Estimating Dynamic Games

Sanjog Misra
Anderson UCLA

Structural Workshop 2013

Estimating Dynamic Discrete Games

- In this workshop we have gone through a lot.
- The main topics include...
  1. Foundations: Structural Models (Reiss), Causality and Identification (Goldfarb), Instruments (Rossi), Data (Mela)
  2. Methods: Static Demand Models (Sudhir), Single Agent Dynamics: Theory and Econometrics (Hitsch), Static Games (Ellickson)
- What you should have gleaned from their talks is that the estimation of structural models requires

$$Data + Theory + Econometrics$$

- The estimation of dynamic games combines elements from all of the above talks and requires considerable expertise in handling data, game theory, econometrics and computational methods.

Estimating Dynamic Discrete Games

- We've already learned how to estimate single-agent (SA) dynamic discrete choice (DDC) models
- Two main approaches
  1. Full solution: NXFP (Rust, 1987)
  2. Two-Step: CCP (Hotz and Miller, 1993)
- In both cases, the underlying SA optimization problem involved agent's solving a dynamic programming (DP) problem
- With games, agent's must solve an inter-related system of DP problems
  - Their actions must be optimal given their beliefs & their beliefs must be correct on average (or at least self-confirming)
- As you can imagine, computing equilibria can be pretty complicated...

Estimating Dynamic Discrete Games

- Ericson & Pakes (ReStud, 95) and Pakes & McGuire (Rand, 94) provide a framework for computing equilibria to dynamic games.
    - More recently Goettler and Gordon (JPE 2012) provide an alternative approach
- However, using the original PM algorithm, solving a reasonably complex EP-style dynamic game even *once* is computationally demanding (if not impossible)
    - Estimating the model using NFXP is essentially intractable
- There's also the issue of multiple equilibria
    - The "incompleteness" this introduces can make it difficult to construct a proper likelihood/objective function, further complicating a NFXP approach
- Two-step "CCP" estimation provides a work-around that circumvents the iterative fixed point calculation, "solving" both problems

Estimating Dynamic Discrete Games

- CCP estimators were first developed by Hotz & Miller (HM, 1993) & Hotz, Miller, Sanders & Smith (HMS^2, 1994) for DDC models
- Four sets of authors (contemporaneously) suggested adapting these methods to games:
  1. Aguirregabiria and Mira (AM) (Ema, 2007),
  2. Bajari, Benkard, and Levin (BBL) (Ema, 2007),
  3. Pakes, Ostrovsky, and Berry (POB) (Rand, 2007), &
  4. Pesendorfer and Schmidt-Dengler (PSD) (ReStud, 2008),
- We'll talk about AM & BBL
  - AM extends HM to games
  - BBL extends HMS^2 to games
- Both are based on the framework suggested in Rust (94)
- More recent contributions are Blevins et al. (2012) and Arcidiacono and Miller (Ema, 2011)

Model

- A dynamic discrete game of incomplete information.
- Motivated by stylized model of retail chain competition.
- Let d_t be a vector of demand shifters in period t.
- N firms operate in the market, indexed by $i \in \{1, 2, ..., N\}$.
- In each period t, firms decide simultaneously how many outlets to operate - choosing from the discrete set $A = \{0, 1, ..., J\}$
  - The decision of firm i in period t is $a_{it} \in A$
  - The vector of all firms' actions is $a_t \equiv (a_{1t}, a_{2t}, ..., a_{Nt})$

Model

- Firms are characterized by two vectors of state variables that affect profitability: x_{it} and ε_{it}
  - $x_t \equiv (d_t, x_{1t}, x_{2t}, ..., x_{Nt})$ is common knowledge, but
  - $\varepsilon_t \equiv (\varepsilon_{1t}, \varepsilon_{2t}, ..., \varepsilon_{Nt})$ is privately observed by firm i
- $\widetilde{\Pi}_i (a_t, x_t, \varepsilon_{it})$ is firm i's per-period profit function.
- Assume $\{x_t, \varepsilon_t\}$ follows a controlled Markov process with transition probability $p(x_{t+1}, \varepsilon_{t+1} | a_t, x_t, \varepsilon_t)$, which is common knowledge

Model

- Each firm chooses its number of outlets to maximize expected discounted intertemporal profits,

$$E\left\{\sum_{s=t}^{\infty}\beta^{s-t}\widetilde{\Pi}_i\left(a_t, x_t, \varepsilon_{it}\right) \mid x_t, \varepsilon_{it}\right\}$$

  where $\beta \in (0, 1)$ is the (known) discount factor.
- The primitives of the model are the profit functions $\widetilde{\Pi}_i(\cdot)$, the transition probability $p(\cdot|\cdot)$, and β
- AM make the following set of assumptions...

Assumptions

ASSUMPTION 1—Additive Separability: *Private information appears additively in the profit function. That is,* $\tilde{\Pi}_i(a_t, x_t, \varepsilon_{it}) = \Pi_i(a_t, x_t) + \varepsilon_{it}(a_{it})$, *where* $\Pi_i(\cdot)$ *is a real-valued function and* $\varepsilon_{it}(a_{it})$ *is the* $(a_{it}+1)$*th component of the* $(J+1)\times 1$ *vector* ε_{it} *with support* R^{J+1}.

ASSUMPTION 2—Conditional Independence: *The transition probability* $p(\cdot|\cdot)$ *factors as* $p(x_{t+1}, \varepsilon_{t+1}|a_t, x_t, \varepsilon_t) = p_\varepsilon(\varepsilon_{t+1}) f(x_{t+1}|a_t, x_t)$. *That is, (i) given the firms' decisions at period* t, *private information variables do not affect the transition of common knowledge variables, and (ii) private information variables are independently and identically distributed over time.*

ASSUMPTION 3—Independent Private Values: *Private information is independently distributed across players,* $p_\varepsilon(\varepsilon_t) = \prod_{i=1}^N g_i(\varepsilon_{it})$, *where, for any player* i, $g_i(\cdot)$ *is a density function that is absolutely continuous with respect to the Lebesgue measure.*

ASSUMPTION 4—Discrete Common Knowledge Variables: *Common knowledge variables have a discrete and finite support* $x_t \in X \equiv \{x^1, x^2, \ldots, x^{|X|}\}$, *where* $|X|$ *is a finite number.*

Strategies and Bellman Equations

- AM also assume that firms play stationary Markov strategies
- Let $\sigma = \{\sigma_i(x, \varepsilon_i)\}$ be a set of strategy functions (decision rules), one for each firm, with

$$\sigma_i : X \times R^{J+1} \rightarrow A$$

- Associated with a set of strategy functions σ, define a set of *conditional choice probabilities* $P^\sigma = \{P_i^\sigma(a_i|x)\}$ such that

$$P_i^\sigma(a_i|x) \equiv \Pr(\sigma_i(x, \varepsilon_i) = a_i|x) = \int I\{\sigma_i(x, \varepsilon_i) = a_i\} g_i(\varepsilon_i) d\varepsilon_i$$

which represent the expected behavior of firm i from the point of view of the rest of the firms (& us!), when firm i follows σ.

Expected profits

- Let $\pi_i^\sigma(a_i, x)$ be firm i's current expected profit from choosing alternative a_i while the other firms follow σ.
- By the independence of private information,

$$\pi_i^\sigma(a_i, x) = \sum_{a_{-i} \in A^{N-1}} \left(\prod_{j \neq i} P_j^\sigma(a_{-i}[\,j] \,|x) \right) \Pi_i(a_i, a_{-i}, x)$$

  where $a_{-i}[\,j]$ is the j^{th} firm's element in the vector of actions players other than i
- Let $\widetilde{V}_i^\sigma(x, \varepsilon_i)$ be the value of firm i if it behaves optimally now and in the future given that the other firms follow σ.

Value Functions

- By Bellman's principle of optimality, we can write

$$\widetilde{V}_i^{\sigma}(x, \varepsilon_i) = \max_{a_i \in A} \left\{ \pi_i^{\sigma}(a_i, x) + \varepsilon_i(a_i) + \beta \sum_{x' \in X} \left[\int \widetilde{V}_i^{\sigma}(x', \varepsilon_i') \, g_i(\varepsilon_i') \, d\varepsilon_i' \right] f_i^{\sigma}(x' | x, a_i) \right\} \tag{1}$$

where $f_i^{\sigma}(x'|x, a_i)$ is the transition probability of x conditional on firm i choosing a_i and the other firms following σ:

$$f_i^{\sigma}(x'|x, a_i) = \sum_{a_{-i} \in A^{N-1}} \left(\prod_{j \neq i} P_j^{\sigma}(a_{-i}[j] | x) \right) f(x'|x, a_i, a_{-i})$$

- AM prefer to work with the value functions integrated over the private information variables.
- Let $V_i^{\sigma}(x)$ be the integrated value function $\int V_i^{\sigma}(x, \varepsilon_i) \, g_i(d\varepsilon_i)$

Choice-Specific Value Functions

Based on this definition

$$V_i^{\sigma}(x) = \int V_i^{\sigma}(x, \varepsilon_i)\, g_i(d\varepsilon_i)$$

and the Bellman equation from above

$$\widetilde{V}_i^{\sigma}(x, \varepsilon_i) = \max_{a_i \in A} \left\{ \pi_i^{\sigma}(a_i, x) + \varepsilon_i(a_i) + \beta \sum_{x' \in X} \left[\int \widetilde{V}_i^{\sigma}(x', \varepsilon_i')\, g_i(\varepsilon_i')\, d\varepsilon_i' \right] f_i^{\sigma}(x'|x, a_i) \right\} \tag{1}$$

we can obtain the *integrated Bellman equation*

$$V_i^{\sigma}(x) = \int \max_{a_i \in A} \left\{ v_i^{\sigma}(a_i, x) + \varepsilon_i(a_i) \right\} g_i(d\varepsilon_i)$$

where

$$v_i^{\sigma}(a_i, x) \equiv \pi_i^{\sigma}(a_i, x) + \beta \sum_{x' \in X} V_i^{\sigma}(x')\, f_i^{\sigma}(x'|x, a_i)$$

which are often referred to as choice-specific value functions.

Markov Perfect Equilibria

- Now it's time to enforce the fact that σ describes equilibrium behavior (i.e. it's a best response)
- DEFINITION: A stationary Markov perfect equilibrium (MPE) in this game is a set of strategy functions σ^* such that for any firm i and any $(x, \varepsilon_i) \in X \times R^{J+1}$

$$\sigma_i^* \left(x, \varepsilon_i\right) = \arg\max_{a_i \in A} \left\{ v_i^{\sigma^*} \left(a_i, x\right) + \varepsilon_i \left(a_i\right) \right\}$$

- Ultimately, an equation (somewhat) like this will be the basis of estimation

Markov Perfect Equilibria

- Note that π_i^σ, V_i^σ, & f_i^σ only depend on σ through P, so AM switch notation to π_i^P, V_i^P, & f_i^P so they can represent the MPE in probability space.
- Let σ^* be a set of MPE strategies and let P^* be the probabilities associated with these strategies

$$P^*\left(a_i|x\right) = \int I\left\{a_i = \sigma_i^*\left(x, \varepsilon_i\right)\right\} g_i\left(\varepsilon_i\right) d\varepsilon_i$$

- Equilibrium probabilities are a fixed point ($P^* = \Lambda\left(P^*\right)$), where, for any vector of probabilities $P, \Lambda\left(P\right) = \left\{\Lambda_i\left(a_i|x; P_{-i}\right)\right\}$ and

$$\Lambda_i\left(a_i|x; P_{-i}\right) = \int I\left(a_i = \arg\max_{a_i \in A}\left\{v_i^{P^*}\left(a, x\right) + \varepsilon_i\left(a\right)\right\}\right) g_i\left(\varepsilon_i\right) d\varepsilon_i$$

- The functions Λ_i are *best response probability functions*

Markov Perfect Equilibria

- The equilibrium probabilities solve the coupled fixed-point problems defined by

$$V_i^{\sigma}(x) = \int \max_{a_i \in A} \left\{ v_i^{\sigma}(a_i, x) + \varepsilon_i(a_i) \right\} g_i(d\varepsilon_i) \tag{2}$$

and

$$\Lambda_i(a_i | x; P_{-i}) = \int I\left(a_i = \arg\max_{a_i \in A} \left\{ v_i^{P^*}(a, x) + \varepsilon_i(a) \right\} \right) g_i(\varepsilon_i) \, d\varepsilon_i \tag{3}$$

- Given a set of probabilities P:
  - The value functions V_i^P are solutions of the N Bellman equations in (2), and
  - Given these value functions, the best response probabilities are defined by the RHS of (3).

An Alternative Best Response Mapping

- It's easier to work with an alternative best response mapping (in probability space) that avoids the solution of the N dynamic programming problems in (2).
- The evaluation of this mapping is computationally much simpler than the evaluation of the mapping $\Lambda(P)$, and it will prove more convenient for the estimation of the model.
- Let P^* be an equilibrium and let $V_1^{P^*}, V_2^{P^*}, ..., V_N^{P^*}$ be firms' value functions associated with this equilibrium.

Alternative Best Response Mapping

- Because equilibrium probabilities are best responses, we can rewrite the Bellman equation

$$V_i^{\sigma}(x) = \int \max_{a_i \in A} \left\{ v_i^{\sigma}(a_i, x) + \varepsilon_i(a_i) \right\} g_i(d\varepsilon_i) \tag{2}$$

as

$$V_i^{P^*}(x) = \sum_{a_i \in A} P_i^*(a_i|x) \left[\pi_i^{P^*}(a_i, x) + e_i^{P^*}(a_i, x) \right] + \beta \sum_{x' \in X} V_i^{P^*}(x') f^{P^*}(x'|x) \tag{4}$$

where $f^{P^*}(x'|x)$ is the transition probability of x induced by P^*

- $e_i^{P^*}(a_i, x)$ is the expectation of $\varepsilon_i(a_i)$ conditional on x and on alternative a_i being optimal for player i

Alternative Best Response Mapping

- Note that $e_i^{P^*}(a_i, x)$ is a function of g_i and $P_i^*(x)$ only
- The functional form depends on the probability distribution g_i
- For example, if the $\varepsilon_i(a_i)$ are *iid* T1EV, then

  $$e_i^{P^*}(a_i, x) \equiv E(\varepsilon_i(a)|x, \sigma_i^*(x, \varepsilon_i) = a_i) = \gamma - \sigma \ln(P_i(a_i|x))$$

  where γ is Euler's constant & σ is the logit dispersion parameter.
- Taking equilibrium probabilities as given, expression (4)

  $$V_i^{P^*}(x) = \sum_{a_i \in A} P_i^*(a_i|x)\left[\pi_i^{P^*}(a_i, x) + e_i^{P^*}(a_i, x)\right] + \beta \sum_{x' \in X} V_i^{P^*}(x') f^{P^*}(x'|x)$$

  describes the vector of values $V_i^{P^*}$ as the solution of a system of linear equations, which can be written in vector form as

  $$\left(I - \beta F^{P^*}\right) V_i^{P^*} = \sum_{a_i \in A} P_i^*(a_i) * \left[\pi_i^{P^*}(a_i) + e_i^{P^*}(a_i)\right]$$

Alternative Best Response Mapping

- Let $\Gamma_i(P^*) \equiv \{\Gamma_i(x; P^*) : x \in X\}$ be the solution to this system of linear equations, such that $V_i^{P^*}(x) = \Gamma_i(x; P^*)$.
- For arbitrary probabilities P, the mapping

$$\Gamma_i(P) = \left(I - \beta F^{P^*}\right)^{-1} \left\{ \sum_{a_i \in A} P_i^*(a_i) * \left[\pi_i^{P^*}(a_i) + e_i^{P^*}(a_i)\right] \right\} \tag{5}$$

  can be interpreted as a valuation operator.
- An MPE is then a fixed point $\Psi(P) \equiv \{\Psi_i(a_i|x; P)\}$ where

$$\Psi_i(a_i|x; P) = \int I\left(a_i = \arg\max_{a_i \in A} \left\{ \pi_i^P(a, x) + \varepsilon_i(a) + \beta \sum_{x' \in X} \Gamma_i(x'; P) f_i^P(x'|x, a) \right\}\right) g_i(\varepsilon_i)\, d\varepsilon_i \tag{6}$$

- By definition, an equilibrium vector P^* is a fixed point of Ψ.
- AM's *Representation Lemma* establishes that the reverse is also true.
- Equation (6) will be the basis of estimation

Estimation

- Assume the researcher observes M geographically separate markets over T periods, where M is large and T is small.
- The primitives $\{\Pi_i, g_i, f, \beta : i \in I\}$ are known to the researcher up to a finite vector of structural parameters $\theta \in \Theta \subset R^{|\Theta|}$.
  - β is assumed known (it's *very* difficult to estimate)
- We now incorporate θ as an explicit argument in the equilibrium mapping Ψ.
- Let $\theta_0 \in \Theta$ be the true value of θ in the population
- Under Assumption 2 (i.e., conditional independence), the transition probability function f can be estimated from transition data using a standard maximum likelihood method and without solving the model.

Assumptions on DGP

Assumption 5: *Let* $P^0_{mt} \equiv \{\Pr(a_{mt} = a | x_{mt} = x) : (a, x) \in A^N \times X\}$ *be the distribution of* a_{mt} *conditional on* x_{mt} *in market* m *at period* t. *(A) For every observation* (m, t) *in the sample,* $P^0_{mt} = P^0$ *and* $P^0 = \Psi(\theta^0, P^0)$. *(B) Players expect* P^0 *to be played in future (out of sample) periods. (C) For any* $\theta \neq \theta^0$ *and* P *that solves* $P = \Psi(\theta, P)$, *it is the case that* $P \neq P^0$. *(D) The observations* $\{a_{mt}, x_{mt} : m = 1, 2, \ldots, M; t = 1, 2, \ldots, T\}$ *are independent across markets and* $\Pr(x_{mt} = x) > 0$ *for all* x *in* X.

Maximum Likelihood Estimation

- Define the pseudo likelihood function

$$Q_M(\theta, P) = \frac{1}{M} \sum_{m=1}^{M} \sum_{t=1}^{T} \sum_{i=1}^{N} \ln \Psi_i \left(a_{imt} | x_{mt}; P, \theta\right)$$

where P is an arbitrary vector of players' choice probabilities.

- Consider first the hypothetical case of a model with a unique equilibrium for each possible value of $\theta \in \Theta$.
- Then the maximum likelihood estimator (MLE) of θ_0 can be defined from the constrained multinomial likelihood

$$\hat{\theta}_{MLE} = \arg\max_{\theta \in \Theta} Q_M(\theta, P) \text{ subject to } P = \Psi(\theta, P)$$

MLE

- However, with multiple equilibria the restriction $P = \Psi(\theta, P)$ does not define a unique vector P, but a set of vectors.
- In this case, the MLE can be defined as

$$\hat{\theta}_{MLE} = \arg\max_{\theta \in \Theta} \left\{ \sup_{P \in (0,1)^{N|X|}} Q_M(\theta, P) \text{ subject to } P = \Psi(\theta, P) \right\}$$

- Thus, for each candidate θ, we need to compute all the vectors P that constitute equilibria (given θ) and select the one with the highest value of $Q_M(\theta, P)$.
- At present, there are no known methods that can do so (robustly).
- Therefore, AM introduce a class of pseudo maximum likelihood estimators.

Pseudo Maximum Likelihood (PML) Estimation

- The PML estimators try to minimize the number of evaluations of Ψ for different vectors of players' probabilities P.
- Suppose that we know the population probabilities P^0 and consider the (infeasible) PML estimator

$$\hat{\theta} = \arg\max_{\theta \in \Theta} Q_M\left(\theta, P^0\right)$$

- Under standard regularity conditions, this estimator is $\sqrt{M}$-CAN, but *infeasible* since P^0 is unknown.
- However, if we can obtain a $\sqrt{M}$-consistent nonparametric estimator of P^0, then we can define the feasible two-step PML estimator

$$\hat{\theta}_{2S} = \arg\max_{\theta \in \Theta} Q_M\left(\theta, \hat{P}^0\right)$$

PML

- $\sqrt{M}$-CAN of $\hat{P}^0$ ($+$ regularity conditions) are sufficient to guarantee that the PML estimator is $\sqrt{M}$-CAN.
- There are two good reasons to care about this estimator
  1. It deals with the indeterminacy problem associated with multiple equilibria (it's robust)
  2. Furthermore, repeated solutions of the dynamic game are avoided, which can result in significant computational gains (it's simple)

PML: Drawbacks

- However, the two-step PML has some drawbacks
  1. Its asymptotic variance depends on the variance Σ of the nonparametric estimator $\hat{P}^0$.
     - Therefore, it can be very inefficient when Σ is large.
  2. For the sample sizes available in actual applications, the nonparametric estimator of P^0 can be very imprecise (small sample bias).
     - There's a curse of dimensionality here...
  3. For some models, it's impossible to obtain consistent nonparametric estimates of P^0.
     - e.g. models with unobserved market characteristics.
- To address these issues, AM introduce the Nested Pseudo Likelihood estimator.

Nested Pseudo Likelihood (NPL) Method

- NPL is a recursive extension of the two-step PML estimator.
- Let $\hat{P}^0$ be a (possibly inconsistent) initial guess of the vector of players' choice probabilities.
- Given $\hat{P}^0$, the NPL algorithm generates a sequence of estimators $\left\{\hat{\theta}_K : K \geq 1\right\}$, where the K-stage estimator is defined as

$$\hat{\theta}_K = \arg\max_{\theta \in \Theta} Q_M\left(\theta, \hat{P}_{K-1}\right)$$

  and the probabilities $\left\{\hat{P}_K : K \geq 1\right\}$ are obtained recursively as

$$\hat{P}_K = \Psi\left(\hat{\theta}_K, \hat{P}_{K-1}\right)$$

NPL

- If the initial guess $\hat{P}^0$ is a consistent estimator, all elements of the sequence of estimators $\left\{\hat{\theta}_K : K \geq 1\right\}$ are consistent
- However, AM are interested in the properties of the estimator in the limit (if it converges)
- If the sequence $\left\{\hat{\theta}_K, \hat{P}_K\right\}$ converges, its limit $\left(\hat{\theta}, \hat{P}\right)$ is such that

$$\hat{\theta} \text{ maximizes } Q_M\left(\theta, \hat{P}\right) \text{ and } \hat{P} = \Psi\left(\hat{\theta}, \hat{P}\right)$$

  and any pair that does so is a NPL fixed point.
- AM show that a NPL fixed point always exists and, if there is more than one, the one with the highest value of the pseudo likelihood is a consistent estimator
- Of course, it may be *very* difficult to find multiple roots

NPL

- NPL preserves the two main advantages of PML
  1. It's feasible in models with multiple equilibria, and
  2. It minimizes the number of evaluations of the mapping Ψ for different values of P
- Furthermore
  1. It's more efficient than either infeasible *or* two-step PML (because it imposes the MPE condition in sample)
  2. It reduces the finite sample bias generated by imprecise estimates of P^0
  3. It doesn't require initially consistent $\hat{P}^0$'s (so it can accommodate unobserved heterogeneity)
- AM show how to extend NPL to settings with permanent unobserved heterogeneity, but let's skip that & look at some Monte Carlos

Monte Carlos

- Consider a simple entry/exit example where 5 firms can operate at most 1 store, so $a_{it} \in \{0, 1\}$
- Variable profit is given by

$$\theta_{RS} \ln (S_{mt}) - \theta_{RN} \ln \left(1 + \sum_{j \neq i} a_{jmt}\right)$$

  where S_{mt} is the size of market m in period t, and θ_{RS} & θ_{RN} are parameters to be estimated
- The profit function of an active firm is

$$\tilde{\Pi}_{imt}(1) = \theta_{RS} \ln (S_{mt}) - \theta_{RN} \ln \left(1 + \sum_{j \neq i} a_{jmt}\right) - \theta_{FC,i} - \theta_{EC}(1 - a_{im,t-1}) + \varepsilon$$

  where $\theta_{FC,i}$ is fixed cost, θ_{EC} is entry cost, and $\varepsilon_{imt} \sim$ T1EV
- The profit function of an inactive firm is simply

$$\tilde{\Pi}_{imt}(0) = \varepsilon_{imt}$$

Monte Carlos

- $\ln(S_{mt})$ follows a discrete first order Markov process, with known transition matrix and finite support $\{1, 2, 3, 4, 5\}$
- Fixed operating costs are

  $$\theta_{FC,1} = -1.9, \theta_{FC,2} = -1.8, \theta_{FC,3} = -1.7, \theta_{FC,4} = -1.6, \theta_{FC,5} = -1.5$$

  so that firm 5 is most efficient and firm 1 is least efficient.
- $\theta_{RS} = 1$ and $\beta = .95$ throughout, but AM will vary θ_{RN} and θ_{EC}

MONTE CARLO EXPERIMENTS: PARAMETERS θ_{EC} AND θ_{RN}

Parameter[a]	Exp. 1 $\theta_{EC} = 1.0$ $\theta_{RN} = 0.0$	Exp. 2 $\theta_{EC} = 1.0$ $\theta_{RN} = 1.0$	Exp. 3 $\theta_{EC} = 1.0$ $\theta_{RN} = 2.0$	Exp. 4 $\theta_{EC} = 0.0$ $\theta_{RN} = 1.0$	Exp. 5 $\theta_{EC} = 2.0$ $\theta_{RN} = 1.0$	Exp. 6 $\theta_{EC} = 4.0$ $\theta_{RN} = 1.0$
$\frac{\theta_{EC}}{\theta_{RS}\ln(3)}$	0.91	0.91	0.91	0.00	1.82	3.64
$100\frac{\theta_{RN}\ln(2)}{\theta_{RS}\ln(3)}$	0.0%	63.1%	126.2%	63.1%	63.1%	63.1%

[a]The parameter $\theta_{EC}/(\theta_{RS}\ln(3))$ represents the ratio between entry costs and the annual variable profit of a monopolist in a market of average size, i.e., $S_{mt} = 3$. The parameter $100(\theta_{RN}\ln(2))/(\theta_{RS}\ln(3))$ represents the percentage reduction in annual variable profits when we go from a monopoly to a duopoly in an average size market, i.e., $S_{mt} = 3$.

Monte Carlos

- The space of common knowledge state variables (S_{mt}, a_{t-1}) has $2^5 \times 5 = 160$ cells.
- There's a different vector of CCPs for each firm, so the dimension of the CCP vector for all firms is $5 \times 160 = 800$.
- For each experiment, they compute a MPE.
- The equilibrium is obtained by iterating the best response probability mapping starting with a 800×1 vector of choice probabilities (guesses) - (e.g. $P_i\left(a_i = 1|x\right) = .5\ \forall x, i$)
- They can then calculate the steady state distribution and generate fake data.
- Table II presents some descriptive statistics associated with the MPE of each experiment.

Markov Perfect Equilibria for each experiment

TABLE II

MONTE CARLO EXPERIMENTS: DESCRIPTION OF THE MARKOV PERFECT EQUILIBRIUM IN THE DGP[a]

Descriptive Statistics	Exp. 1 $\theta_{EC}=1.0$ $\theta_{RN}=0.0$	Exp. 2 $\theta_{EC}=1.0$ $\theta_{RN}=1.0$	Exp. 3 $\theta_{EC}=1.0$ $\theta_{RN}=2.0$	Exp. 4 $\theta_{EC}=0.0$ $\theta_{RN}=1.0$	Exp. 5 $\theta_{EC}=2.0$ $\theta_{RN}=1.0$	Exp. 6 $\theta_{EC}=4.0$ $\theta_{RN}=1.0$
Number of active firms						
Mean	3.676	2.760	1.979	2.729	2.790	2.801
Std. dev.	1.551	1.661	1.426	1.515	1.777	1.905
AR(1) for number of active firms (autoregressive parameter)	0.744	0.709	0.571	0.529	0.818	0.924
Number of entrants (or exits) per period	0.520	0.702	0.748	0.991	0.463	0.206
Excess turnover[b] (in # of firms per period)	0.326	0.470	0.516	0.868	0.211	0.029
Correlation between entries and exits	−0.015	−0.169	−0.220	−0.225	−0.140	−0.110
Prob. of being active						
Firm 1	0.699	0.496	0.319	0.508	0.487	0.455
Firm 2	0.718	0.527	0.356	0.523	0.521	0.501
Firm 3	0.735	0.548	0.397	0.547	0.556	0.550
Firm 4	0.753	0.581	0.434	0.564	0.592	0.610
Firm 5	0.770	0.607	0.475	0.586	0.632	0.686

[a]For all these experiments, the values of the rest of the parameters are $N=5$, $\theta_{FC,1}=-1.9$, $\theta_{FC,2}=-1.8$, $\theta_{FC,3}=-1.7$, $\theta_{FC,4}=-1.6$, $\theta_{FC,5}=-1.5$, $\theta_{RS}=1.0$, $\sigma_\varepsilon=1$, and $\beta=0.95$.

[b]Excess turnover is defined as (#entrants + #exits) – abs(#Entrants – #Exits).

Monte Carlos

- They then calculate the two-step PML and NPL estimators using the following choices for the initial vector of probabilities
  1. The true vector of equilibrium probabilities P^0
  2. Nonparametric frequency estimates
  3. Logits (for each firm) with the log of market size and indicators of incumbency status for all firms as explanatory variables, and
  4. Independent random draws from a $U(0, 1)$ r.v.
- Tables IV and V summarize the results

Monte Carlo Results

TABLE IV

MONTE CARLO EXPERIMENTS: EMPIRICAL MEANS AND EMPIRICAL STANDARD DEVIATIONS OF ESTIMATORS

Experiment	Estimator	Parameters			
		$\theta_{FC,1}$	θ_{RS}	θ_{EC}	θ_{RN}
1	True values	**−1.900**	**1.000**	**1.000**	**0.000**
	2S-True	−1.915 (0.273)	1.007 (0.152)	1.002 (0.139)	0.002 (0.422)
	2S-Freq	−0.458 (0.289)	0.374 (0.141)	1.135 (0.190)	0.200 (0.364)
	2S-Logit	−1.929 (0.279)	1.006 (0.153)	0.997 (0.138)	−0.009 (0.431)
	NPL	−1.902 (0.279)	1.018 (0.157)	0.994 (0.139)	0.036 (0.439)
2	True values	**−1.900**	**1.000**	**1.000**	**1.000**
	2S-True	−1.894 (0.212)	1.002 (0.186)	1.007 (0.118)	1.007 (0.583)
	2S-Freq	−0.919 (0.208)	0.351 (0.119)	0.886 (0.123)	0.095 (0.337)
	2S-Logit	−1.920 (0.226)	0.977 (0.197)	1.000 (0.122)	0.915 (0.597)
	NPL	−1.893 (0.232)	1.016 (0.220)	0.998 (0.121)	1.050 (0.681)
3	True values	**−1.900**	**1.000**	**1.000**	**2.000**
	2S-True	−1.910 (0.183)	1.006 (0.209)	1.000 (0.112)	2.008 (0.783)
	2S-Freq	−1.126 (0.189)	0.286 (0.094)	0.792 (0.107)	0.027 (0.311)
	2S-Logit	−1.919 (0.248)	1.022 (0.305)	0.985 (0.145)	2.070 (1.110)
	NPL	−1.920 (0.232)	0.950 (0.189)	1.007 (0.116)	1.792 (0.667)
4	True values	**−1.900**	**1.000**	**0.000**	**1.000**
	2S-True	−1.890 (0.516)	1.020 (0.329)	0.001 (0.119)	1.063 (1.345)
	2S-Freq	−0.910 (0.243)	0.337 (0.104)	0.239 (0.113)	0.127 (0.354)
	2S-Logit	−2.070 (0.436)	0.903 (0.262)	0.000 (0.119)	0.571 (1.061)
	NPL	−1.891 (0.482)	1.014 (0.291)	0.001 (0.115)	1.047 (1.186)
5	True values	**−1.900**	**1.000**	**2.000**	**1.000**
	2S-True	−1.912 (0.178)	1.007 (0.142)	2.008 (0.132)	1.006 (0.359)
	2S-Freq	−0.840 (0.218)	1.379 (0.130)	1.591 (0.143)	0.181 (0.302)
	2S-Logit	−1.921 (0.204)	0.997 (0.167)	2.002 (0.138)	0.971 (0.405)
	NPL	−1.924 (0.203)	1.018 (0.178)	2.000 (0.137)	1.027 (0.435)
6	True values	**−1.900**	**1.000**	**4.000**	**1.000**
	2S-True	−1.899 (0.206)	1.003 (0.132)	4.050 (0.203)	1.006 (0.238)
	2S-Freq	−0.558 (0.228)	0.332 (0.128)	2.745 (0.211)	0.206 (0.238)
	2S-Logit	−1.895 (0.240)	0.996 (0.147)	4.048 (0.208)	0.992 (0.277)
	NPL	−1.918 (0.239)	1.009 (0.152)	4.044 (0.207)	1.009 (0.285)

Monte Carlo Results

TABLE V

SQUARE-ROOT MEAN SQUARE ERROR RELATIVE TO THE ONE-STAGE PML WITH TRUE P^0

Experiment	Estimator	Parameters $\theta_{FC,1}$	θ_{RS}	θ_{EC}	θ_{RN}
1	2S-Freq	5.380	4.222	1.676	0.983
	2S-Logit	1.027	1.006	1.002	1.022
	NPL	1.019	1.040	0.996	1.044
2	2S-Freq	4.736	3.553	1.415	1.655
	2S-Logit	1.070	1.066	1.029	1.034
	NPL	1.098	1.188	1.020	1.171
3	2S-Freq	4.347	3.440	2.095	2.549
	2S-Logit	1.357	1.462	1.301	1.419
	NPL	1.268	0.935	1.038	0.892
4	2S-Freq	1.977	2.035	2.228	0.699
	2S-Logit	0.906	0.848	1.000	0.850
	NPL	0.935	0.884	0.969	0.881
5	2S-Freq	6.054	4.459	3.279	2.429
	2S-Logit	1.146	1.176	1.043	1.130
	NPL	1.143	1.250	1.037	1.210
6	2S-Freq	6.591	5.589	6.072	3.487
	2S-Logit	1.162	1.209	1.020	1.166
	NPL	1.158	1.248	1.010	1.197

Discussion of Results

- The NPL algorithm always converged to the same estimates (regardless of the value of $\hat{P}_0$)
- The algorithm converged faster when initialized with logit estimates
- The 2S-Freq estimator is highly biased in all experiments, although its variance is sometimes smaller than the NPL and 2S-True estimators.
  - It's main drawback is small sample bias...
- The NPL estimator performs very well relative to the 2S-True both in terms of variance and bias.
- In all the experiments, the most important gains associated with the NPL estimator occur for the entry cost parameter.

Motivation for BBL

- A big drawback of the AM approach (and PSD & POB as well) is that it's designed for discrete controls (and discrete states)
  - A big selling point of BBL is that it can
- However,
  1. It's not completely clear that AM can't be extended to continuous controls and states (see, e.g., Arcidiacono and Miller (2011))
  2. The way BBL "handles" continuous controls is probably infeasible (unless there are no structural shocks to investment)
  3. BBL's objective function can be difficult to optimize, so you might want to mix & match a bit.
- Let's look now at how BBL works...

Background for BBL (HMSS)

- The original idea for the methods proposed in BBL came from Hotz, Miller, Sanders and Smith (1994)
- They proposed a two-step (really 3) approach
  1. Estimate the transition kernels $f\left(s'|s,a\right)$ and CCPs $P\left(a|s\right)$ from the data
  2. Approximate $V\left(s|\hat{P}\right)$ using Monte-Carlo approaches
  3. Construct and maximize an objective function to obtain parameters
- Let's look how this might work...

Forward Simulation (HMSS)

- For each state $(s \in S)$ draw a sequence of R future paths as follows
  1. Draw iid shocks $\varepsilon_0(a)$
  2. Compute optimal policies $\hat{P}(a|s_0)$ and pick action $a_0(\varepsilon_0, s_0)$
  3. Compute payoffs $U_0^r = u(s_0, a_0; \theta) + \varepsilon_0(a_0)$
  4. Draw next period state s_1˜$\hat{f}(s'|s_0, a_0)$ and repeat for T iterations.
- We can use these to construct Value functions...

Value Function (HMSS)

- At some T we stop the forward simulation and use the fact that

$$V^r\left(s_0\right) = \sum_{t=0}^{T} \beta^t U_t^r$$

- To construct

$$\hat{V}\left(s_0|\hat{P}\right) = \frac{1}{R} \sum_{r=1}^{R} V^r\left(s_0\right)$$

- We can then obtain an estimator using
- A Pseudo Likelihood, GMM or Least Squares

$$\max_{\theta} \sum_t \sum_i a_{it} \ln \Psi\left(a_{it}|s_{it}, \hat{V}\right)$$

$$\min_{\theta} \sum_i \left(\sum_t \left(a_{it} - \Psi\left(a_{it}|s_{it}, \hat{V}\right) Z_{it}\right)\right) \Omega^{-1} \left(\sum_t \left(a_{it} - \Psi\left(a_{it}|s_{it}, \hat{V}\right)\right.\right.$$

$$\min_{\theta} \left(\hat{P}\left(a|s\right) - \Psi\left(a|s, \hat{V}\right)\right)' W^{-1} \left(\hat{P}\left(a|s\right) - \Psi\left(a|s, \hat{V}\right)\right)$$

- So what does BBL do?
  - Extends HMSS to games
  - Generalizes the idea to continuous actions
  - Proposes an inequality conditions for estimation (Bounds estimator)
- The key element of BBL is that it (like AM 2007) allows the research to be agnostic about equilibrium selection
- and side-step the multiple equilibria problem.

Notation

- The game is in discrete time with an infinite horizon
- There are N firms, denoted $i = 1, ..., N$ making decisions at times $t = 1, 2, ..., \infty$
- Conditions at time t are summarized by discrete states $\mathbf{s}_t \in S \subset \mathcal{R}^L$
- Given $\mathbf{s}_t$, firms choose actions simultaneously
- Let $a_{it} \in A_i$ denote firm i's action at time t, and $\mathbf{a}_t = (a_{1t}, ..., a_{Nt})$ the vector of time t actions

Notation

- Before choosing its action, each firm i receives a private shock ν_{it} drawn *iid* from $G_i\left(\cdot|\mathbf{s}_t\right)$ with support $\mathcal{V}_i \subset \mathcal{R}^M$
- Denote the vector of private shocks $\boldsymbol{\nu}_t = \left(\nu_{1t}, ..., \nu_{Nt}\right)$
- Firm i's profits are given by $\pi_i(\mathbf{a}_t, \mathbf{s}_t, \nu_{it})$ and firms share a common (& known) discount factor $\beta < 1$
- Given $\mathbf{s}_t$, firm i's expected profit (prior to seeing ν_{it}) is

$$E\left[\sum_{\tau=t}^{\infty} \beta^{\tau-t} \pi_i(\mathbf{a}_\tau, \mathbf{s}_\tau, \nu_{i\tau}) \middle| \mathbf{s}_t\right]$$

where the expectation is over current shocks and actions, as well as future states, actions, and shocks.

State Transitions & Equilibrium

- $\mathbf{s}_{t+1}$ is drawn from a probability distribution $P\left(\mathbf{s}_{t+1}|\mathbf{a}_t, \mathbf{s}_t\right)$
- They focus on pure strategy MPE
- A Markov strategy is a function $\sigma_i : S \times \mathcal{V}_i \rightarrow A_i$
- A profile of Markov strategies is a vector, $\sigma = (\sigma_1, ..., \sigma_N)$, where $\sigma : S \times \mathcal{V}_1 \times ... \times \mathcal{V}_N \rightarrow A$
- Given σ, firm i's expected profit can then be written recursively

$$V_i(\mathbf{s};\sigma) = E_\nu \left[\pi_i(\sigma(\mathbf{s},\nu), \mathbf{s}, \nu_i) + \beta \int V\left(\mathbf{s}';\sigma\right) dP\left(\mathbf{s}'|\sigma(\mathbf{s},\nu), \mathbf{s}\right) | \right.$$

- The profile σ is a MPE if, given opponent profile σ_{-i}, each firm i prefers strategy σ_i to all other alternatives σ_i'

$$V_i(\mathbf{s};\sigma) \geq V_i\left(\mathbf{s};\sigma_i', \sigma_{-i}\right)$$

Structural Parameters

- The structural parameters of the model are the discount factor β, the profit functions $\pi_1, ..., \pi_N$, the transition probabilities P, and the distributions of private shocks $G_1, ..., G_N$.
- Like AM, they treat β as known and estimate P directly from the observed state transitions.
- They assume the profits and shock distributions are known functions of a parameter vector $\theta : \pi_i(\mathbf{a}, \mathbf{s}, \nu_i; \theta)$ and $G_i\left(\nu_i | \mathbf{s}, \theta\right)$.
- The goal is to recover the true θ under the assumption that the data are generated by a MPE.

Example: Dynamic Oligopoly

- Their main (novel) example is based on the EP framework.
- Incumbent firms are heterogeneous, each described by its state $z_{it} \in \{1, 2, ..., \overline{z}\}$; potential entrants have $z_{it} = 0$
- Incumbents can make an investment $I_{it} \geq 0$ to improve their state
- An incumbent firm i in period t earns

$$q_{it} \left(p_{it} - mc \left(q_{it}, s_t; \mu\right)\right) - C \left(I_{it}, \nu_{it}; \xi\right)$$

  where p_{it} is firm i's price, $q_{it} = q_i \left(\mathbf{s}_t, \mathbf{p}_t; \lambda\right)$ is quantity, $mc\left(\cdot\right)$ is marginal cost, and ν_{it} is a shock to the cost of investment.
- $C \left(I_{it}, \nu_{it}; \xi\right)$ is the cost of investment.

Example: Dynamic Oligopoly

- Competition is assumed to be static Nash in prices.
- Firms can also enter and exit.
- Exitors receive ϕ and entrants pay ν_{et}, an *iid* draw from G_e
- In equilibrium, incumbents make investment and exit decisions to maximize expected profits.
- Each incumbent i uses an investment strategy $I_i(\mathbf{s}, \nu_i)$ and exit strategy $\chi_i(\mathbf{s}, \nu_i)$ chosen to maximize expected profits.
- Entrants follow a strategy $\chi_e(\mathbf{s}, \nu_e)$ that calls for them to enter if the expected profit from doing so exceeds its entry cost.

First Stage Estimation

- The goal of the first stage is to estimate the state transition probabilities $P(\mathbf{s}'|\mathbf{a}, \mathbf{s})$ and equilibrium policy functions $\sigma(\mathbf{s}, \nu)$
- The second stage will use the equilibrium conditions from above to estimate the structural parameters θ
- In order to obtain consistent first stage estimates, they must assume that the data are generated by a single MPE profile σ
  - This assumption has a lot of bite if the data come from multiple markets
  - It's quite weak if the data come from only a single market

First Stage Estimation

- Stage "0":
  - The static payoff function will typically be estimated "off line" in a 0^{th} stage (e.g. BLP, Olley-Pakes)
- Stage 1:
  - It's usually fairly straightforward to run the first stage: just "regress" actions on states in a "flexible manner".
  - Since these are not structural objects, you should be as flexible as possible. Why?
  - Of course, if $\mathbf{s}$ is big, you may have to be very parametric here (i.e. OLS regressions and probits).
    - In this case, your second stage estimates will be inconsistent...
  - Continuous actions are especially tricky (hard to be nonparametric here)

Estimating the Value Functions

- After estimating policy functions, firm's value functions are estimated by forward simulation.
- Let $V_i(\mathbf{s}, \sigma; \theta)$ denote the value function of firm i at state s assuming firm i follows the Markov strategy σ_i and rival firms follow σ_{-i}
- Then

$$V_i(\mathbf{s}, \sigma; \theta) = E\left[\sum_{0=t}^{\infty} \beta^t \pi_i(\sigma(\mathbf{s}_t, \nu_t), \mathbf{s}_t, \nu_{it}; \theta) | \mathbf{s}_0 = \mathbf{s}; \theta\right]$$

  where the expectation is over current and future values of $\mathbf{s}_t$ and ν_t
- Given a first-stage estimate $\widehat{P}$ of the transition probabilities, we can simulate the value function $V_i(\mathbf{s}, \sigma; \theta)$ for any strategy profile σ and parameter vector θ.

Estimating the Value Functions

- A single simulated path of play can be obtained as follows:
    1. Starting at state $\mathbf{s}_0 = \mathbf{s}$, draw private shocks ν_{i0} from $G_i(\cdot|\mathbf{s}_0, \theta)$ for each firm i.
    2. Calculate the specified action $a_{i0} = \sigma_i(\mathbf{s}_0, \nu_{i0})$ for each firm i, and the resulting profits $\pi_i(\mathbf{a}_0, \mathbf{s}_0, \nu_{i0}; \theta)$
    3. Draw a new state $\mathbf{s}_1$ using the estimated transition probabilities $\widehat{P}(\cdot|\mathbf{a}_0, \mathbf{s}_0)$
    4. Repeat steps 1-3 for T periods or until each firm reaches a terminal state with known payoff (e.g. exits from the market)
- Averaging firm i's discounted sum of profits over many paths yields an estimate $\widehat{V}_i(\mathbf{s}, \sigma; \theta)$, which can be obtained for any (σ, θ) pair, including both the "true" profile (which you estimated in the first stage) and any alternative you care to construct.

Special Case of Linearity

- Forward simulation yields a low cost estimate of the V's for different σ's given θ, but the procedure must be repeated for each candidate θ.
- One case is simpler.
- If the profit function is linear in the parameters θ so that

$$\pi_i\left(\mathbf{a}, \mathbf{s}, \nu_i; \theta\right) = \psi_i\left(\mathbf{a}, \mathbf{s}, \nu_i\right) \cdot \theta$$

  we can then write the value function as

$$V_i\left(\mathbf{s}, \sigma; \theta\right) = E\left[\sum_{t=0}^{\infty} \beta^t \psi_i(\sigma\left(\mathbf{s}_t, \nu_t\right), \mathbf{s}_t, \nu_{it}) \middle| \mathbf{s}_0 = \mathbf{s}\right] \cdot \theta = \mathbf{W}_i\left(\mathbf{s}; \theta\right) \cdot$$

- In this case, for any strategy profile σ, the forward simulation procedure only needs to be used once to construct each $\mathbf{W}_i$.
- You can then obtain V_i easily for any value of θ.

Second Stage Estimation

- The first stage yields estimates of the policy functions, state transitions, and value functions.
- The second stage uses the model's equilibrium conditions

$$V_i\left(\mathbf{s};\sigma_i,\sigma_{-i};\theta\right) \geq V_i\left(\mathbf{s};\sigma_i',\sigma_{-i};\theta\right)$$

  to recover the parameters θ that rationalize the strategy profile σ observed in the data.
- They show how to do so for both set and point identified models
- We will focus on the point identified case here.

Second Stage Estimation

- To see how the second stage works, define

$$g\left(x;\theta,\alpha\right)=V_i\left(\mathbf{s};\sigma_i,\sigma_{-i};\theta,\alpha\right)-V_i\left(\mathbf{s};\sigma_i',\sigma_{-i};\theta,\alpha\right)$$

  where $x\in\mathcal{X}$ indexes the equilibrium conditions and α represents the first-stage parameter vector.
- The inequality defined by x is satisfied at θ,α if $g\left(x;\theta,\alpha\right)\geq 0$
- Define the function

$$Q\left(\theta,\alpha\right)\equiv\int\left(\min\left\{g\left(x;\theta,\alpha\right),0\right\}\right)^2 dH(x)$$

  where H is a distribution over the set $\mathcal{X}$ of inequalities.

Second Stage Estimation

- The true parameter vector θ_0 satisfies

$$Q\left(\theta_0, \alpha_0\right) = 0 = \min_{\theta \in \Theta} Q(\theta, \alpha_0)$$

  so we can estimate θ by minimizing the sample analog of $Q(\theta, \alpha_0)$
- The most straightforward way to do this is to draw firms and states at random and consider alternative policies σ_i' that are slight perturbations of the estimated policies.

Second Stage Estimation

- We can then use the above forward simulation procedure to construct analogues of each of the V_i terms and construct

$$Q(\theta, \alpha) \equiv \frac{1}{n_I} \sum_{k=1}^{n_I} \left(\min\left\{\widehat{g}\left(X_k; \theta, \alpha\right), 0\right\}\right)^2$$

- How? By drawing n_I different alternative policies, computing their values, finding the difference versus the optimal policy payoff, and using an MD procedure to estimate the parameters that minimize these profitable deviations.
- Their estimator minimizes the objective function at $\alpha = \widehat{\alpha}_n$

$$\theta = \arg\min_{\theta \in \Theta} Q_n\left(\theta, \widehat{\alpha}_n\right)$$

- See the paper for the technical details.

Example: Dynamic Oligopoly

- Let's see how they estimate the EP model.
- First, they have to choose some parameterizations.
- They assume a logit demand system for the product market.
- There are M consumers with consumer r deriving utility U_{ri} from good i

$$U_{ri} = \gamma_0 \ln(z_i) + \gamma_1 \ln\left(y_r - p_i\right) + \varepsilon_{ri}$$

  where z_i is the quality of firm i, p_i is firm i's price, y_r is income, and ε_{ri} is an *iid* logit error
- All firms have identical constant marginal costs of production

$$mc(q_i; \mu) = \mu$$

Example: Dynamic Oligopoly

- Each period, firms choose investment levels $I_{it} \in \mathcal{R}_+$ to increase their quality in the next period.
- Firm i's investment is successful with probability

$$\frac{\rho I_{it}}{(1 + \rho I_{it})}$$

  in which case quality increases by one, otherwise it doesn't change.
- There is also an outside good, whose quality moves up by one with probability δ each period.
- Firm i's cost of investment is

$$C(I_i) = \xi \cdot I_i$$

  so there is no shock to investment (it's deterministic)

Example: Dynamic Oligopoly

- The scrap value ϕ is constant and equal for all firms.
- Each period, the potential entrant draws a private entry cost ν_{et} from a uniform distribution on $\left[\nu^L, \nu^H\right]$
- The state variable $s_t = (N_t, z_{1t}, ..., z_{Nt}, z_{out,t})$ includes the number of incumbent firms and current product qualities.
- The model parameters are $\gamma_0, \gamma_1, \mu, \xi, \phi, \nu^L, \nu^H, \rho, \delta, \beta$, & y
- They assume that β & y are known, ρ & δ are transition parameters estimated in a first stage, γ_0, γ_1, & μ are demand parameters (also estimated in a first stage), so the main (dynamic) parameters are simply $\theta = \left(\xi, \phi, \nu^L, \nu^H\right)$
- Due to the computational burden of the PM algorithm, they consider a setting in which only ≤ 3 firms can be active.
- They generated datasets of length 100-400 periods using PM.

Example: Dynamic Oligopoly

- Here are the parameters they use.

Table 2: Dynamic Oligopoly Monte Carlo Parameters

Parameter	Value	Parameter	Value
Demand:		Investment Cost:	
γ_1	1.5	ξ	1
γ_0	0.1		
M	5	Marginal Cost:	
y	6	μ	3
Investment Evolution		Entry Cost Distribution	
δ	0.7	ν^l	7
ρ	1.25	ν^h	11
Discount Factor		Scrap Value:	
β	0.925	ϕ	6

Example: Dynamic Oligopoly

- The first stage requires estimation of the state transitions and policy functions (as well as the demand and *mc* parameters).
- For the state transitions, they used the observed investment levels and qualities to estimate ρ and δ by MLE.
- They estimated the demand parameters by MLE as well, using quantity, price, and quality data.
- They recover μ from the static mark-up formula.
- They used local linear regressions with a normal kernel to estimate the investment, entry, and exit policies.

Example: Dynamic Oligopoly

- Given strategy profile $\sigma = (I, \chi, \chi_e)$, the incumbent value function is

$$\begin{aligned} V_i(s;\sigma) &= W^1(s;\sigma) + W^2(s;\sigma)\cdot\xi + W^3(s;\sigma)\cdot\phi \\ &= E\left[\sum_{t=0}^{\infty}\beta^t\widetilde{\pi}_i(s_t)|s=s_0\right] - E\left[\sum_{t=0}^{\infty}\beta^t I_i(s_t)|s=s_0\right]\cdot\xi \\ &\quad + E\left[\sum_{t=0}^{\infty}\beta^t\chi_i(s_t)|s=s_0\right]\cdot\phi \end{aligned}$$

  where the first term $\widetilde{\pi}_i(s_t)$ is the static profit of incumbent i given state s_t
- The 2^{nd} term is the expected PV of investment
- The 3^{rd} term is the expected PV of the scrap value earned upon exit.

Example: Dynamic Oligopoly

- To apply the MD estimator, they constructed alternative investment and exit policies by drawing a mean zero normal error and adding it to the estimated first stage investment and exit policies.
- They used $n_s = 2000$ simulation paths, each having length at most 80, to compute the PV W^1, W^2, W^3 terms for these alternative policies.
- They can then estimate ξ & ϕ using their MD procedure.
- It's also straightforward to estimate the entry cost distribution (parametrically or non-parametrically) - see the paper for details.

Example: Dynamic Oligopoly

Table 3: Dynamic Oligopoly With Nonparametric Entry Distribution

	Mean	SE(Real)	5%(Real)	95%(Real)	SE(Subsampling)
$n = 400$, $n_I = 500$					
ξ	1.01	0.05	0.91	1.10	0.03
ϕ	5.38	0.43	4.70	6.06	0.39
$n = 200$, $n_I = 500$					
ξ	1.01	0.08	0.89	1.14	0.05
ϕ	5.32	0.56	4.45	6.33	0.53
$n = 100$, $n_I = 300$					
ξ	1.01	0.10	0.84	1.17	0.06
ϕ	5.30	0.72	4.15	6.48	0.72

- Truth: $\xi = 1$ & $\phi = 6$
- For small sample sizes, there is a slight bias in the estimates of the exit value.
- Investment cost parameters are spot on.

Example: Dynamic Oligopoly

Table 4: Dynamic Oligopoly With Parametric Entry Distribution

	Mean	SE(Real)	5%(Real)	95%(Real)	SE(Subsampling)
$n = 400,\ n_I = 500$					
ξ	1.01	0.06	0.92	1.10	0.04
ϕ	5.38	0.42	4.68	6.03	0.41
ν^l	6.21	1.00	4.22	7.38	0.26
ν^h	11.2	0.67	10.2	12.4	0.30
$n = 200,\ n_I = 500$					
ξ	1.01	0.07	0.89	1.13	0.05
ϕ	5.28	0.66	4.18	6.48	0.53
ν^l	6.20	1.16	3.73	7.69	0.34
ν^h	11.2	0.88	9.99	12.9	0.40
$n = 100,\ n_I = 300$					
ξ	1.01	0.10	0.84	1.17	0.06
ϕ	5.43	0.81	4.26	6.74	0.75
ν^l	6.38	1.42	3.65	8.43	0.51
ν^h	11.4	1.14	9.70	13.3	0.58

- Truth: $\xi = 1, \phi = 6, \nu^l = 7$, & $\nu^h = 11$
- The subsampled standard errors are on average slightly smaller than the true SEs.
- This is likely due to small sample sizes.

Example: Dynamic Oligopoly

Figure 1: Entry Cost Distribution for $n = 400$

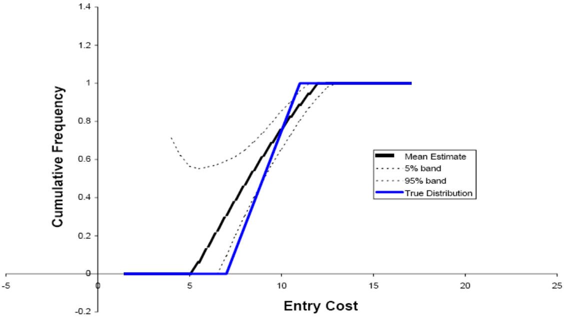

- The entry cost distribution is recovered quite well, despite small sample size (and few entry events).

Conclusions

- Both AM & BBL are based on the same underlying idea (CCP estimation)
- As such, it's quite possible to mix and match from the two approaches
  - e.g. forward simulate the CV terms and use a MNL likelihood
- We have found AM-style approaches easier to implement, but that might be idiosyncratic.
- Applications in marketing are growing: Goettler and Gordon (2012), Ellickson, Misra, Nair (2012), Misra and Nair (2011), Chung et al. (2012) ...
- If you are interested in applying this stuff, you should read *everything* you can get your hands on!