CAP 5993/CAP 4993
Game Theory

Instructor: Sam Ganzfried

sganzfri@cis.fiu.edu

Announcements

- HW 1 due today
- HW 2 out this week (2/2), due 2/14

- 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.

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

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

- 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 $v^\wedge = 3$. 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 $v^{\wedge} = 1$. 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**.

Mixed strategies

- Let G = (N, (S_i) i in N, (u_i) i 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 picking a random number to select his play and player 2 is picking a separate random number 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.

- “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.

- 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.

	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?

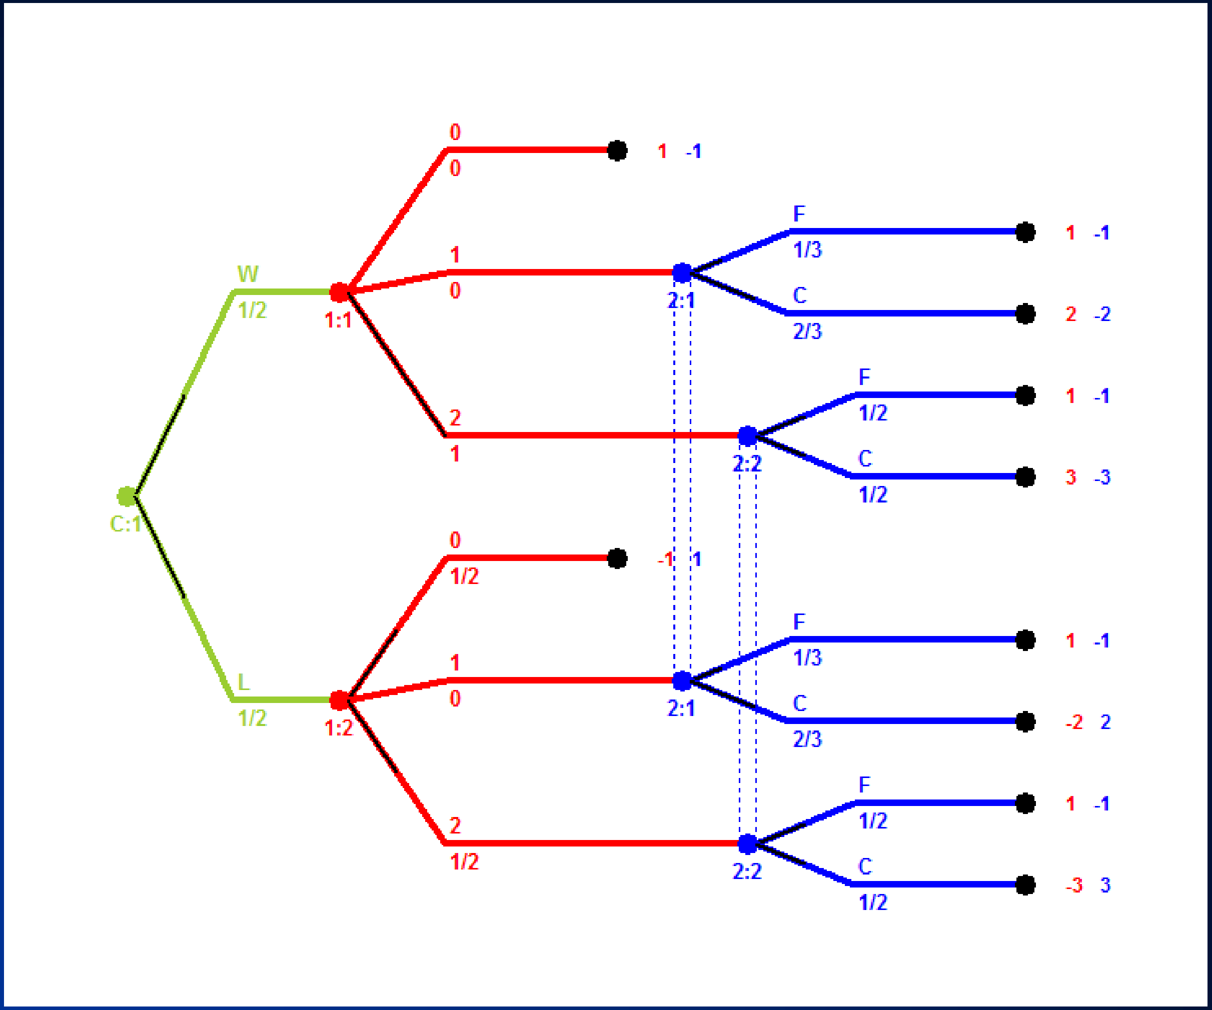

C:1
W
1/2
L
1/2
1:1
1:2
0
0
1 -1
1
0
2:1
F
1/3
1 -1
C
2/3
2 -2
2
1
2:2
F
1/2
1 -1
C
1/2
3 -3
0
1/2
-1 1
1
0
2:1
F
1/3
1 -1
C
2/3
-2 2
2
1/2
2:2
F
1/2
1 -1
C
1/2
-3 3

WL/12	CC	CF	FC	FF
00	0	0	0	0
01	-0.5	-0.5	1	1
02	-1	1	-1	1
10				
11				
12				
20				
21				
22				

Extensive-form game

- A game in **extensive form** is given by a *game tree*, which consists of a directed graph in which the set of vertices represents positions in the game, and a distinguished vertex, called the *root*, represents the starting position of the game. A vertex with no outgoing edges represents a terminal position in which play ends. To each terminal vertex corresponds an outcome that is realized when the play terminates at that vertex. Any nonterminal vertex represents either a chance move (e.g., a toss of a die or a shuffle of a deck of cards) or a move of one of the players. To any chance-move vertex corresponds a probability distribution over edges emanating from that vertex, which correspond to the possible outcomes of the chance move.

Perfect vs. imperfect information

- To describe games with imperfect information, in which players do not necessarily know the full board position (like poker), we introduce the notion of *information sets*. An information set of a player is a set of decision vertices of the player that are indistinguishable by him given his information at that stage of the game. A game of *perfect information* is a game in which all information sets consist of a single vertex. In such a game whenever a player is called to take an action, he knows the exact history of actions and chance moves that led to that position.

- A *strategy* of a player is a function that assigns to each of his information sets an action available to him at that information set. A path from the root to a terminal vertex is called a *play* of the game. When the game has no chance moves, any vector of strategies (one for each player) determines the play of the game, and hence the outcome. In a game with chance moves, any vector of strategies determines a probability distribution over the possible outcomes of the game.

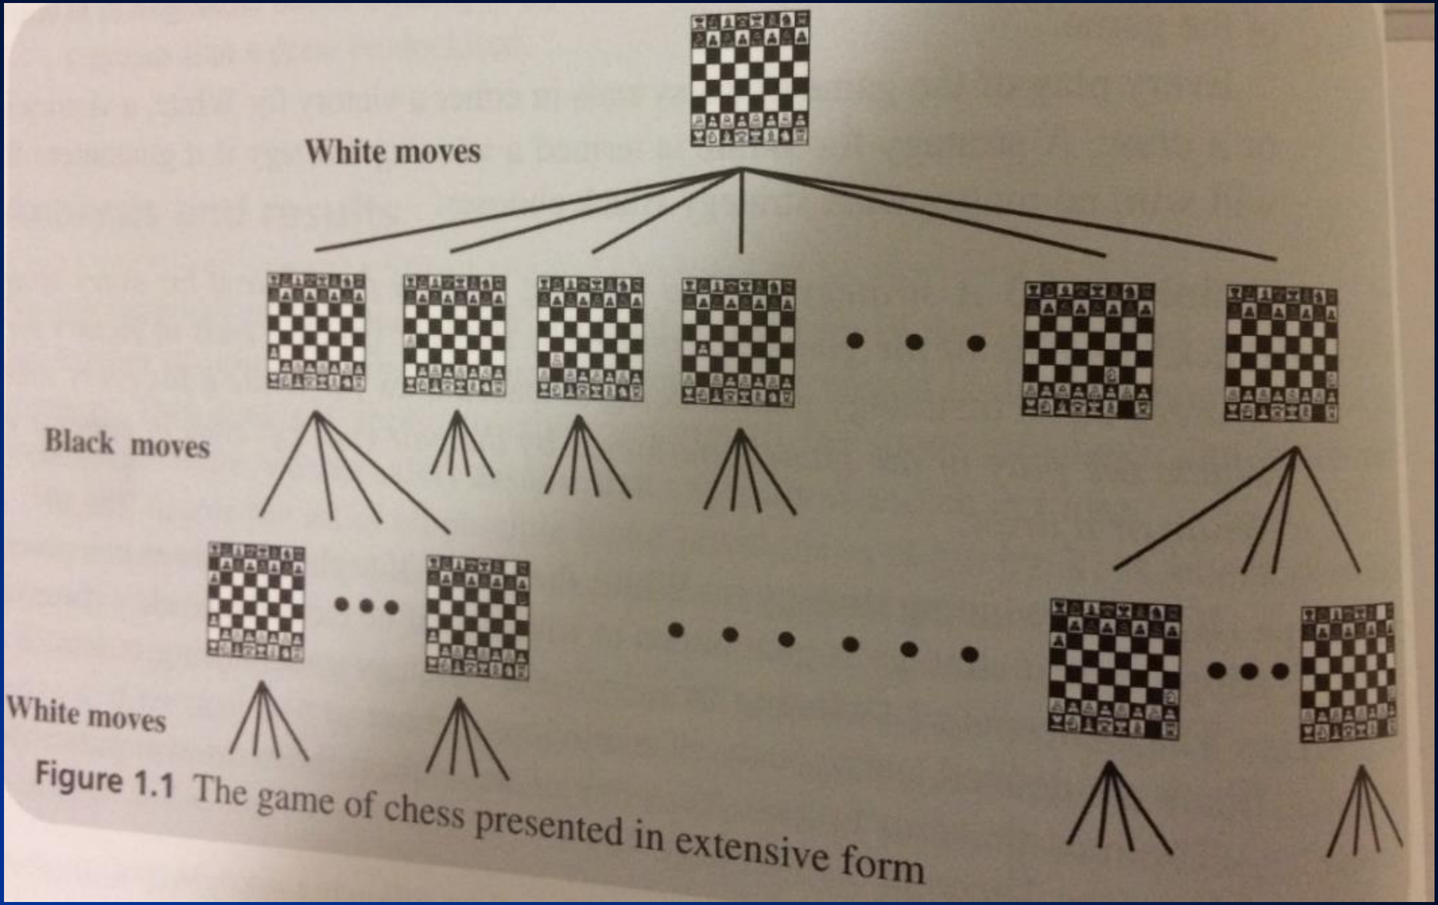

Figure 1.1 The game of chess presented in extensive form

- Every description of a game must include:
  - Set of players
  - The possible actions available to each player
  - Rules determining the order in which players make their moves.
  - A rule determining when the game ends.
  - A rule determining the outcome of every possible game ending.

- A (finite) directed graph is a pair G = (V,E) where:
  - V is a finite set, whose elements are called vertices.
  - E subset of V x V is a finite set of pairs of vertices, whose elements are called edges. Each directed edge is composed of two vertices: the two ends of the edge (it is possible for both ends of a single edge to be the same vertex).

Game trees

- A tree is a triple $G = (V, E, x^0)$ where (V,E) is a directed graph, x^0 in V is a vertex called the *root* of the tree, and for every vertex x in V there is a unique path in the graph from x^0 to x.
- Various games can be represented as trees. When a tree represents a game, the root of the tree corresponds to the *initial position* of the game, and every *game position* is represented by a vertex of the tree. The children of each vertex v are the vertices corresponding to the game positions that can be arrived at from v via one action. In other words, the number of children of a vertex is equal to the number of possible actions in the game position corresponding to that vertex.
  - For every vertex that is not a leaf, we need to specify the player who is to take an action at that vertex
  - At each leaf, we need to describe the outcome of the game.

- A game in extensive form (or extensive-form game) is an ordered vector $\Gamma = (N, V, E, x^0, (V_i)$ i in N, O, u)
  - N is finite set of players
  - (V, E, x^0) is a tree called the *game tree*
  - (V_i) i in N is a partition of the set of vertices that are not leaves.
  - O is the set of possible game outcomes.
  - u is a utility function associating every leaf of the tree with a game outcome in the set O.
- Let B be a nonempty set. A *partition* of B is a collection $B_1, B_2, \ldots, B_K$ of pairwise disjoint and nonempty subsets of B whose union is B.

- By “possible outcome” we mean a detailed description of what happens as a result of the actions undertaken by the players. Some examples of outcomes include:
  - Player 1 is declared the winner of the game, and Player 2 the loser.
  - Player 1 receives $2, player 2 receives $3, player 3 receives $5.
  - Player 1 gets to go out to the cinema with player 2, while player 3 is left at home.

- The requirement that (V_i) i in N be a partition of the set of vertices that are not leaves stems from the fact that at each game situation there is one and only one player who is called upon to take an action. For each vertex x that is not a leaf, there is a single player i in N for whom x in V. The player is called the *decision maker at vertex x*, and denoted by J(x).
- Denote by C(x) the set of all children of non-leaf vertex x.
- Every edge that leads from x to one of its children is called a possible *action* at x. We will associate every action with the child to which it is connected, and denote by A(x) the set of all actions that are possible at the vertex x.

- An extensive-form game proceeds in the following manner:
  - Player $J(x^0)$ initiates the game by choosing a possible action in $A(x^0)$. Equivalently, he chooses an element x^1 in the set $C(x^0)$.
  - If x^1 is not a leaf, player $J(x^1)$ chooses a possible action in $A(x^1)$ (equivalently, an element x^2 in $C(x^1)$).
  - The game continues in this manner, until a leaf vertex x is reached, and then the game ends with outcome u(x).

- By definition, the collection of the vertices of the graph is a finite set, so that the game necessarily ends at a leaf, yielding a sequence of vertices $(x^0, x^1, \ldots x^k)$, where x^0 is the root of the tree, x^k is a leaf, and x^{i+1} in $C(x^i)$ for $l = 0,1,\ldots,k-1$. This sequence is called a *play*. Every play ends at a particular leaf x^k with outcome $u(x^k)$. Similarly, every leaf x^k determines a unique play, which corresponds to the unique path connecting the root x^0 with x^k.

- It follows from the above description that every player who is to take an action knows the current state of the game, meaning that he knows all the actions in the game that led to the current point in the play. This implicit assumption is called *perfect information.*

- A strategy for a player i is a function s_i mapping each vertex x in V_i to an element in A(x) (equivalently, to an element in C(x)).
- According to this definition, a strategy includes instructions on how to behave at each vertex in the game tree, including vertices that previous actions by the player preclude from being reached. For example, in the game of chess, even if White's strategy calls for opening by moving a pawn from c2 to c3, the strategy must include instructions on how White should play his second move if in his first move he instead moved a pawn from c2 to c4, and Black then took his action.

- A strategy vector is a list of strategies $s = (s_i)$ i in N, one for each player. Player i's set of strategies is denoted by S_i, and the set of all strategy vectors is denoted $S = S_1 x\, S_2\, x \ldots x\, S_n$. Every strategy vector determines a unique play from the root to a leaf.
- Let $\Gamma = (N, V, E, x^0, (V_i)$ i in N, O, u) be an extensive-form game (with perfect information), and let x in V be a vertex in the game tree. The *subgame starting at x*, denoted by $\Gamma(x)$, is the extensive-form game $\Gamma(x) = (N, V(x), E(x), x, (V_i(x))$ i in N, O, u).
  - V(x) includes x and all vertices that are descendants of x.

- Theorem (von Neumann [1928]) In every two-player game (with perfect information) in which the set of outcomes is O = {Player 1 wins, Player 2 wins, Draw}, one and only one of the following three alternatives holds:
  1. Player 1 has a winning strategy.
  2. Player 2 has a winning strategy.
  3. Each of the two players has a strategy guaranteeing at least a draw.

Chance moves

- In the games we have seen so far, the transition from one state to another is always accomplished by actions undertaken by the players. Such a model is appropriate for games such as chess and checkers, but not for card games or dice games (such as poker or backgammon) in which the transition from one state to another may depend on a chance process: in card games, the shuffle of the deck, and in backgammon, the toss of the dice. It is possible to come up with situations in which transitions from state to state depend on other chance factors, such as the weather, earthquakes, or the stock market. These sorts of state transitions are called *chance moves*. To accommodate this feature, our model is expanded by labeling some of the vertices in the game tree (V, E, x_0) as chance moves. The edges emanating from vertices corresponding to chance moves represent the possible outcomes of a lottery, and next to each such edge is listed the probability that the outcome it represents will be the result of the lottery.

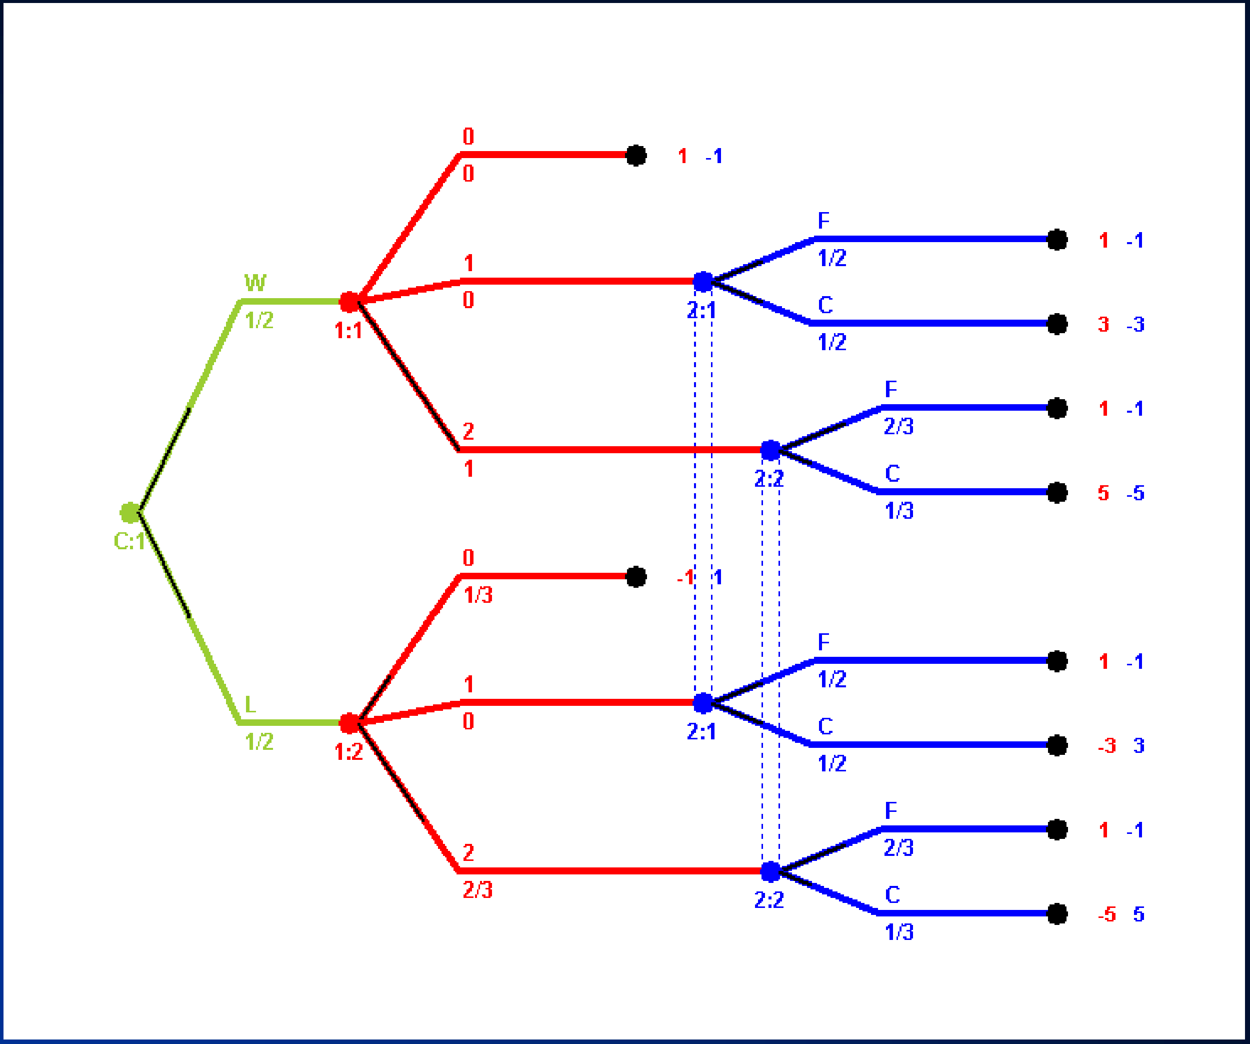
C:1
W
1/2
L
1/2
1:1
1:2
0
0
1 -1
1
0
2
1
2:1
2:2
F
1/2
1 -1
C
1/2
3 -3
F
2/3
1 -1
C
1/3
5 -5
0
1/3
-1 1
1
0
2
2/3
2:1
2:2
F
1/2
1 -1
C
1/2
-3 3
F
2/3
1 -1
C
1/3
-5 5

- Formally, the addition of chance moves to the model proceeds as follows. We add a new player, who is called “Nature,” and denoted by 0. The set of players is thus expanded to N U {0}. For every vertex x at which a chance move is implemented, we denote by p_x the probability vector over the possible outcomes of a lottery that is implemented at vertex x. This leads to the following definition of a game in extensive form.

- A game in extensive form (*with perfect information and chance moves*) is a vector $\Gamma = (N, V, E, x^0, (V_i)$ i in (N U {0}), (p_x) x in V_0, O, u)
  - We have added the set V_0 to the partition, where V_0, is the set of vertices at which a chance move is implemented.
  - For each vertex x in V_0, a vector p_x, which is a probability distribution over the edges emanating from x, has been added to the model.

- One of the distinguishing properties of the games we have seen so far is that at every stage of the game each of the players has perfect knowledge of all of the developments in the game prior to that stage: he knows exactly which actions were taken by all the other players, and if there were chance moves, he knows what the results of the chance moves were. In other words, every player, when it is his turn to take an action, knows precisely at which vertex in the game tree the game is currently at. A game satisfying this condition is called a *game with perfect information*.

- The assumption of perfect information is clearly a very restrictive assumption, limiting the potential scope of analysis. Players often do not know all the actions taken by the other players and/or the results of chance moves (for example, in many card games the hand of cards each player holds is not known to the other players).
- In general, a player's information set consists of a set of vertices that satisfy the property that when play reaches one of these vertices, the player knows that play has reached one of these vertices, but he does not know which vertex has been reached.

Imperfect information

- Let $\Gamma = (N, V, E, x^0, (V_i)$ i in (N U {0}), (p_x) x in V_0, O, u) be a game in extensive form. An *information set* of player i is a pair $(U_i, A(U_i))$ such that
  - $U_i = \{x^1_i, x^2_i, \ldots, x^j_i\}$ is a subset of V_i that satisfies the property that at each vertex in U_i player i has the same number of actions $l_i = l_i(U_i)$, i.e., $|A(x^j_i)| = l_i$ for all j = 1,2,…,m.
  - $A(U_i)$ is a partition of the ml_i edges Union of $A(x^j_i)$ to l_i disjoint sets, each of which contains one element from the sets $A(x^j_i)$. We denote the elements of the partition by a^1_i, $a^2_i, \ldots, a^j_i$. The partition $A(U_i)$ is called the *action set* of player I in the information set U_i.

Rock-paper-scissors

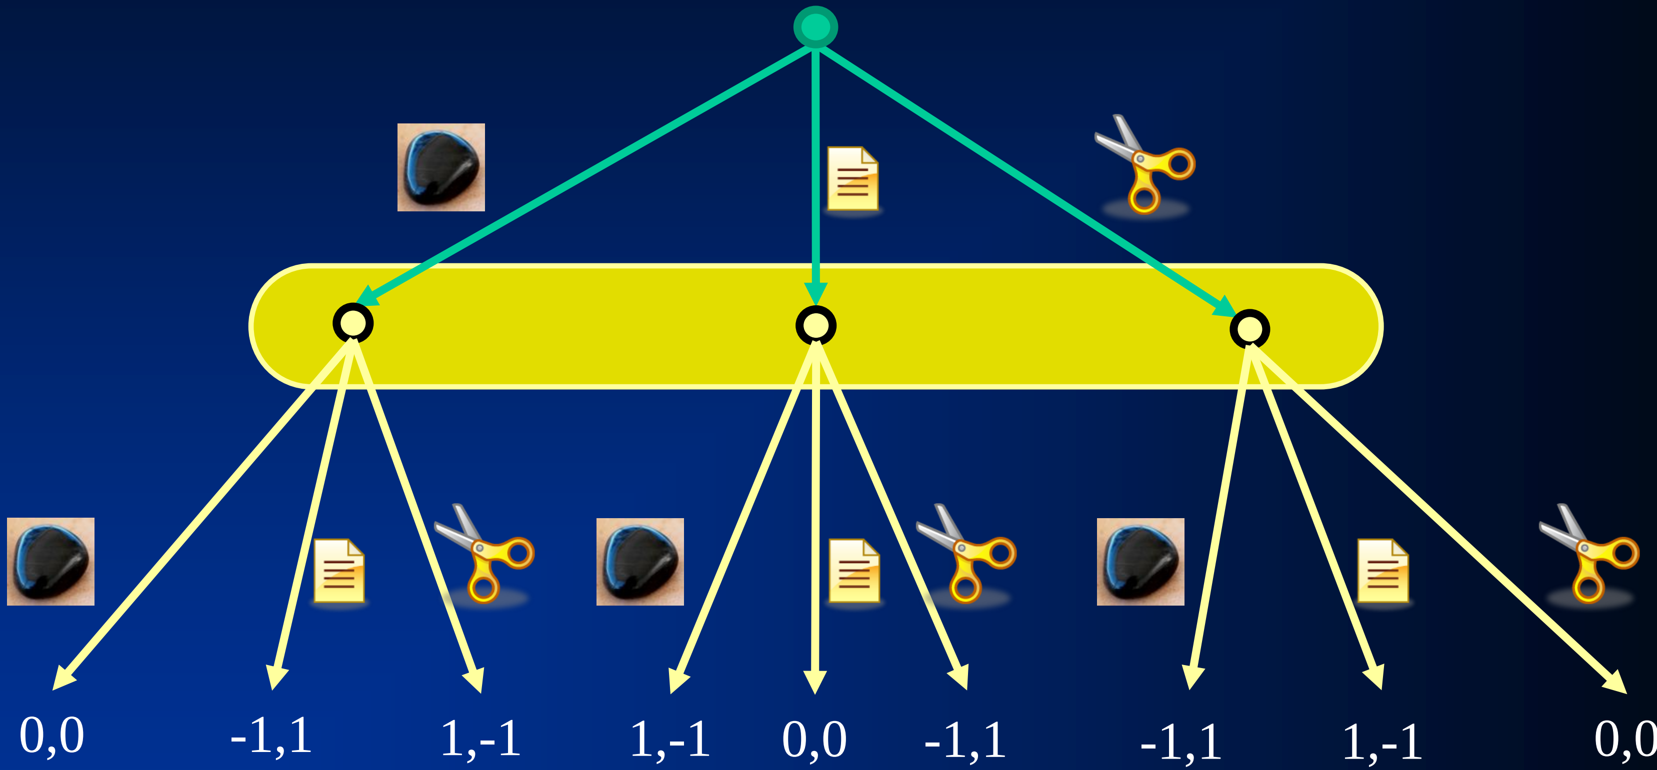

- When the play of the game arrives at vertex x in information set U_i, all that player i knows is that the play has arrived at one of the vertices in this information set. The player therefore cannot choose a particular edge emanating from x. Each element of the partition a^l_i contains m edges, one edge for each vertex in the information set. The partition $a^1_i, a^2_i, \ldots, a^{li}_i$ are the “actions” from which the player can choose; if player i chooses one of the elements from the partition a^l_i, the play continues along the unique edge in the intersection of a^l_i and A(x). For this reason, when we depict games with information sets, we denote edges located in the same partition by the same letter.

- A game in extensive form (with chance moves and with imperfect information) is a vector Γ = (N, V, E, x^0, (V_i) i in (N U {0}), (p_x) x in V_0, (U_i^j) i in N j = 1,…,k_i, O, u), where:
  - For each player i in N, (U_i^j), j = 1,…,k_i is a partition of V_i.
  - For each player i in N and every j in {1,2,… k_i} the pair $(U_i^j, A(U_i^j))$ is an information set of player i.

- In a game with imperfect information, each player i, when choosing an action, does not know at which vertex x the play is located. He only knows the information set U_i^j that contains x, and he chooses an element a in $A(U_{J(x)}^j)$.
- We can now describe many more games in extensive form: various card games such as poker and bridge, games of strategy such as Stratego, and many real-life situations, such as bargaining between two parties.

- An extensive-form game is called a game with perfect information for player i if each information set of player i contains only one vertex. An extensive-form game is called a game with perfect information if it is a game with perfect information for all of the players.

Strategies in imperfect-information games

- A strategy of player i is a function from each of his information sets to the set of actions available at that information set.
- Just as in games with chance moves and perfect information, a strategy vector determines a distribution over the outcomes of a game.

- Every extensive-form game can be converted to an equivalent strategic-form game, and therefore all the prior concepts and theoretical results (e.g., domination, security level, mixed strategies, Nash equilibrium, Minmax Theorem) will apply. However, this conversion produces a strategic-form game that has size that is exponential in the size of the original game tree, and is infeasible for large games. Therefore, we would like do develop algorithms that operate directly on extensive-form games and avoid the conversion to strategic form games.

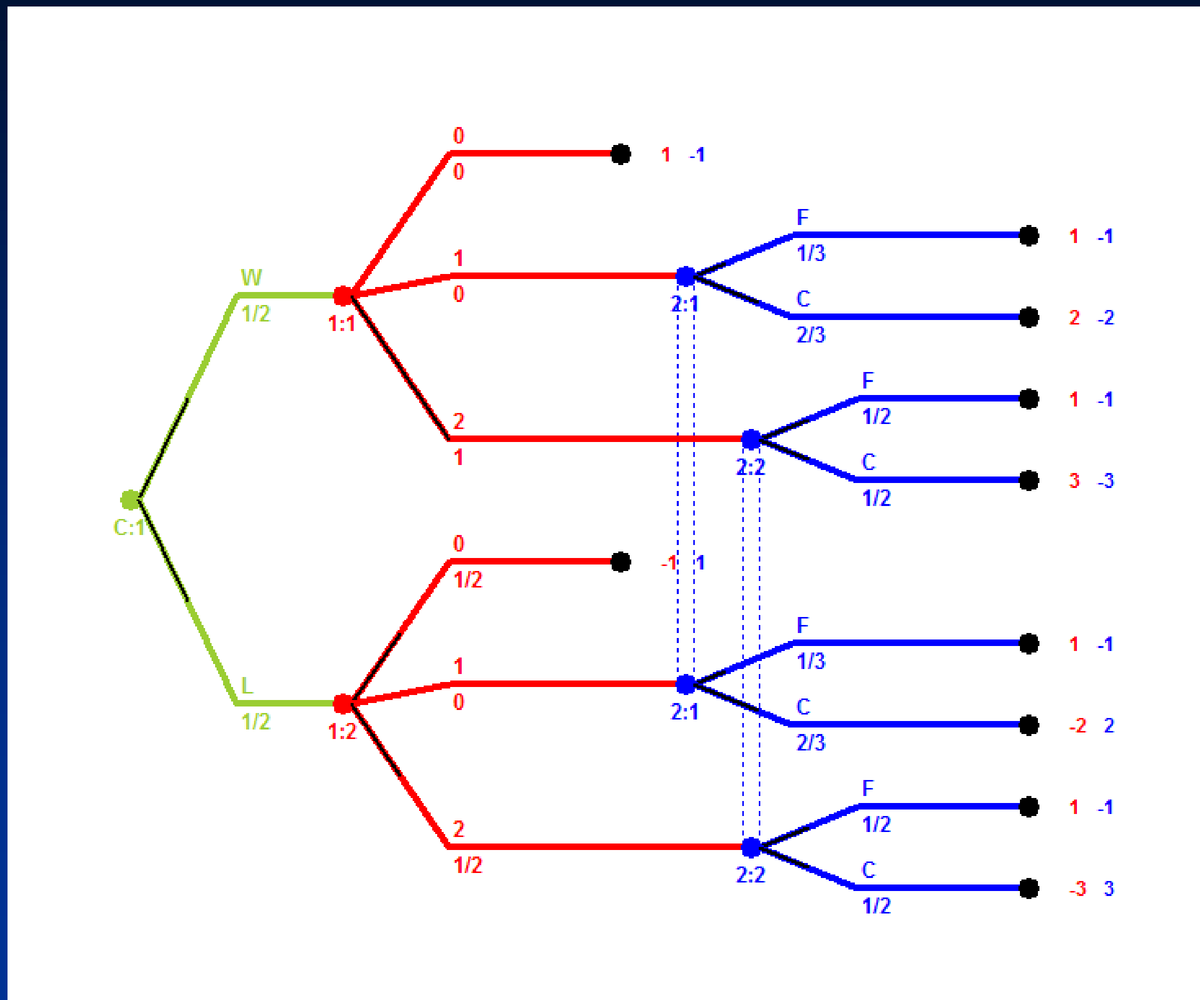
C:1
W
1/2
L
1/2
1:1
1:2
0
0
1 -1
1
0
2:1
F
1/3
1 -1
C
2/3
2 -2
2
1
2:2
F
1/2
1 -1
C
1/2
3 -3
0
1/2
-1 1
1
0
2:1
F
1/3
1 -1
C
2/3
-2 2
2
1/2
2:2
F
1/2
1 -1
C
1/2
-3 3

WL/12	CC	CF	FC	FF
00	0	0	0	0
01	-0.5	-0.5	1	1
02	-1	1	-1	1
10				
11				
12				
20				
21				
22				

- Theorem (Kuhn) Every finite game with perfect information has at least one pure strategy Nash equilibrium.

- Corollary of Nash’s Theorem: Every extensive-form game (of perfect or imperfect information) has an equilibrium in mixed strategies.

Next time

- Algorithms for computing solution concepts in strategic-form and extensive-form games.

Assignment

- HW 2 out this week (2/2), due 2/14
- No class Thursday
- Reading for next class: Chapter 5 from Shoham textbook http://www.masfoundations.org/mas.pdf