Economically Intelligent Learning

Jeff MacKie-Mason

E-Newspaper Wars

- New York Times and Wall Street Journal want to compete in the e-newspaper space
  - What articles should they publish in their e-versions? How many? What editorial niche?
  - How should they price their offerings?
- Suppose NYT sees that emphasis on tech news is working well for WSJ, and mimics
  - Should WSJ try to win with better/more tech news, or try to poach on arts/leisure buyers?
- As new options become available, tastes in hot topics change, customer preferences move
  - Is success change due to competitor strategy shift or preferences shift?

Large problem space

- Suppose NYT has access to N articles
  - can offer 2^N different content bundles
  - can select 2^N different prices, one for each bundle
- How many bundles to offer?
  - Which?
  - At what prices?

Search Problem

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

Learning problem

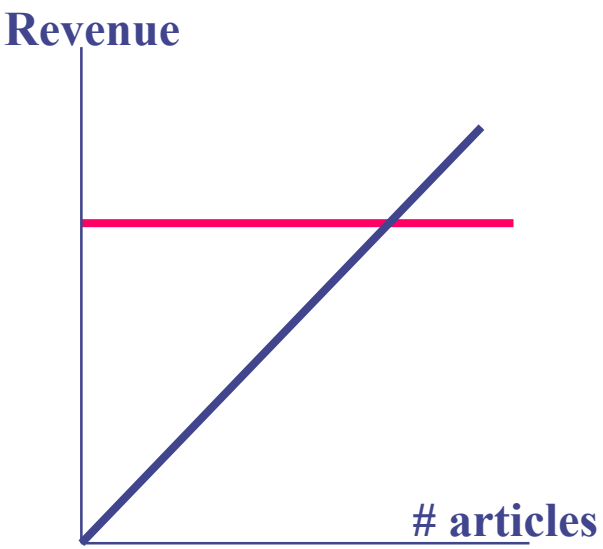
- Don't know everything about customers' preferences for content types and quantity
- Don't know everything about competitors' content selection strategies
- Content, preferences, strategies can change daily
- Problem:
  - For each day,
    - For each of 2^N bundles
      - for each possible bundle/price pair offered by each competitor
        - search for optimal price
    - Choose optimally priced bundle that maximizes expected profit given beliefs about competitor strategy
      - (beliefs may depend on doing above calculations for competitor)

A first look

- Monopolist with 10 information objects that are periodically updated (news stories)
- Large consumer population with stationary but unknown, **heterogeneous** preferences
- Artificial agents assist buyers and sellers search space of possible prices/bundles
  - Brooks, Fay, Das, MacKie-Mason, Kephart, Durfee (EC-99)

Learning Price Schedules

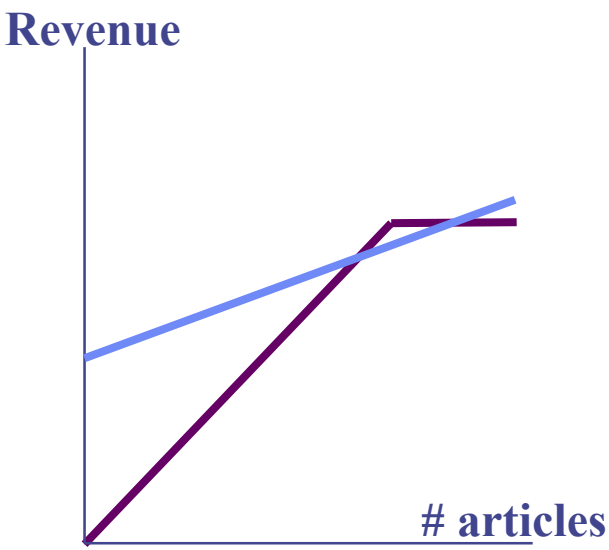
One Parameter



Linear: Each article costs p dollars

Pure Bundling: Entire collection costs b dollars

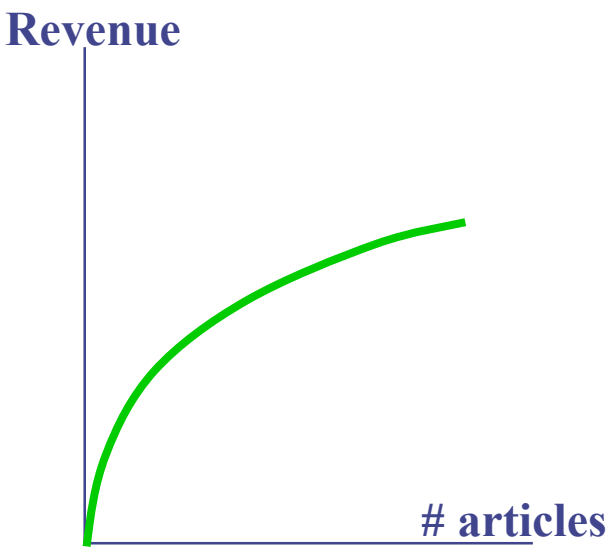
Two Parameter



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

Learning price schedules: Complete information

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

Complete information possibilities

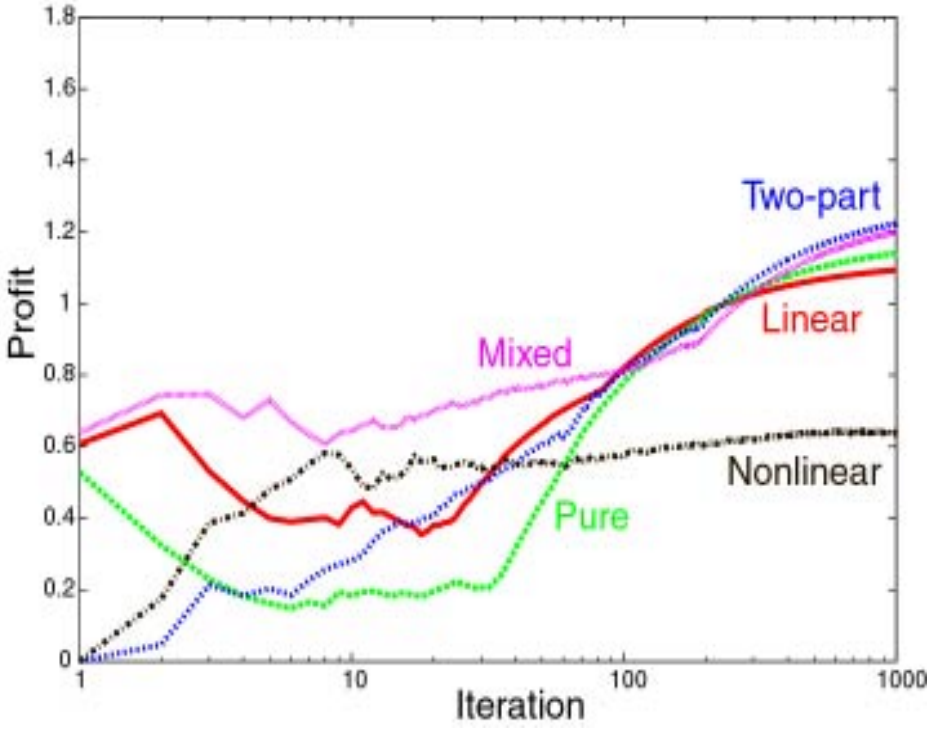
Pricing Scheme	Parameters	Profit
Pure bundling	1	0.88
Linear pricing	1	0.88
Two-part tariffs	2	1.04
Mixed bundling	2	1.04
Block pricing	3	1.09
Non-linear pricing	N	1.15

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

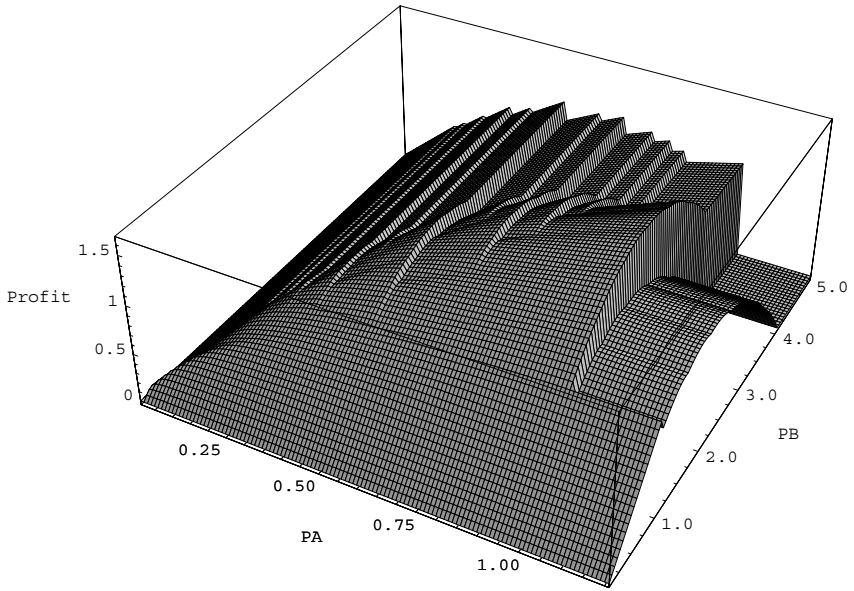
Learning price schedules: Incomplete information

- 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
  - Producer must trade off accuracy of estimate versus time needed to learn.

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.

Broader problem

- How do we develop a principled approach to designing economically rational agents to participate in this environment?
  - What are key features defining environment
  - Which economic behaviors need to be modeled / designed?
  - Given tractability problems (computational & economic), what is "economically intelligent?"

Description of environment (Problem representation)

- Consumer *i* values N goods independently: v^i (**N parameters**)
- Arbitrary pattern of substitution / complementarity. Consumer *i* values 2^N distinct bundles: V(B) (**2^N parameters**)

Problem representation (cont.)

- Suppose all sub / comp effects are bilateral and independent:

$$U^i(B) = \sum_{j \in B} \left(v_j^i - \sum_{\substack{j \in B \\ j \neq i}} \beta_{ij} \right)$$

  Then N^2 elements (v and diagonal elements of β are redundant)

- If β symmetric, then $N(N+1)/2$ unique elements

Learning preferences

- With M consumers, have MN(N+1)/2 parameters to learn…only polynomial
  - N.B.: reduction from exponential required strong restriction on patterns of substitutability
- But: If multiple sellers
  - Consumers on whom you can make observations depends on offerings from other producers
  - If sellers have property rights, firm #1 can only offer a subset N_1 and cannot directly explore (v, β) for other $N-N_1$
  - Indeed, may not be able to directly observe or characterize what competitors are offering

A different representation

- Could pursue learning method further, but first: often cannot track individual purchases, nor offer individual-specific prices
- So, rather than learn M vectors v and matrices β, try to learn the distribution of these parameters
  - (in any case, the distribution, or a close approximation to it, may be more parsimonious)

Preference distributions

- Suppose consumer *i* draws its v^i from $f(\theta)$, with $|\theta|$ parameters
  - $E(v)=\mu_v$, an N vector
  - $vcov(v) = \Sigma_v$, a symmetric NxN matrix
  - possibly higher moments
- And *i* draws its β^i from $g(\Lambda)$
  - $E(\beta) = \mu_B$, a symmetric NxN matrix
  - $vcov(\beta)$ = a big matrix unless structure imposed
- Polynomial number of parameters to learn by imposing restrictions on higher order moments (and still imposing strong restriction on pattern of substitutes and complements)

Learning distributions

- But with multiple sellers have the same problem as before:
  - The observed sample depends systematically on competitor offerings
  - If firm #1 doesn't know **who is**, or **how many** are buying **what from whom**, then each observation is conditional on a sample selection with unknown parameters that must also be learned
  - And, again, **cannot explore directly** the complete space

Alternative representation

- Why try to learn consumer preferences (or their distribution) at all? Ultimately care about profits
- Learn map $\pi(B, p_B) \rightarrow \Pi$ where |B| is 2^N
- Without loss of generality, can set $p_k = \infty$ for bundles not offered, so $\pi : p \rightarrow \Pi$ which is $R_+^{2^N} \rightarrow R$
- If define bundles only in terms of number of items included, so p(B)=p(|B|), then $R_+^N \rightarrow R$

How much easier?

- If C competitors, then the map is

$$\pi : p_1 \times p_2 \times \cdots \times p_C \rightarrow \Pi$$

  so not have CNx1 (or even $C2^N x1$)
- Can't directly search over C(N-1)|p| parameters
- May not be able to observe some things about competitor offerings (at least at reasonable cost)

Equilibrium

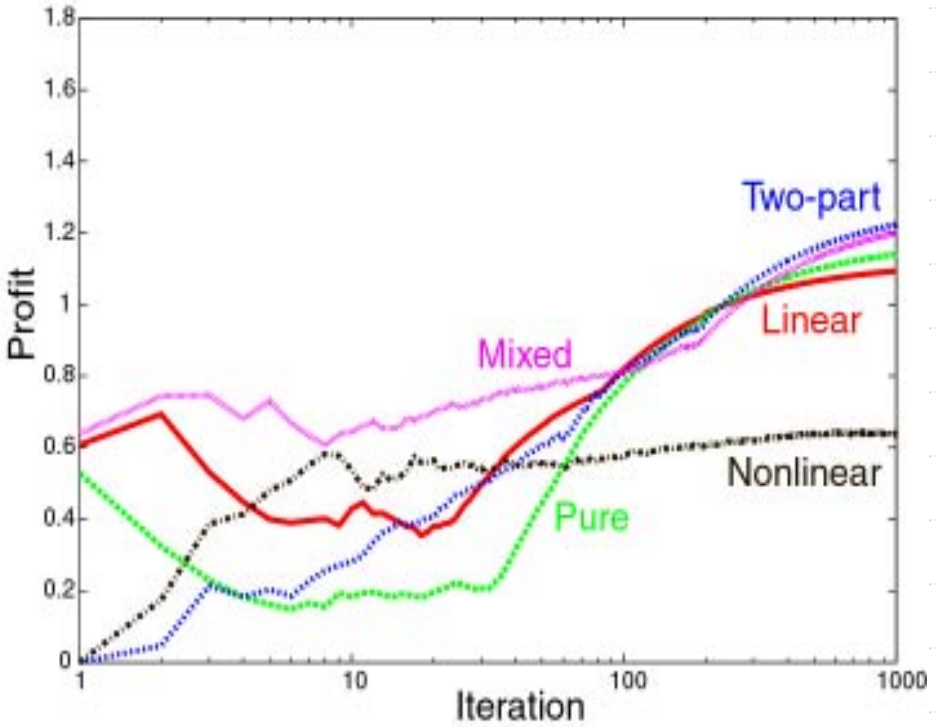
- To solve for Nash, need to observe Π_c for c=1,…C to know full payoff matrix
- Even if known, may have multiple equilibria
- Otherwise, have a dynamic game of incomplete information
  - Convergence?

Learning method

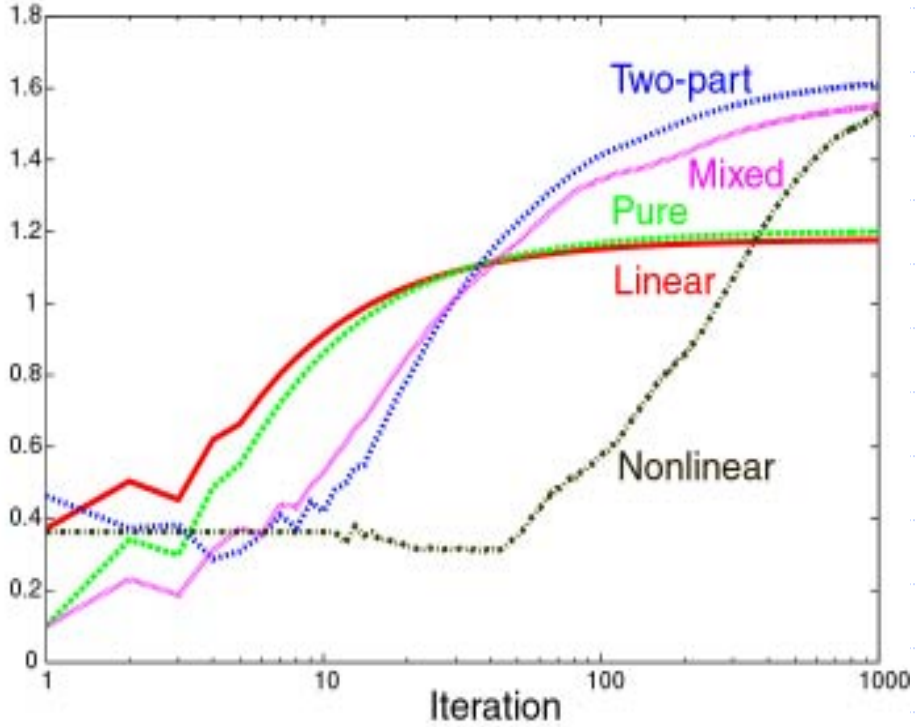
- Want to select / design a learning method that is effective (and efficient) for the problem
- For 1 seller, stationary consumer preferences, hill-climbing beat function-mapping (usually?)

Different learning methods

Neural network



Amoeba



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

But…

- If |C|>1, the "peak" depends on prices offered by others
  - if ignore this and use hill-climbing, will seller find a Nash equilibrium? Get close?
  - would function mapping do better because firm could learn how to change the map as competitors change offerings?
  - different view: On a different map for every matrix of competitors' prices. Should firm try to keep track of which map it is currently on?

Summary: Credit attribution

- Firm may want to solve both inter-agent and intra-agent CAPs
  - CAP: assigning credit for an overall performance change to each of the system activities that contributed to that change
  - **intra**: (consumers) "What knowledge, what inferences and what decisions led to an individual agent's action?"
  - **intra**: (competitors) "What action carried out by what agent contributed to what extent to the performance change?"

Other complications

- If firm knows that other firms trying to learn the map will it try to mislead, or systematically distort its own exploration?

Integrating learning and economic theory

- "The design of the learning element is affected by four major issues:
  - **Which components** of the performance element are **to be improved**?
  - **What representation** is used for those components
  - **What feedback** is available
  - **What prior information** is available" (Russell & Norvig)

The challenge

- The "performance element" of the economically intelligent agent: the element that selects external actions
- Russell & Norvig: "The components [of the performance element] can include the following:
  - A direct mapping from conditions on the current state to actions
  - A means to infer relevant properties of the world from the percept sequence
  - Information about the way the world evolves
  - Information about the results of possible actions the agent can take
  - Utility information indicating the desirability of world states
  - Action-value information indicating the desirability of particular actions in particular states
  - Goals that describe classes of states whose achievement maximizes the agent's utility."
- Each can be described as a function

Economically intelligent learning

- Choosing the representation for each of these functions is the primary challenge
- Then designing a learning strategy that improves our knowledge of those functions