Bayesian Estimation of a Dynamic Game with Endogenous, Partially Observed, Serially Correlated State *

A. Ronald Gallant
Duke University and
New York University

Han Hong
Stanford University

Ahmed Khwaja
Yale University

First draft: June 2010
Current draft: December 2010

Abstract

We consider inference for dynamic games that can have state variables that are partially observed, serially correlated, endogenous, and heterogeneous. We propose a Bayesian method that uses a particle filter to compute an unbiased estimate of the likelihood within a Metropolis chain. Unbiasedness guarantees that the stationary density of the chain is the exact posterior, not an approximation. The remarkable feature of this approach is that the number of particles required is both small and easily determined.

Keywords: Dynamic Games, Partially Observed State, Endogenous State, Serially Correlated State, Particle Filter.

JEL Classification: E00, G12, C51, C52

*Supported by the National Science Foundation and the Alfred P. Sloan Foundation. We thank Neil Shephard, Michael Pitts, and Thomas Flury for helpful discussions. Corresponding author is A. Ronald Gallant, Duke University, DUMC 90120, Durham NC 27708-0120, USA; aronaldg@gmail.com. Most recent version available at http://econ.duke.edu/webfiles/arg/papers/socc.pdf.

1 Introduction

We propose a likelihood based Bayesian method to estimate any dynamic game with partially observed state that has a Markovian representation of the dynamics and an algorithm to solve the game including those with serially correlated, endogenous, heterogeneous, state variables. The method uses sequential importance sampling (particle filter) to compute an unbiased estimate of the likelihood within a Metropolis chain (MCMC chain). Unbiasedness guarantees that the stationary density of the chain is the exact posterior, not an approximation. The remarkable feature of this approach is that the number of particles required is both small and easily determined. The key idea, which directly contradicts most current practice, is that at each repetition of the MCMC chain the initial seed that determines all random draws within the filter is proposed along with the parameter values. An interesting side effect is that the jagged likelihood of a discrete choice game appears smoother to the MCMC algorithm than it would if the seed were fixed thereby making the likelihood easier to explore.

An important class of games to which our methods apply are oligopolistic entry games where costs or other state variables are serially correlated, unobserved, and endogenous. In dynamic, oligopolistic, entry games, the spillover effect that either experience gained due to entering a market (generic pharmaceuticals) or a capacity constraint caused by entering a market (patent law) has on subsequent performance in the market for a similar product is of interest. Entry decisions of firms in a forward looking dynamic environment are drastically different from those in a static competitive environment. A firm might enter a particular product market even if the current opportunity is not profitable as long as the spillovers from entry sufficiently improve the discounted stream of cumulative future profits. Models can incorporate dynamic spillover due to entry by allowing for serially correlated, firm-specific costs that evolve endogenously based on past entry decisions. Endogeneity of costs to past entry decisions induces heterogeneity among firms even if they are identical ex ante. The examples used in our simulation exercise are dynamic, oligopolistic, entry games of this sort and we focus on the spillover effect in the assessment of the estimator that we propose.

Empirical models of static and dynamic games are differentiated by their information

structures. In our examples we focus on complete information games in which all information is publicly available to all firms and individual firms do not have privately observed profit shocks. The complete information assumption allows substantial unobserved heterogeneity at the level of the firms. These games typically require the use of a combinatorial algorithm to search for an equilibrium instead of the continuous fixed point mapping used in incomplete information models to compute equilibria.

While static games of complete information have been estimated by, e.g., Bresnahan and Reiss (1991a), Berry (1992), Tamer (2003), Ciliberto and Tamer (2009) and Bajari, Hong, and Ryan (2010), to our knowledge, Gallant, Hong, and Khwaja (2010) is the first to estimate a dynamic game of complete information. The second is Chen (2010) who applied a sequential Monte Carlo algorithm to the game and data of Gallant, Hong, and Khwaja (2010), arriving at substantially the same results. The present paper marks the third. In contrast to Gallant, Hong, and Khwaja (2010), the models used in the simulation studies in the current paper allow for heterogeneity between the evolution of the unobserved and observed components of the cost variables. The current paper also investigates the statistical properties of the particle filter and estimator. The statistical properties of the estimators are not studied rigorously in the previous two papers. We establish the unbiasedness property of an estimate of the likelihood function obtained through particle filtering simulation in a model that allows for endogenous feedback from the observed measurements to the dynamic state variables. Endogenous feedback is the feature that distinguishes this paper from the bulk of the particle filter literature. We establish our results by means of a recursive setup and an inductive argument that avoids the complexity of ancestor tracing during the resampling steps. This approach allows elegant, compact proofs.

The paper is organized as follows. We begin by discussing the related literature in Section 2. Section 3 describes the general model and the suite of models used in the simulation exercise. The method used to solve the models in the suite is of some independent interest and is discussed in Section 4. An algorithm for unbiased estimation of a likelihood is proposed and unbiasedness is proved in Section 5. The MCMC estimation algorithm is presented in Section 6. Simulation results are reported in Section 7. Section 8 concludes.

2 Related Literature

There is a growing literature on the estimation of games. Static games under the incomplete information assumption have been studied by, e.g., Bjorn and Vuong (1997), Bresnahan and Reiss (1991b), Bresnahan and Reiss (1991c), Haile, Hortaçsu, and Kosenok (2008), Aradillas-Lopez (2010), Ho (2009), Ishii (2005), Pakes, Porter, Ho, and Ishii (2005), Augereau, Greenstein, and Rysman (2006), Seim (2006), Sweeting (2006), Tamer (2003), Rysman (2004), Gowrisankaran and Stavins (2004), Ellickson and Misra (2008) and Bajari, Hong, Krainer, and Nekipelov (2006). Dynamic games of incomplete information have been studied by, e.g., Aguirregabiria and Mira (2007), Bajari, Benkard, and Levin (2007), Pakes, Ostrovsky, and Berry (2007), Ryan (2005), Collard-Wexler (2010), and Bajari, Chernozhukov, Hong, and Nekipelov (2007). The literature on estimating games of incomplete information has mostly relied on a two step estimation strategy building on the conditional choice probability estimator of Hotz and Miller (1993).

The two step estimation strategy requires the assumption that there is no market or firm level unobserved heterogeneity other than a random shock that is independent and identically distributed across both time and players. This assumption is restrictive because it rules out unobserved dynamics in the latent state variables. It also rules out any private information that a player might have about competing firms that the researcher does not have. Arcidiacono and Miller (2008) have extended the literature on two step conditional choice estimation of dynamic discrete models to allow for discrete forms of unobserved heterogeneity using the EM algorithm. In contrast, our method is applicable even when the unobserved variable is continuous. Moreover, while two step methods can be computationally attractive, we think that a likelihood based method, as the one we are employing, has advantages when the model is misspecified. In this case, the likelihood based approach still minimizes a well defined Kullback-Leibler distance between the model and the data whereas two step methods do not.

In the single agent dynamic framework, there is a considerable amount of research that allows for time invariant unobserved heterogeneity, e.g., Keane and Wolpin (1997). However there is very little work that allows for serially correlated unobserved endogenous state

variables. In the context of a finite horizon dynamic discrete choice model, Khwaja (2001) developed a simulation based method to integrate out such state variables from the likelihood exploiting the discrete nature and Markovian dynamic structure of the variables. Bayesian approaches for single agent dynamic discrete choice models with unobserved state variables that are serially correlated over time have been developed by Imai, Jain, and Ching (2009) and Norets (2009). These papers use MCMC for integrating out the unobserved state variables. In contrast, we use sequential importance sampling to integrate out the unobserved state variables. In addition we are the first to apply this method for the estimation of a dynamic game whereas the previous literature has focused on single agent models.

Fernandez-Villaverde and Rubio-Ramirez (2005) used sequential importance sampling methods for estimating macroeconomic dynamic stochastic general equilibrium models. The structure of dynamic stochastic general equilibrium models is closely related to that of dynamic discrete choice models. Blevins (2009) used sequential Monte Carlo to allow for serially correlated unobservable state variables in estimating dynamic single agent models, and dynamic games of incomplete information in a revealed preference framework. In a continuous time setting, Nekipelov (2007) developed a flexible indirect inference estimator for continuous time dynamic games in the context of eBay auctions without requiring the complete solution of the dynamic game. This is a novel approach that has potential applications in dynamic oligopolistic competition models. Ackerberg (2009) has developed a method for using importance sampling coupled with a change of variables technique to provide computational gains in estimating certain game theoretic and dynamic discrete choice models that admit a random coefficient presentation.

The purely methodological papers most closely related to the econometric approach used here are Pitt (2010) and Flury and Shephard (2010).

3 Model

In this section we describe the game and introduce some examples.

There are I players, $i = 1, \ldots, I$, who can choose action a_{it} at each time period t. Let $a_t = (a_{1t}, a_{2t}, \ldots, a_{It})$. In an entry game with firms as players, if firm i enters at time t, $a_{it} = 1$; if not, $a_{it} = 0$. While our examples are entry games, we do not require a_{it}, or

any other variable for that matter, to be discrete. All variables in the game can be either discrete or continuous. To reduce the notational burden, we require each variable to be one or the other so that marginal distributions either put mass on a discrete set of points or on a continuum. Again for notational convenience, we will assume that continuous variables do not have atoms. The game is stationary. Time runs from $-\infty$ to ∞. For the observed data time runs from $-T_0$ to T. The state vector is $x_t = (x_{1t}, x_{2t})$. The state vector is observable by all players. We observe the second component x_{2t} only. The game is indexed by a parameter vector θ that is known to the players but not to us. If the state is x_t, then players choose actions a_t according to the probability density function $p(a_t|x_t, \theta)$. Note that this formulation permits randomized strategies. It also includes the case where players are not, in the end, able to follow through and implement the optimal strategy because, e.g., the necessary regulatory approval was not obtained, which would reverse a decision to enter, or a competing firm was acquired, which could reverse a decision not to enter. The transition density $p(x_t|a_{t-1}, x_{t-1}, \theta)$ governs the evolution of the state vector. We observe a_t.

The methods we develop to estimate the game are generally applicable to dynamic games for which there exists a Markovian representation of the latent dynamics and an algorithm to solve the game. The solution method we use for our examples relies on Bellman equations and a sieve representation of the value function and should, therefore, be broadly applicable.

As the analysis is Bayesian, one has some flexibility in handling nonexistence of equilibria. Because data are at hand, one can presume that at an equilibrium must exist to give rise to the data. Therefore one can impose a prior that assigns zero support to any pairing of history and parameters for which the model being estimated does not have an equilibrium. If MCMC is used to estimate the model, as here, imposing this prior is quite easy: one rejects a pair for which an equilibrium does not exist so that it does not get added to the chain.

3.1 A Suite of Entry Games

In the following we will introduce a sequence of oligopolistic entry game models, each more realistic than its predecessor. We will use a separate notation for the games in the suite because the map from the game's notation to the generic notation is not the same for each game in the suite. Toward the end of the discussion of each game within the suite we provide

the map. Within the suite, we follow the convention that a lower case quantity denotes the logarithm of an upper case quantity, e.g., $c_{it} = \log(C_{it})$.

In all games within the suite, firms maximize profits over an infinite horizon. Each time the market opens counts as one time increment. A market opening is defined to be an entry opportunity that becomes available to firms. Since a time period uniquely identifies a market opening, in what follows t is used interchangeably to denote a market opening or the time period associated with it. One could also think of the dynamics arising from evolution of demand, revenues and costs for a particular product as it diffuses through the market over time (see e.g., Ching (2010)). This would lead to two time indices, one for the sequence of product markets opening over time and the other for profits over time within a product market. For simplicity, we abstract from the latter and assume that once a firm enters a market it realizes all the payoffs associated with that product market as a lump sum at the date of entry. We also assume that within the model market time rather than calendar time is what is relevant to discount factors and serial correlation.

The actions available to firm i when market t opens are to enter, $A_{it} = 1$, or not enter $A_{it} = 0$. There are I firms in total so that the number of entrants in market t is given by

$$N_t = \sum_{i=1}^{I} A_{it} \tag{1}$$

The primary source of dynamics is through costs. The evolution of current costs C_{it} for firm i is determined by past entry decisions and random shocks. Entry can increase the cost of an entry next period by, e.g., constraining capacity or it can reduce cost through, e.g., learning. All firms observe each others' costs and hence this is a game of complete information. Log cost is the sum of two components

$$c_{i,t} \quad = \quad c_{u,i,t} + c_{k,i,t}. \tag{2}$$

We assume throughout that $c_{u,i,t}$, cannot be observed; our assumption regarding $c_{k,i,t}$ varies with the model. The first component follows a stationary autoregressive process of order

one; the second accumulates the consequences of past entry decisions:

$$
\begin{aligned}
c_{u,i,t} &= \mu_c + \rho_c \left(c_{u,i,t-1} - \mu_c \right) + \sigma_c e_{it} && (3) \\
c_{k,i,t} &= \rho_a \, c_{k,i,t-1} + \kappa_a A_{i,t-1} && (4) \\
&= \sum_{j=0}^{\infty} \rho_a^j \kappa_a A_{i,t-j-1}.
\end{aligned}
$$

In the above, e_{it} is a normally distributed shock with mean zero and unit variance, σ_c is a scale parameter, κ_a is the immediate impact on cost at market t if there was entry in market $t-1$, μ_c is a location parameter that represents the unconditional mean of the unobservable portion of log cost; ρ_c and ρ_a are autoregressive parameters that determine persistence. We assume that all firms are ex ante identical, with the effects of current decisions on future costs creating heterogeneity between firms. Hence, none of these parameters are firm specific, i.e., indexed by i. Stated differently, heterogeneity arises endogenously in the model depending on the past actions of the firms.

The total (lump sum) revenue to be divided among firms who enter a market at time t is $R_t = \exp(r_t)$ given by

$$
r_t = \mu_r + \sigma_r e_{I+1,t} \, , \tag{5}
$$

where the $e_{I+1,t}$ are normally and independently distributed with mean zero and unit variance. In (5), μ_r is a location parameter that reflects the average total revenue for all the firms across all market opportunities, and σ_r is a scale parameter. A firm's total discounted profit at time t is

$$
\sum_{j=0}^{\infty} \beta^j A_{i,t+j} \left(R_{t+j}/N_{t+j} - C_{i,t+j} \right) , \tag{6}
$$

where β is the discount factor, $0 < \beta < 1$. A firm's objective is to maximize the present discounted value of its profit at each time period t taking as given the equilibrium action profiles of other firms.

In the following we first describe a baseline model in which the entry decisions of the firms are perfectly predicted by the firms. This model, while theoretically appealing, is not particularly amenable to econometric specification and estimation. This is because of the discrete nature of the likelihood and the observed entry actions. The likelihood function will be identically zero when the profile of firm entry actions does not coincide with the predicted

equilibrium profile of entry decisions at the given parameter values. Therefore, the likelihood function is only nondegenerate when the predicted outcome profile matches perfectly with the observed actions. Running this model through a nonlinear filter to obtain the likelihood of the observed variables can run into a degeneracy issue where all particles are killed.

It is possible to introduce a measurement equation that compounds the predicted outcomes with a misclassification error. However, we argue that a priori such a model does not appear to be suitable for empirical interpretation. We will present two alternative models that avoid this degeneracy issue, after the base model is presented. One has computational advantages while the other has more theoretical appeal.

Throughout we follow the standard convention of suppressing subscripts to indicate grouping, e.g.,

$$c_{i,t} = (c_{u,i,t}, c_{k,i,t})$$

$$c_t = (c_{1t}, c_{2t}, \ldots, c_{I,t}).$$

When these groupings are to be interpreted as vectors, we follow the convention that the leftmost index varies fastest, e.g.

$$c_t = (c_{u,1,t}, c_{k,1,t}, c_{u,2,t}, c_{k,2,t}, \ldots, c_{u,I,t}, c_{k,I,t}).$$

3.2 Baseline Model

In this model, the firms can perfectly predict the observed entry profiles in equilibrium but the econometrician observes the profiles with error.

The Bellman equation for the choice specific value function, $V_i(A_{i,t}, A_{-i,t}, C_{i,t}, C_{-i,t}, R_t)$, for firm i's dynamic problem at time t is

$$\begin{aligned} & V_i(A_{i,t}, A_{-i,t}, C_{i,t}, C_{-i,t}, R_t) \qquad (7) \\ &= A_{it}\,(R_t/N_t - C_{it}) \\ &\quad + \beta\,\mathcal{E}\Big[V_i(A^E_{i,t+1}, A^E_{-i,t+1}, C_{i,t+1}, C_{-i,t+1}, R_{t+1}) \,|\, A_{i,t}, A_{-i,t}, C_{i,t}, C_{-i,t}, R_t,\Big]\,, \end{aligned}$$

where, as is conventional, $-i$ represents the other players. The choice specific value function represents the sum of current and future payoffs to firm i from a choice $A_{i,t}$ at time t explicitly conditioning on the choices that would be made by other firms $A_{-i,t}$ at time t and

with the expectation that firm i and the other firms would be making equilibrium choices, $A^E_{i,t+1}, A^E_{-i,t+1}$, respectively, from period $t+1$ onwards conditional on their current choices. The expectations operator here is over the distribution of the state variables in time period $t+1$ conditional on the realization of the time t state variables and the action profile at time t. Therefore $V_i(A_{i,t}, A_{-i,t}, C_{i,t}, C_{-i,t}, R_t)$ is the payoff to firm i at stage t of the game.

A stationary pure strategy Markov perfect equilibrium of the dynamic game is defined by a best response strategy profile $(A^E_{i,t},\ A^E_{-i,t})$ that satisfies

$$V_i(A^E_{i,t}, A^E_{-i,t}, C_{i,t}, C_{-i,t}, R_t) \geq V_i(A_{i,t}, A^E_{-i,t}, C_{i,t}, C_{-i,t}, R_t) \quad \forall\, i, t. \tag{8}$$

This is a game of complete information. Hence, if the state, which includes the current cost vector of all firms $(C_{i,t}, C_{-i,t})$ and total revenue (R_t), is known, then the equilibrium is known. Therefore, an ex ante value function can be computed from the choice specific value function

$$V_i(C_{i,t}, C_{-i,t}, R_t) = V_i(A^E_{i,t}, A^E_{-i,t}, C_{i,t}, C_{-i,t}, R_t). \tag{9}$$

The ex ante value function satisfies the Bellman equation

$$\begin{aligned} &V_i(C_{it}, C_{-i,t}, R_t) \qquad\qquad (10)\\ &\quad = \ A^E_{it}\left(R_t/N^E_t - C_{it}\right) + \beta\, \mathcal{E}\Big[V_i(C_{i,t+1}, C_{-i,t+1}, R_{t+1}) \,|\, A^E_{i,t}, A^E_{-i,t}, C_{i,t}, C_{-i,t}, R_t\Big], \end{aligned}$$

where N^E_t is the number of firms that enter in equilibrium, which can be computed using equation (1), i.e., $N^E_t = \sum_{i=1}^{I} A^E_{it}$. Equation (10) is different from the Bellman equation associated with the choice specific value function (equation (7)) as it represents the sum of current and future payoffs to firm i from an optimal equilibrium choice $A^E_{i,t}$ at time t explicitly conditioning on the equilibrium choices that would be made by other firms $A^E_{-i,t}$ at time t, and with the expectation that all firms would be making equilibrium choices from period $t+1$ onwards.

A comprehensive discussion of results for existence of equilibria in Markovian games is provided by Dutta and Sundaram (1998). More results on existence of equilibria in dynamic oligopolistic models are to be found in Doraszelski and Satterthwaite (2007). When the state space can only take on a finite set of values, Theorem 3.1 of Dutta and Sundaram (1998) implies that the baseline game just described has a stationary Markov perfect equilibrium in

mixed strategies. Parthasarathy (1973) showed that this support condition can be relaxed to include a state space with countable values. The regularity conditions of Theorem 5.1 of Dutta and Sundaram (1998) come closer to the problem as we have posed it, notably that the revenue and cost do not have to be discrete but they do need to be bounded. The equilibrium strategy profiles provided by Theorem 5.1 may depend on periods t and $t-1$ of the state vector.

We could modify our problem to meet the requirements of Theorem 3.1 that the state space be finite and countable. However we rely on Theorem 5.1 instead as we do not have trouble computing pure strategy equilibria for the problem as posed with a continuous state space. Theorem 3.1 is of interest to us because its proof relies on a dynamic programming approach that motivates our computational strategy, discussed below in Section 5 (see also Rust (2006) for a discussion of a similar computation strategy). We find that we can compute pure strategy equilibria that depend only on period t of the state vector, and hence automatically satisfy the regularity conditions of Theorem 5.1. While the results described above imply that a slightly modified version of the game proposed by us has equilibria, we rely mostly on the fact that we have no difficulty computing equilibria. In fact the key hurdle we face is not the lack of existence of equilibria but instead multiplicity of equilibria. In Sections 5 and 6 we discuss how we resolve this problem.

3.3 Measurement Error Model

To avoid the degeneracy issue associated with likelihood inference, one can assume that the predicted equilibrium outcome $A_{it}^E, \forall i$ is not observed by the econometrician. Instead, the observed entry decision A_{it} of firm i is a Bernoulli random variable taking value A_{it}^E with probability p_a, and taking value $1 - A_{it}^E$ with probability $q_A = 1 - p_A$.

The baseline model with measurement error relates to the generic model as follows. The parameters of the model are

$$\theta = (\mu_c, \rho_c, \sigma_c, \mu_r, \sigma_r, \rho_a, \kappa_a, \beta, p_a). \tag{11}$$

The cost equations become

$$c_{u,i,t} = \mu_c + \rho_c \left(c_{u,i,t-1} - \mu_c\right) + \sigma_c e_{it} \tag{12}$$

$$c_{k,i,t} = \rho_a \, c_{k,i,t-1} + \kappa_a A^E_{i,t-1} \tag{13}$$

The first cost component can not be observed by the econometrician. The equilibrium is given by the deterministic function

$$A^E_t = S_B(c_{u,t}, c_{k,t}, r_t, \theta). \tag{14}$$

The measurement error density described above can be written

$$p(A_t \,|A^E_t, \theta) = \prod_{i=1}^{I} (p_a)^{\delta(A_{it}=A^E_{it})} (1-p_a)^{1-\delta(A_{it}=A^E_{it})}. \tag{15}$$

where $0 < p_a < 1$, $\delta(a = b) = 1$ if $a = b$ (element by element when a and b are vectors as below). The map to the variables and densities of the generic model is

$$\begin{aligned}
x_{1t} &= c_{ut} \\
x_{2t} &= (c_{kt}, r_t) \\
a_t &= A_t \\
p(x_t|a_{t-1}, x_{t-1}, \theta) &= n[c_{ut} \,|\, \mu_c \mathbf{1} + \rho_c(c_{u,t-1} - \mu_c \mathbf{1}), \sigma_c^2 I] \\
&\quad \times \delta[c_{k,t} = c_{k,t-1} + \kappa_a S_B(c_{u,t-1}, c_{k,t-1}, r_{t-1}, \theta)] \\
&\quad \times n(r_t \,|\, \mu_r, \sigma_r^2) \\
p(x_{1,t}|a_{t-1}, x_{t-1}, \theta) &= n[c_{ut} \,|\, \mu_c \mathbf{1} + \rho_c(c_{u,t-1} - \mu_c \mathbf{1}), \sigma_c^2 I] \\
p(a_t \,|\, x_t, \theta) &= p[A_t \,|\, S_B(c_{u,t}, c_{k,t}, r_t, \theta), \theta]
\end{aligned}$$

In the measurement error model, the state variable $x_t = (c_t, A^E_t, r_t)$ evolves autonomously. The observed action profile in the data is a misclassified version of A^E_t. While this model is internally coherent, it is difficult to believe that such nontransitory misclassification errors will persist in a given time series data set of reasonable sample size. Who is making the mistake? Presumably, firms should be able to observe the realized action profile A_t as well as econometricians. It is difficult to believe that firms always perceive A^E_t as the realized market outcome even when it systematically differs from A_t. Conversely, it is equally difficult to believe that the firm observes correctly but that data are perpetually in error regarding something as simple as determining if a firm enters a market or not.

3.4 The Boundedly Rational Model

An alternative model is to assume that the discrepancy between the observed A_t and the equilibrium prediction A_t^E arises structurally from the combination of optimization error and ex post uncertainty rather than merely as a statistical misclassification error. Given the small number of firms in most empirical models and the short horizon of the time series data set, it is more convincing to believe that the entry decisions are recorded correctly in the data. We adopt such a view in the second model. However, the observed action profile A_t can still differ from the equilibrium prediction A_t^E computed by the firms. Consider the generic drug market. The entry decision of firm i is coded as $A_{it}^E = 1$ if the firm desires to submit an application to the FDA. If the A_t in the data set records the approval by FDA instead of the submission of proposal by the firms, each proposal carries a small probability of being rejected by the FDA. On the one hand, firms collectively decide on the equilibrium A_t^E. On the other hand, the econometrician does observe the ex post realization of A_t at the end of each period without misclassification error. Therefore the latent cost state variable will involve according to A_t instead of A_t^E. The cost evolution is observed by all the firms. Conditional on observing the costs, whether the firms observe A_t or not is not important. However, it appears reasonable to assume that they do.

This model relates to the generic model as follows. The parameters of the model are given by (11). The cost equations become

$$c_{u,i,t} = \mu_c + \rho_c \left(c_{u,i,t-1} - \mu_c\right) + \sigma_c e_{it} \tag{16}$$

$$c_{k,i,t} = \rho_a \, c_{k,i,t-1} + \kappa_a A_{i,t-1} \tag{17}$$

The first cost component cannot be observed by the econometrician; the second can be observed. The deterministic function giving the equilibrium is (14). The density for the discrepancy between outcomes and intentions is (15). The map to the variables and densities

of the generic model is

$$
\begin{aligned}
x_{1t} &= c_{ut} \\
x_{2t} &= (c_{kt}, r_t) \\
a_t &= A_t \\
p(x_t \,|\, a_{t-1}, x_{t-1}, \theta) &= n[c_{ut} \,|\, \mu_c \mathbf{1} + \rho_c(c_{u,t-1} - \mu_c \mathbf{1}), \sigma_c^2 I] \\
&\quad \times \delta[c_{k,t} = c_{k,t-1} + \kappa_a a_{t-1}, \theta)] \\
&\quad \times n(r_t \,|\, \mu_r, \sigma_r^2) \qquad (18) \\
p(x_{1,t} \,|\, a_{t-1}, x_{t-1}, \theta) &= n[c_{ut} \,|\, \mu_c \mathbf{1} + \rho_c(c_{u,t-1} - \mu_c \mathbf{1}), \sigma_c^2 I] \qquad (19) \\
p(a_t \,|\, x_t, \theta) &= p[A_t \,|\, S_B(c_{u,t}, c_{k,t}, r_t, \theta), \theta] \qquad (20) \\
p(x_{1,t} \,|\, \theta) &= n[c_{ut} \,|\, \mu_c, \frac{\sigma_c^2}{(1-\rho_c^2)} I] \qquad (21) \\
p(a_t, x_{2t} \,|\, x_{1,t}, \theta) &= p[A_t \,|\, S_B(c_{u,t}, c_{k,t}, r_t, \theta), \theta]\, n(r_t \,|\, \mu_r, \sigma_r^2) \qquad (22)
\end{aligned}
$$

In equations (20) and (22),

$$
p(A_t | A_t^E, \theta) = \prod_{i=1}^{I} p_A^{\delta(A_{i,t} = A_{i,t}^E)} (1 - p_A)^{1 - \delta(A_{i,t} = A_{i,t}^E)}.
$$

Hidden within $S_B(c_{u,t}, c_{k,t}, r_t, \theta)$ of equations (20) and (22) are equations (7) to (10), which describe the computation of the value function and the search algorithm for the Nash equilibrium used to compute the predicted action profiles A_t^E from each c_t and r_t. When firms configure the equilibrium action profile A_t^E according to the Nash condition (8), they act as if they are unaware of the discrepancy between the ex post A_t and A_t^E. The discrepancy $A_t - A_t^E$ appears to the firms as pure ex post random shocks. In other words, at the time of the equilibrium determination firm i does not know either $A_{it} - A_{it}^E$ or any of the $A_{jt} - A_{jt}^E$ for each $j \neq i$. Since the value function iterations in (7) and (10) make use of a single action profile instead of a mixture of action profiles, the computation in (7) and (10) implicitly assumes that firms do not figure in the randomness in $A_t - A_t^E$ when they calculate their value functions in the forward-looking manner. In this sense, firms are *boundedly rational*. While they have perfect observations of the costs when the equilibrium is determined, and presumably are able to observe the history of differences between A_t and A_t^E, they do not

account for this difference in their value function iteration and assume that the future equilibrium path of A_t^E can be perfectly executed in their forward perceptions. However, the evolution of the cost process, as given in the transition equations in (3) and (4), is, in fact, actually determined by the historically realized values of the observed A_t's.

3.5 The Fully Rational Model

Taking the transition equations and the measurement equations in the previous subsection as given, it is possible to amend the value function iterations and equilibrium condition in (7) to (10) to avoid imposing the bounded rationality condition on the firms.

In order for firms to account for the statistical uncertainty of $A_t - A_t^E$ when they decide on their equilibrium behavior, we need to modify the model in (7). This requires redefinition of the choice specific value function (7) as follows:

$$
\begin{aligned}
&V_i^f(a_{i,t}, a_{-i,t}, C_{i,t}, C_{-i,t}, R_t) \\
&\quad = \sum_{l_1=0}^{1} p_A^{\delta(l_1=a_{1t})}(q_A)^{1-\delta(l_1=a_{1t})} \cdots \sum_{l_I=0}^{1} p_A^{\delta(l_I=a_{It})}(q_A)^{1-\delta(l_I=a_{It})} \Bigg\{ l_i \left(\frac{R}{\sum_{j=1}^{I} l_j} - C_{it} \right) \qquad (7') \\
&\qquad + \beta\, \mathcal{E}\Big[V_i^f(A_{i,t+1}^E, A_{-i,t+1}^E, C_{i,t+1}, C_{-i,t+1}, R_{t+1}) \,|\, L_{i,t}, L_{-i,t}, C_{i,t}, C_{-i,t}, R_t \Big] \Bigg\},
\end{aligned}
$$

where $L_{i,t} = l_i$ and $L_{-i,t}$ is $(l_1, \ldots, l_I)$ with l_i deleted. Note that $L_{i,t}$ and $L_{-i,t}$ affect the conditioning information via equations (3) and (4), both for $j \in L_{i,t}$ and $j \in L_{-i,t}$; l_j, like A_j, is one for entry and zero otherwise.

Equations (8) and (9) stay unchanged with this new definition of choice specific function in (7′). A Bellman equation is implied by (7′), (8), and (9) that is easy to code directly from (7′), (8), and (9) without deriving an explicit expression. The firms use the implied Bellman equation to make correct forecasts of the future value function taking into account the uncertainty that the ex post realized action profile A_t might be different from the equilibrium action profile A_t^E. A fully rational stationary pure strategy equilibrium is defined by a best response strategy profile $(A_{i,t}^E, A_{-i,t}^E)$ that satisfies (8) with V^f replacing V. When it is necessary to distinguish a fully rational equilibrium, we use the notation A_t^{FR}.

To summarize, the only difference between the fully rational model and the boundedly rational model is in the deterministic function used to compute the equilibrium. For the

boundedly rational model it is (14). For the fully rational model we write

$$A_t^{FR} = S_R(c_{u,t}, c_{k,t}, r_t, \theta). \tag{23}$$

All that changes in the map to the variables and densities of the generic model are equations (20) and (22), where S_R replaces S_B.

There is yet another possible model with partial bounded rationality of the firms. In this model we replace the choice specific value function with

$$\begin{aligned} &V_i(A_t, C_{i,t}, C_{-i,t}, R_t) \\ &\quad = \sum_{l_1=0}^{1} p_a^{l_1=a_{1t}} q_a^{l_1\neq a_{1t}} \cdots \sum_{l_I=0}^{1} p_a^{l_I=a_{It}} q_a^{l_I\neq a_{It}} l_i \left(R_t / \left(\sum_{j=1}^{I} l_j \right) - C_{it} \right) \\ &\qquad + \beta\, \mathcal{E}\Big[V_i(A^E_{i,t+1}, A^E_{-i,t+1}, C_{i,t+1}, C_{-i,t+1}, R_{t+1}) \,|\, a_t, C_{i,t}, C_{-i,t}, R_t \Big]. \end{aligned} \tag{7''}$$

In such a model, firms are fully rational in computing the current period expected payoff, but use a partial solution method to compute the continuation value function in which they do not account for the discrepancy between the equilibrium predicted action profile and the ex post realized action profile. We will not pursue this alternative model in the rest of paper.

4 Solving the Model

In this section we describe a method for computing the equilibrium of a dynamic game of complete information given the observed and latent state variables and a set of parameter values. Because we consider infinite horizon models, we look for a stationary Markov perfect equilibrium which entails finding the fixed point of Bellman equations. In Section 3 we described three alternative models. The solution method for the baseline model and the bounded rationality model are identical. These two models differ only in the evolution of the latent cost state variable, which only becomes relevant in the likelihood function and parameter estimation sections, which are Sections 5 and 6. The solution for the fully rational model does differ from the baseline model in the definition of the choice specific value function. Therefore in the following we describe the solution method for the baseline model first, and then we discuss how it can be augmented for the fully rational model.

4.1 The Baseline Model

Recall that the entry decisions of all $i = 1, \ldots, I$ firms for a market opening at time t, i.e., the strategy profile of the dynamic game, are denoted by $A_t = (A_{1t}, ..., A_{It})$. As discussed in Section 3, the strategy profile A_t at time t of the dynamic game is a function of the current period state variables $(C_{1t}, ..., C_{It})$ and R_t. The vector of the log of the state variables at time t is

$$s_t = (c_{u1t}, ..., c_{uIt}, c_{k1t}, ..., c_{kIt}, r_t) . \tag{24}$$

In particular, equations (7) and (10) can be expressed in terms of s_t using $C_{uit} = \exp(s_{it})$ and $C_{kit} = \exp(s_{I+i,t})$ for $i = 1, \ldots, I$ and $R_t = \exp(s_{2I+1,t})$. We describe the solution algorithm for a given parameter vector θ and a given state s_t at time t.

We begin by defining a grid on the state space which determines a set of $(2I + 1)$-dimensional hyper-cubes. For each hyper-cube we use its centroid as its index or key K. A state s_t within hyper-cube can be mapped to its key K.[1] Let the vector $V_K(s_t)$ have as its elements the ex ante value functions $V_{i,K}(s_t)$, i.e., $V_K(s_t) = (V_{1,K}(s_t), \ldots, V_{I,K}(s_t))$ (see equations (9) and (10). To each K associate a vector b_K of length I and a matrix B_K of dimension I by $I + 1$. A given state point s_t is mapped to its key K and the value function at state s_t is represented by the affine function $V_K(s_t) = b_K + (B_K)s_t$.[2] A value function $V_K(s_t)$ whose elements satisfy equation (10) is denoted $V_K^*(s_t) = b_K^* + (B_K^*)s_t$.

The game is solved as follows:

1. Given a state point s, get the key K that corresponds to it. (We suppress the subscript t for notational convenience.)[3]

2. Check whether the fixed point $V_K^*(s)$ of the Bellman equations (10) at this key has

[1] Grid increments are chosen to be fractional powers of two so that the key has an exact machine representation. This facilitates efficient computation through compact storage of objects indexed by the key. The rounding rules of the machine resolve which key a state on a grid boundary gets mapped to, although lying on a boundary is a probability zero event in principle. The entire grid itself is never computed because all we require is the mapping $s \mapsto K$, which is determined by the increments and is easily computed as needed.

[2] Keane and Wolpin (1997) adopt a similar approach for a single agent model. Our approach differs from Keane and Wolpin (1997) in that we let the coefficients of the regression depend on the state variables, specifically the key K, whereas Keane and Wolpin (1997) use an OLS regression whose coefficients are not state specific. Thus, our value function, unlike theirs, need not be continuous. Our value function can be thought of as an approximation by a local linear function.

[3] In fact, because it is a stationary game, the subscript t does not really matter

already been computed, i.e., whether the (b_K^*, B_K^*) for the K that corresponds to s has been computed. If not, then use the following steps to compute it.

3. Start with an initial guess of the ex ante value function $V_K^{(0)}(s)$. An initial guess of the value function is represented by the coefficients $(b_K^{(0)}, B_K^{(0)})$ being set to 0.

4. Obtain a set of points s_j, $j = 1, \ldots, J$, that are centered around K. The objective now is to obtain the ex ante value functions associated with these points to use in a regression to recompute (or update) the the coefficients $(b_K^{(0)}, B_K^{(0)})$.

5. Ex ante value functions are evaluated at best response strategies. In order to compute these we must, for each s_j, compute the choice specific value function (7) at as many strategy profiles A as are required to determine whether or not the equilibrium condition in equation (8) is satisfied. In this process we need to take expectations to compute the continuation value $\beta\,\mathcal{E}\left[V_{K,i}^{(0)}(s_{t+1}) \,|\, A_{i,t}, A_{-i,t}, C_{i,t}, C_{-i,t}, R_t,\right]$ that appears in equation (7), where we have used equation (9) to express equation (7) in terms of $V_K^{(0)}(s)$. To compute expectations over the conditional distribution of the random components of next period state variables, we use Gauss-Hermite quadrature. To do this, we obtain another set of points centered around each s_j, i.e., s_{j_l}, $l = 1, \ldots, L$. These points are the abscissae of the Gauss-Hermite quadrature rule which are located relative to s_j but shifted by the actions A under consideration to account for the dynamic effects of current actions on future costs (see equation (4)). Expectations are computed using a weighted sum of the value function evaluated at the abscissae (more details are provided below).

6. We can now compute the continuation value at s_j for each candidate strategy A. We compute the best response strategy profile A_j^E corresponding to s_j by checking the Nash equilibrium condition (8). As just described, the choice specific value function evaluated at $(A_i^E,\ s_j)$ is computed using $V_K^{(0)}(s)$ and equation (7), and denoted by $V_K^{(1)}(A^E, s_j) = (V_{1,K}^{(1)}(A^E, s_j), \ldots, V_{I,K}^{(1)}(A^E, s_j))$.

7. Next we use the "data" $(V_K^{(1)}(A^E, s_j), s_j)_{j=1}^J$ to update the ex ante value function to $V_K^{(1)}(s_j)$. This is done by updating the coefficients of its affine representation to

$(b_K^{(1)}, B_K^{(1)})$ via a multivariate regression on this "data" (as described in detail below).[4]

8. We iterate (go back to step 5) over the ex ante value functions $V_{i,K}^{(0)}(s), V_{i,K}^{(1)}(s), \ldots$ by finding a new equilibrium strategy profile A^E for each s_j until convergence is achieved for the coefficients $(b_K^{(0)}, B_K^{(0)})$, $(b_K^{(1)}, B_K^{(1)}), \ldots, (b_K^{(*)}, B_K^{(*)})$. This gives us $V_K^*(s) = b_K^* + (B_K^*)s$ for every s that maps to key K.

To summarize, the process of solving for the equilibrium begins with a conjecture $(b_K^{(l)} = 0, B_K^{(l)} = 0)$ for the linear approximation of the value functions at a given state at iteration $l = 0$. These guesses are then used in computing the choice specific value functions at iteration $l + 1$ using equation (7). This computation involves taking expectations over the conditional distribution of the future state variables, which is accomplished using Gaussian-Hermite quadrature. Once we have the choice specific value functions we compute the equilibrium strategy profile at iteration $l + 1$ using equation (8). The best response strategy profile at iteration $l + 1$ is then used to compute the iteration $l + 1$ ex ante value functions via a regression that can be viewed as iterating equation (10). The iteration $l + 1$ ex ante value functions are then used to compute the iteration $l + 2$ choice specific value functions using equation (7), and the entire procedure is repeated till a fixed point of equation (10) is obtained. This iterative procedure solves the dynamic game. We next provide additional details about the steps of the algorithm described above to solve the model.

To describe the Gauss-Hermite quadrature procedure used in Step 5, note that if one conditions upon s_t and A_t, then a subset of the elements of s_{t+1} are independently normally distributed with means $\mu_i = \mu_c + \rho_c(c_{it} - \mu_c)$ for the first I elements, mean $\mu_{2I+1} = \mu_R$ for the last element, standard deviations $\sigma_i = \sigma_c$ for the first I elements, and standard deviation $\sigma_{2I+1} = \sigma_R$ for the last. The other I elements of s_{t+1}, $(s_{t+1,I+1}, \ldots, s_{t+1,2I})$ are deterministic functions of s_t and A_{it}. Computing a conditional expectation of functions of the form $f(s_{t+1})$ given $(A_t,\ s_t)$ such as appear in equations (7) and (10) is now a matter of integrating with respect to a normal distribution with these means and variances which can be done by a Gauss-Hermite quadrature rule that has been subjected to location and scale transformations. The weights w_j and abscissae x_j for Gauss-Hermite quadrature may

[4] $V_K^{(1)}(A^E, s_j)$ will not equal $V_K^{(1)}(s_j)$ because the former is "data" and the later is a regression prediction.

be obtained from tables such as Abramowitz and Stegun (1964) or by direct computation using algorithms such as Golub and Welsch (1969) as updated in Golub (1973). To integrate with respect to $s_{i,t+1}$ conditional upon A_t and s_t the abscissae are transformed to $\tilde{s}^j_{t+1,i} = \mu_i + \sqrt{2}\sigma_i x_j$, and the weights are transformed to $\tilde{w}_j = w_j/\sqrt{\pi}$, where $\pi = 3.142$.[5] Then, using a $2L+1$ rule,

$$\mathcal{E}[f(s_{t+1}) \,|\, A_t, s_t] \approx$$
$$\sum_{j_1=-L}^{L} \cdots \sum_{j_I=-L}^{L} \sum_{j_{2I+1}=-L}^{L} f(\tilde{s}^{j_1}_{t+1,1}, \cdots, \tilde{s}^{j_I}_{t+1,I}, \tilde{s}^{j_{2I+1}}_{t+1,2I+1}, s_{t+1,I+1}, \ldots, s_{t+1,2I}) \tilde{w}_{j_1} \cdots \tilde{w}_{j_I} \tilde{w}_{j_{2I+1}}.$$

If, for example, there are three firms and a three point quadrature rule is used, then

$$\mathcal{E}[f(s_{t+1}) \,|\, A_t, s_t]$$
$$\approx \sum_{j_1=-1}^{1} \sum_{j_2=-1}^{1} \sum_{j_3=-1}^{1} \sum_{j_7=-1}^{1} f(\tilde{s}^{j_1}_{t+1,1}, \tilde{s}^{j_2}_{t+1,2}, \tilde{s}^{j_3}_{t+1,3}, \tilde{s}^{j_7}_{t+1,7}, s_{t+1,4}, \ldots, s_{t+1,6}) \tilde{w}_{j_1} \tilde{w}_{j_2} \tilde{w}_{j_3} \tilde{w}_{j_7}.$$

We use three point rules throughout. A three point rule will integrate a polynomial in s_{t+1} up to degree five exactly.[6][7]

Step 7 involves updating the ex ante value function using a regression. We next describe how we do this. As stated above, we have a grid over the state space whose boundaries are fractional powers of two over the state space.[8] We approximate the value function $V(s_t)$ by a locally indexed affine representation as described above. For the the grid increments that determine the index of hyper-cubes we tried a range of values from 4 to 16 times the standard deviation of the state variables rounded to a nearby fractional power of two to scale the grid appropriately. The results are effectively the same. Hence in estimating the

[5] These transformations arise because a Hermite rule integrates $\int_{-\infty}^{\infty} f(x) \exp(-x^2) dx$. Hence we need to do a change of variables to get our integral $\int_{-\infty}^{\infty} g(\sigma z + \mu)(1/\sqrt{2\pi}) \exp(-0.5z^2) dz$ to be of that form. A change of variables puts the equation in the line above in the form $\int_{-\infty}^{\infty} g(\sqrt{2}\sigma x + \mu)(1/\sqrt{\pi}) \exp(-x^2) dx$, which is where the expressions for $\tilde{s}_{t+1,i}$ and $\tilde{w}_i$ come from.

[6] If the $\tilde{s}_{t+1}$ cross a grid boundary when computing (7) in Step 5, we do not recompute K because this would create an impossible circularity due to the fact that the value function at the new K may not yet be available. Our grid increments are large relative to the scatter of abscissae of the quadrature rule so that crossing a boundary will be a rare event, if it happens at all.

[7] Positivity is enforced by setting the value function to zero if a negative value is computed. If this does not happen at any quadrature point, which is easily detected by checking the most extreme point of the quadrature rule, then quadrature is not necessary and the conditional mean can be used as the integral.

[8] Recall that grid increments are chosen to be fractional powers of two so that the key has an exact machine representation. This facilitates efficient computation through compact storage of objects indexed by the key.

model we set the grid increments at 16 times the standard deviation of the state variables.[9] We compute the coefficients b_K and B_K as follows. They are first initialized to zero. We then generate a set of abscissae $\{s_j\}$ clustered about K and solve the game with payoffs (7) to get corresponding equilibria $\{A_j^E\}$. We substitute the (A_j^E, s_j) pairs into equation (7) to get $\{V(A_j^E, s_j)\}_{j=1}^J$. Using the pairs $\{(V(A_j^E, s_j), s_j\}$ as data, we compute b_K and B_K by multivariate least squares. We repeat until the b_K and B_K stabilize. We have found that approximately twenty iterations suffice for three firms and thirty for four firms.[10] The easiest way to get a cluster of points $\{s_j\}$ about a key is to use abscissae from the quadrature rule described above with s set to K and A set to zero. However, one must jiggle the points so that no two firms have exactly the same cost (see next paragraph for the reason for this). Of importance in reducing computational effort is to avoid recomputing the payoff (equation (7)) when checking equilibrium condition (8). Our strategy is to (temporarily) store payoff vectors indexed by A and check for previously computed payoffs before computing new ones in checking condition (8).

There will, at times, be multiple equilibria in solving the game. We therefore adopt an equilibrium selection rule as follows. Multiple equilibria usually take the form of a situation where one or another firm can profitably enter but if both enter they both will incur losses whereas if neither enters then one of them would have an incentive to deviate. We resolve this situation by assuming an explicit equilibrium selection rule. We pick the equilibrium with the lowest total cost. This idea is similar to that used by Berry (1992) and Scott-Morton (1999). That is, the strategy profiles A_t are ordered by increasing aggregate cost, $C = \sum_{i=1}^I A_{it} C_{it}$ and the first A_t that satisfies the equilibrium condition (8) is accepted as the solution. Note that our distributional assumptions on s_t guarantee that no two C can be equal so that this ordering of the A_t is unique. Moreover, none of the C_{it} can equal one another and when that is true failure to compute an equilibrium for a given θ and cost history is extremely rare. At worst all that happens is a few particles in the particle filter are

[9] The set of keys that actually get visited in any MCMC repetition is about the same for grid increments ranging from 4 to 16 times the standard deviation of the state variables in our data. For a three firm game the number of hyper-cubes that actually are visited in any one repetition is about six.

[10] An alternative is to apply a modified Howard acceleration strategy as described in Kuhn (2006); see also Rust (2006) and Howard (1960). The idea is simple: The solution $\{A_t^E\}$ of the game with payoffs (7) will not change much, if at all, for small changes in the value function $V(s)$. Therefore, rather than recompute the solution at every step of the (b_K, B_k) iterations, one can reuse a solution for a few steps.

lost (but replaced at the resampling step). The situation where all particles are lost thereby causing an MCMC proposal to be rejected never occurs.

4.2 The Fully Rational Model

The solution method in the fully rational model is only different from that in the baseline model and the bounded rationality model regarding how the conjectures of the ex ante value functions are being used to compute the choice specific value functions at each iteration using equation (7′) instead of using equation (7). In addition to taking expectations over the conditional distribution of future state variables using Gaussian-Hermite quadrature for each action profile, the computation in the fully rational model also needs to average over all the possible action profiles using the ex post error probabilities p_A and q_A according to the candidate member of the equilibrium action profile.

More specifically, only step 5 of the solution method in the baseline model needs to be modified for the fully rational model. For each s_j, the choice specific value functions (7′) instead of (7) should be computed at as many strategy profiles A as needed to seek an equilibrium that satisfies condition (8). Each of the terms within the curly bracket in (7′) are computed exactly as in the baseline model, for each combination of $A_{it}(l)$ ranging over all possible i and l. However, instead of calculating them only as needed to compute the equilibrium as in the base model, all the terms are now precomputed and stored prior to calculating the weighted sum in (7′). Once these values are precomputed and stored, evaluating the left hand side of (7′) for each candidate equilibrium action profile A_t only requires taking a weighted sum of the stored values, where the weights obviously depend on the action profile A_t.

The issue of multiple equilibria is handled in the same way as in the baseline model. The candidate equilibrium action profiles A_t are pre-sorted on ascending order of total costs. We start search for the Nash equilibrium based on condition (8) from the lowest cost action profile, and stop once the first equilibrium is found.

5 Likelihood Computation

In Section 6 we use MCMC to compute the posterior. If one has an unbiased estimator of the likelihood the posterior is computed exactly.[11] In this section we derive an unbiased particle filter estimate of the likelihood for a Markov process with partially observed state and endogenous feedback that is general enough to accommodate the generic game described at the beginning of Section 3. Because we only require unbiasedness, our regularity conditions are quite weak – much weaker than is standard in the particle filter literature.[12] While the result does not require that the number of particles tend to infinity, the number of particles does affect the rejection rate of the MCMC chain so that the number of particles, like the scale of the proposal density, becomes a tuning parameter of the chain that has to be adjusted. For the examples in Section 7, the requisite number of particles is small.

The essentials of the generic game of Section 3 relative to the requirements of filtering are as follows. The state vector is

$$x_t = (x_{1t}, x_{2t}), \tag{25}$$

where x_{1t} is not observed and x_{2t} is observed. The observation (or measurement) density is

$$p(a_t \,|\, x_t, \theta). \tag{26}$$

The transition density is denoted by

$$p(x_t \,|\, a_{t-1}, x_{t-1}, \theta). \tag{27}$$

Its marginal is

$$p(x_{1t}|a_{t-1}, x_{t-1}, \theta). \tag{28}$$

The stationary density is denoted by

$$p(x_{1t} \,|\, \theta). \tag{29}$$

ASSUMPTION 1 We assume that we can draw from (28) and (29). As to the latter, one way to draw a sample of size N from (29) is to simulate the game and set $x_1^{(k)} = x_{1,\tau+M*k}$

[11] See e.g., Flury and Shephard (2010) and Pitt (2010).

[12] See, e.g. Andrieu, Douced, and Holenstein (2010) and the references therein.

for $k = 1, \ldots, N$ for some τ past the point where transients have died off and some M large enough that the $x_1^{(k)}$ are nearly serially uncorrelated. We do not actually need independence so that one could set $M = 1$ if computational cost is an issue. We can draw from (28) by drawing from (27) and discarding x_{2t}. We assume that there is either an analytic expression or algorithm to compute (26) and (27). We assume the same for (28) but if this is difficult some other importance sampler can be substituted as discussed in Subsection 5.2.

For the examples in Section 3, which generalize an entry game taken from an application, these conditions are met, the analytic expressions are simple, and draws straightforward. In particular, simulation from (28) and (29) is by means of the normal autoregressive process with transition density (19) and stationary density (21). The observation density (26) is binomial with density (20); its evaluation requires solving the dynamic game using the algorithm described in Section 4. Evaluation of two normal densities in (18) gives (27).

In the Bayesian paradigm, θ is random and $\left\{\{a_t, x_t\}_{t=-T_0}^{\infty}, \theta\right\}$ are defined on a common probability space. Let $\mathcal{F}_t = \sigma\left\{\{a_s, x_{2s}\}_{s=-T_0}^{t}, \theta\right\}$. The elements of a_t and x_t may be either real, without atoms, or discrete. No generality is lost by presuming that the discrete elements are positive integers. Let z denote a generic vector some of whose coordinates are real numbers and the others positive integers. Let $\lambda(z)$ denote a product measure whose marginals are either counting measure on the positive integers or Lebesgue ordered as is appropriate to define an integral of the form $\int g(z)\, d\lambda(z)$. We adopt the convention that when a conditional density is represented as the ratio of a joint density to a marginal density, the ratio is zero if the marginal is zero. This is a standard convention.

Particle filters are implemented by drawing independent uniform random variables $u_{t+1}^{(k)}$ and then evaluating a random function[13] of the form $X_{1,t+1}^{(k)}(u)$ and putting $\tilde{x}_{1,t+1}^{(k)} = X_{1,t+1}^{(k)}(u_{t+1}^{(k)})$ for $k = 1, \ldots, N$. Denote integration with respect to $(u_{t+1}^{(1)}, \ldots, u_{t+1}^{(N)})$ with $X_{1,t+1}^{(k)}(u)$ substituted into the integrand by $\tilde{\mathcal{E}}_{1,t+1}$.

Concatenated draws for fixed k that start at time s and end at time t are denoted

$$\tilde{x}_{1,s:t}^{(k)} = (\tilde{x}_{1,s}^{(k)}, \ldots, \tilde{x}_{1,t}^{(k)});$$

$\tilde{x}_{1,0:t}^{(k)}$ is called a particle. Denote expectation with respect to the uniform draws that gener-

[13] E.g., a conditional probability integral transformation, which depends on previous particle draws.

ated the particles by $\tilde{\mathcal{E}}_{1,0:t}$. In the following we present sampling schemes that satisfy certain unbiased properties that we discuss below.

5.1 A Conditionally Unbiased Particle Filter

In this section we discuss a particle implementation that satisfies the following recursive property (30) $\rightarrow$ (31). The result for all t follows by induction from this recursive property. Depending on how initial conditions are handled, the resulting estimate of the likelihood is unbiased for either the full information likelihood or a partial information likelihood.

Given draws $\tilde{x}_{1,0:t}^{(k)}$ and weights $\tilde{w}_t^{(k)}$, $k = 1, \ldots, N$, that approximate the density $p(x_{1,0:t}|\mathcal{F}_t)$ in the sense that

$$\int g(x_{1,0:t})\, dP(x_{1,0:t}|\mathcal{F}_t) = \tilde{\mathcal{E}}_{1,0:t}\left\{\mathcal{E}\left[\sum_{k=1}^{N} \tilde{w}_t^{(k)} g(\tilde{x}_{1,0:t}^{(k)}) \,|\, \mathcal{F}_t\right]\right\} \tag{30}$$

for integrable $g(x_{1,0:t})$, we seek to generate draws $\tilde{x}_{1,t+1}^{(k)}$ and compute weights $\tilde{w}_{t+1}^{(k)}$ that well approximate $p(x_{1,0:t}, x_{1,t+1}|\mathcal{F}_{t+1})$ in the sense that

$$\begin{aligned} &\int g(x_{1,0:t}, x_{1,t+1})\, dP(x_{1,0:t}, x_{1,t+1}|\mathcal{F}_{t+1}) \\ &= \tilde{\mathcal{E}}_{1,t+1}\tilde{\mathcal{E}}_{1,0:t}\left\{\mathcal{E}\left[\sum_{k=1}^{N} \tilde{w}_{t+1}^{(k)}\, g(\tilde{x}_{1,0:t}^{(k)}, \tilde{x}_{1,t+1}^{(k)}) \,|\, \mathcal{F}_{t+1}\right]\right\} \end{aligned} \tag{31}$$

for integrable $g(x_{1,0:t}, x_{1,t+1})$.[14] The notation $\mathcal{E}\left[\sum_{k=1}^{N} \tilde{w}_{t+1}^{(k)}\, g(\tilde{x}_{1,0:t}^{(k)}, \tilde{x}_{1,t+1}^{(k)}) \,|\, \mathcal{F}_{t+1}\right]$ is used to indicate that even with the uniform draws held fixed by the outer expectation the weights $\tilde{w}_{t+1}^{(k)}$ and draws $(\tilde{x}_{1,0:t}^{(k)}, \tilde{x}_{1,t+1}^{(k)})$ are functions of the variables in $\mathcal{F}_{t+1}$.

A specific objective of the paper is to find an unbiased estimator of either the full information or a partial information likelihood. The recursive property (30) $\rightarrow$ (31) is more general than the specific result (35) $\rightarrow$ (36) we require but it is actually easier to establish the general result than the specific result. An outline of the development is as follows. Induction implies (30) holds for all t; put $t = T$. As shown later (43), the weights in (30) are ratios. The unknown true likelihood function enters as the denominator of the weights on the right hand side of (30) and does not depend on k. Putting $g(x_{1,0:T}) \equiv 1$ in (30) implies

[14] An implication of (31) is that N affects the second moment of $\mathcal{E}\left[\sum_{k=1}^{N} \tilde{w}_{t+1}^{(k)}\, g(\tilde{x}_{1,0:t}^{(k)}, \tilde{x}_{1,t+1}^{(k)}) \,|\, \mathcal{F}_{t+1}\right]$ but not the first moment.

that the right hand side of (30) is an unbiased estimate of the constant 1, which, in turn, implies that the sum of the numerators of the weights provides an unbiased estimate of the likelihood, which is a full information likelihood if the expectation of the time $t = 0$ estimate is $p(a_0, x_{0,2}|\theta)$ or a partial information likelihood if the time $t = 0$ estimate is put to 1.

One often resamples particles such as $(\tilde{x}^{(k)}_{1,0:t}, \tilde{x}^{(k)}_{1,t+1})$ in order to prevent the variance of the weights $\tilde{w}^{(k)}_{t+1}$ from increasing with t. This may done by sampling $\left\{(\tilde{x}^{(k)}_{1,0:t}, \tilde{x}^{(k)}_{1,t+1})\right\}_{k=1}^{N}$ with replacement with probability $\frac{\tilde{w}^{(k)}_{t+1}}{\sum_{k=1}^{N} \tilde{w}^{(k)}_{t+1}}$. Some particles will get copied and some particles will not survive. We can represent the outcome of resampling by the number of times $\hat{N}^{(k)}_{t+1}$ that $(\tilde{x}^{(k)}_{1,0:t}, \tilde{x}^{(k)}_{1,t+1})$ is selected, where $\sum_{k=1}^{N} \hat{N}^{(k)}_{t+1} = N$. A particle that does not survive has $\hat{N}^{(k)}_{t+1} = 0$. Denoting expectation with respect to resampling by $\hat{\mathcal{E}}$, note that $\hat{\mathcal{E}}\left(\frac{\hat{N}^{(k)}_{t+1}}{N}\right) = \frac{\tilde{w}^{(k)}_{t+1}}{\sum_{j=1}^{N} \tilde{w}^{(j)}_{t+1}}$.

In the conditionally unbiased particle filter we define weights proportional to the resampled weights as follows

$$\hat{\tilde{w}}^{(k)}_{t+1} = \left(\sum_{j=1}^{N} \tilde{w}^{(j)}_{t+1}\right) \frac{\hat{N}^{(k)}_{t+1}}{N}. \tag{32}$$

Then

$$\begin{aligned}
&\tilde{\mathcal{E}}_{1,t+1}\tilde{\mathcal{E}}_{1,0:t}\,\hat{\mathcal{E}}\left\{\mathcal{E}\left[\sum_{k=1}^{N} \hat{\tilde{w}}^{(k)}_{t+1} g(\tilde{x}^{(k)}_{1,0:t}, \tilde{x}^{(k)}_{1,t+1}) \,|\, \mathcal{F}_{t+1}\right]\right\} \qquad (33)\\
&= \tilde{\mathcal{E}}_{1,t+1}\tilde{\mathcal{E}}_{1,0:t}\left\{\mathcal{E}\left[\hat{\mathcal{E}}\sum_{k=1}^{N} \hat{\tilde{w}}^{(k)}_{t+1} g(\tilde{x}^{(k)}_{1,0:t}, \tilde{x}^{(k)}_{1,t+1}) \,|\, \mathcal{F}_{t+1}\right]\right\}\\
&= \tilde{\mathcal{E}}_{1,t+1}\tilde{\mathcal{E}}_{1,0:t}\left\{\mathcal{E}\left[\hat{\mathcal{E}}\sum_{k=1}^{N}\left(\sum_{j=1}^{N} \tilde{w}^{(j)}_{t+1}\right) \frac{\hat{N}^{(k)}_{t+1}}{N} g(\tilde{x}^{(k)}_{1,0:t}, \tilde{x}^{(k)}_{1,t+1}) \,|\, \mathcal{F}_{t+1}\right]\right\}\\
&= \tilde{\mathcal{E}}_{1,t+1}\tilde{\mathcal{E}}_{1,0:t}\left\{\mathcal{E}\left[\sum_{k=1}^{N}\left(\sum_{j=1}^{N} \tilde{w}^{(j)}_{t+1}\right) \hat{\mathcal{E}}\left(\frac{\hat{N}^{(k)}_{t+1}}{N}\right) g(\tilde{x}^{(k)}_{1,0:t}, \tilde{x}^{(k)}_{1,t+1}) \,|\, \mathcal{F}_{t+1}\right]\right\}\\
&= \tilde{\mathcal{E}}_{1,t+1}\tilde{\mathcal{E}}_{1,0:t}\left\{\mathcal{E}\left[\sum_{k=1}^{N} \tilde{w}^{(k)}_{t+1} g(\tilde{x}^{(k)}_{1,0:t}, \tilde{x}^{(k)}_{1,t+1}) \,|\, \mathcal{F}_{t+1}\right]\right\}\\
&= \iint g(x_{1,0:t}, x_{1,t+1})\, dP(x_{1,0:t}, x_{1,t+1}|\mathcal{F}_{t+1}).
\end{aligned}$$

This shows that it is possible to retain conditional unbiasedness, by which we mean (31) holds, after resampling.

When resampling, it is equivalent to replicate a particle that receives positive weight $\hat{N}_{t+1}^{(k)}$ times, discard those for which $\hat{N}_{t+1}^{(k)} = 0$, renumber the particles, and use the weight

$$\hat{\tilde{w}}_{t+1} = \left(\sum_{j=1}^{N} \tilde{w}_{t+1}^{(j)} \right) \frac{1}{N} \tag{34}$$

for each particle.

Any resampling scheme for which $\hat{\mathcal{E}} \left(\frac{\hat{N}_{t+1}^{(k)}}{N} \right) = \frac{\tilde{w}_{t+1}^{(k)}}{\sum_{j=1}^{N} \tilde{w}_{t+1}^{(j)}}$. will have the property that conditional unbiasedness can be retained. What was described above is usually called multinomial resampling where uniform random numbers are drawn on the interval $(0, 1)$ and the inverse of the distribution function defined by the weights is evaluated. Other resampling schemes seek to improve performance by having one uniform random number in each interval $[(i-1)/N, i/N]$ for $i = 1, \cdots, N$. One approach is stratified resampling where one uniform u is drawn inside each interval. Another is systematic resampling where the same uniform u is placed inside each interval. In a comparison of stratified and systematic resampling, Douc, Cappé, and Moulines (2005) find that their performance is similar.

It might be that the weights $\tilde{w}_{t+1}^{(k)}$ are only known to within a constant that does not depend on k; i.e. $\tilde{w}_{t+1}^{(k)} = \bar{w}_{t+1}^{(k)} / \prod_{s=0}^{t+1} C_s$, where $\bar{w}_{t+1}^{(k)}$ can be computed but $\prod_{s=0}^{t+1} C_s$ cannot. There are two cases: (1) This is a nuisance to be dealt with. (2) $\prod_{s=0}^{T} C_s$ is the object of interest. In the first case it is possible to use cross-validation notions (Hastie, Tibshirani, and Friedman (2009)) and bias reduction formulae to achieve appoximate unbiasedness in (31). However, we are in the second situation where $\prod_{t=0}^{T} C_{t+1}$ is the object of interest and so omit consideration of the first case.

For the case where $\prod_{s=0}^{T} C_s$ is the object of interest, we modify our objective as given by (30) and (31) to the following: Given weights $\bar{w}_t^{(k)}$, $k = 1, \ldots, N$, that satisfy

$$\prod_{s=0}^{t} C_s = \tilde{\mathcal{E}}_{1,0:t} \, \mathcal{E} \left[\sum_{k=1}^{N} \bar{w}_t^{(k)} \mid \mathcal{F}_t \right], \tag{35}$$

we seek to generate weights $\bar{w}_{t+1}^{(k)}$ that satisfy

$$\prod_{s=0}^{t+1} C_s = \tilde{\mathcal{E}}_{1,t+1} \tilde{\mathcal{E}}_{1,0:t} \, \mathcal{E} \left[\sum_{k=1}^{N} \bar{w}_{t+1}^{(k)} \mid \mathcal{F}_{t+1} \right]. \tag{36}$$

Given that we achieve this objective, the same argument as above shows that resampling

does not destroy conditional unbiasedness if we set

$$\hat{\bar{w}}_{t+1}^{(k)} = \left(\sum_{j=1}^{N} \bar{w}_{t+1}^{(j)} \right) \frac{\hat{N}_{t+1}^{(k)}}{N}. \tag{37}$$

Or, if we replicate the particles with $\hat{N}_{t+1}^{(k)} > 0$ exactly $\hat{N}_{t+1}^{(k)}$ times, discard those with $\hat{N}_{t+1}^{(k)} = 0$, and renumber

$$\hat{\bar{w}}_{t+1} = \left(\sum_{j=1}^{N} \bar{w}_{t+1}^{(j)} \right) \frac{1}{N}. \tag{38}$$

As discussed immediately after (31), one observes from the foregoing algebra that by taking $g(\cdot) \equiv 1$ in both (30) and (31), it follows that if we can show that (30) implies (31) for weights of the form

$$\tilde{w}_t^{(k)} = \frac{\bar{w}_t^{(k)}}{\prod_{s=0}^{t} C_s} \qquad \tilde{w}_{t+1}^{(k)} = \frac{\bar{w}_{t+1}^{(k)}}{\prod_{s=0}^{t+1} C_s}, \tag{39}$$

then (35) implies (36) for the weights $\bar{w}_t^{(k)}$, $\bar{w}_{t+1}^{(k)}$.

Bayes Theorem states that

$$p(x_{1,0:t}, x_{1,t+1} | a_{t+1}, x_{2,t+1}, \mathcal{F}_t) = \frac{p(a_{t+1}, x_{2,t+1}, x_{1,0:t}, x_{1,t+1} | \mathcal{F}_t)}{p(a_{t+1}, x_{2,t+1} | \mathcal{F}_t)}. \tag{40}$$

However

$$p(x_{1,0:t}, x_{1,t+1} | a_{t+1}, x_{2,t+1}, \mathcal{F}_t) = p(x_{1,0:t}, x_{1,t+1} | \mathcal{F}_{t+1}) \tag{41}$$

and

$$\begin{aligned} & p(a_{t+1}, x_{2,t+1}, x_{1,0:t}, x_{1,t+1} | \mathcal{F}_t) \\ &= \; p(a_{t+1}, x_{2,t+1} | x_{1,0:t}, x_{1,t+1}, \mathcal{F}_t) p(x_{1,t+1} | x_{1,0:t}, \mathcal{F}_t) p(x_{1,0:t} | \mathcal{F}_t). \end{aligned} \tag{42}$$

We use Bayes theorem to show that (30) implies (31) when weights and draws at time $t+1$ are defined as follows: For $k = 1, \ldots, N$, given $\tilde{w}_t^{(k)}$ and $\tilde{x}_{1,0:t}^{(k)}$ defined by (30), draw $\tilde{x}_{1,t+1}^{(k)}$ from $p(x_{1,t+1} | \tilde{x}_{1,0:t}^{(k)}, \mathcal{F}_t)$ and define

$$\tilde{w}_{t+1}^{(k)} = \frac{p(a_{t+1}, x_{2,t+1} | \tilde{x}_{1,0:t}^{(k)}, \tilde{x}_{1,t+1}^{(k)}, \mathcal{F}_t)}{p(a_{t+1}, x_{2,t+1} | \mathcal{F}_t)} \, \tilde{w}_t^{(k)}. \tag{43}$$

We assume without loss of generality that all N of the $\tilde{w}_t^{(k)}$ are positive because one can, e.g., discard all particles with zero weight then, as often as necessary to get N particles,

replicate the particle with the largest weight and divide that weight half and half between that particle and its replicate. Then

$$\begin{aligned}
&\iint g(x_{1,0:t}, x_{1,t+1})\, dP(x_{1,0:t}, x_{1,t+1}|\mathcal{F}_{t+1}) \\
&\quad = \iint g(x_{1,0:t}, x_{1,t+1}) \frac{p(a_{t+1}, x_{2,t+1}|x_{1,0:t}, x_{1,t+1}, \mathcal{F}_t)}{p(a_{t+1}, x_{2,t+1}|\mathcal{F}_t)} p(x_{1,t+1}|x_{1,0:t}, \mathcal{F}_t) \\
&\qquad \times d\lambda(x_{1,t+1}) dP(x_{1,0:t}|\mathcal{F}_t) \qquad (44) \\
&\quad = \tilde{\mathcal{E}}_{1,0:t}\, \mathcal{E}\left[\int \sum_{k=1}^{N} g(\tilde{x}_{1,0:t}^{(k)}, x_{1,t+1})\ \tilde{w}_{t+1}^{(k)}\Big|_{\tilde{x}_{t+1}^{(k)} = x_{t+1}} p(x_{1,t+1}|\tilde{x}_{1,0:t}^{(k)}, \mathcal{F}_t)\, d\lambda(x_{1,t+1}) \,|\, \mathcal{F}_t \right] \qquad (45) \\
&\quad = \tilde{\mathcal{E}}_{1,t+1} \tilde{\mathcal{E}}_{1,0:t}\, \mathcal{E}\left[\sum_{k=1}^{N} g(\tilde{x}_{1,0:t}^{(k)}, \tilde{x}_{1,t+1}^{(k)})\, \tilde{w}_{t+1}^{(k)} \,|\, \mathcal{F}_{t+1} \right] \qquad (46)
\end{aligned}$$

where (44) is due to (40) after substituting (41) and (42), (45) is due to (30) and (43), and (46) is due to the fact that $\tilde{x}_{1,t+1}^{(k)}$ is a draw from $p(x_{1,t+1}|\tilde{x}_{1,0:t}^{(k)}, \mathcal{F}_t)$.

The denominator $p(a_{t+1}, x_{2,t+1}|\mathcal{F}_t)$ of (43) is C_{t+1}; i.e., one of the components of the object of interest $\prod_{s=0}^{T} C_s$. We need to express the numerator of (43) in terms of the primitives (26), (27), and (28). Now,

$$\begin{aligned}
&p(a_{t+1}, x_{2,t+1}|x_{1,0:t}, x_{1,t+1}, \mathcal{F}_t) \\
&\quad = p(a_{t+1}, x_{2,t+1}|x_{1,0:t}, x_{1,t+1}, x_{2,0:t}, a_{0:t}, \theta) \\
&\quad = \frac{p(a_{t+1}, x_{t+1}, a_{0:t}, x_{0:t}, \theta)}{\int p(a_{t+1}, x_{t+1}, a_{0:t}, x_{0:t}, \theta)\, d\lambda(a_{t+1}, x_{2,t+1})} \\
&\quad = \frac{p(a_{t+1}|x_{t+1}, a_{0:t}, x_{0:t}, \theta) p(x_{t+1}|a_{0:t}, x_{0:t}, \theta) p(a_{0:t}, x_{0:t}, \theta)}{\int p(a_{t+1}|x_{t+1}, a_{0:t}, x_{0:t}, \theta) p(x_{t+1}|a_{0:t}, x_{0:t}, \theta) d\lambda(a_{t+1}, x_{2,t+1})\, p(a_{0:t}, x_{0:t}, \theta)} \\
&\quad = \frac{p(a_{t+1}|x_{t+1}, \theta)\, p(x_{t+1}|a_t, x_t, \theta)}{\int p(a_{t+1}|x_{t+1}, \theta)\, p(x_{t+1}|a_t, x_t, \theta)\, d\lambda(a_{t+1}, x_{2,t+1})} \\
&\quad = \frac{p(a_{t+1}|x_{t+1}, \theta)\, p(x_{t+1}|a_t, x_t, \theta)}{\int p(x_{t+1}|a_t, x_t, \theta)\, d\lambda(x_{2,t+1})} \\
&\quad = \frac{p(a_{t+1}|x_{t+1}, \theta)\, p(x_{t+1}|a_t, x_t, \theta)}{p(x_{1,t+1}|a_t, x_t, \theta)} \qquad (47)
\end{aligned}$$

Therefore,

$$\tilde{w}_{t+1}^{(k)} = \frac{\bar{v}_{t+1}^{(k)}}{C_{t+1}}\, \tilde{w}_t^{(k)} \qquad (48)$$

where

$$\bar{v}_{t+1}^{(k)} = \frac{p\Big(a_{t+1} \,|\, \tilde{x}_{1,t+1}^{(k)}, x_{2,t+1}, \theta\Big)\, p\Big(\tilde{x}_{1,t+1}^{(k)}, x_{2,t+1} \,|\, a_t, \tilde{x}_{1,t}^{(k)}, x_{2,t}, \theta\Big)}{p\Big(\tilde{x}_{1,t+1}^{(k)} \,|\, a_t, \tilde{x}_{1,t}^{(k)}, x_{2,t}, \theta\Big)} \tag{49}$$

and

$$C_{t+1} = p(a_{t+1}, x_{2,t+1} | \mathcal{F}_t). \tag{50}$$

The weights we need to estimate the likelihood follow the recursion

$$\bar{w}_{t+1}^{(k)} = \bar{v}_{t+1}^{(k)}\, \bar{w}_t^{(k)} \tag{51}$$

because (39) follows from

$$\tilde{w}_{t+1}^{(k)} = \frac{\bar{v}_{t+1}^{(k)} \bar{w}_t^{(k)}}{\prod_{s=0}^{t+1} C_s} \tag{52}$$

provided that $\tilde{w}_t^{(k)} = \frac{\bar{w}_t^{(k)}}{\prod_{s=0}^{t} C_s}$.

We have established the recursion (30) $\rightarrow$ (31), now we must think about how to start it. We need an estimator for C_0. Ideally the estimator should be conditionally unbiased for $p(a_0, x_{2,0} \,|\, \theta)$. We are unwilling to impose the additional structure on the game necessary to be able to estimate that value although some games actually do have the requisite structure. Our example is an instance.[15] Therefore, as is routinely done in time series analysis, we discard the information in the stationary density for $(a_0, x_{2,0})$ and set $C_0 = 1$.[16] With this convention, we can start the filter with draws $\left\{\tilde{x}_0^{(k)}\right\}_{k=1}^{N}$ from the stationary density (29) and put the initial weights to $\bar{w}_0^{(k)} = 1/N$.

Consider the case of no resampling where particles do not lose their original labels. In this case

$$\begin{aligned}
\sum_{k=0}^{N} \bar{w}_T^{(k)} &= \left(\frac{\sum_{k=1}^{N} \bar{v}_T^{(k)} \bar{w}_{T-1}^{(k)}}{\sum_{k=1}^{N} \bar{v}_{T-1}^{(k)} \bar{w}_{T-2}^{(k)}}\right)\left(\frac{\sum_{k=1}^{N} \bar{v}_{T-1}^{(k)} \bar{w}_{T-2}^{(k)}}{\sum_{k=1}^{N} \bar{v}_{T-2}^{(k)} \bar{w}_{T-3}^{(k)}}\right) \cdots \left(\frac{\sum_{k=1}^{N} \bar{v}_1^{(k)} \bar{w}_0^{(k)}}{\sum_{k=1}^{N} \bar{w}_0^{(k)}}\right)\left(\sum_{k=1}^{N} \bar{w}_0^{(k)}\right) \\
&= \left(\sum_{k=1}^{N} \bar{v}_T^{(k)} \frac{\bar{w}_{T-1}^{(k)}}{\sum_{k=1}^{N} \bar{w}_{T-1}^{(k)}}\right)\left(\sum_{k=1}^{N} \bar{v}_{T-1}^{(k)} \frac{\bar{w}_{T-2}^{(k)}}{\sum_{k=1}^{N} \bar{w}_{T-2}^{(k)}}\right) \cdots \left(\sum_{k=1}^{N} \bar{v}_1^{(k)} \frac{\bar{w}_0^{(k)}}{\sum_{k=1}^{N} \bar{w}_0^{(k)}}\right)\left(\sum_{k=1}^{N} \bar{w}_0^{(k)}\right) \\
&= \left(\sum_{k=1}^{N} \bar{v}_T^{(k)} \hat{w}_{T-1}^{(k)}\right)\left(\sum_{k=1}^{N} \bar{v}_{T-1}^{(k)} \hat{w}_{T-2}^{(k)}\right) \cdots \left(\sum_{k=1}^{N} \bar{v}_1^{(k)} \hat{w}_0^{(k)}\right)\left(\sum_{k=1}^{N} \bar{w}_0^{(k)}\right)
\end{aligned} \tag{53}$$

[15] Use (22) to evaluate $p(a_t, x_{2t} \,|\, x_{1,t}, \theta)$; use (21) to draw from $p(x_{1,t} \,|\, \theta)$.

[16] In the Bayesian framework the stationary density $p(a_0, x_{2,0} \,|\, \theta)$ can be regarded as (part of) the prior for θ. Putting $C_0 = 1$ replaces this informative prior by an uninformative prior.

where

$$\hat{w}_t^{(k)} = \frac{\bar{w}_t^{(k)}}{\sum_{k=1}^N \bar{w}_t^{(k)}}. \tag{54}$$

The import of (53) is that the time t weights can be normalized to sum to one before proceeding to time $t+1$ because normalization does not affect which $\tilde{x}_t^{(k)}$ get drawn. If weights are normalized, the estimator of the likelihood is

$$\widehat{\prod_{t=0}^T C_t} = \prod_{t=0}^T \hat{C}_t$$

where

$$\hat{C}_t = \sum_{k=1}^N \bar{v}_t^{(k)} \hat{w}_{t-1}^{(k)}$$

The same is true if resampling is used because the telescoping argument (53) shows that the scale factors that appear in (38) cancel, in which case

$$\hat{w}_t^{(k)} = \frac{\bar{v}_t^{(k)}}{\sum_{k=1}^N \bar{v}_t^{(k)}}. \tag{54$'$}$$

5.2 An Alternative Importance Sampler

If computing $p(x_{1,t+1}|a_t, x_t, \theta)$ is costly or drawing from it troublesome, one can substitute an alternative importance sampler. The idea is that one can advance a filter from $(\tilde{\tilde{x}}_t^{(k)}, \tilde{\tilde{w}}_t^{(k)})$ that satisfies (30) to $(\tilde{\tilde{x}}_{t+1}^{(k)}, \tilde{\tilde{w}}_{t+1}^{(k)})$ that satisfies (31) by drawing $\tilde{\tilde{x}}_{t+1}^{(k)}$ from

$$f(x_{1,t+1}|x_{1t}, \mathcal{F}_t) = f(x_{1,t+1}|a_t, x_t, \theta) \tag{55}$$

for $k = 1, \ldots, N$, and setting

$$\tilde{\tilde{w}}_{t+1}^{(k)} = \frac{p\Big(a_{t+1}\,|\tilde{\tilde{x}}_{1,t+1}^{(k)}, x_{2,t+1}, \theta\Big)\, p\Big(\tilde{\tilde{x}}_{1,t+1}^{(k)}, x_{2,t+1}\,|\,a_t, \tilde{\tilde{x}}_{1t}^{(k)}, x_{2t}, \theta\Big)}{C_{t+1}\; f\Big(\tilde{\tilde{x}}_{1,t+1}^{(k)}\,|\,a_t, \tilde{\tilde{x}}_{1t}^{(k)}, x_{2t}, \theta\Big)} \tilde{\tilde{w}}_t^{(k)} \tag{56}$$

as is seen by noting that (44), (45), and (46) can be rewritten as

$$\begin{aligned}
&\iint g(x_{1,0:t}, x_{1,t+1}) \frac{p(a_{t+1}, x_{2,t+1}|x_{1,0:t}, x_{1,t+1}, \mathcal{F}_t)}{p(a_{t+1}, x_{2,t+1}|\mathcal{F}_t)} \frac{p(x_{1,t+1}|x_{1,0:t}, \mathcal{F}_t)}{f(x_{1,t+1}|x_{1,t}, \mathcal{F}_t)} f(x_{1,t+1}|x_{1t}, \mathcal{F}_t) \\
&\qquad \times d\lambda(x_{t+1})\, dP(x_{1,0:t}|\mathcal{F}_t) \qquad (57) \\
&= \tilde{\mathcal{E}}_{1,0:t} \int \mathcal{E} \left[\sum_{k=1}^{N} g(\tilde{\tilde{x}}_{1,0:t}^{(k)}, x_{1,t+1}) \tilde{\tilde{w}}_{t+1}^{(k)} f(x_{1,t+1}|\tilde{\tilde{x}}_t^{(k)}, \mathcal{F}_t)\, d\lambda(x_{1,t+1}) \,|\, \mathcal{F}_t \right] \\
&= \tilde{\mathcal{E}}_{1,t+1} \tilde{\mathcal{E}}_{1,0:t}\, \mathcal{E} \left[\sum_{k=1}^{N} g(\tilde{\tilde{x}}_{1,0:t}^{(k)}, \tilde{\tilde{x}}_{t+1}^{(k)})\, \tilde{\tilde{w}}_{t+1}^{(k)} \,|\, \mathcal{F}_{t+1} \right]
\end{aligned}$$

due to the cancellation $p(x_{1,t+1}|x_{1,0:t}, \mathcal{F}_t)/p(x_{1,t+1}|a_t, x_t, \theta) = 1$ that occurs after the expression for $p(a_{t+1}, x_{2,t+1}|x_{1,0:t}, x_{1,t+1}, \mathcal{F}_t)$ given by (47) is substituted in (57).

The equations that replace (48), (49), and (51) when an alternative importance sampler is used are

$$\tilde{\tilde{w}}_{t+1}^{(k)} = \frac{\bar{v}_{t+1}^{(k)}}{C_{t+1}}\, \tilde{\tilde{w}}_t^{(k)} \qquad (48')$$

$$\bar{v}_{t+1}^{(k)} = \frac{p\Big(a_{t+1} \,|\tilde{\tilde{x}}_{1,t+1}^{(k)}, x_{2,t+1}, \theta\Big)\, p\Big(\tilde{\tilde{x}}_{1,t+1}^{(k)}, x_{2,t+1} \,|\, a_t, \tilde{\tilde{x}}_{1,t}^{(k)}, x_{2,t}, \theta\Big)}{f\Big(\tilde{\tilde{x}}_{1,t+1}^{(k)} \,|\, a_t, \tilde{\tilde{x}}_{1,t}^{(k)}, x_{2,t}, \theta\Big)}. \qquad (49')$$

$$\bar{\tilde{w}}_{t+1}^{(k)} = \bar{\tilde{v}}_{t+1}^{(k)}\, \bar{\tilde{w}}_t^{(k)} \qquad (51')$$

The requisite regularity condition is the following:

ASSUMPTION 2

$$g(x_{1,0:t}, x_{1,t+1}) \frac{p\left(a_{t+1} \,|\, x_{1,t+1}, x_{2,t+1}, \theta\right)\, p\left(x_{1,t+1}, x_{2,t+1} \,|\, a_t, x_{1,t}, x_{2,t}, \theta\right)}{f(x_{1,t+1} \,|\, a_t, x_{1,t}, x_{2,t}, \theta)}$$

is integrable with respect to $f(x_{1,t+1} \,|\, a_t, x_{1,t}, x_{2,t}, \theta)$, the support of which contains the support of $p\left(x_{1,t+1} \,|\, a_t, x_{1,t}, x_{2,t}, \theta\right)$.

Another reason to consider an alternative importance sampler is to improve efficiency. Pitt and Shephard (1999) suggest some adaptive importance samplers that one might consider. In addition to Pitt and Shephard's (1999) suggestions, one can use the notion of reprojection (Gallant and Tauchen (1998)) to construct an adaptive density for (55) as follows. The model can be simulated. Therefore, for given θ^* a large simulation of

$(a_t, x_{1t}, x_{2t}, a_{t+1}, x_{1,t+1}, x_{2,t+1})$ can be generated. Using multivariate regression one can determine the location $\mu(v)$ of $x_{1,t+1}$ as a linear function of

$$v = (a_t, x_{1t}, x_{2t}, a_{t+1}, x_{2,t+1}) \tag{58}$$

and the conditional variance Σ. The simulation can be taken so large that $\mu(v)$ and Σ can be regarded as population quantities. We put

$$h(x_{1,t+1}|x_{1t}, \mathcal{F}_{t+1}) = n(x_{1,t+1}|\mu(v), \Sigma), \tag{59}$$

where $n(\cdot|\mu, \Sigma)$ denotes the multivariate normal density and use (59) in place of (55), which is a slight abuse of notation because the argument lists are different. We also experimented with the multivariate Student-t density on five and six degrees of freedom with the same location and scale but found that it had little effect on results other than increase run times. One can see this same notion of looking ahead one step in order to improve the efficiency of an estimate of a latent variable in a Bayesian context in Jacquier, Polson, and Rossi (1994).

This begs the question of how to choose θ^* for the simulation that determines $\mu(v)$ and Σ. If one is using a strongly informative prior one can use the mean or the mode of the prior for θ^*. Otherwise, one can use the mean or mode of the MCMC chain described in Section 6 using the filter described in Section 5 with draws from (28).

5.3 Computing the Likelihood

A draw from a density $f(v)$ is obtained by drawing a seed s from a uniform density $u(s)$ defined over a finite set of integers and executing an algorithm that evaluates a function $V(s)$ and returns $v' = V(s)$ and s' such that v' has density $f(v)$, s' has density $u(s)$, and s' is independent of s. The next draw from the same or a different density uses s' to return a draw v'' from that density and another new seed s'', and so on. The algorithm that we describe next has sequence of such draws within it but viewed as a whole it has the same flavor as a single draw: One specifies θ and provides a random draw s from $u(s)$. The algorithm evaluates a function $\mathcal{L}(\theta, s)$ and returns $\ell' = \mathcal{L}(\theta, s)$ and a draw s' from $u(s)$ that is independent of s. The crucial fact regarding the algorithm is that $\int \mathcal{L}(\theta, s)\, u(s)\, ds = \mathcal{L}(\theta)$, where $\mathcal{L}(\theta)$ is the likelihood of the game described at the beginning of Section 3. See Flury

and Shephard (2010) for a further discussion of this point and illustration with econometric examples.

Given seed s and parameter θ, the algorithm for evaluating $\mathcal{L}(\theta, s)$ follows. All draws use the seed returned by the previous draw; there are no fixed seeds anywhere within the algorithm.

1. For $t = 0$

   (a) Start N particles by drawing $\tilde{x}_{1,0}^{(k)}$ from $p(x_{1,0} \,|\, \theta)$ using s as the initial seed.

   (b) If $p(a_t, x_{2t} \,|\, x_{1,t}, \theta)$ is available, compute $\hat{C}_0 = \frac{1}{N} \sum_{k=1}^{N} p(a_0, x_{2,0} \,|\, \tilde{x}_{1,0}^{(k)}, \theta)$ otherwise put $\hat{C}_0 = 1$.

   (c) Set $x_{1,0:0}^{(k)} = \tilde{x}_{1,0}^{(k)}$ and $x_{1,0}^{(k)} = \tilde{x}_{1,0}^{(k)}$

2. For $t = 1, \ldots, T$

   (a) For each particle, draw $\tilde{x}_{1t}^{(k)}$ from the transition density

   $$p(x_{1t} \,|\, a_{t-1}, x_{1,t-1}^{(k)}, x_{2,t-1}, \theta). \tag{60}$$

   (b) Compute

   $$\begin{aligned} \bar{v}_t^{(k)} &= \frac{p\Big(a_t \,|\, \tilde{x}_{1,t}^{(k)}, x_{2,t}, \theta\Big)\, p\Big(\tilde{x}_{1,t}^{(k)}, x_{2,t} \,|\, a_{t-1}, x_{1,t-1}^{(k)}, x_{2,t-1}, \theta\Big)}{p\Big(\tilde{x}_{1,t}^{(k)} \,|\, a_{t-1}, x_{1,t-1}^{(k)}, x_{2,t-1}, \theta\Big)} \\ \hat{C}_t &= \frac{1}{N} \sum_{k=1}^{N} \bar{v}_t^{(k)} \end{aligned} \tag{61}$$

   (Note that the draw pair is $(x_{1,t-1}^{(k)}, \tilde{x}_{1,t}^{(k)})$ and the weight is $\bar{v}_t^{(k)} \frac{1}{N}$.)

   (c) Set

   $$\tilde{x}_{1,0:t}^{(k)} = \left(x_{1,0:t-1}^{(k)}, \tilde{x}_{1,t}^{(k)}\right).$$

   (d) Compute the normalized weights

   $$\hat{w}_t^{(k)} = \frac{\bar{v}_t^{(k)}}{\sum_{k=1}^{N} \bar{v}_t^{(k)}}$$

(e) For $k = 1, \ldots, N$ draw $x_{1,0:t}^{(k)}$ by sampling with replacement from the set $\{\tilde{x}_{1,0:t}^{(k)}\}$ according to the weights $\{\hat{w}_t^{(k)}\}$.

(Note the convention: Particles with unequal weights are denoted by $\{\tilde{x}_{0:t}^{(k)}\}$. After resampling the particles are denoted by $\{x_{0:t}^{(k)}\}$.)

(f) Set $x_t^{(k)}$ to the last element of $x_{1,0:t}^{(k)}$.

3. Done

(a) An unbiased estimate of the likelihood is

$$\ell' = \prod_{t=0}^{T} \hat{C}_t \tag{62}$$

(b) s' is the last seed returned in Step 2e.

Systematic or stratified sampling can be used at step 2e instead of multinomial resampling. To use the alternative importance sampler of Section 5.2, replace (60) with $f(x_{1,t+1}|a_t, x_t^{(k)}, \theta)$[17] and replace (61) with

$$\bar{\bar{v}}_t^{(k)} = \frac{p\Big(a_t \,|\, \tilde{x}_{1,t}^{(k)}, x_{2,t}, \theta\Big)\, p\Big(\tilde{x}_{1,t}^{(k)}, x_{2,t} \,|\, a_{t-1}, x_{1,t-1}^{(k)}, x_{2,t-1}, \theta\Big)}{f\Big(x_{1,t}^{(k)} \,|\, a_{t-1}, x_{1,t-1}^{(k)}, x_{2,t-1}, \theta\Big)}. \tag{63}$$

The algorithm may or may not produce a likelihood whose log is the sum of a martingale difference sequence. It depends on the application and the choice of importance sampler. Whether it does or not is irrelevant to Bayesian inference but might be relevant to frequentist inference. Because its output is rough, the algorithm is completely inappropriate for maximum likelihood inference that uses an optimization algorithm that relies on smoothness (Pitt (2002)) but would be a good choice when using an MCMC algorithm (Chernozhukov and Hong (2003)).

5.3.1 Specialization to the Boundedly and Fully Rational Models

For the boundedly rational model, substituting (20) and (18) into the numerator of (61) and (19) into the denominator, we have

$$\bar{v}_t^{(k)} = p[A_t \,|\, S_B(c_{u,t}^{(k)}, c_{k,t}, r_t, \theta), \theta]\, n(r_t \,|\, \mu_r, \sigma_r^2)\, \delta[c_{k,t} = c_{k,t-1} + \kappa_c S_B(c_{u,t-1}^{(k)}, c_{k,t-1}, r_{t-1}, \theta)]$$

[17] See (55).

For the alternative importance sampler,

$$\bar{\bar{v}}_t^{(k)} = \bar{v}_t^{(k)} \frac{n[c_{ut}^{(k)} \mid \mu_c \mathbf{1} + \rho_c(c_{u,t-1}^{(k)} - \mu_c \mathbf{1}), \sigma_c^2 I]}{f\left(c_{ut}^{(k)} \mid c_{u,t-1}^{(k)}, c_{k,t-1}, r_{t-1}, \theta\right)}.$$

The expressions needed at Step 1 to compute $\hat{C}_0$ are given by (21) and (22). For the fully rational model the expressions are the same with S_R replacing S_B.

6 Computing the Posterior

We determine the posterior density of θ using the Metropolis algorithm. The Metropolis algorithm is an iterative scheme that generates a Markov chain whose stationary distribution is the posterior of θ. To implement our version we require the particle filter algorithm for drawing (ℓ, s) described in Section 5.3, a prior $\pi(\theta)$, and a transition density in θ called the proposal density. For a given θ', a proposal density $q(\theta', \theta^*)$ defines a distribution of potential new values θ^*. We use a move-one-at-a-time, random-walk, proposal density which is built in to the public domain software that we use: `http://econ.duke.edu/webfiles/arg/emm`

The algorithm for the Markov chain follows. For the first θ in the chain we also need to write to memory a draw s'' from the uniform density on a finite set of integers used to implement the particle filter in Section 5.3.

Given a current θ' we obtain the next θ'' as follows:

1. Draw θ^* according to $q(\theta', \theta^*)$.
2. Set s^* to s'' retrieved from memory.
3. Compute ℓ^* corresponding to (θ^*, s^*) using the particle filter in Section 5.3 and write to memory the s'' returned by the particle filter.
4. Compute $\alpha = \min\left(1, \frac{\ell^* \, \pi(\theta^*) \, q(\theta^*, \theta')}{\ell' \, \pi(\theta') \, q(\theta', \theta^*)}\right)$.
5. With probability α, set $\theta'' = \theta^*$, otherwise set $\theta'' = \theta'$.
6. Return to 1.

The choice for the parameter N of the particle filter in Section 5.3 does play a role in the performance of the MCMC chain: It influences the rejection rate. If N is too small then

$\mathcal{L}(\theta, s) = \ell'$ given by (62) will be a jittery estimator of $\mathcal{L}(\theta)$ which will increase the chance that the chain gets stuck. Pitt (2010) shows that what is relevant is the variance

$$\text{Var}\left\{\log \mathcal{L}(\theta, s)\right\} = \int \left[\log \mathcal{L}(\theta, s) - \int \log \mathcal{L}(\theta, s)\, ds\right]^2 ds, \tag{64}$$

which can be computed from draws of ℓ' obtained by putting the filter in a loop. It is interesting that for an entry game such as our example, the classification error rate can be so small that one is almost matching 0's and 1's and using the particle filter to solve backwards for $\{x_{1t}\}$ that will allow the match. The consequence is that N can be quite small. For our example, Pitt's charts suggest that $N = 300$ will suffice. We actually use $N = 512$. Even without such charts, one can always determine N empirically by increasing it until the chain is no longer sticky. The fact that N can be quite small is one of the main features of our approach.

It is possible to set forth regularity conditions such that $\lim_{N\to\infty} \sup_\theta \text{Var}\left\{\log \mathcal{L}(\theta, s)\right\} = 0$. They are stringent: see Andrieu, Douced, and Holenstein (2010). One might argue that there is no point to verifying that variance declines with N in an application because N is a tuning parameter that affects the rejection rate of MCMC in much the same manner as the scale parameters of the proposal density. One must still determine the requisite N empirically. If an acceptable N is found, it does not matter if variance declines with N or not. If an affordable N cannot be found, a proof that variance declines with N does not help except to provide support for a request for more computing resources.

7 Simulation Experiment Results

To assess the efficacy of the approach proposed here that directly contradicts current practice in that the seed is random and the number of particles small, we conduct a simulation exercise.

We simulate the fully rational game described in Subsection 3.5 configured to represent the manufacture of a single object where entry constrains capacity. There are three firms. The time increment is one year. We set parameters according to the following considerations. A hurdle rate of 20% is a standard assumption in business which leads to a discount factor of $\beta = 0.83333$. Setting $p_a = 0.95$ seems intuitively reasonable and is in line with the estimates

of Gallant, Hong, and Khwaja (2010) who estimate a similar model from pharmaceutical data, where entry has the effect of reducing rather than increasing costs. We set $\rho_a = 0.5$, which gives the entry effect a half-life of six-months. Costs are usually persistent so $\rho_c = 0.9$ seems reasonable. The remaining parameters scale with μ_r. The parameter μ_r can be chosen arbitrarily because it is the log of the nominal price of the product. We chose $\mu_r = 10$. A gross margin of 30% puts $\mu_c = 9.7$. With $\kappa_a = 0.2$ the immediate impact of entry is to reduce the gross margin to 10%. The two scale parameters σ_c and σ_r are determined by the foregoing because, if one wants a sample that mimics competition to some extent, there is far less freedom in their choice than one might imagine. One can easily produce samples where one firm is dominant for long periods or a monopoly develops. By trial and error, we found $\sigma_c = 0.1$ and $\sigma_r = 2$ to be satisfactory. In general, σ_r must be fairly large, as it is here, to prevent a monopoly.

Gallant, Hong, and Khwaja (2010) reported that p_a was estimated precisely and varying it within reason had little effect on estimates. Because the parameter was of no intrinsic interest, they fixed it to reduce computational cost. We estimated with p_a both fixed and free to see if that held true here.

The firm's discount rate β is extremely difficult to estimate in studies of this sort (see e.g., Magnac and Thesmar (2002) and Rust (1994)). On the other hand it is not difficult to form priors for β. As mentioned above, a common rule of thumb in business is not to undertake a project whose internal rate of return is less than 20%. Theoretically, a firm should not undertake a project whose rate of return is less than its cost of capital. The historical risk premia for various industries are available (e.g., Gebhardt, Lee, and Swaminathan (2001)) to which one can add a nominal borrowing rate of 5% to arrive at a value for β. We estimated with β both fixed and free to assess the value of prior information regarding β.

We also investigated the effect of misspecification by estimating using the boundedly rational model of Subsection 3.4.

All the models in the suite described in Section 3.1 are recursive due to (4). The customary way of dealing with this situation in time series analysis (e.g. GARCH models) is to run the recursion over a few initial lags prior to estimation. We set the number of initial lags to a large value $T_0 = 160$ to reduce effect of the choice of T_0 in our results. The choice of large T_0

was also motivated by the Gallant, Hong, and Khwaja (2010) study where a structural break – a bribery scandal – gave rise to 160 initial lags that could be used to run the recursion (4) but could not be used for estimation. As in Gallant, Hong, and Khwaja (2010), we also pass (5) through the recursion as part of the likelihood which is equivalent to determining a loose prior for μ_r and σ_r from the initial lags. We used three simulated data sets, small, medium, and large, with $T = 40, 120$, and 360 respectively.

What we propose here is computationally intensive. Serial computation on a 2.9 MHz CPU takes about 8 hours per 5,000 MCMC repetitions for the medium size data set. Our code is not particularly efficient because it collects a lot of diagnostic information. Perhaps efficiency could be improved by 20% by removing these subsidiary computations. On the other hand, the computations are trivially parallelizable with linear scaling. The public domain code that we use, `http://econ.duke.edu/webfiles/arg/emm`, has parallelization built in. Machines with 8 cores are nearly standard (two Intel quad core chips). Machines with 48 cores (four AMD twelve core chips) are available at a reasonable price. On a 48 core machine the computational cost would be 10 minutes per 5,000 MCMC repetitions.

We considered four cases

1. The fully rational model is fit to the data using a blind sampler and multinomial resampling. Estimates are in Table 1. Histograms of the marginals of the posterior density are in Figure 1 for the medium sample size. Figure 2 is the same with β constrained. Figure 3 shows the latent cost estimates for the medium sample size and β constrained.

2. The boundedly rational model is fit to the data using a blind sampler and multinomial resampling. Estimates are in Table 2.

3. The fully rational model is fit to the data using an adaptive sampler and multinomial resampling. Estimates are in Table 3.

4. The fully rational model is fit to the data using an adaptive sampler and systematic resampling. Estimates are in Table 4. Figure 4 shows the latent cost estimates for the medium sample size and β constrained.

The key parameter in the study of games of this sort is κ_a so we mainly focus on it although our remarks generally apply to the other parameters as well. In most respects our results are not surprising.

- A large sample size is better. In Tables 1 through 4 the estimates shown in the columns labeled "lg" would not give misleading results in an application.

- In smaller sample sizes the specification error caused by fitting the boundedly rational model to data generated by the fully rational model can be serious: compare columns "sm" and "md" in Tables 1 and 2. The saving in computational time is about 10% relative to the fully rational model so there seems to be no point to using the boundedly rational model unless that is what firms are actually doing, which they are not in this instance.

- Constraining β is beneficial: compare Figures 1 and 2. The constraint reduces the bimodality of the marginal posterior distribution of σ_r and pushes all histograms closer to unimodality. In consequence, the descriptive statistics in the columns labeled "β" and "β & p_a" of Tables 1 through 4 represent the posterior distribution better than those in the columns labled "Unconstrained."

- Constraining p_a is irrelevant except for a small savings in computational cost: compare columns "β" and "β & p_a" in Tables 1 through 4.

- Improvements to the particle filter are helpful. In particular, an adaptive sampler is better than a blind sampler; compare Tables 1 and 3 and compare Figures 3 and 4. Systematic resampling is better than multinomial resampling; compare Tables 3 and 4.

Table 1 about here

Table 2 about here

Table 3 about here

Table 4 about here

Figure 1 about here

Figure 2 about here

Figure 3 about here

Figure 4 about here

8 Conclusions

We propose a method based on sequential importance sampling (particle filtering) to estimate the parameters of a dynamic game that can have state variables that are partially observed, serially correlated, endogenous, and heterogeneous. We illustrated by application to a dynamic oligopolistic model for which a capacity constraint due to entry affects future costs.

The method depends on computing an unbiased estimate of the likelihood that is used within a Metropolis chain to conduct Bayesian inference. Unbiasedness guarantees that the stationary density of the chain is the exact posterior, not an approximation. The remarkable feature of this approach is that the number of particles required is both small and easily determined.

References

Abramowitz, M., and I. A. Stegun (1964): *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. Dover.

Ackerberg, D. (2009): "A new use of importance sampling to reduce computational burden in simulation estimation," *Quantitative Marketing and Economics*, 7(4), 343–376.

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

Andrieu, C., A. Douced, and R. Holenstein (2010): "Particle Markov Chain Monte Carlo Methods," *Journal of the Royal Statistical Society, Series B*, 72, 269–342.

Aradillas-Lopez, A. (2010): "Semiparametric Estimation of a Simultaneous Game with Incomplete Information," *Journal of Econometrics*, 157(2), 409–431.

Arcidiacono, P., and R. Miller (2008): "CCP Estimation of Dynamic Discrete Choice Models with Unobserved Heterogeneity," Working paper, Duke University.

Augereau, A., S. Greenstein, and M. Rysman (2006): "Coordination versus differentiation in a standards war: 56k modems," *RAND Journal of Economics*, 37(4), 887–909.

Bajari, P., L. Benkard, and J. Levin (2007): "Estimating Dynamic Models of Imperfect Competition," *Econometrica*, 75, 1331–1370.

Bajari, P., V. Chernozhukov, H. Hong, and D. Nekipelov (2007): "Semiparametric Estimation of a Dynamic Game of Incomplete Information," Working paper, Stanford University.

Bajari, P., H. Hong, J. Krainer, and D. Nekipelov (2006): "Estimating Static Models of Strategic Interactions," Working paper, Stanford University.

Bajari, P., H. Hong, and S. Ryan (2010): "Identification and Estimation of Discrete Games of Complete Information," *Econometrica*, 78(5), 1529–1568.

Berry, S. (1992): "Estimation of a model of entry in the airline industry," *Econometrica*, 60(4), 889–917.

Bjorn, P. A., and Q. Vuong (1997): "Modéles d'Equations Simultanées pour Variables Endogénes Fictives: Une Formulation par la Théorie des Jeux avec Application á la Participation au Marché du Travail," *L'Actualité Economique*, 73, 161–205, In English at http://www.hss.caltech.edu/SSPapers/sswp537.pdf.

Blevins, J. (2009): "Sequential Monte Carlo Methods for Estimating Dynamic Microeconomic Models," Working paper, Duke University.

Bresnahan, T., and P. Reiss (1991a): "Empirical Models of Discrete Games," *Journal of Econometrics*, 48, 57—81.

——— (1991b): "Empirical Models of Discrete Games," *Journal of Econometrics*, 48, 57—81.

——— (1991c): "Entry and competition in concentrated markets," *Journal of Political Economy*, 99, 977–1009.

Chen, H. (2010): ""Bayesian SMC Estimation of a Dynamic Ologopolistic Game of Complete Information."," Working paper, Duke University.

Chernozhukov, V., and H. Hong (2003): "A MCMC Approach to Classical Estimation," *Journal of Econometrics*, 115(2), 293–346.

Ching, A. (2010): "Consumer learning and heterogeneity: Dynamics of demand for prescription drugs after patent expiration," *International Journal of Industrial Organization*.

Ciliberto, F., and E. Tamer (2009): "Market structure and multiple equilibria in airline markets," *Econometrica*, 77(6), 1791–1828.

Collard-Wexler, A. (2010): "Demand Fluctuations in the Ready-Mix Concrete Industry," Working paper, NYU Stern.

Doraszelski, U., and M. Satterthwaite (2007): "Computable Markov-Perfect Industry Dynamics: Existence, Purification, and Multiplicity," Working paper, Harvard University.

Douc, R., O. Cappé, and E. Moulines (2005): "Comparison of Resampling Schemes for Particle Filtering," in *In 4th International Symposium on Image and Signal Processing and Analysis (ISPA)*, pp. 64–69.

Dutta, P. K., and R. K. Sundaram (1998): "The Equilibrium Existence Problem in General Markovian Games," in *Organizations with Incomplete Information*, ed. by M. Majumdar, pp. 159–207. Cambridge University Press.

Ellickson, P., and S. Misra (2008): "Supermarket pricing strategies," *Marketing Science*, 27(5), 811–828.

Fernandez-Villaverde, J., and J. F. Rubio-Ramirez (2005): "Estimating Dynamic Equilibrium Economies: Linear versus Nonlinear Likelihood," *Journal of Applied Econometrics*, 20, 891–910.

Flury, T., and N. Shephard (2010): "Bayesian Inference Based Only on Simulated Likelihood: Particle Filter Analysis of Dynamic Economic Models," forthcoming in Econometric Theory.

Gallant, A. R., H. Hong, and A. Khwaja (2010): "Dynamic Entry with Cross Product Spillovers: An Application to the Generic Drug Industry," Working paper, Yale University.

Gallant, A. R., and G. Tauchen (1998): "Reprojecting Partially Observed Systems with Application to Interest Rate Diffusions," *Journal of the American Statistical Association*, 93, 10–24.

Gebhardt, W. R., C. M. C. Lee, and B. Swaminathan (2001): "Toward and Implied Cost of Capital," *Journal of Accounting Research*, 39, 135–176.

Golub, G. H. (1973): "Some modified matrix eigenvalue problems," *SIAM Review*, 15, 318–334.

Golub, G. H., and J. H. Welsch (1969): "Calculation of Gaussian quadrature rules," *Mathematics of Computation*, 23, 221–230.

Gowrisankaran, G., and J. Stavins (2004): "Network Externalities and Technology Adoption: Lessons from Electronic Payments," *RAND Journal of Economics*, 35(2), 260–276.

Haile, P., A. Hortaçsu, and G. Kosenok (2008): "On the empirical content of quantal response equilibrium," *American Economic Review*, 98(1), 180–200.

Hastie, T., R. Tibshirani, and J. Friedman (2009): *The Elements of Statistical Learning, Second Edition*. Springer.

Ho, K. (2009): "Insurer-provider networks in the medical care market," *American Economic Review*, 99(1), 393–430.

Hotz, J., and R. Miller (1993): "Conditional Choice Probabilties and the Estimation of Dynamic Models," *Review of Economic Studies*, 60, 497–529.

Howard, R. A. (1960): *Dynamic Programming and Markov Processes*. Wiley.

Imai, S., N. Jain, and A. Ching (2009): "Bayesian Estimation of Dynamic Discrete Choice Models," *Econometrica*, 77(6), 1865–1899.

Ishii, J. (2005): "Interconnection Pricing, Compatibility, and Investment in Network Industries: ATM Networks in the Banking Industry," Working paper, Stanford University.

Jacquier, E., N. G. Polson, and P. E. Rossi (1994): "Bayesian Analysis of Stochastic Volatility Models," *Journal of Business and Economic Statistics*, 12, 371–389.

Keane, M., and K. Wolpin (1997): "The Career Decisions of Young Men," *The Journal of Political Economy*, 105(3), 473–522.

Khwaja, A. (2001): "Health Insurance, Habits and Health Outcomes: A Dynamic Stochastic Model of Investment in Health," Ph.D. Dissertation, University of Minnesota.

Kuhn, M. (2006): "Notes on Numerical Dynamic Programming in Economic Applications," Working paper, CDSEM, University of Mannheim, Germany.

Magnac, T., and D. Thesmar (2002): "Identifying dynamic discrete decision processes," *Econometrica*, 70(2), 801–816.

Nekipelov, D. (2007): "Entry Deterrence and Learning Prevention on eBay," Working paper, Duke University.

Norets, A. (2009): "Inference in dynamic discrete choice models with serially correlated unobserved state variables," *Econometrica*, 77(5), 1665–1682.

Pakes, A., M. Ostrovsky, and S. Berry (2007): "Simple estimators for the parameters of discrete dynamic games (with entry/exit examples)," *RAND Journal of Economics*, pp. 373–399.

Pakes, A., J. Porter, K. Ho, and J. Ishii (2005): "Moment Inequalities and Their Application," Working paper, Harvard University.

Parthasarathy, T. (1973): "Discounted, Positive, and Noncooperative Stochastic Games," *International Journal of Game Theory*, 2(1), 25–37.

Pitt, M. K. (2002): "Smooth particle filters for likelihood evaluation and maximisation," Working paper, The University of Warwick.

——— (2010): "Some Properties of Particle Filter Markov Chain Monte Carlo Methods," Working paper, The University of Warwick.

Pitt, M. K., and N. Shephard (1999): "Filtering via Simulation: Auxiliary Particle Filters," *Journal of the American Statistical Association*, 94, 590–599.

Rust, J. (1994): "Structural Estimation of Markov Decision Processes," in *Handbook of Econometrics, Vol. 4*, ed. by R. Engle, and D. McFadden, pp. 3081–3143. North Holland.

Rust, J. (2006): "Dynamic Programming," Working paper, University of Maryland.

Ryan, S. (2005): "The Costs of Environmental Regulation in a Concentrated Industry," Working paper, MIT.

Rysman, M. (2004): "Competition Between Networks: A Study of the Market for Yellow Pages," *The Review of Economic Studies*, 71(2), 483–512.

Scott-Morton, F. (1999): "Entry Decisions into Generic Pharmaceutical Markets," *RAND Journal of Economics*, 30(3), 421–440.

Seim, K. (2006): "An empirical model of firm entry with endogenous product-type choices," *RAND Journal of Economics*, 37(3), 619–640.

Sweeting, A. (2006): "Coordination Games, Multiple Equilibria and The Timing of Radio Commercials," *Journal of Economics and Management Strategy*, 15(4), 909–942.

Tamer, E. (2003): "Incomplete Simultaneous Discrete Response Model with Multiple Equilibria," *Review of Economic Studies*, 70(1), 147–165.

Table 1. Parameter Estimates for the Fully Rational Model
Blind Proposal, Multinomial Resampling

Parameter		Unconstrained			Constrained: β			Constrained: β & p_a		
	value	sm	md	lg	sm	md	lg	sm	md	lg
μ_c	9.70	10.10	9.72	9.68	9.94	9.67	9.68	9.86	9.72	9.68
		(0.15)	(0.12)	(0.06)	(0.19)	(0.11)	(0.06)	(0.18)	(0.12)	(0.06)
ρ_c	0.90	0.58	0.86	0.92	0.69	0.92	0.91	0.69	0.85	0.91
		(0.25)	(0.09)	(0.03)	(0.26)	(0.05)	(0.03)	(0.25)	(0.11)	(0.03)
σ_c	0.10	0.16	0.09	0.09	0.17	0.08	0.10	0.15	0.09	0.10
		(0.05)	(0.03)	(0.01)	(0.06)	(0.03)	(0.01)	(0.07)	(0.03)	(0.01)
μ_r	10.00	9.87	9.98	9.96	9.88	9.99	9.98	9.84	9.99	9.99
		(0.10)	(0.03)	(0.02)	(0.10)	(0.03)	(0.02)	(0.13)	(0.06)	(0.02)
σ_r	2.00	1.95	1.97	1.98	2.02	2.00	2.02	2.04	2.00	2.03
		(0.09)	(0.05)	(0.01)	(0.08)	(0.02)	(0.02)	(0.10)	(0.03)	(0.01)
ρ_a	0.50	0.76	0.56	0.58	0.59	0.57	0.56	0.76	0.57	0.52
		(0.09)	(0.07)	(0.06)	(0.22)	(0.09)	(0.05)	(0.10)	(0.07)	(0.04)
κ_a	0.20	0.04	0.24	0.19	0.15	0.26	0.20	0.14	0.22	0.22
		(0.05)	(0.05)	(0.02)	(0.07)	(0.05)	(0.03)	(0.06)	(0.06)	(0.03)
β	0.83	0.90	0.95	0.87	0.83	0.83	0.83	0.83	0.83	0.83
		(0.07)	(0.04)	(0.04)						
p_a	0.95	0.97	0.94	0.95	0.96	0.94	0.95	0.95	0.95	0.95
		(0.02)	(0.01)	(0.01)	(0.02)	(0.01)	(0.01)			

The data were generated according to the fully rational model with parameters set as shown in the column labeled "value". For all data sets $T_0 = -160$. For the small data set $T = 40$; for the medium $T = 120$; and for the large $T = 360$. The estimate is the mean of the posterior distribution. The values below each estimate in parentheses are the standard deviation of the posterior. The prior is uninformative except for the following support conditions $|\rho_c| < 1$, $|\rho_a| < 1$, $0 < \beta < 1$, and $0 < p_a < 1$. The likelihood for μ_r and σ_r includes the observations from T_0 to 0. In the columns labeled constrained, the parameters β and p_a are constrained to equal their true values as shown in the table. The number of MCMC repetitions is 240,000 with every 25th retained for use in estimation.

Table 2. Parameter Estimates for the Boundedly Rational Model
Blind Proposal, Multinomial Resampling

Parameter		Unconstrained			Constrained β			Constrained β & p_a		
	value	sm	md	lg	sm	md	lg	sm	md	lg
μ_c	9.70	10.06	9.71	9.69	9.71	9.48	9.64	9.90	9.57	9.66
		(0.18)	(0.10)	(0.06)	(0.18)	(0.13)	(0.06)	(0.23)	(0.14)	(0.05)
ρ_c	0.90	0.80	0.92	0.91	0.83	0.95	0.90	0.73	0.94	0.92
		(0.13)	(0.03)	(0.02)	(0.13)	(0.04)	(0.03)	(0.20)	(0.03)	(0.03)
σ_c	0.10	0.31	0.08	0.09	0.13	0.06	0.09	0.13	0.07	0.09
		(0.13)	(0.02)	(0.01)	(0.06)	(0.02)	(0.02)	(0.05)	(0.02)	(0.01)
μ_r	10.00	9.84	9.96	9.96	9.91	9.99	9.92	9.82	10.00	9.94
		(0.07)	(0.02)	(0.03)	(0.08)	(0.02)	(0.04)	(0.14)	(0.02)	(0.04)
σ_r	2.00	1.91	1.95	1.99	1.93	1.96	1.99	2.00	2.01	2.00
		(0.09)	(0.04)	(0.03)	(0.05)	(0.05)	(0.02)	(0.09)	(0.05)	(0.02)
ρ_a	0.50	0.22	0.47	0.52	0.72	0.56	0.55	0.78	0.55	0.57
		(0.15)	(0.13)	(0.06)	(0.15)	(0.06)	(0.07)	(0.06)	(0.07)	(0.05)
κ_a	0.20	0.01	0.25	0.19	0.19	0.36	0.20	0.10	0.32	0.19
		(0.14)	(0.05)	(0.02)	(0.07)	(0.08)	(0.03)	(0.10)	(0.07)	(0.02)
β	0.83	0.61	0.95	0.85	0.83	0.83	0.83	0.83	0.83	0.83
		(0.28)	(0.04)	(0.06)						
p_a	0.95	0.97	0.93	0.95	0.97	0.94	0.95	0.95	0.95	0.95
		(0.02)	(0.01)	(0.01)	(0.02)	(0.01)	(0.01)			

The data were generated according to the fully rational model with parameters set as shown in the column labeled "value". For all data sets $T_0 = -160$. For the small data set $T = 40$; for the medium $T = 120$; and for the large $T = 360$. The estimate is the mean of the posterior distribution. The values below each estimate in parentheses are the standard deviation of the posterior. The prior is uninformative except for the following support conditions $|\rho_c| < 1$, $|\rho_a| < 1$, $0 < \beta < 1$, and $0 < p_a < 1$. The likelihood for μ_r and σ_r includes the observations from T_0 to 0. In the columns labeled constrained, the parameters β and p_a are constrained to equal their true values as shown in the table. The number of MCMC repetitions is 80,000 with every 25th retained for use in estimation.

Table 3. Parameter Estimates for the Fully Rational Model
Adaptive Proposal, Multinomial Resampling

Parameter		Unconstrained			Constrained					
					β			β & p_a		
	value	sm	md	lg	sm	md	lg	sm	md	lg
μ_c	9.70	10.00	9.82	9.77	9.93	9.74	9.70	9.85	9.73	9.65
		(0.24)	(0.07)	(0.05)	(0.12)	(0.07)	(0.06)	(0.15)	(0.09)	(0.05)
ρ_c	0.90	0.95	0.85	0.87	0.87	0.92	0.93	0.87	0.92	0.94
		(0.03)	(0.07)	(0.05)	(0.08)	(0.04)	(0.03)	(0.09)	(0.04)	(0.02)
σ_c	0.10	0.14	0.09	0.10	0.12	0.08	0.08	0.12	0.09	0.08
		(0.02)	(0.02)	(0.01)	(0.04)	(0.02)	(0.01)	(0.04)	(0.03)	(0.01)
μ_r	10.00	9.93	10.00	10.01	10.00	9.99	9.97	9.94	9.96	9.96
		(0.06)	(0.02)	(0.01)	(0.05)	(0.02)	(0.02)	(0.07)	(0.03)	(0.03)
σ_r	2.00	1.93	1.98	1.99	2.01	1.98	2.00	2.03	1.97	1.99
		(0.10)	(0.02)	(0.02)	(0.09)	(0.01)	(0.01)	(0.09)	(0.02)	(0.02)
ρ_a	0.50	-0.11	0.51	0.47	0.56	0.59	0.57	0.47	0.51	0.61
		(0.21)	(0.09)	(0.06)	(0.17)	(0.06)	(0.06)	(0.20)	(0.07)	(0.05)
κ_a	0.20	0.19	0.20	0.17	0.17	0.21	0.18	0.24	0.20	0.19
		(0.02)	(0.03)	(0.02)	(0.06)	(0.02)	(0.02)	(0.03)	(0.02)	(0.02)
β	0.83	0.87	0.95	0.92	0.83	0.83	0.83	0.83	0.83	0.83
		(0.10)	(0.03)	(0.04)						
p_a	0.95	0.95	0.94	0.95	0.96	0.95	0.95	0.95	0.95	0.95
		(0.01)	(0.01)	(0.01)	(0.02)	(0.01)	(0.01)			

The data were generated according to the fully rational model with parameters set as shown in the column labeled "value". For all data sets $T_0 = -160$. For the small data set $T = 40$; for the medium $T = 120$; and for the large $T = 360$. The estimate is the mean of the posterior distribution. The values below each estimate in parentheses are the standard deviation of the posterior. The prior is uninformative except for the following support conditions $|\rho_c| < 1$, $|\rho_a| < 1$, $0 < \beta < 1$, and $0 < p_a < 1$. The likelihood for μ_r and σ_r includes the observations from T_0 to 0. In the columns labeled constrained, the parameters β and p_a are constrained to equal their true values as shown in the table. The number of MCMC repetitions is 80,000 with every 25th retained for use in estimation.

Table 4. Parameter Estimates for the Fully Rational Model
Adaptive Proposal, Systematic Resampling

Parameter		Unconstrained			Constrained					
					β			β & p_a		
	value	sm	md	lg	sm	md	lg	sm	md	lg
μ_c	9.70	9.87	9.82	9.72	9.81	9.78	9.68	9.78	9.76	9.65
		(0.24)	(0.07)	(0.05)	(0.12)	(0.07)	(0.06)	(0.15)	(0.09)	(0.05)
ρ_c	0.90	0.77	0.82	0.91	0.93	0.94	0.94	0.86	0.92	0.94
		(0.03)	(0.07)	(0.05)	(0.08)	(0.04)	(0.03)	(0.09)	(0.04)	(0.02)
σ_c	0.10	0.14	0.10	0.09	0.14	0.08	0.08	0.11	0.08	0.08
		(0.02)	(0.02)	(0.01)	(0.04)	(0.02)	(0.01)	(0.04)	(0.03)	(0.01)
μ_r	10.00	10.05	10.00	9.97	9.95	9.96	9.94	9.78	9.95	9.96
		(0.06)	(0.02)	(0.01)	(0.05)	(0.02)	(0.02)	(0.07)	(0.03)	(0.03)
σ_r	2.00	1.94	1.99	1.99	1.93	1.97	2.01	2.07	1.98	1.97
		(0.10)	(0.02)	(0.02)	(0.09)	(0.01)	(0.01)	(0.09)	(0.02)	(0.02)
ρ_a	0.50	0.61	0.53	0.56	0.41	0.36	0.61	0.71	0.58	0.64
		(0.21)	(0.09)	(0.06)	(0.17)	(0.06)	(0.06)	(0.20)	(0.07)	(0.05)
κ_a	0.20	0.21	0.22	0.18	0.20	0.18	0.18	0.17	0.19	0.18
		(0.02)	(0.03)	(0.02)	(0.06)	(0.02)	(0.02)	(0.03)	(0.02)	(0.02)
β	0.83	0.93	0.96	0.90	0.83	0.83	0.83	0.83	0.83	0.83
		(0.10)	(0.03)	(0.04)						
p_a	0.95	0.96	0.94	0.95	0.95	0.93	0.95	0.95	0.95	0.95
		(0.01)	(0.01)	(0.01)	(0.02)	(0.01)	(0.01)			

The data were generated according to the fully rational model with parameters set as shown in the column labeled "value". For all data sets $T_0 = -160$. For the small data set $T = 40$; for the medium $T = 120$; and for the large $T = 360$. The estimate is the mean of the posterior distribution. The values below each estimate in parentheses are the standard deviation of the posterior. The prior is uninformative except for the following support conditions $|\rho_c| < 1$, $|\rho_a| < 1$, $0 < \beta < 1$, and $0 < p_a < 1$. The likelihood for μ_r and σ_r includes the observations from T_0 to 0. In the columns labeled constrained, the parameters β and p_a are constrained to equal their true values as shown in the table. The number of MCMC repetitions is 80,000 with every 25th retained for use in estimation.

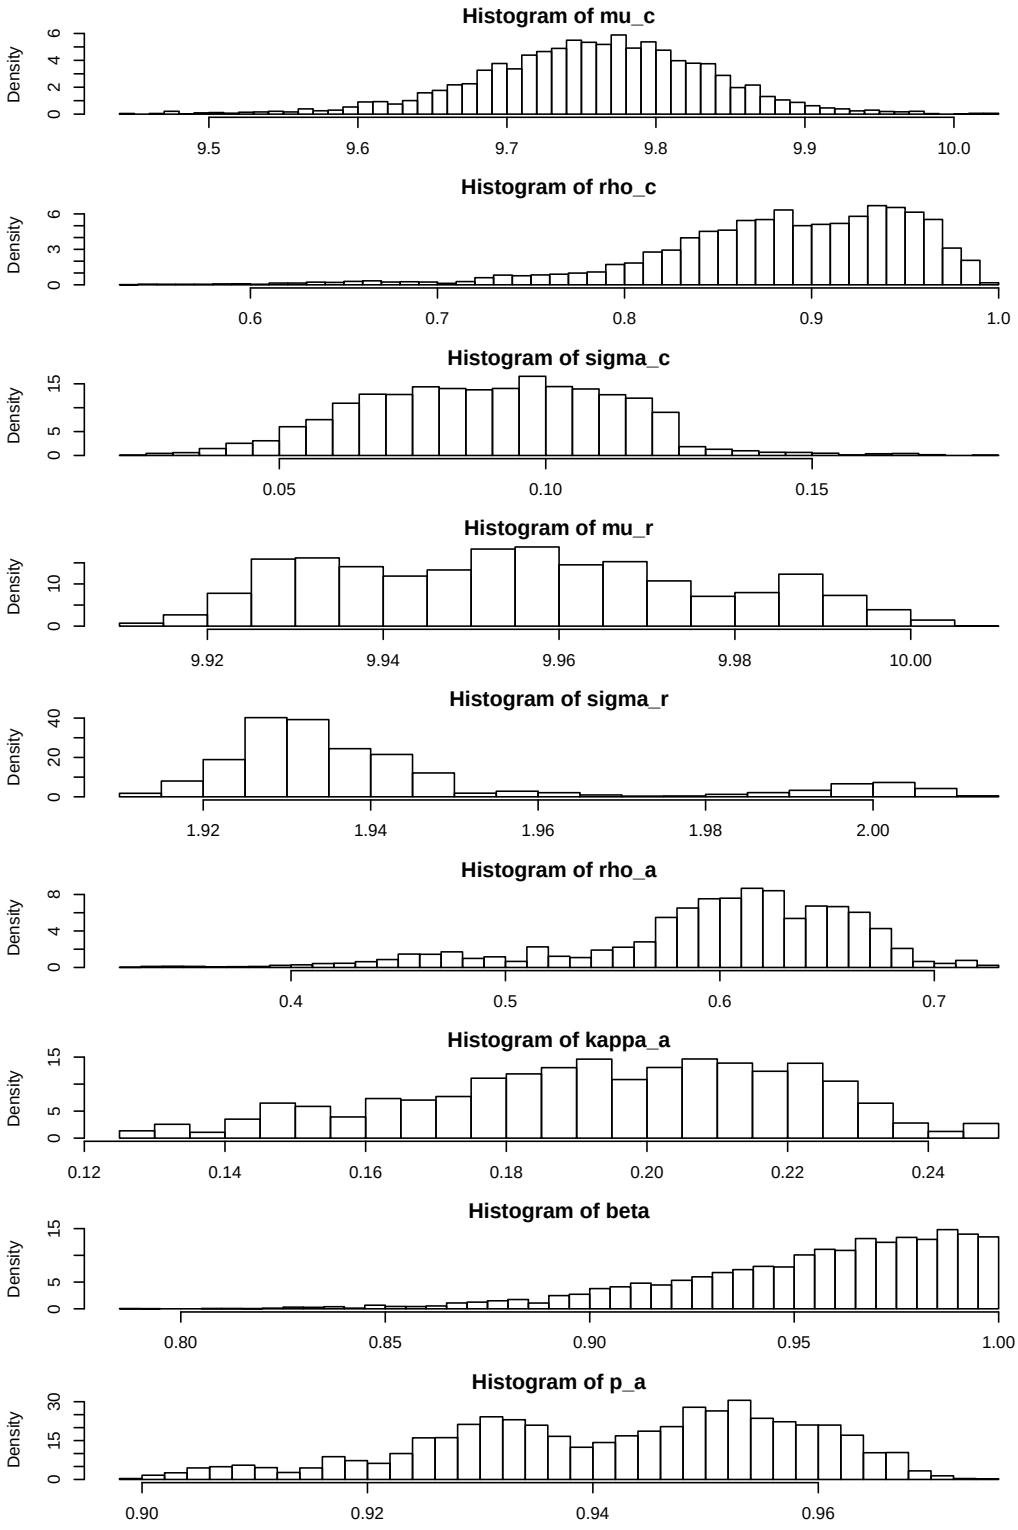

Figure 1. Fully Rational Model, Unconstrained, Blind Proposal. Shown are histograms constructed from the MCMC repetitions for the column labeled "Unconstrained," "md" in Table 1.

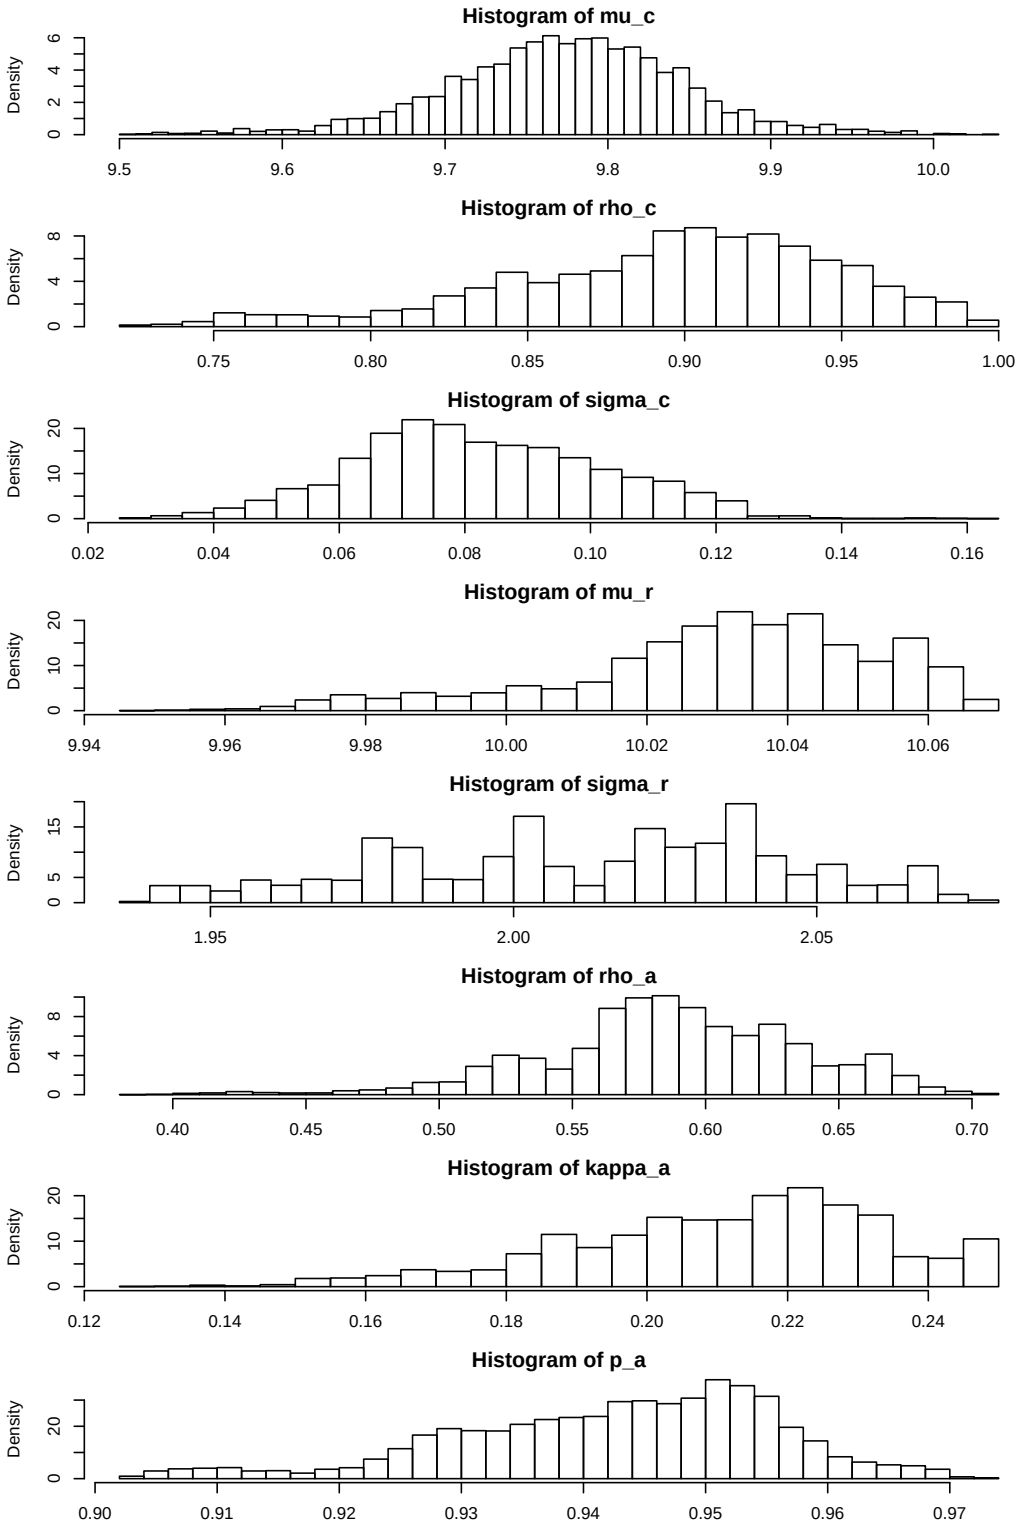

Figure 2. Fully Rational Model, β Constrained, Blind Proposal. Shown are histograms constructed from the MCMC repetitions for the column labeled "Constrained," "β," "md" in Table 1.

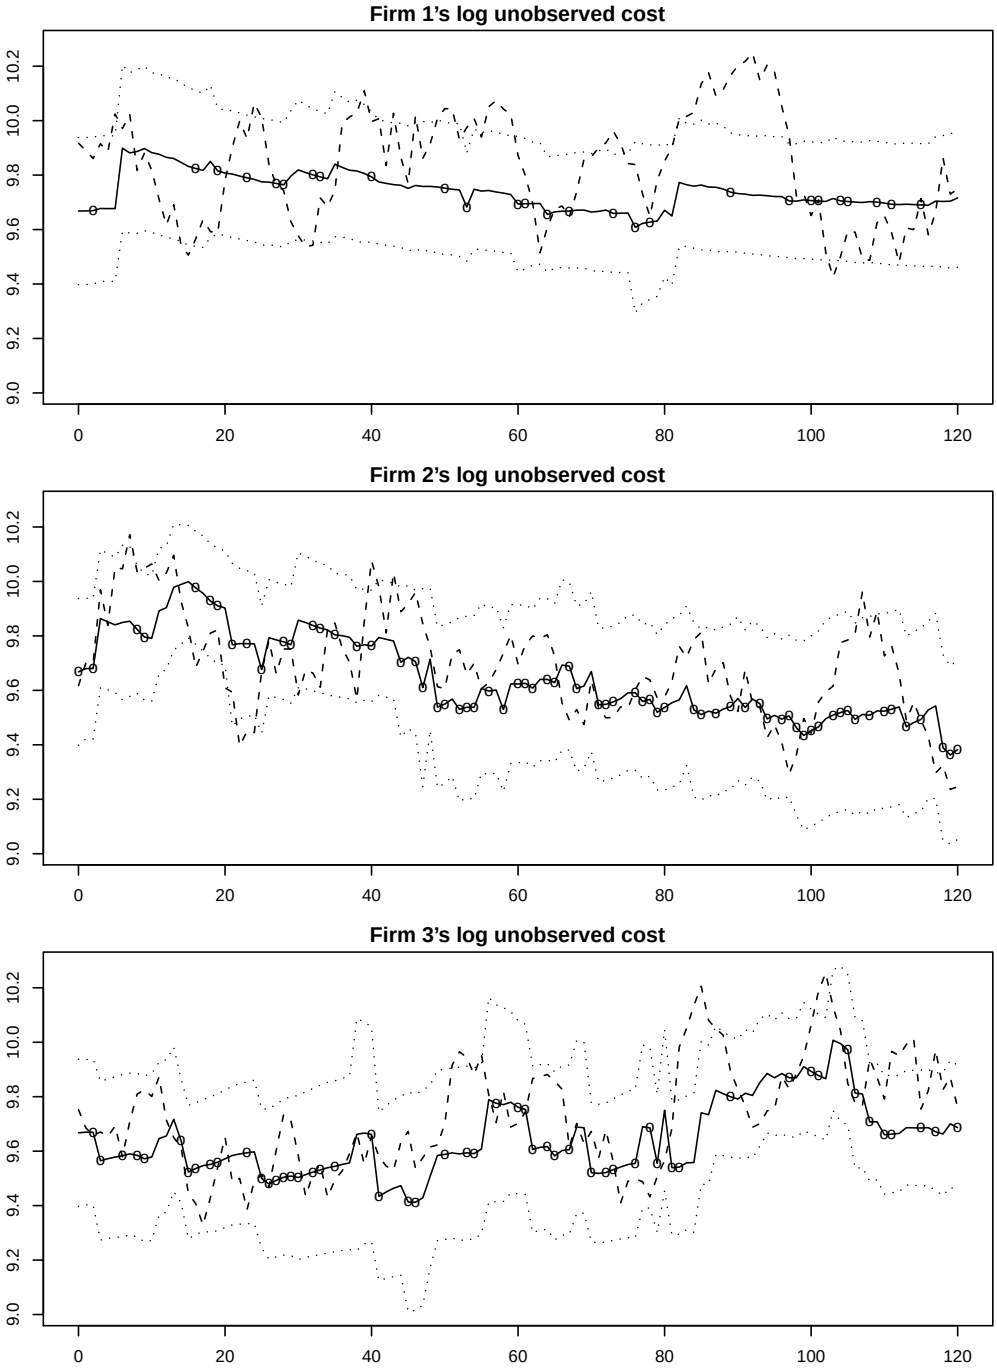

Figure 3. Fully Rational Cost Estimates, β Constrained, Blind Proposal. Shown are the particle filter estimates of each firm's unobserved computed from the MCMC repetitions for the column labeled "Constrained," "β," "md" in Table 1. The dashed line is the true unobserved cost. At each MCMC repetition, $\bar{c}_{ut} = \sum_{k=1}^{N} \tilde{x}_{1t}^{(k)}$ is computed for $t = 0, \ldots, T$; $T = 120$ and $N = 512$. The solid line is the average with a stride of 25 of the $\bar{c}_{ut}$ over 240,000 MCMC repetitions. The dotted lines are $\pm$ 1.96 standard deviations about the solid line. The circles indicate that the firm entered the market at time t. The sum of the norms of the difference between the solid and dashed lines is 0.186146.

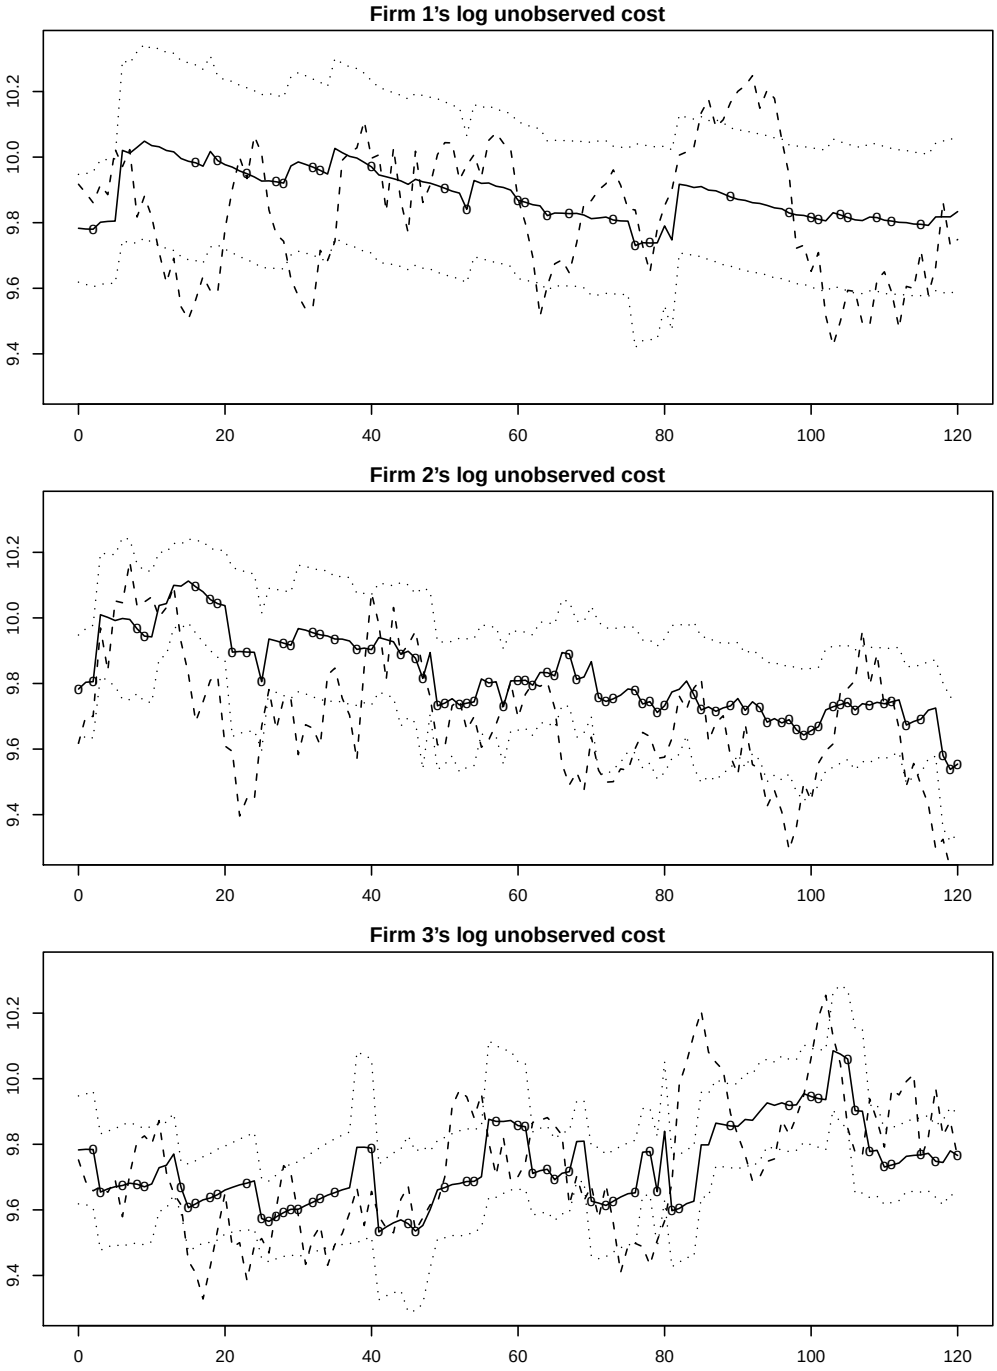

Figure 4. Fully Rational Cost Estimates, β Constrained, Adaptive Proposal. Shown are the particle filter estimates of each firm's unobserved computed from the MCMC repetitions for the column labeled "Constrained," "β," "md" in Table 4. The dashed line is the true unobserved cost. At each MCMC repetition, $\bar{c}_{ut} = \sum_{k=1}^{N} \tilde{x}_{1t}^{(k)}$ is computed for $t = 0, \ldots, T$; $T = 120$ and $N = 512$. The solid line is the average with a stride of 25 of the $\bar{c}_{ut}$ over 80,000 MCMC repetitions. The dotted lines are $\pm$ 1.96 standard deviations about the solid line. The circles indicate that the firm entered the market at time t. The sum of the norms of the difference between the solid and dashed lines is 0.169411.