8. Basic Competition Economics

For Class 11--More to come for Class 12

Merger Antitrust Law

Fall 2017 Georgetown University Law Center

Dale Collins

September 28, 2018

Topics

- Introduction
- Consumer demand and the aggregate consumer demand curve
- Producer profit maximization and the aggregate supply curve
- Perfect market equilibria
  - Perfectly competitive markets
  - Perfect monopoly
  - Incentives for coordination
  - Public policy re monopoly
- Imperfect market equilibria
  - Cournot equilibria
  - Bertrand equilibria
  - Dominant firm with a competitive fringe
- Merger typology, substitutes and complements, and elasticities

Introduction

Basic competition economics

- Two asides
  - A plea

> Do not be put off by the mathematical notation in the slides that follow. All of the notation can be ignored without losing any substance. However, economics is an essential language in modern antitrust law in general and merger antitrust law in particular. As you will see, economists love to use mathematical notation to make things look complicated, but with a small investment of effort you will see that all of this is very simple. Learning the basic economics is an investment that will give you a significant comparative advantage over many other antitrust attorneys.

  - An observation by Dave Berry

> Later on, Newton also invented calculus, which is defined as "the branch of mathematics that is so scary it causes everybody to stop studying mathematics." That's the whole point of calculus. At colleges and universities, on the first day of calculus, professors go to the board and write huge, incomprehensible "equations" that they make up right on the spot, knowing that this will cause all the students to drop the course and never return to the mathematics building. This frees the professors to spend the rest of the semester playing cards and regaling one another with stories about the "mathematical symbols" they've invented over the years. ("Remember the time Professor Hinkwattle drew a 'cosine derivative' that was actually a picture of a squid?" "Yes! Students were diving out the windows! From the fourth floor!")[1]

[1] Dave Berry, *Up in the Air on the Question of Gravity,* Baltimore Sun, Mar. 16, 1997, at 3J.

Price formation models

- Standard assumptions in the neo-classical model
  - Consumers
    - Individually maximize preferences (utility) subject to their individual budget constraints
    - Yields a consumer demand function, which gives the quantity demanded $q_i^{demanded}$ by consumer *i* for a given market price *p*
  - Firms
    - Individually maximize profits subject to their available production technology (production possibility sets)
    - Yields a production function that gives the quantity produced $q_j^{produced}$ by firm *j* for a given market price *p*
  - Equilibrium condition
    - No price discrimination (all purchases are made at the single market price)
    - Market clears at the market price (i.e., demand equals supply):

$$\sum_i q_i^{demanded} = \sum_j q_j^{produced}$$

> $\sum$ simply means to add up the *q*'s. So if $q_1 = 10$. $q_2 = 7$, and $q_3 = 5$, then $\sum q_i = 10 + 7 + 5 = 22$.

Consumers

Consumers

- *Assumption*: Consumers maximize their preference (utility) subject to their individual budget constraints
  - An individual consumer's demand for a product is a function of:
    - The consumer's preferences
    - The price the consumer pays for product
    - Other products and services the consumer may purchase and their respective prices
    - The consumer's budget constraint
  - The relationship between quantity and price is known as the *consumer demand function* or *consumer demand curve*
    - Typically, the consumer will purchase a larger quantity of the product as the price decreases
    - If so, then the consumer demand curve is *downward sloping*
  - The sum of consumer demand functions is known as the *industry (aggregate) demand function*

Almost all antitrust economic analysis takes this as the point of departure. It is a critical assumption. This is often called the *law of demand*.

Demand curves

- Demand curves and inverse demand curves
  - A *demand curve* gives quantity as a function of price
    - The demand curve gives the quantity the consumer will purchase for any given price
    - *Example:* Given my budget constraint, if the price is $4.00, I will buy 12 units, but if the price is $5.00 I will buy only 10 units
  - Linear demand curves
    - Linear demand curves are straight lines
      - Although demand curves need not be straight lines, all of the principles in which we will be interested may be illustrated using linear demand curves
    - A *linear demand curve* has the form $q = a + bp$, where q is the quantity demanded at price p, a is the quantity when $p = 0$, and b gives the change in q for a change in price
      - q and p are called *variables* and are the numbers of interest to us
        - They are related in pairs (p_i, q_i) by the demand curve, that is, each (p_i, q_i) lies on the demand curve so that $q_i = a + bp_i$ for each observation i
      - a and b are constants called *parameters*
        - The parameter a is the quantity demanded when the price is equal to zero
        - The parameter b is the *slope* of the demand curve: it gives the decrease in the quantity demanded for an increase of one unit in price
        - Since demand curves are downward sloping, b will be a negative number (i.e., $b < 0$)

Demand curves

- Demand curves and inverse demand curves
  - Graphing demand curves
    - *Example:* Given my budget constraint, if the price is \$4.00 I will buy 12 units, but if the price is \$5.00 I will buy only 10 units
    - A linear demand curve has the form $q = a + bp$, where q is the quantity demanded at price p, a is the quantity when p =0, and b gives the change in q for a change in price
      - b is also the *slope* of the demand curve
    - If we know two points on a linear demand curve, we can derive the demand function

Notation: Δq means the change in q *and* is read "delta q"

$$b = \frac{\text{Change in quantity}}{\text{Change in price}} \equiv \frac{\Delta q}{\Delta p} = \frac{-2}{1} = -2$$

The change Δq is negative because demand declines as price increases

The symbol "≡" means a definition

Use one point to solve for a (say, p =4; q = 12):

$$12 = a - 2 \times 4 \Rightarrow a = 20$$

So the demand curve is:

$$q = 20 - 2p$$

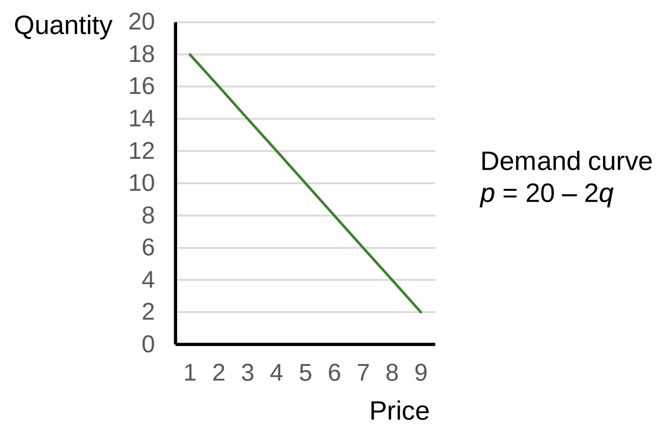

Demand curves

- A more detailed diagram

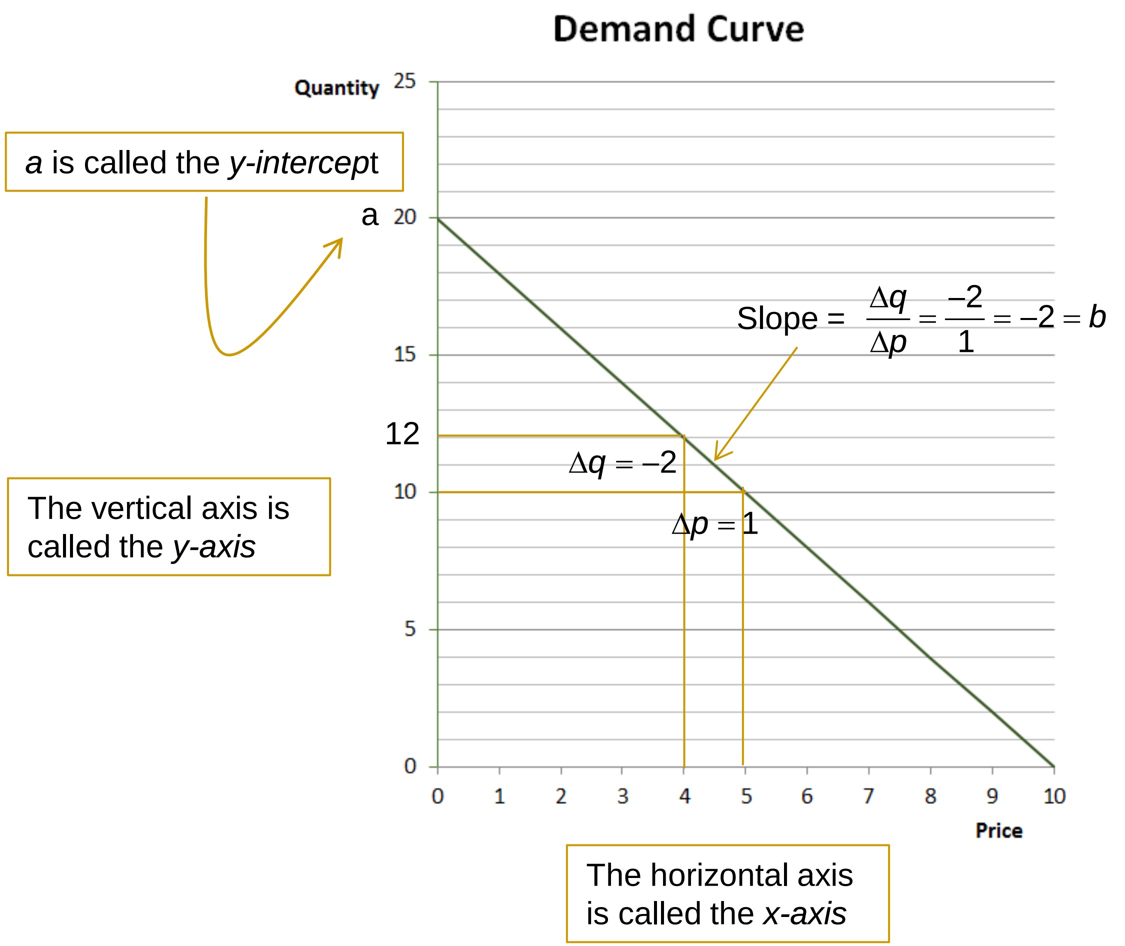

Demand curve: $q = 20 - 2p$

General form: $q = a + bp$
So
$a = 20$
$b = -2$

Demand curves

- Demand curves and inverse demand curves
  - An *inverse demand curve* gives price as a function of quantity
  - So if the demand curve is $q = a + bp$, the inverse demand curve can be derived by solving for p:
    - *Example*: If the demand curve is $q = 20 - 2p$, the inverse demand curve is:

$$p = \frac{20 - 2q}{2} = 10 - \frac{1}{2}q$$

  - Think about the inverse demand curve as the price necessary to *clear the market* given production level q
    - "Clear the market" means that consumers demand no more and no less than q at price p
  - Inverse demand functions put price on the *y*-axis and quantity on the *x*-axis of a graph
    - Just the opposite of the demand curve

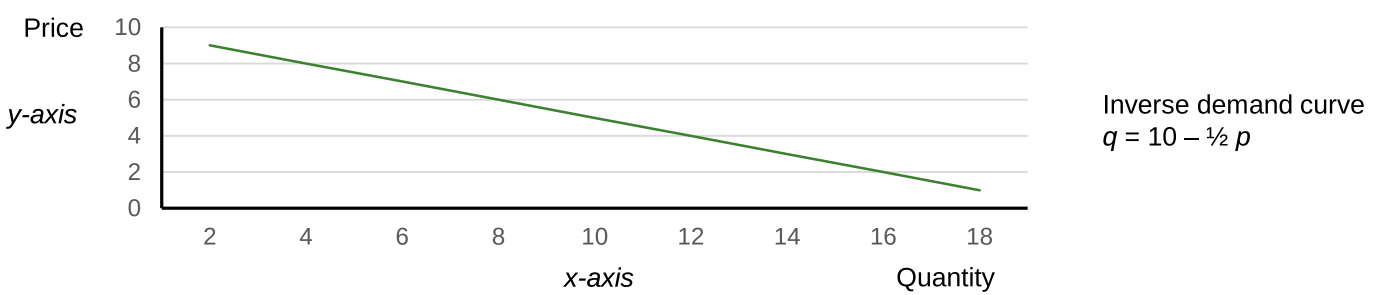

Demand curves

- Linear (inverse) demand curve

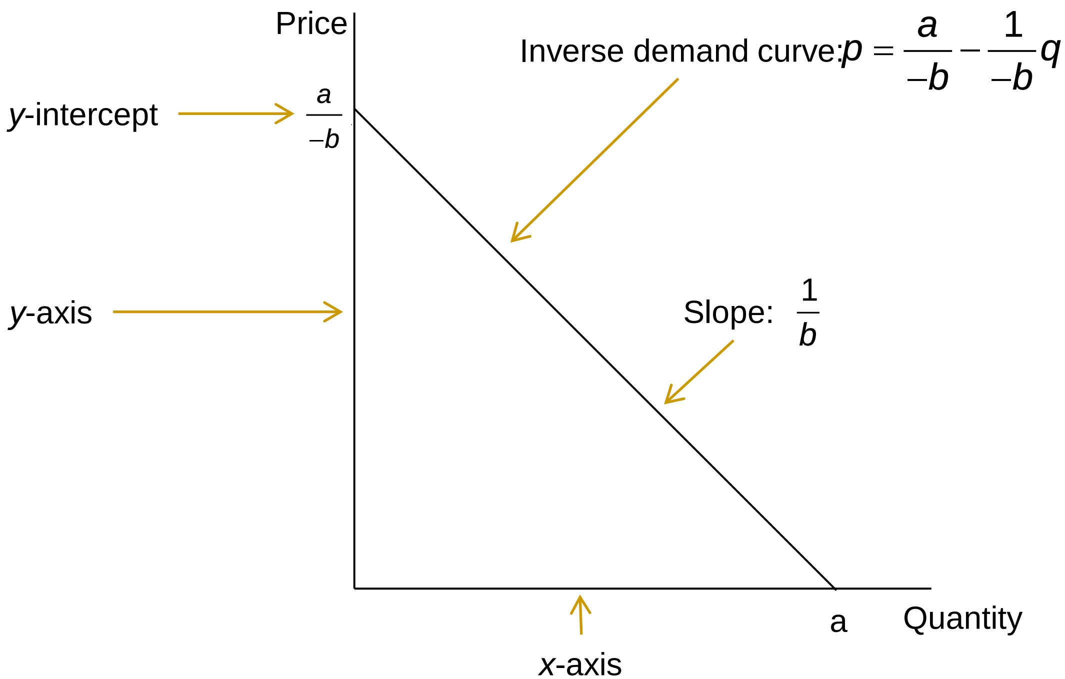

The slope of the demand curve gives the required change in the price to sell one additional unit of the product. So the price needs to drop by -1/*b* to sell one additional unit:.

$$p_1 = \frac{a}{-b} - \frac{1}{-b}(q+1)$$

$$p_0 = \frac{a}{-b} - \frac{1}{-b}(q)$$

$$\text{So } \Delta p = p_1 - p_0 = \frac{1}{b}$$

Notes: The y-intercept *a*/*b* is the price above which there is zero demand.

The *x*-intercept *a* is the quantity demanded when the price is zero.

For linear demand, unless the demand curve is strictly vertical, the *x*- and *y*-intercepts will be finite. Because they can be very large, this usually does not result in any loss of generality.

Demand curves

- Example: Nonlinear inverse demand curve with no *x*-axis intercept

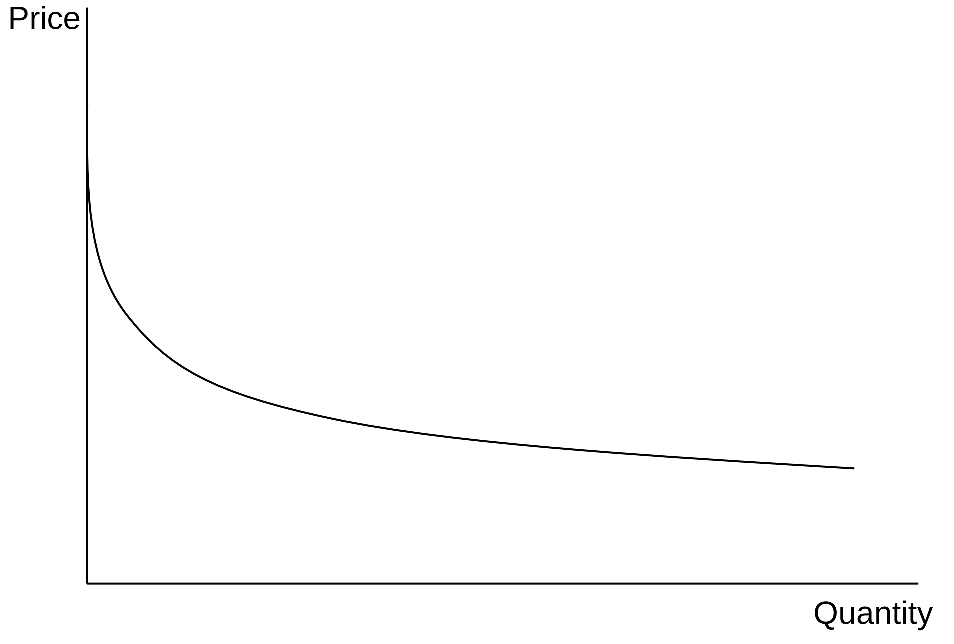

Demand curves

- Example: Nonlinear inverse demand curve with no x-axis intercept

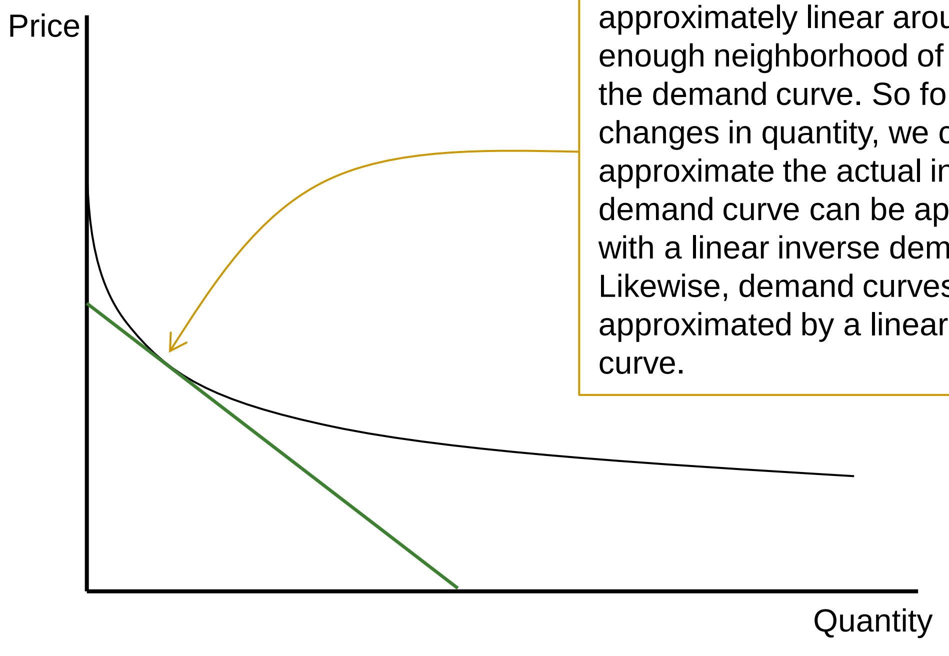

Note: Demand curves are approximately linear around any small enough neighborhood of any point on the demand curve. So for small changes in quantity, we can approximate the actual inverse demand curve can be approximated with a linear inverse demand curve. Likewise, demand curves can be approximated by a linear demand curve.

Demand curves

- Some technical points about demand curves and inverse demand curves
  - When economists and antitrust lawyers refer to the demand curve, they almost always mean the inverse demand curve
    - You can tell the difference in context by whether:
      - *Demand curve*: Quantity is a function of price and on the *y*-axis
      - *Inverse demand curve*: Price is a function of quantity and on the *y*-axis
    - We will follow convention and not draw the distinction
  - What I have called the “demand curve” is really the “demand function”
    - The demand curve is the *graph* of the demand function
    - Distinguishing between the two will qualify you as an irredeemable geek
  - Demand curves are for aggregate demand in the marketplace, that is, the sum of demands by consumers of all firms in the marketplace
    - The demand curve for a single firm in the market is called the firm’s *residual demand curve*
      - Formally, the residual demand faced by any firm is that part of the total demand which is not met by the other firms in the industry:
    - Residual demand is a critical concept in antitrust economics

Demand curves

- Some technical points about demand curves and inverse demand curves
  - Even if we assume linear demand, two observations of prices and quantities demanded (p_1,q_1) and (p_2,q_2) may either be—
    - On the same demand curve, or
    - On different demand curves (when demand has *shifted*)

Two points on the same demand curve

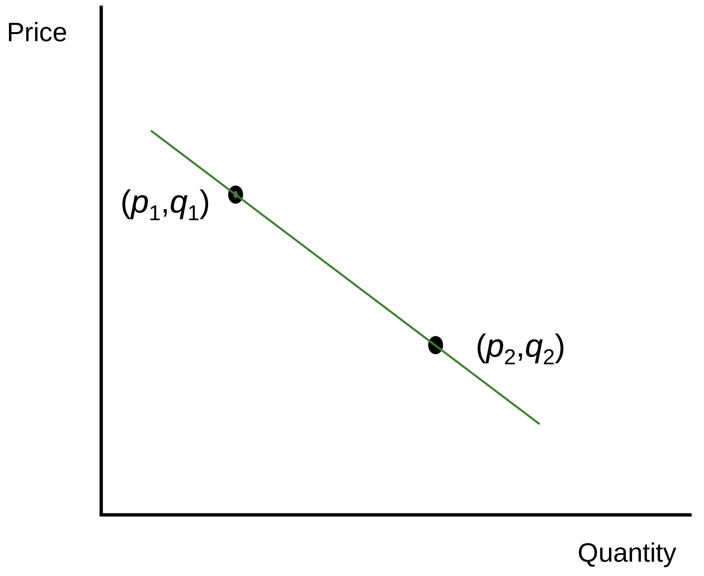

Two points on different demand curves

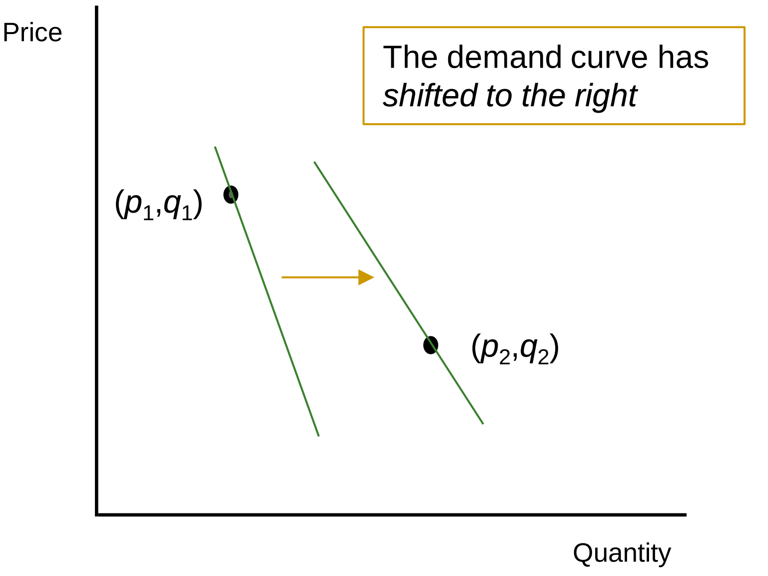

Aggregate consumer demand

- Aggregate consumer demand
  - Sum of individual consumer demands = Aggregate consumer demand (by definition)

$$\sum_i q_i^{demanded}(p) \equiv q(p),$$

where $q(p)$ is aggregate demand at price p

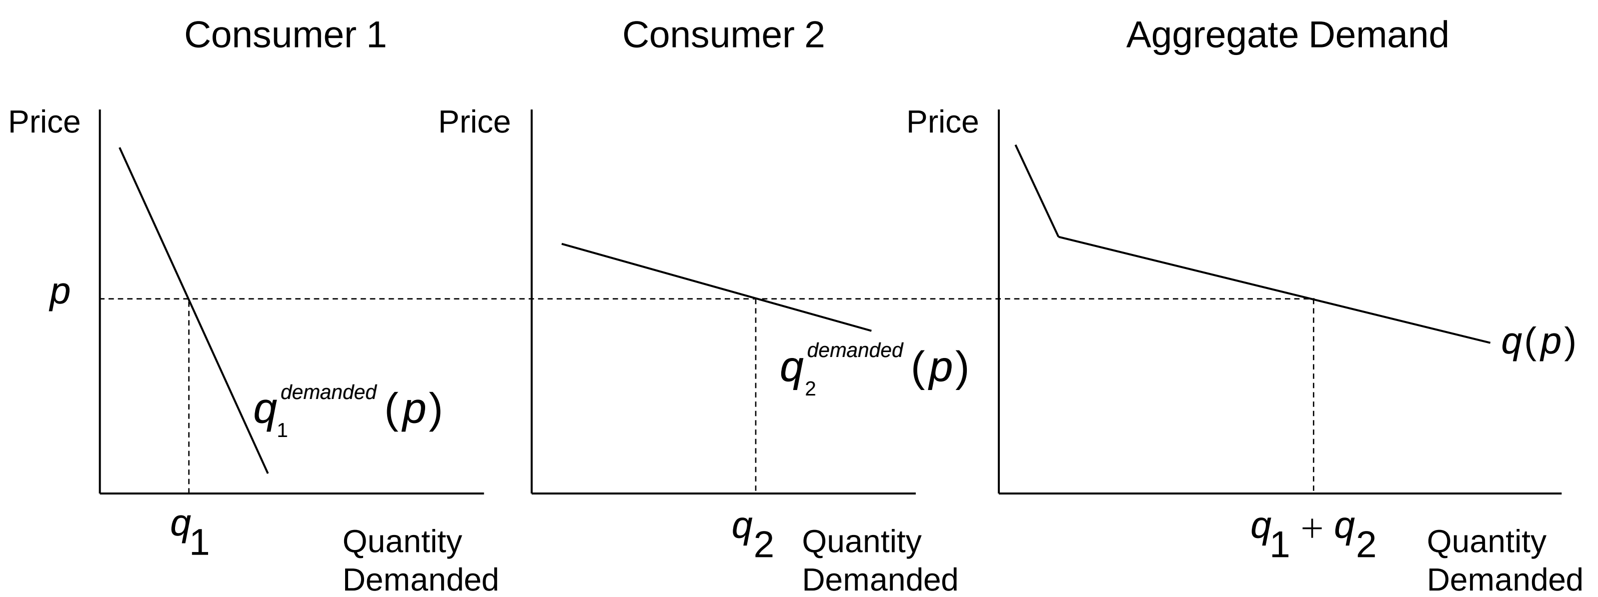

Producers

Producers

- *Assumption*: Firms maximize their profits subject to the technology available to them
  - Profits (π) = Revenues (r) – Costs (c)
- To analyze the conditions under which a firm maximizes it profit, need to look at:
  - Revenues and revenue functions
  - Costs and cost functions
  - The relationship between revenues and costs when the firm maximizes its profit

We will see that a firm maximizes its profit when it sets its marginal revenue equal to its marginal cost

Revenues

Revenue = *p* times *q* (= *pq*)
This is just the area of the rectangle in the chart below

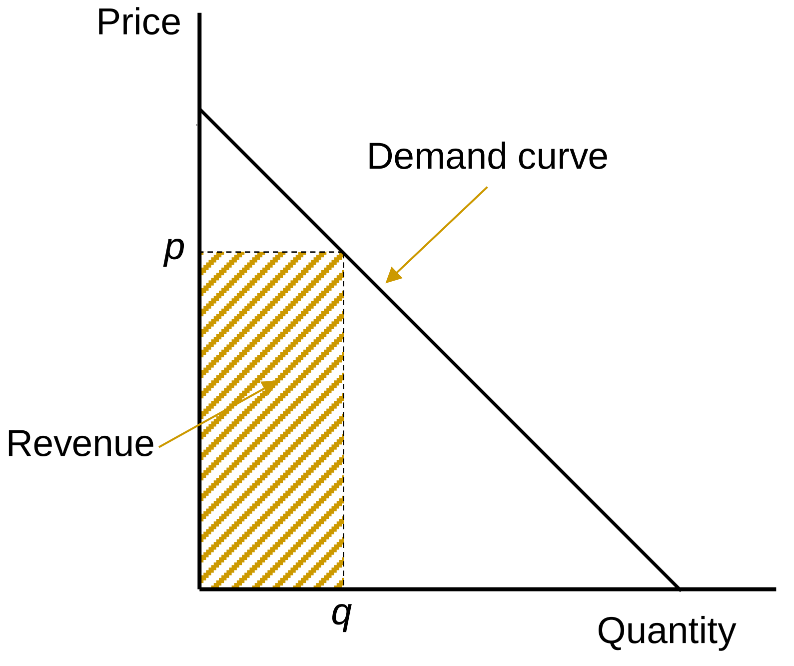

Revenues

Marginal revenue (mr) = Revenue gain from incremental sales (the sale of one additional unit) – revenue loss from lower price on preexisting sales

Notation: Δr means the change in *r and* is read "delta *r*"

$$\equiv \frac{\Delta r}{\Delta q} = \underbrace{(p + \Delta p)\Delta q}_{\text{Gain}} + \underbrace{\Delta p q}_{\text{Loss}}$$

NB: If Δq is positive, the Δp will be negative since demand curves are downward sloping. So Δpq is a negative number, reflecting a loss.

The next slide shows this graphically

Revenues

Marginal revenue (mr) = Revenue gain from incremental sales (the sale of one additional unit) – revenue loss from lower price on preexisting sales

Step 1: Look first at the *revenue gain* from the incremental sales (albeit at a lower price)

$$mr \equiv \frac{\Delta r}{\Delta q} = (p + \Delta p)\Delta q + \Delta pq$$

$\Delta p = p_2 - p$

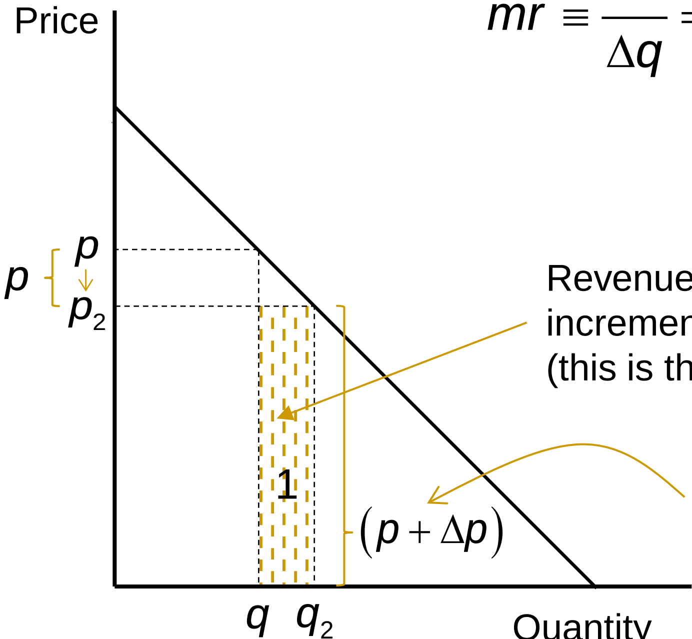

$\Delta q = q_2 - q$

Revenue gain from lower price on incremental sales $= (p + \Delta p)\Delta q$ (this is the area in Box 1)

(Negative) change in price necessary to make the incremental sales

Revenues from each incremental sale. (Remember, Δp is negative, meaning that price is decreasing)

Remember, *pq* are the original revenues from preexisting sales *q* at price *p*

Incremental sales

Revenues

Marginal revenue (mr) = Revenue gain from incremental sales (the sale of one additional unit)
– revenue from lower price on preexisting sales

Step 2: Now look at the *revenue loss* from the lower price on preexisting sales

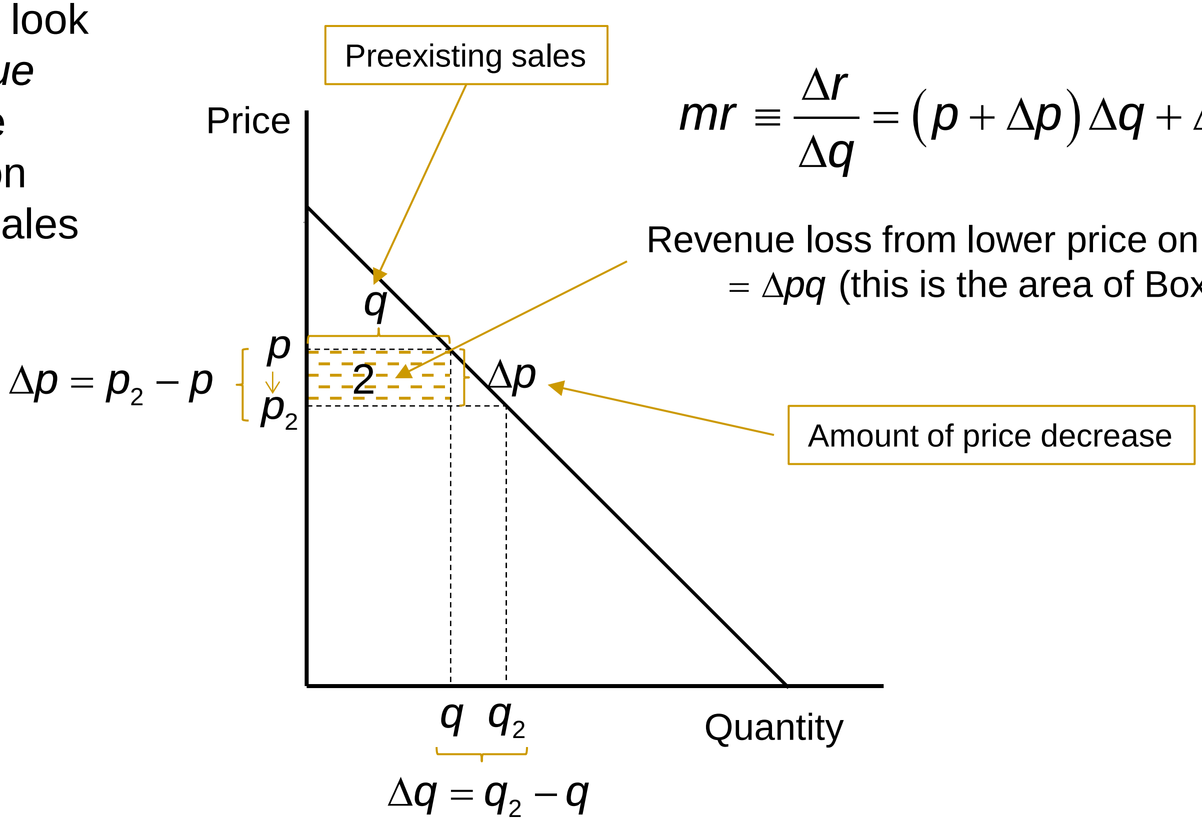

$$mr \equiv \frac{\Delta r}{\Delta q} = (p + \Delta p)\Delta q + \Delta p q$$

Revenue loss from lower price on preexisting sales $= \Delta p q$ (this is the area of Box 2)

Revenues

Marginal revenue (mr) = Revenue gain from incremental sales (the sale of one additional unit)
– revenue loss from lower price on preexisting sales

Step 3: Marginal revenue is the difference between Box 1 and Box 2

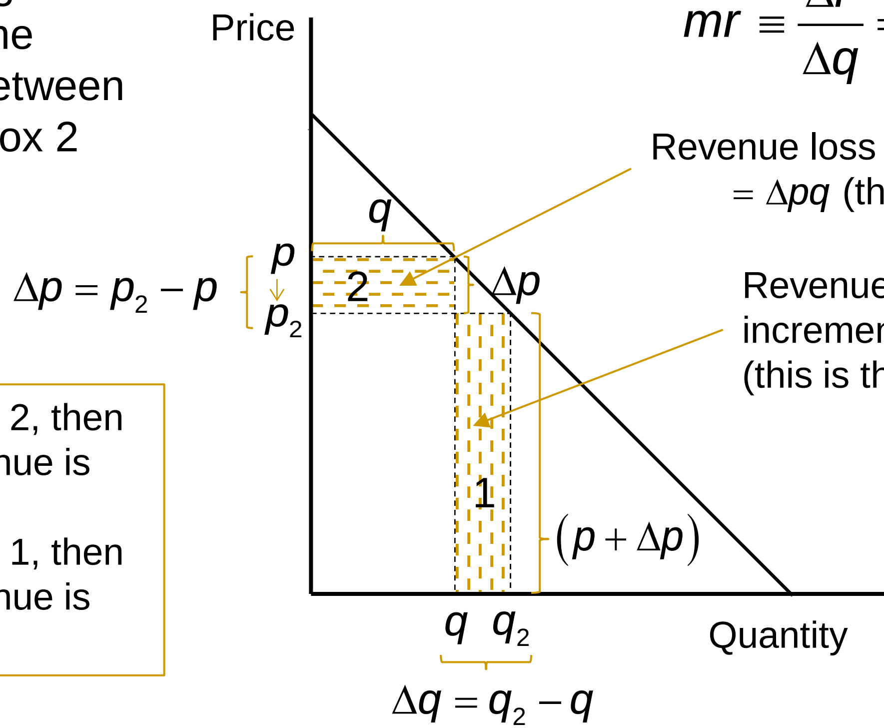

$$mr \equiv \frac{\Delta r}{\Delta q} = (p + \Delta p)\Delta q + \Delta pq$$

Revenue loss from lower price on preexisting sales $= \Delta pq$ (this is the area of Box 2)

Revenue gain from lower price on incremental sales $= (p + \Delta p)\Delta q$ (this is the area of Box 1)

If Box 1 > Box 2, then marginal revenue is positive
If Box 2 > Box 1, then marginal revenue is negative

Revenues

- Relationship between revenues and marginal revenue (discrete case)

Read this " *r* of *q*": This is the revenues *r* at production level *q*. The notation means that *r* is a function of *q*

$$r(q) = \sum_{i=1}^{q} mr_i$$

  - That is, total revenues for a production level *q* is equal to the sum of the marginal revenues for units 1 to *q*

Revenues

$$r(q) = \sum_{i=1}^{q} mr_i$$

- Numerical example
  - Inverse demand: $p = 10 - \frac{1}{2}q$

Quantity q	Price p	Revenue $r = pq$	Marginal Quantity Δq	Change Price Δp	Marginal Gain (p+Δp)Δq	Marginal Loss Δpq	Marginal Revenue (p+Δp)Δq+Δpq	Sum mr
0	10.0	0.0					10.0	
1	9.5	9.5	1	-0.5	10.0	-0.5	9.5	9.5
2	9.0	18.0	1	-0.5	9.5	-1.0	8.5	18.0
3	8.5	25.5	1	-0.5	9.0	-1.5	7.5	25.5
4	8.0	32.0	1	-0.5	8.5	-2.0	6.5	32.0
5	7.5	37.5	1	-0.5	8.0	-2.5	5.5	37.5
6	7.0	42.0	1	-0.5	7.5	-3.0	4.5	42.0
7	6.5	45.5	1	-0.5	7.0	-3.5	3.5	45.5
8	6.0	48.0	1	-0.5	6.5	-4.0	2.5	48.0
9	5.5	49.5	1	-0.5	6.0	-4.5	1.5	49.5
10	5.0	50.0	1	-0.5	5.5	-5.0	0.5	50.0
11	4.5	49.5	1	-0.5	5.0	-5.5	-0.5	49.5
12	4.0	48.0	1	-0.5	4.5	-6.0	-1.5	48.0
13	3.5	45.5	1	-0.5	4.0	-6.5	-2.5	45.5
14	3.0	42.0	1	-0.5	3.5	-7.0	-3.5	42.0
15	2.5	37.5	1	-0.5	3.0	-7.5	-4.5	37.5
16	2.0	32.0	1	-0.5	2.5	-8.0	-5.5	32.0
17	1.5	25.5	1	-0.5	2.0	-8.5	-6.5	25.5
18	1.0	18.0	1	-0.5	1.5	-9.0	-7.5	18.0
19	0.5	9.5	1	-0.5	1.0	-9.5	-8.5	9.5
20	0.0	0.0	1	-0.5	0.5	-10.0	-9.5	0.0

So that you can see that they are the same

Revenues

- Graphing revenue and marginal revenue curves

Demand: $p = 10 - \frac{1}{2}q$ Revenues: $r = pq$ Marginal revenue $mr = 10 - q$
(see the next slide as to why)

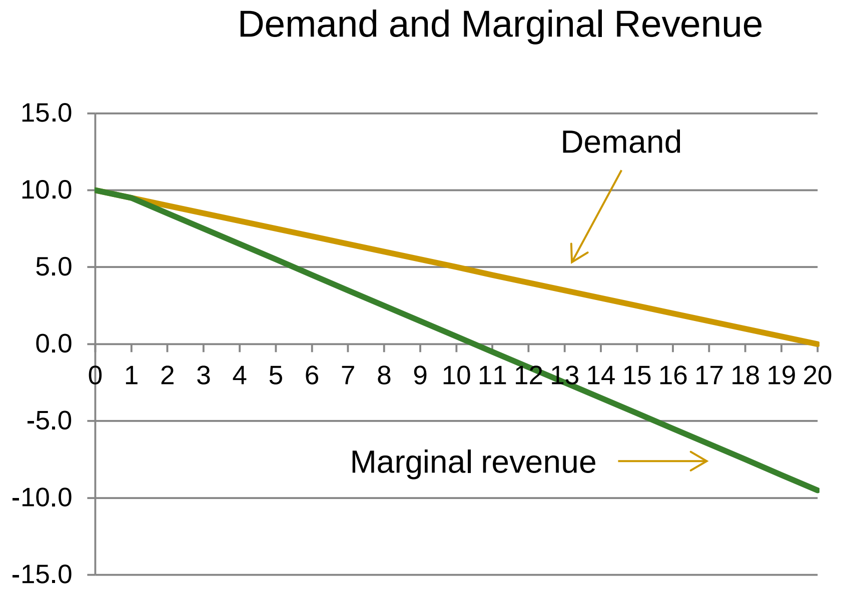

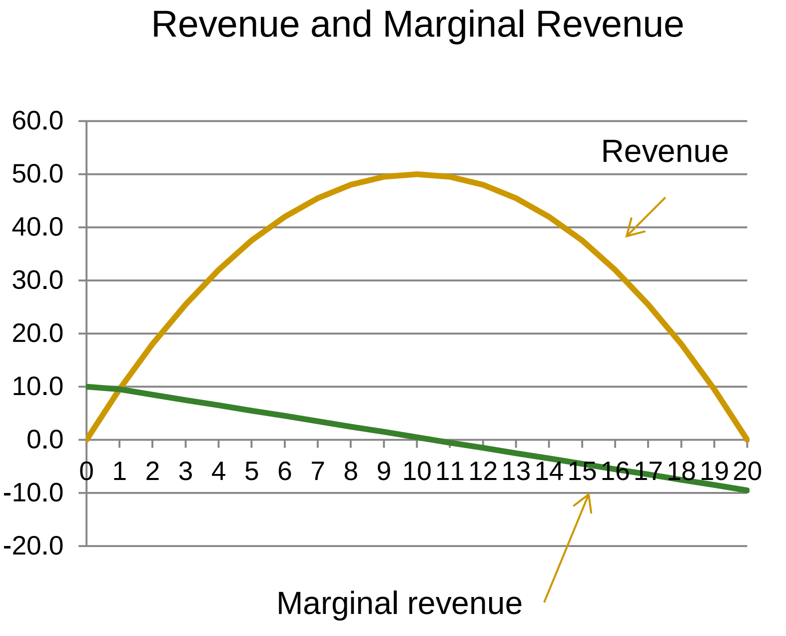

Notes:

1. When demand is linear, the slope of the marginal revenue curve is twice as steep downward as the demand curve.
2. When marginal revenue equals zero (here, a q = 10), revenues are at their maximum.

Revenues

- Marginal revenue with linear demand
  - A linear (inverse) demand function has the form $p = a + bq$
  - Revenue is equal to price times quantity:

$$\begin{aligned} R(q) &= pq \\ &= (a+bq)q \\ &= aq + bq^2 \end{aligned}$$

  - Marginal revenue $MR(q)$ at production level q is then:

$$MR(q) = a + 2bq$$

NB: The marginal revenue curve has the same *y*-intercept *a* as the demand curve and has a slope (2*b*) twice as steep downward as the demand curve

- Proof (optional): (uses calculus)

$$\begin{aligned} MR(q) &= \frac{dr}{dq} \\ &= \frac{d\left(aq+bq^2\right)}{dq} \\ &= a + 2bq \end{aligned}$$

Marginal revenue is the first derivative of the revenue function

What mathematicians use at the of a proof. Use this to impress your friends. It means *quod erat demonstrandum,* which roughly translates into "what was to be demonstrated."

Q.E.D.

Costs

- Cost function
  - The cost to produce output q depends on the costs of the inputs to produce quantity q
  - The *technology* available to the firm provides the relationship between the inputs (including labor and capital) the firm purchases and the output the firm can produce with those inputs
  - The firm's *cost function* $c(q)$ is the minimum cost to the firm of producing quantity q given the firm's technology
    - The firm's cost function c may change as the technology changes

This just means that c is a function of q, that is, cost is a function of the quantity produced

Costs

- Cost function—Some useful definitions
  - *Total cost* (TC)
    - The sum of all costs incurred by the firm to produce output *q*. Total cost is equal to the sum of fixed cost plus variable cost.
  - *Fixed cost* (FC)
    - The cost incurred by the firm that do not depend on the firm's level of production (e.g., the cost of the factory)
  - *Variable cost* (VC)
    - The cost incurred by the firm that depends on the firm's level of production
  - *Average total cost* (ATC)
    - Total cost divided by output
  - *Average variable cost* (AVC)
    - Variable cost divided by output
  - *Marginal cost* (MC)
    - The cost to the firm of producing one incremental unit of output

$$TC(q) = FC + VC(q)$$

$$ATC(q) = \frac{TC(q)}{q}$$

$$AVC(q) = \frac{AC(q)}{q}$$

$$MC(q) = C(q) - C(q-1)$$

$$= \frac{\Delta C}{\Delta q} \quad \text{where } \Delta q = 1$$

$$= \frac{dC}{dq} \quad \text{(in calculus notation)}$$

Marginal costs

Marginal cost (*mc*): The cost *mc* of producing the $(x + 1)^{th}$ unit after producing *x* units
Marginal cost curve: Traces the relationship between *x* and *mc*

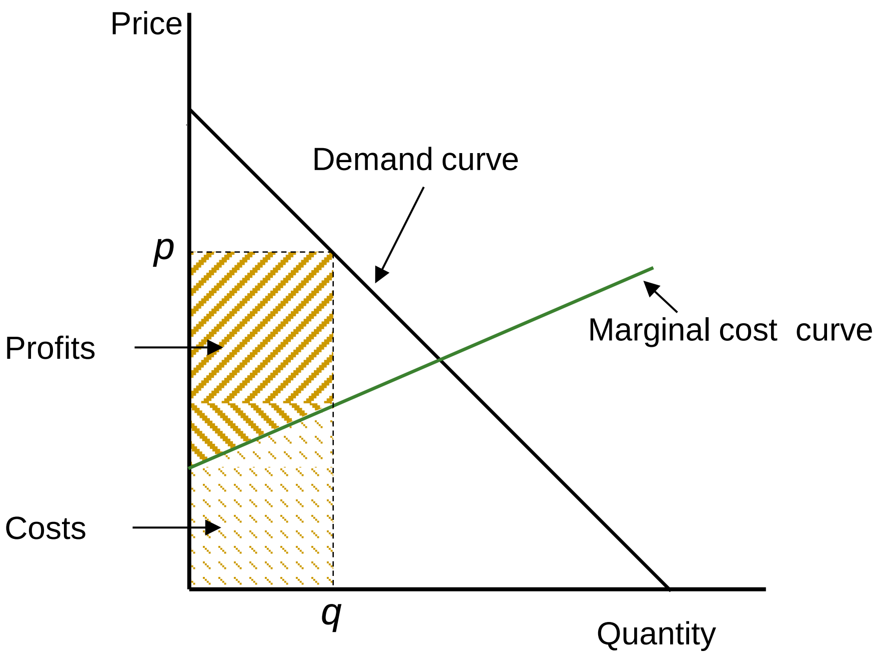

$$c(q) = \sum_{i=1}^{q} mc_i$$

Throughout the course, we will usually assume that marginal cost is constant, that is, that each incremental unit costs the same amount to make. This is a common assumption for small changes in production.

Query: The marginal cost curve is shown upward sloping. Why might that be?
Can the marginal cost curve be flat or even downward sloping?

Total costs

- Recall some definitions
  - *Total cost* (TC): The sum of all costs incurred by the firm to produce output *q*. Total cost is equal to the sum of fixed cost plus variable cost
  - *Fixed cost* (FC): The cost incurred by the firm that do not depend on the firm's level of production (e.g., the cost of the factory)
  - *Variable cost* (VC): The cost incurred by the firm that depends on the firm's level of production
  - *Marginal cost* (MC): The cost to the firm of producing one incremental unit of output
- Some important cost relationships
  - Variable cost is equal to the sum marginal costs to reach production level *q*:

$$VC(q) = \sum_{i=1}^{q} mc_i$$

That is, variable cost is the sum of the marginal costs of producing each successive unit up to production level *q*

    - When marginal costs are constant at a level *k*, the variable costs for a production level *q* is $VC(q) = kq$
  - Total cost is equal to fixed cost plus variable cost (all for producing a level of output *q*):

$$\begin{aligned} TC(q) &= F + VC(q) \\ &= F + kq \text{ (when marginal cost is a constant } k\text{)} \end{aligned}$$

Profit maximization

- Profit maximization
  - Firm's objective function in revenues (with quantity *q* as the control variable):

$$\max_q \ \text{Profits} = \text{Revenues} - \text{Costs}$$

$$= r(q) - c(q)$$

This say to pick q (the control variable) to maximize the function

Models that use q as the control variable are called *Cournot models*. Learn this term.

This equation says pick production level *q* to maximize profits, that is, the difference between the revenues the firms earns when it sells quantity *q* and the costs it incurs to produce quantity *q.*

In this maximization problem, the *objective function* is the function that we are trying to maximize, in this case *r*(*q*) – *c*(*q*).

The *control variable* is the variable the firm gets to pick. In this simple model, the firm can control its production level *q*, but market conditions determine the price at which the sells. Variables that the firm does not control are called as *parameters*.

Alternatively, we could develop a model in which price *p* as the control variable. Models that use p as the control variable are called *Bertrand models*. We will study Bertrand models where *p* is the control variable later.

Profit maximization

- Profit maximization
  - The profit function looks like a hill

    Think about it this way: When price equal zero, there are zero profits. At some high price point, no one will be willing to buy, so there are zero profits. In between, there are positive profits.

  - The profit-maximizing quantity *q** is the quantity at the peak of the profit curve

    Economists typically use an asterisk to denote an optimum, so that *q** is the profit-maximizing level of output and π* is the maximum level of profits.

  - The profit-maximizing quantity *q** occurs when MR = MC

    MR = MC is called the *first order condition* for a profit maximum. This attribute of a profit maximum is invoked frequently in antitrust analysis.

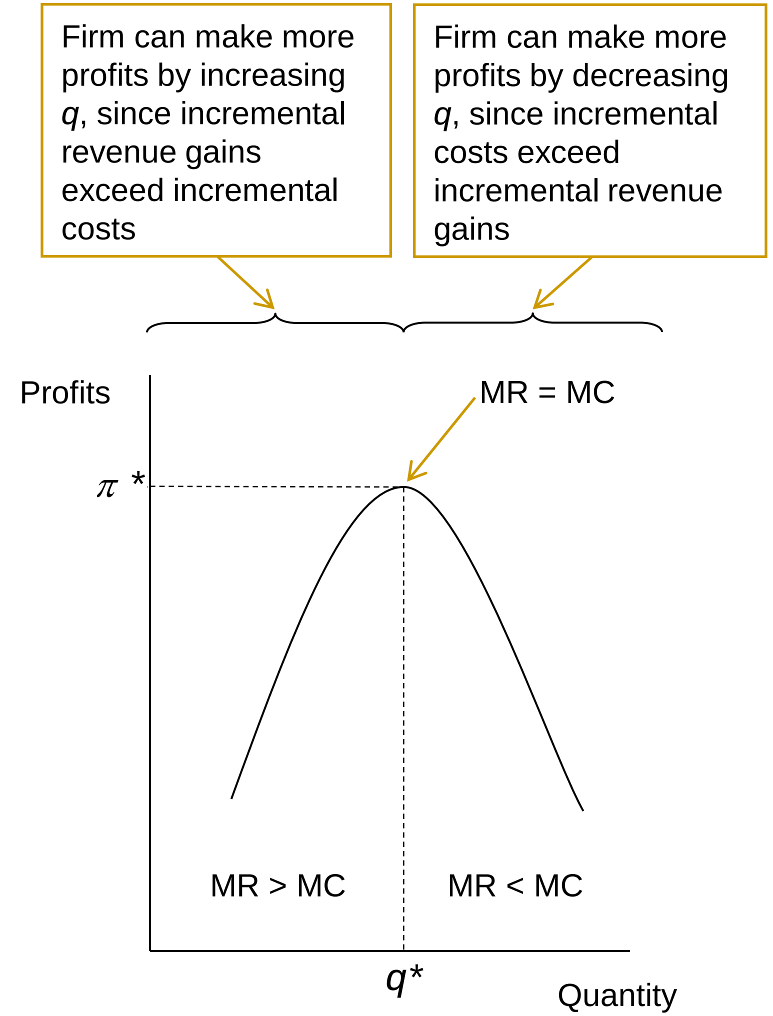

Profit maximization

- Profit maximization (this time with a little calculus)
  - At its peak, the slope of the profit curve is zero, that is, where

$$\frac{\Delta \pi}{\Delta q} = 0$$

  - We get the same result by setting the derivative of the profit function to zero:

$$\frac{d\pi}{dq} = \frac{dr}{dq} - \frac{dc}{dq} = 0$$

  - Rearranging terms yields:

Marginal revenue → $$\frac{dr}{dq} = \frac{dc}{dq}$$ ← Marginal cost

which is just another way of saying marginal revenue equal marginal cost

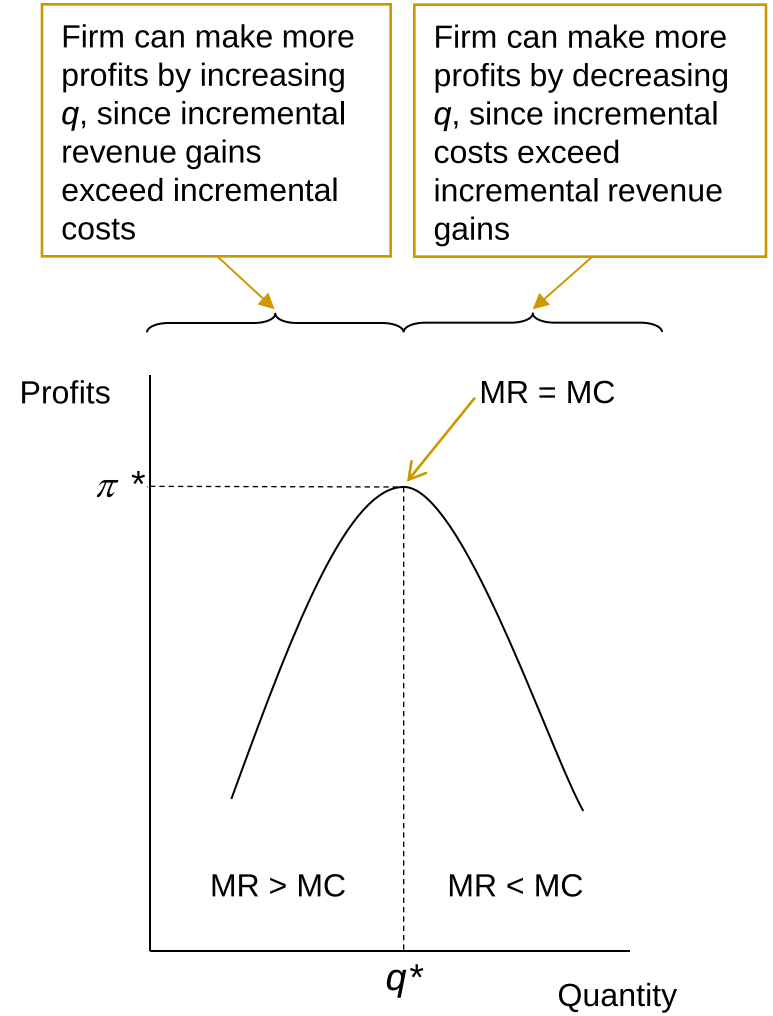

Profit maximization

- In Cournot models, firms maximize profits when they pick *q* so that *mr* = *mc*

Firm can make more profits by increasing *q*, since incremental revenue gains exceed incremental costs

Firm can make more profits by decreasing *q*, since incremental costs exceed incremental revenue gains

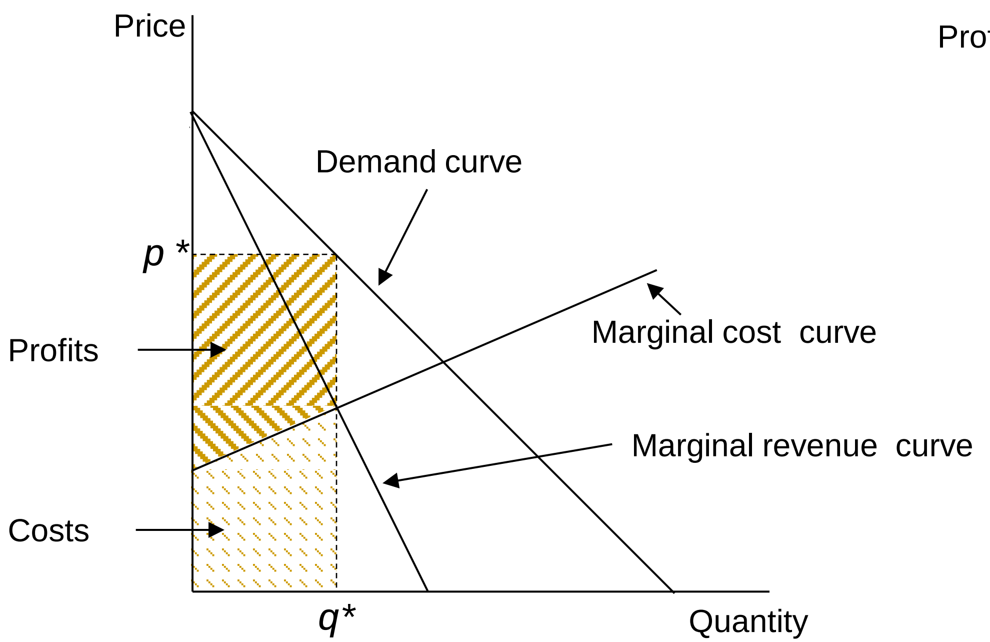

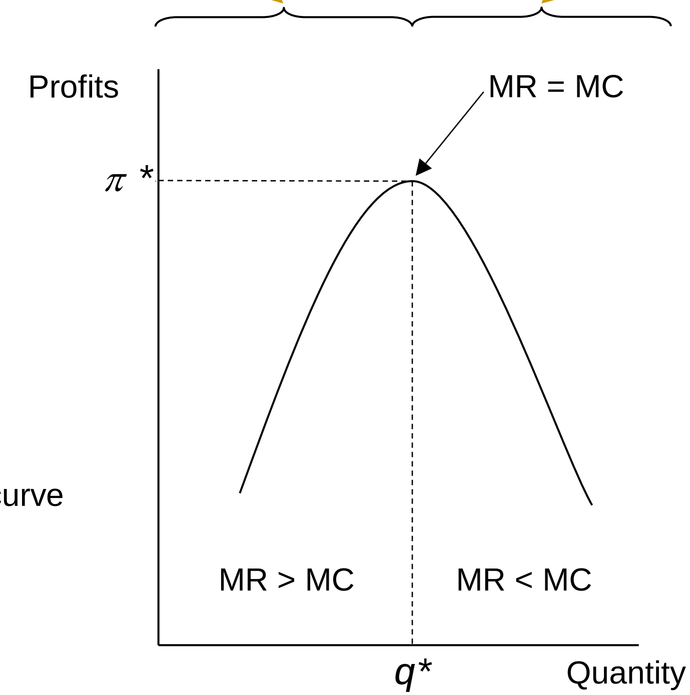

NB: Make sure that you understand what is going on here. The firm picks q* so that it marginal revenue is equal to its marginal cost. It then choses a *p** that clears the market for the quantity *q**. See the next slide for an example.

Profit maximization

- Profit maximization for the individual firm—Example

Assume $q = 20 - 2p$ (firm's residual demand curve)

so $p = 10 - \frac{1}{2}q$ (inverse demand curve)

Revenue $(r) = pq = (10 - \frac{1}{2}q)q = 10q - \frac{1}{2}q^2$

Marginal revenue $(MR) = \frac{dr}{dq} = 10 - q$

Constant marginal cost $(MC) = 4$

Equating marginal revenue and marginal cost

for a profit maximum: $10 - q^* = 6$

So $q^* = 6$ is the firm's profit-maximizing quantity

Plugging q^* into the inverse demand function to obtain

$p^* = 7$ as the firm's profit-maximizing price

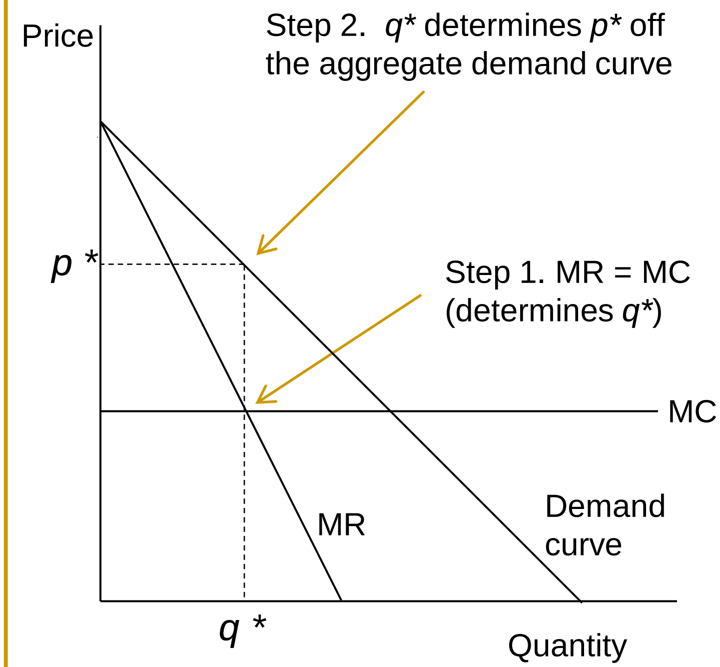

Profit maximization

- Calculus version (another chart)

Demand: $p = 10 - \frac{1}{2}q$

Revenue: $r = pq = 10q - \frac{1}{2}q^2$

Marginal revenue: $mr = \frac{dr}{dq} = 10 - q$

Marginal cost: $mc = 4$

Profit function: $\pi = r - c = \left[10q - \frac{1}{2}q^2\right] - 4q$

$$= 6q - \frac{1}{2}q^2$$

Profit max:
$$mr = mc$$
$$10 - q = 4$$
$$q = 6$$
$$p = 7$$

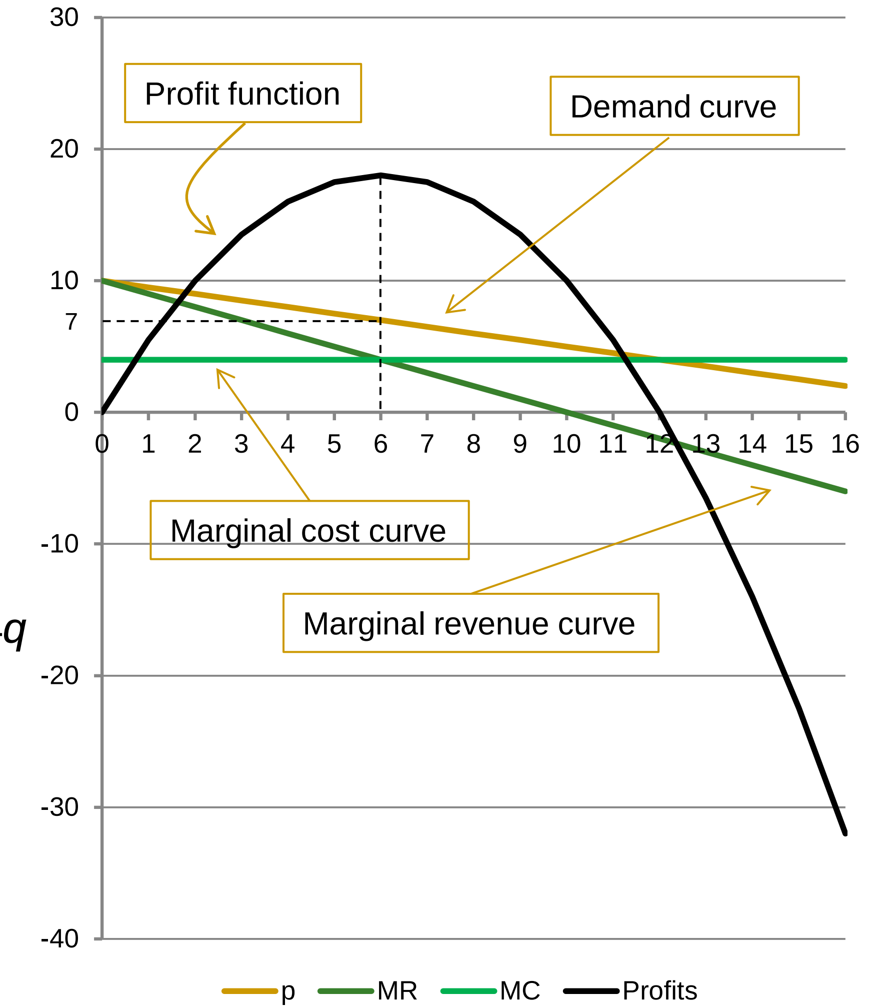

Profit maximization

- Numerical version—Same example

Quantity q	Price p	Revenue r	Marginal Revenue mr	Marginal Costs mc	Total Costs c	Profits Π
0	10.0	0.0			0.0	0.0
1	9.5	9.5	9.5	4.0	4.0	5.5
2	9.0	18.0	8.5	4.0	8.0	10.0
3	8.5	25.5	7.5	4.0	12.0	13.5
4	8.0	32.0	6.5	4.0	16.0	16.0
5	7.5	37.5	5.5	4.0	20.0	17.5
6	7.0	42.0	4.5	4.0	24.0	18.0
7	6.5	45.5	3.5	4.0	28.0	17.5
8	6.0	48.0	2.5	4.0	32.0	16.0
9	5.5	49.5	1.5	4.0	36.0	13.5
10	5.0	50.0	0.5	4.0	40.0	10.0
11	4.5	49.5	-0.5	4.0	44.0	5.5
12	4.0	48.0	-1.5	4.0	48.0	0.0
13	3.5	45.5	-2.5	4.0	52.0	-6.5
14	3.0	42.0	-3.5	4.0	56.0	-14.0
15	2.5	37.5	-4.5	4.0	60.0	-22.5
16	2.0	32.0	-5.5	4.0	64.0	-32.0
17	1.5	25.5	-6.5	4.0	68.0	-42.5
18	1.0	18.0	-7.5	4.0	72.0	-54.0
19	0.5	9.5	-8.5	4.0	76.0	-66.5
20	0.0	0.0	-9.5	4.0	80.0	-80.0

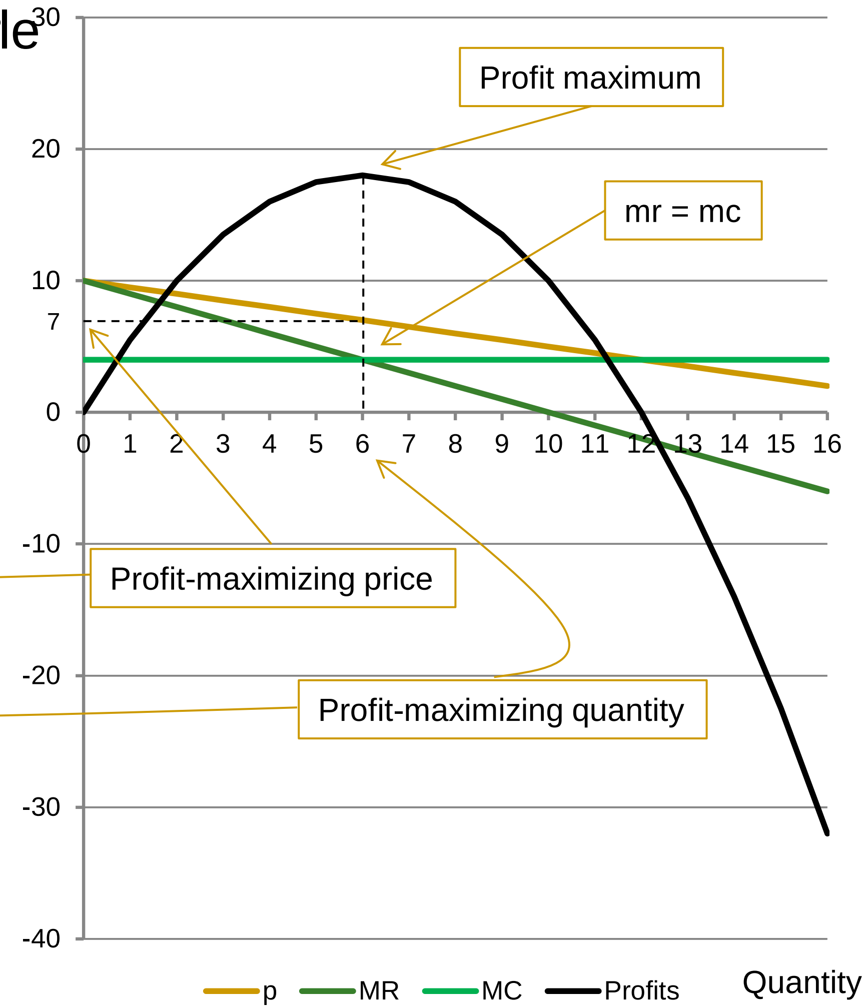

Profit maximization

- Illustration of profit loss from price changes from *p**
  - Assuming no fixed costs

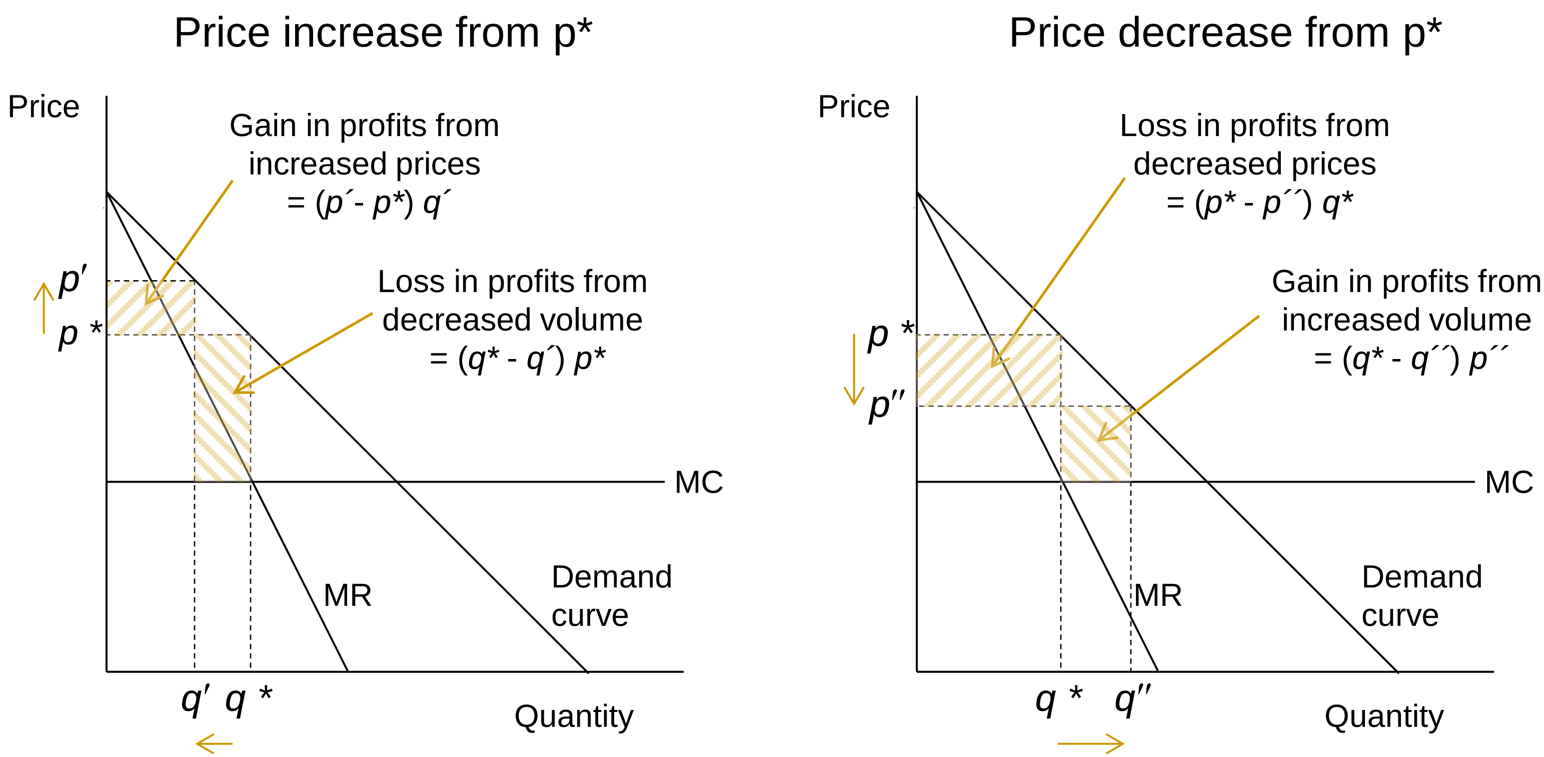

In each case, the loss from the price change exceeds the gain, so that moving away from *p** decreases profits.

Perfect Market Equilibria

Perfectly Competitive Markets

- *Definition*: A market in which no single firm can effect price, meaning:
  - The firm's residual demand curve is horizontal,
  - The firm can sell any amount of product without affecting the market price,
  - $\frac{dp}{dq} = 0$, or
  - $p = \frac{dc}{dq}$ (i.e., price = marginal cost)

These four bullets are just different ways of saying exactly the same thing.

- What could cause a market to be perfectly competitive?
  - *Traditional theory*: Each individual firm's production is very small compared to aggregate demand at any price, so that individual production changes cannot move significantly along the aggregate demand curve
    - This implies that there are a very large number of firms in the market
  - *Modern theory*: Competitors in the market place react strategically but non-collusively to price or quantity changes by a firm in ways that maintain the competitive equilibrium

Competitive firms

- Competitive firms take prices as given
  - → Individual output decisions do not affect the market-clearing price

When price does not change as the firm expands output, the firm will produce every unit for which $p \geq MC$

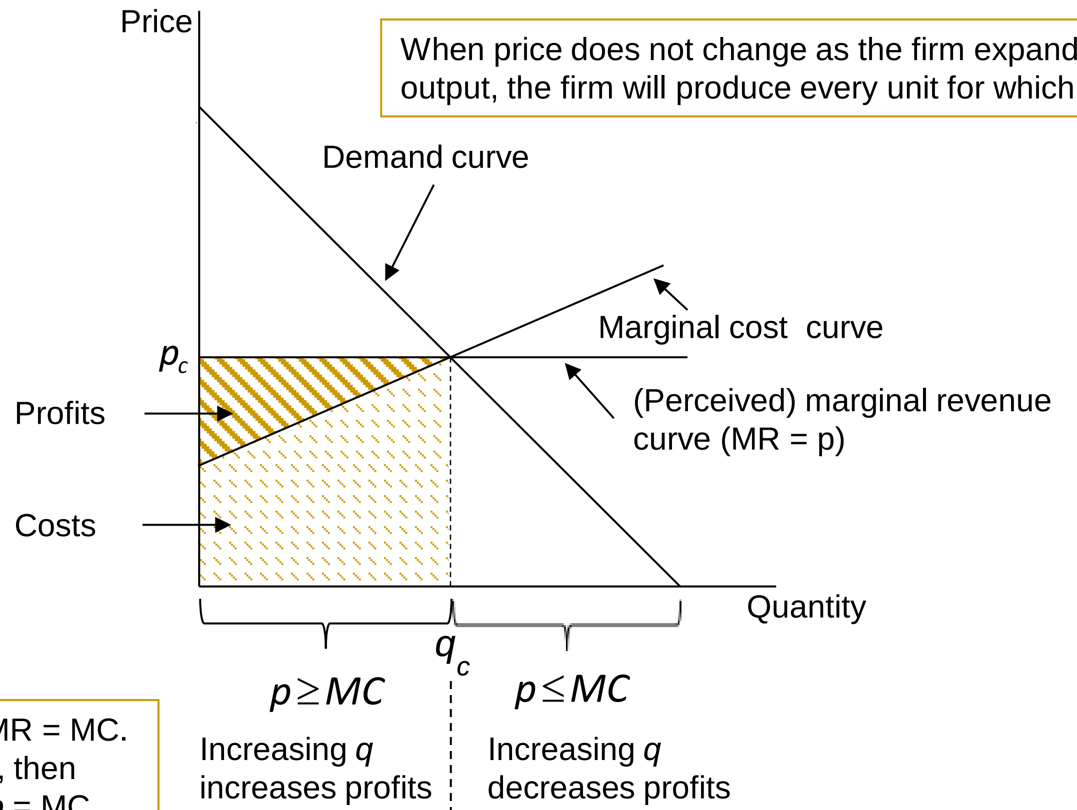

As always, the FOC is MR = MC. If the firm is competitive, then MR = *p and so* FOC is *p* = MC

Perfect Monopoly

- *Economic definition*: A market in which only one firm operates
  - In this case, the firm's residual demand curve is the same as the aggregate demand curve
  - As always, the firm's chooses production so that its marginal cost is equal to marginal revenue

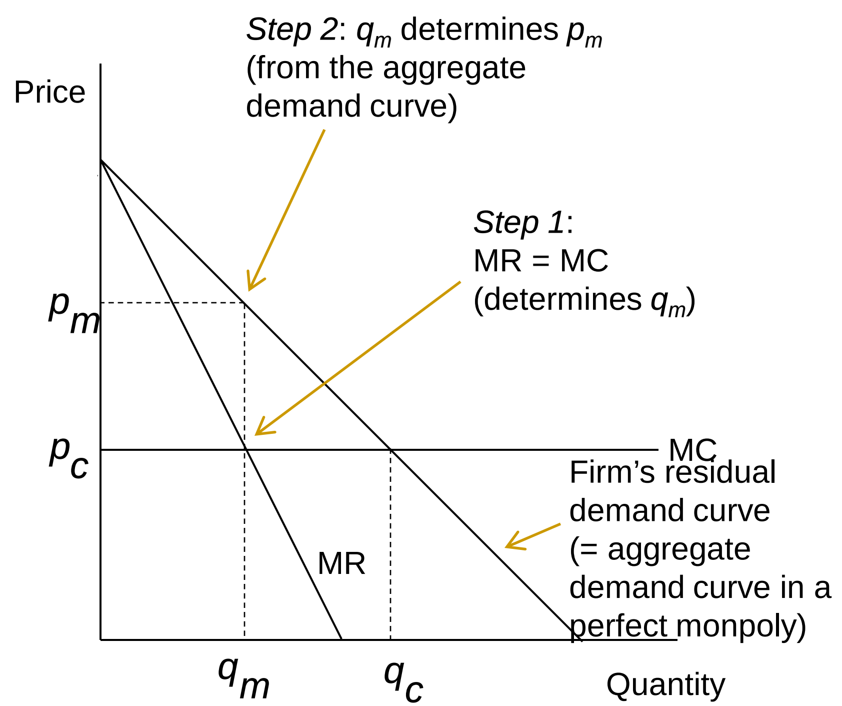

Keep in mind that this is the way economists define "perfect monopoly." As we shall see later, the legal definition of "monopoly" is quite different.

Monopolist Firm

- A monopolist choice of output *q* affects the market-clearing price *p*
  - Contrast with a perfectly competitive firm, which cannot affect p through its choice of q (by definition)

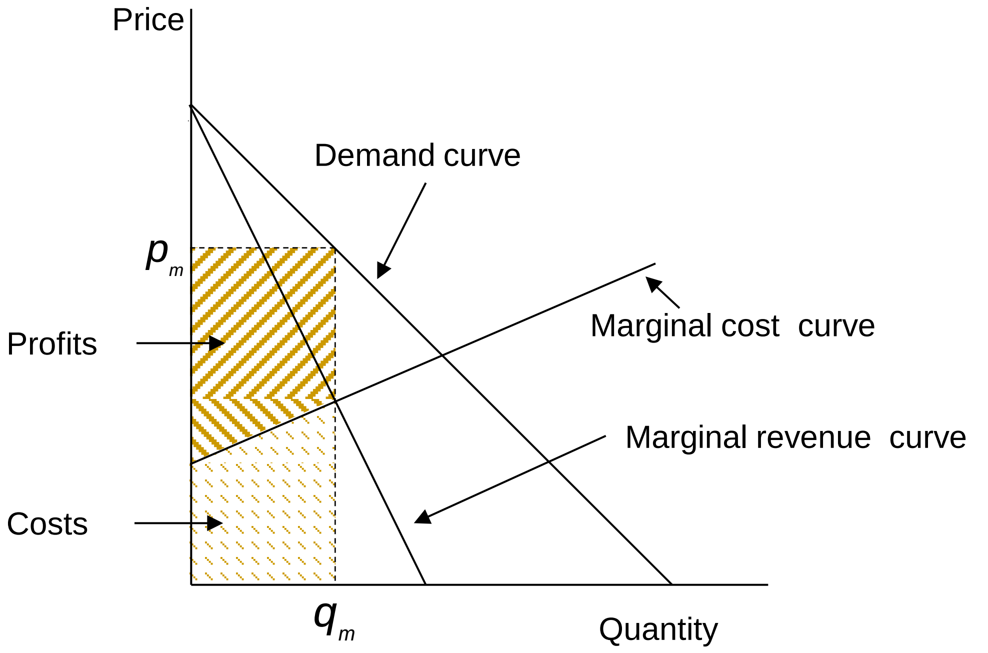

Monopolists (like all profit-maximizing firms) price at MR = MC, where the monopolist’s marginal revenue is determined by the aggregate demand curve

Summary of Key Results So Far

- Profit-maximizing firms choose production levels so that marginal revenue equals marginal cost (MR = MC) (in Cournot models)
  - Step 1: MR = MC determines the firm's profit-maximizing production level q^*
  - Step 2: The firm's residual demand curve determines the firm's profit-maximizing price p^* given q^*
- In a perfectly competitive market, a firm's choice of production level cannot effect market price, so:
  - Marginal revenue is equal to the market unit price (MR = p_{market}),
  - MR = MC implies that MC = p_{market}, so firm picks $q_{competitive}$ to satisfy this condition
    - We have not discussed how the market price is determined in a perfectly competitive market. For our purpose, just take market price as a given.
  - By definition, firm cannot affect market price, so $p_{competitive} = p_{market}$
- In a perfect monopoly market, consumers can only purchase from the monopolist, so:
  - The firm's residual demand curve is the same as the aggregate demand curve
  - MR = MC determines the monopolist's profit-maximizing quantity q_m
  - The aggregate demand curve determines p_m given q_m