Virtual Corporate Reality (VCR) Project

Econ 331 - Industrial Organization

> Business is a good game - lots of competition and a minimum of rules. You keep score with money. *Atari founder Nolan Bushnell*

> An economist is someone who sees something happen [in reality] and wonders whether it would work in theory. *Ronald Reagan*

You will participate in a game this semester as a member of a team and your team will be charged with running a firm in an abstract market. Twice a week, your firm will make decisions on the pricing of its product(s), product introductions, product withdrawals, product relocations, and cost-reducing investment. As a result of the decisions made by your team and the other teams that produce in the same industry, demand for the products will generate revenues and incur costs of production. Your goal will be to maximize the value of your firm.

To better understand the abstract industry in which you compete, we need to specify the nature of demand, costs, entry & exit, some financial details, and communication. A word at the outset: the model of this market is described in detail conceptually *and* mathematically in the next few sections. The "scenarios" area of the web site (described in Section 5) helps you get a better intuitive feel for the market by performing "what if" scenarios.

1 Demand

First, think of yourself as a customer in this market. The customers that buy this product only care about one product feature. Some examples are a cereal's sweetness, Chinese food's spiciness, the tightness in fit of a pair of jeans, or the location of a gas station. We can talk about customers' "location in product space" by, in the Chinese food example, assigning 0 to very bland food and 1 to the most spicy food imaginable. A customer "located at 0.45" means that he prefers his food to have spiciness level 0.45. In our simple model, there is no other way in which customers differentiate Chinese food.

We will also assume that customers are "evenly distributed" among these values so that, continuing the example, if there are 15 people that like foods of spiciness levels from 0.2 to 0.3 then there must also be 15 people that like spiciness in the range 0.8 to 0.9 (or any other interval of length 0.1), and the 15 people in each of these two ranges are evenly distributed in their ranges.

Customers prefer one particular type of product: a customer located at x thus would most prefer a product also located x. But if there is no product where the customer is, the customer will consider buying another product not too far away. The further away the product is, though, the less the customer will be willing to pay, as detailed below. Note that "further away" means "less similar," so it makes sense that the consumer will be willing to pay less for the product. Also, moving left and right are equivalent; all that matters (other than the products' prices) is how far from the customer's location the alternative products are.

For technical reasons, we will assume that the points 0 and 1 are located right on top of each other.[1] Thus it's a circle—not a line—model of demand. See Figure 1, and remember that there is only *one* dimension to product space, so you can't get to the interior of the circle. You can only move around the circle. It might be helpful to think of this as modeling demand for gas stations located on a beltway around a city like Washington, D.C. (But don't take for granted that the prices you set should match those that you see at gas stations in D.C. And ignore the negative externality created for a consumer from having a gasoline station located right around the corner from his home.)

Note that the locations 0 and 1 occupy the same point. If two products at w and y have the *same price*, then a customer at x prefers the one located at w because it is closer to x: $d(w, x) = 1.01 - 0.9 = 0.11$ is smaller than $d(y, x) = 0.2 - 0.01 = 0.19$.
If the prices of w and y are *not* the same, then the customer's choice takes this into account. In this example, if the price of y is *enough* lower than the price of w, then the customer at x will instead prefer y.

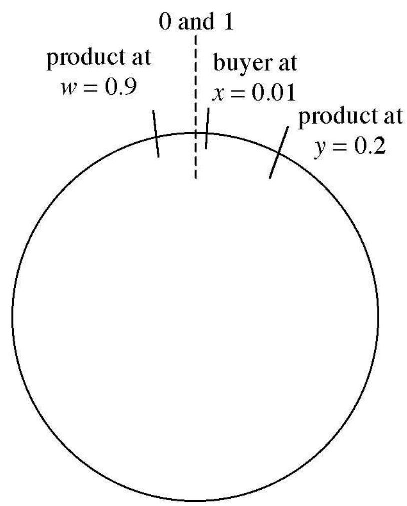

Figure 1: Some points in the circle model of demand.

Figure 1 considers a specific example in the circular model with a customer, located at $x = 0.01$, comparing two identically priced products, located at $w = 0.99$ and $y = 0.20$.

[1]The technical reason for using a circle model instead of a line model is that if the two ends are not attached then the points 0 and 1 are very different from all other points. By attaching the ends of the line to make a circle (and assuming that customers are evenly distributed along the circle), all points are exactly the same as each other.

If the two products are the same price, the customer at x prefers w because it is closer (0.11 units away from x) than y (0.19 units away from x).

How does the consumer take price into account? Read on.

1.1 Consumers' utility and their consumer surplus

Consumers place value V on the good that is exactly what they want. If someone located at x buys a product located at w and pays price p_w, she receives consumer surplus

$$V - 200d(w, x) - p_w \ . \tag{1}$$

The net benefit (utility) that the consumer receives from the product is equal to $V - 200d(w, x)$, where the function $d(w, x)$ is the distance between the product's location, w, and the consumer's location, x. The negative sign ensures that the value the customer receives decreases as the distance to the product located at w grows. The consumer's surplus from purchasing a product is the net benefit she receives minus the price she pays. (Hopefully that sounds familiar from Principles and Micro!) Price is the easy part: the lower the better. The consumer's net benefit (utility), though, depends on how close the product is to her location (her preferred product). The distance function $d(w, x)$ is always a positive number, so it doesn't matter to which side the product is, just how far. The formula for distance, $d(w, x)$, is simple if the shortest way to get from x to w doesn't pass over the 0 and 1 point: it's just $|w - x|$. If the shortest path around the circle from the consumer to the product crosses the 0 and 1 point, add 1 to the smaller of the two locations and subtract the larger. (See the caption to Figure 1 for examples.) V is the same for every product.[2]

You can also think of $d(w, x)$ as the "utility cost" to the customer of not buying her most ideal product. In order to get the benefit V from a product, the consumer pays a price <u>and</u> incurs a utility cost due to the product not being exactly what she prefers.

1.2 A consumer's decision between two products

Consumers only buy one product each period. If a consumer's location is between two products, how does he decide which one he prefers? *He buys the one that gives him the most consumer surplus.* Thus we can see from equation (1) that the decision depends on both the distance from the consumer to each product and on the price of each product. There are three possible cases. If a customer located at x is comparing a product located at w with price p_w to a product located at y with price p_y, then

[2] We don't need to assign a number to V, as you'll see later in equation (2). The details of why V is in equation (1) aren't too important, but note that its size determines whether or not consumer surplus is positive.

he prefers w if $V - 200d(w,x) - p_w > V - 200d(y,x) - p_y$;

he prefers y if $V - 200d(w,x) - p_w < V - 200d(y,x) - p_y$;

and he is indifferent between w and y if

$$V - 200d(w,x) - p_w = V - 200d(y,x) - p_y;$$

Rearranging this last case (indifference between w and y) gives

$$200d(w,x) + p_w = 200d(y,x) + p_y. \tag{2}$$

Thinking of $200d(w,x) + p_w$ as combining both the price of w and the utility cost of buying a less than ideal product, these equations say that the customer located at x buys the product which has the lowest of these *combined* costs. He is indifferent between the two products if equation (2) is satisfied.

Some important features to remember about the model of demand are the following.

- Consumers differ only in their ideal product (represented by the consumer's location on the circle).
- Customers are always located evenly (uniformly distributed) around the circle.
- Consumers don't change their positions—although firms can change products' positions.

1.3 Which consumers buy a product

Now shift perspective to the one which you will have for the rest of the semester.

You are choosing the price of a product located at y, and you want to know how much demand y will receive. The "scenarios" area of the web site you have is going to tell you this, but the description that follows gives the details of how the computer program calculates it.

Say the two products closest to y are located at x and z, as pictured in Figure 2. If the prices of the three products are not too different from each other, the only other two products that affect the demand for y are these two nearest ones.[3]

For example, say that there are six products in the industry, located at points 0.1, 0.3, 0.5, 0.65, 0.8, and 0.95 around the circle. This is not shown in Figure 2, but you can draw it yourself. First consider the product at 0.65: its nearest neighbors are those (going

[3] If one price is much lower than the others, then one or both of the other products can receive zero demand. The math is somewhat different from that presented below, but still follows the principal that each consumer buys the product that provides him or her with the most consumer surplus. The "scenarios" area of the web site takes all of this into account.

The two nearest products are located at x and z, and they may be your firm's or another's. Given these product locations and the prices of the products (p_x, p_y, p_z), we can calculate where the locations u and v are such that all customers between u and v (the bold portion of the circle) decide to purchase y. From the locations of u and v in this particular graph, we can see that $p_y < p_x$ (because more customers between y and x buy y) and $p_y > p_z$ (because more customers between z and y buy z).

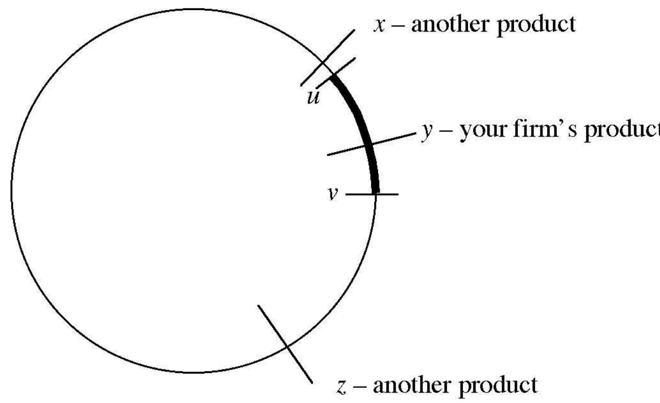

Figure 2: How much demand does a product located at y receive?

counter-clockwise) at 0.5 and (going clockwise) at 0.8. If we consider the product located at 0.1, then its nearest neighbors are located (going counter clockwise) at 0.95 and (going clockwise) at 0.3.

Returning to Figure 2, let p_x, p_y, and p_z be the prices associated with the products at x, y, and z, respectively.

Let's see first how many of the customers between x and y decide to buy y. If p_x and p_y are very different from each other, then all the customers between the two products will buy the lower priced one, but if the prices are not too different from each other, then some customers will buy x while some buy y. Assume that the prices are not too different. We want to find the location u (see Figure 2) at which a customer (located at u) is indifferent between products x and y because then we know that the customers closer to x than u will buy x and the ones closer to y than u will buy y. So in Figure 2, we're finding the bold portion of the circle between y and u: the demand for y that's counterclockwise from y. (Then we will find the rest of of the bold portion, that which is clockwise from y, between y and v.)

We know from equation (2) above that when products are located at x and y, then the location u where a customer is indifferent between them is described by

$$200d(x,u) + p_x = 200d(y,u) + p_y,$$

so solve this for the distance from y to the location of the customer who is indifferent between products x and y:

$$d(y,u) = d(x,u) + \frac{p_x - p_y}{200}. \tag{3}$$

That's the portion of demand counter-clockwise from product y that it receives (the bold area between y and u in Figure 2), but it's still got u in the equation. We can get rid of that by using the fact that the distance from x to y is equal to the distance from x to u plus the distance from u to y, or $d(x,y) = d(x,u) + d(y,u)$. Rearrange that to be

$$d(x,u) = d(x,y) - d(y,u). \tag{4}$$

Now substitute equation (4) into (3) and solve for

$$d(y,u) = \frac{d(x,y)}{2} + \frac{p_x - p_y}{400}, \tag{5}$$

which is the distance from y to the consumer indifferent between products x and y. And it's in terms of the things we know: prices and locations of the two products.

We can do the same thing for demand between x and z by talking about a point v that divides the demand between the two products. All consumers clockwise from y to v (the bold area between y and v in Figure 2) will buy y instead of z. Using a derivation similar to that for equation (5), this formula for the distance between y and v is

$$d(y,v) = \frac{d(y,z)}{2} + \frac{p_z - p_y}{400}. \tag{6}$$

Finally, add together equations (5) and (6) to get the distance around the circle within which consumers all buy y (all of the bold line in Figure 2):

$$d(u,v) = \frac{d(x,y) + d(y,z)}{2} + \frac{p_x + p_z - 2p_y}{400}. \tag{7}$$

In other words, this is the market share of product y.

1.4 The number of consumers in the *total* market

We will assume that the total number of consumers in the market, N, increases as the total number of products available, n, increases, but at a decreasing rate. This is based on the idea that more consumers will enter this market when there is a greater variety of products, but that there is some degree of saturation. The following function increases as n increases, for values of n that are not too large.

$$N = -110,000 + 120,000 \cdot n - 1,500 \cdot n^2 \tag{8}$$

1.5 The demand function

So we're there! We now know the quantity demanded of product y when its nearest neighbors are x and z, with prices p_y, p_x, and p_z, respectively. It's the fraction of the total consumers (market share) captured by y, multiplied by the total market size, using equations (8) and (7):

$$
\begin{aligned}
D(y, p_y; x, z, p_x, p_z) &= N \times d(u, v) \\
&= \left(-110,000 + 120,000n - 1,500n^2\right) \times \left(\frac{d(x,y) + d(y,z)}{2} + \frac{p_x + p_z - 2p_y}{400}\right).
\end{aligned}
$$

Let's check to make sure that this demand function has the properties that we expect all demand functions to satisfy.

- As the price of the product p_y increases, demand for it decreases (because the coefficient on p_y is negative).
- As either price p_x or p_z of the similar products (substitutes) increases, though, the demand for y increases (because the coefficients on p_x and p_z are positive).
- As the substitute x or z gets further from y (becomes less of a substitute), demand for y increases. It can also be shown that y's price elasticity of demand decreases as the distance to other products increases.

No surprises here.

As was mentioned in footnote 3, this demand function is only true when prices are not *too* different. If p_y is very high compared to p_x and/or p_z, there may be no demand for y. If p_y is very low, there may be customers *beyond* locations x and/or z that purchase y.

2 Costs

The firm cares about not only demand for its products, but also about the costs incurred in production as well. (In the discussion below, the terms "location" and "plant" are synonymous.)

Cost function: Each product is produced at constant marginal cost, and there is a fixed cost incurred each period per product. If a product is withdrawn ("srapped") from the market (see below), it no longer incurs the fixed cost. The fixed cost is \$500,000 per period, but the marginal cost depends on the level of investment spent on the product, as described next.

Cost-reducing investment: Marginal cost does not depend on q, the quantity produced, so it is *constant* marginal cost. But it does depend on, X, the *total* investment made in this product's manufacturing plant *in previous periods*. Marginal cost is given by

$$
\$50 \cdot \left(\frac{50,000,000}{X + 50,000,000}\right),
$$

so you can see that increased investment, X, decreases the marginal cost of production. To repeat, X is total investment made in *all* previous periods.

For example, if \$0 dollars were invested in the product during periods 1 and 2 and \$1,000,000 was invested during period 3 and again in period 4, then marginal cost in periods 1, 2, and 3 would be \$50. In period 4, marginal cost would be $50 \cdot \frac{50}{51}$ dollars. In period 5, marginal cost would be $50 \cdot \frac{50}{52}$ dollars.

Thus, the cost function for a product of which q is produced and in which a total of X has been invested in previous periods is its fixed cost plus its variable cost

$$C(q; X) = \$500,000 + \$50 \cdot \left(\frac{50,000,000}{X + 50,000,000} \right) \cdot q \, .$$

Make sure the following facts are clear.

- Investment is specific to a particular product, but you can invest in as many of your products as you like.
- There is a one-period delay between when you invest and when marginal cost falls.
- Investment is "private" information (see below), so you only announce it to me, not to the industry in general.
- Investment is irreversible and sunk, so once you spend the funds on investment they cannot be recovered.
- If a product is relocated (see below), *all investment in that product is lost.*

The initial products start with zero investment (X=0 in period 1), but you may invest in period 1 to bring down marginal cost in period 2. Of course, you may also invest more in later periods to further decrease marginal cost.

3 Product entry, exit, relocation, and auction sales

Initially, each firm in the industry will have one product, and these products will be spaced equally around the circle. So if there are five firms in the industry there will be products located at 0, 0.2, 0.4, 0.6, and 0.8. If there are four firms, the products will be at 0, 0.25, 0.50, and 0.75. Firms may also introduce, withdraw ("scrap"), and trade products during all but the final three periods.

3.1 Entry/Introduction

In any period, a firm can announce the introduction of one or more new products and specify the locations at which they will be available. In the period the announcement is made, these products have no demand (and do not affect other products' demand), so they do not generate revenue, but the firm does incur the cost of introducing them in that period. That cost, \$5,000,000, is half due to the cost of constructing a new plant and half due to advertising costs. The advertising cost is sunk (none of it is recoverable), but the manufacturing plant is assumed to have a resale value (see below) and to never have decreased production capacity.

The firm can make cost-reducing investment along with the introduction announcement, or it can wait until later periods. Thus, new products may have $X > 0$ in their first period of generating revenue (the period after the introduction is announced).

To reiterate:

- Production begins and revenues begin being generated the period *after* the announcement of a new product. So if a new product is announced in period t, the cost of introduction and any initial cost-reducing investment is incurred during period t, but production costs $C(q; X)$ and revenue do not start occurring until period $t + 1$.

3.2 Exit/Withdrawal ("Scrap")

A firm can also choose to withdraw any of its products at any time. Like introduction, a withdrawal announced in period t affects production and demand only starting at $t + 1$.

The manufacturing plant for the withdrawn product can be sold on the open market for \$1,000,000 ("scrapped"), or to another firm in an "auction" on the VCR site for as much as that firm is willing to pay. A general announcement will be made that the plant is available for purchase, and then other teams can make bids. As emphasized below, you can talk about this, but you *cannot* talk about product prices and other decisions. Other teams make bids for the location using the VCR site. The announcment area can also be used to justify proposed bids and an asking price.

The plant does *not* need to be bought by any other firm. If it is not bought by another firm the firm that is withdrawing the plant receives the \$1,000,000 market value.

Remember:

- Although the firm announces in period t that it will withdraw a product in $t + 1$, the firm must *still also set a price* in period t. Of course, in period $t + 1$ and all following periods, the product is no longer sold and no longer priced.

- If you want to ask for buyers of a plant, make the announcement (to me) before the period ends. Another firm must commit to buying the plant during the same period you announce the withdrawal. If the plant is purchased, then in period $t+1$ the *new* firm must set price for this location.

3.3 Relocation

A product can also be relocated. This costs $2,500,000 and all investment in cost-reduction is lost (X goes to zero). Think of it as the need to advertise the new location and to relearn or reinvest in any production techniques that reduce marginal cost. The relocated product, like a newly introduced product, may receive cost-reducing investment in the same period as the announcement of the relocation.

Note again:

- Like entry and exit, relocation occurs with a one-period delay. So along with announcing in period t the new location of a product for period $t+1$ and all future periods, the firm must also set the price at the *old* location in period t. In period $t+1$ and later periods the firm will set price at the *new* location.

4 Financial details

Your team is endowed with $25,000,000 cash (and one product, as stated above) at the beginning of the game. If the firm ever has a negative cash balance, you are declared insolvent and will be forced to restructure. Details will provided if this becomes necessary. There is no market for borrowing, so your only sources of cash are your initial assets, 5% interest received on cash, and your profits. Each period, your cash balance changes according to the rules outlined above and summarized here as the following credits [+] and charges [–].

+ interest of 5% on beginning of period cash balance

+ the revenue from all products

– the cost of production for all products' output

– cost-reducing investment for all products

– the number of product introductions × $5,000,000

+ the number of products withdrawn × $1,000,000 (or revenues received from their sale)

– the number of product relocations × $2,500,000

– the amount paid for any manufacturing plants purchased from other teams

+/– any other charges or credits as appropriate.

5 Procedures, software, and grading

So now you know all there is to know about the industry in which you will be participating and the financial structure of revenues and costs. This final section describes how to report decisions, what you can talk about and with whom, and the incentive to do well.

Communication The antitrust laws apply to your industry and firm. That means that you are not allowed to talk or otherwise communicate with any member of another team about price, investment, and/or product decisions. Any communication about these subjects are *the same as cheating on an exam* (see the course syllabus).

Some communication, though, is allowed. One example is given above in the discussion of selling a manufacturing plant. Note that you must already be committed to selling the plant before discussing it with another firm. Your (legal) discussions with other firms can take place either directly or through me.

I will also consider any other type of proposal that you wish to make to me.

You are not allowed to ask questions of me or anyone outside your group about strategy or any particular decision. Asking me for clarification of rules or procedure is allowed.

Scenarios You can perform "what if" analyses about the industry using the "scenarios" area of the VCR web site. It embodies all of the definitions given in this document about demand and cost. By specifying each firm's product location(s) and pricing, you can see what demand and revenues would be generated for your firm and all others. By specifying investment, you can predict your firm's profits. You can also save set-ups for reference and modification later.

Decisions You will be responsible for reporting your decisions on the VCR web site. You can change those decisions at any time before the deadline for each period.

You will need to specify the following each period.

- The prices of your products (which default to the last period's prices).
- The level of any additional cost-reducing investment for each product (which defaults to zero—but X is still the sum of all previous investment in each particular product).
- The location of any new products, the withdrawal of any existing products, and any relocations.
- Any other notes about changes that you would like to make (which can also be communicated via email before the deadline).

Important: *If you fail to submit any of the information by the required deadline, the "default" values will be used.* "Defaults" are your past pricing information; product entry, exit, relocation, and investment will not occur.

If there is no past pricing information (during the very first period of the game or in the first period after a product is introduced), demand will be zero but *fixed costs will still be incurred.*

Auctions During any period, some team may put a location up for auction. If your team does so, please send me an email message notifying me so that I can let the rest of the class know. Once you have decided to put a location up for auction, you cannot remove it.

If you wish to bid on a location being auctioned, you must do so before the end of the period, and you are responsible for monitoring others bids that may be higher than yours.

Public Results The following information about your industry will be available on the VCR web site as soon as I increment the period. The full history of this information will be available, so you don't have to keep track of it.

- The location and ownership of all products. This includes any announcements of new products that will generate revenue in the next period.
- A list of products that will be withdrawn and will generate no more revenue in the next period.
- Change in ownership will be evident from reviewing the past history of existing products but not indicated directly.
- Prices of all products in production that period.
- The profits and cash holdings of *all* firms

Private Results The following information is made available on the VCR site at the same time as the public results, but other firms cannot see it. (The full private history remains available to you throughout the game.) This private information is your firm's

- investment decisions,
- production costs,
- demand,
- revenues, and
- profits

broken out separately by product.

The final three decision rounds During the last three periods, no changes in products other than their prices will be allowed. No introductions, withdrawals, relocations, changes in investment, etc. Only prices may be changed in the last three periods. This allows future expected profits to be calculated, as described below.

Grading Your grade will be based on your final financial assets compared to levels which I expect you to reach based on the structure of the game.

Your final assets consist of both your cash and the present value of each manufacturing plant's future profit stream.

To estimate the future profit stream of each plant, I will assume that in every future period each product will earn the average of its profit over the final three periods, and I will apply a discount rate of 10% to these future profits.

For example, suppose that during the last three periods a product generated $250,000, $500,000, and $450,000 in profits, respectively. The average is then $(\$250,000 + \$500,000 + \$450,000)/3 = \$400,000$, and the present value of the plant producing this average forever in the future is $\$400,000/0.1 = \$4,000,000$.

If the present value of any plant's future profit stream is less than its $1,000,000 market (scrap) value, I will use the market value—so the value of each plant will be at least $1,000,000.

One final note Your grade does NOT depend on the size of your financial assets relative to the other teams in the industry in which you compete. Causing other firms to earn lower profits is *not* inherently better for your firm.