Game Theory

6. Extensive Games

Albert-Ludwigs-Universität Freiburg

UNI FREIBURG

Bernhard Nebel and Robert Mattmüller

May 9th, 2018

Motivation
Definitions
Solution Concepts
One-Deviation Property
Kuhn's Theorem
Two Extensions
Summary

Motivation

Motivation

- So far: All players move simultaneously, and then the outcome is determined.
- Often in practice: Several moves in sequence (e. g. in chess).
  $\rightsquigarrow$ cannot be directly reflected by strategic games.
- Extensive games (with perfect information) reflect such situations by modeling games as game trees.
- Idea: Players have several decision points where they can decide how to play.
- Strategies: Mappings from decision points in the game tree to actions to be played.

Definitions

Motivation
Definitions
Solution Concepts
One-Deviation Property
Kuhn's Theorem
Two Extensions
Summary

Extensive Games

Definition (Extensive game with perfect information)

An extensive game with perfect information is a tuple $\Gamma = \langle N, H, P, (u_i)_{i \in N} \rangle$ that consists of:

- A finite non-empty set N of players.
- A set H of (finite or infinite) sequences, called histories, such that
  - the empty sequence $\langle\rangle \in H$,
  - H is closed under prefixes: if $\langle a^1, \dots, a^k \rangle \in H$ for some $k \in \mathbb{N} \cup \{\infty\}$, and $l < k$, then also $\langle a^1, \dots, a^l \rangle \in H$, and
  - H is closed under limits: if for some infinite sequence $\langle a^i \rangle_{i=1}^{\infty}$, we have $\langle a^i \rangle_{i=1}^{k} \in H$ for all $k \in \mathbb{N}$, then $\langle a^i \rangle_{i=1}^{\infty} \in H$.

  All infinite histories and all histories $\langle a^i \rangle_{i=1}^{k} \in H$, for which there is no a^{k+1} such that $\langle a^i \rangle_{i=1}^{k+1} \in H$ are called terminal histories Z. Components of a history are called actions.

Motivation
Definitions
Solution Concepts
One-Deviation Property
Kuhn's Theorem
Two Extensions
Summary

Extensive Games

Definition (Extensive game with perfect information, ctd.)

- A player function $P : H \setminus Z \to N$ that determines which player's turn it is to move after a given nonterminal history.
- For each player $i \in N$, a utility function (or payoff function) $u_i : Z \to \mathbb{R}$ defined on the set of terminal histories.

The game is called finite, if H is finite. It has a finite horizon, if the lenght of histories is bounded from above.

Assumption: All ingredients of Γ are common knowledge amongst the players of the game.

Terminology: In the following, we will simply write extensive games instead of extensive games with perfect information.

Example (Division game)

- Two identical objects should be divided among two players.
- Player 1 proposes an allocation.
- Player 2 agrees or rejects.
  - On agreement: Allocation as proposed.
  - On rejection: Nobody gets anything.

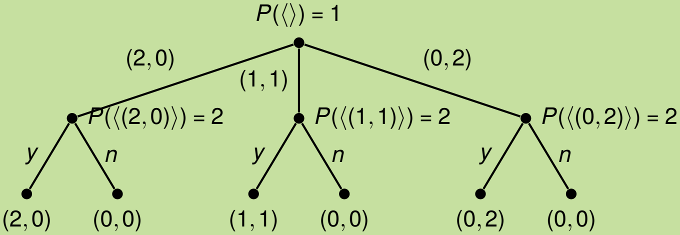

Motivation
Definitions
Solution Concepts
One-Deviation Property
Kuhn's Theorem
Two Extensions
Summary

UNI FREIBURG

Example (Division game, formally)

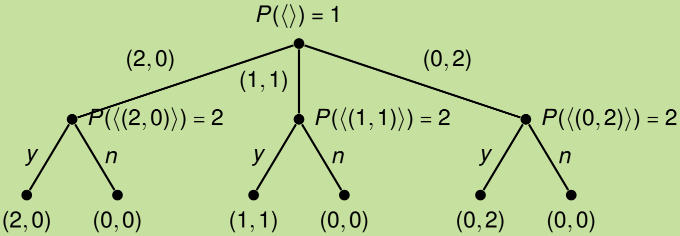

- $N = \{1,2\}$
- $H = \{\langle\rangle, \langle(2,0)\rangle, \langle(1,1)\rangle, \langle(0,2)\rangle, \langle(2,0),y\rangle, \langle(2,0),n\rangle, \dots\}$
- $P(\langle\rangle) = 1$, $P(h) = 2$ for all $h \in H \setminus Z$ with $h \neq \langle\rangle$
- $u_1(\langle(2,0),y\rangle) = 2$, $u_2(\langle(2,0),y\rangle) = 0$, etc.

Motivation
Definitions
Solution Concepts
One-Deviation Property
Kuhn's Theorem
Two Extensions
Summary

Motivation
Definitions
Solution Concepts
One-Deviation Property
Kuhn's Theorem
Two Extensions
Summary

Notation:

Let $h = \langle a^1, \dots, a^k \rangle$ be a history, and a an action.

- Then (h, a) is the history $\langle a^1, \dots, a^k, a \rangle$.
- If $h' = \langle b^1, \dots, b^\ell \rangle$, then (h, h') is the history $\langle a^1, \dots, a^k, b^1, \dots, b^\ell \rangle$.
- The set of actions from which player $P(h)$ can choose after a history $h \in H \setminus Z$ is written as

$$A(h) = \{a \mid (h, a) \in H\}.$$

Strategies

Definition (Strategy in an extensive game)

A strategy of a player i in an extensive game $\Gamma = \langle N, H, P, (u_i)_{i \in N} \rangle$ is a function s_i that assigns to each nonterminal history $h \in H \setminus Z$ with $P(h) = i$ an action $a \in A(h)$. The set of strategies of player i is denoted as S_i.

Remark: Strategies require us to assign actions to histories h, even if it is clear that they will never be played (e. g., because h will never be reached because of some earlier action).

Notation (for finite games): A strategy for a player is written as a string of actions at decision nodes as visited in a breadth-first order.

Motivation
Definitions
Solution Concepts
One-Deviation Property
Kuhn's Theorem
Two Extensions
Summary

Example (Strategies in an extensive game)

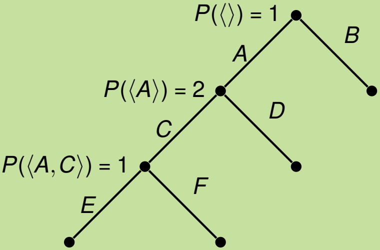

- Strategies for player 1: AE, AF, BE and BF
- Strategies for player 2: C and D.

Outcome

Definition (Outcome)

The outcome $O(s)$ of a strategy profile $s = (s_i)_{i \in N}$ is the (possibly infinite) terminal history $h = \langle a^i \rangle_{i=1}^{k}$, with $k \in \mathbb{N} \cup \{\infty\}$, such that for all $\ell \in \mathbb{N}$ with $0 \leq \ell < k$,

$$s_{P(\langle a^1,\ldots,a^\ell \rangle)}(\langle a^1, \ldots, a^\ell \rangle) = a^{\ell+1}.$$

Example (Outcome)

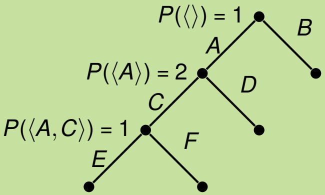

$O(AF, C) = \langle A, C, F \rangle$

$O(AE, D) = \langle A, D \rangle.$

Motivation

Definitions

Solution Concepts

One-Deviation Property

Kuhn's Theorem

Two Extensions

Summary

Solution Concepts

Definition (Nash equilibrium in an extensive game)

A Nash equilibrium in an extensive game $\Gamma = \langle N, H, P, (u_i)_{i \in N} \rangle$ is a strategy profile s^* such that for every player $i \in N$ and for all strategies $s_i \in S_i$,

$$u_i(O(s^*_{-i}, s^*_i)) \geq u_i(O(s^*_{-i}, s_i)).$$

Induced Strategic Game

Definition (Induced strategic game)

The strategic game G induced by an extensive game $\Gamma = \langle N, H, P, (u_i)_{i \in N} \rangle$ is defined by $G = \langle N, (A'_i)_{i \in N}, (u'_i)_{i \in N} \rangle$, where

- $A'_i = S_i$ for all $i \in N$, and
- $u'_i(a) = u_i(O(a))$ for all $i \in N$.

Proposition

The Nash equilibria of an extensive game Γ are exactly the Nash equilibria of the induced strategic game G of Γ. □

Induced Strategic Game

Motivation
Definitions
Solution Concepts
One-Deviation Property
Kuhn's Theorem
Two Extensions
Summary

Remarks:

- Each extensive game can be transformed into a strategic game, but the resulting game can be exponentially larger.
- The other direction does not work, because in extensive games, we do not have simultaneous actions.

Empty Threats

Example (Empty threat)

Extensive game:

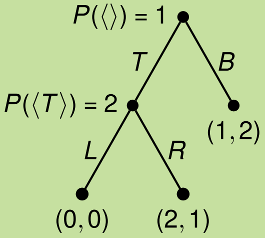

Strategies:

- Player 1: T and B
- Player 2: L and R

Strategic form:

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

Nash equilibria: (B,L) and (T,R).

However, (B,L) is not realistic:

- Player 1 plays B, "fearing" response L to T.
- But player 2 would never play L in the extensive game. $\rightsquigarrow$ (B,L) involves "empty threat".

Subgames

Idea: Exclude empty threats.

How? Demand that a strategy profile is not only a Nash equilibrium in the strategic form, but also in every subgame.

Definition (Subgame)

A subgame of an extensive game $\Gamma = \langle N, H, P, (u_i)_{i \in N} \rangle$, starting after history h, is the game $\Gamma(h) = \langle N, H|_h, P|_h, (u_i|_h)_{i \in N} \rangle$, where

- $H|_h = \{h' \mid (h, h') \in H\}$,
- $P|_h(h') = P(h, h')$ for all $h' \in H|_h$, and
- $u_i|_h(h') = u_i(h, h')$ for all $h' \in H|_h$.

Motivation
Definitions
Solution Concepts
One-Deviation Property
Kuhn's Theorem
Two Extensions
Summary

Definition (Strategy in a subgame)

Let Γ be an extensive game and $\Gamma(h)$ a subgame of Γ starting after some history h.

For each strategy s_i of Γ, let $s_i|_h$ be the strategy induced by s_i for $\Gamma(h)$. Formally, for all $h' \in H|_h$,

$$s_i|_h(h') = s_i(h, h').$$

The outcome function of $\Gamma(h)$ is denoted by O_h.

Subgame-Perfect Equilibria

Motivation
Definitions
Solution Concepts
One-Deviation Property
Kuhn's Theorem
Two Extensions
Summary

Definition (Subgame-perfect equilibrium)

A strategy profile s^* in an extensive game $\Gamma = \langle N, H, P, (u_i)_{i \in N} \rangle$ is a subgame-perfect equilibrium if and only if for every player $i \in N$ and every nonterminal history $h \in H \setminus Z$ with $P(h) = i$,

$$u_i|_h(O_h(s^*_{-i}|_h, s^*_i|_h)) \geq u_i|_h(O_h(s^*_{-i}|_h, s_i))$$

for every strategy $s_i \in S_i$ in subgame $\Gamma(h)$.

Subgame-Perfect Equilibria

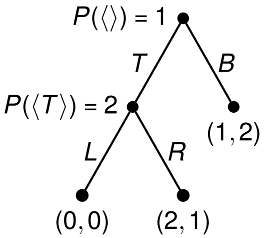

Two Nash equilibria:

- (T,R): subgame-perfect, because:
  - In history $h = \langle T\rangle$: subgame-perfect.
  - In history $h = \langle\rangle$: player 1 obtains utility 1 when choosing B and utility of 2 when choosing T.
- (B,L): not subgame-perfect, since L does not maximize the utility of player 2 in history $h = \langle T\rangle$.

Example (Subgame-perfect equilibria in division game)

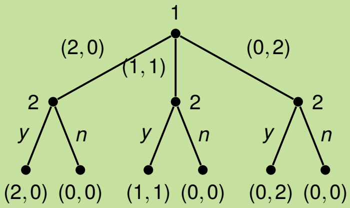

Equilibria in subgames:

- in $\Gamma(\langle(2,0)\rangle)$: y and n
- in $\Gamma(\langle(1,1)\rangle)$: only y
- in $\Gamma(\langle(0,2)\rangle)$: only y
- in $\Gamma(\langle\rangle)$: $((2,0), yyy)$ and $((1,1), nyy)$

Nash equilibria (red: empty threat):

- $((2,0), yyy)$, $((2,0), yyn)$, $((2,0), yny)$, $((2,0), ynn)$, $((2,0), nny)$, $((2,0), nnn)$,
- $((1,1), nyy)$, $((1,1), nyn)$,
- $((0,2), nny)$, $((0,2), nnn)$.

Motivation

Definitions

Solution Concepts

One-Deviation Property

Kuhn's Theorem

Two Extensions

Summary

One-Deviation Property

Motivation

- Existence:
  - Does every extensive game have a subgame-perfect equilibrium?
  - If not, which extensive games do have a subgame-perfect equilibrium?
- Computation:
  - If a subgame-perfect equilibrium exists, how to compute it?
  - How complex is that computation?

Motivation

Motivation

Definitions

Solution Concepts

One-Deviation Property

Kuhn's Theorem

Two Extensions

Summary

Positive case (a subgame-perfect equilibrium exists):

- Step 1: Show that is suffices to consider local deviations from strategies (for finite-horizon games).
- Step 2: Show how to systematically explore such local deviations to find a subgame-perfect equilibrium (for finite games).

Step 1: One-Deviation Property

Definition

Let Γ be a finite-horizon extensive game. Then $\ell(\Gamma)$ denotes the length of the longest history of Γ.

Step 1: One-Deviation Property

Motivation
Definitions
Solution Concepts
One-Deviation Property
Kuhn's Theorem
Two Extensions
Summary

Definition (One-deviation property)

A strategy profile s^* in an extensive game $\Gamma = \langle N, H, P, (u_i)_{i \in N} \rangle$ satisfies the one-deviation property if and only if for every player $i \in N$ and every nonterminal history $h \in H \setminus Z$ with $P(h) = i$,

$$u_i|_h(O_h(s^*_{-i}|_h, s^*_i|_h)) \geq u_i|_h(O_h(s^*_{-i}|_h, s_i))$$

for every strategy $s_i \in S_i$ in subgame $\Gamma(h)$ that differs from $s^*_i|_h$ only in the action it prescribes after the initial history of $\Gamma(h)$.

Note: Without the highlighted parts, this is just the definition of subgame-perfect equilibria!

Step 1: One-Deviation Property

Lemma

Let $\Gamma = \langle N, H, P, (u_i)_{i \in N} \rangle$ be a finite-horizon extensive game. Then a strategy profile s^* is a subgame-perfect equilibrium of Γ if and only if it satisfies the one-deviation property.

Proof

- ($\Rightarrow$) Clear.
- ($\Leftarrow$) By contradiction:

  Suppose that s^* is not a subgame-perfect equilibrium.

  Then there is a history h and a player i such that s_i is a profitable deviation for player i in subgame $\Gamma(h)$.

  ...

Step 1: One-Deviation Property

Motivation
Definitions
Solution Concepts
One-Deviation Property
Kuhn's Theorem
Two Extensions
Summary

Lemma

Let $\Gamma = \langle N, H, P, (u_i)_{i \in N} \rangle$ be a finite-horizon extensive game. Then a strategy profile s^* is a subgame-perfect equilibrium of Γ if and only if it satisfies the one-deviation property.

Proof

- ($\Rightarrow$) Clear.
- ($\Leftarrow$) By contradiction:

  Suppose that s^* is not a subgame-perfect equilibrium.

  Then there is a history h and a player i such that s_i is a profitable deviation for player i in subgame $\Gamma(h)$.

  ...

Step 1: One-Deviation Property

Motivation
Definitions
Solution Concepts
One-Deviation Property
Kuhn's Theorem
Two Extensions
Summary

Lemma

Let $\Gamma = \langle N, H, P, (u_i)_{i \in N} \rangle$ be a finite-horizon extensive game. Then a strategy profile s^* is a subgame-perfect equilibrium of Γ if and only if it satisfies the one-deviation property.

Proof

- ($\Rightarrow$) Clear.
- ($\Leftarrow$) By contradiction:

  Suppose that s^* is not a subgame-perfect equilibrium.

  Then there is a history h and a player i such that s_i is a profitable deviation for player i in subgame $\Gamma(h)$.

  ...

Proof (ctd.)

- ($\Leftarrow$) …WLOG, the number of histories h' with $s_i(h') \neq s_i^*|_h(h')$ is at most $\ell(\Gamma(h))$ and hence finite (finite horizon assumption!), since deviations not on resulting outcome path are irrelevant.

  Illustration: strategies $s_1^*|_h = AGILN$ and $s_2^*|_h = CF$ red:

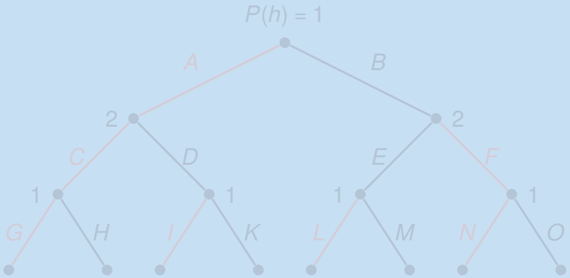

Proof (ctd.)

- ($\Leftarrow$) …WLOG, the number of histories h' with $s_i(h') \neq s_i^*|_h(h')$ is at most $\ell(\Gamma(h))$ and hence finite (finite horizon assumption!), since deviations not on resulting outcome path are irrelevant.

  Illustration: strategies $s_1^*|_h = AGILN$ and $s_2^*|_h = CF$ red:

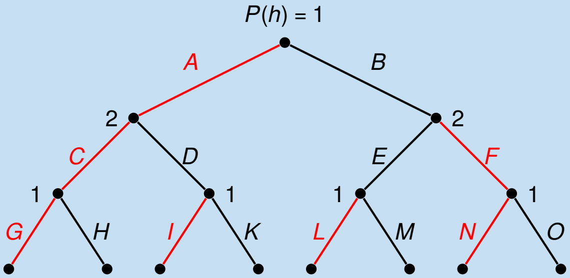

Motivation
Definitions
Solution Concepts
One-Deviation Property
Kuhn's Theorem
Two Extensions
Summary

Proof (ctd.)

- (⇐) …Illustration for WLOG assumption: Assume $s_1 = BHKMO$ (blue) profitable deviation:

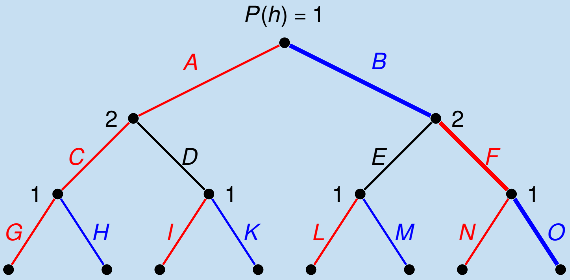

Then only B and O really matter.

Motivation
Definitions
Solution Concepts
One-Deviation Property
Kuhn's Theorem
Two Extensions
Summary

Proof (ctd.)

- (⇐) …Illustration for WLOG assumption: And hence $\tilde{s}_1 = BGILO$ (blue) also profitable deviation:

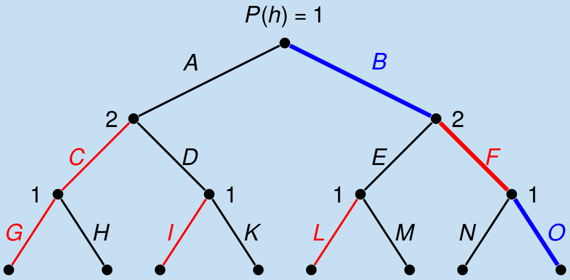

Step 1: One-Deviation Property

Proof (ctd.)

- $(\Leftarrow)$ ...

  Choose profitable deviation s_i in $\Gamma(h)$ with minimal number of deviation points (such s_i must exist).

  Let h^* be the longest history in $\Gamma(h)$ with $s_i(h^*) \neq s_i^*|_h(h^*)$, i.e., "deepest" deviation point for s_i.

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

  Moreover, $s_i|_{h^*}$ is a profitable deviation in $\Gamma(h, h^*)$, since h^* is the *longest* history in $\Gamma(h)$ with $s_i(h^*) \neq s_i^*|_h(h^*)$.

  So, $\Gamma(h, h^*)$ is the desired subgame where a one-step deviation is sufficient to improve utility. □

Motivation
Definitions
Solution Concepts
One-Deviation Property
Kuhn's Theorem
Two Extensions
Summary

Step 1: One-Deviation Property

Proof (ctd.)

- ($\Leftarrow$) ...

  Choose profitable deviation s_i in $\Gamma(h)$ with minimal number of deviation points (such s_i must exist).

  Let h^* be the longest history in $\Gamma(h)$ with $s_i(h^*) \neq s_i^*|_h(h^*)$, i.e., "deepest" deviation point for s_i.

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

  Moreover, $s_i|_{h^*}$ is a profitable deviation in $\Gamma(h,h^*)$, since h^* is the *longest* history in $\Gamma(h)$ with $s_i(h^*) \neq s_i^*|_h(h^*)$.

  So, $\Gamma(h,h^*)$ is the desired subgame where a one-step deviation is sufficient to improve utility. □

Motivation
Definitions
Solution Concepts
One-Deviation Property
Kuhn's Theorem
Two Extensions
Summary

Step 1: One-Deviation Property

Proof (ctd.)

- ($\Leftarrow$) …

  Choose profitable deviation s_i in $\Gamma(h)$ with minimal number of deviation points (such s_i must exist).

  Let h^* be the longest history in $\Gamma(h)$ with $s_i(h^*) \neq s_i^*|_h(h^*)$, i.e., "deepest" deviation point for s_i.

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

  Moreover, $s_i|_{h^*}$ is a profitable deviation in $\Gamma(h, h^*)$, since h^* is the *longest* history in $\Gamma(h)$ with $s_i(h^*) \neq s_i^*|_h(h^*)$.

  So, $\Gamma(h, h^*)$ is the desired subgame where a one-step deviation is sufficient to improve utility. □

Proof (ctd.)

- ($\Leftarrow$) ...

  Choose profitable deviation s_i in $\Gamma(h)$ with minimal number of deviation points (such s_i must exist).

  Let h^* be the longest history in $\Gamma(h)$ with $s_i(h^*) \neq s_i^*|_h(h^*)$, i.e., "deepest" deviation point for s_i.

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

  Moreover, $s_i|_{h^*}$ is a profitable deviation in $\Gamma(h, h^*)$, since h^* is the *longest* history in $\Gamma(h)$ with $s_i(h^*) \neq s_i^*|_h(h^*)$.

  So, $\Gamma(h, h^*)$ is the desired subgame where a one-step deviation is sufficient to improve utility. □

Motivation
Definitions
Solution Concepts
One-Deviation Property
Kuhn's Theorem
Two Extensions
Summary

Step 1: One-Deviation Property

Proof (ctd.)

- (⇐) …

  Choose profitable deviation s_i in $\Gamma(h)$ with minimal number of deviation points (such s_i must exist).

  Let h^* be the longest history in $\Gamma(h)$ with $s_i(h^*) \neq s_i^*|_h(h^*)$, i.e., "deepest" deviation point for s_i.

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

  Moreover, $s_i|_{h^*}$ is a profitable deviation in $\Gamma(h, h^*)$, since h^* is the *longest* history in $\Gamma(h)$ with $s_i(h^*) \neq s_i^*|_h(h^*)$.

  So, $\Gamma(h, h^*)$ is the desired subgame where a one-step deviation is sufficient to improve utility. □

Motivation
Definitions
Solution Concepts
One-Deviation Property
Kuhn's Theorem
Two Extensions
Summary

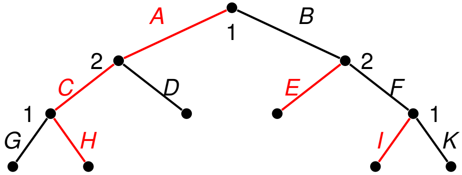

To show that (AHI, CE) is a subgame-perfect equilibrium, it suffices to check these deviating strategies:

Player 1:

- G in subgame $\Gamma(\langle A, C\rangle)$
- K in subgame $\Gamma(\langle B, F\rangle)$
- BHI in Γ

Player 2:

- D in subgame $\Gamma(\langle A\rangle)$
- F in subgame $\Gamma(\langle B\rangle)$

In particular, e.g., no need to check if strategy BGK of player 1 is profitable in Γ.

UNI FREIBURG

The corresponding proposition for infinite-horizon games does not hold.

Counterexample (one-player case):

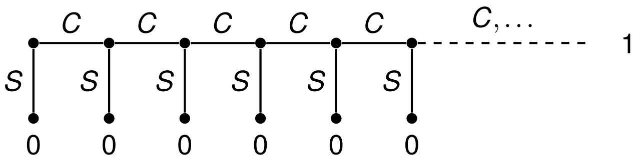

Strategy s_i with $s_i(h) = S$ for all $h \in H \setminus Z$

- satisfies one deviation property, but
- is not a subgame-perfect equilibrium, since it is dominated by s_i^* with $s_i^*(h) = C$ for all $h \in H \setminus Z$.

Kuhn's Theorem

Step 2: Kuhn's Theorem

Motivation
Definitions
Solution Concepts
One-Deviation Property
Kuhn's Theorem
Two Extensions
Summary

Theorem (Kuhn)

Every finite extensive game has a subgame-perfect equilibrium.

Proof idea:

- Proof is constructive and builds a subgame-perfect equilibrium bottom-up (aka backward induction).
- For those familiar with the Foundations of AI lecture: generalization of Minimax algorithm to general-sum games with possibly more than two players.

Example

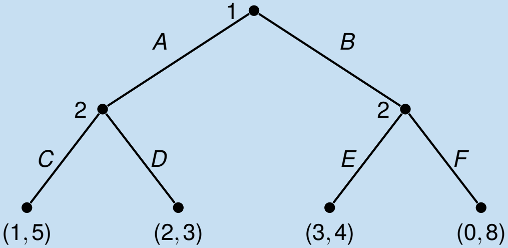

Motivation
Definitions
Solution Concepts
One-Deviation Property
Kuhn's Theorem
Two Extensions
Summary

Step 2: Kuhn's Theorem

Motivation
Definitions
Solution Concepts
One-Deviation Property
Kuhn's Theorem
Two Extensions
Summary

Example

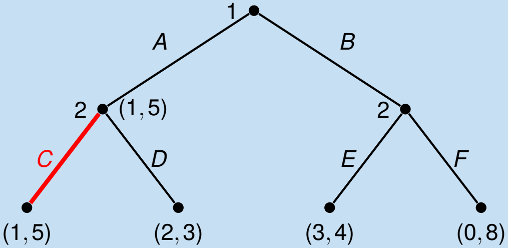

$s_2(\langle A \rangle) = C$ $\quad t_1(\langle A \rangle) = 1$ $\quad t_2(\langle A \rangle) = 5$

Example

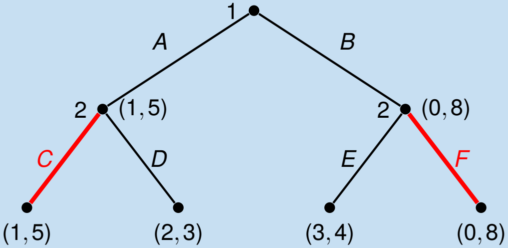

$s_2(\langle A \rangle) = C$ | $t_1(\langle A \rangle) = 1$ | $t_2(\langle A \rangle) = 5$

$s_2(\langle B \rangle) = F$ | $t_1(\langle B \rangle) = 0$ | $t_2(\langle B \rangle) = 8$

Motivation
Definitions
Solution Concepts
One-Deviation Property
Kuhn's Theorem
Two Extensions
Summary

Step 2: Kuhn's Theorem

Example

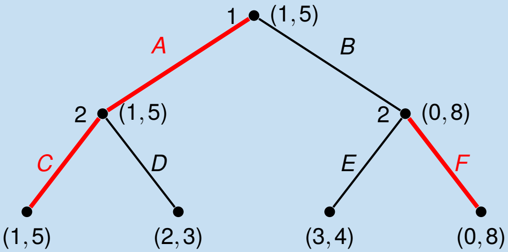

$s_2(\langle A\rangle) = C$ $\quad t_1(\langle A\rangle) = 1$ $\quad t_2(\langle A\rangle) = 5$

$s_2(\langle B\rangle) = F$ $\quad t_1(\langle B\rangle) = 0$ $\quad t_2(\langle B\rangle) = 8$

$s_1(\langle\rangle) = A$ $\quad t_1(\langle\rangle) = 1$ $\quad t_2(\langle\rangle) = 5$

Motivation
Definitions
Solution Concepts
One-Deviation Property
Kuhn's Theorem
Two Extensions
Summary

Step 2: Kuhn's Theorem

A bit more formally:

Proof

Let $\Gamma = \langle N, H, P, (u_i)_{i \in N} \rangle$ be a finite extensive game.

Construct a subgame-perfect equilibrium by induction on $\ell(\Gamma(h))$ for all subgames $\Gamma(h)$. In parallel, construct functions $t_i : H \to \mathbb{R}$ for all players $i \in N$ s. t. $t_i(h)$ is the payoff for player i in a subgame-perfect equilibrium in subgame $\Gamma(h)$.

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

...

Motivation
Definitions
Solution Concepts
One-Deviation Property
Kuhn's Theorem
Two Extensions
Summary

A bit more formally:

Proof

Let $\Gamma = \langle N, H, P, (u_i)_{i \in N} \rangle$ be a finite extensive game.

Construct a subgame-perfect equilibrium by induction on $\ell(\Gamma(h))$ for all subgames $\Gamma(h)$. In parallel, construct functions $t_i : H \to \mathbb{R}$ for all players $i \in N$ s. t. $t_i(h)$ is the payoff for player i in a subgame-perfect equilibrium in subgame $\Gamma(h)$.

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

...

Step 2: Kuhn's Theorem

A bit more formally:

Proof

Let $\Gamma = \langle N, H, P, (u_i)_{i \in N} \rangle$ be a finite extensive game.

Construct a subgame-perfect equilibrium by induction on $\ell(\Gamma(h))$ for all subgames $\Gamma(h)$. In parallel, construct functions $t_i : H \to \mathbb{R}$ for all players $i \in N$ s. t. $t_i(h)$ is the payoff for player i in a subgame-perfect equilibrium in subgame $\Gamma(h)$.

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

...

Motivation

Definitions

Solution Concepts

One-Deviation Property

Kuhn's Theorem

Two Extensions

Summary

Step 2: Kuhn's Theorem

Proof (ctd.)

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

For all $a \in A(h^*)$, $\ell(\Gamma(h^*, a)) \leq k$, let

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

$$t_j(h^*) := t_j(h^*, s_i(h^*)) \quad \text{for all players } j \in N.$$

Inductively, we obtain a strategy profile s that satisfies the one-deviation property.

With the one-deviation property lemma it follows that s is a subgame-perfect equilibrium. □

Motivation
Definitions
Solution Concepts
One-Deviation Property
Kuhn's Theorem
Two Extensions
Summary

Step 2: Kuhn's Theorem

Motivation
Definitions
Solution Concepts
One-Deviation Property
Kuhn's Theorem
Two Extensions
Summary

Proof (ctd.)

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

For all $a \in A(h^*)$, $\ell(\Gamma(h^*, a)) \leq k$, let

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

$$t_j(h^*) := t_j(h^*, s_i(h^*)) \quad \text{for all players } j \in N.$$

Inductively, we obtain a strategy profile s that satisfies the one-deviation property.

With the one-deviation property lemma it follows that s is a subgame-perfect equilibrium. □

Motivation
Definitions
Solution Concepts
One-Deviation Property
Kuhn's Theorem
Two Extensions
Summary

Step 2: Kuhn's Theorem

Proof (ctd.)

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

For all $a \in A(h^*)$, $\ell(\Gamma(h^*, a)) \leq k$, let

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

$$t_j(h^*) := t_j(h^*, s_i(h^*)) \quad \text{for all players } j \in N.$$

Inductively, we obtain a strategy profile s that satisfies the one-deviation property.

With the one-deviation property lemma it follows that s is a subgame-perfect equilibrium. □

Step 2: Kuhn's Theorem

Motivation
Definitions
Solution Concepts
One-Deviation Property
Kuhn's Theorem
Two Extensions
Summary

- In principle: sample subgame-perfect equilibrium effectively computable using the technique from the above proof.
- In practice: often game trees not enumerated in advance, hence unavailable for backward induction.
- E.g., for branching factor b and depth m, procedure needs time $O(b^m)$.

Motivation
Definitions
Solution Concepts
One-Deviation Property
Kuhn's Theorem
Two Extensions
Summary

Corresponding proposition for infinite games does not hold.

Counterexamples (both for one-player case):

A) finite horizon, infinite branching factor:

Infinitely many actions $a \in A = [0, 1)$ with payoffs $u_1(\langle a \rangle) = a$ for all $a \in A$.
There exists no subgame-perfect equilibrium in this game.

Motivation
Definitions
Solution Concepts
One-Deviation Property
Kuhn's Theorem
Two Extensions
Summary

B) infinite horizon, finite branching factor:

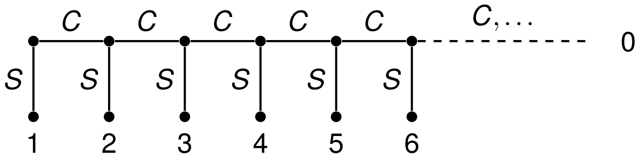

$u_1(CCC\ldots) = 0$ and $u_1(\underbrace{CC\ldots C}_{n}S) = n+1$.

No subgame-perfect equilibrium.

Step 2: Kuhn's Theorem

Motivation
Definitions
Solution Concepts
One-Deviation Property
Kuhn's Theorem
Two Extensions
Summary

Uniqueness:

Kuhn's theorem tells us nothing about uniqueness of subgame-perfect equilibria. However, if no two histories get the same evaluation by any player, then the subgame-perfect equilibrium is unique.

Extended Example: Pirate Game

1. There are 5 *rational* pirates, A, B, C, D and E. They find 100 gold coins. They must decide how to distribute them.
2. The pirates have a strict order of *seniority*: A is senior to B, who is senior to C, who is senior to D, who is senior to E.
3. The pirate world's rules of distribution say that the most senior pirate first *proposes* a distribution of coins. The pirates, including the proposer, then *vote* on whether to accept this distribution (in order from most junior to senior). In case of a tie vote, the proposer has the casting vote. If the distribution is accepted, the coins are disbursed and the *game ends*. If not, the proposer is thrown overboard from the pirate ship and dies, and the next most senior pirate makes a new proposal to apply the method again.

Extended Example: Pirate Game

1. There are 5 *rational* pirates, A, B, C, D and E. They find 100 gold coins. They must decide how to distribute them.
2. The pirates have a strict order of *seniority*: A is senior to B, who is senior to C, who is senior to D, who is senior to E.
3. The pirate world's rules of distribution say that the most senior pirate first *proposes* a distribution of coins. The pirates, including the proposer, then *vote* on whether to accept this distribution (in order from most junior to senior). In case of a tie vote, the proposer has the casting vote. If the distribution is accepted, the coins are disbursed and the *game ends*. If not, the proposer is thrown overboard from the pirate ship and dies, and the next most senior pirate makes a new proposal to apply the method again.

Extended Example: Pirate Game

1. There are 5 *rational* pirates, A, B, C, D and E. They find 100 gold coins. They must decide how to distribute them.
2. The pirates have a strict order of *seniority*: A is senior to B, who is senior to C, who is senior to D, who is senior to E.
3. The pirate world's rules of distribution say that the most senior pirate first *proposes* a distribution of coins. The pirates, including the proposer, then *vote* on whether to accept this distribution (in order from most junior to senior). In case of a tie vote, the proposer has the casting vote. If the distribution is accepted, the coins are disbursed and the *game ends*. If not, the proposer is thrown overboard from the pirate ship and dies, and the next most senior pirate makes a new proposal to apply the method again.

Motivation
Definitions
Solution Concepts
One-Deviation Property
Kuhn's Theorem
Two Extensions
Summary

Pirates: General Setting & Utility

4. The pirates do not trust each other, and will neither make nor honor any promises between pirates apart from a proposed distribution plan that gives a whole number of gold coins to each pirate.
5. Pirates base their decisions on three factors. First of all, each pirate wants to *survive*. Second, everything being equal, each pirate wants to *maximize the number of gold coins* each receives. Third, each pirate would prefer to *throw another overboard*, if all other results would otherwise be equal.

Motivation
Definitions
Solution Concepts
One-Deviation Property
Kuhn's Theorem
Two Extensions
Summary

Pirates: General Setting & Utility

4. The pirates do not trust each other, and will neither make nor honor any promises between pirates apart from a proposed distribution plan that gives a whole number of gold coins to each pirate.
5. Pirates base their decisions on three factors. First of all, each pirate wants to *survive*. Second, everything being equal, each pirate wants to *maximize the number of gold coins* each receives. Third, each pirate would prefer to *throw another overboard*, if all other results would otherwise be equal.

Pirates: Formalization

- Players $N = \{A, B, C, D, E\}$;
- actions are:
  - proposals by a pirate: $\langle A : x_A, B : x_b, C : x_B, D : x_D, E : x_E \rangle$, with $\sum_{i \in \{A,B,C,D,E\}} x_i = 100$;
  - votings: y for accepting, n for rejecting;
- histories are sequences of a proposal, followed by votings of the alive pirates;
- utilities:
  - for pirates who are alive: utilities are according to the accepted proposal plus $x/100$, x being the number of dead pirates;
  - for dead pirates: -100.

Remark: Very large game tree!

Pirates: Analysis by Backward Induction

1. Assume only D and E are still alive. D can propose $\langle A:0, B:0, C:0, D:100, E:0\rangle$, because D has the casting vote!
2. Assume C, D, and E are alive. For C it is enough to offer 1 coin to E to get his vote: $\langle A:0, B:0, C:99, D:0, E:1\rangle$.
3. Assume B, C, D, and E are alive. B offering D one coin is enough because of the casting vote: $\langle A:0, B:99, C:0, D:1, E:0\rangle$.
4. Assume A, B, C, D, and E are alive. A offering C and E each one coin is enough: $\langle A:98, B:0, C:1, D:0, E:1\rangle$ (note that giving 1 to D instaed to E does not help).

Pirates: Analysis by Backward Induction

1. Assume only D and E are still alive. D can propose $\langle A:0, B:0, C:0, D:100, E:0\rangle$, because D has the casting vote!
2. Assume C, D, and E are alive. For C it is enough to offer 1 coin to E to get his vote: $\langle A:0, B:0, C:99, D:0, E:1\rangle$.
3. Assume B, C, D, and E are alive. B offering D one coin is enough because of the casting vote: $\langle A:0, B:99, C:0, D:1, E:0\rangle$.
4. Assume A, B, C, D, and E are alive. A offering C and E each one coin is enough: $\langle A:98, B:0, C:1, D:0, E:1\rangle$ (note that giving 1 to D instaed to E does not help).

Pirates: Analysis by Backward Induction

1. Assume only D and E are still alive. D can propose $\langle A:0, B:0, C:0, D:100, E:0\rangle$, because D has the casting vote!
2. Assume C, D, and E are alive. For C it is enough to offer 1 coin to E to get his vote: $\langle A:0, B:0, C:99, D:0, E:1\rangle$.
3. Assume B, C, D, and E are alive. B offering D one coin is enough because of the casting vote: $\langle A:0, B:99, C:0, D:1, E:0\rangle$.
4. Assume A, B, C, D, and E are alive. A offering C and E each one coin is enough: $\langle A:98, B:0, C:1, D:0, E:1\rangle$ (note that giving 1 to D instaed to E does not help).

Pirates: Analysis by Backward Induction

1. Assume only D and E are still alive. D can propose $\langle A:0, B:0, C:0, D:100, E:0\rangle$, because D has the casting vote!
2. Assume C, D, and E are alive. For C it is enough to offer 1 coin to E to get his vote: $\langle A:0, B:0, C:99, D:0, E:1\rangle$.
3. Assume B, C, D, and E are alive. B offering D one coin is enough because of the casting vote: $\langle A:0, B:99, C:0, D:1, E:0\rangle$.
4. Assume A, B, C, D, and E are alive. A offering C and E each one coin is enough: $\langle A:98, B:0, C:1, D:0, E:1\rangle$ (note that giving 1 to D instaed to E does not help).

Two Extensions

Motivation
Definitions
Solution Concepts
One-Deviation Property
Kuhn's Theorem
Two Extensions
Simultaneous Moves
Chance
Summary

Motivation
Definitions
Solution Concepts
One-Deviation Property
Kuhn's Theorem
Two Extensions
Simultaneous Moves
Chance
Summary

Definition

An extensive game with simultaneous moves is a tuple $\Gamma = \langle N, H, P, (u_i)_{i \in N} \rangle$, where

- N, H, P and (u_i) are defined as before, and
- $P : H \to 2^N$ assigns to each nonterminal history a set of players to move; for all $h \in H \setminus Z$, there exists a family $(A_i(h))_{i \in P(h)}$ such that

$$A(h) = \{a \mid (h,a) \in H\} = \prod_{i \in P(h)} A_i(h).$$

Simultaneous Moves

- Intended meaning of simultaneous moves: All players from $P(h)$ move simultaneously.
- Strategies: Functions $s_i : h \mapsto a_i$ with $a_i \in A_i(h)$.
- Histories: Sequences of vectors of actions.
- Outcome: Terminal history reached when tracing strategy profile.
- Payoffs: Utilities at outcome history.

Motivation
Definitions
Solution Concepts
One-Deviation Property
Kuhn's Theorem
Two Extensions
Simultaneous Moves
Chance
Summary

Remark:

- The one-deviation property still holds for extensive game with perfect information and simultaneous moves.
- Kuhn's theorem does not hold for extensive game with simultaneous moves.

  Example: MATCHING PENNIES can be viewed as extensive game with simultaneous moves. No Nash equilibrium/subgame-perfect equilibrium.

  player 2

		H	T
player 1	H	$1, -1$	$-1, \ 1$
	T	$-1, \ 1$	$1, -1$

  $\rightsquigarrow$ Need more sophisticated solution concepts (cf. mixed strategies). Not covered in this lecture.

Motivation
Definitions
Solution Concepts
One-Deviation Property
Kuhn's Theorem
Two Extensions
Simultaneous Moves
Chance
Summary

Setting:

- Three players have to split a cake fairly.
- Player 1 suggest split: shares $x_1, x_2, x_3 \in [0, 1]$ s.t. $x_1 + x_2 + x_3 = 1$.
- Then players 2 and 3 simultaneously and independently decide whether to accept ("y") or reject ("n") the suggested splitting.
- If both accept, each player i gets his allotted share (utility x_i). Otherwise, no player gets anything (utility 0).

Motivation
Definitions
Solution Concepts
One-Deviation Property
Kuhn's Theorem
Two Extensions
Simultaneous Moves
Chance
Summary

Motivation
Definitions
Solution Concepts
One-Deviation Property
Kuhn's Theorem
Two Extensions
Simultaneous Moves
Chance
Summary

Formally:

$$
\begin{aligned}
N &= \{1,2,3\} \\
X &= \{(x_1,x_2,x_3) \in [0,1]^3 \mid x_1+x_2+x_3 = 1\} \\
H &= \{\langle\rangle\} \cup \{\langle x\rangle \mid x \in X\} \cup \{\langle x,z\rangle \mid x \in X, z \in \{\mathrm{y},\mathrm{n}\} \times \{\mathrm{y},\mathrm{n}\}\} \\
P(\langle\rangle) &= \{1\} \\
P(\langle x\rangle) &= \{2,3\} \text{ for all } x \in X \\
u_i(\langle x,z\rangle) &= \begin{cases} 0 & \text{if } z \in \{(\mathrm{y},\mathrm{n}),(\mathrm{n},\mathrm{y}),(\mathrm{n},\mathrm{n})\} \\ x_i & \text{if } z = (\mathrm{y},\mathrm{y}). \end{cases} \text{ for all } i \in N
\end{aligned}
$$

Motivation
Definitions
Solution Concepts
One-Deviation Property
Kuhn's Theorem
Two Extensions
Simultaneous Moves
Chance
Summary

Subgame-perfect equilibria:

- Subgames after legal split (x_1, x_2, x_3) by player 1:
  - NE (y, y) (both accept)
  - NE (n, n) (neither accepts)
  - If $x_2 = 0$, NE (n, y) (only player 3 accepts)
  - If $x_3 = 0$, NE (y, n) (only player 2 accepts)

Subgame-perfect equilibria (ctd.):

- Entire game:
  Let s_2 and s_3 be any two strategies of players 2 and 3 such that for all splits $x \in X$ the profile $(s_2(\langle x \rangle), s_3(\langle x \rangle))$ is one of the NEs from above.
  Let $X_y = \{x \in X \mid s_2(\langle x \rangle) = s_3(\langle x \rangle) = y\}$ be the set of splits accepted under s_2 and s_3. Distinguish three cases:
  - $X_y = \emptyset$ or $x_1 = 0$ for all $x \in X_y$. Then (s_1, s_2, s_3) is a subgame-perfect equilibrium for any possible s_1.
  - $X_y \neq \emptyset$ and there are splits $x_{max} = (x_1, x_2, x_3) \in X_y$ that maximize $x_1 > 0$. Then (s_1, s_2, s_3) is a subgame-perfect equilibrium if and only if $s_1(\langle\rangle)$ is such a split x_{max}.
  - $X_y \neq \emptyset$ and there are no splits $(x_1, x_2, x_3) \in X_y$ that maximize x_1. Then there is no subgame-perfect equilibrium, in which player 2 follows strategy s_2 and player 3 follows strategy s_3.

Motivation
Definitions
Solution Concepts
One-Deviation Property
Kuhn's Theorem
Two Extensions
Simultaneous Moves
Chance
Summary

Chance Moves

Definition

An extensive game with chance moves is a tuple $\Gamma = \langle N, H, P, f_c, (u_i)_{i \in N} \rangle$, where

- N, A, H and u_i are defined as before,
- the player function $P : H \setminus Z \to N \cup \{c\}$ can also take the value c for a chance node, and
- for each $h \in H \setminus Z$ with $P(h) = c$, the function $f_c(\cdot|h)$ is a probability distribution on $A(h)$ such that the probability distributions for all $h \in H$ are independent of each other.

Motivation
Definitions
Solution Concepts
One-Deviation Property
Kuhn's Theorem
Two Extensions
Simultaneous Moves
Chance
Summary

Chance Moves

- Intended meaning of chance moves: In chance node, an applicable action is chosen randomly with probability according to f_c.
- Strategies: Defined as before.
- Outcome: For a given strategy profile, the outcome is a probability distribution on the set of terminal histories.
- Payoffs: For player i, U_i is the expected payoff (with weights according to outcome probabilities).

Chance Moves

Example

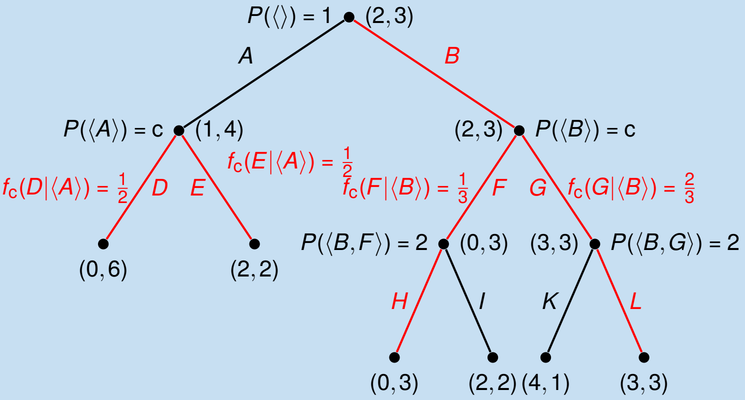

Motivation

Definitions

Solution Concepts

One-Deviation Property

Kuhn's Theorem

Two Extensions

Simultaneous Moves

Chance

Summary

Remark:

The one-deviation property and Kuhn's theorem still hold in the presence of chance moves. When proving Kuhn's theorem, expected utilities have to be used.

Summary

Motivation
Definitions
Solution Concepts
One-Deviation Property
Kuhn's Theorem
Two Extensions
Summary

Summary

- For finite-horizon extensive games, it suffices to consider local deviations when looking for better strategies.
- For infinite-horizon games, this is not true in general.
- Every finite extensive game has a subgame-perfect equilibrium.
- This does not generally hold for infinite games, no matter is game is infinite due to infinite branching factor or infinitely long histories (or both).

- With chance moves, one deviation property and Kuhn's theorem still hold.
- With simultaneous moves, Kuhn's theorem no longer holds.

Motivation
Definitions
Solution Concepts
One-Deviation Property
Kuhn's Theorem
Two Extensions
Summary