Lecture Note 10: Oligopoly

oligopoly: a market structure in which a small number of firms compete.

- Aircraft production (Boeing, Airbus).
- Cell phone service providers (Verizon, AT&T, Sprint, T-Mobile).
- Commercial airlines flying SMF $\leftrightarrow$ SEA (Delta, Alaska Airlines).

Oligopoly (a word that literally means "a few sellers") is a middle ground between monopoly (one seller) and perfect competition (many sellers).

Under oligopoly—unlike monopoly and perfect competition—we have to think carefully about strategic interactions between competing firms.

- When deciding how many 787 Dreamliners to produce, Boeing thinks about how many A350 jets Airbus is going to produce.
- When deciding where to build new cell towers, Verizon thinks about where Sprint's towers are located.
- When deciding what to charge for an economy-class ticket, Delta thinks about the price Alaska will charge.

Economists use a number of different models to understand oligopolies. We will study two of these models:[†]

- <u>Cournot</u>: firms compete by choosing quantities at the same time.
- <u>Stackelberg</u>: similar to Cournot, but one firm gets to move first.

We'll use the Cournot model to understand how a market behaves when there is *some* competition, but not *perfect* competition. Building on our earlier study of dynamic games, we'll use the Stackelberg model to explore the potential advantages of being the first firm to enter a market.

[†] A third, widely-used oligopoly model is the *Bertrand model*, in which firms compete by deciding what price to charge. We will not be covering Bertrand oligopoly this year.

The Cournot model

Two key features distinguish Cournot from other oligopoly models:†

1. Firms compete by choosing the quantities they want to produce.
2. Firms choose their quantities at the same time, without "colluding" (that is, without working together), and without knowing how much the other firm(s) will produce.

Under Cournot, firms don't directly choose their prices. Instead, once quantities are chosen, prices adjust to clear the market (supply = demand).

- Assumption: once each firm picks its quantity, that choice is final.
- This assumption is most reasonable in markets in which firms have to make their production decisions long before a good is actually sold (because the production process takes time)—e.g., the auto market.

We'll focus on markets with two firms (a **Cournot duopoly**), but I'll talk briefly about how the market changes as the number of firms increases.

The Cournot model is a *static game*: the players are the two firms, the strategies are their choices of output, and the payoffs are their profits.

- We'll look for a *Nash equilibrium*: choices of quantity q_1^* and q_2^* such that neither firm can increase its profit by producing more/less.
- One twist: since each firm can pick any non-negative quantity, it has a continuous (infinite) choice of strategies. We have to use calculus.

We'll assume the firms are selling **identical products**. (This assumption is reasonable for some markets, but not others. It's possible to modify the Cournot model to study markets with *differentiated* products.)

† A historical note: the Cournot model is one of the oldest mathematical models in economics (first published in 1838). In fact, Cournot's approach anticipated much of 20th-century game theory.

The Cournot model (continued)

Here is the general setup:

- Firm 1 produces $q_1 \geq 0$ units of output at a total cost $C_1(q_1)$.
- Firm 2 produces $q_2 \geq 0$ units of output at a total cost $C_2(q_2)$.
- The market price (same for both firms) is determined by a downward-sloping market demand curve, which depends on total output:

$$p(Q) = p(q_1 + q_2) \qquad \text{where } Q = q_1 + q_2$$

- Why is p the same for both firms? They're selling identical products.

We have all the information we need to write each firm's profit function:

$$\pi_1(q_1, q_2) = p(q_1 + q_2)q_1 - C_1(q_1)$$
$$\pi_2(q_1, q_2) = p(q_1 + q_2)q_2 - C_2(q_2)$$

Each firm picks q based on its "guess" about what the other firm will do.

- Suppose firm 1 thinks firm 2 will produce $\widehat{q}_2$ (pronounced "q_2 hat").
- Then firm 1 will produce whatever amount q_1^* is the best response against $\widehat{q}_2$. (There will typically be a *unique* best response.) We write firm 1's **best-response function** as $q_1^* = BR_1(\widehat{q}_2)$.
- Likewise: firm 2 will produce its best response $q_2^* = BR_2(\widehat{q}_1)$, where $\widehat{q}_1$ is firm 2's guess about the amount that firm 1 will produce.

A *Nash equilibrium*† is a pair of choices (q_1^*, q_2^*) such that q_1^* is firm 1's best response against q_2^* and q_2^* is firm 2's best response against q_1^*:

$$q_1^* = BR_1(q_2^*) \qquad \text{and} \qquad q_2^* = BR_2(q_1^*)$$

In a Nash equilbrium, each firm's "guess" must be correct: $q_1^* = \widehat{q}_1$ and $q_2^* = \widehat{q}_2$. If either firm guesses wrong, it will regret its choice—but that would mean its choice wasn't a best response. Nash means no regret!

†This is just our usual definition of Nash equilibrium (a pair of strategies that are best responses against each other). The only difference here is that each player has a continuous set of strategies.

Finding the Nash equilibrium in a Cournot problem

Example: Suppose that:

- The market demand curve is $p(Q) = 18 - Q = 18 - q_1 - q_2$.
- The firms have cost functions $C_1(q_1) = 6q_1$ and $C_2(q_2) = 6q_2$.

What are the Nash equilibrium quantities, q_1^* and q_2^*?

(Once we've found the equilibrium quantities, we can also calculate other familiar expressions, like the equilibrium price p^*, each firm's equilibrium profit, the consumer surplus, and the deadweight loss.)

Step 1: find firm 1's best-response function, $BR_1(\widehat{q}_2)$.

- Firm 1 picks q_1 to maximize its profits, taking $\widehat{q}_2$ as "given":

$$\max_{q_1} \ \pi_1 = (18 - q_1 - \widehat{q}_2)q_1 - 6q_1$$

- Key observation: *firm 1 treats firm 2's quantity like a constant.* So we partially differentiate profits with respect to q_1 to get the FOC:

$$18 - 2q_1 - \widehat{q}_2 - 6 = 0$$

- Solving the FOC for q_1 gives us firm 1's best-response function:

$$q_1^* = BR_1(\widehat{q}_2) = 6 - \frac{1}{2}\widehat{q}_2$$

- Notice that firm 1's choice q_1^* is a decreasing function of $\widehat{q}_2$. Why?
  - The more firm 2 sells, the less *residual demand* firm 1 faces.
  - With $\widehat{q}_2 = 3$, it's as though firm 1 faces the demand curve $p(q_1) = 15 - q_1$. With $\widehat{q}_2 = 10$, it's as though firm 1 faces lower demand, $p(q_1) = 8 - q_1$. Its optimal quantity falls as a result.

Step 2: find firm 2's best-response function, $BR_2(\widehat{q}_1)$.

- There are two ways to do this: a long way and a shortcut.
- The long way is to set up firm 2's profit-maximization problem, take the FOC, and solve for q_2^*, just as we did for firm 1.
- The shortcut is to exploit the fact that firm 1 and 2 are treated *symmetrically* in this problem. What do I mean by "symmetric"?
  - In the demand curve $p(Q) = 18 - q_1 - q_2$, the same coefficient (in this case, "1") appears in front of both q_1 and q_2.
  - The firms face identical costs: $C_1(q_1) = 6q_1$, $C_2(q_2) = 6q_2$.
  - The firms move at the same time (as is always true in Cournot).
  - Thus, firm 2's optimization problem looks exactly like firm 1's (with the subscripts "1" and "2" swapped). Therefore, firm 2's best-response function will be a "mirror image" of firm 1's.
- The long way works for all problems; the "symmetry shortcut" only works if the problem is symmetric. If there is an asymmetry between the firms—e.g., the demand curve is $p(q_1, q_2) = 18 - q_1 - 3q_2$ or the firms have different cost functions—we'll have to use the long way.
- Either way, we get the same answer: $q_2^* = BR_2(\widehat{q}_1) = 6 - \frac{1}{2}\widehat{q}_1$.

Step 3: set up the system of equations.

- The two best-response functions give us a system of two equations:

$$q_1^* = 6 - \frac{1}{2}\widehat{q}_2 \qquad \text{and} \qquad q_2^* = 6 - \frac{1}{2}\widehat{q}_1$$

- In a Nash equilibrium, each firm's guess about the other player is correct: $q_1^* = \widehat{q}_1$ and $q_2^* = \widehat{q}_2$. So, let's turn the "hats" into "stars":

$$q_1^* = 6 - \frac{1}{2}q_2^* \qquad \text{and} \qquad q_2^* = 6 - \frac{1}{2}q_1^*$$

Step 4: *solve* the system of equations to find q_1^* and q_2^*.

- Here's where we are again:

$$q_1^* = 6 - \frac{1}{2}q_2^*$$
$$q_2^* = 6 - \frac{1}{2}q_1^*$$

- To solve, we can substitute q_2^* into the first equation and find q_1^*:

$$q_1^* = 6 - \frac{1}{2}\left(6 - \frac{1}{2}q_1^*\right) \implies \frac{3}{4}q_1^* = 3 \implies q_1^* = 4$$

- We can now plug q_1^* into firm 2's best-response function to find q_2^*:

$$q_2^* = 6 - \frac{1}{2} \cdot 4 = 4$$

- Reassuringly, we find that both firms choose the same level of output. Since they produce identical products, face identical costs, and move at the same time, that's exactly what we should expect here.

Step 5: calculate other expressions of interest (if asked to do so).

- Once we know q_1^* and q_2^*, we can calculate all of the expressions we studied in the first part of the course: I'll leave this as an exercise:

  - Find p^*. Answer: since $Q^* = 4 + 4 = 8$, $p^* = 18 - 8 = 10$.
  - Find the producer surplus. Answer: Adding up the two firms' profits, we get $PS = 2 \text{ firms } \times (10 \cdot 4 - 6 \cdot 4) = 32$.
  - Find the consumer surplus. Answer: $CS = \frac{1}{2}(8)(18 - 10) = 32$.
  - Find the deadweight loss. Ans.: $DWL = \frac{1}{2}(12 - 8)(10 - 6) = 8$.

- You will have to calculate some of these things on Midterm 2, so make sure you're comfortable doing so. (As the syllabus says, Midterm 2 is a cumulative exam, and this is the "cumulative" part.)

The Stackelberg model

The Cournot model assumes firms choose quantities at the same time.

The **Stackelberg model** of duopoly assumes that one firm (the **leader**) chooses its quantity first, followed by the other firm (the **follower**).

- We can think of the leader as an *incumbent* firm—one that's already in the market—and the follower as a *potential entrant*.
- Example: Apple (iPad) was the leader in the market for tablets. Microsoft (Surface) was a follower in that market.

Just as Cournot is a *static game*, Stackelberg is a *dynamic game*.

- As with the dynamic games we've already seen, we'll use *backward induction* to find this game's *subgame-perfect Nash equilibrium*.
- However, because each player now has an infinite (and continuous) set of available strategies, we can't use a game tree to represent this game, and we'll have to use calculus to solve it.

Since we'll want to compare the outcome of the Stackelberg game to the Cournot outcome, we'll use the same demand curve and cost functions:

- Firm 1 (the leader) moves first, then firm 2 (the follower).
- Demand curve: $p(Q) = 18 - Q$, where $Q = q_1 + q_2$ is total output.
- Cost functions: $C_1(q_1) = 6q_1 \quad \text{and} \quad C_2(q_2) = 6q_2$.

We'll start by trying to determine the (subgame-perfect) Nash equilibrium quantities produced, q_1^* and q_2^*. Once we know these, we can (as usual) compute the price, profits, consumer surplus, and deadweight loss.

Solving a Stackelberg problem

Since the follower moves last, backward induction tells us to start with the follower's decision. This part will look very similar to the Cournot model.

Step 1: find the follower's best-response function, $q_2^* = BR_2(q_1)$.

In Cournot, firm 2 has to make an "educated guess" about firm 1's choice. In Stackelberg, by contrast, firm 2 actually observes q_1 before choosing its quantity. To reflect the fact that there is no "guesswork" involved, let's write "q_1" instead of "$\widehat{q}_1$". Firm 2's profit-maximization problem is

$$\max_{q_2} \ \pi_2 = (18 - q_1 - q_2)q_2 - 6q_2$$

We find firm 2's best-response function by taking the FOC:

$$18 - q_1 - 2q_2^* - 6 = 0 \implies q_2^* = BR_2(q_1) = 6 - \frac{1}{2}q_1$$

That's the same best-response function we found in the Cournot model. Well, almost: the only difference is that here we have q_1 instead of $\widehat{q}_1$.

This seems like a minor detail, but it has a big impact on the outcome of the game. Here is the big difference between Cournot and Stackelberg:

- In Cournot, firm 2's action depends on what it expects firm 1 to do.
- In Stackelberg, it depends on what firm 1 actually does.

This is a subtle point, because in Cournot's Nash equilibrium firm 2's "guess" is correct: firm 1's expected and actual output are the same.

But the fact that the follower sees what the leader *actually chose* before choosing q_2 allows the leader to influence what the follower chooses. And the leader's ability to do this turns out to be very good for the leader.

To see why, we have to look at the world from the leader's perspective.

Step 2: determine the leader's optimal quantity q_1^*.

The key thing now is that, when solving the leader's problem, we have to take the follower's predictable reaction into account. Since both players are rational, the leader can perfectly predict how the follower react, and the leader will factor this information into its decision-making process.

This has a key implication for the math:

- In Cournot, we find each firm's best-response function by treating the other firm's output as a constant. It's just the firm's best guess.
- In Stackelberg, we find the follower's best-response function $BR_2(q_1)$ by treating the leader's output as a constant, just as we did before. But we find the leader's best-response function by plugging $BR_2(q_1)$ into the leader's profit function before we take the FOC.

When it chooses q_1, the leader *indirectly controls the follower's choice*. Therefore, when maximizing its profits, the leader treats q_2 as a function of q_1—specifically, as firm 2's best-response function—not as a constant:

$$\max_{q_1} \ \pi_1 = (18 - q_1 - BR_2(q_1))q_1 - 6q_1$$

We have to plug in $q_2 = BR_2(q_1)$ before taking the derivative. This ensures that firm 1 accounts for firm 2's (predictable) reaction when it chooses q_1.

Substituting $BR_2(q_1) = 6 - \frac{1}{2}q_1$ into firm 1's profit function:

$$\max_{q_1} \ \pi_1 = \left(18 - q_1 - \left(6 - \frac{1}{2}q_1\right)\right) q_1 - 6q_1$$

In Stackelberg problems, it helps to clean this up before differentiating:

$$\max_{q_1} \ \pi_1 = 6q_1 - \frac{1}{2}q_1^2$$

Now we can take the FOC: $\frac{d\pi_1}{dq_1} = 6 - q_1 = 0 \implies q_1^* = 6$.

Step 3: describe the subgame-perfect Nash equilibrium.

Since firm 1 moves first, its strategy is just a single quantity, $q_1^* = 6$.

By contrast, since firm 2 moves second, its strategy is a complete "if-then" plan indicating how it will respond to every action q_1 that firm 1 could choose. In other words, firm 2's strategy is its best-response function.

So, the pair of strategies $q_1^* = 6$ and $q_2^* = BR_2(q_1) = 6 - \frac{1}{2}q_1$ form a subgame-perfect Nash equilibrium. In equilibrium, since firm 1 produces 6 units, firm 2 will respond by producing $q_2^* = 6 - \frac{1}{2} \cdot 6 = 3$ units.

Comparing the Cournot and Stackelberg outcomes

Under Cournot, the equilibrium quantities were $q_1^{\text{Cour}} = 4$ and $q_2^{\text{Cour}} = 4$.

Under Stackelberg, they are $q_1^{\text{Stack}} = 6$ and $q_2^{\text{Stack}} = 3$.

- The Stackelberg leader produces more than it would under Cournot. The leader "commits" to high production in hopes of getting the follower to choose a low level of production in response.
- The Stackelberg follower produces less than it would under Cournot. The leader's "intimidation tactics" worked: the follower "backs down".
- The Stackelberg leader makes more profit than it would under Cournot.
    - We can show this mathematically, by computing firm 1's profits under both Cournot and Stackelberg. (Exercise!)
    - Alternatively, we can make another **revealed preference** argument: the Stackelberg leader could have chosen $q_1 = 4$ units, which would have resulted in it getting "Cournot profits". Since it chose $q_1 = 6$ instead, doing so must be more profitable.

Like the pharmaceutical entry game we studied in Lecture Note 9, the Stackelberg game has **first-mover advantage**. Fortune favors the bold!