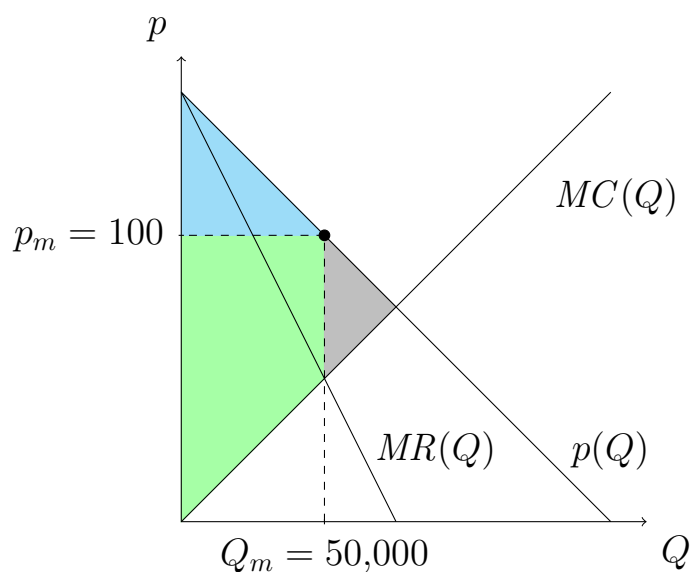# Lecture Note 5: Price Discrimination

Imagine you are in charge of Disneyland's pricing division.

- Congratulations! ECN 100B has served you well.

Your mission: pick the ticket price p to maximize profits.

- Write the profit function, take the FOC, set $MR = MC$.

- End up with $p_m = \$100$ and $Q_m = 50{,}000$ visitors per day.

- Lots of profit, lots of consumer surplus. Everybody wins!

- Green: producer surplus

- Blue: consumer surplus (people who'd pay > $100)

- Gray: DWL (people who'd pay more than MC but < $100)

But something is bugging you . . .

- Lots of people would pay over $100. Can we charge them more?

- Lots of people won't pay $100. Should we charge them less?

Can we turn the CS and DWL into even more profit?

Types of pricing practices

uniform pricing: charging the same price for every unit you sell.

- Example: a gas station charges \$3.75 per gallon of gas, no matter who you are, no matter how many gallons you buy.

price discrimination: charging different per-unit prices for the same product, depending on who is buying it or on how many units they buy.

Economists subdivide price discrimination into three categories.[1]

personalized pricing: a firm offers each consumer her own individual price. We'll focus on an extreme form called *perfect* price discrimination.

- Contractors (e.g., exterminators) give their clients personalized prices.

- So do car dealerships.

group pricing: a firm charges different prices to different groups of customers, but the same price for all members of each group.

- Disneyland gives SoCal residents a better deal.

- Comcast gives a low "teaser" rate to new subscribers.

nonlinear pricing: the unit price depends on the quantity purchased.

- Most consumer goods are cheaper if you purchase in bulk.

- Your first Netflix show costs you \$12.99. The rest are free.

We'll restrict our attention to perfect and group price discrimination.

---

[1]Perfect, nonlinear, and group pricing are also called first-, second-, and third-degree price discrimination, respectively. Firms often combine them: e.g., Disneyland charges lower admission fees for SoCal residents (group pricing), but it also gives multi-day discounts (nonlinear pricing).

The gold standard: perfect price discrimination

Perfect price discrimination is a form of personalized pricing in which a firm sells each unit at the maximum price that the consumer will accept.

- Not only does the firm charge each buyer her own *individual* price— it charges her exactly her **reservation price** (willingness to pay).

Suppose that Davis Chevrolet can buy 2020 Chevy Malibus at a constant (marginal) cost of $20,000 per car. Also suppose that they have top-notch salespeople who can tell <u>exactly</u> how much each customer is willing to pay.

Let's consider three different customers:

- Bonnie has just robbed a bank, and she is trying to get out of town. She's willing to pay $60,000 for a Chevy Malibu.

  - The dealership should sell her a car for $60,000. The dealership makes $40,000 in profit; Bonnie gets zero consumer surplus.

- Jisung would pay $10,000 for a Malibu. (He'd rather drive a hybrid.)

  - They'll never be able to make a deal. Any price that Jisung would accept would generate negative profit for the dealership.

- Luu would pay $20,000 for a Malibu. (After all, it's a decent car.)

  - It doesn't really matter. The only price they'd both accept is $20,000, which yields zero surplus for both parties.

We can extend these intuitions to analyze the entire market.

- Suppose the demand curve is $p(Q) = 60 - Q$.
  The demand curve *ranks consumers by their WTP for the good.*

- Suppose that the marginal cost curve is $MC(Q) = 10 + \frac{1}{4}Q$.
  This *ranks items by the cost of producing/supplying them.*

3

The Malibu market

With $p(Q) = 60 - Q$ and $MC(Q) = 10 + \frac{1}{4}Q$, the market looks like this:
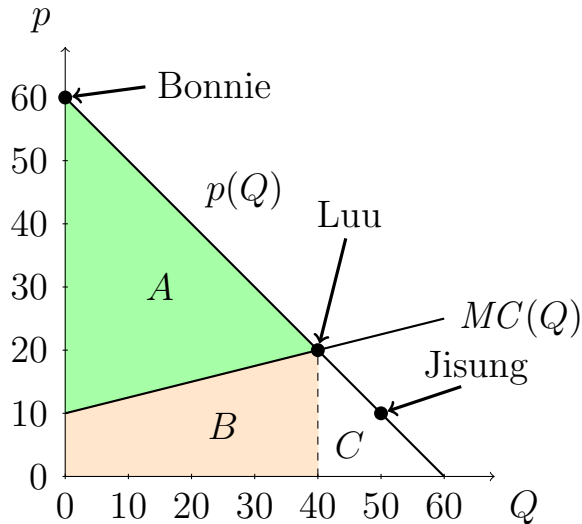
- Total revenue: $A + B$.

- Total (variable) costs: B.

- Profit is revenue $-$ cost: A.

The firm decides what price to offer each consumer, one at a time:

- If $p(Q) > MC(Q)$: charge customer Q the price $p(Q)$.[2]

- If $p(Q) < MC(Q)$: refuse to sell. (We can also think of this as the firm charging $MC(Q)$, knowing that the consumer will walk away.)

- If $p(Q) = MC(Q)$: either charge $p(Q)$ or refuse to sell. (It doesn't really matter whether this marginal transaction happens or not.)

Let's do our usual welfare analysis—calculating everybody's surplus:

- CS $= 0$ (because each consumer is charged their reservation price).

- PS $= 1000$ (which equals the total surplus under perfect comp).

- DWL $= 0$ (there are "no missed opportunities" to make a deal).

We get the Pareto efficient outcome . . . but the firm extracts all the surplus.

---

[2]We're assuming here that, if the customer is indifferent, he buys the car. This assumption isn't very important: the dealership could always charge him $p(Q) - \varepsilon$, where ε is negligibly small. The customer would accept this deal and the dealership would make essentially the same profit.

Group price discrimination

Most firms don't try to charge every customer their own price. And the ones that do are unlikely to achieve truly "perfect" price discrimination.

The main reason is that doing so requires *a ton of information.*[3]

- A uniform-pricing monopolist just needs to know *how many* units it can sell at any given price—what does the demand curve look like?

- A perfect price discriminator has to know *every consumer's WTP*—the *identity* of the customer at each point on the demand curve.

- Some sellers (e.g., pawn shops) charge each buyer a different price, based on their best guess of what price she'll accept. But even good salespeople usually can't extract *all* of the surplus from a trade.[4]

But firms engage in **group price discrimination** all the time.

- Movie theaters give discounts for students, seniors, and veterans.

- UC Davis charges vendors more than visitors for parking.

Disneyland sells tickets to SoCal residents and NorCal residents.

- SoCal demand $= p_S(Q_S)$.

- NorCal demand $= p_N(Q_N)$.

Disneyland's costs depend on its sales to both groups: $C(Q_S, Q_N)$.

The cost of an extra visitor probably doesn't depend on where she's from, so let's assume that costs are $C(Q_S + Q_N)$: they just depend on the sum.

What price should we charge each group? Should we price discriminate?

---

[3]There are other reasons. Price discrimination only works if firms can prevent "resale" from customers with big discounts to customers that face high prices. In some contexts, consumers may view price discrimination as unfair and switch to other sellers. Certain forms of price discrimination (e.g., charging different prices based on a customer's race) are illegal under US law.

[4]Big data may be changing this, by enabling tech firms to better guess buyers' reservation prices.

Maximizing profits with two groups

As always, let's start by writing the profit maximization problem:

$$\max_{Q_S, Q_N} \pi(Q_S, Q_N) = R_S(Q_S) + R_N(Q_N) - C(Q_S + Q_N)$$

Calculus tells us that we need to take one FOC for every choice variable. Partially differentiating profits with respect to Q_S yields the FOC for Q_S:

$$\frac{\partial \pi(Q_S, Q_N)}{\partial Q_S} = MR_S(Q_S) - \frac{\partial}{\partial Q_S} C(Q_S + Q_N) = 0$$

Let's define $Q = Q_S + Q_N$. Using the chain rule,

$$\frac{\partial}{\partial Q_S} C(Q_S + Q_N) = C'(Q_S + Q_N) \frac{\partial (Q_S + Q_N)}{\partial Q_S} = C'(Q) \cdot 1 = MC(Q)$$

So we can write the FOC for Q_S as

$$MR_S(Q_S) = MC(Q)$$

Likewise, we can write the FOC for Q_N as

$$MR_N(Q_N) = MC(Q)$$

Putting it all together, Disneyland chooses Q_S and Q_N to satisfy

$$MR_S(Q_S) = MR_N(Q_N) = MC(Q)$$

To gain intuition for optimality conditions like these, it's helpful to start at the <u>wrong</u> answer and look for a **profitable deviation**.

- Example 1: suppose that $MR_S(Q_S) > MR_N(Q_N)$.
  Raise Q_S by δ, lower Q_N by δ. Revenue rises, costs same $\implies \pi \uparrow$.

- Example 2: suppose that $MR_S(Q_S) < MC(Q)$.
  Lower Q_S by δ. Revenue falls, but costs fall even more $\implies \pi \uparrow$.

Unless both FOCs hold, we aren't optimizing: there's a way to do better.

A numerical example

Suppose that ...

- SoCal's (inverse) demand is $p_S(Q_S) = 90 - Q_S$

- NorCal's (inverse) demand is $p_N(Q_N) = 180 - 2Q_N$

- Costs are $C(Q_S, Q_N) = 60Q_S + 60Q_N$

Disneyland maximizes profits:

$$\max_{Q_S, Q_N} \pi(Q_S, Q_N) = \underbrace{(90 - Q_S)Q_S}_{R_S(Q_S)} + \underbrace{(180 - 2Q_N)Q_N}_{R_N(Q_N)} - 60Q_S - 60Q_N$$

FOC for Q_S : $\quad \dfrac{\partial \pi}{\partial Q_S} = 90 - 2Q_S - 60 = 0 \implies Q_S^* = 15, p_S^* = 75$

FOC for Q_N : $\quad \dfrac{\partial \pi}{\partial Q_N} = 180 - 4Q_N - 60 = 0 \implies Q_N^* = 30, p_N^* = 120$

This example is a bit special: because each FOC just depends on a single choice variable, we can solve the system of equations one line at a time.

What this means is that—in this example—Disneyland's optimal choice of how many tickets to sell to SoCal residents doesn't depend on how many tickets it plans to sell to NorCal customers. It's as though Disneyland is running two separate businesses—one for SoCal and one for NorCal.

We can see this if we rearrange the profit equation:

$$\pi(Q_S, Q_N) = \underbrace{(90 - Q_S)Q_S - 60Q_S}_{\text{profits from SoCal customers}} + \underbrace{(180 - 2Q_N)Q_N - 60Q_N}_{\text{profits from NorCal customers)}}$$

In most cases, however, it's impossible to split up profits like this, and we will have to jointly solve the full system of equations.[5]

---

[5]For example, suppose Disneyland's costs are $C(Q_S, Q_N) = (Q_S + Q_N)^2 = Q_S^2 + 2Q_SQ_N + Q_N^2$. The interaction term $2Q_SQ_N$ depends on both SoCal and NorCal ticket sales. This creates a link between the two choice variables, and Disneyland will have to choose them simultaneously.

The price-sensitive group gets a better deal

$$SoCal: \ p_S(Q_S) = 90 - Q_S \qquad NorCal: \ p_N(Q_N) = 180 - 2Q_N$$
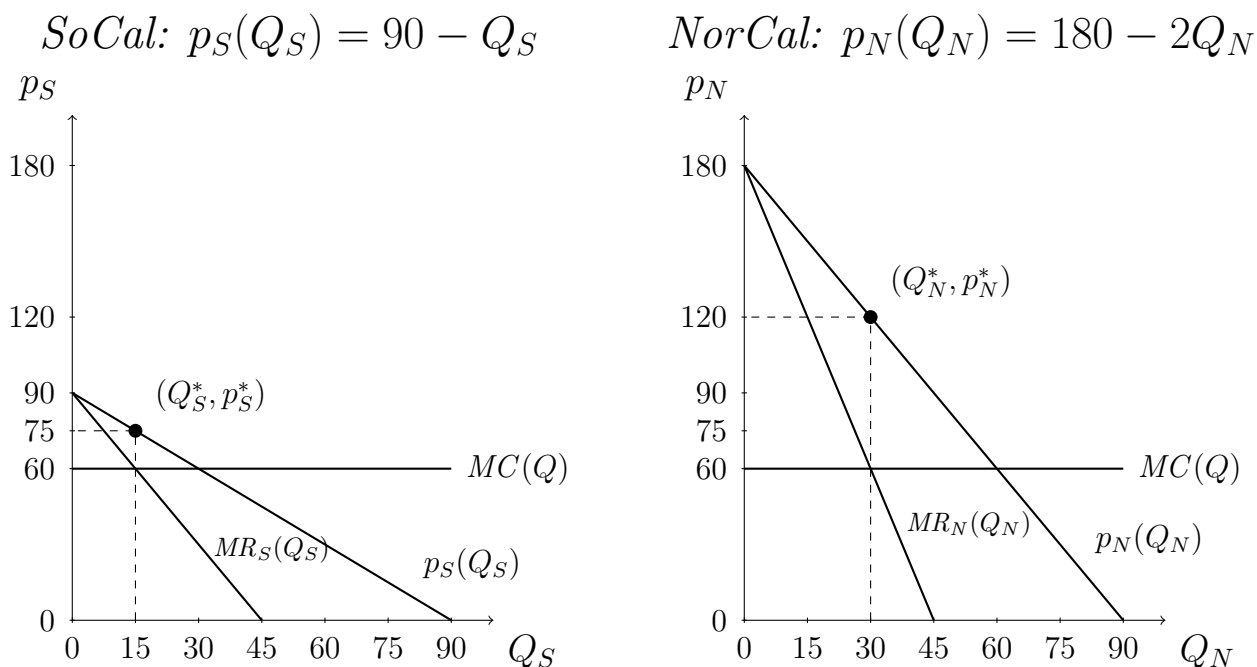
What can we learn from Disneyland's pricing behavior?

- SoCal markup: $\frac{75-60}{60} = 25\%$. NorCal markup: $\frac{120-60}{60} = 100\%$.

- Recall that price markups are bigger when consumers are less elastic.

- For SoCal, $-\frac{1}{1+\varepsilon_S} = 25\% \implies \varepsilon_S = -5$.

- For NorCal, $-\frac{1}{1+\varepsilon_N} = 100\% \implies \varepsilon_N = -2$.

SoCal customers are more elastic, and they face a lower markup.

This is the key point about group price discrimination: *groups that are more price-sensitive get a better deal.* Intuitively, the firm knows it can't rip them off: if it charges too much, price-sensitive consumers walk away.

Group pricing lets a firm "segment the market", charging each group a price tailored to its willingness to pay. Many real-world pricing strategies are designed to segment the market in this manner, giving discounts to price-sensitive consumers and charging high prices to the inelastic ones.

Kinked demand curves

Recall that we obtain the market demand curve by adding up individual consumers' demand curves *horizontally.* We can do the same thing here, taking the horizontal sum of the SoCal and NorCal demand curves.

The market demand curve looks like this:
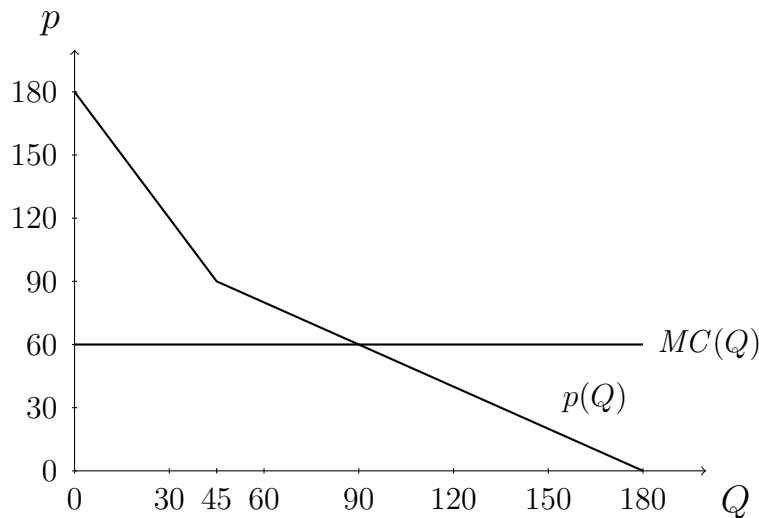
This is a "kinked" demand curve, with a bend or "kink" at $Q = 45$. Drawing kinked demand curves is a bit tricky. Let's do it in two steps.

First, when the price is above 90, nobody from SoCal is willing to buy. Therefore, for $p > 90$, the market demand curve lines up exactly with the NorCal demand curve. That gets us from $Q = 0$ to $Q = 45$:

$$p(Q) = 180 - 2Q \quad \text{for} \quad Q < 45$$

Next, when the price is below 90, some consumers from both SoCal and NorCal are willing to buy. To find out how many consumers *in total* are willing to buy, we have to invert the two demand curves:

$$p_S(Q_S) = 90 - Q_S \quad \Longrightarrow \quad Q_S(p_S) = 90 - p_S$$

$$p_N(Q_N) = 180 - 2Q_N \quad \Longrightarrow \quad Q_N(p_N) = 90 - \frac{1}{2}p_N$$

If Disneyland charges both groups the same price p, then total demand is

$$Q(p) = Q_S(p) + Q_N(p) = 180 - \frac{3}{2}p$$

Inverting this equation, we obtain the market demand curve for $p < 90$, or in other words for $Q \geq 45$:

$$p(Q) = 120 - \frac{2}{3}Q \quad \text{for} \quad Q \geq 45$$

Putting it all together, the market demand curve is defined piecewise by

$$p(Q) = \begin{cases} 180 - 2Q & \text{for} \quad Q < 45 \\ 120 - \frac{2}{3}Q & \text{for} \quad Q \geq 45 \end{cases}$$

That's the function I plotted in the graph above.

What if Disneyland were a uniform-pricing monopolist?

Let's now imagine that Disneyland has to operate as a uniform-pricing monopolist. (Perhaps a California state law has outlawed the practice of charging different prices to families from different parts of the state.) Our goal is to find the pair (Q_m, p_m) that maximizes Disneyland's profits.

Unfortunately, the presence of the kink makes optimization a little more difficult, because the profit function is not differentiable at $Q = 45$. To find the optimal choice Q_m, here's what we have to do:

1. Maximize profits as if demand were $p(Q) = 180 - 2Q$ for all Q. We get a candidate solution $Q_m = 30, p_m = 120$, with $\pi = 1800$.

2. Maximize profits as if demand were $p(Q) = 120 - \frac{2}{3}Q$ for all Q. We get a candidate solution $Q_m = 45, p_m = 90$, with $\pi = 1350$.

3. We now have two candidate solutions. Picking the one that yields higher profits gives us the final answer: $Q_m = 30, p_m = 120$.

So, if Disneyland has to charge everybody the same price, we get this:
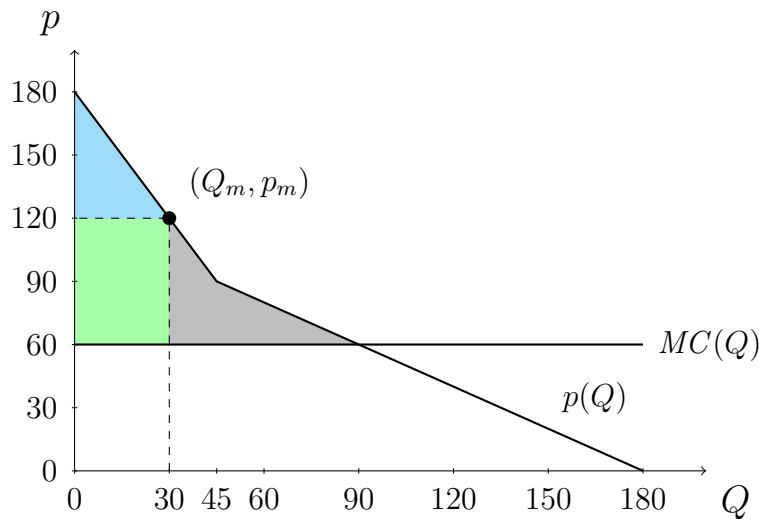
This solution reveals some important insights:

1. Is price discrimination profitable? As a uniform-pricing monopolist, Disneyland would make profit of $\pi = 1800$. Under group price discrimination, Disneyland makes extra profit: $\pi = 2025$.

    This is a general insight: *engaging in price discrimination allows firms to increase their profits.* We'll come back to this shortly.

2. Under group price discrimination, Disneyland sells tickets to both SoCal and NorCal customers. Under uniform pricing, it "gives up" on SoCal: at the uniform price $p_m = 120$, nobody from SoCal buys. Because SoCal residents have much lower WTP than people from NorCal, the only way Disneyland can attract SoCal customers is by lowering the price a lot. Under uniform pricing, Disneyland would have to give a big discount to NorCal, too, and that's not worth it.

    The lesson here is that, when firms have to charge everybody the same price, they sometimes just sell to the top end of the market. Group price discrimination gets them to service the low end, too.

3. Price-sensitive groups tend to benefit from group pricing because firms typically give them discounted prices. In this example, SoCal residents get consumer surplus under group price discrimination, but not under uniform pricing.

4. Price-inelastic groups are generally worse off under group pricing because firms typically charge them higher prices. That doesn't happen in this example, but in the real world Disneyland probably charges NorCal customers more than it would under uniform pricing.

5. Group pricing sometimes reduces the monopoly deadweight loss by getting monopolists to service segments of the market they would otherwise ignore. (Sometimes, however, group pricing raises DWL.)

Revealed preference

I want to conclude by revisiting the question of whether the ability to engage in price discrimination increases Disneyland's profits.

We calculated profits under both uniform and group pricing, and we found that group pricing yields higher profit. But there's a simpler argument:

- Under group pricing, Disneyland <u>could</u> have chosen $p_S = p_N$.

- But they didn't: they maximized profits with $p_S \neq p_N$.

- Thus, uniform pricing ($p_S = p_N$) would have yielded lower profits.

- The <u>ability</u> to price discriminate can never *lower* a firm's profits (they can always choose not to do it) and usually *raises* their profits.

This elegant line of reasoning is called a **revealed preference** argument: Disneyland's behavior "reveals" that it benefits from price discrimination. Whenever we see firms engage in price discrimination, we can safely assume that it's increasing their profits: otherwise they wouldn't be doing it!