CAP 5993/CAP 4993
Game Theory

Instructor: Sam Ganzfried
sganzfri@cis.fiu.edu

HW1

- Deadline extended to 1/31
- HW policy:
  - You can discuss general concepts with other students, but must work on the problems individually.
  - List out all resources consulted.
  - Two late days, then 50% credit, then 0%.
  - Homework due at start of class (3:30 PM). Can be emailed.

Central question of game theory

1. An empirical, descriptive interpretation: How do players, in fact, play in a given game?
2. A normative interpretation: How “should” players play in a given game?
3. A theoretical interpretation: What can we predict will happen in a game given certain assumptions regarding “reasonable” or “rational” behavior on the part of the players?

Prisoner’s dilemma experiments

- First year economics students, and students doing disciplines other than economics, overwhelmingly chose to cooperate. But 4th year students in economics tended to not cooperate. Frank et al. concluded, that with "an eye toward both the social good and the well-being of their own students, economists may wish to stress a broader view of human motivation in their teaching."

Ultimatum game experiments

- This paper reports two experiments involving an ultimatum game, conducted in Japan. There were two treatments in our experiments. One was called a cash session and the other was called a point session. The cash session means introducing cash into the ultimatum game. In other words, in a cash session, subjects bargained money in cash but not points or tokens as most prior experiments did. We found that compared to those in the point sessions, **proposers offered more and responders rejected less in the cash sessions**. These evidences imply that a cash effect does exist in the ultimatum game experiments.
- From Shen/Takahashi 2013 paper

Ultimatum game experiments

- The present study examined how the size of the initial endowment ($10, $3,000, and $250,000); social distance (close friend, acquaintance, and unacquainted person), and whether the responder's identity is made known to the proposer affect the behavior of responders in the Ultimatum Game. The amount of money involved in the game proved to be highly significant. **As the size of the endowment increased, responders were willing to accept proportionally smaller offers**. Social distance had an overall effect, with responders expressing a **greater willingness**
- **to accept proportionally smaller offers from people to whom they were closer**. Responders whose identity was known did not behave significantly differently from responders whose identity was unknown by proposers.

Rock-paper-scissors

- The main strategies it employs are history matching, frequency analysis, and random guessing. Its strongest strategy, history matching, searches for a sequence in the past that matches the last few moves in order to predict the next move of the algorithm. In frequency analysis, the program simply identifies the most frequently played move. The random guess is a fallback method that is used to prevent a devastating loss in the event that the other strategies fail. More than ten years later, the top performing strategies on an ongoing rock–paper–scissors programming competition similarly use metastrategies. However, there have been some innovations, such as using multiple history matching schemes that each match a different aspect of the history – for example, the opponent's moves, the program's own moves, or a combination of both. There have also been other algorithms based on Markov chains.

- Researchers at the University of Tokyo have created a robot hand that has a 100% winning rate playing rock–paper–scissors. Using a high-speed camera, the robot recognizes within one millisecond which shape the human hand is making, then produces the corresponding winning shape.

2/3 the average

- This game is a common demonstration in game theory classes, where even economics graduate students fail to guess 0. When performed among ordinary people it is usually found that the winner guess is much higher than 0, e.g., 21.6 was the winning value in a large internet-based competition organized by the Danish newspaper Politiken. This included 19,196 people and with a prize of 5000 Danish kroner.

Domination

	L	M	R
T	1, 0	1, 2	0, 1
B	0, 3	0, 1	2, 0

- A strategy s_i of player i is **strictly dominated** if there exists another strategy t_i of player i such that for each strategy vector s_{-i} in S_{-i} of the other players, $u_i(s_i, s_{-i}) < u_i(t_i, s_{-i})$

- Assumption: A rational player will not choose a strictly dominated strategy.
- Assumption: All players in a game are rational.
- Can a strictly dominated strategy be eliminated under these two assumptions?

- Assumption: The fact that all players are rational is common knowledge among the players.

Prisoner's dilemma

	C	**D**
C	4, 4	0, 5
D	5, 0	1, 1

- Does it matter if we eliminate player 1's C or player 2's C first?
- Theorem: Whenever iterated elimination of strictly dominated strategies leads to a single strategy vector, that outcome is independent of the order of elimination.

Weakly dominated strategies

- Strategy s_i of player i is **weakly dominated** if there exists another strategy t_i of player i satisfying the following two conditions:
  1. For every strategy vector s_{-i} in S_{-i} of the other players, $u_i(s_i, s_{-i}) <= u_i(t_i, s_{-i})$
  2. There exists a strategy vector t_{-i} in S_{-i} of the other players such that $u_i(s_i, t_{-i}) < u_i(t_i, t_{-i})$
- In this case we say that strategy s_i is **weakly dominated** by strategy t_i, and that strategy t_i **weakly dominates** strategy s_i

Guess 2/3 the average game

- All strategies except guessing 0 are removed by iterated weak domination.

Trembling hand principle

- Suppose player 2 chooses L and R with probabilities x and 1-x respectively, where $0 < x < 1$.
- The expected payoff to player 1 if he chooses T is:
  - $x + 2(1-x) = 2-x$
- The expected payoff to player 1 if he chooses B is 2.
- So strategy B gives him a strictly higher payoff than T, so that a rational player 1 facing player 2 who has a ``trembling hand'' will choose B and not T; i.e., he will not choose the weakly dominated strategy.

Order of elimination

- When only strictly dominated strategies are involved in a process of iterated elimination, the result is independent of the order in which strategies are eliminated.

	L	C	R
T	1, 2	2, 3	0, 3
M	2, 2	2, 1	3, 2
B	2, 1	0, 0	1, 0

	L	**C**	**R**
T	1, 2	2, 3	0, 3
M	2, 2	2, 1	3, 2
B	2, 1	0, 0	1, 0

1. T, R, B, C: ML: 2, 2
2. B, L, C, T: MR: 3, 2
3. T, C, R: ML or BL: 2, 2 or 2, 1

	L	C	R
T	0, 6	6, 0	4, 3
M	6, 0	0, 6	4, 3
B	3, 3	3, 3	5, 5

- M is **best reply (best response**) of Player 1 to L.
  - T is his best reply to C and B is his best reply to R.
- (B,R): each strategy is a best reply to the other strategy
- If both players follow (B,R), neither player has a **profitable deviation**

- A strategy vector $s^* = (s^*_1, \ldots, s^*_n)$ is a **Nash equilibrium** if for each player i in N and each strategy s_i in S_i the following is satisfied:

$$u_i(s^*) >= u_i(s_i, s^*_{-i})$$

- The payoff vector $u(s^*)$ is the **equilibrium payoff** corresponding to the Nash equilibrium s^*.
- No player has a **profitable deviation**.
  - Strategy vector s'_i in S_i with $u_i(s'_i, s_{-i}) > u_i(s)$
- Sometimes just called an **equilibrium**

- Alternatively,

- Let s_{-i} be a strategy vector of all the players not including player i. Player i's strategy is termed a **best reply** to s_{-i} if $u_i(s_i,s_{-i}) = \max u_i(t_i, s_{-i})$

- A strategy vector $s^* = (s^*_1,\ldots, s^*_n)$ is a **Nash equilibrium** if s^*_i, is a best reply to s^*_{-i}, for every player i in N.

	L	**C**	**R**
T	0, 6	6, 0	4, 3
M	6, 0	0, 6	4, 3
B	3, 3	3, 3	5, 5

- (B,R) is the **unique** Nash equilibrium.
- Why is (T,L) not a Nash equilibrium?

Prisoner’s dilemma

	C	**D**
C	4, 4	0, 5
D	5, 0	1, 1

Battle of the sexes

	Opera	Football
Opera	3,2	0,0
Football	0,0	2,3

Battle of the Sexes 1

	Opera	Football
Opera	3,2	1,1
Football	0,0	2,3

Battle of the Sexes 2

Rock-paper-scissors

	rock	**paper**	**scissors**
Rock	0,0	-1, 1	1, -1
Paper	1,-1	0, 0	-1,1
Scissors	-1,1	1,-1	0,0

Security game

- Random strategy:
  - ➧ *Increase cost/uncertainty to attackers*

Adversary

Defender

	Target #1	Target #2
Target #1	4, -3	-1, 1
Target #2	-5, 5	2, -1

Chicken

	Swerve	Straight
Swerve	0, 0	-1, +1
Straight	+1, -1	-10, -10

Fig. 2: Chicken with numerical payoffs

Properties

- Stability: Under Nash equilibrium, each player acts to his best possible advantage with respect to the behavior of the other players.
  - If there were to be any expected "solution concept," seems clear that result should be a Nash equilibrium, or else at least one player would not want to follow it.

	L	R
T	0, 0	4, 2
B	3, 5	0, 0

- Self-fulfilling agreement: if there is an “agreement” to play a particular equilibrium, then, even if the agreement is not binding, it will not be breached; no player will deviate from the equilibrium point, because there is no way to profit from any unilateral violation of the agreement.
- Fulfill “agreement” to play either (B,L) or (T,R).

Equilibrium and evolution

- Darwin’s Theory of Evolution: “survival of the fittest”
  - Expect animal (or plant) to choose those traits that grant the greatest possible advantages in the struggle for survival. Animals, of course, are not typically endowed with rational thought and no animal can choose its own genetic inheritance. What actually happens is that those individuals born with traits that are a poor fit relative to the conditions for survival will pass those same characteristics on to their progeny, and over time their numbers will dwindle.
- In other words, the surviving and prevailing traits are a kind of “best reply” to the environment.

Normative perspective

- Arbitrator or judge recommending a certain course of action based on reasonable and acceptable principles. We should expect the arbitrator's recommendation to be an equilibrium point: otherwise (since it is a recommendation and not a binding agreement) there will be at least one agent who will be tempted to benefit from not following his end of the recommendation.

Problems with Nash equilibrium?

- In some games (e.g., certain infinite games or games with “imperfect recall”) there is no equilibrium.
- In some games there are many equilibria (with different payoffs to the players) – not clear how to select between them.
  - We will study several of the main equilibrium “refinements.”
- Even when there is a unique equilibrium, it may not be “recommended” or predicted. E.g., prisoner’s dilemma.
- We will see examples where it is unclear that an equilibrium will be an outcome (e.g., next example, repeated prisoner’s dilemma, and the centipede game).

Nash equilibrium selection

- Is it too strict?
  - Does not exist in all games
  - Might rule out some more “reasonable” strategies
- Not strict enough?
  - Potentially many equilibria to select through
- Just right?

Maxmin security

	L	R
T	2, 1	2, -20
M	3, 0	-10, 1
B	-100, 2	3, 3

Maxmin security

	L	R
T	2, 1	2, -20
M	3, 0	-10, 1
B	-100, 2	3, 3

- Unique equilibrium is (B,R) with payoff (3,3).
- But would player 1 really choose B?
- Since (B,L) is catastrophic, he may prefer T
  - T guarantees him only 2, but avoids possibility of -100
- If P2 is aware of P1's hesitation and believes reasonable chance P1 will pick T, he will pick L. This increases motivation of P1 to pick T.

Maxmin security

- Players may wish to be "safe" and guarantee the best possible result without "relying" on the rationality of the other players, and even making the most pessimistic assessment of their potential behavior.
- If player i chooses s_i the worst possible payoff he can get is $\min_{t_{-i}} u_i(s_i,t_{-i})$
- Player can choose the strategy s_i that maximizes this value. In other words, disregarding the possible rationality (or irrationality) of the other players, he can guarantee for himself a payoff of:

$$\underline{v}_i = \max_{s_i} \min_{t_{-i}} u_i(s_i,t_{-i})$$

- Maxmin value of player i, sometimes also called the player’s security level. A strategy s^*_i that guarantees this value is called a **maxmin strategy**.

Maxmin security

	L	R
T	2, 1	2, -20
M	3, 0	-10, 1
B	-100, 2	3, 3

- What are the security values for each player?
- What happens if both players choose their maxmin strategies?

Security game

- Random strategy:
  - *Increase cost/uncertainty to attackers*

Adversary

Defender

	Target #1	Target #2
Target #1	4, -3	-1, 1
Target #2	-5, 5	2, -1

Maxmin security

	L	R
T	3, 1	0, 4
B	2, 3	1, 1

- The maxmin value of Player 1 is 1 and his unique maxmin strategy is B. The maxmin value of Player 2 is 1, and both L and R are his maxmin strategies. It follows that when the two players implement maxmin strategies the payoff might be (2,3) or (1,1) depending on which maxmin strategy is implemented by Player 2.

Chicken

	Swerve	Straight
Swerve	0, 0	-1, +1
Straight	+1, -1	-10, -10

Fig. 2: Chicken with numerical payoffs

Battle of the sexes

	Opera	Football
Opera	3,2	0,0
Football	0,0	2,3

Battle of the Sexes 1

	Opera	Football
Opera	3,2	1,1
Football	0,0	2,3

Battle of the Sexes 2

Prisoner's dilemma

	C	D
C	4, 4	0, 5
D	5, 0	1, 1

- Theorem: A strategy of player i that dominates all his other strategies is a maxmin strategy for that player. Such a strategy, furthermore, is a best reply of player i to any strategy vector of the other players.

- Corollary: In a game in which every player has a strategy that dominates all of his other strategies, the vector of dominant strategies is an equilibrium point and a vector of maxmin strategies.

- Theorem: In a game in which every player i has a strategy s^*_i that strictly dominates all of his other strategies, the strategy vector $(s^*_1, \ldots, s^*_n)$ is the unique equilibrium point of the game as well as the unique vector of maxmin strategies.

- Theorem: Every Nash equilibrium σ^* of a strategic-form game satisfies $u_i(s_i,s_{-i}) >= v\_i$ for every player i.

- Proof:

  For every strategy s_i in S_i, we have

  $$u_i(s_i, s^*_{-i}) >= \min_{s-i} u_i(s_i,s_{-i})$$

  Since the definition of an equilibrium implies that

  $u_i(s^*) = \max_{si} u_i(s_i,s^*_{-i})$, we deduce that

  $u_i(s^*) = \max_{si} u_i(s_i,s^*_{-i}) >= \max_{si} \min_{s-i} u_i(s_i,s_{-i}) = v\_i$

Assignment

- HW1 due 1/31
- HW2 out 1/26 due 2/7
- Chapter 3 from Maschler textbook