CUR 412: Game Theory and its Applications, Lecture 5

Prof. Ronaldo CARPIO

March 29, 2016

Administrative Stuff

- The due date for HW 2 is moved back one week, to April 12.
- The midterm will be in class on April 19.
- Midterm will be closed-book, covering Chapters 1-4.
- Previous midterms with solutions are on the class web site.

Review of Last Week

- War of Attrition: a game in which two players dispute an object. Each player chooses a time to concede, and waiting is costly for both players.
- There are two Nash equilibria: one in which player i concedes immediately, and the other player waits long enough to drive player i's payoff to zero, if he decided to wait.
- A *sealed-bid* auction is one where bidders submit their bids simultaneously.
- A *second-price* auction is where the winning bidder, who has the *highest* bid, has to pay the *second-highest* bid.
- Many Nash equilibria. Bidding your true value is a weakly dominant strategy.

Review of Random Variables

- A *random variable* is a variable that can take on different values, according to some probability distribution.
- A *discrete* random variable is random variable that can take on a finite number of values.
- For example, suppose that y is a random variable that can take on two values:
  - $y = 1$ occurs with probability p,
  - $y = 0$ occurs with probability $1 - p$.
- y can represent the outcome of flipping a biased coin that shows 1 on one side and 0 on the other side.
- If we flip this coin a very large number of times, the *frequency* (i.e. the fraction of flips) of showing 1 will be p.

Expected Values

- The *expected value* of a random variable is the weighted sum of all possible outcomes, weighted by the probability of occurrence.
- For the biased coin example, the expected value is:

$$E(y) = Pr(y = 1) \cdot 1 + Pr(y = 0) \cdot 0$$

$$= p \cdot 1 + (1 - p) \cdot 0$$

$$= p$$

- In general, the expected value (EV) of a discrete random variable that can take on n outcomes $y_1, ...y_n$ with probabilities $p_1, ...p_n$ is

$$E(y) = p_1 y_1 + p_2 y_2 + ... + p_n y_n$$

- The sum of probabilities over all outcomes, $p_1 + ... + p_n$, must add up to 1.

Some Properties of Expected Values

- Suppose we have a random variable X that can take two outcomes: x_1 with probability p, and x_2 with probability $1 - p$.
- The expected value of X, denoted $E(X)$, is $px_1 + (1 - p)x_2$.
- Suppose we can choose p. We would like to know how $E(X)$ changes when we change p.
- First, note that p can range between 0 and 1. When $p = 1, E(X) = x_1$. When $p = 0, E(X) = x_2$.
- $E(X)$ is *linear* function of p as p goes from 0 to 1.

Some Properties of Expected Values

- $E(X)$ always takes a value that is *between* x_1 and x_2.
- If $x_1 < x_2$, then $E(X)$ decreases linearly as p goes from 0 to 1. $E(X)$ is maximized at $p = 0$.
- If $x_1 > x_2$, then $E(X)$ increases linearly as p goes from 0 to 1. $E(X)$ is maximized at $p = 1$.
- If $x_1 = x_2$, then $E(X)$ is equal to $x_1 = x_2$ for all values of p. Any value of p maximizes $E(X)$.

Some Properties of Expected Values

- Suppose there are n possible outcomes, $x_1, x_2, ...x_n$, with probabilities $p_1, p_2, ...p_n$, where $\sum_i^n p_i = 1$, and each $p_i \geq 0$.
- $E(X) = p_1 x_1 + ... + p_n x_n$
- As before, the value of $E(X)$ will always be in between the smallest and largest values of x_i.
- If $p_i = 1$ and $p_j = 0$ for $i \neq j$, then $E(X) = x_i$.
- If there is a single largest x_i, then $E(X)$ is maximized when $p_i = 1$ and $p_j = 0$.

Some Properties of Expected Values

- Suppose there are multiple largest values: for example, suppose $x_1 = x_2 = \ldots = x_k > x_{k+1} > \ldots > x_n$.
- Then, $E(X)$ is maximized when all the probability is allocated to $x_1, \ldots x_k$ and zero probability is allocated to the other values:

$$p_1 + p_2 + \ldots + p_k = 1, p_{k+1} = p_{k+2} = \ldots = p_n = 0$$

- Conversely, suppose we know that $E(X)$ is maximized with respect to $p_1, \ldots p_n$, and that $p_1 \ldots p_k$ are nonzero, while the rest of p_i's are zero.
- Then, it must be that the x_i's corresponding to $p_1 \ldots p_k$ are equal, and greater than or equal to the other x_i's.

$$x_1 = x_2 = \ldots = x_k \geq x_j \qquad \text{for } j > k$$

Decision-Making Under Uncertainty

- Previously, we used the concept of preferences (i.e. a ranking) over action profiles.
- Preferences can be represented by a payoff (or utility) function that assigns a higher value to more preferred profiles.
- Now, we will allow profiles to have payoffs that are *uncertain*, i.e. a random variable.
- The player's action might affect the probability distribution of the payoff.
- How can we define preferences over uncertain payoffs?
- We will assume that players rank uncertain payoffs based on the *expected value* of their payoff.

A R&D Problem

- Suppose you are a firm that is deciding whether to spend money on researching a new product.
- Actions: *Go* (spend), *Status Quo* (i.e. don't spend).
- If you spend, the probability of successfully developing a new product is 0.75.
- If you don't spend, the probability of successfully developing a new product is 0.5.
- The cost of spending is 1.
- If you successfully develop a new product, you make a profit of 10.
- If you don't successfully develop a new product, you make a profit of 0.

Decision Tree

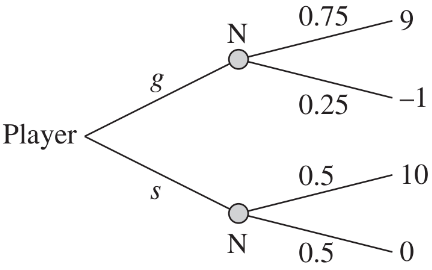

- Each action g, s results in a *lottery*, i.e. a payoff that is a random variable.
- We assume that players prefer (i.e. rank) lotteries based on their expected payoffs.
- Expected payoff of g: $0.75 \cdot 9 + 0.25 \cdot -1 = 6.5$
- Expected payoff of s: $0.5 \cdot 10 + 0.5 \cdot 0 = 5$
- Therefore, a player with our assumed preferences over lotteries will choose g.

von Neumann-Morgenstern Preferences

- Players that value lotteries according to expected payoff are said to have *von Neumann-Morgenstern* preferences, or vNM.
- However, that is not the only way to value lotteries. For example, suppose you have a choice between two lotteries:
- Lottery 1:
  - With probability 1, you win 900,000.
- Lottery 2:
  - With probability 0.9, you win 1,000,000.
  - With probability 0.1, you win 0.
- These two lotteries have the same expected payoff, but most people would prefer the first lottery.
- This is an example of *risk aversion*: if two lotteries have the same expected payoff, the one with less variance is preferred.

Example: Matching Pennies

	Head	*Tail*
Head	1,-1	-1,1
Tail	-1,1	1,-1

- Recall the game of Matching Pennies.
- There is no Nash equilibrium in this game.
- This is a zero-sum game (one in which the sum of payoffs in each outcome adds up to 0).
- Each player wants to "outguess" the other player.
- Other examples:
  - Rock, Paper, Scissors
  - Serving in Tennis

Mixed Strategies

- Before, we assumed a player could only choose one action out of *Head*, *Tail*.
- Now, we will allow players to choose a *probability distribution* over their set of actions (called a *mixed strategy*).
- Suppose each player chooses a probability of playing *Head*.
  - Player 1 chooses p, will play *Head* with probability p, *Tail* with probability $1 - p$.
  - Player 2 chooses q, will play *Head* with probability q, *Tail* with probability $1 - q$.
- A mixed strategy that puts 100% probability on one action is called a *pure strategy*.
- Each player prefers outcomes based on expected payoff.

Nash Equilibrium when Choosing Probabilities

- Before, Nash equilibrium was an action profile where no player had an incentive to deviate.
- Now, when players choose p, q, is there a profile of probability distributions where no player can get a higher expected payoff by changing his probability?
- Consider $p = 0.5, q = 0.5$.
- Player 1's expected payoff for playing *Head*, *Tail*, given $q = 0.5$:
  - EV of *Head*: $0.5 \cdot 1 + 0.5 \cdot (-1) = 0$
  - EV of *Tail*: $0.5 \cdot (-1) + 0.5 \cdot 1 = 0$
- So if $q = 0.5$ is taken as given, any choice of p gives the same expected payoff to Player 1.
- Likewise, for Player 2, if $p = 0.5$ is taken as given, any choice of q gives the same expected payoff.
- Therefore, there is no incentive to deviate.

Nash Equilibrium when Choosing Probabilities

- What about other values of p, q?
- Player 1's expected payoff when using p, taking q as given:

$$EV = p \cdot (\text{EV of playing } Head) + (1-p) \cdot (\text{EV of playing } Tail)$$
$$= p(q \cdot 1 + (1-q) \cdot (-1)) + (1-p)(q \cdot (-1) + (1-q) \cdot 1)$$
$$= p(-1+2q) + (1-p)(1-2q)$$
$$= (1-2p)(1-2q)$$

- If $q < 0.5$, a decrease in p increases EV. If $p > 0$, there is an incentive to deviate.
- If $q > 0.5$, an increase in p increases EV. If $p < 1$, there is an incentive to deviate.
- Likewise, if $p \neq 0.5$, player 2 has an incentive to deviate by changing q.
- $p = q = 0.5$ is the only equilibrium.

Mixed Strategy Nash Equilibrium

- $p = q = 0.5$ is an example of a Nash equilibrium in mixed strategies.
- Before, when players only chose pure strategies, we said a Nash equilibrium was consistent with a steady state of interactions when each player was drawn from a different population.
- With mixed strategies, a Nash equilibrium can be thought of as a steady state when a fraction p of the population that Player 1 is drawn from always plays *Head*, and the other fraction always plays *Tail*.
- Or, we can think of Player 1 as an individual who flips a coin before choosing his action, and plays *Head* with probability p.

Definition of Mixed Strategy

- **Definition**: A *mixed strategy* of a player in a strategic game is a probability distribution over the player's set of actions.
- We will use α to denote a profile of mixed strategies (i.e. a list of the mixed strategies of all players).
- $\alpha_i(a_i)$ is the probability that player i will play action a_i.
- For the Matching Pennies example, with p, q:
  - Player 1's mixed strategy is denoted as: $\alpha_1(Head) = p, \alpha_1(Tail) = 1 - p$
  - Player 2's mixed strategy is denoted as: $\alpha_2(Head) = q, \alpha_2(Tail) = 1 - q$
- If $\alpha_i(a_i) = 1$ for some action a_i, then this is a *pure strategy*.

Definition of Mixed Strategy Nash Equilibrium

- **Definition**: The mixed strategy profile α^* in a strategic game with vNM preferences is a **mixed strategy Nash equilibrium** if, for each player i and every mixed strategy α_i of player i, the expected payoff to player i of α^* is at least as large as the expected payoff of α_i, taking other players' mixed strategies as given.

$$U_i(\alpha^*) \geq U_i(\alpha_i, \alpha_i^*)$$

$$\text{for every mixed strategy } \alpha_i \text{ of player } i, \text{ for all players } i$$

- $U_i(\alpha)$ is the expected payoff to player i of the mixed strategy profile α.

Best Response Functions

- $B_i(\alpha_{-i})$ denotes the set of player i's best mixed strategies, when the list of other players' mixed strategies is given by α_{-i}.
- A profile of mixed strategies α^* is a MSNE if and only if every player's mixed strategy is a best response to the other players' mixed strategies:

$$\alpha_i^* \text{ is in } B_i(\alpha_{-i}^*) \quad \text{for every player } i$$

Two-Player, Two-Action Games

- In Matching Pennies, we saw that the best response was either:
  - A single pure strategy (e.g. if Player 2 plays *Head* with probability 1, the best response is to play *Head* with probability 1)
  - Or, the set of *all* mixed strategies. (if Player 2 plays $q = 0.5$, any value of p is a best response).
- True for any two-player, two-action game.
- This is due to the form of the payoff functions.

2-Player, 2-Action Games

- Consider a generic two-player, two-action game.
  - Player 1's actions: T, B
  - Player 2's actions: L, R
- Let α_1 be Player 1's mixed strategy, in which he chooses T with probability p.
- Similarly, let α_2 be Player 2's mixed strategy, in which he chooses L with probability q.
  - $\alpha_1(T) = p, \alpha_1(B) = 1 - p$
  - $\alpha_2(L) = q, \alpha_2(R) = 1 - q$
- Assume that each player's choices are *independent* of each other.
- The probability of a pair (a_1, a_2) being chosen is the product of the individual probabilities:

$$\text{Prob}(a_1, a_2) = \text{Prob}(a_1) \cdot \text{Prob}(a_2)$$

2-Player, 2-Action Games

- The probability of occurrence for each outcome is given by:

	L	R
T	pq	$p(1-q)$
B	$(1-p)q$	$(1-p)(1-q)$

- Let $u_1(a_1, a_2)$ denote Player 1's Bernoulli payoff function (i.e. the payoff if (a_1, a_2) were to occur with certainty).
- Player 1's expected payoff to the mixed strategy pair (α_1, α_2) is:

$$pq \cdot u_1(T, L) + p(1-q) \cdot u_1(T, R) + (1-p)q \cdot u_1(B, L) + (1-p)(1-q) \cdot u_1(B, R)$$

$$= p \underbrace{\left[q \cdot u_1(T, L) + (1-q) \cdot u_1(T, R)\right]}_{\text{EV of pure strategy } T}$$

$$+ (1-p) \underbrace{\left[q \cdot u_1(B, L) + (1-q) \cdot u_1(B, R)\right]}_{\text{EV of pure strategy } B}$$

Linearity of Expected Payoffs

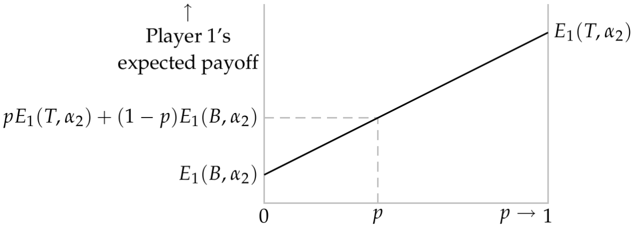

- Let $E_1(T, \alpha_2)$ denote the expected payoff of playing the pure strategy T, when Player 2 plays the mixed strategy α_2.
- Likewise, $E_1(B, \alpha_2)$ denotes the expected payoff of playing the pure strategy B.
- Player 1's expected payoff to (α_1, α_2) is $p \cdot E_1(T, \alpha_2) + (1 - p) \cdot E_1(B, \alpha_2)$.
- This is a linear function of p. For $0 \leq p \leq 1$, the value must lie between $E_1(T, \alpha_2)$ and $E_1(B, \alpha_2)$.

Linearity of Payoffs

- Linearity of the expected payoff implies that one of the following must be true for Player 1:
  - If $E_1(T, \alpha_2) > E_1(B, \alpha_2)$: unique best response is pure strategy T
  - If $E_1(T, \alpha_2) < E_1(B, \alpha_2)$: unique best response is pure strategy B
  - If $E_1(T, \alpha_2) = E_1(B, \alpha_2)$: all mixed strategies (i.e. any value of p) are best responses
- A mixed strategy with $0 < p < 1$ cannot be a *unique* best response; either it is not a best response, or all values of p are best responses.

Matching Pennies: Best Response Function

- Let's construct the best response function for Matching Pennies.
  - $E_1(H, \alpha_2) = q \cdot 1 + (1 - q)(-1) = 2q - 1$
  - $E_1(T, \alpha_2) = q(-1) + (1 - q) \cdot 1 = 1 - 2q$
- If $q < 1/2$, $E_1(T, \alpha_2) > E_1(H, \alpha_2)$: best response is pure strategy T (i.e. $p = 0$)
- If $q > 1/2$, $E_1(T, \alpha_2) < E_1(H, \alpha_2)$: best response is pure strategy H (i.e. $p = 1$)
- If $q = 1/2$, $E_1(T, \alpha_2) = E_1(H, \alpha_2)$: all values of p are a best response

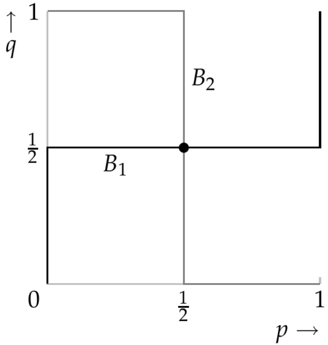

$$B_1(q) = \begin{cases} \{0\} & \text{if } q < 1/2 \\ \{p : 0 \leq p \leq 1\} & \text{if } q = 1/2 \\ \{1\} & \text{if } q > 1/2 \end{cases}$$

- No Nash equilibrium in pure strategies, but there is one in mixed strategies.

Example: BoS

	B	S
B	2,1	0,0
S	0,0	1,2

- Suppose p, q are the probabilities assigned to B by Player 1 and Player 2, respectively.
  - $E_1(B, \alpha_2) = 2 \cdot q + 0 \cdot (1 - q) = 2q$
  - $E_2(S, \alpha_2) = 0 \cdot q + 1 \cdot (1 - q) = 1 - q$
- If $2q > 1 - q \rightarrow q > 1/3$, unique best response is pure strategy B
- If $q < 1/3$, unique best response is pure strategy S
- If $q = 1/3$, all mixed strategies are best responses

BoS: Best Response Function

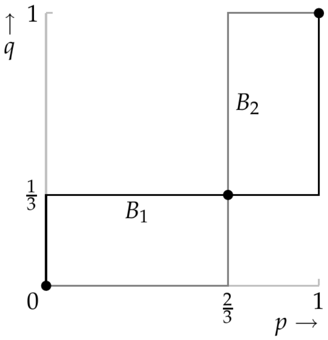

$$B_1(q) = \begin{cases} \{0\} & \text{if } q < 1/3 \\ \{p : 0 \leq p \leq 1\} & \text{if } q = 1/3 \\ \{1\} & \text{if } q > 1/3 \end{cases}$$

- There are two Nash equilibria in pure strategies, and one new equilibrium in mixed strategies.

Next Lecture

- Please read the rest of Chapter 4.
- The due date for HW 2 is moved back one week, to April 12.
- The midterm will be in class on April 19.
- Midterm will be closed-book, covering Chapters 1-4.
- Previous midterms with solutions are on the class web site.