Chapter 5. Dynamic discrete games with incomplete information

Joan Llull

Structural Empirical Methods for Labor Economics (and Industrial Organization)

IDEA PhD Program

I. Introduction

The estimation of oligopolistic discrete games is very popular in industrial organization (IO). Dynamic games are important for the analysis of situations that involve dynamic strategic interactions between multiple agents. The crucial element of these interactions is that player's current decisions affect their own payoffs and those of other players today and in the future. Players make forward looking decisions taking into account the implications of their choices and of the expected behavior of their opponents.

II. Motivating example: market entry and exit

As a motivating example, we consider firms' decision of entry/continuation/exit from a market. Markets are small and isolated, and we focus on the retail sector. Each active firm operates at most in one location or store. We observe a random sample of markets $m = 1, 2, ..., M$. In each market, there are N potentially active and infinitely lived firms which decide simultaneously whether to operate or not. The distinction here between firms and markets is very important. We have $N \times M$ firms (we assume that firms at most operate in one market) whose payoffs are affected by the decisions of other players in the same market. This is the main difference with what we have seen so far.

In our example, the profit function is:

$$\Pi_{it} = \theta_{RS} \ln S_{m(i)t} - \theta_{RN} \ln \left(1 + \sum_{\{j : j \in m(i), j \neq i\}} d_{jt} \right), \tag{1}$$

where $d_{jt} = 1$ if firm j is active in marked $m(j)$ and $S_{m(i)t}$ is the market size. This profit function can be interpreted as the outcome of a static symmetric game. In a competitive setup, there are so many players that individual firm's decisions do not affect decisions of other firms, as the measure of firms operating in the market is unaffected. However, in an oligopolistic market, the number of firms operating in the market is a best-response function with respect to competitor's choices. The fixed cost paid every year by a firm active on a market is:

$$F_{it} = \omega_{Fi} + \varepsilon_{1it}, \tag{2}$$

where ε_{1it} represents an idiosyncratic shock to firm i's fixed cost. The entry cost, paid only when the firm was not active in the previous period, is given by:

$$E_{it} = (1 - d_{it-1})\omega_{Ei}. \tag{3}$$

Hence, the total current profits of an active firm with observable state variables $\boldsymbol{x}_{it} \equiv (S_{m(i)t}, \boldsymbol{d}'_{-i}, d_{it-1})'$, where $\boldsymbol{d}_{-i}$ is the vector of choices for all firms in the same market as i except i itself, are:

$$U_{it}(1, \boldsymbol{x}_{it}, \boldsymbol{\omega}_i, \boldsymbol{\varepsilon}_{it}) = \tag{4}$$
$$\theta_{RS} \ln S_{m(i)t} - \theta_{RN} \ln \left(1 + \sum_{\{j:j\in m(i), j\neq i\}} d_{jt}\right) - \omega_{Fi} - (1 - d_{it-1})\omega_{Ei} + \varepsilon_{1it}.$$

The profit function of a non-active firm is given by the value of the best outside option, that we assume to be:

$$U_{it}(0, \boldsymbol{x}_{it}, \boldsymbol{\omega}_i, \boldsymbol{\varepsilon}_{it}) = \omega_{Ni} + \varepsilon_{0it}, \tag{5}$$

where ω_{Ni} is normalized to zero for identification purposes.

The unobserved state vector $\boldsymbol{\varepsilon}_{it}$ is assumed to be private information of firm i, unknown by other players in the market. We assume that it is normally distributed, i.i.d. across firms and markets, and over time, with zero mean. $\boldsymbol{x}_{it}$ and $\boldsymbol{\omega}_i$ are common knowledge. The researcher observes $\boldsymbol{x}_{it}$ but not $\boldsymbol{\varepsilon}_{it}$ and $\boldsymbol{\omega}_i$.

Note that this model embeds Rust's framework with unobserved heterogeneity whenever $\theta_{RN} = 0$. However, if $\theta_{RN} \neq 0$, $d^*_{it}(\boldsymbol{x}_{it}, \boldsymbol{\omega}_i, \boldsymbol{\varepsilon}_{it})$ is a best response function. The actual decisions are given by the solution of the Nash equilibrium. Full solution methods are often unfeasible in this context. Aguirregabiria and Mira (2007) propose a Hotz-Miller based estimation method for dynamic discrete games. A potential complication of this may be the existence of multiple equilibria, but it can also be handled in estimation.

III. General structure

One of the costly aspects of dynamic discrete games is the need for solving for the equilibrium in each period. If in single agent problems solving for the value functions implies a nested fixed points, in the games context there is a double nesting because on top of solving for the individual's problem, one also needs to solve for the equilibrium conditions. CCP estimations has turned particularly important in developing techniques that allowed us to advance in the possibilities of estimating such problems.

To generalize the setup from the example above, let d_{it} denote the own action of individual i, and let $\boldsymbol{d}_{-it} \equiv (d_{1t}, ..., d_{i-1,t}, d_{i+1,t}, ..., d_{It})'$ denote the actions of all other $I-1$ players. The flow utility of individual i choosing alternative j is thus:

$$u_{ij}(\boldsymbol{x}_t, \boldsymbol{d}_{-it}) + \varepsilon_{ijt}, \tag{6}$$

where $\boldsymbol{\varepsilon}_{it} \equiv (\varepsilon_{i1t}, ..., \varepsilon_{iJt})'$ is an i.i.d. random variable privately observed by individual i but not by the others. The vector $\boldsymbol{x}_t$, on the contrary, is observed by all individuals, and, therefore, includes the state variables of all individuals. The dependence of $u_{ij}(\cdot)$ on i is a reflection of the possibility of different state variables affecting payoffs differently (e.g. own state variables vs other individuals'); the absence of t in it reflects that we are in a stationary environment.

Choices are taken simultaneously in each period. We concentrate on rational stationary Markov perfect equilibria, which imply that, given that $\boldsymbol{\varepsilon}_{it}$ is i.i.d. across individuals, individual i expects other agents to make choices $\boldsymbol{d}_{-it}$ with probabilities:

$$\Pr(\boldsymbol{d}_{-it}|\boldsymbol{x}_t) = \prod_{j \neq i} \Pr(d_{jt}|\boldsymbol{x}_t). \tag{7}$$

These CCPs represent the best-response probability functions. In this type of models, an equilibrium exists, but uniqueness is rather unlikely to hold. However, these CCPs uniquely identify the beliefs of agents.

Taking expectations of $u_{ij}(\boldsymbol{x}_t, \boldsymbol{d}_{-it})$ over $\boldsymbol{d}_{-it}$, we obtain:

$$\tilde{u}_{ij}(\boldsymbol{x}_t) = \sum_{\boldsymbol{d}_{-it} \in \mathcal{D}^{I-1}} \Pr(\boldsymbol{d}_{-it}|\boldsymbol{x}_t) u_{ij}(\boldsymbol{x}_t, \boldsymbol{d}_{-it}). \tag{8}$$

These concentrated payoffs resemble the standard payoffs of the single-agent models seen so far. Now we additionally need to construct the continuation values. The difficulty on that front is that the state variables, which follow a process given by $F_x(\boldsymbol{x}_{t+1}|\boldsymbol{x}_t, d_{it}, \boldsymbol{d}_{-it})$, depend on the unobserved choices by the other individuals. Following a similar idea, we can obtain the concentrated transition functions as:

$$\tilde{F}_i(\boldsymbol{x}_{t+1}|\boldsymbol{x}_t, d_{it}) = \sum_{\boldsymbol{d}_{-it} \in \mathcal{D}^{I-1}} \Pr(\boldsymbol{d}_{-it}|\boldsymbol{x}_t) F_x(\boldsymbol{x}_{t+1}|\boldsymbol{x}_t, d_{it}, \boldsymbol{d}_{-it}). \tag{9}$$

Given all this, the conditional value functions can be expressed as:

$$v_{ij}(\boldsymbol{x}_t) = \tilde{u}_{ij}(\boldsymbol{x}_t) + \beta \int V(\boldsymbol{x}_{t+1}) d\tilde{F}_i(\boldsymbol{x}_{t+1}|\boldsymbol{x}_t, d_{it}), \tag{10}$$

and then $V(\boldsymbol{x}_{t+1})$ can be replaced the standard CCP representation, i.e.:

$$v_{ij}(\boldsymbol{x}_t) = \tilde{u}_{ij}(\boldsymbol{x}_t) + \beta \int \left[v_{ik}(\boldsymbol{x}_t) + \psi_k(\boldsymbol{p}_i(\boldsymbol{x}_t))\right] d\tilde{F}_i(\boldsymbol{x}_{t+1}|\boldsymbol{x}_t, d_{it}). \tag{11}$$

IV. Identification and estimation

The use of CCP estimation methods made the estimation of these models feasible. Full solution maximum likelihood approaches would not be tractable, because it would require solving for the equilibrium of the game (on top of solving for the dynamic problem). All the discussion below is predicated on the following two assumptions: 1) every observation in the sample comes from the same Markov Perfect Equilibrium, and 2) there are no unobserved common-knowledge variables. Given the first assumption, the multiplicity of equilibria in the model does not play any role in the identification of structural parameters. For non-parametric identification of the payoff functions, one further needs to assume a known discount factor, and the presence of an exclusion restriction (i.e., there is some observable variable that is player-specific and that has no direct effect on opponents' payoff, that is, $\boldsymbol{x}_t \equiv (s_{1t}, ..., s_{Nt}, \boldsymbol{w}_t)$ is such that $u_{ij}(\boldsymbol{x}_t, d_{-it}) = u_{ij}(s_{it}, \boldsymbol{w}_t, d_{-it})$.

A. Aguirregabiria and Mira (2007)

Once the model is transformed in terms of $\tilde{u}_{ij}(\boldsymbol{x}_t)$ and $\tilde{F}_i(\boldsymbol{x}_t|\boldsymbol{x}_t, d_{it})$, estimation follows standard CCP procedures, which range from likelihood or GMM versions of Hotz and Miller (1993) to the iterative Nested Pseudo-Likelihood algorithm in Aguirregabiria and Mira (2002). The latter is the approach adopted in the seminal paper by Aguirregabiria and Mira (2007). Kasahara and Shimotsu (2008) show that when the equilibrium in the population is stable, then this recursive procedure reduces significantly the small sample bias generated in the one-step CCP based methods, which, in the case of games increases rapidly both with the number of state variables and in the number of players.

B. Bajari, Benkard, and Levin (2007)

Bajari, Benkard, and Levin (2007) provide a similar alternative, which is based on the Hotz, Miller, Sanders, and Smith (1994) forward simulation method. In particular, these authors propose the following three-step method. First, estimate the transition functions F and CCPs from the data. Second, forward simulate a sequence of optimal choices and states drawing from the distributions dictated by CCPs and transition functions. Finally obtain the parameters maximizing your preferred pseudo-likelihood or GMM criterion function. This method, like Aguirregabiria and Mira (2007), allows the researcher to be agnostic about equilibrium selection, and provides a side-step to the problem of multiple equilibria.