Unit 7. Competition Economics

Part 1. Demand, Costs, and Profit Maximization

Merger Antitrust Law

Georgetown University Law Center
Professor Dale Collins

September 12, 2021

Motivation

- The purpose of merger antitrust law
  - Section 7 of the Clayton Act prohibits mergers and acquisitions whose effect "may be substantially to lessen competition, or to tend to create a monopoly"[1]
  - In modern terms, a transaction may substantially lessen competition when it threatens, with a reasonable probability, to create or facilitate the exercise of market power to the harm of consumers
  - Operationally, a transaction harms consumer when it result in—
    - Higher prices
    - Reduced market output
    - Reduced product or service quality in the market as a whole
    - Reduced rate of technological innovation or product improvement in the market

    compared to what would have been the case in the absence of the transaction (the "but for" world) and without any offsetting consumer benefits

Merger antitrust analysis typically focuses on price effects (see Unit 2)

Consequently, a central focus in merger antitrust law is the effect a merger is likely to have on the profit-maximizing incentives and ability of the merged firm to raise price in the wake of the transaction. In the first instance, this requires us to know how a profit-maximizing firm operates. The basic tools to enable us to do this analysis is the subject of this unit. These same tools are also fundamental to an understanding of merger antitrust law defenses.

[1] 15 U.S.C. § 18.

What you should be able to do after Part 1

For a firm—

- Facing a downward sloping residual (inverse) demand curve $p = a + bq$
- With fixed costs f and constant marginal costs c

1. Determine and graph the profit-maximizing levels of—
   - Output q^*
   - Price p^*
   - Profits π^*[1]
2. Determine and graph the net incremental revenue for a firm increasing output by some amount Δq, including—
   - The gross gain in revenues from the increase in output, and
   - The gross loss in revenues from the reduction of price for sales at the original price
3. Derive and graph an inverse demand curve given a demand curve

[1] The "*" indicates that the variable is at its profit-maximizing level.

The Basic Idea

Profits, revenues, prices, quantities and costs

- *Fundamental assumption #1*: Firms maximize their profits
  - *Profits* are equal to revenues minus costs:

$$\text{Profits } (\pi) = \text{Revenues } (r) \text{ minus total cost } (t)$$
$$\pi = r - t$$

  - *Revenues* are equal to price times the quantity sold[1]

$$\text{Revenues } (r) = \text{Price } (p) \text{ times quantity sold } (q)$$
$$r = pq$$

- *Fundamental assumption #2*: Price and quantity sold are inversely related through a *downward-sloping demand curve*
  - *Corollary*: The lower the price, the greater the quantity sold
  - We usually assume that a firm decides on the quantity it will produce and that the price is given by the demand curve
  - *Notation*: To show that price depends on the quantity produced, we can write $p(q)$ (that is, p "is a function of" q)

Note: q is called an *argument* of p

[1] Throughout the course unless otherwise noted, we will assume that all units of a homogeneous (identical) product sell at the same single price. That is, there is no price discrimination.

Profits, revenues, prices, quantities and costs

- Total cost is equal to the fixed cost plus the variable cost
  - *Fixed costs* do not vary with the quantity produced by the firm (e.g., plant and equipment)
  - *Variable cost* varies with the quantity produced by the firm (e.g., the costs of raw materials)
  - *Total cost*:

$$\text{Total cost } (t) = \text{fixed cost } (f) \text{ plus variable cost } (v)$$
$$t(q) = f + v(q)$$

This notation makes clear that t and v depend on the quantity produced, while f does not

- *Marginal cost* (c) is the cost of producing one additional unit
- When marginal cost is constant for all levels of production ("constant marginal costs"), variable cost equals marginal cost times quantity produced

$$\text{Variable cost } (v) = \text{marginal cost } (c) \text{ times quantity produced } (q)$$
$$v(q) = cq$$

Profits, revenues, prices, quantities and costs

- Putting it all together

Profits (π) = Revenues (r) minus total cost (t)

= Price (p) times quantity (q) minus fixed cost (f) minus variable cost (v)

So:
$$\pi(q) = pq(q) - f - v(q)$$
$$= pq(q) - \left(f + v(q)\right)$$

- Where marginal costs are constant:

$$\pi(q) = pq(q) - f - cq \text{ (when marginal costs are constant)}$$
$$= \left(p(q) - c\right)q - f$$

Note that this is the *dollar gross margin* (*p-c*) times quantity sold minus the fixed costs *f*

Consumers and Demand Curves

Quick introduction: Demand

- The basics
  - At a given price, customers are willing to buy a certain amount of the product
  - "*The law of demand*": The lower the price, the greater the quantity customers are willing to buy
    - *Example*: The lower the price of jelly beans, the more jelly beans customers will buy
    - Conversely, the higher the price, the less consumers will buy
  - *Demand function*: Gives the quantity customers are willing to buy at a given price
    - *Demand curve*: This is the graph of the demand function
    - The law of demand implies that demand curves are downward-sloping
  - *Inverse demand function*: Gives the price that clears the markets for a given level of production (i.e., customers want to buy no more and no less than the quantity produced)
    - The inverse demand function can be derived from the demand function
    - *Example*: If $q = 20 - 2p$ is the demand function, then $p = 10 - 0.5q$ is the demand function (use simple algebra to rearrange terms)

Quick introduction: Demand

- *Example*: Say the *demand curve* for blue widgets is $q = 20 - p$
  - This means that the *inverse demand curve* for blue widgets is $p = 20 - q$

Quantity *q*	Price *p*
0	20
1	19
2	18
3	17
4	16
5	15
6	14
7	13
8	12
9	11
10	10
11	9
12	8
13	7
14	6
15	5
16	4
17	3
18	2
19	1
20	0

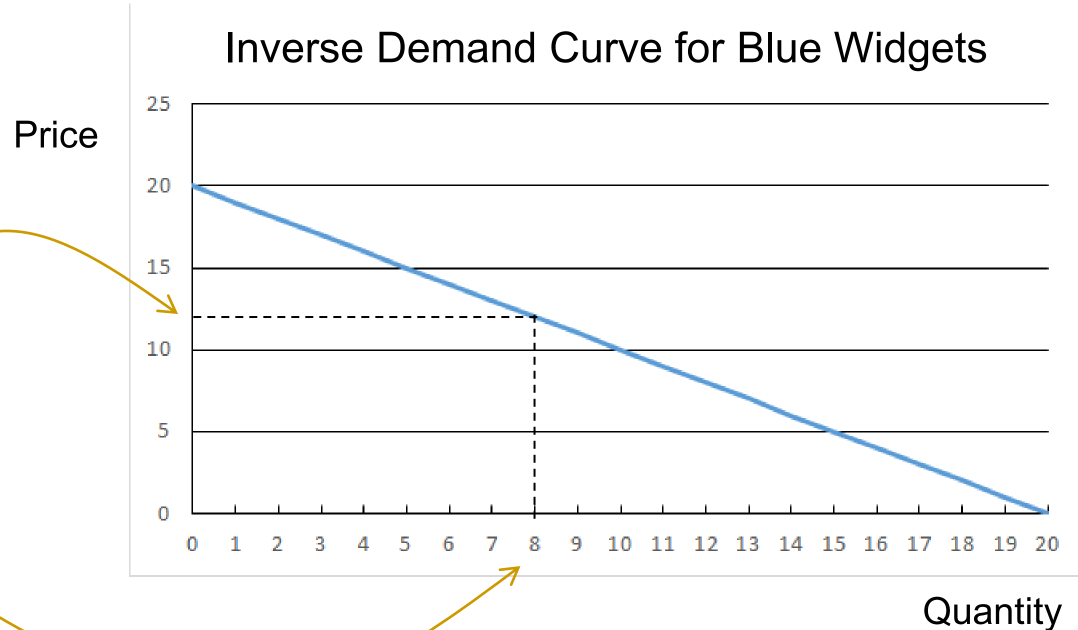

A production level of 8 units has a market-clearing price of $12

Quick introduction: Revenues

- Revenues (*r*) are the quantity sold (*q*) times the price (*p*): *r* = *pq*

Quantity *q*	Price *p*	Revenue *r = p*q*
0	20	0
1	19	19
2	18	36
3	17	51
4	16	64
5	15	75
6	14	84
7	13	91
8	12	96
9	11	99
10	10	100
11	9	99
12	8	96
13	7	91
14	6	84
15	5	75
16	4	64
17	3	51
18	2	36
19	1	19
20	0	0

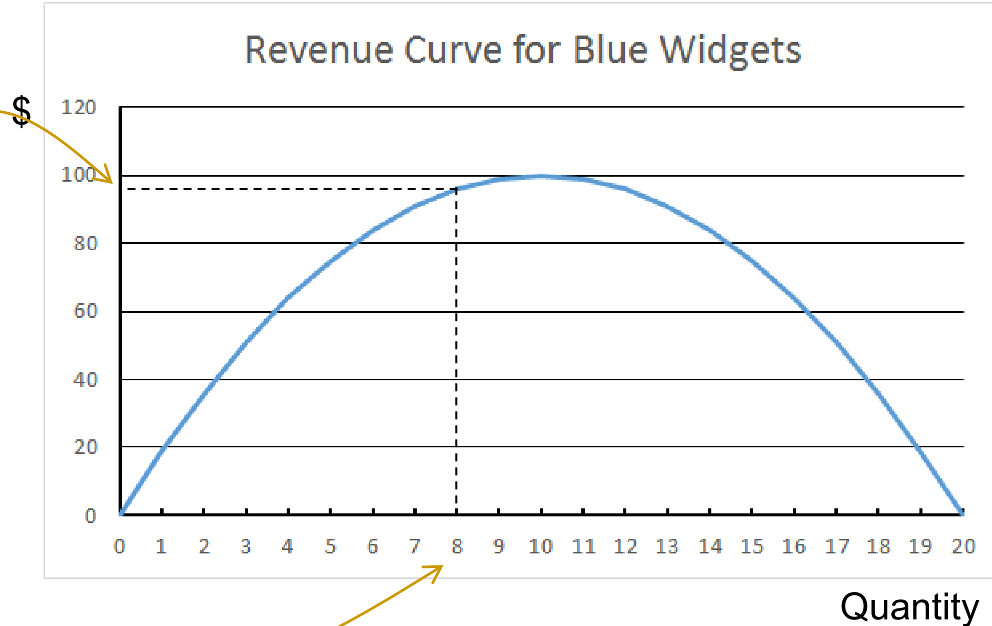

At a production level of 8 units, total revenues are $96

Quick introduction: Marginal revenue

- Marginal revenue (*mr*) is the incremental amount of revenue that could be earned by producing one more unit of output

Quantity *q*	Price *p*	Revenue *r*	Marginal revenue *mr*
0	20	0	19
1	19	19	17
2	18	36	15
3	17	51	13
4	16	64	11
5	15	75	9
6	14	84	7
7	13	91	5
8	12	96	3
9	11	99	1
10	10	100	-1
11	9	99	-3
12	8	96	-5
13	7	91	-7
14	6	84	-9
15	5	75	-11
16	4	64	-13
17	3	51	-15
18	2	36	-17
19	1	19	-19
20	0	0	

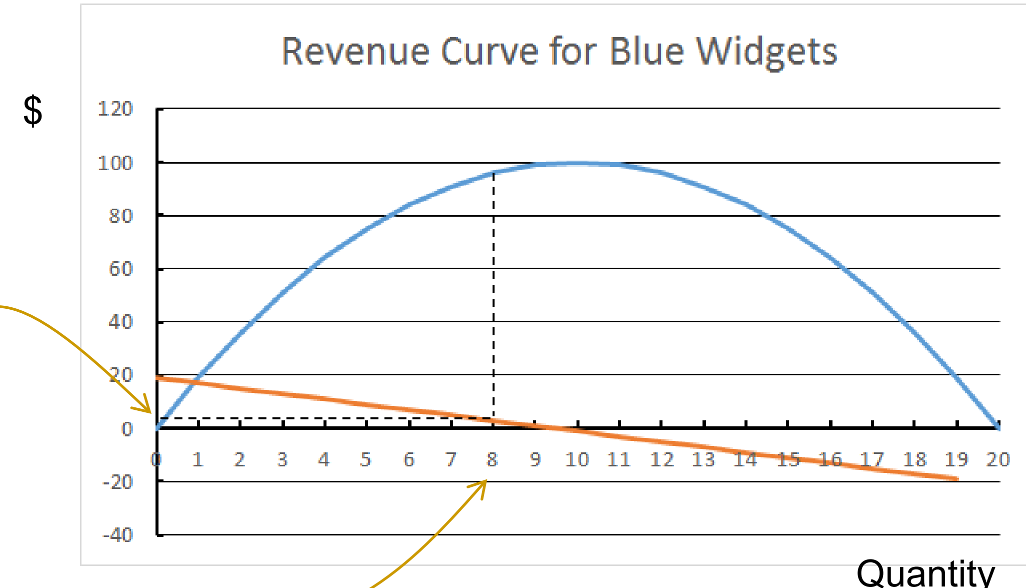

The marginal revenue of increasing sales from 8 units to 9 units is $3 (as revenues go up from $96 to $99)

Note: Marginal revenue can be negative (as the chart shows)

Quick introduction: Marginal revenue

- Query: Why is marginal revenue only $3?
  - The marketing-clearing price for 9 units is $11. Why isn't marginal revenue simply the price of the additional sale ($11)?
  - *Answer*:
    - All blue widgets in the market sell at the same price (no price discrimination), so when the market-clearing price dropped from $12 to $11 to sell the additional unit, the price of the 8 units that customers were willing to buy at $12 (the *inframarginal sales*) also had to drop to $11
    - That price drop of $1 on each of the 8 units causes a loss of $8
    - So the gain of revenue from the additional sale ($11) minus the loss of revenue from the $1 price drop on the eight inframarginal sales ($8) yields a marginal revenue of $3

This concept is critical to most of the economics in this course.
Be sure you understand it.

Quick introduction: Marginal revenue

- We can see this graphically by looking at the demand curve for blue widgets

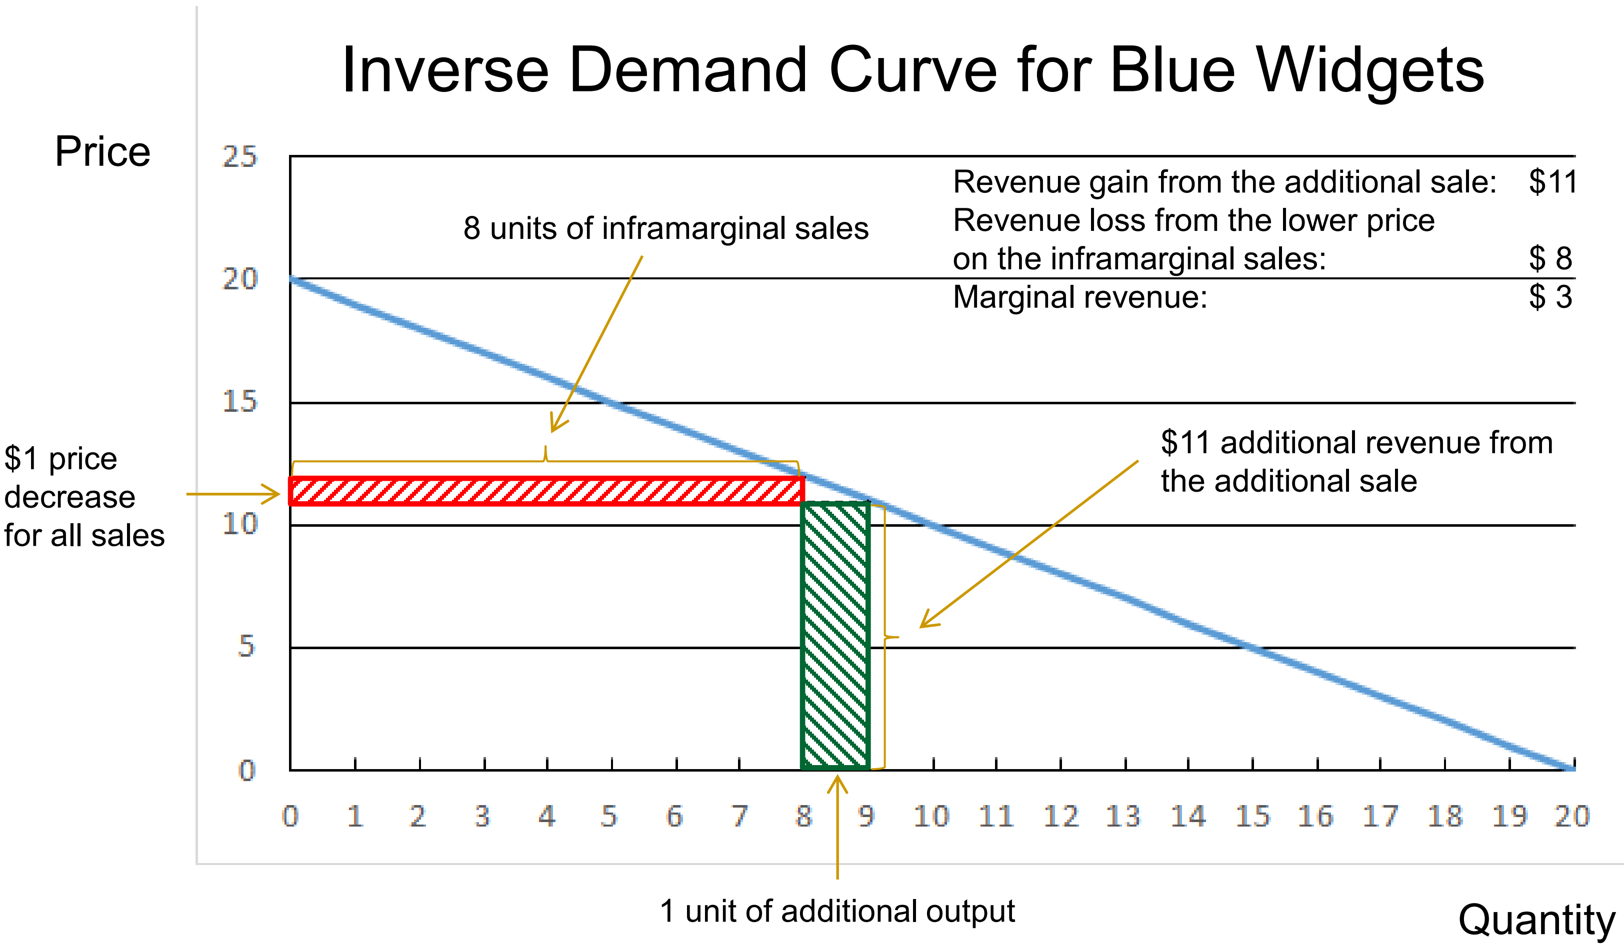

Demand curves

- Demand curves
  - Demand curves are fundamental to an understanding of merger antitrust law
  - A *demand curve* gives quantity consumers will purchase as a function of 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 function* 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 of* prices and production levels
      - 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., $b < 0$)
      - The collection of all of the pairs (p_i, q_i) trace out the *demand curve*

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
    - If we know two points on a linear demand curve, we can derive the demand function

      A linear demand function has the form *q* = *a* + *bp*, where b is the slope of the demand curve

      So:

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

Substituting *b* = -2 into the general function:

$$q = a - 2p$$

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: $q = 20 - 2p$

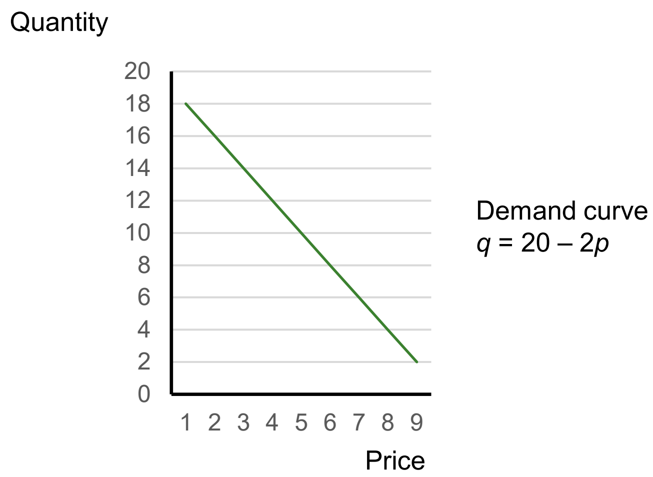

Demand curves

- A more detailed diagram

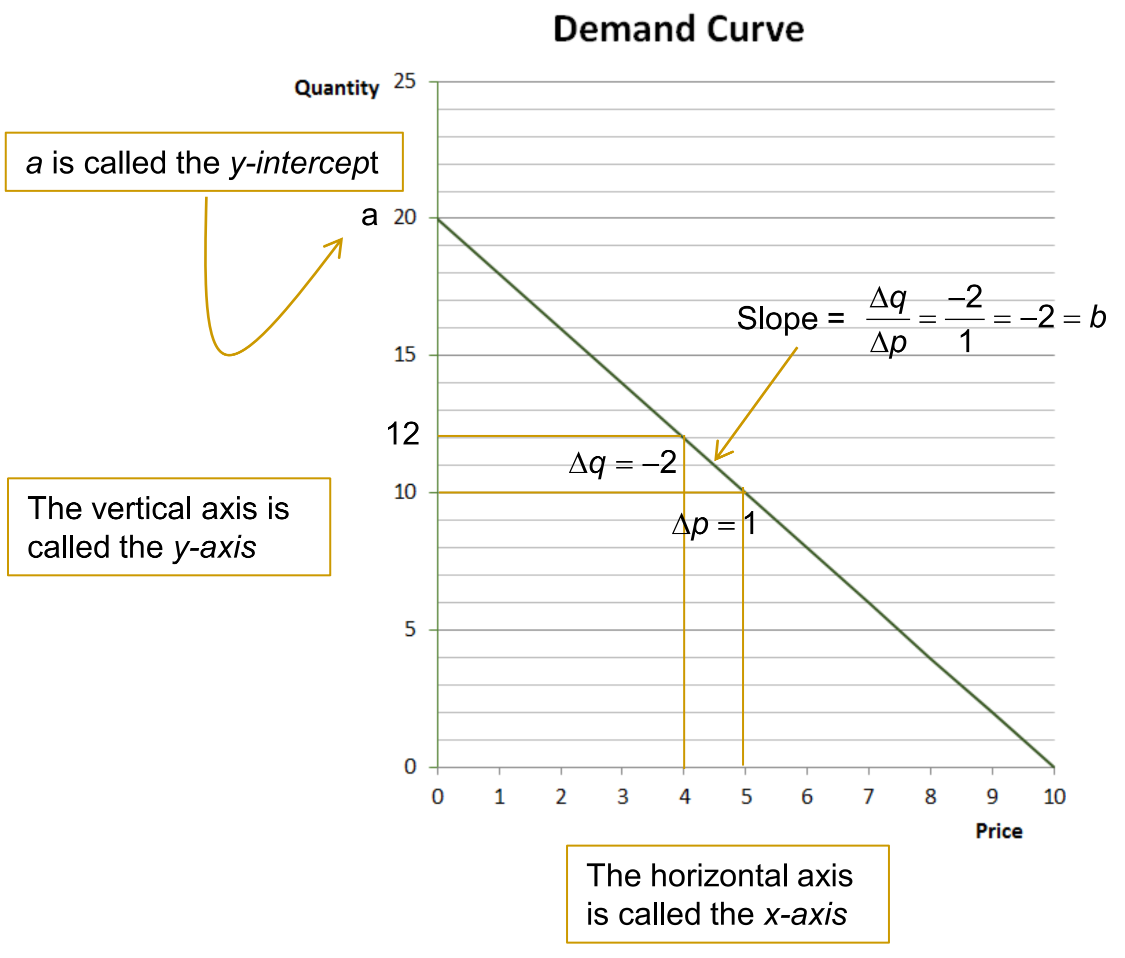

Demand curve: $q = 20 - 2p$

General form: $q = a + bp$

So

$a = 20$

$b = -2$

Inverse demand curves

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

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

  - *Key*: 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
    - Just the opposite of the demand curve

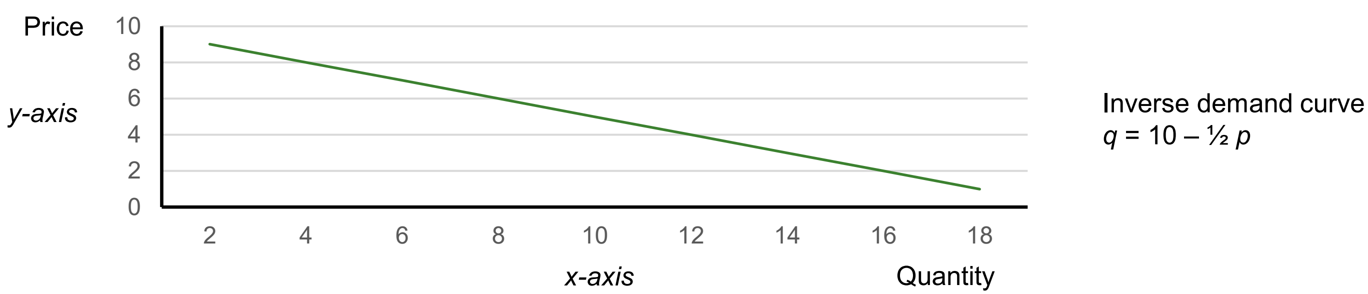

Demand curves

- Linear inverse demand curve for the demand curve $q = a + bq$

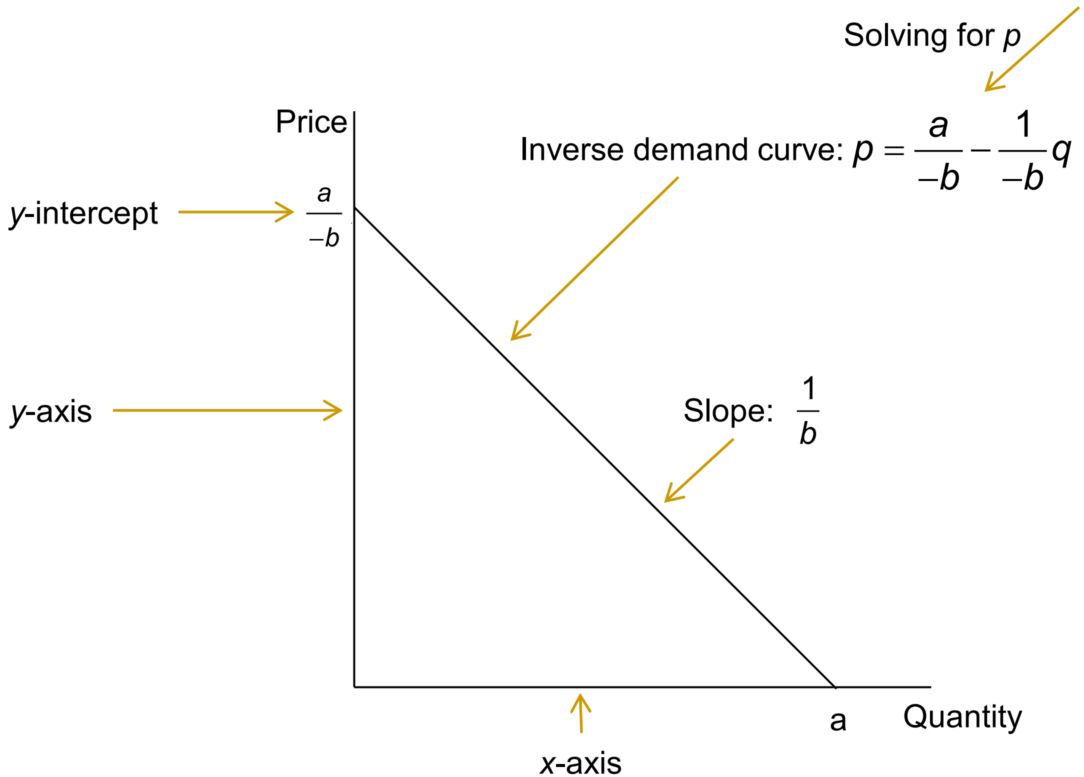

The slope of the demand curve gives the required change in the price to sell one additional unit of the product. So the price needs to drop by -1/*b* to sell one additional unit:.

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

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

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

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.

Nonlinear demand curves

- Demand curves do not need to be linear
  - *Example*: Nonlinear inverse demand curve with no *x*-axis intercept

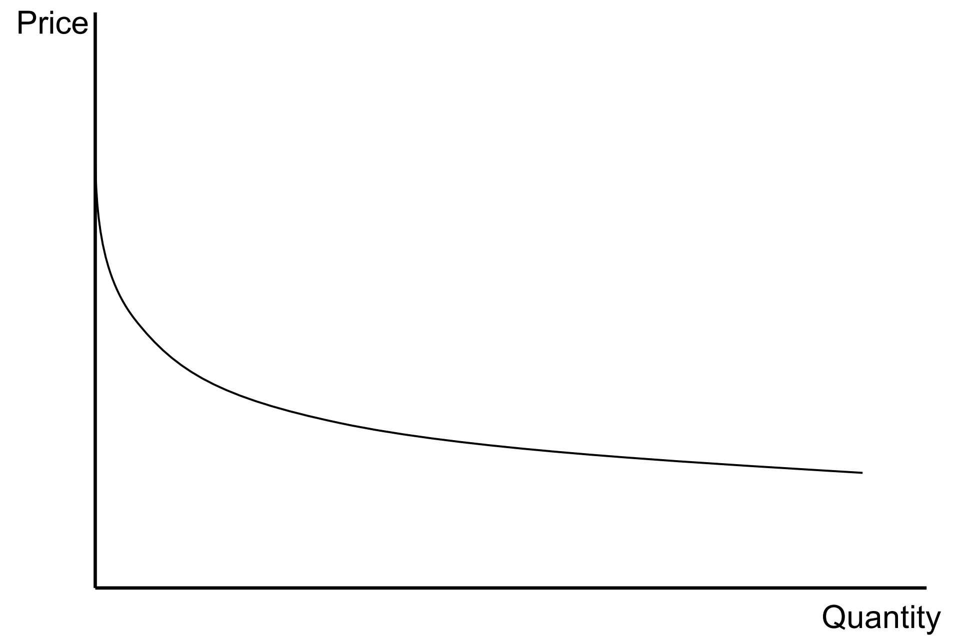

Nonlinear demand curves

- Around a single point, a nonlinear demand curve may be approximated by a linear demand curve

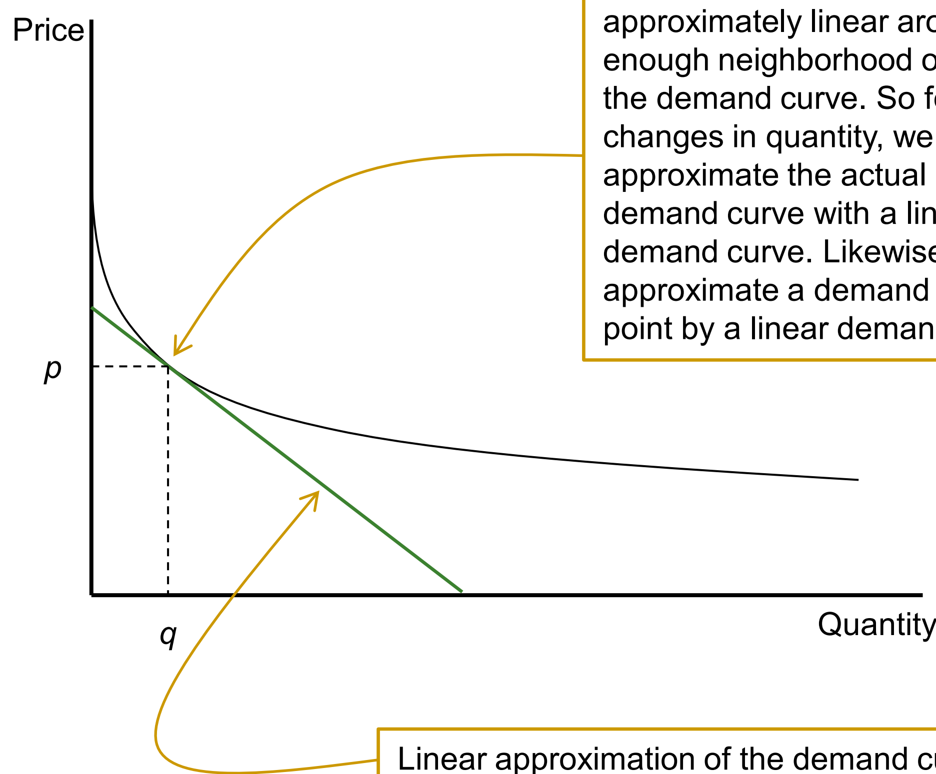

Note: Demand curves are approximately linear around any small enough neighborhood of any point on the demand curve. So for small changes in quantity, we can approximate the actual inverse demand curve with a linear inverse demand curve. Likewise, we can approximate a demand curves at a point by a linear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

Linear approximation of the demand curve at (*p*,*q*)

Aggregate consumer demand

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

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

$\sum_i q_i(p)$ means add the quantity demanded by consumer *i* at price *p* across all consumers

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

- Example

	Demand at $p = 4$	Demand at $p = 6$
Consumer 1	5	3
Consumer 2	3	0
Consumer 3	6	5
Aggregate demand	14	8

  - So $q(4) = 14$ and $q(6) = 8$
  - Consumer demands are independent of one another, but the aggregate demand is always calculated using the same market price for all consumers

Some technical points

- 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:
    - *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
- 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
- Total demand in a market for a product is given by the aggregate demand curve, that is, the sum of demands by consumers of all firms in the marketplace for a given market price
  - 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
    - You will encounter this concept frequently as the course proceeds

This is an important concept

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*)

    You need more information to determine which case applies

Two points on the same demand curve

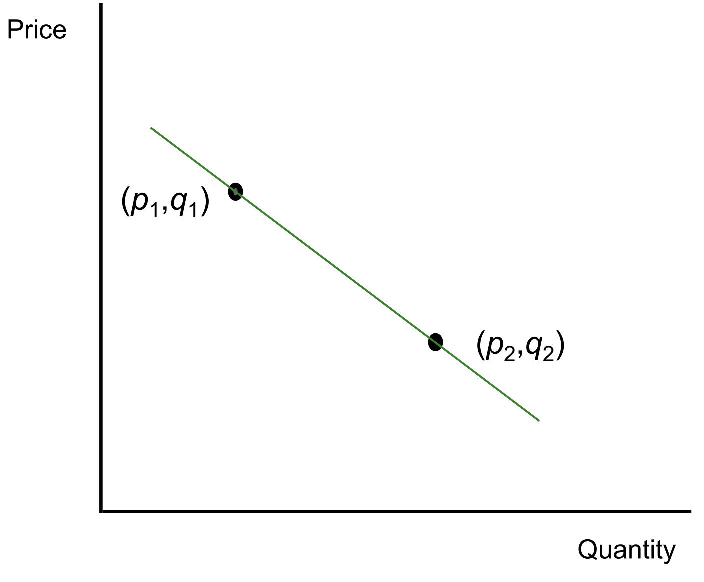

Two points on different demand curves

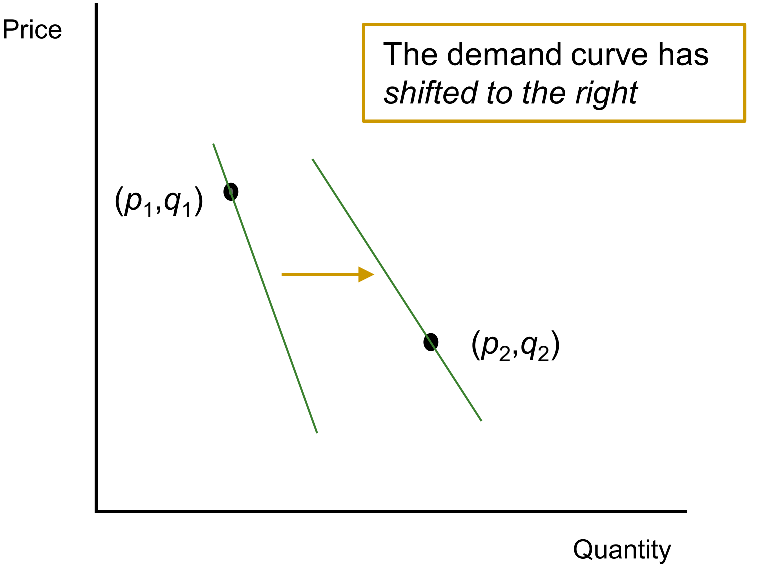

Producers and Revenues

Producers

- *Recall fundamental assumption #1*: Firms maximize their profits
  - Profits (π) = Revenues (r) – Total costs (t)
- To analyze the conditions under which a firm maximizes it profit, you need to look at:
  - Revenues and revenue functions
  - Costs and cost functions
  - The relationship between revenues and costs when the firm maximizes its profit

Key result: We will see that a firm maximizes its profit when the quantity it produces will set its marginal revenue equal to its marginal cost

Expect to see this proposition daily throughout the remainder of the course

Revenues

Revenue = *p* times *q* (= *pq*)

This is just the area of the rectangle in the chart below

The rectangle has a height of *p* and with of *q*

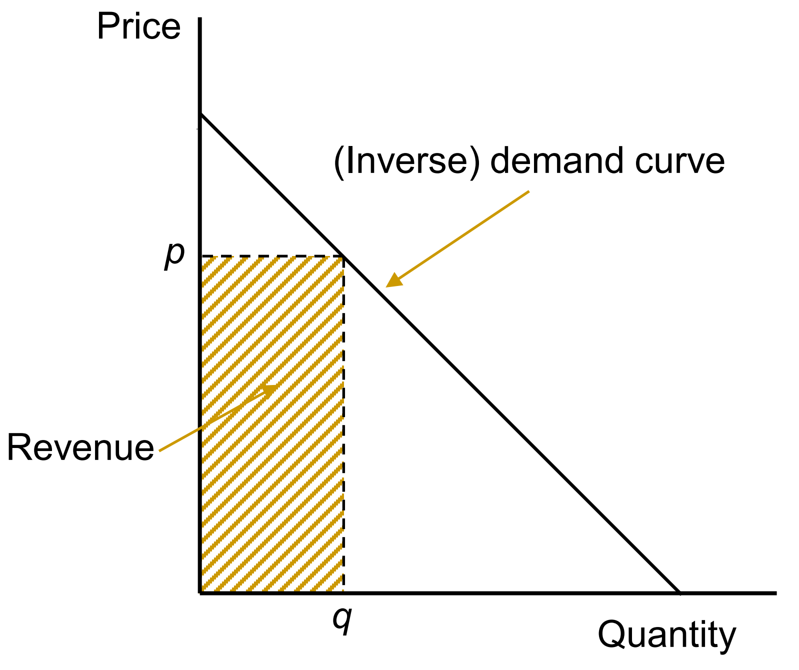

Marginal revenue

- *Definition*: *Marginal revenue* is the net additional revenue the firms earns by producing and selling one unit of *additional* output
  - *NB*: If the firm faces a downward sloping demand curve, marginal revenue will be less than price—the market price p_1 will have to decrease to p_2 after adding the incremental output in order to clear the market
    - This lower price will apply to preexisting sales as well as incremental sales
- Example
  - Say Bob's Widgets sells 4 widgets at $8 each. If Bob's wants to increase its sales to 5 widgets, Bob's has to drop its price to $7.5 per widget. What is Bob's marginal revenue at the current sales of 4 widgets?
    - Brute force calculation
      - Current revenue: 4 units at $8 each equals $32 in total revenue
      - Revenue if sales were increased by one unit: 5 units at $7.5 each for $37.5 in total revenue
      - Marginal revenue: Revenue with one additional unit minus current revenue, which is $37.5 - $32 = $5.5
    - More nuanced calculation
      - Revenue gained on the sale of one more unit at the *new* lower price: Purchase price of $7.5
      - Revenue lost on the sale of existing (inframarginal) units: Price drop of $0.50 times the existing units (4), which is $2
      - Marginal revenue: Revenue gained on the incremental sale minus the revenue lost of existing sales, which is $7.5 – $2 = $5.5

Marginal revenue

- Bob's brute force calculation illustrated:
  - Bob's faces a linear demand curve of
  - This yields an inverse demand curve: $q = 20 - 2p$

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

  - Initial conditions:
    - $q = 4$
    - $p = 8$ (from the inverse demand curve)
    - $r = qp = (4)(8) = 32$
  - Increased production by one unit
    - $q = 5$
    - $p = 7.5$ (from the inverse demand curve)
    - $r = qp = (5)(7.5) = 37.5$
  - So marginal revenue at $q = 4$ is 5.5:

$$\begin{aligned} mr(4) &= r(5) - r(4) \\ &= 37.5 - 32 \\ &= 5.5 \end{aligned}$$

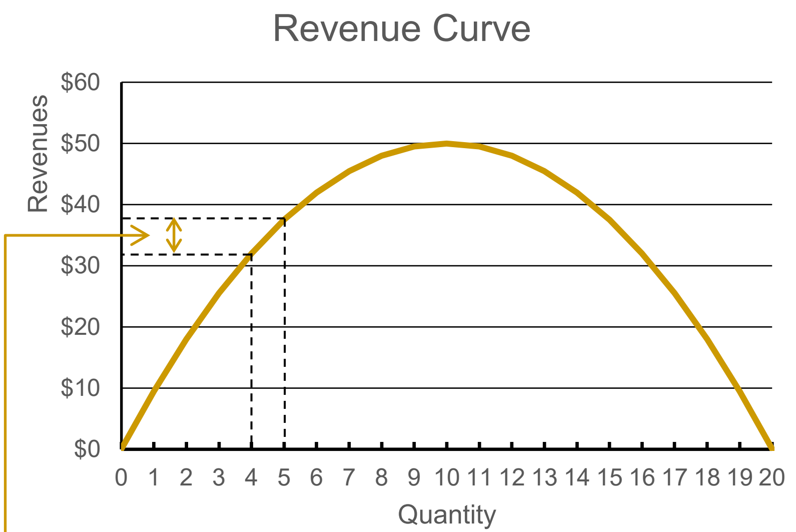

Marginal revenue

- Bob's more nuanced calculation illustrated:
  - Summarize the variables we know from the problem:

$$q_1 = 4 \qquad q_2 = 5 \qquad \Delta q = q_2 - q_1 = 1$$
$$p_1 = 8 \qquad p_2 = 7.5 \qquad \Delta p = p_2 - p_1 = -0.5$$

  - Revenue gained on the sale of one more unit: $\Delta q p_2 = (1)(7.5) = 7.5$ (Area A)
  - Revenue lost on the sale of existing units: $q_1 \Delta p = (4)(0.5) = 2$ (Area B)
  - Marginal revenue: Area A – Area B = 7.5 – 2 = 5.5

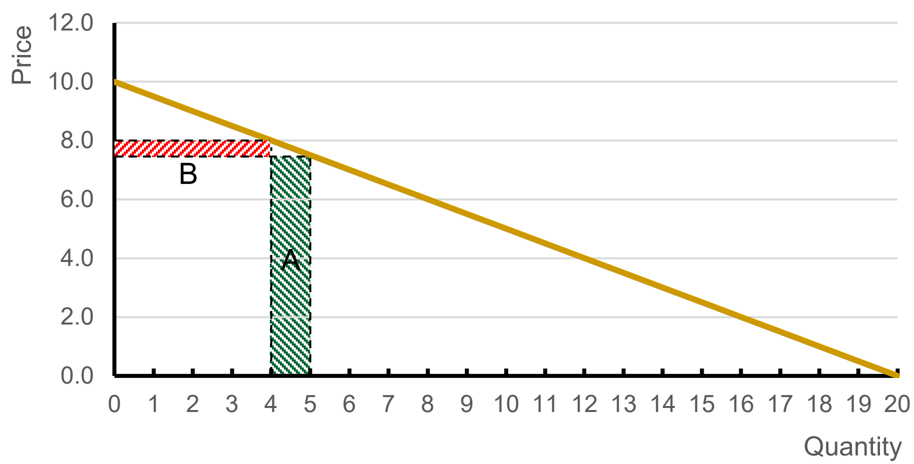

Marginal revenue

- Summary of Bob's two methods of calculating marginal revenue

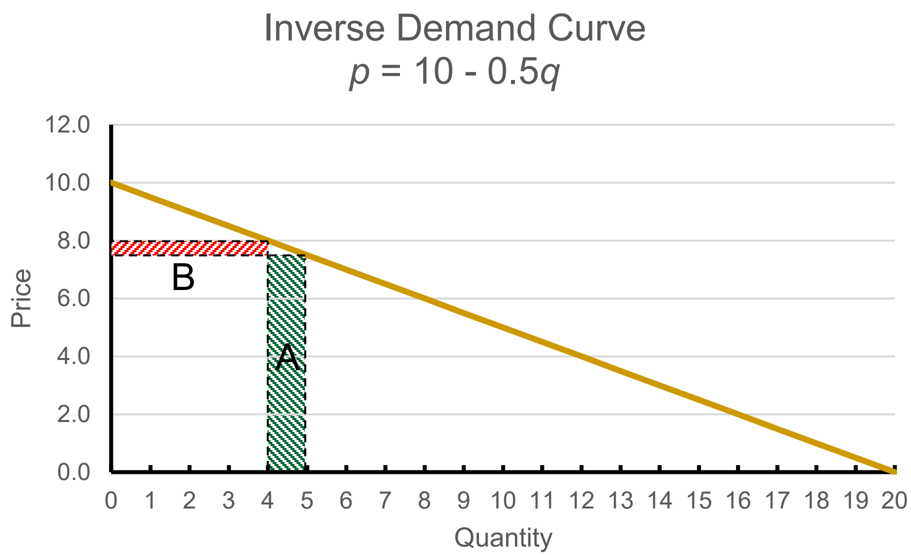

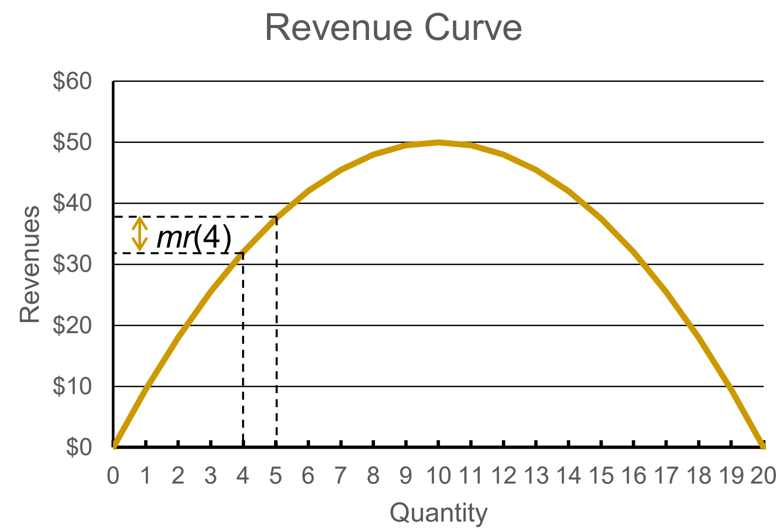

$$q = 4$$
$$p_1 = 8$$
$$p_2 = 7.5$$
$$\Delta p = 0.5$$

$$\text{Area A} = p_2 = 7.5$$
$$\text{Area B} = \Delta p q = 0.5 \times 4 = 2$$
$$mr(4) = \text{Area A - Area B}$$
$$= 7.5 - 2 = 5.5$$

$$p = 10 - 0.5q$$
$$r = 10q - 0.5q^2$$
$$r(5) = 50 - 12.5 = 37.5$$
$$r(4) = 40 - 8 = 32$$
$$mr(4) = r(5) - r(4) = 37.5 - 32 = 5.5$$

Marginal revenue

- Intuitions
  - While the brute force calculation will work, the more nuanced calculation will give you a much better feel for underlying economics
    - This will be invaluable later
- The general formulation of the more nuanced approach

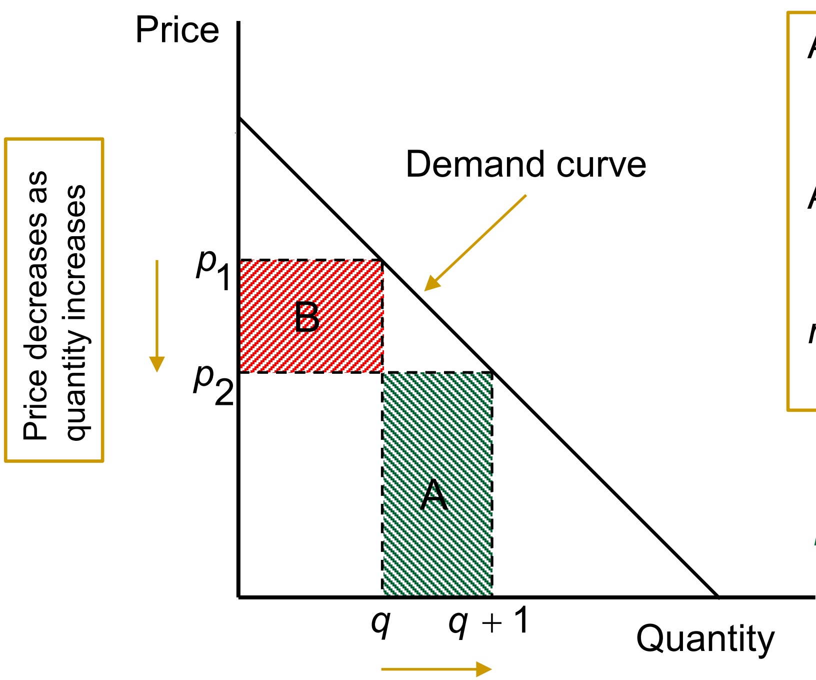

Area A = Gross revenue gain by selling an additional unit
= p_2 times 1

Area B = Gross revenue loss on preexisting sales due to the need to lower price
= $(p_2 - p_1)$ times $q = \Delta pq$

$mr(q)$ = Area A – Area B
= $p_2 - \Delta pq$

BE SURE YOU UNDERSTAND THIS PAGE!

Some technical notes

- Relationship between revenues and marginal revenue
  - *Discrete case* (where one unit of additional output is large compared to total production):

$$mr(q) = r(q+1) - r(q)$$

  - *For diehard calculus fans*: In the *continuous case* (where one unit of additional output is very small compared to total production), marginal revenue is the derivative of the revenue function (that is, the instantaneous rate of change of revenue for an infinitesimal change in output):

Geek note 1: If y is a function of x, so that $y = f(x)$, then the derivative of the function at a point x is simply the slope of the tangent line at that point or equivalently the instantaneous rate of change of the value y with respect to x.

$$mr(q) = \frac{dr(q)}{dq}$$

Geek note 2: the symbol dy/dx is the notation for the *derivative* of $f(x)$ with respect to x. To make this even clearer, we can write the derivative as $df(x)/dx$, as was done here with the revenue function. Think of dy/dx as $\Delta y/\Delta x$, where Δx is very small.

- Linear (inverse) demand curves
  - If $p(q) = a + bq$ is the inverse demand curve, then $r(q) = p(q)q = (a + bq)q = aq + bq^2$
  - *Rule*: Marginal revenue in the continuous case is $mr(q) = a + 2bq$ ← This is a key formula!!
    - If you know calculus, marginal revenue is the *derivative* of the revenue function $r(q)$
    - *If you do not know calculus, just memorize the formula!*

Some technical notes

- There is a difference in the value of marginal revenue in the discrete and continuous cases
- Consider our example in the continuous case
  - In our example, the revenue function is:

$$r(q) = p(q)q = \left(10 - \frac{1}{2}q\right)q = 10q - \frac{1}{2}q^2$$

  - So the marginal revenue function in the continuous case is:

$$mr(q) = \frac{dr}{dq} = 10 - q$$

Using the derivative formula on the prior slide

  - Marginal revenue at $q = 4$ in the continuous case is then 6, not the 5.5 in our discrete case
  - Moreover, inspection of the graph on the next page shows that the slope of the revenue curve is steeper at $q = 4$ than the arc on the revenue curve between $q = 4$ and $q = 5$, confirming that marginal revenue in the continuous case is larger than marginal revenue in the discrete case
    - You should confirm that the slope of the arc is 5.5, that is, the marginal revenue at $q = 4$ in the discrete case

Some technical notes

- Comparing marginal revenue in the discrete and continuous cases

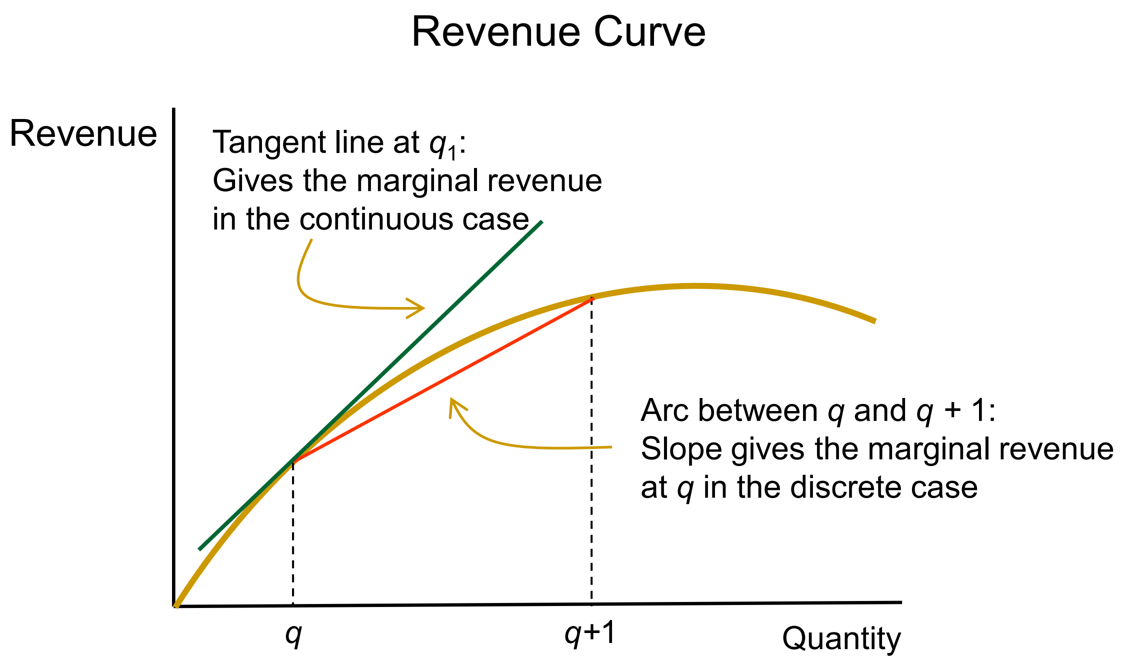

- We will use the continuous case in drawing marginal revenue curves in future slides
  - This will ensure that the profit-maximizing level of output is well-identified
  - Also, will approach marginal revenue in the discrete case as the q becomes large

Drawing demand and marginal revenue curves

- This is simple in the continuous case with linear demand
  - Say inverse demand curve is $p(q) = a + bq$
  - Then the marginal revenue curve is $mr(q) = a + 2bq$ (see Slide 33)
  - Important observations
    - Both the inverse demand curve and the marginal revenue curve intercept the *y*-axis at the same point *a*
    - The slope of the marginal revenue curve (2*b*) falls twice as fast as the slope of the inverse demand curve (b). This means that the marginal revenue curve will intercept the x-axis at half of the distance of the intercept of the inverse demand curve

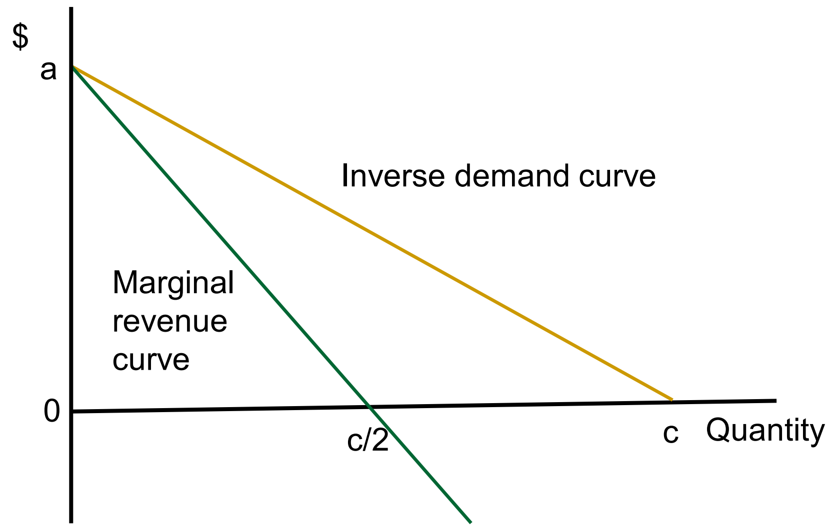

Maximizing revenue

- *Example:*
  - Demand curve: $q = 20 - 2p$
  - This yields an inverse demand curve: $p = 10 - \frac{1}{2}q$
  - Revenues:

$$r(q) = p(q)q$$
$$= \left[10 - \frac{1}{2}q\right]q$$
$$= 10q - \frac{1}{2}q^2.$$

This is a quadratic equation (because it contains a squared variable). Its curve is a *parabola.*

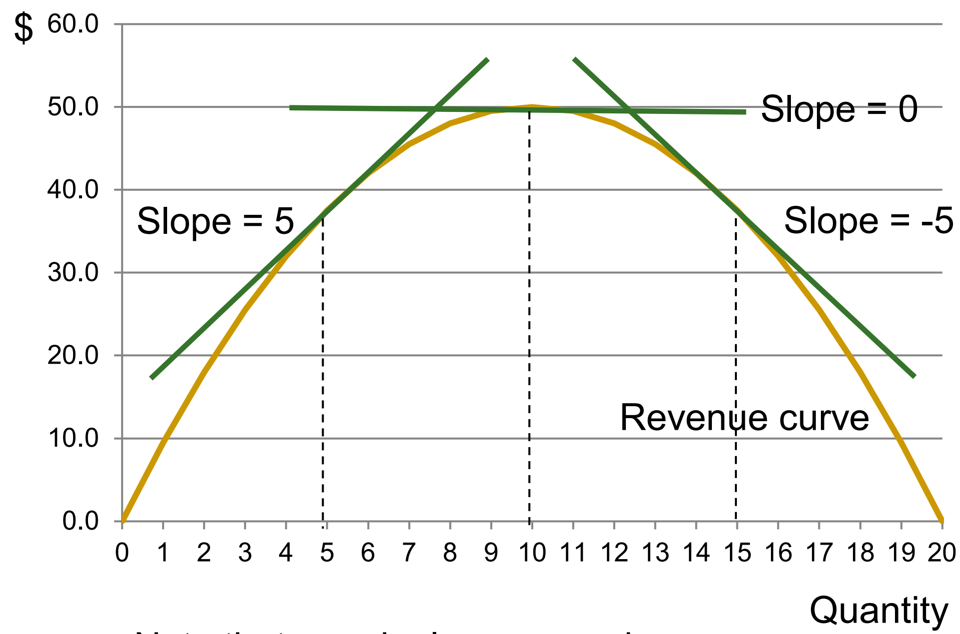

  - As you can see, revenue is maximized at the top of the "hill" where the slope is zero (that is, where q = 10).
    - Said another way, revenue is maximized at the production quantity where the derivative of the revenue function is zero
    - Note that marginal revenues decreases as production quantity increases. It drops to zero at the revenue maximum and then becomes negative.

Graphing revenue and marginal revenue curves

- *Example*:

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

Revenues: $r = pq$

Marginal revenue $mr = 10 - q$ *(continuous case)*

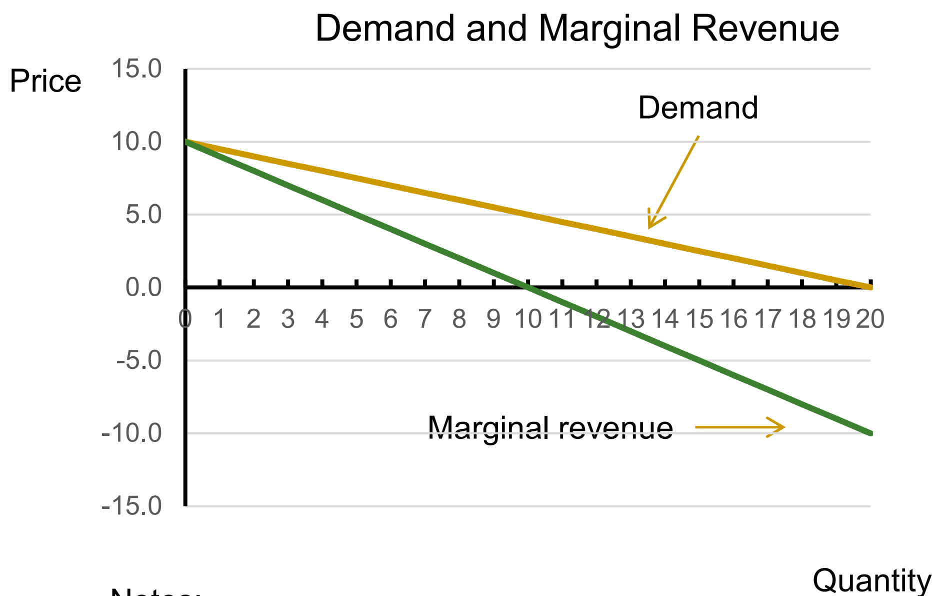

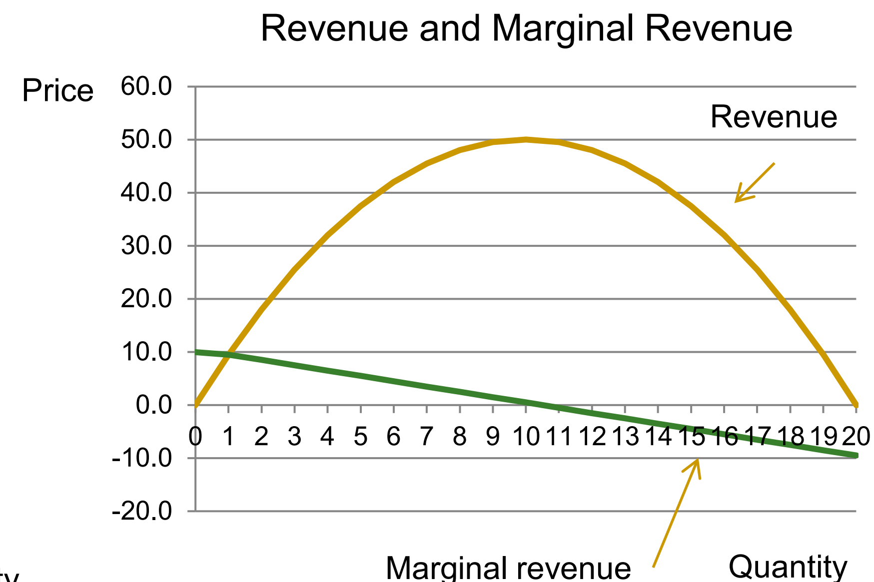

Notes:

1. When demand is linear, the slope of the marginal revenue curve is twice as steep as the demand curve. This means that marginal revenue crosses the *x*-axis at half of the distant to where the demand curve crosses the *x*-axis. In the first chart, the marginal revenue curve crosses the x-axis at 10, half of the distance to where the demand crosses the *x*-axis at 20.
2. When marginal revenue equals zero (here, a $q = 10$), revenues are at their maximum.

Producers and 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* $t(q)$ is the minimum cost to the firm of producing quantity q given the firm's technology
    - The firm's cost function t may change as the technology changes

Costs

- Some basic terms
  - *Revenues* ($r(q)$)
    - Price (p) times quantity (q) sold
    - Evaluated at a production level q
  - *Marginal revenue (mr)*: The net additional revenues that would be earned if the firm produced an additional unit
    - If the firm faces a downward-sloping demand curve for its product, the production of an additional unit will require a decrease in price in order to clear the market of the larger volume
    - Marginal revenue may be positive or negative

Costs

- Some basic terms
  - *Total costs* ($t(q)$)
    - The total cost of producing a production level q
    - Costs $t(q)$ = 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 cost* ($avc(q)$)
    - Variable cost divided by q

$$avc(q) = \frac{v(q)}{q}$$

  - *Marginal cost* ($mc(q)$)
    - The additional costs the firm would incur for producing one additional unit having produced q units

$$mc(q) = t(q+1) - t(q) \text{ in the discrete case}$$

$$= \frac{dr(q)}{dq} \text{ in the continuous case}$$

Costs

- Some basic terms
  - *Profits* ($\pi(q)$)
    - Revenues minus costs earned at a production level *q*
    - $\pi(q) = r(q) - t(q)$
  - *Marginal profit* ($m\pi$)
    - The net additional profit that the firm would make if it produced an additional unit
    - Or equivalently, marginal revenues minus marginal costs:

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

Marginal costs

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

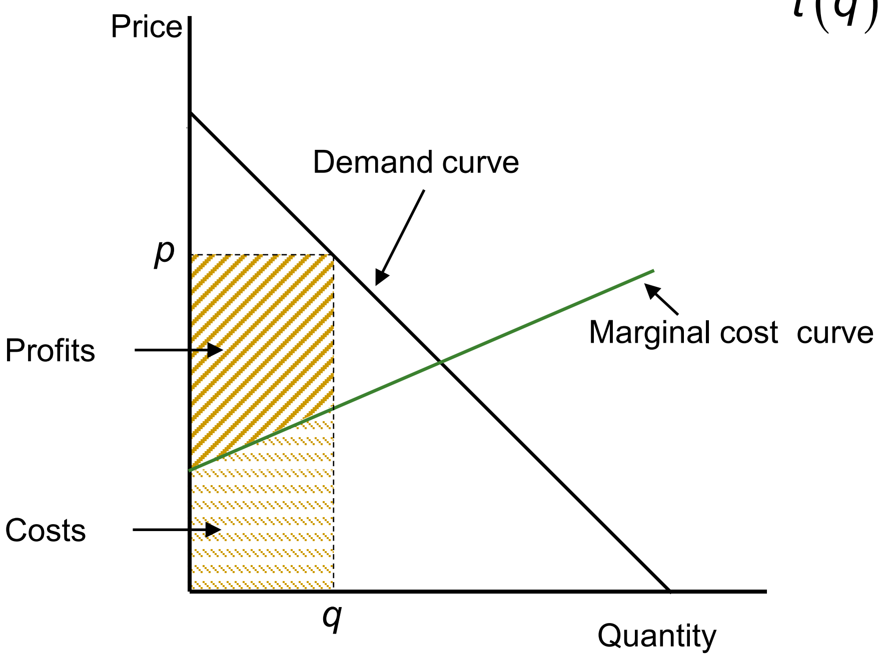

$$t(q) = \sum_{i=1}^{n} mc_i$$

+ *f* if there are fixed costs

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. Marginal costs in this diagram, however, are increasing and *not* constant.

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* ($t(q)$): 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* (f): 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* ($v(q)$): The cost incurred by the firm that depends on the firm's level of production q
  - *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 of the marginal costs to reach production level q:

$$v(q) = \sum_{i=1}^{q} mc(q_i)$$

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 $v(q) = kq$
  - Total cost is equal to fixed cost plus variable cost (all for producing a level of output q):

$$t(q) = f + v(q)$$
$$= f + kq \text{ (when marginal cost is a constant } k\text{)}$$

Profits and Profit Maximization

Profit maximization

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

$$\max_q \text{ Profits} = \text{Revenues} - \text{Costs}$$
$$= r(q) - t(q)$$

    - 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*.
  - Some important definitions
    - In this maximization problem, the *objective function* is the function that we are trying to maximize, in this case $r(q) - t(q)$.
    - 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 *parameters*.
    - Models that use *q* as the control variable are called *Cournot models*. Learn this term. It comes up all the time
    - 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 when we study differentiated product markets.

Profits

- When the firm faces a downward-sloping residual (inverse) demand curve $p = a + bq$:

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

  - The graph of the firm's revenues as a function of q is a parabola:

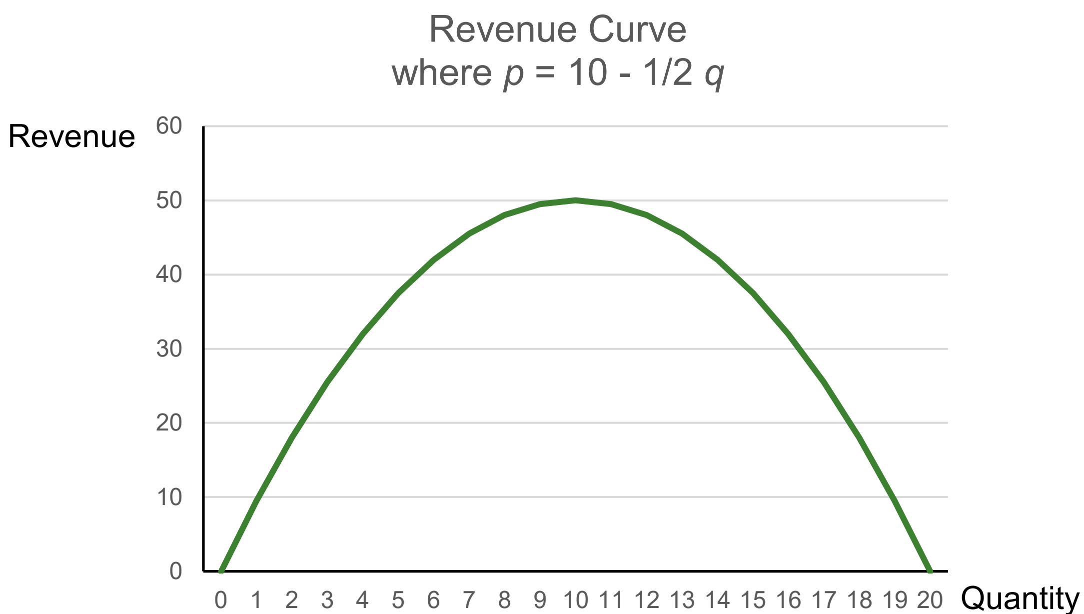

Profits

- At output q, total costs $t(q)$ are equal to fixed costs f plus variable costs $v(q)$:

$$t(q) = f + v(q)$$

  - With constant marginal costs c, variable costs $v(q)$ are equal to marginal cost c times output q:

$$v(q) = cq$$

  - Then total costs $t(q)$ may be expressed as:

$$\begin{aligned} t(q) &= f + v(q) \\ &= f + cq \end{aligned}$$

in the case of constant marginal costs

Profits

- Now we can express total profits $\pi(q)$ as:

$$\pi(p) = r(q) - t(q)$$
$$= (a + bq)q - cq$$
$$= \left[aq + bq^2\right] - cq$$

  - Graphically

where:
$p = 10 - ½\ q$
$f = 0$
$c = 4$

Remember, c here is the constant marginal cost of production

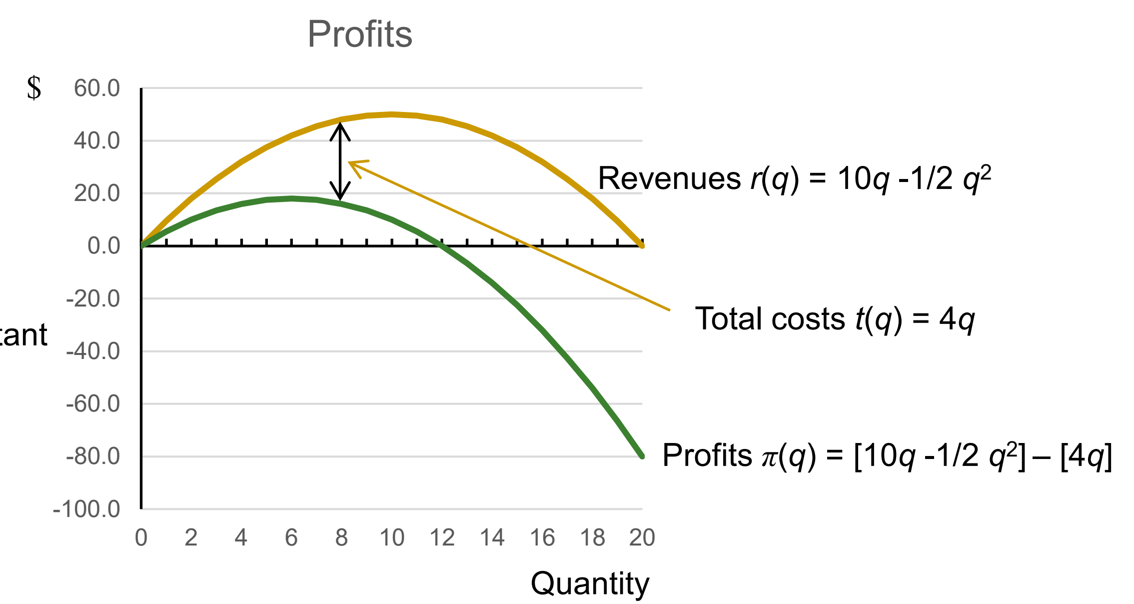

Profit maximization

- Profits are maximized when The slope at the top of the profit “hill” is zero (a horizontal tangent line):

where

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

$f = 0$

$c = 4$

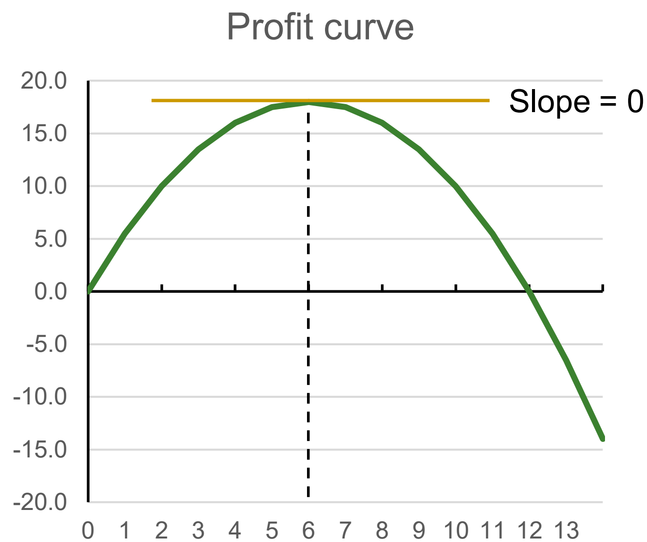

- From the chart we see that the profit-maximizing output q^* is 6.
- From the inverse demand curve, we can calculate $p^*(6) = 10 - (1/2)(6) = 7$
- $r^* = r(6) = p^*q^* = (7)(6) = 42$
- $f = 0$ (from the hypothetical)
- $v^* = v(6) = cq^* = (4)(6) = 24$
- $t^* = t(q^*) = f + v(q^*) = 0 + 24 = 24$
- $\pi^* = \pi(q^*) = r^* - t^* = 42 - 24 = 18$ (as shown in the graph)

Profit maximization

- *Refresher*: Marginal analysis—Some definitions
  - The slope of the revenue curve at an output q is called the *marginal revenue* $mr(q)$
    - Think of marginal revenue as the revenue the firm would earn if it produced one *additional* unit
    - If $r(q) = aq + bq^2$ (the revenue function for a linear inverse demand curve), then:

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

  - The slope of the total cost curve at an output q is called the *marginal cost* $mc(q)$
    - Think of marginal cost as the cost the firm would earn if it produced one additional unit
    - If $t(q) = f + cq$ (total costs with constant marginal costs), then:

$$mc(q) = c$$

  - The slope of the profit curve at an output q is called the *marginal profit* $m\pi(q)$
    - Think of marginal profit as the profit the firm would earn if it produced one additional unit
    - Marginal profit is marginal revenue minus marginal cost:

$$m\pi(q) = mr(q) - mc(q)$$

Optional: The marginal function is the derivative of the primary function. So, for example, the marginal revenue function is the derivative of the revenue function.

Profit maximization

- First order condition (FOC)
  - We know that profits are maximized at the top of the profit "hill," which is where the slope of the profit curve is zero
  - We know that the slope of the profit curve at an output *q* is the marginal profit $m\pi(q)$ evaluated at output *q*.
  - We also know that the marginal profit $m\pi(q)$ is equal to the marginal revenue $mr(q)$ minus the marginal cost $mc(q)$, all evaluated at output *q*, that is:

$$m\pi(q) = mr(q) - mc(q)$$

Be sure to learn this term!

  - The *first order condition* for a profit-maximizing level of output q^* is that the marginal profit at q^* equals zero, that is:

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

or equivalently:

$$mr(q^*) = mc(q^*)$$

This is a critical relationship. We will be using it throughout the course.

Profit maximization

- First order condition—Example

where
$p = 10 - \frac{1}{2} q$
$f = 0$
$c = 4$

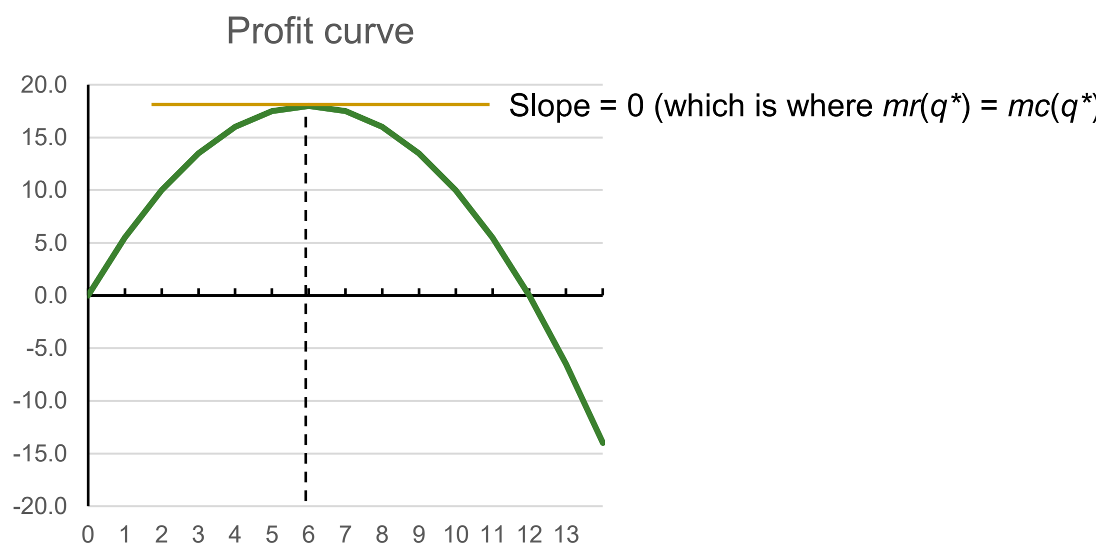

- $mr(q) = 10 - q$ (from the formula on Slide 33)
- $mc(q) = 4$ (from the hypothetical)
- FOC: $mr(q^*) = mc(q^*)$
  So $10 - q^* = 4$, so $q^* = 6$ (as shown in the diagram)
- $p^* = p(q^*) = 10 - \frac{1}{2} q^*$
  $= 10 - (\frac{1}{2})(6) = 7$ (solving for p^* from the inverse demand curve)

Profit maximization

- Marginal revenue/marginal cost diagrams: Step-by-step
  - Will build the diagram in five steps where $p = 10 - \frac{1}{2} q$

a. Draw the (inverse) demand curve

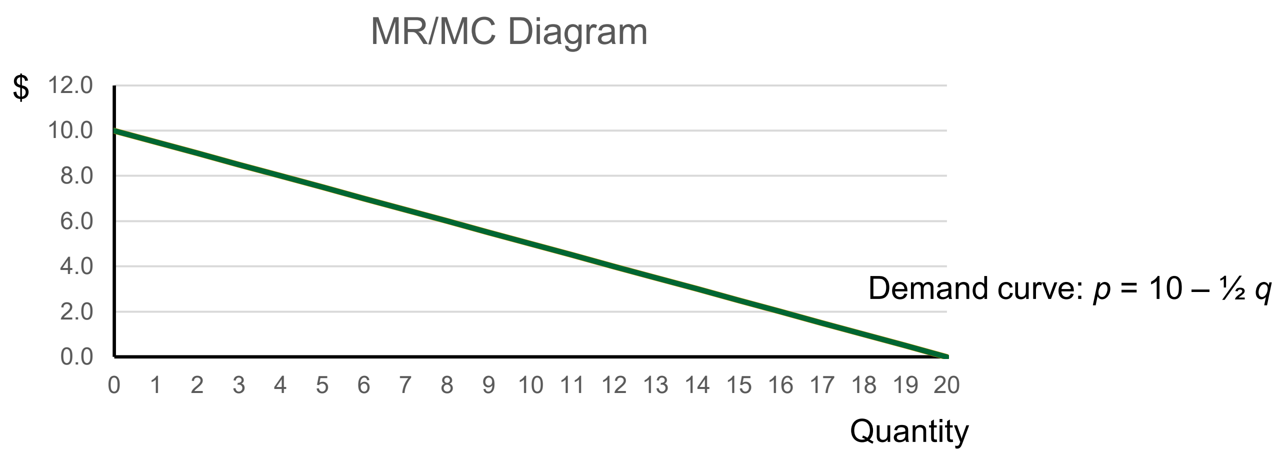

Profit maximization

- Marginal revenue/marginal cost diagrams
  - Will build this step-by-step

a. Draw the (inverse) demand curve: $p = 10 - \frac{1}{2} q$

b. Add the marginal revenue curve: $p = 10 - q$

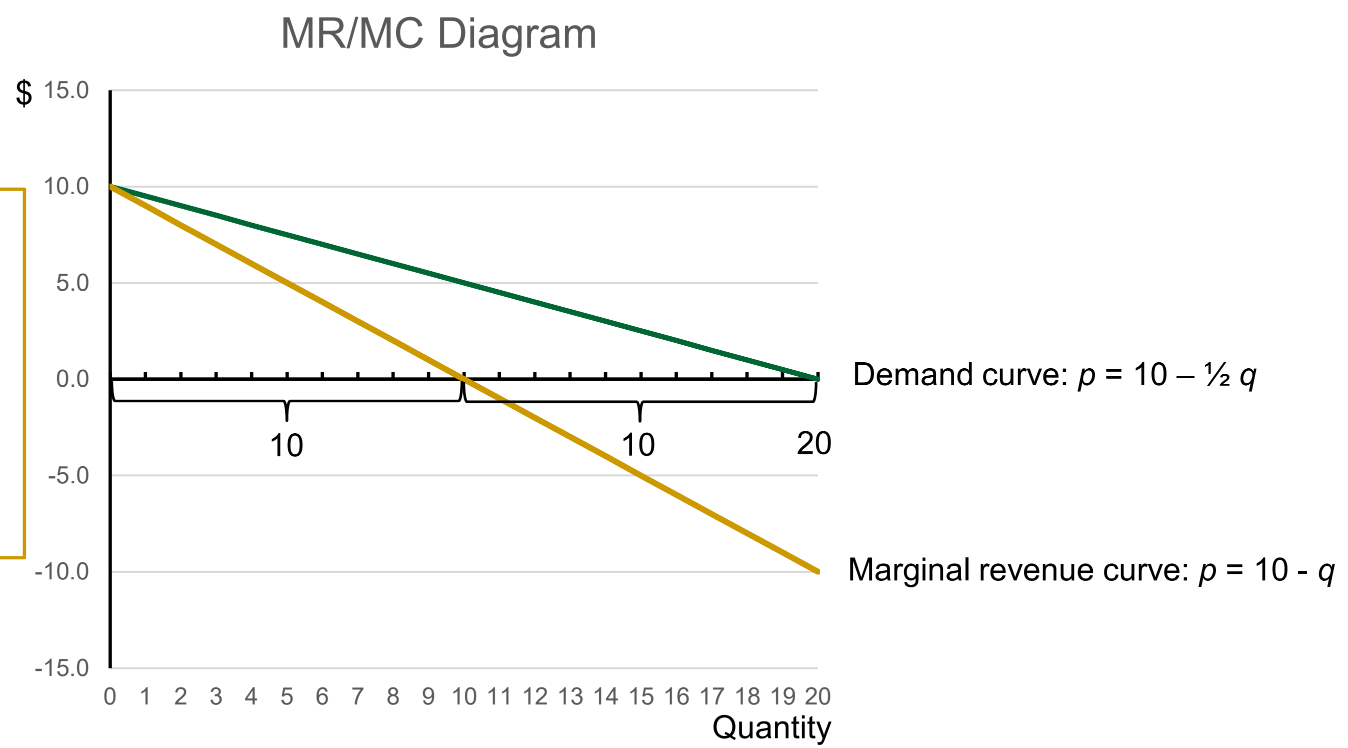

Note: With linear demand, the marginal revenue curve falls twice as fast as the inverse demand curve. Since the inverse demand curve crosses the x-axis at 20, the marginal revenue curve will cross the x-axis at 10.

Profit maximization

- Marginal revenue/marginal cost diagrams
  - Will build this step-by-step

a. Draw the (inverse) demand curve: $p = 10 - \frac{1}{2} q$

b. Add the marginal revenue curve: $p = 10 - q$

→ c. Add the marginal cost curve: $c = 4$ (constant marginal cost)

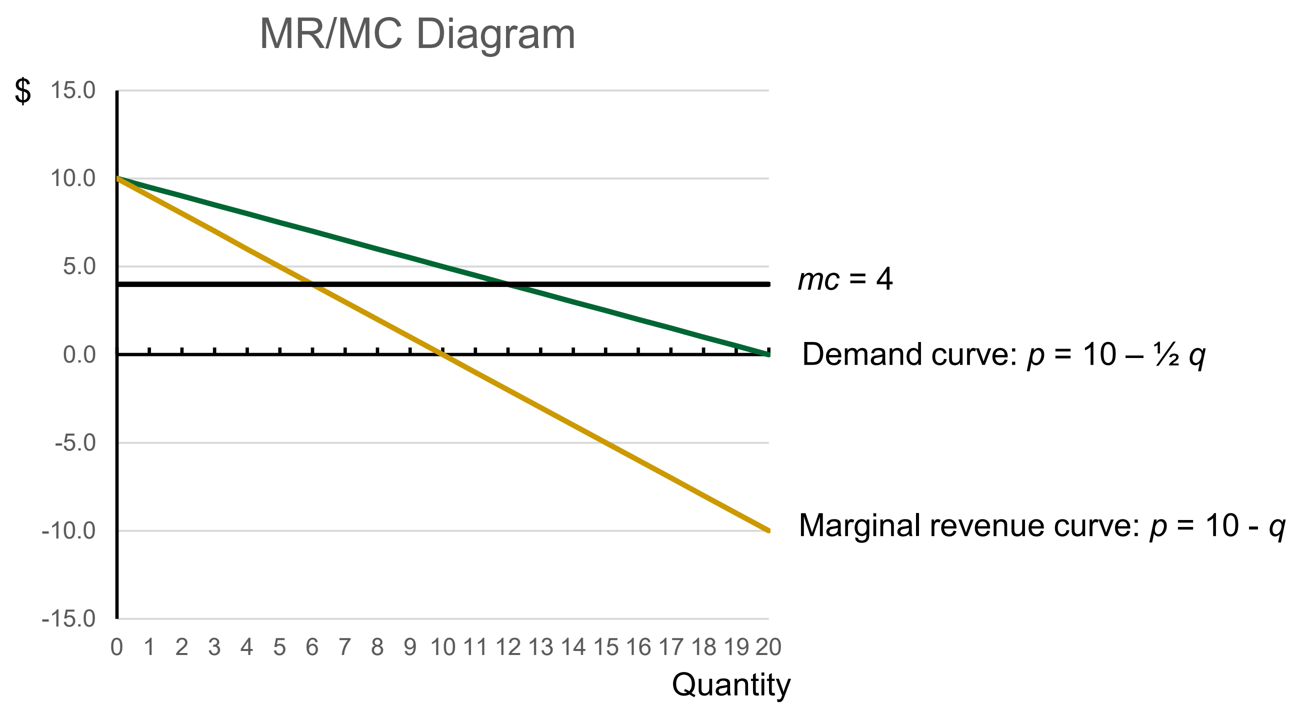

Profit maximization

- Marginal revenue/marginal cost diagrams
  - Will build this step-by-step

a. Draw the (inverse) demand curve: $p = 10 - ½\, q$

b. Add the marginal revenue curve: $p = 10 – q$

c. Add the marginal cost curve: $c = 4$ (constant marginal cost)

→ d. Find intersection of *mr* and *mc* curves to determine profit-maximizing q^* (= 6)

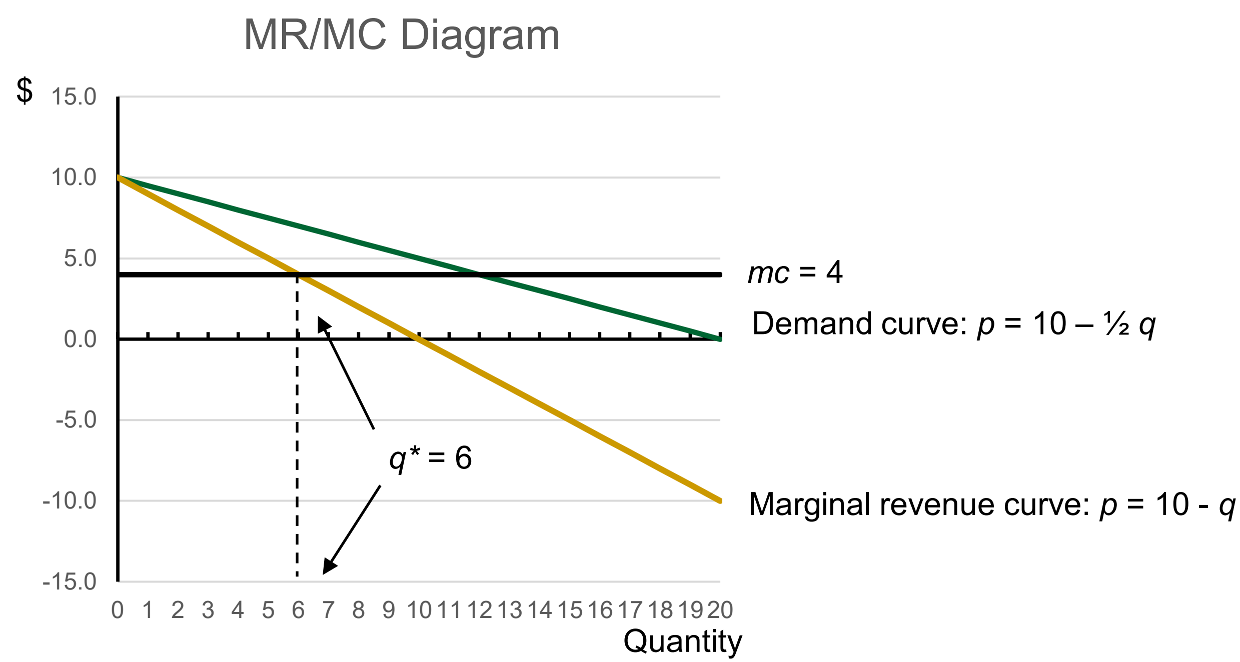

Profit maximization

- Marginal revenue/marginal cost diagrams
  - Will build this step-by-step

a. Draw the (inverse) demand curve: $p = 10 - ½\, q$
b. Add the marginal revenue curve: $p = 10 – q$
c. Add the marginal cost curve: $c = 4$ (constant marginal cost)
d. Find intersection of *mr* and *mc* curves to determine profit-maximizing q^* (= 6)
e. Find $p^* = p(q^*)$ from the inverse demand curve ($p^* = 7$)

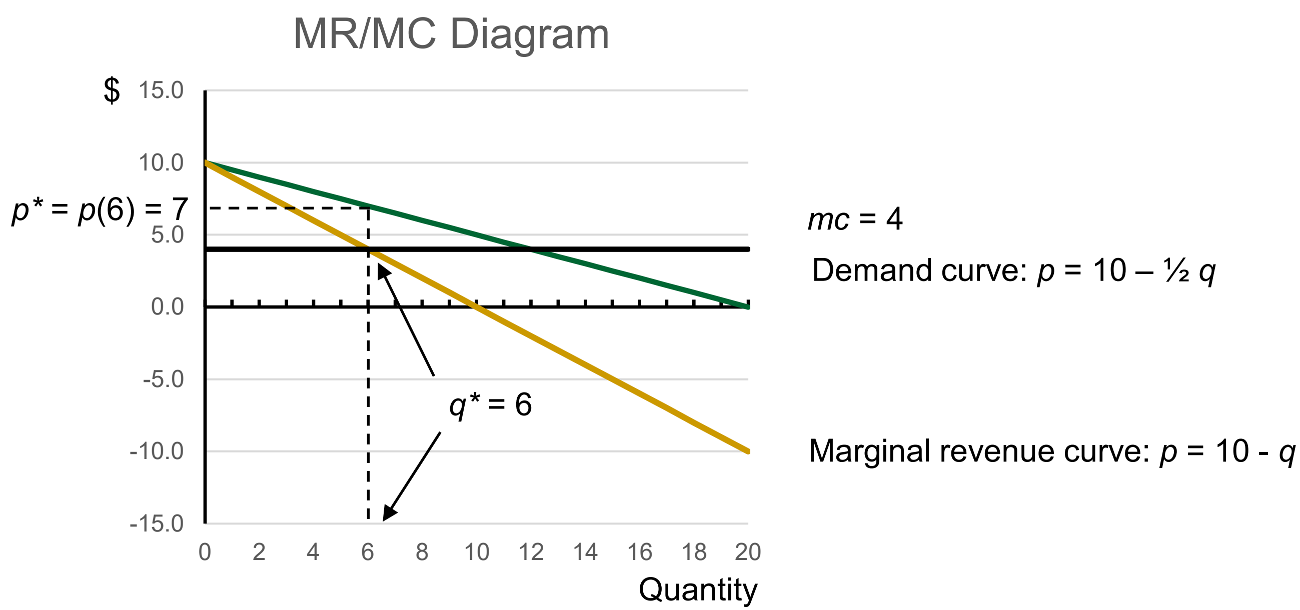

Incremental Revenue

Incremental revenue

- Introduction
  - *Incremental revenue* is the net gain in revenue that a firm could earn if it were to increase its product by some amount Δq
    - Incremental revenue is marginal revenue when $\Delta q = 1$
  - Incremental revenue is important when determining whether a firm should change its output level to increase its profits
  - Incremental revenue can be positive or negative (same as marginal revenue)
    - Moving from q_1 to q_2 increases revenue (incremental revenue is positive)
    - Moving form q_2 to q_3 decreases revenue (incremental revenue is negative)

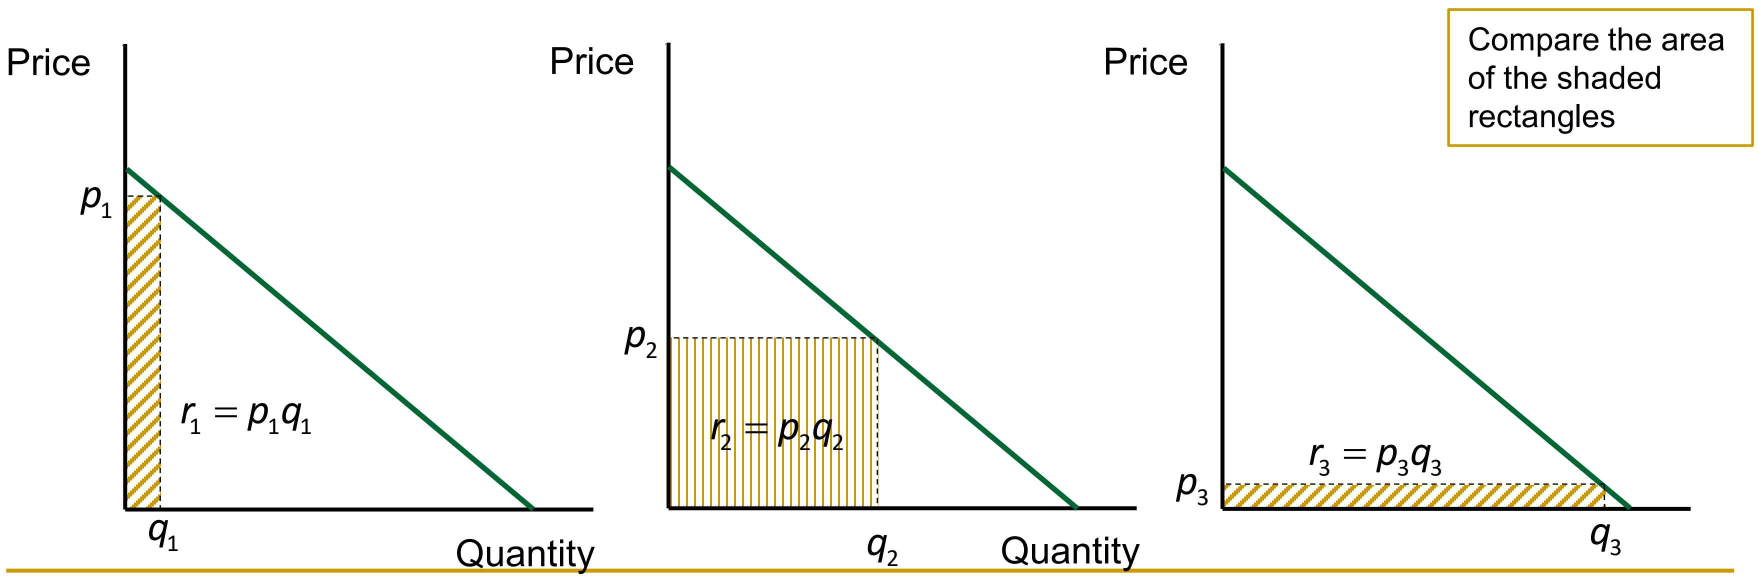

Incremental revenue

- 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 revenue 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 (= $q\Delta p$)
- So

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

When $\Delta q = 1$, this is the formula for marginal revenue in the discrete case

Incremental revenue

- Graphically

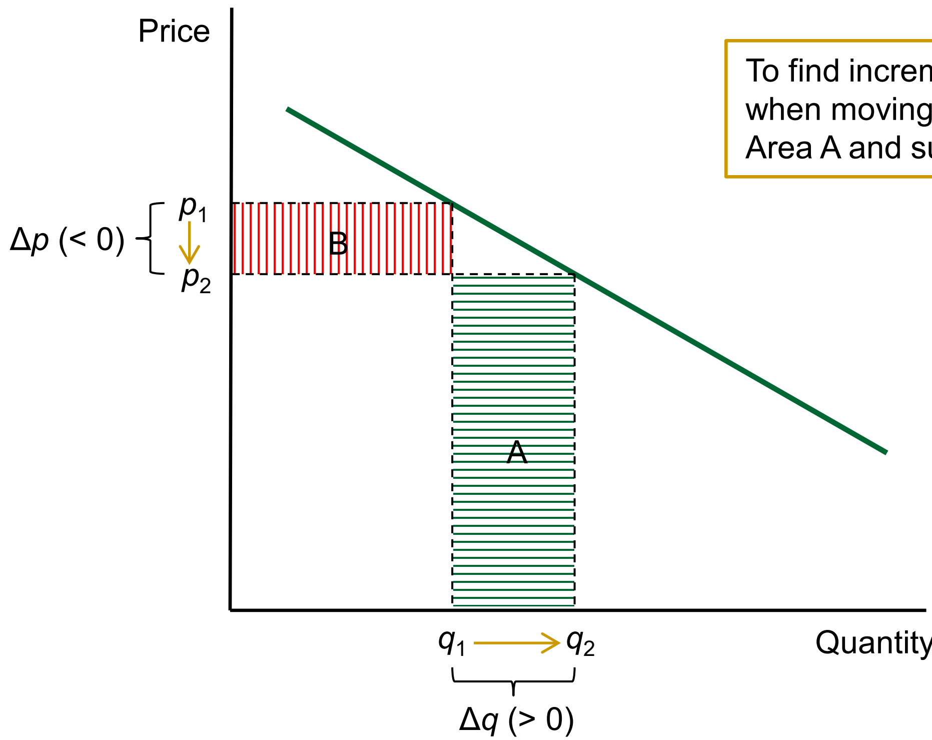

To find incremental revenue *IR* when moving from q_1 to q_2, add Area A and subtract Area B

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

- Example
  - (Inverse) demand: p =10 – ½ q
  - Starting point: $q_1 = 4$
  - End point: $q_2 = 8$

You need to calculate these variables:

So $p_1 = 8$	$\Delta q = q_2 - q_1 = 8 - 4 = 4$
So $p_2 = 6$	$\Delta p = p_2 - p_1 = 6 - 8 = -2$

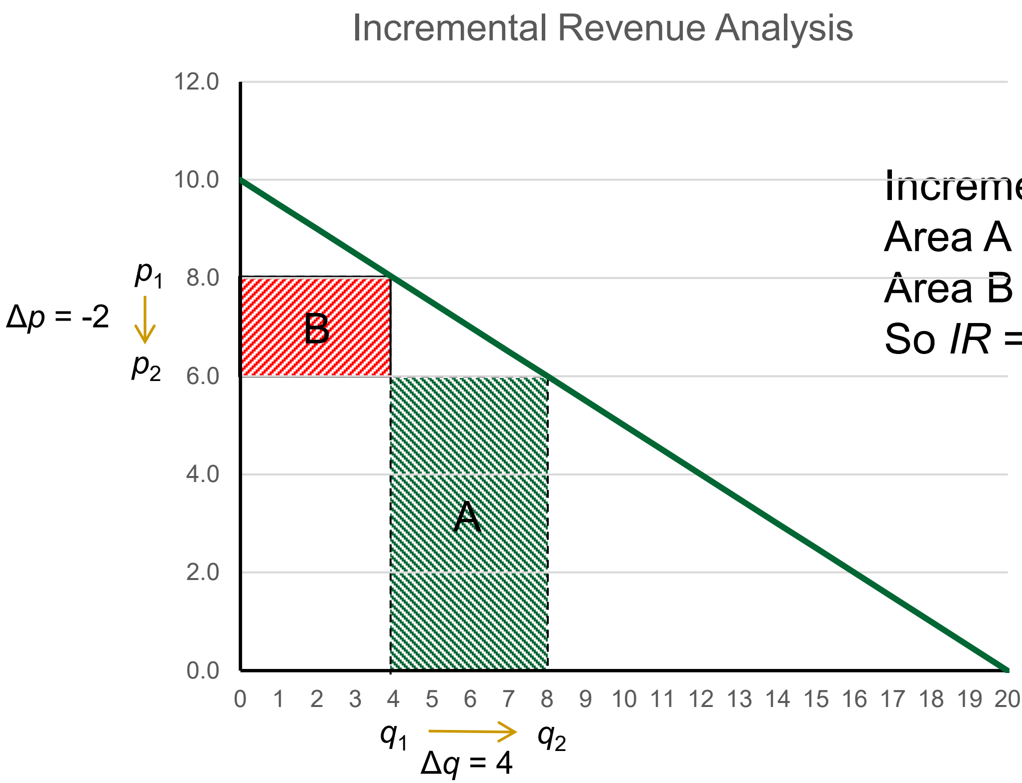

Incremental revenue = Area A – Area B
Area A = $p_2 \Delta q$ = (6)(4) = 24
Area B = $q_1 \Delta p$ = (4)(-2) = -8
So IR = 32 – 8 = 16

That is, the firm makes $16 more in revenues by moving from q_1 to q_2

Incremental profits

- We can easily extend the analysis of incremental revenues to incremental profits—We just have to:
  - Add the costs of additional production if we are adding to output ($\Delta q > 0$), or
  - Subtract the costs of a reduction in output ($\Delta q < 0$)

Example: Adding production

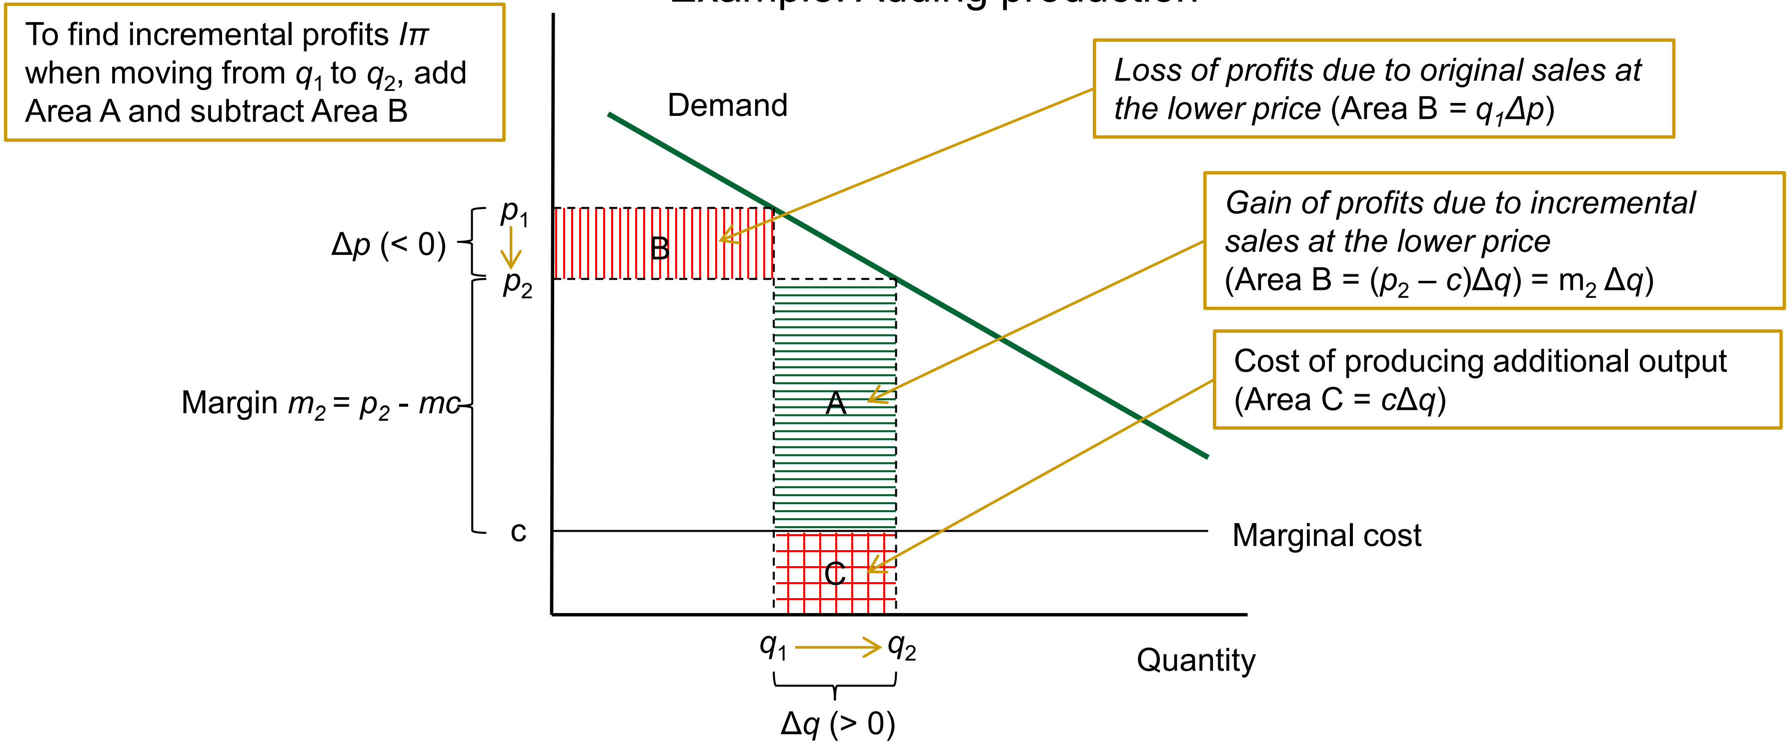

Incremental profits

- Example: Output increase
  - (Inverse) demand: p =10 – ½ q
  - Starting point: q_1 = 2 So p_1 = 9
  - End point: q_2 = 6 So p_2 = 7
  - Constant marginal cost c = 4

$\Delta q = q_2 - q_1 = 6 - 2 = 4$

$\Delta p = p_2 - p_1 = 9 - 7 = -2$

Margin $m_2 = p_2 - c$

$= 7 - 4 = 3$

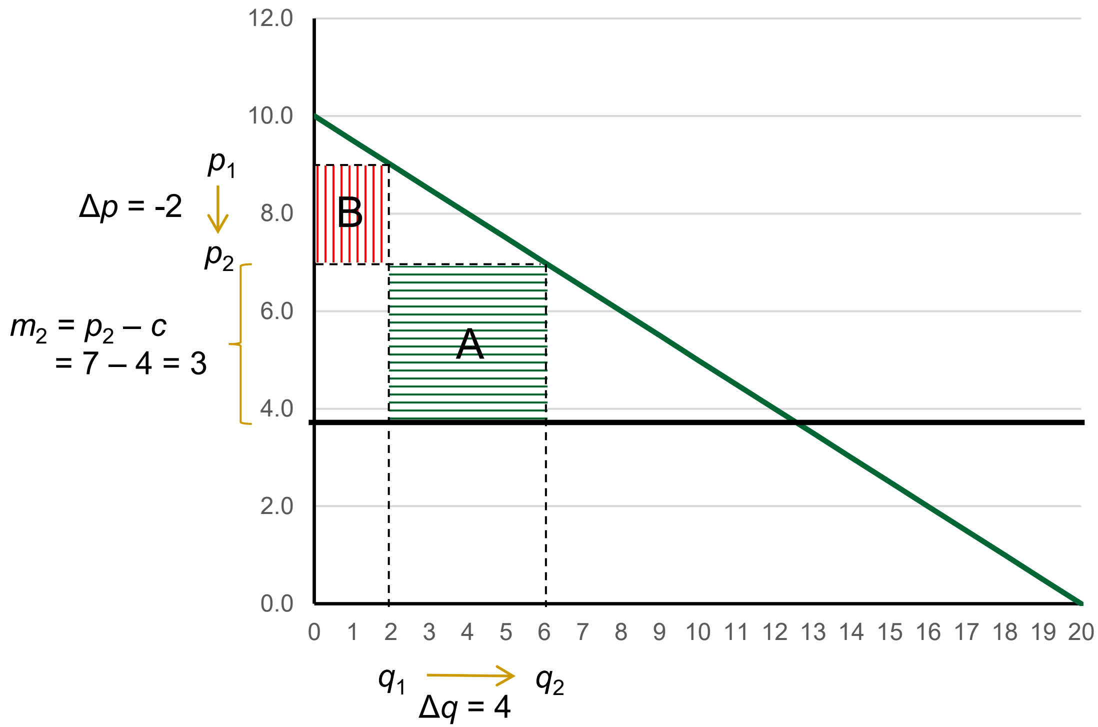

Incremental profits = Area A – Area B

Area A = $m_2\Delta q$ = (3)(4) = 12

Area B = $q_1\Delta p$ = (2)(-2) = 4

So $I\pi$ = 12 – 4 = 8

That is, the firm makes $8 more in profits by moving from q_1 to q_2

Incremental profits

- Example: Price increase
  - (Inverse) demand: $p = 10 - ½ q$ So $q = 20 - 2p$
  - Starting point: $p_1 = 5$ So $q_1 = 10$ $\Delta q = q_2 - q_1 = 9.5 - 10 = -0.5$
    End point: $p_2 = 5.25$ So $q_2 = 9.5$ $\Delta p = p_2 - p_1 = 5.25 - 5 = 0.25$
  - Constant marginal cost $c = 4$

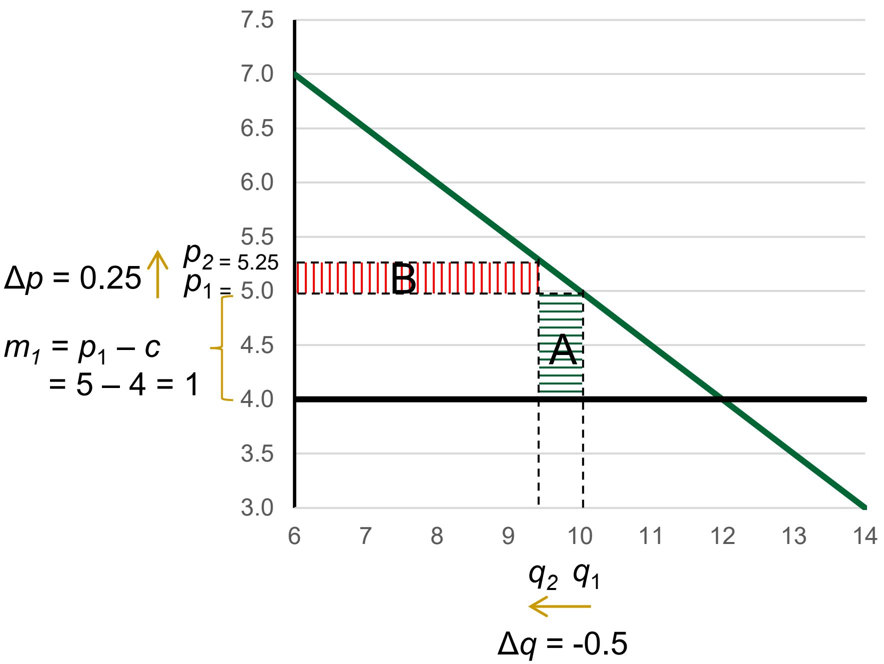

With an increase price and a concomitant *reduction* in output, the roles of Areas A and B are reversed:

Area A now represents the *loss* of profits from lost sales that would have been made at original price p_1 (= $m_1 \Delta q$)

Area B represents the *gain* of profits from the increased price charged on the sales that continue to be made (= $q_2 \Delta p$)

Incremental profits = Area B – Area A

Area B = $q_2 \Delta p$ = (9.5)(0.25) = 2.375
Area A = $m_1 \Delta q$ = (1)(-0.5) = -0.5

So *incremental profits* = 2.375 – 0.5 = 1.875

Incremental profit

- Observations
  - The prior example shows that under the conditions of the hypothetical, a 5 percent price increase would be profitable to the firm

> This is mathematically identical to the exercise required by the *hypothetical monopolist test*, which is the primary analytical tool used by the agencies and the courts to define relevant markets. The hypothetical monopolist test asks whether a hypothetical monopolist of the candidate market could profitably sustain a "small but significant and nontransitory increase in price" (SSNIP), usually taken to be 5 percent. If so, the candidate market is a relevant market. In the prior example, if we assume that the demand curve is for the candidate market as a whole, this will be the residual demand curve for the hypothetical monopolist. If the original market price was $5 (as in the hypothetical), the hypothetical monopolist would find it profitable to reduce output in order to raise price by a 5 percent SSNIP.
>
> We will confront the hypothetical monopolist test in almost every case study going forward, starting with the Sanford/Mid Dakota Clinic case study next week. You will have plenty of opportunities to become familiar with the mechanics of the hypothetical monopolist test.

The Neoclassical Model

Assumptions of the "neoclassical model"

- You may heard references to the "neoclassical model" in economics. Here are the standard assumptions.

1. 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*
   - Aggregate demand is the sum of all individual consumer demands
   - Consumers are *price-takers*, that is, they take the market price as given and something they cannot negotiate
2. 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*
   - Aggregate output is the sum of all individual firm outputs
3. Equilibrium condition
   - No price discrimination (all purchases are made at the single market price)
   - The market clears at the market price (i.e., demand equals supply and there are no inventories):

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

$\sum$ simply means to add up the *q*'s

The neoclassical model

- The neoclassical model drives most of antitrust economics
  - Neoclassical models are essentially *deductivist*, that is, they derive their propositions by deriving the logical consequences of the assumptions
  - This makes mathematics—the key tool in making deductions from economics assumptions—a critical element in manipulating neoclassical models
- Criticisms—We often observe violations of the key assumptions in the real world. For example—
  - Consumers often do not appear to be acting in their own best interests
  - Consumers sometimes have market power to negotiate prices and do not need to accept the market price as given
  - Firms often do not appear to be acting to maximize their profits
- Adjustments
  - Neoclassical models can be modified to deal with some of the reasons why the behavioral assumptions do not appear to be satisfied (e.g., asymmetric information, uncertainty, bounded rationality)
- *Bottom line*:
  - No other general approach has emerged that serves better to tune the intuitions of antitrust policymakers use the neoclassical model
  - So antitrust policymakers use the neoclassical model but modify it as necessary

The Nash Bargaining Model

Bargaining

- The idea
  - An assumption up to this point is that customers are *price-takers*
    - That is, they lack the power to negotiate with suppliers and take the market price as given
  - This assumption is usually justified in market where there are very large number of customers, all of whom want to purchase the supplier's product
  - But when customers are relatively few in number often we see *bargaining* between a supplier and a customer over price
  - Some examples:
    - Negotiations between a seller of a house and a potential buyer over the sale price and other terms of the transaction (e.g., the closing date, repairs to be made to the house)
    - Negotiations between a dealer and a potential purchaser over the sale price of a car
    - Negotiations between a paint manufacturer and a titanium dioxide supplier over the price and terms of a long-term titanium dioxide supply contract

Bargaining

- Reaching agreement
  - Generally, the parties in a negotiation will reach an agreement when—
    1. There is a *gain from trade* from an agreement, that is, when the highest price the purchaser will pay (the purchaser's reservation price) is greater than the lowest price the seller is willing to accept (the seller's reservation price), *and*
    2. Neither party has a better alternative than the negotiated price
  - Two corollaries:
    1. Parties will not reach agreement if there is no gain from trade
    2. Parties will not reach an agreement if at least one party has a better alternative than the best bargain it can negotiate
  - Think about this is the context of a sale of a house
    - Scenario 1
      - The least the seller is willing to take is $500K
      - The most the buyer is willing to pay is $400K
      - → No agreement: The seller's reservation price to sell is higher the buyer's reservation price to buy
    - Scenario 2
      - The seller is willing to sell at $500K
      - Buyer 1 is willing to pay $550K
      - Buyer 2 is willing to pay $560K
      - Although the seller would receive a $50K gain from a sale to Buyer 1 at $550K, the seller would receive a $60K gain from selling to buyer 2 → Seller does not reach agreement with buyer 1 because Seller has a better alternative

Finding agreement on price

- The problem
  - Assume that the conditions for reaching some agreement on price are satisfied:
    - There are gain from trade to both parties from an agreement, and
    - There are no better alternatives to either party
  - *Question*: What is the price on which the parties will ultimately agree?

Finding agreement on price

- Consider this question in three very simple games
  - Game 1:
    - The parties have $1000 to split between them
    - If they agree on the division, each will receive their agreed-upon share
    - If they do not agree, then each get nothing
    - The outcome: When played, most pairs decide on a 50/50 split
  - Game 2:
    - Same as Game 1, except that in the absence of agreement Player 1 gets $600 and Player 2 gets nothing
    - The outcome
      - Not as obvious as in Game 1, except we know that Player 1 will get at least $600 (since it has an no-agreement alternative of $400)
      - Both players know this, so the bargaining is really over how to split the remaining $400
      - One possible outcome: $800/$200 (splitting the $400 equally)
      - Another possible outcome: $600/$400 (but why would Player 1 agree to this?)
  - Game 3
    - Same as Game 1, except that in the absence of agreement Player 1 gets a disagreement payoff of $200 and Player 2 gets nothing: *Now what happens?*

This is known as the "Divide the Dollar" game. With inflation, to make the game serious I have increased the amount to be split to $1000

This is called the *disagreement payoff* to Player 1

This is a *game of perfect information*, so both players know everything about the game prior to playing it (including that Player 1 will get $600 in the absence of an agreement)

Finding agreement on price

- Some observations
  - It is an empirical question of where the players would reach agreement in each of these games
  - There is no natural law of bargaining that requires the parties in any of these games to agree on a particular division (including the “fair” division of Game 1)

The Nash Bargaining Solution[1]

- John Nash (1928 – 2015)
  - Of *A Beautiful Mind* fame
  - The founder of modern game theory
  - Research mathematician at Princeton
  - Winner:
    - Nobel Prize in Economic Sciences
    - Abel Prize (effectively the Nobel Prize for mathematics)

1 John Nash, *The Bargaining Problem*, 18 Econometrica 155 (950).

The Nash Bargaining Solution

- Some definitions
  - *Payoffs*:
    - *Definition*: The possible outcomes of the game to each player
    - The payoffs in Game 1 are the set of pairs (x_1, x_2) such that $x_1 \geq 0$, $x_2 \geq 0$, and $x_1 + x_2 = 1000$ (when there is agreement) plus $x_1 = 0$, $x_2 = 0$ (when there is no agreement)
  - *Utility function*:
    - *Definition*: A function for each player that gives the player's valuation for each payoff she might receive
    - Where the payoffs are in dollars, the utility function is the identify function (i.e., a dollar is valued at $1 in the Divide the Dollar game)
    - In the house selling game, the utility function gives the value to the buyer of owning the house (after accounting for the purchase price)
  - *Strategies*:
    - *Definition*: The choices available to each player (which may differ between players)
    - In the game, each player picks a strategy and the payoffs of the game is determined by the strategies the players chose
    - The strategies in Game 1 for each player are the divisions of $1000 that the player could choose (without regard to what the other player chooses)
      - The choice of strategies determine the payoffs in the game. Let (s_1, s_2) and (t_1, t_2) be the strategy choices of players 1 and 2, respectively. Then if (s_1, s_2) = (t_1, t_2) (agreement), the payoff is (s_1, s_2); If (s_1, s_2) ≠=(t_1, t_2) (no agreement), the payoff is (0,0)

The Nash Bargaining Solution

- Nash equilibrium
  - *Definition*: The equilibrium in a noncooperative game where no player has an incentive to change its strategy choice if that all of the other players continue with their original strategies
  - Most games have many Nash equilibria
  - If x_1 and x_2 are the divisions to players 1 and 2, respectively, of Game 1, then every (x_1, x_2) provided $x_1 \geq 0$, $x_2 \geq 0$, and $x_1 + x_2 = 1000$ is a Nash equilibrium
  - Example: Divide the Dollar game (Game 1 above)
    - *Recall the game*: The parties have $1000 to split between them. If they agree on the division, each will receive their agreed-upon share. If they do not agree, then each get nothing
    - *Nash equilibria*: Every division of the $1000 is a Nash equilibrium
      - Say (x_1, x_2) is the outcome of the game, where $x_1 \geq 0$, $x_2 \geq 0$, and $x_1 + x_2 = 1000$
      - This means that both players agree to the division, that is, $(s_1, s_2) = (t_1, t_2) = (x_1, x_2)$
      - Assuming player 2 says with the division (x_1, x_2), then any other choice of the division by player 1 will result in both players receiving nothing. Hence, under the Nash assumption that Player 2 does not change her strategy, player 1 has no incentive to change its strategy. Since the game is symmetrical, using the same reasoning, player 2 has no incentive to change his strategy. Hence, (x_1, x_2) is a Nash equilibrium.
      - (0,0) is not a Nash equilibrium, since Player 1 could change its strategy to agree with Player 2 and thereby improve Player 1's payoff (and vice versa for Player 2)

The Nash Bargaining Solution

- The Nash bargaining solution
  - *Definition*: A *Nash bargaining solution* is a Nash equilibrium in a bargaining game that also satisfies:
    1. Pareto efficiency
       - Requires all gains for trade to be exhausted (e.g., $x_1 + x_2 = 1000$ and not $x_1 + x_2 < 1000$)
       - An inefficient outcome is unlikely, since it leaves space for renegotiation
    2. Symmetry
       - If the players are indistinguishable, the agreement should not discriminate between them
       - This assures symmetry in payoffs in Game 1 when the players are identical except for their names
    3. Invariance to equivalent payoff representations
       - An affine transformation of the payoffs (including the disagreement payoffs) dos not change the outcome of the game
       - An affine transaction is of the form $a_i + b_i x_i$
         - So if (x_1,x_2) is a Nash bargaining solution of a game, then the new game with transformed payoffs has a Nash bargaining solution of $(a_1 + b_1 x_1 , a_2 + b_2 x_2)$
    4. Independence of irrelevant alternatives
       - If a game has a Nash bargaining solution of (x_1,x_2), then (x_1,x_2) will also be a Nash bargaining solution to a game that contains (x_1,x_2) as a payoff but has fewer strategies available to the players
       - Example: Say (500,500) is a Nash bargaining solution to a Divide the Dollar game where the players can chose any arbitrary dollar division as their strategy. Then (500,500) is also a Nash bargaining solution when the players can chose dollar division only in $100 increments, so (600,400) is a permissible division, but (402,598) is not

The Nash Bargaining Solution

- The Nash bargaining solution
  - *Proposition*: An outcome (x_1, x_2) of a two-person bargaining game is a Nash bargaining solution if it solves the following optimization problem:

$$\max_{x_1, x_2} (x_1 - d_1)(x_2 - d_2)$$

$$\text{subject to} \qquad (x_1, x_2) \geq (d_1, d_2)$$

> The disagreement payoffs (d_1, d_2) are also called the *best alternative to a negotiated agreement* (BATNA)

where d_1 and d_2 are the disagreement payoffs to Players 1 and 2, respectively

  - Observations
    - The Nash bargaining solution requires that the players reach agreement whenever there is some gains to trade for at least one player (that is, there is some $x_i > d_i$ for some player)
      - The gain from trade for player *i* is the payoff from the game minus the disagreement payoff (i.e., $x_i - d_i$)
      - Empirical studies do not provide much support for the Nash bargaining solution predicting real-world outcomes in all but the simplest of games
    - Remember that the Nash bargaining solution is derived from a set of four axioms
      - It is a *normative* solution, not a solution that explains the empirical results of bargaining games
    - The Nash bargaining game is a *game of complete information*—that is, the players know:
      1. All of the available strategies available to all players
      2. All of the available payoffs (including the disagreement payoffs), *and*
      3. How the choice of strategies map into the outcome of the game
    - In a bargaining game, if there exists a player such that there is a payoff to that player greater than the player's disagreement payoff, then the players will reach agreement

The Nash Bargaining Solution

- The Nash bargaining solution
  - *Example 1*: Divide the Dollar game where the amount at stake is $1000 and the disagreement payoffs are zero. The Nash bargaining solution is the outcome (x_1, x_2) that solves the following problem:

$$\max_{x_1,x_2} \left(x_1 - 0\right)\left(x_2 - 0\right)$$

where $$x_1 + x_2 = 1000$$

Mathematically, this is the same as maximizing the area of a rectangle subject to a set perimeter (where x_1 and x_2 are the length of the sides). The maximum area is obtained in a square, so that the sides are equal: $x_1 = x_2 = 500$ is the Nash bargaining solution.

For a slightly more technical approach, consider that $x_2 = 1000 – x_1$. So the problem becomes:

$$\max_{x_1,x_2} x_1\left(1000 - x_1\right) = \max_{x_1,x_2} 1000x_1 - x_1^2$$

The curve of the expression being maximized is a parabola (because of the x^2). To find the maximum, take the derivative and set it equal to zero: $1000 - 2x_1 = 0$

Solving, $x_1 = 500$ and so $x_2 = 500$. So the Nash bargaining solution splits the gains from trade between the two players.

The Nash Bargaining Solution

- The Nash bargaining solution
  - *Example* 2: Same game except (d_1,d_2) = (300, 200). The Nash bargaining solution is the outcome (x_1,x_2) that solves the following problem:

$$\max_{x_1,x_2}(x_1 - 300)(x_2 - 200)$$

where $$x_1 + x_2 = 1000$$

In thinking about this problem, note that

- Player 1 will not agree to accept less than 300
- Player 2 will not agree to accept less than 200
- So the gains from trade in reaching an agreement is 1000 – 300 – 200 = 500

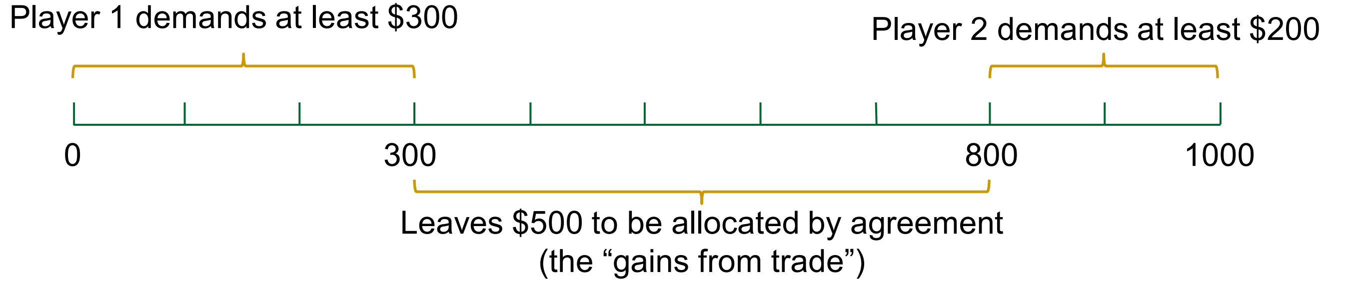

The Nash bargaining solution splits the gains from trade, so each player gets their disagreement payoff plus one-half of the gains from trade (250 each):

- Player 1 gets 300 + 250 = 550
- Player 2 gets 200 + 250 = 450

The Nash Bargaining Solution

- The generalized Nash bargaining solution with asymmetrical bargaining power
  - The Nash bargaining solution assumes symmetry in the players bargaining power
  - We can introduce asymmetrical bargaining power by modifying the Nash solution to the following:

$$\max_{x_1,x_2} \left(x_1 - d_1\right)^{\alpha} \left(x_2 - d_2\right)^{1-\alpha}$$

subject to $\left(x_1, x_2\right) \geq \left(d_1, d_2\right)$

where α is between 0 and 1

  - The idea
    - The more bargaining power Player 1 has relative to Player 2, the larger α is compared to 1 - α
    - This generalizes the concept of the Nash bargaining solution to account for differences in the players, which arise from such factors as—
      - The procedure through which negotiations are conducted (which may favor one player over the other)
      - Differences in the tactics employed by the bargainers
      - Differences in the information available to the players
      - Differences in the discount rate of the players
      - Differences in the risk averseness of the players
    - The generalized Nash bargaining solution in principle can account for these differences
      - BUT any bargaining outcome can be "explained" with the generalized Nash bargaining solution by picking the "right" alpha—that is, the solution lacks testable content since alpha cannot be empirically determined apart from deriving it from the outcome of a bargaining game