ECO 2901
EMPIRICAL INDUSTRIAL ORGANIZATION

Lecture 11:
Dynamic games with biased beliefs and learning

Victor Aguirregabiria (University of Toronto)

March 28, 2019

Dynamic games with biased beliefs and learning: Outline

- 1. Introduction
- 2. Model
- 3. Identification of Beliefs
- 4. Testing learning models

1. Introduction

Introduction

- Firms have uncertainty about future demand, costs, or behavior of competitors.
- They may learn about these elements over time, and this learning process can have substantial implications for their profits and market efficiency.
- The importance of firms' heterogeneous expectations and the implications on firms' performance and market outcomes have been long recognized in economics, at least since the work of Herbert Simon (1958, 1959).
- However, the assumption of rational expectations has been the status quo to represent agents' beliefs in many areas in economics, and in particular in IO.
- It has not been until recently that firms' biased beliefs and learning has received substantial attention in structural models in empirical IO.

Introduction / Outline

- [1] I present a dynamic game of oligopoly competition that allows for biased beliefs and learning, but it is agnostic about the source of the biased beliefs and the form of learning (if any).
- [2] We study nonparametric identification of firms' belief functions in this model.
- [3] Given estimated beliefs, we use them identify different possible forms of firms' learning.

2. Model

Model: Dynamic Game

- N players indexed by i. Every period t, each player takes an action $a_{it} \in \{0, 1, ..., J\}$.
- One-period payoff function is:

$$\Pi_{it} = \pi_{it}(a_{it}, \mathbf{a}_{-it}, \mathbf{x}_t) + \varepsilon_{it}(x_{it})$$

- $\mathbf{x}_t$ is a vector of common knowledge state variables with transition $f_t\left(\mathbf{x}_{t+1} \mid a_{it},\ \mathbf{a}_{-it},\ \mathbf{x}_t\right)$.
- $\varepsilon'_{it}s$ are private info of player i and unobservable to researcher. It is i.i.d. over time and players.

Basic Assumptions

- We maintain some of the assumptions in the concept of Markov Perfect Equilibrium (MPE).

ASSUMPTION 1 (Payoff relevant state variables): *Players' strategy functions depend only on payoff relevant state variables:* $\mathbf{x}_t$ and ε_{it}.

ASSUMPTION 2 (Maximization of expected payoffs): *Players are forward looking and maximize expected intertemporal payoffs.*

ASSUMPTION 3 (Rational beliefs on own future behavior): *Players have rational expectations on their own behavior in the future.*

Strategies, Choice Probabilities, and Beliefs

- Let $\sigma_{it}(\mathbf{x}_t, \varepsilon_{it})$ be the strategy function for player i at period t.
- $P_{it}(a_i|\mathbf{x}_t) \equiv \Pr(\sigma_{it}(\mathbf{x}_t, \varepsilon_{it}) = a_i|\mathbf{x}_t)$ choice probability of player i.
- $B_{it+s}^{(t)}(\mathbf{a}_{-i}|\mathbf{x}_{t+s})$ of player i at period t about the behavior of other players at period $t+s$.
- The model allows the belief functions $B_{it+s}^{(t)}$ to vary freely both over t (i.e., over the period when these beliefs are formed) and over $t+s$ (i.e., over the period of the other players' behavior).
- In particular, the model allows players to update their beliefs and learn (or not) over time t.

Sequence of Beliefs $B_{it+s}^{(t)}$

Beliefs formed (t)	Period of the opponents' behavior $(t+s)$ $t+s=1$	$t+s=2$	$t+s=3$	...	$t+s=T-1$	$t+s$
$t=1$	$B_{i1}^{(1)}$	$B_{i2}^{(1)}$	$B_{i3}^{(1)}$	...	$B_{i,T-1}^{(1)}$	B
$t=2$	-	$B_{i2}^{(2)}$	$B_{i3}^{(2)}$	...	$B_{i,T-1}^{(2)}$	B
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	
$t=T-1$	-	-	-	...	$B_{i,T-1}^{(T-1)}$	B_{iT}
$t=T$	-	-	-	...	-	B

Best Response Functions

- Given her beliefs at period t, $\mathbf{B}_i(t) = \left\{ B_{i,t+s}^{(t)} : s \geq 0 \right\}$, a player best response at period t is the solution of a single-agent Dynamic Programming problem.

- This DP problem can be described in terms of: (1) a discount factor; (2) a sequence of expected one-period payoff functions:

$$\pi_{it+s}^{\mathbf{B}(t)}(a_{it+s}, \mathbf{x}_{t+s}) \equiv \sum_{\mathbf{a}_{-i}} B_{it+s}^{(t)}(\mathbf{a}_{-i}|\mathbf{x}_{t+s}) \ \pi_{it+s}(a_{it+s}, \mathbf{a}_{-i}, \mathbf{x}_{t+s})$$

- And (3) a sequence of transition probability functions:

$$f_{it+s}^{\mathbf{B}(t)}(\mathbf{x}_{t+s+1}|a_{it+s}, \mathbf{x}_{t+s}) \equiv \sum_{\mathbf{a}_{-i}} B_{it+s}^{(t)}(\mathbf{a}_{-i}|\mathbf{x}_{t+s}) \ f_{t+s}(\mathbf{x}_{t+s+1}|a_{it+s}, \mathbf{a}_{-i}, \mathbf{x}_{t}$$

Best Response Functions (2)

- The solution of this DP problem implies the vector of conditional choice value functions at period t:

$$v_{it}^{\mathbf{B}(t)}(\mathbf{x}_t) = \left\{ v_{it}^{\mathbf{B}(t)}(a_i, \mathbf{x}_t) : a_i = 0, 1, ..., J \right\}$$

- And the best response choice probabilities:

$$\begin{aligned} P_{it}(a_i|\mathbf{x}_t) &= \Pr\left(v_{it}^{\mathbf{B}(t)}(a_i, \mathbf{x}_t) + \varepsilon_{it}(a_i) \geq v_{it}^{\mathbf{B}(t)}(a_i', \mathbf{x}_t) + \varepsilon_{it}(a_i') \ \forall a_i' \right. \\ &= \Lambda_i \left(a_i \ ; v_{it}^{\mathbf{B}(t)}(\mathbf{x}_t) \right) \end{aligned}$$

- For instance, in a logit model:

$$P_{it}(a_i|\mathbf{x}_t) = \frac{\exp\left\{ v_{it}^{\mathbf{B}(t)}(a_i, \mathbf{x}_t) \right\}}{\sum_{j=0}^{J} \exp\left\{ v_{it}^{\mathbf{B}(t)}(j, \mathbf{x}_t) \right\}}$$

Restrictions on Beliefs

- In this model, the sequence of beliefs $\mathbf{B}_i(t) = \left\{ B_{i,t+s}^{(t)} : s \geq 0 \right\}$ are completely unrestricted.

- We can think in different type of restrictions on beliefs:
  - Markov Perfect equilibrium, or other equilibrium concepts;
  - Level-K Rationality at each period t (with or without learning);
  - Bayesian Learning about true CCPs of other players;
  - Other forms of learning: Adaptive learning, Reinforced learning, etc.

- Here we consider the following two-step approach:

  [1] identification/estimation of (some) beliefs in $\mathbf{B}_i(t)$ without imposing any restriction;

  [2] testing for different forms of learning.

3. Identification of Beliefs

Data

- We have a random sample of M markets, indexed by m, where we observe

$$\{a_{imt},\ \mathbf{x}_{mt} : i = 1, 2, ..., N;\ t = 1, 2, ..., T^{data}\}$$

- N and T^{data} are small and M is large.

- The payoff functions $\pi_{it}(a_{it}, \mathbf{a}_{-it}, \mathbf{x}_t)$ and the beliefs functions $B_{it+s}^{(t)}(\mathbf{a}_{-i}|\mathbf{x}_{t+s})$ are nonparametrically specified.

- The distribution of the unobservables Λ is assumed known. This can be relaxed if there is a "special" state variable $z_{it}(a_{it})$ that enters additively in π_{it}.

- I focus here in a model with two players, i and j, but the results can be extended to N players.

Inversion of CCPs

- The model is described by the conditions:

$$P_{it}(a_i|\mathbf{x}_t) \quad = \quad \Lambda\left(a_i \; ; v_{it}^{\mathbf{B}(t)}(\mathbf{x}_t)\right)$$

- The CCPs $P_{it}(a_i|\mathbf{x}_t)$ are identified using data from M markets.

- Hotz-Miller inversion theorem implies that we can invert the best response mapping to obtain value differences $\widetilde{v}_{it}^{\mathbf{B}(t)}(a_i, \mathbf{x}_t) \equiv v_{it}^{\mathbf{B}(t)}(a_i, \mathbf{x}_t) - v_{it}^{\mathbf{B}(t)}(a_i, \mathbf{x}_t)$ as functions of CCPs:

$$\widetilde{v}_{it}^{\mathbf{B}(t)}(a_i, \mathbf{x}_t) \quad = \quad \Lambda^{-1}\left(a_i \; ; \mathbf{P}_{it}(\mathbf{x}_t)\right)$$

- The identification problem is to obtain beliefs and payoff functions given that $\Lambda^{-1}\left(a_i \; ; \mathbf{P}_{it}(\mathbf{x}_t)\right)$ are known.

Inversion of CCPs

- The model is described by the conditions:

$$P_{it}(a_i|\mathbf{x}_t) \quad = \quad \Lambda\left(a_i \; ; v_{it}^{\mathbf{B}(t)}(\mathbf{x}_t)\right)$$

- The CCPs $P_{it}(a_i|\mathbf{x}_t)$ are identified using data from M markets.

- Hotz-Miller inversion theorem implies that we can invert the best response mapping to obtain value differences $\widetilde{v}_{it}^{\mathbf{B}(t)}(a_i, \mathbf{x}_t) \equiv v_{it}^{\mathbf{B}(t)}(a_i, \mathbf{x}_t) - v_{it}^{\mathbf{B}(t)}(a_i, \mathbf{x}_t)$ as functions of CCPs:

$$\widetilde{v}_{it}^{\mathbf{B}(t)}(a_i, \mathbf{x}_t) \quad = \quad \Lambda^{-1}\left(a_i \; ; \mathbf{P}_{it}(\mathbf{x}_t)\right)$$

- The identification problem is to obtain beliefs and payoff functions given that $\Lambda^{-1}\left(a_i \; ; \mathbf{P}_{it}(\mathbf{x}_t)\right)$ are known.

Structure of the restrictions

- By definition the value differences $\widetilde{v}_{it}^{\mathbf{B}(t)}(a_i, \mathbf{x}_t)$ have the following structure:

$$\widetilde{v}_{it}^{\mathbf{B}(t)}(a_i, \mathbf{x}_t) \quad = \quad \mathbf{B}_{it}^{(t)}(\mathbf{x}_t)' \, \left[\widetilde{\boldsymbol{\pi}}_{it}(a_{it}, \mathbf{x}_t) + \widetilde{\mathbf{c}}_{it}^{\mathbf{B}(t)}(a_{it}, \mathbf{x}_t) \right]$$

  where $\mathbf{B}_{it}^{(t)}(\mathbf{x}_t)$ is the vector of beliefs $\left[B_{it}^{(t)}(\mathbf{a}_{-i}|\mathbf{x}_t)\right.$ for any value $\left.\mathbf{a}_{-i}\right]$.

- $\widetilde{\boldsymbol{\pi}}_{it}(a_{it}, \mathbf{x}_t)$ is the vector of payoff differences $\left[\pi_{it}(a_{it}, \mathbf{a}_{-i}, \mathbf{x}_t) - \pi_{it}(0, \mathbf{a}_{-i}, \mathbf{x}_t)\right.$ for any value $\left.\mathbf{a}_{-i}\right]$.

- $\widetilde{\mathbf{c}}_{it}^{\mathbf{B}(t)}(a_{it}, \mathbf{x}_t)$ is the vector of **differences of continuation values** $\left[c_{it}^{\mathbf{B}(t)}(a_{it}, \mathbf{a}_{-i}, \mathbf{x}_t) - c_{it}^{\mathbf{B}(t)}(0, \mathbf{a}_{-i}, \mathbf{x}_t)\right.$ for any value $\left.\mathbf{a}_{-i}\right]$ where:

$$c_{it}^{\mathbf{B}(t)}(a_{it}, \mathbf{a}_{-i}, \mathbf{x}_t) = \beta \sum V_{it+1}^{\mathbf{B}(t)}(\mathbf{x}_{t+1}) \; f_t(\mathbf{x}_{t+1}|a_{it}, \mathbf{a}_{-i}, \mathbf{x}_t)$$

Identification Assumptions

- **ASSUMPTION ID-1.** A player has the same beliefs in markets with the same $\mathbf{x}$ variables.

$$B_{imt+s}^{(t)}(.|\mathbf{x}) = B_{it+s}^{(t)}(.|\mathbf{x}) \quad \text{for any market } m$$

- **ASSUMPTION ID-2 (Exclusion Restriction 1)**: $\mathbf{x}_t = (s_{it}, s_{jt}, \mathbf{w}_t)$ such that s_{it} enters in the payoff function of player i but not in the payoff of the other player.

$$\pi_{it}\left(a_{it}, a_{jt}, s_{it}, s_{jt}, \mathbf{w}_t\right) = \pi_{it}\left(a_{it}, a_{jt}, s_{it}, \mathbf{w}_t\right)$$

- **ASSUMPTION ID-3 (Exclusion Restriction 2)**: The transition probability of the state variable s_{it} is such that the value of s_{it+1} does not depend on (s_{it}, s_{jt}):

$$f_t\left(s_{it+1} \mid a_{it}, s_{it}, s_{jt}, \mathbf{w}_t\right) = f_t\left(s_{it+1} \mid a_{it}, s_{it}, \mathbf{w}_t\right)$$

Exclusion restriction in the payoff function (ID-2)

- The exclusion restriction ID-2 in the payoff function π_{it} appears naturally in many applications of dynamic games of oligopoly competition.
- Incumbent status, capacity, capital stock, or product quality of a firm at period $t-1$ are state variables that enter in a firm's payoff function at period t, π_{it}, because there are investment and adjustment costs that depend on these lagged variables.
- A firm's payoff π_{it} depends also on the competitors' values of these variables at period t, but it does not depend on the competitors' values of these variables at $t-1$.
- Importantly, the assumption does not mean that player i does not condition her behavior on those excluded variables. Each player conditions his behavior on all the (common knowledge) state variables that affect the payoff of a player in the game, even if these variables

Exclusion restriction in the transition probability (ID-3)

$$f_t\left(s_{it+1} \mid a_{it}, s_{it}, s_{jt}, \mathbf{w}_t\right) = f_t\left(s_{it+1} \mid a_{it}, s_{it}, \mathbf{w}_t\right)$$

- An important class of models that satisfies this condition is when $s_{it} = a_{i,t-1}$, such that the transition rule is simply:

$$s_{it+1} = a_{it}$$

- Many dynamic games of oligopoly competition belong to this class, e.g., market entry/exit, technology adoption, and some dynamic games of quality or capacity competition, among others.

Example: Quality competition

- Quality ladder dynamic game (Pakes and McGuire, 1994).
- s_{it} is the firm's quality at $t-1$.
- The decision variable a_{it} is the firm's quality at period t, such that:

$$s_{it+1} = a_{it}$$

- The model is dynamic because the payoff function includes a cost of adjusting quality that depends on $a_{it} - s_{it}$:

$$AC_i\left(a_{it} - s_{it}\right)$$

- Given competitors quality at period t, a_{jt}, firm i's profit does not depend on competitors' qualities at $t-1$.

Role of the Exclusion restrictions

$$\ln\left(\frac{P_{it}\left(a_i \mid s_{it}, \mathbf{s}_{-it}\right)}{P_{it}\left(0 \mid s_{it}, \mathbf{s}_{-it}\right)}\right) \quad = \quad \mathbf{B}_{it}^{(t)}(s_{it}, \mathbf{s}_{-it})' \left[\widetilde{\boldsymbol{\pi}}_{it}(a_{it}, s_{it}) + \widetilde{\mathbf{c}}_{it}^{\mathbf{B}(t)}(a_{it}, s_{it})\right]$$

- Under the two exclusion restrictions, the state variables $\mathbf{s}_{-it}$ (the competitors s_j) do not enter in the payoffs $\widetilde{\boldsymbol{\pi}}_{it}(a_{it}, s_{it})$ and on the continuation values $\widetilde{\mathbf{c}}_{it}^{\mathbf{B}(t)}(a_{it}, s_{it})$.
- Note: Though $\widetilde{\mathbf{c}}_{it}^{\mathbf{B}(t)}(a_{it}, s_{it})$ depends on beliefs, these are beliefs at periods $t+s > t$ and therefore depend on $(s_{it+s}, \mathbf{s}_{-it+s})$ for $t+s > t$.
- Therefore, the dependence of $\ln\left(\frac{P_{it}\left(a_i \mid s_{it}, \mathbf{s}_{-it}\right)}{P_{it}\left(0 \mid s_{it}, \mathbf{s}_{-it}\right)}\right)$ with respect to $\mathbf{s}_{-it}$ captures the dependence of beliefs $\mathbf{B}_{it}^{(t)}(s_{it}, \mathbf{s}_{-it})$ with respect to $\mathbf{s}_{-it}$.

Identification of Beliefs

- For any player i, any period t in the data, any value of $(\mathbf{a}_{-i}, s_{it})$, and any combination of three values $\mathbf{s}_{-it}$, say $(\mathbf{s}_{-i}^{(a)}, \mathbf{s}_{-i}^{(b)}, \mathbf{s}_{-i}^{(c)})$, the following function of beliefs is identified:

$$\frac{B_{it}^{(t)}(\mathbf{a}_{-i} \mid s_{it}, \mathbf{s}_{-i}^{(c)}) - B_{it}^{(t)}(\mathbf{a}_{-i} \mid s_{it}, \mathbf{s}_{-i}^{(a)})}{B_{it}^{(t)}(\mathbf{a}_{-i} \mid s_{it}, \mathbf{s}_{-i}^{(b)}) - B_{it}^{(t)}(\mathbf{a}_{-i} \mid s_{it}, \mathbf{s}_{-i}^{(a)})}$$

Identification of Beliefs [2]

- For instance, in a binary choice logit with two-players:

$$\frac{B_{it}^{(t)}(1 \mid s_{it}, \mathbf{s}_{-i}^{(c)}) - B_{it}^{(t)}(1 \mid s_{it}, \mathbf{s}_{-i}^{(a)})}{B_{it}^{(t)}(1 \mid s_{it}, \mathbf{s}_{-i}^{(b)}) - B_{it}^{(t)}(1 \mid s_{it}, \mathbf{s}_{-i}^{(a)})} =$$

$$\frac{\ln\left(\dfrac{P_{it}(1 \mid s_{it}, \mathbf{s}_{-i}^{(c)})}{P_{it}(0 \mid s_{it}, \mathbf{s}_{-i}^{(c)})}\right) - \ln\left(\dfrac{P_{it}(1 \mid s_{it}, \mathbf{s}_{-i}^{(a)})}{P_{it}(0 \mid s_{it}, \mathbf{s}_{-i}^{(a)})}\right)}{\ln\left(\dfrac{P_{it}(1 \mid s_{it}, \mathbf{s}_{-i}^{(b)})}{P_{it}(0 \mid s_{it}, \mathbf{s}_{-i}^{(b)})}\right) - \ln\left(\dfrac{P_{it}(1 \mid s_{it}, \mathbf{s}_{-i}^{(a)})}{P_{it}(0 \mid s_{it}, \mathbf{s}_{-i}^{(a)})}\right)}$$

- Note that we cannot identify beliefs about competitors' behavior at future periods: $B_{it+s}^{(t)}$ for $s > 0$. However, $B_{it}^{(t)}$ can provide substantial information about learning.

4. Testing learning models

Testing learning models

- Suppose that the researcher has identified the sequence of belief functions $B_{it}^{(t)}(a_{-it}|\mathbf{x}_t)$.

$$B_{it}^{(t)}(a_{-it}|\mathbf{x}_t) : t = 1, 2, ..., T^{data}$$

- Given these data on beliefs, we can test for different hypothesis about the evolution of beliefs.

- For notational simplicity, we represent these beliefs as if they were transition probabilities $B_{it}^{(t)}(\mathbf{x}_{t+1}|\mathbf{x}_t)$

Testing for Rational Beliefs

- Let $P_t(\mathbf{x}_{t+1}|\mathbf{x}_t)$ be the actual distribution of $\mathbf{x}_{t+1}$ conditional on $\mathbf{x}_t$ in the data. $P_t(\mathbf{x}_{t+1}|\mathbf{x}_t)$ is identified.

- Testing for Rational Beliefs is equivalent to testing for the restrictions:

$$B_{it}^{(t)}(\mathbf{x}_{t+1}|\mathbf{x}_t) = P_t(\mathbf{x}_{t+1}|\mathbf{x}_t)$$

Testing for Bayesian Learning

- Let $\mathcal{P}_i \equiv \{\psi_{\ell,i}(\mathbf{x}'|\mathbf{x}) : \ell = 1, 2, ..., L\}$ be a collection of L transition probabilities.
- The prior belief function for firm i at period $t = 0$ is a mixture of the distributions in $\mathcal{P}_i$, where $\{\lambda_{\ell,i}^{(0)}\}$ are the mixing probabilities.
- At any period $t \geq 1$, firms observe the new state $\mathbf{x}_t$ and use this information to update their respective beliefs using Bayes rule.

$$B_{it}^{(t)}(\mathbf{x}_{t+1}|\mathbf{x}_t) = \sum_{\ell=1}^{L} \lambda_{\ell,i}^{(t)}(\mathbf{x}', \mathbf{x}) \ \psi_{\ell,i}(\mathbf{x}'|\mathbf{x})$$

where Bayesian updating implies:

$$\lambda_{\ell,i}^{(t)}(\mathbf{x}', \mathbf{x}) = \frac{\psi_{\ell,i}(\mathbf{x}_t|\mathbf{x}_{t-1}) \ \lambda_{\ell,i}^{(t-1)}(\mathbf{x}', \mathbf{x})}{\sum_{\ell'=1}^{L} \psi_{\ell',i}(\mathbf{x}_t|\mathbf{x}_{t-1}) \ \lambda_{\ell',i}^{(t-1)}(\mathbf{x}', \mathbf{x})}$$

Testing for Adaptive Learning

- At period t:

$$B_{it}^{(t)}(\mathbf{x}'|\mathbf{x}) = (1-\delta_i) \;\; B_{it-1}^{(t-1)}(\mathbf{x}'|\mathbf{x}) + \delta_{it} \;\; K([\mathbf{x}_t, \mathbf{x}_{t-1}] - [\mathbf{x}', \mathbf{x}])$$

- $\delta_i \in (0,1)$ is a parameter that determines the speed of learning.
- $K(.)$ is a Kernel function that establishes whether the new information at period t is used to update beliefs only at that point or also at nearby values.

Testing for Fictitious Play

- Fictitious play is a learning rule where each firm believes that rivals' actions are sampled from the empirical distribution of their past actions.

- The belief function of firm i about the choice probability of firm j is:

$$B_{it}^{(t)}(a_j|\mathbf{x}) = \frac{\sum_{s=1}^{t} \omega_{(s,t)} \ 1\{[a_{jt-s}, \mathbf{x}_{t-s}] = [a_j, \mathbf{x}]\}}{\sum_{s=1}^{t} \omega_{(s,t)} \ 1\{\mathbf{x}_{t-s} = \mathbf{x}\}}$$

- $\{\omega_{(s,t)} : s \leq t\}$ are weights non-increasing in the lag index s.

- In its original version (Brown, 1951) the fictitious play model assumes that the weights $\omega_{(s,t)}$ are the same at every period s such that belief $B_{it}^{(t)}(a_j|\mathbf{x})$ is just the empirical frequency of action a_j conditional on state $\mathbf{x}$ during periods 1 to t.

Testing for Rationalizability

- *Rationalizability* (Bernheim, 1984; Pearce, 1984).

- The concept of rationalizability imposes two simple restrictions on firms' beliefs and behavior.
  - [A.1] Every firm is rational in the sense that it maximizes its own expected profit given beliefs.
  - [A.2] This rationality is common knowledge, i.e., every firms knows that all the firms know that it knows ... that all the firms are rational.

- We have impose [A.1] to identify beliefs, but we have not impose [A.2]. We can test for [A.2].

- The set of outcomes of the game that satisfy these conditions (the set of rationalizable outcomes) includes all the MPE Nash equilibria of the game, but it also includes many other outcomes too.

Testing for Level-K Rationality

- *Cognitive Hierarchy and Level-k Rationality*. These models assume that players have different levels of strategic sophistication.
- Every firm (player) maximizes its subjective expected profit given its beliefs.
- Firms are heterogeneous in their beliefs and there is a finite number of belief types.
- Beliefs for each type are determined by a hierarchical structure.
- Level-0 firms believe that strategic interactions are negligible and therefore they behave as in a single-agent model, i.e., as if they were monopolists.
- Level-1 firms believe that the rest of the firms are level-0, and they behave by best responding to these beliefs. And so on.