Lecture 1:

1. <u>Stigler, The economics of information, *Journal of political economy*:</u>
    a. Model set-up:

    Let's abstract from the process that generates prices, assume uniform distribution $p \sim U(0, 1)$

    What is the result of searching for a buyer, given the prices posted by sellers? Let's write down the problem

    Uniform distribution of prices $F(p) = p$

    The buyer samples n sellers

    The cumulative distribution of the minumum price is:

    $$G(p_{min}) = [1 - F(p)]^n = \left[\int_p^1 dx\right]^n$$

    The distribution function of the minimum price is:

    $$g(p_{min}) = n(1 - p)^{n-1}$$

    The average minimum price is:

    $$\frac{1}{n+1}$$

    Graphically:

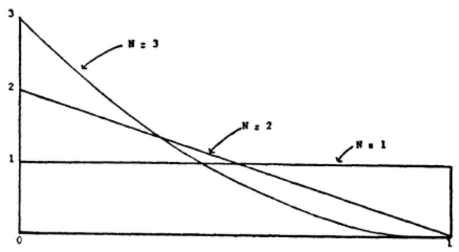

    b. Insights:
        i. Searching for prices reduces the expected price
        ii. Search has diminishing returns
        iii. If the expected gains from search are greater, the greater is the dispersion in prices
    c. Consequences:
        i. searching for prices is costly - ways to reduce the fog surrounding prices:
            1. bringing buyers and sellers to the same place
            2. advertising
            3. have specialized traders
            4. technology enables to put previous categories together

2. <u>Salop and Stiglitz, :</u>
    a. Question:
        i. How far should we go from perfect information to break perfect competition outcomes?
        ii. Previous analysis did not take into account the cost consumers incur to search and the incentive of sellers to offer a certain price
    b. Assumptions:
        i. Consumers know the general distribution of prices
        ii. They can become fully informed by paying a cost
    c. Model set-up:

Large number L of consumer for a homogeneous good purchased once
Maximum price a consumer will pay is u
Finite number of sellers n in the market, asking prices $\underline{p} = p_1, p_2, \ldots, p_n$
Consumers know the prices but do not know the location (or the identity) sellers $\underline{l} = l_1, l_2, \ldots, l_n$

Consumer i can pay the cost c^i to become informed, i.e., to know $\underline{l}$. Then:

- Informed consumers: buy from the cheapest seller
- Uninformed consumers: buy from one seller at random

Only two levels of cost c: a fraction α of consumers has low information cost c^1, and the remaining fraction $(1-\alpha)$ has high information cost c^2, with $c^1 < c^2$

Consumers trade-off between the costs and the benefits of becoming informed

If consumer i pays the cost, the total expenditure will be:

$$E_S^i = p^{min} + c^i \tag{2}$$

If consumer i does not pay the cost, the expected total expenditure will be:

$$E_N^i = \bar{p} = \frac{1}{n}\sum_j^n p_j \tag{3}$$

Total expenditure depends on the search strategy!

When is the consumer buying the information? Assuming risk-neutrality if:

$$E_S^i < E_N^i \iff p^{min} + c^i < \bar{p} \tag{4}$$

What are sellers going to do?

First, they compete with each other on price. Thus, each seller j takes the price of others as given and maximizes profit:

$$\max_{p} \pi^j(p|p_{-j}), \quad \text{where } \underline{p}_{-j} = \{p_1, p_2, \ldots, p_{j-1}, p_{j+1}, \ldots, p_n\} \tag{5}$$

Second, the prices of other firms, and the search decisions of consumers enter into the profit function (through the demand function), because a consumer i searches the info if:

$$c^i < \bar{p} - p^{min} \tag{6}$$

and each firms knows that, with its decision, it is going to affect both:

$$\bar{p} = \frac{1}{n}p_j + \frac{1}{n}\sum_{i \neq j} p_i \tag{7}$$

$$p^{min} = min\{p_j, \underline{p}_{-j}\} \tag{8}$$

d. Outcomes:
   i. 3 possible equilibria
      1. one price at competitive level
      2. one price close to monopoly price
      3. a two-prices equilibrium

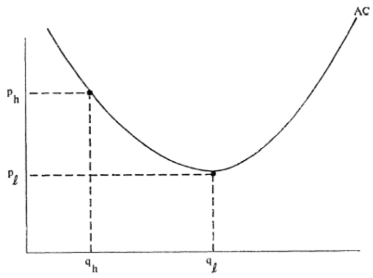

e. Insights:
   i. price dispersion will be constrained between the competitive price and the monopolistically price
   ii. magnitude of price spread depends on cost to acquire info and the degree of scale economies in the supply side
   iii. the U-shaped average cost means that there a too many small firms in equilibrium
   iv. the economy does not produce information efficiently because there is an informational externality at work + production is not efficient

3. <u>Varian:</u>
   a. Question:
      i. sellers also move price dynamically, i.e. product prices alternate between regular prices and sale prices
      ii. imperfect market information? Can we derive an equilibrium price distribution?
   b. Model set-up:
      i. similar to previous
      ii. firms set price for the good weekly
      iii. firms have decreasing cost curves
      iv. free entry in the market, which implies zero profit in equilibrium
   c. Outcomes:
      i. firms have an incentive to offer low price, because you buy some more probability to be the cheapest, but at the cost of reducing your margin
      ii. high prices result in higher profit per unit sold

The characterization of the equilibrium requires to show:

- The firms' equilibrium distribution of prices $F(p)$ is bounded between a low and a high price, which give in expectation the same profit (as any other intermediate price)
- There are no mass points in $F(p)$
- $F(p)$ is increasing in price
- $F(p)$ does not have gaps

The equilibrium $F(p)$ is:

$$F(p) = 1 - \left(\frac{\pi_f(p)}{\pi_s(p) - \pi_f(p)} \right)^{\frac{1}{n-1}} \tag{9}$$

iii.

iv. low prices (sales) and high prices (standard?) are posted more frequently than intermediate prices

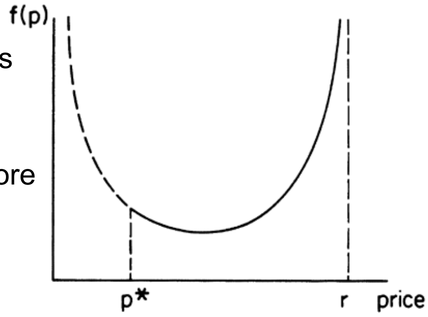

v. what if search is sequential or block search? - the problem is an optimal stopping problem: stops with searching if expected gains of sampling another store is lower than this cost

vi. market equilibrium is price dispersion

4. <u>Consumer search and prices in Automobile industry:</u>
   a. Outcomes I:
      i. full information is not compatible with the evidence
      ii. search costs are important, reduce price elasticity and increase mark-ups
      iii. people do search, but not in the same amount
      iv. drivers also ask for drive test
      v. factors influencing search intensity:
         1. higher income and new cars induce more search
         2. distance to dealers reduces

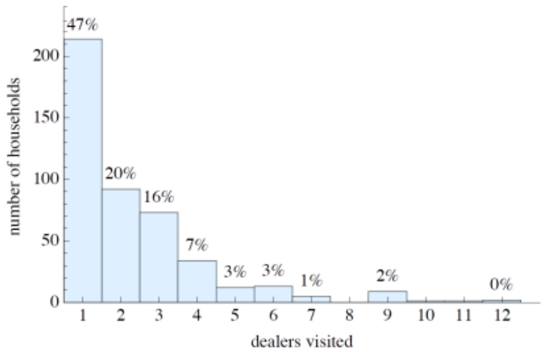

   b. Model set-up:
      i. sequential search that stops when the expected gain by visiting another dealer is lower than the cost of making the visit
      ii. survey takes place after purchase
      iii. questions about type of car, search process, demographic characteristics and dealership locations. Also info about market share of brand and models' characteristics

iv. optimal search:

Utility is at the model level, but search is at the dealer level, thus, after a visit, a consumer updates his/her utility:

$$U_{if} = \max_{j \in G_f} \{u_{ij}\} \tag{11}$$

i.e., by visiting the dealer f, the consumer learns what he/she can get from that dealer, and this utility corresponds to the best model (from the perspective of consumer i)

Searching updates the information possessed by consumer i, who then decides whether to continue with another visit or not

The expected gains from another round of search (given a best value found so far of r) are:

$$H_{if}(r) \equiv \int_r^\infty (z - r) dF_{if}(z) \tag{12}$$

And search continues as long as $H_{if}(r) > c_{if}$, i.e., until the gains are larger than the costs c_{if} of visiting the next dealer

Key idea: by calculating the value at each round of search and using the information of when search stops, one can invert the relation between gains and search costs to get an estimate of the latter!

$$r_{if} = H_{if}^{-1}(c_{if}) \tag{13}$$

v. Estimation requires to solve the search reservation value and then incorporate that into the demand function that delivers the buying probability of each car model (complex estimation)

vi. sellers' equilibrium takes into account the level of competition and the markups also respond to demand elasticity

c. Outcomes II:

i. car manufacturers have different dealer networks which lead to different exposure of consumers to the distribution network: higher distance implies greater search costs and similar households that live in a different location have different costs to visit 1,2,3.. dealers

ii. markups are much higher when we take into account the corrected demand elasticities - this implies that profits are also greater

iii. a reduction in search frictions would reduce price, but not at the cost of lower sales, that would actually increase. Industry profits would go up

Lecture 2:

1. Brown and Goolsbee, *Does the internet make markets more competitive?*:
   a. Empirical study on life insurance market:
      i. goal to estimate effect on Internet diffusion on market price
      ii. market with homogeneous good with easy characteristics
      iii. sizable reduction in search costs

iv. competition between firms and incentive to differentiate and price discriminate
v. only estimate term insurance market as they are easier to compare
vi. online tool, but still offline purchase

b. Predictions:
i. equilibrium price distribution → price dispersion in equilibrium
ii. larger share of buyers online shifts the distribution down and decreases average price
iii. effect of varying share of buyers online on price distribution is non-linear

c. New study including direct information on search activity:
i. model internet usage, compute predicted share of internet search
ii. insurance price comparison website appeared in 1996
iii. online tool, but still offline purchase

d. Outcomes:
i. strong decline in term insurances, but not for whole insurance contracts (no possibility to compare online)
ii. stronger decline in high internet take-up states, in areas with more high skill workers and younger population
iii. in regression model is main driving factor the % use of internet = internet diffusion
iv. increasing share of internet usage by 10 percentage points reduces prices by 1.5 to 4.5%
v. online usage generated an annual increase in consumer surplus which can be larger in the long run
vi. prices going down in these markets is a positive spillover to the rest of society, also to people who are not searching
vii. lower price is compatible with both users getting lower prices and sellers posting lower prices

2. <u>Baye and Morgan, *Information gatekeepers on the internet and the competitiveness of homogeneous product markets*</u>:
a. Model set-up:
i. several local market each served by a local monopolist
ii. virtual markets where firms can list their offers and again consumers from own market and other markets
iii. virtual markets sets fee both to sellers and buyers

Local markets $1, 2, \ldots, n$ populated by consumers (each market has $1/n$ of them)

Each local market has 1 local monopolist, thus n firms in total

Homogenous product, produced at constant marginal cost c (normalized to 0)

Virtual market enables consumers to buy from firms located in other markets at 0 extra-cost

Access to the virtual market subject to an advertising fee ϕ for sellers and a subscription fee κ for consumers

iv.

Timing:

- First, the gatekeeper announces the fees ϕ and κ
- Second, both firms and consumers decide whether to go online
  - If they go online, firms also decide the price to post

v. - Third, consumers shop

vi. consumers decide whether to go online or not (fraction mu)

vii. assumption: if no price is available in the virtual market, then they buy from the local firm

viii. there is a small search cost to visit the local seller e (epsilon), which is small enough to make the monopoly price still a viable price (surplus S(r) > e)

ix. What would a firm do without a virtual market? Exploit their market power

Without a virtual market firms would charge monopoly price r and earn monopoly profit on the share of consumers who live in the city:

$$\pi(r) = (r - c)\frac{q(r)}{n} \tag{2}$$

NOTE: This is also the price that firms not advertising charge in equilibrium (see Proposition 2)

x. Firms' decision:

Suppose $n = 2$, what is the expected profit of the firm?

- This depends on the decisions of firms, consumers and the gatekeeper
- Firms can go online (α is the propensity to do it), and the price p to charge

If it does not advertise:

$$E\pi_i(r, N) = (1-\alpha)\frac{\mu}{2}\pi(r) + \frac{1-\mu}{2}\pi(r) \tag{3}$$

If it does advertise, charging a price $p_i \in [c, r]$:

$$\begin{aligned} E\pi_i(p, A) =& \alpha\left[\mu p\Big(1 - F(p)\Big) + \frac{1-\mu}{2}p\right] + \\ &+ (1-\alpha)\left[\mu p + \frac{1-\mu}{2}p\right] - \phi \end{aligned} \tag{4}$$

More generally, the two expected profits are:

$$E\pi_i(r, N) = (1-\alpha)^{n-1}\frac{\mu}{n}\pi(r) + \frac{1-\mu}{n}\pi(r) \tag{5}$$

and

$$E\pi_i(p, A) = \mu\pi(p)\Big(1 - \alpha F(p)\Big)^{n-1} + \frac{(1-\mu)}{n}\pi(p) - \phi \tag{6}$$

xi. in equilibrium, all posted prices deliver the same profit. Thus the profit in the case of advertising are the same as in the case of not advertising and charging monopoly price

$$E\pi_i(p, A) = E\pi_i(r, N)$$

and thus, combining (6) and (5) and rearranging, we get to the equation that describes the firms' propensity to participate in the virtual market:

$$\alpha = 1 - \left(\frac{n\phi}{(n-1)\mu\pi(r)}\right)^{\frac{1}{n-1}} \quad (7)$$

xii.

xiii. the propensity to join the virtual market
   1. declines with the advertising fee (phi)
   2. increases with the share of consumers who go online
   3. decreases with the number of firms

b. Outcomes:
   i. equilibrium prices posted on the virtual market are distributed according to F(p) and importantly support [p_0, r]: prices online are lower on average
   ii. consumers trade-off the costs and benefits to go online and they depend on: the fee *kappa*, the share *mu* of other consumers online and the share of firms *alpha* that go online
   iii. gatekeeper? the decision on the fees is crucial for the participation of the agents on the virtual market:
      1. a fee to firms so to induce optimal firm participation (not full): does not want to reduce too much price dispersion, because that would eliminate the incentive to search
      2. a fee to the consumers which is as low as possible: keep lowering the fee until full participation is reached

c. Insights:
   i. virtual markets exist when participation costs are not extreme
   ii. prices decrease with the establishment of a virtual market: the possibility to shop outside my city puts pressure on prices as it increases competition between firms
   iii. price dispersion does not disappear: firms participation is not full, due to the incentives of the gatekeeper to raise the advertising fee
   iv. fees for consumers will be as low as possible to induce full participation
   v. pricing will become more competitive and fees charged by the gatekeeper exceed the socially optimal level, thus inducting suboptimal provision online (due to non-full firm participation)

3. <u>Brynjolfsson and Smith, *Frictionless commerce?*</u>:
   a. Question and approach:
      i. goal: compare pricing between online and offline retailers
      ii. homogeneous products: CDs and books
      iii. internet retailers, hybrid retailers and conventional (only physical stores)
      iv. the approach is to select largest retailers and/or representative retailers

        v. first half in each category are best sellers (Billboard, NYT best sellers) and second half is a miscellaneous of random titles, with the constraint of being available on both channels
        vi. question: is price online lower, higher or the same as offline?
    b. Empirical results:
        i. in all cases, price with and without shipping and tax costs (for online) and physical transportation costs for offline
        ii. mean price: online is most of the time cheaper, also if online is not weighted by traffic (just posted prices)
        iii. minimum price: online most of the time cheapest alternative
        iv. mean price (weighted), including ancillary costs: online still cheaper on average
        v. with more than one gatekeeper and local competition between stores offline:
            1. price dispersion should arise in both markets
            2. unclear if dispersion online is larger than offline
        vi. price dispersion is found in both online and offline:
            1. best vs. worst price
            2. trimmed range: second best vs. second worst price
            3. standard deviation
        vii. mixed evidence: books' prices tend to be more dispersed online but the reverse is true for CDs
4. <u>Baye and Morgan, *Price dispersion in the small and in the large*</u>:
    a. Question on online price dispersion:
        i. why is price dispersion different in the two cases? is it a disequilibrium phenomenon or is price dispersion an equilibrium phenomenon that depends on market fundamentals
    b. Predictions:
        i. price dispersion should arise on the platform and is not transient (it is no disequilibrium phenomenon)
        ii. the gap between the two lowest listed prices is positive
        iii. price dispersion is greater in the small (with fewer firms than in the large (when a large number of firms list prices)
    c. Model set-up
        i. data for a price comparison website specialized in consumer electronics (shopper.com)
        ii. each observations consists of rank of the product based on popularity; and price posted by each firm
        iii. information is collected twice a day (1am - 2pm)
        iv. products include popular electronics: digital cameras, software, printers, notebooks…
        v. strength of the data:
            1. homogeneity of products and composition, sample size + frequency, duration of the study
            2. relatively small incidence of shipping costs on total price
            3. fee structure of the gatekeeper: fee to be admitted to the platform + monthly fee + fee for each posted price → no incentive to post unrealistic prices

        4. technology to verify that posted price corresponds to own price
    vi. weakness: lack of information on the quantities → focus on the % gap between two lowest prices

d. Outcomes:
    i. more popular products tend to be cheaper
    ii. more sellers = more price competition = lower price
    iii. single price listings: only one firm posts a price on a given date; multiple price listings: more than 1 firm post a price on a given date
    iv. products with multiple price listings have lower average price and lower average minimum price than procuts with single price listings (not conclusive)
    v. average % gap between two lowest prices is maller for popular products who also have a larger number of sellers

	All Product Ranks	Product Ranks 1 – 250	Product Ranks 251 – 500	Product Ranks 501 – 750	Product Ranks 751 – 1000
Total Number of Prices					
Multi-Price Listings	3,925,947	1,202,912	960,709	904,256	858,070
Single-Price Listings	13,743	2,846	3,416	3,785	3,696
Average Price in					
All Listings	$513.23	$472.73	$494.91	$529.60	$555.64
	(882.8)	(665.2)	(838.3)	(1,039.6)	(941.7)
Multi-Price Listings	$491.64	$461.07	$476.41	$486.56	$543.08
	(760.8)	(590.7)	(706.1)	(820.0)	(892.0)
Average Minimum Price in					
All Listings	$457.62	$417.94	$442.78	$475.77	$493.93
	(818.7)	(611.9)	(781.3)	(980.0)	(855.4)
Multi-Price Listings	$432.47	$403.40	$420.97	$428.91	$477.09
	(678.2)	(525.1)	(630.9)	(733.7)	(792.4)
Average Number of Firms in					
All Listings	17.27	21.17	16.90	15.91	15.12
	(11.7)	(14.1)	(10.8)	(10.4)	(10.0)
Multi-Price Listings	18.32	22.23	17.91	16.97	16.10
	(11.3)	(13.7)	(10.3)	(9.9)	(9.6)

    vi. level of competition influences our values of interest: frequency distribution still high at 20 firms (2.8%)
    vii. online price dispersion trend:
        1. average % gap = difference between cheapest and second cheapest, divided by the baseline price
        2. no trend towards decreasing in dispersion → stays stable
        3. bv 20% of products have 5% gap

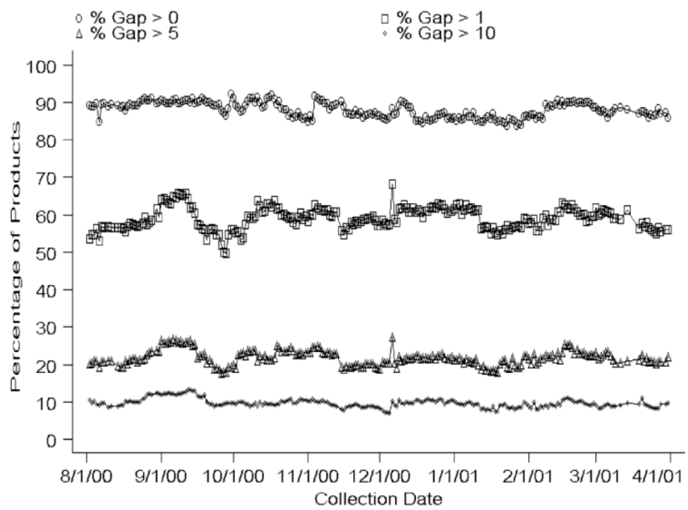

viii. no evidence that price dispersion is a disequilibrium phenomenon
ix. in equilibrium, the amount of price dispersion should depend on market structure
x. on the platform (at least search costs are the same for products with few or many sellers
xi. market structure: showing gap between lowest and second lowest price → lower and lower prices prices when more competitors (MC); price dispersion doesn't go away

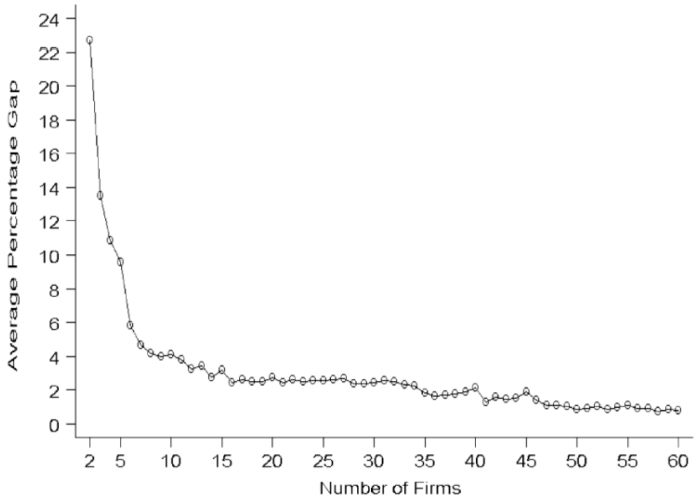

xii. the level of price dispersion seems to be strongly related with the number of firms offering the item on the virtual market (descriptive result - control for confounding effects (popularity, trends in participation of firms))
xiii. the role of market structure in not affected by the inclusion of controls in the regression: results are consistent with the theory
xiv. digitalisation has changed things, competition has increased, but old lessons remain overall valid and competition is important even in online world to benefit consumers

Lecture 3:

1. Brynjolfsson, Hu, and Simester, *goodbye pareto principe, hello long tail*:
   a. Goal:
      i. many industry observers noticing that niche products sell non-trivial amount online
      ii. having more products does not imply more sales for the niche! the pareto principle can still apply, just on a larger inventory
      iii. prove the emergence of the long tail and unveil the role of the demand-side factors and in particular the product search and recommendation tools
   b. Model set-up:
      i. moderate price women's clothing: the catalog (40% of sales) is closer to the offline store where consumers can only see the shelves, while online (60% of sales) they can use the recommendation and discovery tools

        ii. the two channels are identical on all dimensions but the recommendation tool

    c. Outcomes I:

        i. lorenz curve of the internet channel lies above the one of the catalog channel → 60% of catalog sales comes from 20% of the products, whereas this is 25% in internet channel

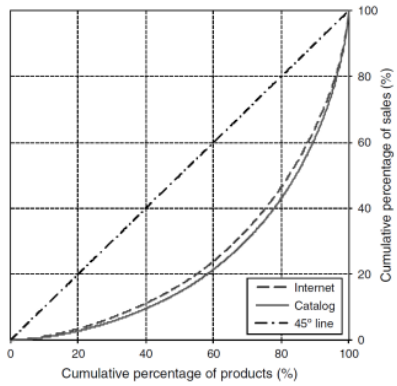

        ii. niche is gaining share and catalog is more concentrated
        iii. is the difference statistically significant?

$$ln(Sales_j) = \beta_0 + \beta_1 ln(Sales\ Rank_j) + \varepsilon_j$$

        iv. sales decrease faster with rank in catalog → top selling products sell more
        v. in pooled data: introduce the interaction between the channel *Internet* and the *Sales rank* to capture a difference in the rank-sales relation between channels

    d. Adjusted model set-up:
        i. shoppers might be different: online shoppers might be more interested in niche products than those who browse the catalog
        ii. solution: propensity score matching → finding counterfactuals (via a synthetic score)
    e. Outcomes II:
        i. internet shoppers are richer, more educated, younger and less likely to be female
        ii. still an increase in niche sales
        iii. must be something related to consumers at work
        iv. the use of search tools and recommendation systems is an explanation:
            1. directed search
            2. indirected search
            3. recommendation
        v. nondirected search and recommendation tools boost the sales of niche products (focus on bottom 50% of products)
    f. Summary:
        i. long tail in digital markets is real
        ii. it is possible because of supply-side factors: inventory costs decrease
        iii. it arises because of demand-side factors: lower search costs

2. <u>Brynjolfsson, Hu and Smith, *Consumer surplus in the digital economy*</u>:
    a. Motivation and goals:
        i. long tail signals the interest of consumers for niche products whose demand increases
        ii. variety online is much larger than offline and consumers make use of digital tools
        iii. can we quantify the associated welfare gains?
    b. Question:
        i. how can we compute the benefit from a larger selection of products?

ii. comparing the consumer surplus with and without the excluded products
iii. change in consumer surplus due to a change in prices
iv. compensating variation (=CV) is the monetary transfer that would compensate the consumer for the change in prices holding utility constant: $CV = e(p_0, u_0) - e(p_1, u_0)$
v. our problem is to compare:
    1. an online store with only the selection of products that are available offline
    2. an online store with all online products

c. Model set-up:

i. In the case of estimating the gains of product variety online, the compensating variation can be written as:

$$CV = e(p_{e0}, p_{n0}, u_1) - (p_{e1}, p_{n1}, u_1) \tag{4}$$

i.e., the expenditure in existing and new (non existing ex-ante) products, minus the expenditure in already existing and new (ex-post) products

The total effect can then be broken in two parts:

$$CV = \underbrace{[e(p_{e1}, p_{n0}, u_1) - (p_{e1}, p_{n1}, u_1)]}_{\text{Introduction of more variety}} + \\ + \underbrace{[e(p_{e0}, p_{n0}, u_1) - (p_{e1}, p_{n0}, u_1)]}_{\text{Price change pre-existing products}} \tag{5}$$

Demand function:

$$x(p, y) = Ap^{\alpha} y^{\delta} \tag{6}$$

where α is price elasticity and δ is income elasticity

Under some assumptions on the price movements

After using the standard results to derive the Hicksian demand curve, the CV is:

$$CV = \left[\frac{1-\delta}{1+\alpha} y^{-\delta}(p_{n0}x_0 - p_{n1}x_1) + y^{1-\delta}\right]^{\frac{1}{1-\delta}} \tag{7}$$

which, after assuming that income elasticity can be ignored for the items that can be bought on this platform (i.e., $\delta = 0$), simplifies to:

$$CV = \frac{p_{n1}x_1}{1+\alpha} \tag{8}$$

ii. empirical analysis is carried on data from a book retailer which has both online and offline stores: focus on obscure books that before the internet where in the tail of the sales distribution
iii. these can be considered as new products because while being readily available through the internet, the cost of acquiring them would be extremely high offline

iv. compute *alpha*:
  1. publishers and retailers have agreements on final price
  2. retailers do not keep unsold copies
  3. assumptions: final price is a markup over the cost and the price elasticity faced by the retailer and publisher is the same
  4. it's an industry structure where publisher-retailer(s) behave as if vertically integrated
  5. retailer has info on *alpha*, but only on publisher level:
     Because the publisher of a particular title has total control over title's price, we can apply the usual formula relating markup and demand elasticity:
     $$\frac{p_i - C_i}{p_i} = -\frac{1}{\alpha_{ii}}$$
     where α_{ii} is the (own) price elasticity for title i

v. compute *x_i*:
  1. $ln(Quantity_i) = \beta_1 + \beta_2 ln(Rank_i) + \varepsilon_i$
  2. the quantity sold of obscure books is estimated: a book with rank 10 gets 5,000 sales per week and a book with rank 100,000 gets 1.6 sales per week
  3. aggregate share of sales of obscure books on the platform at different product ranks: (physical stores stop at threshold, based on superstore offline with 100,000 titles)

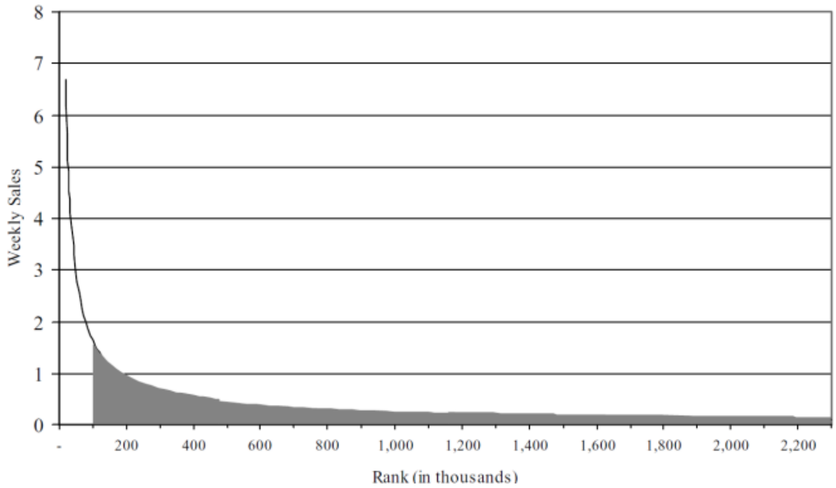

d. Outcomes:
  i. based on formula 8, able to compute the increased consumer welfare
  ii. gains from variety are substantially larger than those from offline vs. online competition

3. <u>Waldfogel, *How digitization has created a golden age of music, movies, books and television* + Aguiar and Waldfogel, *Quality predictability and the welfare benefits from new products*:</u>
  a. Motivation:
    i. many industries have been heavily affected by digitization: lower search costs, lower costs to bring products to market, low costs to make digital copy
    ii. question: can digitization bring more/better products to the market?

b. Empirical study:
    i. lower costs should benefit consumers and producers and bring innovation → long tail in production
    ii. in a deterministic world: a producer would only invest on the projects that are profitable; in this world: a reduction in costs to bring product to market increases variety
    iii. in a world of uncertainty: cost reduction can bring to market products that turn out to be successes while deemed to fail ex-ante; extreme case: the tail of new products has the same quality as the core of old products
    iv. goal: estimate the gain in consumer surplus due to the arrival of products that would have been discarded otherwise but turn out to be valuable
c. Model set-up:
    i. demand for songs:

$$u_{ij} = \underbrace{x_{jc}\beta - \alpha p_{jc} + \xi_{jc}}_{\delta_{jc}} + \varsigma_i + (1-\sigma)\epsilon_{ij}$$

    ii. prediction of quality:
        1. with estimates of demand model → forecasting model (quality ex-ante)
    iii. supply of songs and welfare: determine which songs are going to be supplied in the market under the 3 model of predictability and compute welfare

Perfect foresight. Here the products enter the market based on their actual (realized) quality δ_j (no forecasting error), with the best entering first until the costs to supply equal the revenue:

$$r_k = pMs_k = pM\left[\frac{e^{\delta_k/(1-\sigma)}}{D_k^{\sigma} + D_k}\right] \quad (11)$$

NOTE: revenue depends on demand which depends on quality of the song δ and the taste parameter for appreciating a new song fitting your taste σ

No predictability. Nothing can predict success of a new song. Thus, products have ex-ante the same expected quality. Thus, products enter at random, in a number such that average revenue equals cost:

$$E[r_k] = pMs_k = pME\left[\frac{\frac{D_k}{D_k^{\sigma}+D_k}}{k}\right] \quad (12)$$

NOTE: in this case all songs have the same market share ex-ante, which is the average market share

Imperfect prediction. Here investors have some ability to predict the success of a song, but that is subject to an error. Thus, songs enter in decreasing order of expected quality δ_j', until expected revenues equal costs:

$$E[r_k] = pMs_k = pME[s_k] \quad (13)$$

NOTE: $E[s_k]$ is the expected market share of the last song supplied. Thus, this case is intermediate between the previous two as long as the ordering based on the forecasted value does better than random.

iv. estimation procedure:
    1. computing the correlation: the forecasting model can explain 20% of realized quality
    2. cost to bring a song to market reduced by factor 100
v. consumer surplus:
    1. welfare: $W = CS + Rev + N \cdot FC$

d. Outcomes:
    i. benefits for consumers arising from more variety available
    ii. benefits for consumers arising from lucky losers and variety is large
    iii. the larger the departure from perfect foresight, the highest the return of the long tail for consumers

Lecture 4:

1. <u>Bolton, Katok and Ockenfels, *How effective are electronic reputation mechanisms?*:</u>
    a. Motivation:
        i. reputation mechanisms aim to mitigate the asymmetric information problems in electronic markets by aggregating past experiences of buyers
    b. Insights:
        i. even if info about reputation is shared and reliable, online feedback systems provide fewer incentives to trust than offline
        ii. info provision through the feedback system has the property of a public good → positive externality not internalized
        iii. distinction between direct and indirect reciprocity: when system provides enough info, the two should have the same effectiveness
    c. Model set-up:
        i. lab experiment
        ii. 3 markets: stranger market (meet no more than once, no rating system), feedback market (meet no more than once, but can see seller histories) and partner market (interact in every round)
    d. Outcomes:
        i. efficiency: number of transactions completed.
        ii. trust: share of buy decisions (how many buyers do you trust).
        iii. trustworthiness: share of ship decisions conditional on buy orders (how many sellers are honest).
        iv. → for all three: partners > feedback > strangers, but differences not always statistically significant
        v. building trust (requires a lot of actions) takes a lot more than destroying trust (only one action) → trust strongly depends on last action
    e. Consequences:
        i. feedback markets generate less information than partner markets on sellers, because buyers tend to trust less and sellers cannot prove to be trustworthy
        ii. buyers do not internalize the positive externality they generate on testing sellers

iii. a feedback mechanism can self-sustain and generate information about sellers so to convince buyers to trust, but trust can be eroded quickly and information provision has the nature of a public good

2. <u>Cabral and Hortacsu, *The dynamics of seller reputation*</u>:
    a. Question:
        i. does price/quantity depend on seller rating?
        ii. empirical study following sellers (panel data) = observational study
    b. Model set-up
        i. after auction is completed, both parties can give a grade and leave textual comments + aggregated information on the grades is reported
        ii. 4 products in auction
        iii. whole seller profile is downloaded
        iv. variables: rating information and auction prices
        v. focus on quantity, but no direct information → #grades as proxy
        vi. empirical model:

$$price = \beta Reputation\ measure + \gamma Other\ demand\ factors + \varepsilon \quad (1)$$

        vii. cross sectional regression, but still concerns about unobservables (ex. quality), although negative feedback seems to reduce price
        viii. measure quantity through a proxy: amount of feedback, with the assumption that feedback rate is not influenced by rating
        ix. first: de-trend the sales growth pattern:

$$Sales\ growth_{it} = \beta_1 Total\ sales_{jt} + \beta_2 Total\ sales^2_{it} + Category_j + \varepsilon_{ji} \quad (2)$$

        x. second: average weekly sales growth rate before/after negative feedback
        xi. result: strong negative effect of negative rating on sales, less clear results for 2nd/3rd negative feedback
    c. Question:
        i. how does negative feedback evolve over transactions? Does it accelerate?
        ii. does negative feedback anticipate or trigger exit? (=leaving the platform or changing identity - no activity last 45 days)
        iii. are those who exit more likely to misbehave right before?
    d. Outcomes:
        i. it takes less to get to the second negative than to the first (seems not random/stationary)
        ii. threshold to give negative feedback lowers when others do it
        iii. sellers change approach over time as a response on negative feedback
        iv. if seller receives lot of positive feedback, he is more likely to stay (not other way around)
        v. in the last months before exit, sellers get more negative feedback: negative feedback accumulates and decides to stop / seller exploits its reputation as in static game
        vi. sellers' performance on the platform seems to respond to the rating: price decrease and quantity sold decreases after negative feedback and exit decreases if positive

vii. seller is penalized if perceived less likely to be honest and seller is more likely to be opportunistic before exit

3. Resnick, Zeckhauser, Swanson and Lockwood, *the value of reputation on eBay*:
    a. Question:
        i. previous study possibly plagued by omitted variables biasing the results: better sellers prove better description, more photos, are more careful in packaging items and also get better rating
        ii. design experiment on eBay to test the effect of rating on the price: what is the effect of negative feedback on the price?
    b. Model set-up:
        i. generate new profiles for 7 new sellers and 1 strong seller (via A/B testing)
        ii. differentiate the sellers without changing the actual content, but different lay-out
        iii. start with giving negative feedback to three new sellers
    c. Outcomes:
        i. hypothesis 1: buyers are willing to pay more to a seller with a strong positive reputation
            1. the sign of the ratio should be positive more frequently, which is significant!

            $$\ln\left(\frac{Price_{Strong}}{Price_{New}}\right) = \ln(Price_{Strong}) - \ln(Price_{New})$$

            2. the sign of the difference in sales $Sales_{strong} - Sales_{new}$ should be positive more frequently, significant under *alpha* < 0.10
        ii. hypothesis 2: the new sellers with negative feedback will reap lower profits
            1. not confirmed in the data, because the result is nog significant (only 6% difference in sales)
            2. price difference was instead favoring the sellers with negative

Lecture 5:

1. Dellarocas and Wood, *The sound of silence in online feedback*:
    a. Question:
        i. is there reporting bias? And if so, how should we correct for the bias? (inflation of positive feedback)
        ii. derive unbiased estimates of the distribution of private (79% +) transaction outcomes that produced the public feedback(99% +)
        iii. information about silence and timing
    b. Model set-up I:
        i. auctions of rare coins, info about the auction and about the sellers and buyers
    c. Outcomes:
        i. when feedback is given, most of the time both give feedback (seller moves first mostly)
        ii. important rate of silent buyers and positive sellers
        iii. buyer has reason to delay its positive feedback → waiting until product arrives

iv. when you have negative feedback, buyer takes lead, but still takes time to have effect: mostly about not receiving the item or expectations are not met; seller gives negative feedback mostly about payment

d. Model set-up II:
   i. derive and estimate a model of trade and feedback between a buyer and a seller
   ii. complexity in outcomes has to be reduced and general correspondence between values and feedback has to be restricted (traders cannot report systematically in an untruthful way)
   iii. notation:

      Outcome of the transaction (privately observed): i_b, i_s
      Outcome can be of 3 types: $\Omega = \{G, M, B\}$, (good, mediocre or bad)
      Probabilities of transaction outcomes: π_{i_b,i_s}

      - e.g., π_{G_b,M_s} is the probability that a trade ends up with a good outcome for the buyer and a mediocre outcome for the seller
      - Note: we use the indexing convention that the buyer is first and the seller second, thus we can write π_{G_b,M_s} simply as π_{GM}

      Feedback of the transaction (publicly observed): j_b, j_s
      Feedback can be of 3+1 types: $\Theta = \{+, 0, -, S\}$, (positive, neutral, negative or silence)
      Probabilities of reporting trade outcomes: $\rho^k_{j_k|i_k}$

      - e.g., $\rho^b_{G_b|G_b}$ is the probability that a buyer reports positive feedback after getting good outcome from the trade
      - Note: because of the restriction to the mapping between outcome and feedback, the notation simplifies to ρ^b_G

   iv. assumptions:
      1. there is one-to-one mapping between transaction outcomes and report types
      2. traders either truthfully report the transaction outcome they observe or remain silent
   v. the observed frequencies in the data are:

$$F_{j_b j_s} = \sum_{i_b \in \Omega} \sum_{i_s \in \Omega} \pi_{i_b i_s} \rho^b_{j_b|i_b} \rho^s_{j_s|i_s}$$

   vi. estimation by Maximum Likelihood:

$$\mathcal{L} = \sum_{j_b j_s \in \Theta} N_{j_b j_s} \log(F_{j_b j_s})$$

s.t.

$$\pi_{i_b i_s}, \rho^k_i \in [0, 1] \quad \text{and} \quad \sum_{i_b, i_s} \pi_{i_b, i_s} = 1$$

e. Outcomes II:
   i. satisfied traders (buyers + sellers) have high propensity to report and also dissatisfied traders
   ii. mildly dissatisfied trader prefer to stay silent

f. Model extension I:
   i. timing of feedback enters the model

        ii. traders’ time to feedback correlates with the private outcome: good outcomes are likely to report sooner than bad outcomes
        iii. results: receipt of positive feedback from seller increases the buyer’s propensity to report good and bad outcomes and decrease the mediocre outcomes; similarly for sellers who respond to buyer’s feedback
    g. Model extension II:
        i. take into account the vast heterogeneity in traders and auctions characteristics on eBay
        ii. verify if observable characteristics (ex. rating) is related to fundamental parameters (ex. probability of a good outcome)
        iii. results: rating helps in predicting satisfaction, but is likely to be biased
    h. Summary:
        i. online platforms adopt rating systems to sustain trust and to help buyers to discriminate among sellers
        ii. limitations: public good and possible bias
        iii. correct for inflation of positive feedback to get realistic satisfaction
        iv. conclusions:
            1. ratings are inflated
            2. mediocre and silent feedback when mildly dissatisfied
            3. evidence of reciprocity
            4. rating helps to predict satisfaction
2. <u>Bolton, Greiner and Ockenfels, *Engineering Trust: reciprocity in the production of reputation information*</u>:
    a. Question:
        i. how should a rating system be designed?
        ii. each platform has specificities
            1. unilateral/bilateral feedback
            2. technology might change over time
        iii. implications of different design on behavior on the platform
        iv. motives for providing feedback
    b. Empirical study:
        i. 70% of traders leave feedback
        ii. this would provide 49% cases with mutual feedback, while actual frequency is 64%
        iii. feedback content is expected to be positively correlated, but also symmetric
        iv. 90% correlation when seller gave feedback second (30% for buyers)
        v. timing of feedback is not independent of content - most of mutually positive and problematic feedback below 45-degrees. This means that seller is waiting and then responds
        vi. hard evidence that sellers retaliate negative feedback of buyers: social preference or emotional response, increase chance to mutual withdraw, punish the cause of a loss in future trading opportunities

Figure 1 Content and Timing of Mutual Feedback on eBay

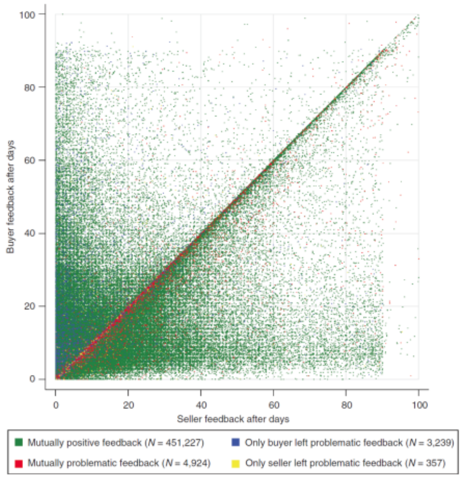

vii. social cost: underprovision of information, bias in the rating and the sound of silence

c. Proposals
    i. double blind conventional feedback: feedback remains bilateral, but is revealed after a deadline
    ii. conventional feedback plus a one-sided detailed seller rating (DSR) by the buyer

d. Lab experiment:
    i. goal: compare moral hazard, signal quality and market efficiency under the three feedback systems (open two-sided, double blind and single-sided)
    ii. pay-off:

$$\text{Seller: } \pi_s = p - 100q_s$$

$$\text{(Winning) Buyer: } \pi_i = q_s v_i - p$$

    iii. comments:
        1. still public good
        2. retaliation from seller is threat to the buyer
        3. focus on moral hazard of seller
    iv. overall feedback rate: baseline > DSR > blind
    v. sellers try to wait before giving feedback
    vi. problematic feedback increases in blind and DSR
    vii. question: does a better system improve economic outcomes?
        1. should improve quality, price and trust
        2. scope for trade

e. Summary:
    i. both alternative feedback systems seem to provide more dispersed rating
        1. field: feedback level does not seems to be reduced
        2. lab: both alternatives help to predict actual quality and increases significantly market efficiently
    ii. overall it is advisable to change the rating system, although it is hard to claim that one alternative is clearly superior to the other
        1. incentives to give feedback stay the same
        2. DSR is used to convey information on seller
    iii. feedback systems are fundamental to boost trust in online markets
    iv. their design is important
    v. designing optimal feedback system mean to address platform's specificities

3. Klein, Lambertz and Stahl, *Market transparency, adverse selection and moral hazard*:
    a. Question:
        i. what would you find if on top of DSR feedback becomes unilateral?
    b. Predictions:
        i. sellers have to take costly actions to ship something in line with the expectations of the buyers
        ii. sellers know that rating matters
        iii. rating is going to be less biased

iv. BUT moral hazard can induce adverse selection, which gives an incentive to underprovide quality and thus the good sellers are going to exit
v. interplay between sellers' response to buyers' feedback
   1. direct effects on the signal (feedback)
   2. first-order effects on the quality of shipped products
   3. second-order effects on market participation, frequency of interactions on the platform, willingness to report a score…
   4. third-order effects on long-term value of the platform

c. Results:
   i. descriptive statistics:
      1. % of positive classic rating declines
      2. DSR rating improves: signs of convergence of biased signal towards a more truthful information and some effects on signal provision
   ii. what happens to quality? Observed DSR are averages over the last 12 months:

$$DSR_{it} = \sum_{\tau=t-12}^{t-1} \omega_{i\tau}^{t} \cdot dsr_{i\tau}$$

   iii. one would like to estimate:

$$dsr_{i\tau} = \alpha + \beta POST_{i\tau} + \alpha_i + \varepsilon_{i\tau}$$

   iv. one can estimate:

$$DSR_{i\tau} = \alpha + \beta \left(\sum_{\tau=t-12}^{t-1} \omega_{i\tau}^{t} POST_{i\tau} \right) + \alpha_i + \left(\sum_{\tau=t-12}^{t-1} \omega_{i\tau}^{t} \varepsilon_{i\tau} \right)$$

   v. when platform provides unilateral DSR scores, the average score goes up from 4.69 to 7.78
   vi. there is a 0.05 increase in customer's satisfaction (DSR score) and the effects become larger after some months
   vii. low quality sellers have an incentive to do better or is there an exit or reduction in activity by low-quality sellers instead (?)
   viii. change in feedback system removed the possibility of sellers to retaliate and buyers can express better their true view on quality → raise in quality
   ix. sellers did not adjust immediately, because they did not fully capture the importance, but once they got punished by buyers for not meeting the standards, they improved quality

Lecture 6:

1. <u>Calvano, Calzolari, Denicolò and Pastorello, *Artificial intelligence, algorithmic pricing and collusion*:</u>
   a. goal:
      i. study pricing behavior (pricing strategies) of AI automated pricing algorithms
      ii. setting: simulated market environment where algo's compete

b. Model set-up:
   i. discrete decision process over time: t= 0,1,2,...
   ii. objective of the problem is to maximize:
   $$E\left[\sum_{t=0}^{\infty}\delta^t\pi_t\right]$$
   this is a standard object in dynamic programming
   iii. Bellman equation:
   $$V(s)=\max_{a\in A}\{E[\pi|s,A]+\delta E[V(s')|s,a]\}$$
   can also be written as:
   $$Q(s,a)=E(\pi|s,A)+\delta E\left[\max_{a\in A}Q(s',a')|s,a\right]$$
   where Q-function is the discounted pay-off of taking action *a* when the state is *s*
   iv. the agent has two options:
      1. finding the optimal policy in closed-form
      2. find a numerical solution to the optimal policy
   v. Q-learning is a procedure to estimate the Q-function with two important features:
      1. agnostic about the underline model → only observes actions and consequences
      2. tackles the problem from a single agent perspective → no instructions on the details of the environment
   vi. features of Q-learning:
      1. exploration of the consequences of actions and of the transition probability function between states
      2. the more information is gathered, the less exploration and the more exploitation
      3. exploitation: given the knowledge accumulated up to time t, the algo picks the best action
      4. exploration: given the knowledge accumulated up to time t, the algo pcis a non-optimal action at random
   vii. exploration/exploitation probability:
   $$\varepsilon_t=e^{-\beta t}$$

   After a_t in state s_t and observing π_t and s_{t+1} it updates:
   $$Q_{t+1}(s,a)=(1-\alpha)\underbrace{Q_t(s,a)}_{\text{Previous value}}+\alpha\left[\overbrace{\pi_t}^{\text{Reward}}+\underbrace{\delta\max_{a\in A}Q_t(s',a)}_{\text{Value new state}}\right] \quad (5)$$
   thus, α is the weight the algo gives to the new information regarding this action/state against the pre-existing information (so called *learning rate*)

   The other cells (not visited in t) are not updated, so $Q_{t+1}=Q_t$

viii. algorithms playing a pricing game with differentiated products; demand for the product of algo *i* at time *t* is:

$$q_{it} = \frac{e^{\frac{a_i - p_{it}}{\mu}}}{\sum_{j=1}^{n} e^{\frac{a_j - p_{jt}}{\mu}} + e^{\frac{a_0}{\mu}}}$$

initialized with a clear Q-matrix: Q_0=0 → the first choice is random

ix. algos run during a session until convergence or until 1billion rounds per session is reached

c. Results:

i. profit gain:

$$\Delta \equiv \frac{\overline{\pi} - \pi^N}{\pi^M - \pi^N}$$

optimal price → full collusion with competitor or perfect competition

ii. for non extreme values of *beta* converge is the best response function. In the remaining cases, the difference to the best-response is minimal: lost of profit within 1% from the best response

iii. algos do learn to play Nash: once the learning process is complete, the algos cannot be exploited, no matter how smart the opponent is

iv. average profit gain is 70%-90% and this is not a jump but playing systematically high prices (partial collusion)

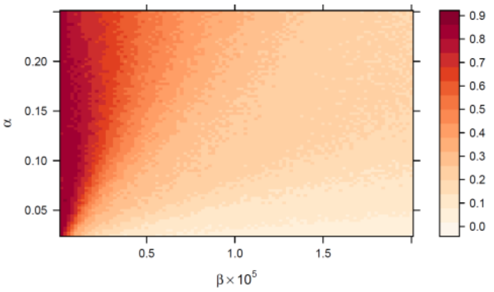

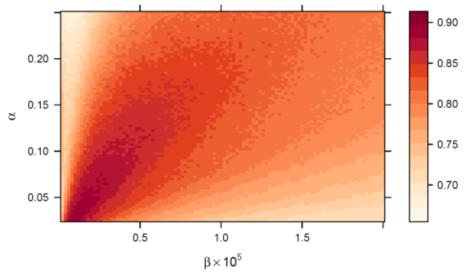

v. reaction to deviation?

1. response to non-collusive prices such that market goes back to the equilibrium prices
2. only situations where there are no collusive strategies are obvious from theoretical perspective
3. when deviating: algos converge back to high price equilibrium

d. Equilibrium strategies:

i. high prices maximize value

ii. if opponent deviates, best action is to punish the deviation

iii. opponent moves back to high prices

iv. partnership (self interested) - punishment - forgiveness

v. the strategies self-sustain:

1. deliver high profits
2. deviations don't pay

vi. best thing to do for a player facing an algo environment is to play the same strategy

vii. strong incentive for collusive behavior to arise

viii. implications for policy?