
Unit 8. Competition Economics

Part 1. Demand, Costs, and Profit Maximization

Professor Dale Collins
Merger Antitrust Law
Georgetown University Law Center

September 10, 2023

A prefatory note

- The focus of this unit
 - We will focus on the basic economics principles that will help us model (predict) the market equilibrium outcomes on a going-forward basis with and without the merger under review
- These tools are central to the consumer welfare approach to antitrust law
 - Under the consumer welfare approach to antitrust law, competitive harm occurs when consumers are worse off with the merger than without the merger
 - Requires a two-step analysis
 1. Predict the equilibrium outcomes of the market—primarily market price and market output—with and without the merger (for reasons we discussed in Unit 2)
 - When the merger antitrust analysis is performed preclosing (the typical case), we cannot empirically observe the market equilibrium with the merger and so we have to predict it
 - Moreover, when markets are dynamic and changing, the market equilibrium in the future without the merger may be different from the equilibrium in the past, so here too the market equilibrium must be predicted rather than observed
 - Economics provides the tools to make these predictions
 2. Compare the consumer welfare implications of the predicted market equilibria
 - The merger is anticompetitive if and only if consumers are worse off with the merger than without it

The consumer welfare approach compares outcomes of the market process; economics supplies the tools to predict these market outcomes

A prefatory note

- Will these tools still be applicable in the Biden administration?

The short answer is “Yes, absolutely!”

- It is true that the Neo-Brandeisian leadership of the Biden DOJ and FTC philosophically rejects the consumer welfare standard and prefers a structural approach focused on decentralizing the economic and political power of firms
 - Recall from Unit 3 that Neo-Brandeisians compare market structures and not the outcomes of market processes
- BUT in the absence of a legislative change, courts will continue to use the consumer welfare approach and the economic methods to apply it that have been developed in precedent
 - This effectively forces the Biden DOJ and FTC to use the consumer welfare standard and its associated economic tools in the cases they are litigating
 - In practice, this is exactly what the Biden agencies are doing in court
 - Moreover, with two exceptions (see the next slide), the Biden agencies—like their predecessors—are focusing on competitive effects in seller (downstream) markets

We will develop the basic economic tools for analyzing downstream competitive effects in this unit

A prefatory note

- Some new ideas that will require some new economic tools
 - Within the context of the consumer welfare approach, the Biden agencies are introducing two (somewhat) new concepts
 1. Monopsony power: Anticompetitive effects in buyer (input) markets
 - *Example:* Penguin/Simon & Schuster¹ (covered in Class 10)
 - *Holding:* The merger between these two major book publishers violated Section 7 because it was likely to reduce the advances paid to authors (i.e., input suppliers) with no offsetting consumer benefits
 - Observations
 - Although there have been few Section 7 litigated cases based on an increase in monopsony power, other areas of antitrust law (which equally employ the consumer welfare standard) see this as an unexceptional application—so it should be easily accepted in Section 7 law
 - BUT the test of anticompetitive effect when consumers benefits from lower prices remains unresolved

¹ See *United States v. Bertelsmann SE & Co. KGaA*, No. CV 21-2886-FYP, 2022 WL 16949715 (D.D.C. Nov. 15, 2022) (permanently enjoining the Penguin/Simon & Schuster merger)

A prefatory note

- Some new ideas that will require some new economic tools
 - Within the context of the consumer welfare approach, the Biden agencies are introducing two (somewhat) new concepts (con't)
 2. Entrenchment: Dominant firms should not be allowed to acquire firms or assets that could further entrench them
 - Entrenchment is a “conglomerate” merger theory, that is, a theory applying to transactions that are neither presently nor in the foreseeable future horizontal nor vertical
 - The idea is that somehow the combination of the products of the merging firms will “entrench” the dominant positions of the some of the products of the merging firms
 - The DOJ and FTC used the theory in a few cases in the 1960s, but the theory never gained any meaningful traction
 - The theory has been essentially dormant in the United States since before the mid-1970s
 - Modern antitrust law almost surely will be hostile to entrenchment theories as too speculative to establish the requisite “reasonable probability” of anticompetitive harm necessary to prove a Section 7 violation
 - *Example: Amgen/Horizon Therapeutics*¹ (covered in Class 26)
 - The FTC alleges that the deal would allow Amgen to leverage its portfolio of blockbuster drugs to entrench the monopoly positions of Horizon medications used to treat two serious conditions, thyroid eye disease and chronic refractory gout
 - The FTC alleges that Amgen to use rebates on its existing blockbuster drugs to pressure insurance companies and pharmacy benefit managers (PBMs) into favoring Horizon’s two monopoly products, thereby reducing demand for alternative drugs and reducing the incentives of other drug companies to develop them.

¹ See [Complaint for a Temporary Restraining Order and Preliminary Injunction Pursuant to Section 13\(b\) of the Federal Trade Commission Act, FTC v. Amgen Inc.](#), No. 23-CV-3053 (N.D. Ill. filed May 16, 2023).

Analyzing Competitive Effects in Downstream Markets

Horizontal merger challenges almost always focus on anticompetitive effects in downstream markets. Despite a major push by both the FTC and DOJ to find mergers they believe will have anticompetitive effects in upstream (input) markets, the agencies have brought only one case (Bertelsmann Penguin Random House/Simon & Schuster). Accordingly, we will focus our attention in this unit on competitive effects in downstream markets.

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 downstream markets, 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 the consumer when it results 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^*

The “*” indicates that the variable is at its profit-maximizing level

2. Determine and graph the net incremental profit for a firm increasing output by some amount Δq , including—

- ❑ The gross gain in profit from the increase in output, and
- ❑ The gross loss in profit from the reduction of price for sales at the original price

3. Derive and graph an inverse demand curve given a demand curve

If q is the quantity produced (output), Δq (read “delta q ”) is an incremental change in q

Consumers and Demand Curves

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 point (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 pairs (p_i, q_i) that satisfy the demand function traces out the *demand curve*

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

■ Notation: Δq means the change in q and is read “delta q ”

So:

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

■ The symbol “ \equiv ” means a definition

■ The change Δq is negative because demand declines as price increases


- 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




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

Inverse demand curves

- *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	Price
q	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
Conversely, if the market price is \$12, then consumers will demand 8 units

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:* A common nonlinear form of the demand curves used in econometrics is the log-linear demand curve: $\ln(q) = a - b \ln(p)$
 - This is equivalent to:


$\ln(x)$ is the natural logarithm of x

$$q = e^a p^{-b},$$

where a and b are constants and e is the natural logarithm

- A mathematically (although not necessarily an empirically) attractive property of a log-linear demand curve is that the own-elasticity is constant at all points on the curve:

$$\varepsilon = \frac{dq}{dp} \frac{p}{q} = \left[-be^a p^{-b-1} \right] \left[\frac{p}{e^a p^{-b}} \right] = -b$$




This is all you need to know on this slide

We will only use linear demand curves in this course!

Nonlinear demand curves

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



Aggregate consumer demand

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

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

$\sum_i q_i(p)$ 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”
 - Technically, 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 market
 - Residual demand is a critical concept in antitrust economics
 - You will encounter this concept frequently as the course proceeds

This is an important concept


Some technical points

- 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

- *Fundamental assumption #1*: Firms maximize their profits
 - Profits (π) = Revenues (r) – Total costs (t)
- To analyze the conditions under which a firm maximizes its 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



Revenues

- Example: Say the inverse demand curve for blue widgets is $p = 20 - q$
 - Revenues: $r = p * q$

Quantity	Price	Revenue
q	p	$r = p * q$
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


Marginal revenue

- *Definition:* Marginal revenue is the net additional revenue the firm 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

Marginal revenue

- *Example 1:* Say the *inverse demand curve* for blue widgets is $p = 20 - q$
 - Marginal revenue (discrete case): $mr(q) = r(q+1) - r(q)$

Quantity	Price	Revenue	Marginal revenue
q	p	r	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 and graph show—the orange line is the *marginal revenue curve*)

Marginal revenue

- Example 1: We can see this graphically by looking at the inverse demand curve for blue widgets



Marginal revenue

- *Example 1: 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 given the downward slope of the demand curve
 - 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

Marginal revenue

■ Example 2: Bob's Widgets

- 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


- *Example 2: Bob's brute force calculation illustrated:*

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

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

- 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

- *Example 2: Bob's more nuanced calculation illustrated:*

- Summarize the variables we know from the problem:

$$\begin{array}{lll} q_1 = 4 & q_2 = 5 & \Delta q = q_2 - q_1 = 1 \\ p_1 = 8 & p_2 = 7.5 & \Delta p = p_2 - p_1 = -0.5 \end{array}$$



- 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 = $7.5 - 2 = 5.5$



In other words, marginal revenue is revenue gained on the marginal sale minus revenue lost on the inframarginal sales

Marginal revenue

- Example 2: Summary of Bob's two methods of calculating marginal revenue




$$\begin{aligned}
 q &= 4 & \text{Area A} &= p_2 = 7.5 \\
 p_1 &= 8 & \text{Area B} &= \Delta p q = 0.5 \times 4 = 2 \\
 p_2 &= 7.5 & mr(4) &= \text{Area A} - \text{Area B} \\
 \Delta p &= 0.5 & &= 7.5 - 2 = 5.5
 \end{aligned}$$

$$\begin{aligned}
 p &= 10 - 0.5q \\
 r &= 10q - 0.5q^2 \\
 r(5) &= 50 - 12.5 = 37.5 \\
 r(4) &= 40 - 8 = 32 \\
 mr(4) &= r(5) - r(4) = 37.5 - 32 = 5.5
 \end{aligned}$$

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 in the course
- The general formulation of the more nuanced approach



Area A = Gross revenue gain by selling an additional unit

$$= p_2 \text{ times } 1$$

Area B = Gross revenue loss on preexisting sales due to the need to lower price

$$= (p_2 - p_1) \text{ times } q = \Delta p q$$

$$mr(q) = \text{Area A} - \text{Area B}$$

$$= p_2 - \Delta p q$$

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 ($2b$) 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

Remember, b is negative, so the curve slopes downward



Maximizing revenue


■ Example:

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

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

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)

Demand and Marginal Revenue



Revenue and Marginal Revenue



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 distance 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 $q = 10$), revenue is at its 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 $t(q) = \text{fixed cost } (f) + \text{variable cost } v(q)$

- *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

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

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

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:

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

Marginal costs

Marginal cost (mc): The cost mc of producing the $(q + 1)^{\text{th}}$ 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):

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

Profits and Profit Maximization

Profit maximization

- Profit maximization

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

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


- 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:

$$\begin{aligned} t(q) &= f + v(q) \\ &= f + cq \end{aligned} \quad \text{in the case of constant marginal costs}$$

Profits

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

$$\begin{aligned}\pi(p) &= r(q) - t(q) \\ &= (a + bq)q - cq \\ &= [aq + bq^2] - cq\end{aligned}$$

where marginal costs are constant and there are no fixed costs

- Graphically


where:

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

$$f = 0$$

$$c = 4$$

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:

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

- 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.

Be sure to learn this term!

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: $p = 10 - q$

MR/MC Diagram



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

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: $p = 10 - \frac{1}{2}q$
- 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 from 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 q times the required price decrease ($= q\Delta p$)
- So

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

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

Incremental revenue

- Graphically



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

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

Incremental revenue

■ Example

- (Inverse) demand: $p = 10 - \frac{1}{2}q$
- Starting point: $q_1 = 4$
- End point: $q_2 = 8$

You need to calculate these variables:


$$\text{So } p_1 = 8$$

$$\Delta q = q_2 - q_1 = 8 - 4 = 4$$

$$\text{So } p_2 = 6$$

$$\Delta p = p_2 - p_1 = 6 - 8 = -2$$

Incremental Revenue Analysis



Incremental revenue = Area A – Area B

$$\text{Area A} = p_2 \Delta q = (6)(4) = 24$$


$$\text{Area B} = q_1 \Delta p = (4)(-2) = -8$$

$$\text{So } IR = 32 - 8 = 16$$

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 ($\Delta q < 0$)



Incremental profits

■ Example: Output increase

- (Inverse) demand: $p = 10 - \frac{1}{2}q$
- Starting point: $q_1 = 2$
- End point: $q_2 = 6$
- Constant marginal cost $c = 4$


So $p_1 = 9$

So $p_2 = 7$

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

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

$$\begin{aligned} \text{Margin } m_2 &= p_2 - c \\ &= 7 - 4 = 3 \end{aligned}$$



Incremental profits = Area A – Area B

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

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


$$\text{So } \Delta \pi = 12 - 4 = 8$$

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

Incremental profits

■ Example: Price increase

- (Inverse) demand: $p = 10 - \frac{1}{2}q$ 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

$$\text{Area B} = q_2 \Delta p = (9.5)(0.25) = 2.375$$

$$\text{Area A} = m_1 \Delta q = (1)(-0.5) = -0.5$$

$$\text{So incremental profits} = 2.375 - 0.5 = 1.875$$

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 have 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):

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

Σ 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