A4M33MAS - Multiagent Systems
Introduction to Game Theory

Michal Pechoucek & Branislav Bosansky

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

- 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

Game Theory

- 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 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

Game Theory

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

http://www.masfoundations.org

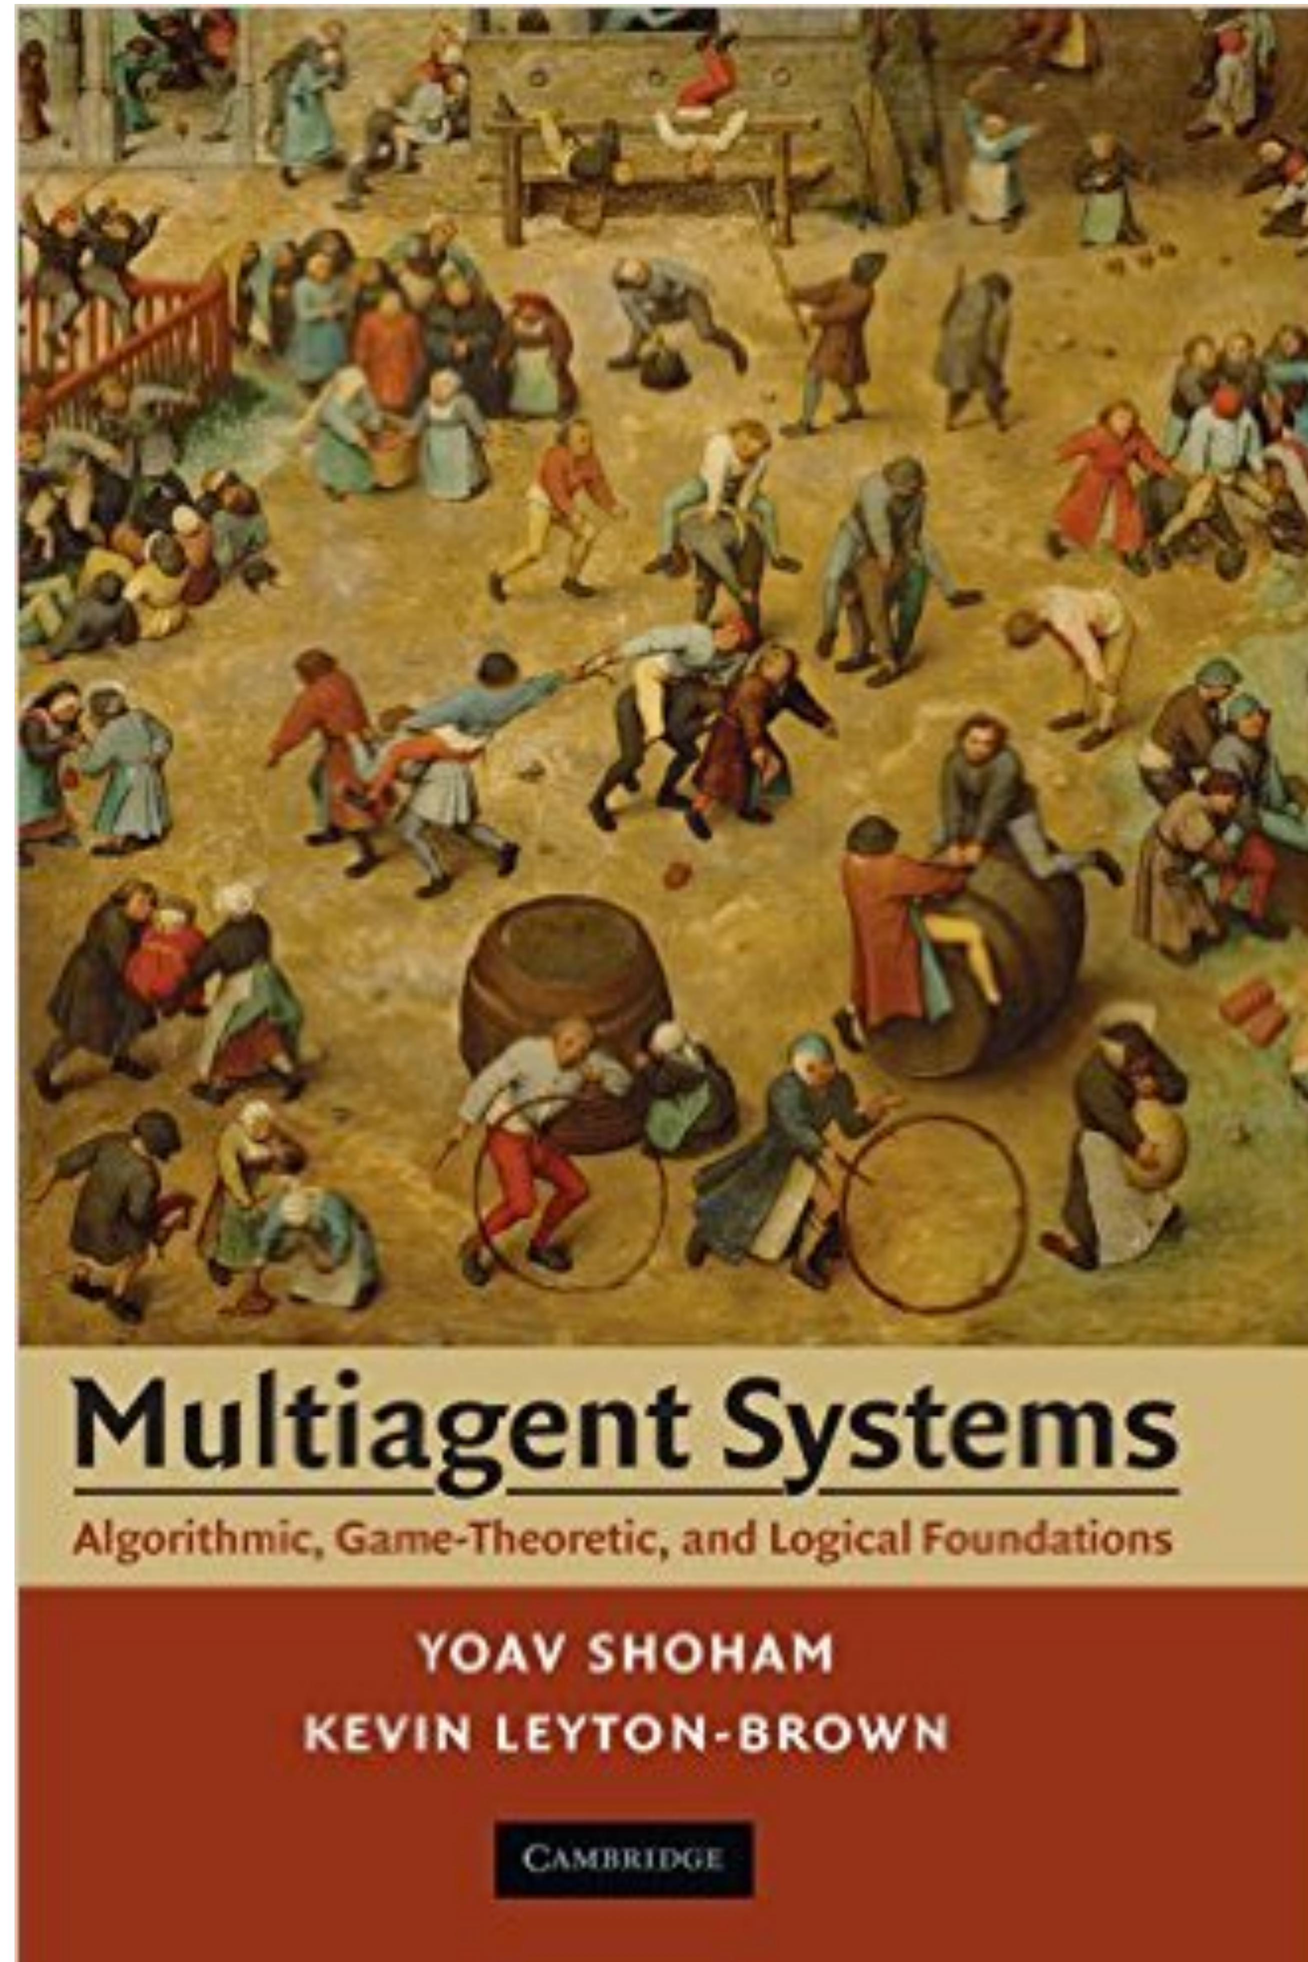

Types of Games

- Cooperative or non-cooperative
- Symmetric and asymmetric
- Zero-sum and non-zero-sum
- Simultaneous and sequential
- Combinatorial games and imperfect information games
- Infinitely long games
- Discrete and continuous games, differential games

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.

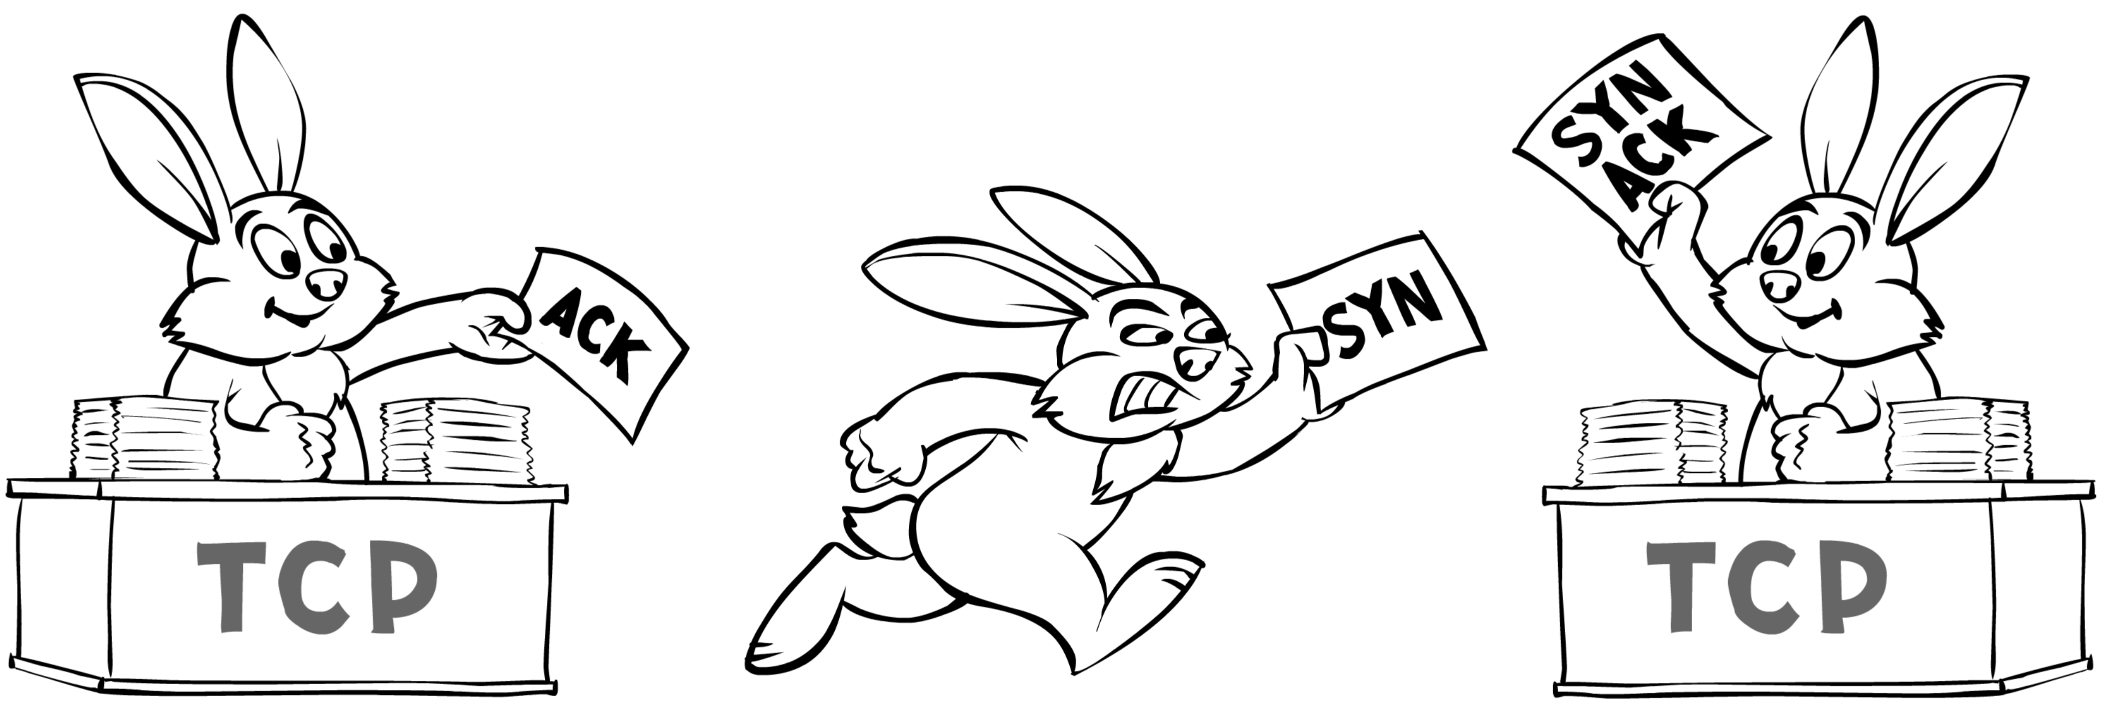

TCP Backoff Game

- 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

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

Other Games: Coordination Games

	Left	Right
Left	1	0
Right	0	1

$driving\ side$

Other Games: Coordination Games

	Left	Right
Left	1	0
Right	0	1

driving side

	Ball	*Football*
Ball	$2, 1$	$0, 0$
Football	$0, 0$	$1, 2$

battle of sexes

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

	Heads	Tails
Heads	1	-1
Tails	-1	1

Other Games: Rock-paper-scissors

01

	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

01

- 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

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 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

23

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

- If you know what everyone else was going to do, it would be easy to pick your own actions

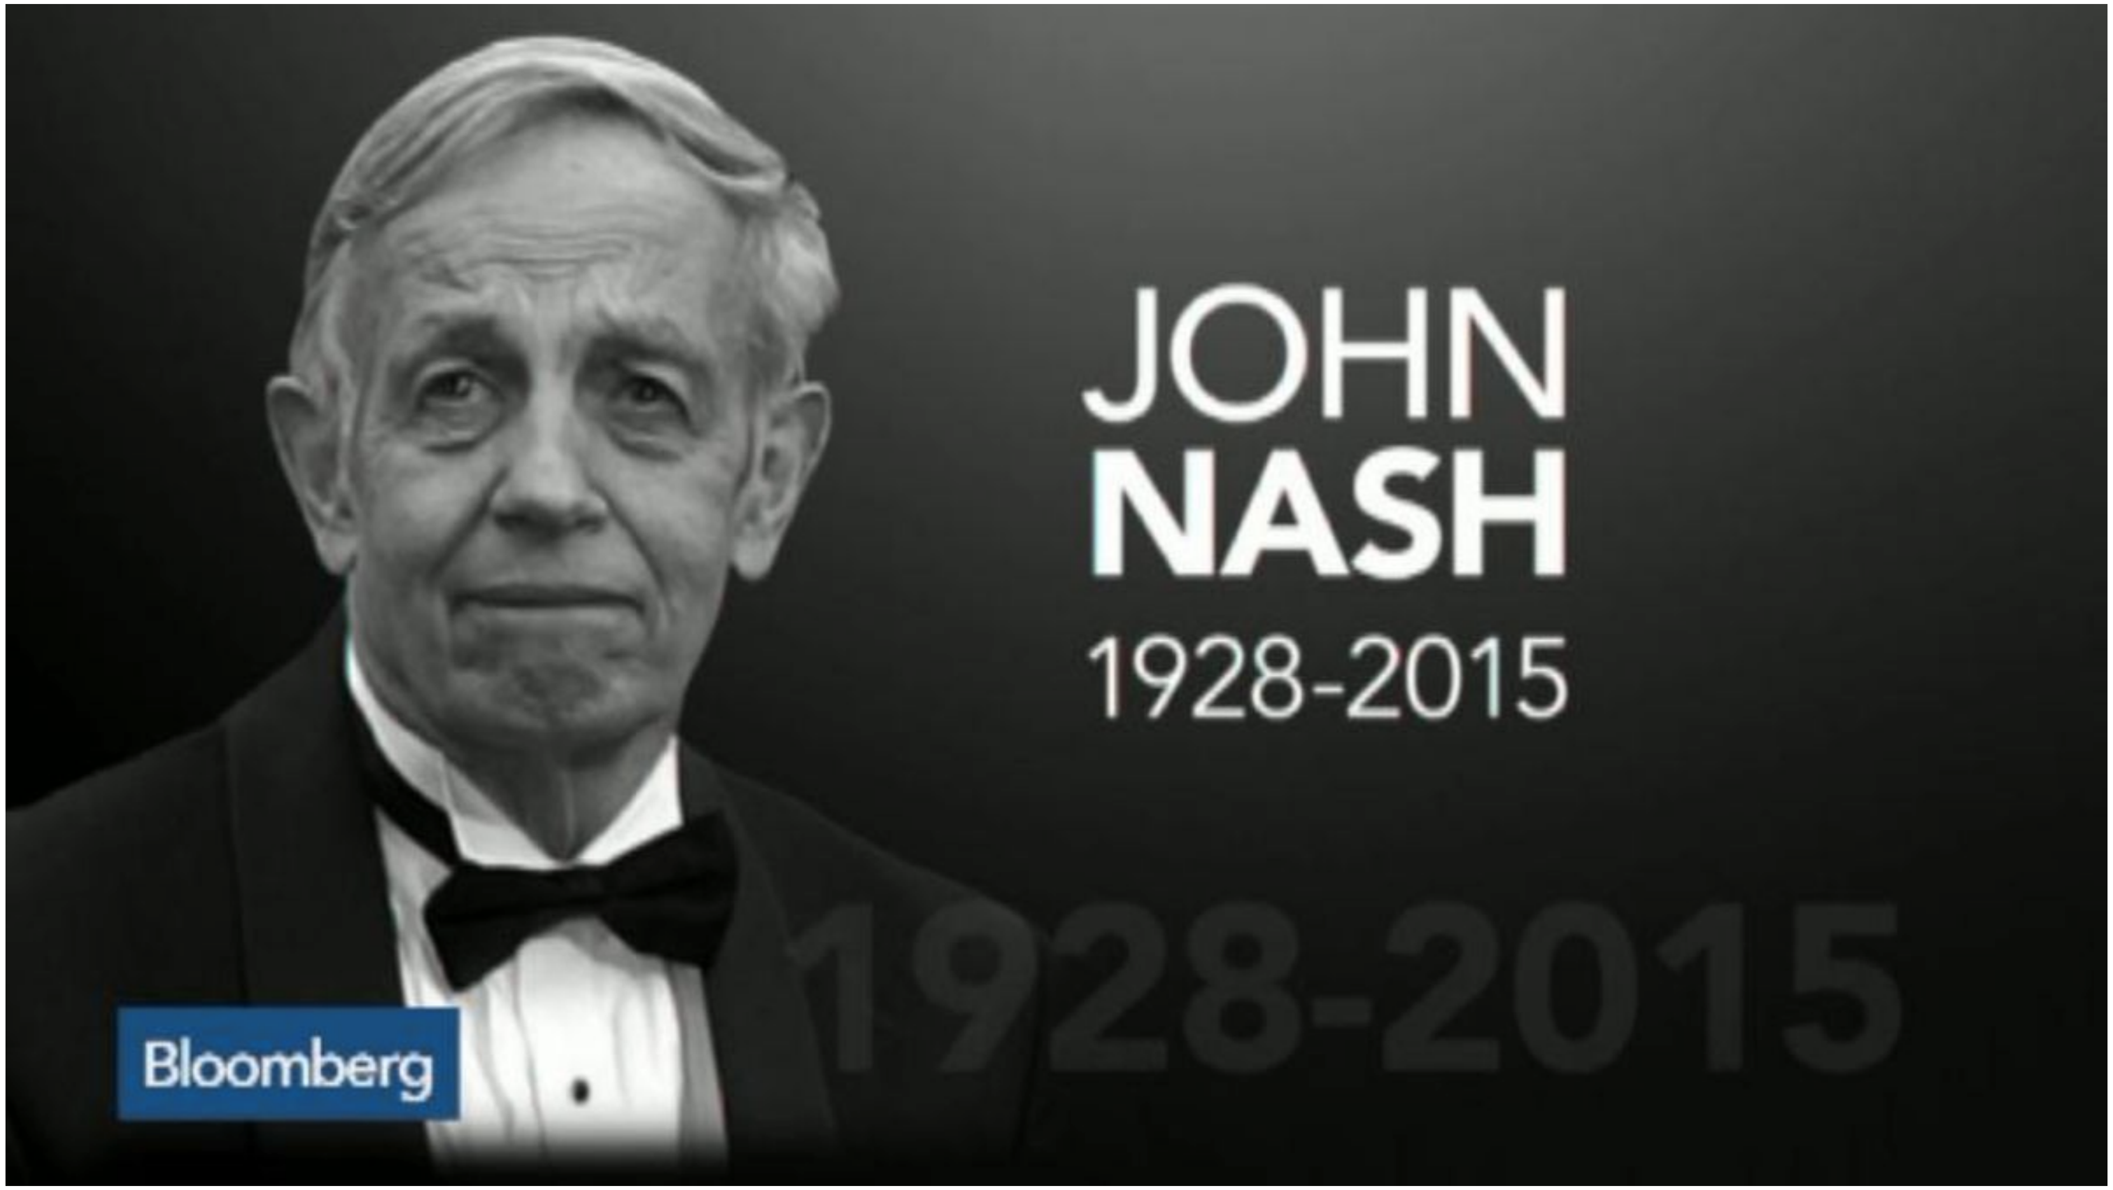

Nash Equilibrium

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

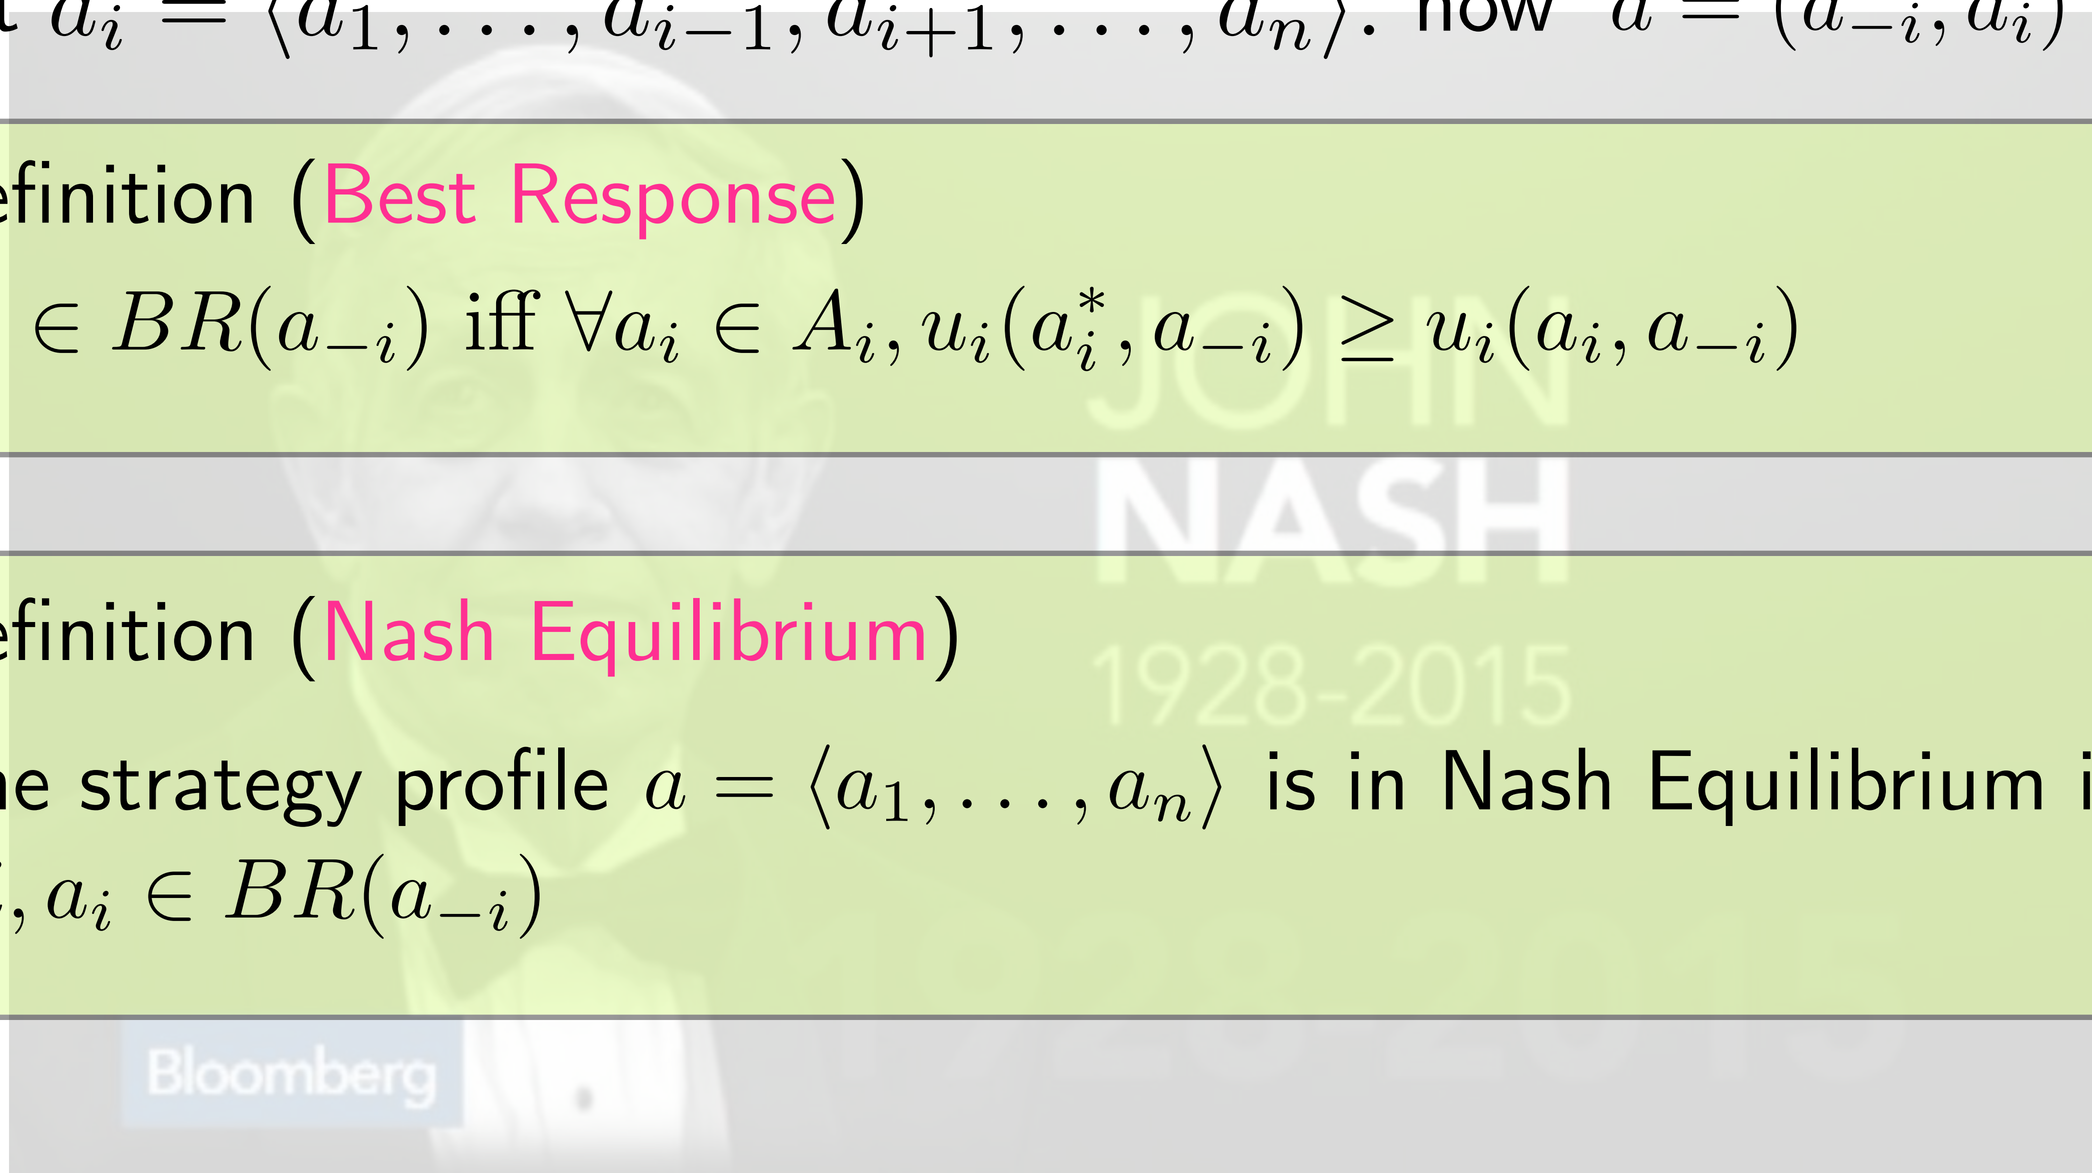

Definition (Best Response)

$a_i^* \in BR(a_{-i})$ 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 (Best Response)

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

- 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 Strict Nash 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

36

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

	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

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

PE

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

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

		B_C	B_D	
social welfare optimal	A_C	1, 1	5, 0	PE
dominant	A_D	0, 5	3, 3	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 $S \to A \to D$ or $S \to B \to D$
- Road from $S \to A$, $B \to D$ is long: $t = 50$ minutes for any $|\text{cars}|$
- Road from $A \to D$, $S \to B$ is shorter but is narrow $t = |\text{cars}|/25$

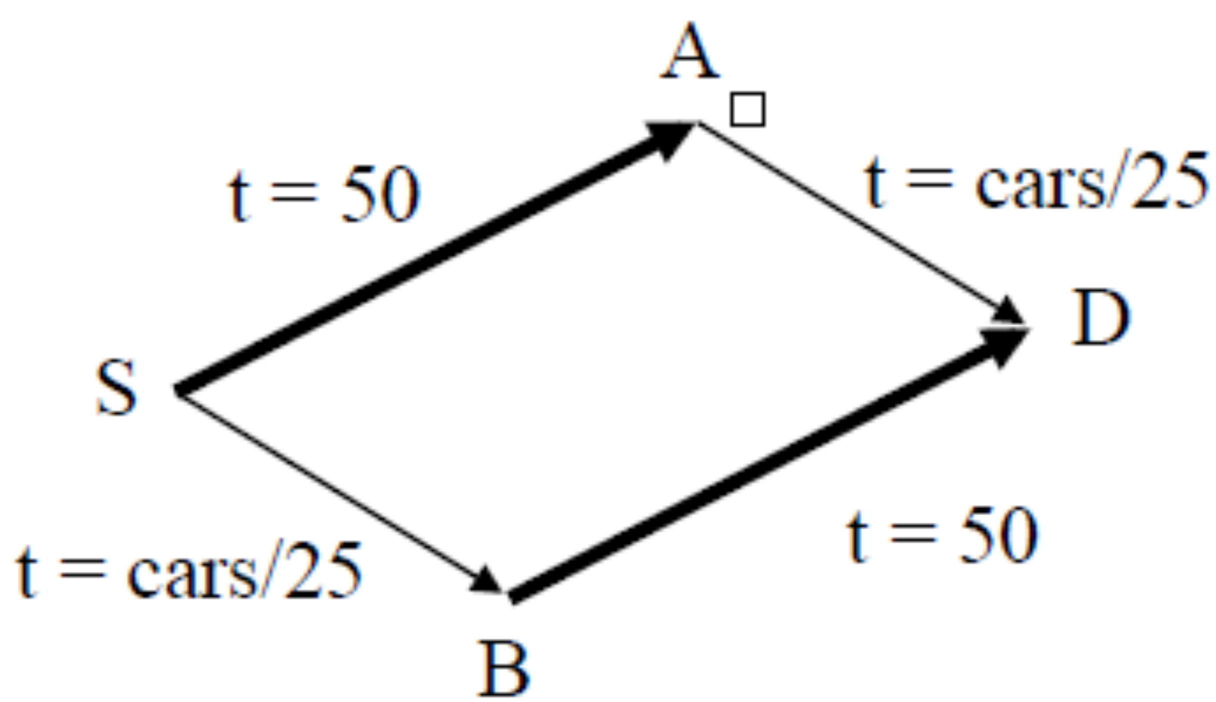

- Nash equilibrium:

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

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

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

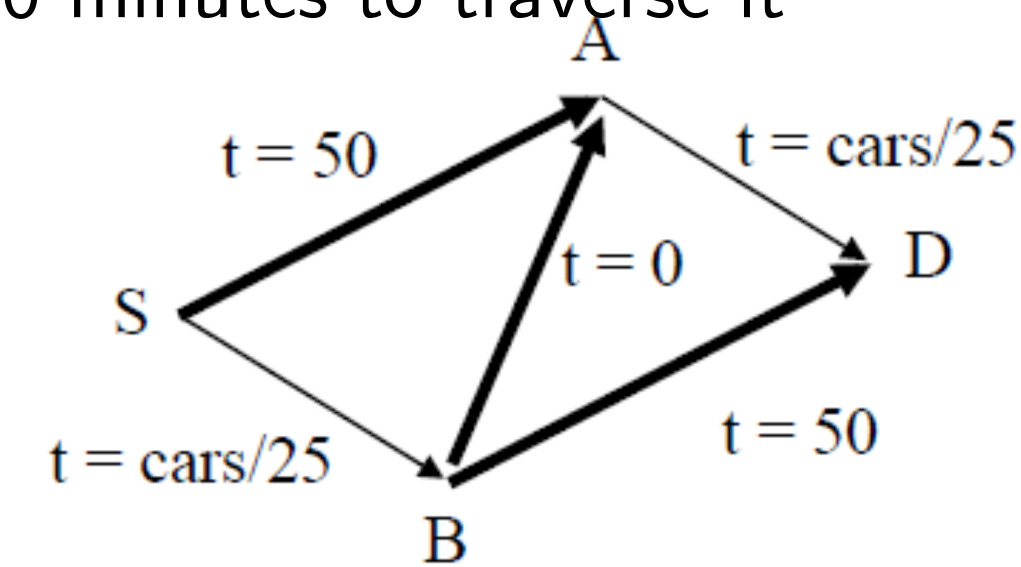

- To see that this is an equilibrium:
  - If driver goes $S \rightarrow A \rightarrow D$, his/her cost is $50 + 40 = 90$ minutes
  - If driver goes $S \rightarrow B \rightarrow D$, his/her cost is $40 + 50 = 90$ minutes
  - Both are dominated by S→B→A→D
- To see that it's the *only* Nash equilibrium: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

Mediated Prisoners Dilemma

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

Mediated Game

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

Mediated Game

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

Mediated Game

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