

Lecture Notes
on
Market Definition, Concentration, and Advertising
(Revised: July 2017)

1 Markets and Market Definition

Much of this course will deal with the analysis of markets, and the consumers and producers that buy and sell in those markets. We will be particularly concerned with market structure and its implications for strategic decision making. Therefore, it is important to remember what we mean when we refer to a market. A market is the “place” where price is determined. In other words, *a market is the collection of buyers and sellers that, through their actual or potential interactions, determine the price of a product or set of products.*

Note the words *potential interactions*. A market can exist even if buyers and sellers don’t meet and interact. A good example is the market for gold, which is a *world market*. If you wanted to buy some gold, it wouldn’t matter whether you bought it in New York, London, or Zurich — after adjusting for exchange rates, the prices will be the same. Why? Because gold is easy to transport. If the price in New York were higher in London than in New York, people could make unlimited amounts of money by arbitraging the difference — they would buy gold in New York, hop on a plane, and sell it in London at the higher price. It is this potential for *arbitrage* that ensures that the market for gold is indeed a world market.

Why should we care about market definition? It is important for a number of reasons. From the point of view of a company, it is crucial to understand who your actual and potential competitors are with respect to the various products you now sell or might sell in the future. It is likewise important to know the product characteristic boundaries and geographical boundaries of one’s market in order to be able to set price, determine advertising budgets,



Figure 1: Market Boundaries

or make capital investment decisions. Market definition is likewise important from a public policy perspective. Should the Justice Department allow a merger or acquisition involving companies that produce similar products, or should it challenge it? The answer depends on the impact of that merger or acquisition on future competition and prices, but often this can best be evaluated in the context of a proper definition of the market.

In defining a market, we must decide what products to include. For example, is there a single market for digital cameras that includes expensive single-lens reflex (SLR) cameras as well as simple point and shoot cameras, or are these separate markets? Likewise, is there a single market for bicycles, or does it make more sense to think of “private label” bicycles (bicycles made by Huffy and Schwinn, for example, that are sold in stores such as Target and Wal-Mart) as belonging to a separate market from the more expensive bicycles (Giant, Cannondale, Trek, etc.) sold by dealers?

In addition to deciding what products should be included in the market, it is often necessary to decide on the geographical boundaries. If we wanted to define a regional market for gasoline centered around MIT, how large an area should we include? If we are interested in pricing and competition in the airline industry, should we consider city-pairs as the relevant markets for airline travel, or should we consider travel across different regions of the country, or, for that matter, travel within the entire United States?

To answer these questions, we need to consider both *demand substitutability* and *supply substitutability*. Remember that a market is the collection of buyers and sellers that determine the price of a product or set of products. If Products A and B have high demand substitutability *or* high supply substitutability, their prices are likely to equilibrate, so that they are part of the same market. It is important to understand why.

Demand Substitutability. When deciding whether Product B is in the same market as Product A, we can start by asking whether buyers of Product A might instead purchase Product B if the price of A rose substantially above that of Product B. In other words, are Products A and B close substitutes from the point of view of consumers—i.e., is there *high demand substitutability*. If the answer is yes, then the prices of Products A and B will be determined together, so that if the price of A falls, the price of B will also fall.

Recall that a measure of demand substitutability is the *cross-price elasticity of demand*. The cross-price elasticity of demand for product A with respect to the price of product B is

$$E_{AB}^D = \frac{P_B}{Q_A} \frac{\partial Q_A}{\partial P_B} = \partial \log Q_A / \partial \log P_B ,$$

where $Q_A(P_A, P_B)$ is the market demand curve for Product A. High demand substitutability means that this cross-price elasticity is high.

Exactly how high does the cross-price elasticity of demand have to be for us to conclude that two products are in the same market? There is no simple answer to this question. From the point of view of product pricing by companies, the question boils down to whether competing products are close enough substitutes so that a price increase of, say, 10 percent would cause enough consumers to switch to make the increase unprofitable. From an antitrust perspective, the question boils down to whether competing products are close enough substitutes to constrain a firm's market power. Thus judgments have to be made, as we will see shortly when we look at some examples.

Supply Substitutability. Even if there is little or no demand substitutability, two products could be in the same market if there is a *high supply substitutability*. Why? Because high supply substitutability between Products A and B means that an increase in the price of Product A would lead producers of Product B to shift some of their production to Product A,

and thereby limit the price increase.

Recall that supply substitutability is measured in the same way: The *cross-price elasticity of supply* for product A with respect to the price of product B is

$$E_{AB}^S = \frac{P_B}{Q_A} \frac{\partial Q_A}{\partial P_B} = \partial \log Q_A / \partial \log P_B ,$$

where in this case $Q_A(P_A, P_B)$ refers to the market supply curve for Product A.

As with demand substitutability, there is no clear threshold for the cross-price elasticity of supply that when crossed would put two products in the same market. Whether supply substitutability is sufficiently high for us to conclude that two products are in the same market is a judgment that has to be made on a case-by-case basis.

This is illustrated in Figure 1. Remember that *either* high demand substitutability *or* high supply substitutability is sufficient for two or more products to be in the same market. Therefore, when deciding on the boundaries of the market, we must decide whether other firms producing different products might enter if the price of the product or products currently in the market were to rise. Likewise, we must ask whether consumers might respond to a price increase by shifting their purchases to other products that we have not included in the market. Sometimes the answers to these questions are ambiguous, as we will see from the following examples.

Example: The Bicycle Industry. Given the fact that one can buy a bicycle for less than \$100 or for more than \$2000, does it make sense to analyze the industry in terms of a single market? Should we instead think of the industry in terms of ten or twenty different markets—for example, a market for bicycles costing up to \$100, a market for bicycles costing \$100 to \$200, and so on?

To answer these questions, we must consider how prices are determined. This, in turn, means examining the extent to which consumers will substitute one type (or brand) of bicycle for another in response to changes in relative prices, and the extent to which producers can shift production from one type of bicycle to another, again in response to changes in relative prices. Analyzing the industry in this way leads to the identification of two distinct markets.

(Management consultants would refer to these markets as “strategic groups.”) The first is the low-end bicycles, typically costing less than \$200, that are mass-merchandised through stores such as Target, Wal-Mart, Sears, etc. The biggest producers of these bicycles are Schwinn, Huffy, Mantis, and Mongoose. The second market or “strategic group” consists of the premium quality bicycles that are sold largely through dealers, i.e, stores that sell only (or mostly) bicycles and bicycle equipment.

Why does it make sense to think of these two groups as separate markets? First, demand substitutability across the groups is fairly low (but demand substitutability within the group is high). Most consumers who go to Target to buy a bicycle are not concerned with the price of a Cannondale or Trek, and likewise, most consumers choosing among high-end bikes at a dealer simply will not consider the Huffy and Schwinn bikes sold in a department store or mass merchandiser. To put this in terms of cross-price elasticities of demand, if the \$800 Trek bike went on sale for \$720, it is unlikely that many people shopping for a \$120 bike at Target will instead go to a dealer to buy the Trek bike. Likewise, if the \$120 bike at Target went on sale for \$100, it is unlikely that many people looking at Trek and Cannondale bikes instead rush over to Target. Think about your own purchase decisions. If you are a serious cyclist (or like to think of yourself as one), you would probably go to a bicycle dealer, where it would be difficult to find a bike costing less than \$400, and you could easily spend far more. (But of course you would have been happy to spend more, because you are serious cyclist.) On the other hand, if you just wanted a functional bicycle to get you from A to B, you would have done well by going to a store like Target or Wal-Mart, where you could easily find a decent bike costing around \$100 to \$200.

Second, supply substitutability across the two groups is low. Huffy and Schwinn compete by trying to shave a dollar or two off their production costs, taking advantage of scale economies wherever possible. The manufacturers of “premium” bicycles, on the other hand, compete along dimensions of technology and features. It would be very difficult for a manufacturer of one type of bicycle to shift production to the other type.

Figure 2 shows how bicycle producers can be grouped along two dimensions—the nature of the distribution channel, and product quality. The difference in distribution channels is

TYPE OF BICYCLE	COMPANIES AND PRICES (2011)
Mass Market Bicycles: Sold by mass merchandisers such as Target, Wal-Mart, Kmart, and Sears.	Huffy: \$90—\$140 Schwinn: \$140—\$240 Mantis: \$129—\$140 Mongoose: \$120—\$280
Dealer Bicycles: Sold by bicycle dealers – stores that sell only (or mostly) bicycles and bicycle equipment.	Trek: \$400—\$2500 Cannondale: \$500—\$2000 Giant: \$500—\$2500 Gary Fisher: \$600—\$2000 Mongoose: \$700—\$2000 Ridley: \$1300—\$2500 Scott: \$1000—\$3000 Ibis: \$2000 and up

Figure 2: Two Distinct Markets for Bicycles

another reason why both demand and supply substitutability are low across the two groups. Note that “mass market” bicycles, made by companies such as Huffy, Schwinn, and Mantis, are priced as low as \$90 and rarely cost more than \$250. These companies are focused on producing functional bicycles as cheaply as possible, and typically do their manufacturing in China. “Dealer” bicycles, the ones sold in your local bicycle store, include such brands as Trek, Cannondale, Giant, Gary Fisher, and Ridley, and are priced from \$400 and up way up. For these companies the emphasis is on performance, as measured by weight and the quality of the brakes, gears, tires, and other hardware. (Mongoose straddles both markets; they produce mass market bicycles costing as little as \$120, but also high-quality dealer bicycles costing \$700 to \$2000.)

Example: Airline Travel. In the case of airline travel, geographical considerations can make market definition fairly complex. On the demand side, consumers usually want to go from one city to another. A consumer who wants to go from Boston to Chicago will not be interested in the price of a ticket from Boston to Atlanta. Hence there is very little demand substitutability across routes, which suggests that city-pairs should be treated as individual markets. For example, based on demand substitutability alone, airline travel between Boston and Chicago and airline travel between Boston and Atlanta would clearly be viewed as two

distinct markets.

This picture can become complicated, however, because in some cases we might want to combine nearby cities. For example, Miami and Ft. Lauderdale are close enough together so that someone traveling from Boston to Miami might fly into Ft. Lauderdale instead if the fare were significantly lower. Hence we would view Boston-Miami and Boston-Ft. Lauderdale as belonging to the same market.

This picture is also complicated by the fact that there is a certain degree of supply substitutability across various regions. This is due in part to the hub-and-spoke system that has evolved over the past few decades. A hub-and-spoke system provides airlines with considerable economies, allowing them to serve a large network of cities more efficiently. (As we will see later in this course, it can sometimes be a source of market power for airlines.) But the hub-and-spoke system also creates supply substitutability. For example, an airline that serves 30 cities out of its hub could easily shift service from one city to another if the fare (and hence profit) differential were sufficiently large. Even without hubs, a high density of routes across certain regions of the country could provide airlines with the ability to substitute service from one route to another at fairly low cost. As a result of this supply substitutability, it is sometimes useful to define markets in terms of traffic flows across regions of the country, or in terms of flows into and out of hubs. Indeed, airlines often look at markets this way for purposes of pricing and strategic planning.

Example: The Sale and Service of High-Technology Equipment. Several different manufacturers sell high-capacity copying machines, several different manufacturers produce and sell diagnostic medical imaging equipment (e.g., CT scanners and MRI scanners), and similarly for other types of high-technology equipment. In most cases, each manufacturer does most of the servicing of its own equipment, although independent service organizations (ISO's) will sometimes provide service as well.

Is there a separate market for the service of each manufacturer's brand of equipment? For example, does the servicing of General Electric's CT scanners comprise a separate and distinct market, as opposed to the servicing of Picker International's CT scanners? Or does

competition among GE, Picker, and other manufacturers determine the prices for the sale and service of each company's scanners? This is of antitrust concern when an ISO sues a manufacturer (e.g., for restrictive practices), but it is also of concern for the manufacturer—it is crucial in the setting of prices for both sales and service.

At issue, once again, is the question of just how price (or prices) are determined. Do customers look only at the sales price when they make a purchase decision, or do they also consider the cost of service down the road and calculate a “life-cycle” cost of owning and operating the equipment? If the latter is true, then the prices for sales and service are intertwined, and competition across manufacturers will determine each company's sales price as well as its service price. In effect, the “product” is capability to do precision CT scans at all hours of the day or night, which means that the “product” includes both the sale *and* service of a CT scanner. In that case, the market would include all of the manufacturers, together with any ISO's that are also involved in the service portion of the business.

Example: Pharmaceuticals. New drug development is expensive. It begins with considerable expenditures on R&D, then requires various stages of laboratory and clinical testing, and if the new drug is finally approved, involves marketing, production, and sales. Then, the company must determine the price of the new drug. Pricing depends to a large extent on the characteristics of the consumers who will be buying the drug, as well as the number and characteristics of *competing* drugs. Hence, pricing a new drug requires a good understanding of the market in which that drug will be sold.

In the pharmaceutical industry, market boundaries are sometimes easy to determine, and sometimes not so easy. Markets are often defined in terms of *therapeutic classes* of drugs. An example is the class of selective serotonin reuptake inhibitor (SSRI) anti-depressants, which includes Prozac (the very first SSRI, manufactured by Eli Lilly), Zoloft (Pfizer) and Paxil (Glaxo).¹ These three drugs all do pretty much the same thing (reduce the body's absorption of serotonin). While their side effects and interactions with other drugs differ somewhat, they are all substitutes. Hence when Pfizer sets the price of Zoloft, it must be

¹All three of these drugs are now off-patent, and thus are also sold in generic form.

concerned not only with the willingness to pay of patients (and their insurance companies), but also with the prices and characteristics of the competing drugs. Likewise, a drug company that is considering whether to develop a new anti-depressant drug knows that if it undertakes R&D investment and succeeds, it will have to compete with at least these existing drugs. The company can use this information to project its potential revenues from the new drug, and thereby evaluate the investment.

Sometimes pharmaceutical market boundaries are more ambiguous. An example is painkillers. There are many types of painkillers, and some are more efficacious for certain types of pain (e.g., headaches, arthritis, muscle aches, etc.) than others. Strength and side effects likewise differ. Examples of painkillers include aspirin, acetaminophen (sold under the brand name Tylenol but also sold generically), ibuprofen (sold under such brand names as Motrin and Advil, but also sold generically), NSAID's (Non-Steroidal Anti-Inflammatory Drugs) which include naproxin (sold by prescription, but also sold over the counter under the brand name Aleve) and Voltaren (a prescription drug), and narcotic painkillers such as codeine. While some types of painkillers are used more frequently than others for certain conditions, there is considerable spillover. For example, depending on the severity of the pain and the pain tolerance of the patient, a toothache might be treated with any of the painkillers listed above. Hence the boundaries of a "painkiller market" or set of "painkiller submarkets" can be hard to define.

2 Measures of Market Concentration

In this course, we will often be concerned with *market concentration*. Usually this will pertain to the supply side of the market — are there just two or three firms that dominate supply, and thus can affect price through their monopoly power, or do many, roughly equal-sized firms compete? In some cases we will also be concerned with concentration on the demand side of the market — are there just a few buyers, for example, that do most or all of the purchasing, and thus can affect the market price through their monopsony power, or are there many small buyers? For simplicity, we will focus here on the supply side of the market.

We want to be able to describe the extent of market concentration in some quantitative way. Generally, there are two types of “statistics” that we use to do this. The first is a *concentration ratio*. It simply measures the total combined market share of some number of the largest firms. Most widely used is the *4-Firm Concentration Ratio*, CR(4), which is the combined market share of the four largest firms in the industry:

$$\text{CR}(4) = \sum_{i=1}^4 S_i ,$$

where S_i is the market share of Firm i . Likewise, the two-firm concentration ratio, CR(2), is the combined market share of the two largest firms in the industry.

Concentration ratios are useful, but somewhat limited in terms of their information content. Consider two industries, A and B, both with a CR(4) of 80 percent. However, the four largest firms in industry A each have a market share of 20%, while in industry B the largest firm has a share of 65% and the next three firms each have shares of 5%. Clearly we would think of industry B as far more concentrated than industry A, even though both industries have the same CR(4).

A more useful statistic (and one that is much more widely used) is the *Herfindahl-Hirschman Index* (HHI). It is simply the sum of the squared market shares for all of the firms in the industry:

$$\text{HHI} = \sum_{i=1}^m S_i^2 ,$$

where once again S_i is the market share of Firm i .

Note that it is always the case that $0 \leq \text{HHI} \leq 1$. In a (theoretical) perfectly competitive industry (thousands of tiny firms), the HHI is approximately zero, and for a monopoly it is 1. Suppose, for example, that there are four firms in the industry with market shares of 40%, 30%, 15%, and 15% respectively. In that case, the HHI is given by: $\text{HHI} = (.4)^2 + (.3)^2 + (.15)^2 + (.15)^2 = .295$.²

To compute an HHI, one must first address the question of market definition — we need to know what firms and products to include when calculating market shares. (We will consider

²In merger and other antitrust applications, the DOJ and FTC usually express the market shares as percentages, so that the HHI ranges from 0 to 10,000, and in this example would be given by $\text{HHI} = (40)^2 + (30)^2 + (15)^2 + (15)^2 = 2,950$.

this problem when we examine competition in the beer industry.) In addition, the most appropriate meaning of “market share” is not always clear. Most often, we use the shares of sales, but in some cases shares of capacity, or something related to capacity, may be more relevant. Consider, for example, an extractive resource industry. Two or three firms might have small shares of current production and sales, but might have relatively large shares of proven resource reserves. Thus the market power of those firms might be small in the short run, but larger in the longer run when their greater reserve base will allow them to dominate the industry. We will discuss issues of this sort as they arise.

3 Optimal Advertising Decisions

Once we have defined the market, and hence know whom we are competing against, we can think about pricing and advertising. Here we will focus on advertising decisions. I will assume that the firm has estimates of two numbers: its price elasticity of demand, e^P , and its advertising elasticity of demand, e^A . Let’s be clear about what these elasticities are.

Price Elasticity of Demand. We are interested here not in the *market elasticity*, but rather the *firm’s price elasticity of demand*. Don’t confuse these two numbers. The market elasticity, which we’ll denote by E^P , is the percentage change in the quantity sold if *all firms in the market* raise their prices by 1 percent. The firm’s price elasticity of demand, e^P is the percentage change in *the firm’s sales* if the firm raises its price by 1 percent, and *all other firms in the market leave their prices unchanged*.

For many products, there are good estimates of the market elasticity of demand. For example, we know from statistical studies that the market elasticity of demand for mass market beer (i.e., Budweiser, Miller, Coors) is about -0.8 . Determining the firm’s elasticity of demand, however, can be more challenging. The reason is that it will depend on how the firms compete with each other. For example, do they compete very aggressively (so that there are frequent price wars), or is competition among firms “softer?” Shortly I will discuss how we can estimate the firm’s elasticity e^P if we know the market elasticity E^P and the number of firms competing. For now, let’s just assume that we have an estimate of e^P .

Advertising Elasticity of Demand. The advertising elasticity, e^A , is the percentage change in the quantity the firm will sell resulting from a 1-percent increase in the firm's advertising expenditure. We can write this as:

$$e^A = \frac{A}{Q} \frac{\partial Q}{\partial A} = \frac{\partial \log Q}{\partial \log A}$$

We might also be interested in the *cross-firm advertising elasticity of demand*, which is the percentage change in the sales of Firm 1 that results from a 1-percent increase in the advertising expenditures of Firm 2. We can write this as:

$$e_{12}^A = \frac{A_2}{Q_1} \frac{\partial Q_1}{\partial A_2} = \frac{\partial \log Q_1}{\partial \log A_2}$$

For many products, we would expect that $e_{12}^A < 0$, i.e., if Firm 2 does more advertising (presumably claiming its product is better than the competing product of Firm 1), it will take sales away from Firm 1. This is usually the case — but not always. Here are two ways in which the advertising of Firm 2 could *increase* the sales of Firm 1, so that $e_{12}^A > 0$:

(1) *Educate the Public About a New Product.* Suppose Merck and Pfizer have both introduced a new generation of anti-cholesterol drugs. The two drugs compete with each other, but both firms face resistance from doctors and patients because this type of drug is new and not yet widely used. Merck might advertise that its drug is safe and very effective, which could boost its sales. But by educating doctors and patients, it will boost overall sales of *both* drugs. This “spillover” effect could easily outweigh the brand preference effect, so that the net effect is to increase the sales not just of Merck's drug, but also the sales of Pfizer's drug.

(2) *Differentiate the Products.* Will Tide laundry detergent clean your clothes better than Cheer (a competing brand)? How would you know? Aren't all laundry detergents essentially the same, in which case you should buy the cheapest one? For most consumers, the answers to these questions come from the advertising they have seen. If they have seen more ads for Cheer, they may believe that Cheer will do a better job in cleaning their clothes, and therefore be willing to spend more for Cheer than a competing (and possibly cheaper) brand. But look what is happening here. If Tide, Cheer, and the other brands of laundry detergent

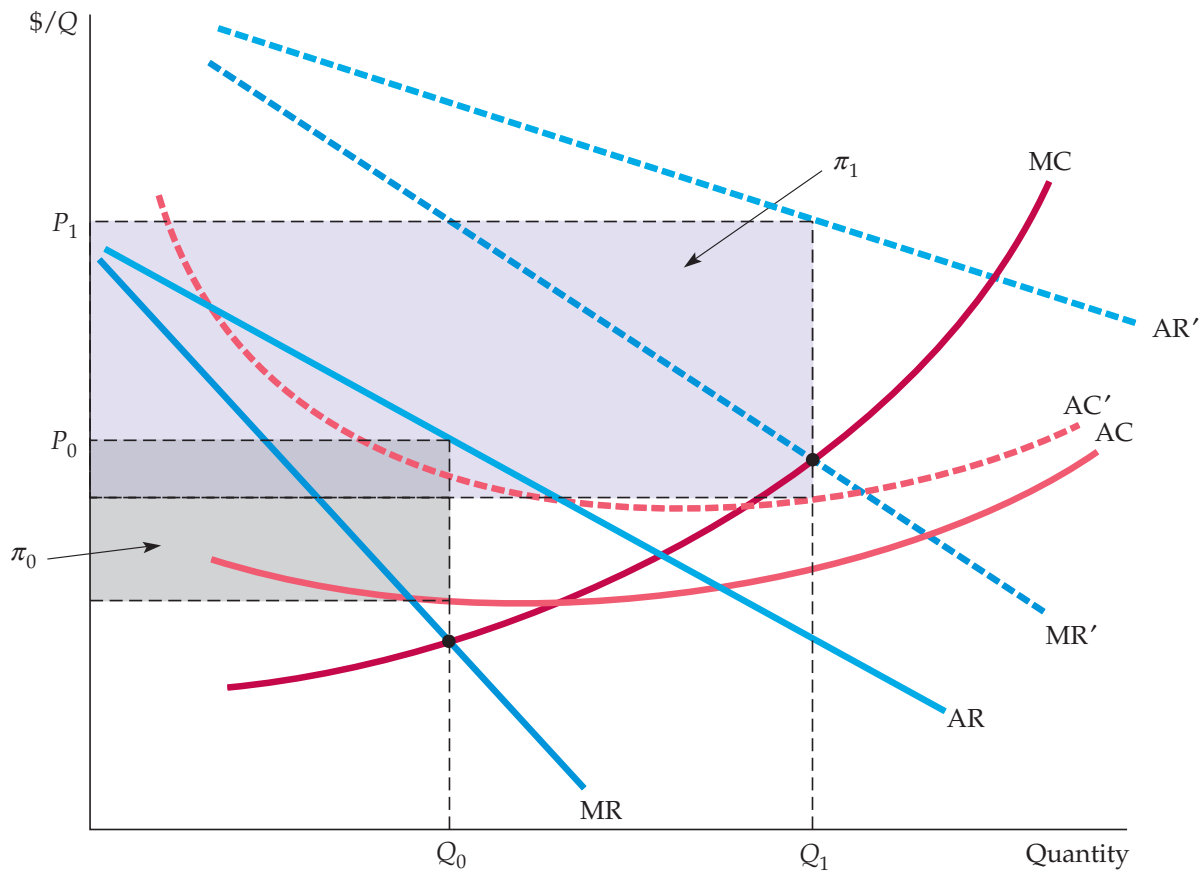


Figure 3: Effects of Advertising

are in fact the same (in terms of how well they work), the only way to get consumers to pay more is to convince them that *they are not the same*, i.e., to differentiate the products. That is what advertising does. The result is that ads for Tide can actually boost sales of Cheer, and vice versa. By convincing consumers that these brands are different, consumers are willing to pay more for their “favorite” brand, whatever that brand happens to be. As we will see, this kind of product differentiation applies to mass market beer.

3.1 How Much to Advertise?

Figure 3 shows how advertising can (if it is successful) shift the firm’s demand curve to the right so that it can charge a higher price, sell a larger quantity, and earn a larger profit. Very nice, but how much advertising should the firm do? This question was addressed

in Section 11.6 of Pindyck and Rubinfeld, *Microeconomics*, where a “rule of thumb” for advertising was derived that maximizes profit, which is now given by:

$$\pi = PQ(P, A) - C(Q) - A$$

Note that the quantity demanded now depends not only on the price P , but also on A , the dollar amount of advertising.

A key point is that the firm should *not* advertise up to the point that an extra dollar of advertising generates an extra dollar of revenue, because that would ignore the fact that increased sales implies increased production cost. Instead, the firm should advertise up to the point that the marginal revenue from advertising should equal the *full* marginal cost of advertising:

$$\text{MR}_{\text{Ads}} = P \frac{dQ}{dA} = 1 + \text{MC} \frac{dQ}{dA}$$

This in turn yields the following rule of thumb for the advertising-to-sales ratio:

$$\frac{A}{PQ} = -\frac{e^A}{e^P}$$

The profit-maximizing advertising-to-sales ratio is just the ratio of the firm’s advertising and price elasticities of demand (with a minus sign because the price elasticity is negative).

It probably seems obvious to you that the advertising-to-sales ratio should be higher if the advertising elasticity of demand is higher, i.e., if advertising is more effective in stimulating demand. But why should A/PQ be higher when demand is less elastic? Shouldn’t a less elastic demand mean that the firm has less need to advertise? The key here is that a less elastic demand implies that the profit-maximizing price will be well above marginal cost, which in turn means that every unit of extra sales generated by advertising will bring in more profit. Put simply, a lower price elasticity of demand implies more “bang for the buck” from advertising.

Of course to apply this rule of thumb for advertising, we need estimates of the two elasticities. Also, note that this rule of thumb ignores possible cross-firm effects of advertising. As we will see in the context of mass-market beer, those cross-firm effects can be important.

3.2 Advertising Mass Market Beer

Have you ever watched the Super Bowl on television? It seems as though it's 50% football and 50% beer commercials. That's an exaggeration, but there is clearly a great deal of advertising for Budweiser and other mass market beers. And remember that there is no other television show with a higher per-minute price for commercials. Why is so much money spent to advertise beer? Is it consistent with the rule of thumb for advertising? To find out, we need to know what are e^A and e^P for a brand like Bud or Miller.

Price Elasticity of Demand. Statistical studies say that the *market* price elasticity is about -0.8 . But what is the price elasticity for a typical brand? The answer depends on the number of competitors, N . If firms compete à la Cournot, and the firms are of equal size, then it turns out that the price elasticity of demand for the firm is N times the market elasticity, i.e., $e^P = N \times E^P$. (See my Lecture Notes on Pricing for a detailed derivation and explanation of this result.)

Competition in attribute space is *local* — e.g., Bud competes with Miller, but not with Sam Adams or Magic Hat #9. This means that for a typical brand (e.g., Miller), there are roughly $N = 3$ or 4 firms competing in attribute space. If we use $N = 3$, the price elasticity for one of these brands is roughly $e^P \approx (3)(-0.8) = -2.4$. This means that if the price of a single brand is increased by 1% and the prices of all other brands remain the same, demand for that brand will fall by 2.4%.

Advertising Elasticity of Demand. What is e^A , the advertising elasticity of demand for a typical brand? It turns out that this elasticity is quite low, somewhere between 0.1 and 0.2. Why is this? Because when people watch a commercial for, say, Miller beer, they go to the refrigerator and take a beer. And if the refrigerator is empty they might go to the local store and buy some beer. But which brand of beer? Whichever brand happens to be in the refrigerator, or if they go to the store, whichever brand they are used to buying. In other words, they are not moved to buy a Miller beer, but rather just to buy (or drink) *any* beer.

Using the middle of the range, $e^A = .15$. This says that A/PQ should be about $.15/2.4 \approx .06$. But A/PQ is in fact considerably higher than this, around .10 or .12. So it seems that

there is *too much advertising*. Why?

The answer is partly the cross-firm elasticity. Bud and Coors benefit when Miller advertises, Coors and Miller benefit when Bud advertises, etc. The reason is that mass market beer is in some sense like laundry detergent — they are all very much the same. (In fact, in blind tasting tests, consumers cannot tell the brands apart.) So advertising helps to differentiate the products. After watching commercials, consumers *think* that Bud and Miller are different. This means they will pay more for their “favorite” brand (even though in a tasting test they can’t tell their favorite brand apart from other brands). If *all* firms advertise, the relevant elasticity is the *market* advertising elasticity of demand, E^A , which will be larger than the brand elasticity e^A .

Hard versus Soft Competition. How should the mass market beer companies compete? Would you advise Miller to cut its prices in an attempt to win market share? If it did, how do you think Bud would respond? Price competition is destructive. Competing via advertising can be costly, but in the long run it is much less destructive.

Game theory tells us how to best play a game. You studied elements of game theory in 15.010, and perhaps in your strategy courses. But firms must also decide *which game to play*. Should they play a pricing game (“hard competition”) or an advertising game (“soft competition”)? What would be your advice?