Lecture Notes: Estimation of Dynamic Games

Jean-François Houde
Cornell University & NBER

April 20, 2018

1. Introduction: Pseudo-Maximum-Likelihood (PML) estimators for discrete games of incomplete information.

2. PML estimators for dynamic discrete games: Aguirregabiria and Mira (2007).

3. Simulated estimators for dynamic games: Bajari, Benkard, and Levin (2007).

4. Related approaches: Pakes, Ostrovsky, and Berry (2007), Pesendorfer and Schmidt-Dengler (2008).

5. Applications: Schmidt-Dengler (2006), Ryan (2012), Collard-Wexler (2013), Sweeting (2013).

6. Alternative approaches based on "nested fixed-point" estimators (i.e. full-solution): Goettler and Gordon (2011), Takahashi (2015), Igami (2017).

Introduction: PML estimators for static entry games

- **Model assumptions/notation:**

1. M independent markets.
2. Players: $i \in \{1, ..., N_m\}$
3. Actions (simultaneous): $a_i \in \{0, 1\}$ (e.g. Enter/Exit)
4. Payoffs:

$$\Pi_i(a_i, a_{-i}, X, \epsilon_i) = (1 - a_i)\epsilon_i(0) + a_i\Big(X\beta_i - \delta \log(1 + \sum_{j \neq i} a_j) + \epsilon_i(1)\Big).$$

5. Private information: $\epsilon_{im}(1) - \epsilon_{im}(0) \sim F(\cdot)$ + iid over players and markets.
6. Strategies in probability space:

$$P_i^\sigma(a_i|X) = \Pr\left(a_i = \arg\max_{a_i} E_{a_{-i}}\Big[\Pi(a_i, a_{-i}, X, \epsilon_i)|P_{-i}^\sigma\Big]\right)$$

7. Expected payoffs:

$$\begin{aligned} \pi(a_i = 1, X) &= X\beta_i - \sum_{a_{-i} \in \{0,1\}^{N-1}} \Big[\prod_{j \neq i} P_j^\sigma(a_j|X)\Big]\delta \log(1 + \sum_{j \neq i} a_j) \\ &= X\beta_i - \delta h(X|P^\sigma) \end{aligned} \tag{1}$$

8. Bayesian Nash Equilibirum: $P^*(X) = \{P_1^*(X), ..., P_N^*(X)\}$ is a BNE if P_i^* is a fixed point of the best-response (choice probability) mapping:

$$P_i^*(1|X) = F(X_i\beta_i - \delta h(X|P^*)) = \Psi_i(1|X, P^*).$$

- **Maximum Likelihood Estimator (MLE):**
  - Data: $\{a_{1m}, ..., a_{N_m m}, X_{1m}, ..., X_{N_m m}\}_{m=1,...,M}$
  - Assumption: $P^*(X)$ is unique for all X.
  - Likelihood function:

$$\begin{aligned} Q^{MLE}(A, X) &= \max_{\theta} \sum_m \sum_i \log P^*(a_{im}|X_m, \theta) \\ &\text{s.t.} \qquad P_i^*(a_{im}|X_m, \theta) = \Psi_i(a_{im}|X_m, P^*, \theta) \end{aligned} \tag{2}$$

  - If the BNE is unique, $Q^{MLE}(A, X)$ is solved using a Nested Fixed Point (NFXP) algorithm (e.g. Seim (2006)).

– If the set BNE is not singleton but finite, two possible representations of the MLE problem:

1. All markets "select" the same equilibrium τ^*:

$$Q^{MLE}(A,X) = \max_{\theta} \left\{ \max_{\tau \in \mathcal{T}(X,\theta)} \sum_m \sum_i \log P_i^*(a_{im}|X_m,\theta,\tau) \right\} \quad (3)$$
$$\text{s.t.} \qquad P_i^*(a_{im}|X_m,\theta,\tau) = \Psi_i(a_{im}|X_m,P^*,\theta,\tau)$$

where τ indexes a specific BNE, and $\mathcal{T}(X,\theta)$ is the set of equilibria.
Note: Not very practical because the set of equilibria typically change with X_m, and therefore across markets.

2. Markets have a common equilibrium selection probability μ_τ:

$$Q^{MLE}(A,X) = \max_{\theta,\mu} \left\{ \sum_m \log \left(\sum_{\tau \in \mathcal{T}(X_m,\theta)} \mu_\tau \sum_i P_i^*(a_{im}|X_m,\theta,\tau) \right) \right\} \quad (4)$$
$$\text{s.t.} \qquad P_i^*(a_{im}|X_m,\theta,\tau) = \Psi_i(a_{im}|X_m,P^*,\theta,\tau)$$

Note: μ is potentially a very large dimensional object, equilibria are difficult to characterize or if the set changes significantly with X_m (e.g. one selection rule per market in the limit).

∗ Can be computationally very difficult: For each θ we must identify all BNE.

- **PML estimator:**
  - **Assumption on the DGP:** For every sampled market, choices are generated by the **same** equilibrium τ:

$$P^o_{im}(X_m) = P^o_i(X_m) = \Psi_i(X_m, P^o, \theta^o, \tau), \quad \forall m$$

    where P^o denotes the population choice probabilities.
  - If we know P^o, the PML estimator $\hat{\theta}^{PML}$ is the solution to the following problem:

$$Q^{PML}(A, X) = \max_{\theta} \sum_m \sum_i \log \Psi_i(a_{im}|X_m, P^o, \theta) \tag{5}$$

  - Intuition: If we know P^o we know the beliefs used by players to generate their action. Since the data is generated by one specific equilibrium, $\Psi_i(a_{im}|X_m, P^o)$ is the equilibrium best-response of player i. Therefore, $\hat{\theta}^{PML} = \hat{\theta}^{MLE}$.
  - Problem: We don't know P^o!

- **Feasible PML estimator:** (a.k.a. 2-Step estimator)
  - Let $\hat{P}^1$ be a consistent estimator of P^o such that

$$\sqrt{M}(\hat{P}^1 - P^o) \to_d N(0, \Sigma).$$

  - **Examples:**
    1. Frequency estimator:

$$\hat{P}^1(1|X) = \frac{1}{M(X)} \sum_{m|X_m=X} \sum_i \frac{a_{im}}{N_m}$$

    2. Series estimator:

$$\hat{P}^1(1|X) = \frac{\exp\left(g(X)\hat{B}\right)}{1 + \exp\left(g(X)\hat{B}\right)},$$

       where $g(X)$ is a flexible polynomial function of the state variables.
  - Using $\hat{P}^1$ we can get a 2-step feasible PML estimator of θ:

$$Q^{2S}(A, X) = \max_\theta \sum_m \sum_i \log \Psi_i(a_{im}|X_m, \hat{P}^1, \theta) \tag{6}$$

  - **Problem:** $\hat{\theta}^{2s}$ is less efficient than $\hat{\theta}^{PML}$ (Aguirregabiria and Mira (2007)).
  - Small sample bias in the estimation of $\hat{P}^1$ can generate severe bias in θ.

- **Nexted-Pseudo Likelihood (NPL) estimator:** (a.k.a. K-Step estimator)
  - **Objective:** Reduce the small sample bias of the 2-step estimator by iterating on the feasible PML estimator.
  - **Algorithm:** Start with initial estimate of the choice probabilities $\hat{P}^1(X)$.
    1. PML step:
    $$\hat{\theta}^{k+1} = \arg\max_{\theta} \sum_m \sum_i \log \Psi_i(a_{im}|X_m, \hat{P}^k, \theta)$$
    2. Policy function step (i.e. best-response):
    $$\hat{P}_i^{k+1}(a_{im}|X_m) = \Psi_i(a_{im}|X_m, \hat{P}^k, \hat{\theta}^{k+1})$$
    3. Stop if $||\hat{P}^{k+1} - \hat{P}^k|| < \delta$, otherwise repeat step 1 and 2.
  - **Advantages:**
    * $\hat{P}^1$ can be a very imprecise estimate of P^o.
    * The best-response iteration step will use the model to improve the estimates about players' beliefs.
    * $\hat{P}^1$ does not need to be a consistent estimate of P^o.
    * Allows the researcher to add unobserved heterogeneity.
    * Computationally easier than the MLE.

– **Problems:**

  * Convergence problems.
  * In practice AM suggest to perform K replications, and set the NPL estimate to the highest likelihood:

$$\hat{\theta}^{NPL} = \arg\max_{\hat{\theta}^k} Q^k(A, X|\hat{\theta}^k).$$

PML estimator for Dynamic Discrete Games: Aguirregabiria and Mira (2007)

- **Maintained assumption:**
  - Conditional independence of $\epsilon_i(a)$ with respect to agents, actions, markets, and time.
  - The DGP select one equilibria for all markets/time.
- **Model:** Entry/Exit w/o absorbing state (i.e. re-entry is possible).
  - Common knowledge state:

$$x = \{a_1, ..., a_N, S\}$$

    where is a_j is the activity state of firm j last period and S is the market size.
  - Static profit function:

$$\Pi_i(x, \epsilon | a_i', a_{-j}') = \begin{cases} \theta_s \ln S - \ln\left(1 + \sum_{j \neq i} a_j'\right)\theta_c + \theta_{fc} - \theta_e 1(a_i = 0) + \omega + \epsilon(1) & \text{if } a_i' = 1 \\ \epsilon(0) & \text{if } a_i' = 0. \end{cases}$$

    Where ω is market-specific unobserved heterogeneity parameter, with $\omega_m \in \{\omega_0, ..., \omega_K\}$. The probability of a market being of type k is $\psi(k)$, with $\sum_k \psi(k) = 1$.

– Expected profits conditional on $a_i' = 1$ and competitors strategies $P^\sigma(a_j|x)$:

$$\begin{aligned}\pi_i(x|a_i'=1,P^\sigma) &= \theta_s \ln S - \sum_{a_j' \in \{0,1\}^{N-1}} \Big[\prod_{j\neq i} P^\sigma(a_j'|x)\Big] \ln\Big(1+\sum_{j\neq i} a_j\Big)\theta_c \\ &\quad +\theta_{fc} - \theta_e 1(a_i=0) + \omega \\ &= \theta_s \ln S - \theta_c h(x|P^\sigma) + \theta_{fc} - \theta_e 1(a_i=0) + \omega \\ &= X_i(P^\sigma)\theta + \omega \qquad (7)\end{aligned}$$

– Value function given strategies P^σ:

$$V_i(x|P^\sigma) = \sum_{a_i'} P_i^\sigma(a_i|x)\Big[\pi_i(a_i|x,P^\sigma) + e_i(a_i',x|P^\sigma) + \beta \sum_{x'} V_i(x'|P^\sigma) q^\sigma(x'|x,a_i')\Big]$$

Where, $\quad q_i^\sigma(x'|x,a_i') = \Big[\prod_{j\neq i} P^\sigma(a_j'|x)\Big] g(S'|S)$

where $e_i(a_i',x|P^\sigma)$ is defined in the single agent model section.

– In Matrix form $V_i(P^\sigma)$ is given by:

$$\begin{aligned}V_i(P^\sigma) &= \sum_{a_i'} P_i^\sigma(a_i') * \Big[\pi_i(a_i',P^\sigma) + e_i(a_i'|P^\sigma) + \beta Q_i^\sigma(a_i') V_i(P^\sigma)\Big] \\ &= \Big(I - \beta \sum_{a_{-i}'} P^\sigma(a_i') * Q_i^\sigma(a_i')\Big)^{-1} \sum_{a'} P(a_i') * \Big[\pi_i(a_i',P^\sigma) + e_i(a_i'|P^\sigma)\Big] \qquad (8)\end{aligned}$$

– A Markov Perfect Equilibirum in probability space is then defined as a fixed point of the CCP mapping:

$$\begin{aligned} P_i^*(1|x) &= F(\epsilon(1) - \epsilon(0) > v_i(x|0, P^*) - v_i(x|1, P^*)) \\ &= \Psi_i(1|x, P^*) \end{aligned}$$

where $v_i(x|a, P^*) = \pi_i(x|a_i, P^*) + \beta \sum_{x'} V(x'|P^*) q_i^*(x'|x, a)$.

– As in the single agent case if $\epsilon(a)$ is distributed according to an EV distribution, if $a = 1$ leads to an absorbing state (e.g. exit in EP) the choice specific value functions can be expressed

$$v_i(x|a, P^*) = \ln P_i^*(a|x) - \ln P_i^*(0|x)$$

• **Data:**

– Panel of markets: $\{a_{imt}, x_{mt}\}$

– Number of players:

$$N_m = \max_t \{n_{mt} + e_{mt}\}$$

– Transition probabilities: $g(S'|S)$ is estimated directly from the data and discretized into $|S|$ grids.

- **Feasible PML estimator (i.e 2-step)**:

  1. Reduced form CCP estimator: $\hat{P}_i^1(a|x_{mt})$ for all (i, m, t).
  2. PML step:
$$\hat{\theta}^1 = \arg\max_{\theta} \sum_t \sum_m \sum_i \log \Psi_i(a_{imt}, |x_{mt}, \hat{P}^1) \tag{9}$$

- **NPL estimator (i.e. K-step)**:

  0. Reduced-form CCP estimator: $\hat{P}_i^1(a|x_{mt})$ for all (i, m, t).
  1. PML step:
$$\hat{\theta}^{k+1} = \arg\max_{\theta} \sum_t \sum_m \sum_i \log \Psi_i(a_{imt}, |x_{mt}, \hat{P}^k) \tag{10}$$
  2. Best-response updating step:
$$\hat{P}_i^{k+1}(1|x_{mt}) = \Psi_i(1|x_{mt}, \hat{P}^k, \hat{\theta}^{k+1})$$
  3. Repeat step 1 and 2 K times, and pick the highest likelihood iteration (or until convergence of $\hat{P}(a)$).

- **Tractable NPL estimator with unobserved market heterogeneity:**
  - Let $\omega_m \sim N(0, \sigma_\omega^2)$, and $\bar{\omega}_l \in \{\bar{\omega}_1, ..., \bar{\omega}_L\}$ be a discrete approximation of ω_m such that:
$$\Pr(\omega_m = \bar{\omega}_l) = \psi(l) = \frac{1}{L}$$
    and $\bar{\omega}_l = \int_{\omega_{l-1}}^{\omega l+1} \phi(\omega/\sigma_\omega) d\omega$.
  - This specification adds only 1 parameter to estimate: σ_ω.
  - The 2-step estimator is not feasible: We cannot get a consistent estimate fo $\hat{P}^1(a|x, \bar{\omega}_l)$. The NPL is feasible however.
  - **Assumption:** The initial states x_{m1} are drawn from the invariant distribution generated by the industry state markov chain $Q(x'|x, P^*, \bar{\omega}_l)$.
  - Modified NPL algorithm:
    0. Reduced-form CCP estimator: $\hat{P}_i^1(a|x_{mt}, \bar{\omega}_l)$ for all (i, m, t) and $\bar{\omega}_l$ (e.g. $\hat{P}_i^1(a|x_{mt}, \bar{\omega}_l) = \hat{P}_i^1(a|x_{mt}, \bar{\omega}_k), \forall k$) .
    1. PML step:
$$\hat{\theta}^{k+1} = \arg\max_\theta \sum_m \log \left[\sum_l \psi(l) \prod_t \prod_i \Psi_i(a_{imt}, |x_{mt}, \hat{P}^k, \bar{\omega}_l) \times p^*(x_{m1}|\bar{\omega}_l, \hat{P}^k) \right] \tag{11}$$

where $p^*(x_{m1}|\bar{\omega}_l, \hat{P}^k)$ is the invariant distribution of states conditional on type $\bar{\omega}_l$ (i.e. solves the initial condition problem):

$$\begin{aligned} p^*(x_{m1}|\bar{\omega}_l, \hat{P}^k) &= \sum_{x_0} Q(x_{m1}|x_0, \hat{P}^k, \bar{\omega}_l) p^*(x_0|\bar{\omega}_l, \hat{P}^k) \\ &= \sum_{x_0} \Big[\prod_i \hat{P}^k(a_{im1}|x_0, \bar{\omega}_l) \Big] g(S_{m1}|S_0) p^*(x_0|\bar{\omega}_l, \hat{P}^k) \end{aligned}$$

where $a_{m0} = \{a_{1m1}, ..., a_{Nm1}\}$ is such that $x_{m1} = \{a_{m1}, S_{m1}\}$.

2. Best-response updating step:

$$\hat{P}_i^{k+1}(1|x_{mt}, \bar{\omega}_l) = \Psi_i(1|x_{mt}, \bar{\omega}_l, \hat{P}^k, \hat{\theta}^{k+1})$$

3. Repeat step 1 and 2 K times, and pick the highest likelihood iteration (or until $\hat{P}^k$ converges).

- **Caveats:**
  - If the model does not have an absorbing state: Computationally demanding to invert $\big(I - \beta \sum_a P(a) * Q(x'|x, a)\big)$ when the state space is large.
  - Convergence is not guaranteed.
- **Application:** Entry and Exit into Chilean retail markets.
  - Data: 189 communas in Chile between 1994-1999 for 5 retail industries.

Simulated Estimators for Dynamic Games: Bajari, Benkard, and Levin (2007)

- **Introduction:** Simulation methods for DDC (Hotz, Miller, Sanders, and Smith (1994)).
- **Two-stage method:**
  1. Estimate $f(x'|x,a)$ and $P(a|x)$ from the data.
  2. Approximate $V(x|\hat{P})$ using monte-carlo integration, instead of inverting the transition matrix.
- **Simulator for dynamic discrete choices:** For each $x \in X$ draw S sequences of future paths. Each sequence is generated as followed:
  1. Draw iid shocks: $\epsilon_0(a) \sim F(.|x_0)$
  2. Compute optimal policies given $\hat{P}(a|x_0)$:
$$a_0 = \hat{P}^{-1}(\epsilon_0|x) \tag{12}$$
  3. Draw next period state: $x_1 \sim \hat{f}(x'|x_0,a_0)$.
  4. Compute payoffs: $U_1^s = u(x_0,a_0) + \epsilon_0(a_0)$
  5. Repeat steps 1 to 4 T times.

6. If t=T, compute the discounted value of sequence s:

$$V^s(x_0) = \sum_t^T \beta^t U_t^s$$

The value function is then approximated by:

$$\hat{V}(x_0|\hat{P}) = \frac{1}{S}\sum_s V^s(x_0)$$

- With $\hat{V}(P)$ in hand, we can construct various estimators of the remaining parameters θ:
  - Pseudo-likelihood estimator:

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

  - GMM estimator:

$$\min_\theta \frac{1}{I}\sum_i \Big(\sum_t (\tilde{v}(a_{it}|x_{it}) - \tilde{v}(a_{it}|x_{it}, \hat{P})) \times Z_{it}\Big) W^{-1} \Big(\sum_t (\tilde{v}(a_{it}|x_{it}) - \tilde{v}(a_{it}|x_{it}, \hat{P})) \times Z_{it}\Big)^T$$

  - Least-Square estimator:

$$\min_\theta \Big(\hat{P}(a|x) - \Psi(a|x, \hat{P})\Big)^T W^{-1} \Big(\hat{P}(a|x) - \Psi(a|x, \hat{P})\Big)$$

    The LS estimator uses the model equilibrium conditions (i.e. $P(a|x) = \Psi(a|x, P)$) to form $(A-1)\cdot|X|$ moment conditions (see **?**)).

- **BBL contribution:**
  - Generalizes the idea of HMSS to continuous strategies.
  - Propose a different estimator based on inequality conditions, which includes a bound estimator for partially identified model.
  - Use the same equilibrium selection argument as AM to deal with the multiple equilibria problem in strategic environment.
- **Example:** Dynamic investment game with entry/exit (i.e. EP).
  - Value function:

$$V_i(s|\sigma) = E_\nu\Big(\pi_i(\sigma(s,\nu), s, \nu_i) + \beta \int V_i(s'|\sigma) dP(s'|s, \sigma(s,\nu))|s\Big)$$

    Where $\sigma_i(s,\nu_i) = \{\chi_i(s,\nu_i), I_i(s,\nu_i)\}$ is the strategy followed by player i.
  - A strategy profile $\sigma(s,\nu) = \{\sigma_i(s,\nu_i\}$ is a MPE if:

$$V_i(s|\sigma) \geq V_i(s|\sigma_{-i}, \sigma_i')$$

  - Primitives of the model:
    - $P(s'|s,\sigma)$ is assumed to be known up to a parameter vector α (estimated separately).
    - Entry cost: $\phi_i^e \sim G^{\phi^e}(\cdot|\theta)$

* Scrap value: ϕ is common across firms.
* Static profits of incumbents:

$$\pi_i(\sigma, s, \nu_i) = \underbrace{(p_i - mc(q_i, s, |\mu))}_{\text{Estimated separately}} - C(I_i, \nu|\xi)$$

where $\nu \sim G^{\nu}(\cdot)$ and iid over time and firms.

- **Two-step estimation procedure:**

1. Estimation of policy functions and transition probabilities:
   - Entry/Continuation probabilities:

$$\hat{r}_i^e(s) = F(s|\hat{\beta})$$

Where $F(S|\beta)$ is typically a smooth and flexible function of s.

   - Investment strategies:

Assumption: For all i and ν, $\partial^2 \pi_i(\sigma, s, \nu_i)/\partial I_i \partial \nu_i > 0$.

Then $I(s, \nu)$ is increasing in ν, and invertible. To see this, let $F_i(I|s)$ be the probability of observing investments lower than I in state s (i.e. observed in the data). Then:

$$F_i(I|s) = \Pr(I_i(s, \nu_i) < I|s) = G_i(\nu_i < I_i^{-1}(I|s)|\theta)$$

For instance if $\nu \sim N(0, \sigma_\nu^2)$, the inverse of the investment policy is given by:

$$I_i^{-1}(I|s) = \sigma_\nu \Phi^{-1}(F(I|S)) = \nu_i$$

Conversely, for a given ν the optimal investment is given by:

$$I_i(s, \nu) = F^{-1}(G_i(\nu|s)|s)$$

In practice $F_i(I|s)$ should be estimated using a flexible semi-parametric form, so that $\hat{I}(s, \nu)$ has full support.

– **Estimation of the Value function:** As in the single agent example, we want to simulate forward S paths. Each path is generated as followed:

0. Set $s_0 = s$
1. Draw private value shocks for all i: $\nu_i \sim G(\cdot|s, \theta)$ and $\phi_i^e \sim G^\psi(\cdot|x, \theta)$.
2. Compute actions given strategies $\sigma_i(s_0, \nu_i)$:

$$I_i(s_0, \nu_i) = F^{-1}(G(\nu_i|x, \theta)|x)$$

$$\chi_i^e = 1, \text{ if } G^\psi(\psi_i^e) < r_i^e(s)$$

3. Draw next period state: $s_1 \sim p(s'|s, \sigma_0)$.
4. Repeat step 1-3 for T periods (or until exit) and compute the discounted value of sequence s:

$$V^s = \sum_t \beta^t \pi_i(\sigma_t, s_t, \nu_{it})$$

5. Compute expected value function:

$$\hat{V}_i(s|\sigma) = \frac{1}{S}\sum_s V^s$$

2. **Estimation of the dynamic parameters:**
   - The estimation of θ is based on the following Nash equilibrium conditions:

$$V_i(s|\sigma_i, \sigma_{-i}, \theta) \geq V_i(s|\sigma_i', \sigma_{-i}, \theta), \quad \forall i, s, \sigma_i'$$

   - Let $(i, s, \sigma_i') = x \in \mathcal{X}$. The parameter estimates solve the following least-square problem:

$$Q(\theta) = \int \big(\min\{0, g(x|\theta, \sigma)\}\big)^2 dH(x)$$

     where $g(x|\theta, \sigma) = V_i(s|\sigma_i', \sigma_{-i}, \theta) - V_i(s|\sigma_i, \sigma_{-i}, \theta)$.
   - In practice, the econometrician chooses n_I tuples $X_k = (i, s, \sigma_i')$, and compute the simulated value function as if agent i would follow alternative strategy σ_i'.
   - For instance, BBL draw n_I iid normal random-variables η_k with variance 0.5, and permute the investment strategy by adding η_k to $I_i(s, \nu)$.
   - In order to ensure consistency and asy. normality of the estimator X_k must be iid draws from $H(x)$, and both n_I and S must go to infinity.

- **Advantage:**
  - Computationally simpler than the PML.
  - Especially important in the dynamic game framework where the state space is too large to inverse or compute exactly the transition probability matrix.
- **Drawbacks:**
  - The simulation approximation can be bias if S or n_I are too small.
  - The moment conditions do not impose completely the Nash conditions (i.e. the inequalities are checked only at a subset of states and for a finite number of permutations).
  - Depending on the structure of the model, we might be able to combine the PML or the LS estimator with the simulated value function approximation (e.g. compute exactly the best-response functions $\Psi_(a|x, \hat{V})$ and iterate on the procedure).

Dynamic product repositioning in differentiated markets, by Sweeting (2013)

- **Objective:** Quantify the dynamic costs of product repositioning:
  - Quality decrease,
  - Drop in advertising revenues,
  - Sunk cost.
- Empirical challenges:
  - Demand estimation: If product location (i.e. radio formats) are endogenously determined, then product characteristics are not exogenous.
  - Large dynamic programming problem (i.e. many possible choices + infinite horizon stochastic problem).
  - Complex strategic interaction between firms.
- Overview of the estimation strategy:
  1. Estimation of demand for radio stations (i.e. BLP)
  2. Reduced form estimation policy functions + state transition.
  3. Structural estimation of dynamic parameters: Forward simulation algorithm (i.e. as in BBL) + Moment inequality approach (i.e. as in PPHI).

1. **Demand specification:**

- Payoff function (time and market subscripts are omitted):

$$\begin{aligned} u_{is} &= \gamma_i^C + F_s\bar{\gamma}_{m_s}^F + F_s\gamma_i^F + X_s\gamma^x + \xi_s + \nu_{is} \\ &= \delta_s + \gamma_i^C + F_s\gamma_i^F + \nu_i \end{aligned} \tag{13}$$

Where,

$$\gamma_i^F = \gamma^D D_i + \Gamma^F \eta_i$$

- Match predicted market share (i.e. from mixed-logit model) with observed shares of listeners for each station/market/time.
- Transition of station's quality:

$$\text{Non-switchers} \quad \xi_{st} = \rho_1\xi_{st-1} + \mu_t + \nu_{1st} \tag{14}$$

$$\text{Switchers} \quad \xi_{st} = \rho_1\xi_{st-1} + \mu_t\mu_2 + \nu_{2st} \tag{15}$$

- Timing assumption:
  - (ν_{1st}, ν_{2st}) are iid overtime and across players/markets.
  - Firms observe (ν_{1st}, ν_{2st}) **after** deciding to switch format.

- Moment conditions:
  - Use the timing assumption to construct orthogonally conditions based on quasi-difference in the quality of stations (i.e. similar to production function estimation):

$$\nu_{1st} = (\delta_{st} - \rho_1 \delta_{st-1}) - (\gamma_i^C + F_s \bar{\gamma}_{m_s}^F + X_{st}\gamma^x - \rho_1 \gamma_i^C - \rho_1 F_s \bar{\gamma}_{m_s}^F - \rho_1 X_{st-1}\gamma^x)$$
$$\nu_{2st} = (\delta_{st} - \rho_2 \delta_{st-1}) - (\gamma_i^C + F_s \bar{\gamma}_{m_s}^F + X_{st}\gamma^x - \rho_2 \gamma_i^C - \rho_2 F_s \bar{\gamma}_{m_s}^F - \rho_2 X_{st-1}\gamma^x)$$

    Then $E(Z_{st}\nu_{st}(\theta)) = 0$.
  - Instruments:
    - Current and lag station characteristics,
    - Current and lag average demographics characteristics,
    - Current and lag local competition measures (i.e. number of FM stations, market coverage).
    - Lag market shares.
  - Additional "micro" moment conditions:
    - Observed share of listeners by demographic groups,
    - Total listening time by blacks and hispanics.

2. **Reduced form revenue equation:**

- Observed station revenue for each quarter and by demographic group (d).

$$R_{smdt} = \alpha_{mt}(1 + W_{smt}\alpha^W)(1 + D_d\alpha^D) + \epsilon^R_{smt} \tag{16}$$

  where W_{smt} are covariates measuring characteristics of the stations and characteristics of competitors, and D_d is a vector of demographic characteristics.

- Limitation: The determination of advertising prices is not modeled explicitly.

3. **Policy functions:**

- State space:
  - Current formats: (F_s, F_{-s})
  - Vector of qualities: (ξ_s, ξ_{-s}).
  - Distribution of demographic characteristics (ethnic/racial groups size).
  - Private information format payoff shocks (i.e. type 1 extreme value): ϵ^F_s.
- A product repositioning strategy is described by a multinomial logit probability function of the common state variables S (too large to estimate consistently). Let $\sigma_s(S)$ be a particular strategy for player s.
- The reduced form of those policy functions is estimated using a parametric multinomial logit model with linear and interaction terms.

Computational problem: The model involves a very larger number of state points, and the standard inversion procedure of A&M is not feasible.

- Two alternative solutions
  - Forward simulation (i.e. BBL)
  - Value function approximation (preferred)

Value Function approximation formulation:

- Let $E(\pi(S)|P)$ denotes the expected profit function conditional on the strategy P.
- Assume that the continuation value can be linearly approximated by K functions $\phi_k(S)$ of the state variables:

$$V(S) \approx \sum_k \gamma_k \phi_k(S)$$

- Select N states: $S_1, S_2, ..., S_N$.
- Stacking the N value function in a vector, the following equality holds at the equilibrium strategy P^*:

$$V \approx \Phi\gamma = E(\pi|P^*) + \beta E\left[\Phi\gamma|P^*\right]$$

where $E\left[\phi(S)\gamma|P^*\right] = \int \phi(S')\gamma g(S'|S, P^*)dS'$ is a particular element of the vector $\Phi\gamma$.

- If the number of functions K is smaller than the number of states N, we can find the vector of approximation parameters $\lambda(P^*)$ that satisfy this Bellman equation on average:

$$\lambda(P^*) = ((\Phi - \beta E\left[\Phi|P^*\right])^T(\Phi - \beta E\left[\Phi|P^*\right]))^{-1}(\Phi - \beta E\left[\Phi|P^*\right])^T E(\pi|P^*)$$

- This approach can be used to solve the model (for counterfactuals), AND to construct best-response choice probabilities conditional on beliefs $\hat{P}$ (for estimation).

- Algorithm steps:

  1. Calculate the polynomial basis functions for N states (i.e. data + perturbations): Φ
  2. Set starting values for P^0
  3. Calculate $E(\pi(S_i)|P^0)$ and $E(\Phi(S)|S_i, P^0)$ for all states $i = 1, .., N$.
  4. Calculate the OLS approximation: $\lambda(P^0)$
  5. Calculate the choice-specific value functions: $v(a|S, P^0) = \pi(a, S)+\beta E[\Phi(S)|P^0, S]\lambda(P^0)$
  6. Update choice-probabilities: $P^1(S_i) = \frac{\exp(v(a|S_i,P^0))}{\sum_{a'}\exp(v(a'|S_i,P^0)}$ for all states
  7. If $||P^1(S_i) - P^0(S_i)|| < \eta$ stop, otherwise update the CCP: $P^1(S_i) = \psi P^0(S_i) + (1 - \psi)P^1(S_i)$ and repeat steps 3 to 6 (i.e. ψ is a step size set to 0.1).

Estimation: Pseudo-Likelihood with value-function approximation

- To evaluate the likelihood we need to solve the following value-function approximation fixed-point:

  1. For candidate parameter θ^i and choice probability P^i, calculate $E(\pi|\theta^i, P^i)$ and $E[\Phi|P^i]$.
  2. Compute the vector of value-function approximation parameters:
  $$\lambda(P^i) = ((\Phi - \beta E\left[\Phi|P^i\right])^T(\Phi - \beta E\left[\Phi|P^i\right]))^{-1}(\Phi - \beta E\left[\Phi|P^i\right])^T E(\pi|\theta^i, P^i)$$
  3. Calculate the continuation value for each state S and player j conditional on choosing action a:
  $$W_j(a, S|P_j^i, P_{-j}) = \sum_{k=1}^{N} \int \phi_{kj} g(dS_{t+1}|S_t = S, P_j^i, P_{-j})\lambda_{kj}^i$$
  4. Calculate the best-response choice-probability:
  $$\Psi_j(a|S, P_j^i, P_{-j}) = \frac{\exp(E[\Phi|P^i] - C(a|\theta^i) + W_j(a, S|P_j^i, P_{-j})}{\sum_{a'} \exp(E[\Phi|P^i] - C(a'|\theta^i) + W_j(a', S|P_j^i, P_{-j})}$$
  5. Pseudo-likelihood for step i:
  $$\theta^i = \arg\max_{\theta} \sum_{j,t} \log \Psi_j(a_{jt}|S_{jt}, P_j^i, P_{-j})$$
  6. Convergence: Repeat steps 1-4 until $|\theta^i - \theta^{i-1}| < \eta$

Forward simulation approximation:

- If $V_s(S|\sigma_s^*, \sigma_{-s}^*)$ is the expected value function (i.e. prior to the realized of ϵ_s^F), σ^* is a Nash equilibrium if:

$$V_s(S|\sigma_s^*, \sigma_{-s}^*) \geq V_s(S|\sigma_s, \sigma_{-s}^*) \quad \forall S, \sigma_s$$

- The value function can be written as follows:

$$\begin{aligned} V_s(S|\sigma_s, \sigma_{-s}^*) &= E_{S,\sigma_s,\sigma_{-s}^*}\left[\sum_{t=0}^{\infty} \beta^t R(L(S_t, \Gamma_t), \alpha)\right] \\ &\quad -\theta_1 E_{S,\sigma_s,\sigma_{-s}^*}\left[\beta^t \sum_{t=0}^{\infty} I(F_{st} \neq F_{st+1}, F_{st+1} \neq 0)\right] \\ &\quad -\theta_2 E_{S,\sigma_s,\sigma_{-s}^*}\left[\beta^t \sum_{t=0}^{\infty} I(F_{st} \neq 0)\right] + \sigma_\epsilon E_{S,\sigma_s,\sigma_{-s}^*}\left[\sum_{t=0}^{\infty} \beta^t \epsilon_{st}^F(F_{st+1})\right] \\ &\equiv R_{S,\sigma_s,\sigma_{-s}^*} - \theta_1 S_{S,\sigma_s,\sigma_{-s}^*} - \theta_2 F_{S,\sigma_s,\sigma_{-s}^*} + \sigma_\epsilon e_{S,\sigma_s,\sigma_{-s}^*} \end{aligned}$$

- Notice that the value function is linear in the three dynamic parameters: $\theta_1, \theta_2, \sigma_\epsilon$.
- As in BBL and Hotz-Miller, the value function at a given state S and policy functions (σ_s, σ_s^*) is approximated using a forward simulation algorithm.

Moment inequality formulation:

- Follows the approach proposed by Pakes, Porter, Ho and Ishii (aka PPHI).
- A sufficient condition for Nash equilibrium can be described by:

$$E(\pi(\sigma_s^*, \sigma_{-s}^*, x) - \pi(\sigma_s^*, \sigma_{-s}^*, x)|I_s) \geq 0$$

Or using an econometric model for $\pi() = r() + \nu_1 + \nu_2$:

$$E(r(\sigma_s^*, \sigma_{-s}^*, x) - r(\sigma_s^*, \sigma_{-s}^*, x)|z_s) \geq 0$$

for any $z_s \in I_s$.

- In this set-up $r()$ corresponds to the simulated approximation of the expected value function, and ν_1 corresponds to measurement and simulation errors.
- The second residual ν_2 is ignored in the analysis (i.e. the econometrician uses the same information as players in the market).
- PPHI show that one can use these necessary conditions to estimate the parameters θ. The empirical analog of the previous expectation becomes:

$$\frac{1}{S}\sum_s \left(r(\sigma_s^*, \sigma_{-s}^*, x) - r(\sigma_s^*, \sigma_{-s}^*, x) \otimes h(z_s)\right) \geq 0,$$

where $h(z)$ is a non-negative function of the instruments z_s and S is the number of markets.

- To apply this method, we need to construct possible deviations from the observed strategy (i.e. one deviation = one moment inequality).
- In this context Sweeting uses four deviations from the estimated choice probabilities:
  1. Deviation that increases (decrease) the expected revenue and increase (decrease) the expected switching cost (lower (upper) bound on θ_1):
     Decrease/Increase the probability of non-switching from current format by 5%, and scale up the other probabilities for all state.
  2. Deviation that reduces the amount of switching and reduces future revenue gives and upper bound on θ_2
     Set switching probabilities to zero for all states
  3. Change conditional choice probabilities such that the expected value of ϵ^F is increased. In the multinomial case this is true if we equalize the choice probabilities. This provides an upper bound on σ_ϵ.
     Leave switching probability constant and equalized the other format probabilities.
- These four deviations plus restrictions that $\theta_2 > 0$ and $\sigma_\epsilon > 0$ identify the upper/lower bounds for all three parameters.
- Estimate these bounds for six types of markets (i.e. by size).

Table 3: Estimates of Format Taste Parameters

	Mean Tastes	γ^{σ}	Age 25-49	Age 50 plus	Female	Black	Hispanic
Radio Listening	-15.8753	16.7749	2.1404	7.7839	2.8197	3.6171	2.8563
	(0.9224)	(1.3946)	(0.2046)	(0.7247)	(0.2098)	(0.3303)	(0.4419)
Format Interactions (AC/CHR excluded)							
Rock	0.6575	-	0.2813	-0.4054	-1.2581	-1.8814	-0.8157
	(0.0580)		(0.0087)	(0.0158)	(0.0043)	(0.0292)	(0.0402)
Country	-0.1187	-	0.4847	1.3039	-0.4428	-1.9726	-1.2733
	(0.0646)		(0.0095)	(0.0175)	(0.0046)	(0.0330)	(0.0524)
Urban	-1.2040	-	-0.7373	-1.0188	-0.3996	3.9375	0.6158
	(0.0832)		(0.0393)	(0.0497)	(0.0086)	(0.0406)	(0.0501)
News/Talk	-1.2918	-	1.6979	3.1485	-1.1171	-0.7935	-1.1071
	(0.1443)		(0.0080)	(0.0144)	(0.0049)	(0.0275)	(0.0385)
Other Programming	-0.9883	-	1.0958	2.4600	-0.5384	-0.4204	-0.2916
	(0.0650)		(0.0079)	(0.0150)	(0.0049)	(0.0275)	(0.0385)
Spanish	-2.7945	-	1.0300	1.1111	-0.3649	-0.5138	3.9489
	(0.1955)		(0.0264)	(0.0506)	(0.0163)	(0.1519)	(0.1694)

Notes: 16,481 observations, GMM objective function 2.90e-12, standard errors in parentheses.
Mean tastes will reflect valuations of a white male aged 12-24.

Table 6: Parameter Estimates from the Dynamic Model

Specification	(1) Modified Procedure P-Likelihood	(2) Modified Procedure Moments	(3) Iterated Procedure P-Likelihood		(1) Modified Procedure P-Likelihood	(2) Modified Procedure Moments	(3) Iterated Procedure P-Likelihood
COSTS OF MOVE TO ACTIVE FORMAT ($ M.)				COST OF MOVING FROM ACTIVE TO DARK ($ M.)			
Markets with pop. 1 m. +	2.524	1.897	1.467	Markets with pop. 1 m +	2.704	3.126	1.572
	(0.380)	(0.441)	(0.409)		(0.744)	(0.964)	(0.503)
* Recent Format Switch	-0.001	-0.001	-0.000	Markets with pop. 0.25- 1 m.	0.636	1.636	0.369
	(0.098)	(0.110)	(0.065)		(0.230)	(0.582)	(0.130)
Markets with pop. 0.25-1 m.	0.669	0.446	0.388	Markets with pop. < 0.25 m.	0.199	0.654	0.115
	(0.157)	(0.167)	(0.102)		(0.054)	(0.322)	(0.022)
* Recent Format Switch	0.077	-0.213	-0.044	Revenue of Switching	1.108	3.105	0.642
	(0.082)	(0.145)	(0.067)	Station	(1.330)	(1.506)	(1.003)
Markets with pop. < 0.25 m.	0.233	0.069	0.135	ECONOMIES OF SCOPE			
	(0.099)	(0.048)	(0.063)	Markets with pop. 1 m. +	0.134	-0.046	0.078
* Recent Format Switch	0.030	-0.0310	-0.017		(0.099)	(0.072)	(0.075)
	(0.023)	(0.105)	(0.054)	Markets with pop. 0.25-1 m.	0.026	0.102	0.015
*Revenue of Switching	0.034	0.116	0.022		(0.018)	(0.021)	(0.010)
Station (all markets)	(0.168)	(0.211)	(0.118)	Markets with pop. < 0.25 m.	-0.006	-0.014	-0.033
					(0.012)	(0.010)	(0.029)
ADDITIONAL COST OF MOVE TO ACTIVE FROM DARK ($ M.)				SCALE OF εs			
Markets with pop. 1 m +	-0.501	-0.445	-0.291	Markets with pop. 1 m. +	0.517	0.408	0.300
	(0.220)	(0.183)	(0.095)		(0.180)	(0.063)	(0.087)
Markets with pop. 0.25- 1 m.	-0.255	-0.495	-0.147	Markets with pop. 0.25-1 m.	0.144	0.091	0.083
	(0.166)	(0.098)	(0.069)		(0.055)	(0.037)	(0.026)
Markets with pop. < 0.25 m.	-0.061	-0.078	-0.035	Markets with pop. < 0.25 m.	0.050	0.014	0.028
	(0.091)	(0.032)	(0.042)		(0.014)	(0.007)	(0.009)

Table 7: Estimates Based on Forward Simulation

	Repositioning Cost	Scope Economy	Scale of Payoff Shock
MARKETS WITH POP. >1 m.			
Moment inequality bounds (PPHI)	[2.194,11.031]	[-0.081,0.056]	[0.074,1.729]
95% CIs	[0.652,13.877]	[-0.126,0.102]	[0.011,1.793]
BBL point estimate	18.668	0.337	3.771
std. error	(1.765)	(0.048)	(0.342)
proportion of inequalities violated	24.4%	-	-
MARKETS WITH POP. 0.25m. -1m.			
Moment inequality bounds (PPHI)	[0.464,3.421]	[-0.071,0.031]	[0.015,0.549]
95% CI	[0.232,4.035]	[-0.082,0.042]	[0,0.568]
BBL point estimate	3.046	0.013	0.630
std. error	(0.190)	(0.006)	(0.043)
proportion of inequalities violated	10.0%	-	-
MARKETS WITH POP. <0.25m.			
Moment inequality bounds (PPHI)	[0.230,1.541]	[-0.022,0.008]	[0.005,0.251]
95% CI	[0.081,1.690]	[-0.27,0.014]	[0,0.258]
BBL point estimate	2.148	0.011	0.455
std. error	(0.242)	(0.004)	(0.051)
% of BBL inequalities violated	22.4%	-	-

Demand fluctuations in the ready-mix concrete industry, Collard-Wexler (2013)

Motivation:

- When firms incur sunk entry costs, uncertainty in demand conditions create a barrier to entry (i.e. there is an option value of waiting to enter).
- This implies that business-cycle fluctuations can have real effects on the structure of markets, and therefore on welfare.
- **Case study:** Ready-mix concrete industry.
- Why? Homogenous products + spatial differentiation = Large number of isolated markets with easily defined market structure.
- What do we need?
  - Estimate of entry and investment costs
  - Measure of market power
  - Census of manufacturer: Over 15,000 plants observed annually from 1976 to 1999.
- Methodology innovation: Extends the A&M methodology to accommodate unobserved market heterogeneity and alleviate the computational cost associated with the inversion of choice-prabilities.

Assumptions and characteristics of the market:

- Homogenous products
- County = Market
- Demand shock = Employment in construction sector
- Governments as consumers (50%): Increase idiosyncratic fluctuations in demand.
- Concentrated local markets: Nearly 80% of counties have 2 or less plants.

Data:

- Longitudinal Business Database: IRS data on business tax records
  - Key variables: Employment, salary, firm identification, entry/exit dates
  - Plant size: Employment category (small, medium, large).
- Census of Manufacturers and Annual Survey of Manufacturers:
  - Key variables: inputs, outputs, and assets
- Period: 1976-1999
  - Stationary industry: About 350 plants exit/enter + 5000 continuing plants
  - Continuing plants are twice as large as either entrants or exitors

Prices and market structure:

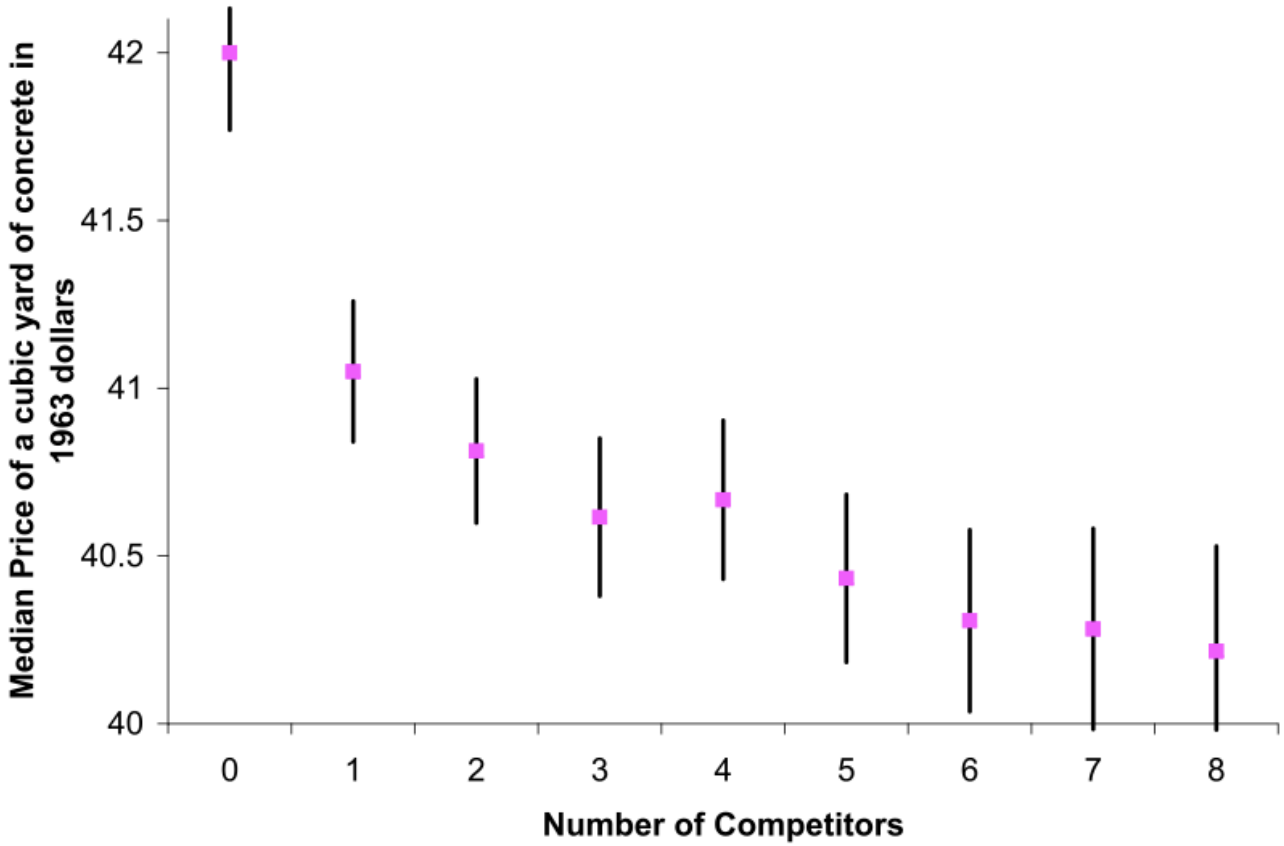


FIGURE 2.—Price declines with the addition of the first competitors, but drops by very little thereafter. Bars represent 95% confidence interval on median price.

Model:

- N = 10: Number of firms in each market (i.e. potential entrants or incumbents)
- Industry state space: $s^t = (s_1^t, s_2^t, ..., s_N^t, M^t)$, where $s_j^t = (x_j^t, \epsilon_j^t) = (\text{Size}_j^t, \epsilon_j^t)$ and ϵ_i^t is a vector of IID extreme-value random variable. Let $x^t = \{x_1^t, x_2^t, ..., x_N^t, M^t\}$ denote the public information state vector.
- Actions: $a_i^t \in \{\emptyset, \text{Small}, \text{Medium}, \text{Large}\}$.
- Payoffs:

$$\Pi_i(x^t, \epsilon_i^t, a_i^t) = r_i(x^{t+1}) + \tau(a_i^t, x_i^t) + \epsilon_{i,a_i^t}^t$$

  Where $r_i(x^t) = r(x_{-i}^t, x_i^t) = \theta_1(x_t^i) + \theta_2(x_i^t)M^t + \theta_3^t g(\sum_{j \neq i} 1(x_j^t \neq \emptyset))$, and $\tau(a_i^t, x_i^t)$ is an adjustment cost equal to zero if $a_i^t = x_i^t$ or if $a_i^t = \emptyset$.

- Timing:
  1. Private information shocks are realized
  2. Firms simultaneously choose actions
  3. Demand shock is realized
  4. Payoffs are realized

- Ex-ante value function:

$$V_i(x^t) = E_{\epsilon_i^t}\left(\max_{a_i} E_{x^{t+1}}\left[r(x_{-i}^{t+1}, x_i^{t+1} = a_i) + \tau(x_i^{t+1} = a_i, x_i^t) + \epsilon_{i,a_i}^t + \beta V(x^{t+1})\middle| x^t, a_i\right]\right)$$

- Choice-specific value function:

$$v_i(a_i, x^t) = E_{x^{t+1}}\left[r(x_{-i}^{t+1}, x_i^{t+1} = a_i) + \tau(x_i^{t+1} = a_i, x_i^t) + \beta V(x^{t+1})\middle| x^t, a_i\right]$$

- Best-response choice-probability:

$$\Psi(a_i|x^t) = \frac{\exp(v_i(a_i, x^t))}{\sum_{a_i'} \exp(v_i(a_i', x^t))}$$

Estimation method:

- Challenges:
  - Multiplicity of Nash equilibrium: Follow the insight of A&M and BBL and assume that the same equilibrium is played over time and across markets.
  - Large state space (over 350,000 points): Extends the simulation-based algorithm of Hotz, Miller, Sanders, and Smith (1994)
  - New second-stage estimation algorithm based on indirect-inference (rather than PML or moment inequalities).

Estimation algorithm

1. Initial steps:
   - Market-size transition probability matrix (10 bins): $\hat{D}(M^{t+1}|M^t)$
   - Conditional choice probabilities: $\hat{P}(a_i|x^t) \approx \Psi(a_i|x^t)$
2. Compute choice-specific value function up to parameter vector, conditional on $\hat{P}(a_i|x^t)$ and $\hat{D}(M^{t+1}|M^t)$.
   Recall that:

$$\begin{aligned} v(a_i|x^0) &= E\left[\sum_{t=0}^{\infty} \beta^t \left(r(x_{-i}^{t+1}, x_i^{t+1} = a_i^t|\theta) - \tau(x_i^{t+1} = a_i^t, x_i^t)\right) \Big| a_i, x^0\right] \\ &= E\left[\sum_{t=0}^{\infty} \beta^t B(x_{-i}^{t+1}, x_i^{t+1} = a_i^t, x_i^t) \Big| a_i x^0\right] \theta = \Gamma(a_i, x^0)\theta \end{aligned}$$

   where $B(x_{-i}^{t+1}, x_i^{t+1} = a_i^t)$ is a function of the state variables, and θ is a vector of parameters.

   Following Hotz et al. (1994), $\Gamma(a_i, x^0)$ can be computed using a forward simulation algorithm: Sample a sequence of actions for each players, according to the reduced-form policy function $\hat{P}(a_i|x_i^t)$ and aggregate transition function $\hat{D}(M^{t+1}|M^t)$. Let $\hat{\Gamma}(a_i, x)$ denote this matrix.

3. Evaluate best-response function:

$$\Psi(a_i^t|x^t,\hat{P}) = \frac{\exp\left(\hat{\Gamma}(a_i,x^t)\theta\right)}{\sum_{a_i'}\exp\left(\hat{\Gamma}(a_i',x^t)\theta\right)}$$

4. Second-stage estimation routine: Indirect inference. Indirect inference is a GMM estimation procedure in which moments are constructed by matching the parameters of reduced-form auxiliary regression models.

   In this case, the auxiliary equations are coming from a linear multinomial probability model. The predicted and observed outcome variables are:

$$\mathbf{y}_n = \begin{bmatrix} 1(a_i^t = \text{small}) \\ 1(a_i^t = \text{medium}) \\ 1(a_i^t = \text{large}) \end{bmatrix}, \quad \tilde{\mathbf{y}}_n = \begin{bmatrix} \Psi(a_i^t = \text{small}|x^t,\hat{P}) \\ \Psi(a_i^t = \text{medium}|x^t,\hat{P}) \\ \Psi(a_i^t = \text{large}|x^t,\hat{P}) \end{bmatrix}$$

   For each outcomes (predicted/observed), the "moments" are obtained by OLS:

$$y_{it} = Z_{it}\beta + u_{it}, \quad \tilde{y}_{it} = Z_{it}\beta(\theta) + u_{it}$$

   where Z_{it} includes: indicators for the firm's current state, the number of competitors in a market, and the log of construction employment in the county. The GMM objective function is constructed as follows:

$$J(\theta) = \left(\hat{\beta} - \hat{\beta}(\theta)\right)^T \mathbf{W} \left(\hat{\beta} - \hat{\beta}(\theta)\right)$$

First-stage CCP and unobserved heterogeneity:

- Problem: The market state vector x^t does not measure all time-persistent state variables. This is a problem, since the model assumes that all the randomness is caused by IID extreme-value shocks.
- If these omitted variables are positively corrected with the number of competitors, this will cause an upward bias in the reduced-form competitive effects, and therefore under-estimate the effect of competition on profits.
- Market-fixed effects are not feasible: Too many parameters + impossible to exploit cross-sectional differences to identify the model.
- Solution: Add a time-invariant control variable that proxies for the market fixed-effects. In this case, the average number of competitors is used as a control.
- $\hat{P}(a_i^t|x^t)$ is estimated via a multinomial logit model.
- The control variable is added to the profit function: $r(x_{-i}^{t+1}, x_i^{t+1} = a_i, \mu)$.

TABLE V

MULTINOMIAL LOGIT ON THE CHOICE TO BE LARGE, MEDIUM, OR SMALL

Dependent Variable	Independent Variable	I		II (Market Category)	
		Coeff.	S.E.	Coeff.	S.E.
Small in $t+1$	Small	6.59	(0.03)	6.42	(0.03)
	Small, Medium in Past	6.45	(0.04)	6.18	(0.04)
	Small, Large in Past	5.94	(0.06)	5.72	(0.06)
	Medium	6.02	(0.05)	5.81	(0.05)
	Medium, Large in Past	5.41	(0.08)	5.16	(0.08)
	Large	4.58	(0.06)	4.37	(0.06)
	Log County Employment	0.13	(0.01)	−0.06	(0.01)
	First Competitor	−1.42	(0.04)	−1.71	(0.04)
	Second Competitor	0.10	(0.03)	−0.46	(0.03)
	Third Competitor	0.16	(0.04)	−0.26	(0.04)
	Log of Competitors Above 3	0.11	(0.03)	−0.04	(0.03)
	Market Category μ			X	
	Constant	−3.94	(0.06)	−3.17	(0.06)
Medium in $t+1$	Small	6.25	(0.05)	6.08	(0.05)
	Small, Medium in Past	6.96	(0.06)	6.70	(0.06)
	Small, Large in Past	6.43	(0.08)	6.22	(0.08)
	Medium	9.16	(0.06)	8.96	(0.06)
	Medium, Large in Past	9.08	(0.08)	8.83	(0.08)
	Large	7.44	(0.07)	7.23	(0.07)
	Log County Employment	0.29	(0.01)	0.12	(0.01)
	First Competitor	−1.54	(0.05)	−1.87	(0.05)
	Second Competitor	0.00	(0.04)	−0.53	(0.04)
	Third Competitor	0.06	(0.05)	−0.32	(0.05)
	Log of Competitors Above 3	0.02	(0.03)	−0.11	(0.03)
	Market Category μ			X	
	Constant	−6.72	(0.08)	−5.99	(0.09)
Large in $t+1$	Small	5.04	(0.08)	4.88	(0.08)
	Small, Medium in Past	4.53	(0.13)	4.28	(0.13)
	Small, Large in Past	5.78	(0.11)	5.58	(0.11)
	Medium	7.46	(0.08)	7.27	(0.08)
	Medium, Large in Past	8.37	(0.09)	8.13	(0.09)
	Large	9.76	(0.07)	9.56	(0.07)
	Log County Employment	0.52	(0.01)	0.34	(0.02)
	First Competitor	−1.61	(0.05)	−1.94	(0.06)
	Second Competitor	−0.03	(0.05)	−0.58	(0.05)
	Third Competitor	−0.02	(0.06)	−0.42	(0.06)
	Log of Competitors Above 3	−0.04	(0.04)	−0.17	(0.04)
	Market Category μ			X	
	Constant	−8.58	(0.11)	−7.83	(0.11)
	Observations	409,850		409,850	
	Log-Likelihood	84,855		83,814	
	Likelihood Ratio	400,760		402,841	

TABLE VI

ESTIMATES FOR THE DYNAMIC MODEL OF ENTRY, EXIT, AND INVESTMENT[a]

		Coeff.	S.E.*
Fixed Cost	Small	−139	(6)
	Medium	−244	(10)
	Large	−285	(6)
Log Construction	Small	20	(1)
Employment	Medium	35	(2)
	Large	45	(1)
1st Competitor	Small	−48	(4)
	Medium	−58	(5)
	Large	−63	(6)
Log Competitors	Small	−17	(3)
(Above 1)	Medium	−44	(4)
	Large	−48	(3)
Transition Costs			
Out → Small		−1002	(11)
Out → Medium[†]		*−2000*	(107)
Out → Large		−1771	(53)
Small → Medium		−332	(7)
Small, Past Medium → Medium		−772	(32)
Small, Past Large → Medium		−325	(8)
Small → Large		−1809	(73)
Small, Past Medium → Large		−608	(19)
Small, Past Large → Large		−343	(16)
Medium → Small		−107	(6)
Medium, Past Large → Small		−314	(6)
Medium → Large		101	(14)
Medium, Past Large → Large		−43	(7)
Large → Small		−254	(7)
Large → Medium		−403	(6)
Standard Deviation of Shock		133	

[a]All estimates in thousands of dollars. [†] The entry costs of a medium-sized plant are calibrated to $ 2 million. * Standard Errors are computed using 100 block bootstrap replications, where I reestimate the demand transition process $\hat{D}^{\mu}$ and the conditional choice probabilities $\hat{P}^{\mu}$, then minimize the criterion function Q to find $\hat{\theta}$. I block bootstrap by market, resampling a market's history from 1976 to 1999, so the computed standard errors account for serial correlation within a market.

Counter-factual experiments: What is the impact of eliminating/reducing demand uncertainty?

- Consider two counter-factual environments:
  - Constant demand
  - Demand is constant for five years
- three effects:
  - Direct effect: Eliminating unanticipated shocks to demand recede industry turnover.
  - Option value: Uncertainty creates an option value of waiting, and therefore reduces the entry probability, and the number of firms active.
  - Market expansion: If profits are concave in demand, expected profits with uncertainty are **lower** than profits without uncertainty (Jensen's inequality).
- Solution procedure: Stochastic algorithm (Pakes and McGuire 2001)

TABLE VIII

DEMAND SMOOTHING, TURNOVER, AND SIZE CHANGING

	Unsmoothed Demand ($\hat{D}^{\mu}$)	5 Years of Smoothing	Constant Demand	Firms Believe Demand is Constant
Turnover				
Entry Rate	2.7%	2.2%	2.2%	4.1%
Exit Rate	2.9%	2.0%	2.1%	4.5%
Change in Size Rate	20%	18%	17%	18%
Investment				
Sunk Entry Costs per Year (in Million $)	132	137	107	155
Size Changing Costs per Year (in Million $)	307	496	407	337
Total Plants	3643	5433	4264	3879

TABLE IX

DEMAND SMOOTHING AND INDUSTRY COMPOSITION

	Unsmoothed Demand	Constant Demand	5 Years of Smoothing
Total Plants	3645	4264	5433
Fixed Costs (per Period in Millions of $)	717	878	1109
Industry Composition			
Small Plants	54%	48%	49%
Medium Plants	23%	23%	24%
Big Plants	23%	29%	28%
Market Structure			
Markets With no Plants	5%	8%	1%
Markets With 1 Plant	43%	36%	25%
Market With 2 Plants	28%	24%	29%
Markets With More Than 2 Plants	25%	32%	46%

TABLE X

WELFARE EFFECTS OF DEMAND-SMOOTHING POLICIES[a]

Change in Net Present Value of	
Consumer Surplus	$860 Million
Producer Surplus for Incumbents	−$609 Million
Producer Surplus for Potential Entrants	−$36 Billion

[a]Numbers in this table refer to the difference in the net present value of surplus (using a 5% discount rate) between five years of smoothing and unsmoothed demand, averaged between 25 and 50 years after the policies were put into place, using 1976 as an initial state.

Goettler and Gordon (2011): Competition and innovation in the microprocessor industry

- **Background:** A central question in IO is wether competition leads to more or less innovation.
  - Schumpeter: *Creative destruction* effects suggest that monopolies can be more innovative
  - Arrow: Perfect competition leads to optimal allocation of resources under uncertainty (i.e. optimal investment)
- Ultimately this is an empirical question: market imperfections can lead to *dynamic inefficiencies* with ambiguous effects.

- Empirical literature: Inverted-U relationship?

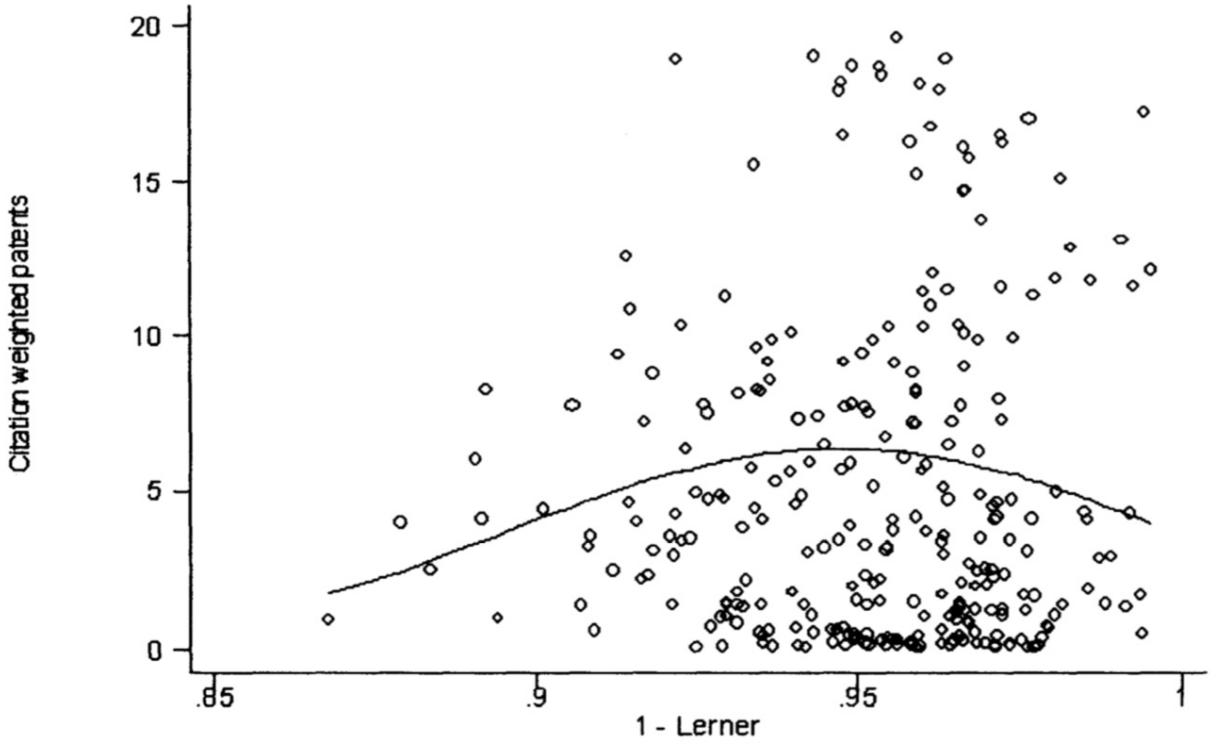
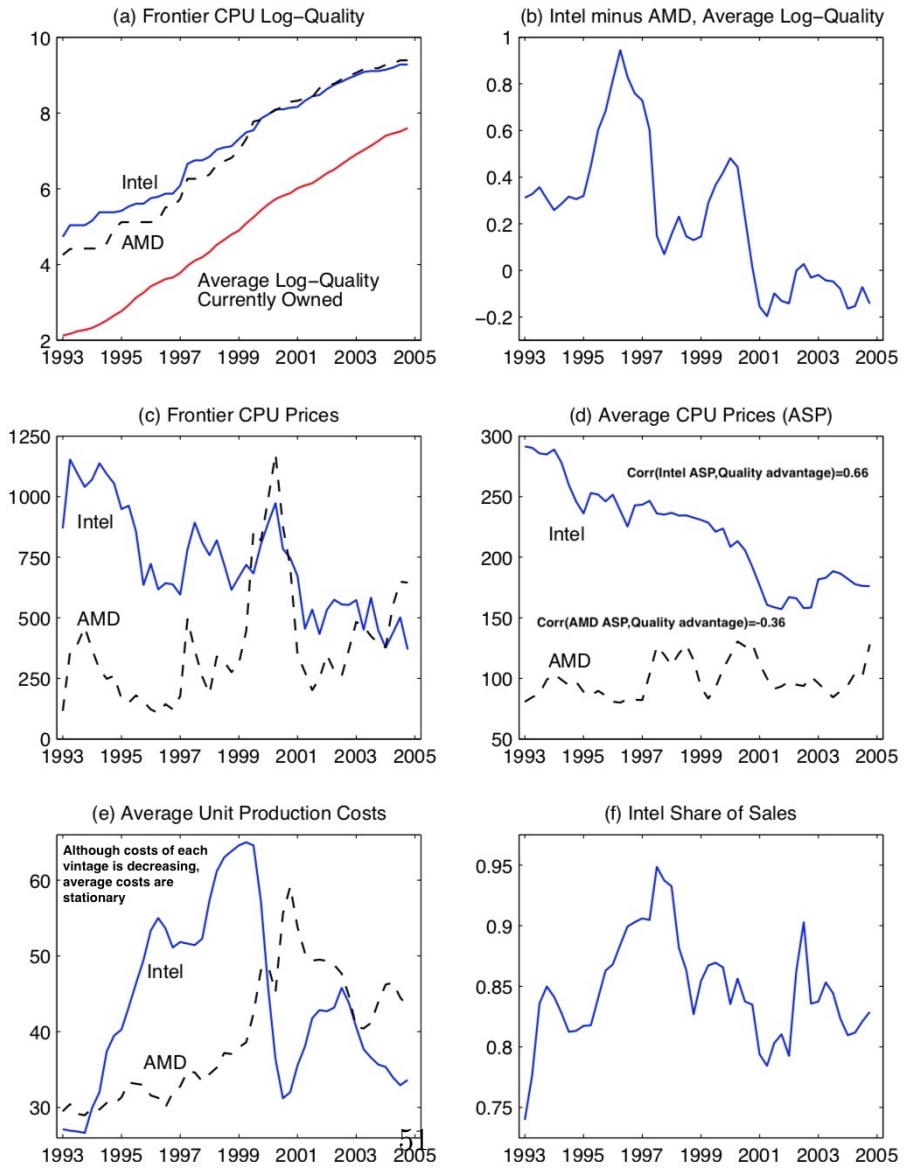


FIGURE I
Scatter Plot of Innovation on Competition

- Source: Aghion, Bloom, Blundell, Griffith, and Howitt (2005)

Figure 1: CPU Qualities, Prices, Costs, and Shares: 1993 Q1 to 2004 Q4



- **Model:** Dynamic innovation game with durable goods
  - Single-product duopoly with quality differentiation
  - Investment stochastically increases quality (as in PM)
  - Quality ladder: $q_{jt} \in \{\ldots, -2\delta, \delta, 0, \delta, 2\delta, \ldots\}$
  - **Durable good:** Distribution of current ownership affect current demand through replacement decision of consumers.

$$\Delta_t = \left(\Delta_{\underline{q}_t,t}, \ldots, \Delta_{k,t}, \ldots, \Delta_{\bar{q}_t,t}\right)$$

  - **Assumption:** Consumers scrap computers that are more than δ_c "quality" levels below the frontier. That is, $\underline{q}_t = \bar{q}_t - \delta_c$
  - **Industry state space:** Current quality levels for both firms (q_{1t}, q_{2t}) and ownership distribution Δ_t.

- **Demand and ownership:**
  - Consumers choose between three options:
    * Replacing with frontier product form firm j:

$$U_{ijt} = \gamma q_{jt} - \alpha p_{jt} + \xi_j + \epsilon_{ijt} = u_j(q, \tilde{q}, \Delta) + \epsilon_{ij}$$

    * Staying with current computer:

$$U_{i0t} = \gamma \tilde{q}_{it} + \epsilon_{i0t} = u_0(q, \tilde{q}, \Delta) + \epsilon_{i0}$$

    * Consumers' choice-specific value-function (ex-ante) for the replacement problem is:

$$v_a(q, \Delta, \tilde{q}) = u_a(q, \Delta, \tilde{q}) + \beta \sum_{q', \Delta'} \bar{V}(q', \Delta', \tilde{q}_a) g(\Delta' | \Delta, q, q') h(q' | q, \Delta)$$

      where $g(\Delta' | \Delta, q, q')$ and $h(q' | q, \Delta)$ are consumers' beliefs relative to the evolution of Δ and firms' qualities (i.e. frontier). And $\bar{V}$ is:

$$\bar{V}(q, \Delta, \tilde{q}) = \log \left(\sum_a \exp(v_a(q, \Delta, \tilde{q}) \right)$$

* The conditional-choice probabilities for a consumer with holding $\tilde{q}$ is:

$$s_j(q, \Delta, \tilde{q}) = \frac{\exp(v_j(q, \Delta, \tilde{q}))}{\sum_a \exp(v_a(q, \Delta, \tilde{q}))}$$

* From this, we can compute the aggregate market shares:

$$s_{jt} = \sum_{\tilde{q}} s_j(q_t, \Delta_t, \tilde{q}) \Delta_{\tilde{q}_t, t}$$

* Finally, the market shares defines a first-order Markov process for the ownership matrix:
  - If the frontier does not advance:

$$\Delta_{k,t+1} = s_{0,k} \Delta_k + \sum_{j=1,2} s_{jt} 1(q_{jt} = q_k)$$

  - If the quality frontier moves up: the second element of Δ' is added to its first element, the third element becomes the new second element, and so on, and the new last element is initialized to zero.

- **Investment and dynamic price competition:**
  - R&D is successful (i.e. $\Delta q_t = \delta$) with probability:

$$\chi_j(\Delta q_t = \delta | x, q) = \frac{a^j(q)x}{1 + a^j(q)x}$$

    where $a_j(q) = a_{0,j} \max(1, a_1 ((\bar{q} - q_j)/\delta)^{0.5})$.
  - Profit:

$$\pi_j(p, q, \Delta) = s_j(q, \Delta)(p_j - mc_j(q_j))$$

    depends on the quality level *relative* to the frontier where $mc_j(q_j) = \lambda_0 + \lambda_1 q_j$
  - Value function:

$$W_j(q_j, q_{-j}, \Delta) = \max_{p_j, x_j} \pi_j(p, q, \Delta) - x_j + \beta \sum_{\tau_j, \tau_{-j}, \Delta'} W_j(q'_j, q'_{-j}, \Delta') \chi_j(\tau_j | x_j, q_j) \chi_{-j}(\tau_{-j} | x_{-j}, q_{-j}) g(\Delta' | \Delta, q, q')$$

  - This leads to two FOCs:

$$x_j : \quad -1 + \sum_{\tau_j \tau_{-j} \Delta'} \beta W_j(q'_j, q'_{-j}, \Delta') \chi_{-j}(\tau_{-j} | x_{-j}, q_{-j}) g(\Delta' | \Delta, q, q') \frac{\partial \chi_j(\tau_j | x_j, q_j)}{\partial x_j} = 0$$

$$p_j : \frac{\partial \pi_j(p, q, \Delta)}{\partial p_j} + \beta \sum_{\tau_j, \tau_{-j}, \Delta'} W_j(q'_j, q'_{-j}, \Delta') \chi_j(\tau_j | x_j, q_j) \chi_{-j}(\tau_{-j} | x_{-j}, q_{-j}) \frac{\partial g(\Delta' | \Delta, q, q')}{\partial p_j} = 0$$

- Computational details:
  - The investment best-response function has a closed-from solution (as in PM). This facilitates the computation.
  - In order to "select" the same equilibrium for every guess of the parameters the GG solve a finite-horizon version of the game:
    * For each T and each state, solve the finite-horizon game by solving the price and investment vector that solves the system of FOCs using backward induction (unique SPNE).
    * In the limit $T \to \infty$, the finite horizon model is equivalent to the infinite horizon game.
  - Not a trivial task since consumers and firms are foward-looking (see Appendix for details)
  - **Important:** To solve the model, it is important that the state space is bounded
    * Up to now the model is allowed to be non-stationary: Investment pushes the frontier upward constantly.
    * However, the state-space can be re-written in *difference* relative to quality frontier since consumers care only about relative quality, and the innovation probability and marginal-cost are expressed relative to the best-available quality.

- **Estimation:** Moments
  - Average prices and the coefficients (other than the constant) from regressing each firm?s price on a constant, $q_{Intel,t} - q_{AMD,t}$, and $q_{own,t} - \bar{q}_t$.
  - Coefficients from regressing Intel' share of sales on a constant and $q_{Intel,t}?q_{AMD,t}$,
  - Rate at which consumers upgrade.
  - Mean innovation rates for each firm
  - Market share difference between intel and AMD
  - Mean investment per unit revenue for each firm

Table 3: Industry Measures under Various Scenarios

	(1) AMD-Intel Duopoly	(2) Symmetric Duopoly	(3) Monopoly	(4) No Spillover Duopoly	(5) Myopic Pricing AMD-Intel	(6) Myopic Pricing Monopoly	(7) Social Planner
Industry Profits ($billions)	408	400	567	382	318	322	-267
Consumer Surplus	2978	3012	2857	3068	2800	2762	4032
CS as Share of Monopoly CS	1.042	1.054	1.000	1.074	0.980	0.967	1.411
Social Surplus	3386	3412	3424	3450	3118	3084	3765
SS as Share of Planner SS	0.929	0.906	0.940	0.916	0.828	0.819	1
Margins	3.434	2.424	5.672	3.478	2.176	2.216	0.000
Price	194.17	146.73	296.98	157.63	140.06	143.16	43.57
Frontier Innovation Rate	0.599	0.501	0.624	0.438	0.447	0.438	0.869
Industry Investment ($millions)	830	652	1672	486	456	787	6672
Mean Quality Upgrade %	261	148	410	187	175	181	97
Intel or Leader Share	0.164	0.135	0.143	0.160	0.203	0.211	0.346
AMD or Laggard Share	0.024	0.125		0.091	0.016		0.014

Reported values are based on 10,000 simulations of 300 periods each. Profits, surplus, and investments are reported in billions of dollars. Profits and surplus are discounted back to period 0. Social surplus is the sum of consumer surplus and industry profits. Symmetric duopoly uses Intel's firm-specific parameters for both firms. Under "myopic pricing" firms choose price ignoring its effect on future demand. The "no spillover" duopoly uses symmetric firms, both with Intel's parameters. The social planner sells two products, but the results are nearly identical for a single-product planner. The monopolist offers one product. Margins are computed as $(p - mc)/mc$. Price and margins are share-weighted averages. In the symmetric duopoly, both firms have Intel's ξ_j and $a_{0,j}$.

References

Aghion, P., N. Bloom, R. Blundell, R. Griffith, and P. Howitt (2005). Competition and innovation: An inverted-u relationship. *Quarterly Journal of Economics 120*(2), 701–728.

Aguirregabiria, V. and P. Mira (2007). Sequential estimation of dynamic discrete games. *Econometrica 75*, 1–53.

Bajari, P., L. Benkard, and J. Levin (2007). Estimating dynamic models of imperfect competition. *Econometrica*.

Collard-Wexler, A. (2013). Demand fluctuations and plant turnover in the ready-mix concrete industry. *Econometrica 81*(3), 1003–1037.

Goettler, R. and B. Gordon (2011). Competition and innovation in the microprocessor industry: Does amd spur intel to innovate more. *Journal of Political Economy 119*(6), 1141–1200.

Hotz, V. J., R. A. Miller, S. Sanders, and J. Smith (1994). A simulation estimator for dynamic models of discrete choice. *The Review of Economic Studies 61*(2), 265–289.

Igami, M. (2017). Estimating the innovator's dilemma: Structural analysis of creative destruction in the hard disk drive industry. *Journal of Political Economy 125*(3), 798–847.

Pakes, A. and P. McGuire (2001). Stochastic algorithms, symmetric markov perfect equilibrium, and the 'curse' of dimensionality. *Econometrica 69*(5), 1261–1281.

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 38*(2), 373–399.

Pesendorfer, M. and P. Schmidt-Dengler (2008). Asymptotic least squares estimators for dynamic games. *Review of Economic Studies 75*, 901–928.

Ryan, S. P. (2012, May). The costs of environmental regulation in a concentrated industry. *Econometrica 80*(3), 1019–1061.

Schmidt-Dengler, P. (2006, June). The timing of new technology adoption: The case of mri. mimeo, LSE.

Seim, K. (2006, 09). An empirical model of firm entry with endogenous producttype choices. *RAND Journal of Economics 37*(3), 619–640.

Sweeting, A. (2013). Dynamic product repositioning in differentiated product industries: The case of format switching in the commercial radio industry. *Econometrica 81*(5).

Takahashi, Y. (2015). Estimating a war of attrition: The case of the us movie theater industry. *American Economic Review 105*(7), 2204–2241.