{nst} = E[s_{nt} = s | d_n, x_n]$$

Again applying the law of iterated expectations we obtain

$$\Pr\{s | x\} = E\{E[s_{nt} = s | d_n, x_n] | x\}$$

Dividing the first identity through by $\Pr\{s | x\}$, and substituting $E[q_{nst} | x]$ for $E[s_{nt} = s | d_n, x_n]$ throughout it now follows that

$$p_{kxs} \equiv \Pr\{d_{nkt} | x, s\} = \frac{E[d_{nkt} q_{nst} | x]}{E[q_{nst} | x]}$$

In words, of the fraction of the total population with characteristic x in state s , the portion choosing the k^{th} action is p_{kxs} . This formulation suggests a second way of updating p , using the weighted empirical likelihood:

$$p_{kxs}^{(m+1)} = \frac{\sum_{t=1}^T \sum_{n=1}^N d_{nkt} q_{nst}^{(m+1)} I(x = x_{nt})}{\sum_{t=1}^T \sum_{n=1}^N q_{nst}^{(m+1)} I(x = x_{nt})} \quad (15)$$

where $I(x = x_{nt})$ is the indicator function for x .

Using (14) to update the CCP's rather than (15) imposes more restrictions from the underlying theory. To prove this claim, first note that the framework is not identified if the dimension of p , denoted $\dim(p)$, is strictly less than $\dim(\theta) + \dim(\pi)$. Typically parameters are used to estimate the process governing unobserved heterogeneity, ensuring $\dim(p) > \dim(\theta)$. (Indeed this strict inequality is met in all practical applications of CCP estimation.) Consequently the number of equations used to determine p from (14), obtained from the first order conditions by maximizing (10), is strictly less than the number used to determine p from (15). Hence the converged values of (14) satisfy overidentifying restrictions that result in greater precision than the converged values of (15), leading to lower standard errors in the structural parameters (θ, π) . However, there may be cases when updating with the data is computationally much simpler than updating from the model. Further, the modified algorithm we propose in the next subsection, for estimating models where not only choices but also other outcomes are observed that are related to the unobserved state variables, builds on the updating method given in (15).

We have now defined all the pieces necessary to implement the algorithm. It is triggered by setting values for the structural parameters, $\theta^{(1)}$, the initial distribution of the unobserved states plus their probability transitions, $\pi^{(1)}$, and the conditional choice probabilities $p^{(1)}$. Natural candidates for $(\theta^{(1)}, \pi^{(1)}, p^{(1)})$ come from estimating a model without any unobserved heterogeneity and perturbing the estimates obtained. Each iteration in the algorithm has four steps. Given $(\theta^{(m)}, \pi^{(m)}, p^{(m)})$ the $(m+1)^{th}$ proceeds as follows:

Step 1 Compute $q_{nst}^{(m+1)}$ and $q_{nst|j}^{(m+1)}$ for each (n, s, t, j) using (11) with parameters $(\theta^{(m)}, \pi^{(m)}, p^{(m)})$.

Step 2 Compute $\pi^{(m+1)}$ from (13) and (12) using $q_{nst}^{(m+1)}$ and $q_{nst|j}^{(m+1)}$.

Step 3 Obtain $\theta^{(m+1)}$ by maximizing (10) with respect to θ evaluated at $\pi^{(m+1)}$, $p^{(m)}$ and $q_{nst}^{(m+1)}$.

Step 4 Update $p^{(m+1)}$, using either (14) or (15).

Let (θ^*, π^*, p^*) denote the converged values of the structural parameters and CCP estimators from the EM algorithm. Following the arguments in Arcidiacono and Jones (2003), the EM solution satisfies the first order conditions derived from maximizing (9) with respect to θ given p^* .

5.2 Auxiliary data on continuous choices and outcomes

When there is auxiliary data that depend upon the unobserved heterogeneity to supplement the discrete choice data, the estimator we have just described can be modified and applied to a broader

class of models than those satisfying finite dependence. This situation arises when the conditional transition probability for the observed state variables depends on the current values of the unobserved state variables, when there is data on a payoff of a choice that depends on the unobserved heterogeneity, when data exists on some other outcome that is determined by the unobserved state variables, or when a first order condition fully characterizes a continuous choice that is affected by the unobserved heterogeneity.

The modified algorithm is implemented by updating the conditional choice probabilities using equation (15), an empirical estimator of the fraction of people in any given state making a particular choice. When information is available on both the individual choices and an outcome, this method for updating the conditional choice probabilities implies that we can substitute the empirical estimator into the likelihood for observing a sequence of outcomes without estimating all the structural parameters that affect the decision itself.

Denote by c_{nt} the outcome observed for individual n at time t . For example c_{nt} might be a continuous choice satisfying a first order condition. Conditional on x_{nt} , the observed exogenous variables and s , the unobserved state, we express the likelihood of choosing c_{nt} by $\mathcal{L}_{1nst} \equiv \mathcal{L}_1(c_{nt} | d_{nt}, x_{nt}, s; \theta_1)$ with parameter vector θ_1 . Appealing to the definition of conditional probability, the joint likelihood for (c_{nt}, d_{nt}, x_{nt}) , can be decomposed multiplicatively into the product $\mathcal{L}_{1nst}\mathcal{L}_{2nst}$, where $\mathcal{L}_{2nst} \equiv \mathcal{L}_2(d_{nt} | x_{nt}, s; \theta_2, \pi, p)$ is now the likelihood associated with the discrete choice, and is parametrized by θ_2 . We permit, but do not require, θ_1 and θ_2 to overlap.

The modified algorithm proceeds in two stages, first adapting the algorithm described above to estimate (θ_1, π, p) , and in a second stage estimating θ_2 (or $\theta_2 - \theta_2 \cap \theta_1$) with standard CCP estimation techniques developed for models where there is no time dependent heterogeneity. The first stage is an EM algorithm for iteratively estimating the structural parameters (θ_1, π, p) that characterize a behavioral model for explaining (c_{nt}, d_{nt}, x_{nt}) . The full structure of the model is imposed on the continuous choices. However, the discrete choices are exogenously generated by a multinomial distribution that depends on the partially observed state variables but is otherwise unrestricted, thus breaking the parametric links provided by the discrete choice optimization. At the m^{th} iteration, θ_1 and p are chosen to maximize the expected log likelihood

$$\sum_{n=1}^{\mathcal{N}} \sum_{s=1}^S \sum_{t=1}^{\mathcal{T}} q_{nst}^{(m)} \left[\sum_{k=1}^K d_{nkt} I(x = x_{nt}) \log(p_{kxs}) + \log \mathcal{L}_1(c_{nt} | d_{nkt}, x_{nt}, s; \theta_1) \right] \quad (16)$$

where as before, $q_{nst}^{(m)}$ is the probability that each individual n is of type s at each time period t conditional on the sample (c_n, d_n, x_n) , defined using (11) evaluated at parameters $(\theta_1^{(m-1)}, \pi^{(m-1)}, p^{(m-1)})$.

Differentiating (16) with respect to p_{kxs} yields the following set of equations from the first order conditions for each (j, k) pair and every s

$$\frac{\sum_{n=1}^{\mathcal{N}} q_{nst}^{(m)} d_{nkt} I(x = x_{nt})}{p_{kxs}^{(m+1)}} = \frac{\sum_{n=1}^{\mathcal{N}} q_{nst}^{(m)} d_{njt} I(x = x_{nt})}{p_{jxs}^{(m+1)}} \quad (17)$$

Multiplying both sides of (17) through by $p_{kxs}^{(m+1)}$, and then summing both sides over $j \in \{1, \dots, K\}$, we obtain (15). The resulting $p^{(m+1)}$, derived from a model where there are no restrictions on discrete choice behavior, is in the same spirit as the second way of updating the CCP's in the original algorithm.

Formally, the $(m + 1)^{th}$ iteration proceeds as follows:

Step 1 After substituting \mathcal{L}_{1nst} for \mathcal{L}_{nst} in (11), compute $q_{nst}^{(m+1)}$ and $q_{nst|j}^{(m+1)}$ for each (n, s, t, j) , given parameters $(\theta_1^{(m)}, \pi^{(m)}, p^{(m)})$.

Step 2 Compute $\pi^{(m+1)}$ from (13) and (12) using $q_{nst}^{(m+1)}$ and $q_{nst|j}^{(m+1)}$.

Step 3 Maximize (16) with respect to θ_1 and p evaluated at $q_{nst}^{(m+1)}$, to obtain $\theta_1^{(m+1)}$ and $p^{(m+1)}$, where the formula for $p^{(m+1)}$ comes from (15).

This estimation procedure is an EM algorithm for an optimally chosen continuous choice, or an exogenous transition outcome, when the parametric restrictions implied by sequentially optimizing over the discrete choices are not imposed in estimation. Appealing to standard properties of the EM algorithm, the algorithm is (globally) monotone increasing.

Having achieved convergence in the first stage, there are several methods for estimating θ_2 , the parameters determining the (remaining) preferences over choices by substituting our stage estimators for (π, p) , denoted $(\hat{\pi}, \hat{p})$, into the second stage econometric criterion function. If the model satisfies finite dependence, then the appropriate representation can be used to express the conditional valuation functions in conjunction with standard optimization methods. Alternatively, the simulation estimators of Hotz et al (1994) or Bajari et al (2007) can be applied directly, regardless of whether the model satisfies the limited dependence property or not. The second-stage estimation problem is the same as when all state variables are observed. That is, from the $\mathcal{N} \times \mathcal{T}$ data set, create a data set that is $\mathcal{N} \times \mathcal{T} \times \mathcal{S}$ where this second data set has, for each observation in each time period, each possible value of the unobserved state. The second-stage estimation then weights each (n, t, s) observation using the first stage estimated probability weights \hat{q}_{nst} .

5.3 Example 6: Simulation Estimation

For example, to implement the algorithm of Hotz et al (1994), we appeal directly to the representation theorem.¹⁵ Namely, for each unobserved state we can stack the $(K - 1)$ mappings from the conditional choice probabilities into the differences in conditional valuation functions for each individual n in each period t :

$$\begin{bmatrix} \psi_1^2 [p_{n1t}] - (v_{n21t} - v_{n11t}) \\ \vdots \\ \psi_1^K [p_{n1t}] - (v_{nK1t} - v_{n11t}) \\ \vdots \\ \psi_1^2 [p_{nSt}] - (v_{n2St} - v_{n1St}) \\ \vdots \\ \psi_1^K [p_{nSt}] - (v_{nKSt} - v_{n1St}) \end{bmatrix} = \begin{bmatrix} 0 \\ \vdots \\ 0 \\ \vdots \\ 0 \\ \vdots \\ 0 \end{bmatrix} \quad (18)$$

where the second to last subscript on both the conditional choice and the conditional valuation functions is the unobserved state. Future paths are simulated by drawing future choices and transition paths of the observed and unobserved state variables for each initial choice and each initial observed and unobserved state. With the future paths in hand, it is possible to form future utility paths given the sequence of choices and these future utility paths can be substituted in for the conditional valuation functions. Estimation can then proceed by minimizing, for example, the weighted sum of each of the squared values of the left hand side of (18) with respect to θ_2 .

An advantage of using this two stage procedure is that it enlarges the class of models which can be estimated. Although the first estimation method described is computationally feasible for many problems with finite time dependence, not all dynamic discrete choice models have that property. Rather than assuming the model exhibits finite time dependence, one could estimate a stationary Markov model lacking this property, by estimating the distribution of unobserved heterogeneity in the first stage. These estimates could then be combined with non-likelihood based estimation methods in the second stage. Because the second method estimates the distribution of unobserved heterogeneity without fully specifying the dynamic optimization problem, another advantage of the second method is that the likelihood function for the discrete choices is not fully parametrically specified. Consequently the structural parameters estimated in the first stage are robust to different specifications of the within period probability distribution for the unobservable variables and the additively separable parts of the utility that are not directly functions of the outcomes and continuous

¹⁵Finger (2007) applies our two-stage estimator to the Bajari, Benkard, and Levin (2007) algorithm.

choices. A third advantage is computational; sequential estimation is usually easier to implement than simultaneous estimation, and the first stage algorithm is monotone increasing. Against these three advantages is the loss in asymptotic efficiency.

6 Large Sample Properties

The defining equations for this CCP estimator come from three sources. First are orthogonality conditions for θ , the parameters defining utility and the probability transition matrix for the observed states, which are analogous to the score for a discrete choice random utility model with nuisance parameters used in defining the payoffs. Second are the orthogonality conditions for the initial distribution of the unobserved heterogeneity and its transition probability matrix π , again computed from the likelihood as in a random effects model. Third are the equations which define the nuisance parameters as estimators of the conditional choice probabilities p . This section, together with accompanying material in the appendix, lays out the equations defining our estimator and discusses its asymptotic properties.

Let (φ^*, p^*) solve our algorithm in the discrete choice model, where $\varphi \equiv (\theta, \pi)$ is the vector of structural parameters. For any fixed set of nuisance parameters p , the solution to the EM algorithm satisfies the first order conditions of the original problem (9). Consequently setting $p = p^*$ in the original problem implies the first order conditions for the original problem are satisfied. It now follows that the large sample properties of our estimator can be derived by analyzing the score for (9) augmented by a set of equations that solve the conditional choice probability nuisance parameter vector p , either the likelihoods or the weighted empirical likelihoods, as discussed in the previous section.

In Section 5 we defined the conditional likelihood of (φ, p) upon observing d_n given x_n , which we now denote as $L(d_n | x_n; \varphi, p) \equiv L(d_n | x_n; \theta, \pi, p)$. The paragraph above implies that (φ^*, p^*) solves

$$\frac{1}{N} \sum_{n=1}^N \frac{\partial \log [L(d_n | x_n; \varphi^*, p^*)]}{\partial \varphi} = 0$$

When the choice specific likelihood is used to update the nuisance parameters, the definition of the algorithm implies that upon convergence, $p_{jxs}^* = \mathcal{L}_j(x, s; \varphi^*, p^*)$ for each (j, x, s) . Stacking $\mathcal{L}_j(x, s; \varphi^*, p^*)$ for each choice j and each value (x, s) of state variables to form $\mathcal{L}(\varphi, p)$, a $J \times X \times S$ vector function of the parameters (φ, p) , our estimator satisfies the JXS additional parametric restrictions $\mathcal{L}(\varphi^*, p^*) = p^*$. When the weighted empirical likelihoods are used instead, this condition

is replaced by the *JSX* equalities

$$p_{jxs}^* \sum_{t=1}^T \sum_{n=1}^{\mathcal{N}} I(x = x_{nt}) q_{st}(d_n, x_n, \varphi^*, p^*) = \sum_{t=1}^T \sum_{n=1}^{\mathcal{N}} d_{njt} I(x = x_{nt}) q_{st}(d_n, x_n, \varphi^*, p^*)$$

Forming the *SX* dimensional vector $q_t(d_n, \varphi, p)$ from stacking the terms $I(x = x_{nt}) q_{st}(d_n, x_n, \varphi^*, p^*)$ for each state (x, s) and the *JSX* dimensional vector $q_{st}^{(n,t)}(d_n, \varphi, p)$ from $I(x = x_{nt}) q_{st}(d_n, x_n, \varphi, p)$, we rewrite this alternative set of restrictions in vector form as

$$\left[\frac{1}{\mathcal{NT}} \sum_{t=1}^T \sum_{n=1}^{\mathcal{N}} q_t(d_n, \varphi^*, p^*) \right] Cp^* = \frac{1}{\mathcal{NT}} \sum_{t=1}^T \sum_{n=1}^{\mathcal{N}} q_{st}^{(n,t)}(d_n, \varphi^*, p^*)$$

where C is the $SX \times JSX$ block diagonal matrix

$$C \equiv \begin{bmatrix} 1 & 1 & 1 & \dots & 0 & 0 & 0 \\ & \dots & \dots & \dots & & & \\ 0 & 0 & 0 & \dots & 1 & 1 & 1 \end{bmatrix}$$

The main result of this section is that if the model is identified under standard regularity conditions, then it can be estimated with a CCP estimator.¹⁶ For the next proposition implies that, unless the model is unidentified, the algorithms described in Section 5 do not asymptotically have multiple limit points. If the algorithm converges to different limits from different starting values for a given sample size, and this persists as the sample size grows, then a consistent estimator does not exist.

Proposition 2 *Suppose the data $\{d_n, x_n\}$ are generated by φ_0 , exhibiting conditional choice probabilities p_0 . If φ_1 satisfies the vector of moment conditions*

$$E \left[\frac{\partial \log [\mathcal{L}(d_n | x_n; \varphi_1, p_1)]}{\partial \varphi} \right] = 0$$

where the expectation is taken over (d_n, x_n) in the sample population and $\mathcal{L}(\varphi_1, p_1) = p_1$, then under standard regularity conditions φ_0 and φ_1 are observationally equivalent.

Turning to the large sample properties of the CCP estimator, if $\varphi_0 \in \Psi$ is identified, then φ^* is consistent, converges at rate $\sqrt{\mathcal{N}}$, and is asymptotically normal, as can be readily established by appealing to well known results in the literature. The asymptotic covariance matrix is laid out in the appendix.

¹⁶Kasahara and Shimotsu (2006) have recently proved that when the unobserved heterogeneity is a finite mixture over a set of time-invariant effects in the utility function (but does not affect state transitions), knowing the time-invariant effects does not help with identification provided the number of observations on each person is of reasonably large.

The extension to continuous choice and other outcomes is straightforward. There are two extra features to account for, the conditional distribution of the continuous choices, and the adjustment of the reduced form utility $u_j(z) \equiv u_j(z; \varphi)$ formed by replacing the expectations operator with its sample average. When there is a first order condition defining the optimality conditions for the continuous choices, we have

$$\varepsilon_0 = \lambda \left(\frac{\partial U_j(c, z, \varepsilon_0)}{\partial c}, j, c, s \right)$$

from which the likelihood for c can be formed directly conditional on the action and the state (since by assumption c is monotone in ε_0). Similarly the parameters entering $\pi_j(s' | c, s; \varphi)$ can be estimated directly from the state transitions after conditioning on the choices and current state. For expositional purposes we assume here both conditional likelihoods are appended to the likelihood defined for the discrete part of the problem to increase the efficiency of the estimator. However in some applications it might be easier to estimate either or both conditional likelihoods separately, in which case the asymptotic corrections would be made in an analogous way to the corrections for p^* .

The likelihood must also be modified because we form approximate sample averages of $U_j(z, c, \varepsilon_0; \varphi)$ using one of the two representations described in Section 4, rather than using its population expectation over ε_0 , namely $u_j(z; \varphi)$, in estimation. Here we analyze the first representation of $u_j(z; \varphi)$ and assume that $G_0(\varepsilon_0)$ and $\pi_j(z' | c, s)$ are parametrically specified by $G_0(\varepsilon_0; \varphi)$ and $\pi_j(z' | c, z; \varphi)$. (Analyzing the second representation proceeds in a similar way.) In this case we approximate the mapping $u_j(z; \varphi)$ with

$$u_j^{(\mathcal{N})}(z; \varphi) = \frac{1}{\mathcal{N}} \sum_{n=1}^{\mathcal{N}} U_j(c_{sj}^o, z, G_0^{-1}[\pi_j(c_{sj}^o | z; \varphi)]; \varphi)$$

To account for the effects of this substitution within the likelihood, we approximate $L(d_n | x_n; \varphi, p)$ with $L(d_n | x_n; \bar{u}, \varphi, p)$, where $\mathcal{L}(\varphi, p)$ with $\mathcal{L}(\bar{u}, \varphi, p)$, where approximating functions such as $u_j^{(\mathcal{N})}(z; \varphi)$, are substituted for $u_j(z; \varphi)$ in the likelihood. The estimator is defined as the two equation vectors

$$\mathcal{L} \left[u^{(\mathcal{N})}(z; \varphi^*), \varphi^*, p^* \right] = p^*$$

and

$$0 = \frac{1}{\mathcal{N}} \sum_{n=1}^{\mathcal{N}} \frac{\partial \log [L(d_n | x_n; u^{(\mathcal{N})}(z; \varphi^*), \varphi^*, p^*)]}{\partial \varphi}$$

The asymptotic covariance matrix, derived in the appendix, accounts for replacing $u_j(z; \varphi)$ with $u_j^{(\mathcal{N})}(z; \varphi)$ in estimation.

7 Small Sample Performance

To evaluate the finite sample performance of our estimators we conducted three Monte Carlo studies with the purpose of illustrating the versatility of the estimators. The Monte Carlos illustrate the performance of the algorithms along a number of dimensions. We compare full information maximum likelihood to CCP estimates with the different ways of updating the CCP's. We show how well the algorithms perform in a dynamic game with incomplete information. We include cases where the probability of the renewal action is small, and test the performance of the algorithm that estimates the parameters governing the unobserved heterogeneity in a first stage. Finally, we examine the performance of the algorithms when individuals make both continuous and discrete choices.

7.1 Monte Carlo 1: Experimenting with drugs

The first Monte Carlo focuses on a simple learning framework where individual preferences are shaped by experience in ways that the econometrician does not observe. In our model youths have repeated opportunities to experiment with drugs. Experimentation leads individuals to discover their preferences for drugs, though there is a withdrawal cost to stop this acquired habit. We compare our estimates from using both methods for updating the probability distribution for the unobservables with the ML estimator, which is relatively cheap to compute because of the simple structure of the model.

In each period t a teenager decides among three alternatives, which following our notational convention are defined by $d_{jt} \in \{0, 1\}$ for $j \in \{0, 1, 2\}$ and $t \in \{1, \dots, T\}$ where $d_{0t} + d_{1t} + d_{2t} = 1$. He or she can drop out of school ($d_{0t} = 1$), stay in school and do drugs ($d_{1t} = 1$), or stay in school and abstain from drugs ($d_{2t} = 1$). There are three types of teenagers, who we characterize by the two indicator variables $\mathcal{A}_t \in \{0, 1\}$ and $\mathcal{B}_t \in \{0, 1\}$. First, those who have never taken drugs, and therefore do not know their preference at time t , denoted by setting $\mathcal{A}_t = 1$. Next, those who have found through experimentation that they have a high preference for drugs, denoted by setting $(\mathcal{A}_t, \mathcal{B}_t) = (0, 1)$; and finally those who have found through experimentation that they have a low preference for drugs, that is $(\mathcal{A}_t, \mathcal{B}_t) = (0, 0)$. Trying drugs for one period fully reveals an individual's type. Amongst those who have not tried drugs, the probability of having a high preference is π . Breaking a drug habit is modeled with a one period withdrawal cost incurred when $(d_{1t-1}, d_{2t}) = (1, 1)$.

The state variables in this model are $(\mathcal{A}_t, \mathcal{B}_t, d_{1t-1})$. Setting as initial values $(\mathcal{A}_0, \mathcal{B}_0) = (1, 0)$,

our discussion implies the law of motion for $(\mathcal{A}_t, \mathcal{B}_t)$ is

$$\begin{bmatrix} \mathcal{A}_{t+1} \\ \mathcal{B}_{t+1} \end{bmatrix} = \begin{bmatrix} \mathcal{A}_t(1 - d_{1t}) \\ (1 - \mathcal{A}_t + \mathcal{A}_t d_{1t})\zeta \end{bmatrix}$$

where ζ is an independently distributed Bernoulli random variable with probability π . Hence, π is the population probability of being in the high state.

We denote the baseline utility of attending school by α_0 , the baseline utility of setting $d_{1t} = 1$ and using drugs by α_1 , the additional utility from having the high preference for drugs ($\mathcal{B}_t = 1$) and using them by α_2 , and we let α_3 denote a one period withdrawal cost incurred when $(d_{1t-1}, d_{2t}) = (1, 1)$. Dropping out of school by setting $d_{0t} = 1$ is a terminal state, with utility normalized to the choice-specific disturbance ε_{0t} . Note that if the individual uses drugs then no withdrawal cost is paid, implying d_{1t-1} is irrelevant. Similarly if the individual does not use drugs, the only relevant state variable for current utility is whether he or she used them last period, not the level of addiction. We assume that $(\varepsilon_{0t}, \varepsilon_{1t}, \varepsilon_{2t})$ are distributed generalized extreme value, with ε_{0t} independent of the nest $(\varepsilon_{1t}, \varepsilon_{2t})$, thus reflecting the idea that options within school are more related to each other than either of them is to dropping out. The nesting parameter is denoted by δ .¹⁷

Given this payoff structure, the flow utilities from the two schooling choices net of the choice-specific disturbance can be expressed as:

$$\begin{aligned} u_1(\mathcal{A}_t, \mathcal{B}_t, \zeta_t) &= \alpha_0 + \alpha_1 + \alpha_2 \zeta \\ u_2(d_{1t-1}) &= \alpha_0 + \alpha_3 d_{1t-1} \end{aligned}$$

From the individual's perspective, the expected flow utility from trying drugs for the first time at t is $\alpha_0 + \alpha_1 + \alpha_2 \pi + \varepsilon_{1t}$. Since dropping out leads to a terminal state, it follows from our discussion in Section 3 that the conditional valuation functions $v_j(\mathcal{A}_t, \mathcal{B}_t, d_{1t-1})$ for $j \in \{1, 2\}$ may be expressed as

$$\begin{aligned} v_1(\mathcal{A}_t, \mathcal{B}_t, d_{1t-1}) &= \alpha_0 + \alpha_1 + \alpha_2(\mathcal{B}_t + \mathcal{A}_t \pi) - (1 - \mathcal{A}_t) \beta \ln p_0(0, \mathcal{B}_t, 1) \\ &\quad - \mathcal{A}_t \beta [\pi \ln p_0(0, 1, 1) - (1 - \pi) \ln p_0(0, 0, 1)] + \beta \gamma \\ v_2(\mathcal{A}_t, \mathcal{B}_t, d_{1t-1}) &= \alpha_0 + \alpha_3 d_{1t-1} - \beta \ln [p_0(\mathcal{A}_t, \mathcal{B}_t, 0)] + \beta \gamma \end{aligned}$$

Note that the expressions above would be exactly the same if the error structure followed a multinomial logit rather than a nested logit. However, a model generated under a multinomial logit would

¹⁷These assumptions correspond to those made in our companion paper, Arcidiacono, Kinsler and Miller (2008), which applies a CCP/EM estimator to the NLSY data on youth to investigate drug abuse and its consequences within a generalization of the prototype model presented here.

yield different values for the true conditional choice probabilities than those of the nested logit.

For each simulation we create 5000 simulated individuals with at most 5 periods of data. Some individuals have less than five observations because no further decisions occur once the simulated individual leaves school. We assume that the data would show drug usage at school d_{1t} , so that \mathcal{A}_t can be simply constructed, but that \mathcal{B}_t would be unobserved, thus violating the conditional independence assumption. We estimated the model using three different methods, namely maximum likelihood, a CCP estimator that updates with the likelihood functions, and a CCP estimator updated by a weighted empirical likelihood. Each simulation was performed 100 times.

Table 1 shows that one of the CCP estimators performs nearly as well as ML, while using the other entails a noticeable efficiency loss. Every estimated coefficient is unbiased, each lying within one standard deviation of its true value. This attractive feature is replicated in all three of our experimental designs. In this design updating the CCP's with the likelihood yields standard errors on each coefficients that are within 10 percent of the standard errors obtained using ML. Thus the efficiency loss in data sets of moderate sizes appears small. Updating the CPP's with the weighted empirical likelihoods generated less precise estimates. Depending on the coefficient, the increase above the ML standard deviation ranges from negligible, for the discount factor β , to a magnitude of almost three, for the withdrawal cost α_3 . This efficiency loss appears to be driven by only using data on discrete choices to estimate the unobserved heterogeneity parameters. As we show in the next Monte Carlo, having additional data on a continuous outcome that is also affected by the unobserved heterogeneity leads to little difference between techniques that use the empirical likelihood to update the CCP's and those that use the model.

7.2 Monte Carlo 2: Entry and exit in oligopoly

Next we analyze a parameterization of the entry/exit game described in Section 4.3. This Monte Carlo has four distinctive features to focus on. First, unobserved heterogeneity affects both the dynamic discrete choice decisions and another outcome. Since this other outcome is also affected by the dynamic discrete choice, we must account for dynamic selection issues in estimation. Second, in contrast to the first experimental design, the unobserved heterogeneity is modeled as a stationary Markov process, an appealing assumption for an unobserved demand process. Third, we evaluate the estimator when the unobserved heterogeneity and the parameters in the outcome equation are estimated in a first stage, and only the parameters of the dynamic discrete choice decisions are estimated in the second stage. Finally, we exploit the finite dependence property of the entry/exit

game, and evaluate the performance of our estimator when the renewal action is a low probability event.

In this model the state of demand for the market, $s_t \in \{0, 1\}$, is unobserved by econometricians but observed by firms when they make their entry and exit decisions. Demand is in the low (high) state at time t when $s_t = 0$ ($s_t = 1$). The probability of a market being in the low state at $t+1$ given it was in the low state at time t is given by π_{LL} , with the corresponding probability of a persisting in the high state given by π_{HH} . Current profits for staying in or entering a market net of the profit shock are given by $u(E_t, M_t, s_t)$, which is linear in the state variables:

$$u(E_t, M_t, s_t) = \alpha_1(1 - s_t) + \alpha_2 s_t + \alpha_3(1 - M_t) + \alpha_4 E_t + \epsilon_t \quad (19)$$

As in section 3.2, E_t is an indicator for entry (versus incumbency), and M_t is a monopoly (versus duopoly) indicator. Substituting (19) into the conditional valuation function for staying in the market given in equation (4) yields:¹⁸

$$\begin{aligned} v_1(E_t, R_t, s_t) = & E_t R_t \left\{ \alpha_1(1 - s_t) + \alpha_2 s_t - \beta \sum_{s_{t+1}=0}^1 \ln[p_0(0, 1, s_{t+1})] \pi(s_{t+1}|s_t) \right\} \\ & + (1 - E_t R_t) \sum_{k=0}^1 p_k(E_t, R_t, s_t) \left\{ \alpha_1(1 - s_t) + \alpha_2 s_t + \alpha_3(1 - k) + \alpha_4 E_t \right. \\ & \left. - \beta \sum_{s_{t+1}=0}^1 \ln[p_0(0, 1 - k, s_{t+1})] \pi(s_{t+1}|s_t) \right\} + \beta \gamma \end{aligned}$$

where, as in Section 3.2, $R_t = 1$ indicates that there is no incumbent rival. The Type I extreme value profit shocks imply that the probability of entering or staying in the market is given by:

$$p(E_t, R_t, s_t) = \frac{\exp(v_1(E_t, R_t, s_t))}{1 + \exp(v_1(E_t, R_t, s_t))} \quad (20)$$

We also assume that the researcher has data on some measure of demand for each market in each time period. Demand in market n at time t is given y_{nt} , with the number of firms in the market given I_{nt} . We assume y_{nt} follows:¹⁹

$$y_{nt} = \gamma_1(1 - s_{nt}) + \gamma_2 s_{nt} + \gamma_3(I_{nt} = 1) + \gamma_4(I_{nt} = 2) + \zeta_{nt} \quad (21)$$

where ζ_{nt} is measurement error which we assume is normal and independently distributed.²⁰

¹⁸Note that the z_t 's and $f(\cdot)$'s have been replaced by s_t 's and $\pi(\cdot)$'s to indicate that the state variable is unobserved to the econometrician.

¹⁹We assume that the measure of demand is observed even when there are no firms in the market. The Monte Carlo were not sensitive to this assumption, though in this case one less parameter would be estimated.

²⁰Alternatively, it is a shock to demand that was unobserved by the firms until after their entry/exit decisions had been made.

Generating the data for the Monte Carlos proceeds in two steps. After setting the structural parameters, we choose an initial set of conditional choice probabilities, and iterated on (20) until a fixed point is reached. The inversion theorem guarantees that every fixed point corresponds to a Markov perfect equilibrium.²¹ Next, we draw paths, both for entry and exit as well as for the y 's and s 's, for particular markets given the probabilities found in the first step as well as the α 's and the γ 's. Each data set had observations on 3000 markets, simulated over 5 periods.²² We conducted 100 simulations.

The estimates from the Monte Carlos are presented in Table 2. The first set of columns corresponds to our base CCP algorithm described in section 5 where the updating of the CCP's comes from the model. The second set involves estimating the parameters governing the unobserved variables, the conditional probabilities of being in each of the unobserved states, and the parameters of the demand equation in a first stage. The likelihood contributions of the entry/exit choices are calculated from the data in this first stage as opposed to solving out the dynamic discrete choice problem. We then use the conditional probabilities of being in a particular state as weights in a second stage maximization of the dynamic discrete choice problem. The exit probabilities used to form the future value terms are then calculated from the data. The data generating process for these first two sets of columns yields probabilities of exiting a market between 4% and 10% for monopolies and between 13% and 22% for duopolies.

The small sample distribution of the base CCP estimator and the CCP estimator where unobserved heterogeneity is estimated in a first stage yield estimates hardly differ. Both sets of structural estimators are unbiased and the standard deviations are low. These results suggest that the efficiency gains from using the base method are not overwhelming when compared to the two stage method.

Of particular interest are the coefficients on monopoly and duopoly in the demand equation. If the unobserved heterogeneity is ignored, and the demand outcome equation is estimated by OLS, the coefficients are -0.18 and -0.41 respectively, significantly biased upward relative to the true values of -0.3 and -0.7. Firm entry has a much greater effect than an OLS estimator would predict. In contrast, CCP techniques that control dynamic selection capture the competitive effects quite precisely. This finding demonstrates the value of the CCP estimator in controlling for dynamic selection, even in situations where estimating product demand is the main empirical objective, as opposed to other goals such as recovering primitives that define the fixed costs of firms which partially determine

²¹An alternative approach would be to use the algorithm developed by Pakes and McGuire (1994).

²²This is roughly the number of U.S. counties.

their entry decisions.

The final set of columns again uses the baseline algorithm with the CCP's updated using the model. Here, however, we significantly lower the probability of exiting by increasing the demand in the low state. We adjusted the data generating process so that the probability of exiting fell to between 1% and 1.6% for monopolists and 2.3% and 5% for duopolists, thereby reducing the probability of the terminating action. As one might anticipate, most of the standard errors increase with the standard error almost doubling for the entry cost coefficient. Overall, however, the estimates are precise and centered around the true values.

We also estimated the model without using any of the data on demand, but in this case our estimator performed poorly when the unobserved demand state was modeled as a Markov process. However when we simplified the model to eliminate demand transitions, by setting $p_{HH} = p_{LL} = 1$, and treated unobserved demand as a time-invariant effect, the resulting estimates were centered around the true parameter values with tight standard deviations.²³ This last set of results suggests that unless some (error ridden) data on the unobserved process is available, estimating the unobserved process from data on discrete choices alone could flounder. Hence, the method performs well either when the unobserved heterogeneity transitions over time and data is available on another outcome measure or when the unobserved heterogeneity is permanent.

7.3 Monte Carlo 3: Scheduling

Our last study is based on the scheduling problem we described in Section 4.3, where there are both continuous and discrete choices. In this design we also investigate the role of observed heterogeneity in the estimation of unobserved processes.²⁴ We treat the state of demand, z_{1t} as unobserved and we therefore label z_{1t} as s_t . Three parameters govern the transitions in the state of demand between L and H , where π_{LL} is the probability of persisting in state L , π_{HH} is the probability of persisting in state H , and π_L is the initial probability of beginning in state L . The state of the plant, z_{2t} is treated as observed and we therefore label z_{2t} as x_t . We allow the state of the plant to take on either two or four values to test the sensitivity of the estimates in increases in the dimensionality of the observed states.

²³These results are available from the authors upon request.

²⁴Kashahara and Shimotsu (2007) analyze how the identification of dynamic discrete choice with unobserved heterogeneity is partly determined the number of states observed by the econometrician.

We can then substitute in for z_{1t} and z_{2t} in equation (6) with s_t and x_t which yields:

$$\varepsilon_{0t} + \alpha_3 s_t + 2\alpha_4 x_t (\ln c_t) = \left(\sum_{s_{t+1}} [\ln(p_1(s_{t+1}, x_t + 1)) - \ln(p_1(s_{t+1}, x_t))] f(s_{t+1}|s_t) \right) \gamma_0 \gamma_1 c_t^{\gamma_1} (\gamma_0 + c_t^{\gamma_1})^{-2} \quad (22)$$

Appealing to the monotonicity between ε_{0t} and c_t^o , the next lemma gives the probability distribution of c_t^o conditional on the state variables and operating the plant.

Lemma 5 *Define the mapping $\Upsilon(s_t, x_t, c_t)$ as*

$$\begin{aligned} \Upsilon(s_t, x_t, c_t) \equiv & \alpha_3 s_t + 2\alpha_4 x_t (\ln c_t) \\ & - \left(\sum_{s_{t+1}} [\ln(p_1(s_{t+1}, x_t + 1)) - \ln(p_1(s_{t+1}, x_t))] f(s_{t+1}|s_t) \right) \gamma_0 \gamma_1 c_t^{\gamma_1} (\gamma_0 + c_t^{\gamma_1})^{-2} \end{aligned}$$

Recalling $\sigma \varepsilon_{0t}$ is distributed as a standard normal random variable with distribution function $\Phi(\cdot)$, then

$$\Pr \{c_t^o \leq \bar{c} | s_t, x_t\} = \Phi[\Upsilon(s_t, x_t, \bar{c}) / \sigma]$$

If the state vector was observed, and the conditional choice probabilities were known as a mapping from the states, then an ML estimator could be obtained from data on (s_t, x_t, c_t) , by maximizing

$$\sum_{n=1}^{\mathcal{N}} \sum_{t=1}^{\mathcal{T}} \left\{ \log \left[\frac{\partial \Upsilon(s_{nt}, x_{nt}, c_t)}{\partial c} \right] + \log \phi[\Upsilon(s_{nt}, x_{nt}, c_t) / \sigma] - \log \sigma \right\}$$

with respect to $(\gamma_0, \gamma_1, \sigma, \alpha_3, \alpha_4)$ where $\phi(\cdot)$ is the probability density function for a standard normal variable. The demand state transitions $(\pi_{LL}, \pi_{HH}, \pi_L,)$ could be separately estimated from data on s_t alone and, in a second stage, a CCP estimator could be used to recover the fixed costs $\alpha_2 - \alpha_1$. However, with s_t unobserved, we can use the methods developed in this paper to estimate the model and investigate its finite sample performance.

We solve the model for 3000 plants each running for 5 periods, and perform 100 Monte Carlo simulations. Estimation follows the first method where the conditional choice probabilities are updated using the model itself. The maximization part of the EM algorithm proceeds in three stages. Having solved and simulated the model, we then estimate the parameters governing the continuous choice decision. Next, the CCP's are updated. Finally, values for the expected value of utility from the continuous choice and also the future utility term, estimated from the first two stages, are substituted into a logit model to estimate the one remaining parameter, the fixed cost of operating the plant.

The true parameters, the average value of the estimates, and the standard deviation of these estimates are given in Table 3. The first set of columns gives the case where the state of the plant, x , takes on only two values, while the second set of columns allows x to take on one of four values. The estimator performs well under both designs; the means of the estimates are well within three significant digits and one standard deviation of the true parameter values. Nevertheless the usefulness of having observed variation in covariates is evident. The parameters governing unobserved heterogeneity are more precisely estimated when x takes on four values than when x takes on only two values both difference between the means and the true values, as well as their standard deviations, fall.

8 Conclusion

Estimation of dynamic discrete choice models is computationally costly, particularly when implementing controls for unobserved heterogeneity. CCP estimators provide a computationally cheap way of estimating dynamic discrete choice problems. In this paper we extend the class of models that are easily adapted to the CCP framework by both exploiting the finite dependence property present in many dynamic discrete choice problems and showing how to incorporate unobserved heterogeneity into CCP estimation. Unobserved states are accounted for via finite mixture distributions. The computational simplicity of the estimator allows the unobserved state variables to follow a Markov chain and allows for the estimation of mixed discrete and continuous choice processes. Our baseline algorithm builds upon the insights of the EM algorithm: it iterates between updating the conditional probabilities of being a particular type, updating the CCP's for any given state (observed and unobserved), forming the expected future conditional valuation functions as mappings of the CCP's (using the finite dependence representation), and maximizing a likelihood function where the future values terms are in large a part a function of the CCP's.

The baseline method forms the likelihood to obtain estimates of the CCP's, but our algorithm can be adapted to outcomes besides the discrete choices themselves, including but not limited to first order conditions for continuous choices. For these cases, we show how to estimate the parameters governing the unobserved heterogeneity in a first stage. Rather than the using likelihood itself, we update the CCP's using the empirical distribution of discrete choices weighted by the estimated probabilities of being in particular unobserved states conditional on the whole sample. As a stand alone, this is a flexible estimator of outcomes that are affected by discrete choices where we need to specifically account for dynamic selection. The estimator does so by neither taking a strong

stand on the idiosyncratic effects that help determine the discrete choices, nor making assumptions about the parametric specification of discrete choice specific costs and benefits. This approach also serves as a first stage estimator for blending unobserved heterogeneity into non-likelihood based approaches such as Hotz et al. (1994) and Bajari, Benkard, and Levin (2007) in a second stage.²⁵ The conditional probabilities of being in particular unobserved states are estimated in the first stage and are then used as weights when estimating the parameters of the second stage.

We illustrated the small sample properties of our estimator using a set of Monte Carlo experiments designed to highlight the wide variety of problems that can be estimated with the algorithm. The first Monte Carlo illustrated both ways of updating the CCP's, using either the likelihoods themselves or the conditional probabilities of being in each of the unobserved states as weights, and compared the estimates to those produced by full information maximum likelihood. The second was a dynamic entry/exit example where the value of being in a market was in part affected by unobserved heterogeneity in the demand levels for particular markets. The unobserved states were allowed to transition over time and the example explicitly incorporated dynamic selection. We estimated the model both by updating the CCP's with the model and by estimating the unobserved heterogeneity in a first stage. Our final Monte Carlo illustrated the performance of our methods in mixed discrete/continuous settings in the presence of unobserved heterogeneity. For all three sets of Monte Carlos, the estimators performed quite well both in terms of the precision of the estimates as well as the speed at which the estimates are obtained.

9 Appendix

9.1 Proof of Lemma 1

Abbreviating, let $v_{kt} \equiv v_k(z_t)$, $\mathcal{H}_{0t} \equiv \mathcal{H}(e^{v_{1t}}, e^{v_{2t}}, \dots, e^{v_{Kt}})$, and $G_k(\varepsilon_t) \equiv \partial G(\varepsilon_t) / \varepsilon_{kt}$. Using the assumption that $\mathcal{H}(Y_1, Y_2, \dots, Y_K)$ is homogeneous of degree one, and therefore the partial derivative $\mathcal{H}_k(Y_1, Y_2, \dots, Y_K)$ is homogenous of degree zero, we obtain the density of ε_{kt} when the k^{th} choice is optimal

$$G_k(v_{kt} + \varepsilon_{kt} - v_{1t}, \dots, v_{kt} + \varepsilon_{kt} - v_{Kt}) = \mathcal{H}_k(e^{v_{1t}}, \dots, e^{v_{Kt}}) \exp[-\mathcal{H}_0 e^{-v_{kt} - \varepsilon_{kt}}] e^{-\varepsilon_{kt}}$$

²⁵See Finger (2007) for an example of our method used in conjunction with Bajari, Benkard, and Levin's method.

Integrating over $G_k(\varepsilon_t)$ yields the conditional choice probability

$$\begin{aligned}
p_k(z_t) &= \int G_k(v_{kt} + \varepsilon_k - v_{1t}, \dots, \varepsilon_k, \dots, v_{kt} + \varepsilon_k - v_{Kt}) d\varepsilon_k \\
&= \mathcal{H}_k(e^{v_{1t}}, \dots, e^{v_{Kt}}) \int \exp[-\mathcal{H}_{0t} e^{-v_{kt} - \varepsilon_k}] e^{-\varepsilon_k} d\varepsilon_k \\
&= e^{v_{kt}} \mathcal{H}_k(e^{v_{1t}}, \dots, e^{v_{Kt}}) / \mathcal{H}_{0t}
\end{aligned} \tag{23}$$

The expected contribution of the disturbance from the k^{th} choice is

$$\begin{aligned}
\int d_{kt} \varepsilon_{kt} dG(\varepsilon_t) &= \int \varepsilon_k G_k(v_{kt} + \varepsilon_k - v_{1t}, \dots, v_{kt} + \varepsilon_k - v_{Kt}) d\varepsilon_k \\
&= \mathcal{H}_k(e^{v_{1t}}, \dots, e^{v_{Kt}}) \int \varepsilon_k \exp[-\mathcal{H}_{0t} e^{-v_{kt} - \varepsilon_k}] e^{-\varepsilon_k} d\varepsilon_k \\
&= e^{v_{kt}} \mathcal{H}_k(e^{v_{1t}}, \dots, e^{v_{Kt}}) [\gamma - v_{kt} + \log \mathcal{H}_{0t}] / \mathcal{H}_{0t}
\end{aligned}$$

where the last line uses the property that the mean of a Type 1 extreme value distribution with parameter a is $(\gamma + \log a)$ and $\gamma = 0.57\dots$ is Euler's constant. Taking the quotient yields $w_k(\psi[p(z_t)])$ as a function of $\psi[p(z_t)]$

$$w_k(\psi[p(z_t)]) = \gamma + \log(\mathcal{H}_{0t} e^{-v_{kt}}) = \gamma + \log \mathcal{H}(e^{\psi_k^1[p(z_t)]}, e^{\psi_k^2[p(z_t)]}, \dots, e^{\psi_k^K[p(z_t)]})$$

9.2 Proof of Lemma 2

From the definitions of p_{jrt} and $\mathcal{H}(Y)$ for the nested logit given in the text, it follows from Equation (23) that

$$p_{jrt} = \mathcal{H}_{0t}^{-1} \left[\sum_{k=1}^{K_r} e^{\delta_r v_{krt}} \right]^{1/\delta_r - 1} e^{\delta_r v_{jrt}}$$

Summing over all choices within the r^{th} cluster we obtain p_{rt} , the conditional probability of making one of the choices in the r^{th} cluster

$$p_{rt} \equiv \sum_{k=1}^{K_r} p_{krt} = \mathcal{H}_{0t}^{-1} \left[\sum_{j=1}^{K_r} e^{\delta_r v_{jrt}} \right]^{1/\delta_r}$$

We use these equations to substitute out the summation ($e^{\delta_r v_{1rt}} + \dots + e^{\delta_r v_{K_r rt}}$) and make \mathcal{H}_{0t} the subject of the resulting formula, proving that for any two choices j and k belonging to clusters r and s

$$\mathcal{H}_{0t} = e^{v_{jrt}} p_{rt}^{(1/\delta_r) - 1} p_{jrt}^{-1/\delta_r} = e^{v_{kst}} p_{st}^{(1/\delta_s) - 1} p_{kst}^{-1/\delta_s}$$

Upon rearrangement, the difference in their conditional valuation functions emerges as

$$v_{jrt} - v_{kst} = \frac{1}{\delta_r} \log(p_{jr}) - \frac{1}{\delta_s} \log(p_{kst}) + \left(1 - \frac{1}{\delta_r}\right) \log(p_{rt}) - \left(1 - \frac{1}{\delta_s}\right) \log(p_{st})$$

Applying the formula for $w_k(\psi[p(z_t)])$ in the statement of Lemma 1

$$\begin{aligned}
E[\varepsilon_{kst} | d_{kst} = 1] &= \gamma + \log \left[\mathcal{H} \left(e^{v_{11t} - v_{kst}}, e^{v_{12t} - v_{kst}}, \dots, e^{v_{KRt} - v_{kst}} \right) \right] \\
&= \gamma + \log \left\{ \sum_{r=1}^R \left[\sum_{j=1}^{K_r} e^{\delta_s (v_{jrt} - v_{kst})} \right]^{1/\delta_s} \right\} \\
&= \gamma - \frac{1}{\delta_s} \log(p_{kst}) - \left(1 - \frac{1}{\delta_s} \right) \log(p_{st}) + \log \left\{ \sum_{r=1}^R p_{rt}^{1-1/\delta_r} \left[\sum_{j=1}^{K_r} p_{jrt}^{\delta_s/\delta_r} \right]^{1/\delta_s} \right\}
\end{aligned}$$

9.3 Proof of Proposition 1

We first show that for all $r \in \{t+1, \dots, \rho\}$ with $\rho < \infty$, and any stochastic choice process $\lambda(j, z_t)$ inducing the associated probability distributions $\kappa_r(z_r | j, z_t)$ as defined in Section 3, the conditional valuation functions $v_j(s_t)$ can be expressed as:

$$\begin{aligned}
v_j(z_t) &= u_j(z_t) + \beta^{r-t} \sum_{z_r} [v_{\lambda_r(j)}(z_r) - u_{\lambda_r(j)}(z_r)] \kappa_r(z_r | j, z_t) \\
&\quad + \sum_{\tau=t+1}^r \sum_{k=1}^K \sum_{z_\tau} \beta^{\tau-t} p_k(z_\tau) \left\{ \psi_{\lambda_r(j)}^k[p(z_\tau)] + u_{\lambda_r(j)}(z_\tau) + w_k[p(z_\tau)] \right\} \kappa_\tau(z_\tau | j, z_t)
\end{aligned} \tag{24}$$

As in the text, we denote by $v_{\lambda_r(j)}(z_\tau)$ the conditional valuation function evaluated at z_τ for the choice dictated by $\lambda_r(j, z_t)$ and so forth for $\psi_{\lambda_r(j)}^k[p(z_\tau)]$ and $u_{\lambda_r(j)}(z_\tau)$. Section 3.1 established this equation holds for $r = t+1$. Proceeding with an induction, we show that if the statement is true for all $\tau \in \{t+1, \dots, r\}$, then it holds true period $r+1$ too. Following equation (1) we telescope $v_{\lambda_r(j)}(z_r)$ forward one period by applying the $\lambda(j, z_t)$ choice process to obtain

$$\begin{aligned}
v_{\lambda_r(j)}(z_r) - u_{\lambda_r(j)}(z_r) &= \beta \sum_{z_{r+1}} v_{\lambda(j)}(z_{r+1}) f_{\lambda_{r+1}(j)}(z_{r+1} | z_r) \\
&\quad + \sum_{k=1}^K \sum_{z_{r+1}} \beta p_k(z_{r+1}) \left\{ \psi_{\lambda_{r+1}(j)}^k[p(z_{r+1})] + w_k[p(z_{r+1})] \right\} f_{\lambda_{r+1}(j)}(z_{r+1} | z_r)
\end{aligned}$$

Substituting this expression into equation (24) for $v_j(z_t)$ and rearranging we have

$$\begin{aligned}
&v_j(z_t) - u_j(z_t) - \sum_{\tau=t+1}^r \sum_{k=1}^K \sum_{z_\tau} \beta^{\tau-t} p_k(z_\tau) \left\{ \psi_{\lambda_\tau(j)}^k[p(z_\tau)] + u_{\lambda_\tau(j)}(z_\tau) + w_k[p(z_\tau)] \right\} \kappa_\tau(z_\tau | j, z_t) \\
&= \beta^{r+1-t} \sum_{z_r} \left[\sum_{z_{r+1}} v_{\lambda_{r+1}(j)}(z_{r+1}) f_{\lambda_{r+1}(j)}(z_{r+1} | z_r) \right] \kappa_r(z_r | j, z_t) \\
&\quad + \beta^{r+1-t} \sum_{z_r} \left[\sum_{k=1}^K \sum_{z_{r+1}} \beta p_k(z_{r+1}) \left\{ \psi_{\lambda_{r+1}(j)}^k[p(z_{r+1})] + w_k[p(z_{r+1})] \right\} f_{\lambda_{r+1}(j)}(z_{r+1} | z_r) \right] \kappa_r(z_r | j, z_t)
\end{aligned}$$

For all $v_{j'}(z_{r+1})$, it follows from the definitions of $\kappa_r(s | j, z_t, \rho)$ and $f_{\lambda_{r+1}(j)}(z_{r+1} | z_r)$ that

$$\sum_{z_r} \sum_{z_{r+1}} v_{j'}(z_{r+1}) f_{\lambda_{r+1}(j)}(z_{r+1} | z_r) \kappa_r(z_r | j, z_t) = \sum_{z_{r+1}} v_{j'}(z_{r+1}) \kappa_{r+1}(z_{r+1} | j, z_t)$$

Substituting for $\kappa_{r+1}(z|j, z_t)$ and rearranging yields

$$\begin{aligned}
v_j(z_t) &= u_j(z_t) + \sum_{\tau=t+1}^r \sum_{k=1}^K \sum_{z_\tau} \beta^{\tau-t} p_k(z_\tau) \left\{ \psi_{\lambda_\tau(j)}^k [p(z_\tau)] + u_{\lambda_\tau(j)}(z_\tau) + w_k [p(z_\tau)] \right\} \kappa_\tau(z_\tau|j, z_t) \\
&+ \beta^{r+1-t} \sum_{z_{r+1}} \left(v_{\lambda_{r+1}(j)}(z_{r+1}) + \sum_{k=1}^K p_k(z_{r+1}) \left\{ \psi_{\lambda_{r+1}(j)}^k [p(z_{r+1})] + w_k [p(z_{r+1})] \right\} \right) \kappa_{r+1}(z_{r+1}|j, z_t) \\
&= u_j(z_t) + \beta^{r+1-t} \sum_{z_{r+1}} [v_{\lambda_{r+1}(j)}(z_{r+1}) - u_{\lambda_{r+1}(j)}(z_{r+1})] \kappa_{r+1}(z_{r+1}|j, z_t) \\
&+ \sum_{\tau=t+1}^{r+1} \sum_{k=1}^K \sum_{z_\tau} \beta^{\tau-t} p_k(z_\tau) \left\{ \psi_{\lambda_\tau(j)}^k [p(z_\tau)] + u_{\lambda_\tau(j)}(z_\tau) + w_k [p(z_\tau)] \right\} \kappa_\tau(z_\tau|j, z_t)
\end{aligned}$$

where the second equality is obtained by adding in and subtracting $u_{\lambda_{r+1}(j)}(z_{r+1})$ to complete the induction. Upon forming the difference $v_j(z_t) - v_{j'}(z_t)$ we thus obtain the expression stated in the proposition from Equation (24) if and only if there is a finite integer ρ and choice process $\lambda(j, z_t)$ such that for two paths satisfying the equality

$$\sum_{z_{t+\rho}} [v_{\lambda_{t+\rho}(j)}(z_{t+\rho}) - u_{\lambda_{t+\rho}(j)}(z_{t+\rho})] \kappa_{t+\rho}(z|j, z_t) = \sum_{z_{t+\rho}} [v_{\lambda_{t+\rho}(j')} (z_{t+\rho}) - u_{\lambda_{t+\rho}(j')} (z_{t+\rho})] \kappa_{t+\rho}(z_{t+\rho}|j', z_t)$$

9.4 Proof of Lemma 3

Weighting by $p_k(z_t)$ the expression for $w_k(\psi[p(z_t)])$ derived in Section 2, summing over all choices $k \in \{1, \dots, K\}$ and rearranging yields an expression for the social surplus function for any model with generalized extreme values, namely

$$\sum_{k=1}^K p_k(z_t) [v_k(z_t) + w_k(\psi[p(z_t)])] = \ln \mathcal{H}(e^{v_1(z_{t+1})}, \dots, e^{v_K(z_{t+1})}) + \gamma$$

which is a restatement of the corollary to Theorem 1 in McFadden (1978). It now follows from Bellman's (1957) equation that the conditional valuation function for the j^{th} choice can be written as

$$\begin{aligned}
v_j(z_t) &= u_j(z_t) + \beta \sum_{z_{t+1}} \sum_{k=1}^K p_k(z_{t+1}) \{v_k(z_{t+1}) + w_k(\psi[p(z_{t+1})])\} f_j(z_{t+1}|z_t) \quad (25) \\
&= u_j(z_t) + \beta \sum_{z_{t+1}} \left[\ln \mathcal{H}(e^{v_1(z_{t+1})}, \dots, e^{v_K(z_{t+1})}) \right] f_j(z_{t+1}|z_t) + \beta\gamma
\end{aligned}$$

Since $\mathcal{H}_1(Y_1, \dots, Y_K) = 1$ under our assumptions for this renewal example, Equation (23) implies

$$\begin{aligned}
\ln \mathcal{H}(e^{v_1(z_t)}, \dots, e^{v_K(z_t)}) &= v_1(z_t) - \ln p_1(z_t) \quad (26) \\
&= u_1(z_t) + \beta V^* - \ln p_1(z_t)
\end{aligned}$$

where the second equality follows from the fact that the first choice is a renewal action. Substituting the expression for $\ln \mathcal{H}(e^{v_1(z_t)}, \dots, e^{v_K(z_t)})$ obtained from Equation (26) into Equation (25) then yields (2) in the text.

9.5 Proof of Lemma 4

Since the disturbances are distributed Type 1 Extreme value, the formulas in Lemma 2 imply

$$\sum_{k=1}^K p_k(z_t) \left\{ \psi_{\lambda(j)}^k [p(z_t)] + u_{\lambda(j)}(z_t) + w_k(\psi[p(z_t)]) \right\} = u_{\lambda(j)}(z_t) + \gamma - \log p_{\lambda(j)}(z_t)$$

Applying Proposition 1 to the differences $v_1(z_t) - v_0(z_t)$ and cancelling the γ terms yields

$$\begin{aligned} v_1(z_t) - v_0(z_t) &= u_1(z_t) - u_0(z_t) \\ &+ \sum_{\tau=t+1}^T \sum_{z_\tau} \beta^{\tau-t} [u_{\lambda(1)}(z_\tau) - \log p_{\lambda(1)}(z_\tau)] \kappa_\tau(z_\tau|1, z_t) \\ &- \sum_{\tau=t+1}^T \sum_{z_\tau} \beta^{\tau-t} [u_{\lambda(0)}(z_\tau) - \log p_{\lambda(0)}(z_\tau)] \kappa_\tau(z_\tau|0, z_t) \end{aligned}$$

The state variables in this example are $z_t = (h_t, d_{t-1})$. We establish $\rho = 2$ in this case, by noting that if the woman participates in period t , and then does not participate in periods $t+1$ and $t+2$, her state variables in period $t+3$ have the same probability distribution as if she does not participate in period t , participates in period $t+1$ instead but does not participate at $t+2$. Let $\lambda(j, z_t)$ denote any predetermined choice process defined by: $\lambda(1, h_t, d_{t-1})$ sets $(d_{1t}, d_{1,t+1}, d_{1,t+2}) = (1, 0, 0)$, and $\lambda(0, h_t, d_{t-1})$ sets $(d_{1t}, d_{1,t+1}, d_{1,t+2}) = (0, 1, 0)$. Also let $f(z)$ denote the probability that human capital increases by z units from participating in any period. Then in period $t+3$ her state will be $(h_t + z, 0)$ regardless of whether she takes the choice path $\lambda(1, h_t, d_{t-1})$ or $\lambda(0, h_t, d_{t-1})$. It now follows that

$$\begin{aligned} &\sum_{\tau=t+1}^{t+2} \sum_{h_\tau} \beta^{\tau-t} [u_{\lambda(1)}(h_\tau, d_{\tau-1}) - \log p_{\lambda(1)}(h_\tau, d_{\tau-1})] \kappa_\tau(h_\tau, d_{\tau-1}|1, h_t, d_{t-1}) \\ &= \sum_z \left\{ \beta [u_0(h_t + z, 1) - \log p_0(h_t + z, 1)] + \beta^2 [u_0(h_t + z, 0) - \log p_0(h_t + z, 0)] \right\} f(z) \end{aligned}$$

and

$$\begin{aligned} &\sum_{\tau=t+1}^{t+2} \sum_{h_\tau} \beta^{tau-t} [u_{\lambda(0)}(h_\tau, d_{\tau-1}) - \log p_{\lambda(0)}(h_\tau, d_{\tau-1})] \kappa_\tau(h_\tau, d_{\tau-1}|0, h_t, d_{t-1}) \\ &= \beta [u_1(h_t, 0) - \log p_1(h_t, 0)] + \sum_z \left\{ \beta^2 [u_0(h_t + z, 1) - \log p_0(h_t + z, 1)] \right\} f(z) \end{aligned}$$

Collecting terms in the expression for $v_1(h_t, d_{t-1}) - v_0(h_t, d_{t-1})$ we obtain the difference in conditional valuation functions between working and not working, which is Equation (5) in the text.

9.6 Proof of Proposition 2

The definition of this estimator ensures (π_1, p_1) define a set of probability distributions that generate (d, x) , the choices and the observed state variables. Thus the data generation process defined by

(φ_0, p_0) is also fully characterized by (φ_1, p_1) . It only remains to show that the restrictions implied by the structural model with true parameter vector (φ_0, p_0) are also satisfied by (φ_1, p_1) . In other words, optimizing behavior by the agents living in a world with structural parameters φ_1 leads to the vector of conditional choice probabilities p_1 .

Following the notation in the text, let $z \equiv (x, s)$ where s denotes the unobserved variables, let $u_k(z) \equiv u_k(z; \theta)$ denote the component of current utility that depends on the state variables when the k^{th} choice is made and given parameterization θ , and let $f_k(z' | z)$ denote the (finite) state space probability transition function for z given choice j . The structural parameters for this problem are denoted by $\varphi \equiv (\theta, \pi)$. We form the expected lifetime value function, $V(z_t) \equiv V(z_t; \varphi, p_1)$, using the finite dependence representation discussed in Section 2, constructed from $u_k(z; \theta)$ parameterized by (φ, p) and evaluated at p_1 . It follows that, conditional on $(z_{nt}, \varepsilon_{nt})$, the value of choosing k is

$$v_k(z_{nt}) + \varepsilon_{knt} \equiv u_k(z_{nt}; \theta) + \varepsilon_{knt} + \sum_{z'} V(z_t; \varphi, p_1) f_k(z' | z_{nt})$$

From the definition of φ_1 and p_1 , formed from generic elements $p_{1k}(z_{nt})$

$$p_{1k}(z_{nt}) = \Pr \left\{ \begin{array}{l} \varepsilon_{kt} - \varepsilon_{lt} \geq u_l(z_{nt}; \varphi) - u_k(z_{nt}; \varphi) + \sum_{z'} V(z_t; \varphi, p_1) [f_l(z' | z_{nt}) - f_k(z' | z_{nt})] \\ \text{for all } l \in \{1, \dots, K\} \end{array} \right\}$$

for all (k, z) where as before ε_t is parameterized by $G(\varepsilon_t; \varphi)$. Appealing to Proposition 1, $V(z_t; \varphi, p_1)$ is the expected valuation function for the dynamic programming problem, and therefore the conditional choice probabilities arise from individuals optimizing in a world with structural parameters φ_1 . We conclude that the choice probabilities conditional on the observed state variables for both φ_0 and φ_1 are the same, thus proving φ_0 and φ_1 are observationally equivalent.

9.7 Proof of Lemma 5

From the definition $\Upsilon(s_{1t}, s_{2t}, c_t)$ first order equation can be expressed

$$\varepsilon_{0t} - \Upsilon(s_{1t}, s_{2t}, c_t) = 0$$

Holding constant the state variables (s_{1t}, s_{2t}) while totally differentiating this equation with respect to c_t and ε_{0t} yields

$$\frac{dc_t}{d\varepsilon_{0t}} = \left[\frac{\partial \Upsilon(s_{1t}, s_{2t}, c_t)}{\partial c_t} \right]^{-1}$$

The second order conditions require $\partial \Upsilon(s_{1t}, s_{2t}, c_t) / \partial c_t > 0$, an inequality that holds globally in our specification because $0 < \alpha_1 < 1$. Therefore c_t is monotone increasing in ε_{0t} , and hence

$$\Pr \{c_t \leq \bar{c} | s_{1t}, s_{2t}, \} = \Pr \{\varepsilon_{0t} \leq \Upsilon(s_{1t}, s_{2t}, \bar{c}) | s_{1t}, s_{2t}, \} \equiv \Phi[\Upsilon(s_{1t}, s_{2t}, \bar{c}) / \sigma]$$

9.8 Asymptotic Covariance Matrix

There are several variations on the CCP estimator depending on whether the likelihood or the weighted empirical likelihood is used to estimate the conditional choice probabilities, whether there are auxiliary outcome equations or not, whether there are continuous choices or not, and if so which of the two reduced form utility representations is used in estimation. Here we derive asymptotic covariance for two main cases. First we account for continuous choices when the second representation of the reduced utility function is used and the likelihood is used to estimate the CCP's. The formula we obtain reduces to the base case when there are no supplementary outcomes. Then we modify the base case by using the empirical likelihood instead of the likelihood to compute the CCP's.

When there are continuous choices augmenting the discrete choices, in estimation we form approximate sample averages of expression from $U_k(c, z, \varepsilon_0; \varphi) \equiv U_k(c, z, \varepsilon_0)$, rather than using its exact expectation over ε_0 when the optimal continuous choice is taken, namely $u_k(z; \varphi, p)$. For brevity we consider only the second representation described in Section 4, but deriving the asymptotic covariance for the first representation is similar. Let c_{kn} denote the continuous choice(s) for an individual observation with characteristic z making discrete choice k , designated by the indicator function $d_{kn} = 1$. Form the KZ dimensional vector function $U_n(\varphi, p)$ for each $n \in \{1, \dots, \mathcal{N}\}$ from the sample points

$$U_{nkz}(\varphi, p) \equiv \left[\mathcal{N}^{-1} \sum_{n'=1}^{\mathcal{N}} I\{z_{n'} = z\} d_{kn'} \right]^{-1} I\{z_n = z\} d_{kn} U_k[c_{kn}, z, \lambda_k(c_{kn}, z, \varphi, p); \varphi]$$

weighted for expositional convenience to reflect the occurrence of events (z, k) in the sample. Also define the sample average as

$$u^{(\mathcal{N})}(\varphi, p) = \mathcal{N}^{-1} \sum_{n=1}^{\mathcal{N}} U_n(\varphi, p)$$

and form its KZ dimensional vector function population analogue $u(\varphi, p)$ with generic components $u_k(z; \varphi, p)$. To make explicit the role of substituting $u^{(\mathcal{N})}(\varphi, p)$ for $u(\varphi, p)$ in the approximate score, denote by $L(d_n | x_n; u^{(\mathcal{N})}(\varphi, p), \varphi, p)$ the approximate likelihood when we substitute $u^{(\mathcal{N})}(\varphi, p)$ for $u(z; \varphi, p)$ into the likelihood for (φ, p) . The CCP estimator is defined by the two equation vectors

$$p^* = \mathcal{L}\left(u^{(\mathcal{N})}(\varphi^*, p^*), \varphi^*, p^*\right)$$

and

$$0 = \frac{1}{\mathcal{N}} \sum_{n=1}^{\mathcal{N}} S_n\left(u^{(\mathcal{N})}(\varphi^*, p^*), \varphi^*, p^*\right)$$

where $\sum_{n=1}^{\mathcal{N}} S_n(u(\varphi, p), \varphi, p)$ is the score for (φ, p) , and more generally

$$S_n\left(u^{(\mathcal{N})}(\varphi, p), \varphi, p\right) \equiv \frac{\partial \log [L(d_n | x_n; u^{(\mathcal{N})}(\varphi, p), \varphi, p)]}{\partial \varphi}$$

Under standard regularity conditions (φ^*, p^*) converges in probability to its true value (φ_0, p_0) at rate $\mathcal{N}^{1/2}$, and $\mathcal{N}^{-1/2} [u^{(\mathcal{N})}(\varphi, p) - u(\varphi, p)]$ satisfies a central limit theorem for all (φ, p) .

The asymptotic covariance matrix for φ^* is constructed from the derivatives of the vector of discrete choice likelihoods conditional on observed and unobserved states, formed over all choices and all states and denoted by

$$\begin{aligned}\mathcal{L}_u &\equiv \frac{\partial \mathcal{L}(u_0, \varphi_0, p_0)}{\partial u} \\ \mathcal{L}_\varphi &\equiv \frac{\partial \mathcal{L}[u(\varphi_0, p_0), \varphi_0, p_0]}{\partial \varphi} \\ \mathcal{L}_p &\equiv \frac{\partial \mathcal{L}[u(\varphi_0, p_0), \varphi_0, p_0]}{\partial p}\end{aligned}$$

where $u_0 \equiv u(\varphi_0, p_0)$, as well as the limit matrices, defined as

$$\begin{aligned}A_u &\equiv \lim_{\mathcal{N} \rightarrow \infty} \left[\frac{1}{\mathcal{N}} \sum_{n=1}^{\mathcal{N}} \frac{\partial S_n(u_0, \varphi_0, p_0)}{\partial u} \right] \\ A_\varphi &\equiv \lim_{\mathcal{N} \rightarrow \infty} \left[\frac{1}{\mathcal{N}} \sum_{n=1}^{\mathcal{N}} \frac{\partial S_n[u(\varphi_0, p_0), \varphi_0, p_0]}{\partial \varphi} \right] \\ A_p &\equiv \lim_{\mathcal{N} \rightarrow \infty} \left[\frac{1}{\mathcal{N}} \sum_{n=1}^{\mathcal{N}} \frac{\partial S_n[u(\varphi_0, p_0), \varphi_0, p_0]}{\partial p} \right]\end{aligned}$$

From the identities

$$0 = \mathcal{L}\left(u^{(\mathcal{N})}(\varphi^*, p^*), \varphi^*, p^*\right) - p^* = \mathcal{L}(u_0, \varphi_0, p_0) - p_0$$

we expand the second equation to the first order and rearrange, obtaining

$$\mathcal{L}_u \sqrt{\mathcal{N}} \left[u^{(\mathcal{N})}(\varphi^*, p^*) - u_0 \right] = (I - \mathcal{L}_p) \sqrt{\mathcal{N}} (p^* - p_0) - \mathcal{L}_\varphi \sqrt{\mathcal{N}} (\varphi^* - \varphi_0) + o_p(1) \quad (27)$$

Similarly the second set of equations can be expanded around its defining sample moments to obtain

$$A_u \sqrt{\mathcal{N}} \left(u^{(\mathcal{N})}(\varphi^*, p^*) - u_0 \right) = N^{-1/2} \sum_{n=1}^{\mathcal{N}} S_n(u_0, \varphi_0, p_0) \quad (28)$$

$$-A_\varphi \sqrt{\mathcal{N}} (\varphi^* - \varphi_0) - A_p \sqrt{\mathcal{N}} (p^* - p_0) + o_p(1) \quad (29)$$

Telescoping the difference $\sqrt{\mathcal{N}} (u^{(\mathcal{N})}(\varphi^*, p^*) - u_0)$ about $u(\varphi^*, p^*)$ and taking a first order Taylor expansion gives

$$\sqrt{\mathcal{N}} \left(u^{(\mathcal{N})}(\varphi^*, p^*) - u_0 \right) = \sqrt{\mathcal{N}} \left[u^{(\mathcal{N})}(\varphi^*, p^*) - u(\varphi^*, p^*) \right] \quad (30)$$

$$-u_\varphi \sqrt{\mathcal{N}} (\varphi^* - \varphi_0) - u_p \sqrt{\mathcal{N}} (p^* - p_0) + o_p(1) \quad (31)$$

Using (27), (28) and (30) we substitute out $\sqrt{\mathcal{N}} (p^* - p_0)$ and $\sqrt{\mathcal{N}} (u^{(\mathcal{N})}(\varphi^*, p^*) - u_0)$ to solve for $\sqrt{\mathcal{N}} (\varphi^* - \varphi_0)$ in terms of $\sqrt{\mathcal{N}} [u^{(\mathcal{N})}(\varphi^*, p^*) - u(\varphi^*, p^*)]$ and $N^{-1/2} \sum_{n=1}^{\mathcal{N}} S_n(u_0, \varphi_0, p_0)$, to obtain

$$\sqrt{\mathcal{N}} (\varphi^* - \varphi_0) = (B'_1 B_1)^{-1} B'_1 N^{-1/2} \sum_{n=1}^{\mathcal{N}} \{ S_n(u_0, \varphi_0, p_0) + B_2 [U_n(\varphi_0, p_0) - u(\varphi_0, p_0)] \}$$

where

$$B_1 \equiv (A_\varphi - A_u u_\varphi) + (A_p - A_u u_p) [(\mathcal{L}_u u_p + I - \mathcal{L}_p)' (\mathcal{L}_u u_p + I - \mathcal{L}_p)]^{-1} (\mathcal{L}_u u_p + I - \mathcal{L}_p)' (\mathcal{L}_\varphi - \mathcal{L}_u u_\varphi)$$

$$B_2 \equiv \left\{ (A_u u_p - A_p) [(\mathcal{L}_u u_p + I - \mathcal{L}_p)' (\mathcal{L}_u u_p + I - \mathcal{L}_p)]^{-1} (\mathcal{L}_u u_p + I - \mathcal{L}_p)' \mathcal{L}_u - A_u \right\}$$

Appealing to the central limit theorem, the asymptotic covariance matrix for $\sqrt{N}(\varphi^* - \varphi_0)$ is thus

$$(B_1' B_1)^{-1} B_1' \Omega B_1 (B_1 B_1')^{-1}$$

where

$$\Omega \equiv E \left[\{S_n(u_0, \varphi_0, p_0) + B_2 U_n(\varphi_0, p_0)\} \{S_n(u_0, \varphi_0, p_0) + B_2 U_n(\varphi_0, p_0)\}' \right] - B_2 u(\varphi_0, p_0) u(\varphi_0, p_0)'$$

is the covariance matrix of

$$S_n(u_0, \varphi_0, p_0) + B_2 [U_n(\varphi_0, p_0) - u(\varphi_0, p_0)]$$

The base case is derived from this formula by setting $U_n(\varphi_0, p_0) \equiv u(\varphi_0, p_0)$ using the fact that identification implies $(I - \mathcal{L}_p)$ is invertible, and that the expected value of the outer product of S_n is just A_φ . In this case the covariance matrix reduces to

$$(B_3' B_3)^{-1} B_3' A_\varphi B_3 (B_3 B_3')^{-1}$$

where

$$B_3 \equiv A_\varphi + A_p (I - \mathcal{L}_p)^{-1} \mathcal{L}_\varphi$$

To cover the case where weighted empirical likelihood are used to estimate conditional choice probabilities, let $S_n(\varphi, p) \equiv S_n(u(\varphi), \varphi, p)$ and define

$$h(d_n, x_n, \varphi, p) \equiv \begin{bmatrix} S_n(\varphi, p) \\ [q(d_n, x_n, \varphi, p) C p - q^{(n,t)}(d_n, x_n, \varphi, p)] \end{bmatrix}$$

Defining the expected outer product and derivative matrices as

$$\begin{aligned} \Omega &= E \left[h(d_n, x_n, \varphi_0, p_0) h(d_n, x_n, \varphi_0, p_0)' \right] \\ \Gamma &= E \left[\frac{\partial h(d_n, x_n, \varphi_0, p_0)}{\partial \varphi} \frac{\partial h(d_n, x_n, \varphi_0, p_0)}{\partial p} \right] \end{aligned}$$

it follows from Hansen (1982, Theorem 3.1) or Newey and McFadden (1994, Theorem 6.1), that $\sqrt{N}(\varphi^* - \varphi_0)$ is asymptotically normally distributed with mean zero and covariance matrix

$$\Sigma = [I_\rho : 0_{JSX}] \Gamma^{-1} \Omega \Gamma^{-1'} [I_\rho : 0_{JSX}]'$$

where I_ρ is the ρ dimensional identity matrix, ρ is the dimension of φ , and 0_{JSX} is a $\rho \times JSX$ matrix of zeros.

References

- [1] Altug, S., and R. A. Miller (1998): “The Effect of Work Experience on Female Wages and Labour Supply,” *Review of Economic Studies*, 62, 45-85.
- [2] Aguirregabiria, V., and P. Mira (2007): “Sequential Estimation of Dynamic Discrete Games,” *Econometrica*, 75, 1-54.
- [3] Aguirregabiria, V., and P. Mira (2007): “Dynamic Discrete Choice Structural Models: A Survey,” working paper.
- [4] Arcidiacono, P. (2005): “Affirmative Action in Higher Education: How do Admission and Financial Aid Rules Affect Future Earnings?” *Econometrica*, 73, 1477-1524.
- [5] Arcidiacono, P., and J. B. Jones (2003): “Finite Mixture Distributions, Sequential Likelihood, and the EM Algorithm,” *Econometrica*, 71, 933-946.
- [6] Arcidiacono, P., Kinsler, J., and R.A. Miller (2008): “How experimenting with drugs affects schooling achievement and early work experience,” working paper.
- [7] Bajari, P., L. Benkard, and J. Levin (2007): “Estimating Dynamic Models of Imperfect Competition,” *Econometrica*, 75, 1331-1371.
- [8] Becker, G., and K. M. Murphy (1988): “A Theory of Rational Addiction,” *Journal of Political Economy*, 96, 675-700.
- [9] Benkard, L., G. Weintraub and B. Van Roy (2007): “Computational Methods for Oblivious Equilibrium,” working paper.
- [10] Beresteanu, A. and P. Ellickson (2006): “The Dynamics of Retail Oligopoly,” working paper.
- [11] Bishop, K.C. (2007): “A Dynamic Model of Location Choice and Hedonic Valuation,” working paper.
- [12] Bresnahan, T. F., S. Stern, and M. Trajtenberg (1997): “Market Segmentation and the Sources of Rents from Innovation: Personal Computers in the Late 1980s,” *RAND Journal of Economics*, 28, 17-44.
- [13] Buchinsky, M., J. Hahn, and V. J. Hotz (2005): “Cluster Analysis: A Tool for Preliminary Structural Analysis,” working paper.

- [14] Collard-Wexler, A. (2006): "Demand Fluctuations and Plant Turnover in the Ready-Mix Concrete Industry," working paper.
- [15] Dempster, A. P., N. M. Laird, and D. B. Rubin (1977): "Maximum Likelihood from Incomplete Data via the EM Algorithm," *Journal of Royal Statistical Society Series B*, 39,1-38
- [16] Dunne, T., Klimek S., Roberts, M. and Y. Xu (2006): "Entry and Exit in Geographic Markets," working paper.
- [17] Ericson, R., and A. Pakes (1995): "Markov-Perfect Industry Dynamics: A Framework for Empirical Work," *Review of Economic Studies*, 62, 53-82.
- [18] Finger, S.R. (2007): "Research and Development Competition in the Chemicals Industry," working paper.
- [19] Gayle, G., and L. Golan (2007): "Estimating a Dynamic Adverse Selection Model: Labor Force Experience and the Changing Gender Earnings Gap 1968-1993," working paper.
- [20] Gayle, G., and R. A. Miller (2006): "Life Cycle Fertility and Human Capital Accumulation," working paper.
- [21] Hansen, L. P. (1982): "Large Sample Properties of Generalized Methods of Moments Estimators," *Econometrica*, 50, 1029-1054.
- [22] Hamilton, J. D. (1990): "Analysis of Time Series Subject to Changes in Regime," *Journal of Econometrics*, 45, 39-70.
- [23] Hotz, V. J., and R. A. Miller (1993): "Conditional Choice Probabilities and Estimation of Dynamic Models," *Review of Economic Studies*, 61, 265-289.
- [24] Hotz, V. J., R. A. Miller, S. Sanders and J. Smith (1994): "A Simulation Estimator for Dynamic Models of Discrete Choice," *Review of Economic Studies*, 60, 265-289.
- [25] Joensen, J.S. (2007): "Academic and Labor Market Success: The Impact of Student Employment, Abilities, and Preferences," working paper.
- [26] Judd, K. L., and C. Su (2007): "Constrained Optimization Approaches to Estimation of Structural Models," working paper.

- [27] Kasahara, H., and K. Shimotsu (2007a): “Nonparametric Identification and Estimation of Finite Mixture Models of Dynamic Discrete Choices,” working paper.
- [28] Kasahara, H. and K. Shimotsu (2007b): “Nested Pseudo-likelihood Estimation and Bootstrap-based Inference for Structural Discrete Markov Decision Models,” working paper.
- [29] Keane M., and K. Wolpin (1994): “The Solution and Estimation of Discrete Choice Dynamic Programming Models by Simulation and Interpolation: Monte Carlo Evidence” *The Review of Economics and Statistics*, 76, 648-672.
- [30] Keane M., and K. Wolpin (1997): “The Career Decisions of Young Men,” *Journal of Political Economy*, 105, 473-522.
- [31] Kennan, J. and J.R. Walker (2006): “The Effect of Expected Income on Individual Migration Decisions,” working paper.
- [32] McFadden, D. (1978): “Modelling the Choice of Residential Location,” *Spatial Interaction and Planning Models*, editors F. Snickars and J. Weibull, North Holland Amsterdam, 75-96.
- [33] Miller, R. A. (1984): “Job Matching and Occupational Choice,” *Journal of Political Economy*, 92, 1086-1020.
- [34] Newey, W. K., and D. McFadden (1994): “Large Sample Estimation and Hypothesis Testing,” Chapter 36 in *Handbook of Econometrics*, 4, editors R. Engle and D. McFadden, Elsevier Science, Amsterdam, 2111-2245.
- [35] Pakes, A. (1986): “Patents as Options: Some Estimates of the Value of Holding European Patent Stocks,” *Econometrica*, 54, 755-784.
- [36] Pakes, A., M. Ostrovsky, and S. Berry (2004): “Simple Estimators for the Parameters of Discrete Dynamic Games (with Entry/Exit Examples),” working paper.
- [37] Pakes, A. and P. McGuire (1994): “Computing Markov Perfect Nash Equilibria: Numerical Implications of a Dynamic Differentiated Product Model” *RAND Journal of Economics*, 25, 555-589.
- [38] Pesendorfer, M., and P. Schmidt-Dengler (2007): “Asymptotic Least Square Estimators for Dynamic Games,” working paper.

- [39] Rust, J. (1987): “Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher,” *Econometrica*, 55, 999-1033.
- [40] Ryan, S. (2006): “The Costs of Environmental Regulation in a Concentrated Industry,” working paper.
- [41] Wolpin, K. (1984): “An Estimable Dynamic Stochastic Model of Fertility and Child Mortality,” *Journal of Political Economy*, 92, 852-874.
- [42] Wu, C. F. (1983): “On the Convergence Properties of the EM Algorithm,” *The Annals of Statistics*, 11, 95-103.

Table 1: School and Drug Choice Monte Carlo[†]

	School	Withdrawal	Low Drug	High Drug	Nesting	Discount	Prob. of
Intercept	Cost	Type	Type	Type	Parameter	Factor	High Type
α_0	α_3	α_1	α_2	δ	β	π	
True Parameters	-0.5	-0.4	-1.0	1.5	0.7	0.9	0.7
ML Estimates	-0.502	-0.400	-1.007	1.502	0.701	0.901	0.700
Standard Error	0.046	0.038	0.100	0.054	0.027	0.021	0.014
CCP Estimates 1 [‡]	-0.502	-0.400	-1.007	1.503	0.702	0.900	0.700
Standard Error	0.049	0.038	0.112	0.055	0.028	0.022	0.015
CCP Estimates 2	-0.517	-0.405	-1.013	1.537	0.718	0.893	0.699
Standard Error	0.059	0.105	0.196	0.078	0.058	0.022	0.018

[†]Listed values are the means and standard errors of the parameter estimates over the 100 simulations.

[‡]CCP Estimates 1 refers to updating the CCP's via the likelihoods while CCP Estimates 2 updates the CCP's using the data directly.

Table 2: Dynamic Entry/Exit Simulations[†]

	True Values	Method 1 [§]			Method 2			Method 1-Low Exit		
		Estimates	Std. Error	Estimates	Std. Error	Estimates	Std. Error	True Values	Estimates	Std. Error
Intercept L	γ_1	6.999	0.042	6.999	0.043	6.998	0.055	7.000	6.998	0.055
Intercept H	γ_2	8.005	0.066	8.004	0.067	8.005	0.080	8.000	8.005	0.080
Equation	1 Firm	-0.299	0.035	-0.304	0.036	-0.302	0.051	-0.300	-0.302	0.051
	2 Firms	-0.702	0.053	-0.702	0.054	-0.698	0.064	-0.700	-0.698	0.064
Profit	Flow Profit L	-0.002	0.032	-0.002	0.040	0.501	0.034	0.500	0.501	0.034
	Flow Profit H	0.516	0.103	0.496	0.111	0.988	0.084	1.000	0.988	0.084
Function	Duopoly Cost	-1.009	0.043	-1.002	0.055	-0.995	0.031	-1.000	-0.995	0.031
	Entry Cost	-1.495	0.027	-1.495	0.029	-3.509	0.053	-3.500	-3.509	0.053
Unobserved	π_{LL}^\ddagger	0.799	0.028	0.799	0.029	0.799	0.026	0.800	0.799	0.026
	π_{HH}	0.702	0.040	0.705	0.037	0.697	0.047	0.700	0.697	0.047
Heterogeneity	$\pi_L^{\dagger\dagger}$	0.799	0.031	0.799	0.032	0.800	0.036	0.800	0.800	0.036

[†] 100 simulations of 3000 markets for 5 periods. β set at 0.9. The third set of columns shows results when the probability of exiting is low ($<2.5\%$ for monopolists, $<5\%$ for duopolists)

[‡] The probability of a market being in the low state in period t conditional on being in the low state at $t-1$.

^{††} Initial probability of a market being assigned the low state.

[§] Method 1 uses our base algorithm. Method 2 estimates unobserved heterogeneity using two-stage method taking the distribution of unobserved heterogeneity as given when solving the dynamic discrete choice problem.

Table 3: Mixed Discrete/Continuous Choice Simulations[†]

		$x_t \in \{1, 4\}$			$x_t \in \{1, 2, 3, 4\}$	
		True	Coefficient	Std. Error	Coefficient	Std. Error
Demand	α_{31}	6.000	5.992	0.160	6.014	0.242
Parameters	α_{32}	6.700	6.704	0.203	6.727	0.293
Variable	$x_t = 1$ α_{41}	1.100	1.094	0.038	1.103	0.050
Cost	$x_t = 2$ α_{42}	1.200			1.203	0.054
Parameters [§]	$x_t = 3$ α_{43}	1.300			1.303	0.058
	$x_t = 4$ α_{44}	1.500	1.489	0.052	1.501	0.065
Fixed Cost	$\alpha_2 - \alpha_1$	5.000	4.946	0.096	4.985	0.100
Unobserved	π_{LL}^{\ddagger}	0.800	0.841	0.083	0.813	0.069
Heterogeneity	π_{HH}	0.700	0.676	0.136	0.707	0.069
	$\pi_L^{\dagger\dagger}$	0.500	0.544	0.115	0.504	0.074

[†] 100 simulations of 3000 firms for 5 periods. β is set at 0.9. CCP's are updated from the model. We do not report estimates of the γ 's because estimation of these is unaffected by the unobserved variables and are estimated in a first stage.

[‡] The probability of a demand being in the low state in period t conditional on being in the low state at $t - 1$.

^{††} Initial probability of demand being in the low state.

[§] Rather than have α_4 times x_t , we set a different value of α_4 for each value of x_t .