UC Berkeley
Haas School of Business
Economic Analysis for Business Decisions
(EWMBA 201A)
Fall 2017

Module II

Competitive markets

– Profit maximization and the firm's supply (PR 8.1-8.6)

– The analysis of competitive markets (PR 9.1-9.6)

Monopolistic markets

– Market power (PR 10.1-10.4)

– Pricing with market power (PR 11.1-11.4)

Competitive markets

Perfectly competitive markets

The *theoretical ideal* of perfect competition rest on three important assumptions:

[1] Price taking

Each individual firm (resp. consumer) sells (resp. buys) a sufficiently small proportion of total market output so its decisions have no impact on market price.

[2] Product homogeneity

The products of all the firms in the market are perfectly substitutable with one another so no firm can raise the price of its product above the price of the other firm without loosing all its business.

$\Rightarrow$ Oil, iron, lumber, cotton and other raw materials and so-called commodities are fairly homogeneous.

[3] Free entry and exit

Firms can easily enter or exit (if cannot make a profit) the market and consumers can easily switch from one firm to another.

$\Rightarrow$ There is fierce competition in the pharmaceutical industry but it cannot be perfectly competitive because firms hold patents that give them unique rights to produce drugs.

Q When a market is competitive?

A Most real-world markets are not perfectly competitive in the sense that each firm faces a horizontal demand curve (more below).

There is no simple rule of thumb to measure the *extent* to which a market is competitive. It is necessary to analyze the strategic interaction among market participants (game theory).

Simple pricing

In the model of perfect competition, each firm must charge the same price per unit to all of its consumers (no matter who the buyer is or how many units the buyer purchases). Simple pricing applies when

– the identity of the buyer cannot be observed or inferred at reasonable cost.

– the firm cannot prevent *arbitrage* among buyers when buyers can purchase multiple units.

Profit maximization by a competitive firm

A firm's profit is the revenue it takes in minus its cost. If we let $R(q)$ the revenue from selling q units, then its profit from selling q units is

$$\pi(q) = R(q) - C(q)$$

where $C(q)$ is the total cost of q units, and if the firm sets a price of p per unit – engages in simple pricing – then $R(q) = pq$.

In choosing the amount to produce and sell, the firm seeks to find the quantity q that maximizes profit $\pi(q)$. We use an asterisk to denote the profit maximizing quantity q^*.

The discrete case

Saying that q^* is the profit-maximizing quantity is the same as saying that

$$\pi(q^*) \geq \pi(q) \text{ for any } q \neq q^*.$$

In particular, consider the quantities $q^* - 1$ and $q^* + 1$. We know that

$$\pi(q^*) \geq \pi(q-1) \text{ and } \pi(q^*) \geq \pi(q+1),$$

and substituting $R(q) - C(q)$ for $\pi(q)$ yields

$$R(q^*) - C(q^*) \geq R(q^* - 1) - C(q^* - 1) \text{ and}$$
$$R(q^*) - C(q^*) \geq R(q^* + 1) - C(q^* + 1).$$

Rearranging,

$$R(q^*) - R(q^* - 1) \geq C(q^*) - C(q^* - 1) \text{ and}$$
$$C(q^* + 1) - C(q^*) \geq R(q^* + 1) - R(q^*).$$

Let $MR(q) = R(q) - R(q - 1)$ and $MC(q) = C(q) - C(q - 1)$ and rewrite this last pair of inequalities as

$$(i)\ MR(q^*) \geq MC(q^*) \text{ and } (ii)\ MR(q + 1^*) \leq MC(q + 1^*).$$

A *necessary* condition for q^* to be the profit-maximizing output is that expressions (i) and (ii) both hold true.

The continuous case

Profit $\pi(q)$ is maximized at a point at which an additional (small) increment to output leave profit unchanged, that is

$$\frac{\Delta \pi}{\Delta q} = \frac{\Delta R}{\Delta q} - \frac{\Delta C}{\Delta q} = 0.$$

Thus, in the continuous case, a *necessary* condition for q^* to be the profit-maximizing output is that

$$\frac{\Delta R}{\Delta q} = \frac{\Delta C}{\Delta q}$$

(the $MR(q) = MC(q)$ rule).

Sufficiency and the shutdown rule

The above results are only necessary conditions; that is, they only identify possible candidates for being the profit-maximizing quantity.

There is a condition, however, that insures that, if the firm should be in business at all, the conditions stated above are also sufficient.

We will establish the sufficiency condition for the continuous case (a similar argument applies to the discrete case).

If the following conditions hold

$$\begin{array}{rl} (i) & MR(q^*) = MC(q^*) \\ (ii) & MR(q) > MC(q) \text{ for all } q < q^* \\ (iii) & MR(q) < MC(q) \text{ for all } q > q^* \end{array}$$

then q^* is the profit-maximizing quantity for the firm to produce (if it should be in business at all).

Another way to view this result is that q^* is the profit-maximizing quantity (if it should be in business at all) if marginal revenue crosses marginal cost once at q^* and does so from above.

The marginal revenue of a competitive firm

In a competitive market, how much output the firm decides to produce and sell have no effect on the market price of the product (price taking). Therefore,

$$R(q) = pq \text{ for all } q.$$

and as a result the marginal revenue, average revenue and price are all equal. As a result, the profit maximizing quantity q^* of a perfectly combative firm satisfies

$$MC(q^*) = p$$

(if it should be in business at all).

– PR Figures 8.3 and 8.4 here –

The firm and market (short-run) supply curves

- The firm's supply curve of the firm specifies how much output the firm will produce at every possible price.

- The firm will produce at a point at which price is equal to marginal cost, but will shut down if price is below average variable cost.

- Therefore, the firm's supply curve is the portion of the marginal cost curve for which marginal cost is greater than average variable cost.

- The industry supply curve is the summation of the supply curves of the individual firms in the market.

– PR Figures 8.6 and 8.9 here –

The important takeaways are

- Marginal revenue equals marginal cost at the optimal quantity produced (this equality may be approximate in the discrete case).

- Marginal revenue comes from an underlying demand curve. Demand curves themselves come from consumer preferences (see below).

The individual and market demand curves

- Consumers tend to buy more of the good that has become cheaper and less of those that become *relatively* more expensive.

- The market demand relate the quantity of a good that all consumers in a market will buy to its price.

- Only factors that influence the demands of many consumers will also effect market demand.

The analysis of competitive markets

- The equilibrium price and quantity in a competitive market maximizes the economic welfare of producers and consumers.

- The model of competitive markets can be used to study the welfare effects of different government polices.

- Next we will evaluate the "gains" and "losses" to consumers and producers from different government polices.

– PR Figures 9.1, 9.2 and 9.5 here –

Monopolistic markets

Monopoly

- In contrast to perfect competition, a **monopoly** is a market that has only one seller but many buyers.

- A **monopsony** is exactly the opposite – is a market that has many sellers but only one buyer.

- Monopoly and monopsony are forms of **market power** – an ability to effect the market price.

- Our goal is to understand how market power works and how it effects producers and consumers.

The theory of monopoly is, on the face of it, simple and straightforward, but behind it lie some deep and interesting questions.

[1] How the monopoly came to be a monopoly, and why it stays that way?

[2] If the monopoly makes profit, why does not the industry attract entrants?

Standard stories, if given at all, get very fuzzy at this point. Hands start to wave, hems give away to haws, and on to the next subject...

Perhaps most importantly, the monopolist is the market so it completely controls the amount of output offered for sale (or the price per unit).

– When the monopolist decides how much to produce, the price per unit that it receives follows directly from the market demand.

– When the monopolist determines a price, the quantity it will sell at that price follows from the market demand.

The standard theory is that the monopoly set a quantity of output $Q \geq 0$ to maximize its profits (but we can also think of the monopoly choosing a price p).

Average revenue and marginal revenue

The monopolist's *average revenue* – the price it receives per unit sold – is the market demand curve.

To see the relationship among total, average, and marginal revenue, consider a monopolist facing a linear demand curve

$$P(Q) = A - Q \text{ where } A > 0.$$

Then,

$$R(Q) = P(Q)Q = AQ - Q^2 \qquad MR = \Delta R/\Delta Q = A - 2Q.$$

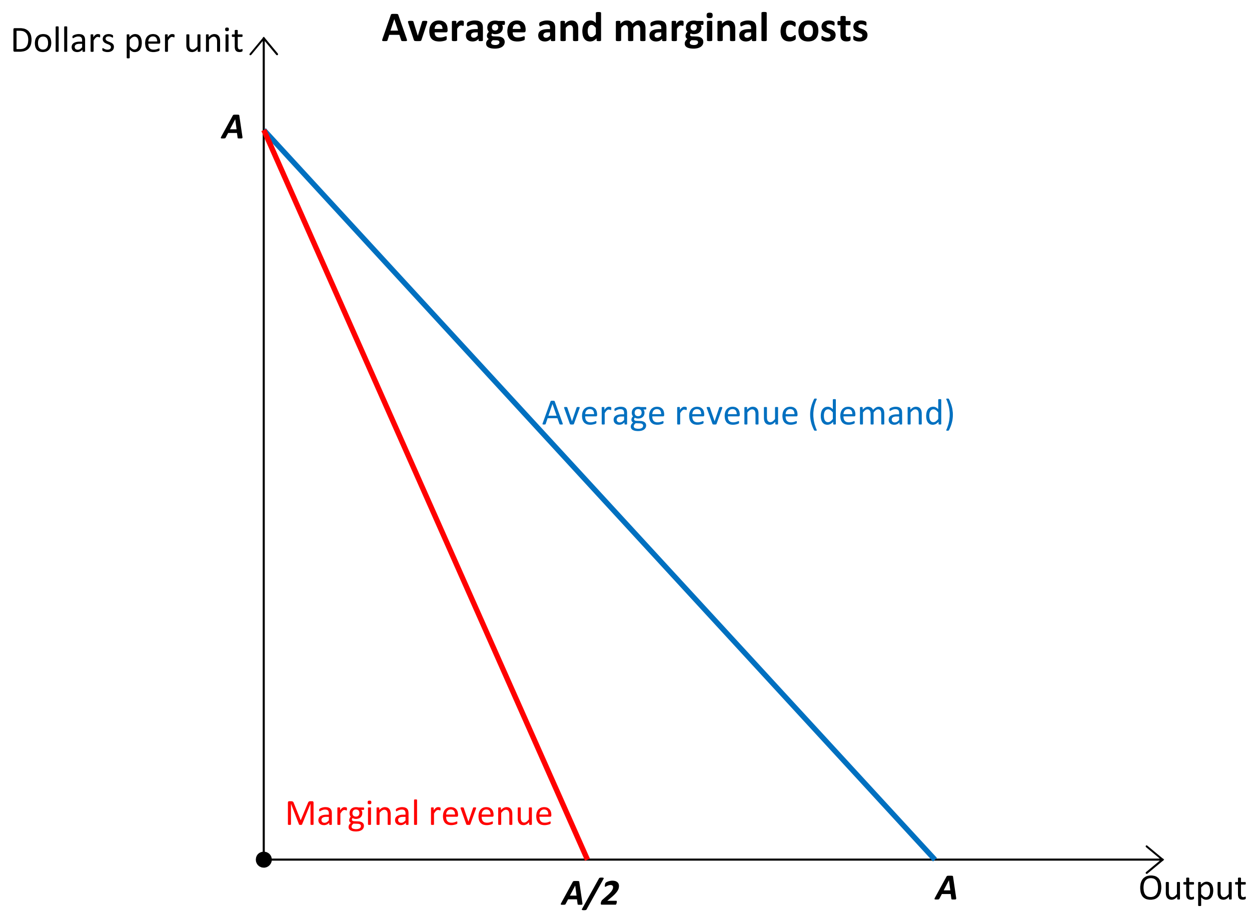
Average and marginal costs
Dollars per unit
A
Average revenue (demand)
Marginal revenue
A/2
A
Output

The monopolist's output decision problem

The monopolist's profit $\pi(Q)$ is the difference between revenue and cost

$$\pi(Q) = R(Q) - C(Q),$$

both of which depend on Q.

As Q increases, π will increase until it reaches a maximum and then start to decrease.

Hence, the profit-maximizing quantity of output Q^* is such that the marginal (incremental) profit resulting from a small increase in Q equals zero.

Algebraically,

$$\Delta\pi/\Delta Q = \Delta R/\Delta Q - \Delta C/\Delta Q = 0,$$

or equivalently,

$$\Delta R/\Delta Q = \Delta C/\Delta Q.$$

That is, we have the slogan that *marginal revenues* MR *equals marginal costs* MC.

An example

Suppose the cost of production is given by

$$C(Q) = 50 + Q^2$$

(a fixed costs of \$50 and a variable costs of and variable costs of Q^2) the demand is given by

$$P(Q) = 40 - Q.$$

Note well that

$$AC = C(Q)/Q = 50/Q + Q \qquad MC = \Delta C/\Delta Q = 2Q$$

$$R(Q) = P(Q)Q = 40Q - Q^2 \qquad MR = 40 - 2Q$$

Setting marginal revenue equal to marginal cost $MR = MC$ gives

$$40 - 2Q = 2Q,$$

or $Q^* = 10$ (reaching the maximum profit of \$150).

Alternatively,

$$\begin{aligned} \pi(Q) &= R(Q) - C(Q) = P(Q)Q - C(Q) \\ &= (40 - Q)Q - 50 - Q^2 = 40Q - Q^2 - 50 - Q^2 \\ &= 40Q - 50 - 2Q^2. \end{aligned}$$

and setting $\Delta\pi/\Delta Q$ equal zero gives $40 - 4Q = 0$, or $Q^* = 10$.

Next we will give a geometrical procedure for doing this.

The monopolist's decision problem

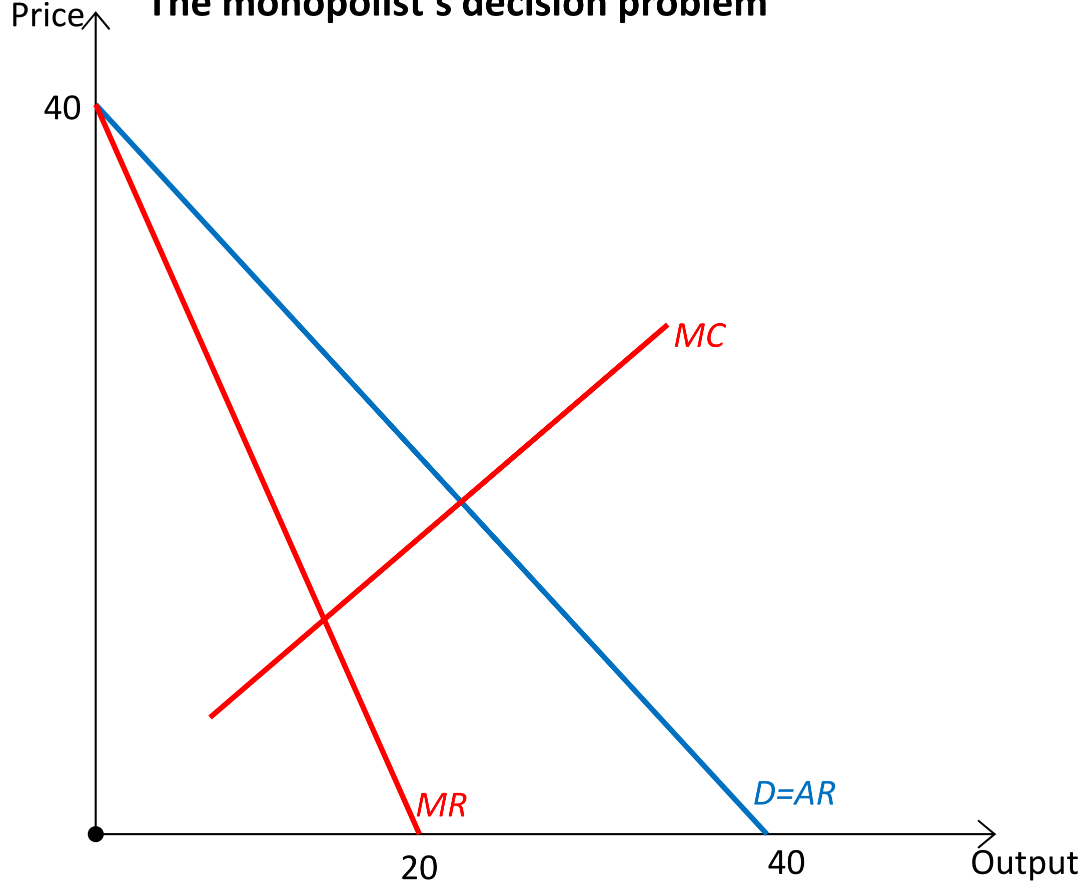

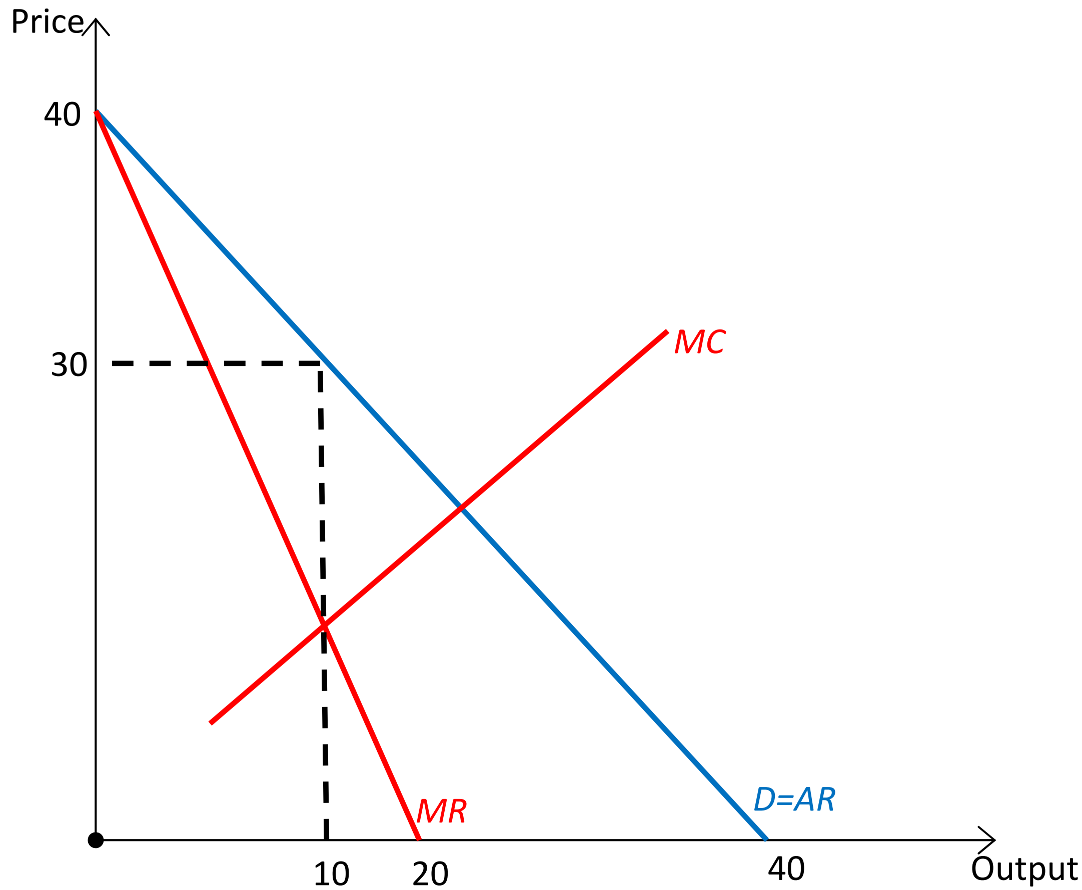
Price
40
30
MC
MR
D=AR
10
20
40
Output

The monopolist's profit

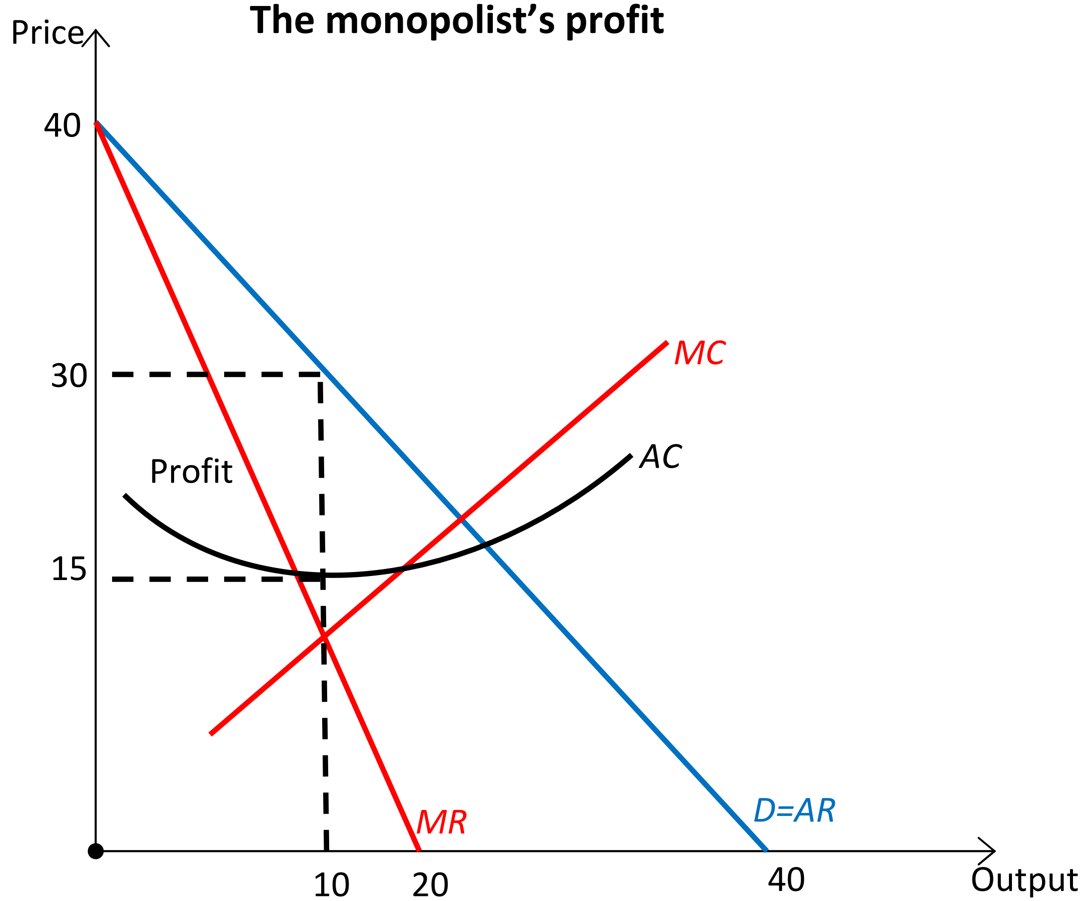

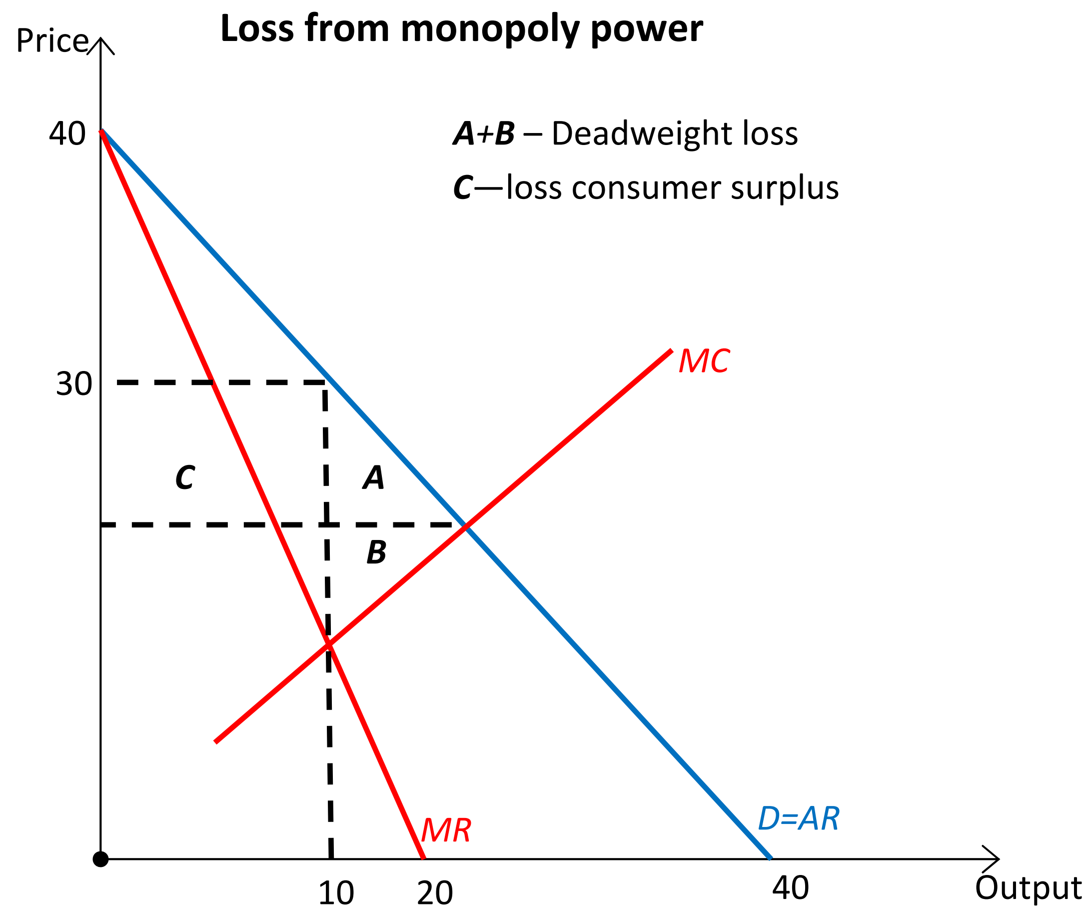
Loss from monopoly power
Price
40
30
A+B – Deadweight loss
C—loss consumer surplus
MC
C
A
B
MR
D=AR
10
20
40
Output

The "rule of thumb" for pricing

But a lot is wrong with the story just told – managers have only limited information of the average and marginal revenue curves facing their firms. To this end, so we need a rule of thumb that can be applied in the real-world.

Note that selling an extra unit *must* result in a small drop in price $\Delta P/\Delta Q$ which reduces the revenue from *all* units sold! We therefore rewrite the marginal revenue as follows

$$\begin{aligned} MR &= P + Q\frac{\Delta P}{\Delta Q} = P + P\left(\frac{Q}{p}\right)\left(\frac{\Delta P}{\Delta Q}\right) \\ &= P + P\frac{1}{E_d}. \end{aligned}$$

Recall that the price elasticity of demand – the percentage change (decrease) in quantity demanded of a good resulting from a 1-percent increase in its price – is given by

$$E_d = \frac{\Delta Q/Q}{\Delta P/P} = \frac{P}{Q}\frac{\Delta Q}{\Delta P}.$$

When we set marginal revenues to marginal costs we get

$$MR = P + P\frac{1}{E_d} = MC.$$

Rearranging,

$$P = \frac{MC}{1 + (1/E_d)}.$$

Monopoly power?

For a competitive firm, price equals marginal costs; for a firm with monopoly power, price exceeds marginal costs.

The Lerner Index of Monopoly Power (1934) given mathematically by

$$\begin{aligned} L &= \frac{P - MC}{P} \\ &= -1/E_d \end{aligned}$$

uses the markup ratio of price minus marginal costs to price to measure the monopoly power.

! Firms prices are sometimes below its optimal price so its monopoly power will not be noted by the Lerner Index.

Sources of market power

The more inelastic its demand curve, the more monopoly power the firm has. These factors determine a firm's demand elasticity:

[1] The elasticity of market demand.

[2] The number of firms in the market.

[3] The interaction among firms.

Maintaining monopoly

$\Longrightarrow$ Differentiated / branded goods.

$\Longrightarrow$ Barriers to entry (e.g., patents).

$\Longrightarrow$ Customer lock-in.

$\Longrightarrow$ Predatory pricing.

Summary

- In a competitive market there are many firms selling an identical product.
  - When one of these firms raises its price above the market price it loses all its customers.

- In a monopolized market, there is only one firm selling a given product.
  - When a monopolist raises its price it loses some, but not all, its customers.

! In reality, most industries are somewhere in between these two extremes.

- If a firm has some degree of monopoly power then it has more strategies than a firm in a perfectly competitive market.

- The problem faced by firms with some monopoly power is how to enhance and exploit their market power most effectively.

- Their objective – capturing more consumer surplus and converting it into additional profits for the firm.

- This goal can be can be achieved using *price discrimination*, that is charging different prices for different consumers.

The deficiencies of simple pricing

[1] If the firm can charge only one price for all it consumers, to maximize profit, it would pick the price P^* and corresponding output Q^* where its marginal cost (MC) and marginal revenue (MR) curves intersect.

[2] But although *simple pricing* or *uniform pricing* is quite prevalent, it is neither the only nor the most desirable form of pricing. If the firm could sell different units of output at different prices, then we have another story.

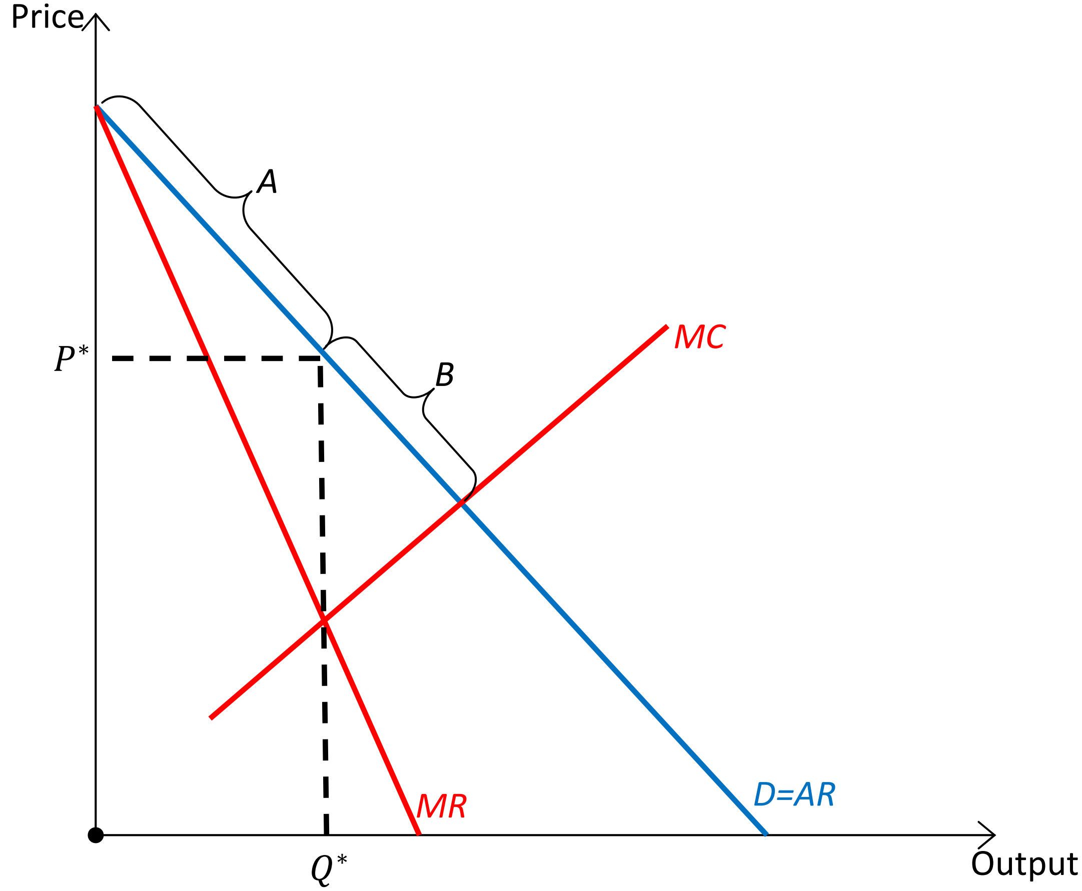

Price
A
B
MC
P^*
MR
D=AR
Q^*
Output

What simple pricing loses?!

1. It leaves potential profit in the hands of consumers in the form of their consumer surplus (the surplus under region A of the demand curve).

2. It leaves "money on the table" (deadweight loss) due to the driving-down-the-price effect (the surplus under region B of the demand curve).

Two observations:

1. Under simple pricing, the firm cannot charge a price for the $Q^* + 1$ unit that is different than the price it charges for the other Q^* units. And the fact that it would, have to lower the price on all Q^* units in order to induce someone to buy the $Q^* + 1$ unit, makes selling that unit unprofitable.

2. The firm would like to charge a higher price to consumers willing to pay more than P^*. The firm would also like to sell to consumers willing to pay prices lower than P^*, but only if doing so does not entail lowering the price for the other consumers.

! This is the basis for price discrimination – charging different prices to different consumers.

Price discrimination

Economists typically consider three degrees of prices discrimination:

1st The firm sells different units of output for different prices and these prices may differ from consumer to consumer.

2nd The firm sells different units of output for different prices but every consumer who buys the same amount of the good pays the same price.

3rd The firm sells output for different consumers at different prices, but every unit of output sold to a given consumer sells for the same price.

Next, we look at each of these types of prices discrimination to see what economics can say about how prices discrimination works!

First-degree (or perfect) price discrimination

$\Longrightarrow$ A firm that is able to perfectly price discriminate will sell each unit of output at the highest price it will command, that is, at each consumer's reservation price.

$\Longrightarrow$ Since the firm can capture all the welfare generated from selling any number of units, it will want to produce to the point at which demand intersects marginal cost Q^{**}.

First-degree price discrimination

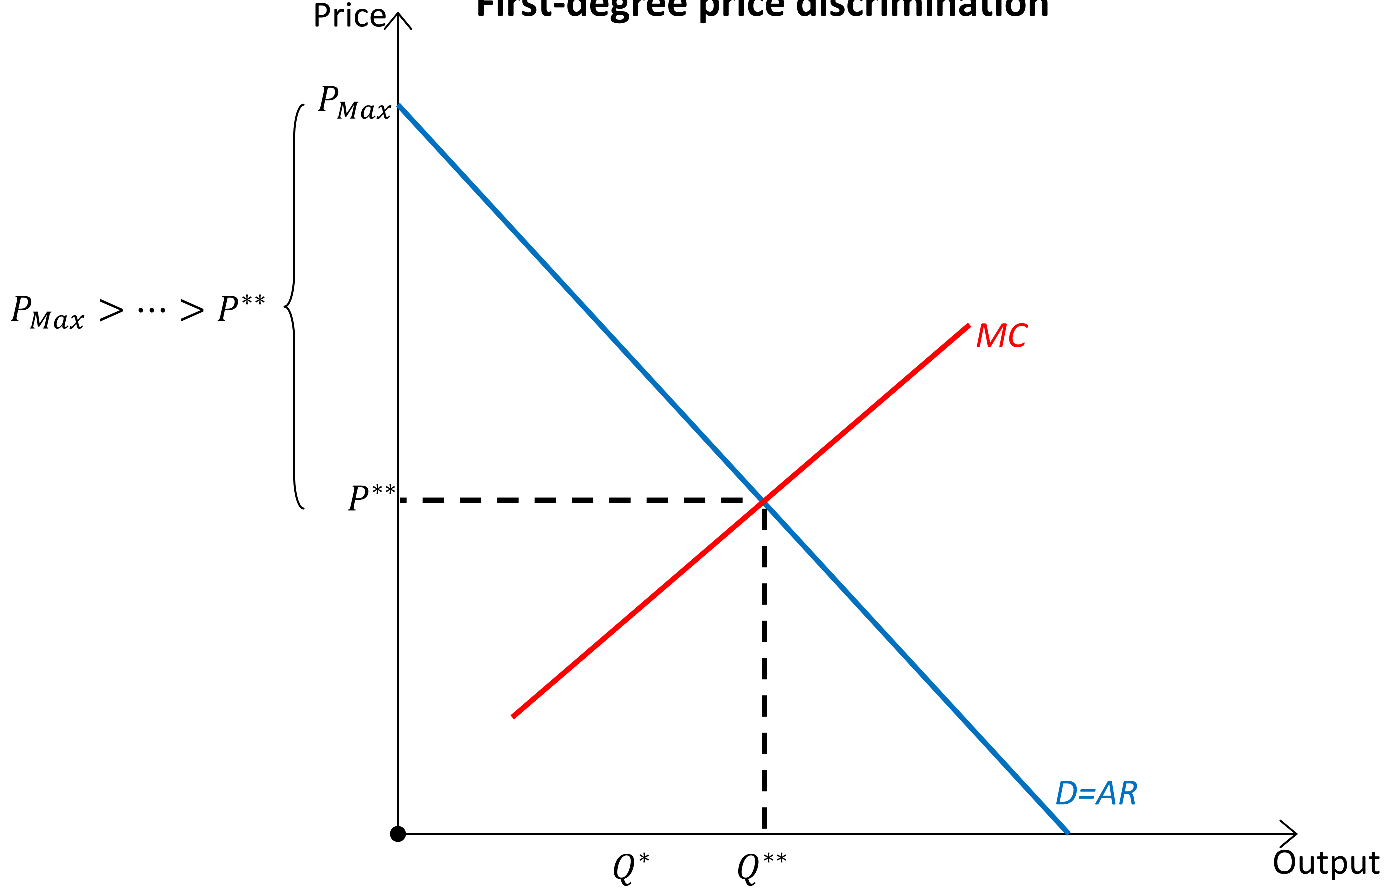

- Producing Q^{**} units and capturing all of welfare is the very best the firm could ever do (the "Holy Grail" of pricing).

- Perfect price discrimination is an idealized concept but it is interesting theoretical because it achieves Pareto efficiency.

- Clearly, firms typically cannot know the reservation price of every consumer, but sometimes reservation prices can be roughly identified.

- There are very few real-life examples. Perhaps to closest example would be tuition rates in Ivy League colleges, based on ability to pay.

Second-degree price discrimination

- Under second-degree price discrimination the price per unit of output is not constant but depends on how much you buy (non-linear pricing).

- One form of second-degree price discrimination is via *quantity discounts* – liter bottle of Coke is less than twice as expensive as the half-liter bottle.

- Typically, each price-quantity package is targeted toward different consumers, giving consumers an incentive to self-select.

- In practice, firms often encourage this self-selection not by adjusting *quantity* of the good, but rather by adjusting the *quality* of the good.

The airline industry has been very successful in prices discrimination (the industry spokesperson prefer to use the term "yield management"):

[1] restricted and unrestricted fares
[2] first-class and couch-class travel
[3] Saturday night stayovers
[4] advance-ticketing

! In the 19$^{\text{th}}$ century, the French railroads removed the roofs from second-class carriages to create third-class carriages...

Second-degree price discrimination (two different price-quantity packages)

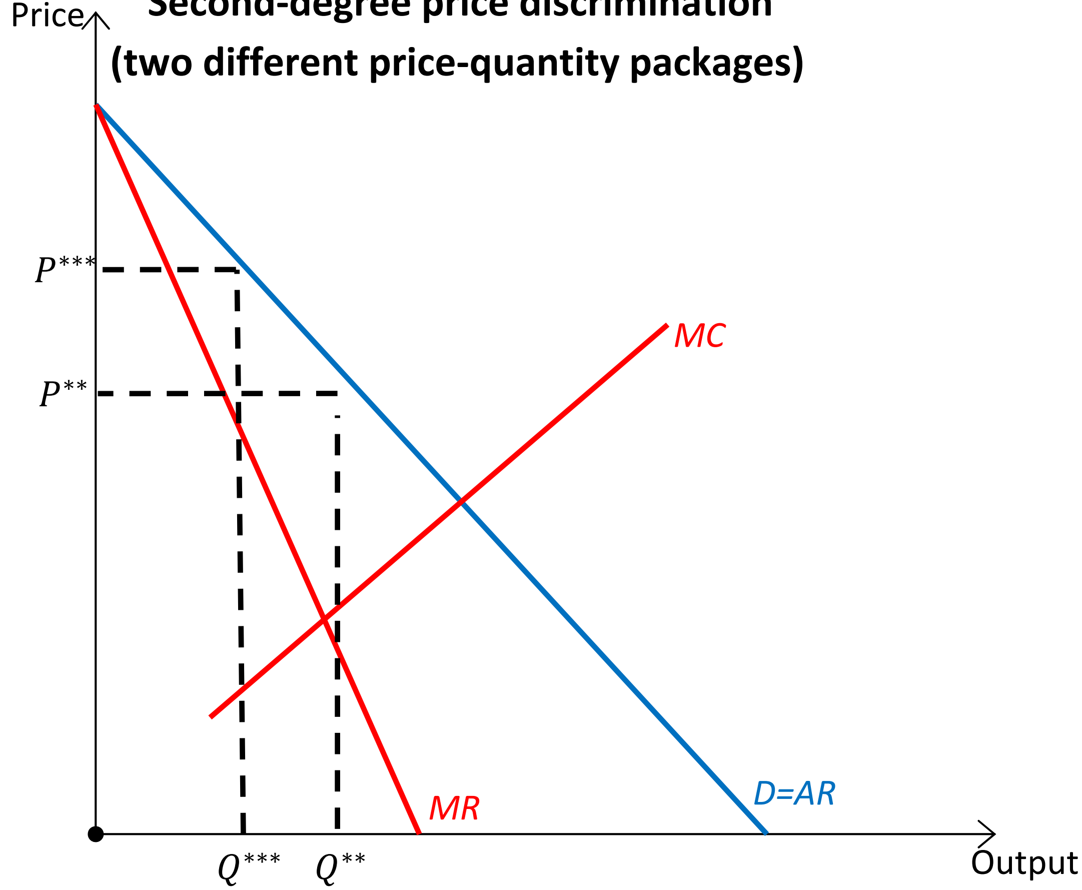

Third-degree price discrimination

- Third-degree price discrimination means charging different prices to different consumers on the basis of *identifiable* characteristics.

- An example of third-degree discrimination is charging different prices on the basis of observed group membership (children, seniors, students, etc).

- Geography-based third-degree price discrimination is quite prevalent – in air travel, a round-trip SFO-JFK has a different price than JFK-SFO...

- Clearly, the market with the higher price must have a lower elasticity of demand.

An example

Let $P_1(Q_1)$ be the demand of students and $P_2(Q_2)$ be the demand of non-students where Q_1 and Q_2 are the amount sold to students and non-students, respectively. Let $C(Q_T)$ for $Q_T = Q_1 + Q_2$ be the firm's cost function.

The firm's profit is its revenue from each population less its costs is given by

$$\pi(Q_T) = P_1(Q_1)Q_1 + P_2(Q_2)Q_2 - C(Q_T).$$

The profit-maximizing quantity of output Q_1^* (resp. Q_2^*) is such that the marginal (incremental) profit resulting from a small increase in Q_1 (resp. Q_2) equals zero, that is,

$$\frac{\Delta\pi}{\Delta Q_1} = \frac{\Delta P_1 Q_1}{\Delta Q_1} - \frac{\Delta C}{\Delta Q_1} = 0,$$

and

$$\frac{\Delta\pi}{\Delta Q_2} = \frac{\Delta P_2 Q_2}{\Delta Q_2} - \frac{\Delta C}{\Delta Q_2} = 0,$$

Thus, at the profit-maximizing quantities Q_1^* and Q_2^*

$$MR_1(Q_1^*) = MR_2(Q_2^*) = MC(Q_1^* + Q_2^*).$$

To determine *relative prices*, we can write the marginal revenues in terms of elasticity of demand:

$$MR_1 = P_1(1 + 1/E_1)$$

and

$$MR_2 = P_2(1 + 1/E_2).$$

By equating MR_1 and MR_2, we must have

$$\frac{P_1}{p_2} = \frac{1 + 1/E_2}{1 + 1/E_1}.$$

! This analysis generalizes to any (finite) number of populations.

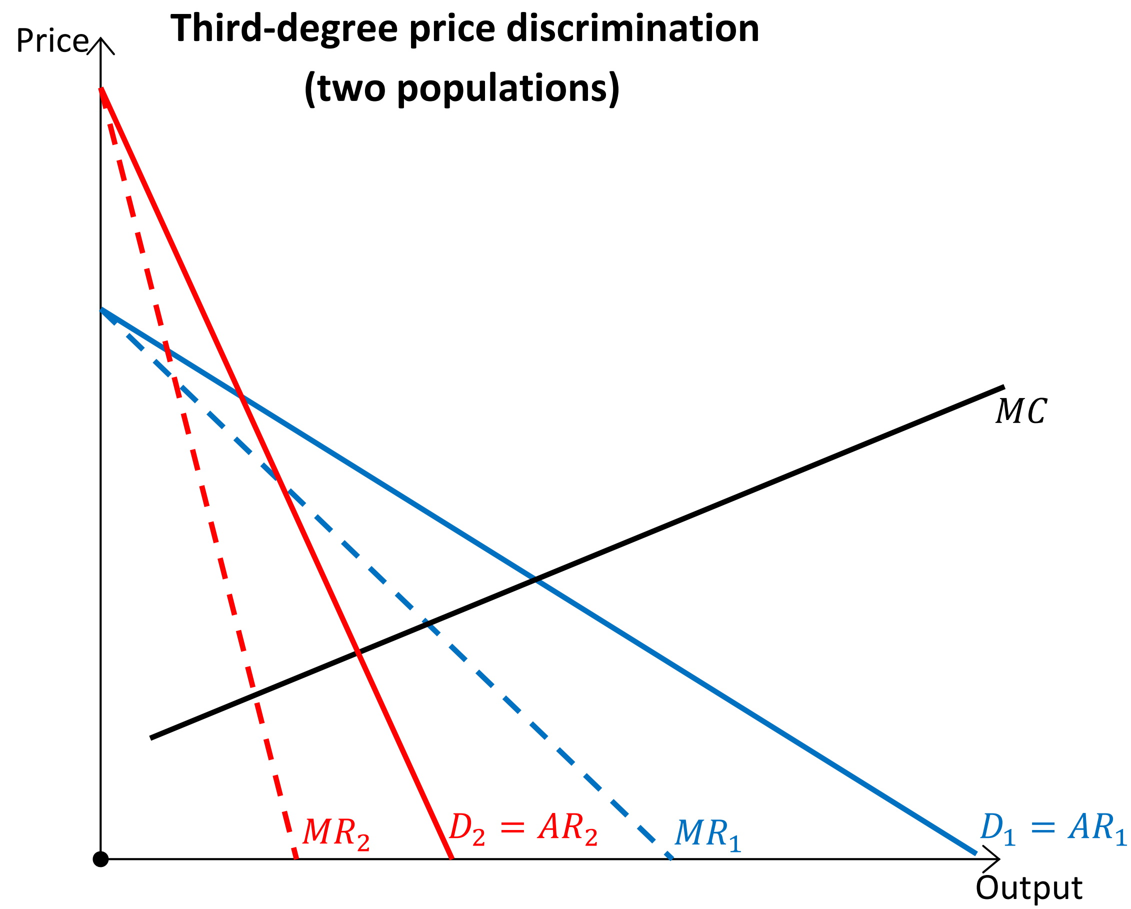
Third-degree price discrimination
(two populations)
Price
MC
MR_2
$D_2 = AR_2$
MR_1
$D_1 = AR_1$
Output

Third-degree price discrimination

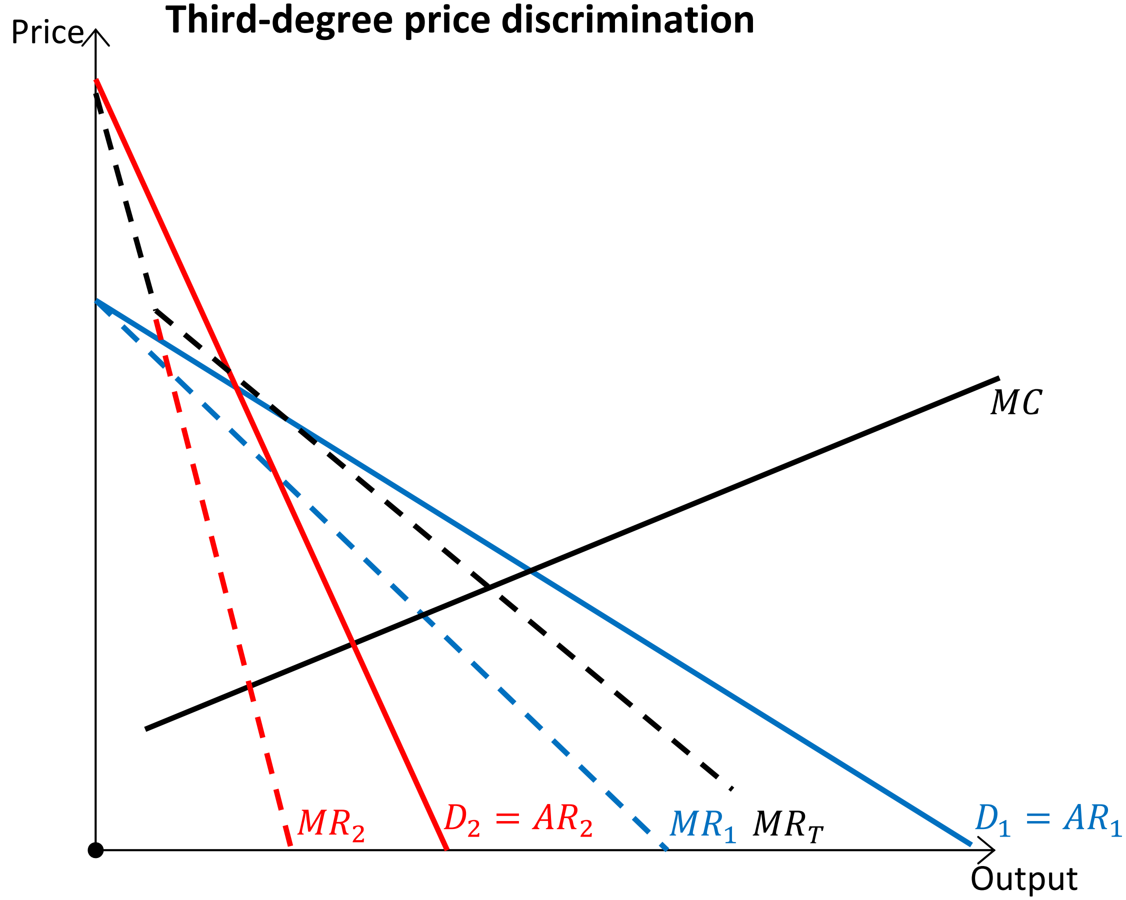

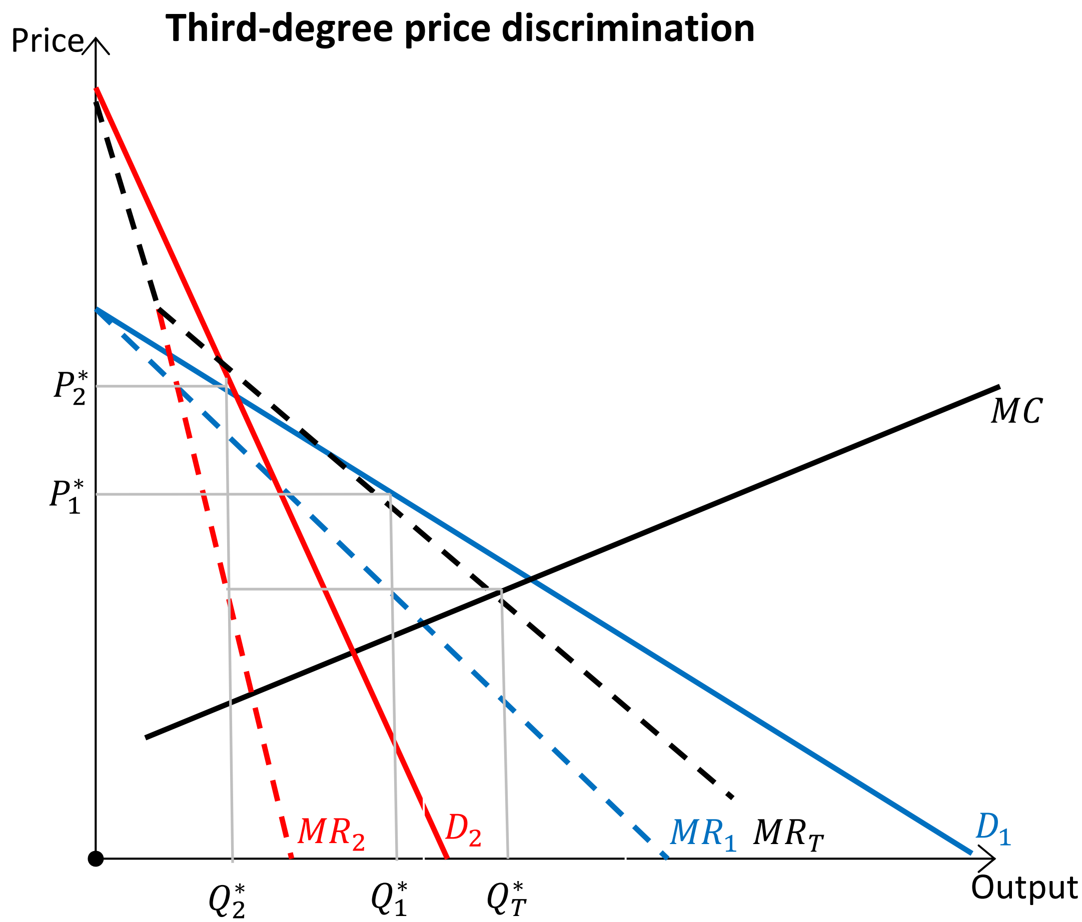

Third-degree price discrimination
Price
P_2^*
P_1^*
MC
MR_2
D_2
MR_1
MR_T
D_1
Output
Q_2^*
Q_1^*
Q_T^*

Two-part tariffs
(The Disneyland Dilemma)

A two-part tariff can help get a firm closer to the Grail than can simple pricing. It is a pricing scheme (tariff) with, as the name indicates, two parts:

I An entry fee – the amount that the consumer must pay before she can buy any units at all (overhead charge).

II A per-unit charge – the amount that the consumer must pay for each unit she chooses to purchase.

! In some instances, as with some – but not all – amusement parks, the per-unit charge might even be set to zero (so-called Disneyland pricing).

$\Longrightarrow$ **Identical consumers**

Suppose consumers have identical demands (are homogeneous). Under the profit-maximizing two-part tariff, the firm

– produces Q^{**} units, where $P(Q^{**}) = MC(Q^{**})$

– sets the per-unit charge equal $P(Q^{**})$

– sets an entry fee F to equal CS/N where N is the number of consumers.

Two-part tariff

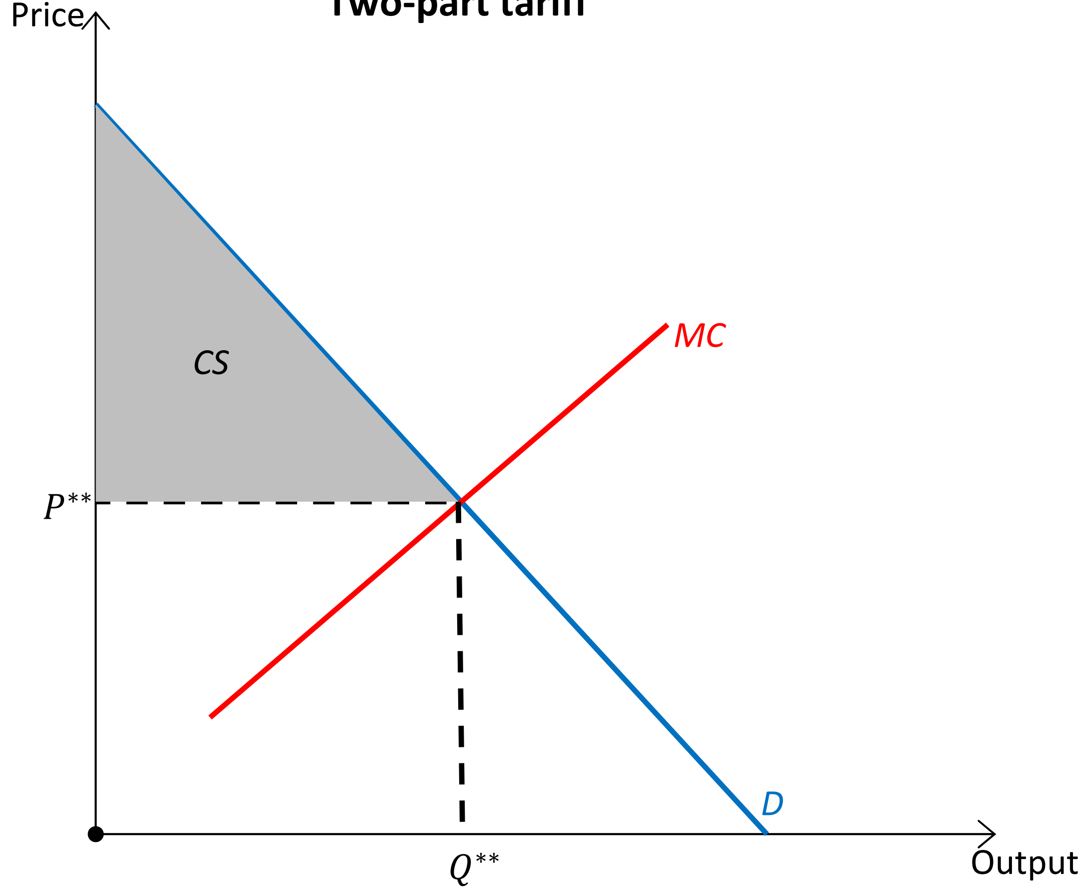

$\Longrightarrow$ **Consumers are heterogeneous**

If consumers do not all have the same demand curves then designing the optimal two-part tariff becomes much more complicated.

One strategy would be to divide customers into homogeneous subgroups (third-degree price discrimination), and then employ the optimal two-part tariff on each group.

! Real-life use of two-part tariffs are hard to recognize initially – club stores, tech support (some has service-call charge), land-line phones, and more.

Takeaways

1. If all consumers are identical, then perfect price discrimination can be achieved by a two-part tariff. When consumers are not identical, then it is typically not possible to achieve perfect discrimination.

2. When consumers are heterogeneous it is sometimes possible to divide them into different populations that are more homogeneous. If so, then the firm can engage in third-degree price discrimination.

3. When the firm cannot freely identify consumers' types, it can try to induce them to reveal their types. This form of price discrimination is known as second-degree price discrimination.

4. Two prevalent forms of second-degree price discrimination are using quality distortions and quantity discounts.

5. Another method of price discrimination is *bundling*. Bundling allows the firm to take advantage of the correlations that exist between consumers' preferences for different products.

6. All discriminatory pricing is, in theory, vulnerable to *arbitrage* – the advantaged reselling to the disadvantaged. Firms seek to deter or reduce arbitrage.