31E11100 - Microeconomics: Pricing

Pauli Murto

Aalto University

Part 1: prices and competitive markets
Lectures on 4.9., 6.9., and 11.9.2023

Content of Part 1

- Lecture 1: Partial equilibrium framework (September 5)
- Lecture 2: Price dispersion in real markets (September 7)
- Lecture 3: Search costs and price dispersion (Septmeber 12)

Lecture 1: prices in competitive partial equilibrium analysis

Modeling framework for this lecture: partial equilibrium analysis

- This is the standard modeling approach to focus on a single market, and will be utilized throughout the course
- Interactions with other markets assumed away
- Underlying assumption is that the market under consideration is only a small part of the overall economy. Then:
  - Prices of other commodities unaffeted by the price of the commodity under consideration
  - Wealth effects can be ignored (the spending in this market is a negligible part of the consumers' budget)
- This should be seen as an approximation
- Examples: grocery items, books, digital goods, banking services, hairdressing, etc.

Modeling tool for partial equilibrium analysis: quasilinear preferences

- The quantity of the good bought by individual i is denoted by $q_i \geq 0$.
- Utility to i from consuming q_i units: $v_i(q_i)$.
- Outside good y that can be thought of as money or a composite good reflecting all other consumption.
- The price of the good is fixed at $p > 0$ per unit of consumption. The composite good is priced at 1, and hence p is also the relative price.
- Initial holdings of m_i units of the composite good or money.
- Quasilinear utility:

$$u_i(y_i, q_i) = v_i(q_i) + y_i.$$

Consumer's problem

- Consumer's problem is to maximize

$$\max_{q_i \geq 0} v_i(q_i) + y_i$$

subject to

$$y_i + pq_i = m_i.$$

- We will assume that the function v_i is increasing and has decreasing marginal utilities:

$$v_i'(q_i) \geq 0 \text{ and } v_i''(q_i) \leq 0,$$

i.e. $v_i(\cdot)$ is an increasing and concave function.

- Plugging the budget constraint $y_i + pq_i = m_i$ into objective function, the first order condition for maximum is

$$v_i'(q_i) = p$$

- The solution to the consumer's problem is the demand function, denoted by

$$q_i(p).$$

- Important simplification by the quasi-linearity: demand is independent of m_i since q_i maximizes $v_i(q_i) + m_i - pq_i$ if and only if it maximizes $v_i(q_i) - pq_i$. We can ignore initial wealth m_i from now on.

- By differentiating the first-order condition, we get the law of demand:

$$q_i'(p) = \frac{1}{v_i''(q_i(p))} \leq 0,$$

or in words, the individual demand is downward sloping

Production side

- Firm j supplies q_j^s units of the good.
- The cost function $c_j\left(q_j^s\right)$ measures the cost of delivering q_j^s units on the market in terms of the composite good (or money).
- The produced good is priced at p in the market, and the firm chooses q_j^s to maximizes its profit.

- Hence, the Firm's problem is:

$$\max_{q_j^s} \; pq_j^s - c_j\left(q_j^s\right),$$

- For the most part, we assume that the cost function is increasing and convex:

$$c_j'\left(q_j^s\right) > 0 \text{ and } c_j''\left(q_j^s\right) \geq 0.$$

- First order condition for Firm's problem in the continuous case:

$$c_j'\left(q_j^s\right) = p.$$

- The solution to this equation is called the individual supply of firm j and it is denoted by $q_j^s\left(p\right)$.
- By differentiation we get the individual law of supply:

$$q_j^{s\prime}\left(p\right) = \frac{1}{c_j''\left(q_j^s\right)} > 0.$$

- In words, more is supplied at higher output prices.

Description of the market place

- In competitive analysis, firms and consumers meet in an anonymous market.
- Anonymity means that the price is the same for all participants and not dependent on the identities i and j.
- Both the buyers and the sellers are price takers: the price is assumed to be given and does not depend on individual demands q_i and supplies q_j^s.
- Price is linear: the cost of buying q_i units is $p \cdot q_i$ rather than a more general function $p(q_i)$.
- All buyers and all sellers know the price.
- We cover here the case with continuous demands and supplies, but the arguments generalize easily to the discrete case too.

- Aggregate market demand Q is obtained by summing over i all individual demand functions:

$$Q(p) = \sum_{i=1}^{I} q_i(p),$$

where I is the total number of consumers in the market.

- By the individual laws of demand, we get

$$Q'(p) = \sum_{i=1}^{I} q_i'(p) < 0.$$

Market demand curve is thus downward sloping.

- Market supply $Q^s(p)$ is obtained by summing over j all individual supply functions:

$$Q^s(p) = \sum_{j=1}^{J} q_j^s(p),$$

where j is the total number of firms in the market.

- By individual laws of supply, we get

$$Q^{s\prime}(p) = \sum_{j=1}^{J} q_j^{s\prime}(p) > 0,$$

so that market supply curve is upward sloping.

- An equilibrium in the market is a pair (p^*, Q^*) such that markets clear:

$$Q^* = Q(p^*) = Q^s(p^*).$$

- Observe that in equilibrium, each firm j supplies quantity $q_j^s(p^*)$ and each consumer i demands $q_i(p^*)$.
- In other words, every single consumer chooses the optimal consumption level, and every firm chooses the profit-maximizing output, given the market price.
- This shows that in partial equilibrium analysis, aggregating individual consumers and firms is really easy.

Markets, Efficiency, and Welfare

- How to define efficiency?
- As you know, in economics this refers to **Pareto-efficiency**:

Definition

A feasible allocation is Pareto-efficient if there is no other feasible allocation where at least one of the agents is better off and none of the agents is worse off.

- Another way of phrasing this: Starting from a Pareto-efficient allocation, you cannot help anyone without hurting someone else.

- What is an Pareto-efficient allocation in the current context?
- Recall, the welfare of a consumer with quasilinear preferences is

$$u_i (y_i, q_i) = v_i (q_i) + y_i.$$

- Similarily, for firms the relevant welfare measure is their profit in terms of the composite good:

$$\pi_j (y_j, q_j) = y_j - c_j (q_j).$$

- The key insight is the following: in any Pareto-efficient allocation,

$$v_i'(q_i) = v_{i'}'(q_{i'}) = c_j'(q_j) = c_{j'}'\left(q_{j'}\right) \text{ for all } i, i', j, j'.$$

- In other words, marginal utilities and costs must be equal across consumers and firms

- To see why this must be true, notice that profitable trade exists between buyers i and i' if $v_i'(q_i) \neq v_{i'}'(q_{i'})$, and between consumer i and producer j if $v_i'(q_i) \neq c_j'(q_j)$.
- Similarly, a profitable reallocation of production exists between j and j' if $c_j'\left(q_j^s\right) \neq c_{j'}'\left(q_{j'}^s\right)$.
- Hence the allocation is Pareto efficient only if $v_i'(q_i) = c_j'(q_j) = p$ for some p.

- Finally, an allocation is feasible only if at least as much is produced as consumed.
- But efficiency requires that total produced should not exceed total amount consumed (no waste). Therefore

$$Q(p) = Q^s(p),$$

and this happens only when $p = p^*$ and $Q(p) = Q^s(p) = Q^*$.

- This is the *First Welfare Theorem* applied in our context: The competitive equilibrium is Pareto efficient.
- Moreover, with quasi-linear preferences, there no are other Pareto-efficient allocations: there is only one allocation that equates marginal utilities and marginal costs across consumers and firms.

- Inverses of aggregate demand and supply have also straight-forward interpretations:
  - Inverse of aggregate supply can be viewed as the industry marginal cost function:
$$C'(\cdot) = (Q^s)^{-1}(\cdot)$$
  - Inverse of aggregate demand (inverse demand function), gives the marginal social benefit of good l :
$$P(\cdot) = (Q^l)^{-1}(\cdot)$$

- There is a "normative" representative consumer
- Equilibrium price equates marginal social benefit and marginal social cost of production of l, and in this way ensures that the total welfare is maximized
- The familiar Consumer's and Producer's surpluses measure the shares of welfare going to the consumers and firms, respectively
- Note that this is true under certain conditions:
  - How would externalities affect the properties of equilibrium?
  - How could prices be corrected to account for externalities?

Learning points from the model

- In competitive anonymous markets:
  - Prices adjust to clear the market
  - A unique price obtains in the market: Law of one price.
  - Competitive equilibrium allocation is efficient, and hence maximizes the sum of producers' and consumers' surplus
  - In other words, price signals simultaneously marginal value and marginal cost of adding output
  - Of course, this hinges on strong assumptions of the model:
    - No market power
    - Perfect information

Further readings

- For the partial equilibrium analysis, see any intermediate level microeconomics book
- To move beyond quasilinear utilities to general equilibrium analysis, see more advanced microeconomics textbooks such as Jehle and Reny: Advanced Microeconomic Theory or Mas-Colell, Whinston and Green: Microeconomic Theory

Lectures 2-3: Price dispersion in real markets

Law of one price in real markets

- Predictions of basic partial equilibrium analysis seem to work quite well in centralized markets such as commodity exchanges
- But fail in many familiar decentralized markets
- A good example: supermarkets
  - Markets do seem to clear
  - Law of one price fails.
- This failure has been documented in numerous markets. (see the review article Baye and Morgan, 2005, for extensive discussion on this)
- Does not seem to be too sensitive to the definition of market
- A broad and systematic analysis is presented by Kaplan & Menzio (2014): The Morphology of Prices

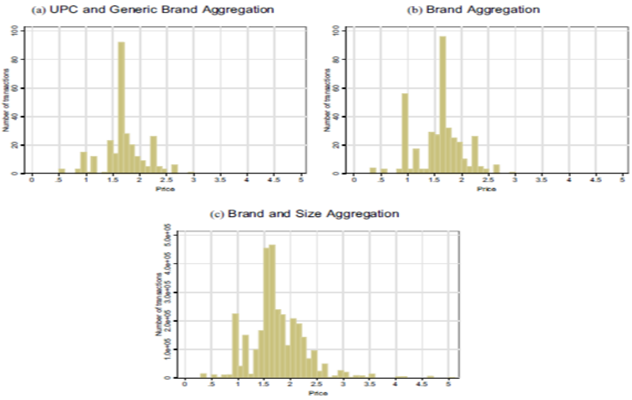

NOTES: Figures show distribution of transaction prices for ketchup in Minneapolis in 2007:Q1. Panel (a) shows prices for 36-oz bottles of Heinz brand ketchup, in accordance with the UPC and Generic Brand Aggregation definitions of a good. Panel (b) shows prices for 36-oz bottles of ketchup from all brands, in accordance with the Brand Aggregation definition of a good. Panel (c) shows prices per 36 ounces of ketchup for bottles of ketchup from all brands in all sizes, in accordance with the Brand and Size Aggregation definition of a good. Transactions include those at all stores in Minneapolis, including stores without a unique identifier.

FIGURE 1

DISTRIBUTION OF PRICES FOR A 36-OZ BOTTLE OF HEINZ KETCHUP

How to explain this?

- Maybe the stores are differentiated
  - Some stores have higher quality
  - Some stores at more attractive locations
  - But according to Kaplan & Menzio variation by store accounts only for 10% of total price variation
- Maybe supermarkets have different costs
  - Maybe they have different wholesale prices
  - But wholesale price depends on the chain to which the store belongs and the chain explains very little of the price variation.

- Maybe supermarkets are special
  - shoppers get a basket of goods, not a single good
  - But Sorensen (2000) shows similar price dispersion for prescription drugs in pharmacies in a small town.
    - In Sorensen (2000), variation in prices seems to be independent in the sense that some drugs are expensive while others are cheap at a given pharmacy.
- Maybe shoppers are not aware of all the prices in their market
  - Maybe searching for price information is costly
  - How should sellers' respond to consumer search frictions?
  - This is in fact the leading explanation
  - We look at this in more detail in the next lecture

How to explain theoretically price dispersion?

- We saw that there is substantial price dispersion in many real markets
  - Failure of the Law of One Price
- We will try to explain this theoretically by looking at a richer model
- The key increadient for the model is that the consumers are imperfectly informed about the prices
- This seems realistic in many situations:
  - Searching for the lowest price is costly (think about opportunity cost of time)
  - Should one bother to search? What is the expected gain in searching? This depends on price dispersion...

- Why would costly search lead to price dispersion?
  - If all the buyers are well informed of all prices, competition should drive all prices low
  - If all the buyers are very poorly informed (say, know only prices in their local store), then the sellers have monopoly power over their local customers and should price high
  - But if some buyers are better informed than others, then perhaps some sellers should price low to attract well informed buyers, while some should price high to sell only to uninformed buyers
  - Is this consistent with buyer's optimal search behavior?
- To analyze this, we need a model

What are economic models?

- Economic models are simplified descriptions of reality
- Their purpose is to isolate key elements of an economic situation to understand causal relationships:
  - How do patent policies affect innovation?
  - How do rental price controls affect housing markets?
  - etc. etc.
- Challenge for modeling: individuals' behavior is not fixed
- Economic models usually build on the following principles:
  - Individual rationality: individuals make choices as if maximizing some objective function.
  - Consistent expectations: what you believe about others is consistent with the others' actual behavior.

Strategic interactions

- Often what is good for you depends on what others do.
  - Do you want to drive on the right-hand side of the road?
  - How much do you want to bid in an auction?
  - For this lecture: Your optimal price depends on prices charged by others.
- Similarly what others want to do depends on what you do.
- Game theory is the tool for analyzing such situations.
- If you are not familiar with game theory, or need a recap, then:
  - Have a look at the APPENDIX at the end of this slide set
  - Have a look at Robert Gibbons: "An Introduction to Applicable Game Theory", Journal of Economic Perspectives, 11(1), 1997
  - Or consult any text book on game theory (such as Gibbons: "Game Theory for Applied Economists", Princeton University Press)

Model with Imperfectly Informed consumers

- Back to modeling price dispersion
- We build a model with many buyers and sellers
- We then vary the degree of information that the buyers have.
- We start with exogenously given information
- How is sellers' equilibrium pricing strategies affected?
- At the end, we ask: if the buyers can choose, how much information will they collect?

Formal model

- Suppose that there are $N = 2$ identical sellers (for the working of this model N does not matter, could be infinite just as well).
- A large number of buyers. Model this as a continuum of mass 1.
- Each seller sells identical goods at production cost c per unit of good.
- Each buyer has a unit demand with reservation value $v > c$ (i.e. demand one unit of the good as long as price is less than v).
- Firms set their prices simultaneously.

Buyers' behavior

- For now, assume that the information of the buyers is exogenous.
- In other words, they don't make an active choice in how much information to collect.
- Suppose the following:
  - Fraction α of the buyers sample at random a single price.
  - Fraction $(1-\alpha)$ sample two prices.
- Buyers choose the lowest of the prices they observe (If the two prices are tied, both sellers are chosen with equal probability)

Game between sellers

- Set of players $i \in \{1, 2\}$.
- Strategy of seller i: choose price $s_i \in [c, v]$.
- To allow for mixed strategies (turns out to be important here!), we may represent a strategy as a cumulative distribution of prices $F(s)$.
  - This means that $F(s) = \Pr\{s_i \leq s\}$ is the probability that i chooses a price at or below s.
  - We will analyze symmetric equilibria, and therefore this describes the behavior of both sellers.
  - $F(s)$ may contain atoms, so this description allows pure strategies.

(Nothing essential in this model would change if we assume N sellers or even a continuum of identical sellers)

- Denote by β the fraction of the buyers visiting firm i that have not sampled another price offer. By Bayes' rule:

$$\beta = \frac{\alpha}{2-\alpha}.$$

- Expected payoff to seller i per customer when choosing s_i and other seller pricing according to $F(s)$:

$$\beta(s_i - c) + (1-\beta)\left((1 - F(s_i)) + \frac{1}{2}p(s_i)\right)(s_i - c),$$

  where $p(s_i) = \Pr\{s_j = s_i\}$ for the other seller j.
- Simple cases:
  - If $\beta = 1$, then $s_i = v$ is optimal.
  - If $\beta = 0$, then all buyers have seen two prices and the equilibria are as in the Bertrand game, and the only equilibrium is $s_i = c$, $i = 1, 2$

Partial price information

- If $0 < \beta < 1$, then there are no single price equilibria in the market.
  - Suppose to the contrary that both firms were pricing at some s^*.
  - If $s^* = c$, firm i gets a higher profit at $s_i = v$ than at $s_i = c$. Hence $s_i = c$ for $i = 1, 2$ is not an equilibrium.
  - If $s^* > c$, then by setting $s_i = s^*$, firm i gets expected profit (per customer):

$$\beta (s^* - c) + (1 - \beta) \frac{1}{2} (s^* - c) = \frac{1}{2} (1 + \beta) (s^* - c)$$

- For any $s_i = s^* - \varepsilon$, (where $\varepsilon > 0$), firm i gets (per customer)

$$(s^* - \varepsilon - c).$$

- This follows since the firm now sells to all buyers with probability 1.
- Profit from $s_i = s^* - \varepsilon$ exceeds profit from $s_i = s^*$ if:

$$\varepsilon < \frac{1}{2}(1-\beta)(s^* - c).$$

- Hence $s_i = s^*$ for all i is not an equilibrium.

- One can show similarly that there are no equilibria where $p(s) > 0$ for some s.
  - The argument is basically the same undercutting argument.
- Hence we have in equilibrium that $p(s) = 0$ for all s.
- But then the expected profit for firm i from price s_i if the others price according to $F(s)$ is:

$$\beta (s_i - c) + (1 - \beta)((1 - F(s_i)))(s_i - c).$$

- A really basic and useful observation is that in a mixed strategy equlibrium, all prices chosen (i.e. all prices in the support of $F(s)$) must yield the same expected profit. (Why?)
- Furthermore, the highest price chosen in equilibrium must be v. (Why?)

- The profit from choosing $s_i = v$ is

$$\beta (v - c) .$$

- The firm with the lowest price $\underline{s}$ sells with probability 1 to all buyers. By equal profit requirement, we get

$$\underline{s} - c = \beta (v - c) .$$

For prices between the highest and the lowest, we have:

$$\beta (s - c) + (1 - \beta) ((1 - F (s))) (s - c) = \beta (v - c) .$$

- Solving for $F(s)$ gives:

$$F(s) = 1 - \frac{\beta(v - s)}{(1 - \beta)(s - c)}.$$

- You can check that $F(\underline{s}) = 0$ and $F(v) = 1$ and so we have a distribution function on $[c + \beta(v - c), v]$.

Analyzing the results

- Do the predictions of the model make sense?
- What are the effects of an increase in information (i.e. a decrease in β)?
    - First of all, as β decreases, the lowest price decreases.
    - Furthermore, since $\frac{-\beta}{1-\beta}$ is decreasing in β,

    $$F(s) = 1 - \frac{\beta(v-s)}{(1-\beta)(s-c)}$$

    is decreasing in β for all s.
    - We say that a distribution G first-order stochastically dominates distribution F if for all s,

    $$F(s) \geq G(s).$$

    - Hence we see that price distributions with less information first-order stochastically dominate distributions with more information.
    - Industry profit is $\alpha(v-c)$ and hence increasing in β (and so decreasing in information).

Is buyers search behavior optimal?

- So far we just assumed that some buyers see a single price and some see two prices
- What is the benefit of observing two rather than one price?
- Expected payoff difference $\Delta(\beta)$ from sampling the prices of two rather than one firm when fraction β of each firm's buyers are uninformed is:

$$\Delta(\beta) := \mathbb{E}s_i - \mathbb{E}\min\{s_i, s_j\},$$

where s_i and s_j are independent random draws from $F^{\beta}(s)$.

- Note that $\Delta(0) = \Delta(1) = 0$.
- For $0 < \beta < 1$, $\Delta(\beta) > 0$.

Detour: How to compute the expectation of the lowest price in a sample of n draws?

- Take a sample of n prices, where each price p_i is independently drawn from a cumulative distribution function $F(s)$ with support $[0,1]$.
- Let $F_n(s)$ denote the cumulative distribution function for the lowest price in a sample of n prices
- We can directly compute

$$\begin{aligned} F_n(s) &= \Pr(\text{"lowest price } \leq s\text{"}) \\ &= 1 - \Pr(\text{"all prices } > s\text{"}) \\ &= 1 - (1 - F(s))^n . \end{aligned}$$

- The corresponding density function is the derivative of cumulative distribution function:

$$f_n(s) = F_n'(s) = n(1 - F(s))^{n-1} F'(s).$$

- Therefore, the expected lowest price can be directly computed as

$$\mathbb{E}(\min \{s_i\}_{i=1}^n) = \int_0^1 s f_n(s)\, ds.$$

- Example: suppose that p_i are drawn from uniform distribution $p_i \sim U[0,1]$. Then, $F(s) = s$, and so

$$\begin{aligned}\mathbb{E}(\min \{s_i\}_{i=1}^n) &= \int_0^1 s f_n(s)\, ds \\ &= \int_0^1 sn(1-s)^{n-1}\, ds \\ &= \frac{1}{1+n}.\end{aligned}$$

Modeling buyers' search behavior

Two possible models for buyers:

- Fixed sample search
  - A buyer must decide at once whether to search for one or two prices (or more, if we have more than two sellers)
- Sequential search
  - A buyer samples sequentially and decides after each sample whether to continue search or to buy at the lowest price found so far

- With fixed sample search, we can rationalize the behavior described above by a suitable cost structure
    - Given cost γ for each sampled price, if $\Delta(\beta) = \gamma$, each buyer indifferent between sampling one or two prices
    - Then, a mixed strategy where each buyer samples one price with probability $\alpha = 2\beta / (1 + \beta)$ is an equilibrium
- The equilibrium description above remains hence valid even if consumers optimize their search sample
- But what about sequential search? Suppose next that the buyers see one price for free, but after observing it they can observe a second price at a cost γ

Diamond's Paradox

- If buyers see one price for free and then must choose whether to sample another price at cost $\gamma > 0$, then one equilibrium in the market is $s_i = v$ for all i, and no buyer pays the cost of becoming informed.
- To see this, notice that given the pricing decisions by the firms, additional price samples bring no benefit.
- Therefore no buyer becomes informed and the equilibrium is as with $\alpha = \beta = 1$ above.

- Moreover, this is the only possible equilibrium. Too see this:
  - Suppose there is an alternative equilibrium where some sellers price below v
  - Take the lowest $s' < v$ in the price support
  - No buyer who oberves price $s \in [s' + \gamma]$ will search for another price
  - So, why wouldn't a seller charging price to s' rather charge $s' + \varepsilon$?
- Diamond's Paradox: even a small search cost results in monopoly profits to the firms.

Discussion

- This model shows that we can have price dispersion even with homogenous product, sellers, and buyers
- But with sequential search, this basic model leads to monopoly prices (Diamond' paradox)
- Introducing heterogeneity for sellers and buyers makes it easier to get price dispersion. One can show that then price dispersion possible even with sequential search.
- Heterogeneities are natural for buyer side: opportunity costs of search are very different for different people (there are also "shoppers" who get positive utility from search)
- This is supported by empirical results that show that some buyers pay consistently much less for their shopping baskets than others

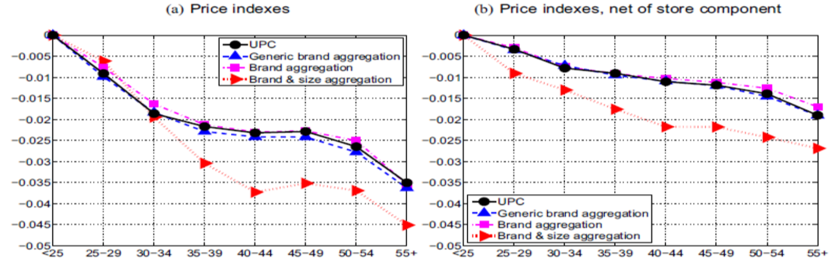

NOTES: Age refers to average age of household head(s).

FIGURE 9

AVERAGE PRICE INDEXES BY AGE

Discussion: on-line shopping and price dispersion

- Buying online should reduce search costs
- Does this mean that price dispersion disappears?
- Empirical findings so far suggest that price dispersion is not going anywhere
- Also theory predictions are ambiguous. There are models that allow substantial price dispersion even with very low marginal search costs

Brynjolfsson and Smith (2000): Frictionless Commerce? A Comparison of Internet and Conventional Retailers, Management Science.
Price dispersion of books and CD:s similar in internet and convential retail outlets.

Figure A5 Kernel Density for De-Meaned Full Prices for Books (Epanechnikov Kernel)

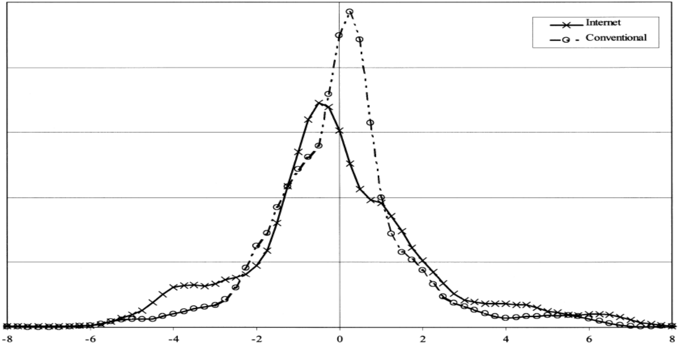

Learning points from this model

- Using a game theoretic model, we can explain price dispersion as resulting from imperfect price information amongst buyers.
- The model predicts higher prices for markets with less informed buyers.
- The degree of price dispersion in the market is determined in equilibrium, and depends on the interaction between buyers and sellers:
  - Buyers' benefit of search
  - Seller's profitability of lowering prices
- Robustness of modeling: similar results obtain with richer models (see other literature)
- Scope of application goes beyond product markets: think about labor markets and wage dispersion, minimum wages etc.

Further readings

- For empirical analysis of price dispersion, see e.g.
  - Kaplan and Menzio (2015), "The Morphology of Price Dispersion", International Economic Review
  - Sorensen (2000), "Equilibrium Price Dispersion in Retail Markets for Prescription Drugs", Journal of Political Economy
- The model framework of this lecture was a simplified version of Burdett and Judd (1983): "Equilibrium Price Dispersion", Econometrica.
- Other classical models of consumer search include:
  - Stigler (1961): "The Economics of Information", Journal of Political Economy.
  - Diamond (1971): "A model of price adjustment", Journal of Economic Theory.
  - Varian (1980): "A model of sales", American Economic Review.
- For a review of different models (and empirics as well), see Baye and Morgan (2005): "Information, Search, and Price Dispersion", Handbook on Economics and Information Systems (available online).

APPENDIX: Very short introduction to game theory

- To specify a game, one needs to specify *players*, *strategies*, and *payoffs*:
- N economic agents called *players* are engaged in an economic interaction. Set of players $\{1, ..., N\}$.
- Each agent i has a set of feasible choices called *strategies* S_i
- Each player has a preference over vectors of choices $(s_1, ..., s_N)$ represented by a *payoff function* $u_i(s_1, ..., s_N)$.
- We call the collection $\{S_i, u_i\}_{i=1}^{N}$ a game.

- All players i choose independently and simultaneously a strategy $s_i \in S_{i.}$.
- Payoffs are realized for the vector of choices $(s_1, ..., s_N)$.
- How should each i choose her action?
- The best action of i depends on her beliefs about the choices of s_j for $i \neq j$.
- Consistency of beliefs: Nash equilibrium.

Definition

A Nash equilibrium is a vector of strategies $s^* = (s_1^*, \ldots, s_N^*)$ such that for each i, s_i^* solves

$$\max_{s_i \in S_i} u_i\left(s_1^*, \ldots, s_{i-1}^*, s_i, s_{i+1}^*, \ldots, s_N^*\right).$$

- This says that at a Nash equilibrium, each player i is maximizing her payoff holding fixed the strategies of the other players.
- Nash equilibrium is a stable situation: starting from a Nash equilibrium, no player has an incentive to change her strategy.
- Conversely, if s is not a Nash equilibrium, then some player has an incentive to change her behavior.
- Notice that we have here N optimization problems that must be solved simultaneously.

- Example: Bertrand Pricing Game:
  - The players are two firms selling identical products. Set of players $\{1, 2\}$.
  - Each firm chooses a positive price s_i at which it agrees to sell its product in the market. The size of the market is 1.
  - $S_i = \mathbb{R}_+$ for $i \in \{1, 2\}$.
  - All buyers in the market know the prices and buy from a firm charging the lowest price. At equal prices, the market is split equally. Production cost c per unit.
  - Payoffs:

$$u_i = \begin{cases} s_i - c \text{ if } s_i < s_j, \\ \frac{1}{2}\left(s_i - c\right) \text{ if } s_i = s_j, \\ 0 \text{ if } s_i > s_j. \end{cases}$$

  - Claim: $s^* = (c, c)$ is the only Nash equilibrium of the game. Hence 2 firms is enough for competitive prices.

- Do all games have a Nash equilibrium? Example: inspection game
  - A firm may commit tax evasion. Gives positive payoff 1 if undetected, but entails loss of 1 due to penalty if detected
  - Tax authority wants to detect fraud, but inspection is costly.
  - Strategies are $S_1 = \{$ Inspect, Do not inspect$\}$ for the tax authority and $S_2 = \{$ Commit evasion, Be honest$\}$ for the firm
  - Assume the following payoffs (The first number in each cell is for the row player = the tax authority, and second for the column player = the firm):

	Commit evasion	Be honest
Inspect	$-1, -1$	$-1, 0$
Do not inspect	$-3, 1$	$0, 0$

- Is there a Nash equilibrium in pure strategies? (i.e. in the sense of our definition above)
- What if the players can randomize over their choices?
- Every finite game has a Nash equilibrium in pure or mixed strategies. See the additional material on game theory for more on this

- Example of a dynamic game: Entry game
  - An entrant considers entry into an industry with a current incumbent firm
  - Entry costs 1 unit
  - Monopoly profit in the industry is 4
  - If entry takes place, the monopolist can either accomodate or fight
  - Accommodation splits monopoly profits, whereas fighting gives zero profit to both firms
  - Will entrant enter, and if so, will incumbent fight or accomodate?
- Dynamic games are normally expressed in *extensive form* (see additional material)
- But can also be be represented in the normal form

- Normal form representation of the entry game:

	Fight if entry	Accommodate if entry
Enter	$-1, 0$	$1, 2$
Stay out	$0, 4$	$0, 4$

- There are now two Nash equilibria: (Enter, Accommodate) and (Stay out, Fight if entry)
- But only (Enter, Accommodate) is a *sub-game perfect* Nash equilibrium
- Sub-game perfect equilibrium can be derived by backward induction: start by analysing the optimal behavior of the incumbent firm once the entry has already taken place