

CCP Estimation of Dynamic Discrete Choice Models with Unobserved Heterogeneity*

Peter Arcidiacono Robert A. Miller
Duke University Carnegie Mellon University

August 13, 2010

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 or follow a Markov chain. By developing a class of problems where difference in future value terms depend on few conditional choice probabilities, we extend the class of dynamic optimization problems where CCP estimators provide a computationally cheap alternative to full solution methods. For this class of problems, we establish partial identification results for cases where the environment is non-stationary and the sample period falls short of the decision-maker's time horizon. Monte Carlo results 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 Victor Aguirregabiria, Esteban Aucejo, Lanier Benkard, Jason Blevins, Paul Ellickson, George-Levi Gayle, Joe Hotz, three anonymous referees and the co-editor for their comments. We have benefited from seminars at UC Berkeley, Duke University, University College London, The Ohio State University, University of Pennsylvania, Stanford University, University of Texas, University of Virginia, Washington University, University of Wisconsin, IZA, and the NASM of the Econometric Society. Andrew Beauchamp, Jon James, and Josh Kinsler 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 full solution methods and have experienced a resurgence in the literature on estimating dynamic games.² The computational gains associated with CCP estimation give researchers considerable latitude to explore different functional forms for their models.

In order to implement CCP estimators, two things are necessary. First, the researcher must know how to formulate the value function—or the differenced value function across two choices—as a function of the conditional choice probabilities. These formulations depend upon the distribution of the structural errors and, except in special cases, forming choice paths far out into the future. Second, CCP estimators require calculating choice probabilities conditional on all state variables. Calculating the conditional choice probabilities can be difficult when some of the state variables are unobserved.

Our first contribution is to broaden the class of models where CCP estimation can be implemented without resorting to matrix inversion or simulation. We prove that the expected value of future utilities from optimal decision making can always be expressed as functions of the flow payoffs and conditional choice probabilities for *any* sequence of future choices, optimal or not. When two choice sequences with different initial decisions lead to the same distribution of states after a few periods, we say there is finite dependence, generalizing Altug and Miller (1998). In such cases the likelihood of a decision can be constructed from current payoffs and conditional choice probabilities that occur a few periods into the future. Finite dependence also facilitates identification. We show cases where the parameters of the utility function are identified even when the sample period falls short of the time horizon and the time horizon is unknown or when the transitions of the state variables are unknown beyond the sample period.

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 (2010) have recently surveyed the literature on estimating dynamic models of discrete choice. For developments of CCP estimators that apply to dynamic games, see Aguirregabiria and Mira (2007), Bajari, Benkard, and Levin (2007), Pakes, Ostrovsky, and Berry (2007), and Pesendorfer and Schmidt-Dengler (2008).

Key to exploiting finite dependence, however, is knowing the mappings between the conditional choice probabilities and the difference in the payoffs between two choices. These mappings depend on the distribution of the structural errors. We show how to obtain the mappings when the structural errors follow any generalized extreme value distribution, substantially broadening the class of error structures that are easily adapted to CCP estimation.

Our second contribution is developing CCP estimators that are capable of handling rich classes of unobserved heterogeneity where there are a finite number of unobserved states.³ Accounting for unobserved heterogeneity, and therefore dynamic selection, is important to many economic problems and has been a standard feature of dynamic discrete choice models in labor economics.⁴ Our estimators can readily be adapted to cases where the unobserved state variables are time-invariant, such as is standard in the dynamic discrete choice literature, as well as to cases where the unobserved states transition over time. We further characterize the large number of moments that can be used to identify the conditional choice probabilities in the presence of unobserved heterogeneity.

To operationalize our estimators, we modify the Expectations-Maximization (EM) algorithm, and in particular its application to sequential estimation as developed in Arcidiacono and Jones (2003), to include updates of the conditional choice probabilities. The EM algorithm iterates on two steps. In the expectations step, the conditional probability of each observation being in each unobserved state is calculated given the data and the structure of the model. In the maximization step, the unobserved states are treated as observed, with the conditional probabilities of being in each unobserved state used as weights. Because the unobserved states are treated as known in the second step of the EM algorithm, we show that there are natural ways of updating the CCP's in the presence of unobserved states. Since the EM algorithm requires solving the maximization step multiple times, it is important that the maximization step be fast. Hence, it is the coupling of CCP estimators—particularly those that exhibit finite dependence—with the EM algorithm which allows for large computational gains despite having to iterate.

Subject to identification issues we discuss in Section 4, we further show how to modify our algorithm to estimate the distribution of unobserved heterogeneity and the conditional choice prob-

³An alternative is to have unobserved continuous variables. Mroz (1999) shows that using finite mixtures when the true model has a persistent unobserved variable with continuous support yields similar estimates to the case when the unobserved variable is treated as continuous in estimation. For Bayesian approaches to this issue, see Imai, Jain, and Ching (2009) and Norets (2009).

⁴For example, see Miller (1984), Keane and Wolpin (1997, 2000, 2001), Eckstein and Wolpin (1999), Arcidiacono (2005), Arcidiacono, Sloan, and Sieg (2007), and Kennan and Walker (forthcoming).

abilities in a first stage. The key insight is to use the empirical distribution of choices—rather than the structural choice probabilities themselves—when updating the conditional probability of being in each unobserved state. 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 estimation. We show how the first stage of this modified algorithm can be paired with non-likelihood based 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 version of the Rust bus engine problem with permanent unobserved heterogeneity. Here we compare computational times and the precision of the estimates with full information maximum likelihood. We further show cases where estimation is only feasible via conditional choice probabilities such as when the time horizon is unknown or when there are time-specific parameters and the data stop short of the full time horizon. The second Monte Carlo is a dynamic game of firm entry and exit. In this example, the unobserved heterogeneity affects the demand levels for particular markets which and, in turn, the value of entering or remaining in the market. The unobserved states are allowed to transition over time and the example explicitly incorporates dynamic selection. We estimate the model using the baseline algorithm as well as the two-stage method. For both sets of Monte Carlos, the estimators perform quite well both in terms of the precision of the estimates as well as computational time.

Our contributions build on some of the points made in the literature on estimating dynamic games. Bajari, Benkard and Levin (2007) build off the approach of Hotz et al. (1994), estimating reduced form policy functions in order to forward simulate the future component of the dynamic discrete choice problem. In principle, their method can be used for any distribution of the structural errors. In practice, this is difficult because the simulations have to be redrawn with each new estimate of the correlation parameters. Aguirregabiria and Mira (2007) show how to incorporate permanent unobserved heterogeneity into stationary dynamic games. Their method requires inverting matrices multiple times where the matrices are dimensioned by the number of states.⁵ Further, their estimator

⁵Kasahara and Shimotsu (2008) propose methods to weaken the computational requirements of Aguirregabiria and Mira (2007), in part by developing a procedure of obtaining non-parametric estimates of the conditional choice probabilities in a first stage. Hu and Shum (2010b) take a similar approach while allowing the unobserved states to

is restricted to the case where the unobserved heterogeneity only affects the payoff functions. This limits the researcher's ability to account for dynamic selection by adopting a selection on observables approach to the transitions of the state variables.

The techniques developed in this paper are being used to estimate structural models in environmental economics, labor economics, industrial organization, and marketing. Bishop (2008) applies the reformulation of the value functions to the migration model of Kennan and Walker (forthcoming) to accommodate state spaces that are computationally intractable using standard techniques. Joensen (2009) incorporates unobserved heterogeneity into a CCP estimator of educational attainment and work decisions. Beresteanu, Ellickson, and Misra (2010) combine our value function reformulation with simulations of the one-period ahead probabilities to estimate a large scale discrete game between retailers. Finally, Chung, Steenburgh, and Sudhir (2009), Beauchamp (2010), and Finger (2008) use our two-stage algorithm to obtain estimates of the unobserved heterogeneity parameters in a first stage, the latter two applying the estimator in a games environment.

The rest of the paper proceeds as follows. Section 2 uses Rust's bus engine problem (1987) as an example of how to apply CCP estimation with unobserved heterogeneity. Section 3 sets up the general framework for our analysis. Section 3 further shows that, for many cases, the differences in conditional value functions only depend upon a small number of conditional choice probabilities and extends the classes of error distributions that can easily be mapped into a CCP framework. Section 4 discusses identification, establishing identification for classes of models that are non-stationary and where the data does not cover the full time horizon. It also shows the moments in the data that can be used to recover conditional choice probabilities in the presence of unobserved heterogeneity. Section 5 develops the estimators, while Section 6 derives the algorithms used to operationalize them. Section 7 shows how the parameters governing the unobserved heterogeneity can sometimes be estimated in a first stage. Section 8 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 9 concludes. All proofs are in the appendix.

transition over time. Buchinsky et al (2005) use the tools of cluster analysis to incorporate permanent unobserved heterogeneity, 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.

2 Motivating Example

To motivate our approach, we first show how the tools developed in this paper apply to the bus engine example considered by Rust (1987) when unobserved heterogeneity is present. This example highlights several features of the paper. First, we show how to characterize the future value term— or more precisely the difference in future value terms across the two choices— as a function of just the one period ahead probability of replacing the engine. Second, we show how to estimate the model when there is time-invariant unobserved heterogeneity. In later sections, we extend the estimation to include more general forms of unobserved heterogeneity as well as showing generally how to characterize differences in future value terms as functions of only the conditional choice probabilities from a few periods ahead.

2.1 Set up

In each period $t \leq \infty$ Harold Zurcher decides whether to replace the existing engine of a bus by choosing $d_{1t} = 1$, or keep it for at least one more period by choosing $d_{2t} = 1$, where $d_{1t} + d_{2t} = 1$. The current period payoff for action j depends upon how much mileage the bus has accumulated since the last replacement, $x_t \in \{1, 2, \dots\}$, and the brand of the bus, $s \in \{1, \dots, S\}$. It is through s that we bring unobserved heterogeneity into the bus replacement problem; both x_t and s are observed by Zurcher but the econometrician only observes x_t .

Mileage advances one unit if Zurcher keeps the current engine and is set to zero if the engine is replaced. Thus $x_{t+1} = x_t + 1$ if $d_{2t} = 1$ and $x_{t+1} = 0$ if $d_{1t} = 1$. There is a choice-specific transitory shock, ϵ_{jt} , that also affects current period payoffs and is independently distributed Type 1 extreme value. The current period payoff for keeping the engine at time t is given by $\theta_1 x_t + \theta_2 s + \epsilon_{2t}$, where $\theta \equiv \{\theta_1, \theta_2\}$ is a set of parameters to be estimated. Since decisions in discrete choice models are unaffected by increasing the payoff to all choices by the same amount, we normalize the current period payoff of the first choice to ϵ_{1t} . This normalization implies $\theta_1 x_t + \theta_2 s + \epsilon_{2t} - \epsilon_{1t}$ measures, for a brand s bus in period t , the cost of maintaining an old bus engine for another period, net of expenditures incurred by purchasing, installing and maintaining a new bus engine.

Zurcher takes into account both the current period payoff, as well as how his decision today will affect the future, with the per-period discount factor given by β . He chooses d_{1t} (and therefore d_{2t}) to sequentially maximize the expected discounted sum of payoffs:

$$E \left\{ \sum_{t=1}^{\infty} \beta^{t-1} [d_{2t}(\theta_1 x_t + \theta_2 s + \epsilon_{2t}) + d_{1t} \epsilon_{1t}] \right\} \quad (2.1)$$

Let $V(x_t, s)$ denote the ex-ante value function at the beginning of period t .⁶ It is the discounted sum of current and future payoffs just before $\epsilon_t \equiv \{\epsilon_{1t}, \epsilon_{2t}\}$ is realized and before the decision at t is made, conditional on making optimal choices at t and every future period, when the bus is brand s , and the mileage is x_t . We also define the conditional value function for choice j as the current period payoff of choice j net of ϵ_{jt} plus the expected future utility from Zurcher behaving optimally in the future:

$$v_j(x, s) \equiv \begin{cases} \beta V(0, s) & \text{if } j = 1 \\ \theta_1 x + \theta_2 s + \beta V(x + 1, s) & \text{if } j = 2 \end{cases} \quad (2.2)$$

Let $p_1(x, s)$ denote the conditional choice probability (CCP) of replacing the engine given x and s . The parametric assumptions about the transitory cost shocks imply:

$$p_1(x, s) = \frac{1}{1 + \exp [v_2(x, s) - v_1(x, s)]} \quad (2.3)$$

2.2 CCP representation of the replacement problem

Rust (1987) showed the conditional value function for keeping the engine, defined in the second line of equation (2.2), can be expressed as:

$$v_2(x, s) = \theta_1 x + \theta_2 s + \beta \ln \left[\exp [v_1(x + 1, s)] + \exp [v_2(x + 1, s)] \right] + \beta \gamma \quad (2.4)$$

where γ is Euler's constant. Multiplying and dividing the expression inside the logarithm of equation (2.4) by $\exp [v_1(x + 1, s)]$, yields:

$$\begin{aligned} v_2(x, s) &= \theta_1 x + \theta_2 s + \beta \ln [\exp [v_1(x + 1, s)] \{1 + \exp [v_2(x + 1, s) - v_1(x + 1, s)]\}] + \beta \gamma \\ &= \theta_1 x + \theta_2 s + \beta v_1(x + 1, s) - \beta \ln [p_1(x + 1, s)] + \beta \gamma \end{aligned} \quad (2.5)$$

where the last line follows from equation (2.3). Equation (2.5) shows that the future value term in the replacement problem can be expressed as the conditional value of replacing at mileage $x + 1$ plus the probability of replacing the engine when the mileage is $x + 1$. Applying the same logic to the conditional value function for engine replacement yields:

$$v_1(x, s) = \beta v_1(0, s) - \beta \ln [p_1(0, s)] + \beta \gamma \quad (2.6)$$

Recall that replacing the engine resets the mileage to zero. Equation (2.2) then implies that

$$v_1(x + 1, s) = v_1(0, s) = \beta V(0, s)$$

⁶Since the optimal decision rule is stationary, subscripting by t is redundant.

for all x . Exploiting this property, we difference equations (2.5) and (2.6) to obtain:⁷

$$v_2(x, s) - v_1(x, s) = \theta_1 x + \theta_2 s + \beta \ln [p_1(0, s)] - \beta \ln [p_1(x + 1, s)] \quad (2.7)$$

Substituting equation (2.7) into equation (2.3) implies that the probability of replacing the engine, $p_1(x, s)$, can be expressed a function of the flow payoff of running the engine, $\theta_1 x + \theta_2 s$, the discount factor, β , and the one-period-ahead probabilities of replacing the engine, $p_1(0, s)$, and $p_1(x + 1, s)$.

2.3 Estimation

To estimate the model, we develop an algorithm that combines key insights from two literatures. The first is the literature on CCP estimation when there is no unobserved heterogeneity. In this literature, estimates of the conditional choice probabilities are obtained in a first stage and substituted into a second stage maximization. The second is the literature on the Expectation-Maximization (EM) algorithm, which provides a way of estimating dynamic discrete choice models when unobserved state variables are present. As we will show, the EM algorithm can be modified to accommodate CCP estimation.

Consider a sample of \mathcal{N} buses over \mathcal{T} time periods where all buses begin with zero mileage. The key insight of Hotz and Miller (1993) is that, when both x and s are observed variables, we can substitute a first stage estimate, $\hat{p}_1(x, s)$, for $p_1(x, s)$ in (2.7), say $\hat{p}_1(x, s)$.⁸ Next, we substitute this expression into equation (2.3) to obtain the likelihood of replacing the engine given the first-stage conditional choice probabilities. Writing \hat{p}_1 as the vector of conditional choice probabilities and $d_{nt} \equiv \{d_{1nt}, d_{2nt}\}$, the likelihood contribution for bus n at time t is then given by:

$$l(d_{nt}|x_{nt}, s_n, \theta, \hat{p}_1) = \frac{d_{1nt} + d_{2nt} \exp(\theta_1 x_{nt} + \theta_2 s_n + \beta \ln [\hat{p}_1(0, s_n)] - \beta \ln [\hat{p}_1(x_{nt} + 1, s_n)])}{1 + \exp(\theta_1 x_{nt} + \theta_2 s_n + \beta \ln [\hat{p}_1(0, s_n)] - \beta \ln [\hat{p}_1(x_{nt} + 1, s_n)])} \quad (2.8)$$

The structural parameters θ_1 and θ_2 can then be estimated in a second stage using a logit.

To illustrate how CCP estimation might be amenable to the EM algorithm, we first demonstrate how to proceed in the infeasible case where s is unobserved but $\hat{p}_1(x, s)$ is known. Let π_s denote the population probability of being in state s . Integrating the unobserved state out of the likelihood

⁷Note the conditional value functions for period $t + 1$ do not cancel if the future value terms are written with respect to the second choice because $v_2(x + 1, s) \neq v_2(0, s)$.

⁸One candidate is a bin estimator where $\hat{p}_1(x, s)$ is given by

$$\hat{p}_1(x, s) = \frac{\sum_n \sum_t I(d_{1nt} = 1)I(x_{nt} = x)I(s_n = s)}{\sum_n \sum_t I(x_{nt} = x)I(s_n = s)}$$

function, the maximum likelihood (ML) estimator for this version of the problem is:

$$\{\hat{\theta}, \hat{\pi}\} = \arg \max_{\theta, \pi} \sum_{n=1}^{\mathcal{N}} \ln \left[\sum_{s=1}^S \pi_s \prod_{t=1}^{\mathcal{T}} l(d_{nt}|x_{nt}, s, \hat{p}_1, \theta) \right] \quad (2.9)$$

Denote by $d_n \equiv (d_{n1}, \dots, d_{n\mathcal{T}})$ and $x_n \equiv (x_{n1}, \dots, x_{n\mathcal{T}})$ the full sequence of choices and mileages observed in the data for bus n . Conditioning on x_n , \hat{p}_1 , and θ , the probability of observing d_n is the expression inside the logarithm in (2.9), while the joint probability of s and d_n is the product of all the terms to the right of the summation over s . Given the ML estimates $(\hat{\theta}, \hat{\pi})$ and using Bayes' rule, we can calculate \hat{q}_{ns} , the probability n is in unobserved state s as:

$$\hat{q}_{ns} = Pr\{s_n = s | d_n, x_n; \hat{\theta}, \hat{\pi}, \hat{p}_1\} = \frac{\hat{\pi}_s \prod_{t=1}^{\mathcal{T}} l(d_{nt}|x_{nt}, s, \hat{p}_1, \hat{\theta})}{\sum_{s'=1}^S \hat{\pi}_{s'} \prod_{t=1}^{\mathcal{T}} l(d_{nt}|x_{nt}, s', \hat{p}_1, \hat{\theta})} \quad (2.10)$$

By definition, $\hat{\pi}_s$ then satisfies:

$$\hat{\pi}_s = \frac{\sum_{n=1}^{\mathcal{N}} \hat{q}_{ns}}{\mathcal{N}} \quad (2.11)$$

The EM algorithm is a computationally attractive alternative to directly maximizing (2.9). At the m^{th} iteration, given values for $\theta^{(m)}$ and $\pi^{(m)}$, update $q_{ns}^{(m+1)}$ by substituting $\theta^{(m)}$ and $\pi^{(m)}$ in for $\hat{\theta}$ and $\hat{\pi}$ in equation (2.10). Next, update $\pi^{(m+1)}$ by substituting $q_{ns}^{(m+1)}$ for \hat{q}_{ns} in equation (2.11). Finally, obtain $\theta^{(m+1)}$ from:

$$\theta^{(m+1)} = \arg \max_{\theta} \sum_{n=1}^{\mathcal{N}} \sum_{s=1}^S \sum_{t=1}^{\mathcal{T}} q_{ns}^{(m+1)} \ln[l(d_{nt}|x_{nt}, s, \hat{p}_1, \theta)] \quad (2.12)$$

Note that at the maximization step $q_{ns}^{(m+1)}$ is taken as given, and that the maximization problem is equivalent to one where s_n is observed and $q_{ns}^{(m+1)}$ are population weights. Under standard regularity conditions $(\theta^{(m)}, \pi^{(m)})$ converges to the ML estimator $(\hat{\theta}, \hat{\pi})$.⁹

We now show how to estimate the structural parameters when both s is unobserved and $\hat{p}(x, s)$ is unknown by building on the features of the estimators discussed above. We modify the EM algorithm, so that instead of just updating $\theta^{(m)}$, $\pi^{(m)}$ and $q_{ns}^{(m)}$ at the m^{th} iteration, we also update the conditional choice probabilities, $p_1^{(m)}(x, s)$.

One way of updating $p_1^{(m)}(x_{nt}, s_n)$ falls naturally out of the EM algorithm. To see this, first note that we can express $p_1(x, s)$ as:

$$p_1(x, s) \equiv Pr\{d_{1nt} = 1 | s_n = s, x_{nt} = x\} = \frac{Pr\{d_{1nt} = 1, s_n = s | x_{nt} = x\}}{Pr\{s_n = s | x_{nt} = x\}} = \frac{E[d_{1nt} I(s_n = s) | x_{nt} = x]}{E[I(s_n = s) | x_{nt} = x]} \quad (2.13)$$

⁹See Dempster, Laird, and Rubin (1977) and Wu (1983) for details.

Applying the law of iterated expectations to both the numerator and the denominator, and using the fact that d_{nt} is a component of d_n , implies (2.13) can be written as:

$$p_1(x, s) = \frac{E[d_{1nt}E\{I(s_n = s)|d_n, x_n\}|x_{nt} = x]}{E[E\{I(s_n = s)|d_n, x_n\}|x_{nt} = x]} \quad (2.14)$$

But the inner expectation in (2.14) is actually q_{ns} as:

$$q_{ns} = Pr\{s_n = s|d_n, x_n\} = E[I(s_n = s)|d_n, x_n] \quad (2.15)$$

It now follows that:

$$p_1(x, s) = \frac{E[d_{1nt}q_{ns}|x_{nt} = x]}{E[q_{ns}|x_{nt} = x]} \quad (2.16)$$

In the algorithm defined below we replace (2.16) with sample analogs in order to update the conditional choice probabilities.

Our algorithm begins by setting initial values for $\theta^{(1)}$, $\pi^{(1)}$, and $p_1^{(1)}$. Estimation then involves iterating on four steps where the m^{th} iteration follows:

Step 1 From (2.10), compute $q_{ns}^{(m+1)}$ as:

$$q_{ns}^{(m+1)} = \frac{\pi_s^{(m)} \prod_{t=1}^{\mathcal{T}} l(d_{nt}|x_{nt}, s, p_1^{(m)}, \theta^{(m)})}{\sum_{s'=1}^S \pi_{s'}^{(m)} \prod_{t=1}^{\mathcal{T}} l(d_{nt}|x_{nt}, s', p_1^{(m)}, \theta^{(m)})} \quad (2.17)$$

Step 2 Using $q_{ns}^{(m+1)}$ compute $\pi_s^{(m+1)}$ according to:

$$\pi_s^{(m+1)} = \frac{\sum_{n=1}^{\mathcal{N}} q_{ns}^{(m+1)}}{\mathcal{N}} \quad (2.18)$$

Step 3 Using $q_{ns}^{(m+1)}$ update $p_1^{(m+1)}(x, s)$ from:

$$p_1^{(m+1)}(x, s) = \frac{\sum_{n=1}^{\mathcal{N}} \sum_{t=1}^{\mathcal{T}} d_{1nt} q_{ns}^{(m+1)} I(x_{nt} = x)}{\sum_{n=1}^{\mathcal{N}} \sum_{t=1}^{\mathcal{T}} q_{ns}^{(m+1)} I(x_{nt} = x)} \quad (2.19)$$

Step 4 Taking $q_{ns}^{(m+1)}$ and $p_1^{(m+1)}(x_{nt}, s_n)$ as given, obtain $\theta^{(m+1)}$ from:

$$\theta^{(m+1)} = \arg \max_{\theta} \sum_{n=1}^{\mathcal{N}} \sum_{s=1}^S \sum_{t=1}^{\mathcal{T}} q_{ns}^{(m+1)} \ln \left[l(d_{nt}|x_{nt}, s_n, p_1^{(m+1)}, \theta) \right] \quad (2.20)$$

Note that Step 3 is a weighted average of decisions to replace conditional on x , where the weights are the conditional probabilities of being in unobserved state s . When s is observed, Step 3 collapses to a bin estimator that could be used in the first stage of CCP estimation.

An alternative to updating the CCP's using a weighted average of the data is based on the identity that the likelihood returns the probability of replacing the engine. Substituting equation (2.7) into equation (2.3) and evaluating it at the relevant values for bus n at time t yields:

$$p_1(x_{nt}, s) = l(d_{1nt} = 1 | x_{nt}, s, p_1, \theta) \quad (2.21)$$

Thus at the m^{th} iteration we could replace Step 3 of the algorithm with:

Step 3A Using $\theta^{(m)}$ and $p_1^{(m)}$, that is the function $p_1^{(m)}(x, s)$, update $p_1^{(m+1)}(x_{nt}, s_n)$ using:

$$p_1^{(m+1)}(x_{nt}, s) = l(d_{1nt} = 1 | x_{nt}, s, p_1^{(m)}, \theta^{(m)}) \quad (2.22)$$

Here, $l(d_{1nt} = 1 | x_{nt}, s, p_1^{(m)}, \theta^{(m)})$ is calculated using (2.8). The CCP updates are tied directly to the structure of the model. We illustrate the tradeoffs of each updating method in the ensuing sections.

3 Framework

This section lays out a general class of dynamic discrete choice models and derives a new representation of the conditional valuation functions that we draw upon in the subsequent sections on identification and estimation. In this section we also use the representation to develop the concept of finite dependence, and determine its functional form for disturbances distributed as generalized extreme value.

3.1 Model

In each period until T , for $T \leq \infty$, an individual chooses among J mutually exclusive actions. Let d_{jt} equal one if action $j \in \{1, \dots, J\}$ is taken at time t and zero otherwise. The current period payoff for action j at time t depends on the state $z_t \in \{1, \dots, Z\}$. In the previous section $z_t \equiv (x_t, s)$ where x_t is observed but s is unobserved to the econometrician. We ignore the distinction in this section because it is not relevant for the agents in the model. If action j is taken at time t , the probability of z_{t+1} occurring in period $t+1$ is denoted by $f_{jt}(z_{t+1} | z_t)$.

The individual's current period payoff from choosing j at time t is also affected by a choice-specific shock, ϵ_{jt} , which is revealed to the individual at the beginning of the period t . We assume the vector $\epsilon_t \equiv (\epsilon_{1t}, \dots, \epsilon_{Jt})$ has continuous support and is drawn from a probability distribution that is independently and identically distributed over time with density function $g(\epsilon_t)$. We model the individual's current period payoff for action j at time t by $u_{jt}(z_t) + \epsilon_{jt}$.

The individual takes into account both the current period payoff as well as how his decision today will affect the future. Denoting the discount factor by $\beta \in (0, 1)$, the individual chooses the vector $d_t \equiv (d_{1t}, \dots, d_{Jt})$ to sequentially maximize the discounted sum of payoffs:

$$E \left\{ \sum_{t=1}^T \sum_{j=1}^J \beta^{t-1} d_{jt} [u_{jt}(z_t) + \epsilon_{jt}] \right\} \quad (3.1)$$

where at each period t the expectation is taken over the future values of z_{t+1}, \dots, z_T and $\epsilon_{t+1}, \dots, \epsilon_T$. Expression (3.1) is maximized by a Markov decision rule which gives the optimal action conditional on t , z_t , and ϵ_t . We denote the optimal decision rule at t as $d_t^o(z_t, \epsilon_t)$, with j th element $d_{jt}^o(z_t, \epsilon_t)$. The probability of choosing j at time t conditional on z_t , $p_{jt}(z_t)$, is found by taking $d_{jt}^o(z_t, \epsilon_t)$ and integrating over ϵ_t :

$$p_{jt}(z_t) \equiv \int d_{jt}^o(z_t, \epsilon_t) g(\epsilon_t) d\epsilon_t \quad (3.2)$$

We then define $p_t(z_t) \equiv (p_{1t}(z_t), \dots, p_{Jt}(z_t))$ as the vector of conditional choice probabilities.

Denote $V_t(z_t)$, the (ex-ante) value function in period t , as the discounted sum of expected future payoffs just before ϵ_t is revealed and conditional on behaving according to the optimal decision rule:

$$V_t(z_t) \equiv E \left\{ \sum_{\tau=t}^T \sum_{j=1}^J \beta^{\tau-t} d_{j\tau}^o(z_\tau, \epsilon_\tau) (u_{j\tau}(z_\tau) + \epsilon_{j\tau}) \right\}$$

Given state variables z_t and choice j in period t , the expected value function in period $t + 1$, discounted one period into the future is:

$$\beta \sum_{z_{t+1}=1}^Z V_{t+1}(z_{t+1}) f_{jt}(z_{t+1}|z_t)$$

Under standard conditions, Bellman's principle applies and $V_t(z_t)$ can be recursively expressed as:

$$\begin{aligned} V_t(z_t) &= E \left\{ \sum_{j=1}^J d_{jt}^o(z_t, \epsilon_t) \left[u_{jt}(z_t) + \epsilon_{jt} + \beta \sum_{z_{t+1}=1}^Z V_{t+1}(z_{t+1}) f_{jt}(z_{t+1}|z_t) \right] \right\} \\ &= \sum_{j=1}^J \int d_{jt}^o(z_t, \epsilon_t) \left[u_{jt}(z_t) + \epsilon_{jt} + \beta \sum_{z_{t+1}=1}^Z V_{t+1}(z_{t+1}) f_{jt}(z_{t+1}|z_t) \right] g(\epsilon_t) d\epsilon_t \quad (3.3) \end{aligned}$$

where the second line integrates out the disturbance vector ϵ_t . We then define the conditional value function, $v_{jt}(z_t)$, as the flow payoff of action j without ϵ_{jt} plus the expected future utility conditional on following the optimal decision rule from period $t + 1$ on:

$$v_{jt}(z_t) = u_{jt}(z_t) + \beta \sum_{z_{t+1}=1}^Z V_{t+1}(z_{t+1}) f_{jt}(z_{t+1}|z_t) \quad (3.4)$$

Hotz and Miller (1993) established that differences in conditional value functions can be expressed as functions of the conditional choice probabilities, $p_{jt}(z_t)$, and the per-period payoffs. Using their results, we show that we can express the value function, $V_t(z_t)$, as a function of one conditional value function, $v_{jt}(z_t)$, plus a function of the conditional choice probabilities, $p_t(z_t)$.

Lemma 1 Define $\boldsymbol{p} \equiv (p_1, \dots, p_J)$ where $\sum_{j=1}^J p_j = 1$ and $p_j > 0$ for all j' . Then there exists a real-valued function $\psi_k(\boldsymbol{p})$ for every $k \in \{1, \dots, J\}$ such that:

$$\psi_k[p_t(z_t)] \equiv V_t(z_t) - v_{kt}(z_t) \quad (3.5)$$

Substituting (3.5) into the right hand side of (3.4) we obtain:

$$v_{jt}(z_t) = u_{jt}(z_t) + \beta \sum_{z_{t+1}=1}^Z [v_{kt+1}(z_{t+1}) + \psi_k[p_{t+1}(z_{t+1})]] f_{jt}(z_{t+1}|z_t) \quad (3.6)$$

Equation (3.6) shows that the conditional value function can be expressed as the flow payoff of the choice plus a function of the one period ahead conditional choice probabilities and the one period ahead conditional value function for any choice. We could repeat this procedure ad infinitum, substituting in for $v_{kt_1}(z_{t+1})$ using (3.4) and then again with (3.5) for any choice k' , ultimately replacing the conditional valuation functions on the right side of (3.6) with a single arbitrary time sequence of current utility terms and conditional value correction terms as defined in (3.5).

To formalize this idea, consider a sequence of decisions from t to T . The first choice in the sequence is the initial choice j which sets $d_{jt}^*(z_t, j) = 1$. For periods $\tau \in \{t+1, \dots, T\}$, the choice sequence maps z_τ and the initial choice j into $d_\tau^*(z_\tau, j) \equiv \{d_{1\tau}^*(z_\tau, j), \dots, d_{J\tau}^*(z_\tau, j)\}$. The choices in the sequence then must satisfy $d_{k\tau}^*(z_\tau, j) \geq 0$ and $\sum_{k=1}^J d_{k\tau}^*(z_\tau, j) = 1$. Note that the choice sequence can depend upon new realizations of the state and may also involve mixing over choices.

Now consider the probability of being in state $z_{\tau+1}$ conditional on following the choices in the sequence. Denote this probability as $\kappa_\tau^*(z_{\tau+1}|z_t, j)$ which is recursively defined by:

$$\kappa_\tau^*(z_{\tau+1}|z_t, j) \equiv \begin{cases} f_{jt}(z_{t+1}|z_t) & \text{for } \tau = t \\ \sum_{z_\tau=1}^Z \sum_{k=1}^J d_{k\tau}^*(z_\tau, j) f_{k\tau}(z_{\tau+1}|z_\tau) \kappa_{\tau-1}^*(z_\tau|z_t, j) & \text{for } \tau = t+1, \dots, T \end{cases} \quad (3.7)$$

The future value term can now be expressed relative to the conditional value functions for the choices in the sequence. Theorem 1 shows that continuing to express the future value term relative to the value of the next choice in the sequence yields an alternative expression for $v_{jt}(z_t)$.

Theorem 1 For any state $z_t \in \{1, \dots, Z\}$, choice $j \in \{1, \dots, J\}$ and decision rule $d_\tau^*(z_\tau, j)$ defined for periods $\tau \in \{t, \dots, T\}$:

$$v_{jt}(z_t) = u_{jt}(z_t) + \sum_{\tau=t+1}^T \sum_{k=1}^J \sum_{z_\tau=1}^Z \beta^{\tau-t} [u_{k\tau}(z_\tau) + \psi_k[p_\tau(z_\tau)]] d_{k\tau}^*(z_\tau, j) \kappa_{\tau-1}^*(z_\tau | z_t, j) \quad (3.8)$$

Theorem 1 shows that future value terms for dynamic discrete choice models can be expressed as functions of flow payoffs and conditional choice probabilities for any sequences of choices out until T and the corresponding transition probabilities associated with the choice sequences. It provides the foundation for the identification results discussed in Section 4 and the estimators developed in Section 5. In this section we use the theorem for deriving conditions to construct estimators that do not depend utility flow terms $u_{k\tau}(z_\tau) + \psi_k[p_\tau(z_\tau)]$ beyond a few periods for each t . Then we show how the functions $\psi_k(p)$ are determined in the generalized extreme value case.

3.2 Finite Dependence

Equation (3.4) implies that j is preferred to k in period t if and only if $v_{jt}(z_t) - v_{kt}(z_t) > \epsilon_{jt} - \epsilon_{kt}$. Consequently conditional valuation functions, such as $v_{jt}(z_t)$ and $v_{kt}(z_t)$, only enter the likelihood function in their differenced form. Substituting (3.8) from Theorem 1 into expressions like $v_{jt}(z_t) - v_{kt}(z_t)$, reveals that all the terms in the sequence after a certain date, say ρ , would cancel out if the state variables had the same probability distribution at ρ , that is if $\kappa_{\rho-1}^*(z_\rho | z_t, j) = \kappa_{\rho-1}^*(z_\rho | z_t, k)$, and the same decisions are selected for all dates beyond ρ .

In the example from Section 2, $z_t \equiv (x_t, s)$, and the act of replacing an engine next period regardless of the choice made in the current period t , thus setting $d_{1,t+1}^*(z_{t+1}, 1) = d_{1,t+1}^*(z_{t+1}, 2) = 1$, restores the state variables to the same value $(x_{t+1}, s) = (0, s)$ at period $t + 1$. Thus any (common) sequence of choices that begins by replacing the engine next period, implies that when considering the difference $v_{2t}(z_t) - v_{1t}(z_t)$ in their telescoped forms using (3.8), all terms beyond the next period disappear.

Exploiting the power of Theorem 1 in this way can be developed within the general framework. Consider using (3.8) to express the conditional value functions for alternative initial choices j and j' . Differencing the two yields:

$$\begin{aligned} v_{jt}(z_t) - v_{j't}(z_t) &= u_{jt}(z_t) - u_{j't}(z_t) \\ &+ \sum_{\tau=t+1}^T \sum_{k=1}^J \sum_{z_\tau=1}^Z \beta^{\tau-t} [u_{k\tau}(z_\tau) + \psi_k[p_\tau(z_\tau)]] [d_{k\tau}^*(z_\tau, j) \kappa_{\tau-1}^*(z_\tau | z_t, j) - d_{k\tau}^*(z_\tau, j') \kappa_{\tau-1}^*(z_\tau | z_t, j')] \end{aligned} \quad (3.9)$$

We say a pair of choices exhibit ρ - **period dependence** if there exists a sequence from initial choice j and a corresponding sequence from initial choice j' such that for all $z_{t+\rho+1}$:

$$\kappa_{t+\rho}^*(z_{t+\rho+1}|z_t, j) = \kappa_{t+\rho}^*(z_{t+\rho+1}|z_t, j') \quad (3.10)$$

The sequence of choices from j and j' then lead to the same state in expectation. When ρ -period dependence holds, the difference in future value terms for j and j' can be expressed as a function of the ρ -period ahead flow payoffs, conditional choice probabilities, and state transition probabilities. Once ρ -period dependence is achieved, the remaining choices in both sequences are set to be the same, implying that equation (3.9) can be written as:

$$\begin{aligned} v_{jt}(z_t) - v_{j't}(z_t) &= u_{jt}(z_t) - u_{j't}(z_t) \\ &+ \sum_{\tau=t+1}^{t+\rho} \sum_{k=1}^J \sum_{z_\tau=1}^Z \beta^{\tau-t+1} \left[u_{k\tau}(z_\tau) + \psi_k[p_\tau(z_\tau)] \right] \left[d_{k\tau}^*(z_\tau, j) \kappa_\tau^*(z_\tau|z_t, j) - d_{k\tau}^*(z_\tau, j') \kappa_\tau^*(z_\tau|z_t, j') \right] \end{aligned} \quad (3.11)$$

as the terms associated with time periods after $t + \rho$ drop out. Conditional on knowing the relevant $\psi_k(\mathbf{p})$ mappings, the CCP's, and the transitions on the state variables, the differenced conditional value function is now a linear function of flow payoffs from t to $t + \rho$. Further, only the $\psi_k(\mathbf{p})$ mappings that are along the two choice paths are needed: the econometrician only needs to know $\psi_k(\mathbf{p})$ if choice k is part of the two decision sequences. We next use some examples to illustrate how to exploit finite dependence in practice.

3.2.1 Example: Renewal Actions and Terminal Choices

We apply the results in the previous section to cases where the differences in future value terms across two choices only depend on one-period-ahead conditional choice probabilities and the flow payoff for a single choice. In particular, we consider renewal problems, such as Miller's (1984) job matching model or Rust's (1987) replacement model, where the individual can nullify the effects of a choice at time t on the state at time $t + 2$ by taking a renewal action at time $t + 1$. For example, if Zurcher replaces the engine at $t + 1$, then the state at time $t + 2$ does not depend upon whether the engine was replaced at time t or not. Let the renewal action be denoted as the first choice in the set $\{1, \dots, J\}$, implying that $d_{1t} = 1$ if the renewal choice is take at time t . Formally, a renewal action at $t + 1$ satisfies:

$$\sum_{z_{t+1}=1}^Z f_{1,t+1}(z_{t+2}|z_{t+1}) f_{jt}(z_{t+1}|z_t) = \sum_{z_{t+1}=1}^Z f_{1,t+1}(z_{t+2}|z_{t+1}) f_{j't}(z_{t+1}|z_t) \quad (3.12)$$

for all (z_{t+2}, j, j') . The state at $t + 2$ may then depend upon the state at $t + 1$, but only through variables that were unaffected by the choice at t .

Since the renewal action at $t + 1$ leads to the same expected state regardless of the choice at t , we define $d_{1,t+1}^*(z_{t+1}, j) = 1$ for all j . Equations (3.7) and (3.12) imply that $\kappa_{t+1}^*(z_{t+2}|z_t, j) = \kappa_{t+1}^*(z_{t+2}|z_t, j')$ for all j and j' . Expressing the future value terms relative to the value of the renewal choice, $v_{jt}(z_t) - v_{j't}(z_t)$ can be written as:

$$v_{jt}(z_t) - v_{j't}(z_t) = u_{jt}(z_t) - u_{j't}(z_t) + \beta \sum_{z_{t+1}=1}^Z [u_{1,t+1}(z_{t+1}) - \psi_1[p_{t+1}(z_{t+1})]] [f_{jt}(z_{t+1}|z_t) - f_{j't}(z_{t+1}|z_t)] \quad (3.13)$$

Hence, the only mapping we need is $\psi_1(\mathbf{p})$, the mapping for the renewal choice.

Hotz and Miller (1993) show another case where only one-period-ahead conditional choice probabilities and the flow payoff for a single choice are needed is when there is a terminal choice—a choice that, when made, implies no further choices. Let the terminal choice be denoted as the first choice. With no further choices being made, the future value term for the terminal choice can be collapsed into the current period payoff. $v_{jt}(z_t) - v_{j't}(z_t)$ then follows the same expression as (3.13) when neither j nor j' are the terminal choice.

3.2.2 Example: Labor Supply

To illustrate how finite dependence works when more than one period is required to eliminate the dependence, we develop the following stylized example.¹⁰ Consider a model of labor supply and human capital. In each of T periods an individual chooses whether to work, $d_{2t} = 1$, or stay home $d_{1t} = 1$. Individuals acquires human capital, z_t , by working, with the payoff to working increasing in human capital. If the individual works in period t , $z_{t+1} = z_t + 2$ with probability 0.5 and $z_{t+1} = z_t + 1$ also with probability 0.5. For periods after t , the human capital gain from working is fixed at one additional unit. When the individual does not work, her human capital remains the same in the next period.

The difference in conditional value functions between working and staying home at period t , $v_{2t}(z_t) - v_{1t}(z_t)$, can be expressed as functions of the two-period-ahead flow utilities and conditional probabilities of working and not working. To see this, consider $v_{1t}(z_t)$ which sets the initial choice to not work ($d_{1t} = 1$). Now set the next two choices to work: $d_{2,t+1}^*(z_{t+1}, 0) = 1$ and $d_{2,t+2}^*(z_{t+2}, 0) = 1$. This sequence of choices (not work, work, work) results in the individual having two additional units of human capital at $t + 3$. Given an initial choice to work ($d_{2t} = 1$) it is also possible to choose a sequence such that the individual will have two additional units of human capital at $t + 3$, but now

¹⁰For an empirical application of finite dependence involving more than one period and multiple discrete choices see Bishop (2008).

the sequence will depend upon the realization of the future states. In particular, if the decision to work at t results in an additional two units of human capital, then set the choice in period $t + 1$ to not work, $d_{1,t+1}(z_t + 2, 1) = 1$. However, if working at t results in only one additional unit, set the choice in period $t + 1$ to work, $d_{2,t+1}(z_t + 1, 1) = 1$. In either case, the third choice in the sequence is set to not work. We can write the future value terms relative to the choices in the sequences. The future value terms after $t + 3$ then cancel out once we difference $v_{1t}(z_t)$ from $v_{2t}(z_t)$.¹¹

3.3 Generalized Extreme Value Distributions

To apply (3.8) in estimation, the functional form of $\psi_j(\boldsymbol{p})$ for some $j \in \{1, \dots, J\}$ must be determined. It is well known that $\psi_j(\boldsymbol{p}) = -\ln(\boldsymbol{p}_j)$ when j is independently distributed as Type 1 Extreme Value. We show how to numerically calculate $\psi_j(\boldsymbol{p})$ for any generalized extreme value (GEV) distribution, and then lay out a class of problems where the mapping $\psi_j(\boldsymbol{p})$ has an analytic solution.

Suppose ϵ_t is drawn from the distribution function $G(\epsilon_{1t}, \epsilon_{2t}, \dots, \epsilon_{Jt})$ where

$$G(\epsilon_{1t}, \epsilon_{2t}, \dots, \epsilon_{Jt}) \equiv \exp[-\mathcal{H}(\exp[-\epsilon_{1t}], \exp[-\epsilon_{2t}], \dots, \exp[-\epsilon_{Jt}])]$$

and $G(\epsilon_{1t}, \epsilon_{2t}, \dots, \epsilon_{Jt})$ satisfies the properties outlined for the GEV distribution in McFadden (1978).¹² Letting $\mathcal{H}_j(Y_1, \dots, Y_J)$ denote the derivative of $\mathcal{H}(Y_1, \dots, Y_J)$ with respect to Y_j , we define $\phi_j(Y)$ as a mapping into the unit interval where

$$\phi_j(Y) = Y_j \mathcal{H}_j(Y_1, \dots, Y_J) / \mathcal{H}(Y_1, \dots, Y_J) \quad (3.14)$$

Note that, since $\mathcal{H}_j(Y_1, \dots, Y_J)$ and $\mathcal{H}(Y_1, \dots, Y_J)$ are homogeneous of degree zero and one respectively, $\phi_j(Y)$ is a probability as $\sum_{j=1}^J \phi_j(Y) = 1$. Indeed, McFadden (1978, page 80) establishes

¹¹There is another sequence that also results in a cancelation occurring after three periods. In this case, we set the sequence that begins with the choice to work ($d_{2t} = 1$) such that the next two choices are to not work regardless of the human capital realizations. In this case, at $t + 3$ the individual will have two additional units of human capital with probability 0.5 and one additional unit with probability 0.5. To have the same distribution of human capital at $t + 3$ given an initial choice not to work ($d_{1t} = 1$) involves mixing. In particular, set the choice at $t + 1$ to not work with probability 0.5, implying that the probability of not working at $t + 1$ is also 0.5. Setting the choice at $t + 2$ to work regardless of the level of human capital implies that an additional two units of human will result from the sequence with probability 0.5 with the probability of one additional unit resulting from the sequence also occurring with probability 0.5. Hence, the distribution of the states is the same given the two initial choices.

¹²The properties are that $\mathcal{H}(Y_1, Y_2, \dots, Y_J)$ is a nonnegative real valued function that is homogeneous of degree one, with $\lim \mathcal{H}(Y_1, Y_2, \dots, Y_J) \rightarrow \infty$ as $Y_k \rightarrow \infty$ for all $j \in \{1, \dots, J\}$, and for any distinct (i_1, i_2, \dots, i_r) , the cross derivative $\partial \mathcal{H}(Y_1, Y_2, \dots, Y_J) / \partial Y_{i_1}, Y_{i_2}, \dots, Y_{i_r}$ is nonnegative for r odd and nonpositive for r even.

that substituting $\exp[v_{jt}(z_t)]$ in for Y_j in equation (3.14) yields the conditional choice probability $p_{jt}(z_t)$:

$$p_{jt}(z_t) = e^{v_{jt}(z_t)} \mathcal{H}_j \left(e^{v_{1t}(z_t)}, \dots, e^{v_{Jt}(z_t)} \right) / \mathcal{H} \left(e^{v_{1t}(z_t)}, \dots, e^{v_{Jt}(z_t)} \right) \quad (3.15)$$

We now establish a relationship between $\phi_j(Y)$ and $\psi_j(\boldsymbol{p})$. Denoting the vector function $\phi(Y) \equiv \{\phi_2(Y), \dots, \phi_J(Y)\}$, lemma 2 shows that $\phi(Y)$ is invertible. Further, lemma 2 establishes that there is a closed form expression for $\psi_j(\boldsymbol{p})$ when $\phi^{-1}(\boldsymbol{p})$ is known.

Lemma 2 *When ϵ_t is drawn from a GEV distribution, the inverse function $\phi^{-1}(\boldsymbol{p})$ exists and $\psi_j(\boldsymbol{p})$ is given by:*

$$\psi_j(\boldsymbol{p}) = \ln \mathcal{H} [1, \phi_2^{-1}(\boldsymbol{p}), \dots, \phi_J^{-1}(\boldsymbol{p})] - \ln \phi_j^{-1}(\boldsymbol{p}) + \gamma \quad (3.16)$$

It is straightforward to use lemma 2 in practice by evaluating $\phi_j^{-1}(\boldsymbol{p})$ at $p_t(z_t)$ for a given z_t . To see this, note that from (3.15) we can express the vector $p_t(z_t)$ as:

$$\begin{bmatrix} p_{2t}(z_t) \\ \vdots \\ p_{Jt}(z_t) \end{bmatrix} = \begin{bmatrix} e^{v_{2t}(z_t) - v_{1t}(z_t)} \mathcal{H}_2(1, \dots, e^{v_{Jt}(z_t) - v_{1t}(z_t)}) / \mathcal{H}(1, \dots, e^{v_{Jt}(z_t) - v_{1t}(z_t)}) \\ \vdots \\ e^{v_{Jt}(z_t) - v_{1t}(z_t)} \mathcal{H}_J(1, \dots, e^{v_{Jt}(z_t) - v_{1t}(z_t)}) / \mathcal{H}(1, \dots, e^{v_{Jt}(z_t) - v_{1t}(z_t)}) \end{bmatrix} \quad (3.17)$$

Making the $(J - 1)$ dimensional vector formed from $\exp[v_{jt}(z_t) - v_{1t}(z_t)]$ the subject of (3.17) for any vector $p_t(z_t)$ solves this $(J - 1)$ equation system to yield the $(J - 1)$ unknowns $\phi_j^{-1}[p_t(z_t)]$.

In some cases, $\psi_j(\boldsymbol{p})$ (and therefore $\phi_j^{-1}(\boldsymbol{p})$) has an analytic solution. For example, consider a case where $G(\epsilon_t)$ factors into two independent distributions, one being a nested logit, and the other a GEV distribution. Let \mathcal{J} denote the set of choices in the nest and let K denote the number of choices that are outside the nest. $G(\epsilon_t)$ can then be expressed as:

$$G(\epsilon_t) \equiv G_0(\epsilon_{1t}, \dots, \epsilon_{Kt}) \exp \left[- \left(\sum_{j \in \mathcal{J}} \exp[-\epsilon_{jt}/\sigma] \right)^\sigma \right] \quad (3.18)$$

where $G_0(Y_1, Y_2, \dots, Y_K)$ satisfies the properties outlined for the GEV distribution as defined by McFadden (1978). The correlation of the errors within the nest is given by $\sigma \in [0, 1]$ and errors within the nest are uncorrelated with errors outside the nest. When $\sigma = 1$, the errors are uncorrelated within the nest, and when $\sigma = 0$ they are perfectly correlated. Lemma (3) then shows the closed form expression for $\psi_j(\boldsymbol{p})$ for all $j \in \mathcal{J}$.

Lemma 3 *If $G(\epsilon_t)$ can be expressed as in (3.18), then*

$$\psi_j(\boldsymbol{p}) = \gamma - \sigma \ln(p_j) - (1 - \sigma) \ln \left(\sum_{k \in \mathcal{J}} p_k \right) \quad (3.19)$$

Note that $\psi_j(\boldsymbol{p})$ only depends on the conditional choice probabilities for choices that are in the nest: the expression is the same no matter how many choices are outside the nest or how those choices are correlated. Hence, $\psi_j(\boldsymbol{p})$ will only depend on $\boldsymbol{p}_{j'}$ if ϵ_{jt} and $\epsilon_{j't}$ are correlated. When $\sigma = 1$, ϵ_{jt} is independent of all other errors and $\psi_j(\boldsymbol{p})$ only depends on \boldsymbol{p}_j .

Lemma 3 is particularly powerful when there is a renewal or terminal choice. Recall from Section 3.2.1 that the only $\psi_j(\boldsymbol{p})$ mappings needed in these cases were for the renewal or terminal choices. The payoffs for these choices may naturally be viewed as having an independent error. For example, in the Rust case, bus engine maintenance actions are more likely to be correlated with each other than with engine replacement. Another example is firm decisions when exit is an option. Choosing how many stores or different levels of product quality are likely to have correlation patterns among the errors that are unrelated to the payoff from exiting. As long as the error associated with the renewal or terminal choice is independent of the other errors, any correlation pattern among the other errors will still result in $\psi_j(\boldsymbol{p}) = -\ln(\boldsymbol{p}_j)$ when j is the renewal or terminal choice.

4 Identification

Having shown how to frame dynamic discrete choice problem in a way that facilitates CCP estimation for broad classes of problems and error structures, we now turn to identification. The literature on the identification of dynamic discrete choice models can be broken down into two questions.¹³ If the conditional choice probabilities and state variable transitions are known, what normalizations suffice to recover the flow payoff functions? Under what conditions are the conditional choice probabilities and the state transitions identified when some elements of z are unobserved? Our discussion follows this format.

4.1 Recovering flow utilities from conditional choice probabilities

Let $\mathcal{T} \leq T$ denoted the date the panel ends. We differentiate between two scenarios. When $\mathcal{T} = T < \infty$, or when $T = \infty$ and the environment is stationary, we extend the results of Magnac and Thesmar (2002) and Pesendorfer and Schmidt-Dengler (2008), who analyze identification when the conditional choice probabilities are known. Identification when $\mathcal{T} < T < \infty$, or $T = \infty$ and the environment is nonstationary, is trickier. In this case the method of solving and imposing the solution of the underlying discrete choice problem on the data generating process is not feasible

¹³See Rust (1994), Magnac and Thesmar (2002), Heckman and Navarro (2007), Pesendorfer and Schmidt-Dengler (2008), Kasahara and Shimotsu (2009), Bajari, Chernozhukov, Hong, and Nekipelov (2009), and Hu and Shum (2010a).

without making strong parametric assumptions about features of the framework that affect the flow payoffs and state variable transitions between dates \mathcal{T} and T . We show, however, that the model is partially identified under much weaker assumptions involving specializations of the finite dependence property.

We adopt several normalizations commonly made in the static and dynamic discrete choice literature. Since only choices are observed rather than utility levels, we normalize the flow payoff function for one of the choices to zero in every time period, setting $u_{1t}(z_t) = 0$ for all z_t and t .¹⁴ Absent observations on utility, we also assume the distribution of the transitory vector of shocks, $G(\epsilon_t)$, is known. Finally because the time-subscripted utility functions depend in an unrestricted way upon the time period, the subjective discount factor β must also be normalized.

4.1.1 When the setting is stationary or $\mathcal{T} = T$

Building on the results of Magnac and Thesmar (2002), we begin by showing how to recover flow payoff functions when the sampling period is the same as the horizon or when the problem is stationary. Theorem 2 and Corollary 3 of Magnac and Thesmar (2002, pages 807 and 808) establish exact identification of the differences in conditional value functions when T is finite. An equivalent result holds for the identification of flow utilities. Let $d_{1\tau}^*(z_\tau) = 1$ for all τ in equation (3.8) and subtract $v_{1t}(z_t)$ from $v_{jt}(z_t)$ to obtain:

$$v_{jt}(z_t) - v_{1t}(z_t) = u_{jt}(z_t) + \sum_{\tau=t+1}^T \sum_{z_\tau=1}^Z \beta^{\tau-t} \psi_1[p_\tau(z_\tau)] [\kappa_{\tau-1}^*(z_\tau|z_t, j) - \kappa_{\tau-1}^*(z_\tau|z_t, 1)] \quad (4.1)$$

An alternative expression for this difference can be obtained by differencing the expressions for $\psi_{1t}(z_t)$ and $\psi_{jt}(z_t)$ given in equation (3.5):

$$v_{jt}(z_t) - v_{1t}(z_t) = \psi_1[p_t(z_t)] - \psi_j[p_t(z_t)] \quad (4.2)$$

As shown in Theorem 2 below, the two expressions for $v_{jt}(z_t) - v_{1t}(z_t)$ can then be used to form expressions for $u_{jt}(z_t)$ as a function of the transition probabilities, the conditional choice probabilities, and the discount factor. Further, Theorem 2 shows how the problem simplifies in the stationary case where the time subscripts are dropped from the flow payoffs and the transition functions and when the time horizon is infinite.

¹⁴As noted by Pesendorfer and Schmidt-Dengler (2008, page 913), in some situations, such when a firm exits an industry, the future value of the choice to exit may be known, providing an empirical justification for the normalization.

Theorem 2 For all j , t , and z_t , the flow payoff $u_{jt}(z_t)$ can be expressed as.

$$u_{jt}(z_t) = \psi_1[p_t(z_t)] - \psi_j[p_t(z_t)] + \sum_{\tau=t+1}^T \sum_{z_\tau=1}^Z \beta^{\tau-t} \psi_1[p_\tau(z_\tau)] [\kappa_{\tau-1}^*(z_\tau|z_t, 1) - \kappa_{\tau-1}^*(z_\tau|z_t, j)] \quad (4.3)$$

When the environment is stationary, let \mathcal{I} denote the Z dimensional identity matrix and define

$$u_j \equiv \begin{bmatrix} u_j(1) \\ \vdots \\ u_j(Z) \end{bmatrix}, \quad F_j \equiv \begin{bmatrix} f_j(1|1) & \dots & f_j(Z|1) \\ \vdots & \ddots & \vdots \\ f_j(1|Z) & \dots & f_j(Z|Z) \end{bmatrix}, \quad \Psi_j \equiv \begin{bmatrix} \psi_j[p(1)] \\ \vdots \\ \psi_j[p(Z)] \end{bmatrix}$$

Then $[\mathcal{I} - \beta F_1]$ is invertible and for all j :

$$u_j = \Psi_j - \Psi_1 + \beta (F_1 - F_j) [\mathcal{I} - \beta F_1]^{-1} \Psi_1 \quad (4.4)$$

Given the assumptions made at the beginning of this section regarding the state transitions, conditional choice probabilities, the discount factor, and the distribution of the structural errors, everything on the right hand side of both (4.3) and (4.4) is known and, therefore, both systems are exactly identified. However, by putting further structure on the flow payoff function, the error distribution can be made more flexible and the discount factor may be identified.¹⁵

4.1.2 When the setting is nonstationary and $\mathcal{T} < T$

When no data on choices or state transitions are available for the last part of the lifecycle, the expressions derived above cannot aid identification, because there are no conditional choice probabilities or transition functions past \mathcal{T} . Nevertheless, under finite dependence there are two special cases when identification can be restored.

First we suppose that finite dependence can be achieved by repeated use of the same action for ρ periods. This is the case of renewal (after one period), or when repeated use of the same action results in the state eventually resetting.¹⁶ Without loss of generality we label that action as the first. By finite dependence we see that:

$$u_{jt}(z_t) = \psi_1[p_t(z_t)] - \psi_j[p_t(z_t)] + \sum_{\tau=t+1}^{t+\rho} \sum_{z_\tau=1}^Z \beta^{\tau-t} \psi_1[p_\tau(z_\tau)] [\kappa_{\tau-1}^*(z_\tau|z_t, 1) - \kappa_{\tau-1}^*(z_\tau|z_t, j)]$$

¹⁵In our first Monte Carlo, we assume there is a variable that affects payoffs only through the transition probabilities of the state variables. This serves as an exclusion restriction permitting estimation of the discount factor. See Miller and Sanders (1997) and Arcidiacono, Sieg, and Sloan (2007) for empirical applications.

¹⁶See Altug and Miller (1998) for an example of the latter case.

Provided there are more than ρ periods of data, it follows that $u_{jt}(z_t)$ can be recovered for $t < \mathcal{T} - \rho$.¹⁷

A second case of interest is when the flow utility functions, but not the state transitions, do not depend on time.¹⁸ In this case, we omit the subscripts on the flow utility functions, setting $u_{jt}(z_t) = u_j(z_t)$ for all t , but retaining the time subscripts for the CCP's and the transitions that define $\kappa_\tau^*(z_\tau|z_t, j)$. Under finite dependence, $u_j(z_t)$ can be expressed as:

$$u_j(z_t) = \psi_1[p_t(z_t)] - \psi_j[p_t(z_t)] + \sum_{\tau=t+1}^{t+\rho} \sum_{k=1}^J \sum_{z_\tau=1}^Z \beta^{\tau-t} \{u_k(z_\tau) + \psi_k[p_\tau(z_\tau)]\} [\kappa_\tau^*(z_\tau|z_t, j) - \kappa_\tau^*(z_\tau|z_t, 1)] \quad (4.5)$$

where the choices in the future sequence generate the finite dependence. Note that, in contrast to previous examples, flow payoff terms are on both sides of (4.5). Note further that there are $(J - 1) \times (\mathcal{T} - \rho) \times Z$ equations but only $(J - 1) \times Z$ unknowns. Hence, if a rank condition on the equations is met and conditional on one flow payoff being normalized, the remaining flow payoff terms may be recovered. Further, since the number of equations exceed the number of unknowns, it may also possible to parameters such as the discount factor or the correlation patterns of the structural errors.

4.2 Identifying conditional choice probabilities and state transitions

To analyze identification of the conditional choice probabilities in the presence of unobserved heterogeneity, we partition z_t , the state variables, into those observed by the econometrician, denoted by $x_t \in \{1, \dots, X\}$, and those that are unobserved, denoted by $s_t \in \{1, \dots, S\}$. We simplify the exposition of the estimator, and reduce the burden of identification, by further assuming that the transition probabilities for the unobserved variables are exogenous. We assume s_t follows a stationary first order Markov process with $\pi(s_{t+1}|s_t)$ giving the probability of transitions from unobserved state s to s' in consecutive periods. The transitions on the observed and unobserved states, $f_{jt}(z_{t+1}|z_t)$, can then be expressed as:

$$f_{jt}(x_{t+1}, s_{t+1}|x_t, s_t) = \pi(s_{t+1}|s_t) f_{jt}(x_{t+1}|x_t, s_t)$$

¹⁷The same flow payoff functions can also be recovered when there is a terminal choice. Normalizing the flow payoff for the terminal choice and expressing the future value function relative to the terminal choice yields the result.

¹⁸An example is the female labor supply case in Section 3.2.2, where data is available for women who are substantially below retirement age. It is possible to recover the flow payoff of working while making no assumptions about how individuals form expectations over social security rules, retirement behavior, or even the time horizon.

The initial distribution of the unobserved state, s_1 , conditional on the observables in the first observed time period, x_1 , is given by $\pi(s_1|x_1)$, permitting an initial conditions problem.¹⁹

If all the state variables are observed, then consistent estimators for the conditional choice probabilities and the state transitions can be computed directly from the data. When s is unobserved, neither the conditional choice probabilities $p_{jt}(z) \equiv p_{jt}(x, s)$, the state transitions represented by $\pi(s_{t+1}|s_t)$ and $f_{jt}(x_{t+1}|x_t, s_t)$, nor the initial conditions denoted by $\pi(s_1|x_1)$ have direct counterparts in the data.

Given $G(\epsilon_t)$ and β , there are $STX(J-1)$ parameters to represent the conditional choice probabilities from the sample period. Without placing structure on the transitions for the observed state variables, $f_{jt}(x_{t+1}|x_t, s_t)$, results in another $J SX(\mathcal{T}-1)(X-1)$ parameters. The transition matrix of the unobserved states, $\pi(s, s')$, requires $S(S-1)$ parameters with an additional $X(S-1)$ parameters characterizing the initial distribution of the unobserved variables, $\pi(s_1|x_1)$. The total number of parameters can then be expressed as:

$$PAR \equiv JSTX^2 + JSX + SX + S^2 - JSX^2 - STX - X - S \quad (4.6)$$

The equations for identifying these parameters come from the joint distribution of the possible decisions and state variables over the sample period from $t = 1$ to \mathcal{T} . Denote the joint probability mass function as $\mathcal{F}(d_1, \dots, d_{\mathcal{T}}, x_1, \dots, x_{\mathcal{T}})$. There are $J^{\mathcal{T}}X^{\mathcal{T}}$ possible sequences of decisions and state variables. Given that the sum of the probabilities associated with each of the sequences must equal one, there are then at most $J^{\mathcal{T}}X^{\mathcal{T}} - 1$ probabilities that are not linear combinations of each other. The probabilities of each of these sequences can then be expressed as functions of $p_{jt}(x, s)$, $f_{jt}(x_{t+1}|x_t, s_t)$, $\pi(s_{t+1}|s_t)$, and $\pi(s_1|x_1)$. In particular, we can express the joint distribution of $(d_1, \dots, d_{\mathcal{T}}, x_1, \dots, x_{\mathcal{T}})$ as:

$$\begin{aligned} \mathcal{F}(d_1, \dots, d_{\mathcal{T}}, x_1, \dots, x_{\mathcal{T}}) &= \sum_{s_1=1}^S \dots \sum_{s_{\mathcal{T}}=1}^S \left(\prod_{j=1}^J \left[\pi(s_1|x_1) p_{j1}(x_1, s_1) f_{j1}(x_2|x_1, s_1) \right]^{d_{j1}} \right) \\ &\times \left(\prod_{t=2}^{\mathcal{T}-1} \prod_{j=1}^J \left[\pi(s_t|s_{t-1}) p_{jt}(x_t, s_t) f_{jt}(x_{t+1}|x_t, s_t) \right]^{d_{jt}} \right) \left(\prod_{j=1}^J \left[\pi(s_{\mathcal{T}}|s_{\mathcal{T}-1}) p_{j\mathcal{T}}(x_{\mathcal{T}}, s_{\mathcal{T}}) \right]^{d_{j\mathcal{T}}} \right) \end{aligned} \quad (4.7)$$

Forming these expressions for each non-linearly dependent sequence yields a system of equations. A necessary condition for identification is that $PAR + 1 \leq J^{\mathcal{T}}X^{\mathcal{T}}$. This restricts the number of unobserved states, S , that can be identified for a given sample length \mathcal{T} . Because the number of

¹⁹Note that the discussion here centers around recovering the distribution of unobserved states and the effect of the unobserved states on the flow payoffs. The unobserved states themselves, be it the individual realizations or what the states correspond to, cannot be recovered.

probabilities grows exponentially with \mathcal{T} , but from (4.6) the number of unknown parameters grows linearly with \mathcal{T} , these necessary conditions are satisfied by moderate \mathcal{T} for quite large S .²⁰

Indeed Kasahara and Shimotsu (2009) and Hu and Shum (2010a) provide rank conditions for identifying conditional choice probabilities and state transitions in the presence of unobserved heterogeneity from very short panels. In a stationary infinite horizon discrete choice model where the unobserved variable does not evolve over time and only affects utility flows, the joint distribution of the observed state variables, denoted by $\mathcal{F}(x_1, x_2, \dots, x_{\mathcal{T}})$, is identified from their transitions. The remaining parameters are the CCP's, numbering $(J - 1) \times X \times S$ in this stationary environment, and the $X \times (S - 1)$ probabilities characterizing the invariant distribution for the unobservables. Factoring (4.7) into $\mathcal{F}(x_1, x_2, \dots, x_{\mathcal{T}})$ and the distribution of choices conditional on the state variables, $\mathcal{F}(d_1, d_2, \dots, d_{\mathcal{T}} | x_1, x_2, \dots, x_{\mathcal{T}})$, the invariance property implies the timing of the observed sequence of sample choices conditional on $(x_1, x_2, \dots, x_{\mathcal{T}})$ conveys no information about their underlying probabilities. In this environment, Kasahara and Shimotsu show that even panels of only three periods are sufficient to identify the CCP's. Hu and Shum extend their work by providing invertibility conditions that identify unobserved Markov processes when the length of the sample is at least five periods.

To summarize, assumptions about the idiosyncratic disturbances defined by $G(\epsilon_t)$, and also the discount factor β , can only be relaxed by imposing economic structure on the utility parameters defined by $u_{jt}(z_t)$. In contrast, the scope for incorporating unobserved state variables is limited by data considerations, namely the length of the sample. Nevertheless, even in very short panels some unobserved heterogeneity is permissible. Further, our discussion of finite dependence and short panels demonstrates that, by imposing parametric structure, it is not necessary to identify all of the unrestricted CCP's and state transitions to recover utility flows.

5 The Estimators and their Asymptotic Properties

In estimation, we parameterize the utility function, transition function, and the probability density function for ϵ_t by a finite dimensional vector $\theta \in \Theta$, where Θ denotes the parameter space for $u_{jt}(x_t, s_t)$, $f_{jt}(x_{t+1} | x_t, s_t)$, $g(\epsilon_t)$ and β and is assumed to be convex and compact. There are then two sets of parameters to be estimated: θ and π , where π includes the initial distribution of the

²⁰Note that when overidentifying equations for the CCPs and the transitions of the state variables exist, they cannot help identify $G(\epsilon_t)$ and β . This is because they are (weakly) less informative than knowing the true CCP's and the true state transitions, to which the exact identification result of Theorem 2 applies.

unobserved heterogeneity, $\pi(s_1|x_1)$, and its transition probability matrix, $\pi(s_{t+1}|s_t)$.²¹

The CCP estimators we propose are derived from two sets of conditions. First are conditions which ensure the estimates of θ and π maximize the likelihood function taking the estimates of the conditional choice probabilities as given. Second are conditions that ensure the estimated conditional choice probabilities are consistent with either the data or the underlying structural model.

5.1 The likelihood

Denote $p_{jt}(x, s)$ as a value for the probability an individual will choose j at t given observed states x and s . Let p indicate the $(J-1) \times T \times X \times S$ vector of conditional choice probabilities with the elements given by $p_{jt}(x, s)$.²² Denote by $l_{jt}(x_{nt}, s_{nt}, \theta, \pi, p)$ the likelihood of observing individual n make choice j at time t , conditional on the state (x_{nt}, s_{nt}) , the parameters θ and π , and the conditional choice probabilities p :

$$l_{jt}(x_{nt}, s_{nt}, \theta, \pi, p) \equiv \Pr \left\{ \arg \max_{j \in \{1, \dots, J\}} [v_{jt}(x_{nt}, s_{nt}, \theta, \pi, p) + \epsilon_{jt}] | x_{nt}, s_{nt}; \theta, \pi, p \right\} \quad (5.1)$$

When $d_{jnt} = 1$, the expression in (5.1) simplifies to (2.8) in the motivating example. The corresponding likelihood of observing (d_{nt}, x_{nt+1}) is then defined as:

$$\mathcal{L}_t(d_{nt}, x_{nt+1} | x_{nt}, s_{nt}; \theta, \pi, p) = \prod_{j=1}^J [l_{jt}(x_{nt}, s_{nt}, \theta, \pi, p) f_{jt}(x_{n,t+1} | x_{nt}, s_{nt}, \theta)]^{d_{jnt}} \quad (5.2)$$

The joint likelihood of any given path of choices, $d_n \equiv (d_{n1}, \dots, d_{n\mathcal{T}})$, and observed states, $x_n \equiv (x_{n1}, \dots, x_{n,\mathcal{T}+1})$, is derived by forming the product of (5.2) over the \mathcal{T} sample periods, and then integrating it over the unobserved states, $(s_{n1}, \dots, s_{n\mathcal{T}})$. Since the probability distribution for the initial unobserved state is $\pi(s_1|x_1)$, and the transition probability is $\pi(s_{t+1}|s_t)$, the likelihood of observing (d_n, x_n) conditional on x_1 for parameters (θ, π, p) is:

$$L(d_n, x_n | x_{n1}; \theta, \pi, p) = \sum_{s_1=1}^S \sum_{s_2=1}^S \dots \sum_{s_{\mathcal{T}}=1}^S \left[\begin{aligned} &\pi(s_1|x_{n1}) \mathcal{L}_1(d_{n1}, x_{n2} | x_{n1}, s_1; \theta, \pi, p) \\ &\times \left(\prod_{t=2}^{\mathcal{T}} \pi(s_t|s_{t-1}) \mathcal{L}_t(d_{nt}, x_{nt+1} | x_{nt}, s_t; \theta, \pi, p) \right) \end{aligned} \right] \quad (5.3)$$

The log likelihood of the sample is then given by:

$$\sum_{n=1}^{\mathcal{N}} \ln L(d_n, x_n | x_{n1}; \theta, \pi, p) \quad (5.4)$$

²¹Note that we assume the number of unobserved states is known. Heckman and Singer (1984) provide conditions for identifying the number of unobserved states in dynamic models. In principle, one could estimate estimate the model separately for different values of S , using the differences in the likelihoods to choose the number of unobserved states. See McLachlan and Peel (2000), chapter 6.

²²In the stationary case the dimension is $(J-1) \times X \times S$.

5.2 The estimators

Following our motivating example in Section 2, we define two estimators of (θ, π, p) . At minimal risk of some ambiguity we label them both by $(\hat{\theta}, \hat{\pi}, \hat{p})$. Both estimators of (θ, π) maximize the expression obtained by substituting an estimator for p into (5.4):

$$(\hat{\theta}, \hat{\pi}) = \arg \max_{\theta, \pi} \sum_{n=1}^{\mathcal{N}} \ln L(d_n, x_n | x_{n1}; \theta, \pi, \hat{p}) \quad (5.5)$$

In one case $\hat{p}_{jt}(x, s)$ satisfies the likelihood expression $l_{jt}(x, s, \theta, \pi, p)$ evaluated at $(\hat{\theta}, \hat{\pi}, \hat{p})$, meaning:

$$\hat{p}_{jt}(x, s) = l_{jt}(x, s, \hat{\theta}, \hat{\pi}, \hat{p}) \quad (5.6)$$

This constraint is motivated by the fact that the model itself generates the conditional choice probabilities. Thus our first estimator solves for (θ, π, p) in the system of equations defined by (5.6) along with satisfying (5.5). To deal with possibility of multiple solutions, we select the one attaining the highest likelihood.

In the second case $\hat{p}_{jt}(x, s)$ is calculated as a weighted average of d_{njt} over the sample. No weight is given to d_{njt} if $x_{nt} \neq x$. When $x_{nt} = x$, the weight corresponds to the probability of individual n being in unobserved state s at time t . To obtain the weight in this case, we define the joint likelihood of both $s_{nt} = s$ and the sequence (d_n, x_n) occurring, as:

$$\begin{aligned} L_{nt}(s_{nt} = s) &\equiv \sum_{s_1=1}^S \dots \sum_{s_{t-1}=1}^S \sum_{s_{t+1}=1}^S \dots \sum_{s_T=1}^S \pi(s_1 | x_{n1}) \mathcal{L}_{n1}(s_1) \left(\prod_{t'=2}^{t-1} \pi(s_{t'} | s_{t'-1}) \mathcal{L}_{nt'}(s_{t'}) \right) \\ &\quad \times \pi(s_t | s_{t-1}) \mathcal{L}_{nt}(s) \pi(s_{t+1} | s) \mathcal{L}_{n,t+1}(s_{t+1}) \left(\prod_{t'=t+2}^T \pi(s_{t'} | s_{t'-1}) \mathcal{L}_{nt'}(s_{t'}) \right) \end{aligned} \quad (5.7)$$

Similarly denote by $L_n \equiv L(d_n, x_n | x_{n1}; \theta, \pi, p)$ the likelihood of observing (d_n, x_n) . From Bayes' rule, the probability that $s_{nt} = s$ conditional on (d_n, x_n) is $L_{nt}(s_{nt} = s) / L_n$. Let $\hat{L}_{nt}(s_{nt} = s)$ denote $L_{nt}(s_{nt} = s)$ evaluated at $(\hat{\theta}, \hat{\pi}, \hat{p})$, and similarly let \hat{L}_n denote L_n evaluated at the parameter estimates. Our second estimator of $\hat{p}_{jt}(x, s)$ satisfies:

$$\begin{aligned} \hat{p}_{jt}(x, s) &= \left[\sum_{n=1}^{\mathcal{N}} d_{njt} I(x_{nt} = x) \frac{\hat{L}_{nt}(s_{nt} = s)}{\hat{L}_n} \right] \bigg/ \left[\sum_{n=1}^{\mathcal{N}} I(x_{nt} = x) \frac{\hat{L}_{nt}(s_{nt} = s)}{\hat{L}_n} \right] \\ &= \sum_{n=1}^{\mathcal{N}} d_{njt} \left[\frac{I(x_{nt} = x) \hat{L}_{nt}(s_{nt} = s) / \hat{L}_n}{\sum_{n'=1}^{\mathcal{N}} I(x_{n't} = x) \hat{L}_{n'}(s_{n't} = s) / \hat{L}_{n'}} \right] \end{aligned} \quad (5.8)$$

Note that the numerator in the first line gives the average probability of $(d_{njt}, x_{nt}, s_{nt}) = (1, x, s)$ and the denominator gives the average probability of $(x_{nt}, s_{nt}) = (x, s)$. Their ratio is then an estimate

of the probability of a sampled person choosing action j at time t conditional on x and s , with the bracketed term in the second line giving the weight placed on d_{njt} . The second estimator $(\hat{\theta}, \hat{\pi}, \hat{p})$ solves (5.5) along with the set of conditions given in (5.8); again in case of multiple solutions, the solution attaining the highest likelihood is selected.

5.3 Large sample properties

Our representation of the conditional value functions implies that any set of conditional choice probabilities \tilde{p} defined for all (j, t, x, s) induces payoffs as a function of (θ, π) . Substituting \tilde{p} for p in (5.4), and then maximizing the resulting expression with respect to (θ, π) , yields estimates of the structural parameters, which we denote by $(\tilde{\theta}, \tilde{\pi})$. If the payoff functional forms were correctly specified, $(\tilde{\theta}, \tilde{\pi})$ would converge to the true parameters under standard regularity conditions for static random utility models. Imposing the condition that $\tilde{p} = \hat{p}$ merely ensures an internal consistency: the conditional valuation functions used in the functional form for utility in (5.4) are based on the same conditional choice probabilities that emerge if the individuals in the sample actually face primitives given by $(\hat{\theta}, \hat{\pi})$. The proof to the following theorem shows that if the model is identified, the true set of parameters satisfy this internal consistency condition. Intuitively this explains why both our estimators are consistent.

Theorem 3 *If the model is identified, then $(\hat{\theta}, \hat{\pi}, \hat{p})$ is consistent in both cases.*

The remaining large sample properties, \sqrt{N} rate of convergence and asymptotic normality, can be established by appealing to well known results in the literature. The covariance matrices of the estimators are given in the appendix.

6 The Algorithm

In order to operationalize our estimators, we modify the EM algorithm. The EM algorithm iterates on two steps. In the expectations step, the conditional probabilities of being in each unobserved state are updated as well as the initial conditions and law of motion for the unobserved states. The maximization step proceeds as if the unobserved state is observed and uses the conditional probabilities of being in each unobserved state as weights.

We show how the m^{th} iteration is updated to the $(m + 1)^{th}$ as well as how to initiate the algorithm. We lay out the expectations step, the maximization step, and then summarize the algorithm.

6.1 Expectations Step

For the sake of the exposition, we break down the expectations step into updating:

1. $q_{nst}^{(m)}$, the probability of n being in unobserved state s at time t ,
2. $\pi^{(m)}(s_1 | x_1)$, the probability distribution over the initial unobserved states conditional on the initial observed states,
3. $\pi^{(m)}(s' | s)$, the transition probabilities of the unobserved states,
4. $p^{(m)}(x, s)$, the conditional choice probabilities.

6.1.1 Updating $q_{nst}^{(m)}$

The first step of the m th iteration is calculating the conditional probability of being in each unobserved state in each time period given the values of the structural parameters and conditional choice probabilities from the m th iteration, $\{\theta^{(m)}, \pi^{(m)}, p^{(m)}\}$. The likelihood of the data on n given the parameters at the m th iteration is found by evaluating (5.3) at $\{\theta^{(m)}, \pi^{(m)}, p^{(m)}\}$:

$$L_n^{(m)} \equiv L(d_n, x_n | x_{n1}; \theta^{(m)}, \pi^{(m)}, p^{(m)}) \quad (6.1)$$

Similarly, we denote by $L_n^{(m)}(s_{nt} = s)$ the joint likelihood of the data and unobserved state s occurring at time t , given the parameter evaluation at iteration m . Evaluating (5.7) at $\{\theta^{(m)}, \pi^{(m)}, p^{(m)}\}$ yields:

$$L_{nt}^{(m)}(s_{nt} = s) \equiv L_{nt}(d_n, x_n, s_{nt} = s | x_{n1}; \theta^{(m)}, \pi^{(m)}, p^{(m)}) \quad (6.2)$$

At iteration $m + 1$, the probability of n being in unobserved state s at time t , $q_{nst}^{(m+1)}$, then follows from Bayes rule:

$$q_{nst}^{(m+1)} = \frac{L_n^{(m)}(s_{nt} = s)}{L_n^{(m)}} \quad (6.3)$$

6.1.2 Updating $\pi^{(m)}(s | x)$

Setting $t = 1$ in (6.3) yields the conditional probability of the n^{th} individual being in unobserved state s in the first time period. When the state variables are exogenous at $t = 1$, we can update the probabilities for the initial states by averaging the conditional probabilities obtained from the previous iteration over the sample population:

$$\pi^{(m+1)}(s) = \frac{1}{\mathcal{N}} \sum_{n=1}^{\mathcal{N}} q_{ns1}^{(m+1)} \quad (6.4)$$

To allow for situations where the distribution of the unobserved states in the first period depends on the values of the observed state variables, we form averages over $q_{nst}^{(m+1)}$ for each value of x . Generalizing (6.4), we set:

$$\pi^{(m+1)}(s|x) = \frac{\sum_{n=1}^{\mathcal{N}} q_{ns1}^{(m+1)} I(x_{n1} = x)}{\sum_{n=1}^{\mathcal{N}} I(x_{n1} = x)} \quad (6.5)$$

6.1.3 Updating $\pi^{(m)}(s'|s)$

Updating the probabilities of transitioning among unobserved states requires calculating the probability of n being in unobserved state s' at time t conditional on the data and also on being in unobserved state s at time $t-1$, $q_{ns't|s}$. The joint probability of n being in states s and s' at time $t-1$ and t can then be expressed as the product of q_{nst-1} and $q_{ns't|s}$. The updating formula for the transitions on the unobserved states is then based on the identities:

$$\begin{aligned} \pi(s'|s) &= \frac{E_n [I(s_{nt-1} = s)I(s_{nt} = s')]}{E_n [I(s_{nt-1} = s)]} \\ &= \frac{E_n \{E [I(s_{nt} = s') | d_n, x_n, s_{nt-1} = s] E [I(s_{nt-1} = s) | d_n, x_n]\}}{E_n \{E [I(s_{nt-1} = s) | d_n, x_n]\}} \\ &= \frac{E_n [q_{ns't|s} q_{nst-1}]}{E_n [q_{nst-1}]} \end{aligned} \quad (6.6)$$

where the n subscript on an expectations operator indicates that the integration is taken over the population. The second line then follows from the law of iterated expectations.

Substituting the relevant sample analogs at the m th iteration for $q_{ns't|s}$ and q_{nst-1} into (6.6) then yields our update of $\pi^{(m)}(s'|s)$. First we compute $q_{ns't|s}^{(m+1)}$ according to:

$$q_{n,s',t|s}^{(m+1)} = \frac{\pi^{(m)}(s'|s) \mathcal{L}_{nt}^{(m)}(s') \left[\sum_{s_{t+1}=1}^S \dots \sum_{s_{\mathcal{T}}=1}^S \left(\pi^{(m)}(s_{t+1}|s') \mathcal{L}_{n,t+1}^{(m)}(s_{t+1}) \prod_{t'=t+2}^{\mathcal{T}} \pi^{(m)}(s_{t'}|s_{t'-1}) \mathcal{L}_{nt'}^{(m)}(s_{t'}) \right) \right]}{\sum_{s_t=1}^S \pi^{(m)}(s_t|s) \mathcal{L}_{nt}^{(m)}(s_t) \left[\sum_{s_{t+1}=1}^S \dots \sum_{s_{\mathcal{T}}=1}^S \left(\prod_{t'=t+1}^{\mathcal{T}} \pi^{(m)}(s_{t'}|s_{t'-1}) \mathcal{L}_{nt'}^{(m)}(s_{t'}) \right) \right]}$$

Then, using the sample analog of (6.6) yields:

$$\pi^{(m+1)}(s'|s) = \frac{\sum_{n=1}^{\mathcal{N}} \sum_{t=2}^{\mathcal{T}} q_{ns't|s}^{(m+1)} q_{nst-1}^{(m+1)}}{\sum_{n=1}^{\mathcal{N}} \sum_{t=2}^{\mathcal{T}} q_{nst-1}^{(m+1)}} \quad (6.7)$$

6.1.4 Updating $p_{jt}^{(m)}(x, s)$

Following our motivating example, we propose two methods of updating the CCP's. One way is to use the current estimates of the model parameters coupled with the current conditional choice

probabilities. Generalizing equation (2.22), the value of $p_{jt}^{(m+1)}(x, s)$ at the $m + 1$ iteration is computed according to:

$$p_{jt}^{(m+1)}(x, s) = l_{jt}\left(x, s, \theta^{(m)}, \pi^{(m)}, p^{(m)}\right) \quad (6.8)$$

An alternative to updating CCP's with the model, is to use the data and the conditional probabilities of being in each of the unobserved states, q_{nst} . Substituting s_{nt} for s_n , and q_{nst} for q_{ns} , we can rewrite equations (2.13) through (2.16) for any choice j at time t to show that at the model's true parameters:

$$p_{jt}(x, s) = \frac{E[d_{njt}q_{nst}I(x_{nt} = x)]}{E[q_{nst}I(x_{nt} = x)]}$$

This formulation suggests a second way of updating $p_{jt}^{(m)}(x, s)$ is to use the weighted empirical likelihood:

$$p_{jt}^{(m+1)}(x, s) = \frac{\sum_{n=1}^{\mathcal{N}} d_{njt}q_{nst}^{(m+1)}I(x_{nt} = x)}{\sum_{n=1}^{\mathcal{N}} q_{nst}^{(m+1)}I(x_{nt} = x)} \quad (6.9)$$

6.2 Maximization Step

The primary benefit of the EM algorithm is that its maximization step is much simpler than the optimization problems defined in equations (5.4) through (5.8). Rather than maximizing a logarithm of weighted summed likelihoods—which requires integrating out over all the unobserved states—the maximization step treats the unobserved states as observed and weights each observation by $q_{nst}^{(m)}$. Thus $\theta^{(m+1)}$ is updated using:

$$\theta^{(m+1)} = \arg \max_{\theta} \sum_{n=1}^{\mathcal{N}} \sum_{t=1}^{\mathcal{T}} \sum_{s=1}^{\mathcal{S}} \sum_{j=1}^{\mathcal{J}} q_{nst}^{(m+1)} \ln \mathcal{L}_t\left(d_{nt}, x_{nt+1} | x_{nt}, s_{nt} = s; \theta, \pi^{(m+1)}, p^{(m+1)}\right) \quad (6.10)$$

6.3 Summary

We have now defined all the pieces necessary to implement the algorithm. It is triggered by setting initial 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 these initial values come from estimating a model without any unobserved heterogeneity and perturbing the estimates. Each iteration in the algorithm has four steps. Given $(\theta^{(m)}, \pi^{(m)}, p^{(m)})$ the $(m + 1)^{th}$ iteration follows:

Step 1 Compute $q_{nst}^{(m+1)}$, the conditional probabilities of being in each unobserved state, using (6.3).

Step 2 Compute $\pi^{(m+1)}(s_1|x_1)$, the set of initial population probabilities of the unobserved states, and $\pi^{(m+1)}(s'|s)$, the transition parameters for the unobserved states, using (6.5) and (6.7).

Step 3 Compute $p^{(m+1)}$, the conditional choice probabilities, using either (6.8) or (6.9).

Step 4 Compute the structural parameters, $\theta^{(m+1)}$, by solving (6.10).

By construction, the converged values satisfy the first order conditions of the estimators we defined in the previous section.

7 A Two Stage Estimator

The estimation approach described above uses information from both the observed state transitions as well as the underlying model of the decision process. When the CCP's are identified from the reduced form alone, we show it is possible to estimate the parameters governing the unobserved heterogeneity and the conditional choice probabilities in a first stage. The structural parameters of the dynamic discrete choice process are then estimated in a second stage.

One advantage of estimating the CCP's in a first stage is that the unrestricted flow utilities can then be computed directly from (4.3) and (4.4); imposing further model based structure can be tested in a straightforward way. Another advantage is that first stage estimates of the CCP's can be paired with any one of several CCP estimators already developed for models where there is no unobserved heterogeneity. For example, at the end of this section we show how to pair our first stage estimates with simulation-based estimators in the second stage.

We begin by partitioning the structural parameters into those that affect the observed state transitions, θ_1 , and those that do not, θ_2 , where $\theta \equiv \{\theta_1, \theta_2\}$. The transitions on the observed state variables are then given by $f_{jt}(x_{t+1}|x_t, s_t, \theta_1)$. We show how to estimate θ_1 , the initial probability distribution of the unobservables, $\pi(s_1|x_1)$, their transitions, $\pi(s'|s)$, and unrestricted estimates of the CCP's, $p_{jt}(x_t, s_t)$, in a first stage. Also estimated in the first stage are the conditional probabilities of being in an unobserved state in a particular time period, q_{nst} . In the second stage, we use the first stage estimates of θ_1 , $\pi(s'|s)$, q_{nst} , and $p_{jt}(x_t, s_t)$, to estimate the remaining parameters, θ_2 , which includes those parameters governing $u_{jt}(x_t, s_t)$, $G(\epsilon_t)$, and the discount factor β .

In the second stage, we use the first stage estimates of θ_1 , $\pi(s'|s)$, q_{nst} , and $p_{jt}(x_t, s_t)$, to estimate the remaining parameters, θ_2 , which includes those parameters governing $u_{jt}(x_t, s_t)$, $G(\epsilon_t)$, and the discount factor β . The two stage estimator presents an attractive option when there are additional outcomes—in this case x_{t+1} —that are affected by the unobserved state.

7.1 First stage

The function $l_{jt}(x_{nt}, s, \theta, \pi, p)$ imposes the structure of the economic model on the probability a particular choice is made. In the first stage we do not impose this structure and replace $l_{jt}(x_{nt}, s, \theta, \pi, p)$ in equation (5.2) with $p_{jt}(x_{nt}, s_{nt})$:

$$\mathcal{L}_t(d_{nt}, x_{nt+1}, |x_{nt}, s_{nt}; \theta_1, \pi, p) = \prod_{j=1}^J [p_{jt}(x_{nt}, s_{nt}) f_{jt}(x_{n,t+1}|x_{nt}, s_{nt}, \theta_1)]^{d_{njt}} \quad (7.1)$$

The updates for $q_{nst}^{(m)}$, $\pi^{(m)}(s_1|x_1)$, and $\pi^{(m)}(s'|s)$ are then defined by the same formulas used in the baseline algorithms, namely (6.3), (6.5) and (6.7). The only difference is that, when calculating the likelihood of (d_{nt}, x_{nt+1}) given x_{nt} , s_{nt} and the m th estimate of the parameters and conditional choice probabilities, we evaluate (7.1) at $(\theta_1^{(m)}, p^{(m)})$ which uses the empirical likelihood for the choices rather than the structural likelihood implied by the model.

At the m^{th} iteration, the maximization step then recovers $\theta_1^{(m+1)}$ and $p^{(m+1)}$ from:

$$\{\theta_1^{(m+1)}, p^{(m+1)}\} = \arg \max_{\theta_1, p} \sum_{n=1}^{\mathcal{N}} \sum_{t=1}^{\mathcal{T}} \sum_{s=1}^{\mathcal{S}} \sum_{j=1}^{\mathcal{J}} q_{nst}^{(m+1)} d_{njt} (\ln p_{jt}(x_{nt}, s) + \ln f_{jt}(x_{nt+1}|x_{nt}, s, \theta_1)) \quad (7.2)$$

This optimization problem is additively separable in θ_1 and p implying the parameters can be estimated in stages. $p_{jt}^{(m+1)}(x, s)$ has the closed form solution given by equation (6.9), the empirical update for the CCP's in the baseline algorithm. To prove this claim, note that for all j and j' in the choice set, the first order conditions can be expressed as:

$$\frac{\sum_{n=1}^{\mathcal{N}} q_{nst}^{(m)} d_{njt} I(x_{nt} = x)}{p_{jt}^{(m+1)}(x, s)} = \frac{\sum_{n=1}^{\mathcal{N}} q_{nst}^{(m)} d_{j'nt} I(x_{nt} = x)}{p_{j't}^{(m+1)}(x, s)}$$

Multiplying both sides by $p_{jt}^{(m+1)} p_{j't}^{(m+1)}$ and summing over all $j' \in \{1, \dots, J\}$, gives the result.

The first stage is then initiated by setting $\theta^{(1)}$, $p^{(1)}$, $\pi^{(1)}(s_1|x_1)$ and $\pi^{(1)}(s'|s)$. At the m th iteration, $q_{nst}^{(m+1)}$, $\pi^{(m+1)}(s_1|x_1)$, and $\pi^{(m+1)}(s'|s)$ are all updated as in the baseline algorithm but using the empirical likelihood for the choices. Then, $p^{(m+1)}$ is updated using (6.9) and $\theta_1^{(m+1)}$ is solved from (7.2). Since this is a standard EM algorithm, the log likelihood increases at each iteration and will converge to a local maximum.

7.2 Second stage

The first stage estimates for p are used in the second stage to obtain estimates for θ_2 , the parameters defining $u_{jt}(x_t, s_t)$, $G(\epsilon_t)$, and the discount factor β . A likelihood function could be formed to

estimate these remaining parameters, but simulation methods, such as those proposed by Hotz, Miller, Sanders, and Smith (1994) and Bajari, Benkard and Levin (2007) are also available.

Consider, for example, a modified version of the estimator of Hotz et al (1994).²³ From equation (4.2), the difference between any two conditional value functions is:

$$v_{jt}(x_{nt}, s_{nt}) - v_{kt}(x_{nt}, s_{nt}) = \psi_k [p_t(x_{nt}, s_{nt})] - \psi_j [p_t(x_{nt}, s_{nt})]$$

For each unobserved state we stack the $(J - 1)$ mappings from the conditional choice probabilities into the differences in conditional value functions for each individual n in each period t :

$$\begin{bmatrix} \psi_2 [p_t(x_{nt}, 1)] - \psi_1 [p_t(x_{nt}, 1)] - v_{2t}(x_{nt}, 1) - v_{1t}(x_{nt}, 1) \\ \vdots \\ \psi_J [p_t(x_{nt}, 1)] - \psi_1 [p_t(x_{nt}, 1)] - v_{Jt}(x_{nt}, 1) - v_{1t}(x_{nt}, 1) \\ \vdots \\ \psi_2 [p_t(x_{nt}, S)] - \psi_1 [p_t(x_{nt}, S)] - v_{2t}(x_{nt}, S) - v_{1t}(x_{nt}, S) \\ \vdots \\ \psi_J [p_t(x_{nt}, S)] - \psi_1 [p_t(x_{nt}, S)] - v_{Jt}(x_{nt}, S) - v_{1t}(x_{nt}, S) \end{bmatrix} = \begin{bmatrix} 0 \\ \vdots \\ 0 \\ \vdots \\ 0 \\ \vdots \\ 0 \end{bmatrix} \quad (7.3)$$

Second stage estimation is based on forward simulating the differences in conditional value functions arrayed in (7.3) to obtain their differences in terms of weighted sums of future utilities $u_{k\tau}(x_{n\tau}, s_{n\tau}) + \psi [p_t(x_{n\tau}, s_{n\tau})]$ for $\tau \geq t$. In the papers we cited above these differences are simulated for all $\tau \leq T$ or in the infinite horizon case, until these differences are negligible due to the combination of discounting and the simulated paths of future choices. A method of moments estimator is then formed by squaring a weighted sum over the sample population and minimizing it with respect to the structural parameters.

The simulation approaches in Hotz, Miller, Sanders, and Smith (1994) and Bajari, Benkard, and Levin (2007) involve simulating both the path of the state variables as well as the decisions significantly out into the future. With finite dependence, an alternative is to use the decision rules such that finite dependence holds and use simulation only for transitions of the state variables.²⁴ Hence, if the problem exhibits ρ -period dependence, only simulations over ρ periods are needed. In this case, forward simulation is particularly powerful as the paths are drawn from where the individual is currently at. For example, Bishop (2008) uses the techniques in this paper for a state

²³See Finger (2008) for an application of our two-stage estimator, where the second stage uses the estimator developed by Bajari, Benkard, and Levin (2007).

²⁴This has the added benefit of weakening the assumptions regarding the time horizon as well as how the state variables transition far out into the future.

space that has $1.12E + 184$ elements with finite dependence possible after three periods. By forward simulating, she is able to evaluate the future value term at a small subset of likely future states.

8 Small Sample Performance

To evaluate the small sample properties and the computational speed of our estimators, we now conduct two Monte Carlo studies. In each design, we illustrate the nature and extent of the problem that our estimators solve, by showing that unobserved heterogeneity would produce biased estimates if left unaccounted. Taken together, the two exercises cover finite and infinite horizon models, single agent problems and dynamic games, and cases where the unobserved state is fixed and when it varies over time.

8.1 Renewal and Finite Horizon

Our first Monte Carlo revisits the renewal problem described in Section 2, where the unobserved state is fixed over time. We describe the experimental design for this study, and then report our Monte Carlo results on the computational gains and the efficiency loss of our estimators relative to FIML. We also consider how well our estimator performs in nonstationary settings when the sample period falls short of the individual’s time horizon.

8.1.1 The bus engine problem revisited

We adapt the model discussed in Section 2 to a finite horizon setting, again normalizing the dependence of flow utility to zero when the engine is replaced. The payoff of keeping the current engine depends on the state s where $s \in \{1, 2\}$ and accumulated mileage, x_{1t} . Maintenance costs increase linearly with accumulated mileage up to 25 and then flatten out. Tracking accumulated mileage beyond 25 is therefore redundant. The flow payoff of keeping the current engine is then specified as:

$$u_2(x_{1t}, s) = \theta_0 + \theta_1 \min \{x_{1t}, 25\} + \theta_2 s \tag{8.1}$$

Mileage accumulates in increments of 0.125. Accumulated mileage depends on the decision to replace the engine, the previous mileage, and a permanent route characteristic of the bus denoted by x_2 . We assume that x_2 is a multiple of 0.01 and drawn from a discrete equiprobability distribution between 0.25 and 1.25. Higher values of x_2 are then associated with shorter trips or less frequent

use. The probability of x_{t+1} conditional on $x_t \equiv \{x_{1t}, x_2\}$ and $d_{jnt} = 1$, $f_j(x_{t+1}|x_t)$, is specified as:

$$f_j(x_{t+1}|x_t) = \begin{cases} \exp[-x_2(x_{1t+1} - x_{1t})] - \exp[-x_2(x_{1t+1} + 0.125 - x_{1t})] & \text{if } j = 2 \text{ and } x_{1t+1} \geq x_{1t} \\ \exp[-x_2(x_{1t+1})] - \exp[-x_2(x_{1t+1} + 0.125)] & \text{if } j = 1 \text{ and } x_{1t+1} \geq 0 \\ 0 & \text{otherwise} \end{cases} \quad (8.2)$$

implying that the mileage transitions follow a discrete analog of an exponential distribution. Since accumulated mileages above 25 are equivalent to mileages at 25 in the payoff function, we collapse all mileage transitions above 25 to 25, implying that $f_1(25|x_t) = \exp[-x_2(25)]$ and $f_2(25|x_t) = \exp[-x_2(25 - x_{1t})]$.

Relative to the simpler version we presented in Section 2, there are then three changes to the conditional value function. First, the conditional value function is now subscripted to reflect the finite horizon. Second, we have an intercept term on the flow payoff of running the engine. Finally, the mileage transitions are now stochastic. These modifications imply that the difference in conditional value function between running and replacing the engine given in (2.7) now become:

$$v_{2t}(x_t, s) - v_{1t}(x_t, s) = \theta_0 + \theta_1 \min \{x_{1t}, 25\} + \theta_2 s + \beta \sum_{x_{t+1}} \ln [p_{1t}(x_{t+1}, s)] [f_1(x_{t+1}|x_t) - f_2(x_{t+1}|x_t)] \quad (8.3)$$

where the sum over x_{t+1} goes from x_t to 25 in increments of 0.125. Since x_2 does not affect flow payoffs but does affect future utility through the mileage transitions, we can estimate β .

We simulate data for a decision-maker who lives for 30 periods and makes decisions on 1000 buses in each period. The data are generated by deriving the value functions at each state using backwards recursion. We then start each bus engine at zero miles and simulate the choices. The data are generated using two unobserved states with the initial of probabilities of each unobserved state set to 0.5. The econometrician is assumed to see only the last 20 periods, implying an initial conditions problem when s is unobserved. Summarizing the dimensions of this problem, there are two choices, twenty periods of data, two unobserved states, two hundred and one possible mileages and one hundred and one observed permanent characteristics. The number of states is therefore $20 \times 2 \times 201 \times 101 = 812,040$. Additional details regarding the data generating process and the estimation methods is in the appendix.

8.1.2 CCP versus FIML

The first column of Table 1 shows the parameters of the model. The next two columns show estimates from 50 simulations using both full information maximum likelihood and conditional

choice probabilities respectively when the type of the bus is observed. The conditional choice probabilities are estimated using a logit that is a flexible function of the state variables.²⁵ Both methods produced estimates centered around the true values. There is some loss of precision using the CCP estimator compared to FIML, but the standard deviations of the CCP estimates are all less than fifty percent higher than the FIML standard deviations and in most cases much less. These results are comparable to those found by Hotz et al (1994) and Aguirregabiria and Mira (2002) in their Monte Carlo studies of a similar problem.²⁶

Weighed against this loss of precision are the computational gains associated with CCP estimation. There are two reasons explaining why it is almost 1700 times faster than FIML. The full solution method solves the dynamic programming problem at each candidate value for the parameter estimates, whereas this estimator pairs a smoothed bin estimator of the CCP's (to handle sparse cells as explained in the appendix) with a logit that estimates the structural parameters. Second, the number of CCP's used to compute the CCP estimator is roughly proportional to the number of data points, because of the finite dependence property. In a typical empirical application, and also here, this number is dwarfed by the size of the state space which is the relevant benchmark for solving the dynamic optimization problems and FIML estimation.

Column 4 shows the results for CCP methods when bus type is unobserved but the heterogeneity in bus type is ignored in estimation. Averaging over the two unobserved states, the expected benefit of running a new engine is 2.5. The estimate of θ_0 when unobserved heterogeneity is ignored is lower than this due to dynamic selection. Since buses with $s = 1$ are less likely to be replaced, they are disproportionately represented at higher accumulated miles. As a result, the parameter on accumulated mileage, θ_1 , is also biased downward. These results suggest that neglecting to account for unobserved heterogeneity can induce serious bias, confirming for this structural dynamic framework early research on unobserved heterogeneity in duration models by Heckman (1981).

In Columns 5 and 6 we estimate the model treating s as unobserved. We used the second CCP estimator, which updates the CCP's using the estimated relative frequencies of being in the unobserved state, updated via a reduced form logit explained in the appendix. To handle the initial

²⁵The state space is too large to use a bin estimator. Alternatives to the flexible logit include nonparametric kernels and basis functions. We use flexible logits because of their computational convenience. See the appendix for details of the terms included in the logit.

²⁶The main difference between our specification and theirs is that we exploit the renewal property. This approach circumvents simulation into the indefinite future taken by Hotz et al (1994), and avoids inverting matrices with dimension equal to the number of elements in the state space (just over 40,000 in our case) followed by Aguirregabiria and Mira (2002).

conditions problem that emerges from only observing engine replacements after Zurcher has been operating his fleet for ten periods, the initial probability of being in an unobserved state (at period eleven) is estimated as a flexible function of the initial observed state variables.²⁷ Again, both FIML and CCP methods yielded estimates centered around the truth. There is surprisingly little precision loss from excluding the brand variable in the data and modeling s as unobserved. Our results also show the loss from using CCP estimator is smaller than FIML; all standard deviations are less than twenty five percent higher than the standard deviations of the corresponding FIML estimates. That the difference in precision shrinks occurs because some of the structure of the model is imposed in the CCP estimation through the probabilities of being in particular unobserved states.

Regardless of which method is used, treating bus brand as an unobserved state variable, rather than observed, increases computing time. The increased computational time in the CCP estimator is fully explained by the iterative nature of the EM algorithm, because each iteration essentially involves estimating the model as if the brand is observed.²⁸ Similarly, though not quite as transparently, FIML does not evaluate the likelihood of each bus history given the actual brand (as in Column 2), but the likelihood for both brands. This explains why estimation time for FIML essentially doubles (because there are two brands), increasing by more than two hours, whereas CCP increases by a factor of eighty, or by six and a half minutes.

8.1.3 Nonstationary problems with short panels

To further explore the properties of our estimators we made the exercises more complex, in the process precluding FIML estimation. As discussed in the Section 4, dynamic discrete choice models are partially identified even when there is no data on the latter periods of the optimization problem. In particular, suppose individuals have information about their future, never captured in the data, that affects their current choices. In life cycle contexts this might include details about inheritances, part time work opportunities following retirement, and prognoses of future health and survival. In this model, Zurcher might know much more about the time horizon and how engine replacement costs will evolve in the future. Both factors affect the optimal policy; he will tend to replace bus engines when they are cheaper, extending the life of some and shortening the life of others to take advantage of sales, and towards the end of the horizon he will become increasingly reluctant to

²⁷This is the same approach used in Keane and Wolpin (1997,2000,2001), Eckstein and Wolpin (1999), and Arcidiacono, Sieg, and Sloan (2007).

²⁸We did not explore the possibility of speeding up the EM algorithm using techniques developed by Jamshidian and Jennrich (1997) and others.

replace engines at all.

We now assume that the cost of running old engines and replacing them varies over time, and substitute a time dependent parameter θ_{0t} for θ_0 in equation (8.1). To emphasize the fact that the time shifters affect the costs for each bus the same way, we subscript bus-specific variables by n . Equation (8.1) becomes:

$$u_{2t}(x_{1nt}, s_n) - u_{1t}(x_{1nt}, s_n) = \theta_{0t} + \theta_1 x_{1nt} + \theta_2 s_n$$

We assume Zurcher knows both the value of T and the sequence $\{\theta_{01}, \dots, \theta_{0T}\}$. In contrast, the econometrician does not observe the values of θ_{0t} , and knows only that the data ends before T .

Implementing FIML amounts to estimating the value of the integer T , the sequence $\{\theta_{01}, \dots, \theta_{0T}\}$, as well as the parameters of interest $\{\theta_2, \theta_3, \beta\}$. If T comes from the (unbounded) set of positive integers, then the model is not identified for panels of fixed length. In practice, strong parametric assumptions must be placed on the length of the horizon, T , and how aggregate effects here modeled by θ_{0t} , evolve over time after the sample ends.²⁹ In contrast, implementing the second of our CCP estimators (that exploits the finite dependence property and updates using the estimated relative frequencies of the unobserved variable), does not require any assumptions about the process driving θ_{0t} or the horizon length. This is because the conditional choice probabilities from the last period of the data convey all the relevant information regarding the time horizon and the future values of θ_{0t} , enabling estimation of the flow payoffs up to one period before the sample ends.³⁰

We again set T to 30, but now assume the econometrician only observes periods 11 through 20, increasing the number of buses from 1000 to 2000 to keep the size of the data sets comparable to the previous Monte Carlos.³¹ Results with both s observed and unobserved are given in the last two columns of Table 1. When s is observed, adding time-varying intercepts has virtually no effect on the precision of the other parameters, or on the computational time. When s is unobserved, computation time almost doubles relative to CCP estimation with no time-varying effects (from 6.59 minutes to 11.31 minutes), and the standard deviations of the parameters increase by up to fifty percent.

²⁹Similarly the first CCP estimator is not feasible. To update $p_{1t}^{(m+1)}(x_{nt}, s_n)$ for period t with $l_1(x_{nt}, s_n, \theta^{(m)}, \pi^{(m)}, p^{(m)})$ using (6.8) we require $p_{1,t+1}^{(m)}(x_{n,t+1}, s_n)$. (See for example equation (2.8) in Section 2.). But an input $p_{1,t+2}^{(m-1)}(x_{n,t+2}, s_n)$ is required in $l_1(x_{n,t+1}, s_n, \theta^{(m-1)}, \pi^{(m-1)}, p^{(m-1)})$ to update $p_{1,t+1}^{(m)}(x_{n,t+1}, s_n)$ for period $t+1$, and so on. The upshot is that the full θ_{0t} sequence and also T is estimated, as in FIML.

³⁰More generally, in problems exhibiting finite dependence of ρ , the second estimator provides estimates up to ρ periods before the sample ends.

³¹The data in period 20 are only used to create the conditional choice probabilities used in the future value term for period 19. This is because the econometrician requires one-period-ahead CCP's to create differences in the conditional value functions for the current period.

However, the estimates are still centered around the truth and reasonably precise, demonstrating that structural estimation in these more demanding environments can be quite informative.

8.2 Dynamic Games

The second Monte Carlo experiment applies our estimators to infinite horizon dynamic games with private information. To motivate the exercise, we first show how our framework can be adapted to stationary infinite horizon games with incomplete information. We then apply our estimators to an entry and exit game.³² Here the unobserved states affect demand and evolve over time according to a Markov chain. Finally, we report on the performance of the baseline estimators as well as the alternative two-stage estimator developed in Section 7.

8.2.1 Adapting our framework to dynamic games

We assume that there are I firms in each of many markets, and that the systematic part of payoffs to the i^{th} firm in a market not only depends on its own choice in period t , denoted by $d_t^{(i)} \equiv (d_{1t}^{(i)}, \dots, d_{Jt}^{(i)})$, the state variables z_t , but also the choices of the other firms in that market, which we now denote by $d_t^{(-i)} \equiv (d_t^{(1)}, \dots, d_t^{(i-1)}, d_t^{(i+1)}, \dots, d_t^{(I)})$. Consistent with the games literature, we assume that the environment is stationary. Denote by $U_j^{(i)}(z_t, d_t^{(-i)}) + \epsilon_{jt}^{(i)}$ the current utility of firm i in period t , where $\epsilon_{jt}^{(i)}$ is an identically and independently distributed random variable that is private information to the firm. Although the firms all face the same observed state variables, these state variables will affect the firms in different ways. For example, a characteristic of firm i will affect the payoff for firm i differently than a characteristic of firm i' . Hence, the payoff function is superscripted by i .

Firms make simultaneous choices in each period. We denote $P(d_t^{(-i)} | z_t)$ as the probability firm i 's competitors choose $d_t^{(-i)}$ conditional on the state variables z_t . Since $\epsilon_t^{(i)}$ is independently distributed across all the firms, $P(d_t^{(-i)} | z_t)$ has the product representation:

$$P(d_t^{(-i)} | z_t) = \prod_{\substack{i'=1 \\ i' \neq i}}^I \left(\sum_{j=1}^J d_{jt}^{(i')} p_j^{(i')}(z_t) \right) \quad (8.4)$$

We impose rational expectations on the firm's beliefs about the choices of its competitors and assume firms are playing stationary Markov-perfect equilibrium strategies. Hence, the beliefs of the

³²The empirical literature on dynamic games with exit begins with Gowrisankaran and Town (1997). Two-step estimation of dynamic games with exit includes Beauchamp (2010), Beresteanu et al. (2010), Collard-Wexler (2008), Dunne et al. (2009), Gowrisankaran et al. (2009), and Ryan (2009), .

Table 1: Monte Carlo for Optimal Stopping Problem[†]

	Ignoring						Time Effects	Time Effects
	<i>s</i> observed			<i>s</i>	<i>s</i> unobserved		<i>s</i> observed	<i>s</i> unobserved
	DGP	FIML	CCP	CCP	FIML	CCP	CCP	CCP
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
θ_0 (Intercept)	2	2.0100 (0.0405)	1.9911 (0.0399)	2.4330 (0.0363)	2.0186 (0.1185)	2.0280 (0.1374)		
θ_1 (Mileage)	-0.15	-0.1488 (0.0074)	-0.1441 (0.0098)	-0.1339 (0.0102)	-0.1504 (0.0091)	-0.1484 (0.0111)	-0.1440 (0.0121)	-0.1514 (0.0136)
θ_2 (Unobs. State)	1	0.9945 (0.0611)	0.9726 (0.0668)		1.0073 (0.0919)	0.9953 (0.0985)	0.9683 (0.0636)	1.0067 (0.1417)
β (Discount Factor)	0.9	0.9102 (0.0411)	0.9099 (0.0554)	0.9115 (0.0591)	0.9004 (0.0473)	0.8979 (0.0585)	0.9172 (0.0639)	0.8870 (0.0752)
Time (Minutes)		130.29 (19.73)	0.078 (0.0041)	0.033 (0.0020)	275.01 (15.23)	6.59 (2.52)	0.079 (0.0047)	11.31 (5.71)

[†]Mean and standard deviations for fifty simulations. For columns (1)-(6), the observed data consist of 1000 buses for 20 periods. For columns (7)-(8), the intercept (θ_0) is allowed to vary over time and the data consist of 2000 buses for 10 periods. See text and appendix for additional details.

firm match the probabilities given in equation (8.4). Taking the expectation of $U_j^{(i)}(z_t, d_t^{(-i)})$ over $d_t^{(-i)}$, we define the systematic component of the current utility of firm i as a function of the firm's state variables as

$$u_j^{(i)}(z_t) = \sum_{d_t^{(-i)} \in J^{I-1}} P(d_t^{(-i)} | z_t) U_j^{(i)}(z_t, d_t^{(-i)}) \quad (8.5)$$

The values of the state variables at period $t+1$ are determined by the period t choices by both the firm and its competitors as well as the period t state variables. Denote $F_j(z_{t+1} | z_t, d_t^{(-i)})$ as the probability of z_{t+1} occurring given action j by firm i in period t , when its state variables are z_t and the other firms in its markets choose $d_t^{(-i)}$. The probability of transitioning from z_t to z_{t+1} given action j by firm i is then given by:

$$f_j^{(i)}(z_{t+1} | z_t) = \sum_{d_t^{(-i)} \in J^{I-1}} P(d_t^{(-i)} | z_t) F_j(z_{t+1} | z_t, d_t^{(-i)}) \quad (8.6)$$

The expressions for the conditional value function for firm i is then no different than what was described in Section 3 subject to the condition that we are now in a stationary environment. For example, equation (3.6) is modified in the stationary games environment to:³³

$$v_j^{(i)}(z_t) = u_j^{(i)}(z_t) + \beta \sum_{z_{t+1}=1}^Z \left[v_k^{(i)}(z_{t+1}) + \psi_k \left[p^{(i)}(z_{t+1}) \right] \right] f_j^{(i)}(z_{t+1} | z_t) \quad (8.7)$$

8.2.2 An entry and exit game

Our second Monte Carlo illustrates the small sample properties of our algorithms in a games environment. We analyze a game of entry and exit. In this game $d_t^{(i)} \equiv (d_{1t}^{(i)}, d_{2t}^{(i)})$ where $d_{1t}^{(i)} = 1$ means i exits the industry in period t , and $d_{2t}^{(i)} = 1$ means the firm is active, either as an entrant (when $d_{1,t-1}^{(i)} = 0$), or as a continuing incumbent (when $d_{1,t-1}^{(i)} = 1$). When a firm exits, it is replaced by a new potential entrant.

Following the notational convention in the rest of the paper, we partition the state variables of the firm, z_t , into those the econometrician observes, x_t , and the unobserved state variables, s_t . The observed state has two components. The first is a permanent market characteristic, denoted by x_1 , and is common across firms in the market. Each market faces an equal probability of drawing any of the possible values of x_1 where $x_1 \in \{1, 2, \dots, 10\}$. The second observed characteristic, x_{2t} , is whether or not each firm is an incumbent, $x_{2t} \equiv \{d_{2t-1}^{(1)}, \dots, d_{2t-1}^{(I)}\}$. Firms who are not incumbents

³³The results described earlier for the single agent case on finite dependence, the structure of the errors, and estimating conditional choice probabilities in the presence of unobserved states apply to games as well.

must pay a start up cost making it less likely that these firms will choose to stay in the market. The observed state variables are then $x_t \equiv \{x_1, x_{2t}\}$.

The unobserved state variable $s_t \in \{1, \dots, 5\}$, which we interpret as a demand shock, follows a first order Markov chain. We assume that the probability of the unobserved state remaining unchanged in successive periods is fixed $\pi \in (0, 1)$. If the state does change, any other state is equally likely to occur implying that the probability of s_{t+1} conditional on s_t when $s_t \neq s_{t+1}$ is $(1 - \pi) / 4$.

The flow payoff of firm i being active, net of private information $\epsilon_{2t}^{(i)}$, is modeled as:

$$U_2^{(i)}(x_t^{(i)}, s_t^{(i)}, d_t^{(-i)}) = \theta_0 + \theta_1 x_1 + \theta_2 s_t + \theta_3 \sum_{i'=1, i' \neq i}^I d_{2t}^{(i')} + \theta_4 d_{1,t-1}^{(i)} \quad (8.8)$$

We normalize $U_1^{(i)}(x_t^{(i)}, s_t^{(i)}, d_t^{(-i)}) = 0$. $U_2^{(i)}(x_t^{(i)}, s_t^{(i)}, d_t^{(-i)})$ is then used to form $u_2^{(i)}(x_t, s_t)$ by way of (8.5).

We assume that the firm's private information, $\epsilon_{jt}^{(i)}$, is distributed Type 1 extreme value. Since exiting is a terminal choice with the exit payoff normalized to zero, the Type 1 extreme value assumption and (8.7) imply that the conditional value function for being active is:

$$v_2^{(i)}(x_t, s_t) = u_2^{(i)}(x_t, s_t) - \beta \sum_{x_{t+1} \in X} \sum_{s_{t+1} \in S} \left(\ln \left[p_1^{(i)}(x_{t+1}, s_{t+1}) \right] \right) f_2^{(i)}(x_{t+1}, s_{t+1} | x_t, s_t) \quad (8.9)$$

The future value term is then expressed as a function solely of the one-period-ahead probabilities of exiting and the transition probabilities of the state variables.

We also generated price data on each market, denoted by y_t to capture the idea that unobserved demand shocks typically affect other outcomes apart from the observed decisions. Prices are a function of the permanent market characteristic, x_1 , the number of firms active in the market, and a shock denoted by η_t . We assume the shock follows a standard normal distribution and is independently distributed across markets and periods. The shock is revealed to each market after the entry and exit decisions are made. The price equation is then specified as:

$$y_t = \alpha_0 + \alpha_1 x_1 + \alpha_2 s_t + \alpha_3 \sum_{i=1}^I d_{1t}^{(i)} + \eta_t$$

The number of firms in each market is set to six and we simulated data for 3,000 markets. The discount factor is set at $\beta = 0.9$. Starting at an initial date with six potential entrants in the market, we ran the simulations forward for twenty periods. To show that our algorithms can easily be adapted to cases where there is an initial conditions problem, we used only the last ten periods to estimate the model. Initial probabilities of being in each unobserved state are again estimated

as a flexible function of the state variables in the first observed period.³⁴ A key difference between this Monte Carlo and the renewal Monte Carlo is that the conditional choice probabilities have an additional effect on both current utility and the transitions on the state variables due to the effect of the choices of the firm’s competitors on profits.

8.2.3 Results

The first column of Table 2 shows the parameters of the data generating process. The next two columns show what happens when s_t is observed and when it is ignored. When s_t is observed, all parameters are centered around the truth with the average estimation time being eight seconds. Column 3 shows that ignoring s_t results in misleading estimates of the effects of competition on prices and profits. The parameters in both the profit function and in the price equation on the number of competitors (θ_3 and α_3) are biased upward, significantly underestimating the effect of competition.³⁵

Column 4 reports our results for the estimator when the conditional choice probabilities are updated using the model. All estimates are centered around the truth with the average computational time being a little over 21 minutes. Column 5 updates the conditional choice probabilities with a reduced form logit of the type analyzed in the renewal problem. The standard deviations of the estimated profit parameters (the θ ’s) increase slightly relative to the case when the CCP’s are updated with the model. Computation time also increases by a little over twenty-five percent.

Column 6 presents results using the two-stage method. Here, the distribution for the unobserved heterogeneity and the CCP’s are estimated in a first stage along with the parameters governing the price process. The only parameters estimated in the second stage are those governing the current flow payoffs. The standard deviations of the coefficient estimates are similar to the case when the CCP’s are updated with the model. Computation time, however, is faster, averaging a little over fifteen minutes.

Finally, we consider the case when no price data is available. When data are only available on the discrete choices, is it still possible to estimate a specification where the unobserved states are allowed to transition over time? Column 7 shows that in some cases the answer is yes. Estimating the model without the price data and updating the CCP’s using the model again produces estimates that are centered around the truth with the standard deviations of the estimates similar to that

³⁴Further details of the reduced form for the initial conditions parameters can be found in the appendix.

³⁵Estimating the price process with market fixed effects did not change this result. In this case, α_3 was estimated to be around -.2, again underestimating the effects of competition.

of the two-stage method. The one parameter that is less precisely estimated is the persistence of the unobserved state. Computation time is also fast at a little less than seventeen minutes. These results would suggest that the additional information on prices is not particularly helpful. Note, however, that there are six entry/exit decisions for every one price observation. In unreported results we reduced the maximum number of firms in a market to two and in this case including price data substantially improved the precision of the estimates.

9 Conclusion

CCP methods can reduce the computational time of estimating dynamic discrete choice models. This paper extends the class of models that are easily adapted to the CCP framework, by broadening the set of dynamic discrete choice problems where few conditional choice probabilities are needed, as well as showing how to incorporate unobserved heterogeneity into CCP estimation.

We establish that future utility terms can always be expressed as function of conditional choice probabilities and the flow payoffs for any choice sequence. When two choice sequences with different initial choices lead to the same distribution of states after a few periods, then estimation requires only conditional choice probabilities for a few periods ahead. This finite dependence property makes it possible to identify utility function parameters in nonstationary environments where the sample period falls short of the time horizon. In this case, identification can be achieved even if the econometrician does not know the time horizon or when the transitions of the state variables are unknown beyond the sample period.

We further show how to accommodate unobserved heterogeneity via finite mixture distributions into CCP estimation. The computational simplicity of the estimator extends to unobserved state variables that follow a Markov chain. Our baseline algorithm iterates between updating the conditional probabilities of being in a particular unobserved state, updating the CCP's for any given state (observed and unobserved), and maximizing a likelihood function where the future values terms are in large part functions of the CCP's.

When the transition on the unobserved states and the CCP's are identified without imposing the structure of the underlying model, it is possible to estimate the parameters governing the unobserved heterogeneity in a first stage. We update the CCP's using the unrestricted distribution of discrete choices weighted by the estimated probabilities of being in particular unobserved states. This approach provides 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

Table 2: Monte Carlo for Entry/Exit Game[†]

	DGP	s_t observed	Ignore s_t	CCP-Model	CCP-Data	2-stage	No Prices
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
θ_0 (Intercept)	0	0.0207	-0.8627	0.0073	0.0126	-0.0251	-0.0086
		(0.0779)	(0.0511)	(0.0812)	(0.0997)	(0.1013)	(0.1083)
θ_1 (Obs. State)	0.05	-0.0505	-0.0118	-0.0500	-0.0502	-0.0487	-0.0495
Profit		(0.0028)	(0.0014)	(0.0029)	(0.0041)	(0.0039)	(0.0038)
Parameters θ_2 (Unobs. State)	0.25	0.2529		0.2502	0.2503	0.2456	0.2477
		(0.0080)		(0.0123)	(0.0148)	(0.0148)	(0.0158)
θ_3 (No. of Competitors)	-0.2	-0.2061	0.1081	-0.2019	-0.2029	-0.1926	-0.1971
		(0.0207)	(0.0115)	(0.0218)	(0.0278)	(0.0270)	(0.0294)
θ_4 (Entry Cost)	-1.5	-1.4992	-1.5715	-1.5014	-1.4992	-1.4995	-1.5007
		(0.0131)	(0.0133)	(0.0116)	(0.0133)	(0.0133)	(0.0139)
α_0 (Intercept)	7	6.9973	6.6571	6.9991	6.9952	6.9946	
		(0.0296)	(0.0281)	(0.0369)	(0.0333)	(0.0335)	
α_1 (Obs. State)	-0.1	-0.0998	-0.0754	-0.0995	-0.0996	-0.0996	
Price		(0.0023)	(0.0025)	(0.0028)	(0.0028)	(0.0028)	
Parameters α_2 (Unobs. State)	0.3	0.2996		0.2982	0.2993	0.2987	
		(0.0045)		(0.0119)	(0.0117)	(0.0116)	
α_3 (No. of Competitors)	-0.4	-0.3995	-0.2211	-0.3994	-0.3989	-0.3984	
		(0.0061)	(0.0051)	(0.0087)	(0.0088)	(0.0089)	
π (Persistence of Unobs. State)	0.7			0.7002	0.7030	0.7032	0.7007
				(0.0122)	(0.0146)	(0.0146)	(0.0184)
Time (Minutes)		0.1354	0.1078	21.54	27.30	15.37	16.92
		(0.0047)	(0.0010)	(1.5278)	(1.9160)	(0.8003)	(1.6467)

[†]Mean and standard deviations for 100 simulations. Observed data consist of 3000 markets for 10 periods with 6 firms in each market. In column (7), the CCP's are updated with the model. See text and appendix for additional details.

stage to recover the remaining structural parameters.

Our estimators are \sqrt{N} consistent and asymptotically normal. We undertake two Monte Carlo studies, modeling a dynamic optimization problem and a dynamic game, to investigate small sample performance. These studies indicate that substantial computational savings can result from using our estimators with little loss of precision.

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 (2002): “Swapping the Nested Fixed Point Algorithm: A Class of Estimators for Discrete Markov Decision Models,” *Econometrica*, 70, 1519-1543.
- [3] Aguirregabiria, V., and P. Mira (2007): “Sequential Estimation of Dynamic Discrete Games,” *Econometrica*, 75, 1-54.
- [4] Aguirregabiria, V., and P. Mira (2010): “Dynamic Discrete Choice Structural Models: A Survey,” *Journal of Econometrics*, 156(1), 38-67.
- [5] Arcidiacono, P. (2005): “Affirmative Action in Higher Education: How do Admission and Financial Aid Rules Affect Future Earnings?” *Econometrica*, 73, 1477-1524.
- [6] Arcidiacono, P., and J. B. Jones (2003): “Finite Mixture Distributions, Sequential Likelihood, and the EM Algorithm,” *Econometrica*, 71, 933-946.
- [7] Arcidiacono, P., H. Seig, and F. Sloan (2007): “Living Rationally Under the Volcano? An Empirical Analysis of Heavy Drinking and Smoking.” *International Economic Review*, 48(1).
- [8] Bajari, P., L. Benkard, and J. Levin (2007): “Estimating Dynamic Models of Imperfect Competition,” *Econometrica*, 75, 1331-1371.
- [9] Bajari, P., V. Chernozhukov, H. Hong, and D. Nekipelov (2009): “Identification and Efficient Semiparametric Estimation of a Dynamic Discrete Game,” working paper.
- [10] Beauchamp, A. (2010): “Abortion Supplier Dynamics”, working paper.
- [11] Beresteanu A., P.B. Ellickson, and S. Misra (2010): “The Dynamics of Retail Oligopoly,” working paper.

- [12] Bishop, K.C. (2008): "A Dynamic Model of Location Choice and Hedonic Valuation," working paper.
- [13] 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.
- [14] Buchinsky, M., J. Hahn, and V. J. Hotz (2005): "Cluster Analysis: A Tool for Preliminary Structural Analysis," working paper.
- [15] Chung, D., T. Steenburgh, and K. Sudhir (2009): "Do Bonuses Enhance Sales Productivity? A Dynamic Structural Analysis of Bonus-Based Compensation Plans," working paper.
- [16] Collard-Wexler, A. (2008): "Demand Fluctuations and Plant Turnover in the Ready-Mix Concrete Industry," working paper.
- [17] 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
- [18] Dunne, T., Klimek S., Roberts, M. and Y. Xu (2009): "Entry, Exit and the Determinants of Market Structure," *NBER Working Paper #15313*.
- [19] Eckstein, Z., and K. Wolpin (1999): "Why Youths Drop out of High School: The Impact of Preferences, Opportunities, and Abilities," *Econometrica*, 67(6), 1295-1339.
- [20] Finger, S.R. (2008): "Research and Development Competition in the Chemicals Industry," working paper.
- [21] Gowrisankaran, G., C. Lucarelli, P. Schmidt-Dengler, and R. Town (2009): "Government policy and the dynamics of market structure: Evidence from Critical Access Hospitals," working paper.
- [22] Gowrisankaran, G. and R. Town (1997): "Dynamic equilibrium in the hospital industry," *Journal of Economics, Management, and Strategy*, 6, 45-74.
- [23] 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.
- [24] Gayle, G., and R. A. Miller (2006): "Life Cycle Fertility and Human Capital Accumulation," working paper.

- [25] Hadley, G. (1961): *Linear Algebra*, Reading MA, Addison-Wesley Publishing Company
- [26] Hansen, L. P. (1982): “Large Sample Properties of Generalized Methods of Moments Estimators,” *Econometrica*, 50, 1029-1054.
- [27] Hamilton, J. D. (1990): “Analysis of Time Series Subject to Changes in Regime,” *Journal of Econometrics*, 45, 39-70.
- [28] Heckman, J.J. (1981): “The Incidental Parameters Problem and the Problem of Initial Conditions in Estimating a Discrete Time-Discrete Data Stochastic Process,” in *Structural Analysis of Discrete Data with Econometric Applications*, ed. C.F. Manski and D. McFadden. Cambridge, MA: MIT Press, 179-195.
- [29] Heckman, J.J., and S. Navarro (2007): “Dynamic Discrete Choice and Dynamic Treatment Effects,” *Journal of Econometrics*, 136, 341-396.
- [30] Heckman, J.J. and B. Singer (1984): “A Method for Minimizing the Impact of Distributional Assumptions in Econometric Models for Duration Data,” *Econometrica*, 52(2), 271-320.
- [31] Hotz, V. J., and R. A. Miller (1993): “Conditional Choice Probabilities and Estimation of Dynamic Models,” *Review of Economic Studies*, 61, 265-289.
- [32] 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.
- [33] Hu, Y. and M. Shum (2010a): “Nonparametric identification of dynamic models with unobserved state variables,” working paper.
- [34] Hu, Y. and M. Shum (2010b): “A Simple Estimator for Dynamic Models with Serially Correlated Unobservables,” working paper.
- [35] Imai, S., N. Jain, and A. Ching (2009): “Bayesian Estimation of Dynamic Discrete Choice Models,” *Econometrica*, 77(6), 1865-1899.
- [36] Jamshidian, Mortaza, and Robert Jennrich (1997): “Acceleration of the EM Algorithm by using Quasi-Newton Methods,” *Journal of the Royal Statistical Society*, B 59, pp. 569-587.
- [37] Joensen, J.S. (2009): “Academic and Labor Market Success: The Impact of Student Employment, Abilities, and Preferences,” working paper.

- [38] Kasahara, H., and K. Shimotsu (2009): “Nonparametric Identification and Estimation of Finite Mixture Models of Dynamic Discrete Choices,” *Econometrica*, 77(1), 135-175.
- [39] Kasahara, H. and K. Shimotsu (2008): “Pseudo-likelihood Estimation and Bootstrap Inference for Structural Discrete Markov Decision Models,” *Journal of Econometrics*, 146(1) 92-106.
- [40] 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.
- [41] Keane, M., and K. Wolpin (1997): “The Career Decisions of Young Men,” *Journal of Political Economy*, 105, 473-522.
- [42] Keane, M., and K. Wolpin (2001): “The Effect of Parental Transfers on Borrowing Constraints,” *International Economic Review*, 42(4), 1051-1103.
- [43] Keane, M., and K. Wolpin (2000): “Eliminating Race Differences in School Attainment and Labor Market Success,” *Journal of Labor Economics*.
- [44] Kennan, J. and J.R. Walker (forthcoming): “The Effect of Expected Income on Individual Migration Decisions,” *Econometrica*.
- [45] Magnac T. and D. Thesmar (2002): “Identifying Dynamic Discrete Decision Processes,” *Econometrica*, 70(2), 801-816.
- [46] 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.
- [47] McLachlan, G.J. and D. Peel (2000): *Finite Mixture Models*, Wiley series in probability and statistics, New York.
- [48] Miller, R. A. (1984): “Job Matching and Occupational Choice,” *Journal of Political Economy*, 92, 1086-1020.
- [49] Miller, R.A. and S. Sanders (1997): “Human capital development and welfare participation,” *Carnegie Rochester Conference Series on Public Policy*, Vol. 46, pp. 1-45.
- [50] Mroz, T. (1999): “Discrete Factor Approximations for Use in Simultaneous Equation Models: Estimating the Impact of a Dummy Endogenous Variable on a Continuous Outcome,” *Journal of Econometrics*, 129, 994-1001.

- [51] 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.
- [52] Norets, A. (2009): “Inference in Dynamic Discrete Choice Models with Serially Correlated Unobserved State Variables”, *Econometrica*, 77(5), 1665-1682.
- [53] Pakes, A. (1986): “Patents as Options: Some Estimates of the Value of Holding European Patent Stocks,” *Econometrica*, 54, 755-784.
- [54] Pakes, A., M. Ostrovsky, and S. Berry (2007): “ Simple Estimators for the Parameters of Discrete Dynamic Games (with Entry/Exit Examples),” *RAND Journal of Economics*, 38(2), 373-399.
- [55] Pesendorfer, M., and P. Schmidt-Dengler (2008): “ Asymptotic Least Square Estimators for Dynamic Games,” *Review of Economic Studies*, 75, 901-908.
- [56] Rust, J. (1987): “Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher,” *Econometrica*, 55, 999-1033.
- [57] Rust, J. (1994): “Structural Estimation of Markov Decision Processes,” in *Handbook of Econometrics, Volume 4*, ed. by R.E. Engle and D. McFadden. Amsterdam: Elsevier-North Holland, Chapter 14.
- [58] Ryan, S. (2009): “The Costs of Environmental Regulation in a Concentrated Industry,” working paper.
- [59] Su, C. and K.L. Judd (2009): “Constrained Optimization Approaches to Estimation of Structural Models,” working paper.
- [60] Wolpin, K. (1984): “An Estimable Dynamic Stochastic Model of Fertility and Child Mortality,” *Journal of Political Economy*, 92, 852-874.
- [61] Wu, C. F. (1983): “On the Convergence Properties of the EM Algorithm,” *The Annals of Statistics*, 11, 95-103.

A Proofs

Proof of Lemma 1. Equation (3.3) implies $V_t(z_t)$ can be written as:

$$\begin{aligned}
V_t(z_t) &= \sum_{j=1}^J \int d_{jt}^o(z_t, \epsilon_t) \left[u_{jt}(z_t) + \epsilon_{jt} + \beta \sum_{z_{t+1}=1}^Z V_{t+1}(z_{t+1}) f_{jt}(z_{t+1}|z_t) \right] g(\epsilon_t) d\epsilon_t \quad (\text{A.1}) \\
&= \sum_{j=1}^J p_{jt}(z_t) \left[u_{jt}(z_t) + \beta \sum_{z_{t+1}=1}^Z V_{t+1}(z_{t+1}) f_{jt}(z_{t+1}|z_t) \right] + \sum_{j=1}^J \int d_{jt}^o(z_t, \epsilon_t) \epsilon_{jt} g(\epsilon_t) d\epsilon_t \\
&= \sum_{j=1}^J p_{jt}(z_t) v_{jt}(z_t) + \sum_{j=1}^J \int d_{jt}^o(z_t, \epsilon_t) \epsilon_{jt} g(\epsilon_t) d\epsilon_t
\end{aligned}$$

Subtracting $v_{kt}(z_t)$ from both sides yields:

$$\begin{aligned}
V_t(z_t) - v_{kt}(z_t) &= \sum_{j=1}^J p_{jt}(z_t) v_{jt}(z_t) + \sum_{j=1}^J \int d_{jt}^o(z_t, \epsilon_t) \epsilon_{jt} g(\epsilon_t) d\epsilon_t - v_{kt}(z_t) \quad (\text{A.2}) \\
&= \sum_{j=1}^J p_{jt}(z_t) [v_{jt}(z_t) - v_{kt}(z_t)] + \sum_{j=1}^J \int d_{jt}^o(z_t, \epsilon_t) \epsilon_{jt} g(\epsilon_t) d\epsilon_t
\end{aligned}$$

From Proposition 1 of Hotz and Miller (1993, page 501), there exists a mapping $\psi_k^{(1)}(\rho)$ for each $j \in \{1, \dots, J\}$ such that:

$$\psi_j^{(1)}[p_t(z_t)] = v_{jt}(z_t) - v_{1t}(z_t) \quad (\text{A.3})$$

which implies:

$$v_{jt}(z_t) - v_{kt}(z_t) = \psi_j^{(1)}[p_t(z_t)] - \psi_k^{(1)}[p_t(z_t)] \quad (\text{A.4})$$

They also prove (A.3) implies there exists a mapping $\psi_j^{(2)}(\rho)$ for each $j \in \{1, \dots, J\}$ such that:

$$\psi_j^{(2)}[p_t(z_t)] = \int d_{jt}^o(z_t, \epsilon_t) \epsilon_{jt} g(\epsilon_t) d\epsilon_t \quad (\text{A.5})$$

Substituting (A.4) and (A.5) into (A.2) completes the proof:

$$\psi_k[p_t(z_t)] \equiv \sum_{j=1}^J p_{jt}(z_t) \left\{ \psi_j^{(1)}[p_t(z_t)] - \psi_k^{(1)}[p_t(z_t)] \right\} + \sum_{j=1}^J \psi_j^{(2)}[p_t(z_t)] = V_t(z_t) - v_{kt}(z_t) \quad (\text{A.6})$$

■

Proof of Theorem 1. The proof is by backward induction. We first establish that it holds when the time horizon is T and where the decision is made at T' and when $T' + 1 = T$. We then show that if it holds for a generic T' where $1 < T' < T$ then it also holds at $T' - 1$, completing the proof. Noting that $v_{kT}(z_T) \equiv u_k(z_T)$ for all $k \in \{1, \dots, J\}$ and $z_T \in \{1, \dots, Z\}$, including those in the decision rule $d_{kT'}^*(z_{T'}, j)$, and noting that when $T' + 1 = T$, equation (3.6) can be expressed as:

$$v_{jT'}(z_{T'}) = u_{jT'}(z_{T'}) + \beta \sum_{z_T=1}^Z \sum_{k=1}^J [u_{kT}(z_T) + \psi_k[p_T(z_T)]] d_{kT'}^*(z_{T'}, j) f_{jT'}(z_T|z_{T'}) \quad (\text{A.7})$$

establishing that the theorem holds for $t = T - 1$.

Setting T' such that $1 < T' < T$ and assuming (3.8) holds implies:

$$v_{jT'}(z_{T'}) = u_{jT'}(z_{T'}) + \sum_{\tau=T'+1}^T \sum_{k=1}^J \sum_{z_\tau=1}^Z \beta^{\tau-T'} [u_{k\tau}(z_\tau) + \psi_k[p_\tau(z_\tau)]] d_{k\tau}^*(z_\tau, j) \kappa_{\tau-1}^*(z_\tau | z_{T'}, j) \quad (\text{A.8})$$

Moving back to $T' - 1$, equation (3.6) implies:

$$v_{jT'-1}(z_{T'-1}) = u_{jT'-1}(z_{T'-1}) + \sum_{z_{T'-1}=1}^Z \sum_{k=1}^J [v_{kT'}(z_{T'} + \psi_k[p_{T'}(z_{T'})])] \kappa_{T'-1}^*(z_{T'} | z_{T'-1}, j) \quad (\text{A.9})$$

Substituting for $v_{kT'}(z_{T'})$ in (A.9) with (A.8) completes the proof.

$$v_{jT'-1}(z_{T'-1}) = u_{jT'-1}(z_{T'-1}) + \sum_{\tau=T'}^T \sum_{k=1}^J \sum_{z_\tau=1}^Z \beta^{\tau-T'-1} [u_{k\tau}(z_\tau) + \psi_k[p_\tau(z_\tau)]] d_{k\tau}^*(z_\tau, j) \kappa_{\tau-1}^*(z_\tau | z_{T'-1}, j) \quad (\text{A.10})$$

Now consider the infinite horizon problem. For $t < T'$, we can express $v_{jt}(z_t)$ as:

$$\begin{aligned} v_{jt}(z_t) &= u_{jt}(z_t) + \sum_{\tau=t+1}^{T'} \sum_{k=1}^J \sum_{z_\tau=1}^Z \beta^{\tau-t} [u_{k\tau}(z_\tau) + \psi_k[p_\tau(z_\tau)]] d_{k\tau}^*(z_\tau, j) \kappa_{\tau-1}^*(z_\tau | z_t, j) \\ &\quad + \sum_{z_{T'+1}=1}^Z V_{T'+1}(z_{T'+1}) \kappa_{T'}^*(z_{T'+1} | z_t, j) \end{aligned} \quad (\text{A.11})$$

We can bound $|V_{T'+1}(z_{T'+1})|$ by \bar{V} , which implies:

$$\left| \sum_{k=1}^J \sum_{z_{T'+1}=1}^Z V_{T'+1}(z_{T'+1}) \kappa_{T'}^*(z_{T'+1} | z_t, j) \right| \leq \sum_{k=1}^J \sum_{z_{T'+1}=1}^Z |V_{T'+1}(z_{T'+1})| \kappa_{T'}^*(z_{T'+1} | z_t, j) \leq \bar{V}$$

since

$$\sum_{z_{T'+1}=1}^Z \kappa_{T'}^*(z_{T'+1} | z_t, j) = 1 \quad \text{and} \quad \kappa_{T'}^*(z_{T'+1} | z_t, j) \geq 0 \quad \text{for all } \{z_{T'+1}, z_t\}$$

It now follows from (A.11) that for all T' :

$$\left| v_{jt}(z_t) - u_{jt}(z_t) - \sum_{\tau=t+1}^{T'} \sum_{k=1}^J \sum_{z_\tau=1}^Z \beta^{\tau-t} [u_{k\tau}(z_\tau) + \psi_k[p_\tau(z_\tau)]] d_{k\tau}^*(z_\tau, j) \kappa_{\tau-1}^*(z_\tau | z_t, j) \right| \leq \beta^{T'-t+1} \bar{V}$$

Since $\beta < 1$, the term $\beta^{T'-t+1} \bar{V} \rightarrow 0$ as $T' \rightarrow \infty$, proving the theorem. ■

Proof of Lemma 2. Define $\mathfrak{y}_j \equiv \ln Y_j$ and let $G_j(\varepsilon) \equiv \partial G(\varepsilon) / \varepsilon_j$. Let $\bar{\mathcal{H}} \equiv \mathcal{H}(e^{\mathfrak{y}_1}, e^{\mathfrak{y}_2}, \dots, e^{\mathfrak{y}_J})$.

Since $\mathcal{H}(Y_1, Y_2, \dots, Y_J)$ is homogeneous of degree one, and therefore the partial derivative $\mathcal{H}_j(Y_1, Y_2, \dots, Y_J)$ is homogenous of degree zero:

$$G_j(\mathfrak{y}_j + \varepsilon_j - \mathfrak{y}_1, \dots, \mathfrak{y}_j + \varepsilon_j - \mathfrak{y}_J) = \mathcal{H}_j(e^{\mathfrak{y}_1}, \dots, e^{\mathfrak{y}_J}) \exp[-\bar{\mathcal{H}} e^{-\mathfrak{y}_j - \varepsilon_j}] e^{-\mathfrak{y}_j - \varepsilon_j}$$

From Theorem 1 of McFadden (1978, page 80), integrating over $G_j(\varepsilon_t)$ yields the conditional choice probability:

$$\begin{aligned} p_j &= \int G_j(\vartheta_j + \varepsilon_j - \vartheta_1, \dots, \varepsilon_j, \dots, \vartheta_j + \varepsilon_j - \vartheta_J) d\varepsilon_j \\ &= e^{\vartheta_j - \vartheta_1} \mathcal{H}_j [1, e^{\vartheta_2 - \vartheta_1}, \dots, e^{\vartheta_J - \vartheta_1}] / \mathcal{H} [1, e^{\vartheta_2 - \vartheta_1}, \dots, e^{\vartheta_J - \vartheta_1}] \\ &\equiv \phi_j (1, e^{\vartheta_2 - \vartheta_1}, \dots, e^{\vartheta_J - \vartheta_1}) \end{aligned} \quad (\text{A.12})$$

By Proposition 1 of Hotz and Miller (1993) we can invert the vector function:

$$\begin{pmatrix} p_2 \\ \vdots \\ p_J \end{pmatrix} = \begin{pmatrix} \phi_2 (1, e^{\vartheta_2 - \vartheta_1}, \dots, e^{\vartheta_J - \vartheta_1}) \\ \vdots \\ \phi_J (1, e^{\vartheta_2 - \vartheta_1}, \dots, e^{\vartheta_J - \vartheta_1}) \end{pmatrix}$$

to make the vector of differences $(\vartheta_j - \vartheta_J)$ the subject of the equation. This is given by $\psi_j^{(1)}(\boldsymbol{p})$ from equation (A.3), implying that $\phi^{-1}(\boldsymbol{p})$ is given by:

$$\phi_j^{-1}(\boldsymbol{p}) = \exp \left[\psi_j^{(1)}(\boldsymbol{p}) \right] \quad (\text{A.13})$$

Noting that $\psi_1^{(1)}(\boldsymbol{p}) = 0$, we define $\phi_1^{-1}(\boldsymbol{p})$ as equal to one.

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

$$\begin{aligned} \int d_j \varepsilon_k dG(\varepsilon) &= \int \varepsilon_j G_j(\vartheta_j + \varepsilon_j - \vartheta_1, \dots, \vartheta_j + \varepsilon_j - \vartheta_J) d\varepsilon_j \\ &= \mathcal{H}_j(e^{\vartheta_1}, \dots, e^{\vartheta_J}) \int \varepsilon_j \exp[-\overline{\mathcal{H}} e^{-\vartheta_j - \varepsilon_j}] e^{-\varepsilon_j} d\varepsilon_j \\ &= e^{\vartheta_j - \vartheta_1} \mathcal{H}_j(1, e^{\vartheta_2 - \vartheta_1}, \dots, e^{\vartheta_J - \vartheta_1}) (\gamma - (\vartheta_j - \vartheta_1) + \ln \overline{\mathcal{H}}) / \overline{\mathcal{H}} \end{aligned}$$

Substituting the formula for p_j and evaluating $\overline{\mathcal{H}}$ at $\phi^{-1}(\boldsymbol{p})$ implies $\psi_j^{(2)}(\boldsymbol{p})$ from (A.5) can now be expressed as:

$$\psi_j^{(2)}(\boldsymbol{p}) = p_j \left[\gamma - \ln \phi_j^{-1}(\boldsymbol{p}) + \ln \mathcal{H} [1, \phi_2^{-1}(\boldsymbol{p}), \dots, \phi_J^{-1}(\boldsymbol{p})] \right] \quad (\text{A.14})$$

Substituting (A.13) and (A.14) into the definition of $\psi_j(\boldsymbol{p})$ given in equation (A.6) yields:

$$\psi_j(\boldsymbol{p}) = \sum_{k=1}^J p_k \left\{ \ln \phi_k^{-1}(\boldsymbol{p}) - \ln \phi_j^{-1}(\boldsymbol{p}) \right\} + \sum_{k=1}^J p_k (\gamma - \ln \phi_k^{-1}(\boldsymbol{p}) + \ln \mathcal{H} [1, \phi_2^{-1}(\boldsymbol{p}), \dots, \phi_J^{-1}(\boldsymbol{p})])$$

Simplifying the expression completes the proof:

$$\psi_j(\boldsymbol{p}) = \ln \mathcal{H} [1, \phi_2^{-1}(\boldsymbol{p}), \dots, \phi_J^{-1}(\boldsymbol{p})] - \ln \phi_j^{-1}(\boldsymbol{p}) + \gamma \quad (\text{A.15})$$

■

Proof of Lemma 3. From (3.18):

$$\mathcal{H}(Y_1, \dots, Y_J) = \mathcal{H}_0(Y_1, \dots, Y_K) + \left(\sum_{j \in \mathcal{J}} Y_j^{1/\sigma} \right)^\sigma \quad (\text{A.16})$$

For all $j \in \mathcal{J}$, the formula for $\phi_j(Y)$ for the nested logit components are:

$$\phi_j(Y) = Y_j^{1/\sigma} \frac{\left(\sum_{j \in \mathcal{J}} Y_j^{1/\sigma} \right)^{\sigma-1}}{\mathcal{H}(Y_1, \dots, Y_J)}$$

Let $\phi^{-1}(\boldsymbol{\rho}) \equiv (\phi_2^{-1}(\boldsymbol{\rho}), \dots, \phi_J^{-1}(\boldsymbol{\rho}))$ denote the inverse of $\phi(Y) \equiv (\phi_1(Y), \dots, \phi_{J-1}(Y))$.

Then from (A.12):

$$\rho_j \equiv \phi_j[\phi^{-1}(\boldsymbol{\rho})] = [\phi_j^{-1}(\boldsymbol{\rho})]^{1/\sigma} \frac{\left(\sum_{k \in \mathcal{J}} [\phi_k^{-1}(\boldsymbol{\rho})]^{1/\sigma} \right)^{\sigma-1}}{\mathcal{H}(1, \phi_2^{-1}(\boldsymbol{\rho}), \dots, \phi_J^{-1}(\boldsymbol{\rho}))} \quad (\text{A.17})$$

Summing over $k \in \mathcal{J}$ and taking the quotient yields:

$$\frac{\rho_j}{\sum_{k \in \mathcal{J}} \rho_k} = \frac{[\phi_j^{-1}(\boldsymbol{\rho})]^{1/\sigma}}{\sum_{k \in \mathcal{J}} [\phi_k^{-1}(\boldsymbol{\rho})]^{1/\sigma}}$$

which implies by direct verification that:

$$\phi_j^{-1}(\boldsymbol{\rho}) = A \rho_j^\sigma \quad (\text{A.18})$$

where A is unknown but greater than zero.

Substituting in for $\phi_j^{-1}(\boldsymbol{\rho})$ in (A.17) with (A.18), we obtain for each choice $j \in \mathcal{J}$:

$$\rho_j = A^{1/\sigma} \rho_j \frac{\left(\sum_{k \in \mathcal{J}} A^{1/\sigma} \rho_k \right)^{\sigma-1}}{\mathcal{H}(1, \phi_2^{-1}(\boldsymbol{\rho}), \dots, \phi_J^{-1}(\boldsymbol{\rho}))} = A \rho_j \frac{\left(\sum_{j \in \mathcal{J}} \rho_j \right)^{\sigma-1}}{\mathcal{H}(1, \phi_2^{-1}(\boldsymbol{\rho}), \dots, \phi_J^{-1}(\boldsymbol{\rho}))}$$

which implies:

$$\mathcal{H}(1, \phi_2^{-1}(\boldsymbol{\rho}), \dots, \phi_J^{-1}(\boldsymbol{\rho})) = A \left(\sum_{k \in \mathcal{J}} \rho_k \right)^{(\sigma-1)} \quad (\text{A.19})$$

We can now substitute in (A.19) and (A.18) into the expression for $\psi_j(\boldsymbol{\rho})$ given in (A.15), completing the proof.

$$\psi_j(\boldsymbol{\rho}) = \ln \left[A \left(\sum_{j \in \mathcal{J}} \rho_j \right)^{(\sigma-1)} \right] - \ln [A \rho_j^\sigma] + \gamma = \gamma - \sigma \ln(\rho_j) - (1 - \sigma) \ln \left(\sum_{k \in \mathcal{J}} \rho_k \right)$$

■

Proof of Theorem 2. Substituting in for $v_{jt}(z_t) - v_{1t}(z_t)$ in (4.1) with the corresponding expression in (4.2) implies:

$$\psi_1[p_t(z_t)] - \psi_j[p_t(z_t)] = u_{jt}(z_t) + \sum_{\tau=t+1}^T \sum_{z_\tau=1}^Z \beta^{\tau-t} \psi_1[p_\tau(z_\tau)] [\kappa_{\tau-1}^*(z_\tau|z_t, j) - \kappa_{\tau-1}^*(z_\tau|z_t, 1)]$$

Solving for $u_{jt}(z_t)$ completes the first part of the theorem:

$$u_{jt}(z_t) = \psi_1[p_t(z_t)] - \psi_j[p_t(z_t)] + \sum_{\tau=t+1}^T \sum_{z_\tau=1}^Z \beta^{\tau-t} \psi_1[p_\tau(z_\tau)] [\kappa_{\tau-1}^*(z_\tau|z_t, 1) - \kappa_{\tau-1}^*(z_\tau|z_t, j)] \quad (\text{A.20})$$

To prove the second part, note that the two decision sequences set the initial choices such that $d_{jt} = 1$ or $d_{1t} = 1$ and then both decision sequences set $d_{1t'} = 1$ for all $t' > t$. From the definition of F_1 , the columns of F_1^τ gives the probabilities of being in each state after τ periods conditional choosing alternative 1 in each of those periods. The rows indicate how these probabilities differ given the initial state. Hence, for $\tau \geq 1$, the (z, z') element of F_1^τ is $\kappa_{t+\tau-1}^*(z'|z, 1)$. Similarly, the (z, z') element of $F_j F_1^\tau$ is $\kappa_{t+\tau-1}^*(z'|z, j)$.

Using the matrix notation defined in the theorem, we can express u_j as:

$$u_j = \Psi_j - \Psi_1 + \sum_{\tau=1}^{\infty} \beta^\tau (F_1 - F_j) F_1^{\tau-1} \Psi_1 = \Psi_j - \Psi_1 + \beta (F_1 - F_j) \left(\sum_{\tau=0}^{\infty} \beta^\tau F_1^\tau \right) \Psi_1 \quad (\text{A.21})$$

Noting that $\beta f_j(z'|z) > 0$ for all (j, z, z') and $\beta \sum_{z'=1}^Z f_j(z'|z) = \beta < 1$ for all (j, z) , the existence of $[\mathcal{I} - \beta F_1]^{-1}$ follows from Hadley (page 118, 1961) with:

$$Q \equiv \sum_{\tau=0}^{\infty} \beta^\tau F_1^\tau = \mathcal{I} + \beta Q F_1 = [\mathcal{I} - \beta F_1]^{-1}$$

Substituting the expression for Q into (A.21) we obtain:

$$u_j = \Psi_j - \Psi_1 + \beta (F_1 - F_j) [\mathcal{I} - \beta F_1]^{-1} \Psi_1$$

which proves the lemma. ■

Proof of Theorem 3. Proof of Part 1. For convenience we consolidate the structural parameters into the vector $\lambda \equiv (\theta, \pi)$. Denote the true parameters and conditional choice probabilities by λ_0 and p_0 respectively. Let $l(\lambda, p)$ denote the corresponding vector of likelihoods associated each choice probability, implying $p_0 = l(\lambda_0, p_0)$. For each \mathcal{N} define $\Lambda_{\mathcal{N}}$ as the set of parameters solving (5.5) at $p = \hat{p}$, where $(\hat{\theta}, \hat{\pi}, \hat{p})$ simultaneously satisfies (5.6):

$$\Lambda_{\mathcal{N}} \equiv \left\{ \lambda_{\mathcal{N}} : \lambda_{\mathcal{N}} = \arg \max_{\lambda} \frac{1}{N} \sum_{n=1}^N \ln [L(d_n, x_n | x_{n1}; \lambda, p_{\mathcal{N}})] \text{ where } p_{\mathcal{N}} = l(\lambda_{\mathcal{N}}, p_{\mathcal{N}}) \right\}$$

Also define the set of parameters that maximize the corresponding expected log likelihood subject to the same constraint as:

$$\Lambda_1 \equiv \left\{ \lambda_1 : \lambda_1 = \arg \max_{\lambda} E \{ \ln [L(d_n, x_n | x_{n1}; \lambda, p_1)] \} \text{ where } p_1 = l(\lambda_1, p_1) \right\}$$

By definition $\lambda_0 \in \Lambda_1$ because (λ_0, p_0) solves:

$$\lambda_0 = \arg \max_{\lambda} E \{ \ln [L(d_n, x_n | x_{n1}; \lambda, p_0)] \} \text{ where } p_0 = l(\lambda_0, p_0)$$

Appealing to the continuity of $L(d_n, x_n | x_{n1}; \lambda, p_N)$ and $p(\lambda_N)$, the weak uniform law of large numbers implies there exists a sequence $\hat{\lambda}_N \in \Lambda_N$ converging to λ_0 . Now consider sequences $\tilde{\lambda}_N \in \Lambda_N$ that converge to other elements in Λ_1 , say $\lambda_1 \neq \lambda_0$. The assumption of identification implies that for all $\lambda_1 \neq \lambda_0$:

$$E \{ \ln [L(d_n, x_n | x_{n1}; \lambda_0, p_0)] \} > E \{ \ln [L(d_n, x_n | x_{n1}; \lambda_1, p_1)] \}$$

By continuity and the law of large numbers:

$$\frac{1}{N} \sum_{n=1}^N \ln [L(d_n, x_n | x_{n1}; \hat{\lambda}_N, \hat{p}_N)] > \frac{1}{N} \sum_{n=1}^N \ln [L(d_n, x_n | x_{n1}; \tilde{\lambda}_N, \tilde{p}_N)] + o_p(1)$$

This proves that choosing the element which maximizes the criterion function, $\hat{\lambda}_N$, from the set of fixed points, Λ_N , selects a weakly consistent estimator for λ_0 . ■

Proof of Part 2. For each t define the joint distribution of (x, s) , induced by the parameter vector (λ, p) and the data, as:

$$P_{Nt}(x, s, \hat{\lambda}, \hat{p}) \equiv \frac{1}{N} \sum_{n=1}^N \left[I(x_{nt} = x) \frac{\hat{L}_{nt}(s_{nt} = s)}{\hat{L}_n} \right]$$

By the law of large numbers, for each x , the $X \times S - 1$ dimensional random variable $P_{Nt}(x, s, \hat{\lambda}, \hat{p})$ converges in probability to:

$$P_t(x, s, \hat{\lambda}, \hat{p}) \equiv E \left[I(x_{nt} = x) \frac{\hat{L}_{nt}(s_{nt} = s)}{\hat{L}_n} \right]$$

Similarly the joint distribution of (j, x, s) is defined at t as:

$$P_{Nt}(j, x, s, \hat{\lambda}, \hat{p}) \equiv \frac{1}{N} \sum_{n=1}^N \left[d_{njt} I(x_{nt} = x) \frac{\hat{L}_{nt}(s_{nt} = s)}{\hat{L}_n} \right]$$

which converges in probability to:

$$P_t(j, x, s, \hat{\lambda}, \hat{p}) \equiv E \left[I(d_{njt} = 1) I(x_{nt} = x) \frac{\hat{L}_{nt}(s_{nt} = s)}{\hat{L}_n} \right]$$

Let $P_N(\lambda, p)$ denote the $T \times (J - 1) \times X \times S$ dimensional vector formed from components $P_{Nt}(j, x, s, \lambda, p) / P_{Nt}(x, s, \lambda, p)$, and let $P(\lambda, p)$ denote the vector of corresponding limit points

$P_t(j, x, s, \lambda, p) / P_t(x, s, \lambda, p)$. Then the parameters solving the fixed point characterized by (5.5) and (5.8) are elements of the set defined by:

$$\Lambda'_N \equiv \left\{ \lambda_N : \lambda_N = \arg \max_{\lambda} \frac{1}{N} \sum_{n=1}^N \ln [L(d_n, x_n | x_{n1}; \lambda, p_N)] \text{ where } p_N = P_N(\lambda_N, p_N) \right\}$$

and similar to Part 1, elements in Λ'_N converge weakly to elements in the set:

$$\Lambda'_1 \equiv \left\{ \lambda_1 : \lambda_1 = \arg \max_{\lambda} E \{ \ln [L(d_n, x_n | x_{n1}; \lambda, p_1)] \} \text{ where } p_1 = P(\lambda_1, p_1) \right\}$$

Noting that $(\lambda_0, p_0) \in \Lambda'_1$, the arguments in Part 1 can be repeated to complete the proof that the fixed point solution in Λ'_N achieving the highest value of (5.4) is consistent. ■

A.1 Asymptotic Covariance Matrix

The asymptotic covariance matrix of our estimators are derived from Taylor expansions of two sets of equations, the first order conditions of (5.5) for λ , and a set of equations that solve the conditional choice probability nuisance parameter vector p . The first order conditions of (5.5) can be written as:

$$\frac{1}{N} \sum_{n=1}^N L_{\lambda n}(\hat{\lambda}, \hat{p}) = 0 \quad (\text{A.22})$$

where

$$L_{\lambda n}(\lambda, p) \equiv \frac{\partial [\ln L(d_n, x_n | x_{n1}; \lambda, p)]}{\partial \lambda}$$

Since $(\hat{\lambda}, \hat{p})$ is consistent and $L_{\lambda n}(\lambda, p)$ is continuously differentiable, we can expand (A.22) around (λ_0, p_0) to obtain:

$$N^{-1/2} \sum_{n=1}^N L_{\lambda n}(\lambda_0, p_0) - A_{\lambda} \sqrt{N} (\hat{\lambda} - \lambda_0) - A_p \sqrt{N} (\hat{p} - p_0) = o_p(1) \quad (\text{A.23})$$

where

$$A_{\lambda} \equiv \lim_{N \rightarrow \infty} \left[\frac{1}{N} \sum_{n=1}^N \frac{\partial L_{\lambda n}(\lambda_0, p_0)}{\partial \lambda} \right], \quad A_p \equiv \lim_{N \rightarrow \infty} \left[\frac{1}{N} \sum_{n=1}^N \frac{\partial L_{\lambda n}(\lambda_0, p_0)}{\partial p} \right]$$

The first estimator sets \hat{p} to solve (5.6) for each (j, t, x, s) . Stacking $l_{jt}(x, s; \lambda, p)$ for each choice (j, t) (time indexed in the nonstationary case), and each value (x, s) of state variables to form $l(\lambda, p)$, the $(J-1) \times T \times X \times S$ vector function of the CCP parameters (λ, p) , our estimator satisfies the $(J-1)TXS$ additional parametric restrictions $l(\hat{\lambda}, \hat{p}) = \hat{p}$. From the identity

$$0 = l(\hat{\lambda}, \hat{p}) - \hat{p} = l(\lambda_0, p_0) - p_0$$

we expand the second equation to the first order and rearrange to obtain:

$$(I - l_p) \sqrt{N} (\hat{p} - p_0) - l_{\lambda} \sqrt{N} (\hat{\lambda} - \lambda_0) = o_p(1) \quad (\text{A.24})$$

where

$$l_\lambda \equiv \frac{\partial l(\lambda_0, p_0)}{\partial \lambda}, \quad l_p \equiv \frac{\partial l(\lambda_0, p_0)}{\partial p}$$

Using (A.24) we substitute out $\sqrt{\mathcal{N}}(\hat{p} - p_0)$ in (A.23), which yields:

$$\sqrt{\mathcal{N}}(\hat{\lambda} - \lambda_0) = (B_1' B_1)^{-1} B_1' N^{-1/2} \sum_{n=1}^{\mathcal{N}} L_{\lambda n}(\lambda_0, p_0) + o_p(1)$$

where

$$B_1 \equiv A_\lambda + A_p (I - l_p)^{-1} l_\lambda$$

Appealing to the central limit theorem, and the fact that:

$$A_\lambda \equiv E [L_{\lambda n}(\lambda_0, p_0) L_{\lambda n}(\lambda_0, p_0)']$$

the asymptotic covariance matrix for $\sqrt{\mathcal{N}}(\hat{\lambda} - \lambda_0)$ is thus:

$$(B_1' B_1)^{-1} B_1' A_\lambda B_1 (B_1 B_1')^{-1}$$

In the second estimator, the condition that $l(\hat{\lambda}, \hat{p}) = \hat{p}$ is replaced by the $(J - 1)$ *TXS* equalities in (5.8). Define:

$$Q_{njtxs}(\lambda, p) \equiv [p_{jt}(x, s) - d_{njt}] I(x = x_{nt}) \frac{L_n(s_{nt} = s)}{L_n}$$

where $L_n \equiv L(d_n, x_n | x_{n1}; \lambda, p)$, and $L_n(s_{nt} = s)$ is given by (5.7) evaluated at (λ, p) . For each sample observation n , stack $Q_{njtxs}(\lambda, p)$ to form the $(J - 1)$ *TXS* dimensional vector $Q_n(\lambda, p)$. In vector form (5.8) can then be expressed as:

$$\frac{1}{\mathcal{N}} \sum_{n=1}^{\mathcal{N}} Q_n(\hat{\lambda}, \hat{p}) = 0$$

We form the vector $h_n(\lambda, p)$, the expected outer product of $h_n(\lambda, p)$, and its square derivative matrix:

$$h_n(\lambda, p) \equiv \begin{bmatrix} L_{\lambda n}(\lambda, p) \\ Q_n(\lambda, p) \end{bmatrix}, \quad \Omega = E [h_n(\lambda_0, p_0) h_n(\lambda_0, p_0)'], \quad \Gamma = E \left[\frac{\partial h_n(\lambda_0, p_0)}{\partial \lambda} \frac{\partial h_n(\lambda_0, p_0)}{\partial p} \right]$$

From Hansen (1982, Theorem 3.1) or Newey and McFadden (1994, Theorem 6.1), it now follows that $\sqrt{\mathcal{N}}(\hat{\lambda} - \lambda_0)$ is asymptotically normally distributed with mean zero and covariance matrix given by the top left square block of $\Gamma^{-1} \Omega \Gamma^{-1'}$ with dimension of λ .

B Additional Information on the Monte Carlo Exercises

All simulations were conducted in Matlab version 7.5 on the Duke economics department 64-bit batch cluster. The code was not parallelized. The cluster and the operating system of Matlab ensures one processor is dedicated to each Matlab job. All non-linear optimization was done using Matlab's canned optimization routine, `fminunc` with the default values used to determine convergence. No derivatives were used in the maximization routines for the structural parameters. Convergence for the EM algorithm was determined by comparing log likelihood values 25 iterations apart. The algorithm was stopped when this difference was less than 10^{-7} for two successive iterations.

B.1 Optimal Stopping

This subsection provides further computational details about the optimal stopping problem. We discuss the data generating process, as well as updating the conditional choice probabilities and the parameters governing the initial conditions.

B.1.1 Data Creation

For the true parameter values and the transition matrix for mileages implied by (8.2) and reported in the first column of Table 1, we obtain the value functions by backwards recursion for every possible mileage, observed permanent characteristic, unobserved state, and time. We draw permanent observed and unobserved characteristics from discrete uniform distributions with support 101 and 2 respectively, and start each bus at zero mileage. Given the parameters of the utility function, the value function, and the permanent observed and unobserved states, we calculate the probability of a replacement occurring in the first period. We then draw from a standard uniform distribution. If the draw is less than the probability of replacement, the decision in the first period is to replace. Otherwise we keep the engine. Conditional on the replacement decision, we draw a mileage transition using (8.2). Continuing this way, decisions and mileage transitions are simulated for thirty periods.

B.1.2 The likelihood

Conditional on the permanent observed state, the mileage and the unobserved state s , the likelihood of a particular decision at time t takes a logit form. The likelihoods for the FIML and CCP cases

are respectively given by:

$$\begin{aligned}\mathcal{L}_t(d_t|x_t, s; \theta) &= \frac{d_{1t} + d_{2t} \exp \left[u_2(x_t, s, \theta) + \beta \sum_{x_{t+1}} V(x_{t+1}, s, \theta)(f_2(x_{t+1}|x_t) - f_1(x_{t+1}|x_t)) \right]}{\exp \left[u_2(x_t, s, \theta) + \beta \sum_{x_{t+1}} V(x_{t+1}, s, \theta)(f_2(x_{t+1}|x_t) - f_1(x_{t+1}|x_t)) \right] + 1} \\ \mathcal{L}_t(d_t|x_t, s, p; \theta) &= \frac{d_{1t} + d_{2t} \exp \left[u_2(x_t, s, \theta) - \beta \sum_{x_{t+1}} \ln [p_{1t+1}(x_{t+1}, s)] (f_2(x_{t+1}|x_t) - f_1(x_{t+1}|x_t)) \right]}{\exp \left[u_2(x_t, s, \theta) - \beta \sum_{x_{t+1}} \ln [p_{1t+1}(x_{t+1}, s)] (f_2(x_{t+1}|x_t) - f_1(x_{t+1}|x_t)) \right] + 1}\end{aligned}$$

When s is unobserved, the log likelihood for a particular bus history is found by first taking the product of the likelihoods over time conditional on type, and then summing across the types inside the logarithm. Thus in the FIML case, the likelihood is:³⁶

$$L(d_n|x_n; \theta, \pi) = \sum_{s=1}^2 \prod_{t=11}^{30} \pi(s|x_1) \mathcal{L}_t(d_t|x_t, s; \theta)$$

In the CCP case, $\mathcal{L}_t(d_t|x_t, s; \theta)$ is replaced by $\mathcal{L}_t(d_t|x_t, s, p; \theta)$.

B.1.3 Conditional choice probability estimates

We approximate (6.9), the second estimator for the conditional choice probabilities with a flexible logit, where the dependent variable is d_{1t} . There are five cases:

1. To obtain the estimates reported in Column 4 of Table 1 (when s is ignored), we estimate the CCP's using $W_{1t} \equiv (1, x_{1t}, x_{1t}^2, x_2, x_2^2, x_{1t}x_2)$ as regressors in a logit.³⁷
2. For the parameters reported in Column 3, W_{1t} is fully interacted with $W_{2t} \equiv (1, s, t, st, t^2, st^2)$, that is 36 parameters to estimate in the logit generating the CCP's. Since s is observed, this flexible logit is estimated once.
3. When s is unobserved, the flexible logit described in the previous case is estimated at each iteration of the EM algorithm; at the m^{th} iteration the conditional probabilities of being in each observed state, $q_s^{(m)}$, are used to weight the flexible logit.
4. For the last two columns, where there are aggregate effects, we fully interact the first set of variables with the s , but not t and t^2 . Instead, we include time dummies, but given the moderate sample size, we did not interact them with the other variables. Hence the logits were estimated with 21 regressors: 12 combinations of x_{1t} , x_2 , and s , as well as 9 time dummies.

³⁶Since we are taking products of potentially small probabilities, numerical issues could arise. These can be solved by scaling up the $\mathcal{L}_t(d_t|x_t, s; \theta, p)$ terms by a constant factor. However, in neither of our Monte Carlos was this an issue.

³⁷For a sufficiently large but finite sample we can saturate the finite set of regressors with a flexible logit that yields numerically identical estimates as the weighted bin estimator presented in the text.

B.1.4 Initial conditions

Initial probabilities are specified as a flexible function of the first period observables, denoted by W_0 . Included in W_0 are the mileage at the first observed time period for the bus, x_{11} , as well as the permanent observed characteristic, x_2 . The prior probability of being in unobserved state 2 during the first observed period in the data, $t = 1$, given the data for n is given by:

$$\pi(2|x_1) = \frac{\exp(W_0\delta)}{1 + \exp(W_0\delta)}$$

At iteration m , we calculate the likelihood for each data point conditional on the unobserved state. Under FIML:

$$L(d_n, s_n = s|x_n; \theta^{(m)}) = \prod_{t=11}^{30} \pi(s|x_1) \mathcal{L}_t(d_t|x_t, s; \theta^{(m)})$$

The iterate $\delta^{(m+1)}$ solves:

$$\delta^{(m+1)} = \arg \max_{\delta} \sum_{n=1}^{\mathcal{N}} \ln \left(\sum_{s=1}^2 \pi(s|x_{n1}) L(d_n, s_n = s|x_n; \theta^{(m)}) \right)$$

In the CCP case, we replace $\mathcal{L}_t(d_t|x_t, s; \theta^{(m)})$ with $\mathcal{L}_t(d_t|x_t, s, p^{(m)}; \theta^{(m)})$, and replace $L(d_n, s_n = s|x_n; \theta^{(m)})$ with $L(d_n, s_n = s|x_n, p^{(m)}; \theta^{(m)})$.³⁸

B.1.5 Creation of time-varying intercepts

In the case where the replacement costs varied over time (Column 8 of Table 1), we create the data by drawing values for the intercept from a normal distribution with standard deviation of 0.5. The value of θ_{0t+1} , is set to $0.7\theta_{0t}$ plus the value drawn at $t + 1$ from the normal distribution.

B.2 Entry/Exit

We now turn to the details of the Monte Carlo for the dynamic game. Again we describe the data creation as well as the variables used in both the conditional choice probabilities and in the reduced form controls for the initial conditions problem.

B.2.1 Data creation

The first step in creating the data is to obtain the probability of entering for every state. Equation (8.8) gives the flow payoff for being in the market conditional on the choices of the other firms. Note that the expected flow payoff of entering depends on the probabilities of other firms entering. Given

³⁸The saturation argument we mentioned in the previous footnote applies here too.

initial guesses for the probability of exiting in each state, we form all the possible combinations of the entry decisions of the other firms using (8.4). We then substitute (8.4) and (8.8) into (8.5) to form the expected flow payoff of staying in or entering the market in every state. Since the transitions on the state variables conditional on the entry/exit decisions are known, we have all the pieces to form (8.9). Given (8.9), the Type 1 extreme value assumption implies the probability of exiting is $1/(1 + \exp(v_2^{(i)}(x_t, s_t)))$. We can then update the entry exit probabilities used to form (8.4). We then iterate on (8.4), (8.5), (8.9), and the logit probability of exiting until a fixed point is reached.³⁹

The observed permanent market characteristics and the initial unobserved states were drawn from a discrete uniform distribution. We then began each market with no incumbents and simulated the model forward. We then removed the first ten periods of data from the sample.

B.2.2 The likelihood

We now derive the likelihood at time t for market n of the observed decisions and price process given the data and the parameters. Note that x_{nt+1} , which includes the permanent market characteristic as well as the incumbency status of each of the firms, is a deterministic function of x_{nt} and y_{nt} . The likelihood contribution for the i^{th} firm at time t conditional on unobserved state s_t is:

$$l^{(i)}(d_t^{(i)}, x_t, s_t; \theta, \pi) = \frac{d_{1t} + d_{2t} \exp[v_2^{(i)}(x_t, s_t, p, \theta, \pi)]}{1 + \exp[v_2^{(i)}(x_t, s_t, p, \theta, \pi)]}$$

Denote $E(y_t) = \alpha_0 + \alpha_1 x_1 + \alpha_2 s_t + \alpha_3 \sum_{i=1}^I d_{2t}^{(i)}$. Denoting n as the market, the likelihood of the data in market n at time t conditional on s_t is:

$$\mathcal{L}_t(d_{nt}, y_{nt} | x_{nt}, s_t; \theta, \alpha, \pi, p) = \phi\left(\frac{y_{nt} - E(y_{nt})}{\sigma}\right) \prod_{i=1}^I l_n^{(i)}(d_{nt}^{(i)}, x_{nt}, s_t; \theta, \pi) \quad (\text{B.1})$$

where $\phi(\cdot)$ is the density function of the standard normal distribution and σ is the standard deviation of η_t .

We can then substitute (B.1) into (5.3) to obtain the likelihood of the data for a particular market:

$$L(d_n, y_n, x_n | x_{n1}; \theta, \alpha, \pi, p) = \sum_{s_1=1}^S \sum_{s_2=1}^S \dots \sum_{s_T=1}^S \left[\begin{aligned} & \pi(s_1 | x_{n1}) \mathcal{L}_1(d_{n1}, y_{n1} | x_{n1}, s_1; \theta, \alpha, \pi, p) \\ & \times \left(\prod_{t=2}^T \pi(s_t | s_{t-1}) \mathcal{L}_t(d_{nt}, y_{nt} | x_{nt}, s_t; \theta, \alpha, \pi, p) \right) \end{aligned} \right] \quad (\text{B.2})$$

To make clear the number of calculations required to form the expression in (B.2) for a particular market, we specify (B.2) using matrix notation. Denote A_{nt} as a $1 \times S$ vector with components given

³⁹Multiple equilibria may be a possibility. This issue did not cause any problems for this set of Monte Carlos.

by the likelihood at $t = 1$ conditional on a particular unobserved states times the initial probability of being in the unobserved state:

$$\mathcal{A}_{n1} = \left[\begin{array}{ccc} \pi(1|x_{n1})\mathcal{L}_1(d_{n1}, y_{n1}|x_{n1}, 1; \theta, \alpha, \pi, p) & \dots & \pi(S|x_{n1})\mathcal{L}_1(d_{n1}, y_{n1}|x_{n1}, S; \theta, \alpha, \pi, p) \end{array} \right] \quad (\text{B.3})$$

If $\mathcal{T} = 1$, summing over the elements of A_{n1} would give $L(d_n, y_n, x_n|x_{n1}; \theta, \alpha, \pi, p)$. For $t > 1$, we form an $S \times S$ matrix where the (i, j) element gives the probability of moving from $s_{t-1} = i$ to $s_t = j$ times the likelihood contribution at t conditional on being in unobserved state j :

$$\mathcal{A}_{nt} = \left[\begin{array}{ccc} \pi(1|1)\mathcal{L}_t(d_{nt}, y_{nt}|x_{nt}, 1; \theta, \alpha, \pi, p) & \dots & \pi(1|S)\mathcal{L}_t(d_{nt}, y_{nt}|x_{nt}, S; \theta, \alpha, \pi, p) \\ \vdots & \ddots & \vdots \\ \pi(S|1)\mathcal{L}_t(d_{nt}, y_{nt}|x_{nt}, 1; \theta, \alpha, \pi, p) & \dots & \pi(S|S)\mathcal{L}_t(d_{nt}, y_{nt}|x_{nt}, S; \theta, \alpha, \pi, p) \end{array} \right] \quad (\text{B.4})$$

Taking A_{n1} times A_{n2} gives a $1 \times S$ vector of the joint likelihood of the data and being in each of the unobserved states. We define A_n as the product of A_{nt} over \mathcal{T} :

$$A_n = \prod_{t=1}^{\mathcal{T}} A_{nt}$$

A_n is then a row vector with S elements with each element giving the joint likelihood of the data and being in a particular unobserved state at \mathcal{T} . To form A_n , an $S \times S$ matrix is multiplied by an $1 \times S$ matrix \mathcal{T} times. Let the s th element be denoted by $A_n(s)$. The likelihood for the n th market is then given by:

$$L(d_n, y_n, x_n | x_{n1}; \theta, \alpha, \pi, p) = \sum_{s=1}^S A_n(s)$$

B.2.3 Obtaining conditional choice probabilities

Four sets of CCP's are used in this Monte Carlo:

1. When s_t is ignored (Column 3 of Table 2), we specify the conditional probability of exiting at $t + 1$ as a flexible function of the observed variables, W_{1ti} (for the i^{th} firm in a given market at time t). The variables included in W_{1ti} are combinations of the permanent market characteristic, x_1 , whether the firm is active in period t , $d_{2t}^{(i)}$, and the number of firms in the market at t :

$$W_{1ti} \equiv \left(1, x_1, x_1^2, d_{2t}^{(i)}, \sum_{i'} d_{2t}^{(i)}, \left(\sum_{i'} d_{2t}^{(i)} \right)^2 \right)$$

We then estimate a logit on the probability of exiting using the variables in W_{1ti} as controls.

2. When s_t is observed (Column 2), we add the variables in W_{2ti} to the logit, where:

$$W_{2ti} \equiv \left(s_t, s_t x_1, s_t d_{2t}^{(i)}, s_t \sum_{i'} d_{2t}^{(i')} \right)$$

implying 10 parameters govern the CCP's.

3. When the conditional choice probabilities are updated with the data (Column 5), and when using the two-stage method (Column 6), we use the variables in both W_{1t} and W_{2t} . In both these cases, the m^{th} iteration uses the conditional probabilities of being in each unobserved state, $q_{nst}^{(m)}$, as weights in the logit estimation.

4. Finally, when the CCP's are updated with the model (Columns 4 and 7), we update the probability of exiting using the logit formula for the likelihood:

$$p_{1t}^{(i,m+1)}(x, s) = \frac{1}{1 + \exp \left[v_{2t}^{(i)}(x, s, p^{(m)}, \theta^{(m)}, \pi^{(m+1)}) \right]}$$

B.2.4 Initial conditions

There is an initial conditions problem in the stationary equilibrium, because the distribution of s_1 depends on the the distribution of the observed states. We estimate this distribution jointly with the other parameters of the model. Since the unobserved state applies at the market level of aggregation, the relevant endogenous variable is the lagged number of firms in the initial period. We regress the lagged number of firms in the initial period on a flexible function of the characteristics of the market, in this exercise, a constant, x_1 , and x_1^2 . Denote the residual from this regression as ζ . We then approximate the initial probability of being in unobserved state s for the n^{th} market using a multinomial logit form:

$$\pi(s|x_{n1}) = \frac{\exp([1 \quad \zeta_n] \delta_s)}{\sum_s \exp([1 \quad \zeta_n] \delta_s)}$$

With δ_1 set to zero, there are 8 parameters to be estimated. We estimate $\pi(s|x_{n1})$ at each iteration using a similar procedure to Section B.1.4, now allowing for the fact that the unobserved states follow a Markov transition. Despite this additional complication, the algorithm is the same: calculate the likelihood given each initial unobserved state and take it as a given when maximizing to update δ .