CSC200: Lecture 10

- Today
  - Continuing game theory: mixed strategy equilibrium (Ch.6.7-6.8), optimality (6.9), start on extensive form games (6.10, Sec. C)
- Next few lectures
  - game theory: Ch.8, Ch.9
- Announcements
  - Quiz 2 on Friday, Oct.23: covers Chapter 4 (closures in a social-affiliation network)
    - see practice question on web page

Mixed Equilibria

		Player 2	
		H	T
Player 1	H	$-1, +1$	$+1, -1$
	T	$+1, -1$	$-1, +1$

Matching Pennies

- Some games have no pure strategy (deterministic) NE
  - MP: 1 heads → 2 heads → 1 tails → 2 tails →1 heads …
- Solution: randomized or *mixed* strategies
  - players randomly choose each action with some probability
  - Nash's result: (mixed) equilibria exist for all normal form games

Mixed Strategies

- Let S_1 be the set of *pure* strategies of player 1
- A *mixed strategy* σ_1 for player 1 is a probability distribution over S_1
  - assigns a probability to each s in S_1 (probabilities must sum to one, must be non-negative)
  - player 1 chooses each *s* with the corresponding probability
  - we say the player is *mixing* between his pure strategies
- Examples (note: pure strategies are a special case):
  - [0.5 heads, 0.5 tails]; [1.0 heads, 0 tails]; [0.01 heads, 0.99 tails]
  - Firm1: [0.5 A, 0.25 B, 0.25 C]; [1.0 A, 0 B, 0 C]
  - Bob: [1.0 uptown, 0 downtown]; [0.2 uptown, 0.8 downtown]
- In two-action games, often just state probability *p* of first action
  - probability of second action is $1-p$
  - so set of strategies σ is just set of real numbers between 0 and 1

Outcomes of Mixed Strategies

- Since players randomize, *outcome and payoffs random are too*
- Consider matching pennies
  - P1 plays heads with probability *p* (tails, *1-p*) *(e.g., p=0.3)*
  - P2 plays heads with probability *q* (tails *1-q*) *(e.g., q=0.8)*
  - Assume they randomize independently (can't coordinate!)
- Probability of various outcomes and payoffs

Outcome	**Prob**	**e.g.**	**Payoff(1)**	**Payoff(2)**
(H,H)	pq	0.24	-1	+1
(H,T)	p(1-q)	0.06	+1	-1
(T,H)	(1-p)q	0.56	+1	-1
(T,T)	(1-p)(1-q)	0.14	-1	+1

$$\begin{array}{cc|c|c|} & & \multicolumn{2}{c}{\text{Player 2}} \\ & & H & T \\ \cline{3-4} \text{Player 1} & H & -1,+1 & +1,-1 \\ \cline{3-4} & T & +1,-1 & -1,+1 \\ \cline{3-4} \end{array}$$

Payoff of Mixed Profile

- Since payoff to each player is random, the value to the player of a profile is her *expected payoff*
  - simply take the (probabilistic weighted) average of her payoffs

Outcome	Prob	e.g.	Payoff(1)	Payoff(2)
(H,H)	pq	0.24	-1	+1
(H,T)	p(1-q)	0.06	+1	-1
(T,H)	(1-p)q	0.56	+1	-1
(T,T)	(1-p)(1-q)	0.14	-1	+1

- Expected payoff (or *expected value*, EV) for player 1:
  - 0.24(-1) + 0.06(1) + 0.56(1) + 0.14(-1) = 0.24
- Expected payoff for player 2:
  - 0.24(1) + 0.06(-1) + 0.56(-1) + 0.14(1) = –0.24
- *Hardly a coincidence that one is minus the other*

Best Responses to Mixed Strategies

- Which strategy should P1 choose if P2 plays q?
  - $EV(H) = q(-1) + (1-q)(1) = 1-2q$
  - $EV(T) = q(1) + (1-q)(-1) = 2q - 1$
  - $EV(p) = p \cdot EV(H) + (1-p) \cdot EV(T) = p(1-2q) + (1-p)(2q - 1)$
- So P1's *best response* to q is:
  - H is better than T: if $1-2q > 2q-1$ (i.e., $q<0.5$)
  - T is better than H: if $1-2q < 2q-1$ (i.e., $q>0.5$)
  - H and T equally good: if $1-2q = 2q-1$ (i.e., $q=0.5$)
  - Mixed strategy p?
    - p must be 1 if $q<0.5$
    - p must be 0 if $q>0.5$
    - p can be anything between 0 and 1 if $q=0.5$

Nash Equilibrium in Mixed Strategies

- Define NE exactly as before, but allowing *mixing*
  - a pair of *mixed* strategies. s for player 1 and t for player 2, such that s is a best response to t, and t is a best response to s
  - generalize to more than two players as before
- What is NE of matching pennies?
  - As before, no pure strategy can be part of NE: *unstable*
  - So P2 must play a mixed strategy q (q can't be 0 or 1)
  - But we've seen P1's best responses are pure unless $q = 0.5$
  - So only NE must have $q = 0.5$
  - We've seen P1's best response to $q = 0.5$ can be *any* p
  - Unless $p = 0.5$, P2's best response would also be pure
  - So we must have $p = 0.5$
- *Only NE:* $([0.5,0.5]_1, [0.5, 0.5]_2)$ i.e., $p = 0.5$ and $q=0.5$

Notes on Mixed Equilibria

- If player has a *mixed* best response that gives *positive probability* to two (or more) pure strategies, *each* such pure strategy must be a deterministic best response
  - Suppose opponent plays σ and responses have EV(s_1) > EV(s_2)
  - Then response $[p\ s_1;\ (1\text{-}p)\ s_2]$ is worse than s_1 unless $p = 1$
  - *Players can only mix between equally good pure strategies!*
  - *Why mix? Prevents other player from exploiting you…*
- In 2-player, 2-action games, a mixed strategy equilibrium requires the mixture of each player to make the other player *indifferent* to playing either of their pure strategies
  - in matching pennies, ensures other player can't "exploit" what you do
  - this is the basis of some algorithms for computing NE
  - applies to general games (more than two players, two actions)
- **Nash's result:** showed the existence of (mixed) Nash equilibrium for any finite game in normal form

Exercise: Rock Paper Scissors

- Determine mixed equilibrium of rock-paper-scissors
  - Just a version of matching pennies with more moves

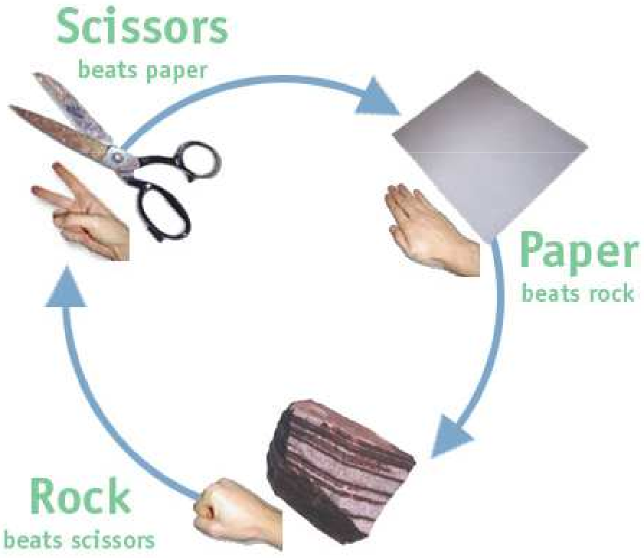

Year	World Champion	Country
2002	Peter Lovering	Canada
2003	Rob Krueger	Canada
2004	Lee Rammage	Canada
2005	Andrew Bergel	Canada
2006	Bob Cooper	United Kingdom
2007	Andrea Farina	USA
2008	Monica Martinez	Canada
2009	Tim Conrad	USA

Exercise: Rock Paper Scissors Lizard Spock

- Or try its more interesting variant:

  rock-paper-scissors-lizard-Spock

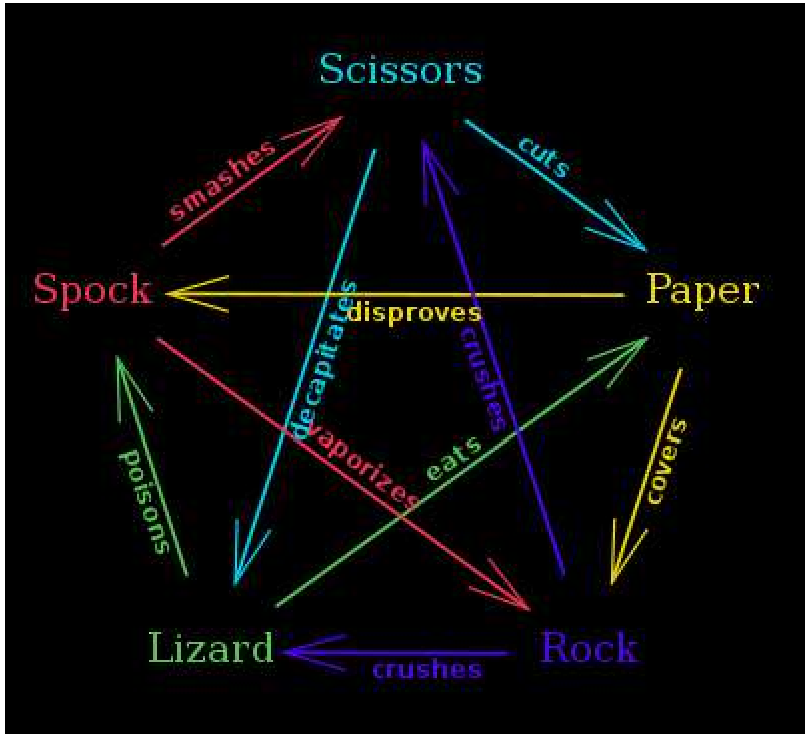

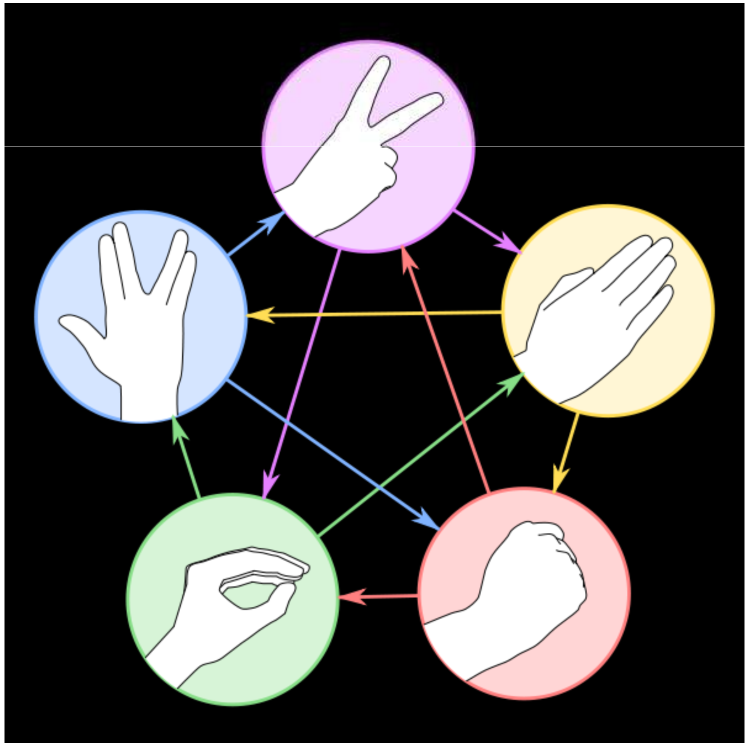

Notes on Mixed Equilibria

- Mixed strategies are rarely “even mixes”
  - Consider run-pass game (American/Canadian football)

		Defense	
		Defend Pass	*Defend Run*
Offense	*Pass*	$0, 0$	$10, -10$
	Run	$5, -5$	$0, 0$

  - No pure NE; only mixed NE requires that defense defends pass with $q = 2/3$; and offence will pass with $p = 1/3$
  - *Why doesn’t offense pass more frequently?*
    - *if it did so, defense would always defend against the pass (and offense would get lower expected payoff)*
    - *why not run more frequently? Same reason.*
    - *threat of passing: defense must defend it more often than it occurs*
  - “Indifference principle” observed (roughly) in NFL
    - avg yds gained per pass attempt close to that of avg. yds. per run

Computing Mixed Equilibria

		P2	
		c	d
P1	a	8 / 6	0 / 10
	b	4 / 4	1 / 2

- We can use principal of indifference to compute mixed Nash equilibria
  - This game has no pure NE, so compute mixed eq:

- *P2's mixture (p_c, p_d) must make P1 indifferent between playing a and b*

$$EV_1(a) = EV_1(b)$$

$$p_c \cdot 8 + (1 - p_c) \cdot 0 = p_c \cdot 4 + (1 - p_c) \cdot 1$$

→ *Recall that $p_d = 1 - p_c$*

$$p_c \cdot 5 = 1$$

$$p_c = {}^{1}/_{5}\,; \;\; p_d = {}^{4}/_{5}$$

- *P1's mixture (p_a, p_b) must make P2 indifferent between playing c and d*

$$EV_2(c) = EV_2(d)$$

$$p_a \cdot 6 + (1 - p_a) \cdot 4 = p_a \cdot 10 + (1 - p_a) \cdot 2$$

→ *Recall that $p_b = 1 - p_a$*

$$p_a \cdot 6 = 2$$

$$p_a = {}^{2}/_{3}\,; \;\; p_b = {}^{1}/_{3}$$

Interpretation of Mixed Equilibria

- **What is meaning of mixed strategies and NE?**

- *Genuine randomization* of choices to prevent exploitation
  - soccer penalty kicks, bluffing in poker, many contests and sports, ...

- Probabilities represent *proportions of strategies in a population*
  - likelihood of two different types of agents (behaviors, species, competing fashions, …) meeting is given by relative proportions
  - payoff interpreted as successful propagation of that agent type
  - style of analysis used a lot in evolutionary biology to explain emergence of many phenomena (altruism, sexual differentiation, etc.); sociology and anthropology to explain evolution of culture, norms, ideas; etc.
  - see Ch.7 of E&K

Interpretation of Mixed Equilibria

- Conditions on *stable beliefs* (or expectations, conjectures) about how opponent will play
  - if I believe you will play q = 0.5, and I believe that you believe I will play p = 0.5, and I believe that you believe that I believe that you will play q = 0.5, and …

- Learning interpretations: beliefs based on experience with opponent (or related opponents)
  - I've (P2) seen you (P1) play heads 47 times and tails 53 times: I assume your strategy is p = 0.47
  - My best response is heads (q = 1.0)
  - But you learn in similar fashion and *also* adopt best response to what you've observed
  - In some cases, observed frequencies converge to NE
  - But actually play may often be deterministic (and quite cyclic, not random at all)

Pareto Optimality

(will come back to this in Ch.8)

	Quiet	Confess
Quiet	-1 / -1	-10 / 0
Confess	0 / -10	-4 / -4

- We called some NE “undesirable” (e.g., Prisoners’s dilemma)
- Need some notion of “optimality” to compare payoff profiles or *strategy profiles*
- *Pareto optimality:* a strategy profile is *Pareto optimal* iff no other strategy profile gives some player(s) a higher (expected) payoff without lowering the payoffs of others
  - Everyone would “agree” to *switch* from a non-Pareto optimal profile to a Pareto-optimal profile (if agreement “enforceable”)
- In PD: the unique NE is not Pareto optimal
  - Cooperative profile is better for all involved (but not enforceable)
  - Note: Q/C and C/Q are also both Pareto optimal (only Quiet player wants to switch)

Social Optimality

(will come back to this in Ch.8)

	Quiet	Confess
Quiet	-1 / -1	-10 / 0
Confess	0 / -10	-4 / -4

- Pareto optimality desirable, but weak
  - and not always achievable in “noncooperative” games
  - cooperative game theory: models commitments/agreements
- *Social optimality:* a strategy profile is *socially optimal* iff it gives rise to the highest *sum* of expected payoffs
  - Total payoff to entire group (*social welfare*) is maximized
  - Assumes that summing payoffs in sensible
  - If a profile is socially optimal, it must be Pareto optimal (if one player better off without hurting others, total payoff must increase), but not vice versa
- In PD: the unique NE is not socially optimal

Extensive Form (Dynamic) Games

- Normal form (matrix) games seem limited
  - Players move simultaneously and outcome determined at once
  - No observation, reaction, etc.
- Most games have a *dynamic* structure (turn taking)
  - Chess, tic-tac-toe, cards games, soccer, corporate decisions, markets…
  - See what "opponent" does before making move
  - Sometimes you only see partial information (won't discuss this)

- Example: Rogers, Telus competing for market in two remote areas
  - Each firm, Rogers and Telus, can tackle one area only
  - Total revenue in Area A: 12, Area B: 6
  - If firm is alone in one area, get all of that area's revenue
  - If both firms target same area, "first mover" gets 2/3, second 1/3
  - Rogers prepped: makes first move, Telus chooses *after* Rogers
  - Critical: Telus *observes* Rogers' choice before moving!

Game Tree (Extensive Form)

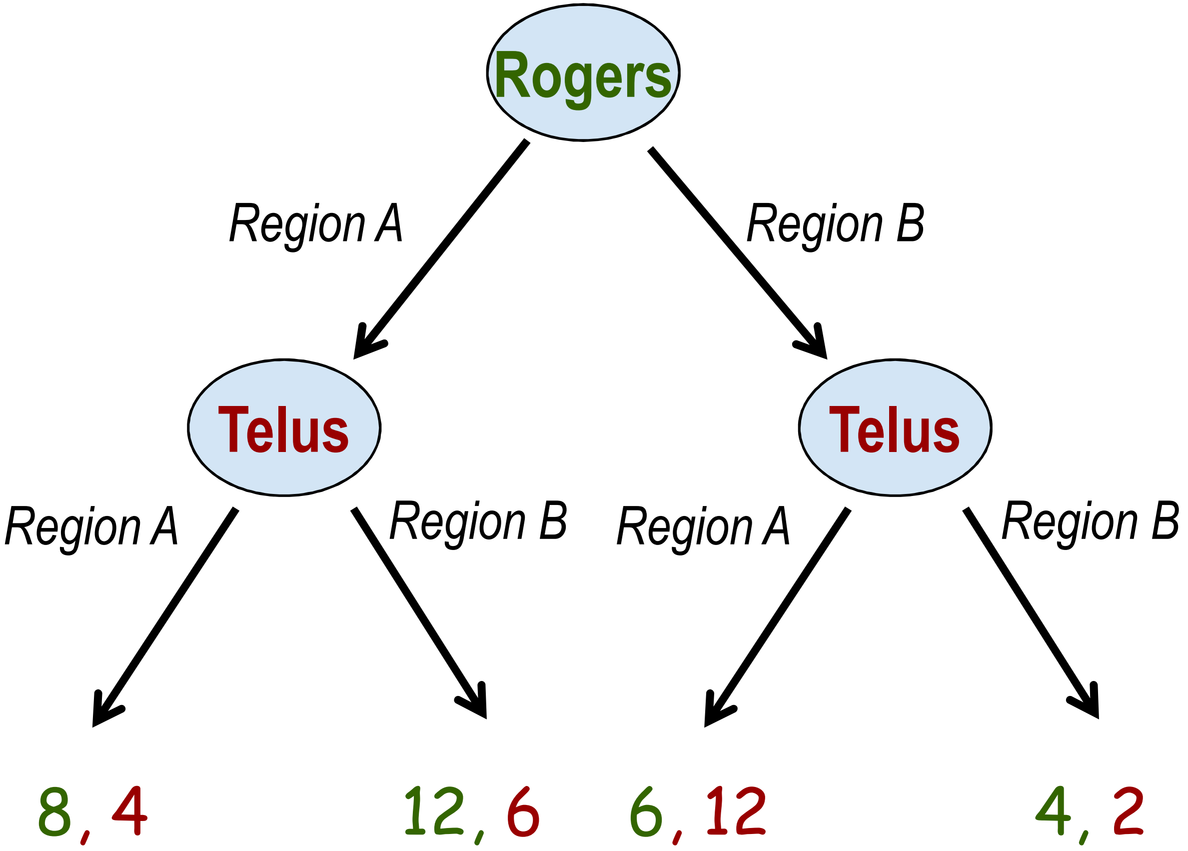

Rogers' choice

Telus' choice

R's payoff, T's payoff

Strategies Adopted

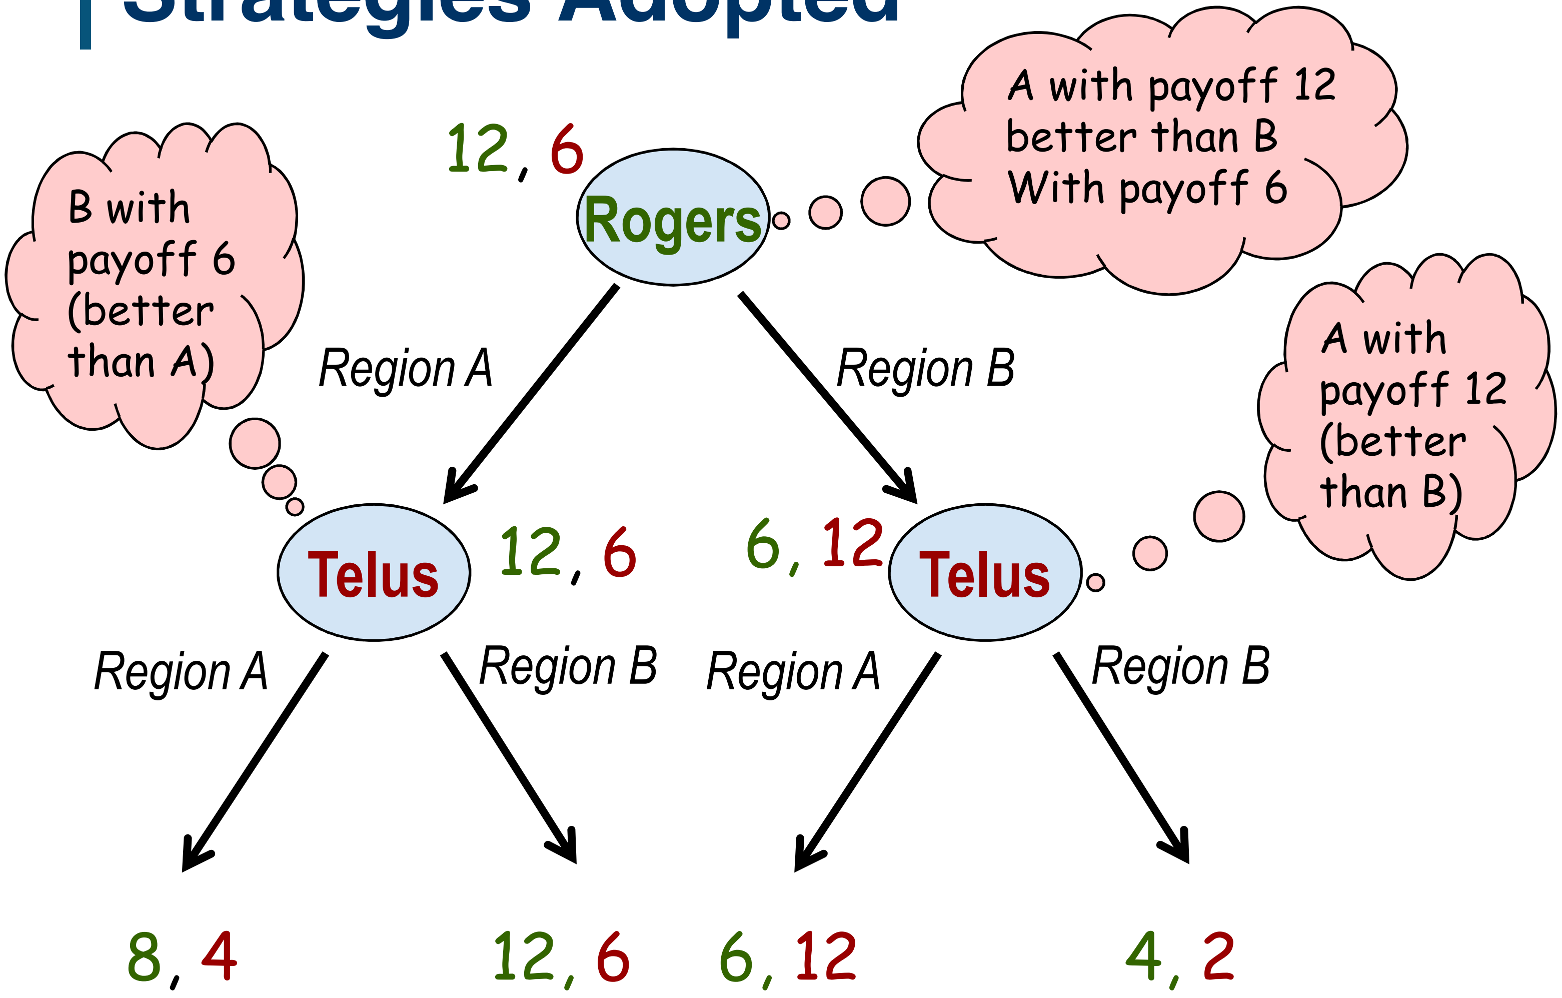

Backward Induction

- Start from choice nodes at bottom of tree
  - Player at that node chooses best move
  - E.g., Telus chooses B at node “Rogers did A”, A at node “Rogers did B”
  - This dictates which terminal node (payoffs) will be reached
  - We can now assign payoffs to that node (from chosen terminal node)
- Work up the tree: compute choices at other nodes once choices/payoffs at all child nodes made
  - Player at that node chooses best move
  - E.g., Rogers chooses A at root node
  - Dictates which child (“Rogers did A”) will be reached, hence which payoff
- At end of procedure:
  - Each choice node labelled with action/choice and payoff vector
  - Payoff for the game is the payoff vector at the root node
  - Path through tree given by choices, tells us how the game will unfold