Learning and Strategic Configuration for Information Goods

Jeff MacKie-Mason

Arthur W. Burks Professor of Information and Computer Science

Professor of Economics

University of Michigan

E-Newspaper Wars: NYT v. WSJ

- Where should they locate in product configuration space?
  - Which articles should they publish? What editorial niche?
  - Suppose NYT mimics WSJ on tech news…
    - Should WSJ try to win with better/more tech news, or poach on arts/leisure buyers?
- How should they locate in price configuration space?
    - Linear? Nonlinear? 3rd degree discrimination?
- Over time, success changes
  - Due to shift in competitor's strategy or shift in consumer preferences?

Large problem space

- Suppose NYT has access to N articles
  - Can offer 2^N-1 different content bundles
  - Can select 2^N-1 different prices, one for each bundle
  - 20 articles > 1,000,000 configurations
- Bundling literature confined to: components, complete bundles, or mixture of these two
- Endogenous product differentiation literature: Hotelling (1-2 dimensions, 1-few locations, usually no uncertainty)

Search Problem

- If know consumer preferences and competitor's strategy:
  - For each of 2^N-1 bundles
    - calculate optimal price
  - Choose bundle/price pairs that give max profit

Learning problem

- Don’t know everything about consumer preferences for content types and quantity
- Don’t know everything about competitors’ content selection strategies
- Content, preferences, strategies can change daily

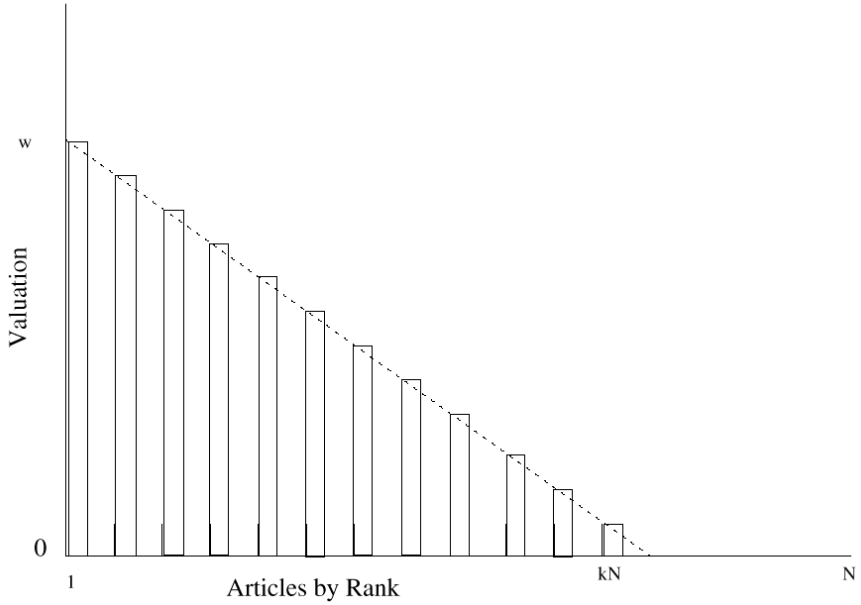
A first look: Monopolist searching price structures

- Monopolist with 10 information objects that are periodically updated (news stories)
- Large consumer population with stationary but unknown, **heterogeneous** preferences
- Machine-learning artificial agents assist buyers and sellers search space of possible prices/bundles
  - Bayesian learning too damn hard!
  - Brooks, Fay, Das, MacKie-Mason, Kephart, Durfee (1999)

Consumer preferences

- A consumer's preferences for articles described by 2 parameters:
  - w_j: value for most preferred article
  - k_j: fraction of available articles with positive value
- Different consumers value various articles differently, but rank-ordered preferences have same (linear) shape

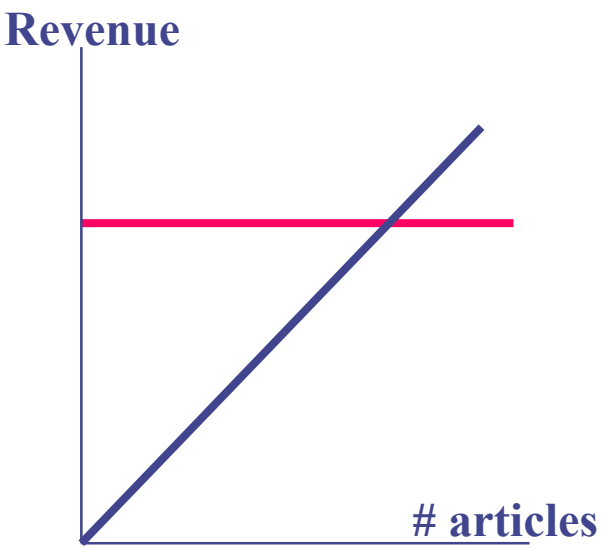
Preference map



$$V(j) = \begin{cases} w_j \left(1 - \dfrac{i-1}{k_j N}\right) & \text{if } i - 1 \le k_j N \\ 0 & \text{otherwise} \end{cases}$$

Possible Price Schedules

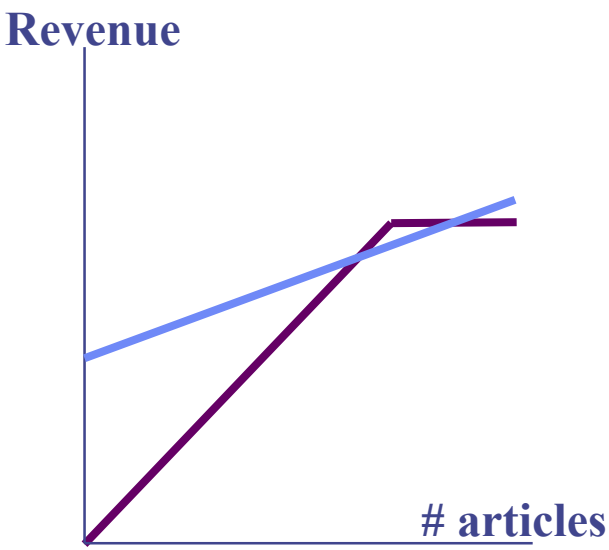
One Parameter



Linear: Each article costs p dollars

Pure Bundling: Entire collection costs b dollars

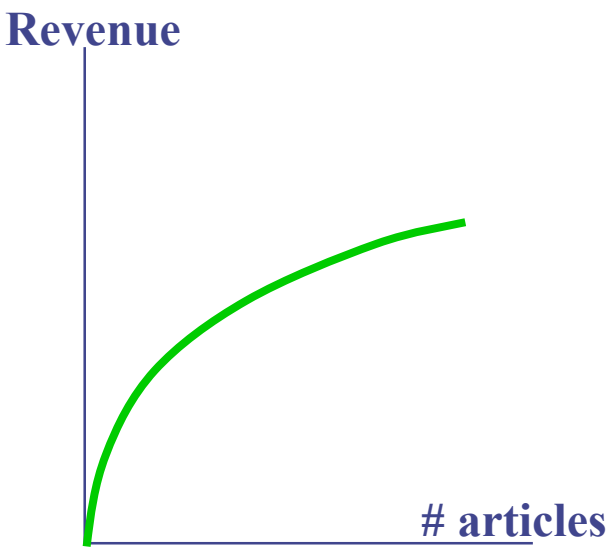
Two Parameter



Two-part tariff: Subscription fee and a per-article cost

Mixed Bundling: Bundle price or per-article price

N Parameter



Nonlinear pricing: different price for each article

Choosing price schedule with complete information

- Seller knows consumer preferences
- Solves one-time optimization
- More parameters lead to greater profit
- Greatest gain: increase from one to two parameters

Complete information possibilities

Pricing Scheme	Parameters	Profit
Pure bundling	1	0.88
Linear pricing	1	0.88
Two-part tariffs	2	1.04
Mixed bundling	2	1.04
Block pricing	3	1.09
Non-linear pricing	N	1.15

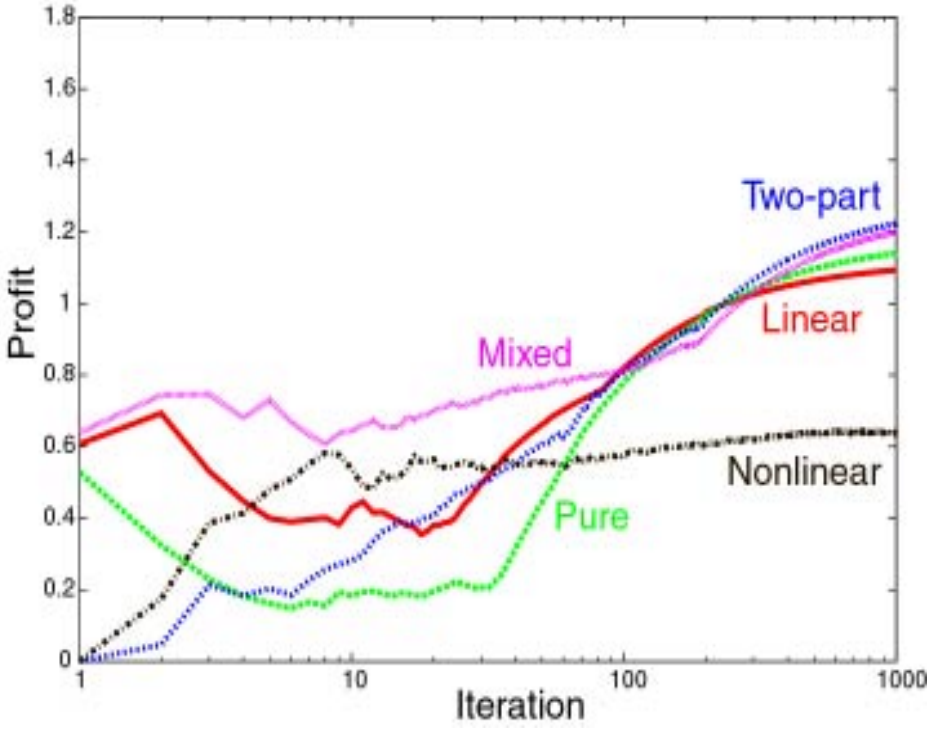
See Brooks et al. (1999) for assumptions about preferences. N=10 items.

- N.B. Profits same for same k, but welfare not

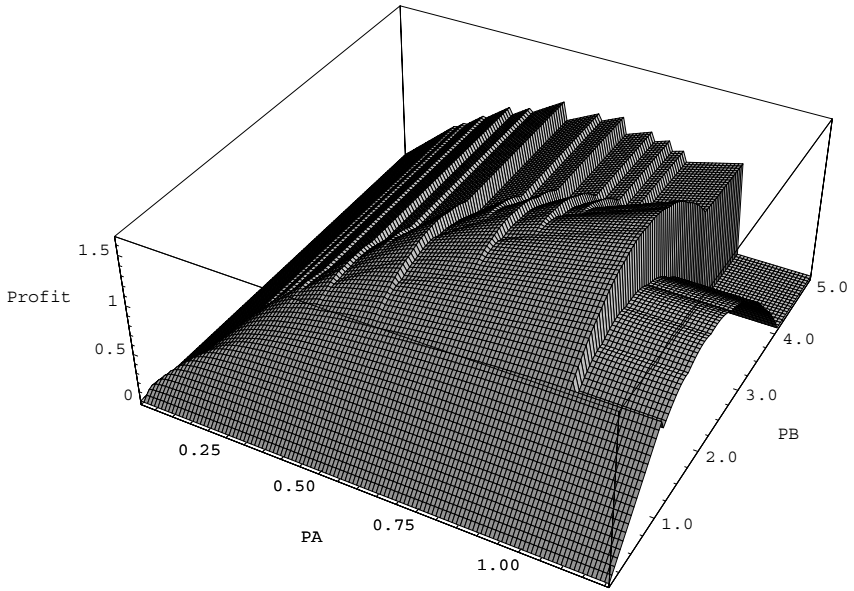
Learning price schedules: Incomplete information

- Producer needs to **learn** an approximation of **consumer preferences** as reflected in a price schedule
- N.B. Learning objective in this problem: **NOT** full information, but maximal present discounted value of profits
  - "Economically intelligent learning"
  - Producer must trade off accuracy of estimate versus time needed to learn
  - Exploration vs. Exploitation

Profits with Learning



Average profits as producer learns preferences



Profit landscape for mixed bundling

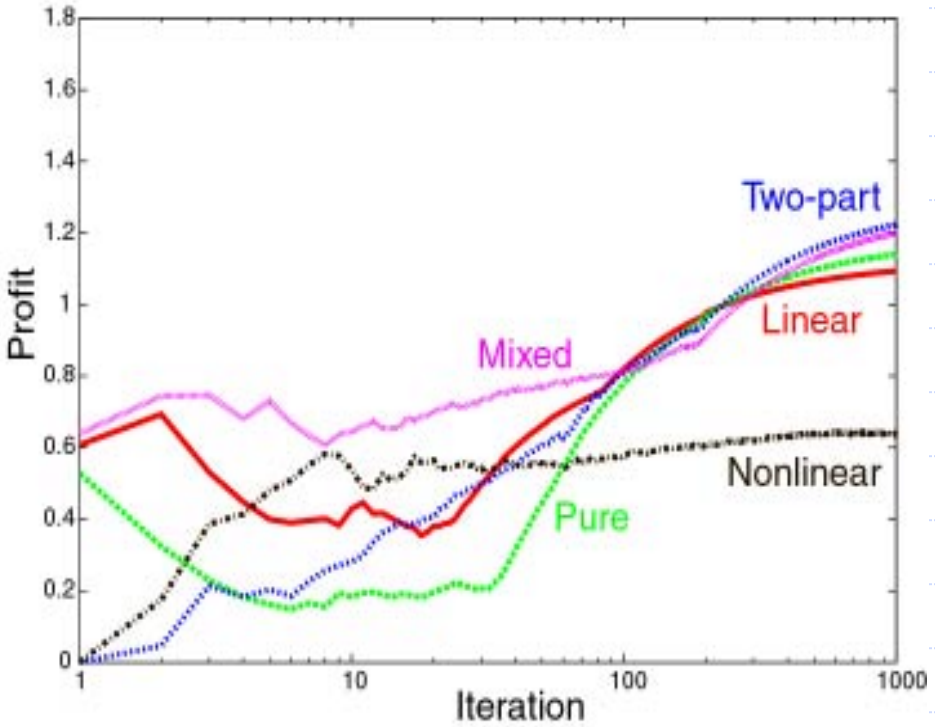
Even after 1000 transactions, nonlinear pricing has reached **less than 50%** of its profit potential.

Learning method matters!

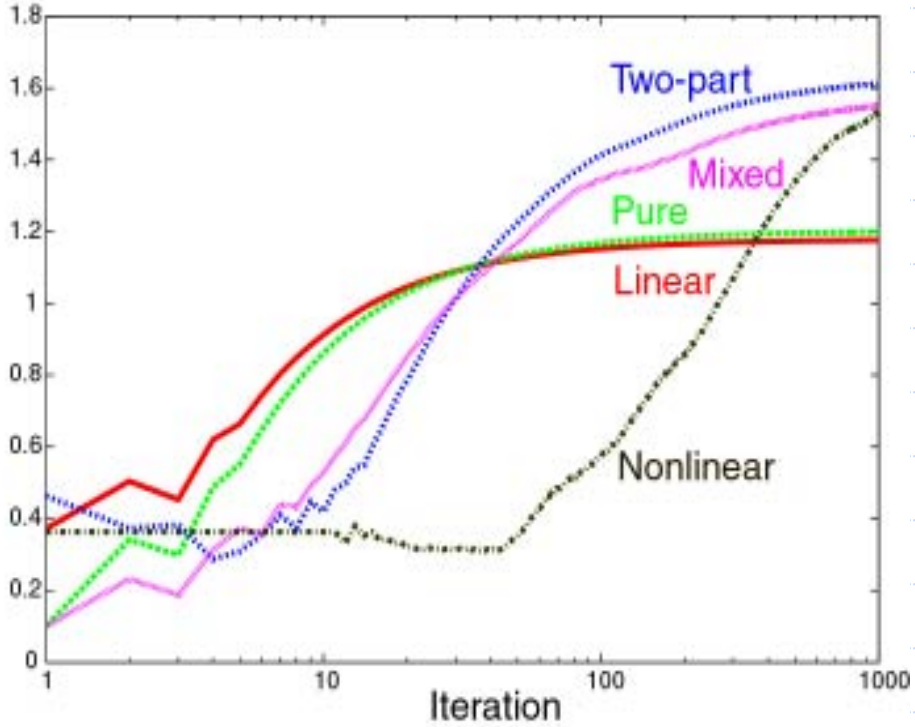
- Want to select / design a learning method that is effective (and efficient) for the problem
- Too complex to solve for optimal learning, so compared different algorithms

Different learning methods

Neural network



Amoeba



For stationary problem, hill-climbing (amoeba) outperforms function mapping (neural network)

Unification: Economically-intelligent learning agents

- We have been studying agents who:
  - Pick a price schedule
  - Use a machine learning algorithm to learn the parameters of that schedule
- But different price schedules are lower-dimensional approximations to the full nonlinear N-dimensional price schedule / profit surface
- The choice of price schedule is part of the learning problem

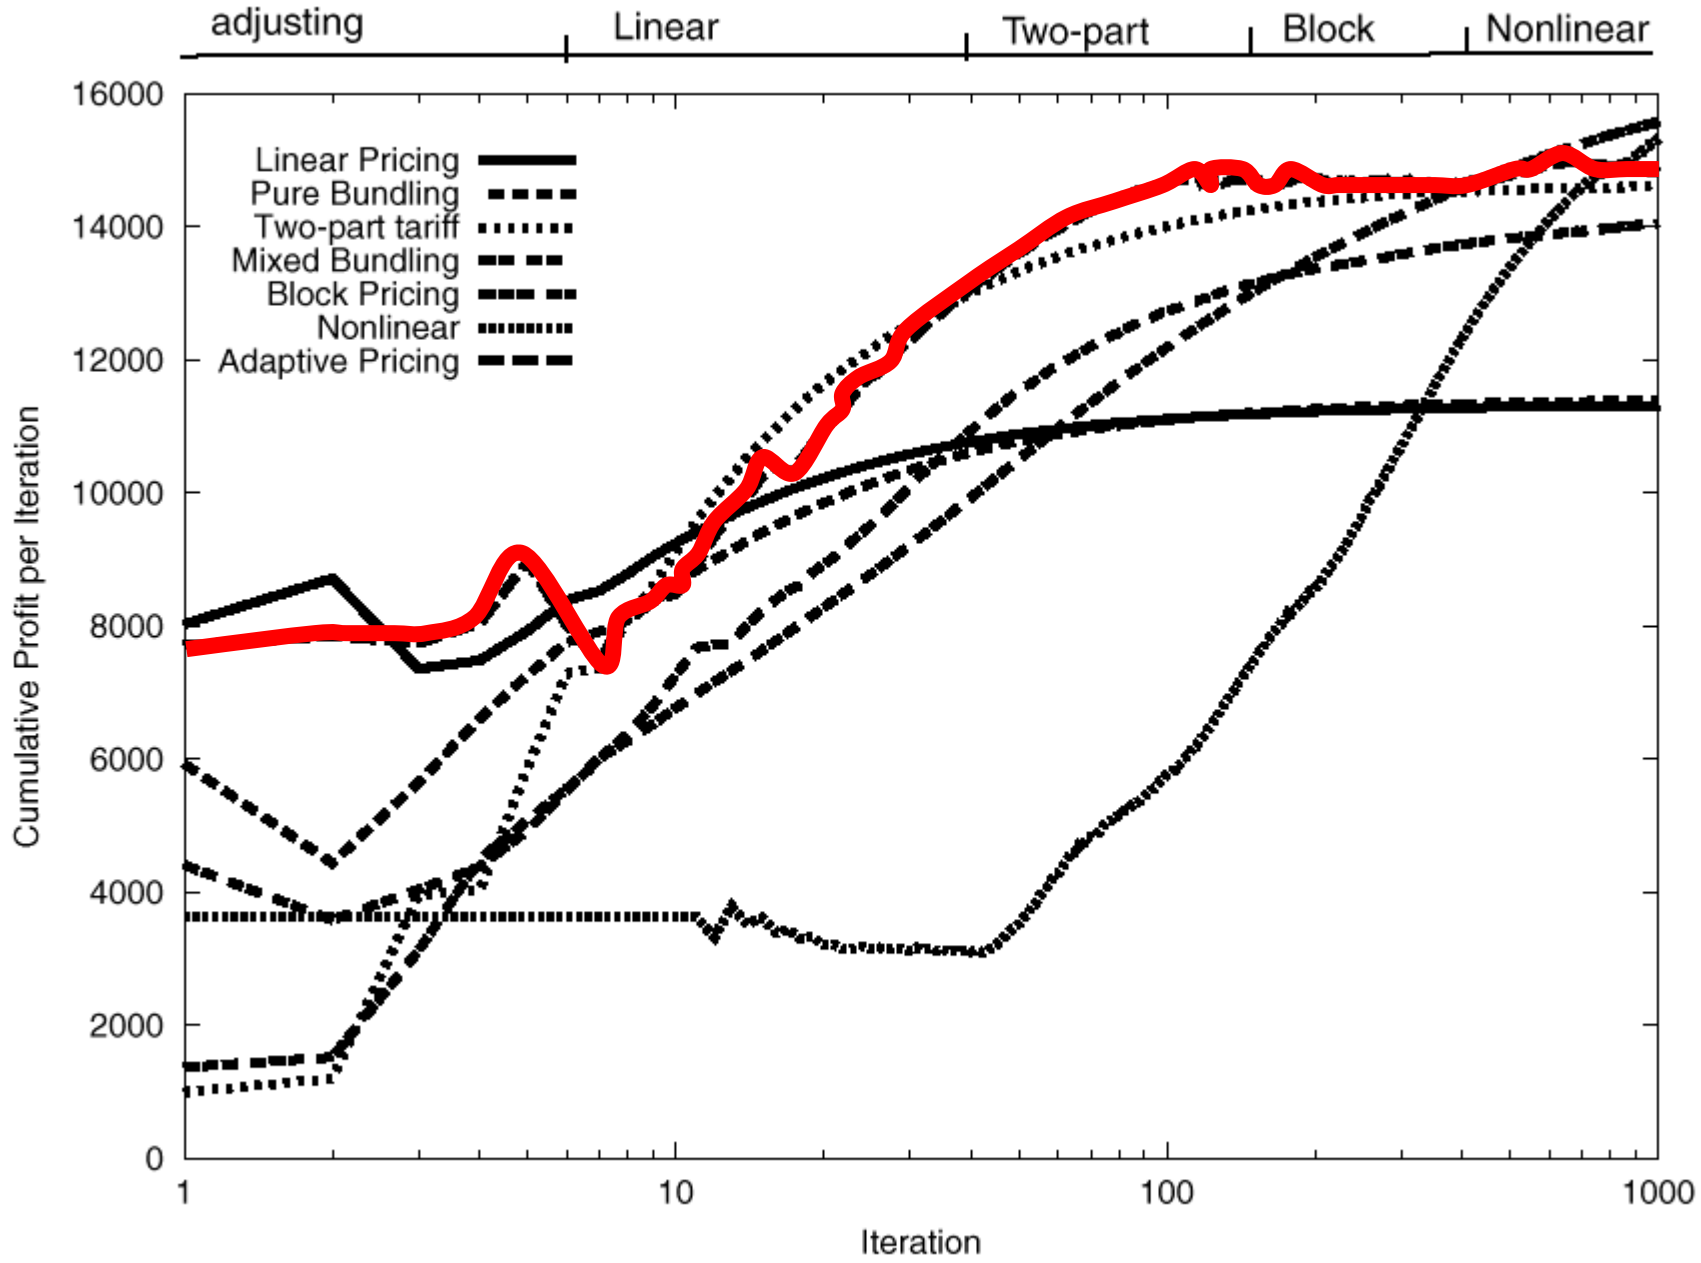
Price schedule learning

- Lower-dimension schedules yield more initial profits because all of the learning effort focused on high order bits
- But they can't extract full profit because they ignore low order bits
- Think of the price schedule / learning method choice as two parts of the learning algorithm
  - Want to focus (costly) learning effort on bits with highest incremental payoff
  - Payoffs:
    - Contribution to current profit
    - Contribution (incl. Learning effect) to expected PDV of future profits

First cut: Adaptive learning

- For a variety of *model families* (price schedules):
  - Calculate *precision*: How closely best model in family MF approximates true profit function
  - Calculate *accuracy*: How close a particular model is to best model in the same family
  - Estimate the learning rate $C(MF,n)=a_m$: after n iterations using model $M \in MF$ the accuracy of M is a_m: $$C(MF,n)=\frac{p_{MF}}{1+e^{(\log_2 n)d_{MF}}}$$
- Each period price using model with current highest expected cumulative profit

"Adaptive pricing"



More generally…

- With an N-dimensional response surface (e.g., for N articles, fully optimal nonlinear pricing has N parameters), want to find a basis Γ such that learning payoff is maximized by first learning on γ_1, then (γ_1, γ_2), then $(\gamma_1, \gamma_2, \gamma_3)$, &c.

Endogenous Product Differentiation

- What happens when there are competing firms choosing a product configuration when uncertain about consumer preferences?
- Strategic variables: (price, config)
- Tradeoff: Exploitation v. exploration
  - Gazzale and MacKie-Mason (2000)

Prior literature

- Most closely related to Harrington (JET 95) and Aghion et al. (ET 93)
- JEH:
  - Price setting only
  - Uncertainty about degree of differentiation
- Us:
  - Price & product config
  - Uncertainty about distrib of consumer tastes

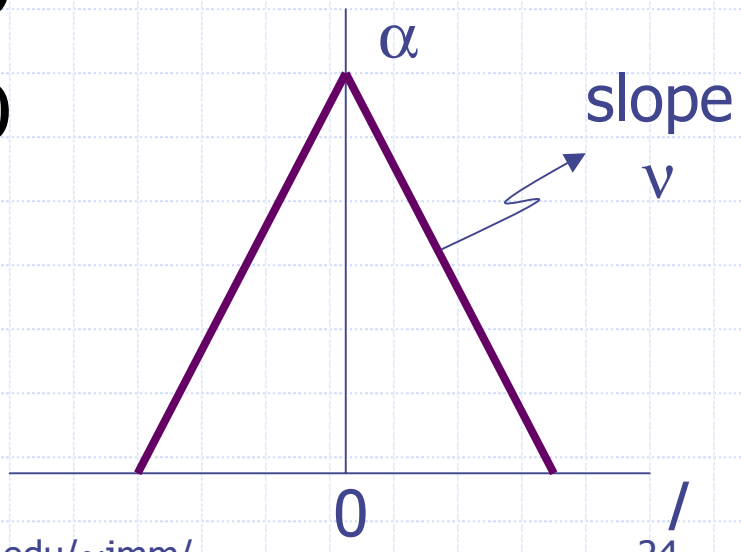
Market

- Differentiable good w/zero cost of relocation (e.g., info good)
  - One dimension
  - E.g., news Web sites: ratio of national / intl news
- Euclidean distance metric over this parameter
- Two periods of competition

Consumers

- Population of heterogeneous consumers
  - Type t_i indicates most preferred good
  - Density of tastes is tent shaped

$$f(t_i) = \begin{cases} \alpha + \nu l & if\ l \le 0 \\ \alpha - \nu l & if\ l > 0 \end{cases}$$



- α is "sweet spot"
  - Mass market and
  - Two niche markets

Consumers (cont.)

- Utility for good j is linearly decreasing in distance from t_i

  $$U_i = r - c(t_j - l_j) - p_j$$

- Hotelling-like
  - Linear distance metric
  - Linear disutility (transport cost)
- But
  - Non-uniform taste distribution
  - Prices matter, not just locations
  - Consume only if $U_i > 0$, not all consumers served

Firms

- Compete each period by simultaneously choosing (p_j, l_j)
  - No cost of moving location
  - Each firm's demand has additive stochastic component $\varepsilon_j \sim G(0, \sigma_\varepsilon)$
- Information
  - α,c,r, distn. of ε common knowledge
  - Shared prior on ν is $\hat{\mu}$
  - (p_i, l_i) and # customers served also common knowledge

Firms (cont.)

- After period 1, firms update beliefs about location of ν
- KEY: How does opportunity to learn about ν affect price and location (product configuration)?

Solution: Preliminaries

- Leftmost firm is #1, at l_1
- Marginal (outside) consumers:

$$t_l = l_1 + \frac{p_1 - r}{c}$$

$$t_r = l_2 + \frac{r - p_2}{c}$$

- Boundary between firms (assume all in-between served), let $t_m \geq 0$ wlog:

$$t_m = \frac{l_1 + l_2}{2} + \frac{p_2 - p_2}{2c}$$

Demands and profit

$$D_1(p_1, p_2, l_1, l_2) = \int_{t_l}^{0} (\alpha + \nu l)dl + \int_{0}^{t_m} (\alpha - \nu l)dl + \varepsilon_1$$

$$D_2(p_1, p_2, l_1, l_2) = \int_{t_m}^{t_r} (\alpha - \nu l)dl + \varepsilon_2$$

$$\pi_i = p_i D_i$$

Second-period equilibrium

- Firms max total expected profit
- Peak in demand -> discontinuity in best response functions (messy)
- Symmetric equilibrium

$$p_1^* = \frac{3c\alpha}{8\hat{\mu}} = p_2^*$$

$$l_1^* = \frac{r}{c} - \frac{7\alpha}{8\hat{\mu}} \qquad l_2^* = \frac{r}{c} + \frac{7\alpha}{8\mu}$$

How does expected profit depend on $\hat{\mu}$?

$$E[\pi_i \mid \hat{\mu}] = p^* \int_0^{l_2^* + \frac{r - p_2}{c}} (\alpha - \mu l)\, dl = \frac{9c\alpha^3}{64\hat{\mu}^2}$$

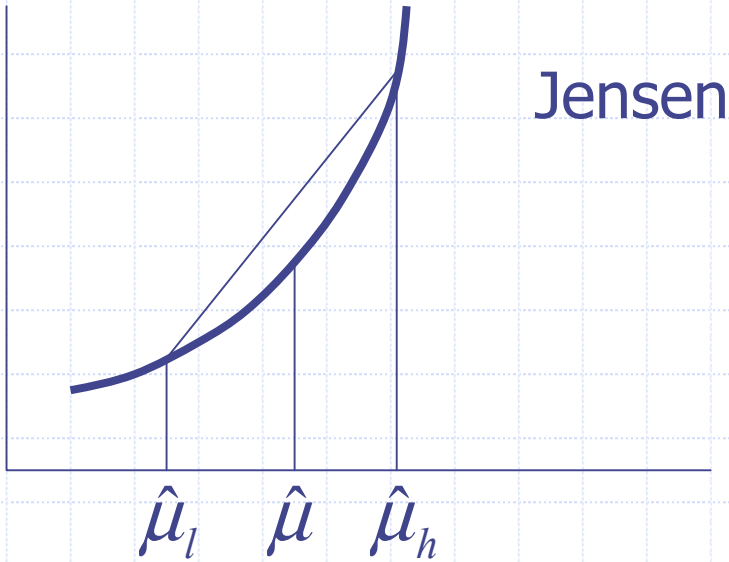
$$\frac{\partial E[\pi_i \mid \hat{\mu}]}{\partial \hat{\mu}} = -\frac{9c\alpha^3}{32\mu^3} < 0 \quad \forall \hat{\mu}$$

$$\frac{\partial^2 E[\pi_i \mid \hat{\mu}]}{\partial \hat{\mu}^2} = \frac{27c\alpha^3}{32\mu^4} > 0 \quad \forall \hat{\mu}$$

- Expected π decreasing in $\hat{\mu}$
  - Higher ν, demand falls more sharply
- Expected π convex in $\hat{\mu}$

Implication of convexity

- Firms are risk-loving in ν
  - Prefer more variability in posterior mean on ν
  - I.e., prefer to learn more about ν



Jensen's Inequality

First-Period Equilibrium

- Firms choose first-period (l_i, p_i)
- Two effects:
  - Expected period 1 profit (exploitation)
  - Learn about ν, affects period 2 profit (exploration)
- Choose (l_i, p_i) taking into account effect on posterior mean on ν

Learning

- Aggregate demand is

$$D = 2\int_0^{t_r}(\alpha - \nu l)dl + \varepsilon$$

$$= 2\alpha(l_2^* + \frac{r - p_2}{c}) - \nu(l_2^* + \frac{r - p_2}{c})^2 + \varepsilon$$

Rearrange:

$$\frac{2\alpha}{l_2^* + \frac{r - p_2}{c}} - \frac{D}{\left(l_2^* + \frac{r - p_2}{c}\right)^2} = \nu - \frac{\varepsilon}{\left(l_2^* + \frac{r - p_2}{c}\right)^2}$$

... an unbiased estimator of ν

Demand experiment

- More informative if it lowers estimator variance
  - Differentiate products more (higher l_2)
  - Lower prices
- Why?
  - Increased reach (further from sweet spot / lower prices) -> demand more affected by ν, less by ε

Solution

- Use Bayes' Rule to calculate posterior cdf on ν
- Find conditional expected second period profit
- Value functions are expected first period profits plus discounted expected second period profits
- Solve for reaction functions and subgame perfect Bayes-Nash equilibrium

Results

- More differentiation: $|l_i|$ greater in period 1
- Lower prices: p_i lower in period 1
- N.B. Exploitation incentive alone is to offer higher price with more product differentiation
  - Holds in period 2 comparative statics: more spread, higher prices wrt α, μ
  - But relative to pure exploitation equilibrium, exploration incentive has first-order effect on both (p,l): more spread, lower prices
  - Both increase sampling from tails, more informative

Comparison to Harrington (1995)

- Information:
  - Harrington: Hotelling line model with firms at extreme ends. Don't know substitutability
  - Ours: Hotelling line but tent distribution of consumers and endogenous firm location. Don't know slope of consumer density
- Results:
  - Harrington: Experiment if not too differentiated
  - Ours: Experiment always

Comparison (cont.)

- Harrington:
  - β is total market wrt price (known)
  - γ is relative demand wrt price (unknown)
- Ours:
  - Market size: $\beta_{GMM} = \left.\frac{\partial D_1}{\partial p_1}\right|_{t_m = t_m^0} = [-\alpha - \nu t_l]/c$
  - Business stealing: $\gamma_{GMM} = \frac{\partial D_1}{\partial p_2} = [-\alpha + \nu t_m]/2c$
  - Both depend on unknown ν

Summary of comparison

- Our producers are uncertain about both business stealing and total market size:
  - learning about market size from public experiment outweighs competitor's harmful behavior

Papers from this project

- Kephart, Das, MacKie-Mason (1999)
  "Two-sided learning in an agent economy for information bundles"
- Brooks, Fay, Das, MacKie-Mason, Kephart, Durfee (1999)
  "Automated strategy searches in an electronic goods market: Learning and complex price schedules"
- Fay, MacKie-Mason (2000)
  "Competition between firms that bundle information goods"
- Kephart, Fay (2000)
  "Competitive bundling of categorized information goods"
- Brooks, Durfee, Das (2000)
  "Price wars and niche discovery in an information economy"
- Gazzale and MacKie-Mason (2000)
  "Endogenous product differentiation under uncertainty"
- Brooks (2001)
  "Learning consumer preferences using model refinement in an information economy"
- Brooks, Das, MacKie-Mason, Kephart, Durfee, Gazzale (2001)
  "Producer learning when consumer preferences for information goods are non-stationary"