9 February 2006. Eric Rasmusen, Erasmuse@indiana.edu.
Http://www.rasmusen.org.

Overheads for Chapter 14, Pricing, of *Games and Information*. These do not cover the entire chapter, just enough for two lectures.

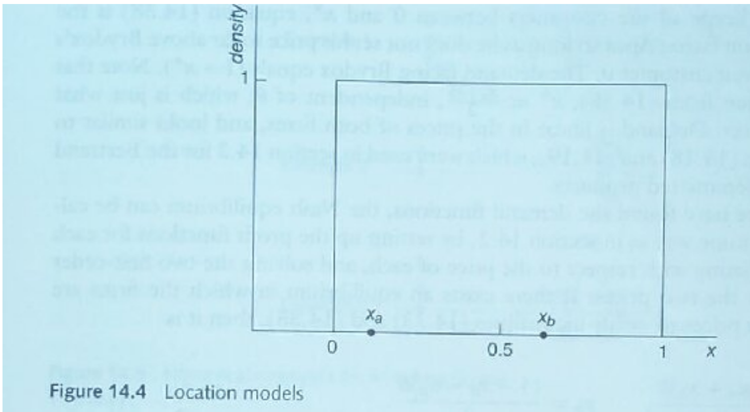

Figure 4: Location Models

The Hotelling Pricing Game

(Hotelling [1929])

Players

Sellers Apex and Brydox, located at x_a and x_b, where $x_a < x_b$, and a continuum of buyers indexed by location $x \in [0, 1]$.

The Order of Play

1 The sellers simultaneously choose prices p_a and p_b.

2 Each buyer chooses a seller.

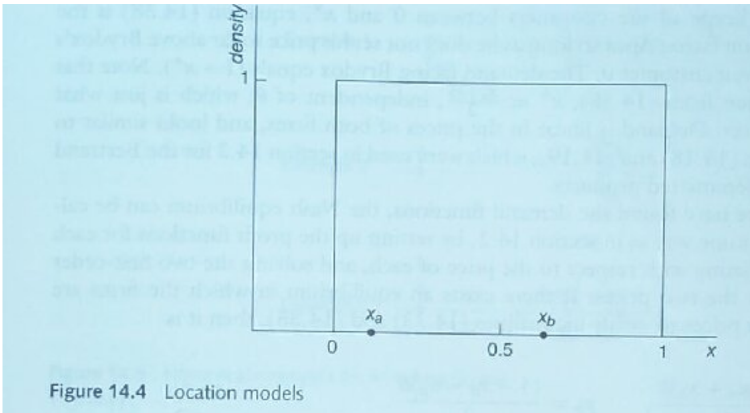

Figure 4: Location Models

Payoffs

Demand is uniformly distributed on the interval [0,1] with a density equal to one (think of each consumer as buying one unit). Production costs are zero. Each consumer always buys, so his problem is to minimize the sum of the price plus the linear transport cost, which is θ per unit distance travelled.

$$\pi_{buyer\ at\ x} = V - Min\{\theta|x_a - x| + p_a,\ \theta|x_b - x| + p_b\}. \qquad (1)$$

$$\pi_a = \begin{cases} p_a(0) = 0 & \text{if } p_a - p_b > \theta(x_b - x_a) \\ & \text{(Brydox captures entire market)} \\ & \\ p_a(1) = p_a & \text{if } p_b - p_a > \theta(x_b - x_a) \\ & \text{(Apex captures entire market)} \\ & \\ p_a(\frac{1}{2\theta}\left[(p_b - p_a) + \theta(x_a + x_b)\right]) & \text{otherwise (the market is divided)} \end{cases} \qquad (2)$$

$$\pi_a = \begin{cases} p_a(0) = 0 & \text{if } p_a - p_b > \theta(x_b - x_a) \\ & \text{(Brydox captures entire market)} \\ & \\ p_a(1) = p_a & \text{if } p_b - p_a > \theta(x_b - x_a) \\ & \text{(Apex captures entire market)} \\ & \\ p_a(\frac{1}{2\theta}\left[(p_b - p_a) + \theta(x_a + x_b)\right]) & \text{otherwise (the market is divided)} \end{cases}$$

A buyer's utility depends on the price he pays and the distance he travels. Price aside, Apex is most attractive of the two sellers to the consumer at $x = 0$ ("consumer 0") and least attractive to the consumer at $x = 1$ ("consumer 1"). Consumer 0 will buy from Apex so long as

$$V - (\theta x_a + p_a) > V - (\theta x_b + p_b), \tag{3}$$

which implies that

$$p_a - p_b < \theta(x_b - x_a), \tag{4}$$

which yields payoff (2a) for Apex. Consumer 1 will buy from Brydox if

$$V - [\theta(1 - x_a) + p_a] < V - [\theta(1 - x_b) + p_b], \tag{5}$$

which implies that

$$p_b - p_a < \theta(x_b - x_a), \tag{6}$$

which yields payoff (2b) for Apex.

Very likely, inequalities (4) and (6) are both satisfied, in which case Consumer 0 goes to Apex and Consumer 1 goes to Brydox. Let consumer x^* be the consumer at the boundary between the two markets, indifferent between Apex and Brydox.

Notice that if Apex attracts Consumer x_b, he also attracts all $x > x_b$, because beyond x_b the consumers' distances from both sellers increase at the same rate. So we know that if there is an indifferent consumer he is between x_a and x_b.

Knowing this, the consumer's payoff equation, (1), tells us that

$$V - [\theta(x^* - x_a) + p_a] = V - [\theta(x_b - x^*) + p_b], \tag{7}$$

so that

$$p_b - p_a = \theta(2x^* - x_a - x_b), \tag{8}$$

and

$$x^* = \frac{1}{2\theta}\left[(p_b - p_a) + \theta(x_a + x_b)\right], \tag{9}$$

which generates demand curve (2c)– a differentiated Bertrand demand curve.

The Nash equilibrium can be calculated by setting up the profit functions for each firm, differentiating with respect to the price of each, and solving the two first-order conditions for the two prices. If there exists an equilibrium in which the firms are willing to pick prices to satisfy inequalities (4) and (6), then it is

$$p_a = \frac{(2 + x_a + x_b)\theta}{3}, \quad p_b = \frac{(4 - x_a - x_b)\theta}{3}. \tag{10}$$

Apex charges a higher price if a large x_a gives it more safe consumers or a large x_b makes the number of contestable consumers greater.

Profits are positive and increasing in the transportation cost.

We cannot rest satisfied with the neat equilibrium of equation (10), because the assumption that there exists an equilibrium in which the firms choose prices so as to split the market on each side of some boundary consumer x^* is often violated.

Vickrey (1964) and D'Aspremont, Gabszewicz & Thisse (1979) have shown that if x_a and x_b are close together, no pure-strategy equilibrium exists, for reasons similar to why none exists in the Bertrand model with capacity constraints.

If both firms charge nonrandom prices, neither would deviate to a slightly different price, but one might deviate to a much lower price that would capture every single consumer.

But if both firms charged that low price, each would deviate by raising his price slightly.

It turns out that if, for example, Apex and Brydox are located symmetrically around the center of the interval, $x_a \geq 0.25$, and $x_b \leq 0.75$, no pure-strategy equilibrium exists (although a mixed-strategy equilibrium does, as Dasgupta & Maskin [1986b] show).

Hotelling should have done some numerical examples.

And he should have thought about the comparative statics. Equation (10) implies that Apex should choose a higher price if both x_a and x_b increase, but it is odd that if the firms are locating closer together, say at 0.90 and 0.91, that Apex should be able to charge a higher price, rather than suffering from more intense competition.

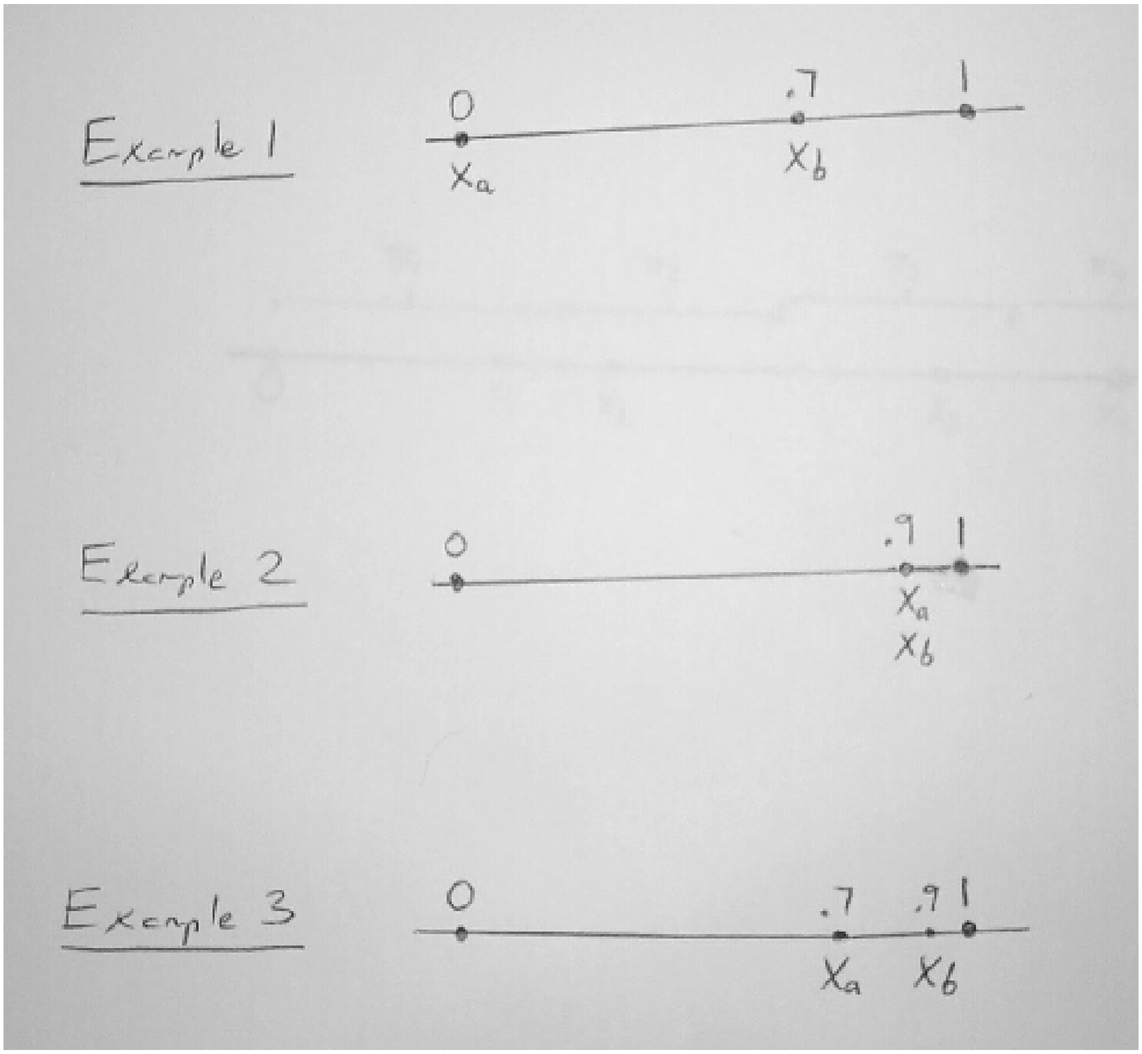

Figure 5: Numerical Examples for Hotelling Pricing

Example 1. Everything works out simply

Try $x_a = 0, x_b = 0.7$ and $\theta = 0.5$. Then equation (10) says $p_a = (2 + 0 + 0.7)0.5/3 = 0.45$ and $p_b = (4 - 0 - 0.7)0.5/3 = 0.55$. Equation (9) says that $x^* = \frac{1}{2*0.5}[(0.55 - 0.45) + 0.5(0.0 + 0.7)] = 0.45$.

In Example 1, there is a pure strategy equilibrium and the equations generated sensible numbers given the parameters we chose.

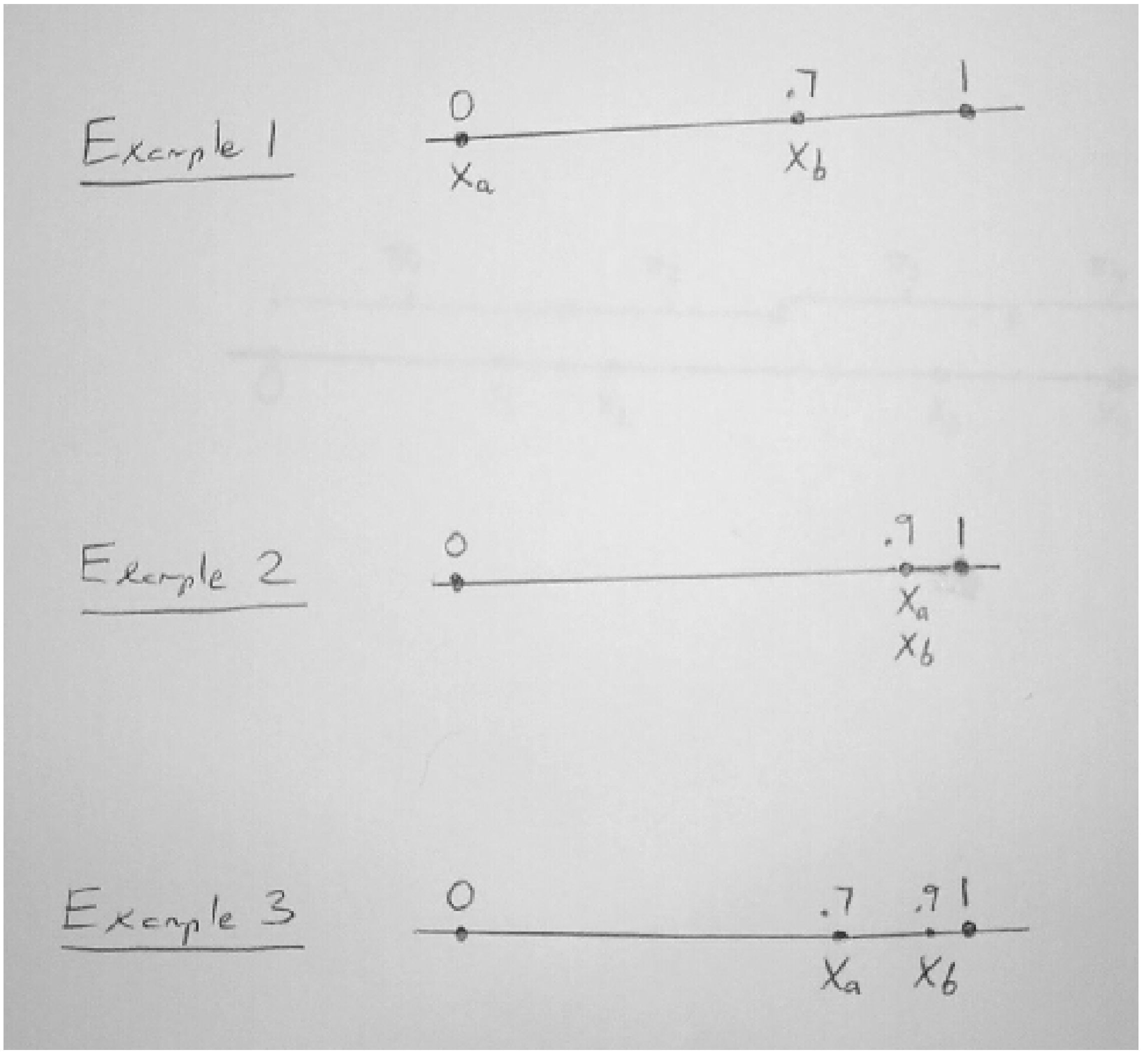

Figure 5: Numerical Examples for Hotelling Pricing

Example 2. Same location – but different prices?

Try $x_a = 0.9, x_b = 0.9$ and $\theta = 0.5$. Then equation (10) says $p_a = (2.0+0.9+0.9)0.5/3 \approx 0.63$ and $p_b = (4.0-0.9-0.9)0.5/3 \approx 0.37$.

The equations generate numbers that seem innocuous until one realizes that if both firms are located at 0.9, but $p_a = 0.63$ and $p_b = 0.37$, then Brydox will capture the entire market!

The result is nonsense, because equation (10)'s derivation relied on the assumption that $x_a < x_b$, which is false in this example.

(draw in Figure 5c by hand)

Example 3. Locations too near each other.

$x^* < x_a < x_b$. Try $x_a = 0.7, x_b = 0.9$ and $\theta = 0.5$. Then equation (10) says that $p_a = (2.0 + 0.7 + 0.9)0.5/3 = 0.6$ and $p_b = (4 - 0.7 - 0.9)0.5/3 = 0.4$. As for the split of the market, equation (9) says that $x^* = \frac{1}{2*0.5}[(0.4 - 0.6) + 0.5(0.7 + 0.9)] = 0.6$.

If the market splits at $x^* = 0.6$ but $x_a = 0.7$ and $x_b = 0.9$, the result violates our implicit assumption that the players split the market.

Equation (9) is based on the premise that there does exist some indifferent consumer, and when that is a false premise, as under the parameters of Example 3, equation (9) will still spit out a value of x^*, but the value will not mean anything.

In fact the consumer at $x = 0.6$ is not really indifferent between Apex and Brydox. He could buy from Apex at a total cost of $0.6 + 0.1(0.5) = 0.65$ or from Brydox, at a total cost of $0.4 + 0.3\ (0.5) = 0.55$. There exists no consumer who strictly prefers Apex.

The Hotelling Location Game
(Hotelling [1929])

Players
n Sellers.

The Order of Play
The sellers simultaneously choose locations $x_i \in [0, 1]$.

Payoffs
Consumers are distributed along the interval [0,1] with a uniform density equal to one. The price equals one, and production costs are zero. The sellers are ordered by their location so $x_1 \leq x_2 \leq \ldots \leq x_n$, $x_0 \equiv 0$ and $x_{n+1} \equiv 1$. Seller i attracts half the consumers from the gaps on each side of him, as shown in Figure 14.6, so that his payoff is

$$\pi_1 = x_1 + \frac{x_2 - x_1}{2}, \tag{11}$$

$$\pi_n = \frac{x_n - x_{n-1}}{2} + 1 - x_n, \tag{12}$$

or, for $i = 2, \ldots n - 1$,

$$\pi_i = \frac{x_i - x_{i-1}}{2} + \frac{x_{i+1} - x_i}{2}. \tag{13}$$

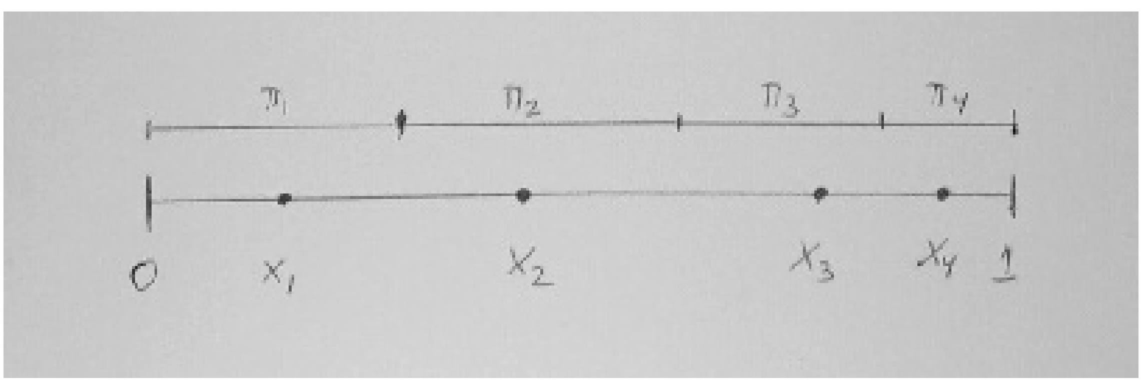

Figure 6: Payoffs in the Hotelling Location Game

With **one seller**, the location does not matter in this model, since the consumers are captive. If price were a choice variable and demand were elastic, we would expect the monopolist to locate at $x = 0.5$.

With **two sellers**, both firms locate at $x = 0.5$, regardless of whether or not demand is elastic. This is a stable Nash equilibrium, as can be seen by inspecting Figure 4 and imagining best responses to each other's location. The best response is always to locate ε closer to the center of the interval than one's rival. When both firms do this, they end up splitting the market since both of them end up exactly at the center.

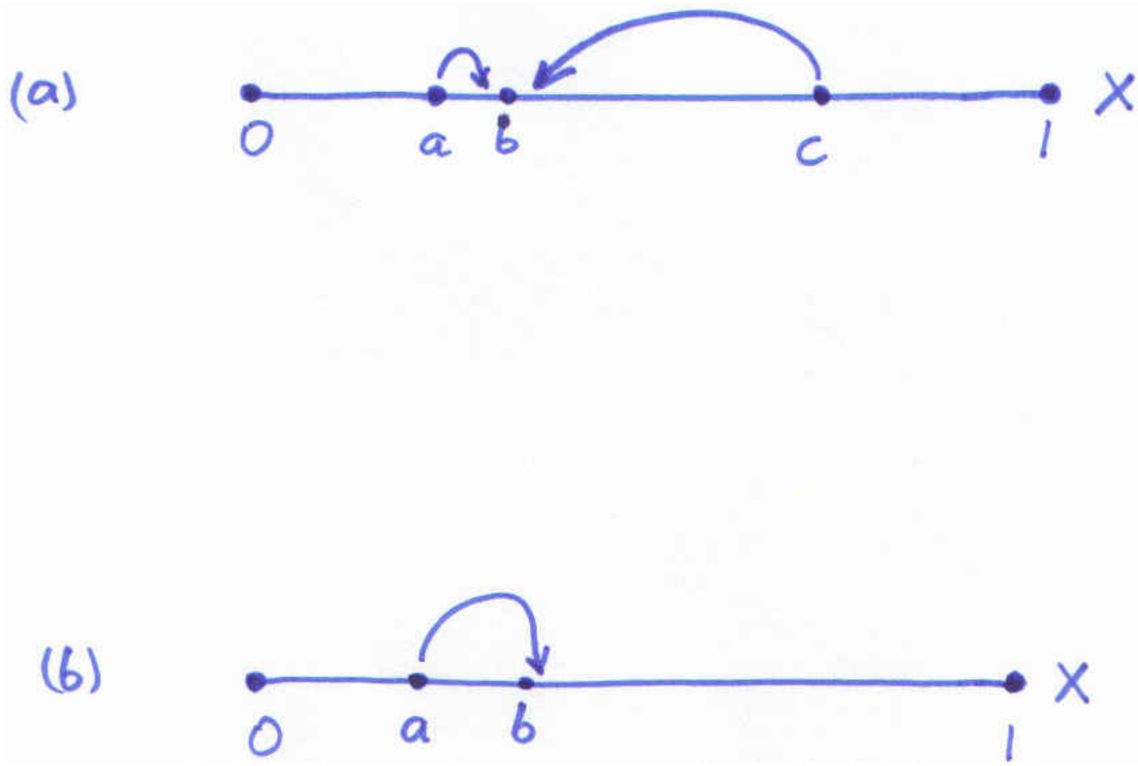

Figure 7: Nonexistence of pure strategies with three players

With **three sellers** the model does not have a Nash equilibrium in pure strategies. Consider any strategy profile in which each player locates at a separate point. Such a strategy profile is not an equilibrium, because the two players nearest the ends would edge in to squeeze the middle player's market share. But if a strategy profile has any two players at the same point a, as in Figure 7, the third player would be able to acquire a share of at least $(0.5 - \epsilon)$ by moving next to them at b; and if the third player's share is that large, one of the doubled-up players would deviate by jumping to his other side and capturing his entire market share. The only equilibrium is in mixed strategies.

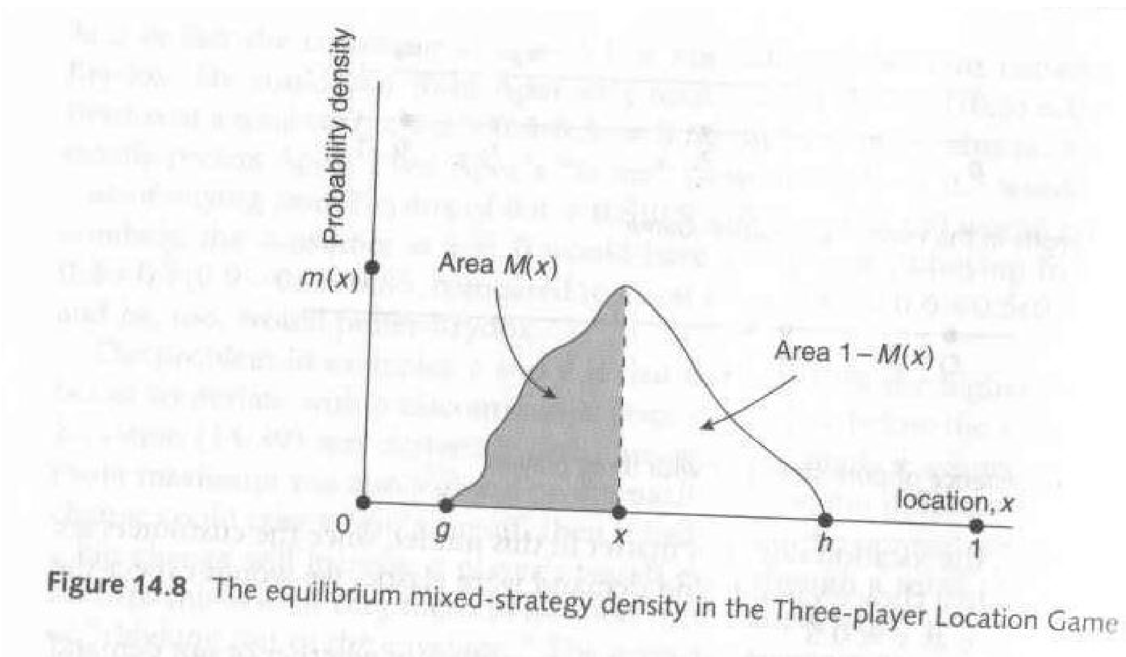

Figure 8: The Equilibrium Mixed-Strategy Density in the Three-Player Location Game

Suppose all three players use the same mixing density, with $m(x)$ the probability density for location x, and positive density on the support $[g, h]$, as depicted in Figure 8.

We will need the density for the distribution of the minimum of the locations of Players 2 and 3.

Player 2 has location x with density $m(x)$, and Player 3's location is greater than that with probability $1 - M(x)$, letting M denote the cumulative distribution, so:

density ($x2 = x$, $x2 < x3 = m(x)[1 - M(x)]$.

The density for either Player 2 or Player 3 choosing x and it being smaller than the other firm's location is then $2m(x)[1 - M(x)]$:

density(Minimum of x2 and x3 equalling x) $= 2m(x)[1 - M(x)]$.

We just found (a) and (b)

(a) density ($x2 = x$, $x2 < x3 = m(x)[1 - M(x)]$.

(b) density(Minimum of x2 and x3 equalling x) $= 2m(x)[1 - M(x)]$.

If Player 1 chooses $x = g$ then his expected payoff is

$$\pi_1(x_1 = g) = g + \int_g^h 2m(x)[1 - M(x)] \left(\frac{x - g}{2}\right) dx, \qquad (14)$$

where g is the safe set of consumers to his left, $2m(x)[1 - M(x)]$ is the density for x being the next biggest location of a firm, and $\frac{x-g}{2}$ is Player 1's share of the consumers between his own location of g and the next biggest location.

If Player 1 chooses $x = h$ his expected payoff is, similarly,

$$\pi_1(x_1 = h) = (1 - h) + \int_g^h 2m(x)M(x) \left(\frac{h - x}{2}\right) dx, \qquad (15)$$

where $(1 - h)$ is the set of safe consumers to his right.

In a mixed strategy equilibrium, Player 1's payoffs from these two pure strategies must be equal, and they are also equal to his payoff from a location of 0.5, which we can plausibly guess is in the support of his mixing distribution. Going on from this point, the algebra and calculus start to become fierce. Shaked (1982) has computed the symmetric mixing probability density $m(x)$ to be:

$$m(x) = \begin{cases} 2 \text{ if } \frac{1}{4} \leq x \leq \frac{3}{4} \\ \\ 0 \text{ otherwise} \end{cases} \qquad (16)$$

Asymmetric equilibria also exist.

Since prices are inflexible, the competitive market does not achieve efficiency.

A benevolent social planner or a monopolist who could charge higher prices if he located his outlets closer to more consumers would choose different locations than competing firms.

In particular, when two competing firms both locate in the center of the line, consumers are no better off than if there were just one firm.

As shown in Figure 10, the average distance of a consumer from a seller would be minimized by setting $x_1 = 0.25$ and $x_2 = 0.75$, the locations that would be chosen either by the social planner or the monopolist.

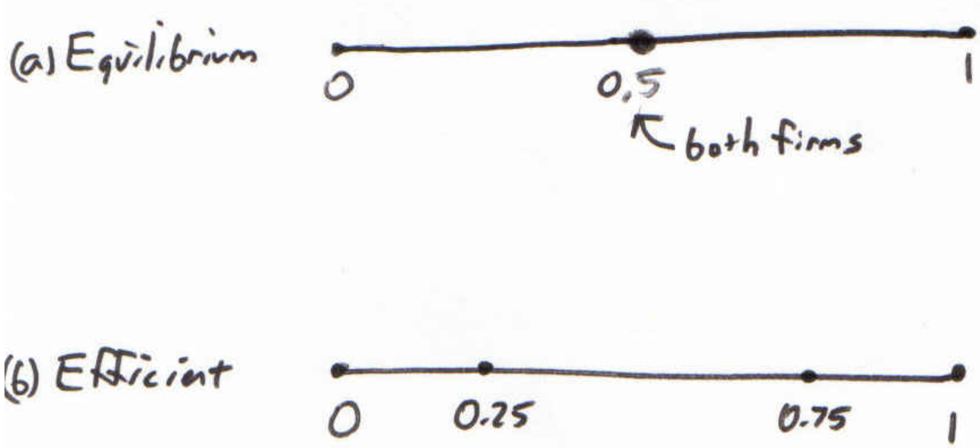

Figure 10: Equilibrium versus Efficiency

The Hotelling Location Model is well suited to politics.

Vertical Differentiation I: Monopoly Quality Choice

Players

A seller and a continuum of buyers.

The Order of Play

0 Nature assigns quality values to a continuum of buyers of length 1. Half of them are "weak" buyers ($\theta = 0$) who value high quality at 20 and low quality at 10. Half of them are "strong" buyers ($\theta = 1$) who value high quality at 50 and low quality at 15.

1 The seller picks quality s to be either $s0$ or 1.

2 The seller picks price p from the interval $[0, \infty)$.

3 Each buyer chooses one unit of a good, or refrains from buying. The seller produces at constant marginal cost $c = 1$, which does not vary with quality.

Payoffs

$$\pi_{seller} = (p - 1)q. \quad (17)$$

The buyer's payoff is zero if he does not buy. If he does buy, it is

$$\pi_{buyer} = (10 + 5\theta) + (10 + 25\theta)s - p. \quad (18)$$

OPTIMAL SELLER PRICE AND QUANTITY

Payoffs

$$\pi_{seller} = (p-1)q. \tag{19}$$

The buyer's payoff is zero if he does not buy. If he does buy, it is

$$\pi_{buyer} = (10+5\theta) + (10+25\theta)s - p. \tag{20}$$

The seller should clearly set the quality to be high,since then he can charge more to the buyer (though note that this runs contrary to a common misimpression that a monopoly will result in lower quality than a competitive market.) The price should be either 50, which is the most the strong buyers would pay, or 20, the most the weak buyers would pay. Since $\pi(50) = 0.5(50-1) = 24.5$ and $\pi(20) = 0.5(20-1) + 0.5(20-1) = 19$, the seller should choose $p = 50$. Separation (by inducing only the strong buyer to buy) is better for the seller than pooling.

Vertical Differentiation II: Crimping the Product

Players

A seller and a continuum of buyers.

The Order of Play

0 Nature assigns quality values to a continuum of buyers of length 1. Half of them are "weak" buyers ($\theta = 0$) who value high quality at 20 and low quality at 10. Half of them are "strong" buyers ($\theta = 1$) who value high quality at 50 and low quality at 15.
1 The seller decides to sell both qualities $s = 0$ and $s = 1$ or just one of them.
2 The seller picks prices p_L and p_H from the interval $[0, \infty)$.
3 Each buyer chooses one unit of a good, or refrains from buying. The seller produces at constant marginal cost $c = 1$, which does not vary with quality.

Payoffs

$$\pi_{seller} = (p_L - 1)q_L + (p_H - 1)q_H. \tag{21}$$

and

$$\pi_{buyer} = (10 + 5\theta) + (10 + 25\theta)s - p. \tag{22}$$

This is a problem of mechanism design. The seller needs to pick p_1, and p_2 to satisfy incentive compatibility and participation constraints if he wants to offer two qualities with positive sales of both, and he also needs to decide if that is more profitable than offering just one quality.

We already solved the one-quality problem in Vertical Differentiation I, yielding profit of 24.5. The monopolist cannot simply add a second, low- quality, low-price good for the weak buyers, because the strong buyers, who derive zero payoff from the high-quality good, would switch to the low-quality good, which would give them a positive payoff. In equilibrium, the monopolist will have to give the strong buyers a positive payoff. Their participation constraint will be non-binding, as we have found so many times before for the "good" type.

Following the usual pattern, the participation constraint for the weak buyers will be binding, so $p_L = 10$. The self-selection constraint for the strong buyers will also be binding, so

$$\pi_{strong}(L) = 15 - p_L = 50 - p_H. \tag{23}$$

Since $p_L = 10$, this results in $p_H = 45$. The price for high quality must be at least 35 higher than the price for low quality to induce separation of the buyer types.

Profits will now be:

$$\pi_{seller} = (10 - 1)(0.5) + (44 - 1)(0.5) = 26. \tag{24}$$

This exceeds the one-quality profit of 24.5, so it is optimal for the seller to sell two qualities.

This result, is, of course, dependent on the parameters chosen, but it is nonetheless a fascinating special case, and one which is perhaps no more special than the other special case, in which the seller finds that profits are maximized with just one quality. The outcome of allowing price discrimination is a pareto improvement. The seller is better off, because profit has risen from 24.5 to 26. The strong buyers are better off, because the price they pay has fallen from 50 to 45. And the weak buyers are no worse off. In Vertical Differentiation I their payoff was zero because they chose not to buy; in Vertical Differentiation I their payoffs are zero because they buy at a price exactly equal to their value for the good.

Indeed, we can go further. Suppose the cost for the low-quality good was actually *higher* than for the high-quality good, e.g., $p_L = 3$ and $p_H = 1$, because the good is normally produced as high quality and needs to be purposely damaged before it becomes low quality. The price-discrimination profit in (24) would then be $\pi_{seller} = (10-3)(0.5) + (44-1)(0.5) = 25$. Since that is still higher than 24.5, the seller would still price- discriminate. The buyers' payoffs would be unaffected. Thus, allowing the seller to damage some of the good at a cost in real resources of 2 per unit, converting it from high to low quality, can result in a pareto improvement!

Vertical Differentiation III: Duopoly Quality Choice

Players

Two sellers and a continuum of buyers.

The Order of Play

0 Nature assigns quality values to a continuum of buyers of length 1. Half of them are "weak" buyers ($\theta = 0$) who value high quality at 20 and low quality at 10. Half of them are "strong" buyers ($\theta = 1$) who value high quality at 50 and low quality at 15.

1 Sellers 1 and 2 simultaneously choose values for s_1 and s_2 from the set $\{s_L = 0, s_H = 1\}$. They may both choose the same value.

2 Sellers 1 and 2 simultaneously choose prices p_1 and p_2 from the interval $[0, \infty)$.

3 Each buyer chooses one unit of a good, or refrains from buying. The sellers produce at constant marginal cost $c = 1$, which does not vary with quality.

Payoffs

$$\pi_{seller} = (p - 1)q \tag{25}$$

and

$$\pi_{buyer} = (10 + 5\theta) + (10 + 25\theta)s - p. \tag{26}$$

If both sellers both choose the same quality level, their profits will be zero, but if they choose different quality levels, profits will be positive. Thus, there are three possible equilibria in the quality stage of the game: (Low, High), (High, Low), and a symmetric mixed-strategy equilibrium. Let us consider the pure-strategy equilibria first, and without loss of generality suppose that Seller 1 is the low-quality seller and Seller 2 is the high- quality seller.

(1) The equilibrium prices of Vertical Differentiation II, ($p_L = 10, p_H = 45$), will no longer be equilibrium prices. The problem is that the low-quality seller would deviate to $p_L = 9$, doubling his sales for a small reduction in price.

(2) Indeed, there is no pure-strategy equilibrium in prices. We have seen that ($p_L = 10, p_H = 45$) is not an equilibrium, even though $p_H = 45$ is the high- quality seller's best response to $p_L = 10$. $P_L > 10$ will attract no buyers, so that cannot be part of an equilibrium. Suppose $P_L \in (1, 10)$. The response of the high-quality seller will be to set $p_H = p_L + 35$, in which case the low-quality seller can increase his profits by slightly reducing p_L and doubling his sales. The only price left for the low-quality seller that does not generate negative profits is $p_L = 1$, but that yields zero profits, and so is worse than $p_L = 10$. So no choice of p_L is part of a pure-strategy equilibrium.

(3) As always, an equilibrium does exist, so it must be in mixed strategies, as shown below.

The Asymmetric Equilibrium: Pure Strategies for Quality, Mixed for Price

The low-quality seller picks p_L on the support $[5.5, 10]$ using the cumulative distribution

$$F(p_L) = 1 - \left(\frac{39.5}{p_L + 34} \right) \tag{27}$$

with an atom of probability $\frac{39.5}{44}$ at $p_L = 10$.

The high-quality seller picks p_H on the support $[40.5, 45]$ using the cumulative distribution

$$G(p_H) = 2 - \left(\frac{9}{p_H - 36} \right) \tag{28}$$

Weak buyers from the low-quality seller if $10 - p_L \geq 20 - p_H$, which is always true in equilibrium. Strong buyers buy from the low-quality seller if $15 - p_L > 50 - p_H$, which has positive probability, and otherwise from the high-quality seller.

This equilibrium is noteworthy because it includes a probability atom in the mixed-strategy distribution, something not uncommon in pricing games.

To start deriving this equilibrium, let us conjecture that the low-quality seller will not include any prices above 10 in his mixing support but will include $p_L = 10$ itself. That is plausible because he would lose all the low-quality buyers at prices above 10, but $p_L = 10$ yields maximal profits whenever p_H is low enough that only weak consumers buy low quality.

The low-quality seller's profit from $p_L = 10$ is $\pi_L(p = 10) = 0.5(10 - 1) = 4.5$. Thus, the lower bound of the support of his mixing distribution (denote it by a_L) must also yield a profit of 4.5. There is no point in charging a price less than the price which would capture even the strong consumers with probability one, in which case

$$\pi_L(a_L) = 0.5(a_L - 1) + 0.5(a_L - 1) = 4.5, \tag{29}$$

and $a_L = 5.5$. Thus, the low-quality seller mixes on $[5.5, 10]$.

On that mixing support, the low-quality seller's profit must equal 4.5 for any price. Thus,

$$\begin{aligned}\pi_L(p_L) = 4.5 &= 0.5(p_L - 1) + 0.5(p_L - 1)Prob(15 - p_L > 50 - p_H) \\ &= 0.5(p_L - 1) + 0.5(p_L - 1)Prob(p_H > 35 + p_L) \\ &= 0.5(p_L - 1) + 0.5(p_L - 1)[1 - G(35 + p_L)]\end{aligned} \tag{30}$$

From the previous page: On that mixing support, the low-quality seller's profit must equal 4.5 for any price. Thus,

$$\pi_L(p_L) = 0.5(p_L - 1) + 0.5(p_L - 1)[1 - G(35 + p_L)]$$

Thus, the $G(p_H)$ function is such that

$$1 - G(35 + p_L) = \frac{4.5}{0.5(p_L - 1)} - 1 \tag{31}$$

and

$$G(35 + p_L) = 2 - \left(\frac{4.5}{0.5(p_L - 1)}\right). \tag{32}$$

We want a G function with the argument p_H, not $(35 + p_L)$, so let's shift the argument by 35:

$$G(p_H) = 2 - \left(\frac{4.5}{0.5([p_H - 35] - 1)}\right) = 2 - \left(\frac{9}{p_H - 36}\right). \tag{33}$$

As explained in Chapter 3, what we have just done is to find the strategy for the high-quality seller that makes the low-quality seller indifferent among all the values of p_L in his mixing support.

We can find the support of the high-quality seller's mixing distribution by finding values a_H and b_H such that $G(a_H) = 0$ and $G(b_H) = 1$, so

$$G(a_H) = 2 - \left(\frac{9}{a_H - 36}\right) = 0, \tag{34}$$

which yields $a_H = 40.5$, and

$$G(b_H) = 2 - \left(\frac{9}{(0) \cdot b_H - 36}\right) = 1, \tag{35}$$

which yields $b_H = 45$. Thus the support of the high-quality seller's mixing distribution is $[40.5, 45]$.

Now let us find the low-quality seller's mixing distribution, $F(p_L)$. At $p_H = 40.5$, the high-quality seller has zero probability of losing the strong buyers to the low-quality seller, so his profit is $0.5(40.5 - 1) = 19.75$. Now comes the tricky step. At $p_h = 45$, if the high-quality seller had probability one of losing the strong buyers to the low-quality seller, his his profit would be zero, and he would strictly prefer $p_H = 40.5$. Thus, it must be that at $p_h = 45$ there is strictly positive probability that $p_L = 10$— not just a positive density. So let us continue, using our finding that the profit of the high-quality seller must be 19.75 from any price in the mixing support. Then,

$$\begin{aligned} \pi_H(p_H) = 19.75 \; &= 0.5(p_H - 1)Prob(15 - p_L < 50 - p_H) \\ &= 0.5(p_H - 1)Prob(p_H - 35 < p_L) \\ &= 0.5(p_H - 1)[1 - F(p_H - 35)] \end{aligned} \tag{36}$$

so

$$F(p_H - 35) = 1 - \left(\frac{19.75}{0.5(p_H - 1)}\right). \tag{37}$$

Using the same substitution trick as in equation (33), putting p_L instead of $(p_H - 35)$ as the argument for F, we get

$$F(p_L) \; = 1 - \left(\frac{19.75}{0.5(p_L + 35 - 1)}\right) = 1 - \left(\frac{39.5}{p_L + 34}\right) \tag{38}$$

In particular, note that

$$F(5.5) = 1 - \left(\frac{39.5}{5.5 + 34}\right) = 0, \tag{39}$$

confirming our earlier finding that the minimum p_L used is 5.5, and

$$F(10) = 1 - \left(\frac{39.5}{10 + 34}\right) = 1 - \frac{39.5}{44} < 1. \tag{40}$$

$$F(10) = 1 - \left(\frac{39.5}{10 + 34} \right) = 1 - \frac{39.5}{44} < 1.$$

Equation (40) shows that at the upper bound of the low-quality seller's mixing support the cumulative mixing distribution does not equal 1, an oddity we usually do not see in mixing distributions. What it implies is that there is an atom of probability at $p_L = 10$, soaking up all the remaining probability beyond what equation (40) yields for the prices below 10. The atom must equal $\frac{39.5}{44} \approx 0.9$.

Happily, this solves our paradox of zero high-quality seller profit at $p_H = 45$. If $p_L = 10$ has probability $\frac{39.5}{44}$, the profit from $p_H = 45$ is $0.5(\frac{39.5}{44})(45 - 1) = 19.75$. Thus, the profit from $p_H = 45$ is the same as from $p_H = 40.5$, and the seller is willing to mix between them.

The duopoly sellers' profits are 4.5 (for low-quality) and 19.75 (for high quality) in the asymmetric equilibrium of Vertical Differentiation III, a total of 24.25 for the industry. This is less than either the 24.5 earned by the nondiscriminating monopolist of Vertical Differentiation I or the 26 earned by the discriminating monopolist of Vertical Differentiation II. But what about the mixed-strategy equilibrium for Vertical Differentiation III?

The Symmetric Equilibrium: Mixed Strategies for Both Quality and Price

Each player chooses low quality with probability $\alpha = 4.5/24.25$ and high quality otherwise. If they choose the same quality, they next both choose a price equal to 1, marginal cost. If they choose different qualities, they choose prices according to the mixing distributions in the asymmetric equilibrium.

This equilibrium is easier to explain. Working back from the end, if they choose the same qualities, the two firms are in undifferentiated price competition and will choose prices equal to marginal cost, with payoffs of zero. If they choose different qualities, they are in the same situation as they would be in the asymmetric equilibrium, with expected payoffs of 4.5 for the low-quality firm and 19.75 for the high-quality firm. As for choice of product quality, the expected payoffs from each quality must be equal in equilibrium, so there must be a higher probability of both choosing high-quality:

$$\pi(Low) = \alpha(0) + (1 - \alpha)4.5 = \pi(High) = \alpha(19.75) + (1 - \alpha)(0). \tag{41}$$

Solving equation (41) yields $\alpha = 4.5/24.25 \approx 0.17$, in which case each player's payoff is about 3.75. Thus, even if a player is stuck in the role of low-quality seller in the pure- strategy equilibrium, with an expected payoff of 4.5, that is better than the expected payoff he would get in the "fairer" symmetric equilibrium.

We can conclude that if the players could somehow arrange what equilibrium would be played out, they would arrange for a pure-strategy equilibrium, perhaps by use of cheap talk and some random focal point variable.

Or, perhaps they could change the rules of the game so that they would choose qualities sequentially. Suppose one seller gets to choose quality first. He would of course choose high quality, for a payoff of 19.75. The second-mover, hwoever, choosing low-quality, would have a payoff of 4.5, better than the expected payoff in the symmetric mixed-strategy equilibrium of the simultaneous quality-choice game. This is the same phenomenon as the pareto superiority of a sequential version of the Battle of the Sexes over the symmetric mixed-strategy equilibrium of the simultaneous-move game.

What if Seller 1 chooses both quality and price first, and Seller 2 responds with quality and price? If Seller 1 chooses low quality, then his optimal price is $p_L = 10$, since the second player will choose high quality and a price low enough to attract the strong buyers— $p_H = 45$, in equilibrium— so Seller 1's payoff would be 0.5(10-1) = 4.5. If Seller 1 chooses high quality, then his optimal price is $p_H = 40.5$, since the second player will choose low quality and would choose a price high enough to lure away the strong buyers if $p_H < 40.5$. If, however, $p_H = 40.5$, Seller 2 would give up on attracting the strong buyers and pick $p_L = 10$. Thus, if Seller 1 chooses both quality and price first, he will choose high quality and $p_H = 40.5$ while Seller 2 will choose low quality and $p_L = 10$, resulting in the same payoffs as in the asymmetric equilibrium of the simultaneous-move game, though no longer in mixed strategies.

What Product Differentiation III shows us is that product differentiation can take place in oligopoly vertically as well as horizontally.

Head-to-head competition reduces profits, so firms will try to differentiate in any way that they can.

This increases their profits, but it can also benefit consumers—though more obviously in the case of horizontal differentiation than in vertical.

Keep in mind, though, that in our games here we have assumed that high quality costs no more than low quality.

Usually high quality is more expensive, which means that having more than one quality level can be efficient.

Often poor people prefer lower quality, given the cost of higher quality, and even a social planner would provide a variety of quality levels.

Here, we see that even when only high quality would be provided in the first-best, it is better that a monopolist provide two qualities than one, and a duopoly is still better for consumers.