

# Sequential Monte Carlo Methods for Estimating Dynamic Microeconomic Models\*

JASON R. BLEVINS

*Department of Economics, Ohio State University*

*Working Paper 11-01*

May 2, 2011<sup>†</sup>

**Abstract.** This paper develops methods for estimating dynamic structural microeconomic models with serially correlated latent state variables. The proposed estimators are based on sequential Monte Carlo methods, or particle filters, and simultaneously estimate both the structural parameters and the trajectory of the unobserved state variables for each observational unit in the dataset. We focus two important special cases: single agent dynamic discrete choice models and dynamic games of incomplete information. The methods are applicable to both discrete and continuous state space models. We first develop a broad nonlinear state space framework which includes as special cases many dynamic structural models commonly used in applied microeconomics. Next, we discuss the nonlinear filtering problem that arises due to the presence of a latent state variable and show how it can be solved using sequential Monte Carlo methods. We then turn to estimation of the structural parameters and consider two approaches: an extension of the standard full-solution maximum likelihood procedure (Rust, 1987) and an extension of the two-step estimation method of Bajari, Benkard, and Levin (2007), in which the structural parameters are estimated using revealed preference conditions. Finally, we introduce an extension of the classic bus engine replacement model of Rust (1987) and use it both to carry out a series of Monte Carlo experiments and to provide empirical results using the original data.

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

**JEL Classification:** C13, C15.

---

\*I am grateful to Han Hong, Shakeeb Khan, Paul Ellickson, and Andrew Sweeting for their guidance and support during the course of this project and many others and to Jesús Fernández-Villaverde and Juan Rubio-Ramírez for an introduction to sequential Monte Carlo methods. I also thank seminar participants at Duke, Ohio State AED Economics, 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.

<sup>†</sup>First version: May 21, 2008. All figures and tables are reproducible using Fortran source code available at <http://jblevins.org/research/smcdmm>.

# 1. Introduction

This paper proposes methods for estimating a class of dynamic microeconomic models with serially correlated latent state variables. The estimators proposed are straightforward 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.

Accounting for unobserved heterogeneity is important in applied work in many fields of economics and latent state variables in a Markov state space model are a very flexible way to model unobserved heterogeneity. Serially-correlated latent state variables 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 include unobserved market-level shocks, firm-specific productivity shocks, or multiple equilibria in industrial organization, unobserved ability or wage offers in labor economics, and latent health status in health economics.

To understand the problem raised by latent state variables, consider the likelihood of an observed sequence  $y_{1:T} = (y_1, \dots, y_T)$ , which is thought to depend on an unobserved sequence of latent state variables  $z_{1:T} = (z_1, \dots, z_T)$  and a vector of parameters  $\theta$ . The likelihood can be written as

$$p(y_{1:T} | \theta) = \int p(y_{1:T}, z_{1:T} | \theta) dz_{1:T} = \int p(y_{1:T} | z_{1:T}, \theta) p(z_{1:T} | \theta) dz_{1:T},$$

but this is potentially a very high-dimensional integral. Even if it is feasible, simulation of paths  $z_{1:T}$  is not a very efficient way to evaluate this integral. Another way to write this integral is

$$\begin{aligned} p(y_{1:T} | \theta) &= \prod_{t=1}^T p(y_t | y_{1:t-1}, \theta) \\ &= \prod_{t=1}^T \int p(y_t | z_t, y_{1:t-1}, \theta) p(z_t | y_{1:t-1}, \theta) dz_t. \end{aligned}$$

Now, if we could only draw from the posterior  $p(z_t | y_{1:t-1}, \theta)$ , we could approximate the overall integral by a product of one-dimensional integrals.

Both of the estimators we propose involve maximum likelihood estimation of parameters, where the likelihood function must be integrated with respect to the posterior distribution of the latent state given the observed data as described above. We show that sequential Monte Carlo (SMC) methods can be used to approximate this posterior distribution, which can be used

to integrate the likelihood function, allowing the parameters of the model to be estimated by maximizing the approximated log-likelihood function.

We first motivate the framework and estimators in detail. Then, in an effort to illustrate the use of SMC methods and to evaluate the performance of the proposed estimators, we perform a series of Monte Carlo experiments using a generalized version of the classic bus engine replacement model of [Rust \(1987\)](#). We also estimate the same model using [Rust's](#) original dataset to highlight the importance of unobserved heterogeneity in this setting.

There are now many methods available for estimating the structural parameters of complex dynamic models, including single agent models and dynamic games. Traditionally these methods require any latent state variables to be serially independent conditional on observed state variables. This precludes many interesting considerations, even agent-specific fixed effects, and rules out dynamic selection on unobservables. Thus, allowing for serially correlated latent state variables is an important, and natural, direction in which to extend existing methods. Several methods for doing so have been introduced recently and are discussed in more detail below. These methods are either geared towards Bayesian analysis or are applicable only to models with discrete state spaces. This paper proposes new methods for classical estimation which are applicable to models with state-spaces and choice sets that are discrete, continuous, or a combination of both.

Serial correlation has been recognized as important topic in the dynamic discrete choice literature from the beginning. [Eckstein and Wolpin \(1989\)](#) survey the early work in this area and emphasize the need to deal with special correlation structures among the unobservable states, such as correlation across time or across choices. [Rust \(1994\)](#) provides another survey of the literature as well as a general econometric framework for the estimation of such models when the unobserved states satisfy a conditional independence assumption. This framework has been widely used, but unfortunately the conditional independence assumption rules out serial correlation in the unobserved state variables. A more recent survey on dynamic discrete choice models by [Aguirregabiria and Mira \(2010\)](#) covers both single agent models and dynamic discrete games and also emphasizes the need to account for potentially serial correlated unobservables.

Until recently, methods for estimating models with non-iid correlation structures were either application-specific or were not applicable to general distributions of the unobserved states. [Pakes \(1994\)](#), and later [Ackerberg, Benkard, Berry, and Pakes \(2007, Section 3.8.1\)](#), outline three methods for estimating models with serial correlation. First, given values of the parameters and an exogenous initial condition for the unobserved state, the model can be simulated in order to obtain some quantities, such as choice probabilities, that can be matched with the data for estimation. [Pakes \(1986\)](#) estimates a model of patent renewal behavior using this approach, which amounts to high-dimensional Monte Carlo integration since the entire path

of the unobservable must be simulated many times. This approach is inefficient because it does not bring information from the data to bear on the simulated paths of the unobservables. Second, models with continuous control variables may satisfy an invertibility condition where the unobserved state variable can be written as a function of parameters and observables. For example, [Timmins \(2002\)](#) estimates a model of the water pricing behavior of a municipal administrator using a combination of the inversion and simulation approaches. Finally, in some situations, particularly in industry studies, it may be reasonable to assume that the process has a long history and that during the observation period, states are drawn from the ergodic distribution. An example of this approach is [Aguirregabiria and Mira \(2007\)](#), who estimate dynamic games while allowing for a time-invariant market-specific fixed effect. They use the implied invariant distribution of the state and fixed effect to find the posterior distribution of the fixed effect.

Identification and estimation of dynamic models with serially correlated unobserved state variables has been a topic of recent interest in the literature. Nonparametric identification of dynamic single-agent models with serially correlated unobserved state variables has been studied by [Hu and Shum \(2010\)](#), who provide function-analytic conditions for nonparametric identification of the Markov transition kernel in such models. [Hu and Shum \(2008\)](#) extend these results to the case of dynamic games. Bayesian methods for estimating single agent dynamic discrete choice models with finite state spaces have been developed by [Imai, Jain, and Ching \(2009\)](#), who propose a Markov chain Monte Carlo (MCMC) algorithm which simultaneously solves the dynamic program and estimates the parameters. [Norets \(2009\)](#) extends this idea, developing Gibbs sampling methods which allow for serially correlated unobserved state variables and compact state spaces. In terms of classical estimation, [Arcidiacono and Miller \(2010\)](#) 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 \(1993\)](#).

This paper is also clearly related to sequential Monte Carlo methods more broadly. These methods originated in the engineering literature and their theoretical properties are being developed in the statistics literature (for a survey, see [Doucet, de Freitas, and Gordon \(2001\)](#) and the citations therein). Sequential Monte Carlo methods have many uses, but in this paper they are used to approximate likelihood functions and to approximate the posterior distributions of latent state variables. The idea of using filter samples to perform integration over latent states in likelihood functions dates back at least to [Kitagawa \(1996\)](#), who estimates the parameters of a simple model in which an observation process depends only on the contemporaneous value of an unobserved Markovian state variable. Recently in the economics literature, this approach has been used used to estimate structural macroeconomic models. [Fernández-Villaverde and](#)

Rubio-Ramírez (2007) use particle filters to estimate macroeconomic models in the presence of stochastic volatility. Particle filters can also be used for Bayesian estimation of parameters and latent state trajectories. Gallant, Hong, and Khwaja (2011) use this approach to estimate a complete-information dynamic discrete model of entry and recover unobserved firm-level costs in the generic pharmaceutical industry.

This remainder of this paper is organized as follows. Section 2 introduces a general nonlinear state space model, which is shown to contain many dynamic microeconomic models. Attention is paid in particular to dynamic discrete choice models and dynamic games of incomplete information. Section 3 motivates the optimal filtering problem that arises in these models due to the latent state variables and gives an overview of sequential Monte Carlo methods, which are used to approximate the posterior distribution of the latent state. In Section 4, we turn our attention to estimation and propose two estimators for dynamic microeconomic models: a full-solution maximum likelihood estimator and a two-step estimator based on revealed preference or equilibrium conditions. We study these estimators in a series of Monte Carlo applications and in an empirical application to the bus engine replacement model of Rust (1987). Section 6 concludes.

## 2. Nonlinear State Space Models

Many dynamic models used in economics and other fields can be restated as simple nonlinear state space models. However, in many cases, some of the states may be unobserved by the researcher. For example, in many signal processing applications, only a noisy observation of an underlying, unobserved signal is available to the researcher. There, the signal is modeled as a hidden state and the noise is a potentially serially-correlated, non-Gaussian, and nonlinear process which contaminates the observed process.

Microeconomic models tend to be more complex due to the fact that they are controlled stochastic processes—they involve a rational decision maker who controls the joint stochastic process in some way. As we will show in the following sections, these models are all members of a class of nonlinear state space models, and that standard sequential Monte Carlo method from the engineering and signal processing literature, with some generalizations, can also be applied to dynamic economic models.

First, we introduce a general state space model and refine it by introducing a few standard assumptions used in dynamic microeconomic models. We then consider two particular special cases: single agent dynamic discrete choice models and dynamic games of incomplete information. We show that both of these models have nonlinear state space representations.

## 2.1. A General State Space Model

Discrete time nonlinear state space models describe the process by which a vector  $s_t$  of state variables on some state space  $\mathcal{S} \subseteq \mathbb{R}^{n_s}$  determine a vector  $a_t$  of outcome or choice variables on some outcome space  $\mathcal{A} \subseteq \mathbb{R}^{n_a}$ . For any sequence  $\{x_t\}$ , let  $x_{j:k}$  denote the vector  $(x_j, \dots, x_k)$ . We begin with a few basic assumptions.<sup>1</sup>

**Assumption 1** (First order Markov process). The joint stochastic process  $\{s_t, a_t\}$  is a first-order Markov process with parameter  $\theta$ :

$$p(a_t, s_t | a_{1:t-1}, s_{1:t-1}, \theta) = p(a_t, s_t | a_{t-1}, s_{t-1}, \theta).$$

This assumption is fairly general in the sense that any  $n$ -th order Markov process can be expressed as a first-order process by appropriately redefining the state variables.

Note that we can always factor the joint density as a product of marginals:

$$p(a_t, s_t | a_{t-1}, s_{t-1}, \theta) = p(a_t | s_t, a_{t-1}, s_{t-1}, \theta) p(s_t | a_{t-1}, s_{t-1}, \theta).$$

However, in order to interpret the model as a state-space model, we require an additional assumption.

**Assumption 2** (State-space model). The contemporaneous state  $s_t$  is a sufficient statistic for the lagged variables  $(a_{t-1}, s_{t-1})$  in the conditional distribution of  $a_t$ :

$$p(a_t | s_t, a_{t-1}, s_{t-1}, \theta) = p(a_t | s_t, \theta).$$

We say that a model for the process  $\{a_t, s_t\}$  is a *nonlinear state space model* if it satisfies Assumptions 1 and 2. In particular, this class of models includes structural models of the form

$$\begin{aligned} a_t &= \sigma(s_t, \varepsilon_t | \theta), \\ s_t &= \kappa(s_{t-1}, a_{t-1}, \eta_t | \theta), \end{aligned}$$

where  $\varepsilon_t$  and  $\eta_t$  are independent random vectors. Depending on the context, the function  $\sigma$  is called the measurement or observation equation and  $\kappa$  is the transition equation. In economic models,  $\sigma$  is usually the policy function and  $\kappa$  is an exogenous state transition function.

---

<sup>1</sup>For simplicity we will use the notation  $x_t$  to denote both the random variable  $X_t$  as well as a particular realization  $X_t = x_t$ . Furthermore, we assume that the distributions of all continuous variables  $X_t$  admit densities with respect to Lebesgue measure. We treat discrete variables analogously with respect to the counting measure. Generically, the density of  $X_t$  will be denoted  $p(x_t)$ .

We are interested in models for which  $s_t$  is only partially observed by the researcher. That is, we can partition the state vector as  $s_t = (x_t, z_t)$  where the researcher can observe  $x_t$  but not  $z_t$ .<sup>2</sup> As before, we can always factor the joint density as

$$p(x_t, z_t | a_{t-1}, x_{t-1}, z_{t-1}, \theta) = p(x_t | z_t, a_{t-1}, x_{t-1}, z_{t-1}, \theta) p(z_t | a_{t-1}, x_{t-1}, z_{t-1}, \theta).$$

In order to apply sequential Monte Carlo methods to this model, we must restrict the nature of the dependence of the joint state with the following limited feedback assumption.

**Assumption 3** (Limited feedback). Conditional on  $z_t$ ,  $x_{t-1}$ , and  $a_{t-1}$ ,  $x_t$  is independent of  $z_{t-1}$ :

$$p(x_t | z_t, x_{t-1}, z_{t-1}, a_{t-1}) = p(x_t | z_t, x_{t-1}, a_{t-1}).$$

Note that this assumption still allows for complex patterns of feedback between the observed and unobserved states and the control variables.

*Remark.* This is where our framework, which is constructed with dynamic microeconomic models in mind, departs from the usual state space models to which SMC methods are applied. Typical models in the particle filtering literature do not have observed state variables  $x_t$ , only an observation process  $a_t$  and an unobserved signal process  $z_t$ . Furthermore,  $a_t$  is usually assumed to be continuous.

In microeconomic applications there are typically several other observable state variables, which can provide more information about the evolution of the latent state, but in many cases the measurement variable is discrete, which makes learning about the distribution of the latent state more difficult. Furthermore, the particle filtering literature usually assumes there is no feedback between the observable variable evolution of the latent state. In economic settings such feedback is usually present since the choice variables may indeed influence the evolution of the latent state.

Thus, the models we consider here are quite general in that they allow for dependence between  $x_t$  and  $z_t$ , but we must limit this dependence to cases under [Assumption 3](#). Although we make this assumption for a different reason, it is worth noting that [Assumption 3](#) is also used by [Hu and Shum \(2010\)](#) for identification purposes.

In light of [Assumption 3](#), if we let  $y_t = (a_t, x_t)$  denote the vector of all variables that are observable by the researcher, then a convenient alternative representation of the model is:

$$p(y_t | y_{t-1}, z_t, \theta) = p(a_t | x_t, z_t, \theta) p(x_t | z_t, a_{t-1}, x_{t-1}, \theta),$$

$$p(z_t | y_{t-1}, z_{t-1}, \theta)$$

---

<sup>2</sup>From the perspective of any decision makers in the model, both  $x_t$  and  $z_t$  are observed.

where  $p(y_t | y_{t-1}, z_t, \theta)$  is the *likelihood*, the conditional density of observables  $y_t = (a_t, x_t)$ , and  $p(z_t | y_{t-1}, z_{t-1}, \theta)$  is the *transition density*.

Our end goal is to perform likelihood-based inference on  $\theta$  without observations on  $z_t$ . To obtain a log-likelihood function in terms of only on the parameters and observables, we must somehow integrate the above likelihood with respect to the posterior distribution of the unobserved state. We return to this problem below, but first we show that many widely-used dynamic microeconomic can be reduced to the simple state-space form above. This will allow us to address issues of estimation in a unified framework without concern for the specific details of each model.

## 2.2. Dynamic Microeconomic Models

It turns out that most dynamic microeconomic models are simply nonlinear state space models of the same kind considered in the previous section, only with very complex likelihoods. Such models typically involve a state vector,  $s_t$ , that evolves according to some Markov process and a vector of choices,  $a_t$ , that arise as the result of the payoff maximizing behavior of the agents being modeled. In structural models, where one is concerned with inference on the parameters of a utility or profit function, the equilibrium conditions of the model imply predictions about the choices. Even in very simple models, this state-space representation can be highly nonlinear and non-Gaussian, preventing the use of methods such as Kalman filtering. For example, in single-agent dynamic discrete choice models the likelihood is the product of  $p(a_t | x_t, z_t, \theta)$ , the conditional choice probability of action  $a_t$  given the state, and  $p(x_t | x_{t-1}, z_t, a_{t-1}, \theta)$ , the transition density of the observed state.

We consider a general class of discrete-time dynamic multi-agent models with  $N$  players, indexed by  $i = 1, \dots, N$  over an infinite time horizon  $t = 1, 2, \dots, \infty$ . The state of the market at time  $t$  can be summarized by a state vector  $(s_t, v_t) \in \mathcal{S} \times \mathcal{N}$  where  $s_t \in \mathcal{S}$  is common knowledge to all players but  $v_t = (v_{1t}, \dots, v_{Nt}) \in \mathcal{N} \equiv \mathcal{N}_1 \times \dots \times \mathcal{N}_N$  is a vector of private shocks where  $v_{it} \in \mathcal{N}_i$  is private information to player  $i$ .

Each period, each player  $i$  observes the state and makes a choice  $a_{it}$  from the choice set  $\mathcal{A}_i$ . We define  $\mathcal{A} \equiv \mathcal{A}_1 \times \dots \times \mathcal{A}_N$  and let  $a_t = (a_{1t}, \dots, a_{Nt}) \in \mathcal{A}$  denote the vector of all actions at time  $t$ . Upon making the choices  $a_t$ , each player  $i$  receives a payoff  $U_i(a_t, s_t, v_{it})$  associated with making choice  $a_{it}$  in state  $s_t$ , given that player  $i$ 's rivals make choices  $a_{-it}$ , where in a slight abuse of notation we define  $a_{-it} = (a_{1t}, \dots, a_{i-1,t}, a_{i+1,t}, \dots, a_{Nt})$ . Although we focus on discrete choice models below for simplicity, in general the choice sets  $\mathcal{A}_i$  and the state space  $\mathcal{S}$  may be multidimensional and may have both continuous and discrete components.

Players are forward looking and discount future payoffs. We assume that players share a common discount factor  $\beta \in [0, 1)$ , which is known by the researcher. Player  $i$ 's discounted



expected future utility when the market is in state  $s_t$  is

$$\mathbb{E} \left[ \sum_{\tau=t}^{\infty} \beta^{\tau-t} U_i(a_{\tau}, s_{\tau}, v_{i\tau}) \mid s_t \right]$$

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

Before proceeding we make several standard assumptions, in addition to Assumptions 1, 2 and 3, to make the model more tractable (cf. Rust, 1994; Aguirregabiria and Mira, 2010).

**Assumption 4** (Conditional independence). The distribution of  $v_t$  is conditionally independent of  $v_{t-1}$  given  $s_t$  and can be factored as follows:

$$p(s_t, v_t \mid s_{t-1}, v_{t-1}, a_{t-1}) = p(s_t \mid s_{t-1}, a_{t-1}) p(v_t \mid s_t).$$

*Remark.* In contrast to the conditional independence assumption of Rust (1994), this is not a conditional independence assumption on any of the persistent state variables  $s_t$ , including any unobserved states, but only on the transient shocks  $v_t$ .

**Assumption 5** (Private Information). The private shocks  $v_{it}$  are independent across  $i$  and follow a known distribution  $G_i(\cdot \mid s_t)$  with support  $\mathcal{N}_i$ .

At this point we diverge from the typical framework and consider the presence of unobserved state variables. We let  $s_t$  be partially observed and write  $s_t = (x_t, z_t)$  where  $x_t \in \mathcal{X}$  is observed by the researcher, along with the choices  $a_t$ , but  $z_t \in \mathcal{Z}$  is an unobserved state. This unobserved state may be used, for example, to control for time-varying unobserved heterogeneity in a very general way. Both states  $x_t$  and  $z_t$  are common knowledge among the players. We shall be interested in estimating both the model primitives as well as the trajectory of the unobserved state  $z_t$ . Note that since both  $x_t$  and  $z_t$  may be multidimensional, this framework allows for both market- and firm-level time-varying unobserved heterogeneity.

Given the framework and assumptions above, a model of this class can be succinctly summarized by the densities  $p(z_t \mid x_{t-1}, z_{t-1}, a_{t-1})$ ,  $p(x_t \mid x_{t-1}, z_t, a_{t-1})$ , and  $p(a_t \mid x_t, z_t)$ . In the sections that follow, we consider two common special cases: dynamic games of incomplete information and single agent dynamic discrete choice models.

### 2.3. Single-Agent Dynamic Discrete Choice Models

Single agent dynamic discrete choice models are an important special case of the more general multi-agent model discussed above. Specifically, 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 player ( $N = 1$ ), we omit the  $i$  subscript from states and payoffs in this section. Each period the player makes a single choice  $a$  from a discrete choice set  $\mathcal{A} = \{0, 1, \dots, K\}$ . In each

period, associated with each choice is a choice specific shock  $\varepsilon_{ta}$ . The support of  $\varepsilon_t = (\varepsilon_{t0}, \dots, \varepsilon_{tK})$  is  $\mathcal{E} = \mathbb{R}^{K+1}$ .

In addition to the assumptions above we make the following standard additive separability assumption.

**Assumption 6** (Additive Separability). The utility function is additively separable in  $\varepsilon$ :

$$U(a, s, \varepsilon) = u(a, s) + \varepsilon_a.$$

The value function for this model can be expressed recursively as

$$V(s, \varepsilon) = \max_{a \in \mathcal{A}} \{u(a, s) + \varepsilon_a + \beta \mathbb{E}[V(s', \varepsilon') \mid s, a]\}.$$

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

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

Under [Assumption 6](#), 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 utility function. Letting  $\sigma(s, \varepsilon)$  denote the optimal policy, or choice of  $a$ , we have

$$\sigma(s, \varepsilon) = \operatorname{argmax}_{a \in \mathcal{A}} [v(a, s) + \varepsilon_a].$$

Under certain distributional assumptions for  $\varepsilon$ , the model admits conditional choice probabilities with known analytical forms. In particular, invoking the following assumption has become routine.

**Assumption 7** (Type I extreme value distribution). The components of  $\varepsilon$  are independent and identically distributed according to 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,

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

and  $v(a, s)$  can be found as the fixed point to the contraction mapping

$$(2) \quad \Gamma(v)(a, s) = u(a, s) + \beta \int \ln \left[ \sum_{j \in \mathcal{A}} \exp(v(j, s')) \right] P(ds' \mid s, a).$$

See Rust (1994) for details.

To summarize, Assumption 6 allows us to restate the problem in the familiar form of a static discrete choice problem, but this problem is still intractable for an arbitrary distribution  $G(d\varepsilon | s)$ . Assumption 7 leads to the convenient closed form above for  $P(a | s)$ , in terms of the choice specific value function  $v(a, s)$ . Importantly, it also allows us to obtain  $v(a, s)$  as the fixed point of the functional equation above. This provides a clear path for calculating it numerically using any number of methods such as value function iteration or projection methods (cf. Judd, 1998). Then, since we can evaluate  $v(a, s)$  for any choice  $a$  and state  $s$ , we can now evaluate  $P(a | s)$ , which is needed to evaluate the likelihood and brings us one step closer to being able to estimate the model.

This derivation also clearly illustrates the distinction between the predictions of the model and estimation, specifically with regard to the choice probabilities  $P(a | s)$ . In terms of computation it is irrelevant that  $s$  is only partially observed. The data is not involved at this stage since the model is defined conditional on both  $x$  and  $z$ , and can be solved as such. It is only in estimation where the distinction between the observed and unobserved states becomes important.

Hence, single agent dynamic discrete models of this form can be written as nonlinear state space models where the likelihood is the product of the conditional choice probabilities and the transition density of the observed states

$$p(y_t | y_{t-1}, z_t, \theta) = P(a_t | x_t, z_t, \theta) p(x_t | a_{t-1}, x_{t-1}, z_t, \theta)$$

Note that the transition densities of both  $x_t$  and  $z_t$  are typically specified as model primitives.

#### 2.4. Dynamic Games of Incomplete Information

In this section we consider the case where  $N > 1$ . In this case, each player's optimal decision depends on the expectations that player holds about the actions of the other players and so we require some sort of equilibrium concept. In particular we assume that players use strategies that are consistent with a Markov perfect equilibrium (MPE).

A Markov strategy for player  $i$  is a mapping  $\sigma_i : \mathcal{S} \times \mathcal{N}_i \rightarrow \mathcal{A}_i$  where  $\sigma_i(s, v_i)$  denotes the decision of player  $i$  given the state  $(s, v_i)$ . Given a profile of Markov strategies  $\sigma = (\sigma_1, \dots, \sigma_N)$ , player  $i$ 's expected discounted future utility in state  $s$  can be expressed recursively in terms of the *ex-ante value function*:

$$\bar{V}_i(s | \sigma) = \mathbb{E} \left[ U_i(\sigma(s, v), s, v_i) + \beta \int \bar{V}_i(s' | \sigma) P(ds' | s, \sigma(s, v)) \middle| s \right].$$

The bar denotes that this is the expected value before  $v$  is realized, hence, the expectation is with respect to the distribution of  $v$ .

**Definition.** A *Markov perfect equilibrium* is a strategy profile  $\sigma$  such that for all  $i = 1, \dots, N$  and  $s \in \mathcal{S}$ ,  $\bar{V}_i(s | \sigma_i, \sigma_{-i}) \geq \bar{V}_i(s | \sigma'_i, \sigma_{-i})$  for all alternative Markov strategies  $\sigma'_i$ .

The primitives of the model are the discount factor  $\beta$ , the distribution of private shocks  $G_i$  for each  $i$ , the utility functions  $U_1, \dots, U_N$ , and the joint state transition kernel  $P(ds' | s, a)$ . As was the case with the single agent model above, dynamic games of this form are simply nonlinear state space models of the kind considered above. Here, the conditional density of choices implied by the Markov strategy  $\sigma$  plays the same role in the likelihood as the discrete choice probabilities in the single agent DDC model. Thus, the likelihood can be written as

$$p(y_t | y_{t-1}, z_t, \theta) = p(a_t | x_t, z_t, \theta) p(x_t | a_{t-1}, x_{t-1}, z_t, \theta)$$

where  $p(a_t | x_t, z_t, \theta)$  is the density of  $a_t$  implied by  $\sigma(s, v)$  given the distribution of  $v$  and  $\theta$ . Again, the transition densities are specified as part of the model.

### 3. Sequential Monte Carlo Methods

Sequential Monte Carlo methods are simulation-based methods for approximating posterior distributions of unobserved state variables in nonlinear and non-Gaussian state space models. They avoid making linearity or normality assumptions required by related methods such as the Kalman filter. For simplicity, we focus on the bootstrap filter introduced by [Gordon, Salmond, and Smith \(1993\)](#). There have been many recent advances in SMC methods, but the bootstrap filter captures the essence of most recent methods while remaining very easy to understand.

In this section, we first introduce the optimal nonlinear filtering problem which arises due to the presence of an unobserved state. Sequential Monte Carlo methods, or particle filters, are then discussed as methods to approximate the optimal filter. This state space model and the methods discussed in this section will provide a foundation for the maximum filtered likelihood estimator 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  $z_t$  given the observed data and estimating the unknown parameters  $\theta$ . From a classical perspective,  $\theta$  is deterministic and these are distinct problems. In this setting, we can first recover the posterior distribution of  $z_t$  then use it form a likelihood function with which we can estimate  $\theta$ . From a Bayesian point of view, inference on  $z_t$  and  $\theta$  are essentially the same problem, since  $\theta$  can be treated as a time invariant component of  $z_t$ , with

the joint posterior of the sequence  $z_{1:T}$  and  $\theta$  being the object of interest. The sequential Monte Carlo methods described in the following sections are able to solve both problems.

In particular, we will need to estimate marginal posterior distributions of the form

$$\pi_{t|s}(dz_t) \equiv P(Z_t \in dz_t \mid Y_{1:s} = y_{1:s}).$$

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. For the purposes of discussion, we omit  $\theta$  as a functional argument and element of the conditioning set in this section.

The filtering distribution  $\pi_{t|t}$  can be represented recursively, starting with the initial distribution  $\pi_0$ , by applying a series of prediction and updating steps using Bayes' theorem and the transition kernel  $Q(dz_t \mid z_{t-1}, y_{t-1})$ . Given the marginal filtering distribution  $\pi_{t-1|t-1}$  at time  $t-1$ , the one-step-ahead prediction distribution is

$$(3) \quad \pi_{t|t-1}(dz_t) = \int Q(dz_t \mid z_{t-1}, y_{t-1}) \pi_{t-1|t-1}(dz_{t-1})$$

and the new filtering distribution is

$$(4) \quad \pi_{t|t}(dz_t) = \frac{p(y_t \mid y_{t-1}, z_t) \pi_{t|t-1}(dz_t)}{\int p(y_t \mid y_{t-1}, z_t) \pi_{t|t-1}(dz_t)}.$$

Analytical solutions for the optimal filtering problem are only known for special cases. In particular, Kalman Filters (Kalman, 1960) have been used very successfully in models that are both linear and Gaussian. Unfortunately, the type of models we wish to study are highly nonlinear and non-Gaussian. We instead employ a class of methods known as sequential Monte Carlo methods to approximate the needed filtering distributions.

### 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  $R$  mass points, or particles,  $\{(z_t^r, w_t^r)\}_{r=1}^R$  that evolve over time. The particles can be used to form an empirical probability measure

$$\pi_{t|t}^R(dz) = \frac{\sum_{r=1}^R w_t^r \delta_{z_t^r}(dz)}{\sum_{r=1}^R w_t^r}$$

which approximates  $\pi_{t|t}$ . Here  $\delta_z$  denotes the Dirac delta measure centered at  $z$ .

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 (3), and use it to form an approximation  $\pi_{t|t}^R$  using

an approach motivated by (4). This process is initialized with a collection  $\{z_0^r\}_{r=1}^R$  of iid draws from a chosen initial distribution  $\pi_0$  and is repeated for each new observation  $y_t$ .

Below, we describe a generic particle filter and a few common variations. For a more thorough description of sequential Monte Carlo methods, their implementation, properties, and potential applications 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).

We begin with an iid sample  $\{(z_0^r, 1)\}_{r=1}^R$  of uniformly weighted draws from  $\pi_0$ . The resulting empirical measure  $\pi_0^R$  approximates  $\pi_0$ . At time  $t$ , we have a uniformly weighted collection of particles  $\{(z_{t-1}^r, 1)\}_{r=1}^R$  distributed approximately according to  $\pi_{t-1|t-1}$ . For each  $r = 1, \dots, R$ , following (3), draw

$$\tilde{z}_t^r \sim Q(dz_t | z_{t-1}^r, y_{t-1})$$

and set  $\tilde{w}_t^r = 1$  to form the uniformly weighted particle system  $\{(\tilde{z}_t^r, \tilde{w}_t^r)\}_{r=1}^R$ . The corresponding empirical measure  $\tilde{\pi}_{t|t-1}^R(dz_t) \equiv R^{-1} \sum_{r=1}^R \delta_{\tilde{z}_t^r}(dz_t)$  approximates  $\pi_{t|t-1}(dz_t)$ .

Following (4), we apply the Bayes operator to  $\pi_{t|t-1}^R$  to obtain

$$\pi_{t|t}^R(dz_t) \equiv \frac{p(y_t | y_{t-1}, z_t) \tilde{\pi}_{t|t-1}^R(dz_t)}{\int p(y_t | y_{t-1}, z_t) \tilde{\pi}_{t|t-1}^R(dz_t)} = \frac{\sum_{i=1}^R p(y_t | y_{t-1}, \tilde{z}_t^i) \delta_{\tilde{z}_t^i}(dz_t)}{\sum_{i=1}^R p(y_t | y_{t-1}, \tilde{z}_t^i)}.$$

The weighted particle representation of this distribution is obtained by setting the weights proportional to the likelihood of the new observation  $y_t$ , with  $w_t^r = p(y_t | y_{t-1}, \tilde{z}_t^r)$ , giving the system  $\{(\tilde{z}_t^r, w_t^r)\}_{r=1}^R$ .

Finally, we obtain a uniformly weighted particle system by resampling  $R$  particles  $\{z_t^r\}_{r=1}^R$  from the empirical distribution  $\tilde{\pi}_{t|t}^R$  and setting  $w_t^r = 1$  for all  $r$ . The resulting approximation is

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

From here, the algorithm can be applied recursively using  $\pi_{t|t}^R$  and a new observation  $y_{t+1}$  to obtain  $\pi_{t+1|t+1}^R$ . The algorithm is perhaps easiest understood by looking at the corresponding pseudocode implementation in Algorithm 1.

One of the benefits of using the particle filter in this context is that it only requires evaluating or sampling from densities that already 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  $\tilde{z}_t^r$  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.

Since the resulting distributions  $\pi_{t|t-1}^R$  and  $\pi_{t|t}^R$  are approximations of the true posterior distributions  $\pi_{t|t-1}$  and  $\pi_{t|t}$ , they can be used to approximate expectations of an arbitrary

---

**Algorithm 1** A Generic Particle Filter

---

- Initialization—Draw  $z_0^r \sim \pi_0(dz_0)$  for each  $r = 1, \dots, R$ .
  - Recursion—Repeat the following steps for each  $t = 1, \dots, T$ .
    - Importance sampling—Draw  $\tilde{z}_t^r \sim Q(dz_t | z_{t-1}^r, y_{t-1})$  and set  $w_t^r = p(y_t | y_{t-1}, \tilde{z}_t^r)$  for each  $r = 1, \dots, R$ .
    - Resampling—For each  $r = 1, \dots, R$ , draw  $z_t^r$  from the collection  $\{\tilde{z}_t^r\}_{r=1}^R$  in proportion to the weights  $\{w_t^r\}_{r=1}^R$ .
- 

function  $\phi$ . Since in both cases, at the respective points in the algorithm, the weights are uniform, we have:

$$\int \phi(z) \pi_{t|t-1}(dz) \approx \int \phi(z) \pi_{t|t-1}^R(dz) = \sum_{r=1}^R \phi(\tilde{z}_t^r),$$
$$\int \phi(z) \pi_{t|t}(dz) \approx \int \phi(z) \pi_{t|t}^R(dz) = \sum_{r=1}^R \phi(z_t^r).$$

### 3.3. Variations

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). See Künsch (2005) for an overview of various resampling methods.

Liu and Chen (1998) recommend against resampling every period and suggest both deterministic and dynamic resampling schemes. The deterministic scheme involves resampling at times  $t_0, 2t_0, \dots$  where  $t_0$  is chosen based on the difficulty of the problem. The dynamic scheme involves monitoring the coefficient of variation of the weights  $c(w_t)$ . When the weights are normalized so that  $\sum_r w_t^r = 1$ , then we have

$$c^2(w_t) = \frac{1}{R} \sum_{r=1}^R \{Rw_t^r - 1\}^2$$

The particles should be resampled whenever  $c(w_t)^2 > c_t$  for some sequence of thresholds  $c_t$ , for example,  $c_t = a + bt^\alpha$ .

Liu and Chen (1995) introduce the “effective sample size”  $R_{\text{eff}}$ , a heuristic measure of the efficiency of the particle system:

$$R_{\text{eff}} = \frac{R}{1 + c^2(w)}$$

Intuitively, when  $R_{\text{eff}}$  is small, or  $c^2(w)$  is large, this indicates that there are too many particles in unimportant areas of the state space. If the particle weights are all relatively equal, then  $c^2(w)$  will be small and the effective sample size is larger because each particle is important.

Most of the recent literature uses a simpler definition of the “effective sample size” (Liu, 2001, p. 35):

$$R_{\text{eff}} = \left( \sum_{r=1}^R (w_t^r)^2 \right)^{-1} .$$

This measure varies between 1 and  $R$  and a typical resampling scheme based on  $R_{\text{eff}}$  resamples whenever it falls below some threshold, usually  $R_{\text{eff}} < R/2$ .

### 3.4. Convergence

Until now particle filters were described as sampling methods but in order to discuss convergence, it is more useful to view them as methods for generating sequences of measures which approximate posterior distributions of the unobserved state. As the number of particles increases, this sequence should converge to the true distribution in some sense. The analysis of the asymptotic behavior of particle systems is difficult because the particles interact at each step and are not independent. Despite these difficulties, significant progress has been made in recent years. Crisan and Doucet (2002) provide a very useful survey which discusses a number of convergence results in the context of a general, unified framework. These results only apply to weak convergence in the space of continuous bounded functions and provide conditions under which  $\int \phi(z_t) \pi_{t|t}^R(dz_t) \rightarrow \int \phi(z_t) \pi_{t|t}(dz_t)$  almost surely as  $R \rightarrow \infty$  for some  $\phi \in C_b(\mathbb{R}^{n_z})$ . Hu, Schön, and Ljung (2008) later generalize many of these results to allow  $\phi$  to be unbounded. Chopin (2004) and Künsch (2005) provide the first attempts at developing central limit theorems. Douc and Moulines (2007) take these ideas further, developing a law of large numbers and central limit theorem for a general class of weighted samples and discussing their relevance for sequential Monte Carlo methods.

## 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 estimating the unknown parameters  $\theta$  given a sample of  $N$  iid observations of length  $T$  denoted  $\{y_{1,1:t}, y_{2,1:t}, \dots, y_{N,1:t}\}$  where  $y_{i,s:t} = \{y_{i,\tau}, s \leq \tau \leq t\}$ . Let  $z_{i,s:t}$  be defined similarly.



**Assumption 8** (Parametric Primitives). The model primitives, and thus the distributions governing the joint stochastic process  $\{a_t, x_t, z_t\}_{t=1}^{\infty}$ , depend on a finite vector of parameters  $\theta \in \Theta$  where  $\Theta$  is a compact subset of Euclidean space.

How one estimates the model and how one solves the model, if at all, are two distinct problems. Here our focus is on estimation. Even then, how and when a particle filter can be applied to recover the distribution of the unobserved heterogeneity differs for each estimation method. We consider both a general full-solution approach, which assumes that the likelihood  $p(y_t | y_{t-1}, z_t, \theta)$  can be evaluated (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 integrated 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 stage to recover both the distribution of the unobserved state as well as the policy functions, which depend on the unobserved state.

#### 4.1. Maximum Filtered Likelihood Estimation

First, consider estimating the general nonlinear state-space model of Section 2.1. Given observations  $\{y_{i,1:T}\}_{i=1}^N$ , the log-likelihood function can be written as

$$\begin{aligned} \ln L(\theta) &= \sum_{i=1}^N \ln p(y_{i,1:T} | \theta) = \sum_{i=1}^N \sum_{t=1}^T \ln p(y_{i,t} | y_{i,1:t-1}, \theta) \\ &= \sum_{i=1}^N \sum_{t=1}^T \ln \int p(y_{i,t} | z_{i,t}, y_{i,t-1}, \theta) p(z_{i,t} | y_{i,1:t-1}, \theta) dz_t. \end{aligned}$$

Note that in the second line, after conditioning on  $z_{i,t}$ , we can use the Markov assumption. What we have is an integral with respect to the step-ahead filtering distribution. If we can evaluate the  $p(y_{i,t} | z_{i,t}, y_{i,t-1}, \theta)$  (e.g., by solving the model) and if we can draw from the transition density of  $z_{i,t}$ , then a particle filter can facilitate maximum likelihood estimation of  $\theta$  since we can readily form the approximation

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

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

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

We take the maximum filtered likelihood estimator to be

$$\tilde{\theta} = \arg \max_{\theta \in \Theta} \ln \tilde{L}(\theta).$$

Although this approach is very similar in spirit to that of maximum simulated likelihood (Lerman and Manski, 1981; Lee, 1992), the source of simulation error is subtly different in that it comes from integrating with respect to an approximate empirical measure rather than from noise due to Monte Carlo integration. Each successive approximate measure is constructed from the previous one, which was itself an approximation, so the approximation error propagates forward. The draws are also not independent since the particles interact each period.

In terms of similar approximate maximum likelihood estimators, Olsson and Rydén (2008) propose an estimator which maximizes a smooth, interpolated version of the log-likelihood function. The log-likelihood function is approximated, using filter samples, at points on a finite grid over the parameter space. An approximation over the entire parameter space is then obtained using splines and maximized to obtain parameter estimates. These parameter estimates are shown to be consistent and asymptotically normal under certain conditions on the number of grid points and the number of particles as the sample size tends to infinity.

However, to the best of the author's knowledge, the theoretical properties of the above approximate maximum likelihood estimator are unknown. This is a clear direction for future work in this area. We provide Monte Carlo evidence in Section 5 which suggests that the estimator performs well even in a standard discrete choice model even when the number of particles is small.

#### *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 stage 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-stage 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 stage, 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 is essentially a computational exercise and proceeds exactly as in [Bajari, Benkard, and Levin \(2007\)](#). Extensions to other two-step methods such as the nested-pseudo likelihood estimator of [Aguirregabiria and Mira \(2007\)](#) should also be possible.

#### 4.2.1. First Stage Estimation

In practice, strategies for estimating the policy functions in the first stage tend to be model-specific and depend on the specific distributional assumptions made. The general goal is to flexibly estimate the policy functions  $\sigma_i(s_t, \eta_{it})$  and the state transition density  $p(s_t | s_{t-1}, a_{t-1})$ . In the first stage we are simply interested in flexibly capturing the reduced form relationships between the states and the actions. No payoff parameters are estimated at this stage.

In order to apply the particle filter, the densities implied by the policy functions  $\sigma_i$  and the distribution of  $\eta_{it}$  must belong to some known parametric family of functions. This rules out the use of many flexible nonparametric techniques in the first stage, however in practice researchers have typically used parametric methods such as probit and logit regressions in the first stage when applying this estimator (cf. [Ryan, 2010](#)).

**Assumption 9** (Parametric first stage). The implied choice densities and the transition densities belong to parametric families of functions indexed by a finite vector of parameters  $\alpha$  and can be written as  $p(a_t | s_t, \alpha)$ ,  $p(x_t | x_{t-1}, z_t, a_{t-1}, \alpha)$ , and  $p(z_t | x_{t-1}, z_{t-1}, a_{t-1}, \alpha)$  respectively.

Note that some of the first stage parameters in  $\alpha$ , namely the parameters of the transition equations, are components of  $\theta$ , the structural parameters of interest. The remaining components of  $\theta$ —payoff parameters and other parameters on which the true policy function depends—will be estimated in the second stage.

For now, we consider estimating the first stage parameters. Although  $z_t$  is unobserved, as before, we can obtain parameter estimates  $\hat{\alpha}$  by using the particle filter to sequentially integrate over the recovered distribution of the latent state and then maximizing the approximated log-likelihood:

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

where  $\tilde{z}_{it}^r$  for  $r = 1, \dots, R$  are the particles drawn at the prediction stage of the algorithm for the  $i$ -th observational unit at period  $t$  given the parameters  $\alpha$ . Note that the transition density for  $z_t$ , which also depends on  $\alpha$ , 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  $\alpha$  and so the algorithm must be repeated for each  $\alpha$ .

After obtaining  $\hat{\alpha}$ , 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 precisely all we need in order to carry out the forward simulation procedure of [Bajari, Benkard, and Levin \(2007\)](#).

#### 4.2.2. Second Stage Estimation

With the first-stage estimated policy and transition equations in hand, estimation of the second stage 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  $\hat{\sigma}(s_t, \varepsilon_t)$  denote the joint policy function associated with the first stage estimates  $\hat{\alpha}$ . Given values for the remaining structural parameters  $\theta$ ,<sup>3</sup> the ex-ante value function implied by these policies is

$$\bar{V}_i(s | \hat{\sigma}, \theta) = \mathbb{E} \left[ \sum_{t=0}^{\infty} \beta^t U_i(\hat{\sigma}(s_t, \varepsilon_t), s_t, \varepsilon_{it}, \theta) \mid s_0 = s \right]$$

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

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

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

and using them to form a minimum distance objective function

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

where  $H$  is a distribution over the set of possible inequalities  $\{(i, s, \sigma'_i)\}$ . By minimizing this function, we minimize the sum of squared violations of the equilibrium conditions. The idea is that under the true parameters  $\theta$ , the true policy should always yield higher discounted future payoffs than any alternative policy  $\sigma'_i$  for each player  $i$ .

In practice, even with fully observed state variables, this procedure is sensitive both to the first stage estimates and the chosen distribution of player 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 Monte Carlo experiments presented in the following section.

---

<sup>3</sup>We simply write the remaining parameters as  $\theta$  with the understanding that some parameters previously in  $\theta$ , such as transition density parameters, have been estimated in the first stage.

## 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) which we use to perform a series of Monte Carlo experiments and obtain results using the original data. The extended model has a two-dimensional continuous state space, rather than a one-dimensional discrete one, and one of these states is a latent state variable.

The agent in the model 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,  $z_t$ , which is observed by the decision maker but not by the researcher and may be serially correlated. Time periods are equal to one month and we assume the discount rate is  $\beta = 0.95$ .

First, we must specify functional forms for the model primitives: the payoff function (in this case, the cost function) and the transition kernels for  $x$  and  $z$ . For simplicity, we assume the cost function is linear:

$$U(a_t, s_t, \varepsilon_t | \theta) = \begin{cases} -c_x x_t - c_z z_t + \varepsilon_{t,0} & \text{if } a_t = 0, \\ -c_0 + \varepsilon_{t,1} & \text{if } a_t = 1. \end{cases}$$

The structural parameters of interest here are the replacement cost,  $c_0$ , the cost of mileage,  $c_x$ ,<sup>4</sup> and the cost associated with the latent state,  $c_z$ .

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

$$p(\Delta x_t | x_{t-1}, z_t, \theta) = \lambda(z_t, \theta) e^{-\lambda(z_t, \theta) \Delta x_t}$$

where  $\lambda(z_t, \theta) = \exp(\lambda_0 + \lambda_z z_t)$ . This exponential functional form ensures that the rate parameter  $\lambda(z_t, \theta)$  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 while  $a_t = 0$ :  $z_{t+1} | x_t, z_t, a_t \sim N(\rho z_t, \sigma^2)$ .<sup>5</sup> When the choice to replace is made ( $a_{t-1} = 1$ ), the engine is replaced and the exponentially distributed mileage increment above is the initial mileage, starting at zero, and the latent state resets to zero.

We must solve for the choice-specific value function  $v_\theta(a, s)$  both in order to simulate data and in order to estimate the model using maximum likelihood below. For any state  $s$  and choice  $d$ , recall that  $v_\theta$  satisfies  $v_\theta - \Gamma(v_\theta) = 0$ , where  $\Gamma$  is the functional operator defined in (2).

<sup>4</sup>Following Rust (1987), we scale  $x_t$ , the 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.

<sup>5</sup>Instead of estimating  $\sigma$  directly, we fix it at  $\sigma = 0.5$  and estimate coefficients on  $z_t$  in the cost function and transition equation.

Although  $\nu_\theta$  is an infinite-dimensional object, we can form a smooth approximation to  $\nu_\theta$  with only a finite number of parameters using Chebyshev polynomials (cf. Judd, 1998). In all cases below we employ the smooth resampling algorithm of Pitt (2002) (Algorithm 2) so that for a fixed simulator sequence used to generate and propagate the particles, the resulting likelihood functions will be smooth. See Appendix A for further computational details.

### 5.1. Monte Carlo Experiments: Maximum Likelihood Estimation


In this section we present the results of a series of Monte Carlo experiments carried out by estimating the continuous state model described above using the approximate maximum likelihood estimator.

The filter-approximated log-likelihood function is a random function which depends on the underlying simulators used to generate the particle swarm. The accuracy of this approximation should increase as the number of particles increases. That is, for a fixed value of  $\theta$ , the variance of the log-likelihood function should decrease.


We illustrate this point through a series of cross-sectional plots, each displaying the median, upper and lower quartiles, and range of 100 replicated log-likelihood values at several regularly spaced values of  $\theta$ . We first generate a single dataset consisting of  $N = 100$  observations of  $T = 100$  periods each. Each replication of the log-likelihood function is then produced by drawing a new sequence of random numbers used to propose and transition particles. Figure 1 displays the function for  $R = 25$  particles. In each panel, the true parameter value is indicated by a vertical line. At regular intervals for each parameter, the minimum, maximum, and all quartiles of the log likelihood values are displayed using box plots. In each dimension, as expected, the log-likelihood function is maximized approximately at the true parameter. Figure 2 presents the corresponding plots for  $R = 100$  particles. The scale is held constant across both figures to illustrate the much improved accuracy of the second plot.

Recall that in addition to estimating the parameters  $\theta$ , the trajectory of the latent state for each observational unit (here, for each bus) can be estimated as well. This is illustrated in Figure 3, which displays a single simulation from the continuous state bus engine replacement model along with the approximate posterior distribution of the unobserved state recovered using the particle filter. The realization of  $x_t$  is plotted in the lower panel. The realization of  $z_t$  is plotted in the upper panel along with the quantiles of the approximate posterior distribution. Figure 4 plots the same weighted particle swarm as sequence of distributions. The initial 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, but quickly assigns more weight to more relevant areas of the state space.


Finally, we generate several datasets consisting of  $N = 100$  bus observations of  $T = 100$




(a)  $\lambda_0$




(b)  $\lambda_z$




(c)  $\rho$



(d)  $c_0$




(e)  $c_x$




(f)  $c_z$


FIGURE 1. Approximated Likelihood Function: 25 Particles, 100 Replications.




(a)  $\lambda_0$




(b)  $\lambda_z$




(c)  $\rho$



(d)  $c_0$



(e)  $c_x$



(f)  $c_z$

FIGURE 2. Approximated Likelihood Function: 100 Particles, 100 Replications.






FIGURE 3. Continuous State Model Simulation and Approximate Posterior

periods each and estimate the model several ways. [Table 1](#) reports the means and standard deviations (in parentheses) 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  $\pi_0 = N(0, \sigma_0^2)$  with  $\sigma_0 = 0.5$ .

First, [Table 1](#) reports the naïve maximum likelihood estimates obtained when ignoring the latent state  $z_t$ . In this case we are only estimating the average mileage transition rate  $\lambda_0$ , without capturing the effects of  $z_t$ . As such, our estimates of  $\lambda_0$  are biased. Similarly, we are significantly underestimating both the replacement cost,  $c_0$ , and the cost of mileage,  $c_x$ . 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 25 particles, the estimates of the cost parameters appear to have a slight downward bias, but after 50 particles, the mean bias is small for all parameters. The bottom panel reports the maximum likelihood estimates obtained when the latent state is fully observed. It is remarkable that the estimates with 100 particles are nearly as good as the estimates with fully observed data.




FIGURE 4. Evolution of Approximated Posterior Distributions

### 5.2. 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  $\sigma$  satisfies the following equilibrium condition

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

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

$$P(\sigma(s, \varepsilon) = 0 \mid s) = \frac{\exp(v(0, s))}{1 + \exp(v(0, s))}$$

and  $P(\sigma(s, \varepsilon) = 1 \mid s) = 1 - P(\sigma(s, \varepsilon) = 0 \mid s)$ .

Now, suppose that we use a simple logit model to approximate the true policy:

$$P(\sigma(s, \varepsilon) = 0 \mid s) \approx \frac{\exp(f(s, \alpha))}{1 + \exp(f(s, \alpha))}$$

where  $f(s, \alpha) = \alpha_0 + \alpha_1 x + \alpha_2 z + \alpha_3 xz$ .

The parametric specifications for the state transition densities are specified as part of the model, and so we can estimate the parameters of those densities and the parameters  $\alpha$  from the

Estimator	$\lambda_0$	$\lambda_z$	$\rho$	$c_0$	$c_x$	$c_z$
Population	0.400	0.300	0.800	7.000	1.000	0.500
Naïve Estimation	0.355 (0.010)	– –	– –	5.342 (0.173)	0.571 (0.054)	– –
Particle Filter ( $R = 25$ )	0.400 (0.011)	0.298 (0.055)	0.800 (0.060)	6.471 (0.753)	0.850 (0.229)	0.409 (0.214)
Particle Filter ( $R = 50$ )	0.402 (0.011)	0.296 (0.055)	0.800 (0.061)	6.945 (1.245)	1.017 (0.456)	0.504 (0.312)
Particle Filter ( $R = 100$ )	0.402 (0.011)	0.301 (0.053)	0.793 (0.056)	6.892 (0.908)	0.987 (0.317)	0.501 (0.227)
Observed Latent State	0.401 (0.009)	0.299 (0.014)	0.798 (0.005)	7.014 (0.303)	1.009 (0.105)	0.501 (0.029)

TABLE 1. Maximum Filtered Likelihood Estimates (25 Replications)

parametric policy approximation using maximum likelihood. As before, we can use the particle filter samples to approximate the likelihood function.

With an estimate  $\hat{\alpha}$  in hand, the approximate policy function is simply

$$\hat{\sigma}(s, \varepsilon) = \begin{cases} 0 & \text{if } f(s, \hat{\alpha}) + \varepsilon_0 - \varepsilon_1 > 0, \\ 1 & \text{otherwise.} \end{cases}$$

Given  $\hat{\alpha}$ , we know  $\hat{\sigma}$  and the estimated transition densities, with which we can use forward simulation to approximate the ex-ante value function  $\bar{V}_i(s | \sigma)$  for any  $s$ .

In the second stage, 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

$$U(a_t, s_t, \varepsilon_t | \theta) = (1 - a_t)(-c_x x_t - c_z z_t + \varepsilon_{t,0}) + a_t(-c_0 + \varepsilon_{t,1}).$$

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

$$\begin{aligned} \bar{V}(s_0 | \sigma, \theta) &= \mathbb{E} \left[ \sum_{t=0}^{\infty} \beta^t U(a_t, s_t, \varepsilon_t | \theta) \right] \\ &= \mathbb{E} \left[ \sum_{t=0}^{\infty} \beta^t (1 - a_t)(-c_x x_t - c_z z_t + \varepsilon_{t,0}) + \beta^t a_t(-c_0 + \varepsilon_{t,1}) \right] \\ &= \mathbb{E} \left[ \sum_{t=0}^{\infty} -c_x \beta^t (1 - a_t) x_t - c_z \beta^t (1 - a_t) z_t + \beta^t (1 - a_t) \varepsilon_{t,0} - c_0 \beta^t a_t + \beta^t a_t \varepsilon_{t,1} \right] \\ &= \mathbb{E} \left[ \sum_{t=0}^{\infty} \beta^t ((1 - a_t) \varepsilon_{t,0} + a_t \varepsilon_{t,1}) \right] - c_0 \mathbb{E} \left[ \sum_{t=0}^{\infty} \beta^t a_t \right] \\ &\quad - c_x \mathbb{E} \left[ \sum_{t=0}^{\infty} \beta^t (1 - a_t) x_t \right] - c_z \mathbb{E} \left[ \sum_{t=0}^{\infty} \beta^t (1 - a_t) z_t \right] \end{aligned}$$

Estimator	$\alpha_0$	$\alpha_1$	$\alpha_2$	$\alpha_3$	$\lambda_0$	$\lambda_z$	$\rho$
Population	–	–	–	–	0.400	0.300	0.800
Particle Filter ( $R = 25$ )	5.569 (0.360)	-6.375 (0.787)	-0.883 (0.755)	-2.665 (3.612)	0.400 (0.012)	0.290 (0.055)	0.807 (0.058)
Particle Filter ( $R = 50$ )	6.002 (0.896)	-6.955 (1.408)	-1.291 (0.779)	-2.845 (1.781)	0.403 (0.010)	0.283 (0.047)	0.807 (0.053)
Particle Filter ( $R = 100$ )	6.218 (0.990)	-6.936 (1.486)	-1.210 (1.606)	-2.201 (2.672)	0.401 (0.014)	0.282 (0.075)	0.761 (0.166)
Observed Latent State	5.908 (0.154)	-7.351 (0.529)	-1.681 (0.124)	0.161 (0.485)	0.401 (0.009)	0.299 (0.014)	0.798 (0.005)

TABLE 2. Two-step estimator: first stage estimates (25 replications)

where  $a_t = \sigma(s_t, \varepsilon_t)$ .

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

$$\begin{aligned}
\hat{V}(s_0 | \sigma, \theta) &= \frac{1}{L} \sum_{l=1}^L \sum_{t=0}^T \beta^t U(\hat{\sigma}(s_t^l, \varepsilon_t^l), s_t^l, \varepsilon_{t+1}^l | \theta) \\
&= \frac{1}{L} \sum_{l=1}^L \sum_{t=0}^T \beta^t \left( (1 - a_t^l) \varepsilon_{t,0}^l + a_t^l \varepsilon_{t,1}^l \right) - c_0 \frac{1}{L} \sum_{l=1}^L \sum_{t=0}^T \beta^t a_t^l \\
&\quad - c_x \frac{1}{L} \sum_{l=1}^L \sum_{t=0}^T \beta^t (1 - a_t^l) x_t^l - c_z \frac{1}{L} \sum_{l=1}^L \sum_{t=0}^T \beta^t (1 - a_t^l) z_t^l.
\end{aligned}$$

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

In each of these experiments, we use 2000 inequalities—state and alternative policy combinations. For each inequality, we simulate  $L = 1000$  paths of length  $\bar{T} = 125$ .<sup>6</sup> Alternative policies are taken to be perturbations of the threshold crossing condition, where the value of the perturbation is

$$\eta_1 + \eta_2 x_t + \eta_3 z_t + \eta_4 x_t z_t$$

where  $\eta_j \sim N(0, \sigma_\eta^2)$ . We vary the standard deviation  $\sigma_\eta$  across experiments to see how the second

<sup>6</sup>This gives  $\beta^{\bar{T}} = 0.95^{125} \approx 0.0016422$ .

stage estimates behave. Initial states for each inequality were drawn uniformly over the state space.

The results of our experiments are summarized in Tables 2 and 3. For each of the estimators considered, Table 2 reports the first stage estimates and Table 3 reports the second stage estimates for  $\sigma_\eta \in \{0.05, 0.10, 0.20\}$ . We compare the two-step estimator using the particle filter with  $R = 25$  and  $R = 50$  with the infeasible two-step estimates obtained when the latent state is fully observable. In addition, we also present values of the infeasible second-stage estimates obtained using the true policy functions (note that there is no data involved in this stage, so the first stage 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.

Our overall impression from these results is that the two-step estimator behaves well on average, but that the estimates are noisier than the full-solution maximum likelihood estimates. The second-stage 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. Except for the infeasible estimator where the true policy functions are used, the estimates are rather sensitive to the choice of  $\sigma_\eta$ . In large scale discrete time dynamic games, however, where full-solution estimation is infeasible, the trade-off is different and the additional noise may be a small price to pay for obtaining estimates that control for the latent state variable.

### 5.3. Empirical Results

We also estimate the model using the original data of Rust (1987). We use the data for bus group number four which consists of 37 buses (1975 GMC A5308) observed over 117 months each for a total of 4,329 observations. Table 4 reports the estimates for several values of  $R$ , the number of particles. In particular, we estimate the model with  $R \in \{50, 100, 200\}$  with nearly identical results for each value of  $R$ . The first row of the table gives the naïve maximum likelihood estimates obtained when ignoring the latent state  $z_t$ .

These estimates were obtained using a  $N(0, \sigma_0^2)$  proposal distribution with  $\sigma_0 = 0.5$ . We also used  $\sigma_0 = 1.0$ ,  $\sigma_0 = 2.0$ , and  $\sigma_0 = 3.0$  with very similar results. The estimates were also robust to different seed values for the underlying random number generator, which result in different filter samples. We only report the results for a single run using a fixed seed value, but again, the results were essentially identical for other seeds.

Overall, we find evidence of a latent state variable that is strongly serially correlated and, while it has no effect on costs, it has a significant effect on mileage transitions. Through our parametrization of the scale parameter of the exponential distribution, higher values of the latent state  $z_t$  increase  $\lambda(z_t, \theta)$  which in turn decreases the mean of the mileage increment. That is,

buses with larger values of  $z_t$  tend to be driven fewer miles each month.

Using the results from [Table 4](#), we can perform a likelihood ratio test for the null hypothesis that the constrained or naïve model is the correct model. Using the value of the log-likelihood function for  $R = 200$  particles, we strongly reject the null since  $LR = 87.36$  and  $P(\chi_3^2 > 87.36) \approx 0$ .

## 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 recover the trajectory 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 provide Monte Carlo evidence to highlight the performance of both estimators in the context of a generalized version of the bus engine replacement model of [Rust \(1987\)](#). We then apply the maximum filtered likelihood estimator using [Rust's](#) original dataset and find evidence for a strongly serially correlated latent state variable that affects mileage transitions but not the cost of operating the bus.

### 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 eight Chebyshev polynomials to approximate the value function and eighth-order quadrature to approximate the double integral in the functional operator  $\Gamma$ , defined in [\(2\)](#), in line with the procedures described below.

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

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 following recurrence relation:

$$\begin{aligned} T_0(x) &= 1, \\ T_1(x) &= x, \\ T_{k+1}(x) &= 2xT_k(x) - T_{k-1}(x). \end{aligned}$$

Let  $\mathcal{H}_k$  denote the set of roots of the  $k$ -th degree Chebyshev polynomial of the first kind,

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

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

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

For each  $a \in \mathcal{A}$ , we approximate  $v_\theta(s, a)$  by a function  $v_\Psi(s, a)$  over the region  $[\underline{x}, \bar{x}] \times [\underline{z}, \bar{z}]$  by taking products of Chebyshev polynomials of degree  $K-1$  in each dimension:

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

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

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

Intuitively, for a given  $\theta$  one wants to choose the values  $\psi_{ij}^a$  in order to make the difference between  $v_\Psi(x, z, a)$  and  $\Gamma(v_\Psi)(x, z, a)$  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  $\theta$ :

$$(5) \quad Q_\theta(\Psi) = \sum_{a \in \mathcal{A}} \sum_{(x,z) \in \mathcal{H}_k^{[\underline{x}, \bar{x}]} \times \mathcal{H}_k^{[\underline{z}, \bar{z}]}} [v_{\theta, \Psi}(x, z, a) - \Gamma(v_{\theta, \Psi})(x, z, a)]^2.$$

Assume for a moment that we can evaluate  $\Gamma$  numerically. Then, for each  $\theta$  the value of  $\Psi$  which minimizes  $Q_\theta(\Psi)$  is used to approximate  $v_\theta$  in the log-likelihood function. In the Monte Carlo experiments and application, we use Newton's method to solve for the coefficients  $\Psi$  by searching for a zero of the residual function in (5).

### A.2. Evaluating $\Gamma$

To evaluate  $\Gamma$  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  $z$ ) and Gauss-Hermite quadrature for the integral with respect to the normal distribution.

Gauss-Laguerre quadrature of order  $n$  provides abscissæ  $\zeta_i$  and weights  $\omega_i$  for  $i = 1, \dots, n$  for the following linear approximation:

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

If  $x$  is an exponential random variable with rate parameter  $\lambda$ , then the expectation of some function  $f(x)$  can be approximated via a simple transformation. Let  $\zeta = \lambda x$  and  $\phi(\cdot) = f(\cdot/\lambda)$ , and so  $d\zeta = \lambda dx$ . Then

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

Similarly, Gauss-Hermite quadrature provides weights  $\omega_i$  and abscissæ  $\zeta_i$  for integral approximations of the form:

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

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

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

### A.3. Optimization

Simulated Annealing was used for optimization in all cases: for the maximum likelihood estimates, for the first stage maximum likelihood policy estimation, and for the second stage minimum distance objective function. Starting values were chosen to be either all ones, or some variation such as all values equal to 0.1, depending on the magnitude of the parameters. 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.



## References

- Ackerberg, D., L. Benkard, S. Berry, and A. Pakes (2007). Econometric tools for analyzing market outcomes. In J. J. Heckman and E. E. Leamer (Eds.), *Handbook of Econometrics*, Volume 6A. North Holland. [3]
- Aguirregabiria, V. and P. Mira (2007). Sequential estimation of dynamic discrete games. *Econometrica* 75, 1–53. [4, 19]
- Aguirregabiria, V. and P. Mira (2010). Dynamic discrete choice structural models: a survey. *Journal of Econometrics* 156(1), 38–67. [3, 9]
- Arcidiacono, P. and R. A. Miller (2010). CCP estimation of dynamic discrete choice models with unobserved heterogeneity. Working paper, Duke University. [4]
- Bajari, P., C. L. Benkard, and J. Levin (2007). Estimating dynamic models of imperfect competition. *Econometrica* 75, 1331–1370. [1, 2, 17, 18, 19, 20, 27, 30]
- Cappé, O., E. Moulines, and T. Ryden (2005). *Inference in Hidden Markov Models*. New York: Springer. [14]
- Chopin, N. (2004). Central limit theorem for sequential monte carlo methods and its application to bayesian inference. *Annals of Statistics* 32, 2385–2411. [16]
- Crisan, D. and A. Doucet (2002). A survey of convergence results on particle filtering methods for practitioners. *IEEE Transactions on Signal Processing* 50, 736–746. [16]
- Douc, R. and E. Moulines (2007). Limit theorems for weighted samples with applications to sequential monte carlo methods. In *ESAIM: Proceedings*, Volume 19, pp. 101–107. [16]
- Doucet, A., N. de Freitas, and N. Gordon (Eds.) (2001). *Sequential Monte Carlo Methods in Practice*. New York: Springer. [4, 14]
- Eckstein, Z. and K. I. Wolpin (1989). The specification and estimation of dynamic stochastic discrete choice models: A survey. *The Journal of Human Resources* 24, 562–598. [3]
- Fernández-Villaverde, J. and J. F. Rubio-Ramírez (2007). Estimating macroeconomic models: A likelihood approach. *Review of Economic Studies* 74, 1059–1087. [4]
- Gallant, A. R., H. Hong, and A. Khwaja (2011). Dynamic entry with cross product spillovers: An application to the generic drug industry. Working paper, Duke University. [5]

- Gordon, N., D. Salmond, and A. Smith (1993). Novel approach to nonlinear/non-gaussian bayesian state estimation. *IEEE Proc. F, Radar and signal processing* 140, 107–113. [12, 15]
- Hotz, V. J. and R. A. Miller (1993). Conditional choice probabilities and the estimation of dynamic models. *Review of Economic Studies* 60, 497–529. [4]
- Hu, X.-L., T. B. Schön, and L. Ljung (2008). A basic convergence result for particle filtering. *IEEE Transactions on Signal Processing* 56, 1337–1348. [16]
- Hu, Y. and M. Shum (2008). Identifying dynamic games with serially-correlated unobservables. Working paper, Johns Hopkins University. [4]
- Hu, Y. and M. Shum (2010). Nonparametric identification of dynamic models with unobserved state variables. Working paper, Johns Hopkins University. [4, 7]
- Imai, S., N. Jain, and A. Ching (2009). Bayesian estimation of dynamic discrete choice models. *Econometrica* 77, 1865–1899. [4]
- Judd, K. L. (1998). *Numerical Methods in Economics*. Cambridge, MA: MIT Press. [11, 22]
- Kalman, R. E. (1960). A new approach to linear filtering and prediction problems. *Transactions of the ASME—Journal of Basic Engineering* 82(Series D), 35–45. [13]
- Kitagawa, G. (1996). Monte carlo filter and smoother for non-Gaussian nonlinear state space models. *Journal of Computational and Graphical Statistics* 5, 1–25. [4]
- Künsch, H. R. (2001). State space and hidden Markov models. In O. E. Barndorff-Nielsen, D. R. Cox, and C. Klüppelberg (Eds.), *Complex Stochastic Systems*, pp. 109–173. Chapman and Hall. [14]
- Künsch, H. R. (2005). Recursive Monte Carlo filters: algorithms and theoretical analysis. *Annals of Statistics* 33, 1983–2021. [15, 16]
- Lee, L.-F. (1992). On efficiency of methods of simulated moments and maximum simulated likelihood estimation of discrete response models. *Econometric Theory* 8, 518–552. [18]
- Lerman, S. R. and C. F. Manski (1981). On the use of simulated frequencies to approximate choice probabilities. In C. F. Manski and D. L. McFadden (Eds.), *Structural Analysis of Discrete Data with Econometric Applications*. MIT Press. [18]
- Liu, J. (2001). *Monte Carlo Strategies in Scientific Computing*. New York: Springer. [14, 16]

- Liu, J. S. and R. Chen (1995). Blind deconvolution via sequential imputations. *Journal of the American Statistical Association* 90, 567–576. [15]
- Liu, J. S. and R. Chen (1998). Sequential Monte Carlo methods for dynamic systems. *Journal of the American Statistical Association* 93, 1032–1044. [15]
- Norets, A. (2009). Inference in dynamic discrete choice models with serially correlated unobserved state variables. *Econometrica* 77, 1665–1682. [4]
- Olsson, J. and T. Rydén (2008). Asymptotic properties of particle filter-based maximum likelihood estimators for state space models. *Stochastic Processes and their Applications* 118, 649–680. [18]
- Pakes, A. (1986). Patents as options: Some estimates of the value of holding European patent stocks. *Econometrica* 54, 755–784. [3]
- Pakes, A. (1994). Estimation of dynamic structural models: Problems and prospects part II: Mixed continuous-discrete models and market interactions. In J. Laffont and C. Sims (Eds.), *Advances in Econometrics: Proceedings of the 6th World Congress of the Econometric Society*. [3]
- Pitt, M. K. (2002). Smooth particle filters for likelihood evaluation and maximisation. Working Paper 651, University of Warwick. [22]
- Pitt, M. K. and N. Shephard (1999). Filtering via simulation: Auxiliary particle filters. *Journal of the American Statistical Association* 94, 590–591. [15]
- Rust, J. (1987). Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher. *Econometrica* 55, 999–1013. [1, 2, 3, 5, 21, 29, 30]
- Rust, J. (1994). Structural estimation of Markov decision processes. In R. F. Engle and D. L. McFadden (Eds.), *Handbook of Econometrics*, Volume 4, Amsterdam. North Holland. [3, 9, 11]
- Ryan, S. (2010). The costs of environmental regulation in a concentrated industry. Working paper, Massachusetts Institute of Technology. [19]
- Srisuma, S. (2010). Estimation of structural optimization models: A note on identification. Working paper, London School of Economics. [20]
- Timmins, C. (2002). Measuring the dynamic efficiency costs of regulators' preferences: Municipal water utilities in the arid west. *Econometrica* 70, 603–629. [4]

Estimator	$\sigma_\eta$	$c_0$	$c_x$	$c_z$
Population	–	7.000	1.000	0.500
Particle Filter ( $R = 25$ )	0.05	6.069 (1.464)	0.780 (0.328)	0.378 (0.192)
	0.10	6.248 (1.487)	0.817 (0.339)	0.396 (0.201)
	0.20	6.487 (1.404)	0.870 (0.326)	0.418 (0.209)
Particle Filter ( $R = 50$ )	0.05	6.507 (1.074)	0.825 (0.255)	0.460 (0.240)
	0.10	6.761 (1.116)	0.875 (0.259)	0.488 (0.261)
	0.20	7.042 (1.212)	0.926 (0.265)	0.523 (0.289)
Particle Filter ( $R = 100$ )	0.05	6.561 (1.277)	0.854 (0.253)	0.516 (0.238)
	0.10	6.852 (1.466)	0.910 (0.283)	0.554 (0.286)
	0.20	7.210 (1.698)	0.979 (0.318)	0.603 (0.335)
Observed Latent State	0.05	7.029 (0.538)	1.134 (0.117)	0.479 (0.068)
	0.10	7.464 (0.488)	1.226 (0.123)	0.514 (0.065)
	0.20	7.900 (0.414)	1.319 (0.133)	0.551 (0.058)
True Policy Function	0.05	7.008 (0.066)	0.908 (0.011)	0.492 (0.004)
	0.10	7.033 (0.035)	0.910 (0.007)	0.494 (0.003)
	0.20	6.988 (0.050)	0.900 (0.011)	0.491 (0.004)

TABLE 3. Two-step estimator: second stage estimates (25 replications)

Estimator	LL	$\lambda_0$	$\lambda_z$	$\rho$	$c_0$	$c_x$	$c_z$
Naïve Estimator	-2393.46	0.481 (0.015)	– –	– –	8.532 (0.841)	0.540 (0.100)	– –
Particle filter ( $R = 50$ )	-2351.77	0.499 (0.028)	0.108 (0.012)	0.985 (0.007)	8.758 (0.959)	0.565 (0.116)	0.000 (0.014)
Particle filter ( $R = 100$ )	-2349.69	0.503 (0.029)	0.106 (0.012)	0.986 (0.008)	8.775 (0.978)	0.567 (0.118)	0.000 (0.014)
Particle filter ( $R = 200$ )	-2349.78	0.498 (0.029)	0.105 (0.012)	0.986 (0.007)	8.759 (0.958)	0.566 (0.116)	0.000 (0.014)

TABLE 4. Estimates for Bus Group 4 of Rust (1987)