

Algorithmic Game Theory

Solution concepts in games

Part I

Georgios Birmpas
birbas@diag.uniroma1.it

Based on slides by Vangelis Markakis and Alexandros Voudouris

Solution concepts

Choosing a strategy...

- Given a game, how should a player choose his strategy?
 - Recall: we assume each player knows the other players' preferences but not what the other players will choose
- The most fundamental question of game theory
 - Clearly, the answer is not always clear
- We will start with 2-player games

Prisoner's Dilemma: The Rational Outcome

- Let's revisit prisoner's dilemma
- Reasoning of pl. 1:
 - If pl. 2 does not confess, then I should confess
 - If pl. 2 confesses, then I should also confess

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

- Similarly for pl. 2
- **Expected outcome for rational players:** they will both confess, and they will go to jail for 3 years each
 - **Observation:** If they had both chosen not to confess, they would go to jail only for 1 year, each of them would have a strictly better utility

Dominant strategies

- Ideally, we would like a strategy that would provide the best possible outcome, regardless of what other players choose

- Definition: A strategy s_i of pl. 1 is *dominant* if

$$u_1(s_i, t_j) \geq u_1(s', t_j)$$

for every strategy $s' \in S^1$ and every strategy $t_j \in S^2$

- Similarly for pl. 2, a strategy t_j is dominant if

$$u_2(s_i, t_j) \geq u_2(s_i, t')$$

for every strategy $t' \in S^2$ and for every strategy $s_i \in S^1$

Dominant strategies

Even better:

- Definition: A strategy s_i of pl. 1 is *strictly dominant* if

$$u_1(s_i, t_j) > u_1(s', t_j)$$

for every strategy $s' \in S^1$ and every strategy $t_j \in S^2$

- Similarly for pl. 2
- In prisoner's dilemma, strategy C (confess) is strictly dominant

Observations:

- There may be more than one dominant strategies for a player, but then they should yield the same utility under all profiles
- Every player can have at most one strictly dominant strategy
- A strictly dominant strategy is also dominant

Existence of dominant strategies

- Few games possess dominant strategies
- It may be too much to ask for
- E.g. in the Bach-or-Stravinsky game, there is no dominant strategy:
 - Strategy B is not dominant for pl. 1:
If pl. 2 chooses S, pl. 1 should choose S
 - Strategy S is also not dominant for pl. 1:
If pl. 2 chooses B, pl. 1 should choose B
- In all the examples we have seen so far, only prisoner's dilemma possesses dominant strategies

	B	S
B	(2, 1)	(0, 0)
S	(0, 0)	(1, 2)

Back to choosing a strategy...

- Hence, the question of how to choose strategies still remains for the majority of games
- **Model of rational choice:** if a player knows or has a strong belief for the choice of the other player, then he should choose the strategy that maximizes his utility
- Suppose that someone suggests to the 2 players the strategy profile (s, t)
- When would the players be willing to follow this profile?
 - For pl. 1 to agree, it should hold that
$$u_1(s, t) \geq u_1(s', t)$$
 for every other strategy s' of pl. 1
 - For pl. 2 to agree, it should hold that
$$u_2(s, t) \geq u_2(s, t')$$
 for every other strategy t' of pl. 2

Nash Equilibria



- Definition (Nash 1950): A strategy profile (s, t) is a **Nash equilibrium**, if no player has a unilateral incentive to deviate, given the other player's choice
- This means that the following conditions should be satisfied:
 1. $u_1(s, t) \geq u_1(s', t)$ for every strategy $s' \in S^1$
 2. $u_2(s, t) \geq u_2(s, t')$ for every strategy $t' \in S^2$
- One of the dominant concepts in game theory from 1950s till now
- Most other concepts in noncooperative game theory are variations/extensions/generalizations of Nash equilibria

Pictorially:

t

	(\quad , \quad)	(\quad , \quad)	(x_1, \quad)	(\quad , \quad)	(\quad , \quad)
	(\quad , \quad)	(\quad , \quad)	(x_2, \quad)	(\quad , \quad)	(\quad , \quad)
	(\quad , \quad)	(\quad , \quad)	(x_3, \quad)	(\quad , \quad)	(\quad , \quad)
s	(\quad , y_1)	(\quad , y_2)	(x, y)	(\quad , y_4)	(\quad , y_5)
	(\quad , \quad)	(\quad , \quad)	(x_5, \quad)	(\quad , \quad)	(\quad , \quad)

In order for (s, t) to be a Nash equilibrium:

- x must be greater than or equal to any x_i in column t
- y must be greater than or equal to any y_j in row s

Nash Equilibria

- We should think of Nash equilibria as “stable” profiles of a game
 - At an equilibrium, each player thinks that if the other player does not change her strategy, then he also does not want to change his own strategy
- Hence, no player would regret for his choice at an equilibrium profile (s, t)
 - If the profile (s, t) is realized, pl. 1 sees that he did the best possible, against strategy t of pl. 2,
 - Similarly, pl. 2 sees that she did the best possible against strategy s of pl. 1
- **Attention:** If both players decide to change simultaneously, then we may have profiles where they are both better off

Strategic games in general

Strategic games in general

- A set of **players**

Strategic games in general

- A set of **players**
- Each player has a set of possible **strategies** (actions)

Strategic games in general

- A set of **players**
- Each player has a set of possible **strategies** (actions)
- Each **state of the game** (defined by a strategy per player) yields a **payoff** (or **utility**) to each player

Strategic games in general

- A set of **players**
- Each player has a set of possible **strategies** (actions)
- Each **state of the game** (defined by a strategy per player) yields a **payoff** (or **utility**) to each player
- Given the strategies of the other players, each player aims to select its strategy in order to maximize its utility
 - Such a strategy is called a **best response**

Strategic games in general

- A set of **players**
- Each player has a set of possible **strategies** (actions)
- Each **state of the game** (defined by a strategy per player) yields a **payoff** (or **utility**) to each player
- Given the strategies of the other players, each player aims to select its strategy in order to maximize its utility
 - Such a strategy is called a **best response**
- A state consisting of best responses is stable, and called a **pure Nash equilibrium**: no player would like to deviate and select a different strategy

Examples of finding Nash equilibria in simple games

Example 1: Prisoner's Dilemma

In small games, we can examine all possible profiles and check if they form an equilibrium

- (D, D): both players have an incentive to deviate to another strategy
- (C, D): pl. 1 has an incentive to deviate
- (D, C): Same for pl. 2
- (C, C): Nobody has an incentive to change


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

Hence: The profile (C, C) is the unique Nash equilibrium of this game

– Recall that C is a dominant strategy for both players in this game

Corollary: If s is a dominant strategy of pl. 1, and t is a dominant strategy for pl. 2, then the profile (s, t) is a Nash equilibrium

Example 2: Bach or Stravinsky (BoS)



		Lucy	
		B	S
Charlie	B	2, 1	0, 0
	S	0, 0	1, 2

2 Nash equilibria:

- (B, B) and (S, S)
- Both derive the same total utility (3 units)
- But each player has a preference for a different equilibrium

Example 2a: Coordination games

Variation of Bach
or Stravinsky

	B	S
B	2, 2	0, 0
S	0, 0	1, 1

Again 2 Nash equilibria:

- (B, B) and (S, S)
- But now (B, B) is clearly the most preferable for both players
- Still the profile (S, S) is a valid equilibrium, no player has a unilateral incentive to deviate
 - At the profile (S, S), both players should deviate together in order to reach a better outcome

Example 3: The Hawk-Dove game

		
	2, 2	0, 4
	4, 0	-1, -1

- The most fair solution (D, D) is not an equilibrium
- 2 Nash equilibria: (D, H), (H, D)
- We have a stable situation only when one population dominates or destroys the other

Example 4: Matching Pennies

	H	T
H	1, -1	-1, 1
T	-1, 1	1, -1

- In every profile, some player has an incentive to deviate
- There is no Nash equilibrium!
- Note: The same is true for Rock-Paper-Scissors

Nash dynamics graph

- An easy way to graphically find Nash equilibria
- Built a graph containing a node per state
- A directed edge between two nodes represents the fact that there exists a player with a profitable unilateral deviation
- A node with only incoming edges corresponds to an equilibrium state: no player would like to deviate from there

Battle of the sexes

		man	
		sports	movie
woman	sports	3, 6	1, 1
	movie	2, 2	6, 3

sports, sports

sports, movie

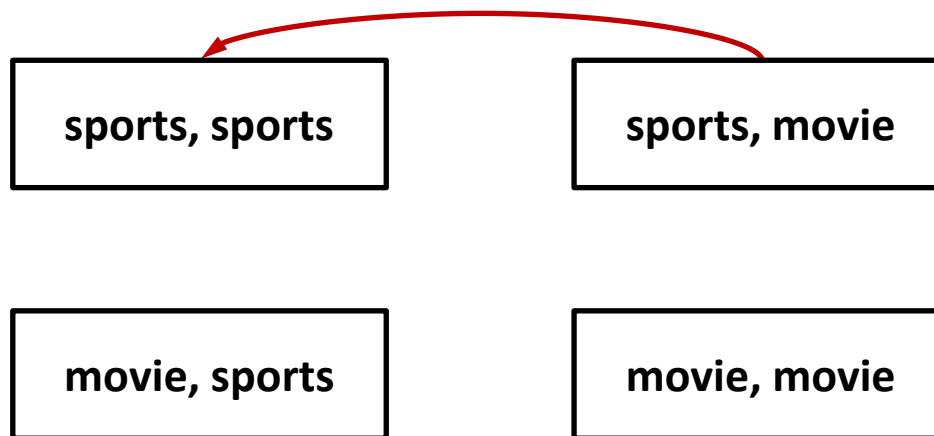
movie, sports

movie, movie

Battle of the sexes

		man	
		sports	movie
woman	sports	3, 6	1, 1
	movie	2, 2	6, 3

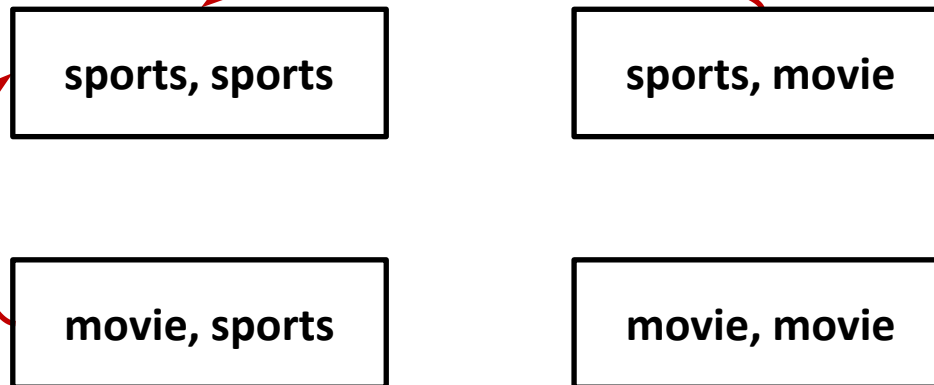
Man improves from 1 to 6



Battle of the sexes

		man	
		sports	movie
woman	sports	3, 6	1, 1
	movie	2, 2	6, 3

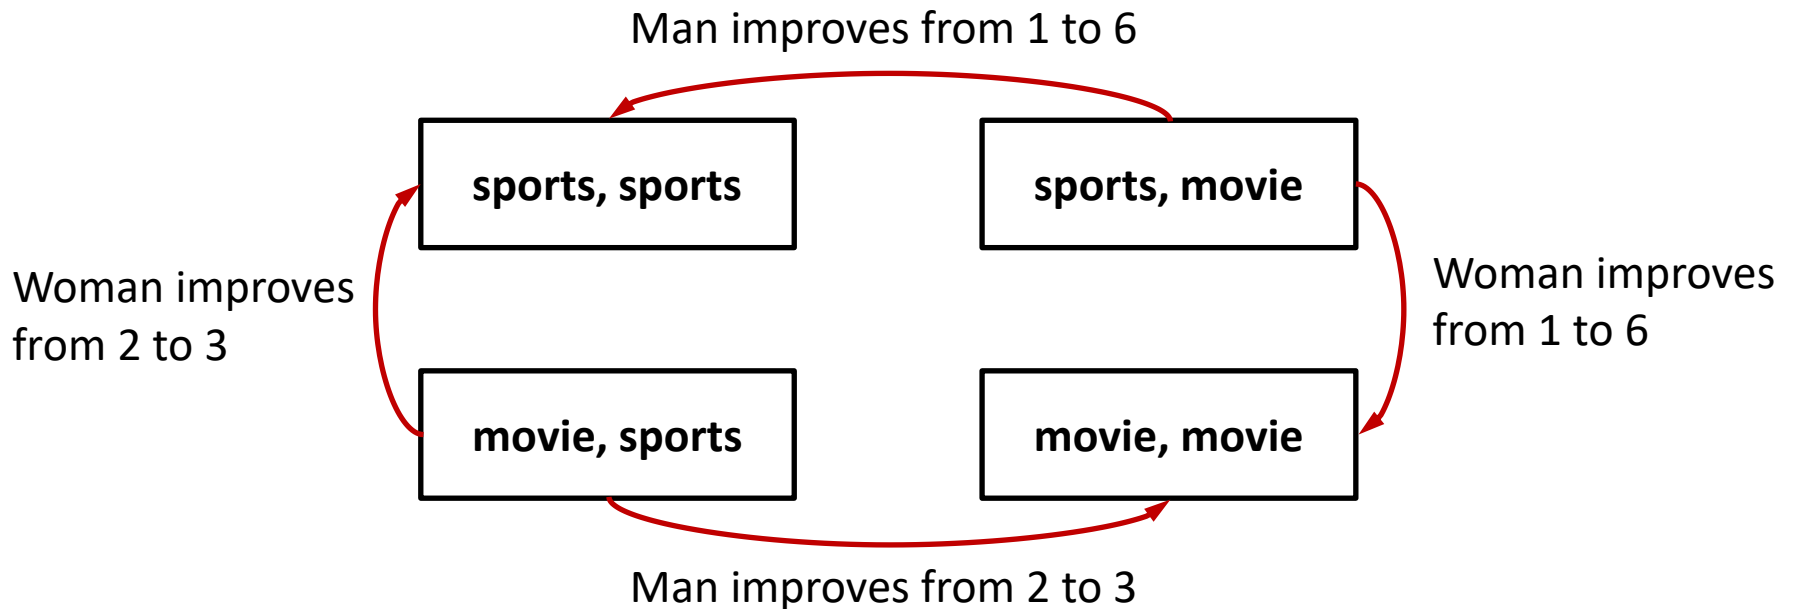
Man improves from 1 to 6



Woman improves from 2 to 3

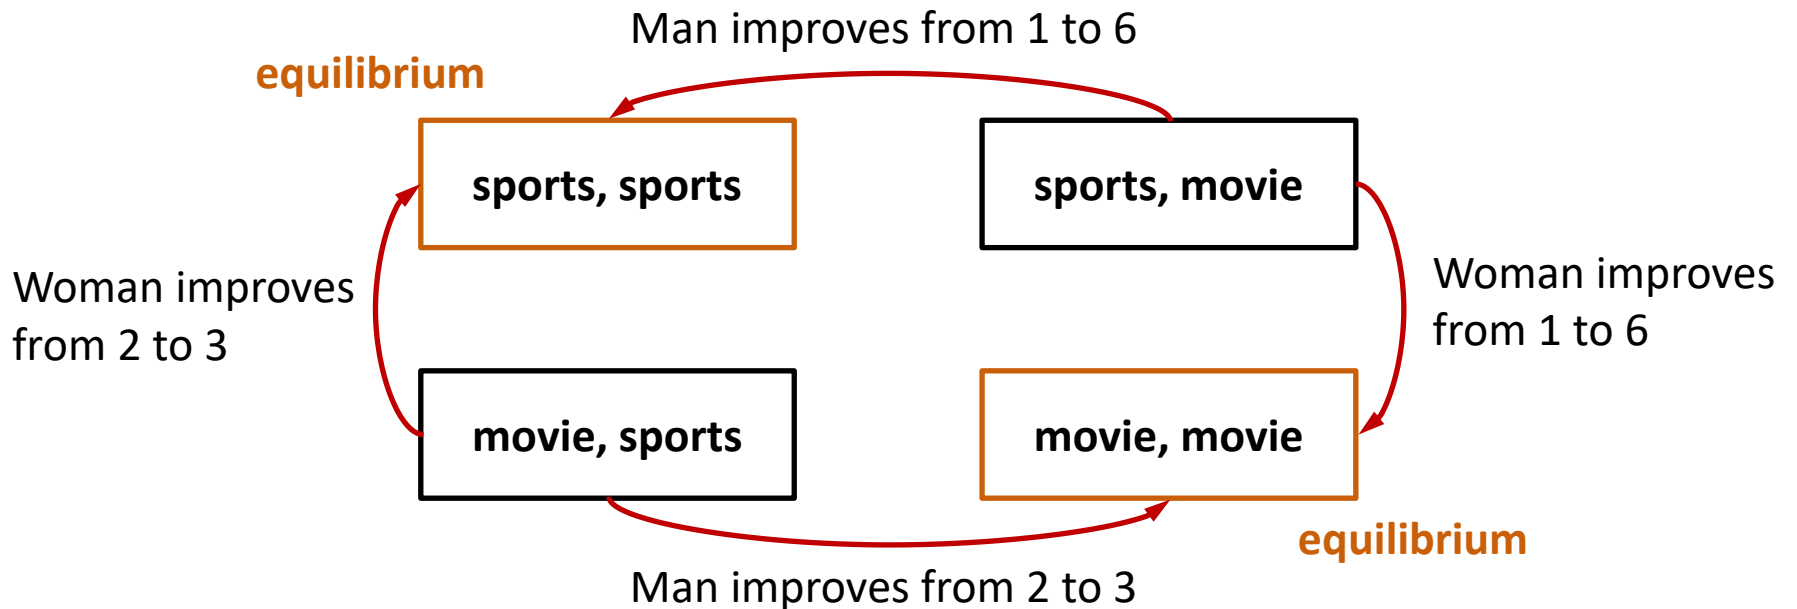
Battle of the sexes

		man	
		sports	movie
woman	sports	3, 6	1, 1
	movie	2, 2	6, 3



Battle of the sexes

		man	
		sports	movie
woman	sports	3, 6	1, 1
	movie	2, 2	6, 3

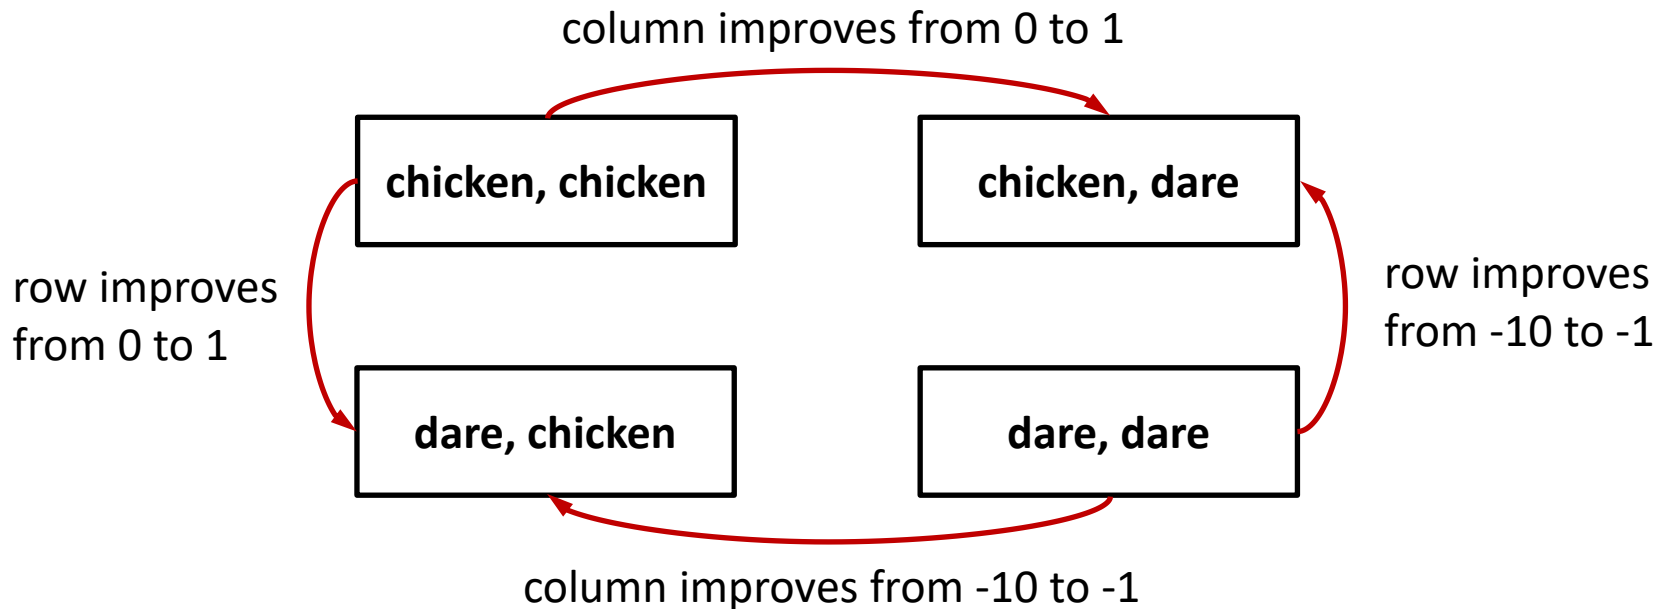


Chicken

		column-driver	
		chicken	dare
row-driver	chicken	0, 0	-1, 1
	dare	1, -1	-10, -10

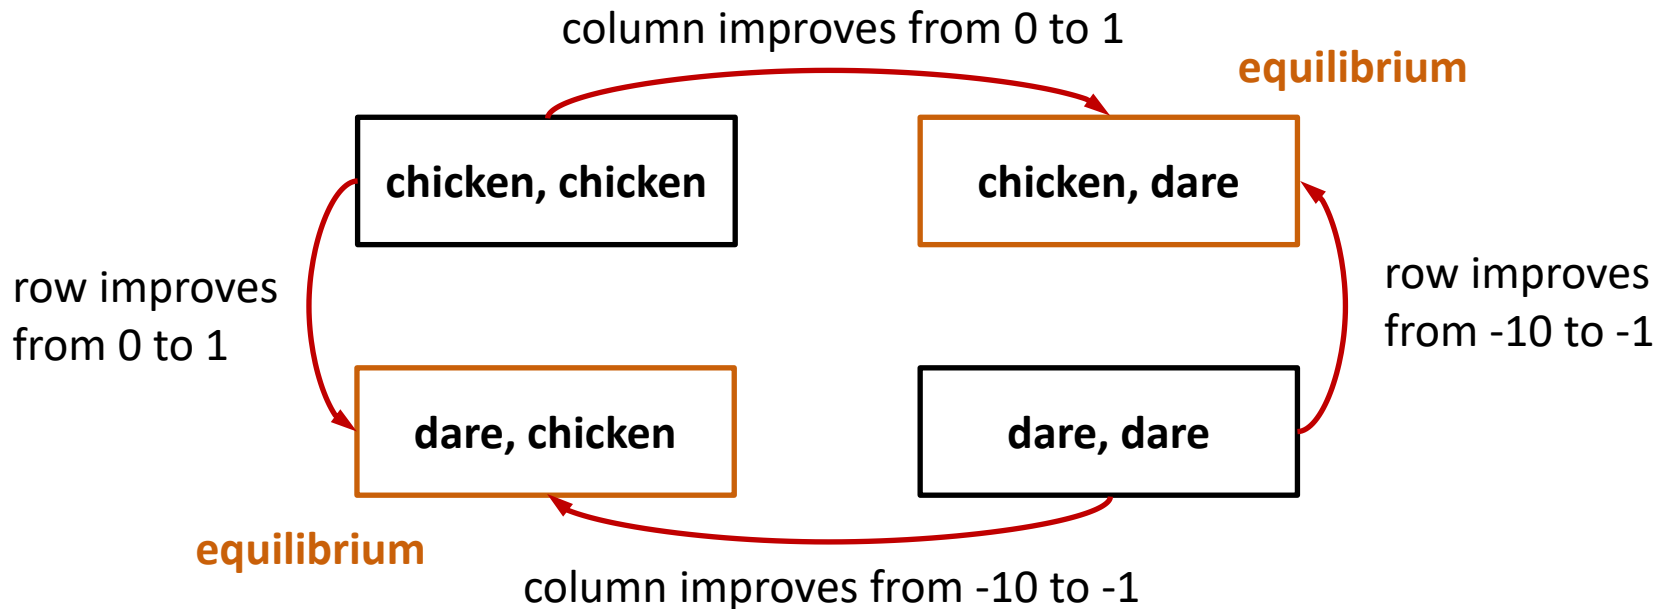
Chicken

		column-driver	
		chicken	dare
row-driver	chicken	0, 0	-1, 1
	dare	1, -1	-10, -10



Chicken

		column-driver	
		chicken	dare
row-driver	chicken	0, 0	-1, 1
	dare	1, -1	-10, -10




Matching pennies

		odd	
		heads	tails
even	heads	1, -1	-1, 1
	tails	-1, 1	1, -1

Matching pennies

		odd	
		heads	tails
even	heads	1, -1	-1, 1
	tails	-1, 1	1, -1



Mixed strategies in games

Existence of Nash equilibria

- We saw that not all games possess Nash equilibria
- E.g. Matching Pennies, Rock-Paper-Scissors, and many others
- What would constitute a good solution in such games?

Example of a game without equilibria: Matching Pennies

	H	T
H	1, -1	-1, 1
T	-1, 1	1, -1

- In every profile, some player has an incentive to change
- Hence, no Nash equilibrium!

Q: How would we play this game in practice?

A: Maybe randomly

Matching Pennies: Randomized strategies

	$\frac{1}{2}$	$\frac{1}{2}$
	H	T
$\frac{1}{2}$ H	1, -1	-1, 1
$\frac{1}{2}$ T	-1, 1	1, -1

- Main idea: Enlarge the strategy space so that players are allowed to play non-deterministically
- Suppose both players play
 - H with probability $\frac{1}{2}$
 - T with probability $\frac{1}{2}$
- Then every outcome has a probability of $\frac{1}{4}$
- For pl. 1:
 - P[win] = P[lose] = $\frac{1}{2}$
 - Average utility = 0
- Similarly for pl. 2

Mixed strategies

- Definition: A **mixed strategy** of a player is a probability distribution on the set of his available choices
- If $S = (s_1, s_2, \dots, s_n)$ is the set of available strategies of a player, then a mixed strategy is a vector in the form
 $\mathbf{p} = (p_1, \dots, p_n)$, where
 $p_i \geq 0$ for $i=1, \dots, n$, and $p_1 + \dots + p_n = 1$
- p_j = probability for selecting the j -th strategy
- We can write it also as $p_j = p(s_j)$ = prob/ty of selecting s_j
- Matching Pennies: the uniform distribution can be written as
 $\mathbf{p} = (1/2, 1/2)$ or $p(H) = p(T) = 1/2$

Pure and mixed strategies

- From now on, we refer to the available choices of a player as *pure strategies* to distinguish them from mixed strategies
- For 2 players with $S^1 = \{s_1, s_2, \dots, s_n\}$ and $S^2 = \{t_1, t_2, \dots, t_m\}$
- Pl. 1 has n pure strategies, Pl. 2 has m pure strategies
- Every pure strategy can also be represented as a mixed strategy that gives probability 1 to only a single choice
- E.g., the pure strategy s_1 can also be written as the mixed strategy $(1, 0, 0, \dots, 0)$
- More generally: strategy s_i can be written in vector form as the mixed strategy $e^i = (0, 0, \dots, 1, 0, \dots, 0)$
 - 1 at position i , 0 everywhere else
 - Some times, it is convenient in the analysis to use the vector form for a pure strategy

Utility under mixed strategies

- Suppose that each player has chosen a mixed strategy in a game
- How does a player now evaluate the outcome of a game?
- We will assume that each player cares for his expected utility
 - Justified when games are played repeatedly
 - Not justified for more risk-averse or risk-seeking players

Expected utility (for 2 players)

- Consider a $n \times m$ game
- Pure strategies of pl. 1: $S^1 = \{s_1, s_2, \dots, s_n\}$
- Pure strategies of pl. 2: $S^2 = \{t_1, t_2, \dots, t_m\}$
- Let $\mathbf{p} = (p_1, \dots, p_n)$ be a mixed strategy of pl. 1 and $\mathbf{q} = (q_1, \dots, q_m)$ be a mixed strategy of pl. 2
- Expected utility of pl. 1:

$$u_1(\mathbf{p}, \mathbf{q}) = \sum_{i=1}^n \sum_{j=1}^m p_i \cdot q_j \cdot u_1(s_i, t_j) = \sum_{i=1}^n \sum_{j=1}^m p(s_i) \cdot q(t_j) \cdot u_1(s_i, t_j)$$

- Similarly for pl. 2 (replace u_1 by u_2)

Example

	B	S
B	2, 1	0, 0
S	0, 0	1, 2

- Let $\mathbf{p} = (4/5, 1/5)$,
 $\mathbf{q} = (1/2, 1/2)$
- $u_1(\mathbf{p}, \mathbf{q}) = 4/5 \times 1/2 \times 2 + 1/5 \times 1/2 \times 1 = 0.9$
- $u_2(\mathbf{p}, \mathbf{q}) = 4/5 \times 1/2 \times 1 + 1/5 \times 1/2 \times 2 = 0.6$
- When can we have an equilibrium with mixed strategies?

Nash equilibria with mixed strategies

- Definition: A profile of mixed strategies (\mathbf{p}, \mathbf{q}) is a **Nash equilibrium** if
 - $u_1(\mathbf{p}, \mathbf{q}) \geq u_1(\mathbf{p}', \mathbf{q})$ for any other mixed strategy \mathbf{p}' of pl. 1
 - $u_2(\mathbf{p}, \mathbf{q}) \geq u_2(\mathbf{p}, \mathbf{q}')$ for any other mixed strategy \mathbf{q}' of pl. 2
- Again, we just demand that no player has a unilateral incentive to deviate to another strategy
- How do we verify that a profile is a Nash equilibrium?
 - There is an infinite number of mixed strategies!
 - Infeasible to check all these deviations

Nash equilibria with mixed strategies

- **Corollary:** It suffices to check only deviations to pure strategies
 - Because each mixed strategy is a convex combination of pure strategies
- Equivalent definition: A profile of mixed strategies (\mathbf{p}, \mathbf{q}) is a **Nash equilibrium** if
 - $u_1(\mathbf{p}, \mathbf{q}) \geq u_1(\mathbf{e}^i, \mathbf{q})$ for every pure strategy \mathbf{e}^i of pl. 1
 - $u_2(\mathbf{p}, \mathbf{q}) \geq u_2(\mathbf{p}, \mathbf{e}^j)$ for every pure strategy \mathbf{e}^j of pl. 2
- Hence, we only need to check $n+m$ inequalities as in the case of pure equilibria

Mixed equilibria

- **Mixed equilibrium:** A profile of *mixed* strategies such that each player maximizes its expected utility, given the strategies of the other players

Mixed equilibria

- **Mixed equilibrium:** A profile of *mixed* strategies such that each player maximizes its expected utility, given the strategies of the other players

Theorem [Nash, 1951]

Every finite strategic game of n players has at least one mixed equilibrium

Mixed equilibria

- **Mixed equilibrium:** A profile of *mixed* strategies such that each player maximizes its expected utility, given the strategies of the other players

Theorem [Nash, 1951]

Every finite strategic game of n players has at least one mixed equilibrium

- Every pure equilibrium is also a mixed equilibrium
 - Every pure strategy can be seen as a probability distribution over all strategies that assigns probability 1 to this one pure strategy