Competitive Demand Learning

Yu-Ching Lee

Department of Industrial Engineering and Engineering Management

National Tsing Hua University

January 17 2022 @ 2021 TMS Annual Meeting

Outline

Education and Affiliation

Assistant Professor, National Tsing Hua University, Hsinchu
- Equilibrium and Data-analytics Laboratory

Ph.D. Industrial Engineering 2012
University of Illinois at Urbana-Champaign
- global optimization, complementarity problems

M.S. Civil Engineering
University of Illinois at Urbana-Champaign
- transportation, airline problems

B.S. Civil Engineering
National Taiwan University

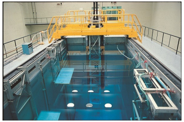

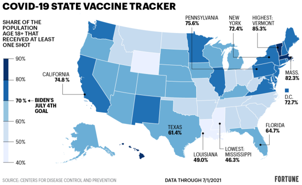

Competitive Demand Learning:
A Coordinated Price Experimentation
and Non-cooperative Pricing Algorithm

Joint work with
Yongge Yang (National Tsing Hua University),
Po-An Chen (National Yang Ming Chiao Tung University).

The spirit of this research is one word: the "teamwork"

Introduction about this work

- Consider a total of F firms selling substitutable products in an oligopolistic market, in which the true underlying demand curve and the presence of demand shocks are unknown. Over a time horizon of T periods, firms make pricing decisions in each period $t = 1, \cdots, T$.
- The price decisions made by other firms will influence the demand for the product of firm i.
- By focusing on competition among the F firms, we do not consider capacity limitation, production cost or marginal cost. Each firm is assumed to be selfish and reacts immediately to price changes made by competitors.

Introduction about this work (conti'd)

- This paper generalizes the work of Besbes and Zeevi (2015), who constructed a dynamic pricing algorithm in a monopoly setting in which a single firm chooses a price to maximize expected revenue without knowledge of the true underlying demand curve.
- We propose a data-driven equilibrium pricing (DDEP) algorithm to solve the dynamic pricing decisions of each firm in competition.
- In an alternative scenario, in which some firms have knowledge of the demand function and the distribution of demand shocks, such firms may be unwilling to engage in price experimentation. Therefore, we propose a modified DDEP algorithm (in the full paper) to account for this.
- The process of learning is often evaluated in terms of *regrets*.
- We also analyze the revenue difference obtained by the algorithm from that obtained by the clairvoyant Nash equilibrium p^* per algorithm iteration.

Related Literature

- **Learning algorithms for pricing models in a monopoly market.**
  - Besbes O, Zeevi A (2015) On the (surprising) sufficiency of linear models for dynamic pricing with demand learning. Management Science 61(4):723–739.
  - Chen B, Chao X, Ahn HS (2019) Coordinating pricing and inventory replenishment with nonparametric demand learning. Operations Research 67(4):1035–1052.
  - Chen B, Chao X, Shi C (2015) Nonparametric algorithms for joint pricing and inventory control with lost sales and censored demand. Mathematics of Operations Research, Major Revision.
- **Dynamic pricing in a competitive environment.**
  - Cooper WL, Homem-de Mello T, Kleywegt AJ (2015) Learning and pricing with models that do not explicitly incorporate competition. Operations research 63(1):86–103.
- **Learning algorithms for pricing models in a competitive environment.**
  - Bertsimas D, Perakis G (2006) Dynamic pricing: A learning approach. Mathematical and computational models for congestion charging, 45–79 (Springer)
  - Gallego G, Talebian M (2012) Demand learning and dynamic pricing for multi-version products. Journal of Revenue and Pricing Management 11(3):303–318.
  - Fisher M, Gallino S, Li J (2018) Competition-based dynamic pricing in online retailing: A methodology validated with field experiments. Management Science 64(6):2496–2514.

Model and Preliminaries

We consider a periodical equilibrium pricing problem for F firms. In each period $t = 1, \cdots, T$, each firm needs to set prices p_t^i, chosen from a feasible and bounded policy set $\mathcal{P}^i = \left[p^{i,\ell}, p^{i,h}\right]$, $p^{i,\ell} < p^{i,h}$, $\forall i = 1, \cdots, F$. The prices set by firms affect the market response of all firms in the competition.

Recall that $\mathrm{p} \equiv (p^i, p^{-i})$ denotes the vector of prices of all firms in the competition. The market response to the price p_t^i for firm i at time t (which is exactly the demand function) is given by $D_t^i(\mathrm{p}_t) = \lambda^i(\mathrm{p}_t) + \varepsilon_t^i$, $\forall i = 1, \cdots, N$, in which $\lambda^i(\mathrm{p}_t)$ is a deterministic twice differentiable function representing the mean demand, conditional on the price p_t

ε_t^i are zero-mean random variables, assumed to be independent and identically distributed.

Model and Preliminaries (conti'd)

Hence, the demand curve $\lambda^i(\mathrm{p})$ of firm i not only depends on the price p^i, chosen by itself, but also on the prices of other firms p^{-i}, where $p^{-i} = \{p^1, \cdots, p^{i-1}, p^{i+1}, \cdots, p^F\}$.

Let $\pi^i = (p^i_1, p^i_2, \cdots)$ denote the sequence of pricing policy of firm i and $\Pi = (\pi^1, \cdots, \pi^F)$ denote the admissible pricing policies of all firms.

The revenue function r^i of firm i obtained from prices p is denoted by $r^i(\mathrm{p}) \equiv p^i \mathbb{E}[D^i(\mathrm{p})]$. Each firm seeks to maximize its revenue in a competitive environment.

Model and Preliminaries (conti'd)

Throughout the paper, we use period t or, equivalently, time t.

Let p_t^{i*} denote the equilibrium price of firm i at time t, which is obtained by the estimated demand curve of firm i at time t, and let p_t^{-i} denote the prices of other competitors.

A *clairvoyant* model implies that a firm has knowledge of the underlying demand curve and the distribution of demand shocks.

The goal of learning is to make p_t^i converge to the clairvoyant equilibrium price of firm i, p^{i*}, as t grows large. Note that a learning scheme in which the difference between p_t^i and p^{i*} will eventually converge to zero is called *complete* learning; otherwise, it is termed *incomplete* learning.

Assumptions

(i) For any $p^i \in \mathcal{P}^i$, $\dfrac{\partial \lambda^i(\cdot, p^{-i})}{\partial p^i} < 0, \forall i = 1, \cdots, F$.

(ii) For any $p^i \in \mathcal{P}^i$, $\dfrac{\partial \lambda^i(p^i, p^{-i \setminus j}, \cdot)}{\partial p^j} > 0, \forall j \neq i, i = 1, \cdots, F$.

(iii) For any $p^i \in \mathcal{P}^i$, $\dfrac{\partial^2 r^i(\mathrm{p})}{\partial^2 \mathrm{p}}, \forall i = 1, \cdots, F$ is a negative definite matrix.

(iv) For every $r^i(\mathrm{p})$,

$$\sum_{j \neq i}^{F} \left| \frac{\partial^2 r^i}{\partial p^i \partial p^j} \right| < \left| \frac{\partial^2 r^i}{\partial p^{i2}} \right|, \forall p^i \in \mathcal{P}^i, p^j \in \mathcal{P}^j.$$

(v) For every firm i, there exists a constant s_0 such that, for all $|s| < s_0$, $\mathbb{E}\left[\exp\left\{s\varepsilon_1^i\right\}\right] < \infty$, and the variance of each firm's ε^i is equal to σ^2.
(vi) For every firm i, given p^{-i}, firm i chooses to price at a $p^i \in \mathcal{P}^i$ satisfying $\mathbb{E}\left[D^i(p^i, p^{-i})\right] \geq 0$.

Competitive Demand Learning (CDL) algorithm

loop n from 0 until a terminal stage, given as period T.

- Step 0. Preparation: If $n = 0$, input I_0, v, and $\hat{p}_1^i, \forall i = 1, \cdots, F$. If $n > 0$, set $I_n = \lfloor I_0 v^n \rfloor$ and $\delta_n = I_n^{-\frac{1}{4}}$.
- Step 1: Setting prices. Loop m from 1 to $F + 1$. The rule of firm i's price p_t^i at time t is

$$\begin{array}{ll} \text{if } m \neq i, & \\ p_t^i = \hat{p}_n^i, & \forall t = t_n + 1, \cdots, t_n + iI_n, t_n + (i+1)I_n + 1, \cdots, t_n + (F+1)I_n, \\ \text{if } m = i, & \\ p_t^i = \hat{p}_n^i + \delta_n, & \forall t = t_n + iI_n + 1, \cdots, t_n + (i+1)I_n. \end{array}$$

  End the m-loop. Set $t_{n+1} = t_n + (F+1)I_n$.
- Step 2. Estimating:

$$(\hat{\alpha}_{n+1}^i, \hat{\beta}_{n+1}^{ij}) = \arg\min_{\alpha^i, \beta^{ij}} \left\{ \sum_{t=t_n+1}^{t=t_n+(F+1)I_n} \left[D_t^i - \left(\alpha^i - \beta^{ii} p_t^i + \sum_{j=1, j\neq i}^{F} \beta^{ij} p_t^j \right) \right]^2 \right\}.$$

- Step 3. Computing the equilibrium: We define the following optimization problem for firm i: $\max_{p^i} r^i_{n+1} \equiv \max_{p^i} G_{n+1}\left\{p^i, p^{-i}, \hat{\alpha}^i_{n+1}, \hat{\beta}^{ij}_{n+1}\right\}$, where $G_{n+1}\left\{p^i, p^{-i}, \hat{\alpha}^i_{n+1}, \hat{\beta}^i_{n+1}\right\}$ $\equiv \left\{ p^i \left(\hat{\alpha}^i_{n+1} - \hat{\beta}^{ii}_{n+1} p^i + \sum_{j=1, j\neq i}^{F} \hat{\beta}^{ij}_{n+1} p^j \right) \middle| \hat{\alpha}^i_{n+1} - \hat{\beta}^{ii}_{n+1} p^i + \sum_{j=1, j\neq i}^{F} \hat{\beta}^{ij}_{n+1} p^j \geq 0, p^i \in \mathcal{P}^i \right\}$. Proceeding to solve the system:

$$\begin{aligned}
&\left[\hat{\alpha}^i_{n+1} - 2\hat{\beta}^{ii}_{n+1} p^i + \sum_{j,j\neq i}^{F} \hat{\beta}^{ij}_{n+1} p^j \right] + \mu^{i,1} \left[-\hat{\beta}^i_{n+1} \right] - \mu^{i,2} + \mu^{i,3} = 0 && \forall i,\\
&\mu^{i,1} \geq 0, \mu^{i,1} \cdot \left(-\hat{\alpha}^i_{n+1} + \hat{\beta}^{ii}_{n+1} p^i - \sum_{j,j\neq i}^{F} \hat{\beta}^{ij}_{n+1} p^j \right) = 0, \hat{\alpha}^i_{n+1} - \hat{\beta}^{ii}_{n+1} p^i + \sum_{j,j\neq i}^{F} \hat{\beta}^{ij}_{n+1} p^j \geq 0 && \forall i,\\
&\mu^{i,2} \geq 0, \mu^{i,2} \cdot \left(p^i - p^{i,h} \right) = 0, p^{i,h} - p^i \geq 0 && \forall i,\\
&\mu^{i,3} \geq 0, \mu^{i,3} \cdot \left(p^{i,l} - p^i \right) = 0, p^i - p^{i,l} \geq 0 && \forall i.
\end{aligned}$$

Then, prices for each firm $\hat{p}^i_{n+1}$ are set to the unique solution of this system. Set $n = n + 1$ and **return to Step 0**.

Assumptions Implications

(i) ensures that for every firm i, the underlying demand function $\lambda^i(\cdot, p^{-i})$ is strictly decreasing on p^i given the prices set by other firms, p^{-i}.
(ii) dictates that $\lambda^i(p^i, p^{-i\setminus j}, \cdot)$ is strictly increasing on p^j with p^i and $p^{-i\setminus j}$ given, in which $p^{-i\setminus j}$ represents the vector constituted by all prices except p^i and p^j.
(iii) dictates that the revenue function $r^i(\mathrm{p})$ is a concave function and thus there exists a unique maximizer for any feasible p.
(iv) is termed as the "diagonal dominance" condition.
(v) ensures that the demand shock ε_t^i of each firm has a light-tailed distribution and the homogeneity of variance.
(vi) ensures that each firm only considers a price within a price interval such that the estimated demand is non-negative.

Assumptions Examples

We give some examples of demand functions from that satisfy Assumption 2.(i).

1. Linear demand: $\lambda^i(\mathrm{p}) = \alpha^i - \beta^{ii} p^i + \sum_{j=1, j\neq i}^{F} \beta^{ij} p^j$, $\beta^{ii} > 0$.
2. Multinomial logit demand: $\lambda^i(\mathrm{p}) = \frac{\exp^{\alpha^i - \beta^i p^i}}{1+\sum_{i=1}^{F} \exp^{\alpha^i - \beta^i p^i}}$, $\alpha^i > 0, \beta^i > 0$ and

$\alpha^i - \beta^i p^i < 0$ for $p^i \in \mathcal{P}^i$.

3. Cobb-Douglas demand: $\lambda^i(\mathrm{p}) = \alpha^i (p^i)^{-\beta^{ii}} \prod_{j\neq i}^{F} (p^j)^{\beta^{ij}}$, $\alpha^i > 0, \beta^{ii} > 1, \beta^{ij} \geq 0$.
4. CES demand: $\lambda^i(\mathrm{p}) = \frac{\gamma (p^i)^{r-1}}{\sum_{j=1}^{F} (p^j)^r}, r < 0, \gamma > 0$.

Lemma 1: Analysis for the Noiseless Case.

Lemma 1

Suppose that $\varepsilon_t^i = 0$, $\forall i$ and t, and that the sequence $\{\hat{\mathrm{p}}_n\}$, assuming nonzero demand and that the price is away from the limits, generated by CDL converges to a limit point $\tilde{\mathrm{p}}$, which satisfies $\tilde{p}^i = -\dfrac{\lambda^i(\tilde{\mathrm{p}})}{\nabla_{p^i}\lambda^i(\tilde{\mathrm{p}})}$. Then, $\tilde{\mathrm{p}}$ is exactly p^*.

Lemma 2: Uniqueness of $\hat{\mathrm{p}}_n$.

Lemma 2

Under the assumptions, $\hat{\mathrm{p}}_n$ is a unique GNE at stage n with high probability.

Proposition 1

Proposition 1

If $p'^i = z^i\left(\breve{\alpha}^i(p^i, p^{-i}), \breve{\beta}^{ii}(p^i, p^{-i}), \breve{\beta}^{ij}(p^i, p^{-i})\right)$ for all i, then there exists a constant $\gamma \in (0,1)$ such that

$$\|\mathrm{p}^* - \mathrm{z}(\hat{\mathrm{p}}_n)\| \leq \gamma \|\mathrm{p}^* - \hat{\mathrm{p}}_n\|.$$

Proposition 1 is based on a deterministic (mean) demand function, and the convergence result follows directly from the property of a contraction mapping. Now, we focus on a randomized demand function and we aim to establish the convergence result as follows.

Proposition 2

Proposition 2

For any given $\hat{p}_n^i \in \mathcal{P}^i$ generated by CDL, with high probability the following inequality holds

$$\|\mathrm{z}(\hat{\mathrm{p}}_n) - \hat{\mathrm{p}}_{n+1}\| \leq \|C_n\|,$$

where $C_n \equiv \left[C_n^1, \cdots, C_n^F\right]$ is a vector of constants.

Proposition 2 shows that the difference between these best response functions is bounded with high probability. The probability lower bound is specified in the proof of Proposition 3. Note that $\mathrm{z}(\hat{\mathrm{p}}_n)$ denotes the collection of all firms' best responses $z^i\left(\breve{\alpha}^i(\hat{\mathrm{p}}_n), \breve{\beta}^{ii}(\hat{\mathrm{p}}_n), \breve{\beta}^{ij}(\hat{\mathrm{p}}_n)\right)$.

Proposition 3

Proposition 3

At stage n, for some suitable constant K_1, the operator $\mathrm{z}(\hat{\mathrm{p}}_n)$ and the CDL generated $\hat{\mathrm{p}}_{n+1}$ satisfy

$$\mathbb{E}\left[\|\mathrm{z}(\hat{\mathrm{p}}_n) - \hat{\mathrm{p}}_{n+1}\|^2\right] \leq F^2 K_1 I_n^{-\frac{1}{2}}.$$

Proposition 3 also provides an upper bound for the deviation between $\mathrm{z}(\hat{\mathrm{p}}_n)$ and $\hat{\mathrm{p}}_{n+1}$. The upper bound is related to the squared number of competitive firms, F^2, and converges to zero as the number of stages increases.

Analysis: Convergence, Revenue Difference, and Regret

Theorem 1: Convergence

Under Assumptions, the GNE, $\hat{\mathrm{p}}_n$ converges in probability to p^* as $n \to \infty$.

The best response function derived by CDL through the quadratic concave function will generate the sequence $\{\mathrm{p}_t\}$ that converges to p^* as t grows large. (Theorem 1)

Analysis: Convergence, Revenue Difference, and Regret

Theorem 2: Revenue Difference

Under Assumptions, the sequence of the generalized Nash equilibrium $\{\hat{\mathrm{p}}_t : t \geq 1\}$ satisfies

$$\mathbb{E}\left[\sum_{t=1}^{T}\left[\left|r^i(p^{i*}, p^{-i*}) - r^i(p_t^i, p_t^{-i})\right|\right]\right] \leq F^2 K_6 T^{\frac{1}{2}},$$

$$\forall i = 1, \cdots, F,$$

for some positive constant K_6, $T \geq 2$, and $F \geq 2$.

The revenue difference converges to zero as time progresses and is related to the quantity of firms in competition. The difference implies that realised revenues are sometimes greater than those revenues obtained by the clairvoyant Nash equilibrium p^*. However, we are unable to predict when this will occur. (Theorem 2)

Analysis: Convergence, Revenue Difference, and Regret

Theorem 3: Regret

Under the defined assumptions, the sequence of optimal decisions $\{p_t^{i*} : t \geq 1\}$ satisfies

$$\mathbb{E}\left[\sum_{t=1}^{T}\left[r^i(p_t^{i*}, p_t^{-i}) - r^i(p_t^i, p_t^{-i})\right]\right] \leq K_7 F T^{\frac{1}{2}},$$

$$\forall i = 1, \cdots, F,$$

for some positive constant K_7, $T \geq 2$ and $F \geq 2$.

As time progresses, the accumulated revenue of each firm generated by the pricing policies of CDL algorithm is asymptotically close to the clairvoyant accumulated revenue. (Theorem 3)

Concluding Remarks

- We designed a mechanism of synchronized dynamic pricing. Such a mechanism ensures that the pricing strategy of each firm is adjusted in a prescribed way to jointly collect demand information and make pricing decisions.
- We asked whether the mechanism may allow prices to reach a stable state and how much regret firms incur by employing such a data-driven pricing algorithm.
- In particular, the effects of noise vanish as n increases and that the fitted linear model can serve as an estimation of the underlying demand model without being affected by F.
- When facing competition, the upper bound of revenue regret, derived in the same way as that of one firm, is scaled by F (i.e., Theorem 3), the upper bound of revenue difference is scaled by F^2 (i.e., Theorem 2), and the deviation between the best responses and the clairvoyant GNE price is upper bounded by a factor of $F^2 I_n^{-\frac{1}{2}}$