CMSC 471
Fall 2012

Class #9

Thurs 9/27/12
Game Theory

Kevin Winner, winnerk1@umbc.edu

Today's class

- Game playing
  - Nim
  - Stochastic games
- Game Theory

Game Playing #2

Still chapter 5

Some material adopted from notes
by Charles R. Dyer, University of
Wisconsin-Madison

Example: Nim

- In Nim, there are a certain number of objects (coins, sticks, etc.) on the table -- we'll play 7-coin Nim
- Each player in turn has to pick up either one or two objects
- Whoever picks up the last object loses

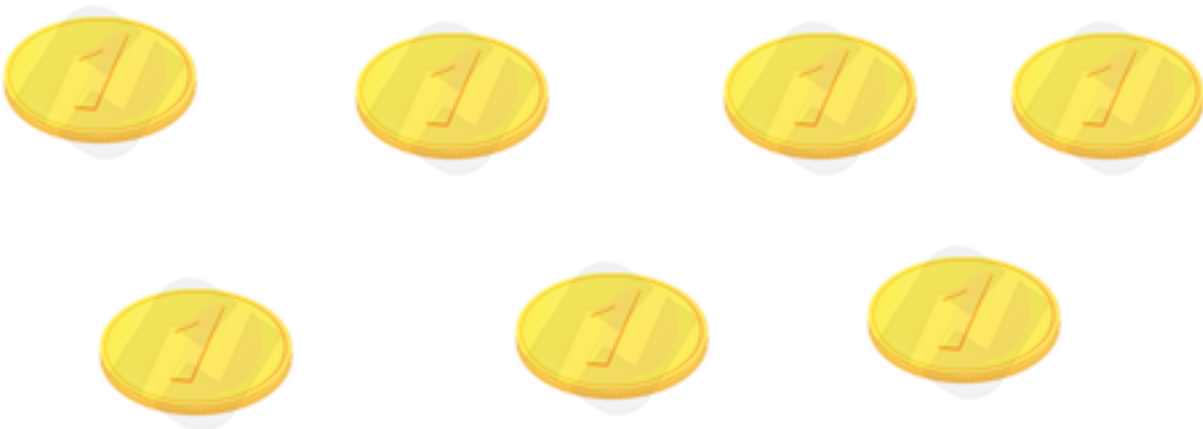

Games of chance

- Backgammon is a two-player game with **uncertainty**.
- Players roll dice to determine what moves to make.
- White has just rolled *5 and 6* and has four legal moves:
  - 5-10, 5-11
  - 5-11, 19-24
  - 5-10, 10-16
  - 5-11, 11-16
- Such games are good for exploring decision making in adversarial problems involving skill and luck.

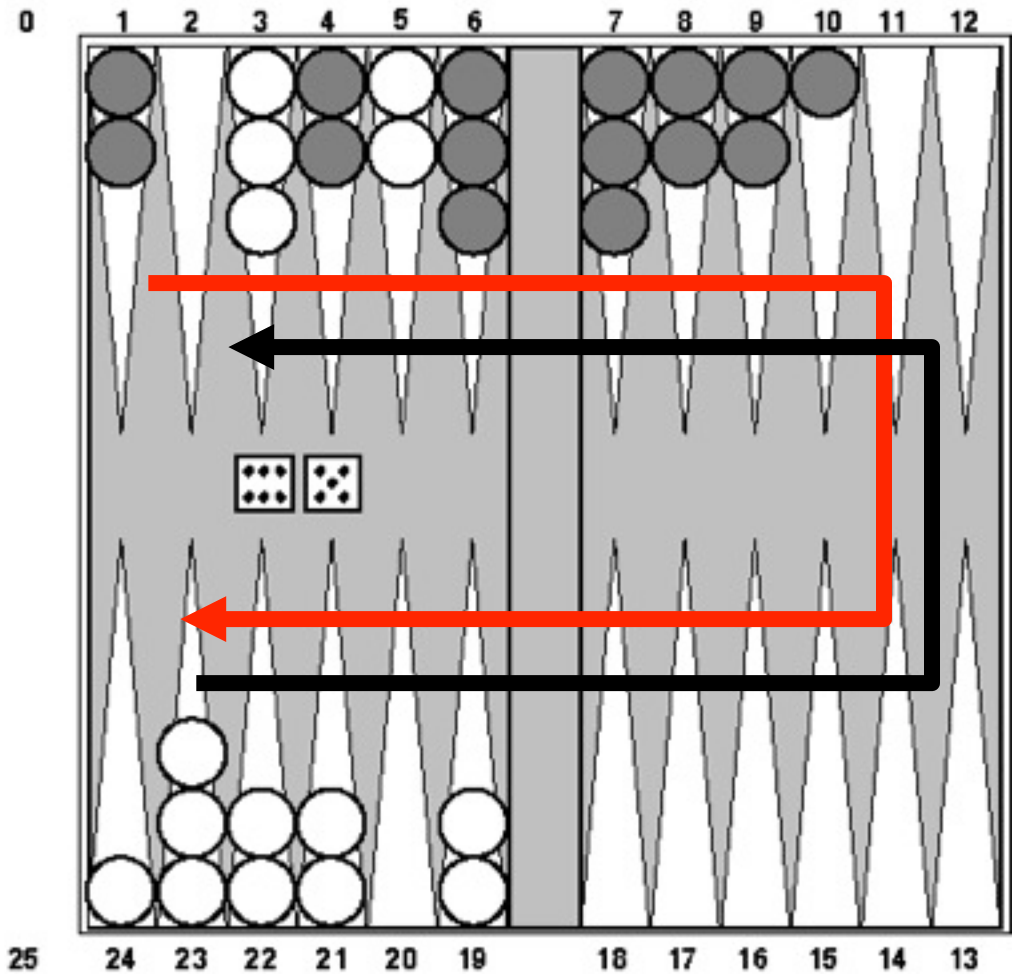

Game trees with chance nodes

- **Chance nodes** (shown as circles) represent random events
- For a random event with N outcomes, each chance node has N distinct children; a probability is associated with each
- (For 2 dice, there are 21 distinct outcomes)
- Use minimax to compute values for MAX and MIN nodes
- Use **expected values** for chance nodes
- For chance nodes over a max node, as in C:

$\text{expectimax}(C) = \sum_i (P(d_i) * \text{maxvalue}(i))$

- For chance nodes over a min node:

$\text{expectimin}(C) = \sum_i (P(d_i) * \text{minvalue}(i))$

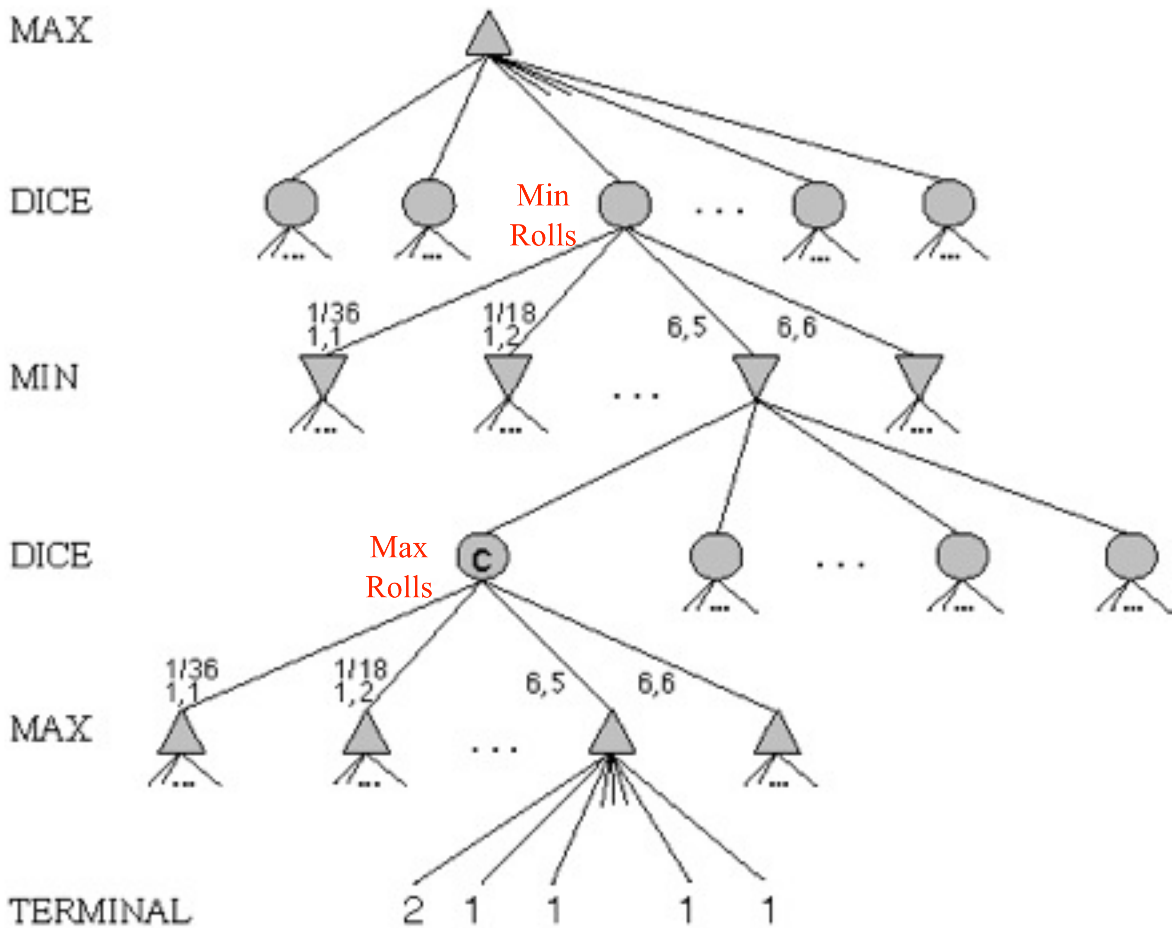

Example: Oopsy-Nim

- Starts out like Nim
- Each player in turn has to pick up either one or two objects
- Sometimes (with probability 0.25), when you try to pick up two objects, you drop them both
- Picking up a single object always works
- Whoever picks up the last object loses

- Question: Why can't we draw the entire game tree?

Game Theory

Not actually in your textbook

Game Theory

- Reasoning about multi-agent interactions
  - Games are episodic (although sometimes we'll track an interaction history)
  - Agents select moves simultaneously and independently
- Describe the game as a table of each permutation of actions
- Based on our predictions of what the other agent(s) will do, what is the optimal strategy for us to play?

Rock-Paper-Scissors

		Player 2		
		R	P	S
Player 1	R	0,0	-1,1	1,-1
	P	1,-1	0,0	-1,1
	S	-1,1	1,-1	0,0

Prisoner’s Dilemma

- From Wikipedia:
  - “Two men are arrested, but the police do not have enough information for a conviction. The police separate the two men, and offer both the same deal: if one testifies against his partner (defects/betrays), and the other remains silent (cooperates with/assists his partner), the betrayer goes free and the one that remains silent gets a one year sentence. If both remain silent, both are sentenced to only one month in jail on a minor charge. If each 'rats out' the other, each receives a three-month sentence. Each prisoner must choose either to betray or remain silent; the decision of each is kept secret from his partner. What should they do?”

Prisoner's Dilemma

		Player 2	
		C	D
Player 1	C	-1,-1	-5,0
	D	0,-5	-12,-12

Prisoner's Dilemma

		Player 2	
		C	D
Player 1	C	3,3	0,5
	D	5,0	1,1

What is the best strategy?

		Player 2	
		C	D
Player 1	C	3,3	0,5
	D	5,0	1,1

Pareto Optimality

- S is a Pareto-optimal solution iff
  - $\forall S' (\exists x\ U_x(S') > U_x(S) \rightarrow \exists y\ U_y(S') < U_y(S))$
  - i.e., if X is better off in S', then some Y must be worse off
- Social welfare, or global utility, is the sum of all agents' utility
  - If S maximizes social welfare, it is also Pareto-optimal (but not vice versa)

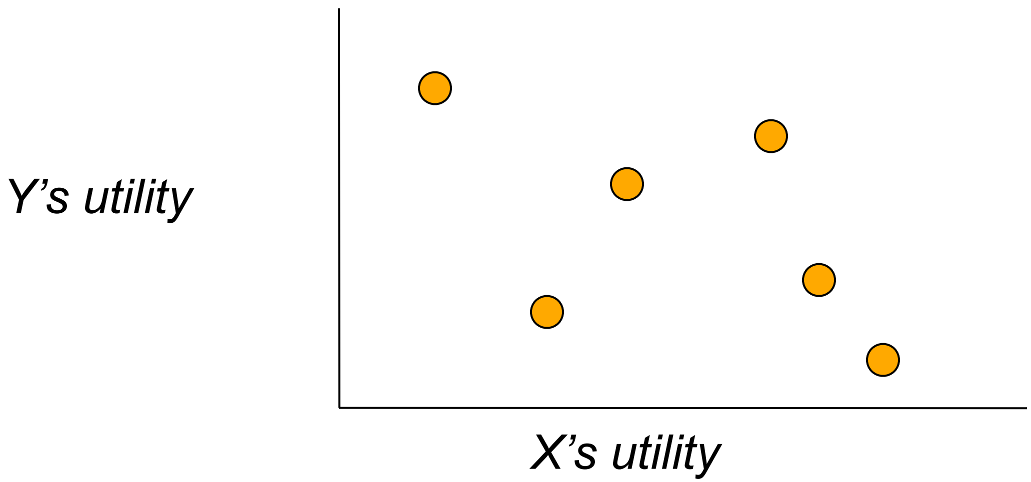

Which solutions are Pareto-optimal?

Which solutions maximize global utility (social welfare)?

Stability

- If an agent can always maximize its utility with a particular strategy (regardless of other agents' behavior) then that strategy is **dominant**
- A set of agent strategies is in **Nash equilibrium** if each agent's strategy Si is locally optimal, given the other agents' strategies
  - No agent has an incentive to change strategies
  - Hence this set of strategies is **locally stable**

Game Analysis

- Which solution(s) maximizes social welfare?
- Which solution(s) are Pareto-optimal?
- Which solution(s) are Nash equilibriums?
- What is the dominant strategy(ies)?

3,2	1,1	4,4	0,5	1,1
1,7	1,2	0,0	7,8	1,5
2,1	2,1	0,0	2,9	2,1
6,6	0,0	3,3	5,6	2,3
4,0	3,10	5,3	2,10	6,1

Iterated Prisoner's Dilemma (IPD)

- Play Prisoner's Dilemma many times with the same opponent

Iterated Prisoner's Dilemma (IPD)

- Play Prisoner's Dilemma many times with the same opponent

IPD Strategies

- Always Cooperate
- Always Defect
- Optimistic defect
- Tit-for-tat
- Tit-for-two-tats
- Tit-for-tat with forgiveness
- Master-slave

- And lots of more complicated strategies

Other games we can play

- Ultimatum Game
  - Two players are splitting $100. One player offers a split to the other. If the split is accepted, both players receive payout equal to the offer. Otherwise, neither player receives anything.
- Guess 2/3 of the average
  - Everyone in the class submits a number from 0-100. The objective is to guess closest to 2/3 of the average of all the guesses.
- Stag hunt

	Stag	Hare
Stag	2,2	0,1
Hare	1,0	1,1