Sequential Monte Carlo Methods for Estimating Dynamic Microeconomic Models*

JASON R. BLEVINS†

Department of Economics, Ohio State University

February 19, 2015‡

Forthcoming at the *Journal of Applied Econometrics*

Abstract. This paper develops estimators for dynamic microeconomic models with serially correlated unobserved state variables using sequential Monte Carlo methods to estimate the parameters and the distribution of the unobservables. If persistent unobservables are ignored, the estimates can be subject to a dynamic form of sample selection bias. We focus on single-agent dynamic discrete choice models and dynamic games of incomplete information. We propose a full-solution maximum likelihood procedure and a two-step method and use them to estimate an extended version of the capital replacement model of Rust with the original data and in a Monte Carlo study.

Keywords: dynamic discrete choice, latent state variables, serial correlation, dynamic selection, sequential Monte Carlo methods, particle filtering.

JEL Classification: C13, C15, C35.

1. Introduction

This paper proposes methods for estimating a class of dynamic microeconomic models with serially correlated latent state variables. Time-varying unobserved heterogeneity of this sort can lead to biased parameter estimates and, in turn, incorrect inferences and predictions about economic outcomes. We propose two estimators that avoid such bias by explicitly accounting for non-iid error structures. They are extensions of two commonly used methods: the full solution maximum likelihood approach of Rust (1987) and the two-step estimator of Bajari, Benkard and Levin (2007). The methods are fairly general, but we focus primarily on two common special cases for simplicity: single-agent dynamic discrete choice (DDC) models and dynamic games of incomplete information. In both cases, the observed and unobserved states may be discrete, continuous, or a combination of both.

*I am grateful to Han Hong, Shakeeb Khan, Paul Ellickson, and Andrew Sweeting for their guidance and support during this project and many others, to Jesús Fernández-Villaverde and Juan Rubio-Ramírez for an introduction to particle filtering, and to Ahmed Khwaja, Mario Miranda, Sanjog Misra, Andriy Norets, K. Sudhir, and Nathan Yang for helpful discussions. Four anonymous referees also provided valuable suggestions. I also thank seminar participants at Duke, Ohio State (AEDE), Rochester (Simon Business School), Yale (School of Management), and the 2008 ERID Conference on Identification Issues in Economics for useful comments. This work was supported in part by an allocation of computing time from the Ohio Supercomputer Center.

†1945 N High Street, 410 Arps Hall, Columbus, OH 43210. Phone: (614) 292-5588. Fax: (614) 292-3906. Email: blevins.141@osu.edu.

‡This paper is a revision of a working paper whose first draft appeared in May 2008.

The source of the bias from ignoring a serially correlated latent state variable can be viewed as a dynamic extension of the bias that arises in static, linear sample selection models (Heckman, 1979). However, the models we consider are dynamic and nonlinear and so the nature of the bias is more complex, the direction of the bias can be ambiguous, and the computational burden faced when accounting for such variables by standard Monte Carlo integration is severe. Nevertheless, accounting for unobserved heterogeneity is important in applied work in many fields of economics and it can take many forms, such as unobserved types of agents, random coefficients, unobserved product or choice characteristics, or unobserved aggregate shocks. As such, there are many potential applications of the methods developed herein to problems in applied microeconomics. Examples of latent state variables from industrial organization include unobserved market-level shocks and firm-specific productivity shocks. In labor economics, ability or wage offers may be unobserved. In health economics, latent health status is important.

To illustrate the proposed methods and to provide intuition for the potential problems that can arise, we revisit the well-known optimal capital replacement problem of Rust (1987) in a generalized form with continuous state variables and where one of these state variables is unobserved and possibly serially correlated. Because the basic model is simple and tractable, but also representative of richer models, the original model has become a benchmark for demonstrating new methods (Aguirregabiria and Mira, 2002; Bajari et al., 2007; Norets, 2009b, 2012; Arcidiacono and Miller, 2011; Su and Judd, 2012). Estimating this optimal renewal model using Rust (1987)'s original data and carrying out Monte Carlo experiments using the extended model will allow us to provide insights into the dynamic sample selection phenomenon and to illustrate our proposed solutions in a familiar setting, but also yields new insights about the data.

The potential dynamic selection problem is straightforward in Rust (1987)'s model. Consider a forward-looking individual facing an optimal replacement investment decision: a machine becomes more and more costly to maintain with each month of use; each month the individual may pay a one-time cost to replace the machine with a new one. The parameter of interest is the cost of replacement relative to the cost of maintenance. In addition to months of use, which is observed, suppose there is an unobserved state variable representing additional costs per month that aren't due directly to usage (e.g., type or intensity of use, conditions, etc.). If these unobserved costs are serially correlated, then there may be dynamic sample selection on unobservables (e.g., if machines that are more unobservably costly are replaced more frequently). If ignored, this can in turn result in biased parameter estimates.

To understand the potential computational problem raised by latent state variables, consider an observed sequence $y_{1:T} = (y_1, \ldots, y_T)$, which is thought to depend on an unobserved sequence of latent state variables $\xi_{1:T} = (\xi_1, \ldots, \xi_T)$ and a vector of parameters θ.[1] The likelihood can be written as[2]

$$p(y_{1:T};\theta) = \int p(y_{1:T}, \xi_{1:T};\theta)\, d\xi_{1:T} = \int p(y_{1:T} \mid \xi_{1:T};\theta)\, p(\xi_{1:T};\theta)\, d\xi_{1:T}. \tag{1}$$

[1] For simplicity we will use the notation y_t to denote both the random variable Y_t as well as a particular realization $Y_t = y_t$. Furthermore, we assume that the distributions of all continuous variables admit densities with respect to Lebesgue measure. We treat discrete variables analogously with respect to the counting measure. Generically, the density of Y_t evaluated at y_t will be denoted $p(y_t)$ as long as there is no ambiguity.

[2] We include θ to indicate that these densities depend on the parameters, but we use a semicolon to separate θ to indicate that it is not a random variable nor is it considered part of the conditioning set; rather, θ is deterministic throughout the paper.

If the latent state variables are exogenous, then in principle one could simulate draws from $p(\xi_{1:T};\theta)$ directly. However, this is potentially a very high-dimensional integral and Monte Carlo integration via simulation of paths $\xi_{1:T}$ is not a very efficient way to evaluate it.

A further complication arises if, as in this paper, we wish to allow for endogenous feedback whereby outcomes of the observable process y_t can influence future values of the latent process ξ_{t+1}. In order to apply maximum simulated likelihood (MSL) one must factor the likelihood into components that can be simulated directly. In such cases we cannot simulate from the unconditional distribution in (1) directly, but an alternative representation is useful:[3]

$$p(y_{1:T};\theta) = \int \cdots \int \left[\prod_{t=1}^{T} p(y_t \mid \xi_t, y_{t-1};\theta)\, p(\xi_t \mid y_{t-1}, \xi_{t-1};\theta) \right] d\xi_T \ldots d\xi_1 \tag{2}$$

Although we can usually now simulate from the distributions $p(\xi_t \mid y_{t-1}, \xi_{t-1};\theta)$ directly, this still involves a high-dimensional integral.

The approach we propose is based on yet a third way of writing the likelihood function:

$$p(y_{1:T};\theta) = \prod_{t=1}^{T} p(y_t \mid y_{1:t-1};\theta) = \prod_{t=1}^{T} \int p(y_t \mid \xi_t, y_{t-1};\theta)\, p(\xi_t \mid y_{1:t-1};\theta)\, d\xi_t. \tag{3}$$

Sequential Monte Carlo (SMC) methods can be used to draw from $p(\xi_t \mid y_{1:t-1};\theta)$, which is the posterior distribution of the unobserved random variable ξ_t conditional on the data $y_{1:t-1}$. This allows us to avoid the high-dimensional integrals in (1) and (2) and instead use a product of one-dimensional integrals as in (3). We propose two estimators which involve maximum likelihood estimation of parameters where the likelihood function is be integrated with respect to the posterior distribution of the latent state as above.

In the dynamic discrete choice literature, serial correlation of unobservables has been a concern from the beginning and remains so at present (Eckstein and Wolpin, 1989; Rust, 1994; Aguirregabiria and Mira, 2010). However, until recently the known methods for estimating models with non-trivial correlation structures were either application-specific or were not applicable to arbitrary distributions of the unobserved states. See Pakes (1994) and Ackerberg, Benkard, Berry and Pakes (2007, Section 3.8.1) for an overview of three such methods and Pakes (1986), Timmins (2002), and Aguirregabiria and Mira (2007) for applications of those methods.

Identification and estimation of dynamic models with serially correlated unobserved state variables is a topic of recent interest in the literature. Hu and Shum (2012, 2013) address nonparametric identification of the Markov transition kernel in dynamic single-agent models and games with serially correlated unobserved state variables. Blevins (2014) provides conditions for nonparametric identification of dynamic models with discrete and continuous controls, such as the dual entry and investment choices encountered in many models in industrial organization, which can be combined with results of Hu and Shum (2012, 2013) to establish nonparametric identification of dual choice models with serially correlated unobservables such as those considered in this paper. Imai, Jain and Ching (2009) develop a Bayesian method for estimating single agent dynamic discrete choice models with finite state spaces using a Markov chain

[3] Here, for simplicity, we have used an additional assumption the process is a first-order Markov process where $p(y_t, \xi_t \mid y_{1:t-1}, \xi_{1:t-1};\theta) = p(y_t, \xi_t \mid y_{t-1}, \xi_{t-1};\theta)$. The models we consider later will satisfy this assumption.

Monte Carlo (MCMC) algorithm which simultaneously solves the dynamic program and estimates the parameters. Norets (2009b) extends this idea, developing Gibbs sampling methods which allow for serially correlated unobserved state variables. In terms of classical estimation, Arcidiacono and Miller (2011) estimate models with discrete, time-varying unobserved states by iteratively applying the EM algorithm in combination with the CCP (conditional choice probability) based estimation approach of Hotz and Miller (1993). Keane and Wolpin (1994) propose a combination of simulation and interpolation and Stinebrickner (2000) considered Gaussian quadrature and interpolation. Finally, Aguirregabiria and Mira (2007) demonstrate how to estimate dynamic games with time-invariant unobserved heterogeneity.

This paper is also related to sequential Monte Carlo methods more broadly, which originated in signal processing and engineering applications. For a survey, see Doucet, de Freitas and Gordon (2001) and the citations therein. In the economics literature, this approach has been used by Fernández-Villaverde and Rubio-Ramírez (2007) to estimate structural macroeconomic models with stochastic volatility. Microeconomic models raise new considerations in that there are numerous unobserved states to consider (e.g., for each agent in each of possibly many markets) and that in many cases the controls are discrete rather than continuous. Furthermore, in the case of dynamic games, solving for the equilibria of the model is not always feasible and so two-step methods can be useful. Creal (2012) surveys the use of SMC methods in economics and finance.

Particle filters can also be used for Bayesian inference about the parameters and latent state variables jointly, as both Flury and Shephard (2011) and Gallant, Hong and Khwaja (2010) have done in work independent of ours. Flury and Shephard (2011) propose constructing unbiased estimates of likelihoods using particle filters for carrying out MCMC using the Metropolis-Hastings algorithm. Gallant et al. (2010) use a Bayesian approach to estimate a complete-information dynamic discrete model of entry and recover unobserved firm-level costs in the generic pharmaceutical industry. In contrast, we consider incomplete information models, focus on classical estimation, and discuss both full-solution and two-step methods.

Finally, the methods proposed herein have subsequently been used in other applications. Fang and Kung (2010) apply the full-solution estimator to estimate a model of life insurance policy lapsation, accounting for unobserved, serially correlated shocks to income, health, and bequest motives. Blevins, Khwaja and Yang (2014) apply the two-step estimator to a dynamic oligopoly model of the Canadian fast food industry in which firm size spillovers are incorporated via firm-specific, time-varying unobservables.

In Section 2, we introduce the class of dynamic discrete choice models and dynamic games of interest and characterize them as nonlinear state space models. Section 3 introduces sequential Monte Carlo methods and shows how to apply them to the models of interest. Section 4 describes the two proposed estimators and in Section 5 we apply them to a generalized version of the capital replacement model of Rust (1987) using the original data and in a series of Monte Carlo experiments. Section 6 concludes.

2. Dynamic Microeconomic Models as Nonlinear State Space Models

Sequential Monte Carlo methods were originally developed for nonlinear state space models. For example, in many signal processing applications, the signal is modeled as an unobserved state and is only observed along with some noise, which is a potentially serially-correlated, non-Gaussian, and nonlinear process

which contaminates the signal. Dynamic microeconomic models tend to be more complex, because they are *controlled* stochastic processes involving rational decision makers who control the joint stochastic process in some way. In this section we show that, with some generalizations, sequential Monte Carlo methods can still be applied successfully to these models. The primary difference is that the likelihoods (conditional on the latent state) are usually more complex and only be evaluated numerically. Even in very simple models, the state-space representation can be highly nonlinear and non-Gaussian, preventing the use of methods such as Kalman filtering or GHK.[4]

Consider a discrete-time dynamic model with N agents, indexed by $i = 1,\ldots,N$, over an infinite time horizon, indexed by $t = 1,2,\ldots,\infty$. The state at time t is summarized by a state vector $(s_t,\eta_t) \in \mathscr{S} \times \mathscr{N}$ where $s_t \in \mathscr{S}$ is common knowledge to all agents but $\eta_t = (\eta_{1t},\ldots,\eta_{Nt}) \in \mathscr{N} \equiv \mathscr{N}_1 \times \cdots \times \mathscr{N}_N$ is a vector of private shocks where $\eta_{it} \in \mathscr{N}_i$ is private information of agent i.

Each period, each agent i observes the state and makes a choice a_{it} from the choice set $\mathscr{A}_i$. Define $\mathscr{A} \equiv \mathscr{A}_1 \times \ldots \times \mathscr{A}_N$ and let $a_t = (a_{1t},\ldots,a_{Nt}) \in \mathscr{A}$ denote the vector of all actions at time t. Upon making the choices a_t, each agent i receives a payoff $U_i(a_t,s_t,\eta_{it},\theta)$ associated with making choice a_{it} in state s_t, given that agent i's rivals make choices $a_{-it} = (a_{1t},\ldots,a_{i-1,t},a_{i+1,t},\ldots,a_{Nt})$.

Agents are forward-looking and share a common discount factor $\beta \in [0,1)$, which is known by the researcher. Agent i's discounted expected future payoff when the market is in state s_t is

$$\mathrm{E}\left[\sum_{\tau=t}^{\infty} \beta^{\tau-t} U_i(a_\tau,s_\tau,\eta_{i\tau},\theta) \,\middle|\, s_t\right],$$

where the expectation is taken over the infinite sequence of actions, states, and private shocks.

Before proceeding we make several standard assumptions to make the model more tractable (cf. Rust, 1994; Aguirregabiria and Mira, 2010). Let $\sigma = (\sigma_1,\ldots,\sigma_N)$ denote a profile of policy functions. A policy function that depends on the past history of the process only via the contemporaneous payoff-relevant state variables s_t and η_{it} is called Markovian.

Assumption 1 (Markovian Policy Functions)**.** Each agent $i = 1,\ldots,N$ follows a Markovian policy function $\sigma_i : \mathscr{S} \times \mathscr{N}_i \to \mathscr{A}_i : (s,\eta_i) \mapsto a = \sigma_i(s,\eta_i)$ that assigns the optimal action a in state s given η_i.

Since we consider the case of stationary decision processes where the transition probabilities, utility functions, and discount factors are time invariant, the optimal Markovian policies will be stationary.

Assumption 2 (Conditional Independence)**.** The state variables follow a first-order controlled Markov process where the joint distribution can be factored as

$$p(s_t,\eta_t \mid a_{1:t-1},s_{1:t-1},\eta_{1:t-1};\theta) = p(s_t \mid a_{t-1},s_{t-1};\theta)\, p(\eta_t \mid s_t;\theta).$$

[4] In linear models with serially-correlated, Gaussian unobserved state variables, the GHK simulator (cf. Hajivassiliou and Ruud, 1994) has been effective in carrying out the required integration. However, the models we consider are highly nonlinear and potentially non-Gaussian and the regions of integration are non-rectangular. Furthermore, in the models we consider the latent variables may be endogenous to the choice variables and there may be feedback to the observed state variables.

Remark. This assumption requires that the transient unobservables η_t are independent of the previous actions and are serially independent conditional on s_t, but it does not restrict either the observed or unobserved components of s_t to be conditionally serially independent. This is in contrast to the conditional independence assumption of Rust (1994), where all unobservables are required to be serially independent.

Remark. Importantly, Assumption 2 allows a_{t-1} to affect the transition from s_{t-1} to s_t. Therefore, variables that depend on actions in the past, such incumbency or potential entrant status in an entry-exit model, are permitted to be elements of s_t under these assumptions.

Assumption 3 (Private Information)**.** The private shocks η_{it} are independent across i in each period t and follow a known distribution $G_i(\cdot \mid s_t;\theta)$ with support $\mathcal{N}_i$.

These assumptions imply that we can write the joint distribution as follows:

$$\begin{aligned} p(a_t, s_t, \eta_t \mid a_{1:t-1}, s_{1:t-1}, \eta_{1:t-1};\theta) &= p(a_t \mid a_{1:t-1}, s_{1:t}, \eta_{1:t};\theta)\, p(s_t, \eta_t \mid a_{1:t-1}, s_{1:t-1}, \eta_{1:t-1};\theta) \\ &= p(a_t \mid s_t, \eta_t;\theta)\, p(s_t \mid s_{t-1}, a_{t-1};\theta)\, p(\eta_t \mid s_t;\theta) \end{aligned}$$

The second line follows immediately Assumptions 1 and 2 above while Assumption 3 further requires that the distribution of unobservables is known.

We now depart from the standard framework and consider the presence of unobserved state variables. We let s_t be partially observed and write $s_t = (x_t, \xi_t)$ where x_t is observed by the researcher, along with the choices a_t, but ξ_t is an unobserved state. Both states x_t and ξ_t are common knowledge among the agents. Since x_t and ξ_t may be multidimensional, this allows for general forms of both market- and agent-level time-varying unobserved heterogeneity. This is also where our framework departs from the state space models to which SMC methods are usually applied; they do not have observed state variables, x_t, only an observation process, a_t, and an unobserved signal, ξ_t, and a_t is continuous but here it may be discrete.

In order to apply sequential Monte Carlo methods to this model, it will be useful to take a stand on the timing of the dependence, if any, between x_t and ξ_t while still allowing for complex patterns of feedback between the observed and unobserved states and the control variables.

Assumption 4 (Feedback Timing)**.** Conditional on ξ_t, x_{t-1}, and a_{t-1}, x_t is independent of ξ_{t-1}:

$$p(x_t \mid \xi_t, x_{t-1}, \xi_{t-1}, a_{t-1}) = p(x_t \mid \xi_t, x_{t-1}, a_{t-1}).$$

Remark. Thus, the models we consider here are quite general in that they allow for dependence between x_t and ξ_t, but we limit this dependence to cases under Assumption 4 in order to simplify the exposition of the algorithm. Without this assumption, one would need to track both ξ_t and ξ_{t-1} when evaluating the transition density of x_t. From a modeling perspective we therefore focus on cases where ξ_t influences x_t but it is equally plausible that one might prefer to model $p(x_t \mid \xi_{t-1}, x_{t-1}, a_{t-1})$ instead. In that case the time indices in the procedures described below simply need to be adjusted accordingly.

A convenient representation of the model will be in the form of the following two densities:

$$\begin{aligned} &p(y_t \mid y_{t-1}, \xi_t;\theta) = p(a_t \mid x_t, \xi_t;\theta)\, p(x_t \mid \xi_t, a_{t-1}, x_{t-1};\theta), \\ &p(\xi_t \mid y_{t-1}, \xi_{t-1};\theta) = p(\xi_t \mid x_{t-1}, a_{t-1}, \xi_{t-1};\theta) \end{aligned}$$

The first density above is the *likelihood*, the conditional density of observables $y_t = (a_t, x_t)$, which is the product of the density (or probability) of the choice variable a_t conditional on the state variables and the transition density for x_t. As is typical, the former is fully specified by the model and the latter is determined by the model and the distribution G of the choice-specific shocks η_t. The second density above is the *transition density* of the unobservables, ξ_t. Relative to standard models, this density is new but is also specified as part of the model along with the transition density of x_t.

In the sections that follow, we consider two common special cases: dynamic games of incomplete information and single agent dynamic discrete choice models. We show, for example, that single agent dynamic discrete models can be written in the nonlinear state space form above where the likelihood is the product of the conditional choice probabilities, $p(a_t \mid x_t, \xi_t; \theta)$, and the transition density of the observed states, $p(x_t \mid a_{t-1}, x_{t-1}, \xi_t; \theta)$. The law of motion for the latent state, $p(\xi_t \mid x_{t-1}, a_{t-1}, \xi_{t-1}; \theta)$, may in turn depend on past values of the discrete choice, a_{t-1}.

2.1. Single-Agent Dynamic Discrete Choice Models

Here we consider a class of models which generalizes the framework of Rust (1994) by incorporating a serially correlated latent state variable. Since there is only a single agent ($N = 1$), we omit the i subscript from states and payoffs in this section. In each period the agent makes a single choice a from a discrete choice set $\mathscr{A} = \{0, 1, \ldots, K\}$ and associated with each choice is a choice-specific shock ε_{ta}. Here, the private information shock is simply $\eta_t = \varepsilon_t$ and the support of $\varepsilon_t = (\varepsilon_{t0}, \ldots, \varepsilon_{tK})$ is $\mathbb{R}^{K+1}$. We make the following standard additive separability assumption.

Assumption 5 (Additive Separability)**.** U additively separable in ε_t: $U(a_t, s_t, \varepsilon_t, \theta) = u(a_t, s_t, \theta) + \varepsilon_{ta}$.

Since the problem is recursive in nature, we will omit the time subscripts on variables when convenient and use (s', ε') to denote the one-period-ahead counterpart of (s, ε). The value function for this model can then be expressed as

$$V(s, \varepsilon, \theta) = \max_{a \in \mathscr{A}} \left\{ u(a, s, \theta) + \varepsilon_a + \beta \, \mathrm{E}\left[V(s', \varepsilon', \theta) \mid s, a \right] \right\}.$$

We also define the choice-specific value function, $v(a, s, \theta)$, which represents the current and future expected payoff from choosing a, net of the idiosyncratic component ε_a:

$$v(a, s, \theta) \equiv u(a, s, \theta) + \beta \, \mathrm{E}\left[V(s', \varepsilon', \theta) \mid s, a \right].$$

Under Assumption 5, we can now express the problem in a more compact form, a form which resembles a static discrete choice problem with the choice-specific value function playing the role of the period payoff function. Letting $\sigma(s, \varepsilon, \theta)$ denote the optimal policy, or choice of a,

$$\sigma(s, \varepsilon, \theta) = \arg\max_{a \in \mathscr{A}} \left[v(a, s, \theta) + \varepsilon_a \right].$$

Under certain distributional assumptions for ε, the model admits conditional choice probabilities with known analytical forms. In particular, in applied work it is often assumed that the components of ε

are iid and follow the type I extreme value distribution. Under this assumption, the conditional choice probabilities have a closed form in terms of the choice-specific value function,

$$P(\sigma(a,\varepsilon,\theta) = a \mid s;\theta) = \frac{\exp(v(a,s,\theta))}{\sum_{j\in\mathscr{A}} \exp(v(j,s,\theta))},$$

and $v(a,s)$ is the unique fixed point to the contraction mapping (Rust, 1994):

$$\Gamma(v)(a,s,\theta) = u(a,s,\theta) + \beta \int \ln\left[\sum_{j\in\mathscr{A}} \exp(v(j,s',\theta))\right] p(s' \mid s,a;\theta)\, ds'. \tag{4}$$

To summarize, Assumption 5 allows us to restate the problem as a simple static discrete choice problem. This problem is still intractable for an arbitrary distribution G_i, but the type I extreme value distribution leads to the convenient closed form above for $P(a \mid s;\theta)$, in terms of the choice specific value function $v(a,s,\theta)$. Importantly, it also allows us to obtain $v(a,s,\theta)$ as the fixed point of the functional equation above. This provides a clear path for evaluating it numerically using any number of methods such as value function iteration or projection methods (cf. Judd, 1998). Then, since we can evaluate $v(a,s,\theta)$ for any choice a and state s, we can now evaluate $P(a \mid s)$, which is needed to evaluate the likelihood and brings us one step closer to being able to estimate the model.

This also illustrates the distinction between computation and estimation with regard to the choice probabilities $P(a \mid s;\theta)$. For computation it is irrelevant that s is only partially observed because the model is defined conditional on $s = (x,\xi)$. The distinction between the observed and unobserved states becomes important at the estimation stage.

2.2. Dynamic Games of Incomplete Information

In the case where $N > 1$, each agent's optimal decision depends on the expectations that agent holds about the actions of the other agents. We assume that agents use Markovian strategies that are consistent with a Markov perfect equilibrium (MPE). Given a strategy profile $\sigma = (\sigma_1,\ldots,\sigma_N)$, agent i's expected discounted future payoff in state s can be expressed recursively in terms of the *ex-ante value function*:

$$\bar{V}_i(s;\sigma,\theta) = \mathrm{E}\left[U_i(\sigma(s,\eta,\theta),s,\eta_i,\theta) + \beta \int \bar{V}_i(s';\sigma,\theta)\, p(s' \mid s,\sigma(s,\eta,\theta);\theta)\, ds' \,\middle|\, s\right].$$

The bar denotes that this is the expected value before η is realized, hence, the expectation is with respect to the distribution of η. Given a strategy σ_i for player i, in equilibrium each rival firm's beliefs about firm i's actions must correspond to the beliefs implied by the strategy σ_i and the distribution G_i of firm i's private information. By including rival strategies inside the expectation above, we are integrating with respect to the implied beliefs about rival firms.

Definition. A *Markov perfect equilibrium (MPE)* is a strategy profile $\sigma = (\sigma_1,\ldots,\sigma_N)$ such that for all $i = 1,\ldots,N$ and $s \in \mathscr{S}$, $\bar{V}_i(s;\sigma_i,\sigma_{-i},\theta) \geq \bar{V}_i(s;\sigma_i',\sigma_{-i},\theta)$ for all alternative Markov strategies σ_i'.

Two-step estimation methods are particularly useful for estimating dynamic games because they do not require one to solve the model to determine the equilibria. Following most papers in the literature on

two-step estimation (e.g., Bajari et al., 2007; Aguirregabiria and Mira, 2010) we require that the data be generated by a single Markov perfect equilibrium and that all players expect the same equilibrium to be played in all periods.

Assumption 6 (Equilibrium Selection)**.** The data are generated by a single MPE profile σ.

Assumption 6 is identical to Assumption ES of Bajari et al. (2007). Under this assumption, dynamic games of this form have a nonlinear state space representation where $p(a_t \mid x_t, \xi_t; \theta)$ is the conditional density of a_t implied by σ given the distribution of η and θ. Thus, the conditional density of choices implied by the Markov strategy σ plays the same role in the likelihood as the discrete choice probabilities in the single agent DDC model.

3. Sequential Monte Carlo Methods

Sequential Monte Carlo methods, or particle filters, are simulation-based methods for approximating posterior distributions of unobserved state variables in nonlinear and non-Gaussian state space models. We focus on the bootstrap filter introduced by Gordon, Salmond and Smith (1993) which is relatively simple but still captures the essence of more recent variations.[5] We first introduce the optimal nonlinear filtering problem arising due to the unobserved states and then introduce SMC methods as approximate solutions to this problem. This state space model and the methods discussed in this section will provide a foundation for the maximum filtered likelihood (MFL) estimators considered in Section 4.

3.1. Optimal Filtering

Given the generic nonlinear state space model from Section 2, the two primary problems of interest to a researcher who has a collection of observations $\{y_t\}_{t=1}^T$ are estimating the posterior distribution of the unobserved state ξ_t given the observed data and estimating the unknown parameters θ. From a classical perspective, θ is deterministic and these are distinct problems. In this setting, we can first recover the posterior distribution of ξ_t then use it form a likelihood function with which we can estimate θ. From a Bayesian point of view, inference on ξ_t and θ are essentially the same problem, since θ can be treated as a time-invariant component of ξ_t, with the joint posterior distribution of $(\xi_{1:T}, \theta)$ being of interest.

Here consider estimation of marginal posterior distributions of the form $\pi_{t|s}(d\xi_t) \equiv P\left(\xi_t \in d\xi_t \mid Y_{1:s} = y_{1:s}\right)$.[6] The cases $s < t$, $s = t$, and $s > t$ are referred to, respectively, as *prediction*, *filtering*, and *smoothing*. Particle filters are methods for approximating sequential filtering distributions $\pi_{t|t}$. Conveniently, the one-step-ahead prediction distribution $\pi_{t|t-1}$ is also approximated in the process.

The filtering distribution $\pi_{t|t}$ can be represented recursively, starting with the initial distribution π_0, by applying a series of prediction and updating steps using Bayes' theorem and the transition kernel

[5] Variations to the basic algorithm usually involve different proposal distributions, such as the auxiliary particle filter (Pitt and Shephard, 1999), or alternative resampling schemes, such as multinomial resampling (Gordon, Salmond and Smith, 1993) and residual resampling (Liu and Chen, 1998).

[6] For simplicity, we let the dependence on θ be implicit in this section.

$Q(d\xi_t \mid \xi_{t-1}, y_{t-1}) = \Pr\left(\xi_t \in d\xi_t \mid \xi_{t-1}, y_{t-1}\right)$, which gives the probability of ξ_t arriving in $d\xi_t$ given ξ_{t-1} and y_{t-1}. Given $\pi_{t-1|t-1}$, the one-step-ahead prediction and filtering distributions are

$$\pi_{t|t-1}(d\xi_t) = \int Q(d\xi_t \mid \xi_{t-1}, y_{t-1})\, \pi_{t-1|t-1}(d\xi_{t-1}) \quad \text{and} \tag{5}$$

$$\pi_{t|t}(d\xi_t) = \frac{p(y_t \mid y_{t-1}, \xi_t)\, \pi_{t|t-1}(d\xi_t)}{\int p(y_t \mid y_{t-1}, \xi_t)\, \pi_{t|t-1}(d\xi_t)}. \tag{6}$$

Analytic solutions for the optimal filtering problem are only known for special cases. Kalman Filters (Kalman, 1960) have been used very successfully in models that are both linear and Gaussian. The models we study, on the other hand, are nonlinear and non-Gaussian.

3.2. A Generic Particle Filter

Sequential Monte Carlo methods, or particle filters, are a class of methods which aim to approximate the sequence of posterior distributions $\pi_{t|t}$ using a weighted collection of R particles, or weighted mass points, $\{(\xi_t^r, w_t^r)\}_{r=1}^R$ that evolve over time. The particles can be used to form an empirical probability measure which approximates $\pi_{t|t}$:

$$\pi_{t|t}^R(d\xi) = \frac{\sum_{r=1}^R w_t^r \delta_{\xi_t^r}(d\xi)}{\sum_{r=1}^R w_t^r},$$

where δ_ξ denotes the measure that assigns mass 1 at ξ and zero elsewhere.

Particle filters operate in a recursive manner: given $\pi_{t-1|t-1}^R$ and a new observation y_t, we form an approximation $\pi_{t|t-1}^R$, motivated by (5), and use it to form an approximation $\pi_{t|t}^R$ using an approach motivated by (6). We describe a generic particle filter here, but for more thorough descriptions of SMC methods there are several very good sources available such as Künsch (2001), Doucet, de Freitas and Gordon (2001), Liu (2001), and Cappé, Moulines and Ryden (2005).

The particle filtering algorithm begins with an iid sample $\{\xi_0^r\}_{r=1}^R$ of draws from a chosen initial distribution π_0 and assigns each draw a weight of 1 to form the uniformly weighted collection $\{(\xi_0^r, 1)\}_{r=1}^R$. The resulting empirical measure π_0^R serves as an approximation to π_0. If the researcher knows the initial distribution of ξ_0, then that distribution should be used as π_0. Otherwise, π_0 is simply the (initial) importance sampling distribution. Relative to static importance sampling, the influence of the choice of this initial distribution is diminished as new information is brought to bear with each additional period of data. The effect of the initial distribution decays exponentially in the number of periods (Whiteley, 2012).

Proceeding recursively, suppose we begin with a uniformly weighted collection of particles $\{(\xi_{t-1}^r, 1)\}_{r=1}^R$ at time t distributed approximately according to $\pi_{t-1|t-1}$. For each $r = 1, \ldots, R$, following (5), draw

$$\tilde{\xi}_t^r \sim Q(d\xi_t \mid \xi_{t-1}^r, y_{t-1})$$

and set $\tilde{w}_t^r = 1$ to form the uniformly weighted particle system $\{(\tilde{\xi}_t^r, \tilde{w}_t^r)\}_{r=1}^R$. The empirical measure $\tilde{\pi}_{t|t-1}^R(d\xi_t) \equiv R^{-1} \sum_{r=1}^R \delta_{\tilde{\xi}_t^r}(d\xi_t)$ approximates $\pi_{t|t-1}(d\xi_t)$. Then following (6), we obtain

$$\tilde{\pi}_{t|t}^R(d\xi_t) \equiv \frac{p(y_t \mid y_{t-1}, \xi_t)\, \tilde{\pi}_{t|t-1}^R(d\xi_t)}{\int p(y_t \mid y_{t-1}, \xi_t)\, \tilde{\pi}_{t|t-1}^R(d\xi_t)} = \frac{\sum_{i=1}^R p(y_t \mid y_{t-1}, \tilde{\xi}_t^i) \delta_{\tilde{\xi}_t^i}(d\xi_t)}{\sum_{i=1}^R p(y_t \mid y_{t-1}, \tilde{\xi}_t^i)}.$$

The weighted particle representation of this distribution is obtained by setting the weights to be proportional to the likelihood of the new observation y_t, with $w_t^r = p(y_t \mid y_{t-1}, \tilde{\xi}_t^r)$, yielding the particle system $\{(\tilde{\xi}_t^r, w_t^r)\}_{r=1}^R$. Finally, we obtain a uniformly weighted particle system by resampling R particles $\{\xi_t^r\}_{r=1}^R$ from the empirical distribution $\tilde{\pi}_{t|t}^R$ and setting $w_t^r = 1$ for all r. The resulting approximation is

$$\pi_{t|t}^R(d\xi_t) \equiv \frac{1}{R}\sum_{r=1}^{R} \delta_{\xi_t^r}(d\xi_t).$$

The process is perhaps easiest to understand in algorithm form:

- *Initialization:* Draw $\xi_0^r \sim \tilde{\pi}_0(d\xi_0)$ for each $r = 1, \ldots, R$.
- *Recursion:* Repeat the following steps for each $t = 1, \ldots, T$.
  - *Importance sampling:* Draw $\tilde{\xi}_t^r \sim Q(d\xi_t \mid \xi_{t-1}^r, y_{t-1})$ and set $w_t^r = p(y_t \mid y_{t-1}, \tilde{\xi}_t^r)$ for each $r = 1, \ldots, R$.
  - *Resampling:* For each $r = 1, \ldots, R$, draw ξ_t^r from the collection $\{\tilde{\xi}_t^r\}_{r=1}^R$ in proportion to the weights $\{w_t^r\}_{r=1}^R$.

One of the benefits of using the particle filter in this context is that it only requires evaluating or sampling from densities that arise naturally as part of the model specification. Note that we also obtain an approximation to the prediction distribution, $\pi_{t|t-1}^R$, as a by-product of the algorithm using the uniformly-weighted particles $\tilde{\xi}_t^r$ drawn from the proposal distribution. This prediction distribution will be useful later, as we use it to integrate the likelihood function over the distribution of the latent state.

4. Estimation

The sequential Monte Carlo methods discussed above solve the problem of drawing from the various posterior distributions. In this section we turn to the problem of estimating the unknown parameters θ given a sample of N iid observations of length T denoted $\{y_{1,1:t}, y_{2,1:t}, \ldots, y_{N,1:T}\}$ where $y_{i,s:t} = \{y_{i,\tau}, s \leq \tau \leq t\}$. Let $\xi_{i,s:t}$ be defined similarly.

How one estimates the model and how one solves the model (if at all) are two distinct problems. Here our focus is on estimation. How and when a particle filter can be applied also differs for each estimation method. We consider both a general full-solution approach, which assumes that it is feasible to evaluate the likelihood $p(y_t \mid \xi_t, y_{t-1}; \theta)$ (e.g., the model can be solved somehow to obtain choice probabilities), and a two-step estimator based on that of Bajari, Benkard and Levin (2007), which flexibly estimates the policy functions in a first step and then estimates the structural parameters using the equilibrium conditions in a second step. In the full solution maximum likelihood approach, we can simply use the particle filter to approximate the likelihood function and maximize it, yielding point estimates for all parameters. In the two step approach, the particle filter must be applied in the first step to recover both the distribution of the unobserved state as well as policy functions which condition on the unobserved state. As is the case without unobserved heterogeneity, the two-step estimator is less efficient than the

full solution approach and requires the researcher to choose certain quantities such as the number of simulation draws, the distribution of alternative policy functions, etc. Hence, when it is computationally feasible the full solution estimator is preferable.

4.1. Maximum Filtered Likelihood (MFL) Estimation

Given observations $\{y_{i,1:T}\}_{i=1}^N$ and following (3) the log-likelihood function is

$$L_N(\theta) = \sum_{i=1}^N \sum_{t=1}^T \ln p(y_{i,t} \mid y_{i,1:t-1};\theta) = \sum_{i=1}^N \sum_{t=1}^T \ln \int p(y_{i,t} \mid \xi_{i,t}, y_{i,t-1};\theta)\, p(\xi_{i,t} \mid y_{i,1:t-1};\theta)\, d\xi_t.$$

This is an integral with respect to the step-ahead filtering distribution. If we can evaluate $p(y_{i,t} \mid \xi_{i,t}, y_{i,t-1};\theta)$ and if we can draw from the transition density of $\xi_{i,t}$, then a particle filter can facilitate maximum likelihood estimation of θ since we can readily form the approximation

$$p(y_{i,t} \mid y_{i,1:t-1};\theta) \approx \frac{1}{R} \sum_{r=1}^R p(y_{i,t} \mid \tilde{\xi}_{i,t}^r, y_{i,t-1};\theta)$$

by using the approximate empirical distribution $\pi_{t|t-1}^R$, formed using the particle system $\{(\tilde{\xi}_{i,t}^r, 1)\}_{r=1}^R$. This leads directly to the following approximation to the log-likelihood function:

$$\hat{L}_{N,R}(\theta) \equiv \frac{1}{N} \sum_{i=1}^N \sum_{t=1}^T \ln \left[\frac{1}{R} \sum_{r=1}^R p(y_{i,t} \mid \tilde{\xi}_{i,t}^r, y_{i,t-1};\theta) \right].$$

We define the *maximum filtered likelihood (MFL) estimator* as

$$\hat{\theta}_{N,R} = \arg\max_{\theta \in \Theta} \hat{L}_{N,R}(\theta).$$

Following the application in Section 5, we report the results of a series of Monte Carlo experiments which suggest that the estimator performs well in dynamic discrete choice models even when the number of particles is small. Here we focus on establishing consistency of the MFL estimator as $N \to \infty$ and $R \to \infty$. Note that we can approximate the log likelihood function as

$$\hat{L}_{N,R}(\theta) = \frac{1}{N} \sum_{i=1}^N \sum_{t=1}^T \ln \left[\frac{1}{R} \sum_{r=1}^R p(y_{i,t} \mid \tilde{\xi}_{i,t}^r, y_{i,t-1};\theta) \right] = \frac{1}{N} \sum_{i=1}^N \ln \hat{h}(y_{i,1:T},\theta)$$

where $\hat{h}(y_{i,1:T},\theta) = \prod_{t=1}^T \frac{1}{R} \sum_{r=1}^R q(\xi_{i,t}^r, y_{i,t-1}, y_{i,t},\theta)$ and where $q(y_{i,t}, y_{i,t-1}, \xi_{i,t}^r,\theta) = p(y_{i,t} \mid \tilde{\xi}_{i,t}^r, y_{i,t-1};\theta)$.

Viewed in empirical process terms, let P_N denote the empirical measure associated with observations $y_{i,1:T}$ for a sample of size N and let P denote the population distribution of $y_{i,1:T}$. Similarly, let $S_R^{y_{i,1:t}}$ denote the empirical measure of the particles $\xi_{i,t}^r$ for $r = 1,\ldots,R$ generated by the particle filtering algorithm conditional on the observations $y_{i,1:t}$ and let $S^{y_{i,1:t}}$ denote the true distribution. Under the assumption that the sample consists of N of independent observations of $y_{i,1:T}$, the behavior of the likelihood function as $N \to \infty$ can be handled by standard empirical process techniques such as those used by Pakes and Pollard (1989). On the other hand, to analyze the behavior as the number of particles $R \to \infty$ we appeal to empirical process results for genetic, interacting particle systems by Del Moral and Ledoux (2000).

Returning to the likelihood function, we have $\hat{L}(\theta) = P_N \ln \hat{h}(\cdot,\theta)$ where $P_N f$ denotes $N^{-1}\sum_{i=1}^{N} f(y_{i,1:T})$ and $\hat{h}(y_{i,1:T},\theta) = \prod_{t=1}^{T} S_R^{y_{i,1:t}} q(\cdot, y_{i,t-1}, y_{i,t}, \theta)$. Likewise, the population likelihood function is $L(\theta) = P \ln h(\cdot,\theta)$ where Pf denotes $\int f\,dP$ and $h(y_{1:T},\theta) = \prod_{t=1}^{T} S^{y_{1:t}} q(\cdot, y_{t-1}, y_t, \theta)$. The functions h and q are, respectively, members of classes of functions $\mathscr{H}$ and $\mathscr{Q}$ indexed by $\theta \in \Theta$ and $(y, y', \theta) \in \mathscr{Y}^2 \times \Theta$. Below we provide conditions on the classes of functions $\mathscr{H}$ and $\mathscr{Q}$ which guarantee that $\hat{\theta}_{N,R}$ is consistent.

Theorem 1 (Consistency of $\hat{\theta}_{N,R}$)**.** *Suppose the following conditions hold:*

C0 *The parameter space, Θ, is compact.*

C1 *The class of functions $\mathscr{H} = \{h(\cdot,\theta) : \theta \in \Theta\} = \{Sq : q \in \mathscr{Q}\}$ are Euclidean in the $L^1(P)$ norm in the sense of Nolan and Pollard (1987, Definition 8) and Pakes and Pollard (1989, Definition 2.7) for the envelope H with $PH^2 < \infty$. Similarly, the class $\mathscr{Q} = \{q(\cdot, y, y', \theta) : (y, y', \theta) \in \mathscr{Y}^2 \times \Theta\}$ is Euclidean in the $L^1(S)$ norm for the envelope Q with $SQ^2 < \infty$.*

C2 *There is an $M < \infty$ such that $\sup_{y_{i,1:T},\theta} \left| \frac{1}{h(y_{i,1:T},\theta)} \right| < M$ and $\sup_{\xi,y,y',\theta} \left| q(\xi, y, y', \theta) \right| < M$.*

C3 *$L(\theta_0) > L(\theta)$ for all $\theta \in \Theta$ with $\theta \neq \theta_0$.*

If $N \to \infty$ and $R \to \infty$, then $\hat{L}_{N,R}(\theta) - \hat{L}_{N,R}(\theta_0)$ converges uniformly in probability over Θ to $L(\theta) - L(\theta_0)$ and $\hat{\theta}_{N,R} \xrightarrow{\text{p}} \theta_0$.

We verify the conditions for consistency below in Section 5 in the context of the application. Several examples of classes of functions satisfying the required uniform entropy conditions of Assumption C1 are discussed in, for example, van der Vaart and Wellner (1996) and Pakes and Pollard (1989).

4.2. Two-Step Estimation

Sequential Monte Carlo methods can also be used to allow for unobserved heterogeneity in two-step estimation methods. Here, we discuss an extension of the estimator of Bajari, Benkard and Levin (2007) which treats the first step policy and transition equation estimation as a joint maximum likelihood problem. As before, once we have specified reduced form policy functions and transition equations that are conditional on the latent state, we can use particle filter samples to integrate the likelihood with respect to the posterior distribution of the latent state. We can then form a joint first-step log-likelihood function and estimate the parameters of the reduced form policy functions and the transition equations.

Since we have controlled for the unobserved state in the first step, these estimated functions can be used to simulate the model in order to approximate the value function. The payoff function and value function in turn depend on the unobserved state since the firms' beliefs about their rivals and state transitions include the unobserved state. With estimated policy and transition equations in hand, estimation of the structural parameters becomes a computational exercise and proceeds almost exactly as in Bajari, Benkard and Levin (2007), apart from the additional state variables.

4.2.1. First Step Estimation

Strategies for estimating the policy functions in the first step tend to be model-specific and will depend on the specific distributional assumptions made. The general goal in the first step is to estimate the policy functions $\sigma_i(s_t, \eta_{it})$ and the state transition density $p(s_t \mid s_{t-1}, a_{t-1})$. Any structural parameters that appear in the transition density will be estimated, but the payoff parameters that determine the structural choice probabilities be estimated in the second step.

In order to apply the particle filter, the densities implied by the policy functions σ_i and the distribution of η_{it} must belong to some known parametric family of functions. This rules out the use of many nonparametric techniques in the first step, however, in practice researchers have typically used parametric methods such as probit and logit regressions in the first step when applying this estimator (cf. Ryan, 2012).

Assumption 7 (Parametric First Step)**.** The implied policy function belongs to a known family of functions $\mathscr{F} = \{f(\cdot;\alpha)\}$ indexed by a finite vector of parameters α. Furthermore, θ can be partitioned as $\theta = (\theta_1, \theta_2)$ so that the transition densities only depend on θ_1 and the payoffs only depend on θ_2:

$$
\begin{aligned}
&p(x_t \mid x_{t-1}, \xi_t, a_{t-1}; \theta) = p(x_t \mid x_{t-1}, \xi_t, a_{t-1}, \theta_1), \\
&p(\xi_t \mid x_{t-1}, \xi_{t-1}, a_{t-1}; \theta) = p(\xi_t \mid x_{t-1}, \xi_{t-1}, a_{t-1}, \theta_1), \\
&U_i(a_t, s_t, \eta_{it}, \theta) = U_i(a_t, s_t, \eta_{it}, \theta_2).
\end{aligned}
$$

The second part of Assumption 7, that θ can be partitioned, is usually not restrictive since the payoff parameters and transition density parameters are typically distinct. Both α and θ_1 will be estimated in the first step, leaving the payoff parameters θ_2 to be estimated in the second step. The first part of the assumption warrants more discussion.

We have maintained throughout that the payoffs and transition densities are parametric, so it follows that the policy functions are parametric; however, it may be difficult to know the correct family $\mathscr{F}$ when the policy functions are analytically intractable and can only be evaluated numerically, so the first part of Assumption 7 is restrictive. It is important to note that the class of parametric functions chosen for the first stage policy functions must be large enough to include the true policy functions in order for the model to be internally consistent. For example, if one only uses a very simple logit model for the first stage policy then most likely the actual policy function implied by the model will not be of this form. In practice, as we demonstrate in Section 5, flexible parametric specifications can yield very reasonable approximations. It also seems likely that one could use a sequence of increasingly flexible parametric families as the number of observations tends to infinity in a sieve-like manner. For consistency of the estimator, the complexity of the sieve space must increase with the sample size so that the sieves are dense in the space of functions being approximated.[7]

In the Monte Carlo experiments below, the first step is essentially sieve maximum likelihood, where the difference in choice-specific value functions is approximated by the sieve basis functions. We explore the behavior of the estimator when increasing the number of basis functions as the sample size increases.

[7] See Chen (2007) for details on sieve estimation, including sieve MLE.

Here, a formal analysis of nonparametric identification in the first step is complicated by the nonlinear and recursive definition of the value function and the corresponding lack of an analytical solution to the dynamic programming problem. It is not as straightforward as in, for example, the case of linear sample selection models, where the first step is a conditional mean (Heckman, 1979). In the present setting, object of interest in the first step is a (potentially) nonlinear function of x and ξ (and possibly η if it is not additively separable). The observable population distribution is a function of only x. However, the joint transition density of the observed and unobserved states and the distribution of the iid errors are fully specified up to finite dimensional parameters. Therefore, consistency of the first step depends on solving the following integral equation for f_0:

$$\begin{aligned} p(a_t, x_t \mid a_{1:t-1}, x_{1:t-1}) &= \int p(a_t \mid x_t, \xi_t; \theta)\, p(x_t \mid \xi_t, x_{t-1}, a_{t-1}; \theta_1)\, \pi_{t|t}(d\xi_t) \\ &= \int \left[\int f_0(a_t, x_t, \xi_t, \eta_t)\, dG(\eta_t \mid x_t, \xi_t; \theta_1) \right] p(x_t \mid \xi_t, x_{t-1}, a_{t-1}; \theta_1)\, \pi_{t|t}(d\xi_t). \end{aligned}$$

The density (or probability) on the left-hand side is potentially observable, G is known, $p(x_t \mid \xi_t, x_{t-1}, a_{t-1}; \theta_1)$ is part of the model specification, and $\pi_{t|t}$ can be determined given f_0 and parameters θ_1 for the transition density of ξ_t, also part of the model specification.

In many cases, the model places more structure on the problem. For example, in the case of a dynamic binary choice model under our assumptions, the payoff function is additively separable in the components of $\eta_t = (\varepsilon_{t0}, \varepsilon_{t1})$. Furthermore, η_t is typically independent of x_t and ξ_t and the distribution may not depend on θ_1. These properties allow us to simplify the problem as

$$\begin{aligned} &\Pr(a_t = 0 \mid a_{1:t-1} x_{1:t}) p(x_t \mid a_{1:t-1} x_{1:t-1}) \\ &\qquad = \int \left[\int 1\{\tilde{f}_0(x_t, \xi_t) + \varepsilon_{t0} - \varepsilon_{t1} \geq 0\}\, dG(\varepsilon_{t0}, \varepsilon_{t1}) \right] p(x_t \mid \xi_t, x_{t-1}, a_{t-1}; \theta_1)\, \pi_{t|t}(d\xi_t). \end{aligned}$$

In practice, the sieve maximum likelihood approach we take does not require explicitly solving the above integral equation. Rather, we approximate the function f_0 (or $\tilde{f}_0$) by a member of $\mathscr{F}$. The problem then becomes choosing α and θ_1 to maximize the integrated likelihood,

$$\iint f(a_t, x_t, \xi_t, \eta_t; \alpha)\, dG(\eta_t \mid x_t, \xi_t; \theta_1) p(x_t \mid \xi_t, x_{t-1}, a_{t-1}; \theta_1)\, \pi_{t|t}(d\xi_t; \alpha, \theta_1).$$

Hence, we maintain the assumption that $f_0 \in \mathscr{F}$ (Assumption 7). As before, structural restrictions such as additive separability or monotonicity in η_t or shape restrictions on f_0 allow us to further restrict the class $\mathscr{F}$ and possibly to limit focus to a lower-dimensional function $\tilde{f}_0$. Identification requires that there be a unique solution f_0, but consistency does not require that there be a unique α such that $f_\alpha = f_0$. We are only interested in α to the extent that estimating α allows us to estimate f_0 consistently.

As before, we obtain estimates $(\hat{\alpha}, \hat{\theta}_1)$ by maximizing the approximated log-likelihood function:

$$\hat{L}(\alpha, \theta_1) = \sum_{i=1}^{N} \sum_{t=1}^{T} \ln \left[\frac{1}{R} \sum_{r=1}^{R} f(a_{it}, x_{it}, \tilde{\xi}_{it}^r; \alpha)\, p(x_{it} \mid x_{i,t-1}, \tilde{\xi}_{it}^r, a_{i,t-1}; \theta_1) \right],$$

where $f(a, x, \xi; \alpha) = \int f(a, y, \xi, \eta; \alpha)\, dG(\eta \mid y, \xi; \theta_1)$ and where $\tilde{\xi}_{it}^r$ for $r = 1, \ldots, R$ are the particles drawn at the prediction stage of the algorithm for the i-th observational unit at period t given the parameters α

and θ_1 (i.e., draws from $\pi^R_{t|t-1}(d\xi_t;\alpha,\theta_1)$). Note that the transition density for ξ_t, which also depends on θ_1, is used for transitioning the particles and that both of the above densities are used for weighting the particles. Thus, the evolution of the particle swarm itself also depends on α and θ_1.

After obtaining $\hat{\alpha}$ and $\hat{\theta}_1$, we can simulate the model from any initial condition by sequentially drawing actions from the estimated policy densities and drawing new states from the estimated transition densities. This is all we need to carry out the forward simulation procedure of Bajari, Benkard and Levin (2007).

4.2.2. Second Step Estimation

With the first-step estimated policy and transition equations in hand, estimation of the second step parameters is computationally unchanged from that of Bajari, Benkard and Levin (2007). The primary conceptual difference is that we have estimated policies and transition equations conditional on the unobserved state. However, given the estimated policy and transition functions, we can still use forward simulation to approximate the value functions.

Let $\hat{\sigma}(s_t,\eta_t)$ denote the joint policy function associated with the estimates $\hat{\alpha}$. Given values for the remaining structural parameters θ_2, the ex-ante value function implied by these policies is

$$\bar{V}_i(s;\hat{\sigma},\hat{\theta}_1,\theta_2) = \mathrm{E}\left[\sum_{t=0}^{\infty}\beta^t U_i\left(\hat{\sigma}(s_t,\eta_t),s_t,\eta_{it},\theta_2\right)\,\middle|\, s_0 = s\right]$$

where the expectation is taken with respect to $\{s_t,\eta_t\}_{t=0}^{\infty}$ under $(\hat{\alpha},\hat{\theta}_1)$.

The structural parameters θ_2 can be estimated as usual by treating the estimated policies $\hat{\sigma}$ as the true policies in the equilibrium conditions,

$$\bar{V}_i(s;\hat{\sigma}_i,\hat{\sigma}_{-i},\hat{\theta}_1,\theta_2) \geq \bar{V}_i(s;\sigma_i',\hat{\sigma}_{-i},\hat{\theta}_1,\theta_2) \quad \forall(i,s,\sigma_i'),$$

and using them to form a minimum distance objective function

$$Q(\theta_2) \equiv \int \left[\min\left\{\bar{V}_i(s;\hat{\sigma}_i,\hat{\sigma}_{-i},\hat{\theta}_1;\theta_2) - \bar{V}_i(s;\sigma_i',\hat{\sigma}_{-i},\hat{\theta}_1;\theta_2),0\right\}\right]^2 dH(i,s,\sigma_i')$$

where H is a distribution over the set of possible inequalities—combinations (i,s,σ_i'). By minimizing this function, we minimize the sum of squared violations of the equilibrium conditions, motivated by revealed preference. Under the true parameters θ_2, the true policy should always yield higher discounted future payoffs than any alternative policy σ_i' for each agent i.

In practice, even with fully observed state variables, this procedure is sensitive both to the first step estimates and the chosen distribution of agent indices, states, and alternative policies. In fact, the structural parameters may not be identified under Q for some choices of the distribution H (Srisuma, 2010). We explore the performance of this estimator in detail, both with fully observed data and with a latent state, in the following section.

5. An Application to the Model of Rust (1987)

In this section, we develop an extension of the classic bus engine replacement model of Rust (1987). The extended model has a two-dimensional continuous state space, rather than a one-dimensional discrete

state, and one of these states is a serially correlated unobserved state variable. We first use the model to obtain parameter estimates using the original data using the full-solution estimator described above in Section 4.1. We then carry out a series of Monte Carlo experiments with true parameter values inspired by our estimates with the actual data to illustrate both the full-solution and two-step estimators.

The agent has two choices each period, to overhaul the engine of a bus, $a_t = 1$, or to do nothing, $a_t = 0$. The two state variables are the observed mileage, x_t, and the latent state of the engine, ξ_t, which is observed by the decision maker but not by the researcher and may be serially correlated. The time period is one month and the discount factor is $\beta = 0.95$.

Here, ξ_t may represent several unobserved factors. Persistent features of the engine (e.g., unobserved quality), the bus (e.g., an inexperienced or abusive driver), or the route which it serves (e.g., difficult terrain, traffic congestion, unobserved changes in the route) may all be relevant. Furthermore, the binary decision may affect the distribution of ξ_t and ξ_t may influence the distribution of x_t. For example, a bus serving a route in an area of the city with heavier traffic (which is unobserved and time-varying) may accrue mileage at a lower rate than buses serving more rural, longer-distance routes. Another example of feedback would be unobserved variation in the route itself—perhaps the route is extended, shortened, or changed completely over the sample period, thus changing the distribution of mileage increments.

First, we specify functional forms for the model primitives: the payoff (cost) function and the transition densities for x and ξ. For simplicity, we assume the cost function is linear:

$$U(a_t, s_t, \varepsilon_t, \theta_2) = \begin{cases} -c_x x_t - c_\xi \xi_t + \varepsilon_{t,0} & \text{if } a_t = 0, \\ -c_0 + \varepsilon_{t,1} & \text{if } a_t = 1. \end{cases}$$

The structural parameters of interest are the replacement cost, c_0, the cost of mileage, c_x,[8] and the cost associated with the latent state, c_ξ. In threshold-crossing models such as this one, the coefficients are only identified relative to the variance of the error term. We will be considering models where the variance of the compound error $-c_\xi \xi_t + \varepsilon_{t,0} - \varepsilon_{t,1}$ depends on the parameter c_ξ, and we consider cases where both $c_\xi > 0$ and $c_\xi = 0$. Therefore, we simply report the estimated cost ratios c_0/c_x and c_ξ/c_x, which are costs relative to the cost of mileage and are comparable across specifications.

When the choice to continue is made in the previous period ($a_{t-1} = 0$), given x_{t-1} and ξ_t, increments to the observed state, $\Delta x_t = x_t - x_{t-1}$, follow an exponential distribution with density

$$p(\Delta x_t \mid x_{t-1}, \xi_t, a_{t-1} = 0; \theta_1) = \lambda(\xi_t, \theta_1)\, e^{-\lambda(\xi_t, \theta_1)\Delta x_t} \quad \text{where} \quad \lambda(\xi_t, \theta_1) = \exp(\lambda_0 + \lambda_\xi \xi_t). \tag{7}$$

This ensures that the rate parameter $\lambda(\xi_t, \theta_1)$ is positive while allowing the unobserved quality of the engine to influence the mileage transition process. The latent state follows a mean-zero AR(1) process when $a_{t-1} = 0$ with $p(\xi_t \mid x_{t-1}, \xi_{t-1}, a_{t-1} = 0; \theta) = \varphi(\xi_t, \rho\xi_{t-1}, \sigma^2)$ where $\varphi(\cdot, \mu, \sigma^2)$ is the pdf of the Normal distribution with mean μ and variance σ^2.[9] When the engine is replaced in the previous period ($a_{t-1} = 1$), the mileage resets to zero and the mileage for period t is exponentially distributed according to the

[8] Following Rust (1987), we scale mileage, by 0.01 so that all parameters are roughly on the same order of magnitude. All coefficients on x_t should be interpreted in light of this scaling.

[9] We use the normalization $\sigma = 0.5$ and estimate coefficients on ξ_t in the cost function and transition equation.

distribution above; the latent state is drawn anew from the stationary distribution: $p(\xi_t \mid x_{t-1}, \xi_{t-1}, a_{t-1} = 1; \theta) = \varphi(\xi_t, 0, \sigma^2/(1-\rho^2))$.

We must solve for the choice-specific value function $v(a, s, \theta)$ both to simulate data and to estimate the model. Recall that v satisfies $v - \Gamma(v) = 0$, where Γ is the functional operator defined in (4). Although v is an infinite-dimensional object, we can form a smooth approximation to v with only a finite number of parameters using Chebyshev polynomials (cf. Judd, 1998). See Appendix A for details.

For consistency of the maximum filtered estimator, we consider the entropy condition in Assumption C1 of Theorem 1. In this model, the class $\mathscr{Q}$ which generates the likelihood function is

$$\mathscr{Q} = \left\{ \exp(\lambda_0 + \lambda_\xi \xi_t) \exp[-\lambda(\xi_t, \theta_1)\Delta x_t] \frac{\exp(v(0, x_t, \xi_t, \theta) - v(1, x_t, \xi_t, \theta))}{1 + \exp(v(0, x_t, \xi_t, \theta) - v(1, x_t, \xi_t, \theta))} \right\}.$$

We first appeal to the fact that exponentiation is Lipschitz on compact sets. We assume the parameter space Θ is compact and in practice we limit the state variables to compact spaces as described in Appendix A. That the class of functions $\mathscr{Q}$ is Euclidean then follows from Lemmas 2.13 and 2.14 of Pakes and Pollard (1989) provided that the choice-specific value function v is well-behaved. Norets (2010) gives conditions for continuity and differentiability of value functions in dynamic discrete choice models.

Finally, we note that for consistency Theorem 1 also requires that the initial distribution of ξ_0 be correctly specified. Otherwise, there could be an initial conditions problem when $T < \infty$ (Heckman, 1981). As a robustness check, we intentionally use a misspecified initial distribution for this model in the Monte Carlo experiments below and find it does not appreciably affect the results.

5.1. Empirical Results

We first estimate the model using the original data of Rust (1987). Keeping in line with the groupings chosen by Rust (1987), we use the data for bus groups 1–3 (pooled) and group 4. These groupings were based on heterogeneity tests of the mileage processes between the groups. Table 1 summarizes the data for the four groups. Groups 1–3 consist of 67 buses, each observed over a period of between 25 and 70 months for a total of 3,931 bus-month observations. Group 4 consists of 37 buses, each observed over 117 months for a total of 4,329 observations. The sample mean and standard deviation of the monthly mileage increments for each group are also reported along with the number of engine replacements and the average cost of replacement including labor and parts.

One distinction that will prove important is that the mean monthly mileage increment for group 4 is much smaller than for groups 1–3, but the standard deviation is nearly twice as large, indicating that there may be more heterogeneity among buses in group 4 than among those in groups 1–3. Rust (1987) tested for heterogeneity within groups and failed to reject the null hypothesis that the distribution of mileage increments for buses within groups 1–3 and within group 4 are, respectively, identical within those groupings. However, he used a discrete-state-space model and based the tests on a multinomial distribution for discretized mileage transitions with three states (increasing by less than 5,000 miles, increasing by 5,000–9,999 miles, or increasing by 10,000 miles or more).

Since we use a continuous-state-space model, we carry out a test for heterogeneity by supposing

that mileage increments are exponentially distributed[10] and estimating an unrestricted model with bus-specific parameters and a restricted model with a common parameter. In other words, this is a reduced-form test for within-group heterogeneity. For each bus group $j = 1,2,3,4$ we suppose mileage increments $x_{it} - x_{i,t-1}$ for each bus i follow an exponential distribution with rate parameter μ_{ij}. We also estimate a homogeneous specification where there is a common rate parameter $\mu_j = \mu_{1j} = \cdots = \mu_{N_j,j}$ shared by all N_j buses in group j. We then carried out a likelihood ratio test for the null hypothesis $H_0 : \mu_{1j} = \mu_{2j} = \cdots = \mu_{N_j,j}$ for each group j. The results given in Table 2 show that we also cannot reject the null hypothesis of homogeneity for any particular group, however, there appears to be substantially more evidence for heterogeneity within group 4 ($p = 0.37$) than for groups 1–3 ($p \geq 0.93$). This is in contrast to the findings for the discretized mileage processes, for which the respective p-values were all above 0.85 (Rust, 1987, Table V, p. 1017). Our findings with the continuous mileage data support the separate analysis of group 4 and also foreshadow the type of heterogeneity reflected by our estimates.

Table 3 reports the estimates obtained using data for bus groups 1–3 and bus group 4. The first row for each grouping gives the naïve estimates obtained when ignoring the latent state ξ_t. Then, we report the estimates obtained using $R = 5,000$ particles. We also used $R = 500$ and $R = 9,000$ with qualitatively very similar results. Furthermore, the Monte Carlo experiments below show that the estimates can be nearly unbiased with as few as $R = 50$ particles.

The estimates were obtained by choosing the initial distribution π_0 to be $\mathrm{N}(0, \sigma_0^2)$ with $\sigma_0 = 2.0$. We also used $\sigma_0 = 0.5$, $\sigma_0 = 1.0$, and $\sigma_0 = 3.0$ with similar results. The estimates were also robust to different seed values for the underlying random number generator, which result in different filter samples. The standard errors were calculated using 250 bootstrap replications. The computational times reported are times in minutes to obtain the "observed" estimate (i.e., not including the bootstrap replications).

Overall, for groups 1–3 we find some evidence of a moderately persistent latent state variable ($\rho = 0.684$) that primarily affects costs but not mileage increments. For comparison, we consider the effect of a one-standard-deviation increase in ξ_t (relative to the stationary distribution, with standard deviation 0.685). Given the estimates, this would decrease the mean of the distribution of mileage increment by 203 miles (1.3%) and increase costs by as much as 1,833 additional miles.

For group 4, we find evidence of a latent state variable that is strongly serially correlated ($\rho = 0.982$) and, while it has a relatively smaller effect on costs, it has a larger, positive effect on mileage transitions. Through our parametrization of the scale parameter of the exponential distribution in (7), when λ_ξ is negative higher values of the latent state ξ_t decrease $\lambda(\xi_t, \theta_1)$ which in turn increases the mean of the mileage increment. That is, buses with larger values of ξ_t tend to be driven more miles each month. A one-standard-deviation increase in ξ_t (2.65) increases the mean of mileage increments by 3,806 miles (36%) and increases costs as much as 384 miles.

Because we use a large number of particles for estimating the model with the real data, the computational times for the particle filter estimator are longer than for the naïve estimator: 38 times longer for groups 1–3 and 54 times longer for group 4. However, even for group 4 the time is under two hours.

[10] The exponential distribution was suggested by Rust (1987) for mileage increments in the general description of the model, before the discrete-state-space model is introduced.

An alternative to the proposed procedure would be to evaluate the likelihood using Monte Carlo integration by simulating entire paths $\xi_{1:T}$ to approximate the T-dimensional integral in (2). On the other hand, the particle filtering procedure yields draws of ξ_t conditional on the observed data and uses them to approximate the one-dimensional integrals in (3). For both procedures, one needs to propagate the draws through the transition kernel and evaluate the likelihood for an observation (conditional on both x_t and ξ_t). The main additional cost of the particle filter is in the resampling step, which is non-trivial but also does not dominate the computational time. Therefore, given an equal number of draws the particle filter uses those draws more effectively.[11]

Using the results from Table 3, we also perform likelihood ratio tests for the null hypothesis that the constrained or naïve model is the correct model. For groups 1–3, we cannot reject the constrained model at any of the usual levels. On the other hand, for group 4 we strongly reject the null. There appear to be strongly persistent differences in mileage transitions that are important for this group, as well as persistent differences in costs. This is in line with the reduced form analysis above, where we found more evidence for heterogeneity in mileage transitions for group 4, as shown in Table 2.

To compare the predictions of the two models, we plot the estimated replacement hazard functions in Figure 1 for groups 1–3 (panel a) and group 4 (panel b). For groups 1–3, the mean mileage at replacement is 200,684. For these groups, the hazard between is slightly higher 100,000 and 200,000 miles for the model including ξ_t but substantially lower for higher mileages. For group 4, although the hazard for the model including ξ_t is slightly larger for mileages over 150,000, the hazards are largely similar as reflected by small estimated cost parameter c_ξ. Importantly, although the estimated hazards are similar for group 4 other implications of the model, such as the replacement demand functions, may differ in important ways.

We can also find differences in economic implications due to unobserved heterogeneity when we use the estimates to derive the demand function for new bus engines. Rust (1987, Table III) reports the average total replacement costs: \$9,499 for groups 1–3 and \$7,513 for group 4. We approximate the demand function by simulating[12] the steady state unconditional replacement probability under different replacement cost values ranging from \$1,000 to \$15,000. Figure 2 plots the demand functions for groups 1–3 and group 4 using both the naïve estimates and the MFL estimates. The vertical lines indicate the actual average engine replacement costs reported by Rust (1987) and which are shown in Table 1. The naïve estimates appear to over-estimate demand at low costs for both groups 1–3 and group 4. Thus, assuming that the unobservables are independent over time may have non-trivial economic implications. Allowing for serially correlated unobserved heterogeneity yields similar implications for demand when costs are above \$8,000, but demand is much lower for low replacement cost values. Annual demand is nearly half for groups 1–3 when the replacement cost is \$1,000 and is nearly one-fifth for group 4.

Finally, we note that Norets (2009a) also estimated a similar model that allowed for serially correlated unobserved heterogeneity using data for bus group 4. He found only moderate evidence for serial correlation while we find strong serial correlation. Two differences in our models explain the difference. First, although Norets (2009a) also allows for serially correlated unobserved costs the model considered

[11] Indeed, we carried out Monte Carlo experiments using MSL using the same model and found appreciable bias with as many as $R = 400$ simulated draws while the bias for the MFL estimates was negligible with $R = 50$ particles.

[12] We report average replacement probabilities using 101,000 simulated periods and discarding the first 1,000 periods.

here additionally allows the serially correlated unobservable to manifest itself in the form of heterogeneity in the distribution of mileage increments. As noted above, there is more evidence of heterogeneity in the mileage process for group 4 than for other groups. Second, we treat mileage as a continuous variable while Norets discretized it as Rust (1987) originally did. Rust (1987) previously carried out likelihood ratio tests for the hypothesis that the discretized mileage process was the same for all buses within a group, but he found no evidence against this hypothesis for any bus group. However, once we use the continuous mileage data there is more evidence of heterogeneity.

5.2. Monte Carlo Experiments: Full Solution Estimation

In this section we explore the results of a series of Monte Carlo experiments carried out by estimating the continuous state model described above using the maximum filtered likelihood estimator. We choose true parameter values that are based on the estimates above obtained using the data from Rust (1987).

Recall that in addition to estimating the parameters θ, the sequence of distributions of the latent state for each observational unit (here, for each bus) can be estimated as well. This is illustrated in Figure 3(a), which displays a single simulation from the continuous state bus engine replacement model along with the distributions of the unobserved state for each period t (conditional on the observed data up to period t) as approximated using the particle filter. The realization of x_t is plotted in the lower panel. The realization of ξ_t is plotted in the upper panel along with the quantiles of the approximate posterior distribution. Figure 3(b) plots the same weighted particle swarm as a sequence of distributions. The parameters used for these simulations are the same those used in the Monte Carlo experiments that follow: $\lambda_0 = 0.4$, $\lambda_\xi = 0.3$, $\rho = 0.8$, $c_0 = 14$, $c_x = 2$, and $c_\xi = 0.5$. The initial proposal distribution here was taken to be normal with mean zero and standard deviation $\sigma_0 = 3.0$. Notice how the particle swarm is initially very dispersed (due to the large value of σ_0), but quickly assigns more weight to more relevant areas of the state space as information from the data is incorporated.

Finally, we generate several datasets consisting of $N = 100$ bus observations of $T = 100$ periods each and estimate the model several ways. Table 4 reports the mean, standard deviation, mean bias, and root mean squared error of the parameter estimates over 25 replications for several different estimators. The panels of this table proceed from less information about the latent state at the top to more information at the bottom. We take the initial distribution of the swarm to be $\tilde{\pi}_0 = \mathrm{N}(0, \sigma_0^2)$ with $\sigma_0 = 0.5$. Although the sample size is fixed throughout, increasing the number of particles results in smaller mean squared errors which approach the ideal but infeasible case where ξ_t is observed and treated as data.

In the data generating process, we generate ξ_0 by drawing from the stationary distribution, $\mathrm{N}(0, \sigma^2/(1-\rho^2))$. The initial distribution π_0 is an importance sampling distribution and is intentionally chosen to be different than the true distribution to illustrate the adaptive behavior of the filtering distribution and robustness of the estimates to misspecification of the initial proposal distribution. This can be seen in Figure 3(a), where the effect of the initial distribution π_0 diminishes quickly.[13]

First, Table 4 reports the naïve estimates obtained when ignoring the latent state ξ_t. In this case we

[13] We also carried out Monte Carlo experiments using datasets with $T = 10$ time periods and $N = 1{,}000$ cross-sectional units and found that the estimator also performed well in this setting despite the increased importance of the initial conditions.

are only estimating the average mileage transition rate λ_0, without capturing the effects of ξ_t. As such, our estimates of λ_0 are biased. Similarly, we are overestimating the replacement cost ratio c_0/c_x by about 21%. Any welfare calculations or policy implications based on these estimates would suffer from this bias.

The following rows present the maximum filtered likelihood estimates, which perform quite well in all cases. With only $R = 25$ particles, the replacement cost ratio estimates and all other parameters have much smaller bias than the naïve estimator. Estimates of the latent cost ratio are quite good for all values of R. The bottom panel reports the maximum likelihood estimates obtained when the latent state is fully observed. It is worth nothing that the estimates with only $R = 50$ particles have as little bias as the estimates with fully observed data.

The computational times for the MFL estimates with $R = 25$ are only 40% longer than for the case of fully observed data. For $R = 100$, the times are a little more than twice as long on average. The computational times for the $R = 25$ case are about seven times longer than for the naïve estimator, but the naïve estimator has over eight times more bias in the relative replacement cost. Note that there is some savings in terms of the total number of functional evaluations required when adding more particles, since the objective function becomes more smooth and easier to optimize. The trade-off, of course, is that using more particles makes each functional evaluation more costly.

5.3. Sources of Bias

In this section, we consider the potential sources of bias when ignoring unobserved heterogeneity in the form of latent, serially correlated state variables. First, we consider two possible sources of bias in the particular optimal renewal model estimated above: time-varying unobserved heterogeneity in the distribution of mileage increments (λ_ξ) and time-varying unobserved costs (c_ξ). Recall that ξ_t can be either positive or negative and that it enters the payoff function as $-c_\xi \xi_t$. We impose the normalization that $c_\xi \geq 0$ so that we interpret positive values of ξ_t as costs and negative values as benefits. Table 5 contains estimates for the naïve estimator for four population parameter choices which are designed to illustrate the roles of the two effects in determining the bias from ignoring the unobserved heterogeneity.

When latent costs are important ($c_\xi > 0$), then the relative replacement cost is biased upward as can be seen in specifications 1–3 in Table 5. In comparing these cases, we can also see that the bias is affected by the degree to which the latent state influences the distribution of mileage transitions, since the bias varies with λ_ξ across specifications 1–3. When there are no unobserved costs ($c_\xi = 0$), then the bias in the relative replacement cost is zero independent of the value of λ_ξ, as seen in specification 4. However, of course the mileage transition parameter λ_0 (for which the mean increment is $1/\lambda_0$) is still biased in this case and in all cases for which $\lambda_\xi \neq 0$.

For this specific model, the estimated replacement cost ratio c_0/c_x tends to be biased upwards. The source of this bias is similar to selection bias in a static model, but there is dynamic selection here on unobservables that are time-varying. In particular, we are more likely to observe a bus accrue additional mileage (the engine is not replaced) when $\xi_t < 0$ (i.e., the bus is unobservably less costly). Unobservably more costly buses tend to have replacements early, so they select out. Therefore, if we could observe ξ_t in the sample then the average value would be negative. So, when using the naïve estimator, since observed

buses are on unobservably *less* costly ($\xi_t < 0$) on average, this is attributed to a lower estimated cost of mileage and a higher cost of replacement. The overall effect is that on average c_0/c_x is biased upwards.

Finally, we note that there are several special features of this model: the choice variable is binary and has a renewal effect, the cost function is linear, the observed state is monotonically increasing unless a replacement occurs, and the unobserved state variable follows an AR(1) process with Normally-distributed innovations. Therefore, it seems unlikely that either the directions of the biases found above or the clean separation of those biases would extend universally to more general models which may not share all of these features. This ambiguity further underscores the need to allow for non-iid error structures.

5.4. Monte Carlo Experiments: Two-Step Estimation

Here we apply the two step estimator discussed above to estimate the model. Under additive separability, the policy function σ satisfies the optimality condition

$$\sigma(s,\varepsilon,\theta) = 0 \iff v(0,s,\theta) + \varepsilon_0 \geq v(1,s,\theta) + \varepsilon_1.$$

If the iid choice-specific errors are distributed according to the type I extreme value distribution, then the corresponding choice probabilities are

$$P(\sigma(s,\varepsilon,\theta) = 0 \mid s;\theta) = \frac{\exp(v(0,s,\theta) - v(1,s,\theta))}{1+\exp(v(0,s,\theta) - v(1,s,\theta))} \quad \text{and} \quad P(\sigma(s,\varepsilon,\theta) = 1 \mid s;\theta) = 1 - P(\sigma(s,\varepsilon,\theta) = 0 \mid s;\theta).$$

Now, suppose that we can approximate the differences $v(0,s,\theta) - v(1,s,\theta)$ by $\tilde{f}(s;\alpha)$. Then it follows that we can use a simple logit model to approximate the true policy:

$$P(\sigma(s,\varepsilon,\theta) = 0 \mid s,\theta) \approx f(a_t = 0, s_t;\alpha) = \frac{\exp(\tilde{f}(s;\alpha))}{1+\exp(\tilde{f}(s;\alpha))}.$$

We choose $\mathscr{F}$ to be the collection of polynomials of the form

$$\tilde{f}(s;\alpha) = \alpha_1 + \alpha_2 x + \alpha_3 \xi + \alpha_4 x\xi + \alpha_4 x^2 + \alpha_6 \xi^2 + \alpha_7 x^2\xi + \alpha_8 x\xi^2 + \alpha_9 x^3 + \alpha_{10}\xi^3. \tag{8}$$

The parametric specifications for the state transition densities are specified as part of the model and so we can estimate the parameters θ_1 of those densities and the parameters α from the parametric policy approximation using maximum likelihood. As before, we can use the particle filter samples to approximate the likelihood function.

With estimates $(\hat{\alpha}, \hat{\theta}_1)$ in hand, the approximate policy function is simply $\hat{\sigma}(s,\varepsilon) = 1\{\tilde{f}(s;\hat{\alpha}) + \varepsilon_0 - \varepsilon_1 \leq 0\}$. Given $\hat{\alpha}$, we know $\hat{\sigma}$ and given $\hat{\theta}_1$, we know the estimated transition densities. We can therefore use forward simulation to approximate the ex-ante value function $\bar{V}_i(s;\sigma,\theta)$ for any s.

In the second step, we follow Bajari, Benkard and Levin (2007) in using linearity of the value function to reduce the computational burden. Note that we can write the payoff function as

$$U(a_t, s_t, \varepsilon_t, \theta_2) = (1-a_t)\left(-c_x x_t - c_\xi \xi_t + \varepsilon_{t,0}\right) + a_t\left(-c_0 + \varepsilon_{t,1}\right).$$

For a given policy function $a_t = \sigma(s_t, \varepsilon_t)$, the corresponding ex-ante value function for some state s_0 is

$$\bar{V}(s_0;\sigma,\theta) = \mathrm{E}\left[\sum_{t=0}^{\infty}\beta^t U(a_t,s_t,\varepsilon_t,\theta_2)\right] = \mathrm{E}\left[\sum_{t=0}^{\infty}\beta^t\left((1-a_t)\varepsilon_{t,0}+a_t\varepsilon_{t,1}\right)\right] - c_0\,\mathrm{E}\left[\sum_{t=0}^{\infty}\beta^t a_t\right] - c_x\,\mathrm{E}\left[\sum_{t=0}^{\infty}\beta^t(1-a_t)x_t\right] - c_\xi\,\mathrm{E}\left[\sum_{t=0}^{\infty}\beta^t(1-a_t)\xi_t\right].$$

We can then approximate $\bar{V}(s_0;\hat{\sigma},\hat{\theta}_1,\theta_2)$ by simulating the L paths of length $\bar{T}$ under $(\hat{\alpha},\hat{\theta}_1)$, with each path starting at s_0. We obtain L sequences $\{a_t^l, x_t^l, \xi_t^l\}_{t=1}^{\bar{T}}$ for $l = 1,\ldots,L$, where $\bar{T}$ is chosen so that $\beta^{\bar{T}}$ is sufficiently small. Hence, given first-step estimates $(\hat{\alpha},\hat{\theta}_1)$, the discounted payoffs can be accumulated to approximate the ex-ante value function at any state s_0 as:

$$\hat{V}(s_0;\hat{\sigma},\hat{\theta}_1,\theta_2) = \frac{1}{L}\sum_{l=1}^{L}\sum_{t=0}^{T}\beta^t U\left(\hat{\sigma}(s_t^l,\varepsilon_t^l),s_t^l,\varepsilon_{it}^l,\theta_2\right) = \left[\frac{1}{L}\sum_{l=1}^{L}\sum_{t=0}^{T}\beta^t\left((1-a_t^l)\varepsilon_{t,0}^l+a_t^l\varepsilon_{t,1}^l\right)\right] - c_0\left[\frac{1}{L}\sum_{l=1}^{L}\sum_{t=0}^{T}\beta^t a_t^l\right] - c_x\left[\frac{1}{L}\sum_{l=1}^{L}\sum_{t=0}^{T}\beta^t(1-a_t^l)x_t^l\right] - c_\xi\left[\frac{1}{L}\sum_{l=1}^{L}\sum_{t=0}^{T}\beta^t(1-a_t^l)\xi_t^l\right].$$

Notice that this approximation to the ex-ante value function is linear in the parameters c_0, c_x, and c_ξ and that the summation terms are independent of the parameters. Thus, they can be pre-calculated so that the value function, and thus the objective function $Q(\theta_2)$, can be quickly calculated for any value of θ_2.

In each of these experiments, we use 2,000 inequalities—state and alternative policy combinations. For each inequality, we simulate $L = 1{,}000$ paths of length $\bar{T} = 250$.Alternative policies are based on perturbations of the first-step parameters α_1, α_2, α_3, and α_4 of up to ±10%. Specifically, we draw γ_1, γ_2, γ_3, and γ_4 as $\gamma_j \sim \mathrm{U}(-\alpha_j/10, \alpha_j/10)$ for $j = 1,\ldots,4$ and perturb the threshold crossing condition by the amount $\gamma_1+\gamma_2 x_t+\gamma_3\xi_t+\gamma_4 x_t\xi_t$. This is equivalent to using the policy function where the first four parameters are $\tilde{\alpha}_1,\tilde{\alpha}_2,\tilde{\alpha}_3,\tilde{\alpha}_4$, with $\tilde{\alpha}_j = \alpha_j+\gamma_j$, and where the remaining parameters are unchanged. Initial states for each inequality were drawn uniformly over the state space.

The results of our experiments are summarized in Tables 6 and 7. In Table 6, we report the means and standard deviations of the first step estimates while considering the sensitivity to the choice of κ, the dimension of the polynomial in (8) used in the first step. We consider $\kappa = 3,4,6,8,10$ where for $\kappa = 3$ only the first three terms are included and for $\kappa = 10$, all terms are included. Then, Table 7 reports the corresponding second step estimates for each κ.

Note that the parameters of interest are λ_0, λ_ξ, and ρ from the first step, shown in Table 6, and the ratios c_0/c_x and c_ξ/c_x from the second step, shown in Table 7. The first step parameters $\alpha_1,\ldots,\alpha_{10}$ are not of direct interest but are a means to obtaining estimates of the choice probabilities as a function of x and ξ. To ensure consistency of the second step estimates, the predicted choice probabilities must be consistent.

It is useful to understand how much variation is inherent in the Bajari et al. (2007) procedure and how much is due to the addition of the particle filter. Therefore, for all experiments we compare the estimator with $R = 100$ particles to the hypothetical case where both x and ξ are fully observed by the econometrician. With the exception of $\kappa = 8$, the difference in the mean estimates is small but the standard deviation of the particle filter estimates is larger, as expected.

In addition, we also present values of the infeasible second-step estimates obtained using the true policy functions (note that there is no data involved in this stage, so the first step is bypassed completely). In this case, the true policy function is obtained by actually solving the model and is then used to simulate the ex-ante value function and form the minimum distance objective function. That there is nearly no bias in this case indicates that any bias present with the estimated policy functions is due to first step estimation error. Importantly, for all values of κ—even the obviously misspecified linear case with $\kappa = 3$—there is less bias in the ratios c_0/c_x and c_ξ/c_x is smaller than for the naïve estimator (see Table 4).

Our overall impression from these results is that the two-step estimator behaves well on average and although the estimates are noisier than the full-solution maximum likelihood estimates, they are still preferable to the naïve estimates. For larger values of κ, the approximation error from including the higher order power series terms seems to be a relatively large component of the standard deviations. However, using orthogonal polynomials such as Hermite, Laguerre, or Chebyshev polynomials might yield better results. The second-step estimates using the true policy function are quite good, indicating that our choices of $\bar{T}$, L, and the form of our alternative policy perturbations are reasonable. In large-scale discrete time dynamic games, where full-solution estimation is infeasible, the additional noise may be a small price to pay for the ability to obtain estimates that control for the latent state variable.

6. Conclusion

This paper has shown that several common dynamic microeconomic models with serially correlated latent state variables can be written in a nonlinear state space form, to which we can apply a particle filter to approximate the distribution of the latent state. We have proposed two estimators, a full-solution maximum filtered likelihood estimator in the spirit of the nested-fixed point estimator of Rust (1987) and a two-step method based on the estimator of Bajari, Benkard and Levin (2007). In both cases, applying a particle filter is straightforward and only requires evaluating and drawing from densities that arise naturally as part of the model specification. We then apply the full-solution estimator to generalized version of the bus engine replacement model of Rust (1987) using Rust's original dataset and find strong evidence of serially correlated latent state variables in group 4 and moderate evidence for groups 1–3. We then provide Monte Carlo evidence to highlight the performance of both estimators.

A. Computational Details of the Empirical Model

This section contains computational details of the empirical model used for the Monte Carlo experiments and empirical results in Section 5. In all cases, we use degree fourteen Chebyshev polynomials to approximate the value function (in each dimension) and seventh-order quadrature to approximate the double integral in the functional operator Γ, defined in (4), in line with the procedures described below.

A.1. Chebyshev Approximation to $v(a,s,\theta)$

Let T_k denote the k-th degree Chebyshev polynomial of the first kind, defined on the interval $[-1,1]$. The values $T_k(x)$ can be calculated using the trigonometric identity $T_k(x) = \cos(k \arccos x)$ or through the

recurrence relation $T_0(x) = 1$, $T_1(x) = x$, $T_{k+1}(x) = 2xT_k(x) - T_{k-1}(x)$. Let $\mathscr{H}_k$ denote the set of roots of the k-th degree Chebyshev polynomial of the first kind,

$$\mathscr{H}_k = \left\{ \cos\left(\frac{\pi}{2} \frac{2j-1}{k} \right) : j = 1, \ldots, k \right\},$$

and let $\mathscr{H}_k^{[a,b]}$ denote the corresponding roots scaled to the $[a,b]$ interval,

$$\mathscr{H}_k^{[a,b]} = \left\{ \frac{x+1}{2}(b-a) + a : x \in \mathscr{H}_k \right\}.$$

For each $a \in \mathscr{A}$ and θ, we approximate $v(a,x,\xi,\theta)$ by a function $v_\psi(a,x,\xi)$ over the region $[\underline{x},\overline{x}] \times [\underline{\xi},\overline{\xi}]$ by taking products of Chebyshev polynomials of degree $K-1$ in each dimension:

$$v_\psi(a,x,\xi) \equiv \sum_{i=1}^{K} \sum_{j=1}^{K} \psi_{ij}^a T_{i-1}\left(2\frac{x-\underline{x}}{\overline{x}-\underline{x}} - 1 \right) T_{j-1}\left(2\frac{\xi-\underline{\xi}}{\overline{\xi}-\underline{\xi}} - 1 \right),$$

where the ψ subscript denotes dependence on a collection Ψ of $K \times K$ coefficient matrices $\Psi^a = \left(\psi_{ij}^a \right)$ for each choice. Note that there is a different matrix of coefficients for each choice a, corresponding to the different functions $v(a,\cdot,\cdot,\theta)$. The dependence of coefficient matrices Ψ^a on θ is implicit. Letting $\tilde{x}$ and $\tilde{\xi}$ denote the values of x and ξ scaled from $[\underline{x},\overline{x}] \times [\underline{\xi},\overline{\xi}]$ to $[-1,1]^2$, we can write this more succinctly as

$$v_\psi(a,x,\xi) = \sum_{i=1}^{K} \sum_{j=1}^{K} \psi_{ij}^a T_{i-1}(\tilde{x}) T_{j-1}(\tilde{\xi}).$$

Intuitively, for a given θ one wants to choose the values ψ_{ij}^a in order to make the difference between $v_\psi(a,x,\xi)$ and $\Gamma(v_\psi)(a,x,\xi)$ small in some sense over the entire state space. We choose the coefficients that minimize the squared residuals over the set of Chebyshev roots for a given θ:

$$Q_\theta(\Psi) = \sum_{a \in \mathscr{A}} \sum_{(x,\xi) \in \mathscr{H}_k^{[\underline{x},\overline{x}]} \times \mathscr{H}_k^{[\underline{\xi},\overline{\xi}]}} \left[v_\psi(a,x,\xi) - \Gamma_\theta(v_\psi)(a,x,\xi) \right]^2. \tag{9}$$

Assume for a moment that we can evaluate Γ numerically. Then, for each θ the value of Ψ which minimizes $Q_\theta(\Psi)$ is used to approximate $v(\cdot,\cdot,\cdot,\theta)$ in the log-likelihood function. In the Monte Carlo experiments and application, we use Newton's method to solve for the coefficients Ψ, which are zeros of (9).

A.2. Evaluating Γ

To evaluate Γ numerically in practice, we use quadrature to approximate the required double integral. In particular, we use Gauss-Laguerre quadrature for the integral with respect to the exponential distribution (conditional on ξ) and Gauss-Hermite quadrature for the integral with respect to the normal distribution.

Gauss-Laguerre quadrature of order n provides abscissæ ζ_i and weights ω_i for $i = 1, \ldots, n$ for the following linear approximation:

$$\int_0^\infty e^{-\zeta} \varphi(\zeta)\, d\zeta \approx \sum_{i=1}^{n} \omega_i \varphi(\zeta_i).$$

This form of quadrature is useful for approximating the expectation of a nonlinear function of an exponentially-distributed random variable. If X is an exponential random variable with rate parameter λ, then the expectation of $f(X)$ can be approximated via a simple transformation. Let $\zeta = \lambda x$ and $\varphi(\cdot) = f(\cdot/\lambda)$ (and note that $d\zeta = \lambda\, dx$). Then, by a change of variables,

$$\mathrm{E}[f(X)] = \int_0^\infty \lambda \mathrm{e}^{-\lambda x} f(x)\, dx = \int_0^\infty \mathrm{e}^{-\zeta} \varphi(\zeta)\, d\zeta \approx \sum_{i=1}^n \omega_i f(\zeta_i/\lambda).$$

Similarly, Gauss-Hermite quadrature provides weights ω_i and abcissæ ζ_i for integrals of the form:

$$\int_{-\infty}^\infty e^{-\zeta^2} \varphi(\zeta)\, d\zeta \approx \sum_{i=1}^n \omega_i \varphi(\zeta_i).$$

If X is a normally distributed random variable with mean μ and variance σ^2, then we can approximate the expectation of $f(X)$ using quadrature by applying the transformation $\zeta = (x-\mu)/(\sqrt{2}\sigma)$, $\varphi(\zeta) = f(\mu + \sqrt{2}\sigma\zeta)/\sqrt{\pi}$, and thus, $d\zeta = dx/\sqrt{2\sigma^2}$. Then,

$$\mathrm{E}[f(X)] = \int_{-\infty}^\infty \frac{1}{\sqrt{2\pi\sigma^2}} \mathrm{e}^{-\frac{(x-\mu)^2}{2\sigma^2}} f(x)\, dx \approx \sum_{i=1}^n \frac{\omega_i}{\sqrt{\pi}} f(\mu + \sqrt{2}\sigma\zeta_i).$$

A.3. Optimization

Simulated Annealing was used for optimization for the maximum likelihood estimates and for the first step sieve maximum likelihood policy estimation. Initial step sizes and temperatures were chosen separately for each estimator, depending on the magnitude of the objective functions and the difficulty in finding the optimum. Some reported runtimes are long due to carrying out a thorough search requiring many thousands of functional evaluations. For the second step minimum distance objective function, we use the Levenberg-Marquardt algorithm for nonlinear least squares.

A.4. Computational Times

Estimation using the actual data and for Monte Carlo experiments was carried out on servers equipped with dual Intel Xeon X5670 or dual X5690 processors, each with six physical cores. Therefore, each server has a total of 12 physical processor cores. The particle filtering routine was carried out in parallel at the "market" level: the log likelihoods for individual buses were calculated on separate processor cores.

B. Proof of Theorem 1

Consider the following decomposition:

$$\begin{aligned}\hat{L}(\theta) - L(\theta) - \hat{L}(\theta_0) + L(\theta_0) &= P_N \ln \hat{h}(\cdot,\theta) - P \ln h(\cdot,\theta) - P_N \ln \hat{h}(\cdot,\theta_0) + P \ln h(\cdot,\theta_0) \\ &= \underbrace{(P_N - P)[\ln h(\cdot,\theta) - \ln h(\cdot,\theta_0)]}_{\equiv A} + \underbrace{P_N[\ln \hat{h}(\cdot,\theta) - \ln \hat{h}(\cdot,\theta_0) - \ln h(\cdot,\theta) + \ln h(\cdot,\theta_0)]}_{\equiv B}\end{aligned}$$

We show below that both terms A and B converge uniformly in probability to 0. Uniform convergence of $\hat{L}_{N,R}(\theta) - \hat{L}_{N,R}(\theta_0)$ to $L(\theta) - L(\theta_0)$ follows from the triangle inequality. Consistency of the maximum filtered likelihood estimator $\hat{\theta}_{N,R}$ follows from Theorem 2.1 of Newey and McFadden (1994).

By Assumption C1 and Lemmas 2.14(i) and 2.15 of Pakes and Pollard (1989) the class of functions $\overline{\mathscr{H}} = \{\ln h(\cdot,\theta) - \ln h(\cdot,\theta_0) : \theta \in \Theta\}$ is Euclidean for the envelope $\overline{H} = \ln H - \ln h(\cdot,\theta_0)$ which is integrable under Assumption C2. Therefore, by Lemma 2.8 of Pakes and Pollard (1989) $\sup_\Theta |A| \xrightarrow{\text{as}} 0$.

By Taylor's theorem,

$$\sup_\Theta |B| \leq 2 \sup_{y,\theta} \left|\ln \hat{h}(y,\theta) - \ln h(y,\theta)\right| = 2 \sup_{y,\theta} \left| \frac{\hat{h}(y,\theta) - h(y,\theta)}{h^*} \right|$$

for some h^* between $\hat{h}(y,\theta)$ and $h(y,\theta)$. By Assumption C2,

$$2 \sup_{y,\theta} \left| \frac{\hat{h}(y,\theta) - h(y,\theta)}{h^*} \right| \leq 2M \sup_{y,\theta} \left|\hat{h}(y,\theta) - h(y,\theta)\right|.$$

Recall the definitions of $\hat{h}$ and h and define h_t and $\hat{h}_t$ as follows:

$$h(y_{1:T},\theta) = \prod_{t=1}^{T} S^{y_{1:t}} q(\cdot, y_{t-1}, y_t, \theta) \equiv \prod_{t=1}^{t} h_t(y_{1:T},\theta),$$
$$\hat{h}(y_{i,1:T},\theta) = \prod_{t=1}^{T} S_R^{y_{i,1:t}} q(\cdot, y_{i,t-1}, y_{i,t}, \theta) \equiv \prod_{t=1}^{t} \hat{h}_t(y_{1:T},\theta).$$

Since the functions in $\mathscr{Q}$ are bounded, uniform convergence of each $\hat{h}_t$ to h_t implies uniform convergence of the product, $\hat{h}$, to the product of the uniform limits, h.

Since the class $\mathscr{Q}$ is Euclidean, there exist positive constants C and V which do not depend on P such that the $L^1(P)$ covering number satisfies $N(\varepsilon, \mathscr{Q}, L^1(P)) \leq C\varepsilon^{-V}$ for all $\varepsilon \in (0,1]$. Since Q is a square-integrable envelope, it follows that the $L^2(P)$ covering number is bounded as $N(\varepsilon, \mathscr{Q}, L^2(P)) \leq C2^{2V}\varepsilon^{-2V}$ for all $\varepsilon \in (0,1]$ (Nolan and Pollard, 1987, p. 789). Hence the uniform L^2 covering number is bounded: $\mathscr{N}(\varepsilon, \mathscr{Q}) \equiv \sup_\mu N(\varepsilon, \mathscr{Q}, L^2(\mu)) < \infty$. We can now apply Theorem 5 of Del Moral and Ledoux (2000) to establish that, for each $t = 1, \ldots, T$, $\sup_{y,\theta} \left|\hat{h}_t(y,\theta) - h_t(y,\theta)\right| \xrightarrow{\text{as}} 0$ as $R \to \infty$.

References

Ackerberg D, Benkard L, Berry S, Pakes A. 2007. Econometric tools for analyzing market outcomes. In Heckman JJ, Leamer EE (eds.) *Handbook of Econometrics*, volume 6A. North Holland.

Aguirregabiria V, Mira P. 2002. Swapping the nested fixed point algorithm: A class of estimators for discrete Markov decision models. *Econometrica* **70**: 1519–1543.

Aguirregabiria V, Mira P. 2007. Sequential estimation of dynamic discrete games. *Econometrica* **75**: 1–53.

Aguirregabiria V, Mira P. 2010. Dynamic discrete choice structural models: a survey. *Journal of Econometrics* **156**: 38–67.

Arcidiacono P, Miller RA. 2011. Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity. *Econometrica* **79**: 1823–1867.

Bajari P, Benkard CL, Levin J. 2007. Estimating dynamic models of imperfect competition. *Econometrica* **75**: 1331–1370.

Blevins JR. 2014. Nonparametric identification of dynamic decision processes with discrete and continuous choices. *Quantitative Economics* **5**: 531–554.

Blevins JR, Khwaja A, Yang N. 2014. Firm expansion, size spillovers and market dominance in retail chain dynamics. Working paper, Yale University.

Cappé O, Moulines E, Ryden T. 2005. *Inference in Hidden Markov Models*. New York: Springer.

Chen X. 2007. Large sample sieve estimation of semi-nonparametric models. In Heckman JJ, Leamer EE (eds.) *Handbook of Econometrics*, volume 6B, chapter 76. Amsterdam: North Holland, 5549–5632.

Creal D. 2012. A survey of sequential Monte Carlo methods for economics and finance. *Econometric Reviews* **31**: 245–296.

Del Moral P, Ledoux M. 2000. Convergence of empirical processes for interacting particle systems with applications to nonlinear filtering. *Journal of Theoretical Probability* **13**.

Doucet A, de Freitas N, Gordon N (eds.) . 2001. *Sequential Monte Carlo Methods in Practice*. New York: Springer.

Eckstein Z, Wolpin KI. 1989. The specification and estimation of dynamic stochastic discrete choice models: A survey. *Journal of Human Resources* **24**: 562–598.

Fang H, Kung E. 2010. Why do life insurance policyholders lapse? The roles of income, health and bequest motive shocks. Working paper, University of Pennsylvania.

Fernández-Villaverde J, Rubio-Ramírez JF. 2007. Estimating macroeconomic models: A likelihood approach. *Review of Economic Studies* **74**: 1059–1087.

Flury T, Shephard N. 2011. Bayesian inference based only on simulated likelihood: Particle filter analysis of dynamic economic models. *Econometric Theory* **27**: 933–956.

Gallant AR, Hong H, Khwaja A. 2010. Dynamic entry with cross product spillovers: An application to the generic drug industry. ERID Working Paper 57, Duke University.

Gordon N, Salmond D, Smith A. 1993. Novel approach to nonlinear/non-gaussian Bayesian state estimation. *IEEE Proc. F, Radar and signal processing* **140**: 107–113.

Hajivassiliou VA, Ruud PA. 1994. Classical estimation methods for LDV models using simulation. In Engle RF, McFadden DL (eds.) *Handbook of Econometrics*, volume 4, chapter 40. Amsterdam: North Holland, 937–976.

Heckman JJ. 1979. Sample selection bias as a specification error. *Econometrica* **47**: 153–161.

Heckman JJ. 1981. The incidental parameters problem and the problem of initial conditions in estimating a discrete timeâĂŞdiscrete data stochastic process. In Manski CF, McFadden DL (eds.) *Structural Analysis of Discrete Data with Econometric Applications*. MIT Press, 179–195.

Hotz VJ, Miller RA. 1993. Conditional choice probabilities and the estimation of dynamic models. *Review of Economic Studies* **60**: 497–529.

Hu Y, Shum M. 2012. Nonparametric identification of dynamic models with unobserved state variables. *Journal of Econometrics* **171**: 32–44.

Hu Y, Shum M. 2013. Identifying dynamic games with serially-correlated unobservables. In *Advances in Econometrics*, volume 31. 97–113.

Imai S, Jain N, Ching A. 2009. Bayesian estimation of dynamic discrete choice models. *Econometrica* **77**: 1865–1899.

Judd KL. 1998. *Numerical Methods in Economics*. Cambridge, MA: MIT Press.

Kalman RE. 1960. A new approach to linear filtering and prediction problems. *Transactions of the ASME–Journal of Basic Engineering* **82**: 35–45.

Keane M, Wolpin KI. 1994. The solution and estimation of discrete choice dynamic programming models by simulation and interpolation: Monte carlo evidence. *Review of Economics and Statistics* **76**: 648–672.

Künsch HR. 2001. State space and hidden Markov models. In Barndorff-Nielsen OE, Cox DR, Klüppelberg C (eds.) *Complex Stochastic Systems*. Chapman and Hall, 109–173.

Liu J. 2001. *Monte Carlo Strategies in Scientific Computing*. New York: Springer.

Liu JS, Chen R. 1998. Sequential Monte Carlo methods for dynamic systems. *Journal of the American Statistical Association* **93**: 1032–1044.

Newey WK, McFadden D. 1994. Large sample estimation and hypothesis testing. In Engle RF, McFadden D (eds.) *Handbook of Econometrics*, volume 4. Amsterdam: North Holland.

Nolan D, Pollard D. 1987. U-Processes: Rates of convergence. *Annals of Statistics* **15**: 780–799.

Norets A. 2009a. Implementation of Bayesian inference in dynamic discrete choice models. Working paper, Princeton University.

Norets A. 2009b. Inference in dynamic discrete choice models with serially correlated unobserved state variables. *Econometrica* **77**: 1665–1682.

Norets A. 2010. Continuity and differentiability of expected value functions in dynamic discrete choice models. *Quantitative Economics* **1**: 305–322.

Norets A. 2012. Estimation of dynamic discrete choice models using artificial neural network approximations. *Econometric Reviews* **31**: 84–106.

Pakes A. 1986. Patents as options: Some estimates of the value of holding European patent stocks. *Econometrica* **54**: 755–784.

Pakes A. 1994. Estimation of dynamic structural models: Problems and prospects part II: Mixed continuous-discrete models and market interactions. In Laffont J, Sims C (eds.) *Advances in Econometrics: Proceedings of the 6th World Congress of the Econometric Society*.

Pakes A, Pollard D. 1989. Simulation and the asymptotics of optimization estimators. *Econometrica* **57**: 1027–1057.

Pitt MK, Shephard N. 1999. Filtering via simulation: Auxiliary particle filters. *Journal of the American Statistical Association* **94**: 590–591.

Rust J. 1987. Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher. *Econometrica* **55**: 999–1013.

Rust J. 1994. Structural estimation of Markov decision processes. In Engle RF, McFadden DL (eds.) *Handbook of Econometrics*, volume 4. Amsterdam: North Holland, 3081–3143.

Ryan S. 2012. The costs of environmental regulation in a concentrated industry. *Econometrica* **80**: 1019–1061.

Srisuma S. 2010. Estimation of structural optimization models: A note on identification. Working paper, London School of Economics.

Stinebrickner TR. 2000. Serially correlated variables in dynamic, discrete choice models. *Journal of Applied Econometrics* **15**: 595–624.

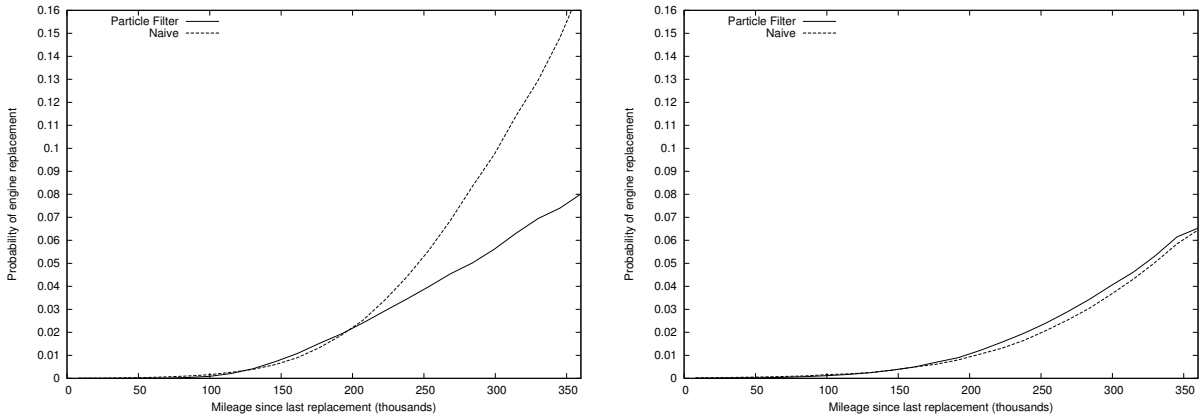

(a) Groups 1–3 (b) Group 4

FIGURE 1. Estimated Annual Replacement Hazard Functions

Su CL, Judd KL. 2012. Constrained optimization approaches to estimation of structural models. *Econometrica* **80**: 2213–2230.

Timmins C. 2002. Measuring the dynamic efficiency costs of regulators' preferences: Municipal water utilities in the arid West. *Econometrica* **70**: 603–629.

van der Vaart AW, Wellner JA. 1996. *Weak Convergence and Empirical Processes.* New York: Springer-Verlag.

Whiteley N. 2012. Sequential Monte Carlo samplers: error bounds and insensitivity to initial conditions. *Stochastic Analysis and Applications* **30**: 774–798.

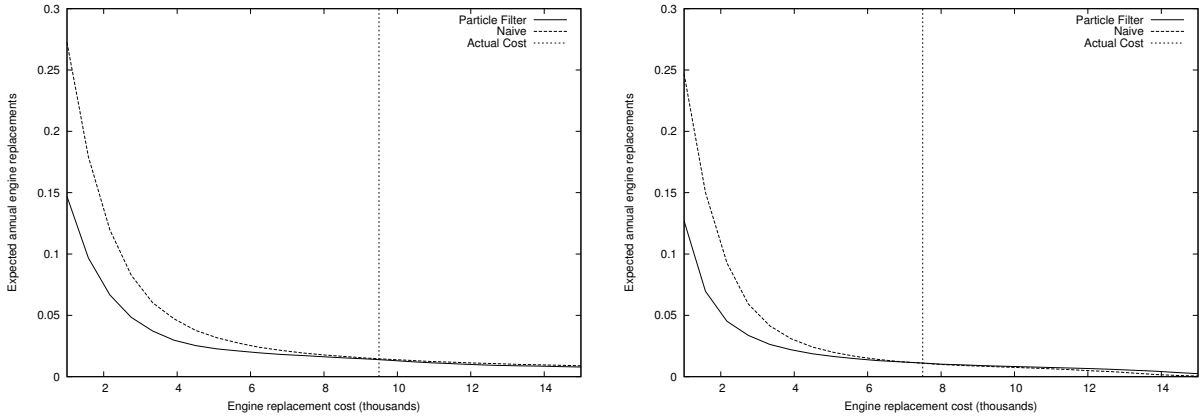

(a) Groups 1–3 (b) Group 4

FIGURE 2. Expected Annual Replacement Demand Functions

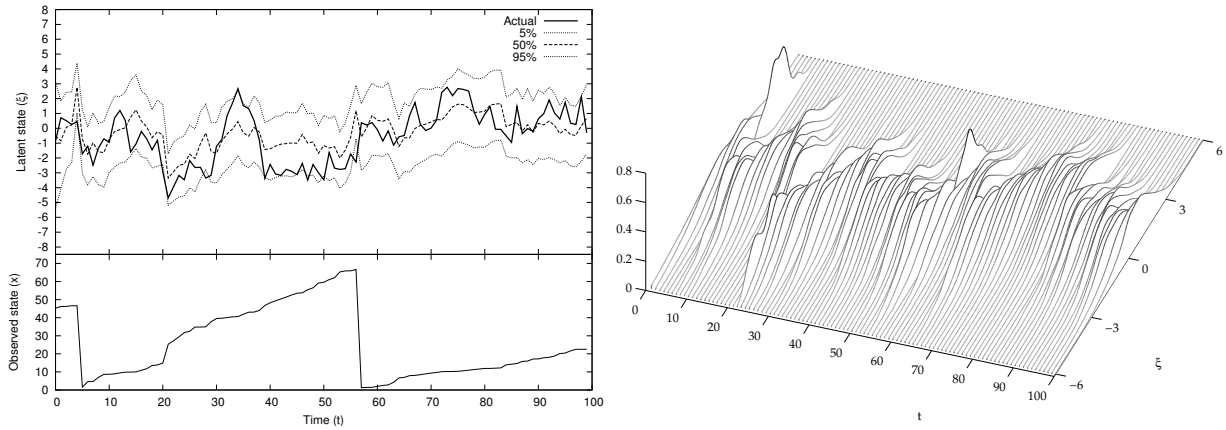

(a) Continuous State Model with Approximate Posterior (b) Evolution of Approximated Posterior Distributions

FIGURE 3. Simulated Sample Path with Posterior Distribution of the Latent State

TABLE 1. Summary of Bus Groups from Rust (1987)

	Sample Size			Mileage		Replacements	
Bus Group	Number of Buses	Months Observed	Total Obs.	Mean	S.D.	Count	Average Cost*
1	15	25	375	4196.29	1148.47	0	–
2	4	49	196	3106.03	937.74	0	–
3	48	70	3360	3708.45	1235.07	27	–
1–3	67	25–70	3931	3552.63	1432.76	27	$9,499
4	37	117	4329	3192.45	2306.53	33	$7,513

*Note: Replacement costs are only available for groups 1–3 combined.

TABLE 2. Within-Group Heterogeneity Tests

	Group 1	Group 2	Group 3	Group 4
Buses	15	4	48	37
Bus-Month Observations	356	192	3138	4229
Restricted Log Likelihood	-864.05	-409.60	-7250.71	-9137.98
Unrestricted Log Likelihood	-866.58	-409.37	-7248.52	-9118.96
LR Statistic	5.04	0.46	4.37	38.05
Degrees of Freedom	14	3	47	36
Marginal Significance Level	0.99	0.93	1.00	0.37

TABLE 3. Estimates for Bus Groups 1–3 and Group 4 of Rust (1987)

Estimator	LL	λ_0	λ_ξ	ρ	c_0/c_x	c_ξ/c_x	Minutes
Groups 1-3							
Naïve Estimator	-2662.88	0.343	–	–	10.363	–	2
		(0.012)	–	–	(0.905)	–	
Particle Filter ($R = 5,000$)	-2660.26	0.341	0.007	0.684	12.025	0.535	76
		(0.013)	(0.005)	(0.093)	(1.834)	(0.252)	
Observations	3931						
LR Statistic (df = 3)	5.242						
Marginal Signficance Level	0.155						
Group 4							
Naïve Estimator	-2393.53	0.477	–	–	15.389	–	2
		(0.013)	–	–	(1.413)	–	
Particle Filter ($R = 5,000$)	-2338.98	0.439	-0.048	0.982	11.464	0.029	107
		(0.023)	(0.038)	(0.168)	(2.657)	(0.573)	
Observations	4329						
LR Statistic (df = 3)	109.107						
Marginal Signficance Level	0.000						

TABLE 4. Maximum Filtered Likelihood Estimates: Monte Carlo Results

		Parameters					Time Per Replication		
Estimator		λ_0	λ_ξ	ρ	c_0/c_x	c_ξ/c_x	Total (m)	Per Eval. (s)	Eval.
Naïve	Mean	0.233	–	–	8.485	–	1	0.046	1849
	S.D.	0.018	–	–	0.326	–			
	Bias	-0.167	–	–	1.485	–			
	RMSE	0.168	–	–	1.521	–			
MFL ($R = 25$)	Mean	0.381	0.289	0.810	7.315	0.231	7	0.052	8353
	S.D.	0.038	0.035	0.036	0.527	0.051			
	Bias	-0.019	-0.011	0.010	0.315	-0.019			
	RMSE	0.043	0.037	0.038	0.614	0.055			
MFL ($R = 50$)	Mean	0.391	0.299	0.802	6.990	0.244	9	0.070	7585
	S.D.	0.031	0.030	0.028	0.405	0.046			
	Bias	-0.009	-0.001	0.002	0.010	-0.006			
	RMSE	0.033	0.030	0.029	0.406	0.046			
MFL ($R = 100$)	Mean	0.389	0.294	0.801	6.824	0.248	11	0.099	6985
	S.D.	0.021	0.023	0.026	0.534	0.035			
	Bias	-0.011	-0.006	0.001	0.176	-0.002			
	RMSE	0.024	0.024	0.026	0.562	0.035			
MFL ($R = 200$)	Mean	0.393	0.306	0.790	6.980	0.262	18	0.157	6937
	S.D.	0.017	0.027	0.031	0.461	0.045			
	Bias	-0.007	0.006	-0.010	0.020	0.012			
	RMSE	0.019	0.028	0.032	0.461	0.046			
MFL ($R = 400$)	Mean	0.390	0.299	0.800	6.987	0.251	28	0.271	6265
	S.D.	0.021	0.020	0.015	0.410	0.020			
	Bias	-0.010	-0.001	-0.000	0.013	0.001			
	RMSE	0.023	0.020	0.015	0.410	0.020			
Observed ξ_t	Mean	0.400	0.299	0.801	6.924	0.249	5	0.050	5977
	S.D.	0.011	0.007	0.006	0.220	0.011			
	Bias	0.000	-0.001	0.001	-0.076	-0.001			
	RMSE	0.011	0.008	0.006	0.233	0.011			

Note: Reported are the mean, standard deviation, mean bias, and root mean squared error over 25 replications. Computational times are average minutes per replication, average seconds per functional evaluation, and average functional evaluations.

TABLE 5. Sources of Replacement Cost Ratio Bias

Specification	Parameters	λ_0	λ_ξ	ρ	c_0/c_x	c_ξ/c_x
	Population	0.400	0.300	0.800	7.000	0.250
1. Baseline	Naïve Estimation	0.233	–	–	8.485	–
		(0.018)	–	–	(0.326)	–
	Population	0.400	0.000	0.800	7.000	0.250
2. $\lambda_\xi = 0$	Naïve Estimation	0.400	–	–	8.163	–
		(0.011)	–	–	(0.334)	–
	Population	0.400	-0.300	0.800	7.000	0.250
3. $\lambda_\xi < 0$	Naïve Estimation	0.341	–	–	8.062	–
		(0.013)	–	–	(0.302)	–
	Population	0.400	0.300	0.800	7.000	0.000
4. $c_\xi = 0$	Naïve Estimation	0.283	–	–	6.963	–
		(0.018)	–	–	(0.188)	–

TABLE 6. Two-Step Estimator: First Step Monte Carlo Results

	Structural Parameters ($\hat{\theta}_1$)			Polynomial Coefficients ($\hat{\alpha}$)										
First-Step Estimator	λ_0	λ_ξ	ρ	1	x	ξ	$x\xi$	x^2	ξ^2	$x^2\xi$	$x\xi^2$	x^3	ξ^3	Minutes
Population	0.400	0.300	0.800	–	–	–	–	–	–	–	–	–	–	–
$\kappa = 3, N = 50$												–	–	
MFL ($R = 100$)	0.393	0.306	0.783	8.043	-10.303	-1.323	–	–	–	–	–	–	–	10
	(0.026)	(0.028)	(0.032)	(0.968)	(1.609)	(0.320)	–	–	–	–	–	–	–	
Observed ξ_t	0.401	0.298	0.801	8.197	-10.740	-1.342	–	–	–	–	–	–	–	1
	(0.013)	(0.012)	(0.008)	(0.417)	(0.955)	(0.106)	–	–	–	–	–	–	–	
$\kappa = 4, N = 100$												–	–	
MFL ($R = 100$)	0.393	0.309	0.787	7.881	-10.073	-1.174	-0.347	–	–	–	–	–	–	18
	(0.021)	(0.024)	(0.028)	(0.785)	(1.088)	(0.293)	(0.053)	–	–	–	–	–	–	
Observed ξ_t	0.400	0.299	0.801	8.011	-10.383	-1.225	-0.374	–	–	–	–	–	–	1
	(0.014)	(0.007)	(0.008)	(0.282)	(0.613)	(0.065)	(0.055)	–	–	–	–	–	–	
$\kappa = 6, N = 400$												–	–	
MFL ($R = 100$)	0.386	0.300	0.796	10.885	-21.596	-1.719	-0.001	10.253	0.022	–	–	–	–	38
	(0.014)	(0.013)	(0.015)	(1.853)	(4.402)	(0.454)	(0.023)	(2.283)	(0.014)	–	–	–	–	
Observed ξ_t	0.399	0.299	0.800	9.534	-18.344	-1.364	0.015	8.567	-0.003	–	–	–	–	1
	(0.008)	(0.004)	(0.005)	(0.243)	(0.864)	(0.056)	(0.034)	(0.767)	(0.013)	–	–	–	–	
$\kappa = 8, N = 1600$												–	–	
MFL ($R = 100$)	0.393	0.299	0.800	10.946	-21.489	-1.623	-0.689	10.073	0.023	0.243	-0.047	–	–	163
	(0.014)	(0.011)	(0.012)	(2.249)	(4.322)	(0.526)	(1.196)	(2.512)	(0.014)	(0.408)	(0.039)	–	–	
Observed ξ_t	0.400	0.300	0.800	9.785	-18.778	-1.560	0.363	8.618	0.023	0.012	-0.036	–	–	3
	(0.007)	(0.003)	(0.003)	(0.426)	(1.533)	(0.153)	(0.465)	(1.259)	(0.008)	(0.275)	(0.009)	–	–	
$\kappa = 10, N = 6400$												–	–	
MFL ($R = 100$)	0.393	0.307	0.786	15.839	-25.788	-8.904	8.256	2.464	2.657	-2.671	-1.775	6.472	-0.271	912
	(0.011)	(0.013)	(0.017)	(6.702)	(10.735)	(5.268)	(6.624)	(20.870)	(1.190)	(4.266)	(1.064)	(13.346)	(0.137)	
Observed ξ_t	0.400	0.299	0.799	12.939	-32.347	-4.745	8.375	25.728	1.010	-4.135	-1.336	-5.949	-0.087	15
	(0.005)	(0.004)	(0.004)	(1.715)	(6.605)	(2.053)	(5.556)	(9.791)	(0.635)	(3.487)	(0.884)	(6.248)	(0.056)	

We let the sample size increase with κ as $N = 50 \times 2^{\kappa-3}$.

TABLE 7. Two-Step Estimator: Second Step Monte Carlo Results

		Structural Parameters ($\hat{\theta}_2$)		
Sample Size	First-Step Estimator	c_0/c_x	c_ξ/c_x	Minutes*
Population		7.000	0.250	–
$\kappa = 3, N = 50$	MFL ($R = 100$)	6.578	0.250	1
		(0.421)	(0.048)	
	Observed ξ_t	6.461	0.234	1
		(0.311)	(0.021)	
$\kappa = 4, N = 100$	MFL ($R = 100$)	6.725	0.259	1
		(0.309)	(0.050)	
	Observed ξ_t	6.653	0.256	1
		(0.200)	(0.014)	
$\kappa = 6, N = 400$	MFL ($R = 100$)	7.528	0.319	1
		(0.274)	(0.029)	
	Observed ξ_t	7.458	0.308	1
		(0.255)	(0.012)	
$\kappa = 8, N = 1600$	MFL ($R = 100$)	7.737	0.316	1
		(0.416)	(0.031)	
	Observed ξ_t	7.151	0.288	1
		(0.234)	(0.016)	
$\kappa = 10, N = 6400$	MFL ($R = 100$)	7.293	0.305	1
		(0.457)	(0.053)	
	Observed ξ_t	6.699	0.249	1
		(0.247)	(0.015)	
$\kappa = \infty, N = \infty$	True Policy	6.976	0.260	1
		(0.041)	(0.003)	

*Computational times include forward simulation of the ex-ante value functions and minimization of Q to determine $\hat{\theta}_2$.