

1 Introduction to structure of dynamic oligopoly models

- Consider a simple two-firm model, and assume that all the dynamics are deterministic.
- Let x_{1t}, x_{2t} , denote the state variables for each firm in each period. Let q_{1t}, q_{2t} denote the control variables. Example: x 's are capacity levels, and q 's are incremental changes to capacity in each period.
- Assume (for now) that $x_{it+1} = g(x_{it}, q_{it})$, $i = 1, 2$, so that next period's state is a *deterministic* function of this period's state and control variable. (Can allow for cross-effects with no problem.)
- Firm i ($=1, 2$) chooses a sequence $q_{i1}, q_{i2}, q_{i3}, \dots$ to maximize its discounted profits:

$$\sum_{t=0}^{\infty} \beta^t \Pi(x_{1t}, x_{2t}, q_{1t}, q_{2t})$$

where $\Pi(\dots)$ denotes single-period profits.

- Because the two firms are duopolists, and they must make these choices recognizing that their choices can affect their rival's choices. We want to consider a dynamic equilibrium of such a model, when (roughly speaking) each firm's sequence of q 's is a "best-response" to its rival's sequence.
- A firm's strategy in period t , q_{it} , can potentially depend on the whole "history" of the game ($\mathcal{H}_{t-1} \equiv \{x_{1t'}, x_{2t'}, q_{1t'}, q_{2t'}\}_{t'=0, \dots, t-1}$), and well as on the time period t itself. This becomes quickly intractable, so we usually make some simplifying regularity conditions:
 - Firms employ *stationary* strategies: so that strategies are not explicitly a function of time t (i.e. they depend on time only indirectly, through the history \mathcal{H}_{t-1}). Given stationarity, we will drop the t subscript, and use primes $'$ to denote next-period values.
 - A dimension-reducing assumption is usually made: for example, we might assume that q_{it} depends only on x_{1t}, x_{2t} , which are the "payoff-relevant"

state variables which directly affect firm i 's profits in period i . This is usually called a "Markov" assumption. With this assumption $q_{it} = q_i(x_{1t}, x_{2t})$, for all t .

- Furthermore, we usually make a *symmetry* assumption, that each firm employs an identical strategy assumption. This implies that $q_1(x_{1t}, x_{2t}) = q_2(x_{2t}, x_{1t})$.
- To characterize the equilibrium further, assume we have an equilibrium strategy function $q^*(\cdot, \cdot)$. For each firm i , then, and at each state vector x_1, x_2 , this optimal policy must satisfy Bellman's equation, in order for the strategy to constitute subgame-perfect behavior:

$$q^*(x_1, x_2) = \operatorname{argmax}_q \{ \Pi(x_1, x_2, q, q^*(x_2, x_1)) + \beta V(x'_1 = g(x_1, q), x'_2 = g(x_2, q^*(x_2, x_1))) \} \quad (1)$$

from firm 1's perspective, and similarly for firm 2. $V(\cdot, \cdot)$ is the value function, defined recursively at all possible state vectors x_1, x_2 via the Bellman equation:

$$V(x_1, x_2) = \max_q \{ \Pi(x_1, x_2, q, q^*(x_2, x_1)) + \beta V(x'_1 = g(x_1, q), x'_2 = g(x_2, q^*(x_2, x_1))) \}. \quad (2)$$

- I have described the simplest case; given this structure, it is clear that the following extensions are straightforward:
 - Cross-effects: $x'_i = g(x_i, x_{-i}, q_i, q_{-i})$
 - Stochastic evolution: $x'_i | x_i, q_i$ is a random variable. In this case, replace last term of Bellman eq. by $E[V(x'_1, x'_2) | x_1, x_2, q, q_2 = q^*(x_2, x_1)]$.
This expectation denotes player 1's *equilibrium beliefs* about the evolution of x_1 and x_2 (equilibrium in the sense that he assumes that player 2 plays the equilibrium strategy $q^*(x_2, x_1)$).
 - > 2 firms
 - Firms employ asymmetric strategies, so that $q_1(x_1, x_2) \neq q_2(x_2, x_1)$
 - ...

- Computing the equilibrium strategy $q^*(\dots)$ consists in iterating over the Bellman equation (1). However, the problem is more complicated than the single-agent case for several reasons:

- The value function itself depends on the optimal strategy function $q^*(\dots)$, via the assumption that the rival firm is always using the optimal strategy.

So value iteration procedure is more complicated:

1. Start with initial guess $V^0(x_1, x_2)$
2. If q 's are continuous controls, we must solve for $q_1^0 \equiv q^0(x_1, x_2)$ and $q_2^0 \equiv q^0(x_2, x_1)$ to satisfy the system of first-order conditions (here subscripts denotes partial derivatives)

$$\begin{aligned} 0 &= \Pi_3(x_1, x_2, q_1^0, q_2^0) + \beta V_1^0(g(x_1, q_1^0), g(x_2, q_2^0)) \cdot g_2(x_1, q_1^0) \\ 0 &= \Pi_3(x_2, x_1, q_2^0, q_1^0) + \beta V_1^0(g(x_2, q_2^0), g(x_1, q_1^0)) \cdot g_2(x_2, q_2^0). \end{aligned} \quad (3)$$

For the discrete control case:

$$\begin{aligned} q_1^0 &= \operatorname{argmax}_{q \in \mathcal{Q}} \{ \Pi(x_1, x_2, q, q_2^0) + \beta V^0(g(x_1, q), g(x_2, q_2^0)) \} \\ q_2^0 &= \operatorname{argmax}_{q \in \mathcal{Q}} \{ \Pi(x_2, x_1, q, q_1^0) + \beta V^0(g(x_2, q), g(x_1, q_1^0)) \}. \end{aligned} \quad (4)$$

3. Update the next iteration of the value function:

$$V^1(x_1, x_2) = \{ \Pi(x_1, x_2, q_1^0, q_2^0) + \beta V^0(g(x_1, q_1^0), g(x_2, q_2^0)) \}. \quad (5)$$

Note: this and the previous step must be done at all points (x_1, x_2) in the discretized grid. As usual, use interpolation or approximation to obtain $V^1(\dots)$ at points not on the grid.

4. Stop when $\sup_{x_1, x_2} \|V^{i+1}(x_1, x_2) - V^i(x_1, x_2)\| \leq \epsilon$.
- In principle, one could estimate a dynamic games model, given time series of $\{x_t, q_t\}$ for both firms, by using a nested fixed-point estimation algorithm. In the outer loop, loop over different values of the parameters θ , and then in the inner loop, you compute the equilibrium of the dynamic game (in the way described above) for each value of θ .
 - However, there is an inherent ‘‘Curse of dimensionality’’ with dynamic games, because the dimensionality of the state vector (x_1, x_2) is equal to

the number of firms. (For instance, if you want to discretize 1000 pts in one dimension, you have to discretize at 1,000,000 pts to maintain the same fineness in two dimensions!)

Some papers provide computational methods to circumvent this problem (Keane and Wolpin (1994), Pakes and McGuire (2001), Imai, Jain, and Ching (2005)). Generally, these papers advocate only computing the value function at a (small) subset of the state points each iteration, and then approximating the value function at the rest of the state points using values calculated during previous iterations.

1.1 Dynamic games with “incomplete information”

- Clearly, it is possible to extend the Hotz-Miller insights to facilitate estimation of dynamic oligopoly models, in the case where q is a discrete control. Advantage, as before, is that you can avoid numerically solving for the value function.

Data directly tell you: the choice probabilities (distribution of $q_1, q_2 | x_1, x_2$); state transitions: (joint distribution of $x'_1 x'_2 | x_1, x_2, q_1, q_2$).

Papers which develop these ideas include: Bajari, Benkard, and Levin (2008), Pesendorfer and Schmidt-Dengler (2003), Aguirregabiria and Mira (2007), Pakes, Ostrovsky, and Berry (2007). Applications include: Collard-Wexler (2006), Ryan (2006), and Dunne, Klimer, Roberts, and Xu (2006). See survey in Pakes (2008), section 3.

- Here I outline the general setting, which can be considered a multi-agent version of Rust setup. (Nitty-gritty details are omitted: you should be able to figure it out!)
- First-order Markov
- Let \vec{X} and \vec{q} denote the N -vector of firms' state variables, and actions.
- Symmetric model: firms are identical, up to the current values of their state variables, and idiosyncratic utility shocks

- Utilities: firm i 's current utility, given her action d_i and her utility shocks ϵ_{i,d_i} is:

$$U(\vec{X}, d_i, \vec{d}_{-i}; \theta) + \epsilon_{i,d_i}.$$

- All utility shocks are *private information* of firm i . Also, they are i.i.d. across firms, and across actions (exactly like the logit errors in Rust's framework).

Because of these assumptions: each firm's period t shocks affect only his period t actions; they do not affect opponents' actions, nor do they affect actions in other periods $t' \neq t$.

In single-agent case, this distinction is nonsensical. But in multi-agent context, it is crucial. Consider several alternatives:

- All the shocks are publicly observed each period, but not observed by econometrician: essentially, we are in the “unobserved state variable” world of the last set of lecture notes, except that in multi-agent context, each firm's actions each period depend on *all* the shocks of all the firms in that period. Likelihood function gets very complicated.
 - Shocks are private information, but *serially correlated* over time: now firms can learn about their opponents' shocks through their actions, and their beliefs may evolve in complicated fashion. Also, firms may strategically try to influence their rivals' beliefs about their shocks (“signalling”). Complicated even from a *theoretical* point of view.
 - Shocks are unobserved by both firms and econometrician. In this case, they are non-structural errors and are useless from the point of view of generating more randomness in the model relative to the data.
- In symmetric equilibrium, firms' optimal policy function takes the form $d_i = d^*(\vec{X}, \epsilon_i)$. Corresponding choice probability: $P(d_i|\vec{X})$ (randomness in ϵ_i). These can be estimated from data.
 - State variables evolve via Markovian law-of-motion $\vec{X}'|\vec{X}, \vec{d}$. This can also be estimated from the data.

- Hence, choice-specific value functions for firm i (identical for all firms) can be forward-simulated:

$$\begin{aligned} \tilde{V}(\vec{X}, d_i = 1; \theta) = & E_{vecd_{-i}|\vec{X}} u(\vec{X}, d_i = 1, \vec{d}_{-i}; \theta) + \beta E_{\vec{X}'|\vec{X}, \vec{d}=(d_i, \vec{d}_{-i})} E_{\vec{d}'|\vec{X}'} E_{\epsilon'|d'_i, \vec{X}'} \left[u(\vec{X}', \vec{d}'; \theta) + \epsilon' \right. \\ & \left. + \beta E_{\vec{X}''|\vec{X}', \vec{d}'} E_{\vec{d}''|\vec{X}''} E_{\epsilon''|d''_i, \vec{X}''} \left[u(\vec{X}'', \vec{d}''; \theta) + \epsilon'' + \beta \dots \right] \right] \end{aligned} \quad (6)$$

When \vec{X} is finite, the matrix representation of value function can also be used here.

- Predicted choice probabilities:

$$\tilde{P}(d_i = 1|\vec{X}) = \frac{\exp(\tilde{V}(\vec{X}, d_i = 1; \theta))}{\exp(\tilde{V}(\vec{X}, d_i = 1; \theta)) + \exp(\tilde{V}(\vec{X}, d_i = 0; \theta))}.$$



2 An example of dynamic oligopoly: automobile market with secondary markets

We go over Esteban and Shum (2007).

In durable goods industries (like car market), secondary markets leads to intertemporal linkages between primary and secondary markets. Used goods of today were new goods of yesterday.

Interesting economic question: Does this harm or benefit producers?

- Intuition different from static markets:
 - Benchmark in DG setting is Coase outcome (firm's inability to commit to low levels of production can erode market power)
- Vs. this benchmark, 2-mkts can benefit producers:
 1. 2-mkt offers substitutes for firms' new production \Rightarrow curtails Coasian tendency to overproduce ("commitment benefit")

2. With heterogeneous consumers, 2-mkts segment market, allow firms to target new goods to high-valuation consumers (“*sorting benefit*”)



Economic Model: Car market

Multiproduct firms producing cars which differ in quality, durability and depreciation schedule.

Empirical model accommodates cost/demand shocks; for simplicity, describe deterministic model.

- Firms $j = 1, \dots, N$. (e.g. Ford, GM, Honda)
- L is total number of brands/models (e.g. Taurus, Accord, Escort).
- Firm j produces L_j models; set of products denoted \mathcal{L}_j .
- Model i lasts T_i periods. There are $K \equiv \sum_{i=1}^L T_i$ “model-years” actively traded during any given period.
- Each model year differs in one-dimensional quality \Rightarrow quality ladder

$$[\alpha_1, \alpha_2, \dots, \alpha_K, \alpha_{K+1} = 0]$$

where α_{K+1} is quality of outside option.

- Notation: depreciation schedules for different models
 - Define: $\eta(i)$ is ranking of model i when new.
 - Define: $v(\eta(i))$ is ranking of 1-yr old; $v^2(\eta(i)) \equiv v(v(\eta(i)))$ is ranking of 2-yr old, etc.
- \Rightarrow Depreciation schedule of model i described by sequence

$$\{\eta(i), v(\eta(i)), \dots, v^{T_i-1}(\eta(i))\}.$$

Note: each model has its own depreciation schedule.

- Note that firms are *asymmetric*. Hence equilibrium is characterized by set of L Bellman equations (as will be derived below).



Economic model: Dynamic demand

Derive from individual-level optimizing behavior.

- A continuum of infinitely-lived consumers who differ in their preference for quality θ (one dimension)
- Quasilinear per-period utility: $U_t = \theta\alpha_k + m - p_t^k$, where m is total income. Assume no liquidity constraints.
- Choice set: model-years $k = 1, \dots, K$, plus outside option (utility normalized =0)
- Consumers incur no **transactions costs**: abstract away from timing issues.

Implies simple form of dynamic decision rule: in period t , consumer θ chooses model-year k yielding maximal “rental utility”:

$$k_t = \operatorname{argmax}_k \left\{ 0, \alpha_k \theta - p_t^k + \delta E_t p_{t+1}^{v(k)}, k = 1, \dots, K \right\}$$

where δ is discount factor and expected rental price is $p_t^k - \delta E_t p_{t+1}^{v(k)}$.

(Drop E_t for convenience: assume perfect foresight, so $E_t p_{t+1} = p_{t+1}$.)



Demand functions

- Given prices $p_t^k, p_{t+1}^{v(k)}$ for all $k = 1, \dots, K$, each period there are K indifferent consumers $\bar{\theta} \geq \tilde{\theta}_t^1 \geq \tilde{\theta}_t^2 \geq \tilde{\theta}_t^3 \geq \dots \geq \tilde{\theta}_t^K \geq 0$, s.t.

- The indifferent consumers solve

$$\alpha_k \tilde{\theta}_t^k - p_t^k + \delta p_{t+1}^{v(k)} = \alpha_{k+1} \tilde{\theta}_t^k - p_t^{k+1} + \delta p_{t+1}^{v(k+1)}, \text{ for } k = 1, \dots, K-1$$

- Consumer heterogeneity: θ is uniformly distributed.
- Derive inverse demand function (subject to non-negativity constraints)

$$p_t^k = (\alpha_k - \alpha_{k+1}) \bar{\theta} \left(1 - \frac{1}{M} \sum_{r=1}^k x_t^r \right) + \delta p_{t+1}^{v(k)} + p_t^{k+1} - \delta p_{t+1}^{v(k+1)}$$



Supply side

- $\mathbf{y}_t = [1, x_t^1, \dots, x_t^K]'$: vector of all cars *transacted* in period t .
- $\mathbf{d}_t \equiv [x_t^{\eta(1)}, x_t^{\eta(2)}, \dots, x_t^{\eta(L)}]'$: vector of all cars *produced* in period t .
- Define matrices A and B, to get law of motion for \mathbf{y}_t :

$$\mathbf{y}_t = \mathbf{A}\mathbf{y}_{t-1} + \mathbf{B}\mathbf{d}_t.$$

- Marginal costs constant; no (dis-)economies of scope: $C_{jt} = \sum_{i \in \mathcal{L}} c_i \cdot x_t^{\eta(i)}$.
- Period t profits for car i is

$$\Pi_t^i(\mathbf{y}_t, \mathbf{y}_{t+1}, \dots, \mathbf{y}_{t+T_i-1}) = (p_t^i(\mathbf{y}_t, \mathbf{y}_{t+1}, \dots, \mathbf{y}_{t+T_i-1}) - c_t^i) \cdot q_t^1 :$$

depends on past, current, future prod'n of car i .

Important: dependence of current profits on future actions leads to a *time-consistency* problem, which is absent from “usual” dynamic problems. Very roughly, time-inconsistency implies that an agent’s optimal action in period t differs depending on whether the agent is deciding in period t , or period $t-1$, or period $t-2$, etc.

Think of durable goods monopoly: in period 1, his optimal period 2 price is the monopoly price (because that would raise his profits in period 1). But when period 2 comes, his optimal period 2 price is actually a lower price (since he wants to sell to people who did not buy in period 1).

- For individual firm: $\forall t, \forall j \in \mathcal{N}, \forall i \in \mathcal{L}_j$, period- t production $x_t^{\eta(i)}$ maximizes

$$(*) \quad \max_{x_t^{\eta(i)}, \forall i \in \mathcal{L}_j} \sum_{\tau=0}^{\infty} \sum_{i \in \mathcal{L}_j} \delta^\tau \underbrace{[\Pi_{t+\tau}^i(\mathbf{y}_{t+\tau}, \mathbf{y}_{t+\tau+1}, \dots, \mathbf{y}_{t+\tau+T_i-1})]}_{\text{period } t + \tau \text{ profits}}, \quad s.t.$$

$$\mathbf{y}_{t+\tau} = \mathbf{A}\mathbf{y}_{t+\tau-1} + \mathbf{B}\mathbf{d}_{t+\tau}, \quad \text{for } \tau = 1, \dots, \infty.$$

- Note: obj fcn different in each period t : usual problem is

$$\max_{\{x_t^{\eta(i)}, i \in \mathcal{L}_j\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \sum_{i \in \mathcal{L}_j} \delta^t \Pi_t^i(\mathbf{y}_t, \mathbf{y}_{t+1}, \dots, \mathbf{y}_{t+T_i-1}).$$

FOC for $x_t^{\eta(i)}$ contains derivative (say) $\frac{\partial \Pi_{t-1}^i}{\partial x_t^{\eta(i)}}$: in choosing period- t prodn, recognize that it affects period- $(t+1)$ profits \Rightarrow time-inconsistent.



Time-consistent Equilibrium production

- Restrict attention to Markov strategies: $\mathbf{A}\mathbf{y}_{t-1}$ is the “payoff-relevant state vector” for period t (stocks of cars produced prior to period t which are still actively traded in period t) \Rightarrow

Therefore consider production rules $x_t^{\eta(i)} = g_i(\mathbf{A}\mathbf{y}_{t-1}), \forall i \in \mathcal{L}_j, \forall j \in \mathcal{N}$.

- In Markov-perfect equilibrium, $g_1(\cdot), \dots, g_L(\cdot)$ satisfy Bellman equation

$$V_j(\mathbf{A}\mathbf{y}_{t-1}) = \max_{x_t^{\eta(i)}, i \in \mathcal{L}_j} \sum_{i \in \mathcal{L}_j} \Pi_j(\mathbf{y}_t, \mathbf{y}_{t+1}, \dots, \mathbf{y}_{t+T_i-1}) + \delta V_j(\mathbf{A}\mathbf{y}_t)$$

for all firms j , and the Markov decision rules

$$x_t^i = g_i(\mathbf{A}\mathbf{y}_{t-1}), \quad \text{for all } i \in \mathcal{L}_j.$$

Value function $V_j(\mathbf{A}\mathbf{y}_{t-1}) = (*)$ (optimal profits from t onwards).



Linear Quadratic (LQ) Specification

- We focus on *linear* equilibrium decision rule $x_{t+h} = \mathbf{G}\mathbf{A}\mathbf{y}_{t+h-1}$
No “trigger” strategies (step functions)
- \Rightarrow Quadratic value function $V(\mathbf{A}\mathbf{y}_{t-1}) = \mathbf{y}'_t \mathbf{A}' \mathbf{S} \mathbf{A} \mathbf{y}_t$.
- Bellman equation can be rewritten in matrix notation

$$\mathbf{y}'_{t-1} \mathbf{A}' \mathbf{S} \mathbf{A} \mathbf{y}_{t-1} = \max_{x_t^i, i \in \mathcal{J}} \left\{ \sum_{i \in \mathcal{J}} \left[\sum_{z=0}^{T_i-1} \mathbf{y}'_{t+z} \delta^z \mathbf{R}_{v^z(i)} \mathbf{y}_t \right] \right\} - \mathbf{y}'_t \mathbf{C}_j \mathbf{y}_t + \mathbf{y}'_t \delta [\mathbf{A}' \mathbf{S}_j \mathbf{A}] \mathbf{y}_t$$

- Recursive substitution yields

$$\begin{aligned} \mathbf{y}'_{t-1} \mathbf{A}' \mathbf{S} \mathbf{A} \mathbf{y}_{t-1} &= \max_{x_t^i, i \in \mathcal{J}} \mathbf{y}'_t \left\{ \left[\sum_{i \in \mathcal{J}} \sum_{z=0}^{T_i-1} (\mathbf{A}')^z [(I + \mathbf{B}\mathbf{G})']^z \delta^z \mathbf{R}_{v^z(i)} \right] - \mathbf{C}_j + \delta [\mathbf{A}' \mathbf{S}_j \mathbf{A}] \right\} \mathbf{y}_t \\ &\equiv \max_{x_t^{\eta(i)}, i \in \mathcal{J}} \mathbf{y}'_t \mathbf{Q}_j \mathbf{y}_t. \end{aligned}$$

■■■

Deriving equilibrium production rules

- Value iteration: solve for \mathbf{S} and \mathbf{G} by iterating over Bellman equation.
- For each \mathbf{S} , derive corresponding \mathbf{G} via FOC of right-hand side:

$$\mathbf{B}'_j (\mathbf{Q}_j + \mathbf{Q}'_j) \mathbf{y}_t = \mathbf{B}'_j (\mathbf{Q}_j + \mathbf{Q}'_j) \mathbf{A} \mathbf{y}_{t-1} + \mathbf{B}'_j (\mathbf{Q}_j + \mathbf{Q}'_j) \mathbf{B} \mathbf{d}_t = 0.$$

This system of FOC's corresponds to Eq. (3) in dynamic games handout.

Rearranging, we get:

$$\mathbf{d}_{t+h} = -(\mathbf{W}\mathbf{B})^{-1} (\mathbf{W}\mathbf{A}) \mathbf{y}_{t-1},$$

where $\mathbf{W}_j \equiv \mathbf{B}'_j (\mathbf{Q}_j + \mathbf{Q}'_j)$ for each firm j and $\mathbf{W} \equiv [\mathbf{W}_1, \dots, \mathbf{W}_N]'$.

Basis for estimating supply side of model.



Estimation

θ is not identified, set to constant.

To generate estimating equations, introduce shocks to firms' marginal costs:

$$C(x_t^{\eta(i)}) = x_t^{\eta(i)} (c_i + \epsilon_{it}).$$

Assumptions: $\epsilon_t \equiv [\epsilon_{1t}, \dots, \epsilon_{Lt}]'$ has zero-mean, *i.i.d.* across t . The vector ϵ_t is known to all firms when they make their period t choices (no asymmetric information).

From "certainty equivalence" properties of linear-quadratic model, optimal firm strategies are the same as in deterministic model, but with an additive shock:

$$\mathbf{d}_t = \mathbf{G}\mathbf{A}\mathbf{y}_{t-1} + \mathbf{w}_t, \text{ where } E(\mathbf{w}_t) = \mathbf{0}.$$

where $E(\mathbf{w}_t) = \mathbf{0}$ and independent of \mathbf{y}_{t-1} .

With the cost shocks, then, production (and also prices) will be random over time. However, due to independence of shocks across time, the innovations in prices will have mean zero at time t :

$$p_{t+1} = E_t p_{t+1} + \omega_{t+1}$$

where $E_t \omega_{t+1} = 0$.

This implies that the realized "demand function residual" will also have mean zero, conditional on formation known at time t :

$$\begin{aligned} 0 &= E \left[p_t^{\eta(i)} - (\alpha_{\eta(i)} - \alpha_{\eta(i)+1}) \left(1 - \frac{1}{M} \sum_{r=1}^{\eta(i)} x_t^r \right) - \delta p_{t+1}^{v(\eta(i))} - p_t^{\eta(i)+1} + \delta p_{t+1}^{v(\eta(i)+1)} \middle| \Omega_t \right] \\ &= E \left[(1 - \delta L^{-1}) (p_t^{\eta(i)} - p_t^{\eta(i)+1}) - (\alpha_{\eta(i)} - \alpha_{\eta(i)+1}) \left(1 - \frac{1}{M} \sum_{r=1}^{\eta(i)} x_t^r \right) \middle| \Omega_t \right] \end{aligned}$$

These form the basis for the moment conditions which we use for estimation, which we denote $\boldsymbol{\mu}_T(\psi)$.



GMM Estimation

- Source of identification: co-movements time series of prices and production.
-
- GMM estimator $\psi \equiv \operatorname{argmin}_{\psi} \mu_T(\psi)' \Omega_T^{-1} \mu_T(\psi)$.
- Nested GMM procedure: for each value of parameters ψ , solve LQ dynamic programming problem for coefficients $\mathbf{G}(\psi)$ of optimal production rules. LQ dynamic programming problem is very easy and quick to solve.
- Results

References

- AGUIRREGABIRIA, V., AND P. MIRA (2007): “Sequential Estimation of Dynamic Discrete Games,” *Econometrica*.
- BAJARI, P., L. BENKARD, AND J. LEVIN (2008): “Estimating Dynamic Models of Imperfect Competition,” *Econometrica*.
- COLLARD-WEXLER, A. (2006): “Demand Fluctuations and Plant Turnover in the Ready-to-Mix Concrete Industry,” manuscript, New York University.
- DUNNE, T., S. KLIMER, M. ROBERTS, AND D. XU (2006): “Entry and Exit in Geographic Markets,” manuscript, Penn State University.
- ESTEBAN, S., AND M. SHUM (2007): “Durable Goods Oligopoly with Secondary Markets: the Case of Automobiles,” Forthcoming, *RAND Journal of Economics*.
- IMAI, S., N. JAIN, AND A. CHING (2005): “Bayesian Estimation of Dynamic Discrete Choice Models,” Queen’s University, mimeo.
- KEANE, M., AND K. WOLPIN (1994): “The Solution and Estimation of Discrete Choice Dynamic Programming Models by Simulation and Interpolation: Monte Carlo Evidence,” *Review of Economics and Statistics*, 76, 648–672.
- PAKES, A. (2008): “Theory and Empirical Work in Imperfectly Competitive Markets,” Fisher-Schultz Lecture at 2005 Econometric Society World Congress (London, England).
- PAKES, A., AND P. MCGUIRE (2001): “Stochastic Algorithms, Symmetric Markov Perfect Equilibrium, and the ‘Curse’ of Dimensionality,” *Econometrica*, 69, 1261–1282.
- PAKES, A., M. OSTROVSKY, AND S. BERRY (2007): “Simple Estimators for the Parameters of Discrete Dynamic Games (with Entry/Exit Examples),” *RAND Journal of Economics*.
- PESENDORFER, M., AND P. SCHMIDT-DENGLER (2003): “Identification and Estimation of Dynamic Games,” NBER Working Paper, #9726.
- RYAN, S. (2006): “The Costs of Environmental Regulation in a Concentrated Industry,” manuscript, MIT.