A4M33MAS - Multiagent Systems
Introduction to Game Theory

Michal Pechoucek & Michal Jakob

Department of Computer Science
Czech Technical University in Prague

OTEVŘENÁ INFORMATIKA

In parts based on Kevin Leyton-Brown: Foundations of Multiagent Systems an introduction to algorithmic game theory, mechanism design and auctions

Game Theory

01

- Game theory is the study of strategic decision making, the study of mathematical models of conflict and cooperation between intelligent rational decision-makers, interactive decision theory
- Given the *rule of the game*, game theory studies strategic behaviour of the agents in the form of a mixed/pure strategy (e.g. optimality, stability)
- Given the *strategic behavior of the agents*, mechanism design (reverse game theory) studies(designs) the rule of games with respect to a specific outcome of the game

Yoav Shoham, Kevin Leyton-Brown, Multiagent Systems: ***Algorithmic, Game-Theoretic, and Logical Foundations*** **Cambridge University Press, 2009**

http://www.masfoundations.org

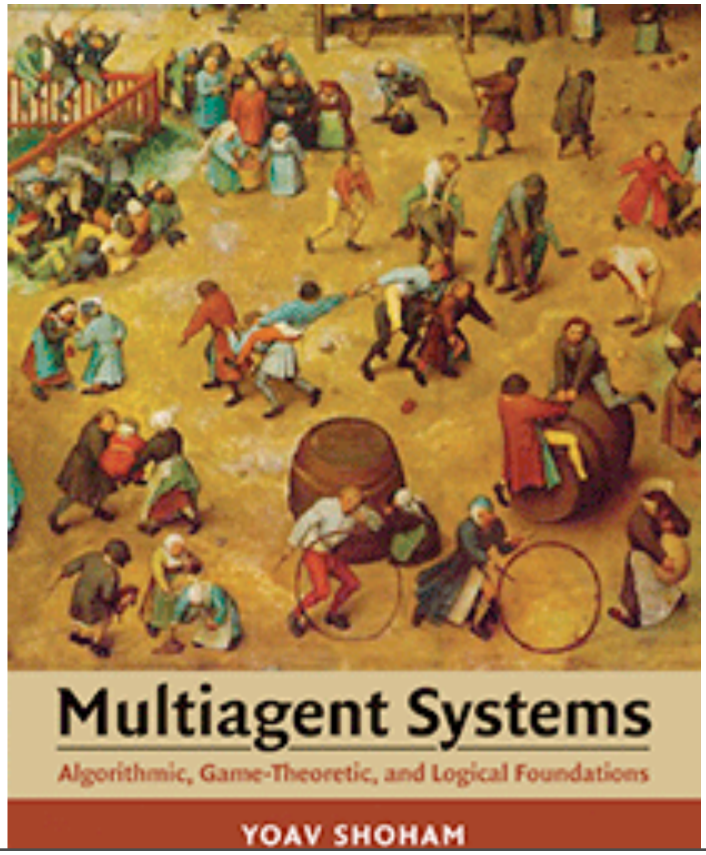

Types of Games

- Cooperative or non-cooperative

Types of Games

- Cooperative or non-cooperative
- Symmetric and asymmetric

Types of Games

- Cooperative or non-cooperative
- Symmetric and asymmetric
- Zero-sum and non-zero-sum

Types of Games

- Cooperative or non-cooperative
- Symmetric and asymmetric
- Zero-sum and non-zero-sum
- Simultaneous and sequential

Types of Games

- Cooperative or non-cooperative
- Symmetric and asymmetric
- Zero-sum and non-zero-sum
- Simultaneous and sequential
- Perfect information and imperfect information (complete info. games)

Types of Games

- Cooperative or non-cooperative
- Symmetric and asymmetric
- Zero-sum and non-zero-sum
- Simultaneous and sequential
- Perfect information and imperfect information (complete info. games)
- Combinatorial games

Types of Games

- Cooperative or non-cooperative
- Symmetric and asymmetric
- Zero-sum and non-zero-sum
- Simultaneous and sequential
- Perfect information and imperfect information (complete info. games)
- Combinatorial games
- Infinitely long games

Types of Games

- Cooperative or non-cooperative
- Symmetric and asymmetric
- Zero-sum and non-zero-sum
- Simultaneous and sequential
- Perfect information and imperfect information (complete info. games)
- Combinatorial games
- Infinitely long games
- Discrete and continuous games, differential games

TCP Backoff Game

TCP Backoff Game

- Consider this situation as a two-player game:
  - both use a correct implementation: both get 1 ms delay
  - one correct, one defective: 4 ms delay for correct, 0 ms for defective
  - both defective: both get a 3 ms delay.

TCP Backoff Game

- Consider this situation as a two-player game:
  - both use a correct implementation: both get 1 ms delay
  - one correct, one defective: 4 ms delay for correct, 0 ms for defective
  - both defective: both get a 3 ms delay.

- Questions:
  - What action should a player of the game take?
  - Would all users behave the same in this scenario?
  - What global patterns of behaviour should the system designer expect?
  - Under what changes to the delay numbers would behavior be the same?
  - What effect would communication have?
  - Repetitions? (finite? infinite?)
  - Does it matter if I believe that my opponent is rational?

Game definition

- Finite, n-person game: $\langle N, A, u \rangle$:
    - N is a finite set of n players, indexed by i
    - $A = A_1 \times \ldots \times A_n$, where A_i is the action set for player i
        - $a \in A$ is an action profile, and so A is the space of action profiles
    - $u = \langle u_1, \ldots, u_n \rangle$, a utility function for each player, where $u_i : A \mapsto \mathbb{R}$

- Writing a 2-player game as a matrix:
    - row player is player 1, column player is player 2
    - rows are actions $a \in A_1$, columns are $a' \in A_2$
    - cells are outcomes, written as a tuple of utility values for each player

Game definition

- Writing a 2-player game as a matrix:
  - row player is player 1, column player is player 2
  - rows are actions $a \in A_1$, columns are $a' \in A_2$
  - cells are outcomes, written as a tuple of utility values for each player

Game definition

- Writing a 2-player game as a matrix:
  - row player is player 1, column player is player 2
  - rows are actions $a \in A_1$, columns are $a' \in A_2$
  - cells are outcomes, written as a tuple of utility values for each player

Game definition

- Writing a 2-player game as a matrix:
  - row player is player 1, column player is player 2
  - rows are actions $a \in A_1$, columns are $a' \in A_2$
  - cells are outcomes, written as a tuple of utility values for each player

	C	*D*
C	$-1, -1$	$-4, 0$
D	$0, -4$	$-3, -3$

Other Games: Coordination Games

	Left	Right
Left	1	0
Right	0	1

driving side

	B	*F*
B	$2, 1$	$0, 0$
F	$0, 0$	$1, 2$

battle of sexes

Other Games: Coordination Games

Players have exactly the same interests.

- no conflict: all players want the same things
- $\forall a \in A, \forall i, j,\ u_i(a) = u_j(a)$
- we often write such games with a single payoff per cell
- why are such games "noncooperative"?

Other Games: Prisoners Dilemma

	B_C	B_D
A_C	1 ,1	5, 0
A_D	0, 5	3, 3

$$(A_d, B_c)^0 \preccurlyeq (A_c, B_c)^1 \preccurlyeq (A_d, B_d)^3 \preccurlyeq (A_c, B_d)^5$$

$$(A_c, B_d)^0 \preccurlyeq (A_c, B_c)^1 \preccurlyeq (A_d, B_d)^3 \preccurlyeq (A_d, B_c)^5$$

Other Games: Prisoners Dilemma

	B_C	B_D
A_C	a, a	b, c
A_D	c, b	d, d

any game where $c \succeq a \succeq d \succeq b$

Other Games: Matching Pennies

	Heads	Tails
Heads	1, -1	-1, 1
Tails	-1, 1	1, -1

Other Games: Matching Pennies

	Heads	Tails
Heads	1, -1	-1, 1
Tails	-1, 1	1, -1

Other Games: Matching Pennies

	Heads	Tails
Heads	1, -1	-1, 1
Tails	-1, 1	1, -1

	Heads	Tails
Heads	1	-1
Tails	-1	1

Other Games: Rock-paper-scissors

	Rock	Paper	Scissors
Rock	0	-1	1
Paper	1	0	-1
Scissors	-1	1	0

Properties of the games

- Finite, n-person game: $\langle N, A, u \rangle$:
  - N is a finite set of n players, indexed by i
  - $A = A_1 \times \ldots \times A_n$, where A_i is the action set for player i
    - $a \in A$ is an action profile, and so A is the space of action profiles
  - $u = \langle u_1, \ldots, u_n \rangle$, a utility function for each player, where $u_i : A \mapsto \mathbb{R}$

- strategy s_i refers to a decision (about action choice) at each stage of the game that the agent i makes and which leads to an outcome
- outcome is the set of possible states resulting from agent's decision making
- strategy profile refers to the set of strategies played by the agents. Set of strategy profiles: $S = S_1 \times \ldots \times S_n$.

Properties of the games

- Social welfare (Collective utility):

$$U(a) = \sum_{\forall i} u_i(a_i)$$

- Cooperative agents choose such a_i that maximizes $U(a)$
- Self-interested (*individually rational*) agents choose such a_i that maximizes $u_i(a_i)$

- When designing a multiagent system designers worry about:
  – individual rationality of each agent
  – social welfare and welfare efficiency
  – stability of the strategy (action) profile

Solution Concepts

- Pareto Efficiency
- Social welfare optimality
- Nash equilibrium
- Maxmin
- Dominant strategies
- Correlated equilibrium
- Minimax regret
- Stackelberg equilibrium
- Perfect equilibrium
- ϵ - Nash equilibrium

Solution Concepts

- Pareto Efficiency
- Social welfare optimality
- Nash equilibrium
- Maxmin
- Dominant strategies
- Correlated equilibrium
- Minimax regret
- Stackelberg equilibrium
- Perfect equilibrium
- ϵ - Nash equilibrium

Solution Concepts

- Pareto Efficiency
- Social welfare optimality
- Nash equilibrium
- Maxmin
- Dominant strategies
- Correlated equilibrium
- Minimax regret
- Stackelberg equilibrium
- Perfect equilibrium
- ϵ - Nash equilibrium

Pareto Efficiency

- Pareto Efficiency:
  - action (strategy) profile is Pareto optimal if there is no other action that at least one agent is better off and no other agent is worse off than in the given profile

- Dominance:
  - measure comparing two strategies. b dominates weakly a as follows:

$$a \preceq b \text{ iff } \forall i: \; u_i(a_i) \leq u_i(b_i)$$

  - dominant strategy: strategy that is not dominated by any other strategy

Pareto efficient strategy is such a strategy that is not weakly dominated by any other strategy

Pareto Efficiency

	C	*D*
C	-1, -1	-4, 0
D	0, -4	-3, -3

Pareto Efficiency

	C	*D*
C	-1, -1	-4, 0
D	0, -4	-3, -3

Pareto Efficiency

	C	D
C	-1, -1	-4, 0
D	0, -4	-3, -3

	Heads	Tails
Heads	1	-1
Tails	-1	1

Pareto Efficiency

	C	*D*
C	-1, -1	-4, 0
D	0, -4	-3, -3

	Heads	Tails
Heads	1	-1
Tails	-1	1

Pareto Efficiency

	C	*D*
C	-1, -1	-4, 0
D	0, -4	-3, -3

	Heads	Tails
Heads	1	-1
Tails	-1	1

	Left	Right
Left	1	0
Right	0	1

Pareto Efficiency

	C	*D*
C	-1, -1	-4, 0
D	0, -4	-3, -3

	Heads	Tails
Heads	1	-1
Tails	-1	1

	Left	Right
Left	1	0
Right	0	1

Pareto Efficiency

	C	*D*
C	-1, -1	-4, 0
D	0, -4	-3, -3

	Heads	Tails
Heads	1	-1
Tails	-1	1

	Left	Right
Left	1	0
Right	0	1

	B	*F*
B	2, 1	0, 0
F	0, 0	1, 2

Pareto Efficiency

	C	*D*
C	-1, -1	-4, 0
D	0, -4	-3, -3

	Heads	Tails
Heads	1	-1
Tails	-1	1

	Left	Right
Left	1	0
Right	0	1

	B	*F*
B	2, 1	0, 0
F	0, 0	1, 2

Solution Concepts

- Pareto Efficiency
- Social welfare optimality
- Nash equilibrium
- Maxmin
- Dominant strategies
- Correlated equilibrium
- Minimax regret
- Stackelberg equilibrium
- Perfect equilibrium
- ϵ - Nash equilibrium

Nash Equilibrium

01

- If you know what everyone else was going to do, it would easy to pick your own actions
- Let $a_i = \langle a_1, \ldots, a_{i-1}, a_{i+1}, \ldots, a_n \rangle$. now $a = (a_{-i}, a_i)$

Definition (Best Response)

$$a_i^* \in BR(a_{-i}) \text{ iff } \forall a_i \in A_i, u_i(a_i^*, a_{-i}) \geq u_i(a_i, a_{-i})$$

Definition (Nash Equilibrium)

The strategy profile $a = \langle a_1, \ldots, a_n \rangle$ is in Nash Equilibrium iff

$$\forall i, a_i \in BR(a_{-i})$$

Idea: look for stable action profiles.

Nash Equilibrium

Definition (Strict Nash Equilibrium)

The strategy profile $a = \langle a_1, \ldots, a_n \rangle$ is in Nash Equilibrium iff $\forall i, a_i \in BR(a_{-i})$ where $|BR(a_{-i})| = 1$

Definition (Weak Nash Equilibrium)

The strategy profile is in Weak NE iff it is not Strict NE

Nash Equilibrium

- Nash equilibrium, is a set of strategies, one for each player, such that no player has incentive to unilaterally change her action. Players are in equilibrium if a change in strategies by any one of them would lead that player to earn less than if she remained with her current strategy.

Definition (Strict Nash Equilibrium)

The strategy profile $a = \langle a_1, \ldots, a_n \rangle$ is in Nash Equilibrium iff $\forall i, a_i \in BR(a_{-i})$ where $|BR(a_{-i})| = 1$

Definition (Weak Nash Equilibrium)

The strategy profile is in Weak NE iff it is not Strict NE

Nash Equilibrium

- Nash equilibrium, is a set of strategies, one for each player, such that no player has incentive to unilaterally change her action. Players are in equilibrium if a change in strategies by any one of them would lead that player to earn less than if she remained with her current strategy.
- Strong Nash Equilibrium is such an equilibrium that is stable against deviations by cooperation.

Definition (Strict Nash Equilibrium)

The strategy profile $a = \langle a_1, \ldots, a_n \rangle$ is in Nash Equilibrium iff $\forall i, a_i \in BR(a_{-i})$ where $|BR(a_{-i})| = 1$

Definition (Weak Nash Equilibrium)

The strategy profile is in Weak NE iff it is not Strict NE

Nash Equilibrium

	C	*D*
C	-1, -1	-4, 0
D	0, -4	-3, -3

Nash Equilibrium

	C	*D*
C	-1, -1	-4, 0
D	0, -4	-3, -3

Nash Equilibrium

	C	*D*
C	-1, -1	-4, 0
D	0, -4	-3, -3

	Heads	Tails
Heads	1	-1
Tails	-1	1

Nash Equilibrium

	C	*D*
C	-1, -1	-4, 0
D	0, -4	-3, -3

	Heads	Tails
Heads	1	-1
Tails	-1	1

	Left	Right
Left	1	0
Right	0	1

Nash Equilibrium

	C	*D*
C	-1, -1	-4, 0
D	0, -4	-3, -3

	Heads	Tails
Heads	1	-1
Tails	-1	1

	Left	Right
Left	1	0
Right	0	1

Nash Equilibrium

	C	*D*
C	-1, -1	-4, 0
D	0, -4	-3, -3

	Heads	Tails
Heads	1	-1
Tails	-1	1

	Left	Right
Left	1	0
Right	0	1

	B	*F*
B	2, 1	0, 0
F	0, 0	1, 2

Nash Equilibrium

	C	*D*
C	-1, -1	-4, 0
D	0, -4	-3, -3

	Heads	Tails
Heads	1	-1
Tails	-1	1

	Left	Right
Left	1	0
Right	0	1

	B	*F*
B	2, 1	0, 0
F	0, 0	1, 2

Strong Nash Equilibrium

01

	C	*D*
C	-1, -1	-4, 0
D	0, -4	-3, -3

	Heads	Tails
Heads	1	-1
Tails	-1	1

	Left	Right
Left	1	0
Right	0	1

	B	*F*
B	2, 1	0, 0
F	0, 0	1, 2

Strong Nash Equilibrium

	C	*D*
C	-1, -1	-4, 0
D	0, -4	-3, -3

	Heads	Tails
Heads	1	-1
Tails	-1	1

	Left	Right
Left	1	0
Right	0	1

	B	*F*
B	2, 1	0, 0
F	0, 0	1, 2

Nash Equilibrium

- Nash equilibrium, is a set of strategies, one for each player, such that no player has incentive to unilaterally change her action. Players are in equilibrium if a change in strategies by any one of them would lead that player to earn less than if she remained with her current strategy.
- Strong Nash Equilibrium is such an equilibrium that is stable against deviations by cooperation.

Nash Equilibrium

- Nash equilibrium, is a set of strategies, one for each player, such that no player has incentive to unilaterally change her action. Players are in equilibrium if a change in strategies by any one of them would lead that player to earn less than if she remained with her current strategy.
- Strong Nash Equilibrium is such an equilibrium that is stable against deviations by cooperation.

Definition (Strict Nash Equilibrium)

The strategy profile $a = \langle a_1, \ldots, a_n \rangle$ is in Nash Equilibrium iff $\forall i, a_i \in BR(a_{-i})$ where $|BR(a_{-i})| = 1$

Definition (Weak Nash Equilibrium)

The strategy profile is in Weak NE iff it is not Strict NE

Nash Equilibrium

Definition (Strict Nash Equilibrium)

The strategy profile $a = \langle a_1, \ldots, a_n \rangle$ is in Nash Equilibrium iff $\forall i, a_i \in BR(a_{-i})$ where $|BR(a_{-i})| = 1$

Definition (Weak Nash Equilibrium)

The strategy profile is in Weak NE iff it is not Strict NE

Nash Equilibrium

Definition (Strict Nash Equilibrium)

The strategy profile $a = \langle a_1, \ldots, a_n \rangle$ is in Nash Equilibrium iff $\forall i, a_i \in BR(a_{-i})$ where $|BR(a_{-i})| = 1$

Definition (Weak Nash Equilibrium)

The strategy profile is in Weak NE iff it is not Strict NE

Nash Equilibrium

01

Definition (Strict Nash Equilibrium)

The strategy profile $a = \langle a_1, \ldots, a_n \rangle$ is in Nash Equilibrium iff $\forall i, a_i \in BR(a_{-i})$ where $|BR(a_{-i})| = 1$

Definition (Weak Nash Equilibrium)

The strategy profile is in Weak NE iff it is not Strict NE

- Weak NE are less stable than Strong NE.
  - in Weak NE there is at least 1 agent with > 1 BR, while 1 is in NE.Pure NE can be either Weak or Strong NE

Nash Equilibrium

Definition (Strict Nash Equilibrium)

The strategy profile $a = \langle a_1, \ldots, a_n \rangle$ is in Nash Equilibrium iff $\forall i, a_i \in BR(a_{-i})$ where $|BR(a_{-i})| = 1$

Definition (Weak Nash Equilibrium)

The strategy profile is in Weak NE iff it is not Strict NE

- Weak NE are less stable than Strong NE.
  - in Weak NE there is at least 1 agent with > 1 BR, while 1 is in NE.Pure NE can be either Weak or Strong NE
- Mixed NE are always Weak NE
  - if there are at least 2 pure strategies in BR, any combination of them is in NE

Nash Equilibrium

Definition (Strict Nash Equilibrium)

The strategy profile $a = \langle a_1, \ldots, a_n \rangle$ is in Nash Equilibrium iff $\forall i, a_i \in BR(a_{-i})$ where $|BR(a_{-i})| = 1$

Definition (Weak Nash Equilibrium)

The strategy profile is in Weak NE iff it is not Strict NE

- Weak NE are less stable than Strong NE.
  - in Weak NE there is at least 1 agent with > 1 BR, while 1 is in NE.Pure NE can be either Weak or Strong NE
- Mixed NE are always Weak NE
  - if there are at least 2 pure strategies in BR, any combination of them is in NE
- Strict NE are always Pure NE

Prisoners Dilemma: PE, NE

	B_C	B_D
A_C	1, 1	5, 0
A_D	0, 5	3, 3

$$\xi_A = (A_d, B_c)^0 \preccurlyeq (A_c, B_c)^1 \preccurlyeq (A_d, B_d)^3 \preccurlyeq (A_c, B_d)^5$$

$$\xi_B = (A_c, B_d)^0 \preccurlyeq (A_c, B_c)^1 \preccurlyeq (A_d, B_d)^3 \preccurlyeq (A_d, B_c)^5$$

Prisoners Dilemma: PE, NE

PE

	B_C	B_D
A_C	1, 1	5, 0
A_D	0, 5	3, 3

$$\xi_A = (A_d, B_c)^0 \preccurlyeq (A_c, B_c)^1 \preccurlyeq (A_d, B_d)^3 \preccurlyeq (A_c, B_d)^5$$

$$\xi_B = (A_c, B_d)^0 \preccurlyeq (A_c, B_c)^1 \preccurlyeq (A_d, B_d)^3 \preccurlyeq (A_d, B_c)^5$$

Prisoners Dilemma: PE, NE

	B_C	B_D
A_C	1, 1	5, 0
A_D	0, 5	3, 3

PE

NE

The paradox of Prisoner's Dilemma: the Nash equilibrium is the only non-Pareto-optimal outcome

Prisoners Dilemma: PE, NE

PE

NE

dominant

	B_C	B_D
A_C	1, 1	5, 0
A_D	0, 5	3, 3

The paradox of Prisoner's Dilemma: the Nash equilibrium is the only non-Pareto-optimal outcome

Prisoners Dilemma: PE, NE

social welfare optimal

dominant

	B_C	B_D
A_C	1, 1	5, 0
A_D	0, 5	3, 3

PE

NE

The paradox of Prisoner's Dilemma: the Nash equilibrium is the only non-Pareto-optimal outcome

Example: Routing

- 1,000 drivers travel from S to D on either $S \rightarrow A \rightarrow D$ or $S \rightarrow B \rightarrow D$
- Road from $S \rightarrow A$, $B \rightarrow D$ is long: $t = 50$ minutes for any $|\text{cars}|$
- Road from $A \rightarrow D$, $S \rightarrow B$ is shorter but is narrow $t = |\text{cars}|/25$

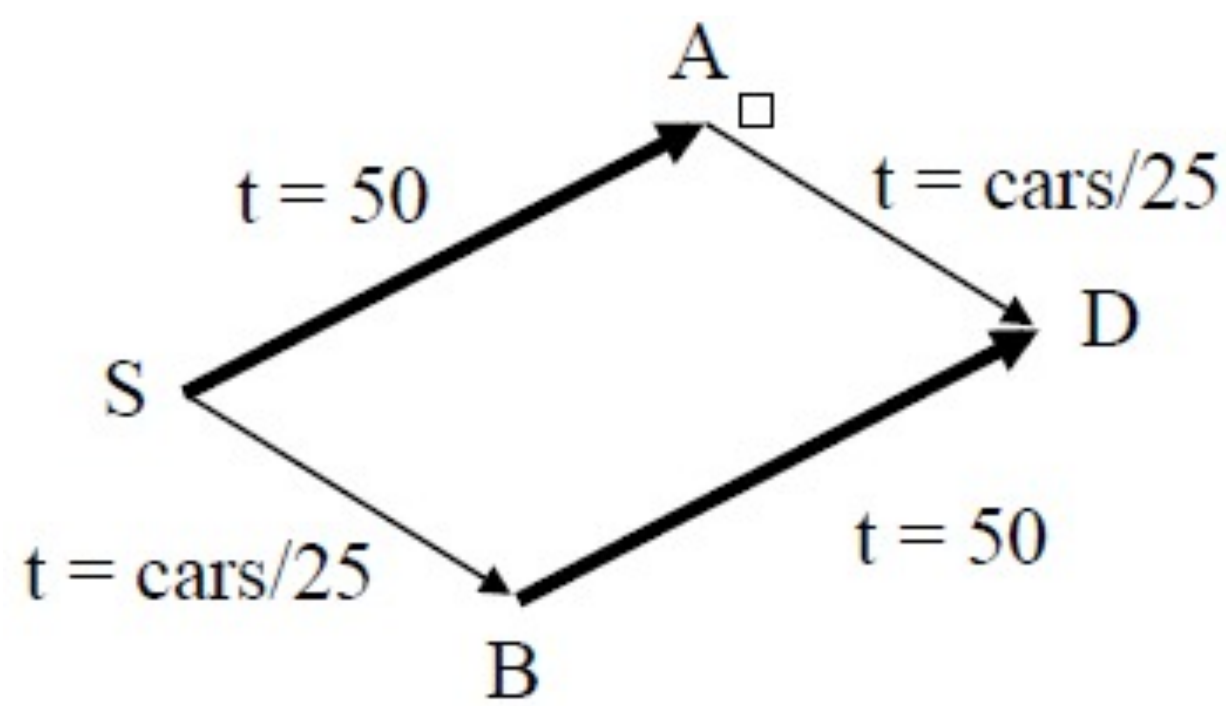

- Nash equilibrium:

– 500 cars go through A, 500 through B with time is 50 + 500/25 = 70m

– If a single driver changes the route, there are 501 cars on that route: time ↑

Braess's Paradox

- Suppose we add a new road from B to A
- The road is so wide and short that it takes 0 minutes to traverse it
- Nash equilibrium:
  - All 1000 cars go S→B→A→D
  - Time for S→B is 1000/25 = 40 minutes
  - Total time is 80 minutes

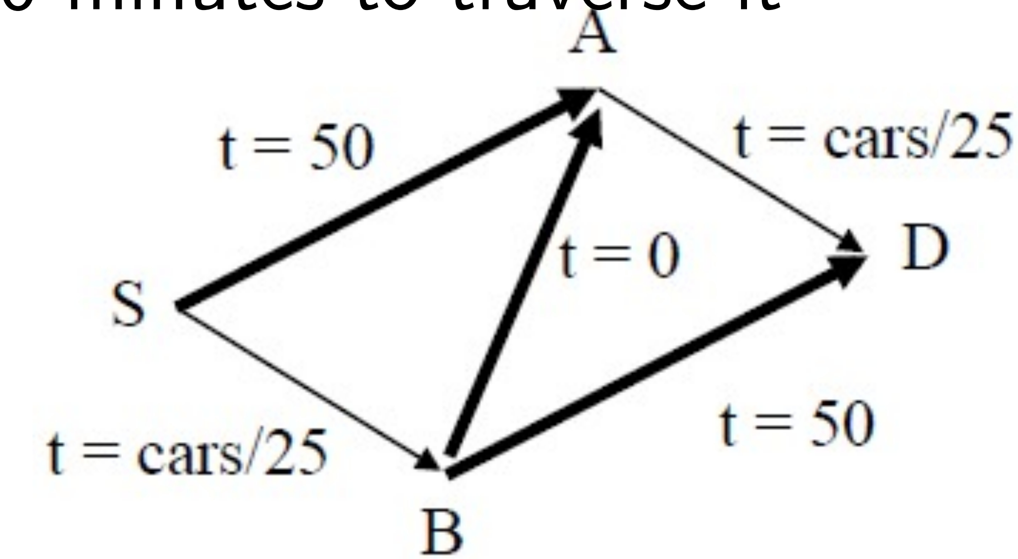

- To see that this is an equilibrium:
  - If driver goes $S \to A \to D$, his/her cost is $50 + 40 = 90$ minutes
  - If driver goes $S \to B \to D$, his/her cost is $40 + 50 = 90$ minutes
  - Both are dominated by S→B→A→D
- To see that it's the ***only*** Nash equilibrium:
  - For every traffic pattern, S→B→A→D dominates $S \to A \to D$ and $S \to B \to D$

Mediated Prisoners Dilemma

	Cooperate	Defect
Cooperate	$1, 1$	$5, 0$
Defect	$0, 5$	$3, 3$

Mediated Game

	Mediator	Cooperate	Defect
Mediator	1, 1	0, 5	2, 2
Cooperate	5, 0	1, 1	5, 0
Defect	2, 2	0, 5	3, 3

Mediated Equilibrium

	Mediator	Cooperate	Defect
Mediator	1, 1	0, 5	2, 2
Cooperate	5, 0	1, 1	5, 0
Defect	2, 2	0, 5	3, 3

Iterated Prisoner Dilemma

- The problem of repeatedly played PD game. Optimization for total count of each player outcome. Sometimes IPD can be played against a range of different; opponents (or even several at the same time).
  - motives for cooperation: (i) if you know you will be meeting your opponent again, then the incentive to defect appears to evaporate. (ii) defection may be punished in the future round,
  - motives for defection: (i) you can test the water by defection (ii) cooperative defection is the rational choice in the infinitely repeated prisoner's dilemma

Iterated Prisoner Dilemma

- What strategy to choose, so as to maximize your overall payoff?
- Axelrod (1984) investigated this problem, with a computer tournament for programs playing the iterated prisoner's dilemma:

ALLD	Always defect. the Hawk or Free rider strategy
ALLC	Always cooperate
TITforTAT	first cooperate, than do what your opponent did
TF2T	Same as above, but requires TWO consecutive defections for a defection to be returned
STFT	Suspicious TFT - first, defect. If the opponent retaliated, then play TITforTAT. Otherwise intersperse cooperation & defection.
JOSS	As TIT-FOR-TAT, except periodically defect.

Mixed Strategy

- In many games, deterministic strategy is very inefficient.
  - example: matching pennies, security games
- Solution: randomize selection of an action
- Pure strategy
  - agents makes the decision to play one action
- Mixed strategy
  - agents chose to play more actions with positive probabilities
  - support of the mixed strategy is the set of all selected actions
- Payoff
  - given the strategy profile $s \in S$ for all agents, the utility for the agent i

$$u_i(s) = \sum_{a \in A} u_i(a_i) Pr(a|s) \qquad Pr(a|s) = \prod_{j \in N} s_j(a_i)$$

Mixed Strategy

Let us generalize the NE concepts for strategy profiles:

- Best response:
  - $s_i^* \in BR(s_{-i})$ iff $\forall s_i \in S_i$, $u_i(s_i^*, s_{-i}) \geq u_i(s_i, s_{-i})$
- Nash equilibrium:
  - $s = \langle s_1, \ldots, s_n \rangle$ is a Nash equilibrium iff $\forall i$, $s_i \in BR(s_{-i})$
- Every finite game has a Nash equilibrium! [Nash, 1950]
  - e.g., matching pennies:

Mixed Strategy

Let us generalize the NE concepts for strategy profiles:

- Best response:
  - $s_i^* \in BR(s_{-i})$ iff $\forall s_i \in S_i,\ u_i(s_i^*, s_{-i}) \geq u_i(s_i, s_{-i})$
- Nash equilibrium:
  - $s = \langle s_1, \ldots, s_n \rangle$ is a Nash equilibrium iff $\forall i,\ s_i \in BR(s_{-i})$
- Every finite game has a Nash equilibrium! [Nash, 1950]
  - e.g., matching pennies: both players play heads/tails 50%/50%

Mixed Strategy

- It's hard in general to compute Nash equilibria, but it's easy when you can guess the support (*actions played with non-null probability*). For BoS, let's look for an equilibrium where all actions are part of the support.

	B	*F*
B	2, 1	0, 0
F	0, 0	1, 2

Mixed Strategy

- It's hard in general to compute Nash equilibria, but it's easy when you can guess the support (*actions played with non-null probability*). For BoS, let's look for an equilibrium where all actions are part of the support.

	B	*F*
B	2, 1	0, 0
F	0, 0	1, 2

- Let player 2 play *B* with p, *F* with $1 - p$.
- If player 1 best-responds with a mixed strategy, player 2 must make him indifferent between *F* and *B* (why?)

Mixed Strategy

	B	*F*
B	$2, 1$	$0, 0$
F	$0, 0$	$1, 2$

- Let player 2 play *B* with p, *F* with $1 - p$.
- If player 1 best-responds with a mixed strategy, player 2 must make him indifferent between *F* and *B* (why?)

Mixed Strategy

	B	*F*
B	$2, 1$	$0, 0$
F	$0, 0$	$1, 2$

- Let player 2 play *B* with p, *F* with $1 - p$.
- If player 1 best-responds with a mixed strategy, player 2 must make him indifferent between *F* and *B* (why?)

Mixed Strategy

	B	*F*
B	$2, 1$	$0, 0$
F	$0, 0$	$1, 2$

- Let player 2 play *B* with p, *F* with $1 - p$.
- If player 1 best-responds with a mixed strategy, player 2 must make him indifferent between *F* and *B* (why?)

$$u_1(B) = u_1(F)$$
$$2p + 0(1-p) = 0p + 1(1-p)$$
$$p = \frac{1}{3}$$

Mixed Strategy

	B	F
B	2, 1	0, 0
F	0, 0	1, 2

- Likewise: Let player I play *B* with q, *F* with $1-q$.

$$u_2(B) = u_2(F)$$
$$q + 0(1-q) = 0q + 2(1-q)$$
$$q = \frac{2}{3}$$

- Thus the mixed strategy $(\frac{2}{3}, \frac{1}{3}), (\frac{1}{3}, \frac{2}{3})$ is in Nash Equilibrium

Interpreting Mixed Strategy

- What does it mean to play a mixed strategy? Different interpretations:
  - Randomize to confuse your opponent
    - *consider the matching pennies example*
  - Randomize when they are uncertain about the other's action
    - *consider battle of the sexes*
  - Randomize when they you allocated limited resources
- Mixed strategies are a concise description of what might happen in repeated play: count of pure strategies in the limit
- Mixed strategies describe population dynamics: 2 agents chosen from a population, all having deterministic strategies.

Solution Concepts

- Pareto Efficiency
- Social welfare optimality
- Nash equilibrium
- Maxmin
- Dominant strategies
- Correlated equilibrium
- Minimax regret
- Stackelberg equilibrium
- Perfect equilibrium
- ϵ - Nash equilibrium

Deductive vs. Stable State

- Deductive:
  - game treated as an isolated "one-shot" event
  - equilibrium reached by deductive process
- Steady-state:
  - player optimizes his strategy based on his experience with the game
  - equilibrium reached through adaptation/learning

Maxmin Strategies

- Player's i maxmin strategy is a strategy that maximizes her worst-case payoff, in the situation where the opponent(s) $-i$ plays a strategy that minimizes i's payoff.
- The `maximum` value (or safety level) of the game for *i* is the minimum amount of payoff guaranteed by `maxmin` strategy
- Good conservative strategy:
  - *maximize his/her expected utility without making any assumption about the other players*

Maxmin Strategies

01

- Player's i maxmin strategy is a strategy that maximizes her worst-case payoff, in the situation where the opponent(s) $-i$ plays a strategy that minimizes i's payoff.
- The `maximum` value (or safety level) of the game for *i* is the minimum amount of payoff guaranteed by `maxmin` strategy
- Good conservative strategy:
  - *maximize his/her expected utility without making any assumption about the other players*

Definition (Maxmin)

The maxmin strategy for player i is $\arg\max_{s_i} \min_{s_{-i}} u_i(s_1, s_2)$, and the maxmin value for player i is $\max_{s_i} \min_{s_{-i}} u_i(s_1, s_2)$.

Minmax Strategies

- Player i's minmax strategy against player $-i$ in a 2-player game is a strategy that minimizes $-i$'s best-case payoff, and the `minmax` value for i against $-i$ is payoff.

Minmax Strategies

- Player i's minmax strategy against player $-i$ in a 2-player game is a strategy that minimizes $-i$'s best-case payoff, and the `minmax` value for i against $-i$ is payoff.

Definition (Minmax, 2-player)

In a two-player game, the minmax strategy for player i against player $-i$ is $\arg\min_{s_i} \max_{s_{-i}} u_{-i}(s_i, s_{-i})$, and player $-i$'s minmax value is $\min_{s_i} \max_{s_{-i}} u_{-i}(s_i, s_{-i})$.

Generalization to n-players possible

Minmax Strategies

- Player i's minmax strategy against player $-i$ in a 2-player game is a strategy that minimizes $-i$'s best-case payoff, and the `minmax` value for i against $-i$ is payoff.

Definition (Minmax, 2-player)

In a two-player game, the minmax strategy for player i against player $-i$ is $\arg\min_{s_i} \max_{s_{-i}} u_{-i}(s_i, s_{-i})$, and player $-i$'s minmax value is $\min_{s_i} \max_{s_{-i}} u_{-i}(s_i, s_{-i})$.

Generalization to n-players possible

- Why would i want to play a `minmax` strategy?
  - *to punish the other agent as much as possible*
- Minmax profile: Minmax strategy for each player

Minmax Theorem

Theorem (Minimax theorem (von Neumann, 1928))

In any finite, two-player, zero-sum game, in any Nash equilibrium each player receives a payoff that is equal to both his maxmin value and his minmax value.

1. Each player's **`maxmin`** value is equal to his `minmax` value. By convention, the `maxmin` value for player 1 is called the **value of the game**.
2. For both players, the set of `maxmin` strategies coincides with the set of **`minmax`** strategies.
3. Any `maxmin` strategy profile (or, equivalently, `minmax` strategy profile) is a Nash equilibrium. Furthermore, these are all the Nash equilibria. Consequently, **all Nash equilibria have the same payoff vector** (namely, those in which player 1 gets the value of the game).

Finding EQ in 2-PL Zero-Sum Game

1. Determine whether the equilibrium point is associated with pure strategies:
   - determine if the row player's `maxmin` strategy and the column player's `minmax` strategy coincide in the same outcome
   - if this is true, then the associated strategies are the equilibrium point of the game
2. If pure strategies do not produce an equilibrium point
   - define variables that represent the probability each player will play each available strategy.
   - for each player we find the probabilities that will provide the lowest expected payoff for the other player

Example

- Game setting:

		Pitcher: *Fastball*	Pitcher: *Curve*
Batter	*Fastball*	0.3, -0.3	0.2, -0.2
Batter	*Curve*	0.1, -0.1	0.4, -0.4

Example

- Game setting:

		Pitcher	
		Fastball	*Curve*
Batter	*Fastball*	0.3	0.2
	Curve	0.1	0.4

Example

- maximin for batter:

Pitcher

		Fastball	*Curve*
Batter	*Fastball*	0.3	0.2
	Curve	0.1	0.4

Example

- maximin for batter:

		Pitcher		
		Fastball	*Curve*	min
Batter	*Fastball*	0.3	0.2	0.2
	Curve	0.1	0.4	0.1

Example

- maximin for batter:

Pitcher

Batter		*Fastball*	*Curve*	min
	Fastball	0.3	0.2	0.2
	Curve	0.1	0.4	0.1

Example

- minimax for pitcher:

		Pitcher	
		Fastball	*Curve*
Batter	*Fastball*	0.3	0.2
	Curve	0.1	0.4

Example

- minimax for pitcher:

		Pitcher	
		Fastball	*Curve*
Batter	*Fastball*	0.3	0.2
	Curve	0.1	0.4
	max	0.3	0.4

Example

- minimax for pitcher:

		Pitcher	
		Fastball	*Curve*
Batter	*Fastball*	0.3	0.2
	Curve	0.1	0.4
	max	0.3	0.4

Example

- Because the maximin and minimax strategies do not equal, the equilibrium point is not found with pure strategies.
- What strategy should each adopt assuming best play by the opponent?

$$u_B(f) = 0.3 * p + 0.2(1-p)$$

$$u_B(c) = 0.1 * p + 0.4(1-p)$$

- minimax for Pitcher? such p that minimizes $\max\left(u_B(f), u_B(c)\right)$
- this happens for

$$u_B(f) = u_B(c)$$

$$0.1p + 0.2 = -0.3p + 0.4$$

$$p = 0.5$$

$$u_B = 0.25$$

Pitcher

Batter		Fastball	Curve
	Fastball	0.3	0.2
	Curve	0.1	0.4
		p	$1-p$

Example

- Because the maximin and minimax strategies do not equal, the equilibrium point is not found with pure strategies.
- What strategy should each adopt assuming best play by the opponent?

$$u_P(f) = 0.3 * q + 0.1(1 - q)$$
$$u_P(c) = 0.2 * q + 0.4(1 - q)$$

- maximin for Batter? such q that maximizes

$$\min\left(u_P(f), u_P(c)\right)$$

- this happens for

$$u_P(f) = u_P(c)$$
$$0.2q + 0.1 = -0.2q + 0.4$$
$$q = 0.75$$
$$u_B = 0.25$$

Pitcher

Batter		Fastball	Curve	
	Fastball	0.3	0.2	q
	Curve	0.1	0.4	1- q

Saddle point: Matching Pennies

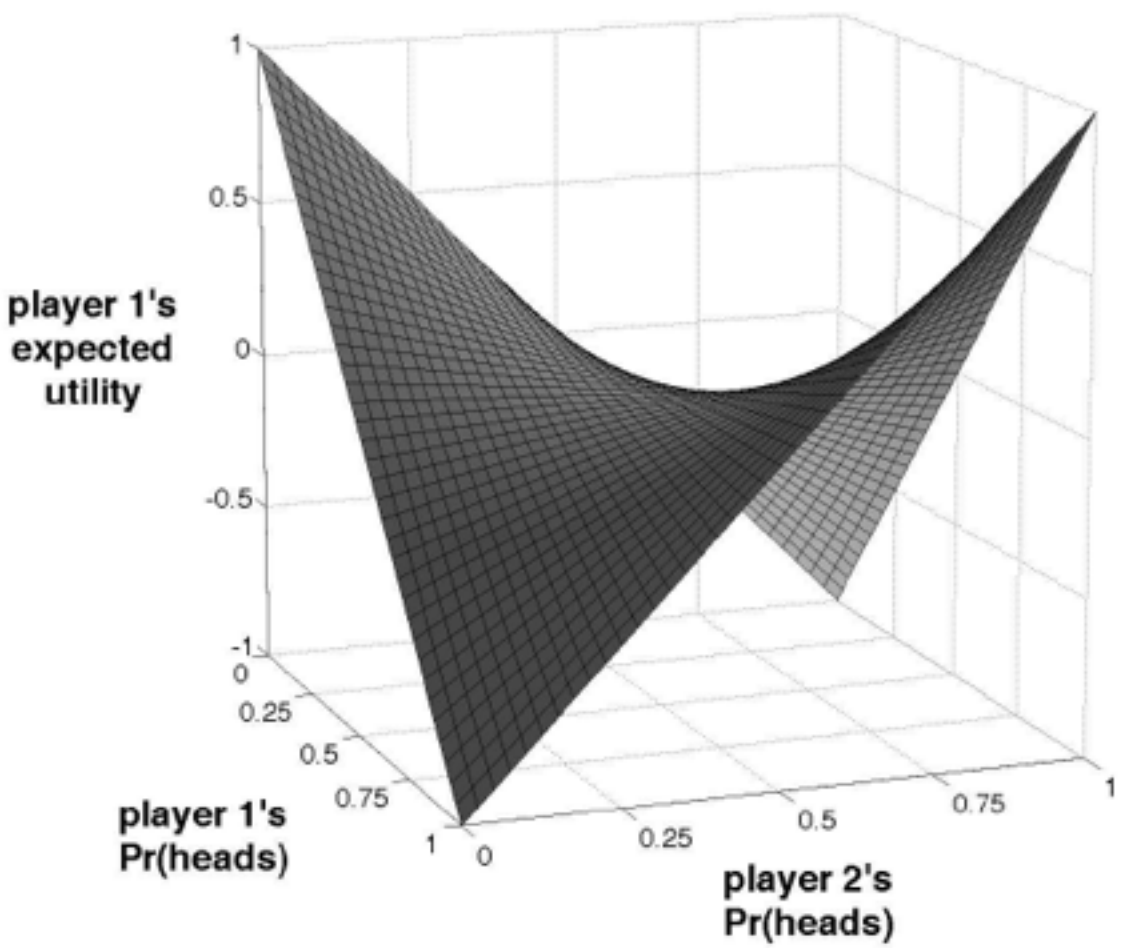

Solution Concepts

- Pareto Efficiency
- Social welfare optimality
- Nash equilibrium
- Maxmin
- Dominant strategies
- Correlated equilibrium
- Minimax regret
- Stackelberg equilibrium
- Perfect equilibrium
- ϵ - Nash equilibrium

Domination

Definition

s_i strictly dominates s_i' if $\forall s_{-i} \in S_{-i}$, $u_i(s_i, s_{-i}) > u_i(s_i', s_{-i})$

Definition

s_i weakly dominates s_i' if $\forall s_{-i} \in S_{-i}$, $u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i})$ and $\exists s_{-i} \in S_{-i}$, $u_i(s_i, s_{-i}) > u_i(s_i', s_{-i})$

Definition

s_i very weakly dominates s_i' if $\forall s_{-i} \in S_{-i}$, $u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i})$

- *dominant* dominates any other strategy, *dominated* is dominated by at least one other strategy

Domination

Definition

s_i strictly dominates s'_i if $\forall s_{-i} \in S_{-i}$, $u_i(s_i, s_{-i}) > u_i(s'_i, s_{-i})$

Definition

s_i weakly dominates s'_i if $\forall s_{-i} \in S_{-i}$, $u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i})$ and $\exists s_{-i} \in S_{-i}$, $u_i(s_i, s_{-i}) > u_i(s'_i, s_{-i})$

Definition

s_i very weakly dominates s'_i if $\forall s_{-i} \in S_{-i}$, $u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i})$

- A strategy profile consisting of dominant strategies for every player must be a Nash equilibrium.
- An equilibrium in strictly dominant strategies must be unique.

Example

	B_C	B_D
A_C	1, 1	5, 0
A_D	0, 5	3, 3

social welfare optimal

dominant

PE

NE

- Defect (D) is strongly dominant for the row player
- Defect (D) is strongly dominant for the column player
- So (D, D) is a Nash equilibrium in dominant strategies
- Ironically, of the pure strategy profiles, (D,D) is the only one that's not Pareto optimal

Example

	Heads	Tails
Heads	1, -1	-1, 1
Tails	-1, 1	1, -1

- Heads isn't dominant for the row player
- Tails isn't dominant for the row player either
- Row player (and column player too) doesn't have a dominant strategy => No Nash equilibrium in dominant strategies
- Dominant strategy does not always exist

Minimax Regret

Definition: Minimax regret actions for agent i are defined as:

$$\arg\min_{a_i \in A_i} \left[\max_{a_{-i} \in A_{-i}} \left(\left[\max_{a'_i \in A_i} u_i(a'_i, a_{-i}) \right] - u_i(a_i, a_{-i}) \right) \right]$$

	L	R
T	$100, a$	$1-\epsilon, b$
B	$2, c$	$1, d$

Minimax Regret

OI

Definition: Minimax regret actions for agent i are defined as:

$$\arg\min_{a_i \in A_i} \left[\max_{a_{-i} \in A_{-i}} \left(\left[\max_{a'_i \in A_i} u_i(a'_i, a_{-i}) \right] - u_i(a_i, a_{-i}) \right) \right]$$

	L	R
T	$100, a$	$1\text{-}\epsilon, b$
B	$2, c$	$1, d$

maxmin (safety level) for row player is to play B, but if not malicious (and we do not know the values of a, b, c, d), minimax regret for row player is to play T to minimize worst case loss (100-2 for playing B > than ϵ for playing T)

Correlated Equilibrium

- Consider again Battle of the Sexes.
  - The best outcome seems a 50-50 split between (F; F) and (B;B).
  - But there's no way to achieve this, so either someone loses out (unfair) or both players often miscoordinate

Correlated Equilibrium

- Consider again Battle of the Sexes.
  - The best outcome seems a 50-50 split between (F; F) and (B;B).
  - But there's no way to achieve this, so either someone loses out (unfair) or both players often miscoordinate
- Another classic example: traffic game

	GO	WAIT
GO	-100, -100	10, 0
WAIT	0, 10	-100, -100

Correlated Equilibrium

- Consider again Battle of the Sexes.
  - The best outcome seems a 50-50 split between (F; F) and (B;B).
  - But there's no way to achieve this, so either someone loses out (unfair) or both players often miscoordinate
- Another classic example: traffic game

	GO	WAIT
GO	-100, -100	10, 0
WAIT	0, 10	-100, -100

- A traffic light: a fair randomizing device that tells one of the agents to go and the other to wait.

Correlated Equilibrium

0I

Definition (Correlated equilibrium)

Given an n-agent game $G = (N, A, u)$, a correlated equilibrium is a tuple (v, π, σ), where v is a tuple of random variables $v = (v_1, \dots, v_n)$ with respective domains $D = (D_1, \dots, D_n)$, π is a joint distribution over v, $\sigma = (\sigma_1, \dots, \sigma_n)$ is a vector of mappings $\sigma_i : D_i \mapsto A_i$, and for each agent i and every mapping $\sigma_i' : D_i \mapsto A_i$ it is the case that

$$\sum_{d \in D} \pi(d) u_i \left(\sigma_1(d_1), \dots, \sigma_n(d_n)\right) \geq \sum_{d \in D} \pi(d) u_i \left(\sigma_1'(d_1), \dots, \sigma_n'(d_n)\right).$$

- For every Nash equilibrium there exists a corresponding correlated equilibrium
- Not every correlated equilibrium is a Nash equilibrium

  => weaker notion

Stackelberg Equilibrium

- A game theoretic equilibrium in which one player acts as a **leader** and another as a **follower**, the leader setting strategy taking account of the follower's optimal response.

	L	R
T	$1, 0$	$3, 2$
B	$2, 1$	$4, 0$

Stackelberg equilibrium: $3, 2$

Nash equilibrium: $2, 1$

- Stackleberg equilibrium is studied in the context of security games.

Computing Solution Concepts

Note on Linear Programing

- Set of real-valued variables
- Linear objective function
  - a weighted sum of the variables
- Set of linear constraints
  - a weighted sum of the variables must be greater than or equal to some constant

} linear program

Note on Linear Programming

01

- Given n variables and m constraints, variables x and constants w, a, b:

$$\text{maximize} \quad \sum_{i=1}^{n} w_i x_i$$

$$\text{subject to} \quad \sum_{i=1}^{n} a_{ij} x_i \le b_j \qquad \forall j = 1 \ldots m$$

$$x_i \in \mathbb{R} \qquad \forall i = 1 \ldots n$$

- Can be solved in polynomial time using interior point methods.
  - Interestingly, the (worst-case exponential) simplex method is often faster in practice.

Computing Equilibria of Zero-Sum Games

0I

$$
\begin{aligned}
\text{minimize} \quad & U_1^* \\
\text{subject to} \quad & \sum_{a_2 \in A_2} u_1(a_1, a_2) \cdot s_2^{a_2} \leq U_1^* & \forall a_1 \in A_1 \\
& \sum_{a_2 \in A_2} s_2^{a_2} = 1 \\
& s_2^{a_2} \geq 0 & \forall a_2 \in A_2
\end{aligned}
$$

- variables:
  - U_1^* is the expected utility for player 1
  - $s_2^{a_2}$ is player 2's probability of playing action a_2 under his mixed strategy
- each $u_1(a_1, a_2)$ is a constant
- we want to minimize player 1's max utility U_1^* (i.e. we get minmax strategy for player 2)

Computing Equilibria of Zero-Sum Games

U_1^* as small as possible

$$\begin{aligned}
\text{minimize} \quad & U_1^* \\
\text{subject to} \quad & \sum_{a_2 \in A_2} u_1(a_1, a_2) \cdot s_2^{a_2} \leq U_1^* && \forall a_1 \in A_1 \\
& \sum_{a_2 \in A_2} s_2^{a_2} = 1 \\
& s_2^{a_2} \geq 0 && \forall a_2 \in A_2
\end{aligned}$$

s_2 is a valid probability distribution

player 1's expected utility for playing each of his actions under player 2's mixed strategy is no more than U_1^*

- *because U_1^* is minimized, this constraint will be tight for some actions: the support of player 1's mixed strategy.*

Computing Equilibria of Zero-Sum Games

- This formulation gives us the minmax strategy for player 2.
- To get the minmax strategy for player 1, we need to solve a second (analogous) LP.

Computing Equilibria of General-Sum Games

- Computing NE in general-sum has exponential worst-case complexity
- Solution using Lemke-Howson algorithm: Formulates the problem as a *linear complementarity problem* (LCP)

- For $n \geq 3$, the problem can no longer be represented even as a linear complementarity problem
  - nonlinear complementarity problem formulation possible but such problems hopelessly impractical to solve exactly
- Approaches used
  - approximate the solution using a **sequence of linear complementarity problems** (SLCP)
  - formulate as a **contraint optimization problem**

Computing Maxmin in General-Sum Games

- Let's say we want to compute a maxmin strategy for player 1 in an arbitrary 2-player game G.
- Create a new game G' where player 2's payoffs are just the negatives of player 1's payoffs.
- The maxmin strategy for player 1 in G does not depend on
- player 2's payoffs
  - Thus, the maxmin strategy for player 1 in G **is the same as** the maxmin strategy for player 1 in G'
- By the minmax theorem, equilibrium strategies for player 1 in G' **are equivalent** to a maxmin strategies
- Thus, to find a maxmin strategy for G, **find an equilibrium** strategy for G'.

Removal of Dominated Strategies

- No equilibrium can involve a strictly dominated strategy
- Thus we can remove it, and end up with a strategically equivalent game
  - This might allow us to remove another strategy that wasn't dominated before
- Running this process to termination is called iterated removal of dominated strategies

Iterated Removal of Dominated Strategies

	L	C	R
U	$3, 1$	$0, 1$	$0, 0$
M	$1, 1$	$1, 1$	$5, 0$
D	$0, 1$	$4, 1$	$0, 0$

Iterated Removal of Dominated Strategies

	L	C	R
U	3, 1	0, 1	0, 0
M	1, 1	1, 1	5, 0
D	0, 1	4, 1	0, 0

R is dominated by L

Iterated Removal of Dominated Strategies

	L	C	R
U	3, 1	0, 1	0, 0
M	1, 1	1, 1	5, 0
D	0, 1	4, 1	0, 0

Iterated Removal of Dominated Strategies

	L	C	R
U	3, 1	0, 1	0, 0
M	1, 1	1, 1	5, 0
D	0, 1	4, 1	0, 0

R is dominated by L

Iterated Removal of Dominated Strategies

	L	C
U	3, 1	0, 1
M	1, 1	1, 1
D	0, 1	4, 1

Iterated Removal of Dominated Strategies

	L	C
U	$3, 1$	$0, 1$
M	$1, 1$	$1, 1$
D	$0, 1$	$4, 1$

M is dominated by the mixed strategy that selects U and D with equal probability

Iterated Removal of Dominated Strategies

	L	C
U	3, 1	0, 1
M	1, 1	1, 1
D	0, 1	4, 1

Iterated Removal of Dominated Strategies

	L	C
U	3, 1	0, 1
M	1, 1	1, 1
D	0, 1	4, 1

M is dominated by the mixed strategy that selects U and D with equal probability

Iterated Removal of Dominated Strategies

	L	C
U	$3, 1$	$0, 1$
D	$0, 1$	$4, 1$

No others strategies are dominated

Iterated Removal of Dominated

- Preserves Nash equilibria
  - strict dominance => all equilibria preserved.
  - weak or very weak dominance => at least one equilibrium preserved.
- Used as a preprocessing step before computing an equilibrium
  - Some games are solvable using this technique (e.g. Traveler's Dilemma)
- Order of removal in the case of multiple dominated strategies
  - strict dominance => doesn't matter.
  - weak or very weak dominance => can affect which equilibria are preserved.

Practical Implications of Solution

- What to do when faced with a game of certain type?
- Zero-sum game => play any maxmin / equilibirum strategy
- General-sum game =>
  - single unique equilibrium: play the equilibrium
  - multiple equilibria:
    - *conservative player: play a maxmin strategy*
    - *otherwise need additional assumptions on how the other player chooses between multiple equillibria*