Game Theory

Russel & Norvig Chapter 6

Russel & Norvig Section 17.6

Types of Games (informal)

	Deterministic	Chance
Perfect Information	Chess, Checkers Go	Backgammon, Monopoly
Imperfect Information	Battleship	Bridge, Poker, Scrabble, wargames

Definitions

- *Two-player game*: Player A and B. Player A starts.
- *Deterministic*: None of the moves/states are subject to chance (no random draws).
- *Zero-sum*: Player's A gain is exactly equal to player B's loss. One of the player's must win or there is a draw (both gains are equal).

Example

- Initially a stack of pennies stands between two players
- Each player divides one of the current stacks into two unequal stacks.
- The game ends when every stack contains one or two pennies
- The first player who cannot play loses

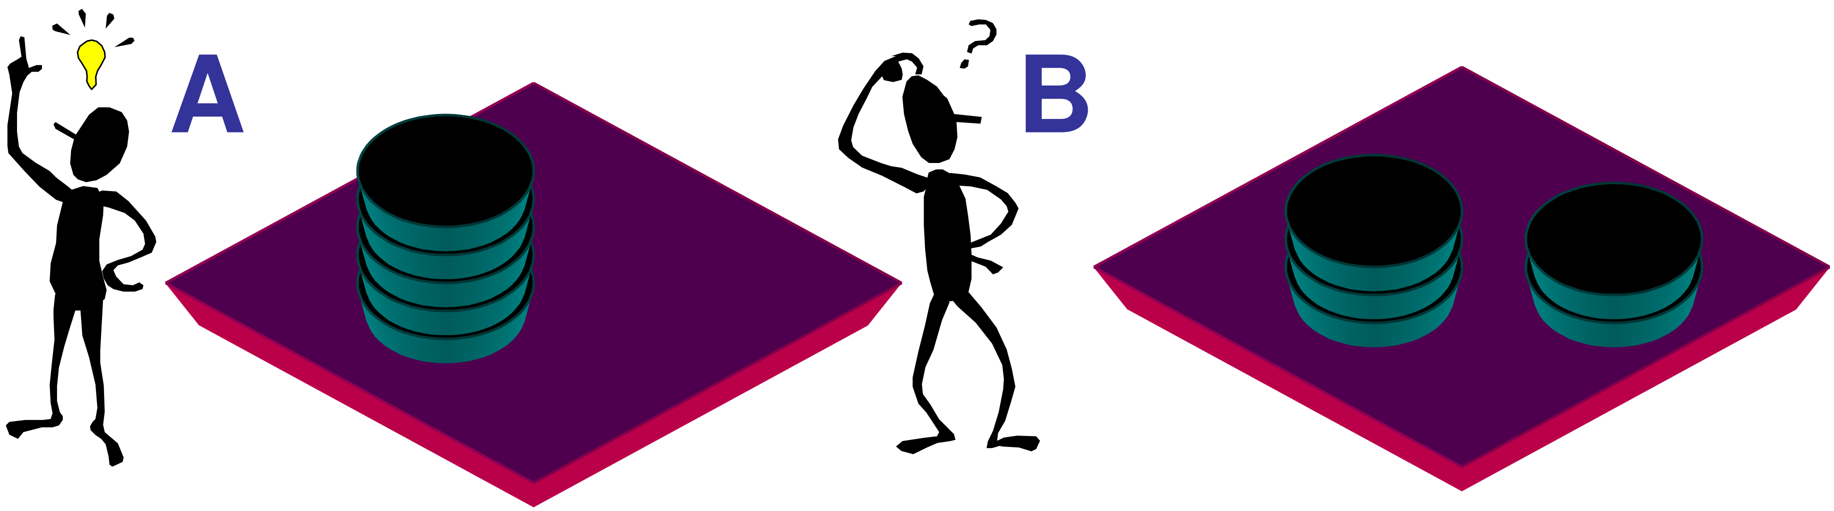

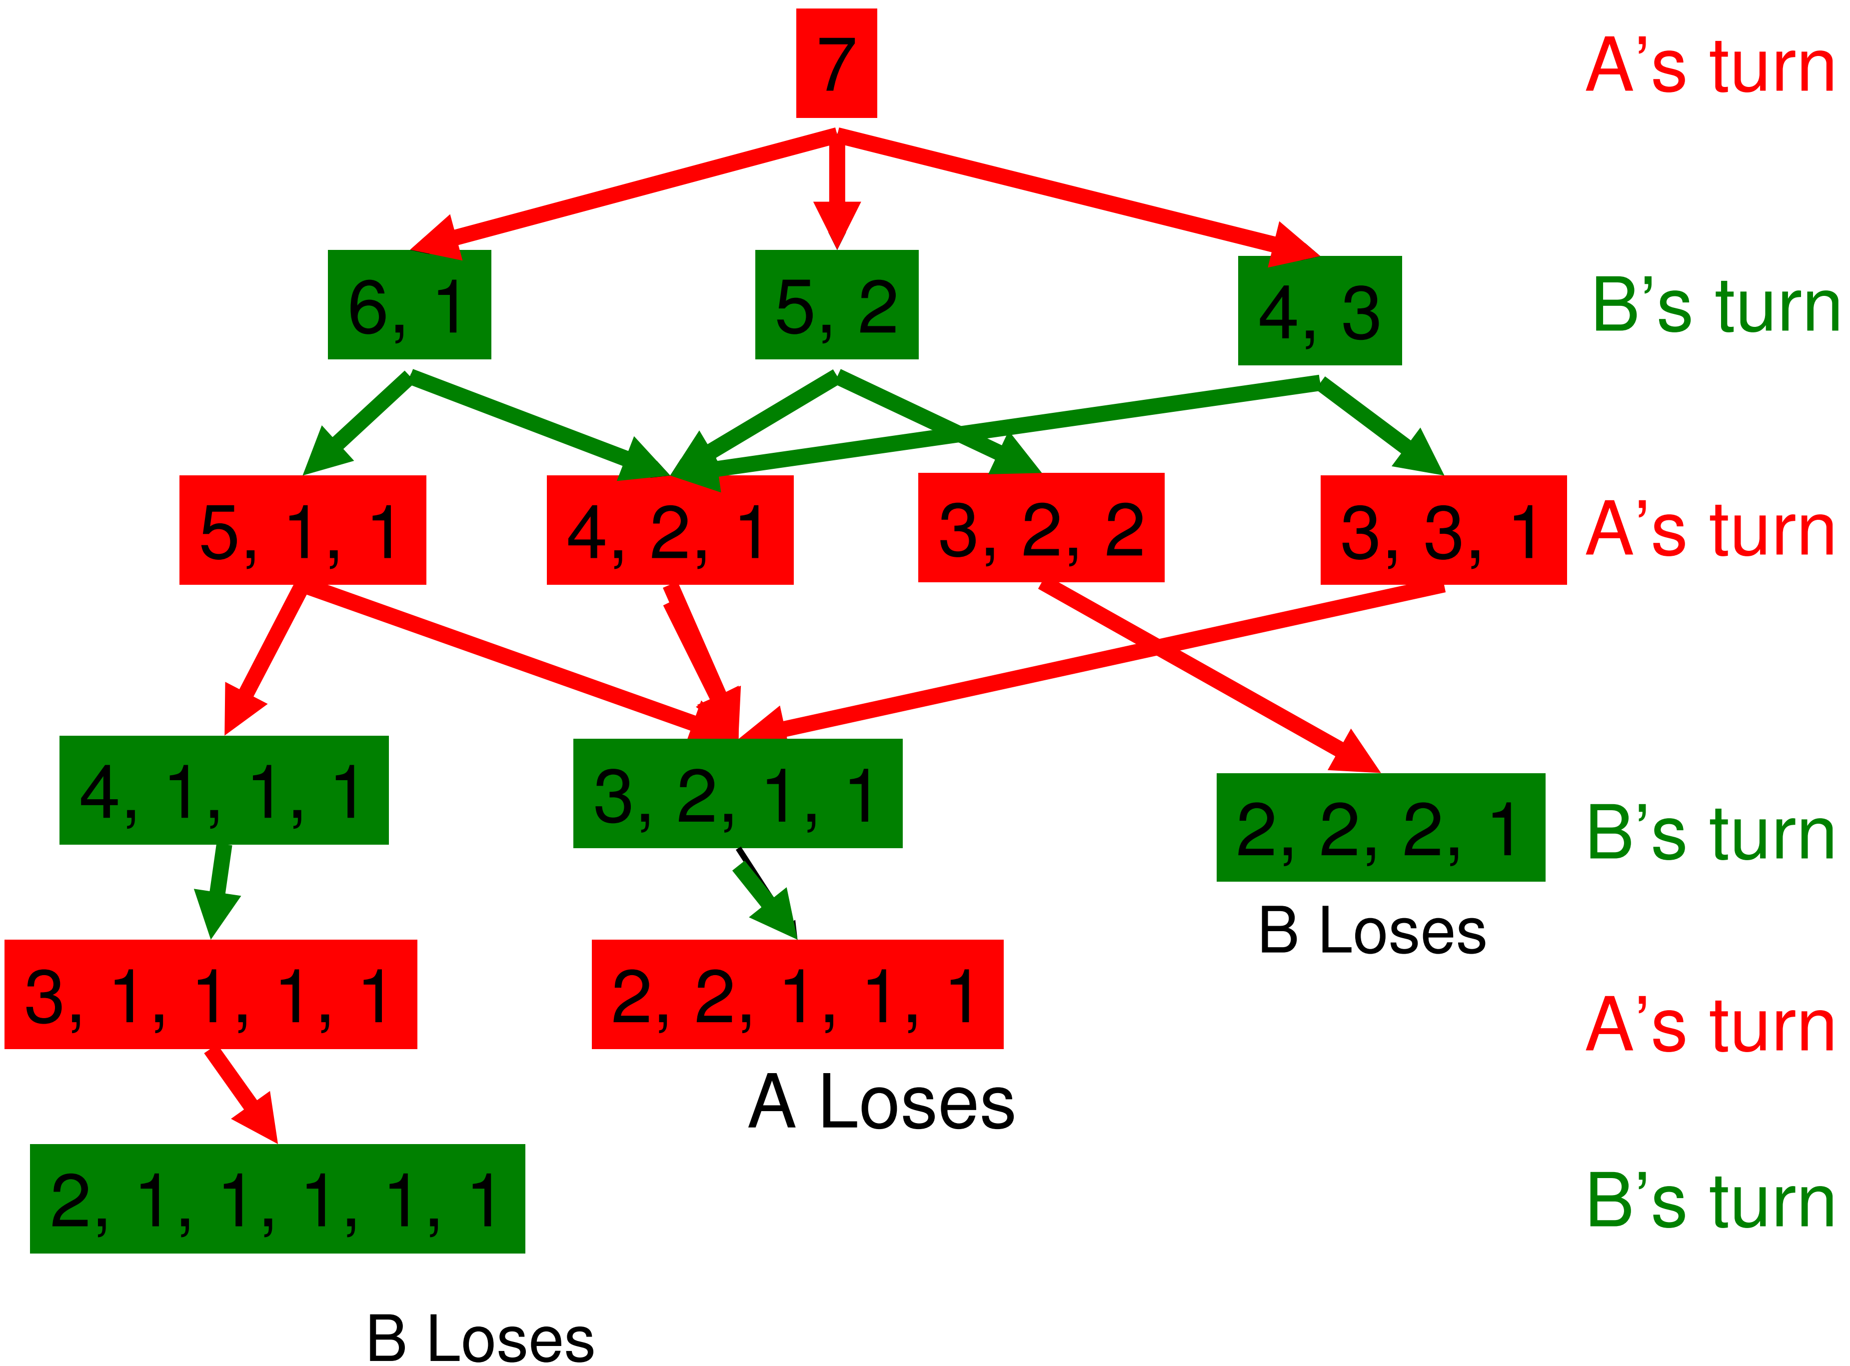

7
A's turn
6, 1
5, 2
4, 3
B's turn
5, 1, 1
4, 2, 1
3, 2, 2
3, 3, 1
A's turn
4, 1, 1, 1
3, 2, 1, 1
2, 2, 2, 1
B's turn
B Loses
3, 1, 1, 1, 1
2, 2, 1, 1, 1
A's turn
A Loses
2, 1, 1, 1, 1, 1
B's turn
B Loses

Search Problem

- *Payoff/Utility*: Numerical value assigned to each terminal state. Example:
  - $U(s) = +1$ for A win, -1 for B win
- *Game value:* The value of a terminal that will be reached assuming optimal strategies from both players (*minimax* value)
- *Search*: Find move that maximizes game value from current state

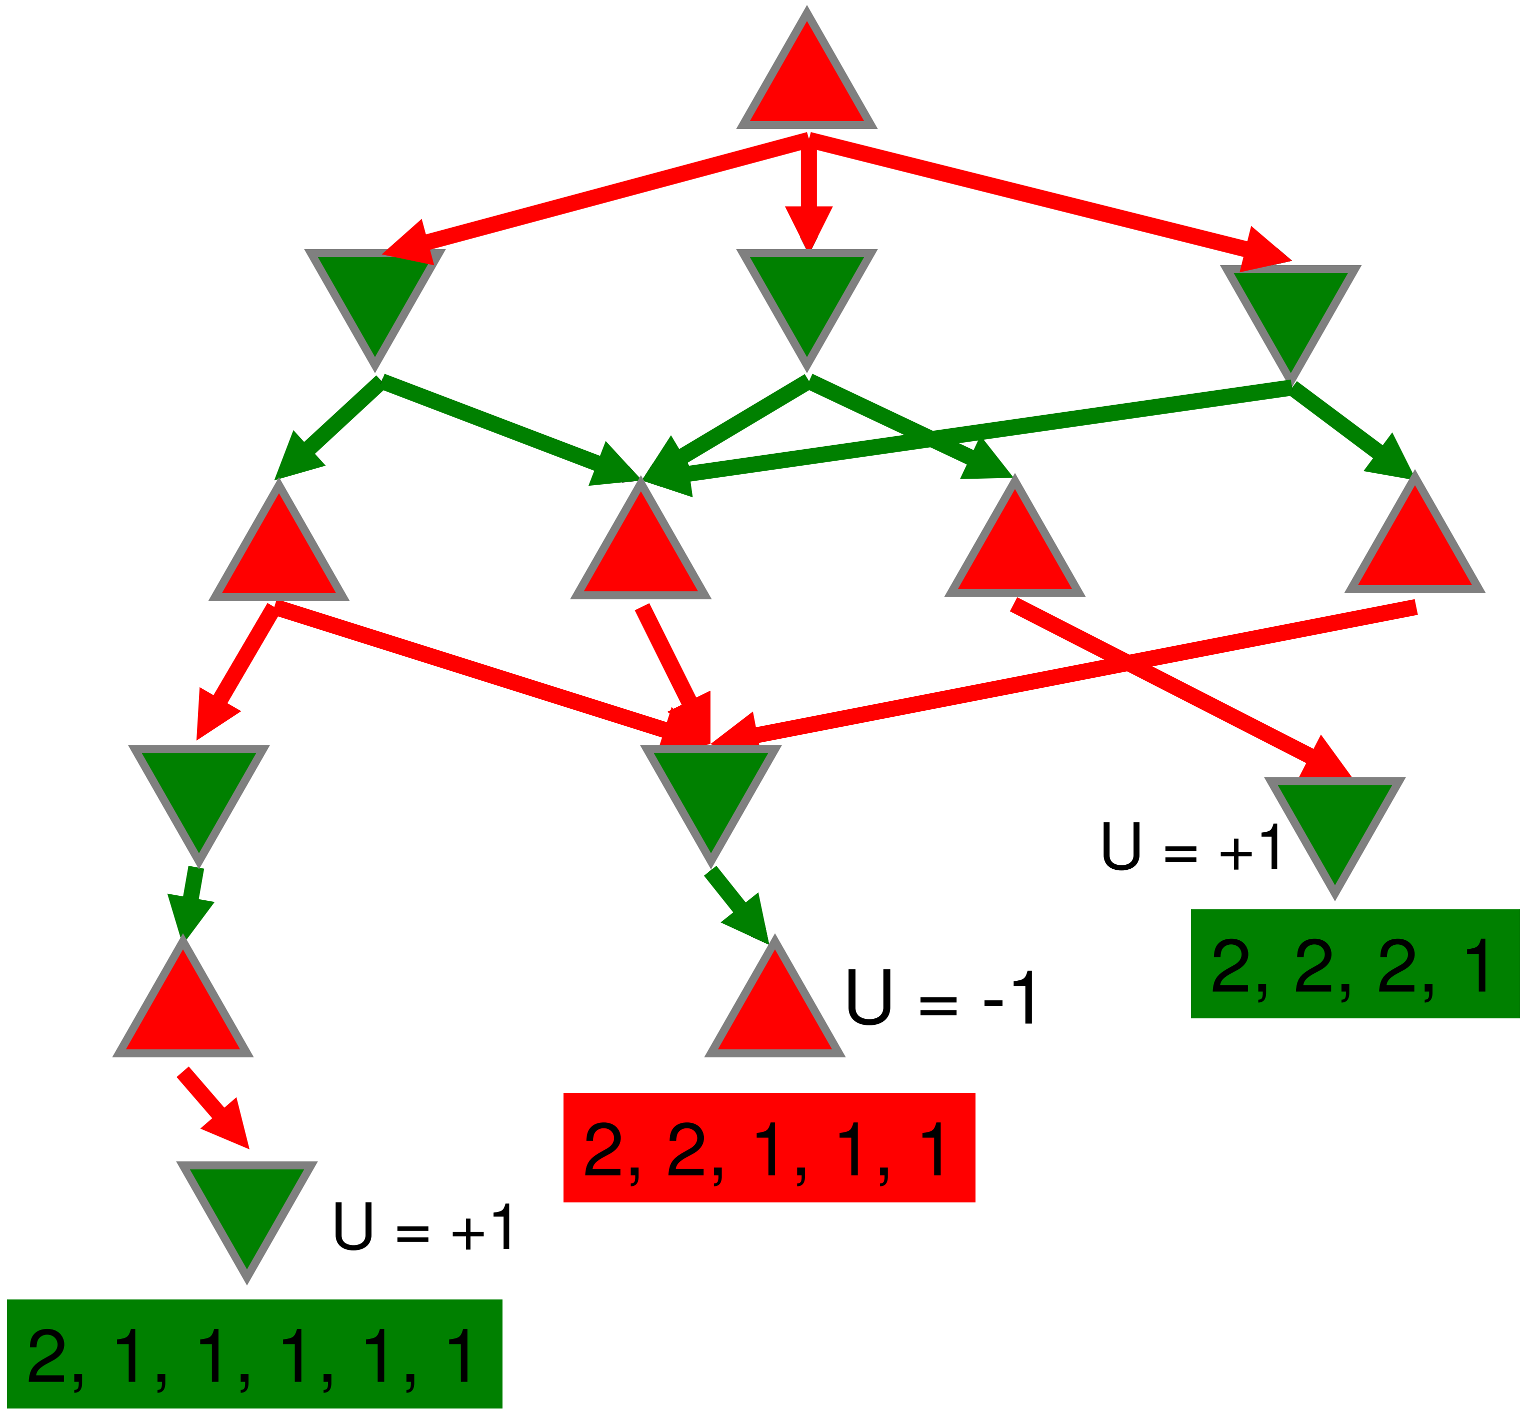
U = +1
2, 2, 2, 1
U = -1
2, 2, 1, 1, 1
U = +1
2, 1, 1, 1, 1, 1

Optimal (minimax) Strategies

- Search the game tree such that:
  - A's turn to move → find the move that yields maximum payoff from the corresponding subtree → This is the move most favorable to A
  - B's turn to move → find the move that yields minimum payoff (best for B) from the corresponding subtree → This is the move most favorable to B

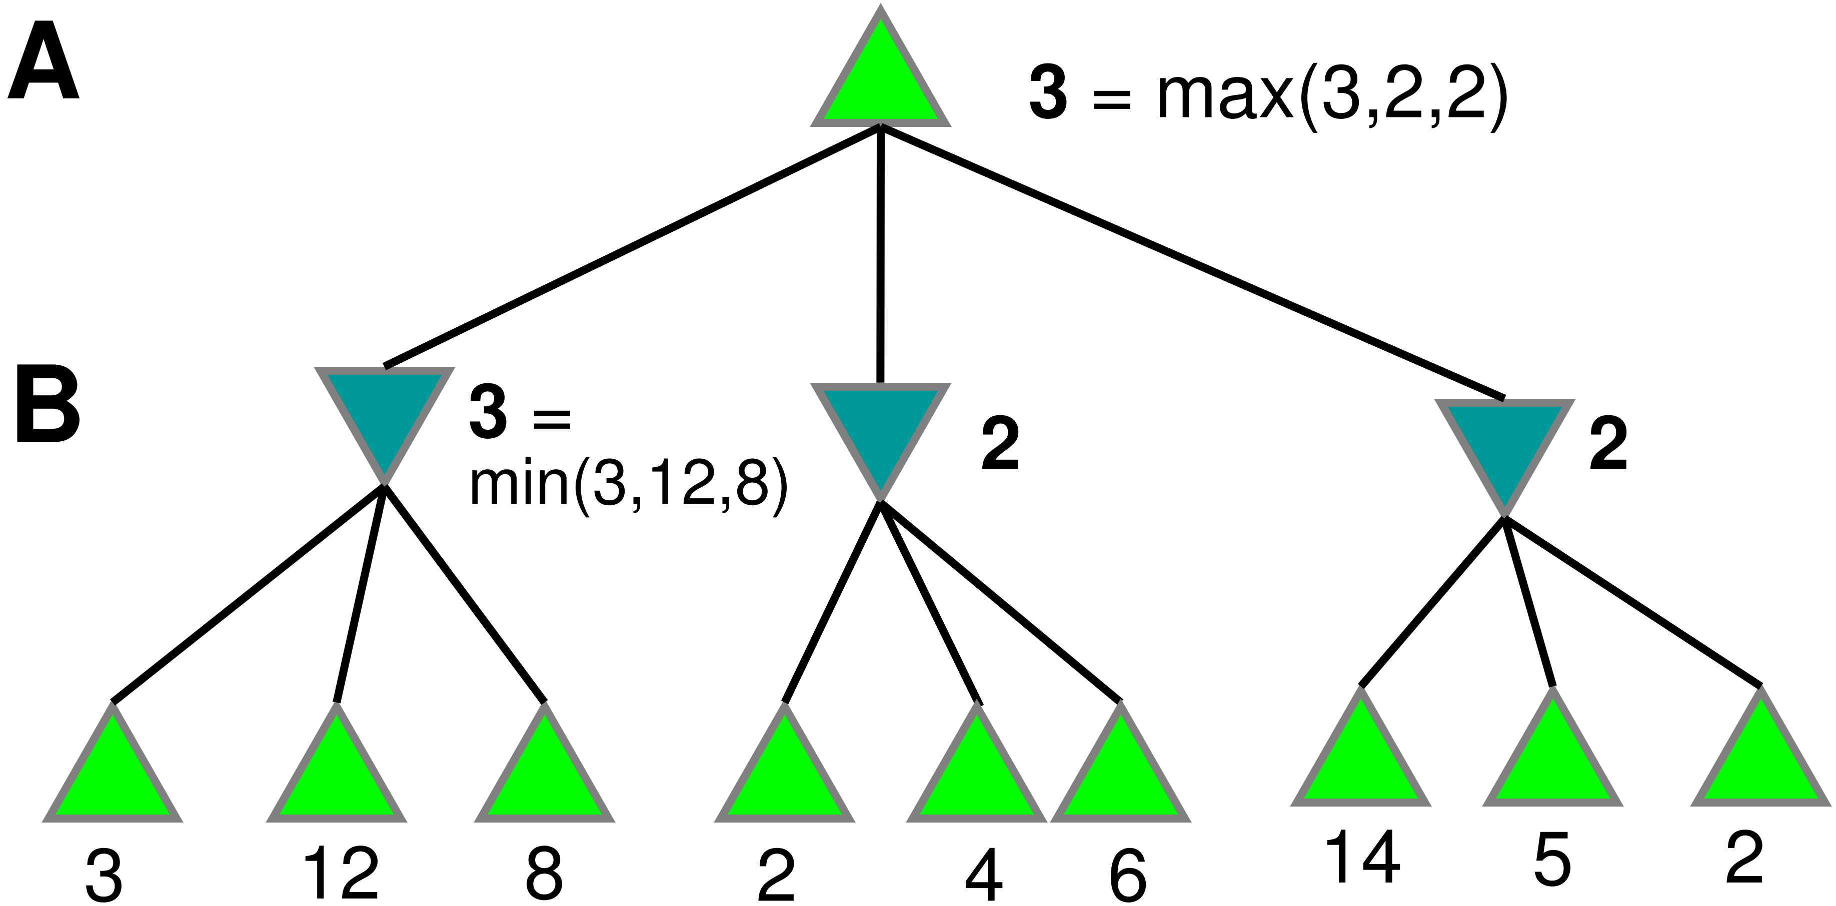
A
3 = max(3,2,2)
B
3 =
min(3,12,8)
2
2
3
12
8
2
4
6
14
5
2

Minimax

Minimax (s)

If s is terminal

Return $U(s)$

If next move is A

Return $\boxed{\max_{s' \in Succs(s)} \text{Minimax}(s')}$

Else

Return $\boxed{\min_{s' \in Succs(s)} \text{Minimax}(s')}$

Minimax Properties

- *Complete:* If finite game
- *Optimal:* If opponent plays optimally
- Essentially DFS

Matrix Form of Games

R&N Chapter 6

R&N Section 17.6

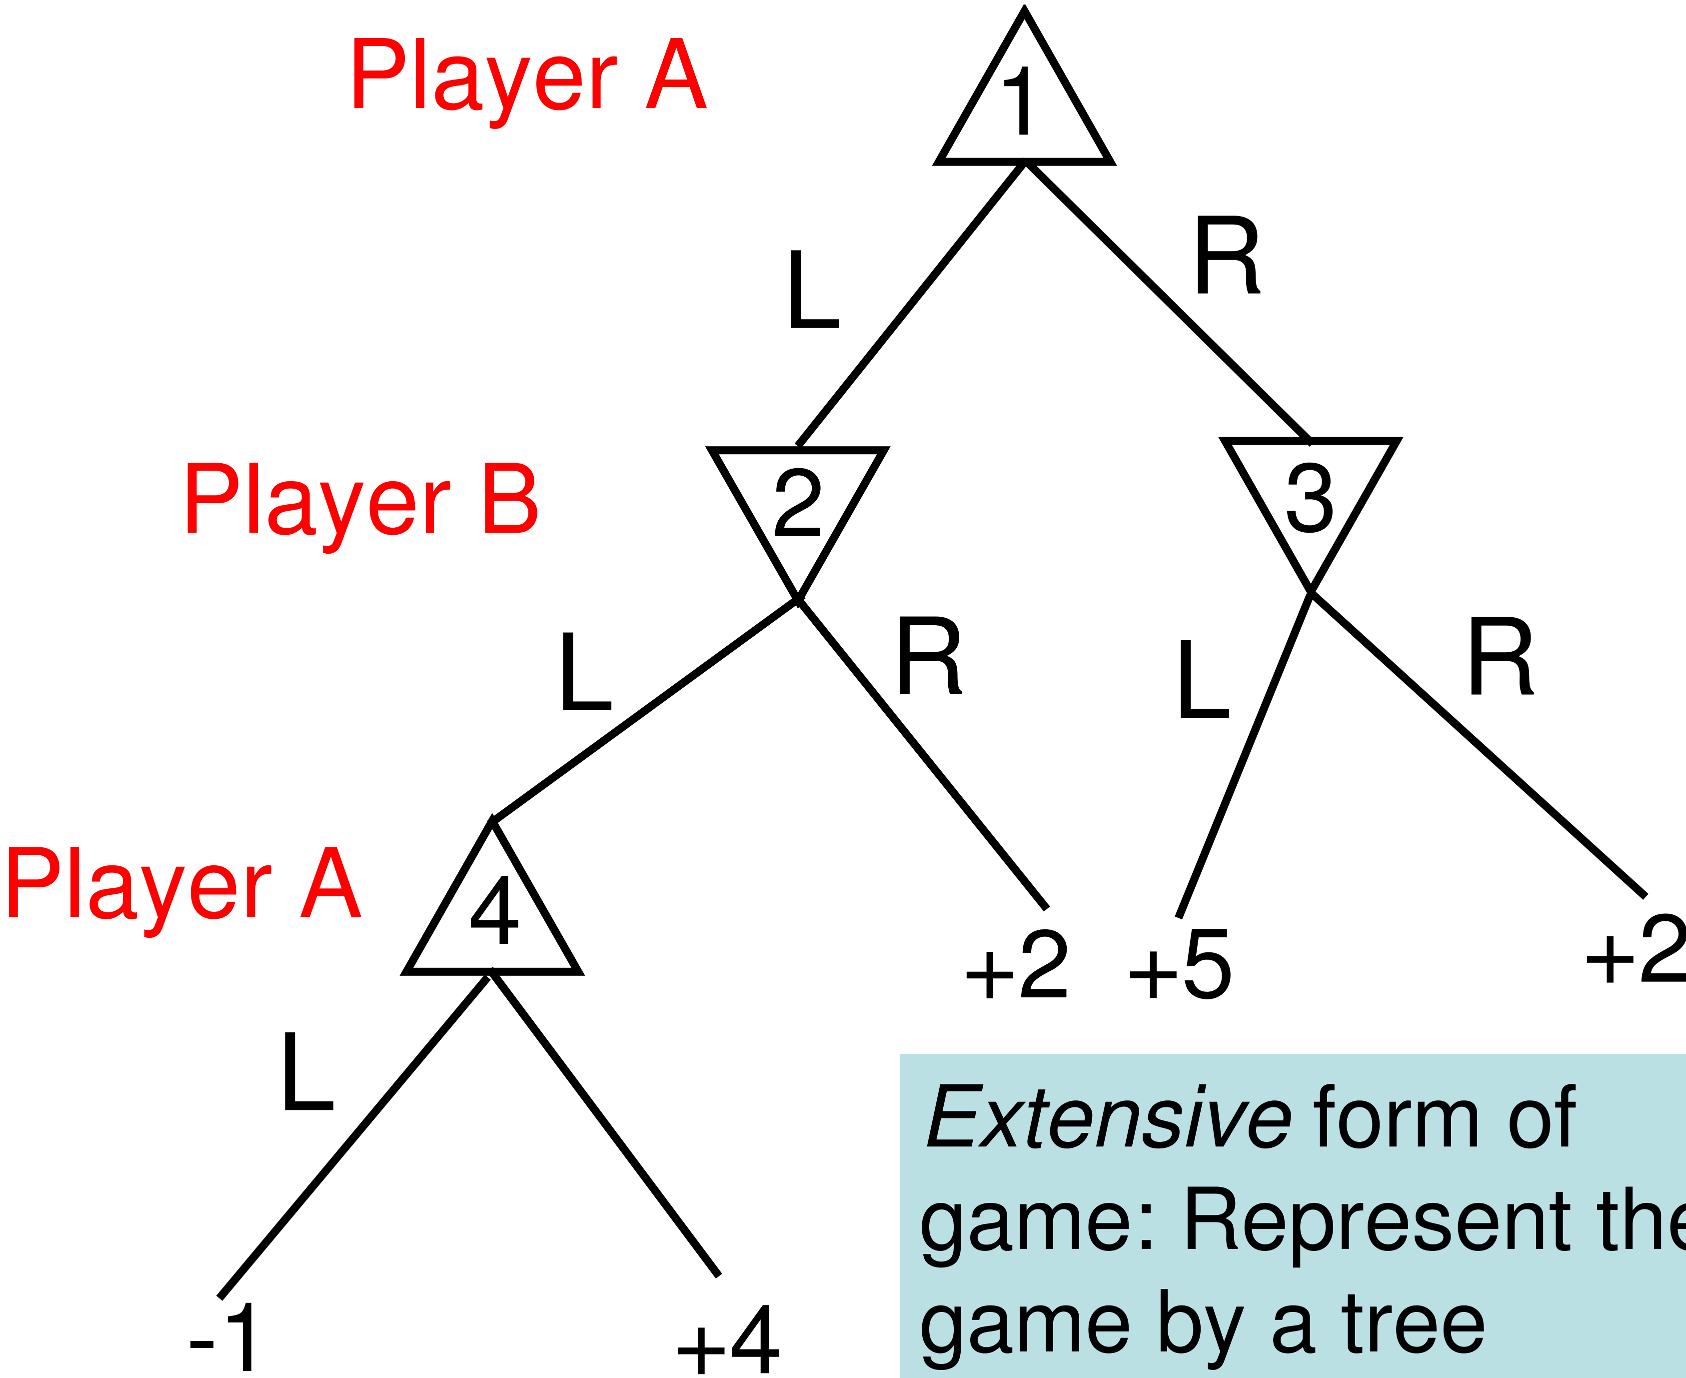

Extensive form of game: Represent the game by a tree

A *pure strategy* for a player defines the move that the player would make for every possible state that the player would see.

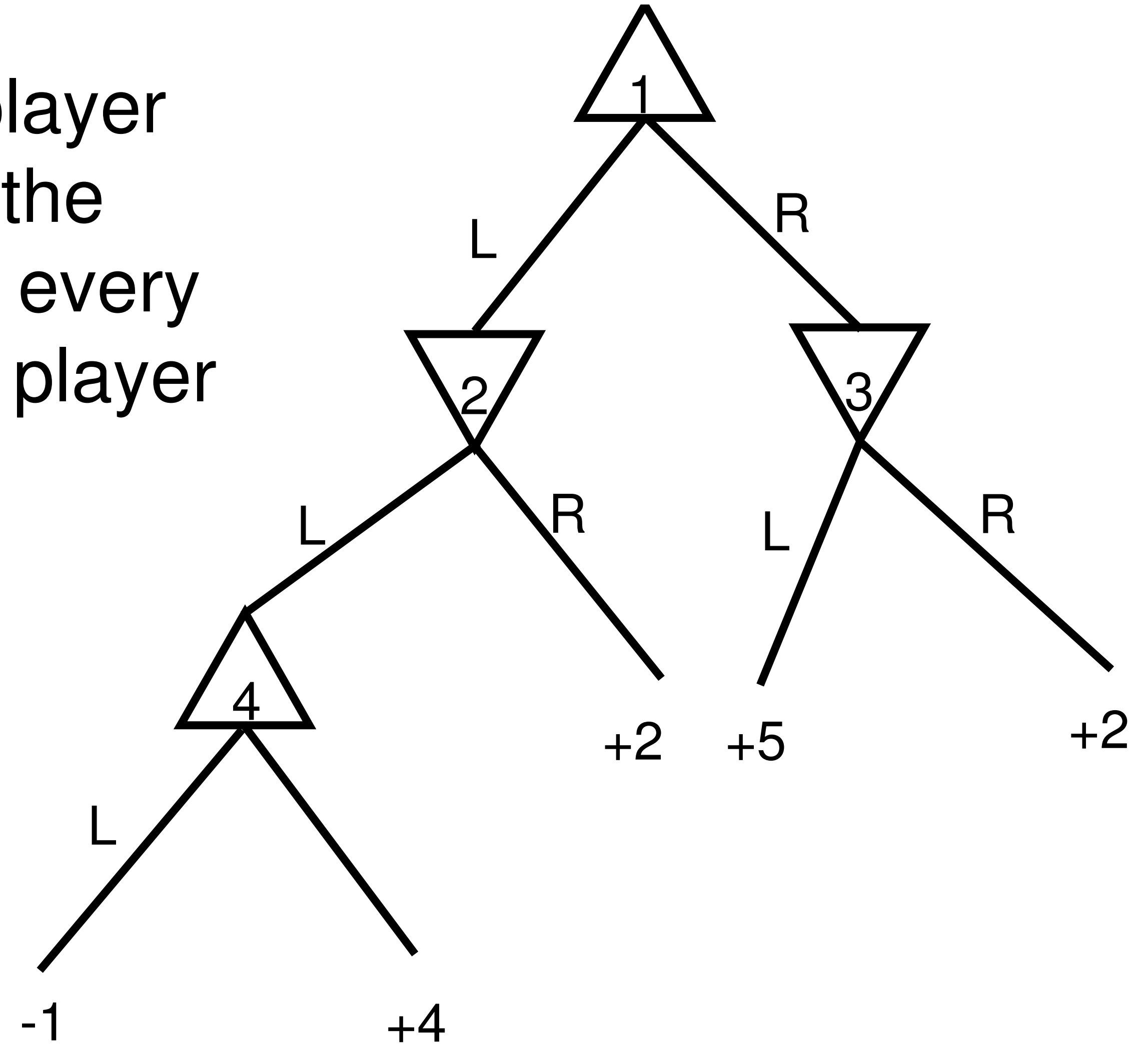

Pure strategies for A:

Strategy I: (1→L,4→L)

Strategy II: (1→L,4→R)

Strategy III: (1→R,4→L)

Strategy IV: (1→R,4→R)

Pure strategies for B:

Strategy I: (2→L,3→L)

Strategy II: (2→L,3→R)

Strategy III: (2→R,3→L)

Strategy IV: (2→R,3→R)

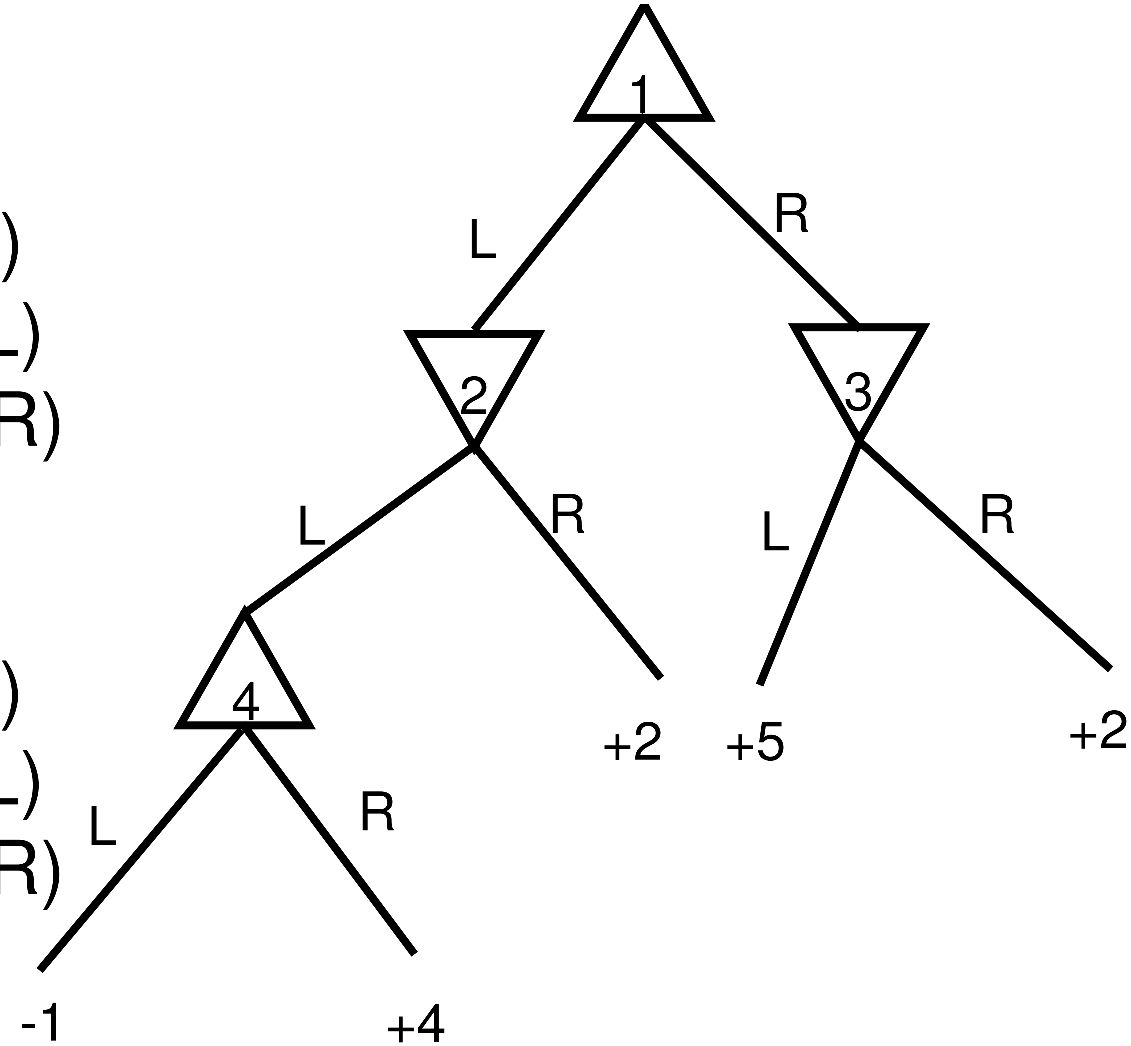

In general: If *N* states and *B* moves, how many pure strategies exist?

Matrix form of games

Pure strategies for A:
Strategy I: (1→L,4→L)
Strategy II: (1→L,4→R)
Strategy III: (1→R,4→L)
Strategy IV: (1→R,4→R)

Pure strategies for B:
Strategy I: (2→L,3→L)
Strategy II: (2→L,3→R)
Strategy III: (2→R,3→L)
Strategy IV: (2→R,3→R)

	I	II	III	IV
I	-1	-1	+2	+2
II	+4	+4	+2	+2
III	+5	+1	+5	+1
IV	+5	+1	+5	+1

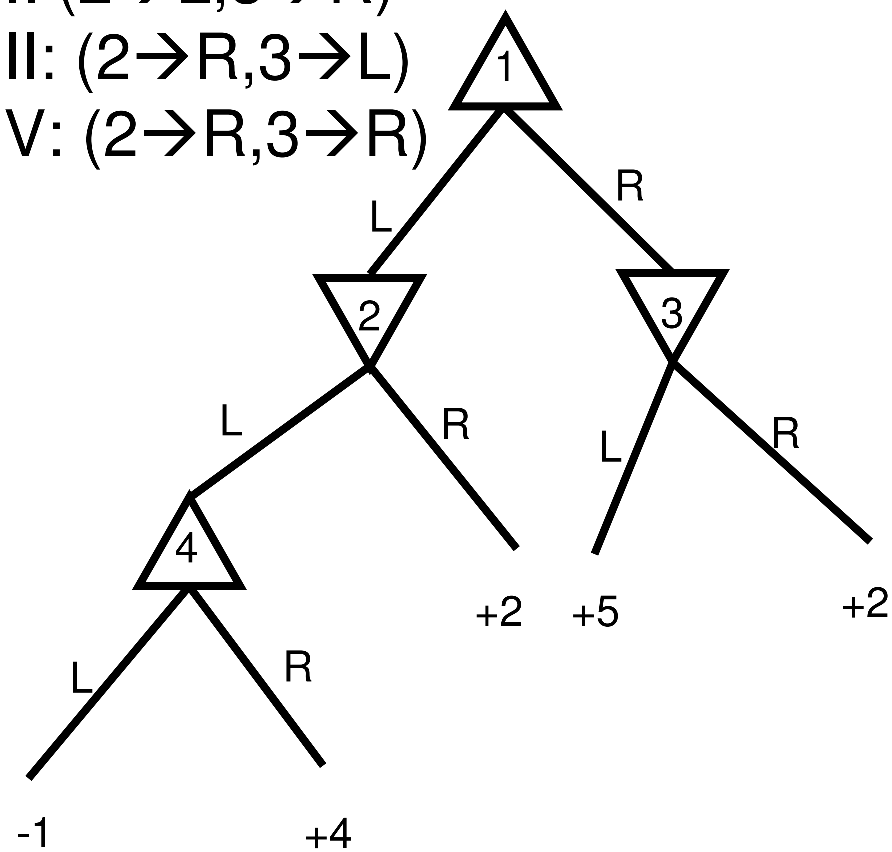

Pure strategies for Player B

Pure strategies for Player A

	I	II	III	IV
I	-1	-1	+2	+2
II	+4	+4	+2	+2
III	+5	+1	+5	+1
IV	+5	+1	+5	+1

Player A's payoff if game is played with strategy I by Player A and strategy III by Player B

- *Matrix normal form* of games: The table contains the payoffs for all the possible combinations of pure strategies for Player A and Player B
- The table characterizes the game completely, there is no need for any additional information about rules, etc.
- Although, in many cases, the number of pure strategies may be too large for the table to be represented explicitly, the matrix representation is the basic representation that is used for deriving fundamental properties of games.

Minimax →Matrix version

Max value of all the rows

		I	II	III	IV
-1	I	-1	-1	+2	+2
+2	II	+4	+4	+2	+2
+1	III	+5	+1	+5	+1
+1	IV	+5	+1	+5	+1

Min value across each row

$$\underset{Rows\ i}{Max}\ \underset{Columns\ j}{Min}\ \mathbf{M}(i,j)$$

Minimax →Matrix version

Max value =
game value = +2

- For each strategy (each row of the game matrix), Player A should assume that Player B will use the optimal strategy given Player A's strategy (the strategy with the minimum value in the row of the matrix). Therefore the best value that Player can achieve is the maximum over all the rows of the minimum values across each of the rows:

$$\boxed{\underset{Rows\ i}{Max}\ \underset{Columns\ j}{Min}\ \mathbf{M}(\boldsymbol{i},\boldsymbol{j})}$$

- The corresponding pure strategy is the optimal solution for this game → It is the optimal strategy for A assuming that B plays optimally.

		I	II	III	IV
-1	I	-1	-1	+2	+2
+2	II	+4	+4	+2	+2
+1	III	+5	+1	+5	+1
+1	IV	+5	+1	+5	+1

Min value across each row

Max value across each column

	I	II	III	IV
I	-1	-1	+2	+2
II	+4	+4	+2	+2
III	+5	+1	+5	+1
IV	+5	+1	+5	+1
	+5	+4	+5	$\boxed{+2}$

Min of all the columns

$$\underset{Columns\ j}{Min}\ \underset{Rows\ i}{Max}\ \mathbf{M}(i,j)$$

Minimax or Maximin?

- But we could have used the opposite argument:
- For each strategy (each column of the game matrix), Player B should assume that Player A will use the optimal strategy given Player B's strategy (the strategy with the maximum value in the column of the matrix):

$$\boxed{\underset{Columns\ j}{Min}\ \underset{Rows\ i}{Max}\ \mathbf{M}(i, j)}$$

- Therefore the best value that Player B can achieve is the minimum over all the columns of the maximum values across each of the columns
- Problem: Do we get to the same result??
- Is there always a solution?

Max value across each column

	I	II	III	IV
I	-1	-1	+2	+2
II	+4	+4	+2	+2
III	+5	+1	+5	+1
IV	+5	+1	+5	+1

+5 +4 +5 +2

Min value =
game value = +2

Max value =
game value = +2

Note that we find the same value and same strategies in both cases. Is that always the case?

Min value across each row		I	II	III	IV
-1	I	-1	-1	+2	+2
+2	II	+4	+4	+2	$\boxed{+2}$
+1	III	+5	+1	+5	+1
+1	IV	+5	+1	+5	+1

Min value across each row

$$\underset{Rows\ i}{Max}\ \underset{Columns\ j}{Min}\ \mathbf{M}(i,j)$$

Max value across each column

	I	II	III	IV
I	-1	-1	+2	+2
II	+4	+4	+2	$\boxed{+2}$
III	+5	+1	+5	+1
IV	+5	+1	+5	+1
	+5	+4	+5	+2

Min value =
game value = +2

$$\underset{Columns\ j}{Min}\ \underset{Rows\ i}{Max}\ \mathbf{M}(i,j)$$

Minimax vs. Maximin

- Fundamental Theorem I (von Neumann):
  - For a two-player, zero-sum game with perfect information:
    - *There always exists an optimal pure strategy for each player*
    - *Minimax = Maximin*

Games with Hidden Information

R&N Chapter 6

R&N Section 17.6

Another (Seemingly Simple) Game

- The two Players A and B each have a coin
- They show each other their coin, choosing to show either head or tail.
- If they both choose head → Player B pays Player A \$2
- If they both choose tail → Player B pays Player A \$1
- If they choose different sides → Player A pays Player B \$1

Side Note about all toy examples

- If you find this kind of toy example annoying, it models a large number of real-life situations.
- For example: Player A is a business owner (e.g., a restaurant, a plant..) and Player B is an inspector. The inspector picks a day to conduct the inspection; the owner picks a day to hide the bad stuff. Player B wins if the days are different; Player A wins if the days are the same.
- This class of problems can be reduced to the “coin game” (with different payoff distributions, perhaps).

Extensive Form

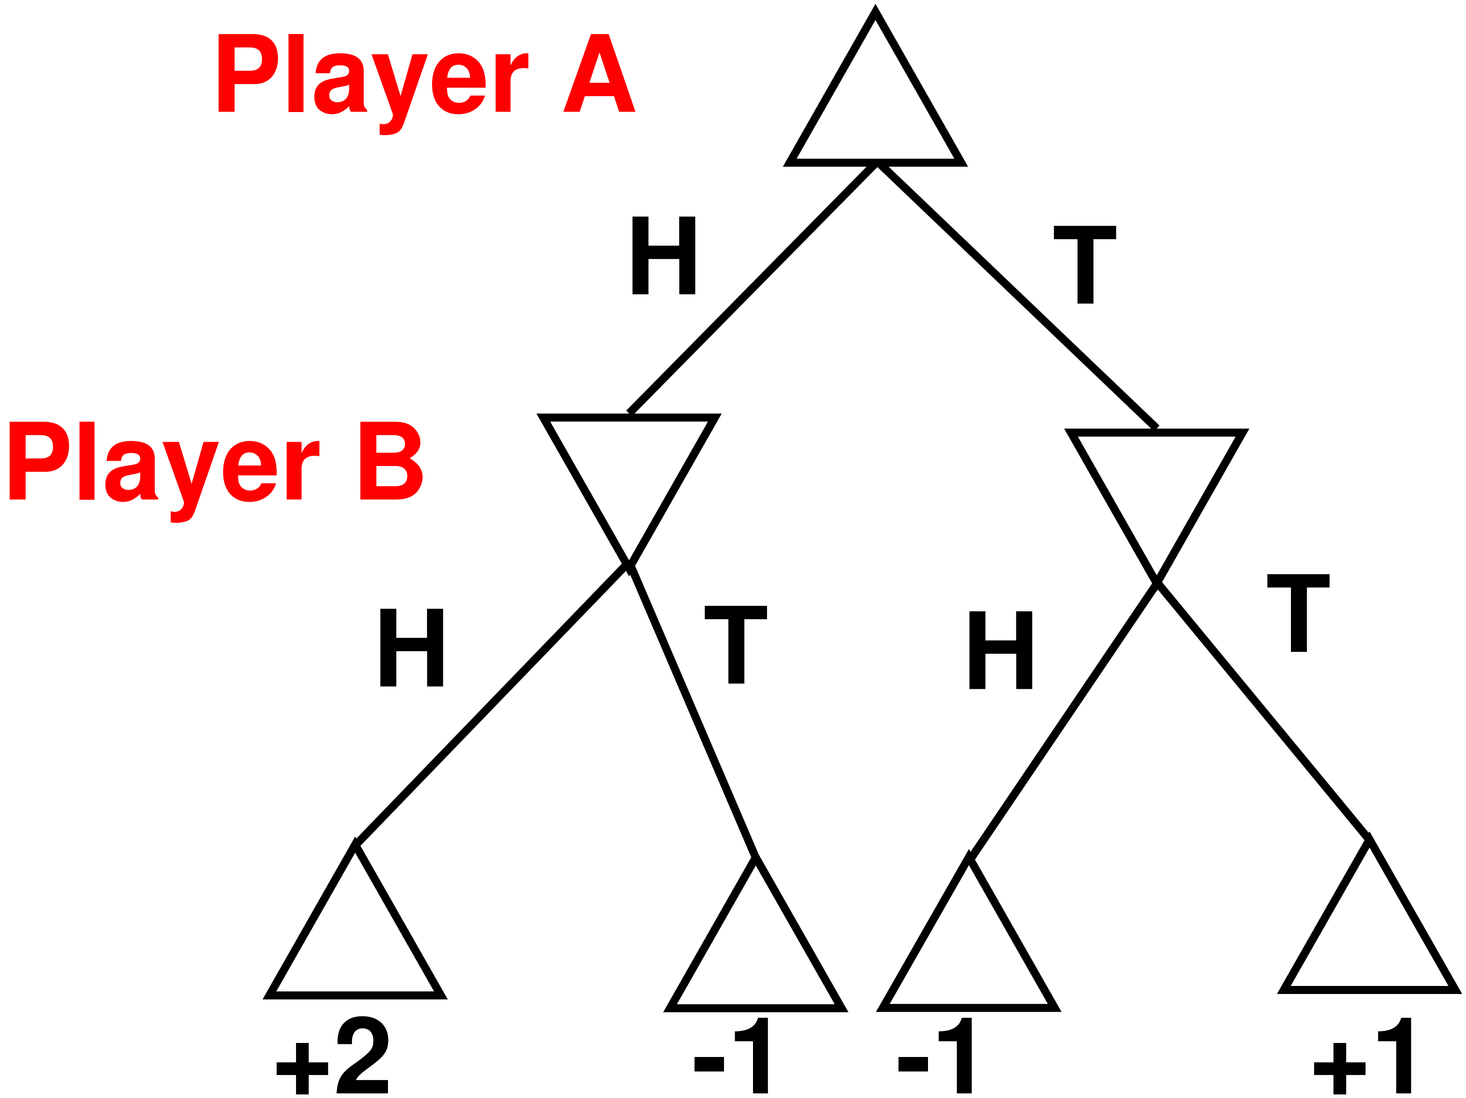

Extensive Form

Problem: Since the moves are simultaneous, Player B does not know which move Player A chose → This is no longer a game with prefect information → we have to deal with *hidden information*

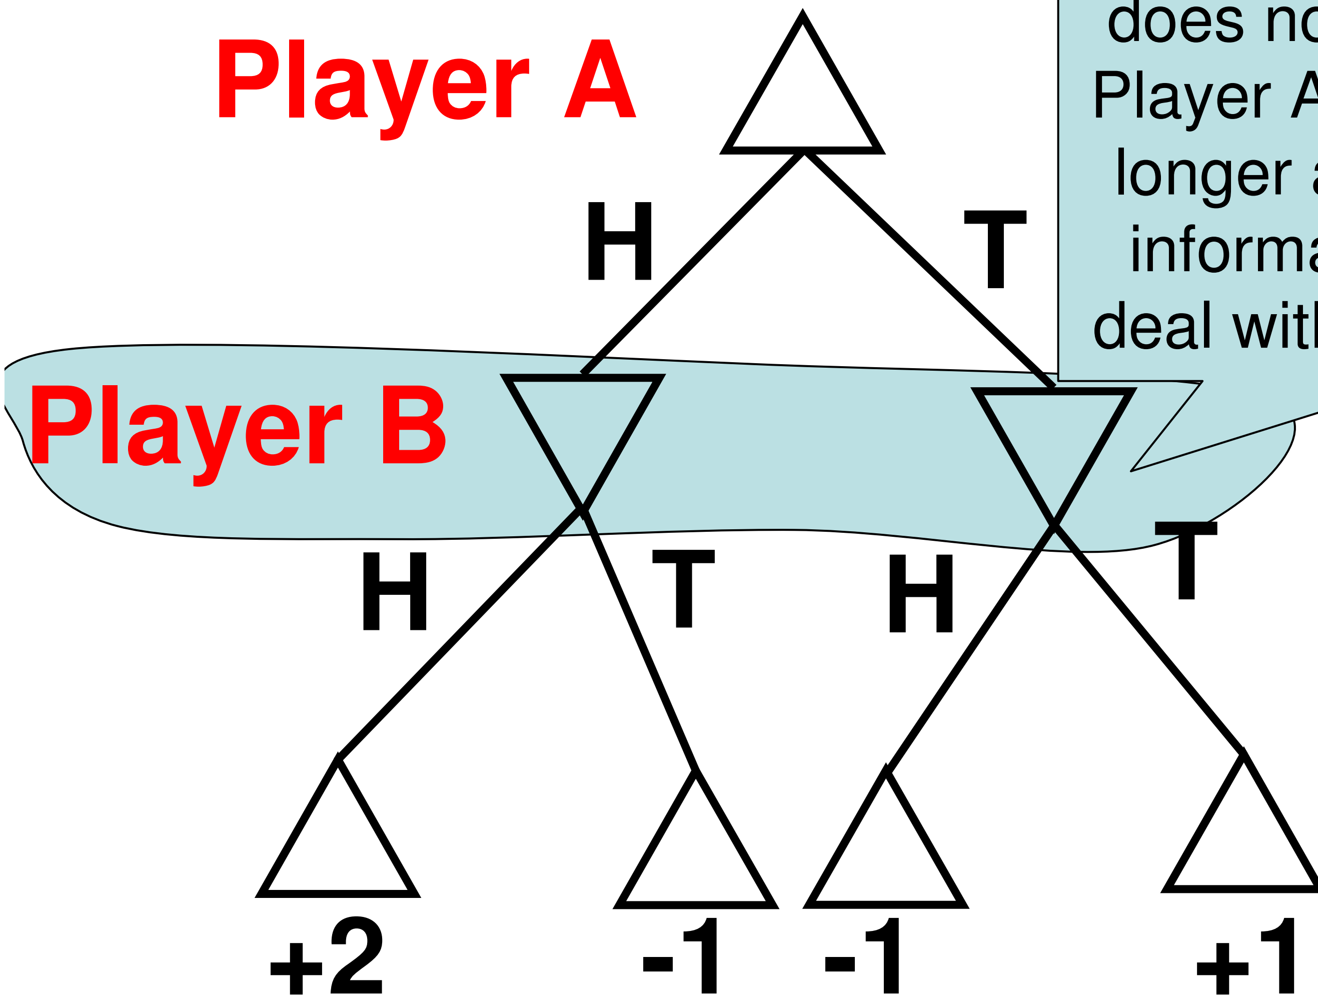

Player B

Player A

	H	T
H	+2	-1
T	-1	+1

Matrix Normal Form

Player B

Player A

	H	T
H	**+2**	**-1**
T	**-1**	**+1**

- It is no longer the case that maximin = minimax (easy to verify: -1 vs. +1)
- Therefore: It appears that there is no pure strategy solution
- In fact, in general, *none of the pure strategies* are solutions to a zero-sum game with *hidden information*

Mixed strategy

Player B

Player A

	H	T
H	**+2**	**-1**
T	**-1**	**+1**

- Solution:
  - Player A chooses
    - strategy **H** with probability p
    - Strategy **H** with probability $1-p$
  - Optimal: $p = 0.4$

Minimax with Mixed Strategies

- Theorem II (von Neumann):
  - For a two-player, zero-sum game with hidden information:
    - *An optimal mixed strategy always exists*

- *Where the matrix form of the game is:*

m_{11}	m_{12}
m_{21}	m_{22}

- *Note: This is a direct generalization of the minimax result to mixed strategies.*

Non-Zero Sum Games

R&N Section 17.6

Matrix Form of Zero-Sum Games

m_{11}	m_{12}
m_{21}	m_{22}

m_{ij}= Player A's payoff if Player A follows pure strategy *i* and Player B follows pure strategy *j*

Prisoner's Dilemma

- Two persons (A and B) are arrested with enough evidence for a minor crime, but not enough for a major crime.
- If they *both* confess to the crime, they each know that they will serve 5 years in prison.
- If only one of them testify, he will go free and the other prisoner will serve 10 years.
- If neither of them confess, they'll each spend 1 year in prison

Matrix normal form for non-zero-sum games

Player B

Player A

	Testify	**Refuse**
Testify	**-5,-5**	**0,-10**
Refuse	**-10,0**	**-1,-1**

Matrix normal form for non-zero-sum games

Player B

Player A

	Testify	**Refuse**
Testify	**-5,-5**	**0,-10**
Refuse	**10,0**	

Player B's payoff for the pair of strategies A: Testify, B: Refuse

Player A's payoff for the pair of strategies A:Testify, B: Testify

Why this example?

- Although simple, this example models a huge variety of situations in which participants have similar rewards as in this game.
- *Joint work*: Two persons are working on a project. Each person can choose to either work hard or rest. If A works hard then prefers to rest, but the outcome of both working is preferable to the outcome of both resting (the project does not get done).
- *Duopoly*: Two firms compete for producing the same product and they both try to maximize profit. They can set two prices, “High” and “Low”. If both firms choose High, they both make a profit of $1000. If they both choose Low, they both make a lower profit of $600. Otherwise, the High firm makes a profit of $1200 and the Low firm takes a loss of $200.
- *Arms race*, *Robot detection*, *Use of common property*……

Matrix normal form for non-zero-sum games

Player B

Player A

	Testify	**Refuse**
Testify	**-5,-5**	**0,-10**
Refuse	**-10,0**	**-1,-1**

- This not a zero-sum game → The interests (payoffs) of the “players” are no longer opposite of each other
- What is the best strategy to follow for each player, assuming that they are both rational

Player B

Player A

	Testify	Refuse
Testify	-5,-5	~~0,-10~~
~~Refuse~~	~~-10,0~~	~~-1,-1~~

Dominant Strategies

Player A \ Player B	Testify	Refuse
Testify	-5,-5	0,-10
~~Refuse~~	~~-10,0~~	~~-1,-1~~

- Player A's payoff is greater if he testifies than if he refuses, no matter what strategy B chooses
- Therefore Player A does not need to consider strategy "refuse" since it cannot possibly yield a higher payoff

Dominant Strategies

Player B

Player A

	Testify	~~Refuse~~
Testify	**-5,-5**	~~0,-10~~
~~Refuse~~	~~-10,0~~	~~-1,-1~~

- The same reasoning can be applied to Player B:
  - Player B's payoff is greater if he testifies than if he refuses, no matter what strategy A chooses
  - Therefore Player B does not need to consider strategy "refuse" since it cannot possibly yield a higher payoff

Dominant Strategies

Player B

Player A

	Testify	Refuse
Testify	**-5,-5**	**0,-10**
Refuse	**-10,0**	**-1,-1**

- We say that a strategy *strictly dominates* if it yields a higher payoff than any other strategy for *every* one of the possible actions of the other player.
- Key result → If both players have strictly dominating strategies, they provide a solution for the game (i.e., predict the outcome of the game) → a *dominant strategy equilibrium*
  - Testify is a strictly dominant strategy for A
  - Testify is a strictly dominant strategy for B
  - Therefore (Testify, Testify) is the solution

	IB	IIB	IIIB
IA	-1,6	6,-1	5,4
IIA	6,-1	-1,6	5,4
IIIA	4,5	4,5	7,7

How would the two players play this game?

	IB	IIB	IIIB
IA	**-1,6**	**6,-1**	**5,4**
IIA	**6,-1**	**-1,6**	**5,4**
IIIA	**4,5**	**4,5**	**7,7**

$u_A(\mathrm{IIIA}, \mathrm{IIIB}) \geq u_A(\mathrm{X}, \mathrm{IIIB})$ For any strategy X of Player A

$u_B(\mathrm{IIIA}, \mathrm{IIIB}) \geq u_B(\mathrm{IIIA}, \mathrm{Y})$ For any strategy Y of Player B

(IIIA,IIIB) is an equilibrium because:

- Player A cannot find a better strategy given that Player B uses strategy IIIB
- Conversely, Player B cannot find a better strategy given that Player A uses strategy IIIA

Side Note: More than 2 Players?

- The formalism extends directly to more than 2 players.
- If we have *n* players, we need to define *n* payoff functions u_i, i=1,..,n.
- Payoff function u_i maps a tuple of *n* strategies to the corresponding payoff for player *i*
- $u_i(s_1,..,s_n)$ = payoff for player *i* if players 1,..,n use pure strategy $s_1,\ldots,s_n$.
- Everything else (definition of dominating strategies, etc. remains the same)

More formal definition

- A tuple of pure strategies $(s_1^*, s_2^*, .., s_n^*)$ is a pure equilibrium if, for all *i*'s:

$$u_i\left(s_1^*, \cdots, s_{i-1}^*, s_i, s_{i+1}^*, \cdots, s_n^*\right) \leq u_i\left(s_1^*, \cdots, s_{i-1}^*, s_i^*, s_{i+1}^*, \cdots, s_n^*\right)$$

for any strategy s_i.

- In words: Player *i* cannot find a better strategy than s_i^* if the other player use the remaining strategies in the equilibrium
- Technically, called a *pure Nash Equilibrium (NE)*

More formal definition (equivalent)

- A tuple of pure strategies (s_1*,s_2*,..,s_n*) is a pure equilibrium if, for all *i*'s:

$$s_i^* = \arg\max_{s_i} u_i\left(s_1^*, \cdots, s_{i-1}^*, s_i\ , s_{i+1}^*, \cdots, s_n^*\right)$$

- In words: Player *i* cannot find a better strategy than s_i* if the other player use the remaining strategies in the equilibrium
- Technically, called a pure Nash Equilibrium (NE)

Player B

Player A	Testify	Refuse
Testify	-5,-5	0,-10
Refuse	-10,0	-1,-1

- Rationality
- No cooperation

Examples and Question

- So, we've generalized our concepts for solving games to non zero-sum games → NEs
- Basic questions:
  - Is there always a NE?
  - Is it unique?

Example with multiple NEs

	Left	**Right**
Left	**+1,+1**	**-1,-1**
Right	**-1,-1**	**+1,+1**

- Two vehicles are driving toward each other; they have 2 choices: Move right or move left.
- Why is having multiple NEs a problem?

Example with multiple NEs

	Hockey	**Movie**
Hockey	**+2,+1**	**0,0**
Movie	**0,0**	**+1,+2**

- Two friends have different tastes, A likes to watch hockey games but B prefers to go see a movie. Neither likes to go to his preferred choice alone; each would rather go the other's preferred choice rather than go alone to its own.

Example with no pure NE

	I	II
I	0,1	1,0
II	1,0	0,1

- Even very simple games may not have a pure strategy equilibrium → This is not surprising since we saw earlier that we had a similar problem with zero-sum games, which did not necessarily have a pure strategy solution
- Solution: Same trick as with zero-sum games → Allow the players to randomize and to use *mixed* strategies

Mixed Strategy Equilibrium

- The concept of equilibrium can be extended to mixed strategies.
- In that case, a mixed strategy for each player i is a vector of probabilities $\boldsymbol{p}_{\mathrm{i}} = (p_{ij})$, such that player i chooses pure strategy j with probability p_{ij}
- A set of mixed strategies $(\boldsymbol{p}^*_1,..,\boldsymbol{p}^*_n)$ if player i (for any i) gets a lower payoff by changing $\boldsymbol{p}^*_{\mathrm{i}}$ to any other mixed strategy $\boldsymbol{p}_{\mathrm{i}}$

Example

	Hockey	Movie
Hockey	+2,+1	0,0
Movie	0,0	+1,+2

- An example mixed strategy is:
  - A chooses Hockey with probability: $p = 2/3$
  - B chooses Hockey with probability: $q = 1/3$
- In fact, this is a mixed strategy equilibrium for this game
- The expected payoff is 2/3 for both A and B

Summary

- Matrix form of non-zero-sum games and basic concepts for those games
- Strict dominance and its use
- Definition of game equilibrium
- Key result: Existence of (possibly mixed) equilibrium for any finite game
- Understand the difference between cooperating and non-cooperating situations
- Continuous games and corresponding recipes