Oligopoly Dynamic Pricing: A Repeated Game with Incomplete Information

Yixuan Zhai — University of California, Davis

Qing Zhao — Cornell University

Oligopoly Dynamic Pricing

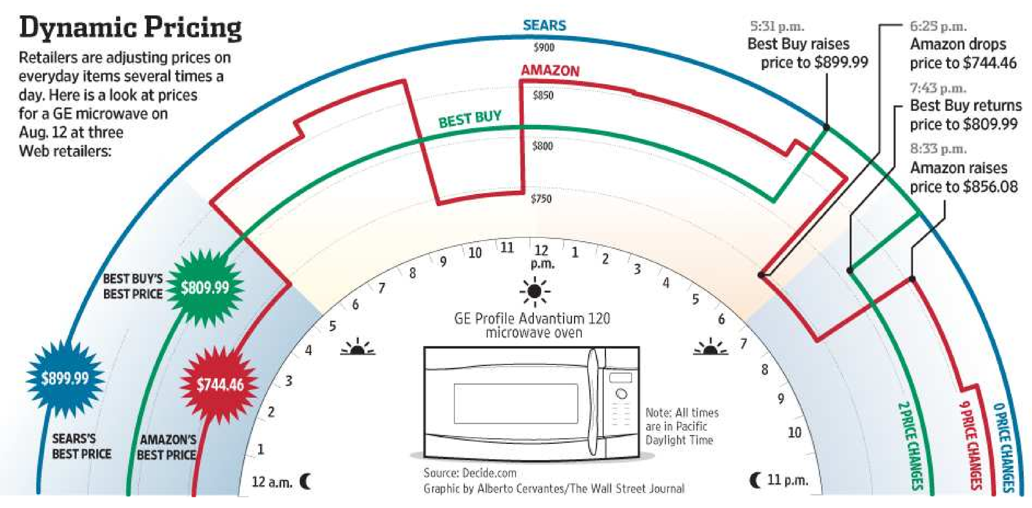

Single Seller Dynamic Pricing

- The *profit-maximization* problem for a seller with an *unlimited* supply of identical goods.
- The seller offers prices sequentially to a stream of potential customers.
- The marginal cost is c.
- For the t-th customer, the seller chooses a price $p(t) \in [c, p_u]$.
- The seller experiences either success or failure.
- The probability of success at price p at any given time is $\rho(p)$.

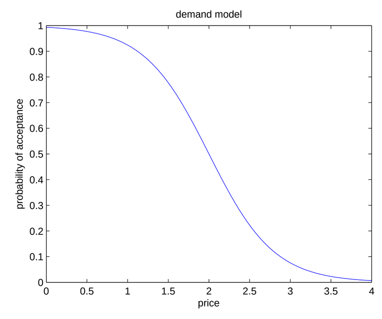

Single Seller Dynamic Pricing

Unknown Demand Model with Finite Space Uncertainty

- The unknown underlying demand model is $\rho(p)$.
- $\rho(p)$ is unknown but belongs to a known finite set $\{\rho^{(m)}(p)\}_{m=1}^{M}$.
- $\rho(p)$ is strictly decreasing and satisfies the increasing generalized failure rate (IGFR).
- Intuitively, IGFR means that given a seller can sell a product at price p, the probability of sale at $p + \Delta p$ is decreasing in p.

Profit

- The expected profit: $r^{(\omega)}(p, c) = (p - c)\rho^{(\omega)}(p)$.
- The profit-maximizing price $p^{(\omega)}(c) = \arg\max_{p \in [c, p_u]} r^{(\omega)}(p, c)$.

Lemma 1. *Suppose that $\rho(p)$ is strictly decreasing and satisfies IGFR. Then $p^{(\omega)}(c)$ is unique, and $r^{(\omega)}(p, c)$ is continuous and strictly increasing with p over $[c, p^{(\omega)}(c)]$.*

Oligopoly Dynamic Pricing

- N sellers.
- The marginal cost for seller i is c_i.
- Without loss of generality, we assume $c_1 < c_2 < \ldots < c_N$.
- Unknown demand model $\{\rho^{(m)}(p)\}_{m=1}^{M}$.
- Sellers propose their prices simultaneously.
- The customer accepts the lowest price p with probability $\rho(p)$.

An Infinitely Repeated Game

Public and private history

- Sellers' price offers are public.
- Each seller's sale history is private.

Strategy of Seller i: A mapping σ from public information and seller i's private information to a price offer at current time.

Seller i's one-shot payoff $u_i(p_i)$

- If a single seller i offers the lowest price p_i:
  - $u_i(p_i) = (p_i - c_i)\rho(p_i)$.
  - $u_j(p_j) = 0 \; \forall j \neq i$.
- If K sellers offer the same lowest price p:
  - $u_i(p) = \frac{1}{K}(p - c_i)\rho(p)$.

Seller i's payoff for the infinitely repeated game:

$$U_i^{(\omega)}(\sigma) = \liminf_{T \to \infty} \frac{1}{T} \sum_{t=1}^{T} u_i^{(\omega)}(a^t(\sigma)).$$

Equilibria and Efficiency

Equilibria

- Nash Equilibrium (NE) σ
  - $U_i(\sigma) \geq U_i(\sigma'_i, \sigma_{-i})$ for every seller i, every demand model $\rho^{(m)}$ and all strategies σ'.
- Subgame perfect equilibrium
  - NE may not be sequentially rational.
  - For every subgame of the original game, the induced continuation strategy is a NE of the subgame.

Efficiency

- Pareto Efficiency:
  - The payoff of any player cannot be increased without reducing the payoff of at least one other player.
- Learning efficiency: Regret
  - Regret is defined as the accumulated profit loss in the unknown demand case as compared to the known demand case.

Dynamic Pricing under Known Demand Model

The colluding strategy σ_C under known demand model:

- Seller 1 forms the optimal collusion of K sellers to maximize its own profit.

$$K = \arg\max_{k} \frac{1}{k} r^{(\omega)}(c_{k+1}, c_1)$$

- Sellers with $c_i < c_{K+1}$ offer the collusive price c_{K+1}.
- Sellers with $c_i \geq c_{K+1}$ will not participate.
- Any deviations will trigger a punishment that seller 1 offers $c_2 - \epsilon$ and seller $i \neq 1$ offers c_i forever.

Theorem 1. *The colluding strategy σ_C is a subgame-perfect and Pareto-efficient Nash equilibrium.*

Demand Learning under Collusion

- The time horizon is partitioned into ***fixed length epochs***.
- Each epoch starts with a ***declaration*** time slot followed by ***cooperation*** time slots.

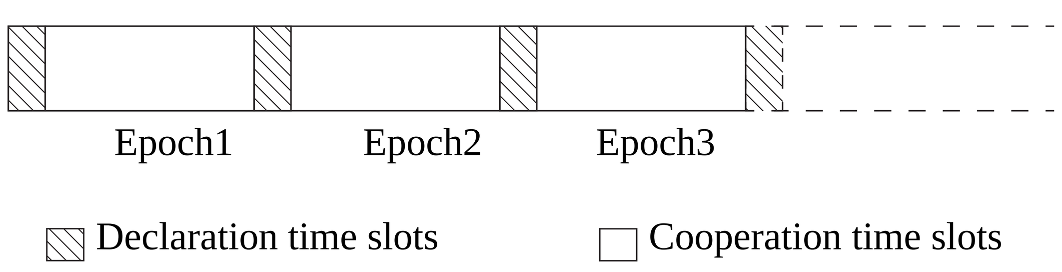

Demand Learning under Collusion

Demand Learning under Collusion

In the declaration slot of epoch t:

- Seller 1 carries out a maximum likelihood estimate $\hat{\omega}(t)$ of the underlying demand model using its private history.
- Seller 1 then offers the profit-maximizing colluding price $\hat{p}^{(\hat{\omega}(t))}$.
- All other sellers offer the same price they offered in the cooperation slots in the previous epoch.

In each cooperation time slot of epoch t

- Seller i with $c_i < \hat{p}^{\hat{\omega}(t)}$ offers $\hat{p}^{\hat{\omega}(t)}$.
- Not participate if $c_i \geq \hat{p}^{\hat{\omega}(t)}$.

Trigger strategy for punishing any deviations

- Any deviations in declaration slots from seller $2, \ldots, N$ and any deviations in cooperation time slots will trigger a everlasting punishment.
- Punishment is that seller 1 offers $c_2 - \epsilon$ and seller $i \neq 1$ offers c_i forever.

Properties of DLC

Theorem 2.

- *DLC is a subgame-perfect Nash equilibrium.*
- *DLC is a Pareto-efficient Nash equilibrium.*
- *DLC achieves a bounded regret, i.e., under any demand model $\rho^{(\omega)} \in \{\rho^{(\omega)}\}_{\omega=1}^{M}$, there exists a positive constant C such that*

$$R_{DLC} \leq C.$$

Implications of bounded regret

- The action profiles converges to the optimal action profile played as under the known demand model.

Simulation

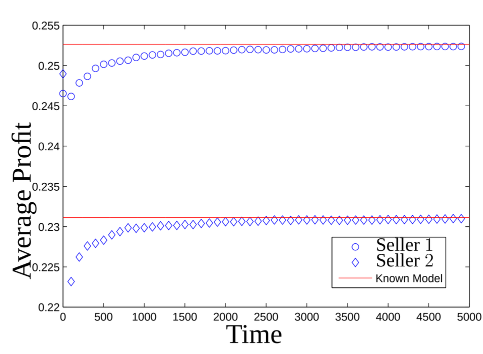

Average Profit when demand model is $\rho^{(2)}$ ($c_1 = 0.5, c_2 = 0.6$)

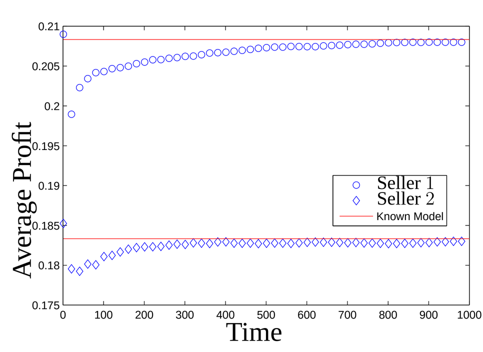

Average Profit when demand model is $\rho^{(1)}$ ($c_1 = 0.5, c_2 = 0.6$)

Conclusion

Oligopoly dynamic pricing

- Infinitely repeated game with private observations.
- Incomplete information: payoff determined by an unknown demand model.
- The optimal collusion with a subset of sellers.

Demand Learning under Collusion (DCL):

- Subgame perfect Nash equilibrium.
- Pareto Efficient.
- Efficient online learning with bounded regret.