Unit 8. Competition Economics

Part 1. Demand, Costs, and Profits

Merger Antitrust Law

Georgetown University Law Center
Dale Collins

October 5, 2019

Topics

- Demand
- Supply
- Profits and profit maximization

Terms to know

Demand	Demand curves	The Demand Curve
		A Deeper Look at the Demand Curve
	Shift in the demand curve	The Demand Curve Shifts
	Aggregate demand curve	
	Residual demand curve	
	Inverse demand curves	From Demand to Inverse Demand
	Revenues	
	Revenue functions	
	Incremental revenue	
	Marginal revenue	Marginal Revenue for a Monopolist Facing Linear Demand

Terms to know

Supply	Total costs	
	Fixed costs	
	Variable costs	
	Marginal costs	
	Constant marginal cost	
Profit maximization	Profits	
	Profit maximization	Finding a Revenue Function from a Linear Demand Function
	Control variables	
	First order condition	

Basic competition economics

- A plea on notation

Do not be put off by the mathematical notation in the slides that follow. For the most part, the notation is just a language that can be used to express concepts precisely that otherwise would take a lot of words. 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 it is very simple. Learning the basic economics and the mathematics notation in which it is expressed is an investment that will give you a significant comparative advantage over many other antitrust attorneys (and the antitrust economists will love you for it).

Demand

Demand functions

- The *demand function* gives the quantity consumers will purchase for any given price
  - *Example:* If the price is \$4.00, consumers will buy 12 units, but if the price is \$5.00 consumers will buy only 10 units
- Let q be the quantity demanded and p the market price
  - We say that *quantity is a function of price*
  - The demand function can be written as $q = d(p)$, where
    - $d(\cdot)$ is the demand function,
    - p is the market price, and
    - q is the quantity demanded by purchasers
  - *Example*: Say $p = 20 – 2p$.
    - Then if the market price p is 4, the quantity demanded is 12 (= 20 - 2·4)
    - If the market price is 5, then the quantity demand is 10 (= 20 - 2·5)
    - The *demand curve* is the graph of the demand function

The demand curve

- Say the demand function is $q = 20 - 2p$
- Then we can create a table that shows q for every p and plot the points in a chart to create the demand curve:

p	q
0	20
1	18
2	16
3	14
4	12
5	10
6	8
7	6
8	4
9	2
10	0

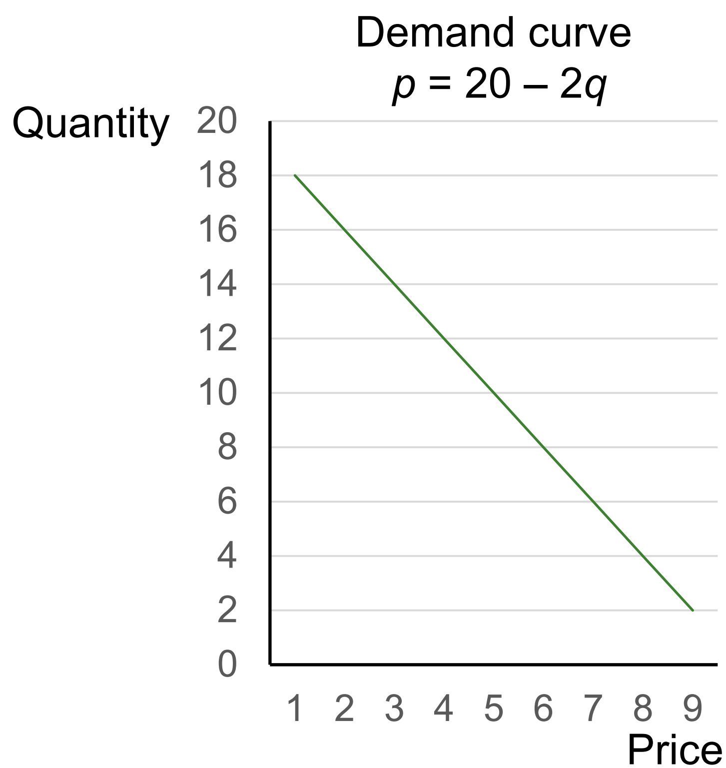

The law of demand

- *The law of demand*: Demand for a product decreases as price increases
  - Equivalently, the demand curve is downward sloping

The downward-sloping demand curve drives almost all of the results in antitrust economics

Linear demand curves

- *Linear demand curves* are straight lines
  - Although demand curves are downward-sloping, they need not be straight lines
  - All of the principles in which we will be interested, however, may be illustrated using linear demand curves
  - For simplicity, we will only use linear demand curves in this course

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
- Some notes
  - q and p are called *variables* and are the numbers of interest to us
  - 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., $b < 0$)

NB: This is important

Graphing a demand curve

a is called the *y-intercept.* It is the maximum quantity purchasers will buy at a zero price.

Demand curve: $q = 20 - 2p$
General form: $q = a + bp$
So
$a = 20$ (the y-intercept)
$b = -2$ (the slope)

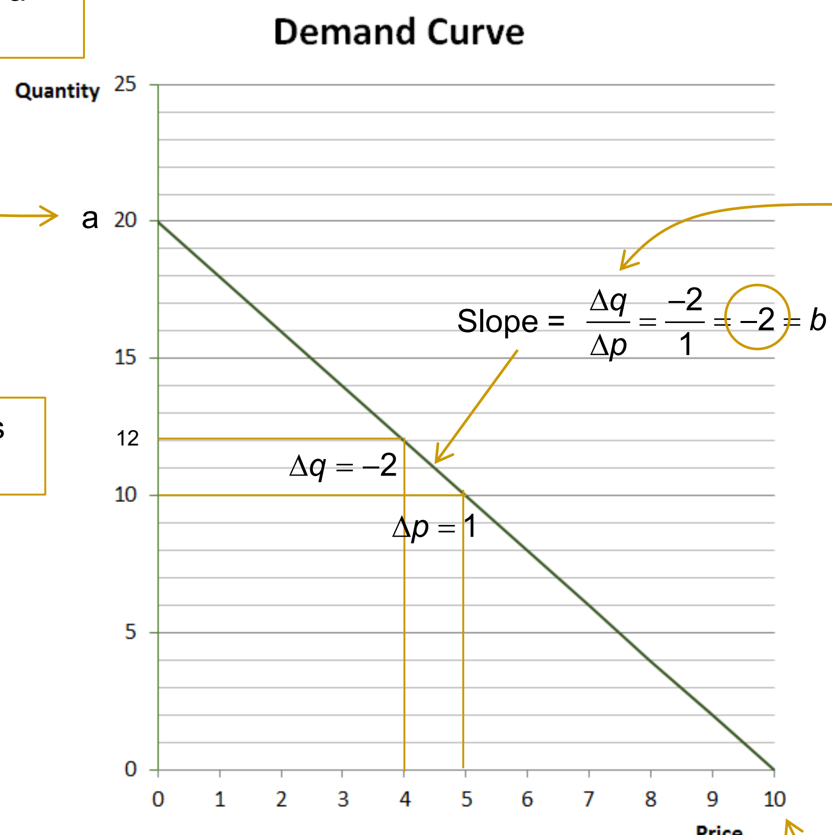

The Greek letter Δ (delta) means change, so Δp means $p_2 - p_1$

The vertical axis is called the *y-axis*

The slope *b* of the demand curve gives the rate of change in demand for an increase in price. So, in this example, $b = -2$, so demand decreases by 2 units for every dollar of price increase.

The horizontal axis is called the *x-axis*

The is the "choke price," that is, the price above which no customer is will to purchase

A nonlinear demand curve

- Demand curves do not have to be straight lines as long as they are downward-sloping
- BUT—
  - 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 demand curve with a linear demand curve

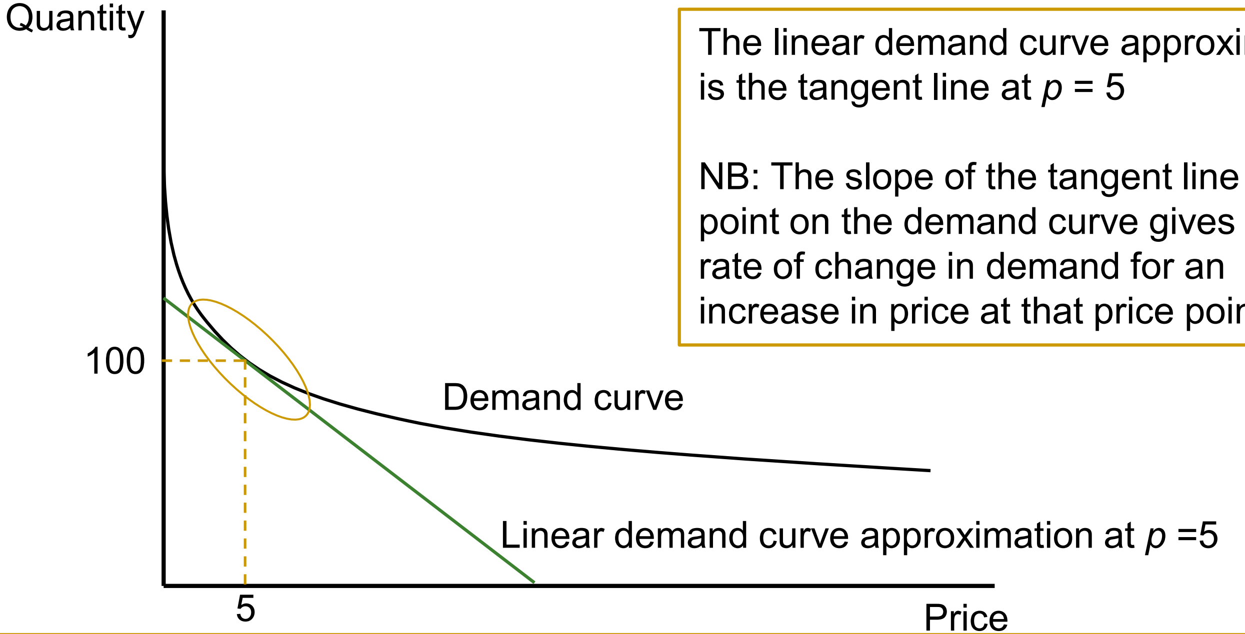

The linear demand curve approximation is the tangent line at *p* = 5

NB: The slope of the tangent line at a point on the demand curve gives the rate of change in demand for an increase in price at that price point

A shift in demand

- 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

Two points on the same demand curve

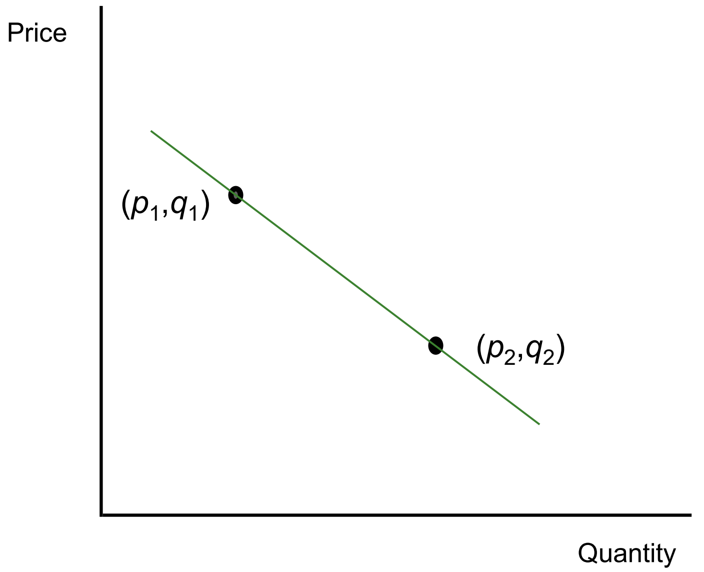

Two points on different demand curves

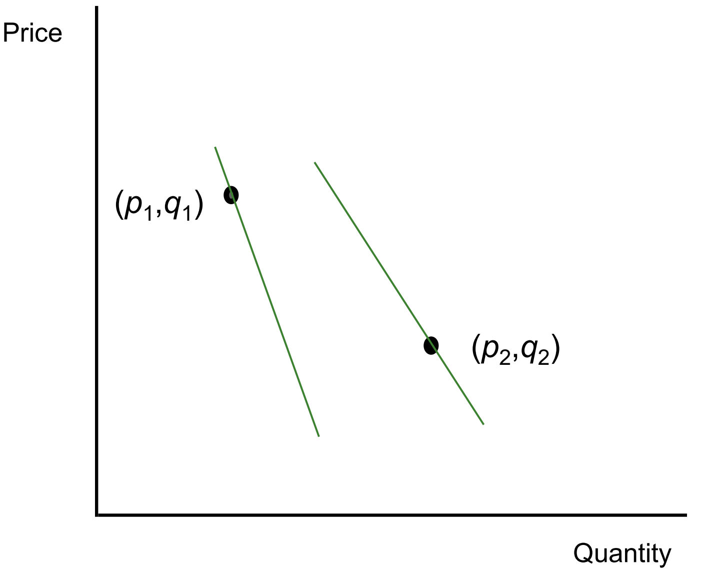

NB: It is important to distinguish between changes along a demand curve and shifts in demand

A shift in demand

- Demand shifts to the right
  - When demand *shifts to the right,* purchasers—
    - will be willing to pay a higher price to clear the market of the same quantity, or
    - equivalently, demand more quantity at a given price

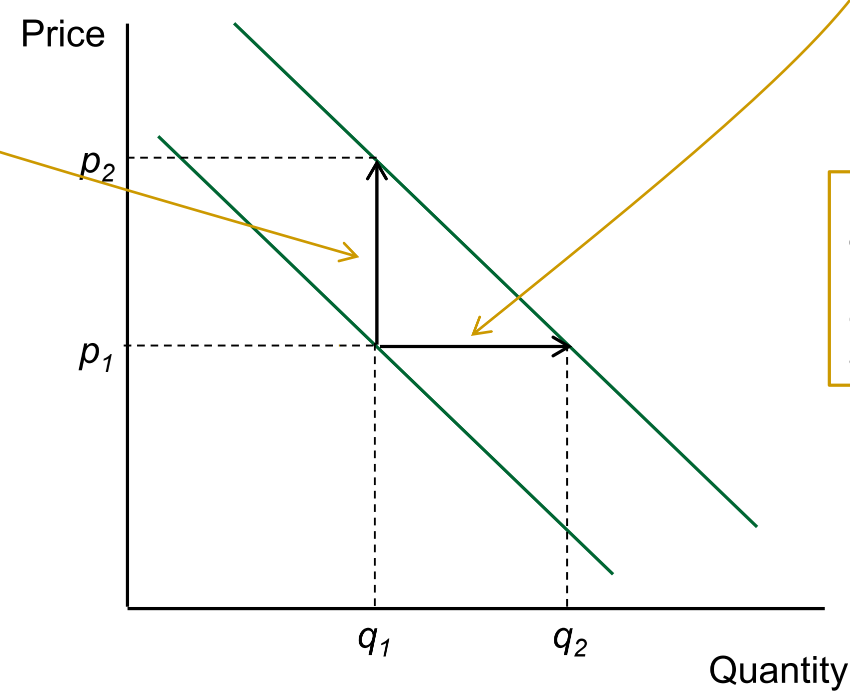

NB: Shifting the demand curve to the right improves consumer welfare at a given price level

Aggregate and residual demand curves

- We can define a demand curve for a single purchaser, all purchasers in the market from all firms in the market, or all purchasers for a single firm in the market
  - Say
    - There are a total of *N* purchasers in the market
    - Each purchaser *i* (where *i* can be any number from 1 to *N*) has a demand function of the form:

$$q_i = d_i(p)$$

      where *p* is the price, d_i is the i^{th} purchaser's demand function, and q_i is the quantity the i^{th} purchaser's demand at price *p*
  - *Aggregate demand curve*: Describes the quantity demanded by all purchasers from all firms in the market at each given price
    - The aggregate demand curve is the sum of the individual demand curves
    - In other words:

$$Q(p) = \sum_{i=1}^{N} d_i(p),$$

> The Greek capital letter sigma just says to add all of the demands d_i from purchasers *i* through *N* at a given price *p*

      where *Q*(*p*) is the total quantity demanded by all *N* purchasers from all firms in the market at price *p*

Aggregate and residual demand curves

- *Residual demand curve*: Describes the quantity demanded by all purchasers from a given single firm in the market at each given price
  - The residual demand curve is the sum of the quantities demanded from the firm from all purchasers in the market
    - In general, some purchasers may buy some or all of their products from other firms in the market; this demand for purchases from other firms is not include in the residual demand curve for the firm in question
  - In other words:

$$Q^j(p) = \sum_{i=1}^{N} d_i^f(p),$$

This is the individual demand function of purchaser *i* from firm *f* at price *p*

where $Q^j(p)$ is the total quantity demanded by all *N* purchasers from firm *j* at price *p* (i.e., the residual demand for firm *f* at price *p*)

Inverse demand curves

Demand curve: Traces the relationship between *q* and *p* (where *p* is a function of *q*)
Inverse demand curve: Traces the relationship between *p* and *q* (where q is a function of p)

Demand Curve | **Inverse Demand Curve**

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

(a little algebra)

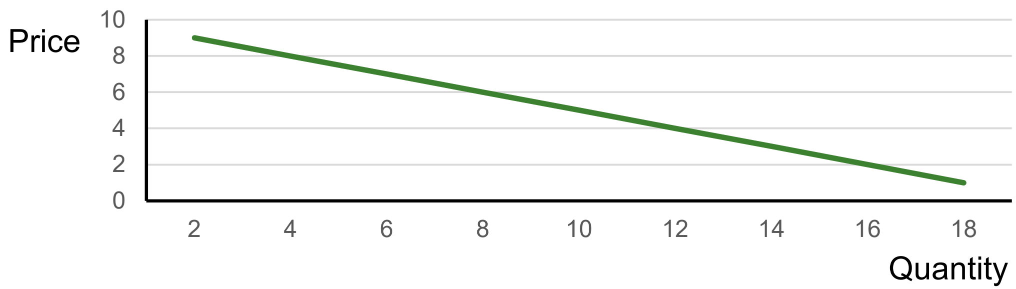

Note: If the demand curve is downward sloping, then the inverse demand curve will be downward sloping

Inverse demand curves

- Another way to think about the inverse demand curve
  - All we have done is switch the axes:
    - The horizontal x-axis in the demand curve (price) becomes the vertical y-axis in the inverse demand curve
    - The vertical y-axis in the demand curve (quantity) becomes the horizontal x-axis in the inverse demand curve

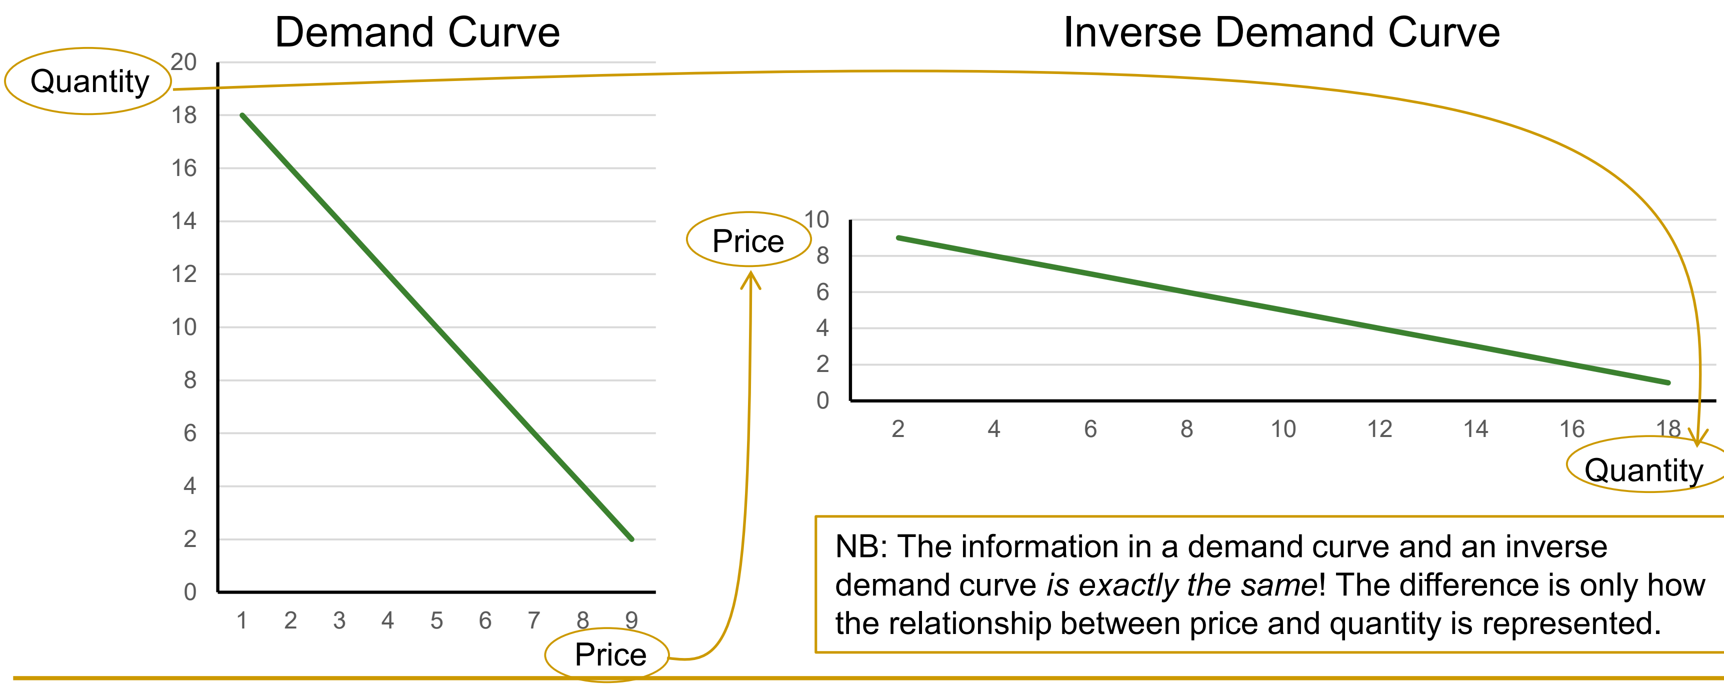

NB: The information in a demand curve and an inverse demand curve *is exactly the same*! The difference is only how the relationship between price and quantity is represented.

More on demand and inverse demand curves

- How to read demand and inverse demand curves
  - *Demand curves* say what quantity *q* would be purchased at a given price *p*
  - *Inverse demand curves* say what price *p* would "clear the market" of a given quantity *q*
    - The *market-clearing price* for a quantity *q* is the price *p* at which the demand for the product is exactly *q*—no more and no less
- Important!
  - The convention among economists is to use inverse demand curves but call them demand curves
    - Both demand curves and inverse demand curves contain the same information, so being precise usually does not matter
    - If it is important to know the difference, the reader will know by whether the variable to be determined is quantity as a function of price (a demand function) or market-clearing price as a function of quantity (an inverse demand function)
    - We will follow the convention in this course

The slope of an inverse demand curve

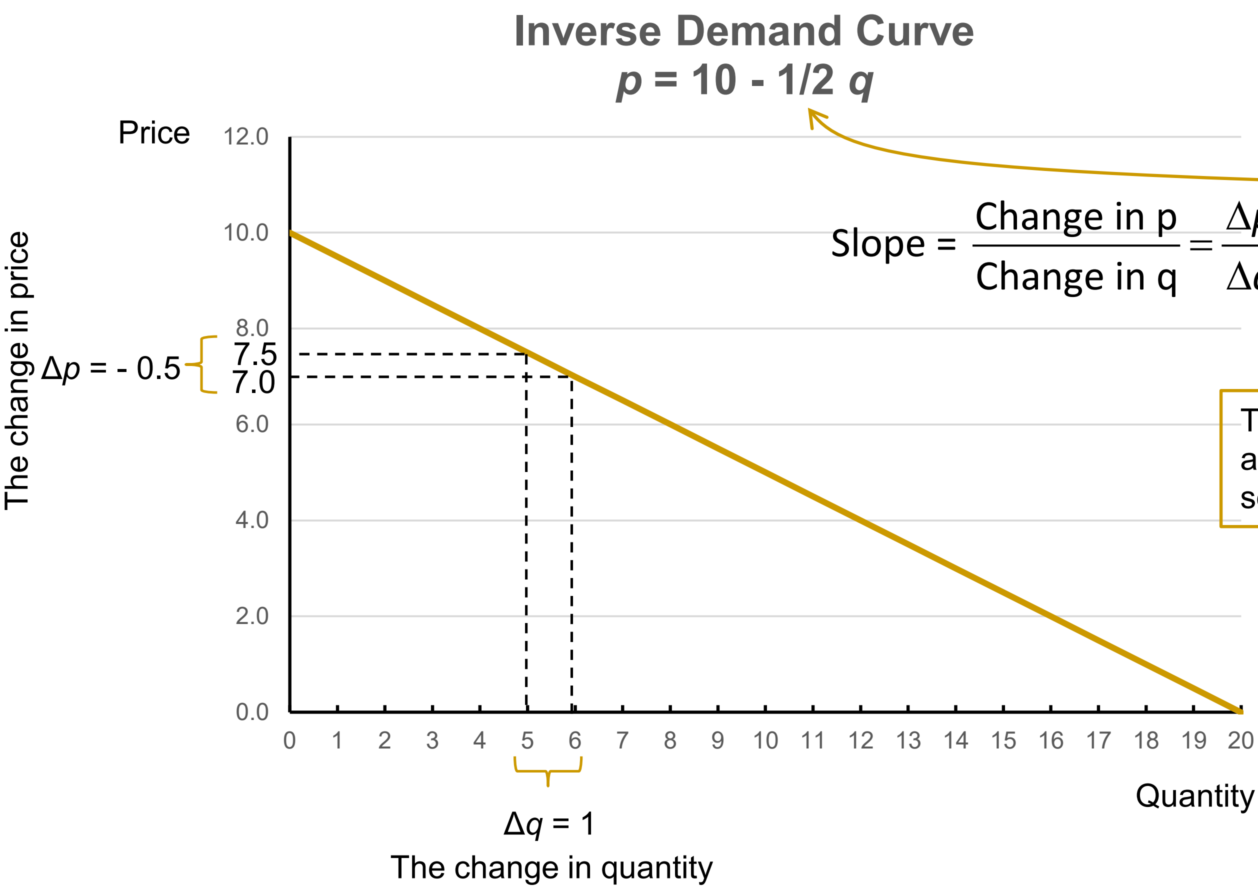

$$\text{Slope} = \frac{\text{Change in p}}{\text{Change in q}} = \frac{\Delta p}{\Delta q} = \frac{7.5 - 7.0}{5 - 6} = \frac{0.5}{-1} = -\frac{1}{2}$$

The slope of -1/2 says that for one additional unit of the product to be sold, the price must drop by $0.5.

Revenue

- *Revenue* = Price times quantity (that is, $r = pq$)
  - Revenue may be depicted as a box in the graph of the inverse demand function

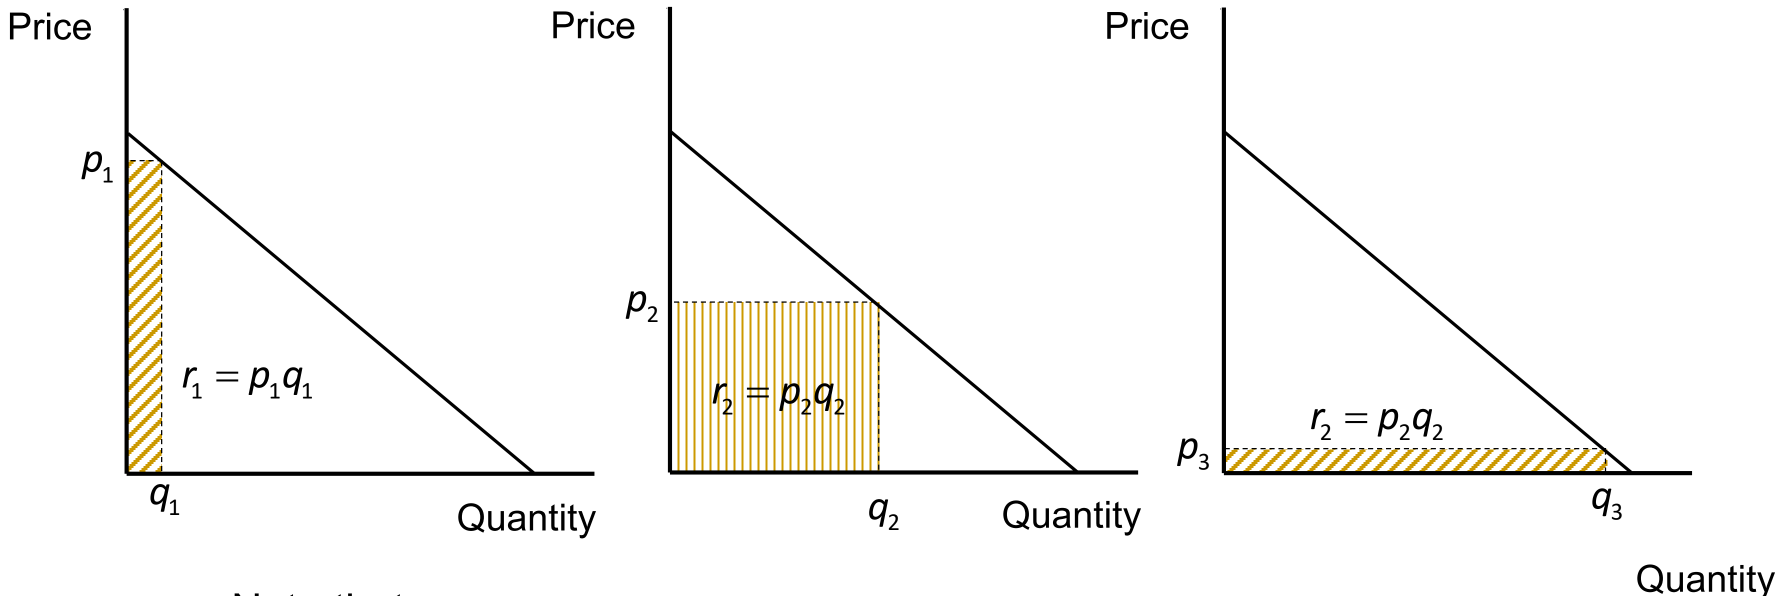

- Note that revenues—
  - Are small when the quantity sold is small and price is high (p_1, q_1)
  - Grow larger when the quantity sold become larger and the price drops (p_2, q_2)
  - But then grow smaller past a certain point as the quantity sold becomes even larger (p_3, q_3)

Revenue

- We can show this relationship between *q* and *r* in a graph

Inverse demand curve $p = 10 - 1/2\ q$		
p	q	r
10.0	0	0.0
9.5	1	9.5
9.0	2	18.0
8.5	3	25.5
8.0	4	32.0
7.5	5	37.5
7.0	6	42.0
6.5	7	45.5
6.0	8	48.0
5.5	9	49.5
5.0	10	50.0
4.5	11	49.5
4.0	12	48.0
3.5	13	45.5
3.0	14	42.0
2.5	15	37.5
2.0	16	32.0
1.5	17	25.5
1.0	18	18.0
0.5	19	9.5
0.0	20	0.0

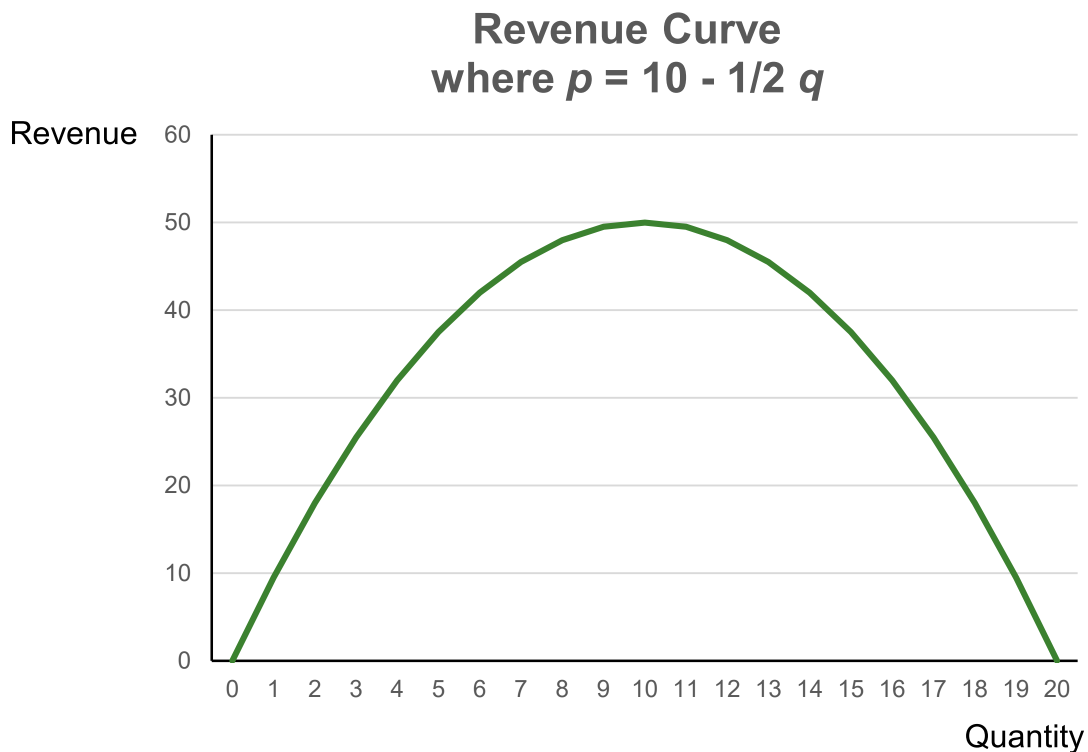

Mathematical aside: The revenue curve as a function of quantity is a parabola described by the quadratic equation $r = pq = (10 - ½\ q)\ q = 10q - ½\ q^2$

Incremental revenue

- Think about incremental revenue in two parts:
  1. The gain in revenue due to the sale of the additional units at the lower market-clearing price
     - Since there are more units to sell and demand is downward-sloping, the price will drop to clear the market
     - The gain in revenue is equal to $\Delta q \times (p - \Delta p)$, where
       - Δq is the additional quantity to be sold
       - Δp is the market price decrease necessary to clear the market with the sale of an additional unit
  2. Minus the loss of margin on prior units sold due to the decrease in the market-clearing price
     - This loss of margin is the prior quantity p times the required price decrease, or [$q\Delta p$]
- So

$$IR = \Delta q\left(p - \Delta p\right) - q\Delta p$$

Incremental revenue

- Graphically

To find incremental revenue *IR* when moving from q_1 to q_2, add Area A and subtract Area B

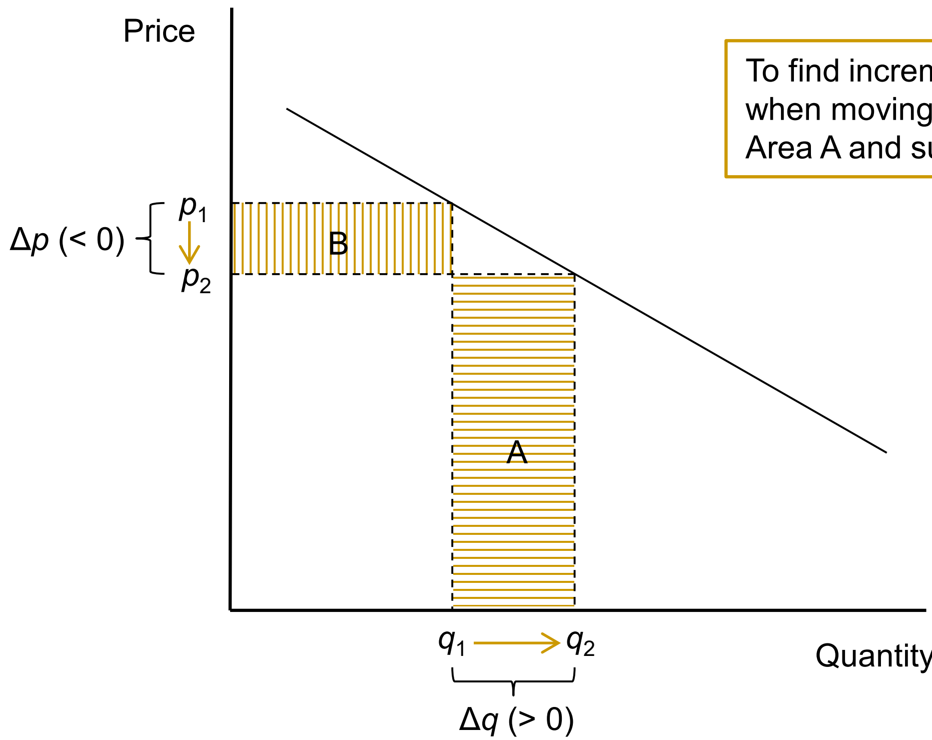

Area A = $\Delta q(p_1 - \Delta p)$ is the gain in revenue from the additional sales Δq at the lower price $p_2 = p_1 - \Delta p$

Area B = $q_1 \Delta p_1$ is the loss in revenue due to the sales of q_1 at the lower price p_2

So

$$IR = \underbrace{\Delta q\left(p - \Delta p\right)}_{\text{Area A}} - \underbrace{q \Delta p}_{\text{Area B}}$$

Incremental revenue

- Another graphical interpretation

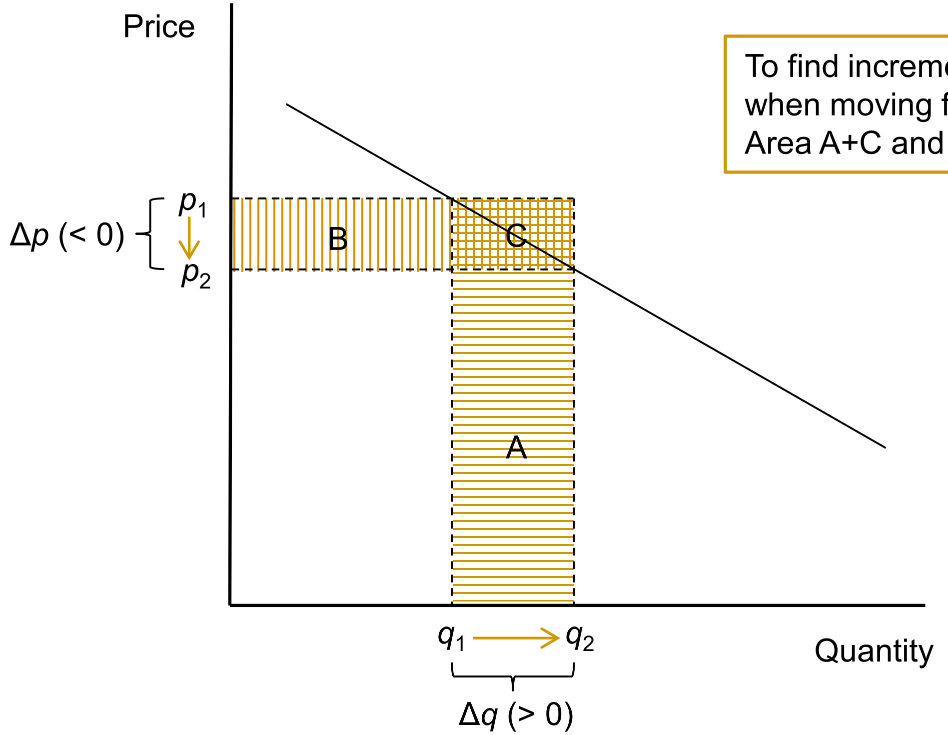

To find incremental revenue *IR* when moving from q_1 to q_2, add Area A+C and subtract Area B+C

Area A + C = $\Delta q(p_1)$ is the gain in revenue from the additional sales Δq at the original price p_1

Area B + C = $(q_1 + \Delta q)\Delta p_1$ is the loss in revenue due to the sales of $q_2 = q_1 + \Delta q$ at the lower price p_2

So

$$IR = \Delta q\left(p - \Delta p\right) - q\Delta p = \underbrace{p\Delta q}_{\text{Area A+C}} - \underbrace{\Delta p\left(q + \Delta q\right)}_{\text{Area B+C}}$$

A little algebra

Marginal revenue

- Some definitions
  - *Marginal revenue* $MR(q)$ is the slope of the revenue curve at a quantity q
  - In other words, it is the rate of which price has to change to clear the market when quantity is increased by a very small amount
- An approximation
  - Marginal revenue $MR(q)$ can be *approximated* by incremental unit revenue when Δq is small compared to q:

$$IUR = \frac{IR}{\Delta q} = \frac{r(q + \Delta q) - r(q)}{\Delta q} \approx MR(q)$$

    - This approximation becomes more accurate as q becomes larger, so that an increase of one unit of sales is an increasingly small percentage of total sales
  - Common interpretation
    - The heuristic interpretation of marginal revenue is the net additional revenue the firms would earn *if* it increased its output by one unit

$$MR(q) \approx r(q + 1) - r(q)$$

Marginal revenue

- A formula for marginal revenue
  - There is a formula for exactly calculating the slope of the tangent line to the revenue curve at a quantity where inverse demand is linear
  - Say $p = a + bq$ (a linear inverse demand curve, with $b < 0$)

    So $r = pq = (a + bq)q = aq + bq^2$

    Then the slope of the tangent line at q is $a + 2bq$. *This is also the marginal revenue at q*
  - Example
    - Suppose inverse demand is $p = 10 – ½ q$
    - Therefore, $r = pq = (10 – ½ q)q = 10q – ½ q^2$
    - Marginal revenue $MR(q) = 10 – q$ (from the formula)
      - $MR(5) = 10 – 5 = 5$
      - $MR(10) = 10 – 10 = 0$
      - $MR (15) = 10 – 15 = -5$

Marginal revenue

- Two observations
  - Marginal revenue is *less* than price when firms face a downward-sloping demand curve (that is, $MR(q) < p$) because:
    - The market price will have to *decrease* after adding the incremental output in order to clear the market
    - This lower price will apply to preexisting sales as well as incremental sales, making marginal revenue less than price
  - Marginal revenue is the *slope of the tangent line* on the revenue curve
    - *Example*: Where $p = 10 - \frac{1}{2} q$:
      - $MR(5) = 5$ $(p = 7.5)$
      - $MR(10) = 0$ $(p = 5)$
      - $MR(15) = -5$ $(p = 2.5)$

Note: $MR(q)$ means evaluate MR at quantity q

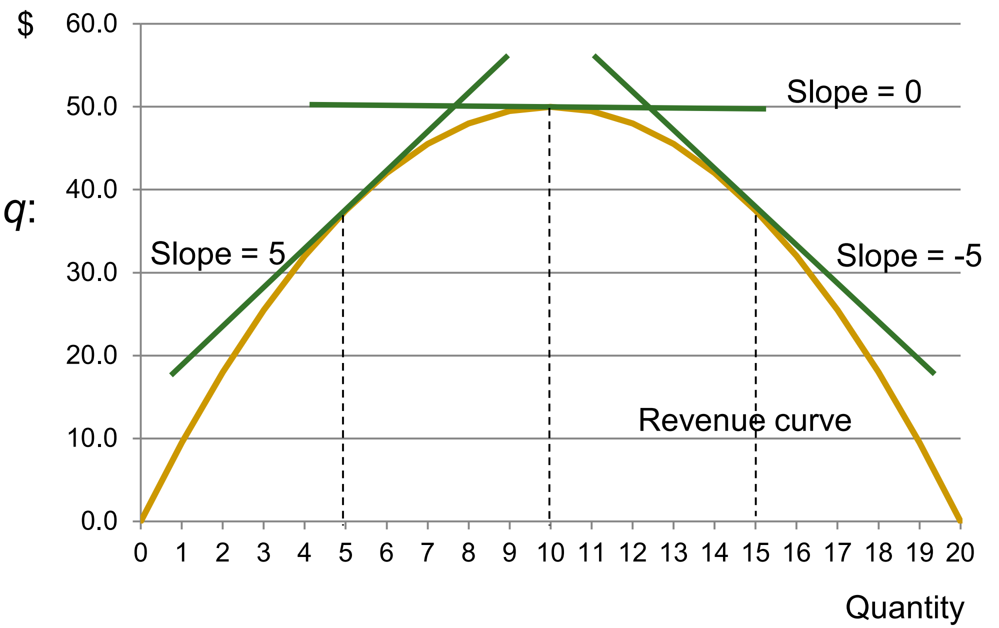

Graphing marginal revenue

- If the inverse demand curve is $p = a + bq$ (with $b < 0$).
  then marginal revenue is $MR = a + 2bq$
  - Graphically, the marginal revenue curve has the same y-intercept as the demand curve (a) and falls at twice the rate ($2b$):

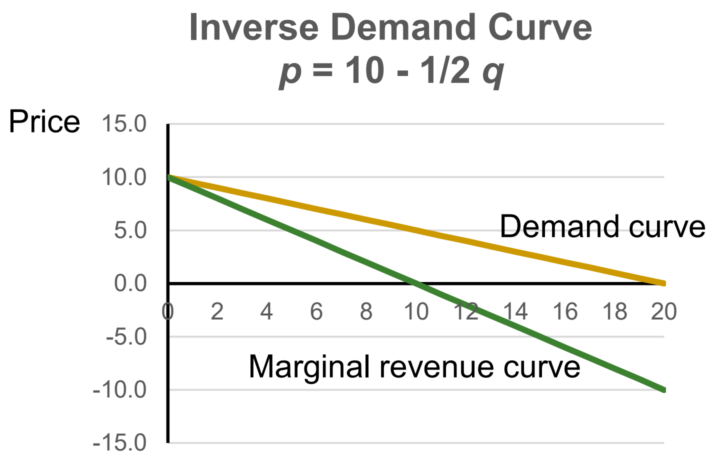

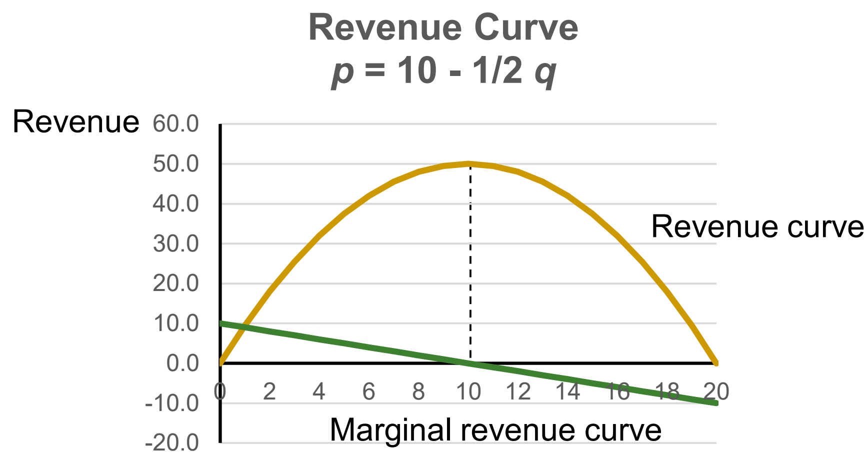

The MR curve crosses the x axis at 10, which is half of the quantity of demand when $p = 0$

The MR curve crosses the x axis at 10, which is the same quantity that maximizes revenue.

A mathematical aside (optional)

- Marginal revenue is the derivative of the revenue curve:

$$MR = \lim_{\Delta q \to 0} \frac{\Delta r}{\Delta q} = \frac{dr}{dq}$$

  - Heuristically, think of Δq approaching zero as 1 unit as market output becomes increasing larger
  - Note that the slopes of the secants $\Delta r/\Delta q$ approach the slope of the tangent line as Δq approaches zero

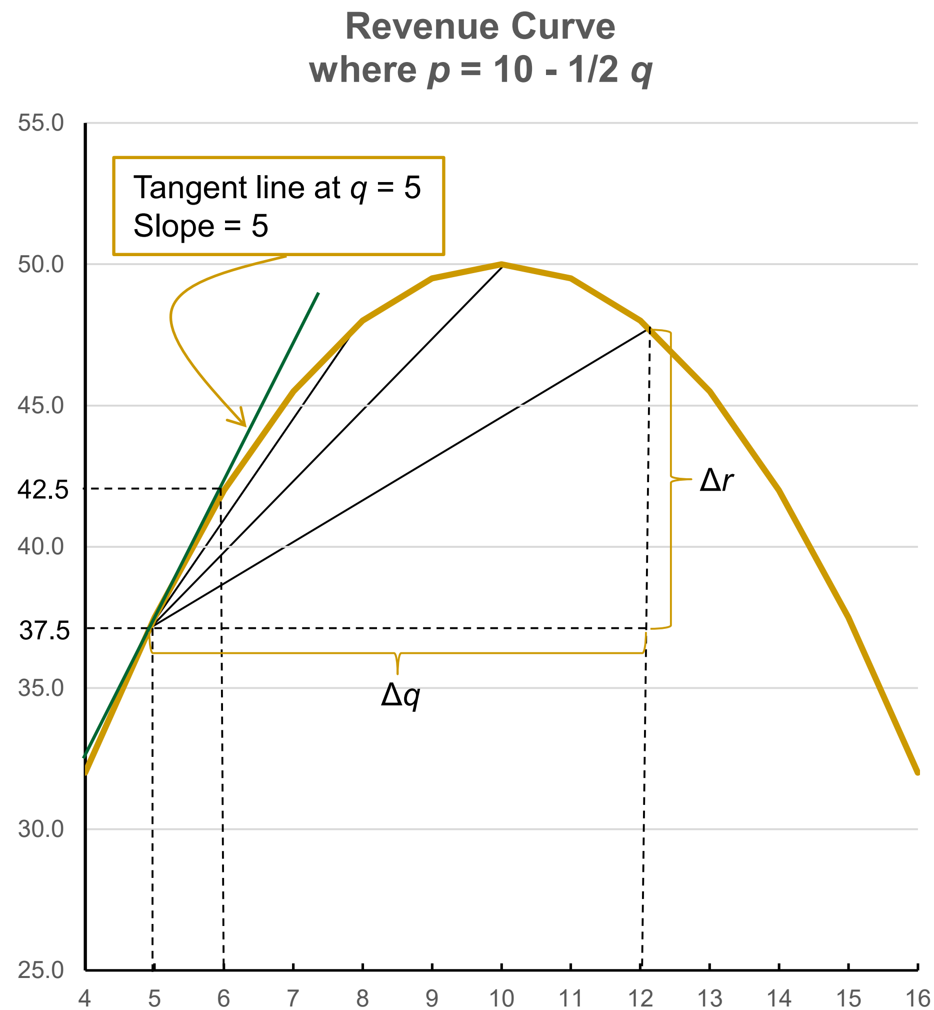

A mathematical aside (optional)

- Some derivative rules
  - The derivative of a constant is zero:

$$y = c \longrightarrow \frac{dy}{dx} = 0$$

  - The derivative of power function:

$$y = cx^n \longrightarrow \frac{dy}{dx} = nax^{n-1}$$

  - The derivative of a polynomial is the sum of the derivatives of the individual terms

Example: $y = a + bx + cx^2 \longrightarrow \frac{dy}{dx} = b + 2cx$

- First-order condition (FOC) for a (local) maximum point where $y = f(x)$:

$$\frac{dy}{dx} = 0$$

The slope of the function at the maximum point is zero. Put another way, you are at the top of the hill.

Production Costs

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

Costs

- Some definitions
  - *Total costs (C*(*q*))
    - The total cost of producing a production level *q*
    - Costs $C(q)$ = fixed cost (F) + variable cost $V(q)$)
  - *Fixed costs* (F)
    - Costs of production that do not vary with the quantity produced
  - *Variable costs* ($V(q)$)
    - Costs of production that vary with the production level and that are incurred producing a level q
  - *Average variable costs* ($AV(q)$)
    - Variable costs divided by quantity
  - *Marginal cost (MC*(*q*))
    - Technically, the slope of the cost curve at a quantity *q*
    - Heuristically, the additional costs the firm would incur for producing one additional unit having produced *q* units (i.e., $MC(q) \approx C(q+1) + C(q)$)

Marginal cost curve

Marginal cost curve: Traces the relationship between *q* and *MC*

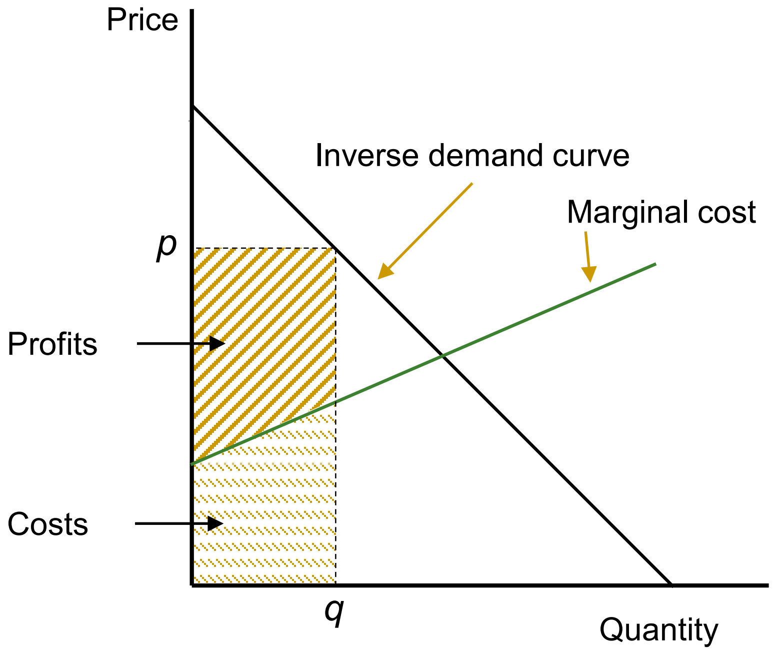

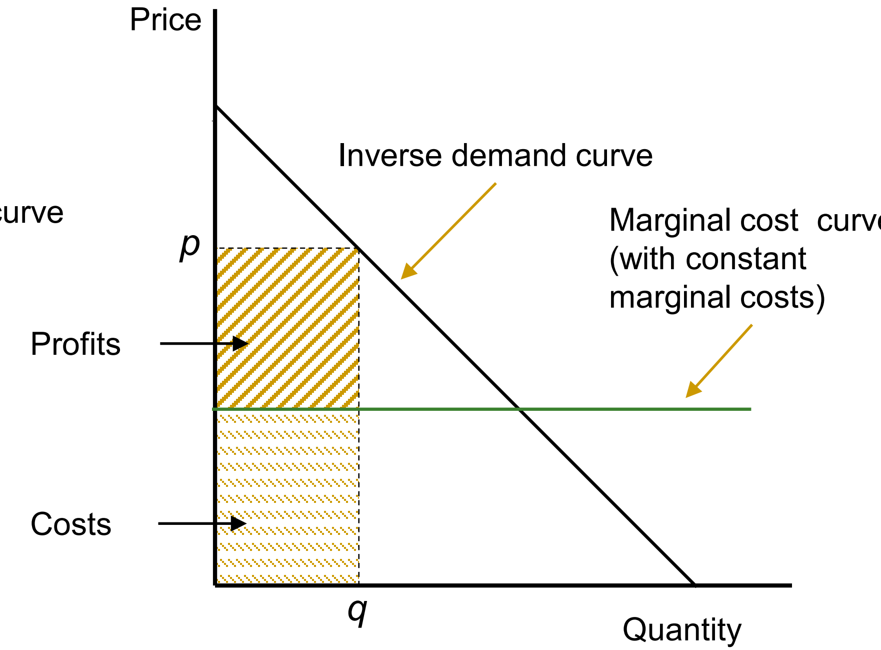

The marginal cost curve is in the first graph is upward sloping. Why might that be? The marginal cost curve in the second graph is flat, meaning that each incremental unit can be produced at the same marginal cost regardless of the level of production.

Profits and Profit Maximization

Profits

- Some basic terms
  - *Profits ($\pi(q)$)*
    - Revenues minus costs earned at a production level *q*:

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

  - *Marginal profit ($m\pi$)*
    - Technically, the slope of the profit curve at a quantity *q*
    - Heuristically, the net additional profit that the firm would make if it produced an additional unit
    - Or equivalently, marginal revenues minus marginal costs:

$$\begin{aligned} m\pi(q) &= \pi(q+1) - \pi(q) \\ &= \left[R(q+1) - C(q+1)\right] - \left[R(q) - C(q)\right] \\ &= \left[R(q+1) - R(q)\right] - \left[C(q+1) - C(q)\right] \\ &= mr(q) - mc(q) \end{aligned}$$

    - This also implies: $\pi(q+1) = \pi(q) + m\pi(q)$

That is, the firm's profits at $q + 1$ is the firm's profit at *q* plus the marginal profit (positive or negative) the firm would earn if it produced another unit

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 the slope of the profit function is zero, that is:

    $$m\pi(q^*) = mr(q^*) - mc(q^*) = 0$$

  - This implies that $mr = mc$ at the profit maximum

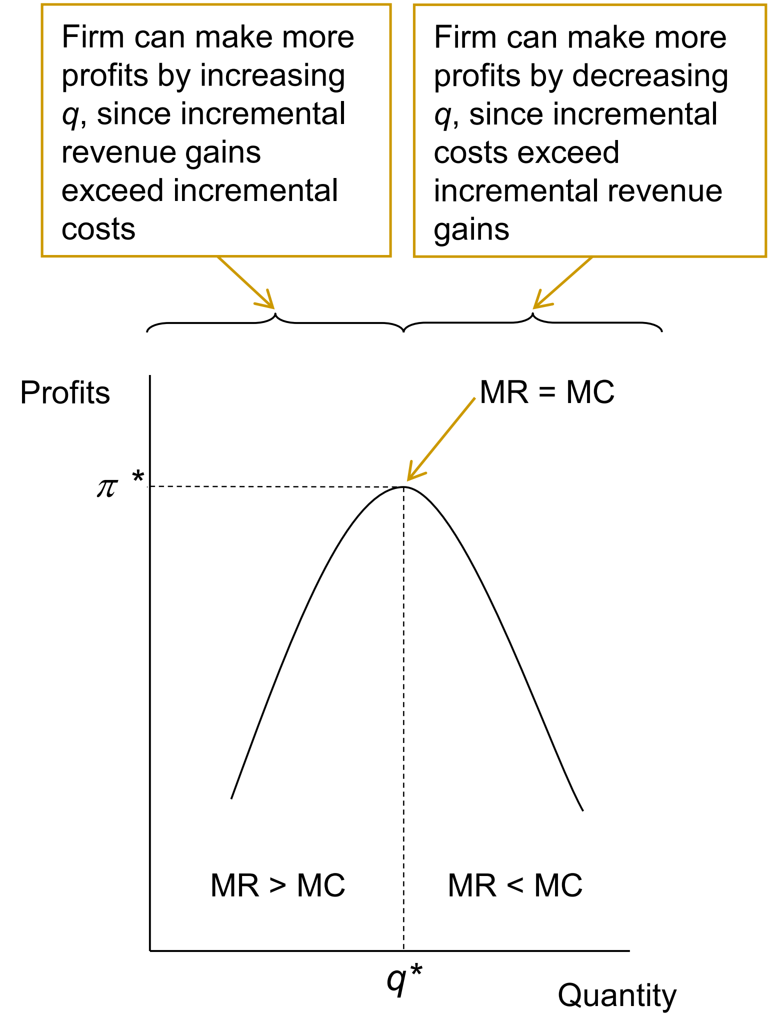

Profit maximization

- Determining the profit-maximizing quantity and price
  - Step 1: Find the q^* where marginal revenue equals marginal cost
  - Step 2: Find p^* for q^* from the inverse demand curve

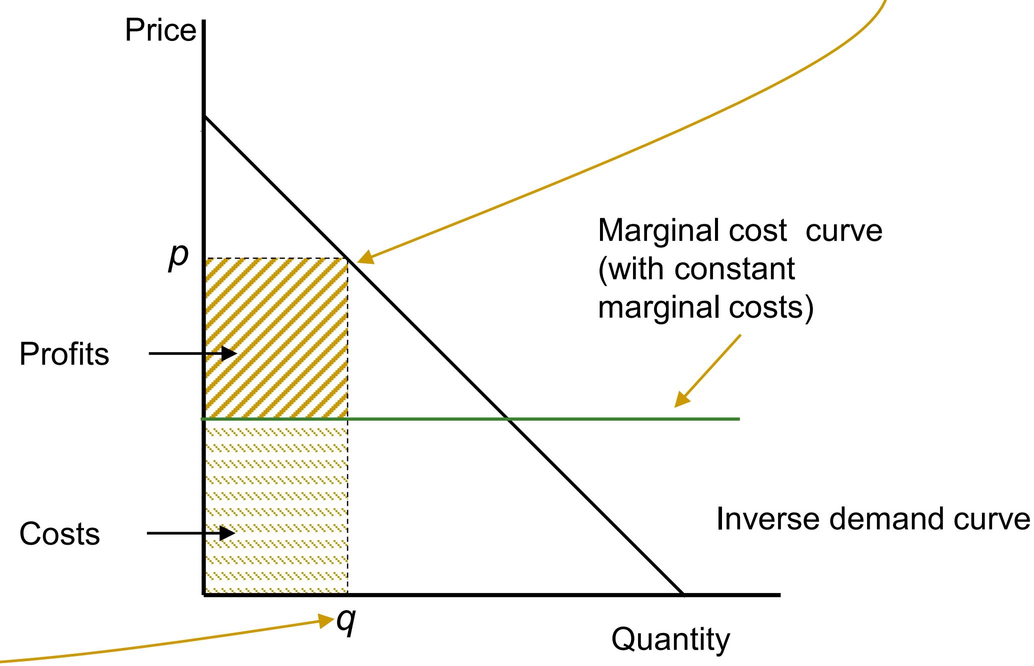

Competitive firms

- Competitive firms take prices as given
  - → Each individual firm perceives that its output decision does not affect the market-clearing price
  - This means that the firm acts as if $MR = p_c$

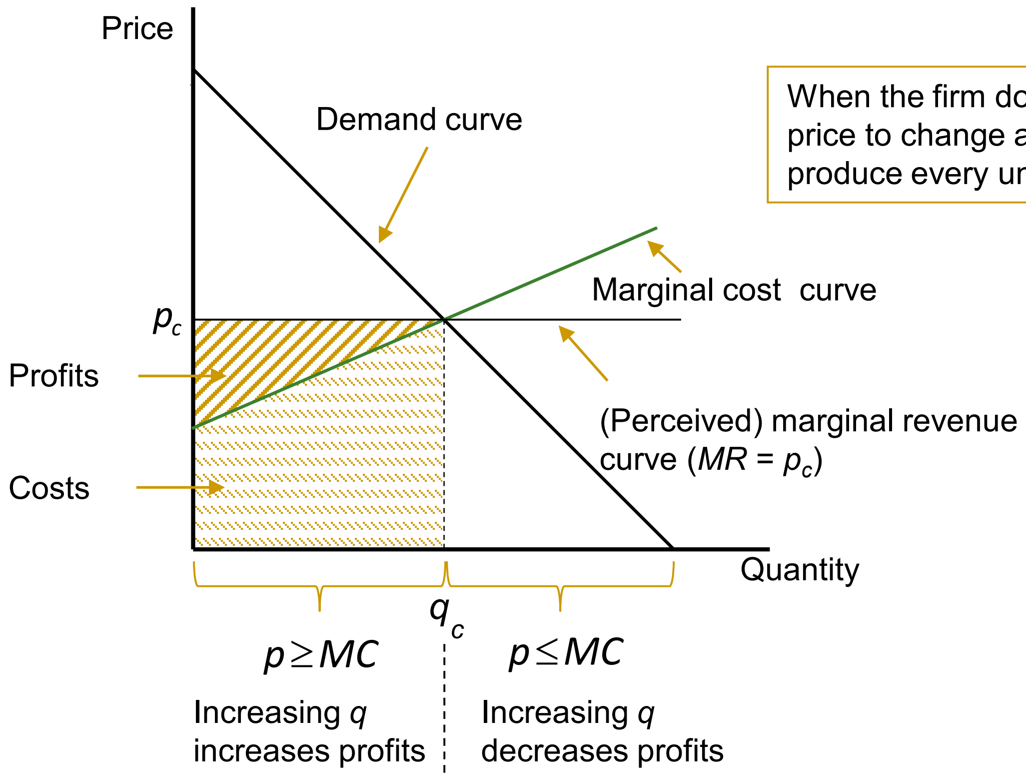

When the firm does not expect the market-clearing price to change as the firm expands output, the firm will produce every unit for which $p \geq MC$

Rule: As always, the FOC is $MR = MC$. If the firm is competitive, then $MR = p_c$ and so FOC is $p_c = MC$.

Profit maximization

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

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

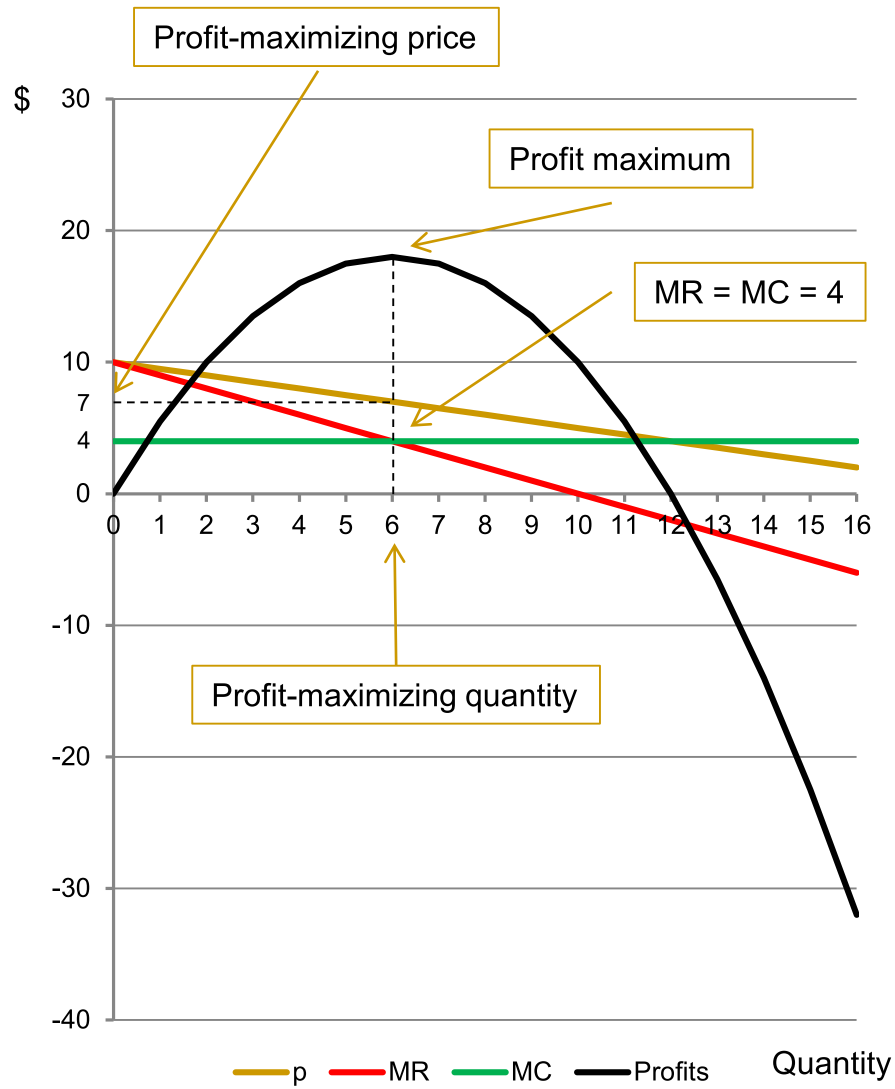

Profit maximization

- The mechanics

1. Hypothetical gives demand and marginal cost

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

Marginal cost: $mc = 4$

2. Calculate :

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

(from marginal revenue formula: 2 times the slope of the inverse demand function)

3. Solve for profit maximum

Set $mr = mc$: $10 - q^* = 4$

Solving: $q^* = 6$

$p^* = 7$

$\pi^* = 18$

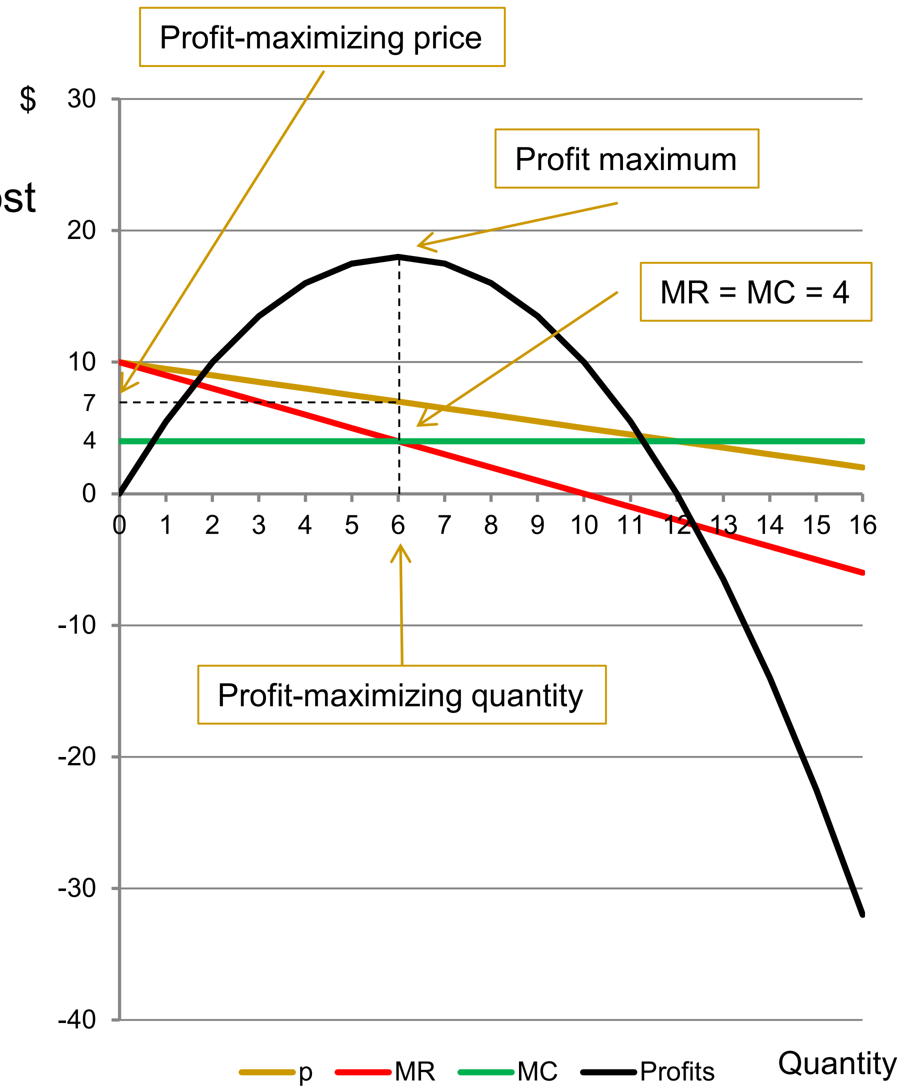

Profit maximization (optional)

- 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

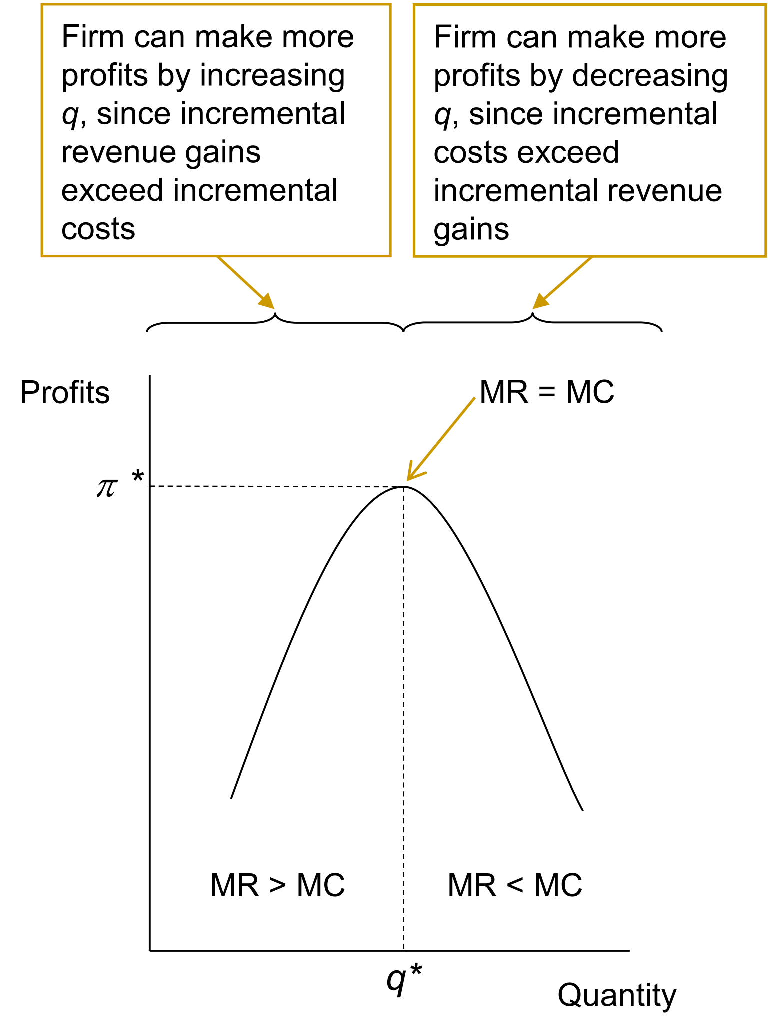

Profit maximization

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

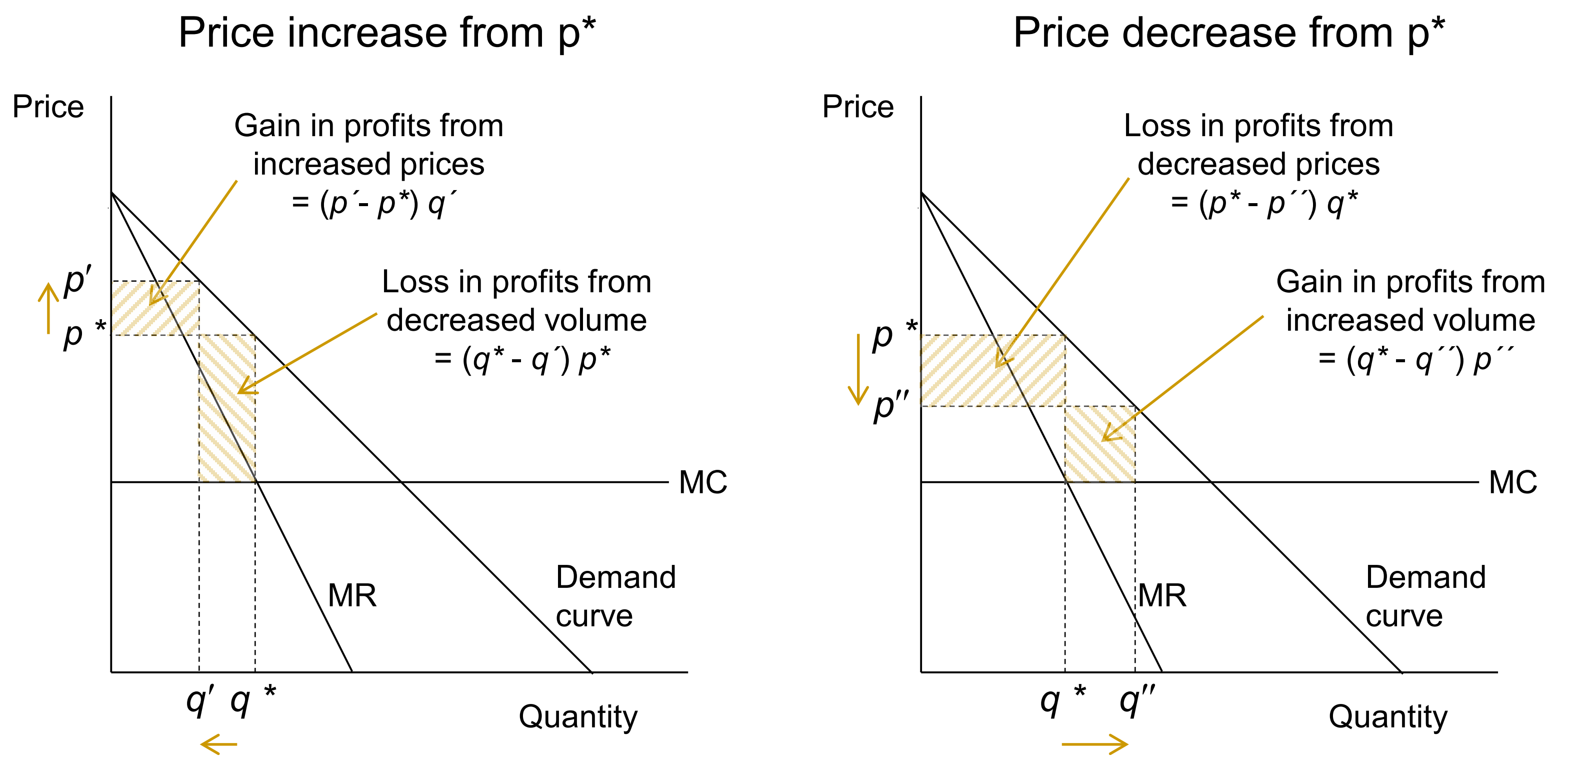

In each case, the loss from the price change exceeds the gain, so that moving away from p^* decreases profits.