Dynamic Competition

Glenn Ellison

Collusion with Perfect Observability

- N identical firms with constant marginal cost *c*. Discount factor δ.
- Bertrand competition at *t* = 1,2,… with market demand *Q(p)*.
- Firms observe all prices at end of each period:
  - History: $h^{t-1} = \{(p_{1\tau}, \dots, p_{N\tau})\}_{\tau \le t-1}$
  - Firm *i*'s strategy: $p_{it}(h^{t-1})$ for *t = 1,2,…*

In 14.122 I noted that

- The model has an SPE that looks like perfect competition:

$$p_{it}\left(h^{t-1}\right) = c \;\; \forall i, t, h^{t-1}$$

- For $\delta \ge 1 - \frac{1}{N}$ the model also has a SPE that looks like perfect collusion:

$$p_{it}\left(h^{t-1}\right) = \begin{cases} p^m \text{ if } h^{t-1} = (p^m, \dots, p^m) \\ c \quad \text{otherwise} \end{cases}$$

Collusion with Perfect Observability

- The model has an SPE that looks like perfect competition:

$$p_{it}\left(h^{t-1}\right) = c \;\; \forall i, t, h^{t-1}$$

- For $\delta \geq 1 - \frac{1}{N}$ the model also has a SPE that looks like perfect collusion:

$$p_{it}\left(h^{t-1}\right) = \begin{cases} p^m \text{ if } h^{t-1} = (p^m, \ldots, p^m) \\ c \quad \text{otherwise} \end{cases}$$

Remarks:

1. The model identifies another factor that can lead to positive markups.
2. It is easier to sustain collusion when δ is larger.
   - Could explain how detection lags, smooth/lumpy demands, growth, and technological change affect markets.
3. It is easier to sustain collusion when N is smaller.
4. Given the equilibrium multiplicity we must believe in some equilibrium selection or accept our inability to forecast.

Factors Limiting Collusive Pricing

We think that tacit collusion is much less common than the $\delta \geq 1 - \frac{1}{N}$ repeated game model suggests. Factors that make collusion more difficult include:

1. Imperfect observation of actions
2. Variable demand
3. Cost shocks
4. Antitrust enforcement

There are also factors that could make collusion easier than the model suggests:

1. Multimarket contact

Collusion with Imperfect Monitoring

Green-Porter (1984)

- Two firms with marginal cost *c* compete as in Bertrand at $t = 1, 2, \ldots$
- Market demand is noisy:

$$Q_t(p) = \begin{cases} Q(p) \text{ with probability } 1 - \alpha \\ 0 \quad\quad \text{ with probability } \alpha \end{cases}$$

Demand goes to the lower-priced firm (or splits 50-50 if prices equal).

- Imperfect observability: firms see their own demand, but not rival's price or quantity.

Note: a firm that received zero demand does not know if the low-demand state arose or if the rival set a lower price.

Proposition: This model does not have an SPE in which both firms set $p_{it} = p^m$ in each period on the equilibrium path.

Proof sketch: In such a SPE firm 1 would need to keep charging p^m even if it got zero demand in the first 1,000,000 periods. Given this, firm 2 will want to cut its price to $p^m - \varepsilon$.

Collusion with Imperfect Monitoring

Green-Porter (1984)

- Market demand is noisy:

$$Q_t(p) = \begin{cases} Q(p) & \text{with probability } 1-\alpha \\ 0 & \text{with probability } \alpha \end{cases}$$

Proposition: For $\alpha < \frac{1}{2}$ and δ sufficiently close to one the model does have a partially collusive equilibrium in which firms initially set $p_{it} = p^m$, switch to $p_{it} = c$ for T periods every time some firm gets zero demand, and then go back to $p_{it} = p^m$ after the T periods are over (for some T).

Proof sketch: Write V^m for the PDV of payoffs at t=0 and V^p for the PDV of payoffs at the start of the punishment phase. We need to show that there are no profitable single period deviations in any states of this process.

Showing that there is no profitable deviation during the punishment phase is easy: a firm that deviates can't earn positive profits in the period in which in deviates and it does not affect pricing in any other period.

It remains just to show that firms don't want to deviate at t with $p_{it} = p^m$.

Collusion with Imperfect Monitoring

Green-Porter (1984)

Proof sketch (cont'd): To show that firms don't want to deviate in the cooperate phase we want to show that $V^m \geq \pi^m(1-\alpha) + \delta V^p$.

The value functions are the solution to

$$V^m = (1-\alpha)\left(\frac{\pi^m}{2} + \delta V^m\right) + \alpha\delta V^p \qquad\qquad V^p = \delta^T V^m$$

Substituting the second expression into the first the solution is

$$V^m = \frac{(1-\alpha)\pi^m}{2}\left(\frac{1}{1-((1-\alpha)\delta+\alpha\delta^{T+1})}\right) \qquad\qquad V^p = \delta^T V^m$$

The required inequality is then $(1-\delta^{T+1})V^m \geq \pi^m(1-\alpha)$. Cancelling the common $\pi^m(1-\alpha)$ terms and multiplying by the denominators gives

$$(1-\delta^{T+1}) \geq 2(1-((1-\alpha)\delta + \alpha\delta^{T+1})) \;\; \Leftrightarrow 2(1-\alpha)\delta + (2\alpha-1)\delta^{T+1} \geq 1.$$

For $\alpha < \frac{1}{2}$ this will hold if δ is close to 1 and T is large.

Collusion with Imperfect Monitoring

Green-Porter (1984)

Proposition: For $\alpha < \frac{1}{2}$ and δ sufficiently close to one the model does have a partially collusive equilibrium in which firms initially set $p_{it} = p^m$, switch to $p_{it} = c$ for T periods every time some firm gets zero demand, and then go back to $p_{it} = p^m$ after the T periods are over (for some T).

Remarks:

1. We see price wars in equilibrium. They can be part of a well functioning cartel.
2. Optimal collusion may have T finite.
3. Price wars are triggered by random demand shocks that resemble firms cheating. In this model, low demand triggers price wars.
4. In a model with continuous demand shocks players will use cutoffs in the observed variable and would sometimes get away with cheating.
5. Sustaining collusion is more difficult when demand is noisier. (Here, when α is larger.)

Collusion with Imperfect Monitoring

More general analyses

Green and Porter (*Econometrica* 1984) study a Cournot-like model with continuous demand shocks.

- Firms choose quantities $q_1, q_{2,} \cdots, q_N$
- $p(Q) = (1 + \varepsilon)P(Q)$ where $\varepsilon \sim F(.)$ has $E(\varepsilon) = 1$
- G-P focus on symmetric "trigger price" equilibria – produce q^* but revert to static Cournot equilibrium for T periods if realized price is below $\hat{p}$.

Abreu, Pearce, and Stachetti (*Econometrica* 1990) provide more general result on optimal strongly symmetric equilibria in models of this type. They focus on the equilibrium payoff set and use dynamic programming arguments to characterize extreme points. In the GP model it can involve two-sided triggers to enter and exit price wars.

Fudenberg, Levine, and Maskin (*Econometrica* 1994) show that in many models we can avoid the inefficiency in the GP and APS equilibria by using asymmetric strategies in which we punish firms that appear to have cheated by transferring market share from firms that appear to have cheated to other firms, avoiding the inefficiency of price wars.

Recent private monitoring papers include Awaya-Krishna (*AER* 2016) and Sugaya-Wolitzky (*JPE* 2018). The latter notes that cartels may sometimes prefer imperfect monitoring: it makes it harder to detect deviations, but can also make it harder to identify deviation opportunities.

Collusion with Cyclical Demand

Rotemberg-Saloner (1986)

- N firms with marginal cost *c* compete as in Bertrand at $t = 1, 2, \ldots$
- Market demand is noisy: $Q_t(p)$ is $Q_L(p)$ or $Q_H(p)$, each with prob. ½.
- Firms observe period *t* demand state before choosing p_{it}.

Consider possible collusive equilibria in which firms charge p^H if demand is high and no one has cheated, p^L if demand is low and no one has cheated, and *c* if anyone has ever deviated.

Firms won't deviate in the high state if

$$\frac{N-1}{N}\pi_H(p_H) \leq \frac{\delta}{1-\delta}\frac{1}{N}\left(\frac{1}{2}\pi_H(p_H)+\frac{1}{2}\pi_L(p_L)\right)$$

Firms won't deviate in the low state if

$$\frac{N-1}{N}\pi_L(p_L) \leq \frac{\delta}{1-\delta}\frac{1}{N}\left(\frac{1}{2}\pi_H(p_H)+\frac{1}{2}\pi_L(p_L)\right)$$

Note: The continuation payoff is state-independent. The short-run gain from deviating is affected both directly by the demand state and by the state-dependent price the firms are charging.

Collusion with Cyclical Demand

Rotemberg-Saloner (1986)

The constraints required for SPE are:

$$\frac{N-1}{N}\pi_H(p_H) \le \frac{\delta}{1-\delta}\frac{1}{N}\left(\frac{1}{2}\pi_H(p_H)+\frac{1}{2}\pi_L(p_L)\right)$$

$$\frac{N-1}{N}\pi_L(p_L) \le \frac{\delta}{1-\delta}\frac{1}{N}\left(\frac{1}{2}\pi_H(p_H)+\frac{1}{2}\pi_L(p_L)\right)$$

- When $\delta \approx 1$ neither constraint is binding and firms can collude on p_L^m, p_H^m.
- When δ is smaller, the first binds. The best SPE has $p_L = p_L^m, p_H < p_H^m$.
- When $\delta = 1 - \frac{1}{N}$ both bind. $\pi_H(p_H) = \pi_L(p_L) \Rightarrow p_L > p_H$.
- For smaller δ no collusion is possible and both firms set $p_{it} = c$.

Remarks:

1. For intermediate δ the model predicts that optimal markups are countercyclical.
2. If we added imperfect observation, the intuition that collusion is more difficult when demand is high should carry over. Whether this results in lower markups, price wars being more likely, or both will be model dependent.
3. Haltiwanger-Harrington and Bagwell-Staiger discuss other cyclical models.

Private Cost Shocks

Athey-Bagwell (*RAND* 2001)

Suppose that the members of a cartel have private cost shocks in each period. Efficiency requires that the low-cost firm produce in each period. With no limits on contracting, side payments to losing bidders could be used to induce firms to reveal costs.

Without side payments we can consider dynamic equilibria.

Suppose N=2 and costs $c_{it} \in \{c_L, c_H\}$ are iid across firms and time.

1. When δ is close to one firms can achieve full collusion via strategies in which announcing low costs today gives the other firm priority to produce in the next period when both firms costs are low.
2. When δ is not as close to one, firms will price below the monopoly price and not produce efficiently. If a firm has lost too much future priority, it will be tempted to deviate from the collusive price. To make deviating less attractive, prices are reduced and production can be assigned to the high priority firm regardless of the cost realization.

Antitrust Authorities

Harrington (*RAND* 2004)

Another constraint on tacit collusion is that firms must coordinate on one of many equilibria and discussing pricing can violate antitrust laws.

Antitrust rules, however, can also have unintended consequences.

Suppose that the probability that collusion will be detected is

$\varphi(\vec{p}_1, \vec{p}_2, \ldots, \vec{p}_t) = \sum_{js} g(p_{js} - p_{js-1})$ with convex g minimized at 0.

Suppose that damages accrue according to

$$x_{jt} = \beta x_{jt-1} + \gamma f(p_{jt})$$

1. Long run prices can be higher with antitrust penalties than without because penalties provide an additional reason not to deviate.
2. Collusion may feature prices that build up over time.
3. More patient firms raise prices more slowly and reach higher price levels.

Multimarket Contact

Bernheim-Whinston (*RAND* 1990)

When firms compete in many different markets does this make collusion easier?

- Two firms choose prices in *N* markets at $t = 0, 1, 2, \cdots$.
- Prices are chosen simultaneously for all markets. All period *t* outcomes observed before choosing period *t* prices.

Firms can be punished in all markets in response to a deviation. But they can also gain by deviating simultaneously in all markets.

<u>Observations</u>:

1. When markets are identical (or with Bertrand competition in all markets) multimarket contact does not make collusion easier or harder.
2. When markets are asymmetric, firms can use spare punishment capacity in the easy-to-collude markets, to enable collusion in other markets.

Markov Models of Price Competition

In some markets in which price is the main strategic variable, and is chosen repeatedly, there are important dynamic considerations, e.g. staggered/costly price changes, investments in quality improvement/cost reduction, etc.

Researchers often wish to ignore repeated game equilibria and focus on something more like static Nash equilibrium. The standard way in which this is done is to assume the firms play a ***Markov perfect equilibrium (MPE).***

This is especially common in empirical IO.

In a MPE all player's strategies are assumed to depend only on payoff-relevant variables, e.g. they can depend on current costs, but not on the price that some other firm charged seven periods ago.

The MPE concept is often treated as a means of obtaining an unique equilibrium, there is no guarantee of uniqueness.

Markov Models of Price Competition

Maskin and Tirole provide interesting examples of MPE that can arise in an alternating move price competition game.

- Firm 1 chooses prices at t=1, 3, 5, … and firm 2 chooses prices at t=2, 4, 6, … All prices are observable.
- Suppose $D(p) = 1 - p$ and prices are restricted to $\left\{0, \frac{1}{6}, \frac{2}{6}, \frac{3}{6}, \frac{4}{6}, \frac{5}{6}\right\}$.

<u>Observations</u>:

1. There is no MPE in which the firms set p=c. Hence, all MPEs have positive profits.

2. If δ is close enough to one the model has a fully collusive equilibrium with price war punishments: s(5/6)=s(4/6)=s(3/6)=1/2, s(2/6)=1/6, s(1/6)=$\alpha \cdot {}^{1}/_{2} + (1 - \alpha) \cdot {}^{1}/_{6}$. s(0)=1/2.

3. There are also other partially collusive equilibria. For example, an “Edgeworth Cycle” is: s(5/6)=4/6, s(4/6)=3/6, s(3/6)=2/6, s(2/6)=1/6, s(1/6)=0, s(0)=$\alpha \cdot {}^{5}/_{6} + (1 - \alpha) \cdot 0$.

Algorithmic Collusion

Calvano, Calzolari, Denicolo, Pastorello (*AER* 2020)

Many online businesses use automated pricing tools to update many prices: hotels, airlines, and car rentals selling though travel sites, Amazon sellers, Amazon vs. Target, etc.

- With complete information there are many repeated-game equilibria, whether we model as simultaneous move or Markov with occasional price changes.
- Both demand and rivals' repeated game strategies are naturally unknown. Firms will want to design pricing algorithms to learn to profit-maximize.

There is a large theory literature from the 1990s on whether naïve learning algorithms lead play to converge to some equilibrium and which equilibria are selected in games with multiple equilibria. See Fudenberg and Levine (MIT Press, 1998).

CCDP use simulations to investigate which repeated game equilibria (if any) emerge when firms adopt "Q-learning" pricing algorithms.

Algorithmic Collusion

Calvano, Calzolari, Denicolo, Pastorello (*AER* 2020)

Q-learning is a standard approach for unknown dynamic environments:

- Profits $\pi(s,a)$ depend on observable state s, action a, and random shocks.
- State transition process $Prob(s'|s,a)$ also unknown.

Dynamic programming suggests we can learn to play optimally against exogenous uncertainty by learning the discounted value function Q defined by

$$Q(s,a) = E\big(\pi(s,a)\big) + \delta E_{s'|s,a}\left(\max_{a'} Q(s',a')\right)$$

When $|S|$ and $|A|$ are finite this is learning the elements of an $|S| \times |A|$ matrix.

A standard approach for this problem is:

- Conjecture some initial matrix Q_0.
- At each t choose $a_t = arg\max_{a} Q_t(s_t,a)$ with prob. $1-\varepsilon_t$. Otherwise a_t random.
- Update $Q_{t+1}(s_t,a_t) = (1-\alpha)Q_t(s_t,a_t) + \alpha\left(\pi_t + \delta \max_{a'} Q_t\,(s_{t+1},a')\right)$ using observed π_t and s_{t+1}. Rest of matrix unchanged.

With ε_t that decline appropriately, Q-learning converges to optimal play against exogenous uncertainty.

Algorithmic Collusion

Calvano, Calzolari, Denicolo, Pastorello (*AER* 2020)

CCDP consider a repeated price competition game.

- Consumers have discrete choice preferences with $u_{ij} = v_j - p_j + \mu\varepsilon_{ij}$.
- Base model has 2 firms, identical costs, 15 possible prices, logit errors, $\delta = 0.95$.
- Q learning models use last period's prices as s, consider 10000 possible choices for algorithm parameters (α, β), and declare learning to have converged when optimal play is unchanged for 100,000 consecutive periods.

Main results include:

1. Profits are usually 70-90% of the way from static Nash to monopoly.
2. Terminal strategies ε Nash, not SPE. Punishments usually finite price wars.
3. Prices are above static Nash even when $\delta \approx 0$. Increase in δ for $\delta > 0.5$.
4. Profits still over 50% of way to monopoly for N=4.

Algorithmic Collusion

Calvano, Calzolari, Denicolo, Pastorello (*AER* 2020)

Paper is thoughtful about presenting results from many simulated models.

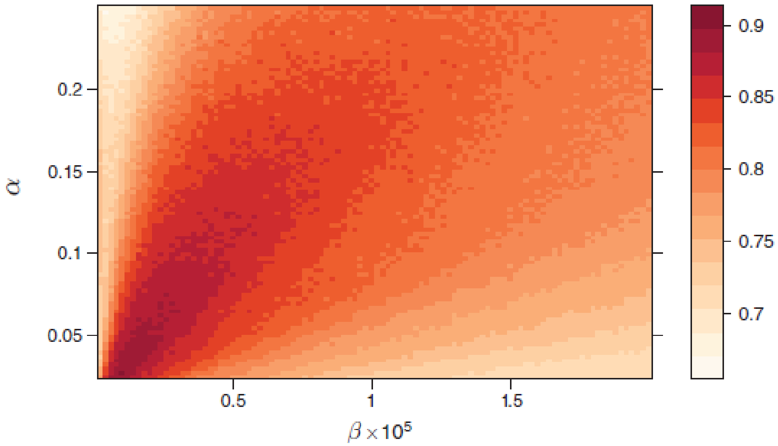

FIGURE 1. AVERAGE PROFIT GAIN Δ FOR A GRID OF VALUES OF α AND β

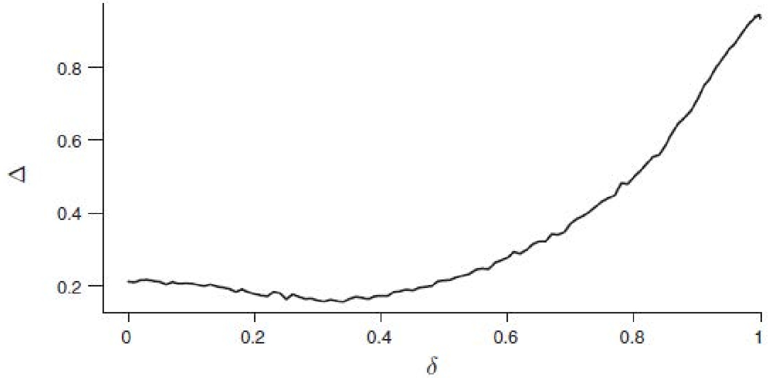

FIGURE 3. THE AVERAGE PROFIT GAIN Δ AS A FUNCTION OF THE DISCOUNT FACTOR δ IN OUR REPRESENTATIVE EXPERIMENT

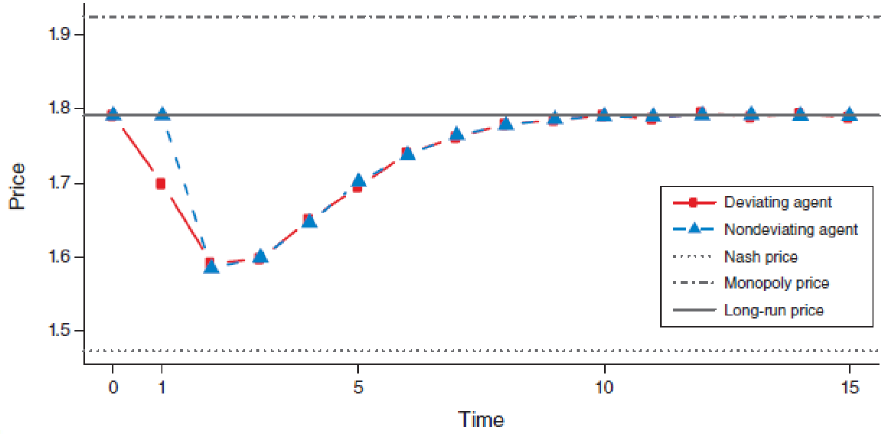

Notes: Prices charged by the two algorithms in period τ after an exogenous price cut by one of them in period $\tau = 1$. The forced cheater deviates to the static best response, and the deviation lasts for one period only. The figure plots the average prices across the 1,000 sessions. For sessions leading to a price cycle, we consider deviations starting from every point of the cycle and take the average of all of them. This counts as one observation in the calculation of the overall average.

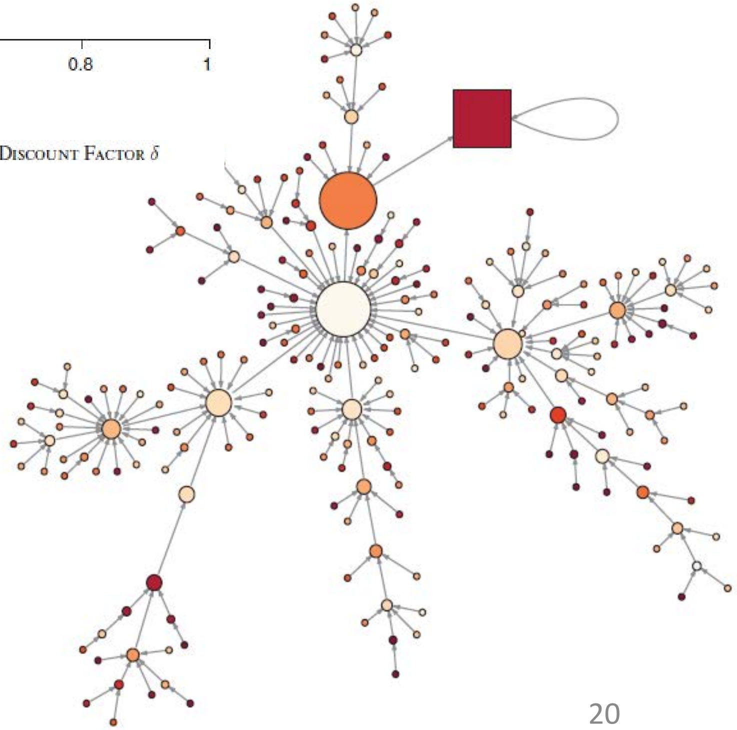

Algorithmic Collusion

Calvano, Calzolari, Denicolo, Pastorello (*AER* 2020)

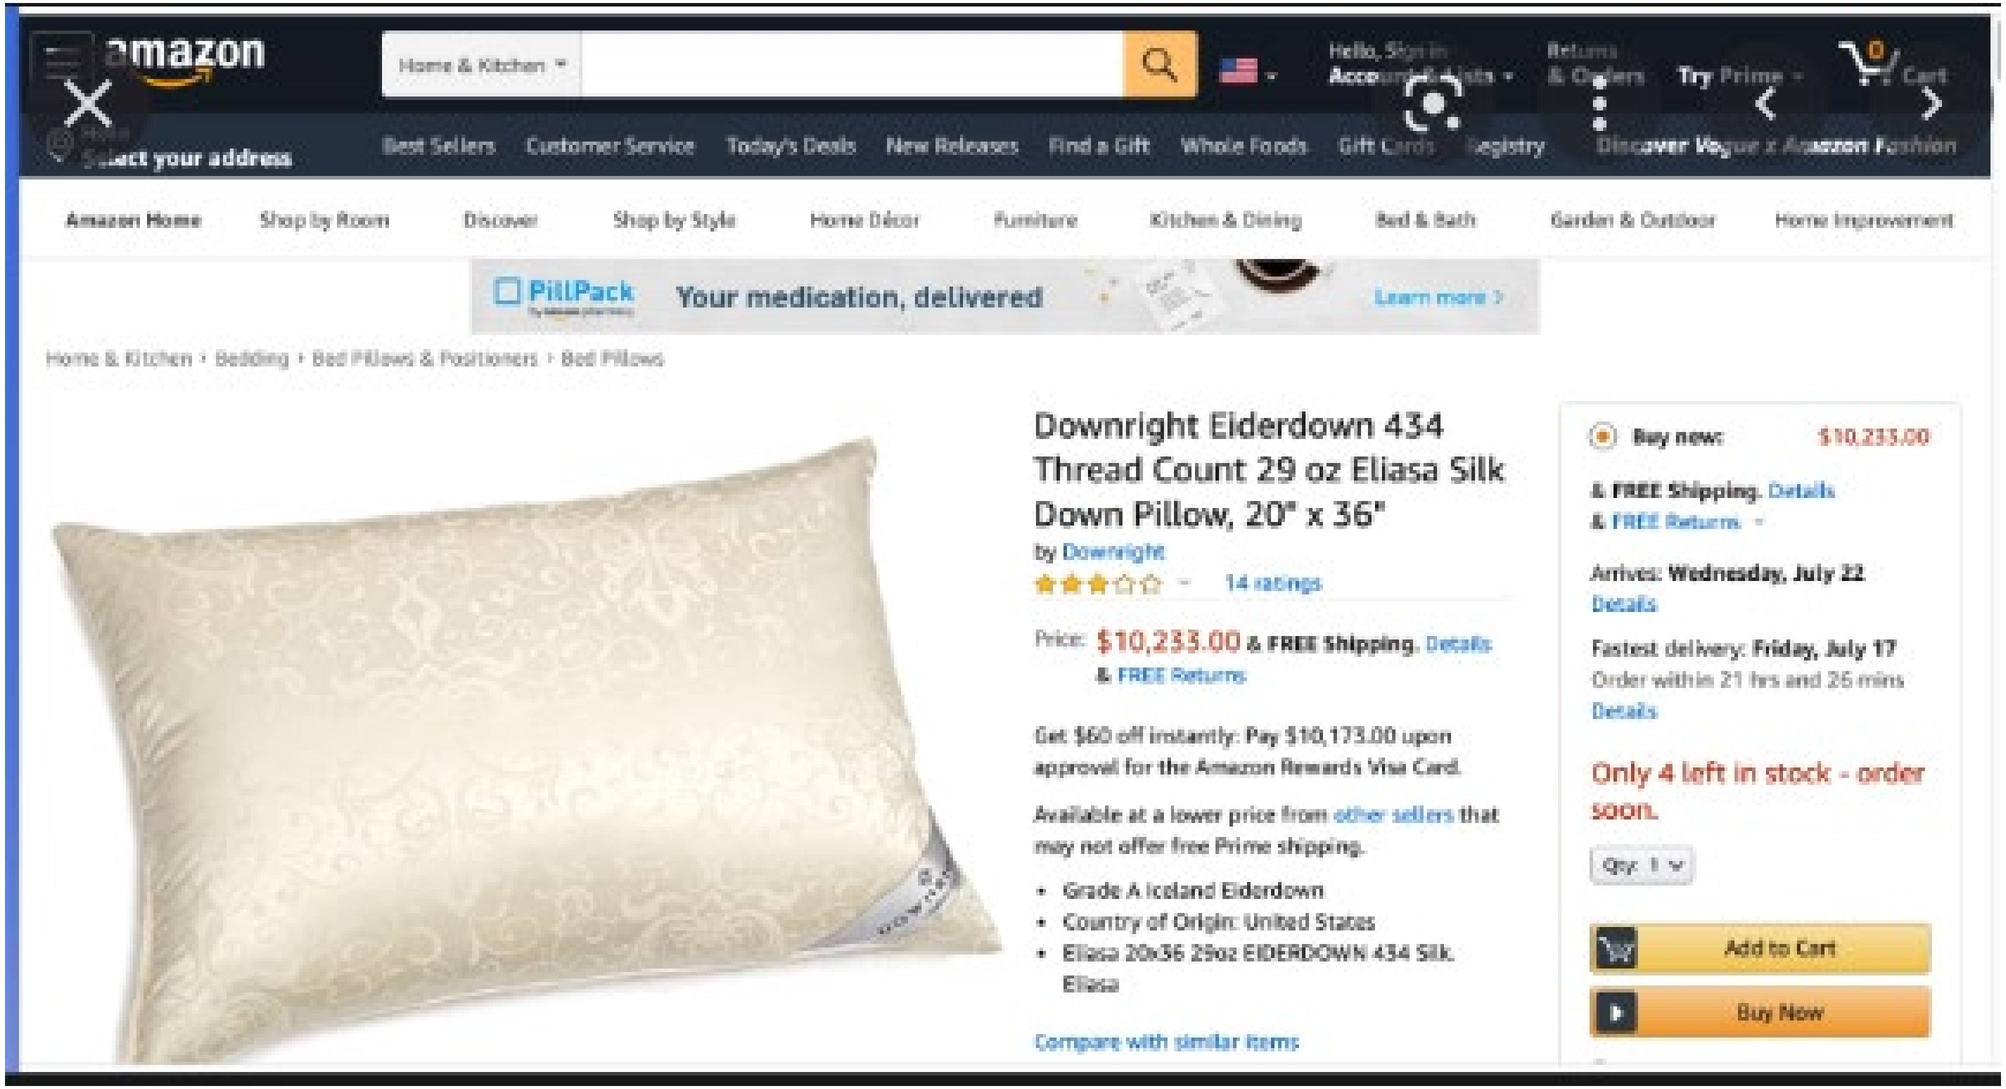

On Wednesday I’ll discuss some empirical papers on dynamic collusion including

- Porter
- Ellison
- Wang

See you then!

Markov Equilibrium

- Consider a multi-stage game with observed actions – i.e., players simultaneously choose actions at each stage, and then all actions are observed prior to the next stage*:*

  Two histories h_{t-1} and h'_{t-1} are ***Markov equivalent*** if for any two sequences $\{x_\tau\}_{\tau\geq t}$ and $\{x'_\tau\}_{\tau\geq t}$ of present and future action profiles for the players, for all *i* we have

  $$u_{it}(\{x_\tau\}_{\tau\geq t}, h_{t-1}) \geq u_{it}(\{x'_\tau\}_{\tau\geq t}, h_{t-1}) \Leftrightarrow u_{it}(\{x_\tau\}_{\tau\geq t}, h'_{t-1}) \geq u_{it}(\{x'_\tau\}_{\tau\geq t}, h'_{t-1})$$

  where $u_{it}(\cdot)$ is player *i*'s continuation payoff.

- A ***Markov strategy*** specifies the same action for any two Markov equivalent histories.

- A ***Markov perfect equilibrium*** is a subgame perfect NE in which all players use Markov strategies.