Solving Normal-form Games

Michal Jakob
Agent Technology Center, Dept. of Computer Science and Engineering, FEE, Czech Technical University

AE4M36MAS Autumn 2012 - Lect. 6

OTEVŘENÁ
INFORMATIKA

Outline

- Recapitulation
- Solution concepts for normal-from games
- Computing solution concepts for normal-form games

Recapitulation

Game Theory

- Mathematical study of **interaction** between **rational, self-interested agents**
- Allows
  - **analyzing properties** of such interactions
  - **developing strategies** for agents participating in such interactions

Types of Games

- Cooperative or **non-cooperative**
- **Simultaneous** and sequential
- **Complete** vs. incomplete information
- **Perfect** vs. imperfect information
- **Zero-sum** and **non-zero-sum**
- **Discrete** and continuous games (differential games)

normal form game

Normal-form games

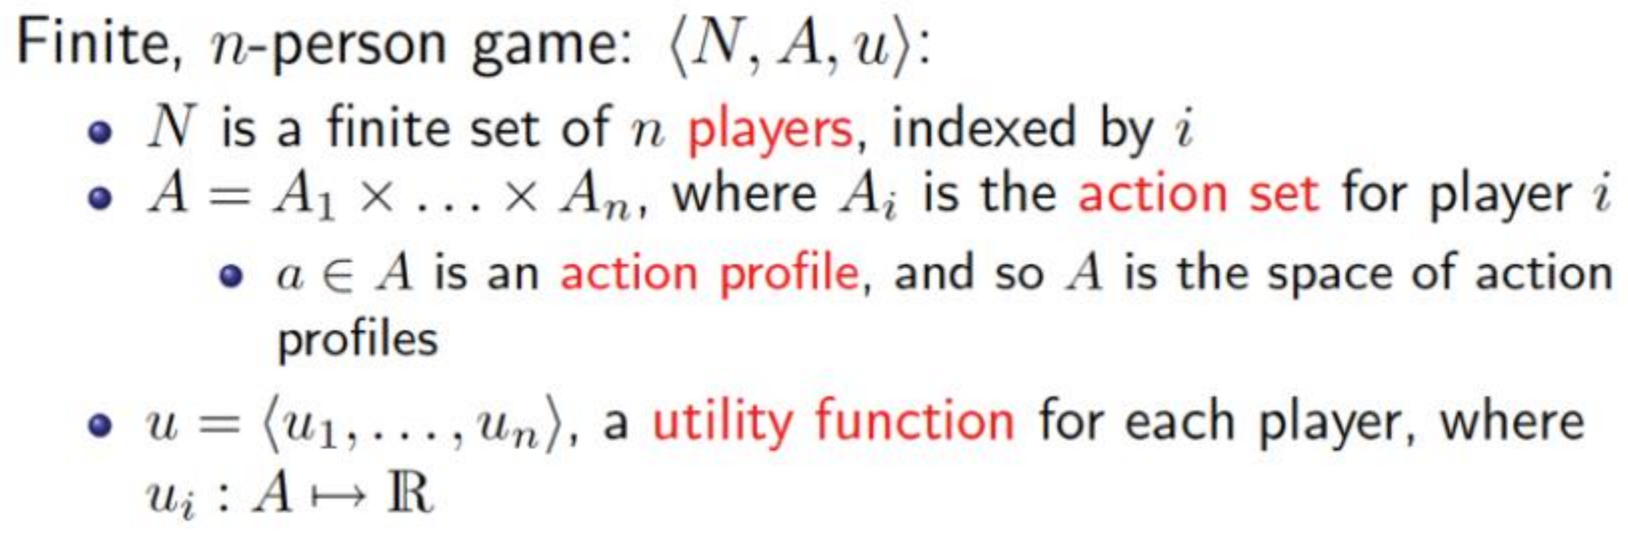

Finite, n-person game: $\langle N, A, u \rangle$:

- 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 for agent i is as any probability distribution over the actions A_i
- Given the strategy profile $s \in S$ for all agents, the utility for the agent i

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

2-player Normal Form Game

- Can be written 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	3, 3	0, 5
D	5, 0	1, 1

Example: Prisoner’s dilemma

Solution Concepts

Reasoning about Games

- In single-agent decision theory, look at an **optimal strategy**
  - Maximize the agent's expected payoff in its environment
- With multiple agents, this is not possible <= **the best strategy depends on others**' choices
- Deal with this by identifying certain subsets of outcomes called **solution concepts**
- **Solution concept**: a subset of game outcomes that are somehow interesting.

Solution Concepts

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

Existence? Uniqueness? Complexity?
Depend on game properties.

Nash Equilibrium

- Definition (**Best Reponse**)

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

- Definition (**Nash Equilibrium**)

  The strategy profile $a = \langle a_1, \dots, a_n \rangle$ is a Nash Equilibrium if a_i profile is a best response to a_{-i} for every i

- Definition (**Strict Nash Equilibrium**)

  The strategy profile $a = \langle a_1, \dots, a_n \rangle$ is a Strict Nash Equilibrium if a_i profile is the only best response to a_{-i} for every i

- Definition (**Weak Nash Equilibrium**)

  Weak Nash equilibrium is a Nash equilibrium which is not strict

Properties of Nash Equilibria

- **Weak Nash** equilibria are **less stable** than **strict Nash** equilibria
  - In the weak case, at least one agent has > 1 best responses, and only one of them is in the Nash equilibrium
- **Pure-strategy Nash** equilibria can be either **strict** or **weak**
- **Mixed-strategy Nash** equilibria are **always weak**
  - Reason: if there are ≥ 2 pure strategies that are best responses to a_{-i} then any mixture of them is also a best response
- Thus in a strict Nash equilibrium, **all strategies are pure**

Example: Routing

- 1,000 drivers wish to travel from *S* (start) to *D* (destination)
- Two possible paths: *S*→*A*→*D* and *S*→*B*→*D*
- The road from *S* to *A* is long: *t* = 50 minutes
  - but it's also very wide: *t* = 50 minutes, no matter how many drivers
- Same for road from *B* to *D*
- Road from *A* to *E* is shorter but is narrow
  - Time = (number of cars)/25
- Nash equilibrium:
  - 500 cars go through A, 500 cars through B
  - Everyone's time is 50 + 500/25 = 70 minutes
  - If a single driver changes to the other route
    - There now are 501 cars on that route, so his/her time goes up

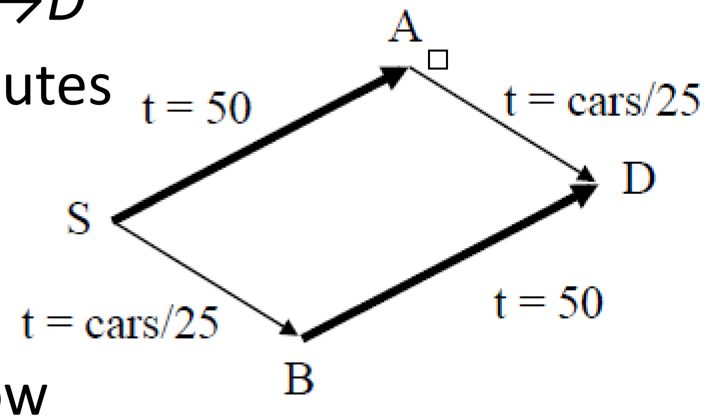

Braess's Paradox

- 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

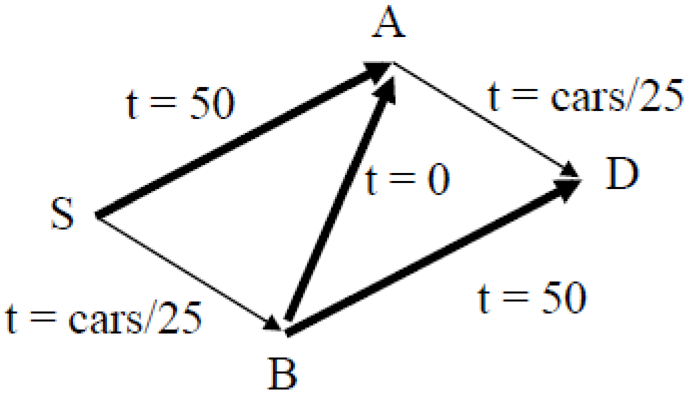

- To see that this is an equilibrium:
  - If driver goes *S→A→D*, his/her cost is 50 + 40 = 90 minutes
  - If driver goes *S→B→D*, his/her cost is 40 + 50 = 90 minutes
  - Both are dominated by S→B→A→D
- To see that it's the ***only*** Nash equilibrium:
  - For every traffic pattern, S→B→A→D dominates *S→A→D* and *S→B→D*
  - Choose any traffic pattern, and compute the times a driver would get on all three routes
- Carelessly adding capacity can actually be hurtful!

Deductive vs. Stable State Interpretation

- 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 i's **maxmin strategy** is a strategy that maximizes i's **worst-case payoff**, in the situation where all the other players (whom we denote $-i$) happen to play the strategies which cause the greatest harm to i.
- The **maxmin value** (or **safety level**) of the game for player i is that minimum amount of payoff guaranteed by a maxmin strategy.
- Good choice for a conservative agent: Maximize his/her expected utility without any assumptions about the others

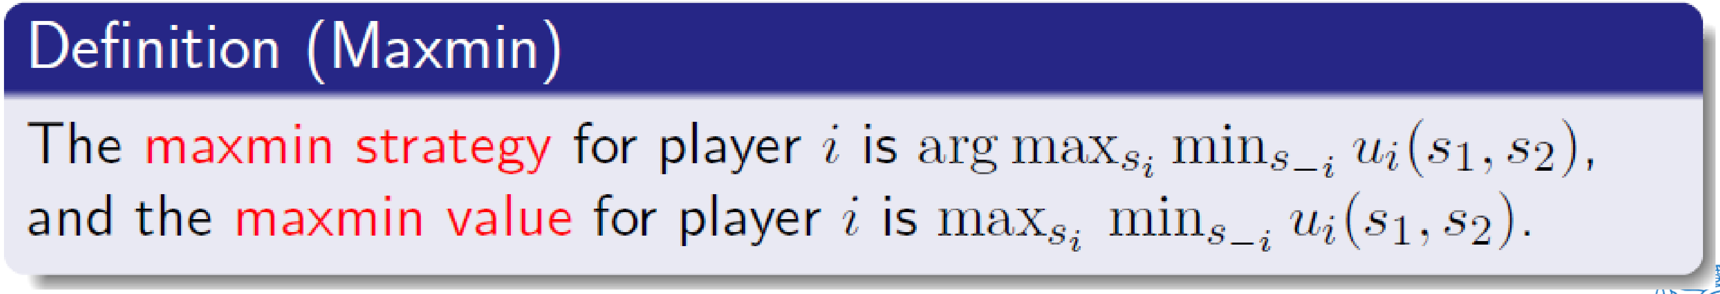

Definition (Maxmin)

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.

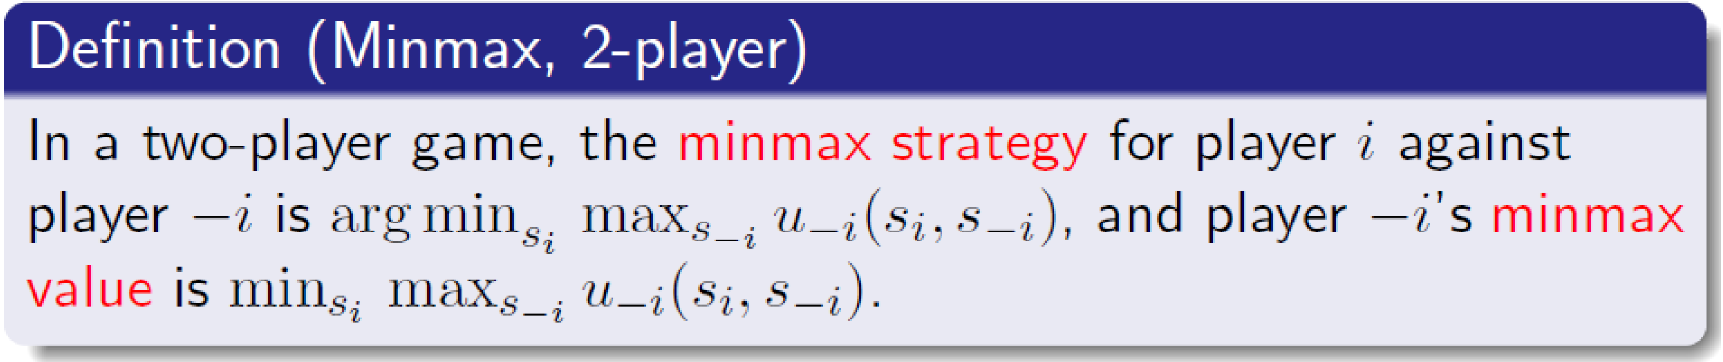

Definition (Minmax, 2-player)

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 Equilibrium in 2-player Zero-Sum Game

- 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
- 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

Finding an Equilibrium Point – 2x2 Zero Sum Game

- *(Following 20 slides adapted from: http://faculty.mc3.edu/cvaughen/mgf1107/game_theory/part3b.ppt)*
- The approach to find mixed strategy equilibrium points is based on the same reasoning as determining an equilibrium point in pure strategies.
- For each player, we are finding the best payoff each player can expect *assuming best play by the opponent*.
- The strategies associated with the best payoff each player can expect assuming best play by the opponent *is the equilibrium point.*
- In a zero sum game, the value of the game is the expected payoff to the row player at the equilibrium point.

Finding an Equilibrium Point – 2x2 Zero Sum Game

- Let’s find the equilibrium point for the following 2x2 zero sum game:

- When a pitcher and batter face each other in a baseball game, we can consider each pitch as a simultaneous move zero-sum game.

- To make this an easier introductory example, let’s suppose the pitcher only has two strategies: throw a curve or a fastball. And, assuming the batter commits to swing, let’s assume the batter only has two strategies: swing expecting a fastball or swing expecting a curve ball. One final assumption that is not entirely unrealistic: This is a simultaneous move game because even though the pitcher commits to his strategy before the batter, we’ll assume the batter does not have time to change strategies once the pitcher commits.

- The outcomes we associate with each pair of strategies is not whether the ball is hit or not but rather the probability the batter has of hitting the ball. Of course, the pitcher wants to minimize this probability and the batter wants to increase it.

Finding an Equilibrium Point – 2x2 Zero Sum Game

- For this example, for a given pitcher and batter, suppose the following outcomes for each strategy:
  - If the pitcher chooses to throw a curve, and the batter expects a curve, the batter will hit the ball 40% of the time. (Which is a 0.400 average)
  - If the pitcher chooses to throw a curve, but the batter is expecting a fastball, then the batter will hit the ball only 20% of the time (a .200 average)
  - If the pitcher chooses fastball while the batter is expecting fastball, then the batter has a 30% chance of connecting (hits for an average of .300).
  - When the pitcher chooses fastball but the batter is expecting a curveball, the batter has only a .100 average (hits the ball 10% of the time).

Finding an Equilibrium Point – 2x2 Zero Sum Game

- We put this game into strategic form as follows:

Pitcher

		Fastball	Curve
Batter	Fastball	.300	.200
	Curve	.100	.400

- Notice that this game is not fair. Because every payoff is positive it is impossible (even with mixed strategies) that the value of the game could be zero. We'd say the game is not fair in the *technical sense of game theory*. Of course the game is fair in a general sense, first because this is based on the rules of the game, and second the game would be very boring if every player had a zero batting average.

Finding an Equilibrium Point – 2x2 Zero Sum Game

- What strategy (mixed or pure) should each player (batter and pitcher) adopt to *optimize* their payoff?

Pitcher

Batter

	Fastball	Curve
Fastball	.300	.200
Curve	.100	.400

- We can find the equilibrium point quickly and easily if it corresponds to pure strategies – so we check that possibility first …

- Determine the maximin strategy for the row player.

- Determine the minimax strategy for the column player.

Finding an Equilibrium Point – 2x2 Zero Sum Game

- We are first checking for an equilibrium point in pure strategies …

Pitcher

		Fastball	Curve	
Batter	Fastball	.300	.200	.200
	Curve	.100	.400	.100
		.300	.400	

First, find the minima for the row player. Now what is the best he can do assuming the pitcher will play to achieve these values? The *maximin* strategy for the row player is .200. This is the “best of the worst outcomes.”

Finding an Equilibrium Point – 2x2 Zero Sum Game

Pitcher

		Fastball	Curve	
Batter	Fastball	.300	.200	.200
	Curve	.100	.400	.100
		.300	.400	

Now we search for the best strategy for the column player. Actually, it would not matter if we had done the column player first, and then the row player.

Find the column maxima. Then find the minimum value of the column maxima. This outcome is associated with the *minimax* strategy for the column player.

Finding an Equilibrium Point – 2x2 Zero Sum Game

Pitcher

Batter

	Fastball	Curve	
Fastball	.300	.200	.200
Curve	.100	.400	.100
	.300	.400	

Notice that the minimax and maximin strategies are not equal.

For the batter, by choosing always to swing for a fastball, the worst he can do is average a .200.

For the pitcher, by choosing always to throw a fastball, the worst he can do is give up a .300 average for this batter.

Finding an Equilibrium Point – 2x2 Zero Sum Game

Pitcher

Batter

	Fastball	Curve	
Fastball	.300	.200	.200
Curve	.100	.400	.100
	.300	.400	

Because the maximin and minimax strategies do not equal, the equilibrium point is not found with pure strategies.

That the equilibrium is not in pure strategies is clear because, if the worst that the batter could do, by always swinging for a fastball, is a .200 average, but would actually get a .300 average if the pitcher actually always threw fastballs, we can see that the pitcher would benefit from throwing some curves.

Finding an Equilibrium Point – 2x2 Zero Sum Game

Pitcher

Batter

	Fastball	Curve	
Fastball	.300	.200	.200
Curve	.100	.400	.100
	.300	.400	

The question becomes: How often should the pitcher throw curves to lower the batter’s probability of hitting? And then: What is the best strategy for the batter assuming best play by the pitcher?

The batter’s worst case average is .200 and the pitcher’s worst case average is .300 using pure strategies. Now by mixing strategies, each player will try to get as much of that difference as possible.

What strategy should each adopt assuming best play by the opponent?

Finding an Equilibrium Point – 2x2 Zero Sum Game

		Pitcher	
		p	1-p
		Fastball	Curve
Batter	Fastball	.300	.200
	Curve	.100	.400

The following procedure finds the equilibrium point resulting from mixed strategies from one or both players…

Let’s find the optimal mixed strategy for the pitcher first.

Let p equal the probability the pitcher will throw a fastball thus 1-p equals the probability he will throw a curve.

We write the *expected value* (which is the average payoff in repeated trials) for the batter based on the probabilities p and 1-p, for each of the batter’s strategies.

$E_F = .300p + .200(1-p)$ is the expected value the batter would receive choosing fastball.

$E_C = .100p + .400(1-p)$ is the expected value the batter would receive choosing to swing for a curveball.

Finding an Equilibrium Point – 2x2 Zero Sum Game

E_F = .300p + .200(1-p) is the expected value the batter would receive choosing fastball.

E_C = .100p + .400(1-p) is the expected value the batter would receive choosing to swing for a curveball.

The pitcher wants to find the probability p, for throwing a fastball, that will minimize the batter's expected value.

We simplify each expression above, which gives expected value (for each strategy) as a function of probability, p, graph each function, and find a minimum value, as follows:

$E_F = .300p + .200(1-p) = .3p + .2 - .2p = .1p + .2$

Likewise,

$E_C = .100p + .400(1-p) = .1p + .4 - .4p = -.3p + .4$

Note that these expected value functions are always linear in p, and are thus easy to graph …

Finding an Equilibrium Point – 2x2 Zero Sum Game

We have expected value functions for each strategy of the batter, for a given choice of p for the pitcher. These functions are: $E_F = .1p + .2$ and $E_C = -.3p + .4$

Graphing these functions, we find …

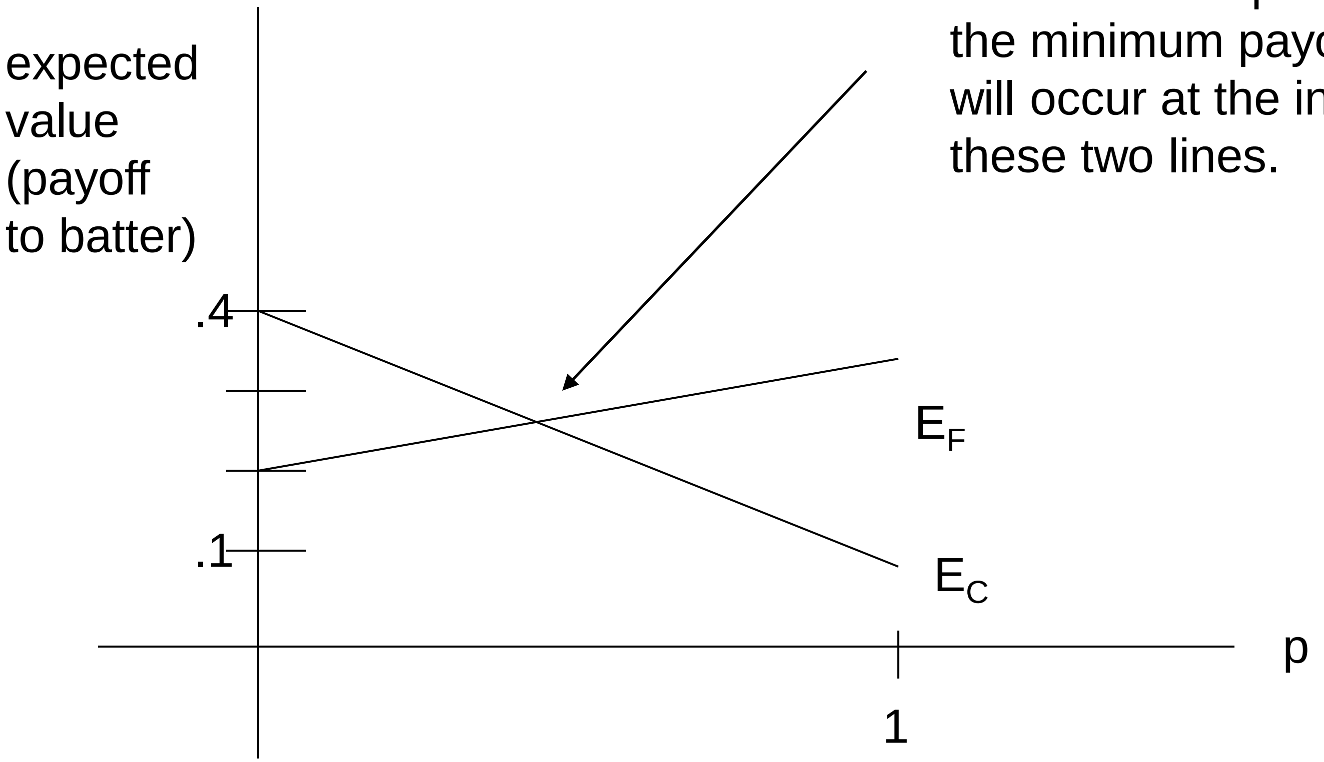

The choice of p that corresponds to the minimum payoff for the batter will occur at the intersection of these two lines.

Finding an Equilibrium Point – 2x2 Zero Sum Game

We have expected value functions for each strategy of the batter, for a given choice of p for the pitcher. These functions are: $E_F = .1p + .2$ and $E_C = -.3p + .4$

Graphing these functions, we find …

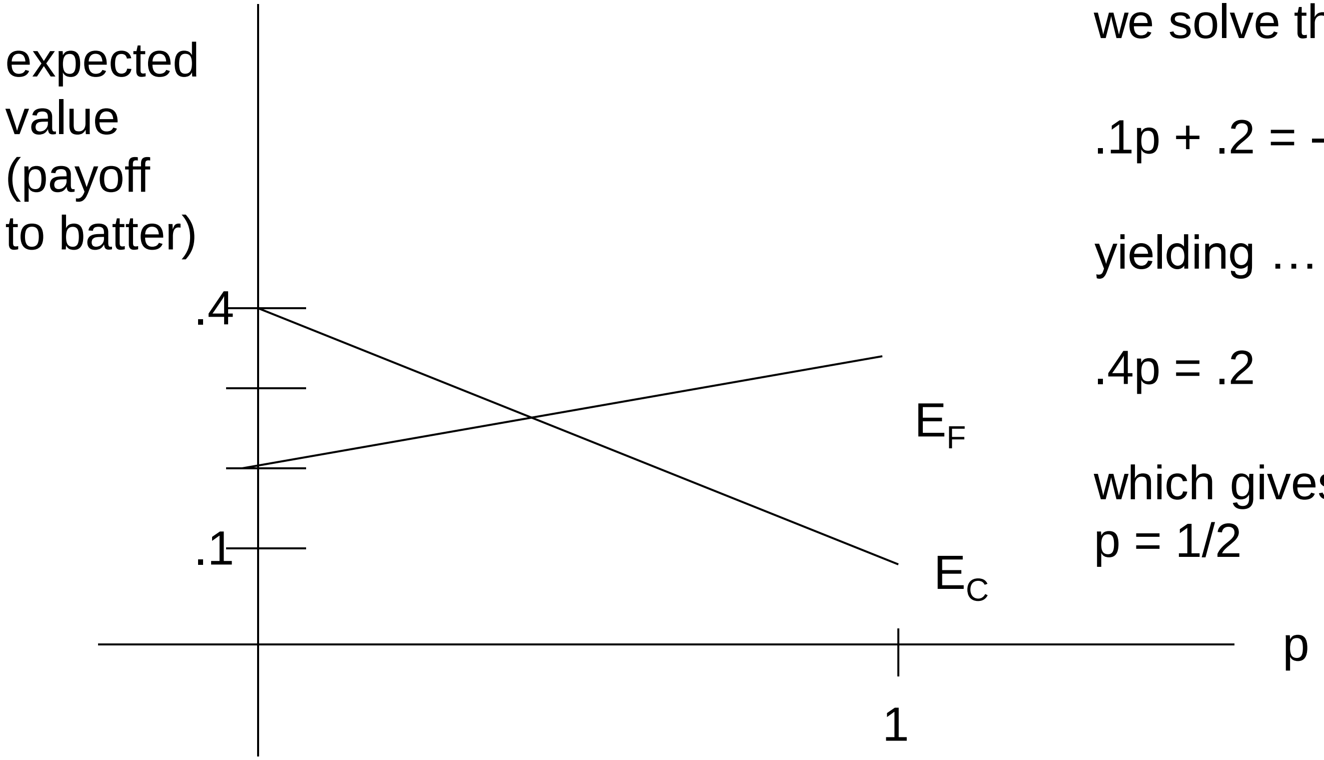

to find the intersection point, we solve the equation:

$.1p + .2 = -.3p + .4$

yielding …

$.4p = .2$

which gives a solution of $p = 1/2$

Finding an Equilibrium Point – 2x2 Zero Sum Game

We have expected value functions for each strategy of the batter, for a given choice of p for the pitcher. These functions are: $E_F = .1p + .2$ and $E_C = -.3p + .4$

Graphing these functions, we find …

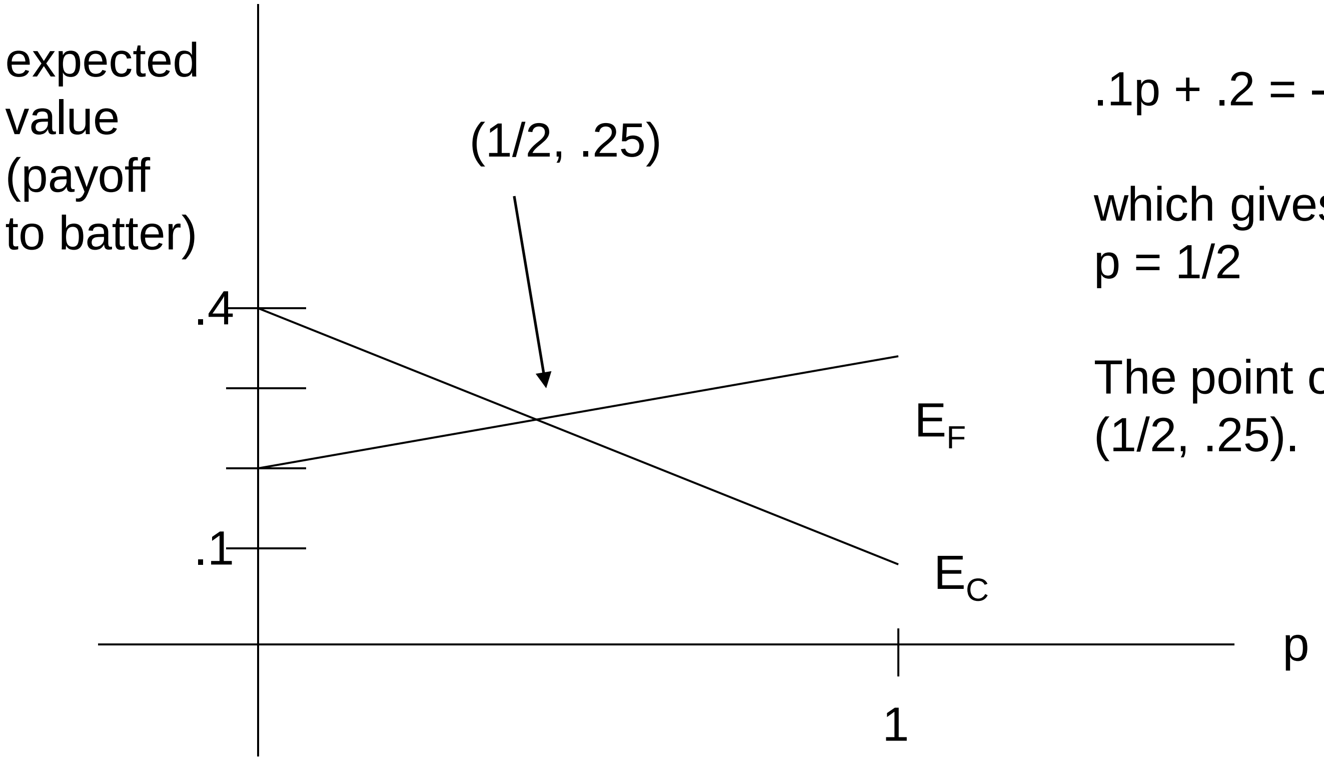

to find the intersection point, we solve the equation:

$.1p + .2 = -.3p + .4$

which gives a solution of
$p = 1/2$

The point of intersection is (1/2, .25).

Because the payoff of .25 (a batting average of .250) corresponds to the batter's (row player's) payoff, this is the value of the game.

Finding an Equilibrium Point – 2x2 Zero Sum Game

The conclusion we have is that the pitcher should throw fastballs with probability p = 1/2 and curveballs with probability 1-p = 1-1/2 = 1/2 to minimize the batter’s average payoff.

We also have found the value of the game, which is .250.

Put another way, we found a mixed strategy of (1/2, 1/2) for fastballs and curveballs that the pitcher can use that will guarantee as low a payoff for the batter assuming best play by the batter.

But we have not yet determined the strategy the batter should use to maximize payoff assuming best play by the pitcher.

This is done the same way for the batter as it was done for the pitcher but the probabilities associated with an optimal mixed strategy may be different for each player. Nevertheless, assuming best play by each player, there is exactly one outcome, which is the value of the game.

We’ll now find the optimal mixed strategy for the batter…

Finding an Equilibrium Point – 2x2 Zero Sum Game

Now for the batter's best-play strategy…

Let q be the probability for the batter of swinging for a fast ball and thus 1-q will be the probability for the batter of swinging for a curve.

Pitcher

			Fastball	Curve
batter	q	Fastball	.300	.200
	1-q	Curve	.100	.400

We'll calculate expected value functions for the pitcher for each strategy which depend on the value of q. Then we'll determine the value of q that is best for the batter assuming best play by the pitcher.

Finding an Equilibrium Point – 2x2 Zero Sum Game

Pitcher

batter

		Fastball	Curve
q	Fastball	.300	.200
1-q	Curve	.100	.400

If the pitcher chooses the fastball strategy, the pitcher’s expected payoff will be

$$E_F = .300q + .100(1-q) = .3q + .1 - .1q = .2q + .1$$

And if the pitcher chooses the curveball strategy, the pitcher’s resulting expected value will be

$$E_C = .200q + .400(1-q) = .2q + .4 - .4q = -.2q + .4$$

Finding an Equilibrium Point – 2x2 Zero Sum Game

Expected payoff value's for the pitcher, for each pure strategy, are as follows:

$E_F = .2q + .1$ and $E_C = -.2q + .4$

We seek the value of q that maximizes payoff to the batter (the worst case for the pitcher). That is, what should the batter choose assuming best play by the pitcher.

Graphing these functions, over values of q from q = 0 to q = 1, we have…

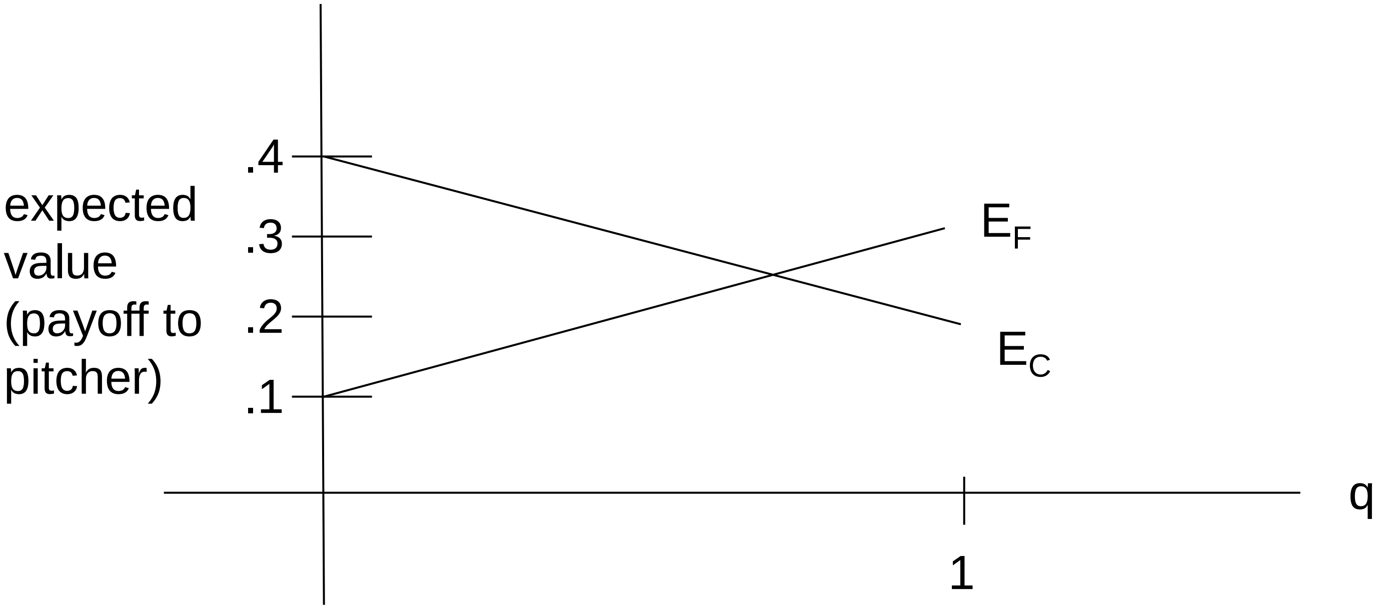

Finding an Equilibrium Point – 2x2 Zero Sum Game

Again, the choice of q that will maximize payoff for the batter assuming best play by the pitcher will occur at the intersection of these lines. Solving .2q + .1 = -.2q + .4 we get .4q = .3 and thus q = 3/4.

Therefore, the strategy for swinging for a curve is 1 – q = 1 – 3/4 = 1/4.

Conclusion: The batter should mix strategies of swinging for a fastball and a curve ball with probability 3/4 and 1/4, respectively.

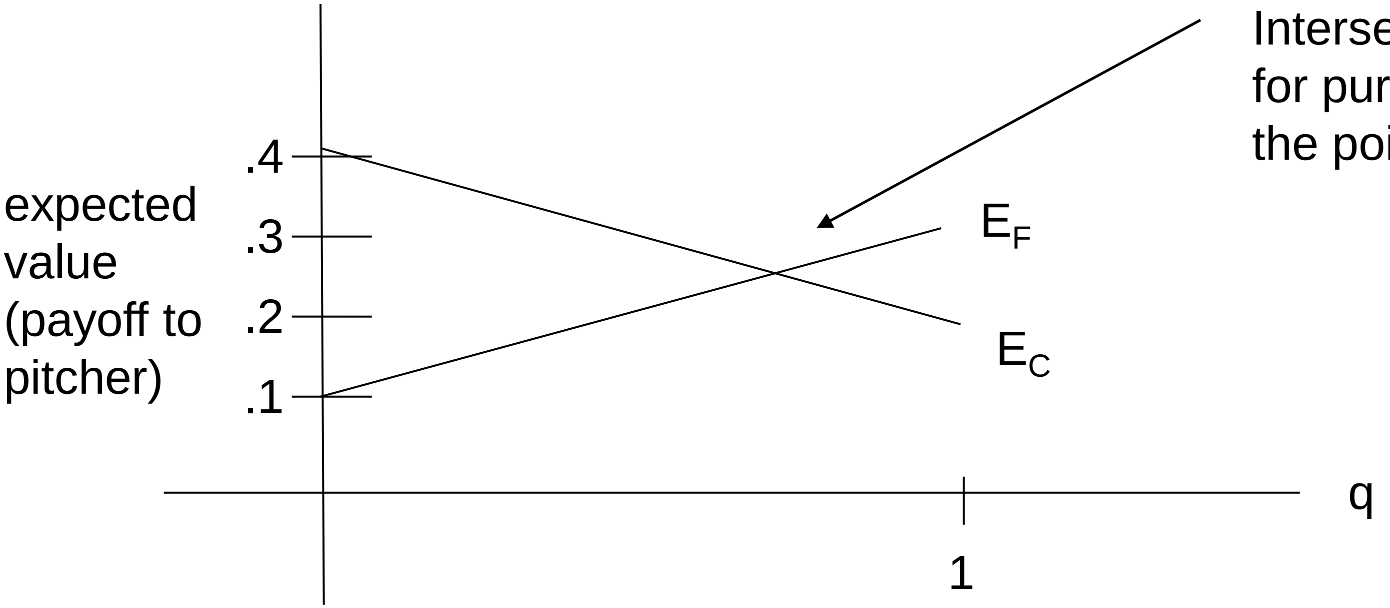

Intersection of payoffs for pure strategies at the point (3/4, 1/4).

Finding an Equilibrium Point – 2x2 Zero Sum Game

We have found what the Minimax Theorem guaranteed exists: the equilibrium point for a 2x2 zero-sum matrix game.

In this case that equilibrium point is found at mixed strategies for both players.

The equilibrium point is the combination of strategies with probabilities (1/2, 1/2) for the pitcher and (3/4, 1/4) for the batter of pure strategies of fastball and curveball, respectively.

That is, these probabilities will yield each player the best possible payoff assuming best play by the opponent. The pitcher should mix fastballs and curveballs with probabilities 1/2 and 1/2 for each. The batter will maximize his average by swinging for fastballs 3/4 of the time and swinging for curves 1/4 of the time.

The value of this particular game is the payoff to the row player (the batter) which is .250.

Domination

- Let s_i and s_i' be two strategies for player i, and let S_{-i} be the set of all possible strategy profiles for the other players

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

Dominated / Dominating Strategy

- A strategy is strictly (resp. weakly; very weakly) **dominant** for an agent if it strictly (weakly; very weakly) dominates any other strategy for that agent
- A strategy s_i is strictly (weakly; very weakly) **dominated** for an agent i if some other strategy s_i' strictly (weakly; very weakly) dominates s_i

Dominant Strategy and NE

- 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 1: Prisoners Dilema

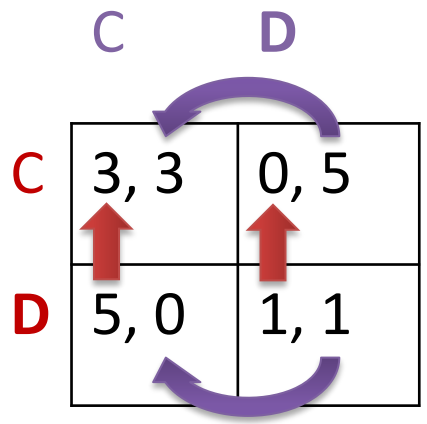

	C	D
C	3, 3	0, 5
D	5, 0	1, 1

- 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 2: Matching Pennies

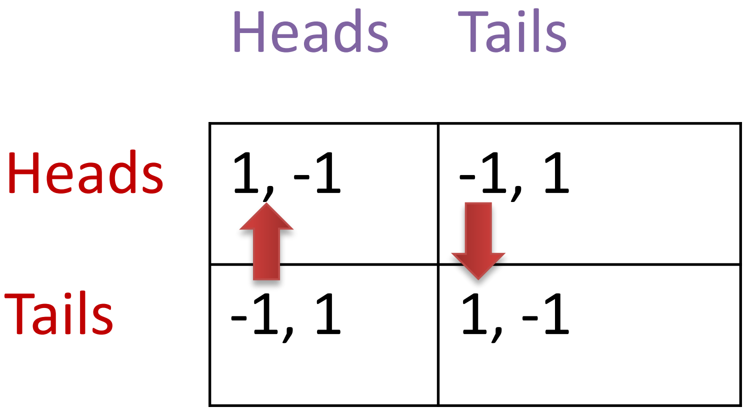

	Heads	Tails
Heads	1, -1	-1, 1
Tails	-1, 1	1, -1

- 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

Correlated Equilibrium

- Consider again Battle of the Sexes.
  - Intuitively, 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

- What is the natural solution here?
  - A traffic light: a fair randomizing device that tells one of the agents to go and the other to wait.
- Benefits:
  - the negative payoff outcomes are completely avoided
  - fairness is achieved

Correlated Equilibrium

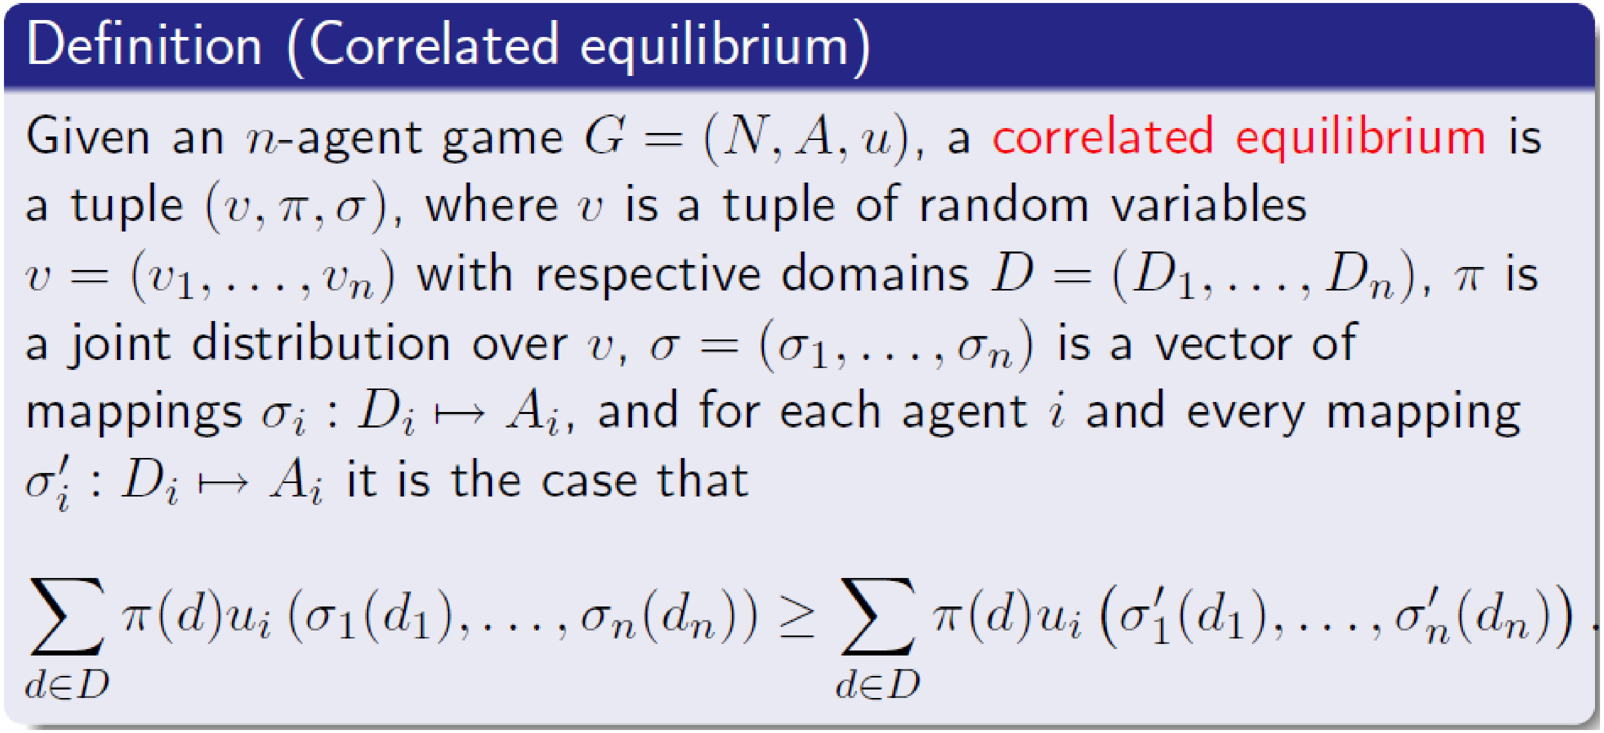

Definition (Correlated equilibrium)

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 $v = (v_1, \ldots, v_n)$ with respective domains $D = (D_1, \ldots, D_n)$, π is a joint distribution over v, $\sigma = (\sigma_1, \ldots, \sigma_n)$ 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

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

- 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.

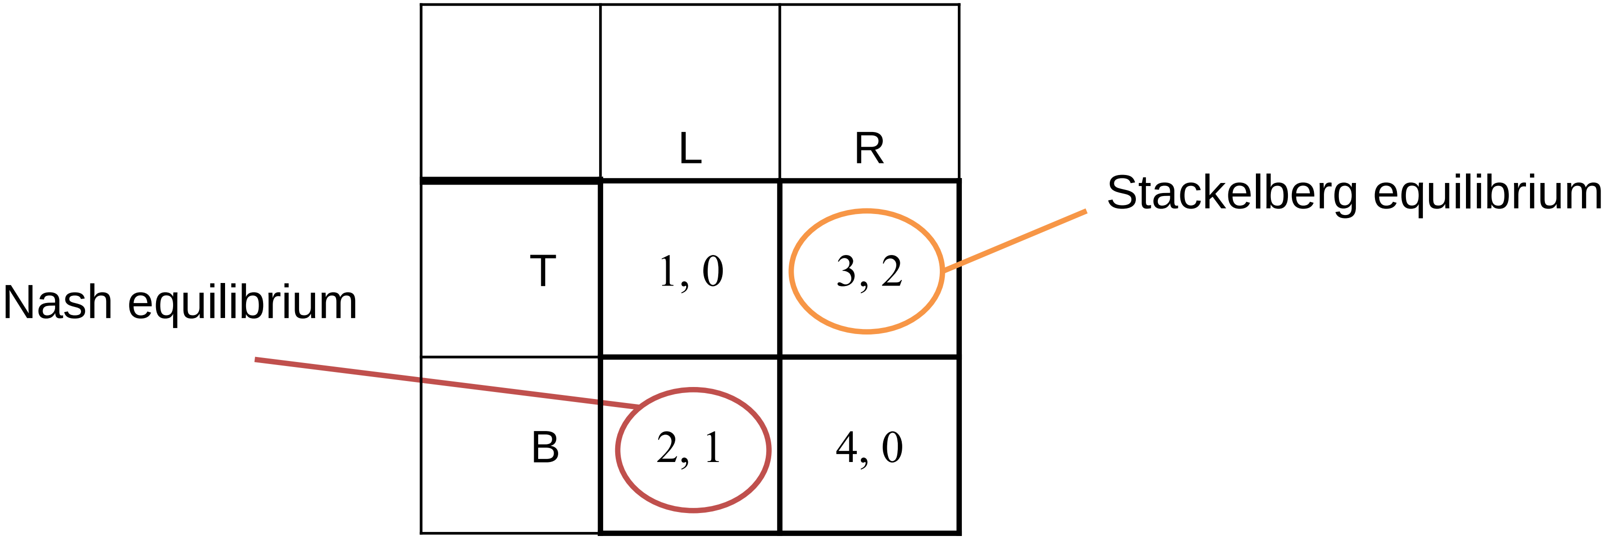

	L	R
T	1, 0	3, 2
B	2, 1	4, 0

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

Additional Solution Concepts

- ϵ-Nash equilibrium
- perfect (trembling hand) equilibrium
- rationalizable strategies
- …

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

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

$$\begin{aligned} &\text{maximize} && \sum_{i=1}^{n} w_i x_i && \\ &\text{subject to} && \sum_{i=1}^{n} a_{ij} x_i \leq b_j && \forall j = 1 \ldots m \\ & && x_i \in \mathbb{R} && \forall i = 1 \ldots n \end{aligned}$$

- 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

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

- 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 Nash 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)

Computing Equilibria of General-Sum *n*-player Games

- 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 Strategies 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~~

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

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 Strategies

- **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 Concepts

- 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

Summary

- Optimal strategy cannot be easily defined in a mutli-agent case => solution concepts: interesting subset out outcomes
  - Nash equilibrium, Maxmin, Dominant strategies, Correlated Nash equilibrium (and other exist)
- Existence and uniqueness of solution concepts differ and depend on the game
- NE provides good guidance for what two player in zero-sum games; in general-sum games, situation is more difficult
- NE in zero-sum games can be found in polynomial time using an LP formulation
- Removal of dominated strategies can reduce the game without changing its strategic properties
- Reading: Shoham 3.3., 3.4, 4.1, 4.5