Bayesian Estimation of Games with Incomplete Information

Christopher Ferrall *
Jean-François Houde †
Susumu Imai ‡, and
Maxwell Pak §

Very preliminary draft. Comments are very welcome.

June 30, 2008

Abstract

In this paper, we propose a new methodological framework to estimate a general class of games of incomplete information using Bayesian techniques. The central contribution of our approach is to construct an MCMC estimation algorithm which is immune to the problem multiple Nash equilibria, in the sense that it imposes very little assumptions about the equilibrium selection process despite using all the restrictions imposed by the theory. The only assumption required with respect to the equilibrium selection is that each observed market is repeatedly playing the same equilibrium over time. Contrary to most of the previous literature, we therefore allow observationally equivalent markets to "select" different equilibria.

We show that our estimation algorithm can easily accommodate heterogeneity parameters that are unobserved to the econometrician and have continuous distribution function.

After the estimation, our estimation algorithm only needs to be slightly modified to work as a simulation algorithm for policy experiments.

*Department of Economics, Queen's University, 94 University Avanue, Kingston, Ontario, K7L 3N6, Canada, e-mail:ferrallc@post.queensu.ca

†Department of Economics, University of Wisconsin-Madison, 1180 Observatory Drive, Madison, WI 83706-1393, e-mail:houdejf@ssc.wisc.edu

‡Department of Economics, Queen's University, 94 University Avenue, Kingston, Ontario, K7L 3N6, Canada, e-mail:imais@econ.queensu.ca

§Department of Economics, Queen's University, 94 University Avenue, Kingston, Ontario, K7L 3N6, Canada, e-mail:pakm@econ.queensu.ca

1 Introduction

In this paper, we propose a new methodological framework to estimate a general class of games of incomplete information using Bayesian techniques. The central contribution of our approach is to construct an MCMC estimation algorithm which is immune to the problem multiple Nash equilibria, in the sense that it imposes very little assumptions about the equilibrium selection process despite using all the restrictions imposed by the theory. The only assumption required with respect to the equilibrium selection is that each observed market is repeatedly playing the same equilibrium over time.

There are two key insights of our approach. The first is to add to the structural model random payoffs parameters. By doing so, we propose a new Markovian stochastic algorithm which simulates from the distribution of Nash equilibrium solution of the game without solving for the Nash equilibrium. We then embed this algorithm into a Bayesian Metropolis-Hastings algorithm we can simultaneously construct the posterior distribution of Nash equilibrium strategies that is consistent with the actions selected by players in the data, and estimate the structural parameters of the model. Here, the second key insight of our algorithm is about how we draw the parameters from the posterior distribution. Conventionally, if we were to use the Metropolis-Hastings algorithm, we would first draw the parameters from the candidate distribution, and then, solve for all the multiple equilibria and the equilibrium selection probabilities, and then based on them evaluate the likelihood to derive the acceptance probability of the new parameter draw. Instead, we would draw the choice probability from a proposal density, then derive the parameters of the model that makes the choice probabilities drawn to be the Nash equilibrium, and then evaluate the Metropolis-Hastings acceptance rate of the choice probability draw. The reason why the procedure is simple is that essentially the problem of multiple equilibria arise because the relationship from the parameter of the model to the equilibrium choice probabilities is not a function, but a correspondence. But, it turns out that the relationship from the equilibrium choice probability to the parameters including the heterogeneity parameters is a function. In other words, the equilibrium choice probability is a sufficient statistics for the parameters including the heterogeneity parameters. Therfore, once the equilibrium choice probabilities are drawn, then given those, the equilibrium is uniquely determined. Furthermore the parameters that generates the equilibrium choice probability can be easily generated by the modified version of the inversion theorem used by Hotz-Miller and many others, without solving for any equilibria of the game. In the classical two step or finite k-step approach, it is necessary that the initial choice probability is an accurate estimator of the true equilibrium choice probability, thus the need for the large sample data and/or the kernel techniques to try to overcome the small sample problems. In Bayesian approach we propose, since we draw the choice probabilities from a proposal density, as part of the MCMC iterations, there is no need for any consistent initial guess of the choice probabilities at all. And as long as the model and the MCMC algorithm is appropriately constructed to satisfy all the standard condition for convegence, it is fairly straightforward to prove that the Markov Chain is going to converge to the true posterior.

Moreover, like many MCMC estimation algorithm, we show that our estimation algorithm can easily accommodate heterogeneity parameters that unobserved to the econometrician and have continuous distribution function. The method is very similar in spirit to those by McCullogh and Rossi (1994) and many others, which draw different random effects parameters for each market to construct the MCMC loop, except for the fact that for each marekt we draw different choice probability for each market.

In contrast, the previous literature addressed the problem of multiple equilibria either by imposing additional equilibrium selection restrictions on the game (e.g. Bresnhan and Reiss (1990, 1991), Berry (1992), Bajari, Hong and Ryan(2007)), by assuming that the data is generated by a single equilibrium (e.g. Moro and Norman (2005), Aguirregabiria and Mira (2006), Bajari, Benkard and Levin (2006)), or by using only necessary equilibrium conditions of the model to partially identify the parameters (e.g. Ciliberto and Tamer (2007), Pakes, Porter, Ho, Ishii (2007)).

We contribute to this literature by proposing an estimation algorithm that is substantially less restrictive about the (unobserved) equilibrium selection rule used by players, while still using the full solution of the theoretical model to identify the parameters.

It is also well known among Bayesian econometricians that the MCMC algorithm is very effective in drawing from the high dimentional posterior distribution. The primary reason is that researchers can very easily draw the parameters sequentially from blocks of small number of parameters given other parameters fixed.

In the similar spirit, that our MCMC estimation will be computationally very effective when the number of choice probability that needs to be drawn becomes large, because we can also use the Metropolis-within Gibbs algorithm to sequentially estimate the model with large state spaces in blocks of draws where during each block we only need to draw a small number of choice probabilities. Thus, just like the way MCMC algorithm has dramatically reduced the computational burden of evaluating a high dimentional posterior distribution in Bayesina estimation, our algorithm would substantially reduce the curse of dimensionality in estimating the game with many players and state variables.

Another advantate of our algorithm is that after we estimate the parameters of the model, it is very straightforward to simulate the model. We would repeat the same MCMC routine until convergence with only a very minor modification. That is, we only draw the choice probabilities and take the parameters as given. Also, the "likelihood" which is used to evaluat the Metropolis-Hastings acceptance probability is only a function of the random effects parameters, not of the choice probability. If we then are interested in making a policy recommendation, all we need to do is to run the above MCMC routine again with the perturbed policy parameter. In short, our algorithm provides a comprehensive package of straightforward estimation, simulation and policy experiments of the games of incomplete information and unobserved heterogeneity, both static and dynamic.

Therefore, common criticism towards the choice probabilty based estimation algorithm proposed by Hotz, Miller and others, that they do not provide any solution for simulating the game after the estimation does not apply to our algorithm.

The remaining of the paper is organized as followed. The next section describes the class of models that we are concerned with. Section 3 and **??** introduces the MCMC algorithm and discusses the theoretical properties of the estimates. In Section 4 we discuss two numerical examples to illustrate the algorithm. Finally, the last section discusses extensions and the general applicability of our approach.

2 Model notation and assumptions

We first discuss the estimation of discrete games of incomplete information, similar to the ones studied in ?), ?), ?), ?), ?), ?), and ?). We then extend the model to include continuous market level random effects as well.

The structure of the game is described as follows. A market is composed of I players who face N options, $A = \{a^1, ..., a^N\}$, in periods $t \in \mathbb{Z}_+$. Each player i bases her action $a_{it} \in A$ on a publicly observed variable $x_t \in X$, where X is discrete, i.e. $X = \left\{x^1, ..., x^S\right\}$ and a privately observed utility shock $\epsilon_{it} = (\epsilon_{it}^1, ..., \epsilon_{it}^N) \in \mathbb{R}^N$. Letting $\epsilon_t = (\epsilon_{1t}, ..., \epsilon_{It})$, the state of the system is given by the realization (x_t, ϵ_t). We assume that x_t's evolve in a Markovian fashion and use $g(x, x'|a)$ to denote the 1-step transition probability from state x to x' when action a is chosen. In particular, we assume that the transition probabilities depend only on the current state and the current action choice. We assume that ϵ_{it}^n's are i.i.d. and have full support on the real line, as well as being independent from x_t's. We use f_ϵ to denote the density function of ϵ_t. When convenient, we also use ϵ_{a_i} to denote the random shock associated with action a_i.

Denote a_{-it} to be the action of players other than player i at period t. That is,

$$a_{-it} \equiv (a_{1t}, ..., a_{i-1\ t}, a_{i+1\ t}, ..., a_{It})$$

The period t utility player i receives from action profile $a_t = (a_{it}, a_{-it})$ in state (x_t, ϵ_t) is assumed to be additive in the random shock, and is given by

$$u_i(a_t, x_t, \epsilon_t, \theta) \equiv \pi_i(a_t, x_t, \theta) + \epsilon_{a_{it}},$$

where $\theta \in \mathbb{R}^K$ is a vector of structural parameters. Since θ is fixed and known to the players, it will be suppressed in the notation unless its role needs to be emphasized.

Player i payoff from the outcome path $(\mathbf{a}, \mathbf{x}, \boldsymbol{\epsilon}) \equiv (a_t, x_t, \epsilon_t : t \in \mathbb{Z}_+)$, is given by

$$v_i(\mathbf{a}, \mathbf{x}, \boldsymbol{\epsilon}) \equiv \sum_{t=0}^{\infty} \beta^t u_i(a_t, x_t, \epsilon_t) = \sum_{t=0}^{\infty} \beta^t \left(\pi_i(a_t, x_t) + \epsilon_{a_{it}}\right),$$

where $\beta \in [0, 1)$ represents the players' common discount factor.[1]

Now, we define the Markovian (pure) strategy. We first introduce notations for the steady state. That is, let x be the state vector, i.e. $x \equiv (x_1, ..., x_I)$ and x' be the next period state vector defined similarly. Similarly, let $\epsilon \equiv (\epsilon_1, ..., \epsilon_I)$ where ϵ_i is a vector whose element is ϵ_{ixa_i} A Markovian (pure) strategy of player i is a mapping $s_i : X \times \mathbb{R}^N \to A$ that specifies the player's choice in each observable state (x, ϵ_{ix}).

Given the Markovian assumptions on the evolution of the states, any Markovian strategy profile $s = (s_1, ..., s_I)$, together with an initial measure, induces a Markov chain, $\boldsymbol{Z^s} = \left(\boldsymbol{Z_0^s}, \boldsymbol{Z_1^s}, \boldsymbol{Z_2^s}, ...\right)$ that takes values in $\left(A^I \times X \times \mathbb{R}^{IN}\right)^\infty$. Let $\boldsymbol{Z^s}(a_0, x_0, \epsilon_0)$ be the chain $\boldsymbol{Z^s}$ started with initial condition $\boldsymbol{Z_0^s} = (a_0, x_0, \epsilon_0)$, and let $\mu(a_0, x_0, \epsilon_0)$ denote the probability measure for $(\mathbf{a}, \mathbf{x}, \boldsymbol{\epsilon})$ conditinal on (a_0, x_0, ϵ_0). Player i's expected utility given s and an initial condition $\boldsymbol{Z_0^s} = (a_0, x_0, \epsilon_0)$ is

$$\begin{aligned} v_i(s, a_0, x_0, \epsilon_0) &\equiv E_{\mu(a_0,x_0,\epsilon_0)}\left[\sum_{t=0}^{\infty} \beta^t u_i(\boldsymbol{Z_t^s})\right] \\ &= u_i(a_0, x_0, \epsilon_0) + E_{\mu(a_0,x_0,\epsilon_0)}\left[\sum_{t=1}^{\infty} \beta^t u_i(a_t, x_t, \epsilon_t)\right]. \end{aligned}$$

The game has a Markov Structure. Here we focus on the Markov equilibria, that is, we assume that the firms play Markov strategies[2].

Assumption 1. *The private information ϵ_{ixa_i} is distributed jointly normally with zero mean and positive definite variance covariance matrix, with its support being on the real line. The distribution function of ϵ is continuously differentiable.*

Given the random strategy $s_i(x, \epsilon_i)$ for each player i, we can define a conditional choice

[1]When $\beta = 0$, $\sum_{t=0}^{\infty} \beta^t u_i(a_t, x_t, \epsilon_t) \equiv u_i(a_0, x_0, \epsilon_0)$, and the game is interpreted to mean a static game.

[2]We heavily borrow the notations and the definition of the equilibrium in the choice space from Aguirreagabiria and Mira (2007)

probability associated with it. They are defined as follows:

$$p_i^s(a_i|x) = \Pr\left(s_i(x,\epsilon_i) = a_i\right) = \int \mathbf{1}\left(s_i(x,\epsilon_i) = a_i\right) df_\epsilon\left(\epsilon_i\right)$$

and

$$p_i^s\left(x\right) = \left(p_i^s(a^1|x), p_i^s(a^2|\,x), ..., p_i^s(a^N|x)\right), \quad p_i^s = \left(p_i^s(x^1), ..., p_i^s(x^S)\right)$$

Then, given ϵ_{-i} being private information unknown for player i, the deterministic part of per period profit can be expressed as follows.

$$\pi_i^s(a_i, p_{-i}^s, x) \equiv \sum_{a_{-i}\in A^{I-1}} \left[\prod_{j\neq i} p_j^s(a_j|x)\right] \pi_i(a, x)$$

Similarly, the transition probability for player i can also be expressed as follows.

$$G_i^s\left(x'|x, p_{-i}^s, a_i\right) = \sum_{a_{-i}\in A^{I-1}} \left[\prod_{j\neq i} p_j^s(a_j|\,x)\right] g\left(x'|x, a_i\right)$$

The Bellman equation of the dynamic discrete choice problem of the individual can be expressed as follows:

$$V_i^s\left(x, \epsilon_i, p_{-i}^s\right) = max_{a_i\in A}\left\{\pi_i^s(a_i, p_{-i}^s, x) + \epsilon_{ia_ix} + \beta\sum_{x'\in X}\left[\int V_i^s\left(x', \epsilon_i', p_{-i}^s\right) f_\epsilon\left(\epsilon_i'\right) d\epsilon_i'\right] G_i^s\left(x'|x, p_{-i}^s, a_i\right)\right\}$$

where $V_i^s\left(x, \epsilon_i, p_{-i}^s\right)$ is the value function of the optimal decision problem of individual i. Let the choice specific value function be defined as follows.

$$V_i^s\left(a_i, p_{-i}^s, x, \epsilon_i\right) = \pi_i^s(a_i, p_{-i}^s, x) + \epsilon_{ia_ix} + \beta\sum_{x'\in X}\left[\int V_i^s\left(x', \epsilon_i', p_{-i}^s\right) f_\epsilon\left(\epsilon_i'\right) d\epsilon_i'\right] G_i^s\left(x'|x, p_{-i}^s, a_i\right)$$

Then,

$$V_i^s\left(x, \epsilon_i, p_{-i}^s\right) = \max_{a_i\in A}\left\{V_i^s\left(a_i, x, \epsilon_i, p_{-i}^s\right)\right\}$$

The best response of player i to her belief p_{-i}^s is given by

$$R_i(x, \epsilon_i, p_{-i}^s) = \arg\max_{a_i\in A}\ V_i^s(a_i, p_{-i}^s, x, \epsilon_i).$$

A strategy profile $(s_1^*, ..., s_I^*)$ is a Markov Bayesian Nash equilibrium of this game if for all (x, ϵ_i),

$$s_i^*(x, \epsilon_i) = R_i(x, p_{-i}^s, \epsilon_i),$$

where $p_j^s(a_j, x) = Pr(s_j^*(x, \epsilon_j) = a_j)$ for all j, x, ϵ_j and a_j. That is, if every player in every state is best responding to her belief and the beliefs are consistent with the players' actual strategies.

Since the utility shocks are unobservable to a modeler, a more relevant definition of an equilibrium is given below.

Definition 1. *A vector of choice probabilities $p = (p_1, ..., p_I)$ is a Markov Perfect Bayesian Nash equilibrium (MPE) if for all i, x, and $a \in A$,*

$$p_i(a|x) = Prob\left(a = R_i(x, \epsilon_i, p_{-i})\right).$$

Thus, we follow Milgrom and Weber (1985) and Aguirregabiria and Mira (2007) and express the Markov Perfect Bayesian Nash Equilibrium in the probability space. That is, MPE is a fixed point of the folloing probability operator, i.e.

$$\Lambda_i\left(a_i, x; p_{-i}\right) = \int 1\left(a_i = \arg\max_{a \in A}\ V_i^s(a, x, \epsilon_i, p_{-i})\right) f_\epsilon(\epsilon_i) d\epsilon_i$$

$$\Lambda\left(p\right) = \left\{\Lambda_i\left(a_i, x; p_{-i}\right)\right\}$$

where Λ_i is called the best response probability function.

The main difficulty associated with the estimation of this model is the presence of multiple equilibria. Under some regularity conditions, there exists an equilibrium in pure strategies (see **?**) for more details). However, its uniqueness is not guaranteed, and multiple equilibria can be very prevalent in some cases (e.g. **?**)). We let $\mathcal{E}(x)$ denotes the set of equilibria, which is assumed to be countable. Assumption 2 describes the data generating process selecting the type of equilibrium being played.

Assumption 2. *In each market, a unique equilibrium is played.*

The previous assumption highlights the fact that the observed choices can be generated from multiple equilibria. For example, if repeated observations are available for one geographically isolated market and the game is dynamic, Assumption 2 requires only that one equilibrium is being played for every time period in this market. Therefore, contrary to the approach advocated by ?), we are not assuming that all cross-sectional markets are playing the same equilibria. Moreover, if the game is static and a market corresponds to a time period, this assumption does not rule out the possibility that two different equilibria are played in two subsequent time periods.

A typical panel data set generated from the above model is described by

$$y^d = (y_1^d, ..., y_M^d) = \left((x_m^d, a_{1m}^d, ..., a_{Im}^d) : m = 1, ..., M\right)$$

and includes M observed sequences of covariates and choices, each of length $T_m \geq 1$. For example, $a_{im}^d = (a_{im1}^d, ..., a_{imT_m}^d) \in A^{T_m}$, where a_{imt}^d is the action choice of player i in market m at time t. Emphasizing the role of the structural parameter in the notation, the likelihood contribution of market m is given by:

$$l_m(y_m^d|\theta) = \sum_{p^* \in \mathcal{E}} \pi(p^*) \left[\prod_i \left(\prod_t p_i^*(a_{imt}^d|x_{mt}^d, \theta) F_i^s(x_{mt+1}^d|x_{mt}^d, a_{imt}^d, p^*)^{I(1 \leq t \leq T_m - 1)}\right) H_i(x_{m1}^d, p^*)\right]. \quad (1)$$

where $H_i(x, p^*)$ is the stationary distribution of state x for player i, which satisfies the following condition.

$$H_i(p^*) = G_i H_i(p^*)$$

where

$$G_i(x, x') = \sum_a \left[p_i^*(a|x, \theta) F_i^s(x'|x, a, p^*)\right]$$

The presence of multiple equilibria is akin to the presence of finite-mixture market-level unobserved heterogeneity, where the mixture types relate to different equilibria. A key difference however is that the presence of multiple equilibria typically depends on the structural parameters θ. This feature seriously restricts the ability of the econometrician to specify the number of mixtures ex-ante and maximize the likelihood function[3].

[3] By pre-specifying the number of equilibria, the finite mixture approach requires that the model generates the same number of equilibria for all parameter values which is false in general. If this assumption is violated, the gradients of the likelihood with respect to the mixture parameters will not be defined everywhere and the estimation algorithm will fail. Another idea recently proposed by Kasahara and Shimotsu, and Houde and Imai is to extend the Hotz-Miller two-stage approach and estimate the number of equilibria and the corresponding choice probabilities in the first stage based on a finite mixture model. Then, in the second stage recover the structural parameters. This would require even more sample size than the original Hotz-Miller approach.

Rather than maximizing directly the likelihood function corresponding to the structural model, our approach consists of estimating a "perturbed" version of the model which turns out to be more tractable. In particular, we augment the model with a vector which represents unobserved heterogeneity , i.e. $c = (c_1, ..., c_I)$, where $c_i = (c_{ixa} : x \in X, a \in A)$. It is publicly observable (to the players). These additional parameters satisfy the following assumption.

Assumption 3. *The heterogeneity parameter vector c has mean zero and finite covariance matrix Σ_c, and its joint distribution is given by a parametric function $g_c(\,\cdot \mid \Sigma_c)$. Furthermore, for all i and x, $c_{ixa^N} = 0$.*

Here, we have normalized c_{ixa^N} to be zero to identify the per period payoff function in the discrete choice model, which will be discussed later. We define the perturbed current period deterministic payoff function as follows: Let

$$\hat{\pi}_i : A^I \times X \times \mathbb{R}^K \times \mathbb{R} \rightarrow \mathbb{R}$$

be an extension of $\pi_i : A^I \times X \times \mathbb{R}^K \rightarrow \mathbb{R}$ such that $\hat{\pi}_i(a, x, 0) = \pi_i(a, x)$ for all a and x. Now, let

$$u_i(a, x, \epsilon, c) \equiv \hat{\pi}_i(a, x, c_{ixa_i}) + \epsilon_{a_i}. \tag{2}$$

be the utility function for the perturbed game. To economize on the notation, we drop the hat from the notation and use $\pi_i(a, x, c_{ixa_i})$ rather than $\hat{\pi}_i(a, x, c_{ixa_i})$. Except for this extension of π_i, the earlier defined functions remain the same for the perturbed game.

We then assume that the relationship between the heterogeneity term and the per period profit function is monotonic. That is,

Assumption 4. *The profit function is continuously differentiable in c, and the derivative is uniformly positive* [4]*. That is,*

$$\frac{\partial \pi_i(a, x, c_{ixa_i})}{\partial c_{ixa_i}} > \eta_\pi$$

for some small $\eta_\pi > 0$

[4] What we need to assume is that the derivative is either strictly positive or negative, i.e. the profit function is strictly monotonic. In this case, without loss of generality, we assume that it is strictly increasing.

Let $c_{ix} = (c_{ixa^1}, .., c_{ixa^{N-1}})$ and $p_i(x) = (p_i(a^1|x), ..., p_i(a^{N-1}|x))$. The following assumption guarantees that there is an well-defined function mapping p to c.

Assumption 5. *For the choice probabilities p_{ixa}, we assume that for any i, x, and a, $p_{ixa} \in [\delta_p, 1 - \delta_p]$ for small $\delta_p > 0$.*

Assumptions 4 and 5 together imply that

$$\frac{\partial \pi_i^s(a, x, c_{ixa_i})}{\partial c_{ixa_i}} > \delta_p \eta_\pi$$

This assumption is somewhat stronger than those in Hotz and Miller and others where only $p_{ixa} \in (0, 1)$ is assumed. If we assume a compact parameter space and the model specification in such a way that the deterministic component of the static per period return is bounded, then the above assumption is automatically guaranteed. Otherwise, we put the above restriction to the prior to the choice probability. That is, we assume the prior distribution of the choice probability to be uniformly distributed, i.e. $U[\delta_p, 1 - \delta_p]$

Assumption 6. *We assume that the transition probability to the next period state x' conditional on the current state x and action a does not depend on the current private information shock ϵ.*

The following proposition is similar to the Inversion Theorem by Hotz and Miller, and others.

Proposition 1. *Suppose Assumptions 1 to 6 hold. Then, given a vector of choice probabilities $p = (p_1, ..., p_I)$ there exists a unique vector of heterogeneity parameters c, with $c_{ixa^N} = 0$ for all i and x, such that p is a Bayesian-Nash equilibrium of the perturbed game. Also the function $c(p)$ is continuous. Suppose in addition, Assumption 4 holds. Then the function $c(p)$ is differentiable in p.*

Proof. Let $\Delta = \{(p_1, ..., p_N) : p_n > 0$ for all n and $\sum_n p_n = 1\}$ be the interior of a $N - 1$ dimensional simplex. Let $\widetilde{\Delta} = \{(p_1, ..., p_{N-1}) : (p_1, ..., p_{N-1}, 1 - \sum_{n=1}^{N-1} p_n) \in \Delta\}$ be the set consisting of first $N - 1$ components of elements of Δ. For each player i and state $x \in X$, suppose we fix $v_{ix}^N = 0$ as normalization. Define a vector-valued function $h_{ix} : \mathbb{R}^{N-1} \to \widetilde{\Delta}$, by

$$h_{ix}^n(v_{ix1}, ..., v_{ix,N-1}) = \text{Prob}(\epsilon_{in} - \epsilon_{ik} \geq v_{ixk} - v_{ixn} \text{ for all } k = 1, ..., N-1 \text{ and } \epsilon_{in} - \epsilon_{iN} \geq -v_{ixn})$$

Then, by construction,

$$p_i(a^n|x) = h_{ix}^n(v_{ix1}, ..., v_{ix,N-1}),\ n = 1, ..., N-1$$

Notice that

$$\begin{aligned} 1 - \sum_{n=1}^{N-1} h_{ix}^n &= 1 - \sum_{n=1}^{N-1} \Pr\left(n = \arg\max_{k=1,..,N} \{v_{ixk} + \epsilon_{ixk}\}\right) \\ &= \Pr\left(N = \arg\max_{k=1,..,N} \{v_{ixk} + \epsilon_{ixk}\}\right) \end{aligned}$$

By Hotz and Miller's inversion theorem, h_{ix} is invertible and its inverse function h_{ix}^{-1} is differentiable. This means that for any choice probabilities $p_{ix} = (p_{ix1}, ..., p_{ix,N-1}, 1 - \sum_{n=1}^{N-1} p_{ixn})$, $v_{ix} = h^{-1}(p_{ix1}, ..., p_{ix,N-1})$ is the unique relative utility values that rationalize the choice probabilities p_{ix}.

Given $dF(\epsilon_{ix})$ being continuous,

$$\begin{aligned} & \int \left[\max_{n=1,...,N} \{v_{ixn} + \epsilon_{ixn}\}\right] dF(\epsilon_{ix}) \\ = & \sum_{n=1,...,N} p_i(a^n|x) \left[v_{ixn} + \int \epsilon_{ixn} \mathbf{1}(\epsilon_{ixn} - \epsilon_{ixk} > v_{ixk} - v_{ixn}; \forall k \neq n) dF_\epsilon(\epsilon_{ix})\right] \\ \equiv & \lambda(v_{ix}) \end{aligned}$$

is also continuously differentiable. We denote

$$\lambda(v_i) = [\lambda(v_{ix^1}), ..., \lambda(v_{ix^S})]' = \left[\lambda(h^{-1}\left(p_{-N}\left(x^1\right)\right)), ..., \lambda\left(h^{-1}\left(p_{-N}\left(x^S\right)\right)\right)\right]'$$

Now, let $g(a^N)$ be a matrix whose (l, k) th element is $g(x^l|x^k, a^N)$. Similarly, let $\pi_i(a^N, p_{-i}, 0)$ be a vector whose lth element is $\pi_i(a^N, x^l, p_{-i}, 0)$ Now, consider the following function

$$\widehat{V}\left(p_{i,-N}, p_{-i}\right) = \left[I - \beta g(a^N)\right]^{-1} \left[\pi_i(a^N, p_{-i}, 0) + \beta\lambda(v_i)\right].$$

where

$$\widehat{V}\left(p_{i,-N},p_{-i}\right)\equiv(\widehat{V}\left(x^{1},p_{i,-N},p_{-i}\right),\ldots,\widehat{V}\left(x^{S},p_{i,-N},p_{-i}\right))'$$

Then,

$$\widehat{V}\left(x,p_{i,-N},p_{-i}\right)=\pi_{i}(a^{N},x,p_{-i},0)+\beta\sum_{x'}\left[\widehat{V}\left(x',p_{i,-N},p_{-i}\right)+\lambda(v_{ix'})\right]g(x'|x,a^{N}).$$

Furthermore, for any $n=1,...,N-1$, let

$$\widetilde{\pi}_{xa^{n}}=\widehat{V}\left(x,p_{i,-N},p_{-i}\right)+v_{ixa^{n}}\left(p_{i,-N}\right)-\beta\sum_{x'\in X}\left[\widehat{V}\left(x,p_{i,-N},p_{-i}\right)+\lambda(v_{ix'}))\right]g\left(x'|x,a_{i},p_{-i}\right)$$

Then, if we set $\overline{V}(a^{n},x,p_{i})$ as follows:

$$\begin{aligned}\overline{V}(a^{n},x,p_{i,-N},p_{-i}) &\equiv \widehat{V}\left(x,p_{i,-N},p_{-i}\right)+v_{ixa^{n}}\left(p_{i,-N}\right),\text{ for } i=1,...,N-1\\ \overline{V}(a^{N},x,p_{i,-N},p_{-i}) &\equiv \widehat{V}\left(x,p_{i,-N},p_{-i}\right)\end{aligned}$$

Then, we will show below that $\overline{V}(a^{n},x,p_{i})$ is the derministic value of choosing a^{n} given the current deterministic return being $\pi_{ixa^{N}}=\pi_{i}(a^{N},p_{-i},x,0)$ for the choice a^{N} and $\pi_{ixa^{n}}=\widetilde{\pi}_{xa^{n}}$ for $n=1,...,N-1$. This is simply because from the definition of λ, we get

$$\widehat{V}\left(x,p_{i,-N},p_{-i}\right)+\lambda(v_{ix})=\int\max_{n}\left\{\overline{V}(a^{n},x,p_{i})+\epsilon_{ixn}\right\}dF(\epsilon_{ix})$$

and thus the following equation holds

$$\overline{V}(a^{n},x,p_{i})=\pi_{ixa^{n}}+\beta\sum_{x'\in X}\left[\int\max_{n}\left\{\overline{V}(a^{n},x',p_{i})+\epsilon_{ix'n}\right\}dF(\epsilon_{ix'})\right]g\left(x'|x,a_{i},p_{-i}\right)$$

That is, the value function defined as

$$V^{*}\left(x,\epsilon_{ix},p_{i}\right)=\max_{n}\left\{\overline{V}(a^{n},x,p_{i})+\epsilon_{ix'n}\right\}$$

is the fixed point of the following Bellman operator

$$T\left(V\right)=Max_{n}\left\{\pi_{ixa^{n}}+\epsilon_{ixn}+\beta\sum_{x'\in X}\left[\int V\left(x',\epsilon_{ix'},p_{i}\right)dF(\epsilon_{ix'})\right]g\left(x'|x,a^{n},p_{-i}\right)\right\}$$

Now, because only $\pi_{i}^{s}\left(a,x,c_{ixa},p_{-i}\right)$ is a function of c_{ixa}, and π_{i}^{s} has a strictly positive derivative with respect to c_{ixa}, for any p_{i}, there exists a unique c_{ixa} such that

$$\pi_{i}^{s}\left(a_{i},x,c_{ixa_{i}},p_{-i}\right)=\pi_{ixa_{i}}\text{ for } a_{i}\neq N$$

Because the above claim holds for any i given p_{-i}, for any p there exists c with $c_N = 0$ for which p is a Bayesian-Nash equilibrium. Now,

$$\frac{\partial \pi_i^s\left(x, c_i, p_{-i}\right)}{\partial c_{i,-N}} dc_{i,-N}$$
$$= \left\{ \frac{\partial \widehat{V}\left(x, p_{i,-N}, p_{-i}\right)}{\partial p_{i,-N}} + \frac{\partial v_{ixa_i}\left(p_{i,-N}\right)}{\partial p_{i,-N}} - \beta \sum_{x' \in X} \left[\frac{\partial \widehat{V}\left(x', p_{i,-N}, p_{-i}\right)}{\partial p_{i,-N}} + \lambda' \frac{\partial v_{ix'}\left(p_{i,-N}\right)}{\partial p_{i,-N}} \right] g\left(x'|x, a_i, p_{-i}\right) \right\}$$
$$dp_{i,-N}$$

Hence

$$\frac{dc_{i,-N}}{dp_{i,-N}} = \left[\frac{\partial \pi_i^s\left(c_i, p_{-i}\right)}{\partial c_{i,-N}} \right]^{-1}$$
$$\left[\frac{\partial \widehat{V}\left(p_{i,-N}, p_{-i}\right)}{\partial p_{i,-N}} + \frac{\partial v_i\left(p_{i,-N}\right)}{\partial p_{i,-N}} - \beta \left[\frac{\partial \widehat{V}\left(p_{i,-N}, p_{-i}\right)}{\partial p_{i,-N}} + \lambda' \frac{\partial v_i\left(p_{i,-N}\right)}{\partial p_{i,-N}}) \right]' g\left(a_i, p_{-i}\right) \right]$$

Since $\frac{\partial \pi_i^s}{\partial c_{i,-N}}$ is assumed to be diagonal and its derivative to be strictly postive and bounded, it is always invertible. Therefore, $\frac{dc_{i,-N}}{dp_{i,-N}}$ is well defined.

Next, we show that $dc_{i,-N}/dp_{i,-N}$ is a continuous function of p_{-i}. First, since $\frac{\partial \pi_i^s(c_i, p_{-i})}{\partial c_{i,-N}}$ is a continuous function of p_{-i}, the denominator is continuous in p_{-i}. Similarly, from the definition of $\widehat{V}\left(p_{i,-N}, p_{-i}\right)$, $\frac{\partial \widehat{V}\left(p_{i,-N}, p_{-i}\right)}{\partial p_{i,-N}}$ is continuous in p_{-i}. So is $\frac{\partial v_i\left(p_{i,-N}\right)}{\partial p_{i,-N}}$ and $\lambda' = \frac{d\lambda}{dv_i}$ and $g\left(x'|x, p_{-i}\right)$. Therefore, both denominator and numerator is continuous in p_{-i} and thus, claim holds. Now, for any vector Δp_{-N}, let $\Delta p_{-N}^j \equiv \left(\Delta p_{1,-N}, ..., \Delta p_{j,-N}, 0, , , 0\right)$, $j = 1, ..., N-1$ and $\Delta p_{-N}^0 = (0, , , 0)$.

$$\begin{aligned} c_{-N}\left(p_{-N} + \Delta p_{-N}\right) - c_{-N}\left(p_{-N}\right) &= \sum_{j=1}^{I} \left[c_{-N}\left(p_{-N} + \Delta p_{-N}^j\right) - c_{-N}\left(p_{-N} + \Delta p_{-N}^{j-1}\right) \right] \\ &= \sum_{j=1}^{I} \left\{ \left[\frac{\partial c_{j,-N}\left(p_j, p_{-j} + \Delta p_{-N}^{j-1}\right)}{\partial p_{j,-N}} \right] \Delta p_{j,-N} + r\left(\Delta p_{j,-N}\right) \right\} \end{aligned}$$

where

$$\frac{r\left(\Delta p_{j,-N}\right)}{\left|\Delta p_{j,-N}\right|} \to 0 \text{ as } \left|\Delta p_{j,-N}\right| \to 0$$

and because of continuity

$$\frac{\partial c_{j,-N}\left(p_j, p_{-j} + \Delta p_{-N}^{j-1}\right)}{\partial p_{j,-N}} \to \frac{\partial c_{j,-N}\left(p_j, p_{-j}\right)}{\partial p_{j,-N}} \text{ as } \left|\Delta p_{-N}^{j-1}\right| \to 0$$

Together, we obtain

$$\begin{aligned}
&\frac{\left|c_{-N}\left(p_{-N}+\Delta p_{-N}\right)-c_{-N}\left(p_{-N}\right)-\left[\partial c_{-N}\left(p\right)/\partial p_{-N}\right]\Delta p_{-N}\right|}{\left|\Delta p_{j,-N}\right|} \\
\leq\ & \sum_{j=1}^{I}\left|\frac{\partial c_{j,-N}\left(p_j,p_{-j}+\Delta p_{-N}^{j-1}\right)}{\partial p_{j,-N}}-\frac{\partial c_{j,-N}\left(p_j,p_{-j}\right)}{\partial p_{j,-N}}\right|+\sum_{j=1}^{I}\frac{\left|r\left(\Delta p_{j,-N}\right)\right|}{\left|\Delta p_{j,-N}\right|} \\
\rightarrow\ & 0 \text{ as } \left|\Delta p_{-N}\right|\rightarrow 0
\end{aligned}$$

Therefore, $c_{-N}\left(p_{-N}\right)$ is differentiable. □

Following Proposition 1, we can define the inverse function of the best-response mapping as follows.

$$c_{i,-N} \equiv B_i(p_i, p_{-i}, \theta)$$

Proposition 1 ensures that even though there is not a unique mapping from c to p, there is a **unique** mapping relating Nash equilibrium choice probabilities to heterogeneity parameters, i.e., from p to c. In other words, there can be multiple Nash equilibria associated by c, but there is only one vector c associated with each Nash equilibrium choice probabilities.

3 MCMC algorithm

The conditional likelihood function of market m of the perturbed model is

$$\begin{aligned}
&l_m(y_m^d|\theta) = \\
&\int_c \sum_{p^*(c,\theta)\in\mathcal{E}(c,\theta)} \pi^*(p^*(c,\theta))\left[\prod_i\left(\prod_t p_i^*(a_{imt}^d|x_{mt}^d,c,\theta)G_i^s(x_{mt+1}^d|x_{mt}^d,a_{imt}^d,p^*)^{I(1\leq t\leq T_m-1)}\right)\right. \\
&\left. H_i\left(x_{m1}^d,p^*\right)\right] df(c,\Sigma_c). \qquad (3)
\end{aligned}$$

where π^* is the 'equilibrium selection rule', which is the probability that one of the multiple equilibria is chosen. The formula of the likelihood function above illustrates the computational difficulty of evaluating it when there are unobserved heterogeneities. That is, to integrate over the random effects term c, at each simulation draw of the vector, one has to evaluate choice probabilities of all the possible equilibria. Also, it is important to notice that the number of equilibria, i.e. the number of elements in $\mathcal{E}(c,\theta)$ varies both with respect to c as well. This makes any kind of first step estimation of equilibrium choice probabilities impossible.

The existence of an inverse function from p to c implies that we can write the conditional likelihood function of the perturbed model without specifying the equilibrium selection probability. To see this, note that conditional on a vector Nash equilibrium choice probabilities p, the likelihood contribution of market m is given by:

$$\begin{aligned} & l_m(y_m^d|c,\theta)dc\Big|_{c=B(p^*,\theta)} \equiv \\ & \quad \left[\prod_i \left(\prod_t p_i^*(a_{imt}^d|x_{mt}^d,c,\theta)G_i^s(x_{mt+1}^d|x_{mt}^d,a_{imt}^d,p^*)^{I(1\le t\le T_m-1)}\right)\right. \\ & \quad \left. H_i\left(x_{m1}^d,p^*\right)\right] f(c|\Sigma_c)dc\Big|_{c=B(p^*,\theta)} \\ & = \left[\prod_i \left(\prod_t p_i(a_{imt}^d|x_{mt}^d)G_i^s(x_{mt+1}^d|x_{mt}^d,a_{imt}^d,p)^{I(1\le t\le T_m-1)}\right)\right. \\ & \quad \left. H_i\left(x_{m1}^d,p\right)\right] f(B(p,\theta)|\Sigma_c)\left|B_p(p,\theta)\right| dp. \end{aligned} \tag{4}$$

where we denote the RHS to be $l_m(y_m^d|p,\theta)dp$ The previous representation allows us to treat the Nash-equilibrium choice probabilities as a latent random variable and estimate the model using a MCMC algorithm with data-augmentation (see for instance **?**)). In particular, if $\pi_\theta(\theta)$ denote the prior distribution of the parameters, our objective is to learn the posterior distribution of $\{\theta\}$:

$$\begin{aligned} & \mathcal{P}(\theta|Y) \propto \\ & \quad \pi_\theta(\theta)\int_p \left[\prod_i \left(\prod_t p_i(a_{imt}^d|x_{mt}^d)G_i^s(x_{mt+1}|x_{mt}^d,a_{imt}^d,p)^{I(1\le t\le T_m-1)}\right)H_i(x_{m1}^d,p)\right] \\ & \quad f(B(p,\theta)|\Sigma_c)\left|B_p(p,\theta)\right| dp \end{aligned} \tag{5}$$

In order to simulate $\mathcal{P}(\theta|Y)$, we propose the following Metropolis-Hastings algorithm with data-augmentation.

Algorithm 1 (Metropolis-Hastings)**.** *Start the chain at* $\{p^{(0)},\theta^{(0)}\}$ *. At iteration* k *iterate over the following steps:*

1. *Data augmentation step. For each market m:*

   (a) *Sample candidate choice probability p^* from the proposal density q_p, i.e. $p^*_m \sim q_p(p^*_m|p^{(k)}_m)$.*

   (b) *Evaluate the likelihood contribution of market m at p^*_m using equation 4.*

   (c) *Update choice probabilities:*

$$p_m^{(k+1)} = \begin{cases} p^*_m & \text{with probability } \alpha(p^*_m, p^{(k)}_m) \\ p^{(k)}_m & \text{with probability } 1-\alpha(p^*_m, p^{(k)}_m) \end{cases}$$

   Where the acceptance probability is given by:

$$\alpha(p^*_m, p^k_m) = \min\left(\frac{l_m(y^d_m|p^*_m, \theta^{(k)})q_p(p^*_m|p^{(k)}_m)}{l_m(y^d_m|p^{(k)}_m, \theta^{(k)})q_p(p^{(k)}_m|p^*_m)}, 1\right) \tag{6}$$

2. *Parameter updating step. Conditional on $p^{(k+1)}$,*

   (a) *Sample $\theta^* \sim q_\theta(\theta^*|\theta^{(k)})$*

   (b) *Evaluate the conditional likelihood function at θ^*:*

$$L(y^d|\theta^*, p^{(k+1)}) = \pi_\theta(\theta)\prod_m l_m(y^d_m|p^{(k+1)}_m, \theta^*) \tag{7}$$

   (c) *Update parameter vector:*

$$\theta^{(k+1)} = \begin{cases} \theta^* & \text{with probability } \alpha(\theta^*, \theta^{(k)}) \\ \theta^{(k)} & \text{with probability } 1-\alpha(\theta^*, \theta^{(k)}) \end{cases}$$

   Where the acceptance probability is given by:

$$\alpha(\theta^*, \theta^{(k)}_v) = \min\left(\frac{L(y^d|\theta^*, p^{(k+1)})q_\theta(\theta^*|\theta^{(k)})\pi_\theta(\theta^*)}{L(y^d|\theta^{(k)}, p^{(k+1)})q_\theta(\theta^{(k)}|\theta^*)\pi_\theta(\theta^{(k)})}, 1\right) \tag{8}$$

Heuristically, the data-augmentation step plays two roles: (i) trace the distribution of Nash-equilibrium choice probabilities of the perturbed game, and (ii) select which equilibrium is more likely in each market.

The first is accomplished by imposing a common distribution for the heterogeneity parameters across markets and players. Since c_{iax} is mean zero, the algorithm is likely to reject any candidate choice probability which does not satisfy the Nash equilibrium condition of the original game. Moreover, the likelihood of accepting a candidate strategy which violates the true Nash equilibrium condition is decreasing in the variance Σ_c.

The second task is accomplished by using the information contained in the observed sequence of choices played. In particular, since the acceptance probability is function of the likelihood contribution of each market, a candidate choice probability is more likely to be accepted if it mimic the observed choices made by firms in the data. Consequently, the posterior distribution of choice probabilities will be differ across markets that are in different equilibria, and have more weight around the selected equilibria.

Proposition 2. *The sequence of parameters $\theta^{(m)}$ generated by the above MCMC algorithm converges to the true posterior.*

Proof. The invariant distribution from which we draw is proportional to

$$\begin{aligned}
&\mathcal{P}(\theta|Y) \propto \\
&\quad \pi_\theta(\theta) \int_p \left[\prod_i \left(\prod_t p_i(a^d_{imt}|x^d_{mt}) G^s_i(x_{mt+1}|x^d_{mt}, a^d_{imt}, p)^{I(1\leq t\leq T_m-1)} \right) H_i(x^d_{m1}, p) \right] \\
&\quad f(B(p,\theta)|\Sigma_c) \left| B_p(p,\theta) \right| dp
\end{aligned}$$

What the above Metropolis-Hastings algorithm does, is draw from the complete data version of the RHS function

$$\begin{aligned}
&\pi\left(p,\theta\right) \propto \\
&\quad \pi_\theta(\theta) \left[\prod_i \left(\prod_t p_i(a^d_{imt}|x^d_{mt}) G^s_i(x_{mt+1}|x^d_{mt}, a^d_{imt}, p)^{I(1\leq t\leq T_m-1)} \right) H_i(x^d_{m1}, p) \right] \\
&\quad f(B(p,\theta)|\Sigma_c) \left| B_p(p,\theta) \right|
\end{aligned}$$

Then, the chain produced is irreducible and aperiodic.

Denote $u = (u_1, ..., u_m, u_{m+1}) = (p_1, ..., p_m, \theta)$ and $f\left(u\right) = \pi\left(p,\theta\right)$. This proof extends Lemma 7.6 of Robert and Casella. Let the parameter space of u be Ξ. Consider $x^{(0)}$ an arbitrary starting point and $A \subset \Xi$ an arbitrary measurable set. Then, the connectedness of Ξ implies that there exist $m > 0$ and a sequence $x^{(i)} \in \Xi$ such that $x^{(m)} \in A$ and $\left\| x^{(i+1)} - x^{(i)} \right\| < \delta$, where the norm $\|\|$ is defined to be the sup norm, i.e. It is therefore possible to link $x^{(0)}$ and A through as sequence of open sets $B_\delta^{(i)}$such that for all $x, x' \in B_\delta^{(i)}$,$\|x - x'\| < \delta$. Now, let B be the compact set that contains $\left\{B_\delta^{(i)}\right\}_{i=1}^m$. Then, the likelihood is bounded and strictly positive in the set B. Now,

$$\begin{aligned} & P\left(X^{(i+1)} \in B_\delta^{(i+1)} | X^{(i)}\right) \\ = & \prod_k P\left(\left(X_{k-}^{(i)}, X_k^{(i+1)}, X_{k+}^{(i+1)}\right), X_k^{(i+1)} \in B_{\delta k}^{(i+1)} | X_{k-}^{(i)}, X_k^{(i)}, X_{k+}^{(i+1)}\right) \end{aligned}$$

Because the proposal density is at least partly random walk, for the proposal density of the kth Metropolis-Hastings draw, whose support is either the real line or $[\delta_p, 1 - \delta_p]$ for any positive $\widehat{\delta} > 0$ there exists positive number $\widehat{\varepsilon}$ such that for any x, $(x_{-k}, u) \in \Xi$

$$q\left(x_{k-}, u, x_{k+} | x_{k-}, x_k, x_{k+}\right) > \widehat{\varepsilon} \text{ if } |x_k - u| < \widehat{\delta}$$

Hence, if we let δ to satisfy $\delta < \widehat{\delta}/2$,then,

$$\begin{aligned} & P\left(\left(X_{k-}^{(i)}, X_k^{(i+1)}, X_{k+}^{(i+1)}\right), X_k^{(i+1)} \in B_{\delta k}^{(i+1)} | X_{k-}^{(i)}, X_k^{(i)}, X_{k+}^{(i+1)}\right) \\ \geq & \int_{B_{\delta k}^{(i+1)}} \min\left\{\frac{f\left(X_{k-}^{(i)}, X_k^{(i+1)}, X_{k+}^{(i+1)}\right)}{f\left(X_{k-}^{(i)}, X_k^{(i)}, X_{k+}^{(i+1)}\right)}, 1\right\} q\left(X_{k-}^{(i)}, X_k^{(i+1)}, X_{k+}^{(i+1)} | X_{k-}^{(i)}, X_k^{(i)}, X_{k+}^{(i+1)}\right) dX_k^{(i+1)} \\ \geq & \widehat{\varepsilon} \int_{B_{\delta k}^{(i+1)}} \min\left\{\frac{f\left(X_{k-}^{(i)}, X_k^{(i+1)}, X_{k+}^{(i+1)}\right)}{f\left(X_{k-}^{(i)}, X_k^{(i)}, X_{k+}^{(i+1)}\right)}, 1\right\} dX_k^{(i+1)} \\ \geq & \widehat{\varepsilon} \frac{\inf_B f(x)}{\sup_B f(x)} \lambda\left(B_{\delta k}^{(i+1)}\right) > \widehat{\varepsilon} \frac{\inf_B f(x)}{\sup_B f(x)} 2\delta > 0 \end{aligned}$$

Therefore,

$$P\left(X^{(i+1)} \in B_\delta^{(i+1)} | X^{(i)}\right) > 0$$

and

$$P\left(X^{(m)} \in A | X^{(0)}\right) \geq \prod_{i=1}^{m-1} P\left(X^{(i+1)} \in B_\delta^{(i+1)} | X^{(i)}\right) > 0$$

thus the Metropolis-Hastings within Gibbs algorithm is irreducible.

Furthermore, for any x and $X \in B_\delta(x)$,

$$P(A|X) \geq \prod_{k=1}^{p} \left[\widehat{\varepsilon} \frac{\inf_B f(x)}{\sup_B f(x)} \lambda\left(B_{\delta k} A_k\right)\right]$$

Hence, B_δ is small associated with the product measure of the uniform distribution over $B_\delta(x)$. Thus, the chain is aperiodic. Furthermore, because all conditional samplers are irreducible for any values of the fixed variables, Theorem 1 of Chan and Geyer (1994) applies and thus the chain is Harris recurrent.

Therefore, all conditions for Theorem 6.51 of Robert and Casella (2006) are satisfied, and thus

$$\lim_{n \longrightarrow \infty} \left\| K^n(x,.)\,\mu(dx) - \pi \right\|_{TV} = 0$$

for every initial distribution μ, where K is defined to be the transition probability of the Markov Chain.

Finally, notice that the posterior distribution of the parameter θ is the integration of the distribution $\pi(p,\theta)$, thus to derive the posterior distribution of θ we only need to marginalize the joint distribution of (p,θ) generated from the MCMC draws over p. This procedure is commonly called data augmentation. See Tanner and Wong (1987) for more details. □

Next we describe a slightly modified version of the above MCMC algorithm which allows researchers to conduct policy experiments given the parameter vector θ^*. Notice that the above MCMC algorithm had two steps. In the first step the equilibrium choice probabilities were drawn given the parameters of the model. In the second step, parameter vector θ was drawn given the equilibrium choice probabilities. The policy simulation algorithm essentially uses only the first step of the MCMC algorithm and omits the second step by keeping the parameter to be fixed. Furthermore, since the policy simulation does not involve fitting the data, the component of the likelihood where the actual choices in the data are incorporated, are dropped.

Algorithm 2 (Metropolis-Hastings2)**.** *Start the chain at* $\{p^{(0)}\}$ *. At iteration* k *iterate over the following steps:*

1. *Sample candidate choice probability* p^* *from the proposal density* q_p*, i.e.* $p^*_m \sim q_p(p^*_m | p^{(k)}_m)$*.*

2. *Evaluate the modified likelihood of market* m *at* p^*_m *where the modified likelihood is:*

$$l_m(c,\theta^*)dc\Big|_{c=B(p^*,\theta^*)} \equiv f(B(p,\theta^*)|\Sigma_c)\,|B_p(p,\theta^*)|\,dp. \tag{9}$$

Notice that in this modified likelihood the data $\vec{y}_m$ is dropped.

3. *Update choice probabilities:*

$$p_m^{(k+1)} = \begin{cases} p_m^* & \text{with probability } \alpha(p_m^*, p_m^{(k)}) \\ p_m^{(k)} & \text{with probability } 1-\alpha(p_m^*, p_m^{(k)}) \end{cases}$$

Where the acceptance probability is given by:

$$\alpha(p_m^*, p_m^k) = \min\left(\frac{l_m(p_m^*,\theta^*)q_p(p_m^*|p_m^{(k)})}{l_m(p_m^{(k)},\theta^*)q_p(p_m^{(k)}|p_m^*)}, 1\right) \tag{10}$$

4 Examples

4.1 Static coordination game

In order to understand the mechanic of the algorithm it is useful to consider the following two-players coordination game studied in **?**). Both players simultaneously choose a binary action $\{a_1, a_2\} \in \{0,1\}^2$ after privately observing a vector of utility shocks $\{\epsilon_{i1}, \epsilon_{i2}\}$. The expected payoff of player i, conditional on choosing action a_i is given by:

$$V_{ia}(\epsilon_{ia}) = a_i\big[\beta + \Pr(a_j=1)\alpha\big] + (1-a_i)\alpha\Pr(a_j=0) + \epsilon_{ia}. \tag{11}$$

If $\{\epsilon_{ia}\}$ are distributed according to a Type-1 extreme value distribution, the set of symmetric Nash equilibria corresponds to the choice probabilities p^* that solves the following equation:

$$p^* = \frac{\exp(\beta+\alpha p^*)}{\exp(\alpha(1-p^*)) + \exp(\beta+\alpha p^*)} = \frac{1}{1+\exp(-\beta+\alpha(1-2p^*))}. \tag{12}$$

By restricting ourselves to symmetric equilibria, the perturbed version of the game is obtained by adding a single random heterogeneity parameter c to the previous payoff function. Conditional on c, a Nash equilibrium of the perturbed game is given by:

$$p^*(c) = \frac{1}{1+\exp(c-\beta+\alpha(1-2p^*))}. \tag{13}$$

Figure 1: Illustration of the Metropolis-Hastings algorithm used as a solution method

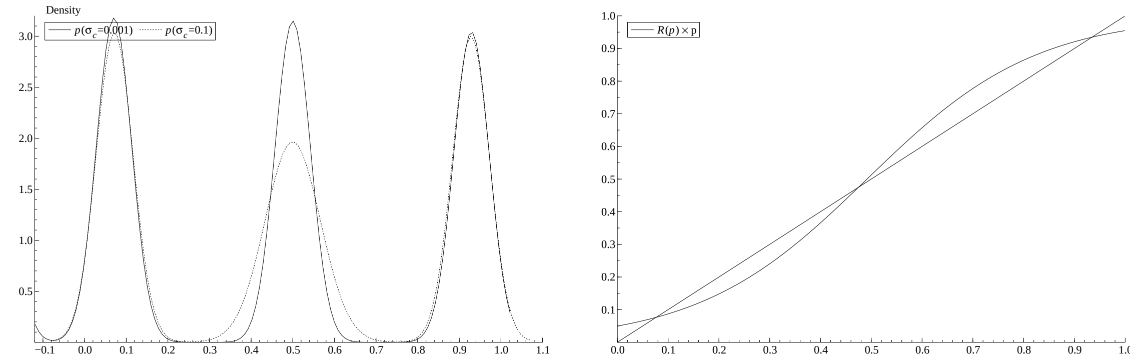

(a) Kernel density of accepted choice probabilities for two values of the variance of the perturbation parameter

(b) Best-response mapping of the true game

Parameters: $\beta = 0.05$, $\alpha = 3$, $\sigma_c \in \{1/1000, 1/10\}$. Number of MCMC replications: $1M$.

From the previous equation, it is easy to see that the inverse best-response probability mapping is unique and can be written as:

$$c = B(p) = \log\left(\frac{1-p}{p}\right) - \alpha(1-2p) + \beta. \tag{14}$$

For large enough values of the strategic interaction parameter α, the game admits three Nash equilibria. Figure 2(b) illustrates the best-response mapping of the true game. The two extreme equilibrium points are stable, and the middle one is unstable.

To illustrate how the MCMC algorithm can approximate the distribution of Nash equilibria, one can apply the Metropolis-Hastings algorithm using only the restrictions imposed by the distribution function $f(c)$. In particular, we repeatedly sample candidate Nash equilibrium choice probabilities, holding fixed the structural parameters. In this case, the acceptance probability for a candidate choice probability p^* is given by:

$$\mu(p^*, p^k) = \min\Big(1, \frac{f(B(p^*))B'^*)}{f(B(p^k))B'^k)}\Big). \tag{15}$$

where $f(B(p))$ is the density of the perturbation parameter $c = B(p)$.

Figure 2(a) shows the kernel density of accepted choice probabilities generated from the Metropolis-Hastings algorithm, using a normal density with variance $\sigma_c = 1/1000$ and

$\sigma_c = 1/10$ to perturb the model. With a small variance of the perturbation parameter, the algorithm is unlikely to accept candidate strategies that do not meet the equilibrium conditions of the true game. As a result, the posterior distribution of strategies is very tight around the three Nash equilibria. The dotted line, which corresponds to a larger variance, is a lot smoother because strategies further away from the Nash equilibria correspond to implicit c's that have larger probability densities. Since the Metropolis-Hastings algorithm cycles between equilibria, perturbing the model with a large variance will typically converge faster to a stationary distribution. However, the perturbed model with a high variance will introduce a bias, since the posterior distribution of strategies will have positive mass in regions that are not Nash equilibria of the true game. This trade-off between bias and speed of convergence will be important when choosing the size of the model perturbation (i.e. σ_c). The choice of σ_c is thus similar to the choice of a bandwidth in kernel density estimation.

Next, we use the full algorithm to estimate the model from two artificial data-sets. Each data-set is generated assuming that each equilibrium is equally likely to be picked (i.e. $\pi_e = 1/3$). The first data-set is a cross-section of 600 markets in which the game is played only once. The second one is a panel, where the same 600 market repeatedly play the game for 10 periods. To simulate the posterior-distribution of parameters and strategies, we use $50,000$ replications of the Metropolis-Hastings algorithm and fix the variance of the normal perturbation parameter to $\sigma_c = 1/10^5$.

The results of the two estimation experiments are reproduced in Figure 3. The first two figures plot the simulated Markov chains. Although the cross-sectional data-set provides a less precise estimate of the parameters, the median parameters are reasonably close to the truth in both cases. The panel data-set contains a lot more information about which equilibria is being played however. As a result, the median simulated parameters are less biased and the posterior distribution of strategies are centered around the selected Nash

[5] An alternative strategy is to treat σ_c as a parameter and estimate the degree of perturbation jointly with the other parameters. While this procedure typically converges quickly to the true value, the value of σ_c eventually becomes too small and the procedure stops accepting candidate parameters. This causes convergence problem as our procedure requires a non-zero value for σ_c. To avoid this problem we can bound the value of σ_c away from zero in the MCMC algorithm.

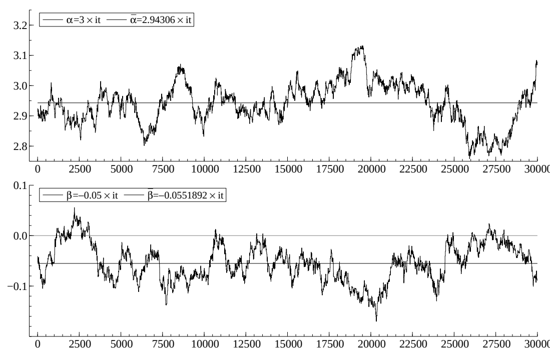

(c) MCMC parameter draws (cross-section)

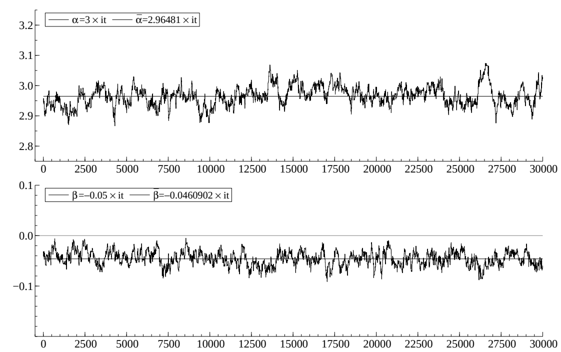

(d) MCMC parameter draws (panel)

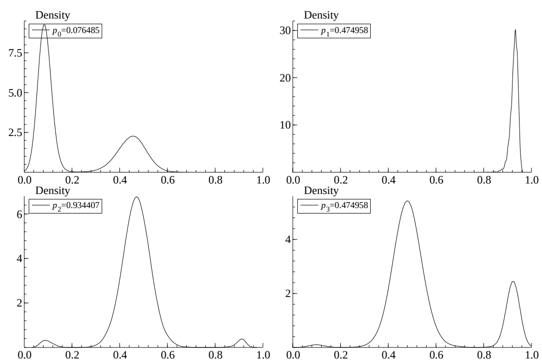

(e) Posterior distribution of choice probabilities for four markets (cross- section)

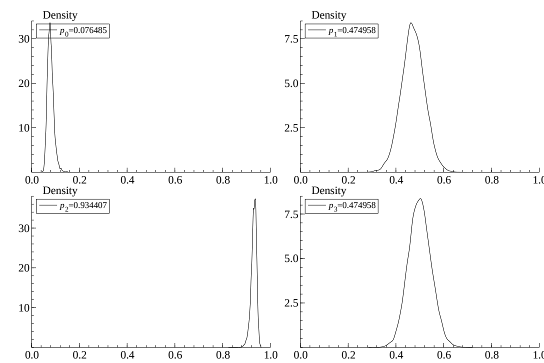

(f) Posterior distribution of choice probabilities for four markets (panel)

Figure 2: Two simulated markov chains from the static coordination game

True parameters: $\beta = 0.05$, $\alpha = 3$. Variance of perturbation parameter: $\sigma_c = 1/10$. Number of MCMC replications: $50,000$. Burn-in period: $20,000$. Sample size: $N = 600$, $T \in \{1, 10\}$. MCMC tuning parameters: $\sigma_p = 0.04$, $\sigma_\alpha = 0.03$, $\sigma_\beta = 0.01$.

equilibrium. Figures 2(e) and 2(f) illustrates the ability of the algorithm to predict the distribution of Nash equilibria for four randomly selected markets. Figure 2(e) suggests that observing only two choices is not enough to classify equilibrium strategies perfectly. However, with repeated observations the posterior distribution of strategies is tightly centered around the selected one. This suggests that our estimation algorithm will work better when repeated time-periods are available or when more than two players play same static game.

Next we consider the impact of ignoring the unstable equilibrium on the estimation results. Figure **??** illustrates the two simulated markov chains for the sample without the

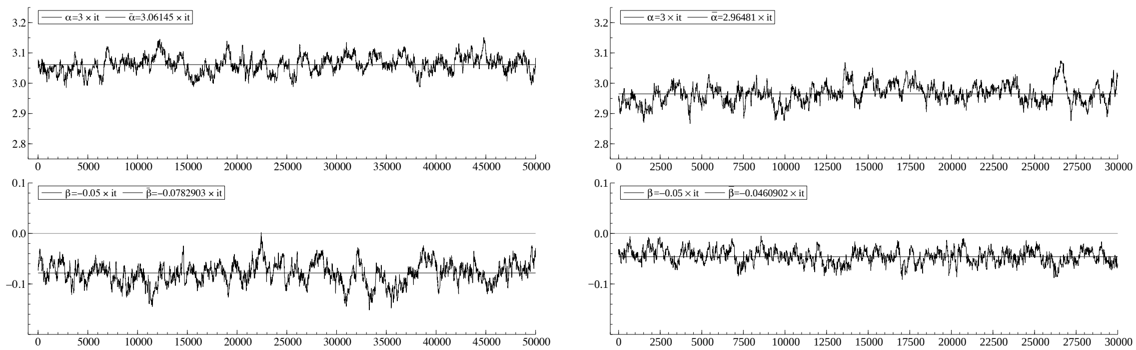

(a) MCMC parameter draws with two equilibria (b) MCMC parameter draws with three equilibria

Figure 3: Simulated markov chains from the static coordination game with two alternative equilibrium selection rules

---

True parameters: $\beta = 0.05$, $\alpha = 3$. Variance of perturbation parameter: $\sigma_c = 1/10$. Number of MCMC replications: $50,000$. Burn-in period: $20,000$. Sample size: $N = 600$, $T = 10$. MCMC tuning parameters: $\sigma_p = 0.04$, $\sigma_\alpha = 0.03$, $\sigma_\beta = 0.01$.

unstable equilibrium (the others are equally likely), and Figure **??** illustrates the simulated chains generated from the sample with three equally likely equilibria. Although the scale of the graphs are different, it is clear that ignoring one equilibrium in the data generating process does not significantly affect the consistency of the estimates. The results from Figure **??** suggests however that the sample with two equilibria is more biased than the one with three equilibria. This is consistent with the fact that in this class of model the presence of multiple equilibria generates exogenous variation that allows the parameters to be identified (see **?**) and **?**)). In fact if the data set was generated with only one equilibria the slope could not be separately identified from the intercept.

Finally, we investigate the size of the bias generated by the heterogeneity added in the model. To do so, we generated simulated Markov chains of equal length with different values of σ_c. Figure 4 plots the median parameters against the variance of the random intercept (i.e. $\sigma_c \in [1/100, 1/2]$). One thing to note from this figure is that both parameters are consistently estimated when the size of the added heterogeneity is below 0.15. The previous simulation results were generated with $\sigma_c = 1/10$ which allows the markov chain to rapidly converge. Interestingly for large values of σ_c the bias of coordination payoff parameter α

Figure 4: Illustration of the estimation bias. Median simulated parameters for various values of the variance of the random intercept: $\sigma_c \in [1/100, 1/2]$.

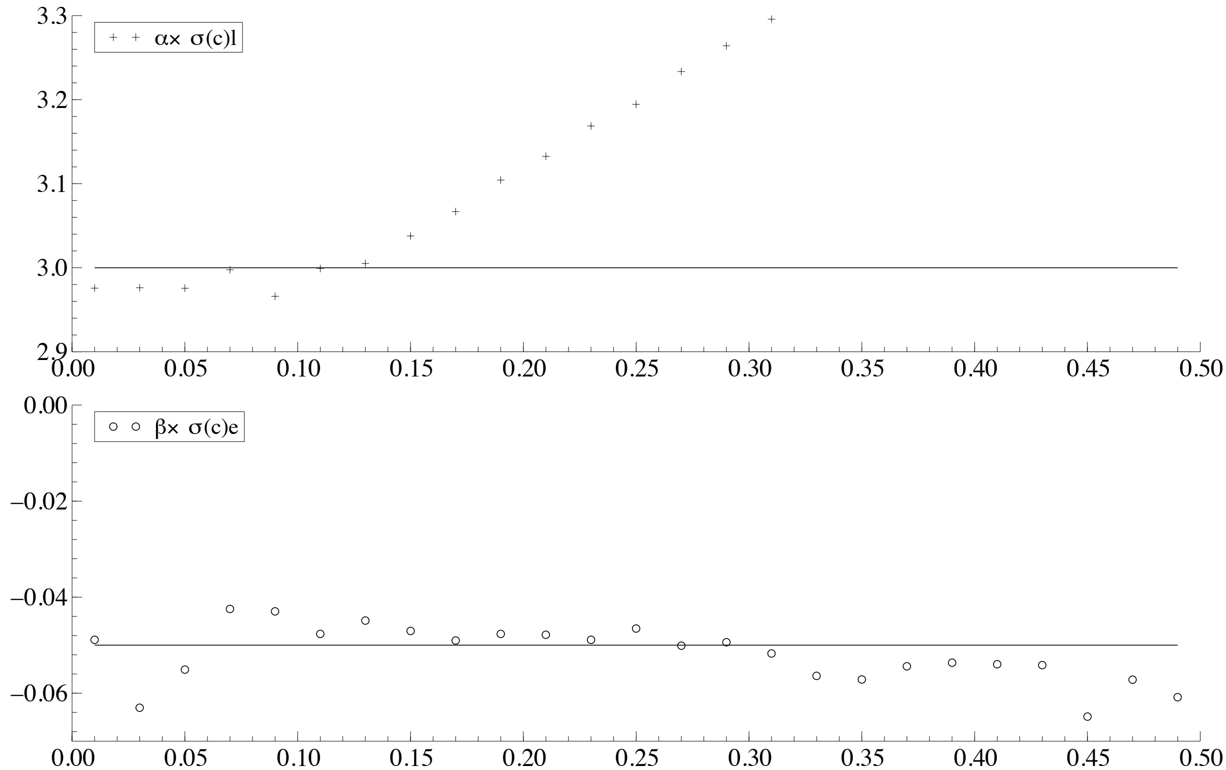

True parameters: $\beta = 0.05$, $\alpha = 3$. Number of MCMC replications: $50,000$. Burn-in period: $20,000$. Sample size: $N = 600$, $T = 10$. MCMC tuning parameters: $\sigma_p = 0.04$, $\sigma_\alpha = 0.03$, $\sigma_\beta = 0.01$.

increases with σ_c. This is not the case however for the intercept which does not seem to be affected by the size of the perturbation. One possible explanation for this is that the coordination payoff parameter affects directly the slope of the best-response function (i.e. steeper), which in turns determines the size of c in the inversion procedure. When σ_c is big, the density of larger c (in absolute value) increases which requires the value of α to raise as well.

4.2 Dynamic game

We analyze an infinitely repeated game with 2 firms and 2 choices. We denote $x \in \mathcal{X}$ to be the state which indicates the action chosen by each player in the previous period. That is,

$$x_{1t} = \{a_{1t-1}, a_{2t-1}\}.$$

The private information state variables are denoted by ϵ_i^a, for all $i = \{1, 2\}$ and $a \in \{0, 1\}$. As before, we consider symmetric markovian strategies that depends solely on the current period state variables. This leads to a stationary markov process, so we define current period actions and states by $\{a'_1, a'_2\}$ and $\{a_1, a_2\}$ respectively.

As before we define strategies as conditional choice probabilities $p_i(x)$. A Nash equilibrium of this game is thus represented by a four (4) dimensional vector of choice probabilities $\vec{p}$. The static payoff of the model with the additional heterogeneity parameters $\vec{c} = \{c(0,0), c(0,1), c(1,0), c(1,1)\}$ is given by:

$$u_i(a', x, c, \epsilon_i, p_{-i}) = \begin{cases} c_x - \kappa 1(a_i = 0) + \alpha p_{-i}(x) \\ \quad -\alpha(1 - p_{-i}(x)) + \epsilon_i^1 & \text{if } a'_i = 1 \\ \epsilon_i^0 & \text{if } a'_i = 0 \end{cases}$$

The parameter vector is $\theta = \{\alpha, \kappa\}$, where $\alpha > 0$ is a coordination payoff and κ is a switching cost. Note that the dimension of the vector $\vec{c}$ is $|X| \times |A| - 1 = 4$ because we restrict our attention to symmetric strategies. To allow for asymmetric strategies and meet the requirements for the existence of the inverse function describe in Section 2 we would need increase the dimension of $\vec{c}$ to $N \times |X| \times |A| - 1 = 8$.

We define two 4×4 actions specific state transition probability matrices that depends directly on players' beliefs $F_1^{p_{-i}} = \left[F^{p_{-i}}(z'|z, 1)\right]$ and $F_0^{p_{-i}} = \left[F^{p_{-i}}(z'|z, 0)\right]$ such that,

$$F^{p_{-i}}(x'|x, a') = \begin{cases} p_{-i}(x) & \text{if } x' = (1, a') \\ 1 - p_{-i}(x) & \text{if } x' = (0, a') \\ 0 & \text{else.} \end{cases}$$

The choice specific value function (without ϵ_i) is given by:

$$\bar{V}(a'_i, x, c_x, p_{-i}) = \pi_i(a'_i, x, c_x, p_{-i}) + \delta \sum_{x'} EV(x', c_{x'}, p_{-i}) F^{p_{-i}}(x'|x, a'_i)$$

where $\pi_i(a', x, c_x, p_{-i}) \equiv u_i(a'_i, x, c_x, \epsilon_i, p_{-i}) - \epsilon_i^{a'}$. We assume that $\vec{\epsilon}$ is distributed according to an extreme value distribution. The expected continuation value is then given by:

$$\begin{aligned} EV(x, c_x, p_{-i}) &= \int \Big[\max_{a'} \bar{V}(a', x, c_x, p_{-i}) + \epsilon_i^{a'} \Big] dg_\epsilon(\epsilon) \\ &= \log \big(\exp(\bar{V}(0, x, c_x, p_{-i})) + \exp(\bar{V}(1, x, c_x, p_{-i})) \big) \end{aligned} \quad (16)$$

A symmetric Markov Perfect Nash equilibrium is a choice probability $p^*(x, c)$ which is a fixed point of the best-response probability mapping:

$$\begin{aligned} p^*(x, c) &= \Lambda(x, c, p^*) = \frac{\exp(\bar{V}(1, x, c_x, p^*))}{\exp(\bar{V}(0, x, c_x, p^*)) + \exp(\bar{V}(1, x, c_x, p^*))}, \\ & \forall x \in \big\{(0,0), (0,1), (1,0), (1,1)\big\}. \end{aligned} \quad (17)$$

To obtain the inverse function $c_x(p)$ we first define the value function difference mapping $\hat{V}(x, c_x, p_{-i})$. As in Hotz-Miller, we can use the extreme value distribution to write the value function difference as:

$$\hat{V}(x, c_x, p_{-i}) \equiv \bar{V}(1, x, c_x, p_{-i}) - \bar{V}(0, x, c_x, p_{-i}) = \log \left(\frac{p_x}{1 - p_x} \right) \quad (18)$$

where p_x denotes element x of the vector $\vec{p}$. Note that $\hat{V}(\vec{p})$ is a 4×1 vector which is function solely of choice probabilities.

Moreover the expected continuation value at the Nash equilibrium choice probability p is[6]:

$$\begin{aligned} EV(x, c_x, p) &= \log \big(\exp(\bar{V}(1, x, c_x, p)) + \exp(\bar{V}(0, x, c_x, p)) \big) \\ &= \bar{V}(1, x, c_x, p) - \log(p_x) \end{aligned}$$

Using the previous equality, the choice specific value function can be written as:

$$\bar{V}(1, x, c_x, p) = \pi(a', x, c_x, p) + \delta \sum_{x'} \Big[\bar{V}(a', x, c_x, p) - \log(p_x) \Big] F^p(x'|x, a')$$

[6]To reduce the notation burden p and c are understood to be vectors of dimension 4×1.

In matrix form, the vectors choice specific value function are:

$$
\begin{aligned}
\bar{V}(0,c,p) &= (I-\delta F_0^p)^{-1}\Big[\pi(0,p)-\delta F_0^p\log(1-p)\Big] \\
&= -(I-\delta F_0^p)^{-1}\delta F_0^p\log(1-p) \\
\bar{V}(1,c,p) &= (I-\delta F_1^p)^{-1}\Big[\pi(1,c,p)-\delta F_1^p\log(p)\Big]
\end{aligned}
$$

Where the second line follows from the normalization of the payoff of action 0. Joining the last two equalities with the expression for the vector of value function differences $\hat{V}(p)$:

$$
\begin{aligned}
\hat{V}(p) &= (I-\delta F_1^p)^{-1}\Big[\pi(1,c,p)-\delta F_0^p\log(p)\Big] \\
&\quad +(I-\delta F_0^p)^{-1}\delta F_0^p\log(1-p)
\end{aligned}
$$

$$
\begin{aligned}
\pi(1,c,p) &= (I-\delta F_1^p)\hat{V}(p)+\delta F_1^p\log(p) \\
&\quad -(I-\delta F_1^p)(I-\delta F_0^p)^{-1}\delta F_0^p\log(1-p)
\end{aligned}
$$

Let $\tilde{\pi}(1,x,p_{-i})=\kappa 1(a_i=0)+\alpha p_{-i}(x)-\alpha(1-p_{-i}(x))$ be the unperturbed payoff function (i.e. without heterogeneity). Then the vector c which rationalize p as a Nash equilibrium is given by:

$$
\begin{aligned}
c(p) &= -\tilde{\pi}(1,p)+(I-\delta F_1^p)\hat{V}(p)+\delta F_1^p\log(p) \\
&\quad -(I-\delta F_1^p)(I-\delta F_0^p)^{-1}\delta F_0^p\log(1-p)
\end{aligned}
$$

As in the static coordination game example, the function $c(p)$ is continuous and differentiable for all $p(x)\in(0,1)$.

To analyze the performance of our estimator in a dynamic setting we generated a random sample of 200 markets over 20 time periods. Note that we generated the data without heterogeneity (i.e. $\sigma_c=0$). We chose the remaining parameter values such that the game generates potentially three Nash equilibria (1 being unstable). We use two selection rules to generate the data: one with equal selection probability and the other assigning zero probability on the unstable equilibrium. The three Nash equilibrium strategies are reproduced in Tabel 1.

Table 1: Nash equilibrium strategies of the dynamic game for $\alpha = 1.5$ and $\kappa = 1/10$.

States	Eq. 1	Eq. 2	Eq. 3 (unstable)
(0,0)	0.501	0.152	0.180
(0,1)	0.334	0.838	0.609
(1,0)	0.379	0.864	0.645
(1,1)	0.261	0.821	0.929

Before discussing the estimation results, we first present Monte-Carlo simulations used as a solution algorithm. In particular we set the joint density of $\vec{c}$ to $N(0, \sigma_c^2 I)$ and perform 1 million replications of the Metropolis algorithm using only the density of c as the "likelihood". This procedure is akin to a random search algorithm which simulates various candidate Nash equilibrium choice probabilities, and keep the ones for which the density of $\vec{c}$ is large (i.e. $c_x \sim 0$ for all x).

Figure 5 plots the simulated markov chain of accepted choice probabilities for the four states over the full set of iterations. Figure 6 presents the density of accepted strategies, where the vertical lines indicates the stable Nash equilibria. With a value of $\sigma_c = 0.05$ the MCMC algorithm easily cycles between all three equilibria. This simulation process successfully approximates all Nash equilibria, including the unstable one. In particular, the density of strategies in Figure 6 has multiple local modes corresponding to each Nash equilibria. Note that the upper left and lower right corner graphs have only two local modes since the unstable equilibrium strategy is numerically close to the first Nash equilibrium in those two states.

We now turn to the estimation results. For this exercise, we estimate the size of the perturbation by imposing a strong prior on σ_c. In particular we assume that $\sigma_c \in [0.02, 0.2]$. This strategy seems to yield better convergence properties than fixing σ_c. We used $50,000$ Monte-Carlo replications and drop the first $20,000$ iterations.

Figure 4.2 presents the results for the sample with 3 equally likely equilibria. The mean estimated parameters are $\hat{\alpha} = 1.479$ and $\hat{\kappa} = 0.36$. These estimates are very close to the

Figure 5: MCMC solution algorithm applied to the dynamic game: Markov chain of choice probabilities for the four states.

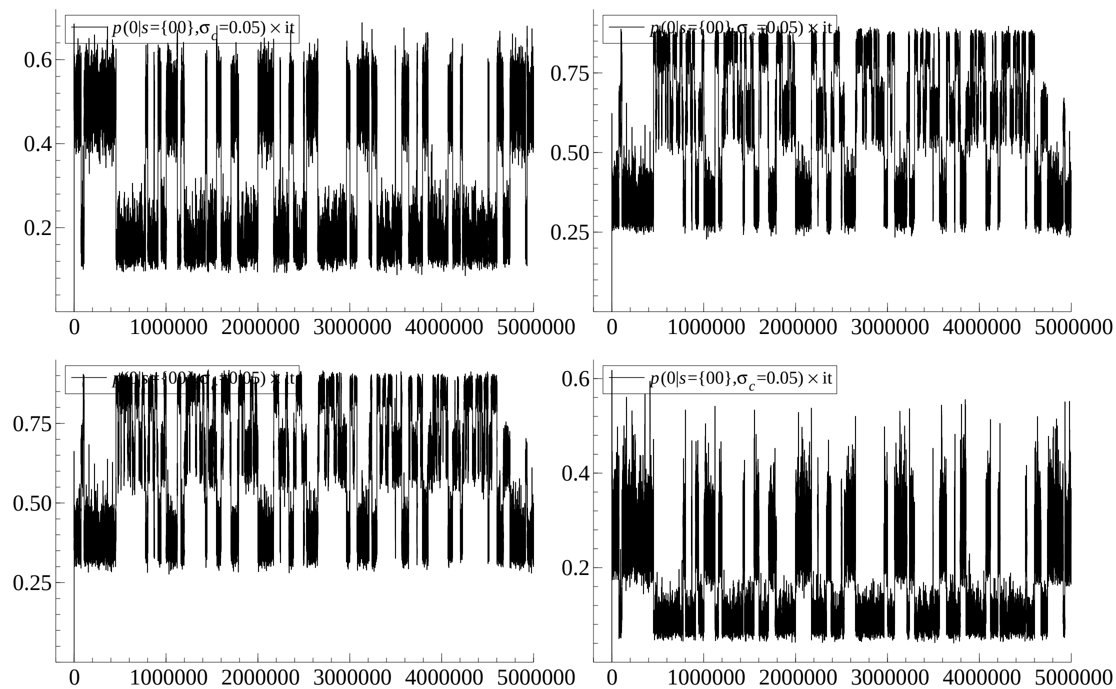

Figure 6: MCMC solution algorithm applied to the dynamic game: Density of accepted probabilities for the four states.

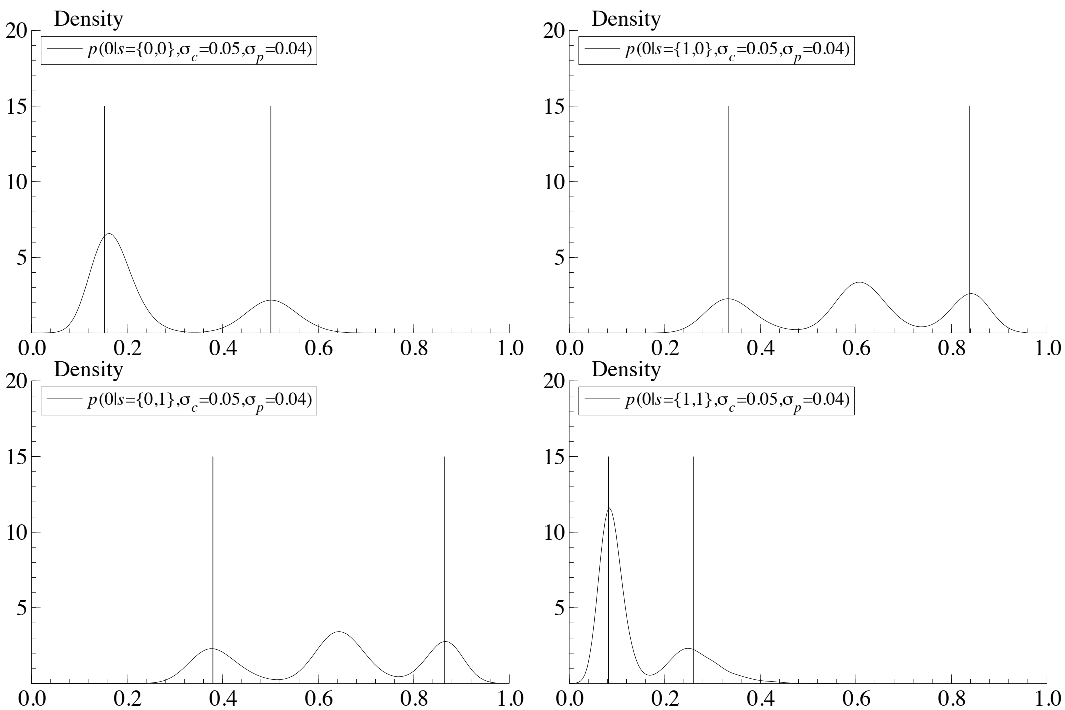

true parameters used to generate the data, despite the fact that the data does not contain any heterogeneity. The bottom figure shows that the simulated values for the size of the perturbation stays very close to the lower bound.

Next, we illustrate the ability of the estimation algorithm to sort markets into different equilibria through the data-augmentation step. Recall that the Metropolis-Hastings algorithm allows each market to be playing a different candidate Nash equilibria, and that its acceptance probability depends only the observed sequence of choices made in that market. Therefore, if the data is rich enough the distribution of accepted choice probabilities will be different across markets that select different equilibria. That is the posterior distribution of Nash equilibria will differ across markets.

Figure 8 illustrates this point by constructing the distribution of accepted choice probabilities for the full sample (top left corner) and for the three sub-samples corresponding to each equilibria. The vertical lines again indicate the true Nash equilibria. The density for the full sample corresponds more or less to the density generated without the data (i.e. Figure 6). The other figures however clearly suggest that the algorithm successfully sort markets in the "right" equilibrium category. In particular, for each sub-sample the distribution of accepted choice probabilities has a mode corresponding to the equilibrium played in the data. This is especially true for the first equilibria which is very different for the two others. For the second and third equilibria, the mode is less pronounced because the strategies are similar.

5 Extensions and future work

To be done.

Figure 7: Markov Chain of the parameters for the dynamic model (Sample with equally likely 3 Nash equilibria)

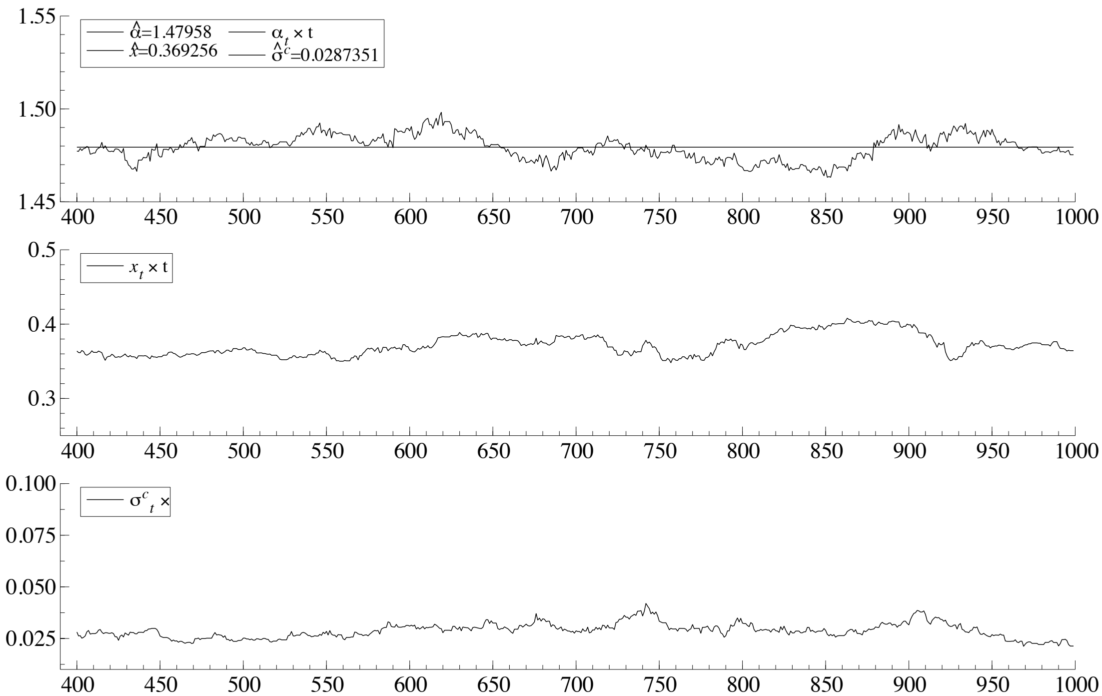

Figure 8: Density of accepted choice probabilities for the dynamic model with three Nash equilibria.