MERGER ANTITRUST LAW

LAWJ/G-1469-05 — Tuesdays and Thursdays, 3:00-5:00 pm
Georgetown University Law Center — Dale Collins
Fall 2021 — wdc30@georgetown.edu
www.appliedantitrust.com

READING GUIDANCE

Class 9 (September 28): Competition Economics (Unit 7)

This week we start the formal analysis part of the course with two classes on competition economics. Economics is central to merger antitrust analysis for predicting the effects on competition of a pending merger. As you know from earlier in the course, most significant mergers are subject to the reporting and waiting period requirements of the HSR Act. This means the likely competitive effects of a merger must be predicted in advance of the transaction rather than observed after the closing. It is economics that provides the tools for predicting these effects.

There is no required reading for this class. However, you should study the class notes carefully, and the homework assignment will test your understanding of some of the basic concepts.[1] As always, read this reading guidance first.

Fundamental assumptions

Two fundamental assumptions drive predictive economics in merger antitrust law:

1. Each firm maximizes its profit
2. Each firm faces a downward-sloping demand curve for its products, so that the higher the price, the less quantity the firm sells

Profit maximization. A firm's profit is the difference between the revenues it earns from the sale of its products and services minus the costs the firm must pay to produce these products and services. Antitrust law assumes the firms in the economy maximize their respective profits. Indeed, it is at least arguable—and the DOJ has taken this position in its briefs—that the Supreme Court has made this a conclusive presumption, so that courts cannot entertain evidence or make findings to the contrary.[2] In the same vein, courts have held that not-for-profit enterprises, such as some hospitals, are presumed (at least rebuttably) to operate to maximize

[1] For those of you who have taken a basic course in microeconomics, the concepts in this unit should be straightforward. For those of you for whom the concepts are new, they make take a little more study. If you do not understand a concept after reading the reading guidance and the class notes, you might want to search the Internet (particularly YouTube) for additional explanations. The Internet is replete with helpful materials. I have included a number of links in the supplemental materials. You might find a different approach to the explanation useful. We will also have an optional economics review session on this coming Friday.

[2] *See* Final, Corrected Brief of Appellant United States of America at 49-53, United States v. AT&T, Inc., 916 F.3d 1029 (D.C. Cir. 2019) (No. 18-5214) (citing Copperweld Corp. v. Indep. Tube Corp., 467 U.S. 752 (1984)).

their profits, even if they reinvest their free cash flow into the business and do not distribute it as dividends to their shareholders.[3]

Courts, enforcement agencies, and the parties can use the profit maximization assumption to predict the likely competitive effects of a merger. In the first instance, the analysis can focus on the profit-maximizing incentives of the merging firms separately and ask whether the incentives of the combined firm are likely to differ. We will see later in the course that in the standard economic models used in merger antitrust analysis, when two firms that compete with one another for customers in some product line, it will be in the profit-maximizing incentive of the combined firm to increase price above premerger levels absent some offsetting price-reducing forces. Of course, there remains the important question of how significant this price-increasing incentive will be, since merger antitrust laws prohibit only those mergers that are reasonably likely to *substantially* lessen competition. Often, even in these standard models, the magnitude of the merged firm's profit-maximizing price increase will be so small as to be competitively immaterial.

But we are getting ahead of ourselves. Before turning to the mechanics of profit maximization, we need to examine the second assumption: demand is downward sloping.

Downward-sloping demand. This is often called the *law of demand*. The idea is simple: all other things being equal, the higher the price for a given product or service, the less customers will purchase. Conversely, again all other things being equal, the lower the price, the more customers will purchase.[4]

So, for example, if the price of cupcakes were to increase, consumers would buy fewer cupcakes. In this case, the reduction in demand resulting from a price increase comes from two sources:

---

3 *See, e.g.*, FTC v. Univ. Health, Inc., 938 F.2d 1206, 1224 (11th Cir. 1991); United States v. Rockford Mem'l Corp., 898 F.2d 1278, 1285 (7th Cir. 1990); Hosp. Corp. of Am. v. FTC, 807 F.2d 1381, 1390 (7th Cir. 1986); FTC v. Butterworth Health Corp., 946 F. Supp. 1285, 1296 (W.D. Mich. 1996) (but considering evidence related to performance of nonprofit hospitals in denying a preliminary injunction in a hospital merger), *aff'd*, 121 F.3d 708 (6th Cir. 1997); United States v. Mercy Health Servs., 902 F. Supp. 968, 989 (N.D. Iowa 1995), *vacated*, 107 F.3d 632 (8th Cir. 1997)

4 For the economists among you, you will recognize the "all other things being equal" (*ceteris paribus*) assumption means that this is a partial equilibrium analysis where the prices of all other products remain constant. In some situations, as customers shift out of spending money one product because of a price increase, they begin buying more of another product. In a more general equilibrium analysis, this increase in demand for the second product will result in a price increase of that product. This superficially gives the appearance of upward-sloping demand, but what is really happening is that the demand curve for the second product is shifting to the right as a result of the substitution effect with the product whose price is being increased. If the price of the first product had been held constant, the second product would exhibit downward-sloping demand per the standard assumption.

In theory, however, it is possible in a partial equilibrium model for a product to exhibit upward-sloping demand. These products are called *Giffen goods* and whether they actually exist has long been a subject of debate. In any event, if they exist, their occurrence is very rare. The product must be an inferior good with no close substitutes and consume a significant portion (but not all) of the consumer's income. As the price of the inferior good increases and so takes more of the consumer's fixed income, it squeezes out the consumer's demand for higher priced, normal goods. The classic example is an inferior quality staple good. As the price of the inferior staple increases, consumers can no longer afford to supplement their diets as much with higher quality foodstuffs (whose price has not changed), and to maintain their caloric intake they purchase more of the inferior good. In formal economic terms, the income effect is greater than the substitution effect.

1. Some customers will simply stop buying cupcakes altogether and spend their money on something else.
2. Other customers will continue to buy some cupcakes, just not as many as they would have bought at the lower price.

A few definitions will be helpful to our discussion of demand.

The demand for the product that would be lost with a price increase is called the *marginal demand*. The quantity of the product that customers still demand after the price increase is called the *inframarginal demand*. If customers only purchase at most one unit of the product (say, the product is an automobile), then the *marginal customers* are those that would purchase the product at the original price but would not purchase the product at the higher price. Likewise, *inframarginal customers* are those that would purchase the product at either the original price or the higher price. Economists often speak of marginal customers when they mean marginal demand, and inframarginal customers when they mean inframarginal demand, relying on the context to make clear what they really mean.

The maximum price a given customer is willing to pay for a given quantity of product i is called the customer's *reservation price*. If q_i is offered to a customer on a take-it-or-leave-it basis, then the customer will buy q_i if and only if the offering price is less than or equal to the customer's reservation price

For each quantity q_i of product, there is a price p_i that clears the market. A market *clears* when the demand for the product equals the amount of product the firms in the market produce. In other words, if p_i is the *market-clearing price* for q_i, then customers want to buy no more and no less than q_i of the product, and profit-maximizing firms in the market want to produce no more and no less than q_i. The market-clearing price is often called the *equilibrium price*.

The relationship between the market-clearing price p_i and the associated quantity q_i demanded is called the *demand function*. We can write this as $q_i = d(p_i)$. This equation simply says that if we know the market price p_i, then the demand function $d(\ )$ will tell us the demand q_i that customers will demand at that price. (This is independent of what firms would produce at this price.) The collection of all pairs (p_i, q_i) of the price and its associated quantity demanded traces out the *demand curve*.[5] The second fundamental assumption implies that the demand curve is downward sloping.

The Demand Curve

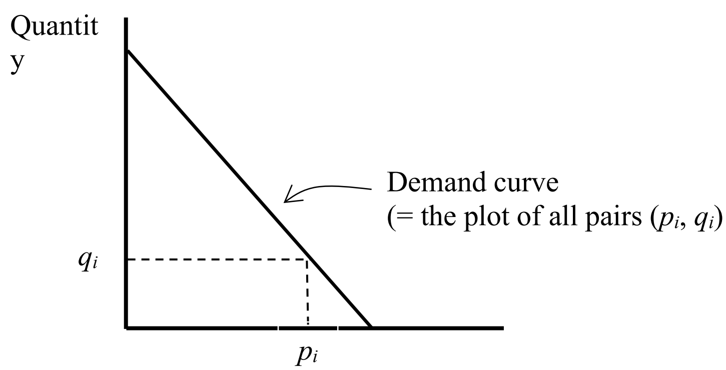

[5] The pair (x,y) is known as *Cartesian coordinates* and are plotted with the first number on the horizontal or x-axis and the second number plotted on the vertical or y-axis.

Pric
e

More often than not, it is more productive to look at the *inverse demand function*, which identifies the market-clearing price p_i for a given quantity q_i. This is written $p_i = p(q_i)$. The demand function and the inverse demand function contain exactly the same information. The *inverse demand curve*, like the demand curve, traces out the price and its associated quantity demanded, but the axes on the graph are reversed, with quantity rather than price on the y-axis and price rather than quantity on the x-axis. If the demand curve is downward sloping, then the inverse demand curve will also be downward sloping.

The Inverse Demand Curve

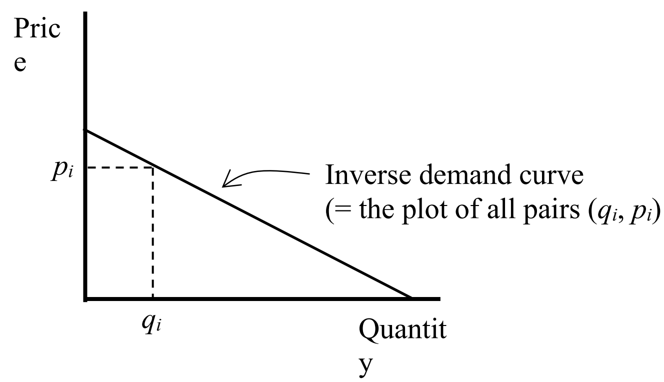

With customary sloppiness, economists often say the demand function when they mean the inverse demand function. I surely will use the wrong terms repeatedly throughout the course, but you should be able to tell by the context what I mean. In most cases, it is likely to be the inverse demand curve.

The demand function for the market as a whole is called the *aggregate demand function*. Each firm has its own demand function called the *residual demand function*.

One other important point. How sharply the inverse demand curve declines depends on the attractiveness of available products that customers are willing to substitute for the product in question as the price of that product goes up. If there are no good substitutes available, then even a moderate increase in price will result in only a slight decrease in the quantity demanded, so the slope of the inverse demand curve will be steeper. Conversely, if good substitutes are available, then even a small increase in price will result in a large decrease in the quantity demanded, making the inverse demand curve less steep. As we shall see later in the course, the reason this property is important is that much anticompetitive behavior turns on making substitutes less available or attractive, and so makes the firm's residual inverse demand curve steeper.

Other assumptions

To keep the analysis and exposition simple, we will make some other assumptions. However, these assumptions do not generally affect the qualitative results of the analysis and may be relaxed in particular circumstances to better fit the facts of the case.

1. The products in the market produced by different firms are fungible, that is, they are indistinguishable from one another in the minds of customers. Gasoline or ready-mix concrete are examples. These are often called *homogeneous products*. By contrast, *differentiated products* are products that differ from one another in the minds of customers, typically due to differences in features, quality attributes, or the seller's brand reputation. Automobiles, which vary by make, model, and features, are an example. We will examine models of differentiated products later in the course.
2. In a homogenous product market, all firms charge the same price. These are called *single-price markets*. If there are search costs to finding the product, however, even homogeneous product markets can support multiple prices.[6] We assume that there are no search costs in our simple models.
3. All economic activity takes place in a single period. These are called *static models*. By contrast, models that take place over time and involve multiple periods are called *dynamic models*. Since the antitrust analysis of mergers typically focuses on short-run competitive effects, most merger antitrust economics involve static models.
4. All markets clear and firms do not hold inventory of products. This means that firms sell everything they produce in the single-period model.

Profit-maximizing mechanics

We will spend most of the class on the mechanics of profit maximization for a single firm. The firm's profits are its *revenues* from the sale of the product minus its *total costs* of producing the product. We can write this algebraically as:

$$\pi(q) = r(q) - t(q),$$

where

q	is the quantity produced by the firm
$\pi(q)$	are the profits of the firm producing quantity q
$r(q)$	are the revenues of the firm producing quantity q
$t(q)$	are the total costs of the firm producing quantity q

Given our assumptions of a single-price market, revenues are the market price of the product times the quantity of the product sold by the firm:

$$r(q) = p(q)q,$$

where

$p(q)$	is the price at which the firm clears its production q (given by the inverse demand function)

The firm's total cost is the sum of its *fixed costs* for plant, equipment, personnel, and other costs that do not vary in the short run depending on the firm's production level and its *variable costs* that depend on the quantity of the product the firm produces:

---

6 *See, e.g.*, Kathy Baylis & Jeffrey M. Perloff, *Price Dispersion on the Internet: Good Firms and Bad Firms*, 21 Rev. Indus. Org. 305 (2002).

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

where

f	the fixed costs of the firm
$v(q)$	the variable costs of the firm producing quantity q

Marginal cost is the cost to the firm of producing one *additional* unit of output.[7] As we will see, marginal cost is a critical concept in antitrust economics. Algebraically:

$$\begin{aligned} m(q) &= \left[f + v(q+1)\right] - \left[f + v(q)\right] \\ &= v(q+1) - v(q), \end{aligned}$$

where

$m(q)$	is the firm's marginal cost when producing quantity q
$v(q+1)$	is the firm's variable cost when producing quantity $q+1$
$v(q)$	is the firm's variable cost when producing quantity q

Two things to note about marginal cost. First, the marginal cost at q is the firm's incremental cost of producing the (q+1)th unit, even though the firm has not produced this unit. Second, as we can see, the firm's fixed costs cancel, so that marginal cost is simply the difference of the variable costs.[8]

An important assumption we will make in the course is that a given firm's marginal costs are constant throughout the range of production of interest. We can write this as:

$$m(q) = c,$$

where

c	is the firm's constant marginal cost

Putting this all together, we see:

$$\begin{aligned} \pi(q) &= r(q) - t(q) \\ &= p(q)q - [f + cq], \end{aligned}$$

Now, without any significant loss of generality, we can make one other simplifying assumption: the firm faces a *linear residual demand curve*, that is, the demand function is a straight line. We can show that if the demand function is a straight line, then the inverse demand function is also a straight line. The function form for a straight line is $y = \mathrm{a} + \mathrm{b}x$, so a linear inverse demand curve has the following form:

$$p(q) = a + bq.$$

[7] Do not confuse this with the cost of producing the *last* unit. When marginal costs are not constant, the cost of producing the last unit and the cost of producing one more additional unit can differ.

[8] For the mathematically inclined, you will recognize marginal cost as the derivative of the total cost function at the quantity q where marginal cost is being evaluated. Equivalently, it is the slope of the total cost curve at quantity q.

Since demand is downward sloping, the parameter b is a negative number. Substituting the linear inverse demand function into the profit equation yields:

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

The details of the above algebraic simplification are not important. What is important is that the curve for final equation is a parabola:

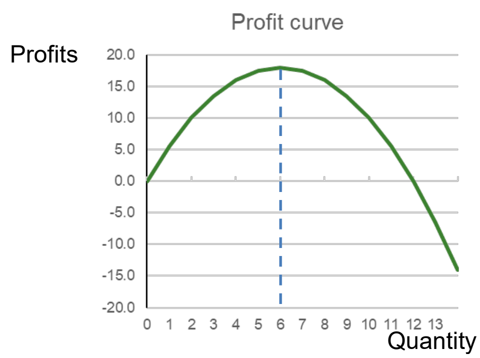

where (in this example)
$p(q) = 10 - \frac{1}{2} q$
$f = 0$
$c = 4$

Note that the profit curve has a maximum at the top of the curve. In this example, maximum profits occur when the firm produces six units. Note that if the firm produces less than six units, it can increase its profits by producing an additional unit. Conversely, if the firm produces more than six units, it can increase its profits by producing one less unit. Define *marginal profit* to be the profit made by producing one more unit (or more generally as the slope of the profit curve). As we can see from the graph, marginal profit is positive at production levels less than six units and negative at production levels of more than six units. The marginal profit is zero at the profit-maximizing production level of six units. Thus, the *first-order condition* for a profit maximum is that marginal profit is zero. We will invoke this property of a profit maximum throughout the course.

Key result #1: At the profit-maximizing level of output, the marginal profit of the firm is zero.

We can unpack this further. Define *marginal revenue* to be the incremental revenue earned by the firm if it sells one additional unit.[9] Marginal profit, then, is simply marginal revenue minus marginal cost. From the first-order condition, marginal profit is zero at the profit-maximizing level of output. Accordingly, at the profit-maximizing level of output, marginal revenue is equal to marginal cost.

[9] More precisely, it is the first derivative of the revenue function.

Key result #2: At the profit-maximizing level of output, marginal revenue is equal to marginal cost.[10]

Operationally, this means that the firm chooses a level of output so that its marginal revenue is equal to its marginal costs.

Finally, we can disaggregate marginal revenue into two components. First, the firm would gain additional revenue from the sale of the additional unit. But since the demand curve is downward-sloping, the firm cannot sell this additional unit at the original price when the firm was producing output *q*. Rather, it must lower its price to clear the market with the firm's expanded production of $q+1$. So the gain from the sale of the additional unit is the lower price. Second, since we assume that markets support only a single price, the firm must also lower the price of the *q* units the firm originally sold. This gives us our third key result:

Key result #3: Marginal revenue is equal to the gain in revenue from selling an additional unit at the new lower price minus the loss in revenues resulting from lowering the price of the existing units.

An equivalent way of saying this is that marginal revenue is equal to the gain in revenue from the sale of the marginal unit minus the loss in revenue from the lower price charged for the inframarginal units.

We can generalize this analysis in two important ways. First, whether or not the firm is operating at its profit maximum, the relationship of the firm's marginal revenue to its marginal cost will indicate whether an increase or a decrease in production would be profitable. If marginal revenue is more than marginal cost, the firm should increase production. If marginal revenue is less than marginal cost, the firm should decrease production:

Key result #4: If $mr(q) > mc(q)$, then the firm should increase production
If $mr(q) < mc(q)$, then the firm should decrease production

Second, we can ask the same question and get an analogous result for changes in *incremental revenue* and *incremental costs* if the firm is thinking about changing its production level by something more than one unit.

An example may be helpful. Say Sam's Cupcakes can sell 120 cupcakes at \$2.50 per cupcake, Sam's can only make cupcakes in batches of 12, and each batch has a constant marginal cost of \$18 per tray or \$1.50 per cupcake (these are spectacular cupcakes!) If Sam's increases its production to 132 cupcakes, it would have to lower the price to \$2.40 to sell out. Should Sam's increase its production to 132 cupcakes? Assume that the fixed cost of production is zero (Sam's uses the home kitchen).

---

[10] In calculus terms:

$$\pi(q) = r(q) - t(q)$$

$$\frac{d\pi}{dq} = \frac{dr}{dq} - \frac{dt}{dq} = 0 \quad \text{as the first-order condition for a profit maximum}$$

$$\text{So } \frac{dr}{dq} = \frac{dt}{dq} \quad \text{that is, marginal revenue equal marginal cost at the profit maximum}$$

We can determine the answer in one of two ways. One way, which I will call the *brute force method*, is to simply calculate the profits at each production level and compare them:

Profit = Revenue minus cost
Profits at 120 units = 120 × $2.50 minus 120/12 × $18 = $120.00
Profits at 132 units = 132 × $2.40 minus 132/12 × $18 = $118.80

So the firm would lose $1.20 from increasing production from 120 to 132 cupcakes. It should not increase production.[11]

The second way, which I will call the *incremental method*, is a bit more cumbersome but better illustrates the underlying economics:

Incremental profit = Incremental revenue minus the incremental cost

= Revenue gain from the incremental sales at the new price
minus the revenue loss from lowering the price for the inframarginal sales
minus the incremental cost

Revenue gain from the incremental sales at the new price = 12 × $2.40 = $28.80 (Area A)
Revenue loss from the price for the inframarginal sales = ($2.50 - $2.40) × 120 = $12 (Area B)
Incremental cost of producing additional units = 12/12 × $18 = $18 (Area C)
So incremental profit = $28.80 - $12.00 - $18.00 = -$1.20 (= Area A – Area B – Area C)

We see that the firm loses $1.20 in profit by increasing production (which the brute force method also indicates), so the firm should not increase production. The diagram below illustrates the incremental method:

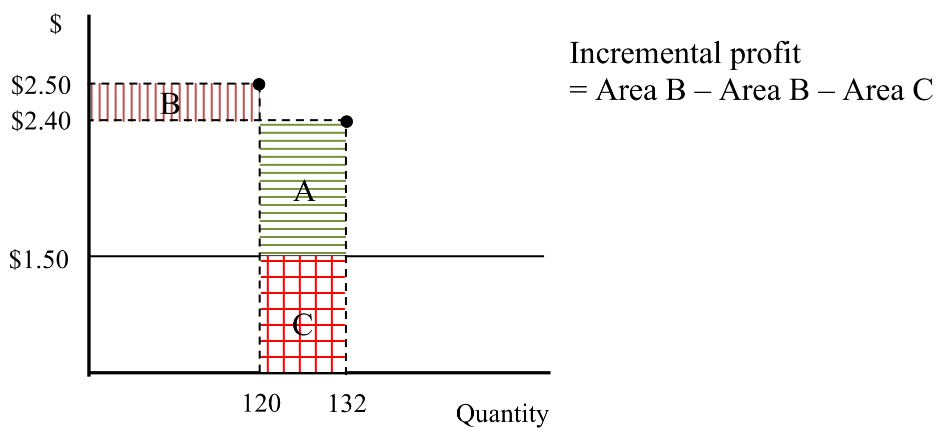

Now read the part 1 of the Unit 7 class notes: Demand, Costs, and Profit Maximization. The notes contain more detail, but if you understand the concepts in this reading guidance, you will be fine.

If you have any questions or comments, send me an e-mail.

---

[11] Say if the firm decreased production by one tray, it could increase its price to $2.60 per cupcakes. Should the firm decrease production?