Advanced Microeconomic Analysis, Lecture 9

Prof. Ronaldo CARPIO

May 15, 2017

Administrative Stuff

- HW #3 is due at the end of lecture today.
- I will post the solutions and HW #4, due in two weeks, to the website after class.

Review of Last Week

- A Walrasian equilibrium is a price vector that sets $\boldsymbol{z}(\boldsymbol{p}) = 0$, i.e. supply = demand for all markets simultaneously.
- We can extend the concept of a "good" to incorporate time, uncertainty, etc.
- A *contingent claim* or *contingent commodity* is a good that is indexed by time or by a random *state of the world* (i.e. a random outcome).
- For example, we can have "1 ton of grain 6 months in the future", "1 umbrella when it is raining", "1 unit of income in a recession", etc.
- Consumers can trade these goods just like ordinary goods, and reach equilibrium at some price vector.

Review of Last Week

- Suppose there are S states of the world (i.e. possible outcomes of a random variable).
- We can define a *financial asset* by how much income it gives to its owner in each possible state of the world.
- For example, cash is $(1, 1, ..., 1)$: 1 unit of cash will pay exactly 1 in every state of the world.
- *Riskless* or *risk-free* assets are those that give the same payoff in every state of the world.
- A *riskless bond* with net yield $1 + r$ is $(1 + r, ..., 1 + r)$. 1 unit of this bond will pay $1 + r$ in every state of the world.

Review of Last Week

- A *risky asset* is an asset that can give different payoffs in different states of the world.
- For example, suppose a driver buys car insurance, and there are two states of the world: "accident" is $s = 1$, "no accident" is $s = 2$.
- Then, an insurance policy would be an asset described by $(p, 0)$.
  - In $s = 1$ (there is an accident), the asset pays amount p.
  - In $s = 2$ (there is no accident), the asset pays 0.

Review of Last Week

- A special type of asset are the *Arrow securities* $e^i = (0, 0, ..., 1, 0, ..., 0)$, which pays 1 in state $s = i$ and zero in all other states.
- The payoffs of any financial asset can be replicated by a linear combination of the Arrow securities $e^1, ..., e^S$.
- Suppose we have a competitive market economy with I agents with risk-averse utility functions over wealth, and there are S states.
- The *mutuality principle* states that if there is no *aggregate* uncertainty (i.e. in each state, the total income of all agents is the same), then in equilibrium, all agents will trade so they achieve *perfect insurance*.
- That is, every agent's income will be the same in every state.

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, your choice may affect the parameters of my optimization problem.
- I need to take your actions into account when making my choice.

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

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

- 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.

Best Response Correspondence

- Suppose we have a game $G = (S_i, u_i)_{i=1}^N$. Let s_{-i} be any joint strategy of the players except for Player i.
- Player i's *best response correspondence* $B_i(s_{-i})$ is the set of strategies of Player i that give the highest possible payoff when s_{-i} are played:

$$B_i(s_{-i}) = \{\hat{s}_i \in S_i | u_i(\hat{s}_i, s_{-i} \geq u_i(s_i, s_{-i}) \qquad \text{for all } s_i \in S_i\}$$

- A member of the set $B_i(s_{-i})$ is called a *best response* of Player i to s_{-i}.
- If $B_i(s_{-i})$ is single-valued, i.e. there is always a unique payoff-maximizing strategy of Player i in response to s_{-i}, we also call this the *best response function*.
- A joint strategy s is a Nash equilibrium if every player is playing a best response to the joint strategy of the other players.
- In 2-player games, one way to find NE is to plot each player's best response correspondence, and find the intersection.

Example: Prisoner's Dilemma

	Cooperate	*Defect*
Cooperate	2,2	0,$\overline{3}$
Defect	$\underline{3}$,0	$\underline{1}$,$\overline{1}$

- The best response correspondences of Player 1 are: $B_1(Cooperate) = \{Defect\}, B_1(Defect) = \{Defect\}$
- The best response correspondences of Player 2 are: $B_2(Cooperate) = \{Defect\}, B_2(Defect) = \{Defect\}$
- We can underline Player 1's best response, and overline Player 2's best response.
- There is one intersection $(Defect, Defect)$, which is therefore the unique Nash equilibrium.

Example: A Joint Project

- Suppose 2 players are working on a joint project.
- Player i chooses the amount of work x_i to contribute, where $x_i \geq 0$.
- The payoff to each Player i is determined by the contributions of both players:

$$u_i = x_i(c + x_j - x_i)$$

- where x_j is the contribution of the other player, and $c > 0$ is a constant.
- Note that in this game, each player's strategy set is infinite: $S_i = [0, \infty)$.

Example: A Joint Project

$$u_i = x_i(c + x_j - x_i) = -x_i^2 + (c + x_j)x_i$$

- The payoff function $u_i(x_i)$ is concave in x_i, so we can use calculus to maximize it.
- Setting $u_i'(x_i) = -2x_i + c + x_j = 0$, we get Player i's best response function: $B_i(x_j) = (c + x_j)/2$.
- Player 1 and Player 2's best response functions are:

$$B_1(x_2) = (c + x_2)/2, B_2(x_1) = (c + x_1)/2$$

- At a Nash equilibrium, each player must be playing a best response to the other's strategy:

$$x_1 = (c + x_2)/2, x_2 = (c + x_1)/2$$

- Solving for x_1 and x_2, we get $x_1 = x_2 = c$.

Ch. 4.2.1: Oligopoly

- Let's go back to the problem of the firm.
- We saw two types of industries: perfect competition, where all firms are price-takers, and monopoly, where one firm faces the entire market.
- In between the cases of perfect competition and pure monopoly, there will be multiple firms, each one with some market power.
- Firms must behave *strategically*, that is, they must take other firms' actions into account when choosing their own behavior.
- We need some theory of how firms will behave in this situation. One possibility is that firms will *collude*, that is, they will work together to extract as much total profit as possible.
- However, each individual firm always has an incentive to increase its own profits and cheat on the collusive agreement.

- Suppose there are J firms, each firm producing output q^j.
- Each firm's profit is negatively affected by an increase in the output of any other firm:

$$\Pi^j(q^1, ..., q^j, ..., q^J), \frac{\partial \Pi}{\partial q^k} < 0 \qquad \text{for } j \neq k$$

- Let $\overline{\boldsymbol{q}} = (q^1, ..., q^j, ..., q^J)$. Joint profits are maximized when the first-order condition is satisfied:

$$\frac{\partial \Pi^k(\overline{\boldsymbol{q}})}{\partial q^k} + \sum_{j \neq k} \frac{\partial \Pi^j(\overline{\boldsymbol{q}})}{\partial q^k} = 0 \qquad \text{for } k = 1, ..., J$$

- Each firm's profit must be increasing in its own quantity produced:

$$\frac{\partial \Pi^k(\overline{\boldsymbol{q}})}{\partial q^k} \geq 0$$

Oligopoly

$$\frac{\partial \Pi^k(\overline{\boldsymbol{q}})}{\partial q^k} \geq 0$$

- Therefore, each firm can increase its own profits by cheating on the collusive agreement, while decreasing profits for everyone else.
- We will assume that firms follow the *Nash equilibrium* solution concept: every agent maximizes his own payoff, given the actions of all other agents.
- If all agents' actions form a Nash equilibrium, then no agent has an incentive to deviate by acting alone.

Oligopoly: Nash Equilibrium Solution

- Applying the Nash equilibrium concept to the market, all firms' choices of output is a Nash equilibrium if:
- Each firm is maximizing its own profit, given the profit-maximizing actions of all other firms.
- That is, each firm's first-order condition must be satisfied individually, for all firms simultaneously.
- The collusive outcome does not satisfy this condition, since all firms can still increase profits by increasing output.
- For the output vector $\boldsymbol{q}^*$ to be a Nash equilibrium:

$$\frac{\partial \Pi^k(\boldsymbol{q}^*)}{\partial q^k} = 0 \qquad \text{for } k = 1, ..., J$$

Cournot Oligopoly

- In this model of oligopoly, each firm chooses its *quantity*.
- Suppose there are J identical firms, and there is no entry of additional firms.
- Each firm has identical costs: $C(q^j) = cq^j, c \geq 0$
- Market price depends on the total output sold by all firms. Assume inverse market demand is:

$$p = a - b\sum_{j=1}^{J} q^j, a > 0, b > 0, a > c$$

- Profit for each firm j is:

$$\Pi^j(q^1, ..., q^j) = \left(a - b\sum_{j=1}^{J} q^j\right) q^j - cq^j$$

Cournot Oligopoly

- We want to find a vector of outputs $(\overline{q}^1, ..., \overline{q}^J)$ that maximizes each firm's individual profits simultaneously.
- This is called a *Cournot-Nash equilibrium*.

$$\frac{\partial \Pi^j}{\partial q^j} = a - 2b\overline{q}^j - b\sum_{k \neq j} \overline{q}^k - c = 0$$

$$b\overline{q}^j = a - c - b\sum_{k=1}^{J} \overline{q}^k$$

- All firms will produce the same amount in equilibrium.

Cournot Oligopoly

$$b\overline{q}^j = a - c - b\sum_{k=1}^{J}\overline{q}^k$$

- All firms will produce the same amount in equilibrium, denote it as $\overline{q}$.

$$b\overline{q} = a - c - Jb\overline{q}$$

$$\overline{q} = \frac{a-c}{b(J+1)}$$

  - Each firm's output: $\overline{q}^j = \frac{a-c}{b(J+1)}$
  - Total output: $\sum_{j=1}^{J}\overline{q}^j = \frac{J(a-c)}{b(J+1)}$
  - Market price: $\overline{p} = a - \frac{J(a-c)}{J+1} < a$
  - Each firm's profits: $\overline{\Pi}^j = \frac{(a-c)^2}{b(J+1)^2}$

Cournot Oligopoly

- The deviation of price from marginal cost is:

$$\overline{p} - c = \frac{a - c}{J + 1} > 0$$

- As $J \to \infty$, the deviation goes to zero.
- As the number of competitors becomes large, the outcome approaches the perfect competition outcome.

Bertrand Oligopoly

- Suppose that instead of firms choosing quantities, they choose *prices*.
- Assume there are two firms producing a homogeneous good.
- Each firm has the same marginal cost $c > 0$ and no fixed costs.
- As before, suppose market demand is linear in total output $Q = q^1 + q^2 = \alpha - \beta p$.
- Firms simultaneously choose the prices they will charge, then produce all quantity demanded at that price.
- Consumers will *only* buy from the cheapest firm. If both firms have the same price, they split demand.

Bertrand Oligopoly

- Each firm's profit depends on its own price, as well as the price of the other firm.

$$\Pi^1(p^1, p^2) = \begin{cases} (p^1 - c)(\alpha - \beta p^1) & \text{if } c < p^1 < p^2 \\ \frac{1}{2}(p^1 - c)(\alpha - \beta p^1) & \text{if } c < p^1 = p^2 \\ 0 & \text{otherwise} \end{cases}$$

- Firm 1's profit is positive as long as $p^1 > MC = c$.
- The profit function is not differentiable in this case, so we can't use calculus methods.
- Let's consider possible combinations of p^1, p^2 and see if either firm can increase its profits by changing its price, while holding the other firm's price constant.
- If no firm can increase its profits, then p^1, p^2 is a Nash equilibrium.

Bertrand Oligopoly

- Case 1: either p^1 or $p^2 < c$.
  - The firm with the lowest price is getting customers, but is making negative profits. It can increase profits by choosing $p^j = c$. Not a Nash equilibrium.
- Case 2: $p^1 = p^2 = c$.
  - Both firms make zero profit. If firm j increases its price, it gets no customers and zero profits. If it decreases its price, it gets all the customers, but makes negative profits. This is a Nash equilibrium.
- Case 3: $p^1 \geq c, p^2 \geq c$, at least one firm has $p^j > c$.
  - The firm offering the lowest price is at best, splitting the customers and making a positive profit. It can increase profits by changing its price to just below the other firm's price. Not a Nash equilibrium.
- The only Nash equilibrium is when $p^1 = p^2 = c$. Both firms make zero profits.

Bertrand Oligopoly

- Note that with only two firms, the outcome is the same as in perfect competition.
- This is due to the assumption that consumers only buy from the lowest price firm.
- In this case, there is complete substitutability between the goods produced by firms 1 and 2.
- If substitutability is not perfect (e.g. if goods are differentiated), then the higher-price firm retains some customers.

Administrative Stuff

- HW #3 is due at the end of lecture today.
- I will post the solutions and HW #4, due in two weeks, to the website after class.