CAP 5993/CAP 4993
Game Theory

Instructor: Sam Ganzfried

sganzfri@cis.fiu.edu

HW1

- Due 1/31
- HW 2: out 1/31, due 2/9
- No class on Thursday 2/2

- Nash equilibrium: stability
- Maxmin strategy: security

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

	L	**C**	**R**
T	3, -3	-5, 5	-2, 2
M	1, -1	4, -4	1, -1
B	6, -6	-3, 3	-5, 5

- Definition: A two-player game is a **zero-sum game** if for each pair of strategies (s_1, s_2) one has $u_1(s_1, s_2) + u_2(s_1, s_2) = 0$.

- In other words, a two-player game is a zero-sum game if it is a closed system from the perspective of the payoffs: each player gains what the other player loses. It is clear that in such a game the two players have diametrically opposed interests.

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

- Many real-life situations analyzed using game theory are not two-player zero-sum games because, even though the interests of the players diverge in many cases, they are often not completely diametrically opposed.
- Despite this, two-player zero-sum games have a special importance that justifies studying them carefully, for several reasons:

1. Many classical games, such as chess, backgammon, checkers, and many dice games, are two-player zero-sum games. These were the first games to be studied mathematically and yield formal results, results that spawned and shaped game theory as a young field of study in the early part of the twentieth century.

2. Given their special and highly restrictive properties, these games are generally simpler and easier to analyze mathematically than many other games. As is usually the case in mathematics, this makes them convenient objects for the initial exploration of ideas and possible directions for research in game theory.
   – Also simpler computationally. Can be solved in polynomial time while for other game classes it is PPAD-complete.

3. Because of the fact that two-player zero-sum games leave no room for cooperation between the players, they are useful for isolating certain aspects of games and checking which results stem from cooperative considerations and which stem from other aspects of the game (information flows, repetitions, and so on).

4. In every situation, no matter how complicated, a natural benchmark for each player is his “security level”: what he can guarantee for himself based solely on his own efforts, without relying on the behavior of other players. In practice, calculating the security level means assuming a worst-case scenario in which all other players are acting as an adversary. This means that the player is considering an auxiliary zero-sum game, in which all the other players act as if they were one opponent whose payoff is the opposite of his own payoff. In other words, even when analyzing a game that is non-zero-sum, the analysis of auxiliary zero-sum games can prove useful.
    - HW1 – compress players Blue and Green into one player “BlueGreen”

5. Two-player zero-sum games emerge naturally in other models. One example is games involving only a single player, which are often termed *decision problems*. They involve a decision maker choosing an action from among a set of alternatives, with the resultant payoff dependent both on his choice of action and on certain, often unknown, parameters over which he has no control. To calculate what the decision maker can guarantee for himself, we model the player's environment as if it were a second player who controls the unknown parameters and whose intent is to minimize the decision maker's payoff. This in effect yields a two-player zero-sum game. This approach is used in statistics.
   – Used to apply algorithm developed for computer poker program to robust diabetes management (Chen/Bowling NIPS 2012).

Zero-sum games

- Since payoffs u_1 and u_2, satisfy $u_1 + u_2 = 0$, we can confine attention to one function, $u_1 = u$, with $u_2 = -u$. The function u is the *payoff function* of the game, and represents the payment that Player 2 makes to Player 1.

	L	**C**	**R**
T	3	-5	-2
M	1	4	1
B	6	-3	-5

Matching Pennies

	H	**T**
H	1, -1	-1, 1
T	-1, 1	1, -1

	H	**T**
H	1	-1
T	-1	1

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

- v_$_1$ = $\max_{s1} \min_{s2} u(s_1,s_2)$
- v_$_2$ = $\max_{s2} \min_{s1} (-u(s_1,s_2)) = -\min_{s2} \max_{s1} u(s_1,s_2)$

- Denote v_ = $\max_{s1} \min_{s2} u(s_1,s_2)$
- Denote v^ = $\min_{s2} \max_{s1} u(s_1,s_2)$

- The value v_ is called the *maxmin value* of the game, and v^ is called the *minmax value*. Player 1 can guarantee that he will get at least v_, and player 2 can guarantee that he will pay no more than v^.

	L	C	R
T	3, -3	-5, 5	-2, 2
M	1, -1	4, -4	1, -1
B	6, -6	-3, 3	-5, 5

- v_ = 1 and v^ = 1. Player 1 can guarantee that he will get a payoff of a least 1 (using the maxmin strategy M), while player 2 can guarantee that he will pay at most 1 (by way of minmax strategy R).

	L	**R**
T	-2	5
B	3	0

- v_ = 0 but v^ = 3. Player 1 cannot guarantee that he will get a payoff higher than 0 (which he can guarantee by using his maxmin strategy B) and player 2 cannot guarantee that he will pay less than 3 (which he can guarantee using his minmax strategy L).

Matching pennies

	H	**T**
H	1	-1
T	-1	1

- v_ = -1 and v^ = 1. Neither player can guarantee a result that is better than the loss of one dollar.

- These examples show that v_ and v^ can be unequal, but it is always the case that v_ <= v^.
  - Player 1 can guarantee that he will get at least v_, while player 2 can guarantee that he will not pay more than v^. As the game is a zero-sum game, the inequality v_ <= v^ must hold (formal proof as exercise).
- A two-player game has a value if v_ = v^. The quantity v = v_= v^ is then called the **value** of the game. Any maxmin and minmax strategies of player 1 and player 2 respectively are then called **optimal strategies**.

	L	**C**	**R**
T	3, -3	-5, 5	-2, 2
M	1, -1	4, -4	1, -1
B	6, -6	-3, 3	-5, 5

- v_ = 1 and v^ = 1. Player 1 can guarantee that he will get a payoff of a least 1 (using the maxmin strategy M), while player 2 can guarantee that he will pay at most 1 (by way of minmax strategy R).
- So the value v=1.

Game of chess

- u(White wins) = 1, u(Black wins) = -1, u(Draw) = 0
- Theorem: In chess, one and only one of the following must be true:
  i. White has a strategy guaranteeing a payoff of 1.
  ii. Black has a strategy guaranteeing a payoff of -1.
  iii. Each of the two players has a strategy guaranteeing a payoff of at least 0; that is, White can guarantee payoff 0 or 1, and Black can guarantee payoff 0 or -1.
- Case i will imply v_ = v^ = 1.
- Case ii will imply v_ = v^ = -1.
- Case iii will imply v_ = v^ = 0.

- Theorem: Every finite two-player zero-sum extensive-form game with perfect information has a value.
- Theorem: If a two-player zero-sum game has a value v, and if s^*_1 and s^*_2 are optimal strategies of the two players, then $s^* = (s^*_1, s^*_2)$ is an equilibrium with payoff (v, -v).
- Theorem: If $s^* = (s^*_1, s^*_2)$ is an equilibrium of a two-player zero-sum game, then the game has a value $v = u(s^*_1, s^*_2)$, and the strategies s^*_1 and s^*_2 are optimal strategies.

- In many situations we would like to be unpredictable
  - If a baseball pitcher throws a waist-high fastball on every pitch, the other team's batters will have an easy time hitting the ball.
  - If a tennis player always serves the ball to the same side of the court, his opponent will have an advantage in returning the serve.
  - If a candidate for political office predictably issues announcements on particular dates, his opponents can adjust their campaign messages ahead of time to pre-empt him and gain valuable points on the polls.
  - If a traffic police car is placed at the same junction at the same time every day, its effectiveness is reduced.

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

	L	**R**
T	4	1
B	2	3

	L	**R**
T	4	1
B	2	3

- Player 1's security level is 2, Player 2's is 3.
- So the game has no value.
- Suppose player 1 tosses a coin that comes heads with probability ¼ and tails with probability ¾. Plays T if heads, B if tails.
- What is player 1's utility if player 2 plays L/R?

Mixed strategies

- Let $G = (N, (S_i)\ i \text{ in } N, (u_i)\ i \text{ in } N)$ be a game in strategic form in which the strategies S_i of each player is finite. A **mixed strategy** of player i is a probability distribution over his set of strategies S_i.
- Probability distribution: function that assigns each value in [0,1] to each element of S_i, and the sum of the probabilities equals 1.
- *Pure strategy* is special case where all probabilities are 0 or 1.

Mixed extension of a strategic-form game

- Need to define utilities of mixed strategies.
- If Player 1 plays 0.2 R, 0.3 P, 0.5 S vs. Player 2 who plays P, (expected) utility is

  0.2*u(R,P) + 0.3 * u(P,P) + 0.5*u(S,P)

  = 0.2*(-1) + 0.3*(0) + 0.5*1 = 0.3.

- If Player 1 plays this strategy against Player 2 who plays 0.1 R, 0.7 P, 0.2 S, then it is:
- 0.2*0.1*u(R,R) + 0.2*0.7*u(R,P) + ….

- Note that the mixed strategies of the players are statistically independent – they are doing their own randomization independently. That is, player 1 is tossing a coin to select his play and player 2 is tossing a separate coin for his.

- Concepts of dominant strategy, security level, and equilibrium are also defined for the mixed extension of a game.

- Theorem [Nash 1950]: Every game in strategic form G, with a finite number of players and in which every player has a finite number of pure strategies, has an equilibrium in mixed strategies.

- http://www.princeton.edu/mudd/news/faq/topics/Non-Cooperative_Games_Nash.pdf

- “That’s just a fixed point theorem.”

- Theorem [von Neumann’s Minmax Theorem 1928]: Every two-player zero-sum game in which every player has a finite number of pure strategies has a value in mixed strategies.

- He listened carefully, with his head cocked slightly to one side and his fingers tapping. Nash started to describe the proof he had in mind… But before he had gotten out more than a few disjointed sentences, von Neumann interrupted, jumped ahead to the as yet unstated conclusion of Nash's argument, and said abruptly, "That's trivial, you know. That's just a fixed point theorem."
- https://mattbaker.blog/2015/05/26/john-nash-and-the-theory-of-games/
- https://en.wikipedia.org/wiki/John_Forbes_Nash_Jr.
- https://en.wikipedia.org/wiki/John_von_Neumann

Non zero-sum game

	L	R
T	1, -1	0, 2
B	0, 1	2, 0

	L	R
T	1, -1	0, 2
B	0, 1	2, 0

- No equilibrium in pure strategies.
- Is there an equilibrium in mixed strategies?

Choosing the largest number

- Two players simultaneously and independently choose a positive integer. The player who chooses the smaller number pays a dollar to the person who chooses the largest number. If the two players choose the same integer, no exchange of money occurs.

- Maxmin value? Minmax value?

- Minmax = -1, maxmin = 1.
- There exists a sufficiently large natural number k such that $\sigma_1(\{1,2,\ldots,k\}) > 1-\varepsilon$
  - Probability that Player 1 chooses a number that is less than or equal to k is greater than 1- ε.
- But then if player 2 chooses the pure strategy k+1 we will have
  - $U(\sigma_1,k+1) < (1-\varepsilon)(-1) + \varepsilon(1) = -1 + 2\varepsilon$
- Since this is true for any ε in (0,1), the minmax value is -1. Since the minmax value does not equal the maxmin value, the game has no value in mixed strategies.

Next time

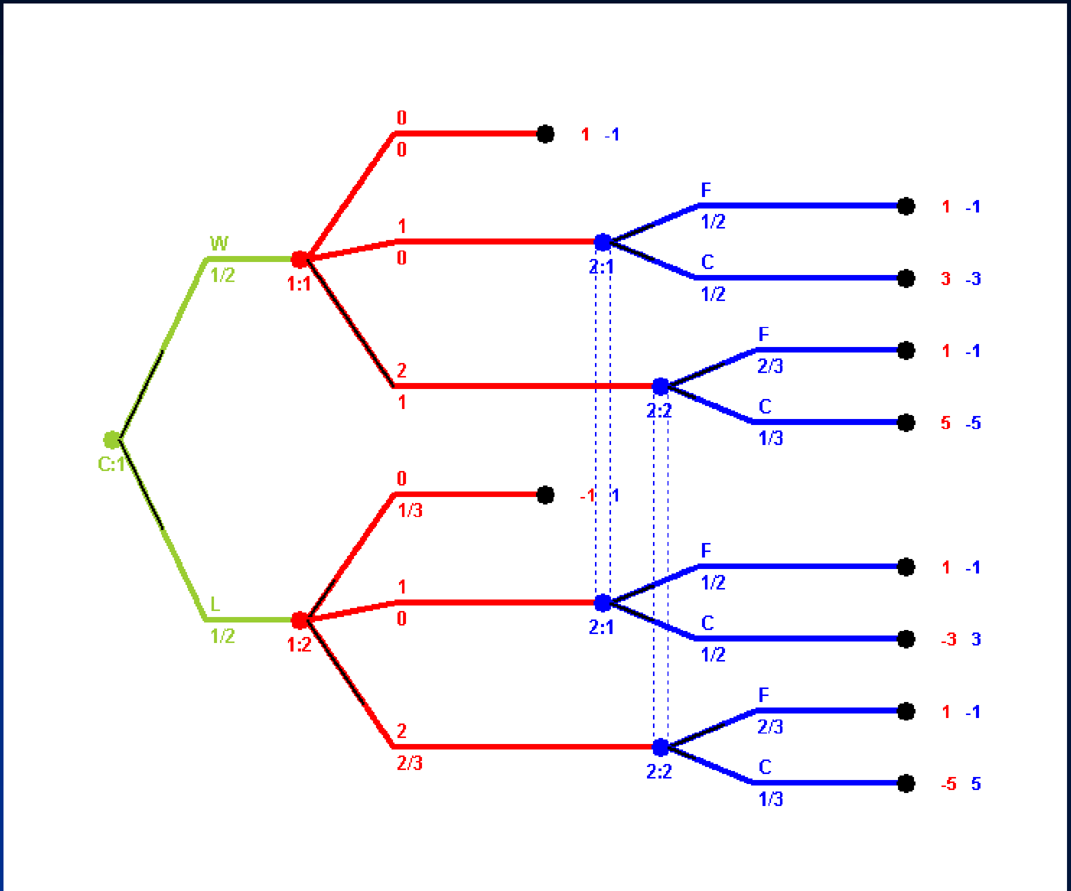

Algorithms!!!

Assignment

- HW1 due 1/31
- HW 2: out 1/31, due 2/9
- Reading for next class: Chapter 4 from Shoham textbook http://www.masfoundations.org/mas.pdf