Advanced Microeconomic Analysis, Lecture 5

Prof. Ronaldo CARPIO

October 14, 2014

Homework #2

- Homework #2 is due today.
- I will post HW #3 and the solutions to HW #2 on the course website later today.
- HW #3 is due in two weeks (Oct. 28).
- The midterm will be on Nov. 4.
- Midterm will be open-book.
- Chapters 1, 2.1, 3, and 7 will be covered.

Properties of Short-Run Costs

- For a given level of output, long-run costs (where the firm can set all input levels) cannot be greater than short-run costs (where the firm cannot set all input levels).
- This is because the feasible region for $\boldsymbol{z}$ of the short-run problem is a *subset* of the feasible region of the long-run problem.
- Any feasible $\boldsymbol{z}$ in the short-run problem is feasible in the long-run problem, but not vice versa.
- In general, if set $A \subset B$, then for any function g:

$$\max_{x \in A} g(x) \leq \max_{x \in B} g(x)$$

$$\min_{x \in A} g(x) \geq \min_{x \in B} g(x)$$

Duality

- In consumer theory, we saw the duality between utility and expenditure.
- Likewise, there is a duality between production and cost.
- If we begin with a production function, we can derive the cost function.
- If we begin with a cost function, we can generate a production function. If the original production function is quasiconcave, the derived production function will be identical.
- This is important for applied work. When analyzing real-world data, it is difficult to know what the "true" production function is.
- Instead, we can observe market input prices and levels of output, and estimate a cost function.
- Then, "recover" the underlying production function.

Recovering Production Function from Cost Function

- Theorem 3.5: Suppose $c(\boldsymbol{w}, y)$ satisfies the properties of a cost function (including differentiability). Then the function $f : \mathbb{R}^n_+ \rightarrow \mathbb{R}_+$ defined by:

$$f(\boldsymbol{x}) = max\{y \geq 0 | \boldsymbol{w} \cdot \boldsymbol{x} \geq c(\boldsymbol{w}, y) \qquad \forall \boldsymbol{x} >> 0\}$$

- is increasing, unbounded above, and quasiconcave. Moreover, the cost function generated by f is c.

The Competitive Firm

- Now, let's assume the firm is in perfect competition on both input and output markets, i.e. it is a price taker for both input and output prices.
- Let p denote the output market price.
- Revenues are: $R(y) = py$
- For a given output level y, profits are: $py - c(\boldsymbol{w}, y)$
- We will consider the problem where the firm can choose both the level of output, and the inputs used to produce it.

Profit Maximization Problem

$$\max_{\mathbf{x},y} py - \mathbf{w} \cdot \mathbf{x} \qquad \text{s.t.} \qquad f(\mathbf{x}) \geq y$$

- Assuming f is strictly increasing, the constraint will be satisfied with equality. Therefore, y is completely determined by $\mathbf{x}$, so the problem becomes:

$$\max_{\mathbf{x}} pf(\mathbf{x}) - \mathbf{w} \cdot \mathbf{x}$$

- Assume this problem has an interior solution at some strictly positive $\mathbf{x}^*$.

Profit Maximization Problem

- First-order conditions require that the gradient be zero, since there are no constraints:

$$p\frac{\partial f(\boldsymbol{x}^*)}{\partial x_i} = w_i \qquad \text{for } i = 1, ..., n$$

- The term on the left-hand side is called the *marginal revenue product* of input i.
- At optimality, marginal revenue product = marginal cost of input i.

Profit Maximization Problem

- Assuming all input prices w_i are positive, we get that the MRTS must equal the ratio of prices:

$$MRTS_{ij} = \frac{\partial f(\boldsymbol{x}^*)/\partial x_i}{\partial f(\boldsymbol{x}^*)/\partial x_j} = \frac{w_i}{w_j} \qquad \text{for all } i, j$$

- This is the same condition as in the cost-minimization problem.
- Input demands will be the same as in the cost-minimization problem.

Profit Function

- Suppose f satisfies the usual conditions, and in addition, is *strictly concave*.
- Then, solutions to the profit maximization problem will be unique for each $(p, \boldsymbol{w})$.
- The optimal choice of output, $y^* = y(p, \boldsymbol{w})$ is called the firm's *output supply function*.
- The optimal choice of inputs, $\boldsymbol{x}^* = \boldsymbol{x}(p, \boldsymbol{w})$ is called the firm's *input demand function*.
- The *profit function* depends only on input and output prices:

$$\pi(p, \boldsymbol{w}) = \max_{\boldsymbol{x}, y} py - \boldsymbol{w} \cdot \boldsymbol{x} \qquad \text{s.t. } f(\boldsymbol{x}) \geq y$$

Profit Function

$$\pi(p, \boldsymbol{w}) = \max_{\boldsymbol{x}, y} py - \boldsymbol{w} \cdot \boldsymbol{x} \qquad \text{s.t. } f(\boldsymbol{x}) \geq y$$

- First, we must be sure that a maximum of profits actually exists (i.e. is finite).
- It may be that there is no finite maximum. For example, suppose that production has increasing returns.
- Then, starting from any choice of $\boldsymbol{x}$ and $y = f(\boldsymbol{x})$, it will always be possible to choose $t\boldsymbol{x}, t > 0$, that gives higher profits.
- If production has constant returns, then a maximum profit may exist. However, the *scale* is indeterminate: the input levels $\boldsymbol{x}$ and $t\boldsymbol{x}$ give the same profits for all $t > 0$.

Properties of the Profit Function

- Theorem 3.7: Assume f is continuous, strictly increasing, and strictly quasiconcave. For $p \geq 0, \boldsymbol{w} \geq 0$, the profit function $\pi(p, \boldsymbol{w})$ (where well-defined) is continuous and:
  - Increasing in p;
  - Decreasing in $\boldsymbol{w}$;
  - Homogeneous of degree one in $(p, \boldsymbol{w})$;
  - Convex in $(p, \boldsymbol{w})$;
  - Differentiable in $(p, \boldsymbol{w})$ for strictly positive $(p, \boldsymbol{w})$. If f is strictly concave, then

$$\frac{\partial \pi(p, \boldsymbol{w})}{\partial p} = y(p, \boldsymbol{w}), -\frac{\partial \pi(p, \boldsymbol{w})}{\partial w_i} = x_i(p, \boldsymbol{w}) \qquad \text{for } i = 1, \ldots, n$$

- The last property is known as Hotelling's Lemma.

Properties of Output Supply and Input Demand Functions

- Theorem 3.8: Assume f is continuous, strictly increasing, and strictly concave, and assume $\pi(p, \boldsymbol{w})$ is twice continuously differentiable. Then, for all $p > 0, \boldsymbol{w} >> 0$ where it is well-defined:
  - Homogeneity of degree zero:

$$y(tp, t\boldsymbol{w}) = y(p, \boldsymbol{w}) \qquad \text{for all } t > 0$$

$$x_i(tp, t\boldsymbol{w}) = x_i(p, \boldsymbol{w}) \qquad \text{for all } t > 0, i = 1, ..., n$$

  - Own-price effects:

$$\frac{\partial y(p, \boldsymbol{w})}{\partial p} \geq 0, \frac{\partial x_i(p, \boldsymbol{w})}{\partial w_i} \leq 0 \qquad \text{for } i = 1, ..., n$$

Properties of Output Supply and Input Demand Functions

- Substitution matrix is symmetric and positive semidefinite:

$$\begin{pmatrix} \frac{\partial y(p,\mathbf{w})}{\partial p} & \frac{\partial y(p,\mathbf{w})}{\partial w_1} & \cdots & \frac{\partial y(p,\mathbf{w})}{\partial w_n} \\ \frac{-\partial x_1(p,\mathbf{w})}{\partial p} & \frac{-\partial x_1(p,\mathbf{w})}{\partial w_1} & \cdots & \frac{-\partial x_1(p,\mathbf{w})}{\partial w_n} \\ \vdots & \vdots & & \vdots \\ \frac{-\partial x_n(p,\mathbf{w})}{\partial p} & \frac{-\partial x_n(p,\mathbf{w})}{\partial w_1} & \cdots & \frac{-\partial x_n(p,\mathbf{w})}{\partial w_n} \end{pmatrix}$$

Example 3.5

- Assume CES production: $y = (x_1^\rho + x_2^\rho)^{\frac{\beta}{\rho}}$.
- Note the parameter β. This determines the scale of production; if $\beta < 1$, there is decreasing returns to scale. Suppose $\beta < 1$.
- Assume an interior solution. First-order conditions:

$$w_1 + p\beta(x_1^\rho + x_2^\rho)^{\frac{\beta-\rho}{\rho}} x_1^{\rho-1} = 0$$

$$w_2 + p\beta(x_1^\rho + x_2^\rho)^{\frac{\beta-\rho}{\rho}} x_2^{\rho-1} = 0$$

$$x_1 = x_2(w_1/w_2)^{\frac{1}{\rho-1}}$$

- The supply function is:

$$y = (p\beta)^{\frac{-\beta}{\beta-1}}\left(w_1^{\frac{\rho}{\rho-1}} + w_2^{\frac{\rho}{\rho-1}}\right)^{\frac{\beta(\rho-1)}{\rho(\beta-1)}}$$

Example 3.5

- The input demand functions are:

$$x_i = w_i^{\frac{1}{\rho-1}} (p\beta)^{\frac{-1}{\beta-1}} \left(w_1^{\frac{\rho}{\rho-1}} + w_2^{\frac{\rho}{\rho-1}}\right)^{\frac{\rho-\beta}{\rho(\beta-1)}}$$

- The profit function is:

$$\pi(p, \boldsymbol{w}) = p^{\frac{-1}{\beta-1}} (w_1^r + w_2^r)^{\frac{\beta}{r(\beta-1)}} \beta^{\frac{-\beta}{\beta-1}} (1-\beta), \qquad r = \frac{\rho}{\rho-1}$$

- If $\beta = 1$, production has constant returns to scale and the profit function is undefined.
- If $\beta > 1$, production has increasing returns to scale. The FOC give conditions for a minimum instead of a maximum.

Short-Run Profit Maximization

- As before, the long-run profit function is when all inputs can be changed. The short-run profit function is when some inputs must be fixed.
- Theorem 3.9: Suppose that f is continuous, strictly increasing, and strictly concave.
  - For $k < n$, let $\overline{\boldsymbol{x}} \in \mathbb{R}_+^k$ be a subvector of fixed inputs
  - Consider $f(\boldsymbol{x}, \overline{\boldsymbol{x}})$ as a function of the subvector of variable inputs $\boldsymbol{x} \in \mathbb{R}_+^{n-k}$.
  - Let $\boldsymbol{w}, \overline{\boldsymbol{w}}$ be the vector of prices for the variable and fixed inputs, respectively.
  - The *short-run*, or *restricted*, profit function is:

$$\pi(p, \boldsymbol{w}, \overline{\boldsymbol{w}}, \overline{\boldsymbol{x}}) = \max_{y, \boldsymbol{x}} py - \boldsymbol{w} \cdot \boldsymbol{x} - \overline{\boldsymbol{w}} \cdot \overline{\boldsymbol{x}} \qquad \text{s.t. } f(\boldsymbol{x}, \overline{\boldsymbol{x}}) \geq y$$

- The solutions $y(p, \boldsymbol{w}, \overline{\boldsymbol{w}}, \overline{\boldsymbol{x}}), \boldsymbol{x}(p, \boldsymbol{w}, \overline{\boldsymbol{w}}, \overline{\boldsymbol{x}})$ are called the short-run output supply and variable input demand functions.

Short-Run (or Restricted) Profit Function

- For all $p > 0$ and $\boldsymbol{w} >> 0, \pi(p, \boldsymbol{w}, \overline{\boldsymbol{w}}, \overline{\boldsymbol{x}})$, where well-defined, is:
  - continuous in p and $\boldsymbol{w}$,
  - increasing in p
  - decreasing in $\boldsymbol{w}$
  - convex in $(p, \boldsymbol{w})$
  - If $\pi(p, \boldsymbol{w}, \overline{\boldsymbol{w}}, \overline{\boldsymbol{x}})$ is twice continuously differentiable, the short-run output supply and variable input demand functions have the same properties as in Theorem 3.8 (homogeneity of degree zero, own-price effects, and positive semidefinite substitution matrix)

Optimal Shutdown

- Recall that $sc(p, \mathbf{w}, \overline{\mathbf{w}}, \overline{\mathbf{x}})$ is the short-run cost function. Consider the short-run profit function

$$\pi(p, \mathbf{w}, \overline{\mathbf{w}}, \overline{\mathbf{x}}) = max_y py - sc(p, \mathbf{w}, \overline{\mathbf{w}}, \overline{\mathbf{x}})$$

- The first-order condition tells us that for optimal $y^* > 0$,

$$p = \frac{d\ sc(y^*)}{dy}$$

- that is, price should equal short-run marginal cost. Suppose this is true at some y^1.
- Let $tvc(y)$ denote the total variable cost, and let tfc denote the total fixed cost. Then

$$\pi^1 = py^1 - tvc(y^1) - tfc$$

Optimal Shutdown

- If π^1 is negative, the firm is better off shutting down and producing nothing ($y = 0$). Let π^0 denote profits when $y = 0$:

$$\pi^0 = p \cdot 0 - tvc(0) - tfc = -tfc < 0$$

- The firm will produce $y^1 > 0$ only if $\pi^1 \geq \pi^0$, or

$$py^1 - tvc(y^1) \geq 0$$

$$p \geq \frac{tvc(y^1)}{y^1} = avc(y^1)$$

- Thus, the firm will shut down if the output price p is less than the average variable cost of y^1.

Chapter 7: Game Theory

- *Game Theory* is the mathematical study of *strategic* situations, i.e. where there is more than one decision-maker, and each decision-maker can affect the outcome.
- So far, we have studied *single-person* problems. For example:
  - How much of each good to consume, in order to maximize my utility?
  - How much output should a firm produce, in order to maximize profits?
- Rational behavior: choose the level that maximizes utility (or profits, or payoffs).
- However, in multi-agent situations, my choice may change your problem.
- We need a method that takes everyone's choices into account.

Strategic Decision Making

- In a single-person decision making problem, assuming the decision-maker is rational, we predict the outcome will be that the decision-maker will choose the action (e.g. a bundle of goods, or set of inputs) that maximizes payoff (e.g. utility, profits).
- Assuming differentiability and concavity, we can then find this optimal choice with first-order conditions.
- However, in a strategic situation, the utility-maximizing choice of one agent may change, depending on what other agents do.

Example: Football Penalty Kick

- Suppose we have the following strategic situation with two agents.
- A football player is kicking a penalty kick against a goalie.
- The kicker can choose from two possible actions: kick *Left* or *Right*.
- The goalie also has two possible options: dive *Left* or *Right*.
- If the kicker and the goalie choose the same direction, the goalie wins. Otherwise, the kicker wins.
- Let's suppose the winning player gets a payoff of 1, while the losing player gets a payoff of -1.

Example: Football Penalty Kick

		Kicker	
		L	R
Goalie	L	1,-1	-1,1
	R	-1,1	1,-1

- We can summarize this situation in a 2×2 matrix.
- Each row corresponds to an action of the *Goalie* player, and each column corresponds to an action of the *Kicker* player.
- In each cell, the first number is the payoff to the row player, and the second number is the payoff to the column player.

Example: Football Penalty Kick

		Kicker	
		L	*R*
Goalie	*L*	1,-1	-1,1
	R	-1,1	1,-1

- We want to predict what the outcome of this situation will be.
- If we could fix the action of one player, then we could predict the other player's action.
- For example, assume the *Kicker* chooses *Left*. Then, we predict *Goalie* will choose his payoff-maximizing choice, *Left*.
- However, if we assume *Goalie* chooses *Left*, then *Kicker*'s payoff-maximizing choice is *Right*.
- Optimization alone cannot predict what the outcome will be.

Strategic Form Game

- We will formally define a strategic situation as follows.
- **Def. 7.1**: A *strategic form game* with N players is a tuple $G = (S_i, u_i)_{i=1}^{N}$, where for each player $i = 1, ..., N$:
  - S_i is the set of actions (or *strategies*) available to player i
  - $u_i(\cdot)$ is a *payoff function* that gives the payoff to player i, given the strategies chosen by all players
- A strategic form game is *finite* if each player's strategy set S_i is finite.

Strategic Form Game

- For the Football Penalty Kick game, the definition is as follows.
- By convention, Player 1 is the row player (*Goalie*), and Player 2 is the column player (*Kicker*).
  - $S_1 = S_2 = \{Left, Right\}$
  - $u_1(L, L) = u_1(R, R) = 1$
  - $u_1(L, R) = u_1(R, L) = -1$
  - $u_2(L, L) = u_2(R, R) = -1$
  - $u_2(L, R) = u_2(R, L) = 1$

Dominant Strategies

- Intuitively, in a strategic situation, each player has some belief about what the other players will do.
- For example, the *Kicker* might believe the *Goalie* has a tendency to choose *Left*, and vice versa.
- However, modeling beliefs can become very difficult.
- *Kicker* might believe that *Goalie*'s behavior depends on what *Goalie* believes about *Kicker*, and so on...
- In order to avoid such problems, economists use a variety of simplifying assumptions.
- The simplest is to consider problems where a player has a strategy that is payoff-maximizing in *every* situation.

Strictly Dominant Strategy

	L	R
U	3,0	0,-4
D	2,4	-1,8

- Consider this two-person strategic form game.
- Player 2's payoff-maximizing strategy depends on Player 1's choice.
  - If Player 1 chooses U, then Player 2 should choose L.
  - If Player 1 chooses D, then Player 2 should choose R.

Strictly Dominant Strategy

	L	R
U	3,0	0,-4
D	2,4	-1,8

- However, Player 1's payoff-maximizing strategy is the *same* (U), no matter what Player 2 chooses.
- If Player 1 is rational (i.e. payoff-maximizing), then it doesn't matter what his beliefs about Player 2 are: he will choose U.
- If we assume Player 2 realizes this, then Player 2 will assume Player 1 chooses U, and Player 2 will choose L.
- We say the *outcome* is (U, L), which gives a payoff vector of $(3, 0)$.
- Thus, we have arrived at a prediction, assuming only that Player 1 is rational, and Player 2 knows Player 1 is rational.

Strictly Dominant Strategy

- Let $S = S_1 \times ... \times S_N$ denote the set of joint strategies.
- We will use the symbol i to denote Player i, and $-i$ to denote all players *except for* Player i.
- So, s_i denotes a strategy in S_i, which is Player i's set of strategies.
- s_{-i} denotes a joint strategy of all players except Player i, which is an element of S_{-i}.
- **Def 7.2**: A strategy $\hat{s}_i$ for Player i is *strictly dominant* if

$$u_i(\hat{s}_i, s_{-i}) > u_i(s_i, s_{-i}) \qquad \text{for all } (s_i, s_{-i}) \in S, s_i \neq \hat{s}_i$$

Strictly Dominant Strategy

- If a rational player has a strictly dominant strategy, he will choose it.
- In a 2-player game, this determines the outcome.
- However, in most games there won't be a strictly dominant strategy.
- We can also find strategies that a rational player will *not* play, and eliminate them. This may determine the ouctome.

Strictly Dominated Strategy

- **Def 7.3** A strategy $\hat{s}_i$ of Player i is said to *strictly dominate* another strategy $\overline{s_i}$ if:

$$u_i(\hat{s}_i, s_{-i}) > u_i(\overline{s_i}, s_{-i}) \qquad \text{for all} s_{-i} \in S_{-i}$$

- We also say that $\overline{s_i}$ is *strictly dominated* in S.
- A rational player will never play a strictly dominated strategy, since there is some other strategy that gives a higher payoff in all situations.
- Assuming that players know other players are rational, we can *iteratively eliminate* strictly dominated strategies, which may reduce the number of possible outcomes to determine a solution.

Eliminating Dominated Strategies

	L	M	R
U	3,0	0,-5	0,-4
C	1,-1	3,3	-2,4
D	2,4	4,1	-1,8

- This game does not have a strictly dominant strategy.
- However, for Player 1, D strictly dominates C.
- For Player 2, R strictly dominates M.
- We can eliminate these strategies, since rational players will not choose them.
- The game then is reduced to the previous situation, which had the outcome (U, L).

Iteratively Strictly Undominated Strategies

- Suppose for each player i, we start with $S_i^0 = S_i$, then eliminate the strictly dominated strategies to get S_i^1.
- Then eliminate again to get $S_i^2, S_i^3, \ldots$
- Let S_i^n denote the strategies of Player i that *survive* after n rounds of elimination.
- $s_i \in S_i^n$ if $s_i \in S_i^{n-1}$ is not strictly dominated in S^{n-1}.
- **Def 7.4**: A strategy s_i of Player i is *iteratively strictly undominated* in S if $s_i \in S_i^n$ for all $n \geq 1$.

Weakly Dominated Strategies

- We can also define notions of *weak dominance*, where one strategy may be equal to another, except in one case.
- **Def 7.5** Player i's strategy $\hat{s}_i$ *weakly dominates* another strategy $\overline{s_i}$, if

$$u_i(\hat{s}_i, s_{-i}) \geq u_i(\overline{s_i}, s_{-i}) \qquad \text{for all } s_{-i} \in S_{-i}$$

- with at least one strict inequality. We also say that $\overline{s_i}$ is *weakly dominated* in S.

Weakly Dominated Strategies

	L	R
U	1,1	0,0
D	0,0	0,0

- In this game, neither player has a strictly dominated strategy.
- D is weakly dominated by U and R is weakly dominated by L.
- Eliminating weakly dominated strategies results in the unique outcome (U, L).

Iteratively Weakly Undominated Strategies

- Let W_i^n denote the strategies of Player i that *survive* after n rounds of elimination, with $W_i^0 = S_i$.
- $s_i \in S_i^n$ if $s_i \in S_i^{n-1}$ is not weakly dominated in S^{n-1}.
- **Def 7.6**: A strategy s_i of Player i is *iteratively weakly undominated* in S if $s_i \in S_i^n$ for all $n \geq 1$.
- The set of strategies remaining after removing weakly dominated strategies is a subset of those remaining after removing strictly dominated strategies.

Nash Equilibrium

- An equilibrium is a situation where no agent changes his behavior.
- When making predictions about strategic situations, equilibria are an attractive concept, since players would move away from non-equilibria.
- We want a concept of equilibrium where players are rational, and they know that all players are rational.
- This leads to the equilibrium concept of Nash equilibrium, in which each player is fully aware of all other players' behavior, and has no incentive to change is own behavior.

Pure Strategy Nash Equilibrium

- **Def 7.7** Given a strategic form game $G = (S_i, u_i)_{i=1}^N$, the joint strategy $\hat{s} \in S$ is a *pure strategy Nash equilibrium* of G is: for each player i,

$$u_i(\hat{s}) \geq u_i(s_i, \hat{s}_{-i}) \qquad \text{for all } s_i \in S_i$$

- Each player cannot find an alternative action that would give him a strictly higher payoff, keeping all other players' strategies constant.

Pure Strategy Nash Equilibrium

	L	R
U	1,1	0,0
D	0,0	0,0

- There are two pure strategy Nash equilibria: (U, L) and (D, R).

Pure Strategy Nash Equilibrium

	L	R
L	1,-1	-1,1
R	-1,1	1,-1

- This game has no pure strategy Nash equilibria.

Homework #2

- Homework #2 is due today.
- I will post HW #3 and the solutions to HW #2 on the course website later today.
- HW #3 is due in two weeks (Oct. 28).
- The midterm will be on Nov. 4.
- Midterm will be open-book.
- Chapters 1, 2.1, 3, and 7 will be covered.