Game Theory

Kuhn's Theorem

Albert-Ludwigs-Universität Freiburg

UNI FREIBURG

Bernhard Nebel, Robert Mattmüller,
Stefan Wölfl, Christian Becker-Asano

Research Group Foundations of Artificial Intelligence

June 17, 2013

One Deviation Property

Motivation

- Existence:
  - Does every extensive game with perfect information have an SPE?
  - If not, which extensive games with perfect information do have an SPE?
- Computation:
  - If an SPE exists, how to compute it?
  - How complex is that computation?

Motivation

Positive Case (an SPE 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 an SPE (for finite games).

Definition

Let G be a finite-horizon extensive game with perfect information. Then $\ell(G)$ denotes the length of the longest history of G.

Step 1: One Deviation Property

Lemma (One Deviation Property)

Let $G = \langle N, A, H, \rho, (u_i) \rangle$ be a finite-horizon extensive game with perfect information. Then a strategy profile s^* is a subgame perfect equilibrium of G if and only if for every player $i \in N$ and every history $h \in H$ for which $\rho(h) = i$, we have

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

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

Note: Without the highlighted parts, this is just the definition of SPEs!

Step 1: One Deviation Property

Proof

- (⇒) Clear.
- (⇐) By contradiction:

  Suppose that s^* is not an SPE.

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

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

  ...

Step 1: One Deviation Property

Proof

- (⇒) Clear.
- (⇐) By contradiction:

  Suppose that s^* is not an SPE.

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

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

  ...

Step 1: One Deviation Property

Proof

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

  Suppose that s^* is not an SPE.

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

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

  ...

Step 1: One Deviation Property

Proof

- (⇒) Clear.
- (⇐) By contradiction:

  Suppose that s^* is not an SPE.

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

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

  ...

Proof (ctd.)

- (⇐) . . . Illustration for WLOG assumption: Strategies $s_1^*|_h = AGILN$ and $s_2^*|_h = CF$ red:

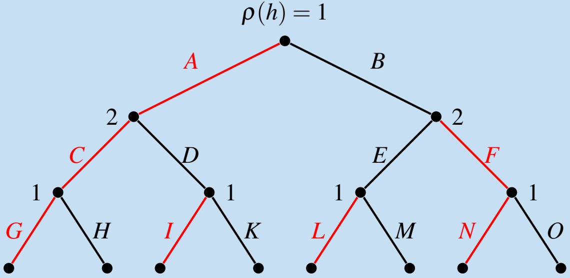

Step 1: One Deviation Property

Proof (ctd.)

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

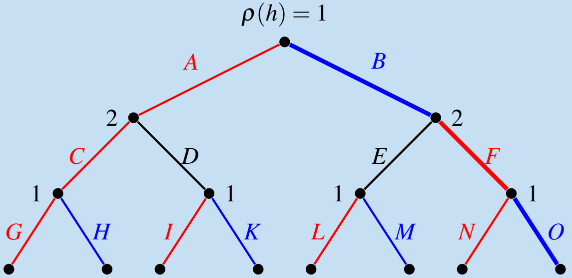

Then only B and O really matter.

Proof (ctd.)

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

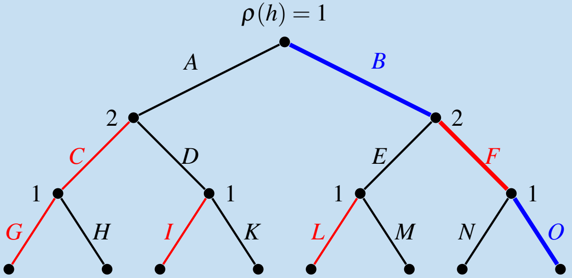

Step 1: One Deviation Property

Proof (ctd.)

- (⇐) …

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

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

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

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

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

Step 1: One Deviation Property

Proof (ctd.)

- ($\Leftarrow$) ...

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

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

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

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

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

Step 1: One Deviation Property

Proof (ctd.)

- ($\Leftarrow$) ...

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

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

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

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

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

Proof (ctd.)

- ($\Leftarrow$) ...

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

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

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

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

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

Step 1: One Deviation Property

Proof (ctd.)

- ($\Leftarrow$) ...

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

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

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

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

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

Step 1: One Deviation Property

Example

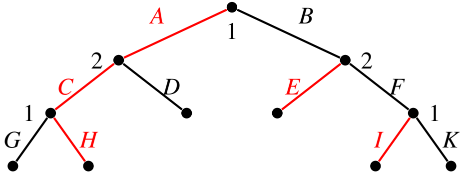

To show that (AHI, CE) is an SPE, it suffices to check these deviant strategies:

Player 1:

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

Player 2:

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

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

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

Counterexample (one-player case):

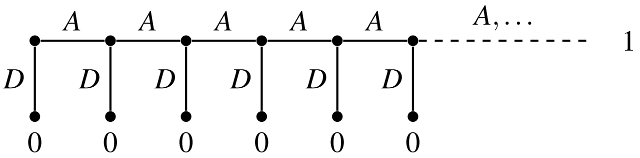

Strategy s_i with $s_i(h) = D$ for all $h \in H \setminus Z$

- satisfies one deviation property, but
- is not an SPE, since it is dominated by s_i^* with $s_i^*(h) = A$ for all $h \in H \setminus Z$.

Kuhn's Theorem

Step 2: Kuhn's Theorem

Theorem (Kuhn)

Every finite extensive game with perfect information has a subgame perfect equilibrium.

Proof idea:

- Proof is constructive and builds an SPE 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

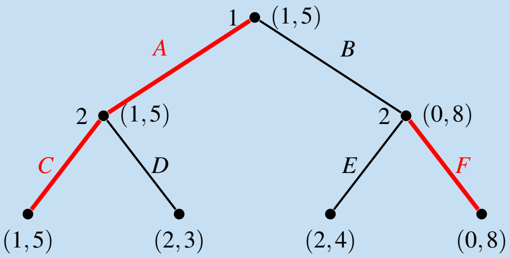

$$s_2(\langle A\rangle) = C \qquad t_1(\langle A\rangle) = 1 \qquad t_2(\langle A\rangle) = 5$$
$$s_2(\langle B\rangle) = F \qquad t_1(\langle B\rangle) = 0 \qquad t_2(\langle B\rangle) = 8$$
$$s_1(\langle\rangle) = A \qquad t_1(\langle\rangle) = 1 \qquad t_2(\langle\rangle) = 5$$

Step 2: Kuhn's Theorem

A bit more formally:

Proof

Let $G = \langle N, A, H, \rho, (u_i) \rangle$ be a finite extensive game with perfect information.

Construct an SPE by induction on $\ell(G(h))$ for all subgames $G(h)$. 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 an SPE in subgame $G(h)$.

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

...

UNI FREIBURG

A bit more formally:

Proof

Let $G = \langle N, A, H, \rho, (u_i) \rangle$ be a finite extensive game with perfect information.

Construct an SPE by induction on $\ell(G(h))$ for all subgames $G(h)$. 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 an SPE in subgame $G(h)$.

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

...

Step 2: Kuhn's Theorem

A bit more formally:

Proof

Let $G = \langle N, A, H, \rho, (u_i) \rangle$ be a finite extensive game with perfect information.

Construct an SPE by induction on $\ell(G(h))$ for all subgames $G(h)$. 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 an SPE in subgame $G(h)$.

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

...

Proof (ctd.)

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

For all $a \in A(h^*)$, $\ell(G(h^*, a)) \leq k$. Let

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

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

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

With the one deviation property lemma it follows that s is an SPE. □

UNI FREIBURG

Proof (ctd.)

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

For all $a \in A(h^*)$, $\ell(G(h^*, a)) \leq k$. 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 an SPE. □

Step 2: Kuhn's Theorem

Proof (ctd.)

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

For all $a \in A(h^*)$, $\ell(G(h^*, a)) \leq k$. 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 an SPE. □

Step 2: Kuhn's Theorem

- In principle: sample SPE 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 $O(b^m)$.

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 SPE in this game.

B) infinite horizon, finite branching factor:

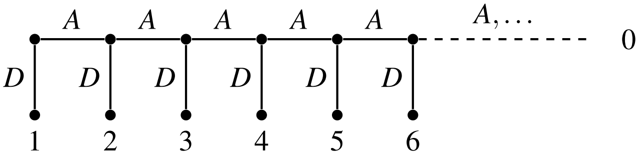

$u_1(AAA\ldots) = 0$ and $u_1(\underbrace{AA\ldots A}_{n}D) = n+1$. No SPE.

Step 2: Kuhn's Theorem

Uniqueness:

Kuhn's theorem tells us nothing about uniqueness of SPEs. However, if no two histories get the same evaluation by any player, the SPE is unique.

Two Extensions

Chance Moves

UNI FREIBURG

Definition

An extensive game with perfect information and chance moves is a tuple $G = \langle N, A, H, \rho, f_c, (u_i) \rangle$, where

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

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 measure on the set of terminal histories.
- Payoffs: For player i, U_i is expected payoff (with weights according to outcome probabilities).

Chance Moves

Example

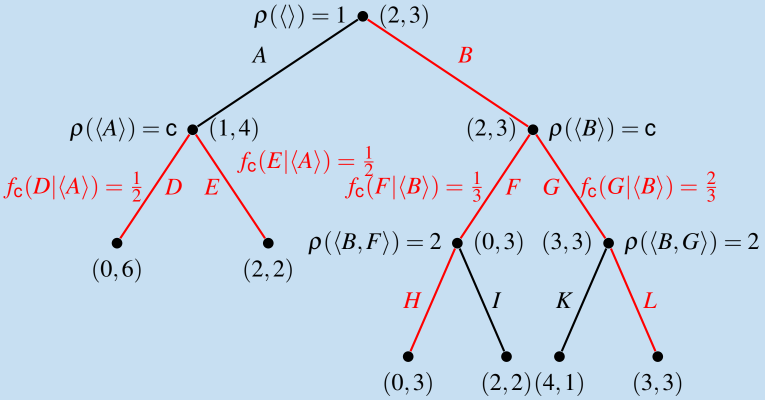

UNI FREIBURG

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.

Simultaneous Moves

Definition

An extensive game with perfect information and simultaneous moves is a tuple $G = \langle N, A, H, \rho, (u_i) \rangle$, where

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

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

Simultaneous Moves

- Intended meaning of simultaneous moves: All players from $\rho(h)$ 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.

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 perfect information and simultaneous moves.

  Example: Matching Pennies can be viewed as extensive game with simultaneous moves. No NE/SPE.

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.

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 deny ("n") the suggested splitting.
- If both accept, each player i gets his allotted share (utility x_i). Otherwise, no player gets anything (utility 0).

UNI FREIBURG

Formally:

$$N = \{1,2,3\}$$

$$X = \{(x_1,x_2,x_3) \in [0,1]^3 \mid x_1 + x_2 + x_3 = 1\}$$

$$H = \{\langle\rangle\} \cup \{\langle x\rangle \mid x \in X\} \cup \{\langle x,z\rangle \mid x \in X, z \in \{\mathsf{y},\mathsf{n}\} \times \{\mathsf{y},\mathsf{n}\}\}$$

$$\rho(\langle\rangle) = \{1\}$$

$$\rho(\langle x\rangle) = \{2,3\} \text{ for all } x \in X$$

$$u_i(\langle x,z\rangle) = \begin{cases} 0 & \text{if } z \in \{(\mathsf{y},\mathsf{n}),(\mathsf{n},\mathsf{y}),(\mathsf{n},\mathsf{n})\} \\ x_i & \text{if } z = (\mathsf{y},\mathsf{y}). \end{cases} \text{ for all } i \in N$$

SPEs:

- 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)

SPEs (ctd.):

- Whole game:
  Let s_2 and s_3 be any 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 $X_\mathrm{y} = \{x \in X \mid s_2(\langle x \rangle) = s_3(\langle x \rangle) = \mathrm{y}\}$ be the set of splits accepted under s_2 and s_3.
  Distinguish three cases:
  - $X_\mathrm{y} = \emptyset$ or $x_1 = 0$ for all $x \in X_\mathrm{y}$. Then (s_1, s_2, s_3) is an SPE for any possible s_1.
  - $X_\mathrm{y} \neq \emptyset$ and there are splits $x_{\max} = (x_1, x_2, x_3) \in X_\mathrm{y}$ that maximize $x_1 > 0$. Then (s_1, s_2, s_3) is an SPE iff $s_1(\langle \rangle)$ is such a split $x_{\max}$.
  - $X_\mathrm{y} \neq \emptyset$ and there are no splits $(x_1, x_2, x_3) \in X_\mathrm{y}$ that maximize x_1. Then there is no SPE, in which player 2 follows strategy s_2 and player 3 follows strategy s_3.

Summary

Summary

- For finite-horizon extensive games with perfect information, 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 with perfect information 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.