

Proof:

\Rightarrow : \checkmark (follows from Def. of SPE)

\Leftarrow : We will prove that if s^* is not a SPE, then there exists one-step deviation in a subgame directly at the root.

Suppose s^* is not a SPE. There is a history h and a player i such that s_i is a profitable deviation for player i in the subgame $\Gamma(h)$.

WLOG, the number of histories h' with

$s_i(h') \neq s_i^*(h')$ is at most $l(\Gamma(h))$.

Hence, the number of changes is finite (because of the finite horizon assumption).

\Rightarrow So $\Gamma(h, h^*)$ is the desired subgame, where one-step deviation improves the utility. \square

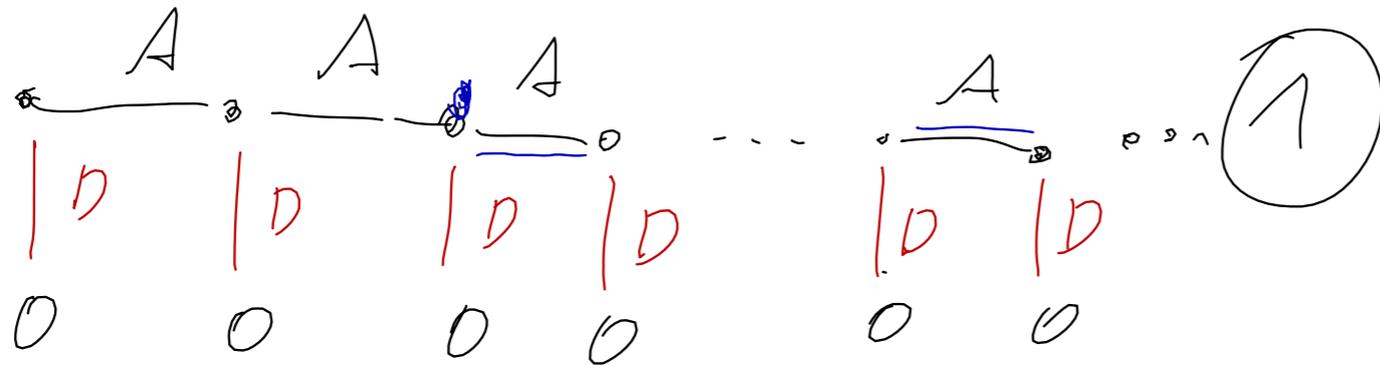
Choose a deviant strategy s_i with a minimal number of changes (this strategy must exist). Let h^* be the longest history in $T(h)$ with $s_i(h^*) \neq s_i|_h(h^*)$, i.e., the "deepest" deviation point of s_i .

Then in $T(h, h^*)$, $s_i|_{h^*}$ differs ^{from} $s_i|_{h, h^*}$ only in the initial history.

Because we assumed that s_i has the minimal number of deviations, it follows that $s_i|_{h^*}$ is really profitable.

In infinite horizon games the one deviation property does not hold!

Can we ever play (1 player game):



Strategy $s_1(h) = D$ for all $h \in H \setminus z$

Is this an SPE? \rightarrow it is not!

For all histories h' , there is no profitable deviation from s_1/h' that gives less than 0.

Strategy $s_1^*(h) = A$ for all $h \in H \setminus z$
 s_1^* is a SPE!

Theorem (Kuhn)

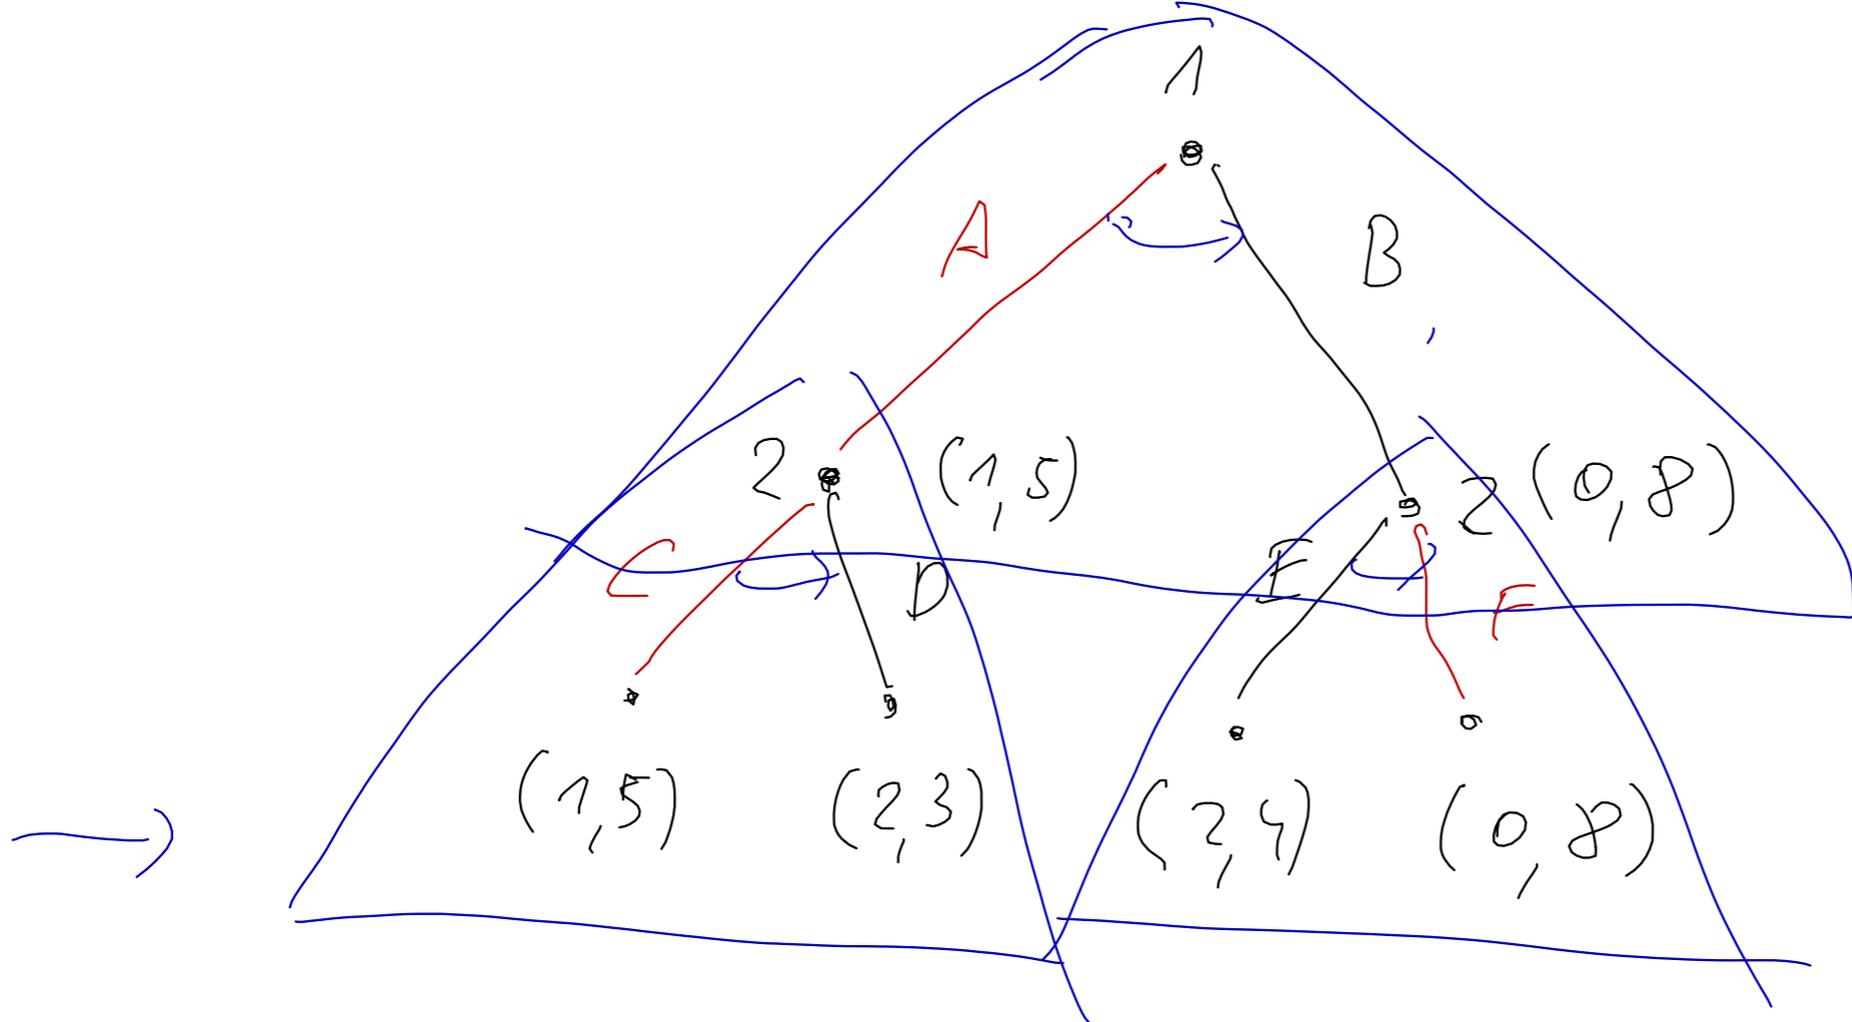
Every finite extensive with perfect information has a SPE.

Proof idea:

- It's constructive. We build an SPE bottom-up (aka backward-induction)

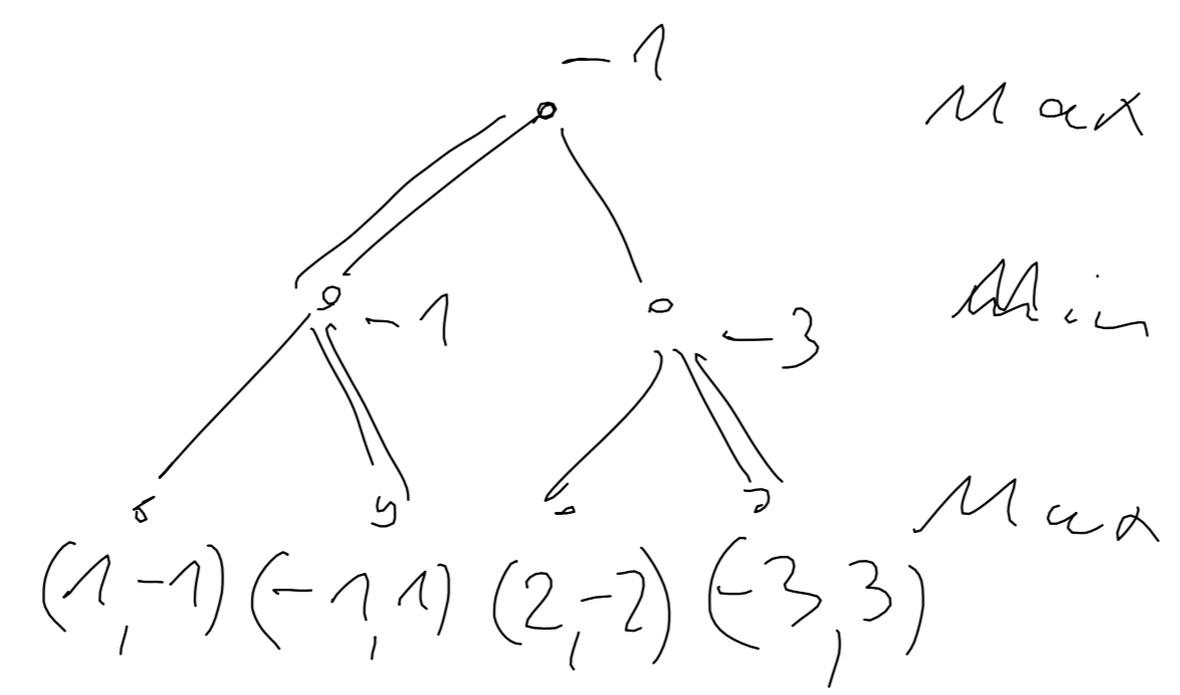
- For those familiar with the RT course:

It is a generalization of the minimax procedure.



backward induction

Two-Person
Zero-sum
Game



Proof: Let $T = \{N, A, H, P, (u_i)\}$ be finite

EGUPI. Construct an SPE by induction on $l(T(h))$ for all subgames $T(h)$. In parallel construct functions $t_i: H \rightarrow \mathbb{R}$ for all players $i \in N$ s.t. $t_i(h)$ is the payoff for player i in an SPE in subgame $T(h)$.

Base case: If $l(T(h)) = 0$ then $t_i(h) = u_i(h)$ for all $i \in N$.

Inductive case: If $t_i(h)$ is already defined for all $h \in H$ with $l(T(h)) \leq k$, consider $h^* \in H$ $l(T(h^*)) = k+1$ and $P(h^*) = i$:

For all $a \in A(h^*)$, let

$$s_i(h^*) := \underset{a \in A(h^*)}{\operatorname{argmax}} t_i(h^*, a) \text{ and}$$

$$t_j(h^*) = \underset{v}{t_j(h^*, s_i(h^*))} \text{ for all players } j \in X.$$

Inductively, we obtain a strategy profile that satisfies the one deviation property (because of the construction). From the one deviation property Lemma it follows that the constructed strategy profile is an SPE. \square

Remarks

- In principle, an SPE can be easily computed - once the full game tree is there

- Unfortunately, often this is not the case.

In practice, you use approximation.

- The problem is: For branching factor of b and a game tree of depth d , the game tree has a size of $O(b^d)$.

8^{50}

Pirate game

There are $n \leq 100$ pirates, who have found 100 gold coins that have distributed amongst them. Pirates have strict seniority principle: pirate i is superior to pirate $i+1$.

The most senior pirate proposes a distribution. Then all pirates vote on whether they accept it. In case of a tie, the proposing pirate has casting vote. If they accept, then game ends with distributing the coins. If not, the proposing pirate is thrown overboard.

Firstly, pirates prefer to survive. Secondly, they want to maximize coins. Thirdly, all the being the same, they prefer to throw pirates overboard.

What happens when 5 rational pirates play this game?

Pirate 5 is the only surviving one: 100 coins \rightarrow pirate 5

pirate 4, 5 are alive: pirate 4: (100, 0)

pirate 3, 4, 5 are alive: pirate 3: (99, 0, 1)

pirate 2, 3, 4, 5 are alive: pirate 2: (99, 0, 1, 0)

pirate 1, 2, 3, 4, 5 are alive: pirate 1: (98, 0, 1, 0, 1)

