

8. Oligopoly

Some of this material can be found in Varian Chapter 15-16; see especially 16.5-16.10.

I. Overview

A. So far we have only looked at two extreme types of markets.

1. Competitive markets have only price-taking firms (presumably lots of them).
2. Monopolist markets have one firm with unilateral pricing power.

B. We now look at markets with firms that have some pricing power, but not unilateral.

Oligopoly: from Greek, more than one but less than “many.”

C. We use game theory to study behavior in oligopolies.

1. Firm decisions affect one another \iff strategic interaction.
 - a. Need an equilibrium concept that describes multiple agents trying to optimize.
2. Game theory yields many surprising conclusions. Here’s the first.
 - a. We solved the monopolist’s problem by describing its choice of quantities.
 - b. We could have just as easily (and with the same result) had the monopolist choose a price.
 - c. This symmetry disappears in our study of oligopoly: the equilibria turn out to be quite different.

II. (Normal Form) Game theory

A. A normal form (simultaneous play) game (NFG) is defined by three elements

1. A list of N players
2. A set of strategies for each player, $s_i \in S_i$. E.g., $S_i = \{s_{i1}, s_{i2}, \dots, s_{ik}\}$.

3. A payoff function for each player, $\pi_i(\mathbf{s})$, where the **profile** of all players' strategies is $\mathbf{s} = (s_i, \mathbf{s}_{-i})$.
- B. Let \mathbf{s}_{-i} be a vector of strategies of all players other than i . The **best response function** (or correspondence) is $BR_i(\mathbf{s}_{-i}) = \operatorname{argmax}_{s_i \in S_i} \pi_i(s_i, \mathbf{s}_{-i})$. In words, for a given profile \mathbf{s}_{-i} of other players' strategies, player i 's best response is the strategy (or subset of strategies) that maximizes his payoff.
- C. A **Nash equilibrium** is a strategy profile \mathbf{s}^* in which every player is making a best response to the other players' strategies, i.e.,

$$s_i^* \in BR_i(\mathbf{s}_{-i}^*), \quad i = 1, \dots, n. \quad (1)$$

- D. These definitions are quite general and apply in politics, biology, business, traffic engineering, etc. etc. Here we will apply them to oligopoly.

III. Quantity Setting: Cournot Markets

A. The Duopoly NFG

1. $N = 2$ players, called firms.
2. Strategy is the output quantity $y_i \in [0, \infty) = S_i$.
3. The choices y_1, y_2 are made simultaneously (logically speaking).
4. We'll keep things simple in computing the payoff functions (profit functions).
Set $Y = \sum_{i=1}^N y_i$ to be total output, and assume a linear inverse demand curve (for perfect substitutes)

$$p = a - bY$$

5. We'll also assume a linear cost curve, i.e., identical marginal cost c for all firms and zero fixed cost.

6. Then the profit to any firm i is:

$$\pi_i = (a - bY)y_i - cy_i = (a - c - bY)y_i. \quad (2)$$

7. N.B. Aggregate output quantity of other firms $Y_{-i} = Y - y_i$ affects firm i 's profits and therefore his optimal choice.

B. Best Response Function

1. Here the best response function BR_i describes firm i 's best choice of quantity y_i as a function of the quantity choices of everyone else.
 - a. Note this doesn't imply that firms actually *know* the quantity choice of others.
 - b. Action is simultaneous.
2. We will see that in this quantity setting game, $\frac{\partial BR_i}{\partial y_j} < 0$
 - a. If your rivals' total output increases, then you want to *decrease* your own output.
 - b. Quantity is a *strategic substitute*.

Ex: The FOC for (2) is

$$0 = \frac{\partial \pi_i}{\partial y_i} = a - c - 2by_i - bY_{-i}. \quad (3)$$

Solving for y_i , we get the BR function

$$BR_i(Y_{-i}) = \left[\frac{a - c}{2b} - \frac{1}{2}Y_{-i} \right]_+ \quad (4)$$

where $[z]_+ = \max\{0, z\}$. See Figure 1.

C. Nash Equilibrium

1. A Nash equilibrium is a profile of strategies at which no player has an incentive to change their behavior given what others are doing.

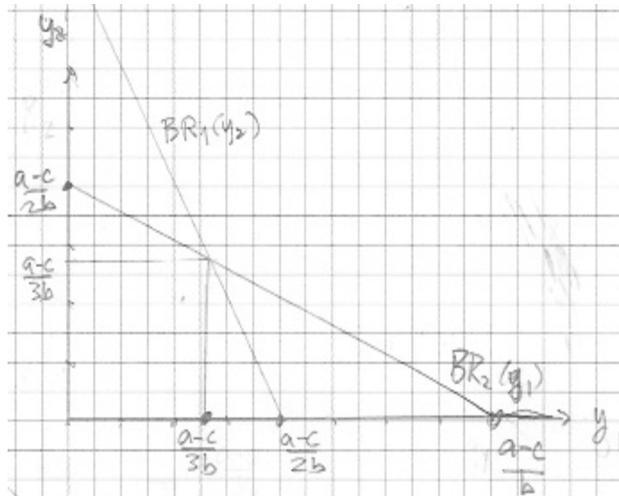


Figure 1: BRs and NE in Cournot duopoly.

2. Or, all firms are simultaneously playing their best response.
3. Adding the FOCs (3) for $i = 1, \dots, N$, we get

$$0 = N(a - c) - 2bY - b(N - 1)Y,$$

so total output in Cournot Nash equilibrium is

$$Y^* = \frac{N}{(N+1)b}(a - c), \text{ and by symmetry,}$$

$$y_i^* = \frac{Y^*}{N} = \frac{a-c}{(N+1)b}.$$

For duopoly, $y_i^* = \frac{a-c}{3b}$, and NE payoff is

$$\pi_i^* = (a - c - bY^*)y_i^* = (by_i^*)y_i^* = \frac{(a-c)^2}{9b}.$$

D. Asymptotics

1. Can we describe the equilibrium behavior of Cournot firms in terms of the number of firms (N)?
 - a. By doing this we can look at the relationship between oligopolies and both monopolies and competition.
2. Denoting s_i as $\frac{y_i}{Y}$, profit maximization gives us Marginal Revenue = Marginal Cost. Using familiar tricks on Marginal Revenue, we get

$$p(Y)(1 + \frac{s_i}{\epsilon}) = c'_i(y_i)$$

3. If all firms have the same constant marginal cost c and fixed costs that they can cover in Nash equilibrium, then $s_i = 1/N$ and

$$p(Y)[1 + \frac{1}{N\epsilon}] = c.$$

4. The main result is a price somewhere between competition and monopoly:
 - a. If $N = 1$ this price is just the monopoly price.
 - b. As N approaches infinity, price converges to the competitive level.
 - c. With N in between, price remains above marginal cost, but below the monopoly level.

E. A problem with Cournot Analysis

1. We usually think of firms setting price – not quantity.
2. Where do prices come from in Cournot markets?
 - a. They come from the inverse demand curve, but what does that mean?
 - b. The Cournot model implicitly assumes that firms just dump their output on the market and accept the market clearing price.
3. Bertrand (1883) criticized that assumption of the Cournot (1838) model, and advised assuming directly that firms set prices. As we will now see, the predicted (equilibrium) outcome is quite different.

IV. Price Setting: Bertrand Markets

A. Describing a Duopoly

1. Two firms simultaneously choose price, given a demand curve $D(p)$.
2. Assume constant marginal costs c_i , possibly different across firms.
3. Consumers will buy from the lower priced firm, since they produce perfect substitutes.

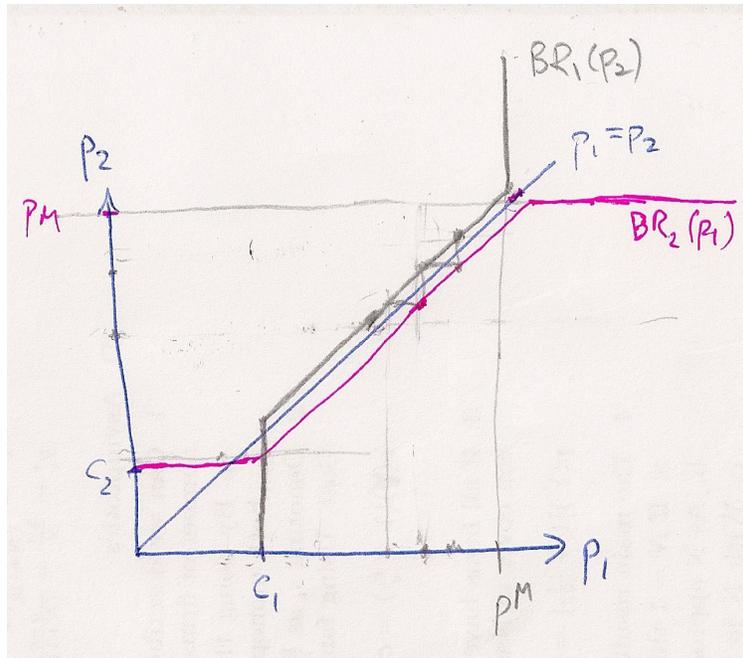


Figure 2: NE in Bertrand Model. Monopoly price is p_M , firms have marginal costs c_1, c_2 .

4. Firm i 's demand is given by

- a. $D(p_i)$ if $p_i < p_j$
- b. $D(p_i)/2$ if $p_i = p_j$
- c. 0 if $p_i > p_j$

B. Best Response

1. That discontinuity in demand leads to a discontinuity in payoff (profit) functions. Taking FOC's won't work well here.
2. The BR is to undercut the rival firm's price very slightly, but never price above the monopoly price or below own MC. See Figure 2.
3. Note now that the strategic variables – p_i – are *strategic complements*: the lower a rival's price, the lower you'd like your price to be.
4. By contrast, Figure 1 shows that in Cournot duopoly, the lower your rival's output, the *higher* is your best response output — a case of *strategic substitutes*.

C. Nash Equilibrium

1. A Nash Equilibrium in the Bertrand game is a set of prices at which no firm has an incentive to change his or her price.
2. Assume, with little loss of generality, that $c_j > c_i$.
3. Then in a Nash equilibrium:
 - a. $p_i = c_j$ (actually, a tiny bit lower)
 - b. $p_j \geq c_j$
4. The price in the market is competitive – it is equal to (the second lowest) marginal cost.
 - a. This is especially striking if we assume (as we often do) that firms share a common marginal cost.
 - b. If $c_i = c_j = c$ then we have $p_i = c$.
5. In the price setting game, then, the prediction is the competitive outcome *even with only 2 firms*.

D. The Bertrand Paradox

1. Some call that last result the “Bertrand paradox.”
2. Intuition tells us that, say, five firms should compete much more fiercely than two firms.
 - a. Indeed there is experimental evidence to this effect.
3. One way of to get less competitive behavior in pricing games with few firms is to feature repeated games.
4. Another is to have firms selling slightly different products, as we now consider.

V. Differentiated goods variants of Cournot and Bertrand

- A. More realistically, the goods are not *perfect* substitutes.

B. The Cournot variant writes the inverse demand functions for two goods like this:

$$p_1 = A_1 - B_1 y_1 - C y_2 \quad (5)$$

$$p_2 = A_2 - C y_1 - B_2 y_2 \quad (6)$$

- The cross price effect C must be the same for both firms by an obscure aspect of demand theory.
- Imperfect substitutes if $C > 0$, complements if $C < 0$, independent if $C = 0$.
- The special case $A_1 = A_2, B_1 = B_2 = C > 0$ reverts to perfect substitutes.
- The special case $C = 0$ reverts to two monopolies (in unrelated goods).

C. For constant marginal cost c for everyone, the usual optimization problem for BR has FOC that yields $y_1^* = \frac{A_1 - c - C y_2}{2B_1}$, and similarly $y_2^* = \frac{A_2 - c - C y_1}{2B_2}$.

D. Comparing the usual case $0 < C < B_1, B_2$ to perfect substitutes, you can see that the BR lines rotate outward, and NE outputs and profits increase, as substitutability C decreases.

E. Moral of story: firms prefer to differentiate their products.

F. There is a parallel analysis for Bertrand, using the direct demand functions. It is still true that prices are strategic complements while outputs are strategic substitutes.

G. These variants are more useful in applied work than the original pure substitute models.

VI. (Extensive Form) Game Theory: Basic ideas, assuming perfect information.

VII. Quantity Setting With A Leader: Stackelberg Markets.

- Draw "tree" for 2 player sequential game, each with three possible output levels. Solve by backward induction.

- Draw tree for 2 player sequential game, each with an interval $[0, p^M]$ of possible output levels. Solve by backward induction: last mover chooses BR. Previous mover predicts this, and optimizes.
- Take the duopoly example with linear demand and constant marginal cost, say $N = 2, a = 30, b = 1, c = 6$.
- For comparison, Cournot-Nash equilibrium is $y_1^* = y_2^* = \frac{30-6}{3} = 8$, price is $p^* = a - bY = 30 - 16 = 14$, and profits are $\pi_1^* = \pi_2^* = (p^* - c)y_i^* = 8 * 8 = 64$.
- The Stackelberg leader, firm 1, chooses output y_1 to maximize her profit, knowing how the follower will react. That is, she assumes that $y_2 = BR_2(y_1)$.
- Using equation (4), we see that $BR_2(y_1) = \frac{a-c}{2b} - \frac{1}{2}y_1 = 12 - \frac{1}{2}y_1$.
- Hence she solves

$$\begin{aligned} \max_{y_1 \geq 0} \pi_1(y_1, BR_2(y_1)) &= (a - c - b(y_1 + BR_2(y_1)))y_1 = (30 - 6 - (y_1 + 12 - \frac{1}{2}y_1))y_1 \\ &= (12 - \frac{1}{2}y_1)y_1, \end{aligned} \quad (7)$$

- which is easily seen to have solution $y_1^{SB} = 12$.
- Hence $y_2^{SB} = BR_2(y_1^{SB}) = 12 - \frac{1}{2}y_1^{SB} = 6$, and $p = 30 - (12 + 6) = 12$, so $\pi_1^{SB} = (12 - c)12 = 72$ and $\pi_2^{SB} = (12 - c)6 = 36$.
- Compared to Cournot NE: price is lower, profit for leader is higher, but follower profit and total profit are lower in Stackelberg NE.

VIII. Collusion and cartels

- If all the firms in an industry could agree on the total output level Y , what would they choose?
- This is just a restatement of the basic monopoly problem.
- But there are difficulties in implementation.

- Firms might disagree on how to split up Y into quotas.
 - Even if they agree in principle, it is true in practice that they all have an incentive to exceed the quota – the firm producing the excess gets all the benefit (pdy) but absorbs only a fraction of the cost Ydp in MR.
 - This is true whether or not the other firms stick to their quotas, as long as Y is less than in Cournot equilibrium.
- To enforce quotas, cartels need to reliably detect violations and punish them.
 - They also need to prevent new entry.
 - Theory of repeated games (not included in this course) says that a key is the discount factor, based on likelihood that the market will continue, and interest rates.
 - Typically even less efficient than simple monopoly; cartel members usually will not minimize total cost (or even equalize marginal costs).
 - Of course, cartels are discouraged by anti-trust laws, at least in major industrialized economies during the last century.
 - Historically speaking, most successful cartels needed strong government support.

IX. Kinked Demand Curve and Sticky Prices

- Paul Sweezy (1939) argued that most oligopolies work differently than in previous models.
- In industries ranging from autos to zinc [banks, cement, tobacco, tractors, ...], there was an established price \bar{p} that seldom moved.
- Firms anticipated that if they cut price below \bar{p} , their rivals would match them (as in the steeper part of the demand curve in Figure 3). But if they raised their price, nobody would follow them, so they would lose share rapidly (as in the flatter part of the demand curve in the Figure).

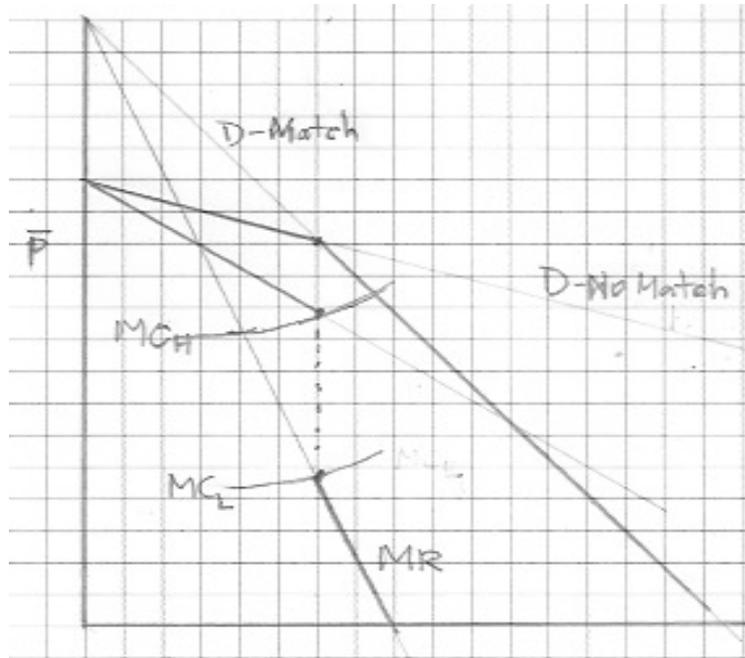


Figure 3: Demand and MR in KDC model.

- Consequently, MR has a discontinuous drop at $D(\bar{p})$.
- Profit-maximization therefore entails not changing quantity as long as MC lies between MC_L and MC_H in Figure 3). That is, prices are “sticky” — not responsive to moderate cost shocks.
- The theory is less popular today, partly because fewer industries now are like that, and partly because Sweezy’s followers never explained where \bar{p} comes from.
- My early paper “Producers’ Markets: A Model of Oligopoly with Sales Costs,” *Journal Economic Behavior and Organization*, 11:3 (May 1989) 381-398, offers an explanation of \bar{p} . Given its cost function, the battle for market share (via ads and other cost of sales), and the rational assumption of price matching on price decreases but not increases, the model shows that each each firm has a profit-maximizing value of \bar{p} . The lowest one among existing firms is the equilibrium value of \bar{p} , and that firm is the price leader.

X. Conjectural Variations.

Recall that, for homogeneous goods with inverse demand $p(Y) = p(y_1 + y_{-1})$, firm 1's problem can be written as $\max_{[y_1 \geq 0]} y_1 p(y_1 + y_{-1}) - c(y_1)$. The first order condition fully written out is

$$c'(y_1) = p(Y) + y_1 p'(Y) \left[\frac{dy_1}{dy_1} + \frac{dy_{-1}}{dy_1} \right] \quad (8)$$

$$= p(Y) + y_1 p'(Y) [1 + \nu], \quad (9)$$

where $\nu = \frac{dy_{-1}}{dy_1}$ is firm 1's *conjectural variation* — her belief about how a change in her output y_1 will affect the total output y_{-1} of all rivals.

- $\nu = 0$ is the Cournot conjectural variation. She takes as given her rivals' output level, and (incorrectly!) assumes that she can't affect it.
- $\nu = -0.5$ is the Stackelberg leader's conjectural variation in the simple linear duopoly. More generally, it can be the slope of other firms' summed reaction functions.
- $\nu = -1.0$ is the competitive or Bertrand conjectural variation. It ensures that price = MC, and says that the firm believes that other firms will replace any units it withholds from the market.
- $\nu = y_{-1}/y_1$ is the collusion conjectural variation – other firms will maintain their current share.
- “consistent” conjectural variations equate ν to the actual comparative statics of the model for each firm.

Economic theorists no longer find it fashionable to write down arbitrary expressions for ν , and the idea of consistent conjectures never got much empirical support. (But McGinty, 2016, may give it new life in the context of greenhouse gas abatement treaties.) Masters students now might find CVs useful in keeping track of various models of imperfect competition.

XI. Spatial Competition: Hotelling location models.

Let us now take a deeper look at imperfect substitutes. So far, we have taken as given substitution elasticities in utility functions and demand functions. We also noted (from the differentiated good Bertrand model) that firms tend to be more profitable when their products are less substitutable. How can we model that dimension of competition? Hotelling (1929) apparently was the first to take up that challenge. His “Main Street” model used a spatial metaphor to describe the substitutability among products. Think of the producers choosing the products’ characteristics (e.g., fuel economy, acceleration, and seating capacity of cars) in order to fill niches of the market that are relatively undersupplied. That is, firms choose location in the space of characteristics.

Hotelling considered a very special characteristic space: location within the interval $[0, 1]$. This could be taken literally as an address on a small town’s Main Street, or metaphorically as in characteristic space.

We begin with a duopoly where firms choose location but not price; for simplicity we assume that price is fixed at $p_1 = p_2 = p > c$, where $c = c_1 = c_2$ is the constant marginal cost faced by both firms.

- Label the firms so that the location choices satisfy $z_1 \leq z_2 \in [0, 1]$.
- For simplicity, assume that consumers’ preferred locations are uniformly distributed along $[0, 1]$.
- Also, for simplicity, assume linear transportation (or transformation) cost $t > 0$. Thus the delivered price at location z for firm j is $p_j(z) = p + t|z - z_j|$. The assumption is that consumers buy at the lowest delivered price.
- Under current simplifications, this means that firm 1 gets all customers in $[0, \hat{z}]$ and firm 2 gets those in $(\hat{z}, 1]$, where \hat{z} solves $p_1(z) = p_2(z)$. In other words, the market shares are \hat{z} and $1 - \hat{z}$, where the customer at $\hat{z} = 0.5(z_1 + z_2)$ faces the

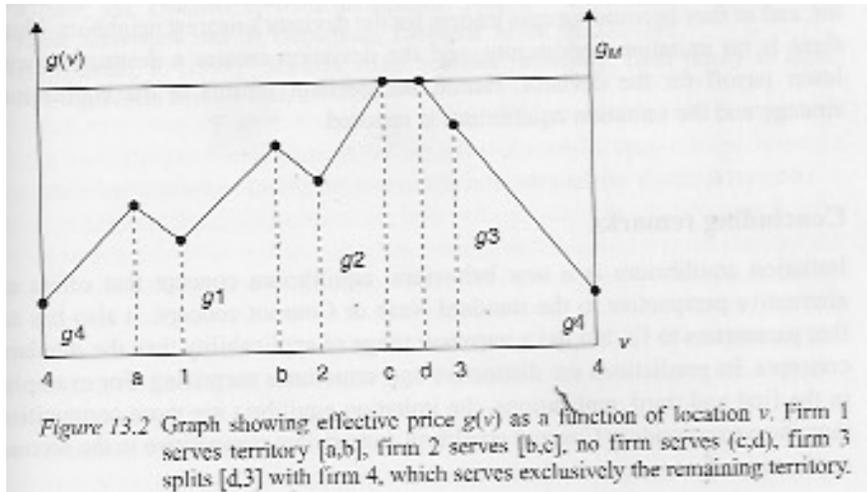


Figure 4: From Selten chapter in Friedman and Cassar (2004). Here $g_i = p_i(z) - c$, and i is location z_i (denoted v in the Figure).

same delivered price from both firms.

- Since $p > c$, firms maximize profit by maximizing market share.
- What is firm i 's BR to location choice z_j of firm j ? If $z_j < 0.5$, it is to locate a tiny bit to the right, at $z_j + \epsilon$. If $z_j > 0.5$, it is to locate a tiny bit to the left, at $z_j - \epsilon$. This is how i maximizes market share.
- So the unique NE of this Hotelling location game is for both firms to locate back-to-back at $z = 0.5$.

This simple game is sometimes used to explain (in part) why firms in a similar line of business tend to locate next door to each other, and why political parties used to adopt very similar platforms.

There are many, many extensions of the model. Expanding the duopoly location problem above to triopoly yields no NE in pure strategies. With 4 firms, the pure strategy NE all have 2 firms back to back at $z = .25$ and the other 2 at $z = .75$.

What if firms first commit to specific locations and then pick prices? Figure 4 illustrates how the shares are determined from an arbitrary set of locations and prices for 4 firms.

The analysis is a bit tricky because the endpoints of the line segment play a special role. To make the location space more homogeneous (the technical word is ‘isotropic’), one can join the endpoints to make a circle, and this is assumed in the Figure. It can be shown in this case that the unique NE in pure strategies is for firms to space themselves equally around the circle (maximum differentiation) and to all charge the price $p_i = c + t$.

Other variants of the Hotelling location game consider two dimensional locations, on a rectangle or (to make it isotropic) a torus and various numbers of firms. One can also consider nonlinear transportation (or transformation) costs. There seems to be room for applied work here, but I’ve not seen much published recently. There are ongoing laboratory experiments by UCSC PhD Curtis Kephart.

XII. Monopolistic Competition

Typically in 1-d space each product has two direct competitors, and perhaps lots of indirect competitors. In 2-d, there are typically at least 3 direct competitors, and often 5 or 6. In higher dimensions, i.e., if many characteristics are relevant, then there could be so many direct competitors it is not worth keeping track of them individually.

How to model competition against large numbers of direct competitors whose products are not especially close substitutes? It sounds complicated, but there actually is a pretty simple way to do it. Invented before game theory by Edwin Chamberlin and (apparently independently by) Joan Robinson in the early 1930’s, it combines the free-entry assumption of perfect competition with markup pricing as in monopoly. Here’s how the undergrad textbook version of Monopolistic Competition goes.

- Consider an industry, like restaurants or breakfast cereal, where each product is unique but has lots of not-very-close substitutes.
- Thus each firm faces a fairly elastic, but not infinitely elastic, demand curve.

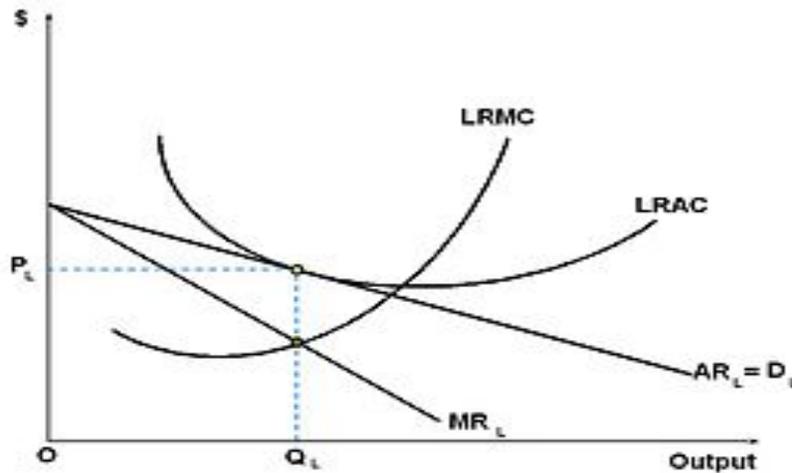


Figure 5: LR equilibrium for a Monopolistic Competitive firm. From Wikipedia article (not Wikip's finest).

- Solve the standard monopoly problem: the profit maximizing quantity is where $MR=MC$, and price is the inverse demand at that quantity.
- If economic profits are positive at that point, then more firms will be attracted into the industry, offering more fairly close substitutes. This will lower and flatten the demand curve, and lower the firm's maximum profit.
- Similarly, if profits are negative, then firms with similar products will exit, raising and steeping the demand curve, thus increasing its profit.
- In LR equilibrium, economic profit will be zero, as in Figure 5, since LRAC is equal to LRAR.
- Note that actually LRAC is tangent to inverse demand, aka LRAR. Since D is downward sloping, the firm is operating at less than efficient scale. In that sense, there is over-entry and inefficiency in monopolistically competitive industries.

There are also a grad level versions of monopolistic competition, which try to be more formal about the number of firms in the industry and where the demand curves come from. I can't find a really good version on the web, but you could start with

<http://www.foltyn.net/wp-content/uploads/2009/12/dixitstiglitz.pdf> .Here is my brief sketch, which might come in handy in your Macro class next quarter.

- Assume quasi-linear preferences for n goods, separable in good 0 (money): $u(x_0, x_1, \dots, x_n) = x_0 + G(f(x_1) + \dots + f(x_n))$, where G and f are concave/convex functions.
- A typical choice is the CES function $u = x_0 + [x_1^\rho + \dots + x_n^\rho]^{1/\rho}$. These are sometimes called Dixit-Stiglitz preferences, after their 1977 originators.
- Using what you learned earlier this quarter, normalize price so $p_0 = 1$, and maximize u subject to a budget constraint. By quasilinearity, we can write the first argument as initial cash less expenditure. The resulting FOCs give the inverse demand equations

$$p_j = G'(\bar{x})f'(x_j), j = 1, \dots, n, \quad (10)$$

where $\bar{x} = f(x_1) + \dots + f(x_n)$.

- This is nice, but for any finite, or even countable, number n of products, there is strategic interdependence, and so we have oligopoly, say a particular sort of Cournot oligopoly.
- To get true monopolistic competition, we have to take the continuum limit, so now $j \in [0, 1]$, and $\bar{x} = \int_0^1 f(x_j) dj$. The point is that (10) still works, so we can write direct demand as

$$x_j(p_j, \bar{x}) = \psi(p_j/G'(\bar{x})), j = 1, \dots, n, \quad (11)$$

where $\psi = f'^{-1}$ is the inverse of f' .

- That is, in the limit we have pricing power, with a monopoly markup independent of choices of other individual firms, but depending on an aggregate production index \bar{x} .