An Estimable Dynamic Model of Entry, Exit, and Growth in Oligopoly Retail Markets

By Victor Aguirregabiria, Pedro Mira, and Hernan Roman*

Retail industries are an important part of today's economy. They employ a large fraction of the labor force, play a key role in the adoption and diffusion of new information technologies, and are closely related to the development and configuration of urban life. This paper presents an estimable dynamic structural model of an oligopoly retail industry. The model can be estimated using panel data of local retail markets with information on new entries, exits, and the size and growth of incumbent firms. In our model, retail firms are vertically and horizontally differentiated, compete in prices, make investments to improve the quality of their businesses, and decide to exit or continue in the market. The model extends the entry-exit model estimated in Aguirregabiria and Mira (2007, AM hereafter) in two important ways. First, it includes firm size and growth as endogenous variables. Second, the empirical model has two sources of permanent unobserved heterogeneity: local-market heterogeneity and firm heterogeneity. This allows the researcher to control for potentially important sources of bias when using firm panel data from many local markets and several time periods. Not accounting for market unobserved heterogeneity can lead to biases in the estimation of those structural parameters that represent strategic interactions between firms' decisions. Market heterogeneity induces a positive correlation between firms' decisions that can be spuriously confounded with positive strategic interactions. It is also well known in panel data econometrics that ignoring unobserved heterogeneity across firms induces bias in the parameters that generate structural state dependence, e.g., investment costs. We extend the *Nested Pseudo Likelihood* estimation method in AM to deal with both forms of permanent unobserved heterogeneity. This note contributes to the emerging literature on the estimation of empirical games of industry dynamics (see also Patrick Bajari, Lanier Benkard, and Jonathan Levin 2007; Martin Pesendorfer and Philipp Schmidt-Dengler 2002; and Ariel Pakes, Michael Ostrovsky, and Steven Berry 2007).

In a related paper (Aguirregabiria, Mira, and Roman 2006, AMR hereafter), we use this model to study the sources of cross-industry heterogeneity in the dynamics of market structure. We use annual panel data from a census of Chilean firms collected by the Chilean *Servicio de Impuestos Internos* (Internal Revenue Service) for the period 1994–2000. For every establishment paying sales tax, this dataset reports its industry at the five-digit level, its annual revenue, and the district where the establishment is located. Competition in retail industries occurs at the local level, and we consider districts as local markets. We find large cross-industry heterogeneity along several dimensions of market structure and industry dynamics, e.g., entry and exit rates, Herfindahl index, the relationship between firm size and market size, or the relationship between firm growth and firm size. We estimate the model separately for each retail industry and use these estimates to evaluate the role that product differentiation, economies of scale, exogenous entry costs, and endogenous sunk costs play in explaining the observed cross-industry heterogeneity.

I. The Model

Consider a local market in a particular retail industry, e.g., hotels, car dealers, or supermarkets. The market is populated by consumers and firms. We index firms by i. Time is discrete and indexed by t. At period t there are S_t consumers, N_t^{in} incumbent firms operating in the market, and

* Aguirregabiria: Department of Economics, University of Toronto, 100 Saint George St., Toronto, Ontario M5S 3G3 (e-mail: victor.aguirregabiria@utoronto.ca); Mira: Centro de Estudios Monetarios y Financieros (CEMFI), Casado del Alisal, 5, 28014 Madrid, Spain (e-mail: mira@cemfi.es); Roman: Escuela de Negocios, Universidad de los Andes, Avd. La Plaza 1905, Santiago, Chile (e-mail: hroman.ese@uandes.cl).

N_t^{out} firms which are not active but are potential entrants, i.e., they may choose to operate in the market. We refer to S_t as the *market size*, and we assume that it evolves exogenously according to a Markov process with transition probability function $f_S(S_{t+1}|S_t)$. The number of potential entrants is proportional to market size, $N_t^{out} = int(\delta S_t)$, where δ is a parameter. The number of incumbent firms is endogenously determined in the equilibrium of the model. Each incumbent firm sells a differentiated product. Firms compete in prices but also in the quality of their products.

We consider a logit model of demand for differentiated products, as in Simon P. Anderson, Andre de Palma, and Jacques-François Thisse (1992, 264–66). Consumers value product attributes such as quality and price, and every period firms compete in prices à la Bertrand. The equilibrium of this static model of demand and price competition results in an indirect variable profit function that depends on a firm's own product quality, the competitors' qualities, and market size. Let w_{it} be the quality of firm i at period t, and let Ω be the set of possible product qualities, which is a discrete and finite set. Define the vector $\boldsymbol{n}_t = \{n_t(w) : w \in \Omega\}$, where $n_t(w)$ is the number of incumbent firms at period t with quality w. The equilibrium indirect variable profit of firm i is $\theta_v S_t m(w_{it}, \boldsymbol{n}_t) s(w_{it}, \boldsymbol{n}_t)$, where $m(w_{it}, \boldsymbol{n}_t)$ and $s(w_{it}, \boldsymbol{n}_t)$ are the retailer's equilibrium price-cost margin and market share, respectively, and θ_v is a parameter that represents the degree of horizontal product differentiation.

A retailer's stock of quality evolves as a controlled Markov chain. A firm can improve its quality by making additional investments. If the firm does not invest in quality improvement, its quality depreciates as the result of increases in the value of the consumers' outside alternative. An incumbent firm's investment decision is a binary choice $i_{it} \in \{0, 1\}$, where $i_{it} = 1$ means positive investment. The transition probability of quality for an incumbent firm is $f_{w1}(w_{i,t+1}|w_{it}, i_{it})$.

Incumbent firms are also heterogeneous in their investment costs. This source of heterogeneity is important in some retail industries where the cost of financing new investments varies significantly across firms. Let γ_i be the firm-specific and time-invariant component of firm i's investment cost. A firm's "type" γ_i is drawn from the probability function f_γ with discrete and finite support Γ. Define the vector $\boldsymbol{n}_t^* = \{n_t^*(w, \gamma) : (w, \gamma) \in \Omega \times \Gamma\}$, where $n_t^*(w, \gamma)$ is the number of incumbent firms at period t with quality w and investment cost γ. We refer to $\boldsymbol{n}_t^*$ as the *latent market structure* because it depends on the distribution of the latent (i.e., unobserved for the researcher) investment costs.

Every period t, incumbent firms decide their respective investments in quality improvements and whether to remain in the market or exit at the next period. Potential entrants decide whether to enter the market. Firms make these decisions to maximize expected discounted intertemporal profits $E_t(\sum_{j=0}^{\infty} \beta^j \Pi_{i,t+j})$, where $\beta \in (0, 1)$ is the discount factor, and Π_{it} is the firm's current profit. A potential entrant who decides to stay out of the market gets zero profits. If he decides to enter, he has to pay an entry cost $\theta_{EC} + \varepsilon_{E,it}$ at period t, where θ_{EC} is the component of entry costs that is common to all the firms in the market; and $\varepsilon_{E,it}$ is a firm-specific component which is private information of the firm, has zero mean, and is i.i.d. over firms and over time. The new entrant is not active until the next period. Furthermore, the firm-specific investment cost and the initial quality of a new entrant are uncertain when the firm makes its entry decision, and they are not realized until the next period. The initial quality is a random draw from the probability function f_{w0}.

Current profits of an incumbent firm that stays in the market are $\theta_v S_t m(w_{it}, \boldsymbol{n}_t) s(w_{it}, \boldsymbol{n}_t) - FC_{it} - IC_{it}$. The first term is the variable profit. The second term is the fixed operating cost $FC_{it} = \theta_{FC} + \varepsilon_{FC,it}$, where θ_{FC} is a parameter and $\varepsilon_{FC,it}$ is a private information shock. The last term, IC_{it}, represents investment costs for quality improvement $IC_{it} = (\theta_{IC} + \sigma_\gamma \gamma_i + \varepsilon_{I,it}) i_{it}$, where θ_{IC} and $\sigma_\gamma > 0$ are parameters; γ_i is the firm-specific component of the investment cost, which has zero mean and unit variance; and $\varepsilon_{I,it}$ is private information of firm i. The profit of an incumbent firm that decides to exit from the market at the end of period t is $\theta_v S_t m(w_{it}, \boldsymbol{n}_t) s(w_{it}, \boldsymbol{n}_t) - FC_{it} + \varepsilon_{X,it}$. The exiting firm is operative during period t. It obtains its variable operating profits and it has to pay fixed costs. It also receives an exit value $\varepsilon_{X,it}$ which is private information. The shocks $\{\varepsilon_{FC,it}, \varepsilon_{I,it}, \varepsilon_{X,it}\}$ have zero mean and are i.i.d. over time and across firms. We assume that the pool of potential entrants is renewed every period. Therefore, exiting firms do not become potential entrants the following period and their continuation value is the value of the best outside alternative, which we normalize to zero. Likewise, the continuation value of potential entrants that choose to stay out of the market upon drawing a large entry cost is also set to zero.

We assume that firms' entry, exit, and investment strategies depend only on payoff-relevant state variables, i.e., Markov Perfect Equilibrium (MPE). The set $\{w_{it}, \gamma_i : i = 1,\ldots,N_t^{in}\}$ is common knowledge. The vector of payoff-relevant state variables of firm i is $\tilde{x}_{it} = (w_{it}, \gamma_i, \boldsymbol{n}_t^*, S_t, \varepsilon_{E,it}, \varepsilon_{FC,it}, \varepsilon_{I,it}, \varepsilon_{X,it})$. The continuation value of an entrant or a staying incumbent as of next period is $\beta \mathrm{E}[V(\tilde{x}_{it+1})|\tilde{x}_{it}]$, where the value function V is the solution of a Bellman equation. Following AM, we can show that an MPE can be represented in terms of players' choice probabilities conditional on common knowledge state variables. In this model, there are three free-choice probability functions: for a potential entrant, the probability of entry $P_E(\boldsymbol{n}_t^*, S_t)$; and, for an incumbent firm, the probability of exiting the market, $P_X(w_{it}, \gamma_i, \boldsymbol{n}_t^*, S_t)$, and the probability of staying and investing in quality, $P_I(w_{it}, \gamma_i, \boldsymbol{n}_t^*, S_t)$. Let $\boldsymbol{P} = (P_E, P_X, P_I)$ denote the vector of choice probability functions. These functions describe a firm's behavior as well as its beliefs about the behavior of its opponents. Given these beliefs, one can interpret each firm's problem as a game against nature with a unique optimal decision rule in probability space, which is the firm's best response. The equilibrium probability function is a fixed point of this best response mapping. Suppose that the private information shock $\varepsilon_{E,it}$ has a logistic distribution with dispersion parameter σ_E, and that the shocks $\{\varepsilon_{FC,it}, \varepsilon_{I,it}, \varepsilon_{X,it}\}$ have a type I extreme value distribution with dispersion parameter σ_X. Let $\tilde{\boldsymbol{\theta}}_\pi$ be the vector of structural parameters $(\theta_v, \theta_{FC}, \theta_{EC}, \theta_{IC}, \sigma_\gamma)'$. Then, the best response of a firm with beliefs $\boldsymbol{P}$ is given by the following probability functions:

$$(1) \qquad \Psi_E(\boldsymbol{n}_t^*, S_t; \boldsymbol{P}) = exp\{\mathbf{z}_t^E(\tilde{\boldsymbol{\theta}}_\pi/\sigma_E) + (\sigma_X/\sigma_E)\,\lambda_t^E\}/\textstyle\sum_t^E$$

$$\Psi_I(w_{it}, \gamma_i, \boldsymbol{n}_t^*, S_t; \boldsymbol{P}) = exp\{\mathbf{z}_{it}^I(\tilde{\boldsymbol{\theta}}_\pi/\sigma_X) + \lambda_t^I\}/\textstyle\sum_{it}^X$$

$$\Psi_X(w_{it}, \gamma_i, \boldsymbol{n}_t^*, S_t; \boldsymbol{P}) = exp\{z_{it}^X(\tilde{\boldsymbol{\theta}}_\pi/\sigma_X)\}/\textstyle\sum_{it}^X,$$

where $\sum_t^E \equiv 1 + exp\{\mathbf{z}_t^E(\tilde{\boldsymbol{\theta}}_\pi/\sigma_E) + (\sigma_X/\sigma_E)\,\lambda_t^E\}$ and $\sum_{it}^X \equiv exp\{\mathbf{z}_{it}^I(\tilde{\boldsymbol{\theta}}_\pi/\sigma_X) + \lambda_{it}^I\} + exp\{\mathbf{z}_{it}^{NI}(\tilde{\boldsymbol{\theta}}_\pi/\sigma_X) + \lambda_{it}^{NI}\} + exp\{\mathbf{z}_{it}^X(\tilde{\boldsymbol{\theta}}_\pi/\sigma_X)\}$. The vectors $\mathbf{z}_t^E$, $\mathbf{z}_{it}^I$, $\mathbf{z}_{it}^{NI}$, and $\mathbf{z}_{it}^X$ and the scalars λ_t^E, λ_{it}^I, and λ_{it}^{NI} are functions of the state $(\boldsymbol{n}_t^*, S_t)$ and collect the infinite sum of expected payoffs along all possible future histories originating from that state. These expected payoffs are obtained using beliefs $\mathbf{P}$ about current and future behavior and the "primitive" probabilities $\{f_{w0}, f_{w1}, f_\gamma\}$, and they are discounted by β. Because z's and λ's depend on choice probabilities $\mathbf{P}$, the expressions in (1) describe a fixed-point mapping, and the equilibrium probability functions $\{P_E, P_X, P_I\}$ are a fixed point of this mapping. Further details including expressions for z's and λ's can be found in AM and AMR. The model implies that market structure $\boldsymbol{n}_t^*$ follows a first-order Markov process.

II. Estimation Method

Suppose we have a sample of M isolated retail markets, where M is large. Our asymptotic estimation results apply when the number of markets M goes to infinity. For each market m we observe all the firms that are active in the market between periods 1 and T. For each firm i we observe $\{e_{im}, x_{im}, R_{imt} : t = e_{im},\ldots,x_{im}\}$, where e_{im} and x_{im} are the entry and exit periods of firm i, respectively; and R_{imt} is the revenue of this firm at period t. The actual entry period of incumbents at $t = 1$ is unknown. We use $e_{im} = 0$ to denote these "left-censored" entry periods since there is a one-period lag between entry decisions and entry. Likewise, we do not know the actual exit period of firms that are still active at the end of period T; and we write $x_{im} = T + 1$ to denote these "right-censored" exit periods. We also observe a measure of market size, S_{mt}. In AMR, we show that, under the assumption that all operating costs are fixed costs, the logit model of demand and price competition can be used to obtain firms' qualities $\{w_{it}\}$ from firms' revenue data. Therefore, we treat qualities as observable variables. We assume that: (a) a firm cannot increase its quality if it does not invest, i.e., $f_{w1}(w_{i,t+1}|w_{it}, i_{it} = 0) = 0$ for $w_{i,t+1} \geq w_{it}$; and (b) if a firm invests, then its quality does not depreciate, i.e., $f_{w1}(w_{i,t+1}|w_{it}, i_{it} = 1) = 0$ for $w_{i,t+1} < w_{it}$. Then, a firm's quality grows if and only if the firm invests in quality improvement, i.e., $i_{it} = I\{w_{i,t+1} \geq w_{it}\}$. Therefore, investment is also observable for every period $t \leq T - 1$. The number of potential entrants is estimated as $\hat{N}_{mt}^{out} = int(\hat{\delta} S_{mt})$, where $\hat{\delta} = \max_{m,t}\{\bar{e}_{mt}/S_{mt}\}$ is a consistent estimator of δ, and $\bar{e}_{mt}$ is the number of entrants at period t.

The structural parameters in the profit function are identified only up to scale. We use $\boldsymbol{\theta}_{\pi}$ to denote the vector normalized parameters which are identified, $\boldsymbol{\theta}_{\pi} \equiv (\tilde{\boldsymbol{\theta}}_{\pi}'/\sigma_X, \sigma_E/\sigma_X)$. Let $\boldsymbol{\theta}_f$ be the vector of structural parameters that characterizes the probability functions f_S, f_{w0}, and f_{w1}. Recall that γ_i is a standardized random variable, i.e., the mean and the variance of the permanent component of investment costs are θ_{IC} and σ_{γ}^2, respectively. We assume that the distribution f_{γ} of the standardized heterogeneity γ_i is known to the econometrician, e.g., it is a discretized version of the density of a standard normal. That is, we consider distributions of unobserved firm heterogeneity which are known up to the mean and variance parameters. This assumption contributes to the effectiveness of our estimation procedure because it implies global concavity of the pseudo likelihood.

The likelihood function for this model and data has the form $\prod_{m=1}^{M} L_m(\boldsymbol{\theta}_{\pi}, \boldsymbol{\theta}_f)$, where $L_m(\boldsymbol{\theta}_{\pi}, \boldsymbol{\theta}_f)$ is the contribution of market m to the likelihood function. Here, we derive the expression of the likelihood L_m. For notational simplicity, we omit the market subindex m. Also, for the sake of space, we ignore unobserved market heterogeneity and focus on firm heterogeneity (see AMR for the description of the NPL method with both forms of unobserved heterogeneity). Equilibrium probabilities depend on the firm's own type and on the types of all incumbent firms. Therefore, in order to obtain the likelihood function, we have to integrate over the distribution of firms' types. Let $\tilde{\boldsymbol{\gamma}} \equiv \{\gamma_i : i = 1, 2, \ldots, N\}$ be the vector with the (unobserved) type of every firm observed in this market during the sample period. Then, $L(\boldsymbol{\theta}_{\pi}, \boldsymbol{\theta}_f) = \sum_{\tilde{\boldsymbol{\gamma}} \in \boldsymbol{\Gamma}^N} \Pr(Data|\tilde{\boldsymbol{\gamma}})$, and $\Pr(Data|\tilde{\boldsymbol{\gamma}}) = L_1(\boldsymbol{\theta}_f) L_2(\boldsymbol{\theta}_{\pi}, \boldsymbol{\theta}_f|\tilde{\boldsymbol{\gamma}})$. The function $L_1(\boldsymbol{\theta}_f)$ is the likelihood of the history of market sizes and firms' qualities (conditional on investment decisions, on initial conditions, and on firm types):

$$(2) \quad L_1(\boldsymbol{\theta}_f|\tilde{\boldsymbol{\gamma}}) = \left[\prod_{i:e_i \neq 0} f_{w0}(w_{ie_i+1}) \prod_{i=1}^{N} \prod_{t=e_i+2}^{\min[x_i,T]} f_{w1}(w_{it}|w_{i,t-1}, i_{i,t-1})\right] \times \left[\prod_{t=2}^{T} f_S(S_t\,|\,S_{t-1})\right].$$

This likelihood depends only on the structural parameters in subvector $\boldsymbol{\theta}_f$, and in this model it does not depend on the unobserved types $\tilde{\boldsymbol{\gamma}}$. The function $L_2(\boldsymbol{\theta}_{\pi}, \boldsymbol{\theta}_f|\tilde{\boldsymbol{\gamma}})$ is the joint likelihood of the initial conditions $(\boldsymbol{n}_1, S_1)$ and of the history of firms' entry, exit, and investment decisions conditional on the vector of firm types $\tilde{\boldsymbol{\gamma}}$:

$$(3) \quad L_2(\boldsymbol{\theta}_{\pi}, \boldsymbol{\theta}_f|\tilde{\boldsymbol{\gamma}}) = \Pr(\boldsymbol{n}_1, \mathrm{S}_1\,|\,\tilde{\boldsymbol{\gamma}}_1) \times \left[\prod_{t=1}^{T-1} P_E(\boldsymbol{n}_t^*, S_t)^{\bar{e}_t}(1 - P_E(\boldsymbol{n}_t^*, S_t))^{\hat{N}_t^{out} - \bar{e}_t}\right] \times \left[\prod_{i=1}^{N} \prod_{t=e_i+1}^{\min[x_i,T]} P_{NI}(w_{it}, \gamma_i, \boldsymbol{n}_t^*, S_t)^{1-i_{it}} \times P_I(w_{it}, \gamma_i, \boldsymbol{n}_t^*, S_t)^{1-i_{it}} \times \left[\frac{P_X(w_{ix_i}, \gamma_i, \boldsymbol{n}_{x_i}^*, S_{x_i})}{P_{NI}(w_{ix_i}, \gamma_i, \boldsymbol{n}_{x_i}^*, S_{x_i})}\right]^{I\{x_i \leq T\}}\right],$$

where P_{NI} is the probability of staying in the market without new investments, and $\tilde{\boldsymbol{\gamma}}_1$ is the vector of types of the incumbent firms at the beginning of $t = 1$. The probability $\Pr(\mathbf{n}_1, \mathrm{S}_1\,|\,\tilde{\boldsymbol{\gamma}}_1)$ represents the so-called *initial conditions problem*. If initial conditions are the outcome of equilibrium play under the same MPE that is played in the sample period, this probability can be obtained easily from the steady-state equilibrium distribution of the state variables $(\boldsymbol{n}_t^*, S_t)$, i.e., $p^*(\boldsymbol{n}_t^*, S_t)$. This steady-state distribution depends on the equilibrium probabilities $\mathbf{P}$ and on the primitives $\{f_S, f_{w0}, f_{w1}\}$, and it can be derived by solving a system of linear equations (see AM for details). A property of

$\Pr(\boldsymbol{n}_1, S_1 | \tilde{\boldsymbol{\gamma}}_1)$ that is key for the global concavity of the pseudo likelihood that we define below is that it depends on the parameters $\boldsymbol{\theta}_\pi$ only through the equilibrium probabilities.

Given this structure of the full likelihood, we can write the log likelihood as the sum of two components: $\log L(\boldsymbol{\theta}_\pi, \boldsymbol{\theta}_f) = \log L_1(\boldsymbol{\theta}_f) + \log\left\{\Sigma_{\tilde{\boldsymbol{\gamma}}} \Pr(\tilde{\boldsymbol{\gamma}}) L_2(\boldsymbol{\theta}_\pi, \boldsymbol{\theta}_f, \tilde{\boldsymbol{\gamma}})\right\}$. Following a standard approach in the estimation of dynamic structural models, we consider partial maximum likelihood estimation. First, we obtain the (partial) maximum likelihood estimator of $\boldsymbol{\theta}_f$ that maximizes $\log L_1(\boldsymbol{\theta}_f)$. This is a standard maximum likelihood estimation problem, and it does not require that we solve the dynamic game. Then, given $\hat{\boldsymbol{\theta}}_f$, we consider the estimation of $\boldsymbol{\theta}_\pi$ using the partial likelihood $\log\left\{\Sigma_{\tilde{\boldsymbol{\gamma}}} \Pr(\tilde{\boldsymbol{\gamma}}) L_2(\boldsymbol{\theta}_\pi, \hat{\boldsymbol{\theta}}_{\mathbf{f}}, \tilde{\boldsymbol{\gamma}})\right\}$. There are three main econometric and computational issues that we have to deal with: multiple equilibria; the computational cost associated with the repeated computation of an equilibrium of the model; and integration over all possible values of $\tilde{\boldsymbol{\gamma}}$. The first two problems were the main concern in AM. To deal with these problems they proposed a procedure, which they called Nested Pseudo Likelihood (NPL), that avoids the repeated solution of the game, and that can be used when the model has multiple equilibria. The main idea of this procedure is relatively simple. While equilibrium probabilities are not unique functions of structural parameters, the best response probabilities $\boldsymbol{\Psi} \equiv \{\Psi_E, \Psi_I, \Psi_X\}$ in equation (1) are always unique functions of structural parameters and firms' beliefs. We use these best response functions to construct a pseudo likelihood function $Q(\boldsymbol{\theta}_\pi, \mathbf{P}) \equiv \log\left\{\Sigma_{\tilde{\boldsymbol{\gamma}}} \Pr(\tilde{\boldsymbol{\gamma}}) L_2^{\Psi}(\boldsymbol{\theta}_\pi, \hat{\boldsymbol{\theta}}_{\mathbf{f}}, \mathbf{P} | \tilde{\boldsymbol{\gamma}})\right\}$, where $L_2^{\Psi}(\cdot)$ is the likelihood in equation (3), where the best response probabilities $\boldsymbol{\Psi}$ replace the equilibrium probabilities, and $\mathbf{P}$ is the vector of beliefs that we use to evaluate firms' best responses. If the pseudo likelihood function is based on a consistent nonparametric estimator $\hat{\mathbf{P}}$ of the equilibrium beliefs in the population, we can get a two-step estimator that is consistent and asymptotically normal. This two-step method cannot be applied to models with market or firm unobserved heterogeneity, however, because consistent nonparametric estimates of choice probabilities are not available for these models. Instead, AM proposed the NPL, which is a recursive extension of this two-step method. The NPL procedure starts with an arbitrary vector $\hat{\mathbf{P}}_0$. Given these initial probabilities, we generate a sequence of $\{\hat{\boldsymbol{\theta}}_{\pi,k}, \hat{\mathbf{P}}_k\}$ such that (1) $\hat{\boldsymbol{\theta}}_{\pi,k} = \arg\max_{\boldsymbol{\theta}_\pi} Q(\boldsymbol{\theta}_\pi, \hat{\mathbf{P}}_{k-1})$; and (2) $\hat{\mathbf{P}}_k = \boldsymbol{\Psi}(\hat{\boldsymbol{\theta}}_{\pi,k}, \hat{\mathbf{P}}_{k-1})$. Upon convergence, this procedure provides a consistent estimator of the structural parameters. Step (1) is a pseudo ML estimation and it is a very simple task in our model. Given $\hat{\mathbf{P}}_{k-1}$, we can construct the $\boldsymbol{z}$'s and λ's that appear in the best response probability functions. Furthermore, we use these probabilities to construct the steady-state distribution of $(\boldsymbol{n}_t^*, S_t)$ and then the probability of the initial conditions $\Pr(\boldsymbol{n}_1, S_1 | \tilde{\boldsymbol{\gamma}}_1)$, which are fixed during the pseudo-ML estimation. Step (2) is a policy iteration, and it consists of the evaluation of the expressions on the right-hand side of equation (1) using the $\boldsymbol{z}$'s and λ's from the previous iteration $(k - 1)$ and the new value of the structural parameters $\hat{\boldsymbol{\theta}}_{\pi,k}$.

Note that our pseudo-likelihood function $Q(\boldsymbol{\theta}_\pi, \mathbf{P})$ is globally concave in $\boldsymbol{\theta}_\pi$ for any value of $\mathbf{P}$. This is important because it facilitates the computation of the procedure and guarantees the consistency of the NPL estimator. Essentially, global concavity results from the linearity of the profit function in the structural parameters $\boldsymbol{\theta}_\pi$ combined with our assumption that the distribution of unobserved firm types is known up to mean and variance. Given these assumptions, it is possible to incorporate firm heterogeneity in any of the other parameters of the profit function (e.g., θ_v, θ_{FC} or θ_{EC}) keeping the global concavity of the pseudo likelihood. Furthermore, it is also possible to extend this method to allow for permanent unobserved heterogeneity in the transition probabilities of market size or firm quality, as long as the partial likelihood L_1 remains concave in $\boldsymbol{\theta}_f$, and a full information pseudo likelihood including L_1 is used in step (1) of each NPL iteration.

A new computational issue that appears when dealing with unobserved firm heterogeneity is that we have to integrate over $\tilde{\boldsymbol{\gamma}}$. The dimension of the set $\boldsymbol{\Gamma}^N$ of possible values of $\tilde{\boldsymbol{\gamma}}$ increases exponentially with the number of firms that we observe in a market, N. For most oligopoly markets, N is small and integration over all $\tilde{\boldsymbol{\gamma}}$'s may be computationally inexpensive. For instance, with $N = 10$ and two types, the number of possible values of $\tilde{\boldsymbol{\gamma}}$ is just $2^{10} = 1{,}024$. With two types and $N \geq 20$, however, we have more than one million possible values for $\tilde{\boldsymbol{\gamma}}$. In these cases, simulation techniques can be used to integrate over the distribution of types. Simulation also can be used to approximate the solution to the system of linear equations that defines the z's and λ's in firms' best response probabilities, as in Bajari, Benkard, and Levin (2007). The combination of these simulation techniques, that further reduces the cost of estimating these models, with the NPL method, and which allows for permanent unobserved heterogeneity, is an interesting topic for further research.

REFERENCES

Aguirregabiria, Victor, and Pedro Mira. 2007. "Sequential Estimation of Dynamic Discrete Games." *Econometrica*, 75(1): 1–53.

Aguirregabiria, Victor, Pedro Mira, and Hernan Roman. 2006. "Entry, Survival, and Growth in Oligopoly Retail Markets: Explaining Cross-Industry Heterogeneity." Unpublished.

Anderson, Simon P., André de Palma, and Jacques-François Thisse. 1992. *Discrete Choice Theory of Product Differentiation*. Cambridge, MA, and London: MIT Press.

Bajari, Patrick, C. Lanier Benkard, and Jonathan Levin. Forthcoming. "Estimating Dynamic Models of Imperfect Competition." *Econometrica*.

Pakes, Ariel, Michael Ostrovsky, and Steve Berry. Forthcoming. "Simple Estimators for the Parameters of Discrete Dynamic Games (with Entry/Exit Examples)." *RAND Journal of Economics*.

Pesendorfer, Martin, and Philipp Schmidt-Dengler. 2002. "Asymptotic Least Squares Estimators for Dynamic Games." Unpublished.