		Player 2		
		ℓ	m	r
Player 1	U	5, 3	0, 4	3, 5
	M	4, 0	5, 5	4, 0
	D	3, 5	0, 4	5, 3

Gambit

GAMUT

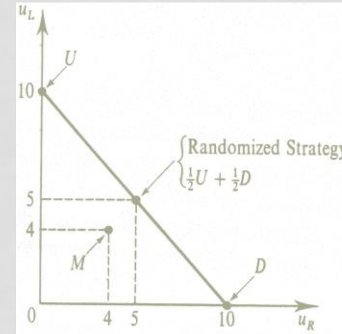

INTRODUCTION TO MULTI-AGENT DECISION MAKING: GAME THEORY

PRITHVIRAJ (RAJ) DASGUPTA

NAVAL RESEARCH LABORATORY, WASHINGTON, DC

1

TJHSST AI/ML Club
Lecture 1

OUTLINE

- Brief history of game theory
- Normal form games
- Main concepts:
  - Actions, strategies
  - Payoff, payoff matrix
- Finding game outcomes
  - Best response, Nash equilibrium
    - Pure and mixed strategies
- Game theory software: MCLP, Gambit, Gamut, NashPy
- Additional Topics:
  - Strictly dominated strategies
  - Weakly dominated strategies
  - Iterated removal of dominated strategies

HISTORY OF GAME THEORY

- 19th century and earlier: mathematical formulations to solve taxation problems, profits, etc.
- First half of 20th century:
  - Von Neumann formalizes utility theory, lays down mathematical foundations for analyzing two player games; early work starts
- 1950s onwards: Nash theorem, analysis of different types of games, beyond two players, relaxing simplifying assumptions (e.g., complete knowledge), more complex settings
- 1970s onwards: evolutionary game theory (applying biological concepts in games), learning in games, mechanism design
- 1990s onwards:
  - computational implementation of game theory algorithms, complexity results (n players),
  - game theory software, programming competitions,
  - applications to real-life domains (auctions, network bandwidth sharing, resource allocation, etc.)

SOME NOTABLE GAME THEORISTS

- John Von Neumann – Founder of field
- John Nash – Nash equilibrium
- John Harsanyi – Incomplete Information (Bayesian Games)
- Roger Myerson – Mechanism Design
- Many others: Morgenstern, Selten, Maynard Smith, Aumann…

Von Neumann

Nash

Harsanyi

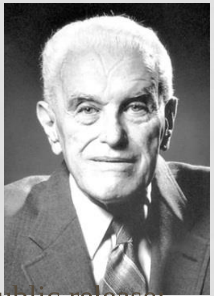

Myerson

SIMPLE 2-PLAYER GAME

OUTLINE

- Brief history of game theory
- Normal form games
- Main concepts:
  - Actions, strategies
  - Payoff, payoff matrix
- Finding game outcomes
  - Best response, Nash equilibrium
    - Pure and mixed strategies
- Game theory software: MCLP, Gambit, Gamut, NashPy
- Additional Topics:
  - Strictly dominated strategies
  - Weakly dominated strategies
  - Iterated removal of dominated strategies

A SIMPLE GAME EXAMPLE

- Simplest case: 2 persons (called players)
- Objective: Cross the bridge
  - Bridge can hold only one person (player) at a time
  - Players cannot see each other from across the bridge
- Each player has a set of (2) actions: Stop or Go
- Each player has to select one action from its action set
- By doing action, the player gets a payoff or utility
  - Rationality assumption: Each player selects action that gives it highest payoff

Actions: Go, Stop

Payoff(Go) = 1
Payoff (Stop) = -1

Actions: Go, Stop

Payoff(Go) = 1
Payoff (Stop) = -1

A SIMPLE GAME EXAMPLE

- Simplest case: 2 persons (called players)
- Objective: Cross the bridge
  - Bridge can hold only one person (player) at a time
  - Players cannot see each other from across the bridge
- Each player has a set of (2) actions: Stop or Go

- Each player has to select one action from its action set
- By doing action, the player gets a payoff or utility
  - Rationality assumption: Each player selects action that gives it highest payoff
- A player's decision (selected action) affects the other player's decision (selected action)
  - And other player's payoff, and in turn its own payoff

Actions: Go, Stop

Joint Actions

Payoff(Go, Go) = -2, -2
Payoff (Go, Stop) = 2, -1
Payoff (Stop, Go) = -1, 2
Payoff (Stop, Stop) = -1, -1

Actions: Go, Stop

Joint Utilities

Payoff(Go, Go) = -2, -2
Payoff (Go, Stop) = 2, -1
Payoff (Stop, Go) = -1, 2
Payoff (Stop, Stop) = -1, -1

A SIMPLE GAME EXAMPLE

- Simplest case: 2 persons (called players)
- Objective: Cross the bridge
  - Bridge can hold only one person (player) at a time
  - Players cannot see each other from across the bridge
- Each player has a set of (2) actions: Stop or Go

- Each player has to select one action from its action set
- By doing action, the player gets a payoff or utility
  - Rationality assumption: Each player selects action that gives it highest payoff
- A player's decision (selected action) affects the other player's decision (selected action)
  - And other player's payoff, and in turn its own payoff

Actions: Go, Stop

	Go	Stop
Go	-2, -2	2, -1
Stop	-1, 2	-1, -1

Actions: Go, Stop

	Go	Stop
Go	-2, -2	2, -1
Stop	-1, 2	-1, -1

THE MAIN PROBLEM IN GAME THEORY…SAID SIMPLY

	Go	Stop
Go	-2, -2	2, -1
Stop	-1, 2	-1, -1

- Main Problem: Given that each player knows the actions and payoffs of each other, how can each player individually make a decision (select an action) that will be 'best' for itself
- 'best' is loosely defined
  - Minimize regret
  - Maximize sum of payoffs to all (both) players
  - Stable or equilibrium action – if player deviates by itself from that action, it will end up lowering its own payoff … should hold for every player

Nash Equilibrium

MULTI-AGENT DECISION MAKING

- For solving game theory computationally, players are implemented as software agents
- Main Problem (restated in terms of agents): When two agents interact, what action should each agent take?
- Simplifying (natural) assumption: agents are self-interested or behave rationally
  - each agent takes an action that maximizes its own benefit

GAME DEFINITION

- Normal form game G = (N, A, U)
- N: set of players
  - For a 2-player game, N = {P1, P2}
  - Convention: P1: row player, P2: column player
- A: set of actions
  - For a 2-player game:
    - A_1: set of actions of P1, e.g., (go, stop)
    - A_2: set of actions of P2, e.g., (go, stop)
    - $A = A_1 \times A_2$ = (go, go), (go, stop), (stop, go), (stop, stop)
- U: set of utilities
  - For a 2 player game:
    - $U_1(a_1, a_2)$: utility to player 1 from joint action (a_1, a_2), e.g., U_1(go, stop) = 2
    - $U_2(a_1, a_2)$: utility to player 2 from joint action (a_1, a_2), e.g., U_2(go, stop) = -1
- Normal form games are usually represented (on paper) using a utility or payoff matrix

P2

P1	Go	Stop
Go	-2, -2	2, -1
Stop	-1, 2	-1, -1

Payoff Matrix

GAME TERMINOLOGY

- Game is
  - Single-shot: only played once (not repeated)
  - Simultaneous move: players move (select action) at the same time
    - Not sequential – each player cannot see opponent's move before selecting own move
- Each (joint) action results in a different outcome
  - Selecting a joint action ≡ selecting an outcome

Outcomes

	Go	Stop
Go	-2, -2	2, -1
Stop	-1, 2	-1, -1

NASH EQUILIBRIUM

OUTLINE

- Brief history of game theory
- Normal form games
- Main concepts:
  - Actions, strategies
  - Payoff, payoff matrix
- Finding game outcomes
  - Best response, Nash equilibrium
    - Pure and mixed strategies
- Game theory software: MCLP, Gambit, Gamut, NashPy
- Additional Topics:
  - Strictly dominated strategies
  - Weakly dominated strategies
  - Iterated removal of dominated strategies

PRISONER'S DILEMMA GAME

- Two persons charged with the same crime and held in separate cells. They have no way of communicating with each other, or, making an agreement beforehand. They are told (at the same time) that:
  - if one of them confesses the crime and the other does not, the confessor will be freed while the other person will get a term of 3 years
  - if both confess the crime, each will get term of 2 years
  - if neither confess the crime, each will get a 1-year term

PD GAME: PAYOFF MATRIX

Player 2

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

C: Confess
D: Do not confess

- If one of them confesses the crime and the other does not, the confessor will be freed while the other person will get a term of 3 years
- If both confess the crime, each will get term of 2 years
- If neither confess the crime, each will get a 1-year term

PD GAME REASONING

- P1 reasons:
  - I don't know what P2 is going to play, but I know P2 will play either C or D
    - If P2 plays C, I will get a payoff of –2 if I play C, and a payoff of -3 if I play D
    - If P2 plays D, I will get a payoff of 0 if I play C, and a payoff of -1 if I play D
  - Irrespective of what P2 plays, I'm better off by playing C
- P2 reasons in a similar manner (since the game is symmetric)
- Both end up playing (C, C)

P1 is calculating its best response to each possible action of P2

	C	D
C	-2, -2	0, -3
D	-3, 0	-1, -1

NASH EQUILIBRIUM

- A joint strategy (s_1^*, s_2^*) is an NE when each player is playing a <u>best response</u> (at that joint strategy)
- What this means:
  - When P1 plays s_1^*, P2 cannot improve its utility by playing another strategy besides s_2^*
    - Mathematically: $U_2(s_1^*, s_2^*) >= U_2(s_1^*, s_2')$, for all s_2' not equal to s_2^*
    - s_2^* is P2's best response to s_1
  - When P2 plays s_2^*, P1 cannot improve its utility by playing another strategy besides s_1^*
    - Mathematically: $U_1(s_1^*, s_2^*) >= U_1(s_1', s_2^*)$, for all s_1' not equal to s_1^*
    - s_1^* is P1's best response to s_2^*
- Neither agent has any incentive to deviate from a NE: <u>NE is a stable outcome of the game</u>

EXAMPLE: PRISONER'S DILEMMA

		P2	
		D	C
P1	D	-2,-2	-10,-1
	C	-1,-10	-5,-5

- Utility values are changed from last example (but their relative values are the same)
- Order or rows/columns are interchanged
- Outcome of the game still remains same (C, C)
- <u>Representation of the game does not change its outcome as long as relative utility values remain same</u>

NASH EQUILIBRIUM CHECK

	C	D
C	-2, -2	0, -3
D	-3, 0	-1, -1

- Is (C, C) a NE of the game?
- When P2 is playing C, P1 can do no better than play C
  - C is a best response of P1 when P2 plays C
- When P1 is playing C, P2 can do no better than play C
  - C is a best response of P2 when P1 plays C
- (C, C) is a NE

- Is (D, D) a NE of the game?
- When P2 is playing D, P1 can do better by playing C
  - D is <u>not a best response</u> for P1 when P2 plays D
- When P1 is playing D, P2 can do better by playing C
  - D is <u>not a best response</u> for P2 when P1 plays D
- (D, D) is not a NE

EXAMPLE: NASH EQUILIBRIUM

Player 2

		ℓ	m	r
	U	5, 3	0, 4	3, 5
Player 1	M	4, 0	5, 5	4, 0
	D	3, 5	0, 4	5, 3

GAMES WITH MULTIPLE NASH EQUILIBRIUM

- Game of chicken (Earlier bridge crossing example)
- Stag-hunt

	Go	Stop
Go	-2, -2	2, -1
Stop	-1, 2	-1, -1

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

COMMON PAYOFF GAME

- Both (all) agents get equal utility in every outcome (e.g., two drivers coming from opposite sides have to choose which side of road to drive on)

		P2: L	P2: R
P1	L	1, 1	0, 0
P1	R	0, 0	1, 1

BATTLE OF SEXES GAME

- Husband prefers going to a Football game, which wife dislikes
- Wife prefers going to the Opera, which husband dislikes
- They like each other's company

		W	
		F	O
H	F	2, 1	0, 0
	O	0, 0	1, 2

MIXED STRATEGIES

ZERO-SUM GAME: MATCHING PENNIES

- Two coins are tossed at the same time
  - Both turn up same face (H or T): P1 wins \$1, P2 loses \$1
  - Each turns up with a different face: P1 loses \$1, P2 wins \$1

		P2: H	P2: T
P1	H	1, -1	-1, 1
P1	T	-1, 1	1, -1

Is there a Nash equilibrium in actions (pure strategies)?

STRATEGY: MIXED AND PURE

- Strategy is a probability distribution over actions
- For example:
  - P1 plays H w/ prob. 0.7 and T w/ prob. 0.3
  - P2 plays H w/ prob. 0.6 and T w/ prob. 0.4
- We call this action profile a mixed strategy
  - Written as ((0.7, 0.3) (0.6, 0.4))
- Action can be looked upon as a strategy as well
  - P1 playing action H ≡ P1 select H w/ prob 1 and T w/ prob 0
  - Called pure strategy

		P2	
		H	T
P1	H	1, -1	-1,1
	T	-1, 1	1, -1

Nash Theorem: Every finite game must have a Nash equilibrium in either pure or mixed strategies

MIXED STRATEGY NASH EQUILIBRIUM

- Idea:
  - I do not know what my opponent is going to play
  - My opponent selects one of their actions over another because the selected action gives a higher utility than the not selected action
  - But I can make my opponent not prefer one of its actions over another
    - by making the utility that they get from both actions equal
    - via choosing probabilities for my own actions
  - In other words, opponent will not have any incentive to deviate from one action to another – Nash equilibrium

Each player randomizes (defines a probability distribution) over its own actions in a way such that its opponents' choice becomes independent of its actions

SOLVING MIXED STRATEGY NASH EQUILIBRIUM (1)

P2

			H	T
P1	p	H	1, -1	-1, 1
	1-p	T	-1, 1	1, -1

- P1 tries to solve: What probability should I play H and T with so that my (expected) utility is independent of whether P2 plays H or T
- Suppose P1 plays H with probability 'p' and T with probability (1-p)
  - Utility to P2 for playing H (column H): -1.p + 1. (1-p) = 1-2p
  - Utility to P2 for playing T (column T): 1.p + (-1).(1-p) = 2p -1
- To find its mixed strategy, P1 solves for p to make P2's utility from H and T equal. That is:

$$2p - 1 = 1 - 2p,$$

$$\text{or, } p = 0.5$$

SOLVING MIXED STRATEGY NASH EQUILIBRIUM (2)

			0.5	0.5
		P2	q	1-q
			H	T
P1	0.5	H	1, -1	-1, 1
	0.5	T	-1, 1	1, -1

- P2 solves similarly denotes q as probability of playing H and (1-q) as probability for playing T
  - Solving for q in the same manner as for P1 gives q = 0.5
- Mixed strategy Nash equilibrium is:
  - ((0.5, 0.5) (0.5, 0.5))

EXAMPLE: MEETING IN NY GAME: MIXED STRATEGY NASH EQUILIBRIUM

		Mr. Schelling	
		Empire State	Grand Central
Mr. Thomas	Empire State	1000, 1000	0, 0
	Grand Central	0, 0	100, 100

EXAMPLE: MEETING IN NY GAME: MIXED STRATEGY NASH EQUILIBRIUM (2)

- T reasons:
  - Let me play ES with probability σ_T and GC with probability $(1- \sigma_T)$
  - Then,
    - Utility to S from ES: $1000\ \sigma_T + 0(1- \sigma_T)$
    - Utility to S from GC: $0\ \sigma_T + 100\ (1- \sigma_T)$
- Recall:
  - T will solve for σ_T that makes the utility that S gets from ES and GC equal
- Therefore,

$$1000\ \sigma_T + 0(1- \sigma_T) = 0\ \sigma_T + 100\ (1- \sigma_T)$$

$$\text{or, } 1100\sigma_T = 100, \text{ or, } \sigma_T = 1/11$$

- By similar reasoning, S plays ES with probability σ_S = <u>?</u>
- And the mixed strategy NE of the game is <u>?</u>

BATTLE OF SEXES GAME

- Husband prefers going to a football game, which wife hates
- Wife prefers going to the opera, which husband hates
- They like each other's company

		W: F	O
h	F	2, 1	0, 0
	O	0, 0	1, 2

BATTLE OF THE SEXES MIXED STRATEGY

			w	
		F		O
h	F	2, 1		0, 0
	O	0, 0		1, 2

- Suppose husband's probability of playing F is p and for playing O is 1-p
- Husband calculates
  - Utility to wife when she plays F is: 1.p + 0 .(1-p)
  - Utility to wife when she plays O is: 0.p + 2 .(1-p)
- Husband will select (or solve for) p in a way such that above two utilities to wife are the same.
- That is, 1p + 0 (1-p) = 0p + 2 (1-p), or, p = 2/3

BATTLE OF THE SEXES MIXED STRATEGY

W

h		F	O
	F	2, 1	0, 0
	O	0, 0	1, 2

- Similarly, suppose wife's probability of playing F is q and for playing O is 1-q
- Wife calculates
  - Utility to husband when he plays F is: 2q + 0 (1-q)
  - Utility to husband when he plays O is: 0q + 1 (1-q)
- Husband will select (or solve for) p in a way such that above two utilities to wife are the same.
- That is, 2q + 0 (1-q) = 0q + 1 (1-q), or, q = 1/3

BATTLE OF THE SEXES MIXED STRATEGY

		W	
		F	O
h	F	2, 1	0, 0
	O	0, 0	1, 2

- So, the mixed strategy NE of the game is ((2/3, 1/3) (1/3, 2/3))
- Denotes player 1 (husband) plays his two actions with probability (2/3, 1/3)
- Similarly for player 2 (wife) plays her two actions with probability (1/3, 2/3)

ROCK PAPER SCISSORS

- Can you find the pure and mixed strategy equilibrium of the classic rock paper scissors game?
- The payoff matrix is given below: +1 means ‘beats’, -1 means ‘is beaten by’, ‘0’ mean ‘draw’

	Rock	**Paper**	**Scissors**
Rock	0, 0	-1, 1	1, -1
Paper	1, -1	0, 0	-1, 1
Scissors	-1, 1	1, -1	0, 0

GAME THEORY SOFTWARE

OUTLINE

- Brief history of game theory
- Normal form games
- Main concepts:
  - Actions, strategies
  - Payoff, payoff matrix
- Finding game outcomes
  - Best response, Nash equilibrium
    - Pure and mixed strategies
- Game theory software: Gambit, Gamut, NashPy
- Additional Topics:
  - Strictly dominated strategies
  - Weakly dominated strategies
  - Iterated removal of dominated strategies

SOLVING FOR NASH EQUILIBRIUM

- Nash equilibrium can be solved as a Mixed Integer Linear Program (MILP)
- MILP is an optimization problem with linear constraints
  - One of the constraint variables can take only integer values
- Gurobi and IBM CPLEX: very fast MILP solvers (licensed, but free for academic use)
- Open source libraries: COIN-OR and its Python interface CyLP
  - Coin-OR:https://www.coin-or.org/
  - CyLP: https://github.com/coin-or/CyLP

GAMUT

- Software to generate games
  - Vary number of players, actions,
  - Different games – restricted to set of well-known games like Prisonsers Dilemma, Battle of Sexes, Common Payoff, etc.
- http://gamut.Stanford.edu
- java -jar gamut.jar -g PrisonersDilemma –f PD.out

GAMUT PARAMETERS

- -g <name of game>
- -f <name of output file>
- -players <n>: no. of players
- -action <n>: no. of actions (each player gets same no. of actions)
- Other parameters that can be set
  - min_payoff, max_payoff, int_payoffs, normalize
- See full documentation in User's Guide http://gamut.stanford.edu/userdoc.pdf

GAMBIT

- Software to specify and solve games
  - http://gambit.sourceforge.net/

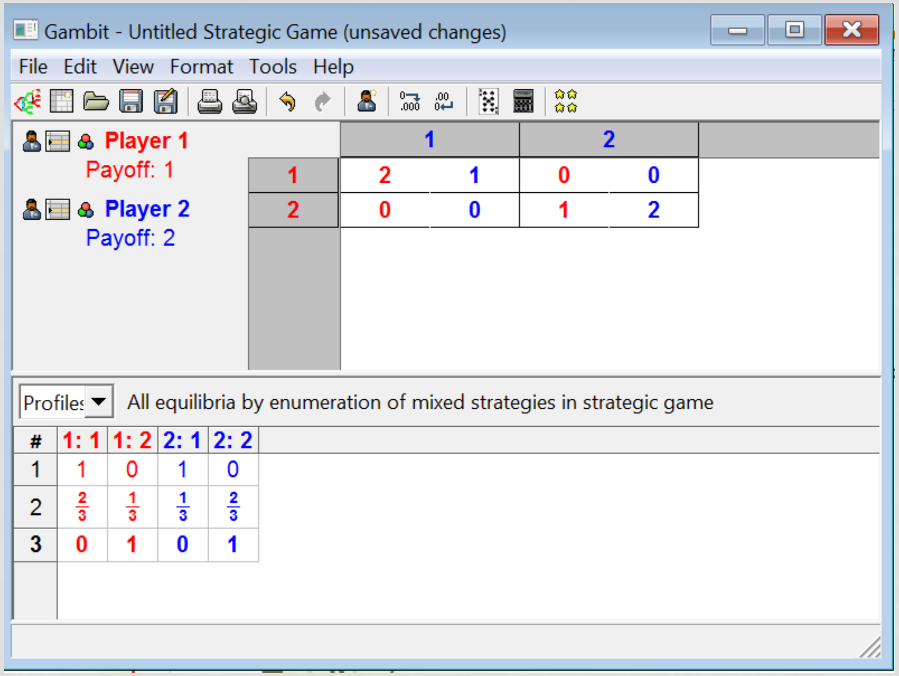

NASHPY

- Python library to create normal form games and solve for Nash equilibrium
- https://nashpy.readthedocs.io/en/latest/

ADDITIONAL RESOURCES

- Textbook on Game Theory (written by computer scientists)
  - Y. Shoham and K. Leyton-Brown, "Multi-agent Systems: Algorithmic, Game-Theoretic and Logical Foundations" (advanced undergrad), online: http://www.masfoundations.org/mas.pdf
- Many textbooks on introductory game theory (written by economists):
  - Osborne and Rubenstein "A Course in Game Theory" (advanced undergrad, grad)
  - Binmore: "Playing for Real" (advanced undergrad, grad), "Fun and Games" (undergrad), "A Very Short Introduction to Game Theory" (undergrad)
  - Gibbons, "Game Theory for Applied Economists" (advanced undergrad)
  - Tadelis "Game Theory: An Introduction" (advanced undergrad)

ADDITIONAL INFORMATION SLIDES

OUTLINE

- Strictly Dominated Strategies
  - Removal of strictly dominated strategies
  - Iterated removal of strictly dominated strategies
- Weakly Dominated Strategies

STRICTLY DOMINATED STRATEGIES

- Strictly Dominated Strategy: an action that can be removed by a player without considering best response
- In example game:
  - $U_1\ (C, C) > U_1(D, C)$
  - $U_1\ (C, D) > U_1(D, D)$
  - That is, $U_1(C, *) > U_1\ (D, *)$: makes D a <u>strictly dominated strategy</u> for P1
- P1 can eliminate D without considering best response
- Similarly, P2 can eliminate D
- (C, C) remains as only outcome of game

	C	D
C	4,4	4,1
D	1,4	1,1

<u>Remember</u>: First each player should remove strictly dominated strategies and then calculate best responses for Nash eq. (as required)

EXAMPLE: STRICT DOMINANCE

P2

		L	R
P1	U	1, -1	-1, 1
	M	-1, 1	1, -1
	D	-2, 5	-3, 2

- Note: All games might not have strictly dominated strategies (e.g., prisoners dilemma does not have)

ITERATED DOMINANCE: MODIFIED PRISONERS DILEMMA

- Recall: Given a game, to solve it, first remove all strictly dominated strategies
- Problem: All games might not have strictly dominated strategies
- Try iterated removal of dominated strategies

	C	D
C	-2, -2	0, -3
D	-3, 0	0, -1

- P1's best response is not D any more
- P1 can calculate that P2's best response is C when P1 plays C or D
- P1 can reason that P2 will never play D
- With D removed for P2 (column D removed), P1's best response is C and game outcome is (C, C)

WEAKLY DOMINATED STRATEGY

P2

		L	R
P1	U	5,1	4,0
	M	6,0	3,1
	D	6,4	4,4

- Weak dominance: when a player's utility from one action is greater than <u>or equal to</u> its utility from another action
    - Recall: Strong dominance requires <u>strict</u> greater than
- Example: for player 1, U and M are weakly dominated by D

DEFINITION: WEAKLY DOMINATED STRATEGY

- A strategy $s_i \in S_i$ is weakly dominated for player I in game (N, $\{A_i\}$, $\{u_i(\cdot)\}$) if there exists another strategy $s_i' \in S_i$ such that for all $s_{-i} \in S_{-i}$

$$u_i(s_i', s_{-i}) >= u_i(s_i, s_{-i})$$

- with strict inequality (>) for some s_{-i}.
- In this case, we say that strategy s_i' weakly dominates strategy s_i
- A strategy is a *weakly dominant strategy* for player I in game (N, $\{A_i\}$, $\{u_i(\cdot)\}$) if it weakly dominates every other strategy in S_i

BACKUP SLIDES

FORMAL DEFINITIONS

Definition 3.3.3 (Best response) *Player i's* best response *to the strategy profile s_{-i} is a mixed strategy $s_i^* \in S_i$ such that $u_i(s_i^*, s_{-i}) \geq u_i(s_i, s_{-i})$ for all strategies $s_i \in S_i$.*

Definition 3.3.4 (Nash equilibrium) *A strategy profile $s = (s_1, \ldots, s_n)$ is a* Nash equilibrium *if, for all agents i, s_i is a best response to s_{-i}.*