Advanced Microeconomic Analysis, Lecture 10

Prof. Ronaldo CARPIO

May 22, 2017

Administrative Stuff

- I will be traveling on June 5, so there will be no class that week.
- For the three remaining classes (including today), lecture will be from 17:00 - 20:30.
- Next week's class has been moved to Saturday (May 27).
- HW #4 is due on Saturday.

Review of Last Week

- Previously, we have looked at situations with a single decision-maker, which can be solved through optimization (consumer's problem, firm's problem).
- A *strategic* situation is one with more than one decision-maker, and each decision-maker's actions have to be taken into account.
- A *strategic form game* is a way to model a strategic situation:
  - n players, $i = 1, ..., n$
  - each player i has a set of strategies, S_i
  - each player has a payoff function $u^i(\cdot)$
- We want to be able to predict what the outcome of a game will be.
- This requires making some kind of assumption about how the players behave and what they know.

Review of Last Week

- The first kind of prediction we can make is based on assuming players are *rational* (i.e. payoff-maximizing).
- A *strictly dominant* strategy is a strategy that gives the highest payoff, no matter what other players do.
- A rational player will always play a strictly dominant strategy, if it exists.
- Strategy s^1 *strictly dominates* s^2 if s^1 gives a strictly higher payoff than s^2, no matter what other players do.
- A strategy s is *strictly dominated*, if some other strategy strictly dominates s.
- A rational player will never play a strictly dominated strategy.
- If we assume that players know other players are rational, we can eliminate strictly dominated strategies from the game.

Review of Last Week

- We can *iterativly* continue eliminating strictly dominated strategies until there are none left.
- This process is called *iterative elimination of strictly dominated strategies* (IESDS).
- If a single outcome remains, then we have made a prediction about the outcome of the game.
- Similarly, strategy s^1 *weakly dominates* s^2 if s^1 gives a strictly higher payoff than s^2 in at least one situation, and never gives a lower payoff than s^2.
- Weakly dominant and weakly dominated strategies are defined similarly.
- We can iteratively eliminate weakly dominated strategies as well.

Review of Last Week

- Many games will not have a single outcome after iterative elimination.
- Another type of solution concept, that requires more assumptions, is *Nash equilibrium* (NE).
- A *best response* of player i to the strategies s_{-i} chosen by the other players is a strategy that gives the highest possible payoff when others play s_{-i}.
- A NE is a strategy profile (i.e. the combination of every player's chosen strategy) such that each player is playing a best response to the strategies of the other players.
- At a NE, no player has an *incentive to deviate*, (that is, cannot get a strictly higher payoff with a different choice) when the other players' strategies are kept constant.

- We saw two types of oligopoly models: Cournot and Bertrand.
- In Cournot oligopoly, firms choose their output quantity.
- As the number of firms goes to infinity, the outcome approaches perfect competition ($P = MC$, zero profits).
- In Bertrand oligopoly, firms choose their price.
- If the goods produced by each firm are perfect substitutes (i.e. consumers will buy based on lowest price only), then two firms are enough to reach the perfect competition outcome.
- Let's look at a model of price competition with imperfect substitution.

Exercise 4.13: Duopoly with Imperfect Substitutes

- There are two firms $i = 1, 2$ producing goods q_1, q_2.
- Each firm has a constant MC of 20 and no fixed cost.
- Each firm faces its own demand curve:

$$p_1 = 20 + \frac{1}{2}p_2 - q_1, \qquad p_2 = 20 + \frac{1}{2}p_1 - q_2$$

- Note that the demand of each good depends on the price of the other good, but does not go to zero if the other good is cheaper.
- Thus, q_1 and q_2 are imperfect substitutes for each other.
- We want to find the Cournot equilibrium.

- Rearrange the demand curve equation to get $q_i = 20 + \frac{1}{2}p_j - p_i$.
- Since each firm faces its own demand curve, choosing q_i is equivalent to choosing p_i. We will use p_1, p_2 as the choice variables
- Each firm's profit is given by: $\pi_i = p_i q_i - 20q_i = (p_i - 20)q_i$

$$\pi_1(p_1, p_2) = (p_1 - 20)(20 + \frac{1}{2}p_2 - p_1) = -p_1^2 + (40 + \frac{p_2}{2})p_1 - (400 + 10p_2)$$

- We want to find the best response. Since the profit function is concave, we can set the derivative to 0:

$$\frac{\partial \pi_1}{\partial p_1} = -2p_1 + 40 + \frac{p_2}{2} = 0$$

$$p_1^*(p_2) = \frac{80 + p_2}{4}, \qquad p_2^*(p_1) = \frac{80 + p_1}{4}$$

$$p_1^* = p_2^* = \frac{80}{3}, \qquad q_1^* = q_2^* = \frac{20}{3}$$

Some Notable Games

- Let's look at some notable applications of game theory.
- Imagine this situation:
    - There are two suspects in a crime.
    - Each suspect can be convicted of a minor offense, but can only be convicted of a major offense if the other suspect *Confesses*.
    - Each suspect can choose to be *Quiet* or *Confess*.
    - If both stay quiet, each gets 1 year in prison.
    - If only one suspect confesses, he goes free while the other suspect gets 4 years.
    - If both suspects confess, they both get 3 years.

Modeling the Prisoner's Dilemma

- Each player's set of strategies is Q, C.
- Suspect 1's preferences, from best to worst:
  - $(C, Q) > (Q, Q) > (C, C) > (Q, C)$
- Suspect 2's preferences, from best to worst:
  - $(Q, C) > (Q, Q) > (C, C) > (C, Q)$
- We can represent preferences with payoff function that assigns a utility to each outcome:
  - Suspect 1:
    $u_1(C, Q) = 3, u_1(Q, Q) = 2, u_1(C, C) = 1, u_1(Q, C) = 0$
  - Suspect 2:
    $u_2(C, Q) = 0, u_2(Q, Q) = 2, u_2(C, C) = 1, u_2(Q, C) = 3$
- Any set of numbers will work, as long as they maintain the same ranking of outcomes.

Prisoner's Dilemma

		Player 2	
		Q	C
Player 1	Q	2,2	0,3
	C	3,0	1,1

- The unique Nash equilibrium is (C, C).
- Therefore, the prediction is that both players will choose *Confess*.
- This results in an outcome that is sub-optimal for both players.

Prisoner's Dilemma

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

- The Prisoner's Dilemma has been applied to many other situations (e.g. an arms race between two countries, sharing a common resource).
- The strategies of each player are generally called "cooperate" or "defect".
- "Defect" is the selfish choice, which results in a higher payoff for the defecting player.
- "Cooperate" is the altruistic choice, which results in a higher payoff for the other player.
- Both players would be better off if they both chose *Cooperate*, but individual self-interest leads them to choose *Defect*, which makes both players worse off.

Prisoner's Dilemma

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

- Thus, game theory would seem to predict that people should behave selfishly in the real world.
- However, in the real world, we observe that cooperative behavior can be sustained.
- One possible explanation is that in the real world, people repeatedly interact, while this game is a *one-shot* game.
- The theory of repeated games shows that it is theoretically possible to sustain cooperation, even with completely self-interested players.

Hotelling's Model of Electoral Competition

- This is a widely used model in political science and industrial organization, Hotelling's "linear city" model.
- Players choose a location on a line; payoffs are determined by how much of the line is closer to them than other players.
- Here, location represents a position on a *one-dimensional* political spectrum, but it can also represent physical space or product space.

Location on Political Spectrum

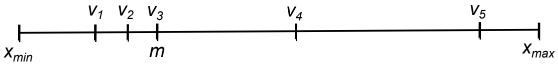

- Political position is measured by position on an interval of numbers
- x_{min} is the most "left-wing" position, x_{max} is the most "right-wing" position
- Voters are located at fixed positions somewhere on the line. This position represents their "favorite position".
- In this example, there are five voters with favorite positions at $v_1 \dots v_5$.
- The *median* position m is the position such that half of voters are to the left or equal to m, and the other half are to the right or equal to m.
- Voters dislike positions that are farther away from them on the line. They are indifferent between positions to their left and right that have the same distance.

Attracting Voters Based on Position

- Candidates can choose their position.
- Assume that voters vote for a candidate based only on distance to the voter's position. They always vote for the closest candidate.
- If there is a tie (two candidates with the same distance), the candidates will split the vote.
- Therefore, each candidate will attract all voters who are closer to him than any other candidate.

Attracting Voters Based on Position

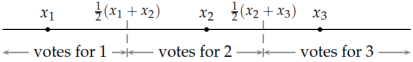

- Suppose there are three candidates who choose positions at x_1, x_2, x_3.
- All voters to the left of x_1 will vote for x_1. Likewise, all voters to the right of x_3 will vote for x_3.
- Between candidates x_1 and x_2, each candidate will attract voters up to the midpoint $(x_1 + x_2)/2$.
- The candidate that attracts the most votes wins. Ties are possible.
- Candidates' most preferred outcome is to win. A tie is less preferable; the more the tie is split, the less preferred.
- Losing is the least preferable outcome.

Candidates' Payoff Function

- Payoffs can be represented by this function:

$$u_i(x_1, \ldots x_n) = \begin{cases} n & \text{if candidate wins} \\ k & \text{if candidate ties with } n-k \text{ other candidates} \\ 0 & \text{if candidate loses} \end{cases}$$

- Definition of Hotelling's Game of Electoral Competition:
  - Players: the candidates
  - Strategies: each candidate can choose a position (a number) on the line
  - Preferences: Each candidate's payoff is given by the function above.

Two Candidates

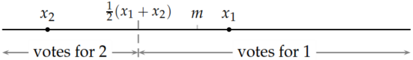

- Suppose there are two candidates that choose positions x_1, x_2.
- The median position (half of voters are on the left, half on the right) is m.
- Let's examine the best response function of player 1 to x_2.
- Case 1: $x_2 < m$
  - Player 1 wins if $x_1 > x_2$ and $(x_1 + x_2)/2 < m$. Every position between x_j and $2m - x_j$ is a best response.
- Case 2: $x_2 > m$
  - By the same reasoning, every position between $2m - x_j$ and x_j is a best response.
- Case 3: $x_2 = m$
  - Choosing m results in a tie; any other choice results in a loss. Therefore, $x_1 = m$ is the best response.

Best Response Function

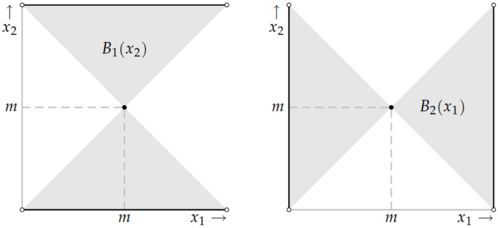

- Best response function is:

$$B_1(x_2) = \begin{cases} \{x_1 : x_2 < x1 < 2m - x2\} & \text{if } x_2 < m \\ \{m\} & \text{if } x_2 = m \\ \{x_1 : 2m - x_2 < x1 < x2\} & \text{if } x_2 > m \end{cases}$$

- Unique Nash equilibrium is when both candidates choose m.

Direct Argument for Nash Equilibrium

- At (m, m), any deviation results in a loss.
- At any other position:
  - If one candidate loses, he can get a better payoff by switching to m.
  - If there is a tie, either candidate can get a better payoff by switching to m.

Implications of Equilibrium

- Conclusion: competition between candidates drives them to take similar positions at the median favorite position of voters
- In physical or product space: competing firms are driven to locate at the same position, or offer similar products
- This is known as "Hotelling's Law" or "principle of minimum differentiation"
- Requires the one-dimensional assumption on voter/consumer preferences.
- If there is more than one dimension (e.g. consumers care about both price and quality), this result may not hold

Mixed Strategies

- Now, suppose that instead of simply choosing one strategy from their strategy set, players were allowed to choose a *probability distribution* over their strategy set.
- Assume each player's strategy set $S_i = \{s^1, s^2, ...s^I\}$ is finite. A *mixed strategy* m_i for Player i is a probability distribution over S_i:
- a set of probabilities $\{m_i(s^1), m_i(s^2), ..., m_i(s^I)\}$ such that
    - $m_i(s^j) \in [0, 1]$ for $j = 1, ..., I$
    - $\sum_{j=1}^{I} m_i(s^j) = 1$
- Let M_i denote the set of all possible mixed strategies for Player i.
- M_i is the $(I - 1)$-dimensional unit simplex.
- A *pure* strategy is a special case of m_i that assigns probability 1 to a single element of S_i, and zero probability to the other elements of S_i.

Mixed Strategies

- Let $M = \times_{i=1}^{N} M_i$ denote the set of joint mixed strategies.
- Then, an element of M, $m \in M$, is a joint mixed strategy.
- We assume that players rank $m \in M$ based on their *expected utility*:

$$u_i(m) = \sum_{s \in S} m_1(s_1)...m_N(s_N)u_i(s)$$

- where $s = (s_1, ..., s_N)$, and $m_1(s_1)...m_N(s_N)$ is the joint probability that s is played, based on m.

Mixed Strategy Nash Equilibrium

- A *mixed strategy Nash equilibrium* is an equilibrium where no player has an incentive to deviate by changing his mixed strategy.
- **Def 7.9**: Given a finite strategic form game $G = (S_i, u_i)_{i=1}^N$, a mixed joint strategy $\hat{m}$ is a mixed strategy Nash equilibrium if, for each player i:

$$u_i(\hat{m}) \geq u_i(m_i, \hat{m}_{-i}) \qquad \text{for all } m_i \in M_i$$

- Now, we will drop the word "mixed". When we say "strategy", we mean "mixed strategy" and when we say "Nash equilibrium", we mean "mixed strategy Nash equilibrium".
- If we want to specifically refer to a situation without randomization, we will say "pure strategy" and "pure strategy Nash equilibrium".

Example: Football Penalty Kick

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

- Suppose Player 1's strategy is: play L with probability p and R with probability $1 - p$.
- Suppose Player 2's strategy is: play L with probability q and R with probability $1 - q$.
- Player 1's expected payoff to the pure strategy L is $q \cdot 1 + (1 - q) \cdot (-1) = 2q - 1$
- Player 1's expected payoff to the pure strategy R is $q \cdot (-1) + (1 - q) \cdot 1 = 1 - 2q$
- Player 1's expected payoff to the joint strategy ((p, 1-p), (q, 1-q)) is $p(2q - 1) + (1 - p)(1 - 2q) = (2p - 1)(2q - 1)$

Example: Football Penalty Kick

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

- Player 1's expected payoff to the joint strategy $((p, 1-p), (q, 1-q))$ is $p(2q-1) + (1-p)(1-2q) = (2p-1)(2q-1)$
- Suppose $q < 0.5$. Then Player 1's unique best response is $p = 0$.
- Suppose $q > 0.5$. Then Player 1's unique best response is $p = 1$.
- Suppose $q = 0.5$. Then Player 1's set of best responses is $p \in [0, 1]$.

Example: Football Penalty Kick

- Player 2's expected payoff to the pure strategy L is $p \cdot (-1) + (1-p) \cdot 1 = 1 - 2p$
- Player 2's expected payoff to the pure strategy R is $p \cdot 1 + (1-p) \cdot (-1) = 2p - 1$
- Player 2's expected payoff to the joint strategy ((p, 1-p), (q, 1-q)) is $q(1-2p) + (1-q)(2p-1) = -(2p-1)(2q-1)$
- Suppose $p < 0.5$. Then Player 2's unique best response is $q = 1$.
- Suppose $p > 0.5$. Then Player 2's unique best response is $q = 0$.
- Suppose $p = 0.5$. Then Player 2's set of best responses is $q \in [0,1]$.

Example: Football Penalty Kick

- The intersection of best responses is $p = 0.5, q = 0.5$.
- For each player, the expected payoff to each of their pure strategies is equal.
  - For Player 1, $1 - 2q = 2q - 1$
  - For Player 2, $1 - 2p = 2p - 1$
- Note that Player 1's expected payoffs are determined by Player 2, and vice versa.
- A game can have a (mixed) NE even if no pure NE exist.

Tests of Nash Equilibrium

- **Theorem 7.1**: The following statements are true if and only if $\hat{m}$ is a Nash equilibrium:
- **1:** For every player i:
  - $u_i(\hat{m}) = u_i(s_i, \hat{m}_{-i})$ for every $s_i \in S_i$ with positive probability in $\hat{m}_i$
  - $u_i(\hat{m}) \geq u_i(s_i, \hat{m}_{-i})$ for every $s_i \in S_i$ with zero probability in $\hat{m}_i$
- **2:** For every player i, $u_i(\hat{m}) \geq u_i(s_i, \hat{m}_{-i})$ for every $s_i \in S_i$.
- For each player's strategy $\hat{m}_i$, the expected payoffs to elements of S_i played with positive probability must be equal.
- The expected payoffs to elements of S_i played with zero probability must be no greater than that for elements played with positive probability.

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 $px_1 + (1 - p)x_2$.
- $E(X)$ 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 $x_1 = x_2 = ... = x_k > x_{k+1} > ... > x_n$.
- Then, $E(X)$ is maximized when all the probability is allocated to $x_1, ...x_k$ and zero probability is allocated to the other values:

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

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

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

Example 7.1: A Coordination Game

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

- There are two pure strategy NE: (B, B) and (S, S).
- Let Player 1's probability of playing S be p, and Player 2's probability of playing B be q.
- Let's check that $p = q = 1/3$ is a NE.
- Let $E_i(B)$ denote the expected payoff to player i of playing pure strategy B.
- Player 1's expected payoff to pure strategies:
  - $E_1(B) = q \cdot 2 + (1 - q) \cdot 0 = 2q = 2/3$
  - $E_1(S) = q \cdot 0 + (1 - q) \cdot 1 = 1 - q = 2/3$

Example 7.1: A Coordination Game

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

- Player 2's expected payoff to pure strategies:
  - $E_2(B) = (1-p)\cdot 1 + p\cdot 0 = 1-p = 2/3$
  - $E_2(S) = (1-p)\cdot 0 + p\cdot 2 = 2p = 2/3$
- Therefore, the test is satisfied and this is a NE.
- Note that in this game, the expected payoff in the mixed NE is lower than in either of the pure strategy NE.
- In this game, the players would prefer to not behave unpredictably.

Example A 3x3 game

	L	C	R
T	$\cdot, 2$	$3, 3$	$1, 1$
M	$\cdot, \cdot$	$0, \cdot$	$2, \cdot$
B	$\cdot, 4$	$5, 1$	$0, 7$

- Is the strategy pair $((\frac{3}{4}, 0, \frac{1}{4}), (0, \frac{1}{3}, \frac{2}{3}))$ a MSNE?
- The dots indicate irrelevant payoffs (they occur with zero probability).
- Given Player 2's mixed strategy: $p(L) = 0, p(C) = \frac{1}{3}, p(R) = \frac{2}{3}$:
  - $E_1(T) = 0 + \frac{1}{3} \cdot 3 + \frac{2}{3} \cdot 1 = \frac{5}{3}$
  - $E_1(M) = 0 + \frac{1}{3} \cdot 0 + \frac{2}{3} \cdot 2 = \frac{4}{3}$
  - $E_1(B) = 0 + \frac{1}{3} \cdot 5 + \frac{2}{3} \cdot 0 = \frac{5}{3}$
- T, B occur with positive probability in $(\frac{3}{4}, 0, \frac{1}{4})$, and have the same expected payoff when Player 2 plays $(0, \frac{1}{3}, \frac{2}{3})$.
- M occurs with zero probability in $(\frac{3}{4}, 0, \frac{1}{4})$, and has an expected payoff not greater than the expected payoffs to T, B.

Example: A 3x3 game

	L	C	R
T	$\cdot, 2$	$3, 3$	$1, 1$
M	$\cdot, \cdot$	$0, \cdot$	$2, \cdot$
B	$\cdot, 4$	$5, 1$	$0, 7$

- Given Player 1's mixed strategy: $p(T) = \frac{3}{4}, p(M) = 0, p(B) = \frac{1}{4}$:
  - $E_2(L) = \frac{3}{4} \cdot 2 + 0 + \frac{1}{4} \cdot 4 = \frac{5}{2}$
  - $E_2(C) = \frac{3}{4} \cdot 3 + 0 + \frac{1}{4} \cdot 1 = \frac{5}{2}$
  - $E_2(R) = \frac{3}{4} \cdot 1 + 0 + \frac{1}{4} \cdot 7 = \frac{5}{2}$
- C, R occur with positive probability in $(0, \frac{1}{3}, \frac{2}{3})$, and have the same expected payoff when Player 1 plays $(\frac{3}{4}, 0, \frac{1}{4})$.
- L occurs with zero probability in $(0, \frac{1}{3}, \frac{2}{3})$, and has an expected payoff not greater than the expected payoffs to C, R.
- Note: the fact that $E_2(L) = \frac{5}{2}$ does not imply anything about the existence of a MSNE that has a positive probability on L.

Example: Choosing Numbers

- Players 1 and 2 choose a positive integer from $1 \ldots K$.
- If the players choose the same number, Player 2 gets a payoff of -1 and Player 1 gets a payoff of 1.
- Otherwise, both players get a payoff of 0.
- First, show that one NE is if both players choose each integer with equal probability $1/K$.

Example: Choosing Numbers

- Denote Player i's expected payoff to pure strategy j as $E_i(j)$.
- Player 1's expected payoffs to his pure strategies $1, \ldots, K$ are:
- $E_1(1) = E_1(2) = \ldots = E_1(K) = 1/K$
- Player 2's expected payoffs to his pure strategies $1, \ldots, K$ are:
- $E_2(1) = E_2(2) = \ldots = E_2(K) = -1/K$
- All actions with positive probability have the same payoff, so the condition for a NE is satisfied.

Example: Choosing Numbers

- Show there is no other NE.
- Let Player 1's mixed strategy be $(p_1, ..., p_K)$.
- Let Player 2's mixed strategy be $(q_1, ..., q_K)$.
- Player 1's expected payoffs to his pure strategies $1, ..., K$ are:
- $E_1(1) = q_1, E_1(2) = q_2, ..., E_1(K) = q_K$
- Player 2's expected payoffs to his pure strategies $1, ..., K$ are:
- $E_2(1) = -p_1, E_2(2) = -p_2, ..., E_2(K) = -p_K$

Example: Choosing Numbers

- Player 1's expected payoff, given both players' mixed strategies, is:

$$p_1 q_1 + p_2 q_2 + \ldots + p_K q_K$$

- Player 2's expected payoff, given both players' mixed strategies, is the negative of Player 1's expected payoff:

$$-p_1 q_1 - p_2 q_2 - \ldots - p_K q_K$$

- Suppose that Player 1 does not place equal probability on each number: there exists a number i such that p_i that is strictly greater than the other p's.
- Player 2's expected payoff to playing i is $-p_i$, so Player 2 will put zero probability on $i : q_i = 0$.
- However, if $q_i = 0$, then Player 1's expected payoff to i is 0, and Player 1's best response is to put zero probability on i, a contradiction.

What Does It Mean To Play A Mixed Strategy?

- The concept of mixed strategy Nash equilibrium makes some assumptions that may or may not hold in the real world.
- First, players are assumed to know the mixed strategies of other players. How do they know this?
- If we assume a situation that is repeatedly played many times, then we can observe the frequency of each action.
- However, in a one-shot situation, there are no past examples to learn from.
- We can interpret the mixed strategy of other players as a *belief* about their behavior, rather than an empirical frequency.

What Does It Mean To Play A Mixed Strategy?

- Second, do people really randomize their actions in important situations?
- Again, this is plausible in repeated situations, but perhaps not in one-shot situations.
- It can be argued that in games with conflict, I want my actions to be unpredictable by the other player.
- However, in games with cooperation (e.g. a coordination game), it is beneficial to me if the other player can predict my actions.
- Mixed strategies can be interpreted as applying to populations, instead of individuals.

Existence of Nash Equilibrium

- **Theorem 7.2**: Every finite strategic form game has at least one Nash equilibrium.
- We will prove this by constructing a function with a *fixed point* that is the Nash equilibrium.
- A fixed point of a function $F()$ is a point x such that $x = F(x)$.

Brouwer's Fixed Point Theorem

- The proof of existence of Nash equilibrium (any many other equilibria in economics) uses Brouwer's fixed point theorem.
- **Thm A1.11 (Brouwer's fixed point theorem)**: Let $S \subset \mathbb{R}^n$ be a non-empty, closed, bounded, and convex set, and let $f : S \to S$ be a continuous function. Then there exists at least one $\boldsymbol{x}^*$ such that $f(\boldsymbol{x}^*) = \boldsymbol{x}^*$.
- $\boldsymbol{x}^*$ is called a *fixed point* of f, since it does not move under the effect of f.
- In economics, an equilibrium is a type of fixed point in behavior, since agents do not change their behavior at an equilibrium.
- For example, recall that Walrasian equilibrium is defined as a price vector that balances aggregate supply and demand in all markets simultaneously.

Brouwer's Fixed Point Theorem

- Recall that $\boldsymbol{z}(\boldsymbol{p})$ denotes aggregate excess demand, i.e. total demand minus total supply for each good.
- We can define a function $f(\cdot)$ that takes a price vector, and returns another price vector $f(\boldsymbol{p})$:

$$f(\boldsymbol{p}) = \boldsymbol{p} + \boldsymbol{z}(\boldsymbol{p})$$

- That is, if there is more demand than supply for some good $(z_j(\boldsymbol{p}) > 0)$, then $f(\boldsymbol{p})$ increases the price of good j. This should reduce demand.
- A Walrasian equilibrium is a fixed point of f, since $f(\boldsymbol{p}) = \boldsymbol{p}$ implies $z(\boldsymbol{p}) = 0$.
- If we can define an appropriate set of prices that satisfies the conditions (convex, closed, boudned, non-empty), then Brouwer's fixed point theorem guarantees that a Walrasian equilibrium exists.

Existence of Nash Equilibrium

- Let $G = (S_i, u_i)_{i=1}^N$ denote a strategic form game with N players.
- Assume each player has the same number of pure strategies, n.
- Then, for each player, his set of strategies can be labeled $S_i = 1, 2, 3, ..., n$.
- $u_i(j_1, j_2, ..., j_n)$ denotes the payoff to player i when player 1 chooses j_1, player 2 chooses j_2, etc.
- Player i's set of mixed strategies is denoted $M_i = \{(m_{i1}, ..., m_{in}) | m_{ij} \geq 0, \sum_j m_{ij} = 1\}$
- m_{ij} is the probability assigned to player i's pure strategy j.
- Note that M_i is a non-empty, compact (i.e. closed and bounded), and convex set.

Existence of Nash Equilibrium

- Define $f : M \to M$ as follows: for each joint strategy $m \in M$, player i and his pure strategy j:

$$f_{ij}(m) = \frac{m_{ij} + \max(0, u_i(j, m_{-i}) - u_i(m))}{1 + \sum_{j'=1}^{n} \max(0, u_i(j', m_{-i}) - u_i(m))}$$

- m is a joint strategy of all the players.
- $u_i(j, m_{-i}) - u_i(m))$ is the change in payoff that player i can get by *unilaterally* deviating to strategy j, while other players' strategies remain constant.
- $\max(0, u_i(j, m_{-i}) - u_i(m))$ ensures that this change in payoff is not negative.
- Note that for every player i, $\sum_{j=1}^{n} f_{ij}(m) = 1$ and $f_{ij}(m) \geq 0$ for any j.
- Let $f_i(m) = (f_{i1}(m), \ldots, f_{in}(m))$. Therefore, $f_i(m)$ is a probability distribution over n elements, and $f_i(m) \in M_i$.

Existence of Nash Equilibrium

$$f_{ij}(m) = \frac{m_{ij} + \max(0, u_i(j, m_{-i}) - u_i(m))}{1 + \sum_{j'=1}^{n} \max(0, u_i(j', m_{-i}) - u_i(m))}$$

- The effect of $f_{ij}(m)$ is to increase the probability on action j of player i, if it gives a higher payoff than not deviating.
- If action j gives the same or a negative payoff, then the probability of action j is unchanged.
- Let $f(m) = (f_1(m), \ldots, f_N(m))$. Therefore, $f(m)$ is a joint probability distribution, and $f(m) \in M$.
- $f(m)$ is continuous in m.
- By Brouwer's fixed point theorem, f has a fixed point, $\hat{m}$.

Existence of Nash Equilibrium

$$\hat{m}_{ij} = \frac{\hat{m}_{ij} + \max(0, u_i(j, \hat{m}_{-i}) - u_i(\hat{m}))}{1 + \sum_{j'=1}^{n} \max(0, u_i(j', \hat{m}_{-i}) - u_i(\hat{m}))}$$

- We can show that $f(\hat{m}) = \hat{m}$ is actually the Nash equilibrium.
- Since every probability of every player's strategy is unchanged from $\hat{m}$, that must mean that every player's strategy gives the same or a lower payoff.
- Therefore, no player has an incentive to deviate, and $\hat{m}$ must be a (mixed-strategy) Nash equilibrium.
- No matter how many players and how many strategies they have, as long as they are finite, a NE exists.
- If we found that NE did not exist in a large number of games, that would be reason to doubt its usefulness.
- However, since at least one NE must exist in any finite game, that is a reason to believe the concept of NE is useful.

Incomplete Information

- So far, we have assumed that players are perfectly informed about the payoffs of all other players.
- However, in many real-life situations, there is uncertainty about the opponents' payoffs. For example:
  - When you buy an item from a seller, the seller knows the item's quality, but you do not.
  - When two people get into a competition, each person knows his own strength, but not the other person's.
- We will show how to specify this situation as a strategic form game by adding two additional elements.

Player Types

- First, for each player i, we introduce a finite set of "types", T_i, that the player might be.
- For example:
  - A firm might have two types, "low production cost" and "high production cost".
  - A competitor might have two types, "strong" and "weak".
- A player's payoffs for a given joint pure strategy now also depend on his type.
- Let $T = \times_{i=1}^{N} T_i$, the set of joint types.
- Player i's payoff function u_i maps $S \times T$ to a real number.

Beliefs as Probability Distributions over Types

- Second, each player has *beliefs* about what all other players' type may be.
- A belief is a probability distribution over the set of possible types.
- For example, suppose there are 2 players, and Player 2 has two possible types: "weak" and "strong".
- Player 1 has a belief about Player 2's type, $(p, 1-p)$, where p is Player 1's *subjective probability* that Player 2 is "weak".
  - If $p = 0$, then Player 1 is certain that Player 2 is "strong".
  - If $p = 0.5$, then Player 1 thinks that it is equally likely that Player 2 is "weak" or "strong".
  - If $p = 1$, then Player 1 is certain that Player 2 is "weak".
- If Player 1 also has more than one type, we need to specify (possibly) different beliefs about Player 2, for each of Player 1's type.

Review of Probability and Bayes' Theorem

- Suppose we have a random experiment with a range of possible outcomes.
- For example: we roll a 4-sided dice, with sides labeled 1, 2, 3, 4. Call the result X.
- The *universe* U is the set of all possible outcomes of the world.
- In this case, $U = \{1, 2, 3, 4\}$.
- An *event* is a *set* of outcomes.
- Examples of events
  - The set of all possible outcomes, U.
  - The empty set $\emptyset$.
  - The set of outcomes where $X \leq 2$: $\{1, 2\}$.
  - The set of outcomes where X is even:: $\{2, 4\}$.

Laws of Probability

- For any event A, $P(A) \geq 0$.
- $P(U) = 1$.
- $\neg A$ is the set of all outcomes that are *not* in A. $P(\neg A) = 1 - P(A)$
- $A \cup B$ is the event that the outcome is in A *or* in B.
- If A and B are disjoint (i.e. no outcome is in both A and B), then $P(A \cup B) = P(A) + P(B)$.
- This extends to $A_1, A_2, ...A_n$: if these are all disjoint, the $P(A_1 \cup ... \cup A_n) = P(A_1) + ...P(A_n)$.
- $A \cap B$ is the event that the outcome is in A *and* in B.
- Two events are *independent* if $P(A \cap B) = P(A)P(B)$.

Examples

- Consider our 4-sided dice roll, with result X. Assume the dice is unbiased: each side occurs with equal probability $\frac{1}{4}$.
- Let A be the event that $X \leq 2$. Then $P(A) = \frac{1}{2}$.
- Let B be the event that X is even. Then $P(A) = P(2) + P(4) = \frac{1}{2}$.
- $P(A \cap B) = P(2) = \frac{1}{4}$, which is equal to $P(A)P(B)$. Therefore, events A and B are independent.

Conditional Probability

- If we know one event A has occurred, does that affect the probability that another event B has occurred?
- Suppose we are told that event A has occurred. Then $P(A) > 0$, and every outcome outside of A is no longer possible.
- Therefore, the universe is reduced to A. The only part of B which can occur is $A \cap B$.
- Since total probability of the universe must equal 1, the probability of $A \cap B$ must be scaled by $\frac{1}{P(A)}$

$$P(B|A) = \frac{P(B \cap A)}{P(A)}, P(B \cap A) = P(A)P(B|A)$$

Conditional Probability

$$P(B|A) = \frac{P(B \cap A)}{P(A)}, P(B \cap A) = P(A)P(B|A)$$

- $B \cap A$ and $\neg B \cap A$ are disjoint, and their union is A. Therefore, $P(B \cap A) + P(\neg B \cap A) = P(A)$.

$$P(B|A) = \frac{P(B \cap A)}{P(A)} = \frac{P(B \cap A)}{P(B \cap A) + P(\neg B \cap A)}$$

- $P(B|A)$ is the probability of B conditional on A occurring. Equivalently,

$$P(A|B) = \frac{P(A \cap B)}{P(B)} = \frac{P(A \cap B)}{P(A \cap B) + P(\neg A \cap B)}, P(A \cap B) = P(B)P(A|B)$$

Bayes' Theorem

$$P(B|A) = \frac{P(B \cap A)}{P(A)} = \frac{P(B \cap A)}{P(B \cap A) + P(\neg B \cap A)}, P(B \cap A) = P(A)P(B|A)$$

- Using the result $P(B \cap A) = P(B)P(A|B)$ and $P(\neg B \cap A) = P(\neg B)P(A|\neg B)$, we get

$$P(B|A) = \frac{P(A|B)P(B)}{P(A|B)P(B) + P(A|\neg B)P(\neg B)}$$

- This is known as Bayes' Theorem. It is simply substituting in different formulas for $\frac{P(B \cap A)}{P(A)}$.
- Using this formula, we can calculate the probability of any event conditional on any other event, if we know the all of the other probabilities in the formula.

Example

- Consider again the 4-sided dice that gives value $X \in \{1, 2, 3, 4\}$. Suppose that the dice is no longer unbiased:

$$P(1) = P(2) = P(3) = 0.2, P(4) = 0.4$$

- A is the event that $X \leq 2$. $P(A) = 0.4$
- B is the event that X is even. $P(B) = 0.6$
- $P(A \cap B) = 0.2$, which is different from $P(A)P(B) = 0.24$. Therefore, A and B are not independent.

$$P(A|B) = \frac{P(A \cap B)}{P(B)} = \frac{0.2}{0.6} = \frac{1}{3}$$

$$P(B|A) = \frac{P(A \cap B)}{P(A)} = \frac{0.2}{0.4} = \frac{1}{2}$$

Example

- Here's an real-world example of using Bayes' Theorem:
- Suppose in the general population, we observe that 1.4% of women whose age 40-49 develop breast cancer.
- A mammogram is a diagnostic test to predict breast cancer. We know that it has a 10% false positive rate: when the test is used on someone who *does not* have cancer, it says they do 10% of the time.
- The test has a 25% false negative rate: when the test is used on someone who *does* have cancer, it says they do 75% of the time.
- Suppose a woman gets a mammogram and it gives a positive result. What is the probability that she has cancer, knowing nothing else?
- Let B be the event that she has cancer, and A the event that the test gives a positive result.
- We want to find $P(B|A)$, which is:

$$P(B|A) = \frac{P(A|B)P(B)}{P(A|B)P(B) + P(A|\neg B)P(\neg B)} = \frac{1.4\% \times 0.75}{1.4\% \times 0.75 + 98.6\% \times 0.1}$$

Example

$$P(B|A) = \frac{P(A|B)P(B)}{P(A|B)P(B) + P(A|\neg B)P(\neg B)} = \frac{1.4\% \times 0.75}{1.4\% \times 0.75 + 98.6\% \times 0.1}$$

- $P(B|A)$ is 0.096242, about 10 %.
- B is our *prior distribution*: our beliefs *before* we learn any new information.
- After we observe new information from A, we update our beliefs, giving us a *posterior distribution*.
- "Prior" means "before", "posterior" means "after".

Example: A Coordination Game with Different Types

- Consider the coordination game we saw earlier, but now suppose that Player 2 can have two different types.
- Type-1 of Player 2 is the same as before, and prefers to choose the same strategy as Player 1.
- Type-2 of Player 2, on the other hand, prefers to choose a different strategy as Player 1.
- If Player 2 is Type-1, then the payoff matrix is:

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

Example: A Coordination Game with Different Types

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

- If Player 2 is Type-2, then the payoff matrix is as above.
- Let $p_1(t_1)$ denote Player 1's subjective probability that Player 2 is of Type-1. Then, $p_1(t_2) = 1 - p_1(t_1)$.
- We will assume that all players know their own types with certainty.
- Suppose $p_1(t_1) = p_1(t_2) = 0.5$.
- We will only look at pure strategies for what follows. However, Player 1 still uses expected payoffs, where the uncertainty now comes from the type of Player 2.

Example: A Coordination Game with Different Types

- We will treat the 2 types of Player 2 as separate players; a joint strategy for all players has 3 strategies, one each for Player 1, Player 2 Type-1, and Player 2 Type-2.
- We will denote a joint strategy as e.g. $(B, (B, S))$, where the first element is Player 1's strategy, and the second element is (Type 1's strategy, Type 2's strategy).
- Player 1's expected payoff to each joint strategy is:
  - $(B, (B, B))$: 0.5*2 + 0.5*2 = 2
  - $(B, (B, S))$: 0.5*2 + 0.5*0 = 1
  - $(B, (S, B))$: 0.5*0 + 0.5*2 = 1
  - $(B, (S, S))$: 0.5*0 + 0.5*0 = 0
  - $(S, (B, B))$: 0.5*0 + 0.5*0 = 0
  - $(S, (B, S))$: 0.5*2 + 0.5*1 = 0.5
  - $(S, (S, B))$: 0.5*1 + 0.5*0 = 0.5
  - $(S, (S, S))$: 0.5*1 + 0.5*1 = 1

Example: A Coordination Game with Different Types

- We will solve this as a three-player game. We can compute the best response functions and see if there is an intersection.
- The best responses of Player 1 are:
  $B_1(B, B) = B, B_1(B, S) = B, B_1(S, B) = B, B_1(S, S) = S$.
- The best responses of Player 2, both types, only depend on Player 1's strategy.
  - Type 1: $B_{21}(B) = B, B_{21}(S) = S$
  - Type 2: $B_{22}(B) = S, B_{22}(S) = B$
- $(B, (B, S))$ is a Nash equilibrium, since all players are playing best responses to the other players.
- In this equilibrium, Player 1 plays B, Player 2-Type 1 plays B, and Player 2-Type 2 plays S.

Game of Incomplete Information (Bayesian Game)

- **Def 7.10**: A game of incomplete information (also called a Bayesian game) is a tuple $G = (p_i, T_i, S_i, u_i)_{i=1}^N$, where for each player i, the set of types T_i is finite, $u_i : S \times T \to \mathbb{R}$, and for each $t_i \in T_i, p_i(\cdot|t_i)$ is a probability distribution on T_{-i}.
- Here, $p_i(t_{-i}|t_i)$ is the probability that the players aside from i have joint type t_{-i}, conditional on Player i's own type being t_i.
- This allows for the possibility that Player i's own type is not independent of the other players' types: if Player i knows he is a specific type, this may give him more information on the distribution of other players' types.

The Associated Strategic Form Game

- Let $G = (p_i, T_i, S_i, u_i)_{i=1}^N$ be a game of incomplete information.
- We will construct a strategic form game G^* in which each player type in G is a separate player.
- For each player $i \in \{1, ..., N\}$ and each Player i-type $t_i \in T_i$, let t_i be a player in G^* whose finite set of pure strategies is S_i.
- Let $s_i(t_i) \in S_i$ denote the a pure strategy chosen by player $t_i \in T_i$.
- The payoff to player t_i from the joint pure strategy $s^* = (s_1(t_1), ..., s_N(t_N))$ is:

$$v_{t_i}(s^*) = \sum_{t_{-i} \in T_{-i}} p_i(t_{-i}|t_i) u_i(s_1(t_1), ..., s_N(t_N), t_1, ..., t_N)$$

Bayesian-Nash Equilibrium

- A *Bayesian-Nash equilibrium* of a game of incomplete information is a Nash equilibrium of the associated strategic form game.
- By the existence of Nash equilibrium in finite strategic form games, every finite incomplete information game has at least one Bayesian-Nash equilibrium.

Administrative Stuff

- I will be traveling on June 5, so there will be no class that week.
- For the three remaining classes (including today), lecture will be from 17:00 - 20:30.
- Next week's class has been moved to Saturday (May 27).
- HW #4 is due on Saturday.