Elements of Game Theory

S. Pinchinat

Master2 RI 2011-2012

Introduction

- Economy
- Biology
- Synthesis and Control of reactive Systems
- Checking and Realizability of Formal Specifications
- Compatibility of Interfaces
- Simulation Relations between Systems
- Test Cases Generation
- ...

In this Course

- Strategic Games (2h)

- Extensive Games (2h)

Bibliography

- R.D. Luce and H. Raiffa: "Games and Decisions" (1957) [LR57]
- K. Binmore: "Fun and Games." (1991) [Bin91]
- R. Myerson: "Game Theory: Analysis of Conflict." (1997) [Mye97]
- M.J. Osborne and A. Rubinstein: "A Course in Game Theory." (1994) [OR94]
- Also the very good lecture notes from Prof Bernhard von Stengel (search the web).

Representions

	ℓ	r
T	w_1, w_2	x_1, x_2
B	y_1, y_2	z_1, z_2

- Player 1 has the rows (Top or Bottom) and Player 2 has the columns (left or right):

$$S_1 = \{T, B\} \text{ and } S_2 = \{\ell, r\}$$

- For example, when Player 1 chooses T and Player 2 chooses ℓ, the payoff for Player 1 (resp. Player 2) is w_1 (resp. w_2), that is

$$u_1(T, \ell) = w_1 \text{ and } u_2(T, \ell) = w_2$$

Example

- The Battle of Sexes

	Bach	Stravinsky
Bach	2, 1	0, 0
Stravinsky	0, 0	1, 2

  Strategic Interaction = players wish to coordinate their behaviors but have conflicting interests.

- A Coordination Game

	Mozart	Mahler
Mozart	1, 1	0, 0
Mahler	0, 0	2, 2

  Strategic Interaction = players wish to coordinate their behaviors and have mutual interests.

The Prisoner's Dilemma

The story behind the name "prisoner's dilemma" is that of two prisoners held suspect of a serious crime. There is no judicial evidence for this crime except if one of the prisoners confesses against the other. If one of them confesses, he will be rewarded with immunity from prosecution (payoff 0), whereas the other will serve a long prison sentence (payoff -3). If both confess, their punishment will be less severe (payoff -2 for each). However, if they both "cooperate" with each other by not confessing at all, they will only be imprisoned briefly for some minor charge that can be held against them (payoff -1 for each). The "defection" from that mutually beneficial outcome is to confess, which gives a higher payoff no matter what the other prisoner does, which makes "confess" a dominating strategy (see later). However, the resulting payoff is lower to both. This constitutes their "dilemma".

	confess	don't Confess
Confess	$-2, -2$	$0, -3$
Don't Confess	$-3, 0$	$-1, -1$

The Prisoner's Dilemma

	confess	don't Confess
Confess	$-2, -2$	$0, -3$
Don't Confess	$-3, 0$	$-1, -1$

- The best outcome for both players is that neither confess.
- Each player is inclined to be a "free rider" and to confess.

A 3-player Games

	M_1 L	M_1 R	M_2 L	M_2 R	M_3 L	M_3 R	M_4 L	M_4 R
T	8	0	4	0	0	0	3	3
B	0	0	0	4	0	8	3	3

Player 1 chooses one of the two rows;
Player 2 chooses one of the two columns;
Player 3 chooses one of the four tables.

(Finite) Strategic Games

A finite strategic game with n-players is $\Gamma = (N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N})$ where

- $N = \{1, \ldots, n\}$ is the set of players.
- $S_i = \{1, \ldots, m_i\}$ is a set of pure strategies (or actions) of player i.
- $u_i : S \rightarrow R$ is the payoff (utility) function.

$$S := S_1 \times S_2 \times \ldots \times S_n \text{ is the set of profiles.}$$

$s = (s_1, s_2, \ldots, s_n)$

Instead of u_i, use preference relations:

$$s' \succsim_i s \text{ for Player } i \text{ prefers profile } s' \text{ than profile } s$$

$$\Gamma = (N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N}) \text{ or } \Gamma = (N, \{S_i\}_{i \in N}, \{\succsim_i\}_{i \in N})$$

Comments on Interpretation

A strategic game describes a situation where

- we have a one-shot even
- each player knows
  - the details of the game.
  - the fact that all players are rational (see futher)
- the players choose their actions “simultaneously” and independently.

Comments on Interpretation

A strategic game describes a situation where

- we have a one-shot even
- each player knows
  - the details of the game.
  - the fact that all players are rational (see futher)
- the players choose their actions “simultaneously” and independently.

Rationality: Every player wants to maximize its own payoff.

Notations

- Use $s_i \in S_i$, or simply $j \in S_i$ where $1 \leq j \leq m_i$.
- Given a profile $s = (s_1, s_2, \ldots, s_n) \in S$, we let a counter profile be an element like

$$s_{-i} := (s_1, s_2, \ldots, s_{i-1}, \text{empty}, s_{i+1}, \ldots, s_n)$$

  which denotes everybody's strategy except that of Player i, and write S_{-i} for the set of such elements.
- For $r_i \in S_i$, let $(s_{-i}, r_i) := (s_1, s_2, \ldots, s_{i-1}, r_i, s_{i+1}, \ldots, s_n)$ denote the new profile where Player i has switched from strategy s_i to strategy r_i.

Dominance

Let $s_i, s_i' \in S_i$.

s_i strongly dominates s_i' if

$$u_i(s_{-i}, s_i) > u_i(s_{-i}, s_i') \text{ for all } s_{-i} \in S_{-i},$$

s_i (weakly) dominates s_i' if

$$\begin{cases} u_i(s_{-i}, s_i) \geq u_i(s_{-i}, s_i'), \text{ for all } s_{-i} \in S_{-i}, \\ u_i(s_{-i}, s_i) > u_i(s_{-i}, s_i'), \text{ for some } s_{-i} \in S_{-i}. \end{cases}$$

Example of Dominance

The Prisoner's Dilemma

	c	d
C	$-2, -2$	$0, -3$
D	$-3, 0$	$-1, -1$

- Strategy C of Player 1 strongly dominates strategy D.
- Because the game is symmetric, strategy c of Player 2 strongly dominates strategy d.
- Note also that $u_1(D, d) > u_1(C, c)$ (and also $u_2(D, d) > u_2(C, c)$), so that
  "C dominates D" does not mean "C is always better than D".

Example of Weak Dominance

	ℓ	r
T	$1,3$	$1,3$
B	$1,1$	$1,0$

ℓ (weakly) dominates r.

Elimination of Dominated Strategies

- If a strategy is dominated, the player can always improve his payoff by choosing a better one (this player considers the strategies of the other players as fixed).
- Turn to the game where dominated strategies are eliminated.

	c	d
C	$-2, -2$	$0, -3$
D	$-3, 0$	$-1, -1$

The game becomes simpler.

- Eliminating D and d, shows (C,c) as the "solution" of the game, i.e. a recommandation of a strategy for each player.

Iterated Elimination of Dominated Strategies

- We consider iterated elimination of dominated strategies.
- The result does not depend on the order of elimination:
  If s_i (strongly) dominates s'_i, it still does in a game where some strategies (other than s'_i) are eliminated.
- In contrast, for iterated elimination of weakly dominated strategies
  the order of elimination may matter
  EXERCISE: find examples, in books

A game is dominance solvable if the Iterated Elimination of Dominated Strategies ends in a single strategy profile.

Motivations

- Not every game is dominance solvable, e.g. Battle of Sexes.

	Bach	Stravinsky
Bach	$2, 1$	$0, 0$
Stravinsky	$0, 0$	$1, 2$

- The central concept is that of Nash Equilibrium [Nas50].

Best Response and Nash Equilibrium

- Informally, a Nash equilibrium is a strategy profile where each player's strategy is a best response to the counter profile.
- Formally:
  Given a strategy profile $s = (s_1, \ldots, s_n)$ in a strategic game $\Gamma = (N, \{S_i\}_{i \in N}, \{\succsim_i\}_{i \in N})$, a strategy s_i is a best response (to s_{-i}) if

$$(s_{-i}, s_i) \succsim_i (s_{-i}, s_i'), \text{ for all } s_i' \in S_i$$

  A Nash Equilibrium in a profile $s^* = (s_1^*, \ldots, s_n^*) \in S$ such that s_i^* is best response to s_{-i}^*, for all $i = 1, \ldots, n$.

Player cannot gain by changing unilateraly her strategy, i.e. with the remained strategies kept fixed.

Illustration

	ℓ	r
T	2 2	1 0
B	0 3	1 1

Illustration

	ℓ	r
T	2, 2	0, 1
B	3, 0	1, 1

draw best reponse in boxes

	ℓ	r
T	2, 2	0, 1
B	3, 0	1, 1

Examples

- Battle of Sexes: two Nash Equilibria.

	Bach	Stravinsky
Bach	2,1	0, 0
Stravinsky	0, 0	1,2

- Mozart-Mahler: two Nash Equilibria.

	Mozart	Mahler
Mozart	1,1	0, 0
Mahler	0, 0	2,2

- Prisoner's Dilemma: EXERCISE

	c	d
C	$-2, -2$	$0, -3$
D	$-3, 0$	$-1, -1$

Dominance and Nash Equilibrium (NE)

- A dominated strategy is never a best response
  $\Rightarrow$ it cannot be part of a NE.
- We can eliminate dominated strategy without loosing any NE.
- Elimination does not create new NE (best response remains when adding a dominated strategy).

Proposition

If a game is dominance solvable, its solution is the only NE of the game.

Nash Equilibria May Not Exist

Matching Pennies: a strictly competitive game.

	Head	Tail
Head	$1, -1$	$-1, 1$
Tail	$-1, 1$	$1, -1$

- Each player chooses either Head or Tail.
- If the choices differ, Player 1 pays Player 2 a dollar;
  if they are the same, Player 2 pays Player 1 a dollar.
- The interests of the players are diametrically opposed.

Those games are often called zero-sum games.

"Extended" Strategic Games

- We have seen that NE need not exist when players deterministically choose one of their strategies (e.g. Matching Pennies)
- If we allow players to non-deterministically choose, then

  NE always exist (Nash Theorem)
- By "non-deterministically" we mean that the player randomises his own choice.
- We consider mixed strategies instead of only pure strategies considered so far.

Mixed Strategies

A mixed (randomized) strategy x_i for Player i is a probability distribution over S_i. Formally, it is a vector $x_i = (x_i(1), \dots, x_i(m_i))$ with

$$\begin{cases} x_i(j) \geq 0, & \text{for all } j \in S_i \text{, and} \\ x_i(1) + \ldots + x_i(m_i) = 1 \end{cases}$$

Let X_i be the set of mixed strategies for Player i.

$$X := X_1 \times X_2 \times \ldots \times X_n \text{ is the set of (mixed) profiles.}$$

- A mixed strategy x_i is pure if $x_i(j) = 1$ for some $j \in S_i$ and $x_i(j') = 0$ for all $j' \neq j$. We use $\pi_{i,j}$ to denote such a pure strategy.
- The support of the mixed strategy x_i is

$$\text{support}(x_i) := \{\pi_{i,j} \mid x_i(j) > 0\}$$

We use e.g. x_{-i} (counter (mixed) profile), X_{-i} (counter profiles), etc.

Interpretation of Mixed Strategies

- Assume given a mixed strategy (probability distribution) of a player.
- This player uses a lottery device with the given probabilities to pick each pure strategy according to its probability.
- The other players are not supposed to know the outcome of the lottery.
- A player bases his own decision on the resulting distribution of payoffs, which represents the player's preference.

Expected Payoffs

- Given a profile $x = (x_1, \dots, x_n)$ of mixed strategies, and a combination of pure strategies $s = (s_1, \dots, s_n)$, the probability of combinaison s under profile x is

$$x(s) := x_1(s_1) * x_2(s_2) * \dots * x_n(s_n)$$

- The expected payoff of Player i is the mapping $U_i : X \to R$ defined by

$$U_i(x) := \sum_{s \in S} x(s) * u_i(s)$$

It coincides with the notion of payoff when only pure strategies are considered.

Mixed Extension of Strategic Games and Mixed Strategy Nash Equilibrium

- The mixed extension of $\Gamma = (N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N})$ is the strategic game

  $$(N, \{X_i\}_{i \in N}, \{U_i\}_{i \in N})$$

  (CONVENTION: we still use u_i instead U_i.)

- A mixed strategy Nash Equilibrium of a strategic game $\Gamma = (N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N})$ is a Nash equilibrium of its mixed extension.

Focus on 2-player Games

- Sets of strategies $\{1, \ldots, n\}$ and $\{1, \ldots, m\}$ respectively;
- Payoffs functions u_1 and u_2 are described by $n \times m$ matrices:

$$U_1 := \begin{bmatrix} u_1(1,1) & u_1(1,2) & \ldots & u_1(1,m) \\ u_1(2,1) & u_1(2,2) & \ldots & u_1(2,m) \\ \ldots & \ldots & \ldots & \ldots \\ u_1(n,1) & u_1(n,2) & \ldots & u_1(n,m) \end{bmatrix} \quad U_2 \ldots$$

- $x_1 = (p_1, \ldots, p_n) \in X_1$ and $x_2 = (q_1, \ldots, q_m) \in X_2$ for mixed strategies; that is $p_j = x_1(j)$ and $q_k = x_2(k)$.
- Write $(U_1 x_2)_j$ the j-th component of vector $U_1 x_2$.

$(U_1 x_2)_j = \sum_{k=1}^{m} u_1(j,k) * q_k$	is the expected payoff of Player 1 when playing row j.

- we shall write $x_1(U_1 x_2)$ instead of $x_1^T(U_1 x_2)$.

Given a mixed profile $x = (x_1, x_2)$,

- The expected payoff of Player 1 is

$$x_1(U_1 x_2) = \sum_{j=1}^{n} \sum_{k=1}^{m} p_j * u_1(j,k) * q_k$$

- The expected payoff of Player 2 is

$$(x_1 U_2) x_2 = \sum_{j=1}^{n} \sum_{k=1}^{m} p_j * u_2(j,k) * q_k$$

Existence: the Nash Theorem

Theorem (Nash 1950)

Every finite strategic game has a mixed strategy NE.

We omit the proof in this course and refer to the literature.

Remarks

- The assumption that each S_i is finite is essential for the proof.
- The proof uses Brouwer's Fixed Point Theorem.

  Theorem

  Let Z be a subset of some space R^N that is convex and compact, and let f be a continuous function from Z to Z. Then f has at least one fixed point, that is, a point $z \in Z$ so that $f(z) = z$.

The Best Response Property

Theorem

Let $x_1 = (p_1, \dots, p_n)$ and $x_2 = (q_1, \dots, q_m)$ be mixed strategies of Player 1 and Player 2 respectively. Then x_1 is a best response to x_2 iff for all pure strategies j of Player 1,

$$p_j > 0 \ \Rightarrow \ (U_1 x_2)_j = max\{(U_1 x_2)_k \,|\, 1 \leq k \leq n\}$$

Proof Recall that $(U_1 x_2)_j$ is the the expected payoff of Player 1 when playing row j. Let $u := max\{(U_1 x_2)_j \,|\, 1 \leq j \leq n\}$. Then

$$\begin{aligned} x_1 U_1 x_2 &= \sum_{j=1}^{n} p_j (U_1 x_2) j = \sum_{j=1}^{n} p_j [u - [u - (U_1 x_2) j]] \\ &= \sum_{j=1}^{n} p_j * u - \sum_{j=1}^{n} p_j [u - (U_1 x_2) j] = u - \sum_{j=1}^{n} p_j [u - (U_1 x_2) j] \end{aligned}$$

Since $p_j \geq 0$ and $u - (U_1 x_2)_j \geq 0$, we have $x_1 U_1 x_2 \leq u$. The expected payoff $x_1 U_1 x_2$ achieves the maximum u iff $\sum_{j=1}^{n} p_j [u - (U_1 x_2) j = 0$, that is $p_j > 0$ implies $(U_1 x_2)_j = u$.

Consequences of the Best Response Property

Only pure strategies that get maximum, and hence equal, expected payoff can be played with postive probability in NE.

Proposition

A propfile (x_1^*, x_2^*) is a NE if and only if there exists $w_1, w_2 \in R$ such that

- *for every $j \in support(x_1^*)$, $(U_1 x_2^*)_j = w_1$, and*
  for every $j \notin support(x_1^*)$, $(U_1 x_2^*)_j \leq w_1$.
- *for every $k \in support(x_2^*)$, $(x_1^* U_2)_k = w_2$. and*
  for every $k \notin support(x_2^*)$, $(x_1^* U_2)_k \leq w_2$.

Finding Mixed Nash Equilibria (1)

We characterize NE: suppose we know the supports $support(x_1^*) \subseteq S_1$ and $support(x_2^*) \subseteq S_2$ of some NE (x_1^*, x_2^*). We consider the following system of constraints over the variables $p_1, \ldots, p_n, q_1, \ldots, q_m, w_1, w_2$:

Write $x_1^* = (p_1, \ldots, p_n)$ and $x_2^* = (q_1, \ldots, q_m)$.

$$\begin{cases} (U_1 x_2^*)_j = w_1 & \text{for all } 1 \leq j \leq n \text{ with } p_j \neq 0 \\ (x_1^* U_2)_k = w_2 & \text{for all } 1 \leq k \leq m \text{ with } q_k \neq 0 \\ \sum_{j=1}^{n} p_j = 1 & \\ \sum_{k=1}^{m} q_k = 1 & \\ \text{EXERCISE : what is missing ?} & \end{cases}$$

Notice that this system is a Linear Programming system – however the system is no longer linear if $n > 2$ –. Use e.g. Simplex algorithm

Finding Mixed Nash Equilibria (1)

We characterize NE: suppose we know the supports $support(x_1^*) \subseteq S_1$ and $support(x_2^*) \subseteq S_2$ of some NE (x_1^*, x_2^*). We consider the following system of constraints over the variables $p_1, \ldots, p_n, q_1, \ldots, q_m, w_1, w_2$:

Write $x_1^* = (p_1, \ldots, p_n)$ and $x_2^* = (q_1, \ldots, q_m)$.

$$\begin{cases} (U_1 x_2^*)_j = w_1 & \text{for all } 1 \leq j \leq n \text{ with } p_j \neq 0 \\ (x_1^* U_2)_k = w_2 & \text{for all } 1 \leq k \leq m \text{ with } q_k \neq 0 \\ \sum_{j=1}^{n} p_j = 1 & \\ \sum_{k=1}^{m} q_k = 1 & \\ \text{EXERCISE : what is missing ?} & \end{cases}$$

Notice that this system is a Linear Programming system – however the system is no longer linear if $n > 2$ –. Use e.g. Simplex algorithm

For each possible $support(x_1^*) \subseteq S_1$ and $support(x_2^*) \subseteq S_2$ solve the system

$\Rightarrow$ Worst-case exponential time (2^{m+n} possibilities for the supports).

Finding Mixed Nash Equilibria (2)

Battle of Sexes

	b	s
B	1, 2	0, 0
S	0, 0	2, 1

As seen before, we already have two NE with pure strategies (B, b) and (S, s).

We now determine the mixed strategy probabilities of a player so as to make the other player indifferent between his or her pure strategies, because only then that player will mix between these strategies. This is a consequence of the Best Response Theorem: only pure strategies that get maximum, and hence equal, expected payoff can be played with positive probability in equilibrium.

Finding Mixed Nash Equilibria (2)

Battle of Sexes

	b	s
B	$\boxed{1}$, $\boxed{2}$	0, 0
S	0, 0	$\boxed{2}$, $\boxed{1}$

Suppose Player I plays B with prob. $1 - p$ and S with prob. p. The best response for Player II is b when $p \longrightarrow 0$, whereas it is s when $p \longrightarrow 1$.

$\Rightarrow$ There is some probability so that Player II is indifferent.

Finding Mixed Nash Equilibria (2)

Battle of Sexes

	b	s
B	$\boxed{2}$ $\boxed{1}$	0 0
S	0 0	$\boxed{1}$ $\boxed{2}$

Fix p (probability) the mixed strategy of Player I. The expected payoff of Player II when she plays b is $2(1-p)$, and it is p when she plays s.

She is indifferent whenever $2(1-p) = p$, that is $p = 2/3$.

If Player I plays B with prob $1/3$ and S with prob $2/3$, Player II has an expected payoff of $2/3$ for both strategies.

Finding Mixed Nash Equilibria (2)

Battle of Sexes

	b	s
B	2 1	0 0
S	0 0	1 2

If Player I plays the mixed strategy $(1/3, 2/3)$, Player II has an expected payoff of $2/3$ for both strategies.

Finding Mixed Nash Equilibria (2)

Battle of Sexes

	b	s
B	[1], [2]	0, 0
S	0, 0	[2], [1]

If Player I plays the mixed strategy $(1/3, 2/3)$, Player II has an expected payoff of $2/3$ for both strategies.

Then Player II can mix between b and s. A similar calculation shows that Player I is indifferent between C and S if Player II uses the mixed strategy $(2/3, 1/3)$, and Player I has an expected payoff of $2/3$ for both strategies.

Finding Mixed Nash Equilibria (2)

Battle of Sexes

	b	s
B	$\boxed{2}$ $\boxed{1}$	0 0
S	0 0	$\boxed{1}$ $\boxed{2}$

If Player I plays the mixed strategy $(1/3, 2/3)$, Player II has an expected payoff of $2/3$ for both strategies.

Then Player II can mix between b and s. A similar calculation shows that Player I is indifferent between C and S if Player II uses the mixed strategy $(2/3, 1/3)$, and Player I has an expected payoff of $2/3$ for both strategies.

The profile of mixed strategies $((1/3, 2/3), (2/3, 1/3))$ is the mixed NE.

Finding Mixed Nash Equilibria (3)

- The difference trick method
- The upper envelope method

Matching Pennies

	Head	Tail
Head	$1, -1$	$-1, 1$
Tail	$-1, 1$	$1, -1$

$$x_1^*(Head) = x_1^*(Tail) = x_2^*(Head) = x_2^*(Tail) = \frac{1}{2}$$

is the unique mixed strategy NE.

Here, support$(x_1^*) = \{\pi_{1Head}, \pi_{1Tail}\}$ and support$(x_2^*) = \{\pi_{2Head}, \pi_{2Tail}\}$.

EXERCISE: apply previous techniques to find it yourself.

Complements

- Particular classes of games, e.g. symmetric games, degenerated games, ...
- Bayesian games for Games with imperfect information.
- Solutions for $N \geq 3$: there are examples where the NE has irrational values. See Nash 1951 for a Poker game.
- ...

- Strategic Games (2h)

- Extensive Games (2h)

By Extensive Games,

we implicitly mean

"Extensive Games with Perfect Information".

Example

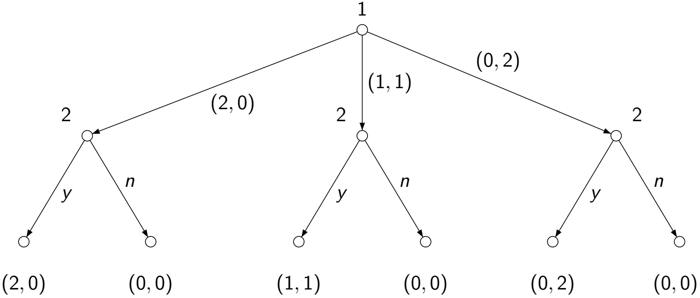

$H = \{\epsilon, (2, 0), (1, 1), (0, 2), ((2, 0), y), \ldots\}$.
Each history denotes a unique node in the game tree, hence we often use the terminology decision nodes instead.

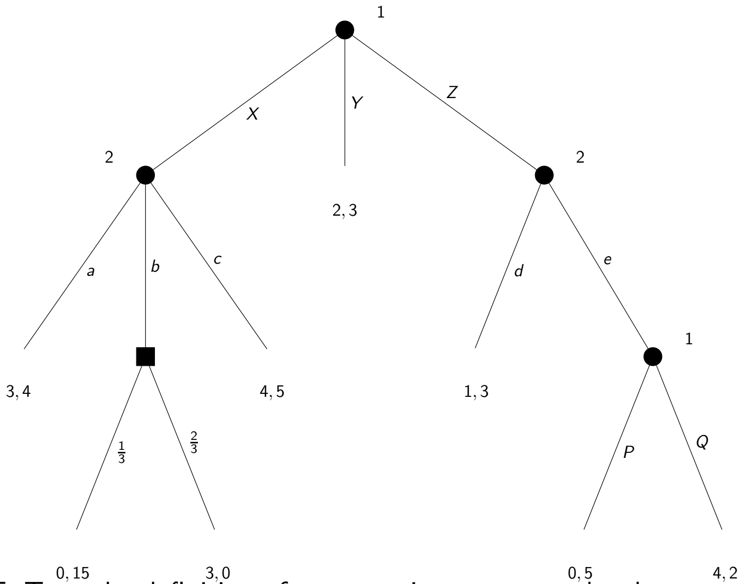

EXERCISE: Tune the definition of an extensive game to take chance nodes into account.

Extensive Games (with Perfect Information)

An extensive game (with perfect information) is a tuple $\mathcal{G} = (N, A, H, P, \{\succsim_i\}_{i=1,\dots,n})$ where

- $N = \{1, \dots, n\}$ is a set of players.
- Write $A = \cup_i A_i$, where A_i is the set of action of Player i.
- $H \subseteq A^*$ is a set of (finite) sequences of actions s.t.
  - The empty sequence $\epsilon \in H$.
  - H is prefix-closed.

  The elements of H are histories; we identify histories with the decision nodes they lead to. A decision node is terminal whenever it is (reached by a history) of the form $h = a^1 a^2 \dots a^K$ and there is no $a^{K+1} \in A$ such that $a^1 a^2 \dots a^K a^{K+1} \in H$. We denote by Z the set of terminal histories.
- $P : H \setminus Z \to N$ indicates whose turn it is to play in a given non-terminal decision node.
- Each $\succsim_i \subseteq Z \times Z$ is a preference relation.

We write $A(h) \subseteq A_{P(h)}$ for the set of actions available to Player $P(h)$ at decision node h.

Strategies and Strategy Profiles

A strategy of a player in an extensive game is a "plan".
Formally, let $\mathcal{G} = (N, H, P, \{\succsim_i\}_{i=1,\dots,n})$ be an extensive game (from now on, we omit the set A of actions).

A strategy of Player i is (a partial mapping) $s_i : H \setminus Z \to A$ whose domain is $\{h \in H \setminus Z \mid P(h) = i\}$ and such that $s_i(h) \in A(h)$.
Write S_i for the set of strategies of Player i.

Note that the definition of a strategy only depends on the game tree (N, H, P), and not on the preferences of the players.

Example

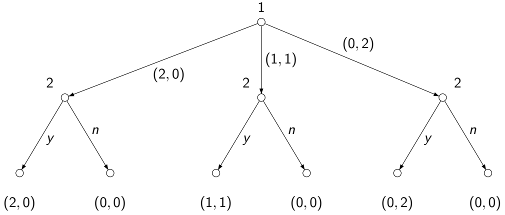

- Player 1 plays at decision node ϵ (she starts the game), and this is the only one; she has 3 strategies $s_1 = (2, 0), s'_1 = (1, 1), s"_1 = (0, 2)$.
- Player 2 takes an action after each of the three histories, and in each case it has 2 possible actions (y or n); we write this as abc ($a, b, c \in \{y, n\}$), meaning that after history $(2, 0)$ Player 2 chooses action a, after history $(1, 1)$ Player 2 chooses action b, and after history $(0, 2)$ Player 2 chooses action c.

Example

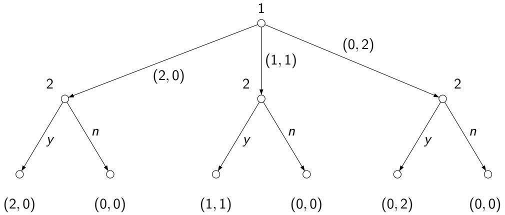

- Player 1 plays at decision node ϵ (she starts the game), and this is the only one; she has 3 strategies $s_1 = (2,0), s_1' = (1,1), s"_1 = (0,2)$.
- Player 2 takes an action after each of the three histories, and in each case it has 2 possible actions (y or n); we write this as abc ($a, b, c \in \{y, n\}$), meaning that after history $(2,0)$ Player 2 chooses action a, after history $(1,1)$ Player 2 chooses action b, and after history $(0,2)$ Player 2 chooses action c. Possible strategies for Player 2 are yyy, yyn, yny, ynn, ... There are 2^3 possibilities.

Example

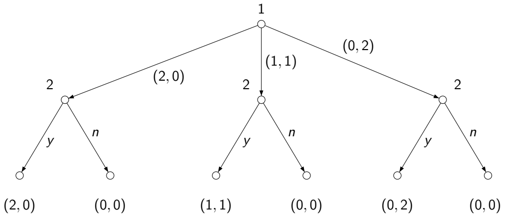

- Player 1 plays at decision node ϵ (she starts the game), and this is the only one; she has 3 strategies $s_1 = (2, 0), s_1' = (1, 1), s"_1 = (0, 2)$.
- Possible strategies for Player 2 are *yyy*, *yyn*, *yny*, *ynn*, ... There are 2^3 possibilities.
- In general the number of strategies of a given Player i is in $O(|A_i|^m)$ where m is the number of decision nodes where Player i plays.

Remarks on the Definition of Strategies

- Strategies of players are defined even for histories that are not reachable if the strategy is followed.

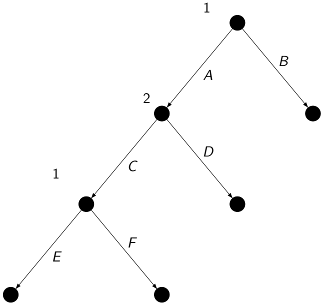

- In this game, Player 1 has four strategies AE, AF, BE, BF: By BE, we specify a strategy after history e.g. AC even if it is specified that she chooses action B at the beginning of the game.

Reduced Strategies

- A reduced strategy of a player specifies a move for the decision nodes of that player, but unreachable nodes due to an earlier own choice.
- It is important to discard a decision node only on the basis of the earlier own choices of the player only.
- A reduced profile is a tuple of reduced strategies.

Outcomes

The outcome of a strategy profile $s = (s_1, \dots, s_n)$ in an extensive game $\mathcal{G} = (N, H, P, \{\succsim_i\}_{i \in N})$, written O(s), is the terminal decision node that results when each Player i follows the precepts s_i:

$O(s)$ is the history $a^1 a^2 \in Z$ s.t. for all (relevant) $k > 1$, we have

$$s_{P(a^1 \dots a^{k-1})}(a^1 \dots a^{k-1}) = a^k$$

Example of Outcomes

$O(AE, C)$ has payoff $v_!$

...

$O(BE, C)$ has payoff v_4

...

$O(BF, D)$ has payoff v_4

...

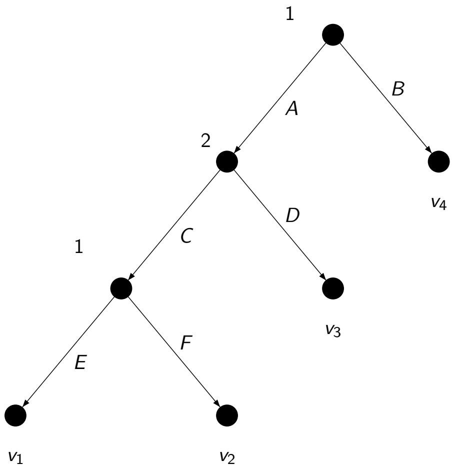

Strategic Form of an Extensive Game

The strategic form of an extensive game $\mathcal{G} = (N, H, P, \{\succsim_i\}_{i \in N})$ is the strategic game $(N, \{S_i\}_{i \in N}, \{\succsim'_i\}_{i \in N})$ where $\succsim'_i \subseteq S_i \times S_i$ is defined by

$$s \succsim'_i s' \text{ whenever } O(s) \succsim_i O(s')$$

As the number of strategies grows exponentially with the number m of decision nodes in the game tree – recall it is in $O(|A_i|^m)$ –, strategic forms of extensive games are in general big objects.

Extensive Games as Strategic Games

	C	D
AE	v_1	v_3
AF	v_2	v_3
BE	v_4	v_4
BF	v_4	v_4

Reduced Strategies:

	C	D
AE	v_1	v_3
AF	v_2	v_3
B	v_4	v_4

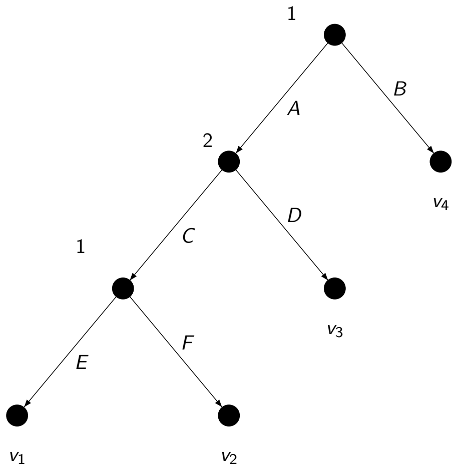

Nash Equilibrium in Extensive Games

Simply use the definition of NE for the strategic game form of the extensive game.

Examples of Nash Equilibria

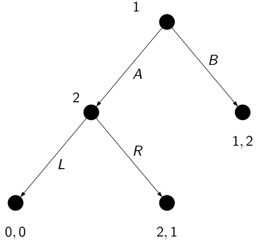

- Two NE (A, R) and (B, L) with payoffs are $(2, 1)$ and $(1, 2)$ resp.
- (B, L) is a NE because:
  - given that Player 2 chooses L after history A, it is always optimal for Player 1 to choose B at the beginning of the game – if she does not, then given Player 2's choice, she obtains 0 rather than 1.

Examples of Nash Equilibria

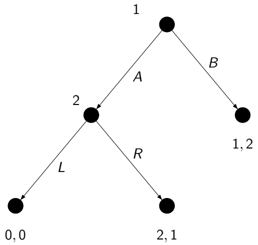

- Two NE (A, R) and (B, L) with payoffs are $(2, 1)$ and $(1, 2)$ resp.
- (B, L) is a NE because:
  - given that Player 2 chooses L after history A, it is always optimal for Player 1 to choose B at the beginning of the game – if she does not, then given Player 2's choice, she obtains 0 rather than 1.
  - given Player 1's choice of B, it is always optimal for Player 2 to play L.
- The equilibrium (B, L) lacks plausibility.

The good notion is the Subgame Perfect Equilibrium

Subgame Perfect Equilibrium

We define the subgame of an extensive game $\mathcal{G} = (N, H, P, \{\succsim_i\}_{i \in N})$ that follows a history h as the extensive game

$$\mathcal{G}(h) = (N, h^{-1}H, P|_h, \{\succsim_i |_h\}_{i \in N})$$

where

- $h^{-1}H := \{h' \mid hh' \in H\}$
- $P|_h(h') := P(hh')$ ...
- $h' \succsim_i |_h h''$ whenever $hh' \succsim_i hh''$

$\mathcal{G}(h)$ is the extensive game which starts at decision node h.

$\mathcal{G}$ and $\mathcal{G}(h)$, for $h = A$

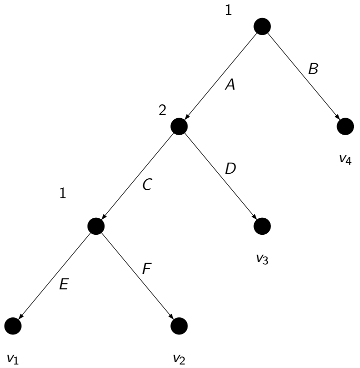

$\rightsquigarrow$

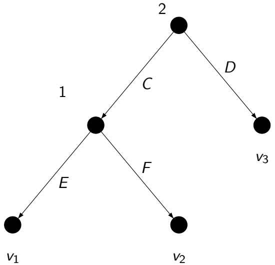

Subgame Perfect (Nash) Equilibrium

A subgame perfect equilibrium represents a Nash equilibrium of every subgame of the original game. More formally,

A subgame perfect equilibrium (SPE) of $\mathcal{G} = (N, H, P, \{\succsim_i\}_{i \in N})$ is a strategy profile s^* s.t. for every Player $i \in N$ and every non-terminal history $h \in H \setminus Z$ for which $P(h) = i$, we have

$$O|_h(s^*_{-i}|_h, s^*_i|_h) \succsim_i |_h O|_h(s^*_{-i}|_h, s_i)$$

for every strategy s_i of Player i in the subgame $\mathcal{G}(h)$.

Example

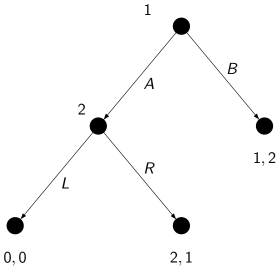

NE were (A, R) and (B, L), but the only SPE is (A, R).

Another Example

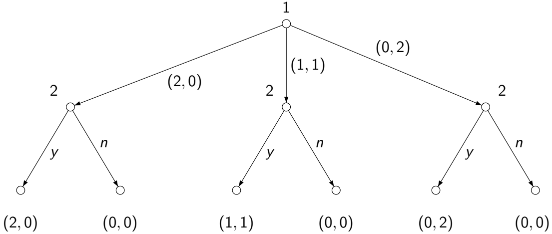

Nash Equilibrium are:

- $((2,0), yyy)$, $((2,0), yyn)$, $((2,0), yny)$, and $((2,0), ynn)$ for the division $(2,0)$.
- $((1,1), nyy)$, $((1,1), nyn)$, and $((1,1), nyn)$ for the division $(1,1)$.
- ...

EXERCISE: What are the SPE?

Computing SPE of Finite Extensive Games (with Perfect Information): Backward Induction

Backward Induction is a procedure to construct a strategy profile which is a SPE.

Start with the decision nodes that are closest to the leaves, consider a history h in a subgame with the assumption that a strategy profile has already been selected in all the subgames $\mathcal{G}(h, a)$, with $a \in A(h)$. Among the actions of $A(h)$, select an action a that maximises the (expected) payoff of Player $P(h)$.

This way, an action is specified for each history of the game $\mathcal{G}$, which determines an entire strategy profile.

Backward Induction Example

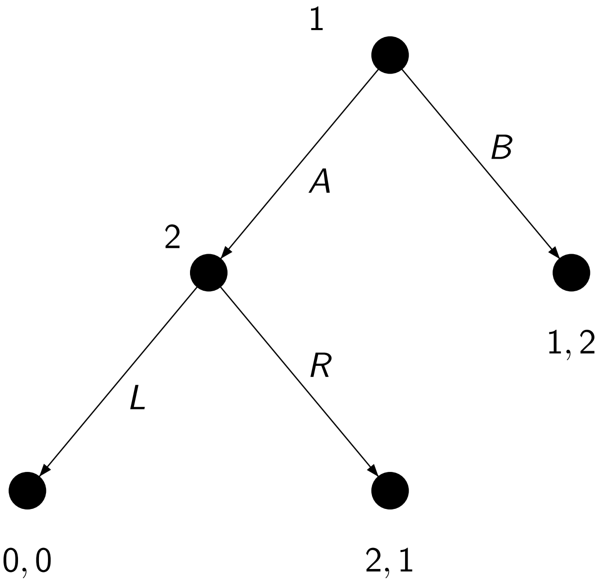

Backward Induction Example

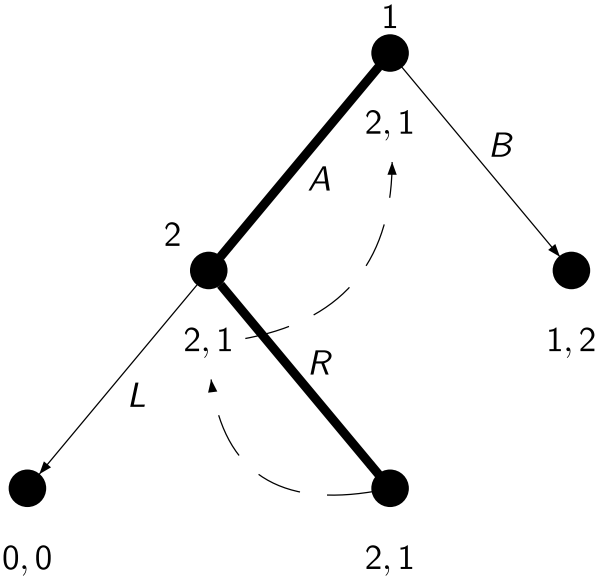

Theorem: Backward Induction defines an SPE.

We prove inductively (and also consider chance nodes): consider a non-terminal history h. Suppose that $A(h) = \{a_1, a_2, \ldots, a_m\}$ each a_j leading to the subgame $\mathcal{G}_j$, and assume, as inductive hypothesis, that the strategy profiles that have been selected by the procedure so far (in the subgames $\mathcal{G}_j$'s) define a SPE.

Theorem: Backward Induction defines an SPE.

We prove inductively (and also consider chance nodes): consider a non-terminal history h. Suppose that $A(h) = \{a_1, a_2, \ldots, a_m\}$ each a_j leading to the subgame $\mathcal{G}_j$, and assume, as inductive hypothesis, that the strategy profiles that have been selected by the procedure so far (in the subgames $\mathcal{G}_j$'s) define a SPE.

If h is a chance node, then the BI procedure does not select a particular action from h. For every player, the expected payoff in the subgame $\mathcal{G}(h)$ is the expectation of the payoffs in each subgame $\mathcal{G}_j$ (weighted with probability to move to $\mathcal{G}_j$). If a player could improve on that payoff, she would have to do so by changing her strategy in one subgame $\mathcal{G}_j$ which contradicts the induction hypothesis.

Theorem: Backward Induction defines an SPE.

We prove inductively (and also consider chance nodes): consider a non-terminal history h. Suppose that $A(h) = \{a_1, a_2, \ldots, a_m\}$ each a_j leading to the subgame $\mathcal{G}_j$, and assume, as inductive hypothesis, that the strategy profiles that have been selected by the procedure so far (in the subgames $\mathcal{G}_j$'s) define a SPE.

Assume now that h is a decision node ($P(h) \in N$).

Every Player $i \neq P(h)$ can improve his payoff only by changing his strategy in a subgame, which contradicts the induction hypothesis.

Theorem: Backward Induction defines an SPE.

We prove inductively (and also consider chance nodes): consider a non-terminal history h. Suppose that $A(h) = \{a_1, a_2, \ldots, a_m\}$ each a_j leading to the subgame $\mathcal{G}_j$, and assume, as inductive hypothesis, that the strategy profiles that have been selected by the procedure so far (in the subgames $\mathcal{G}_j$'s) define a SPE.

Assume now that h is a decision node ($P(h) \in N$).

Player $P(h)$ can improve her payoff, she has to do it by changing her strategy: the only way is to change her local choice to a_j together with changes in the subgame $\mathcal{G}_j$. But the resulting improved payoff would only be the improved payoff in $\mathcal{G}_j$, itself, which contradicts the induction hypothesis.

Consequences of the Theorem

Backward Induction defines an SPE.

(In extensive games with perfect information)

Consequences of the Theorem

Backward Induction defines an SPE.

(In extensive games with perfect information)

Corollary

By the BI procedure, each player's action can be chosen deterministically, so that pure strategies suffice. It is not necessary, as in strategic games, to consider mixed strategies.

Consequences of the Theorem

Backward Induction defines an SPE.

(In extensive games with perfect information)

Corollary

By the BI procedure, each player's action can be chosen deterministically, so that pure strategies suffice. It is not necessary, as in strategic games, to consider mixed strategies.

Corollary

Subgame Perfect Equilibrium always exist.
For game trees, we can use "SPE" synonymously with "strategy profile obtained by backward induction".

An example of an exam (2009)

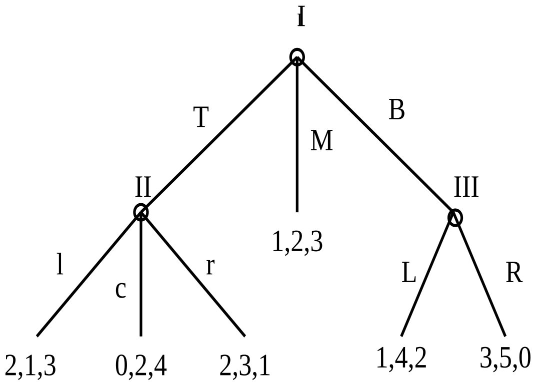

1. How many strategy profiles does this game have?
2. Identify all pairs of strategies where one strategy strictly, or weakly, dominates the other.
3. Find all Nash equilibria in pure strategies. Which of these are subgame perfect?

What we have not seen

- Combinatorial game theory
  A mathematical theory that studies two-player games which have a position in which the players take turns changing in defined ways or moves to achieve a defined winning condition.

What we have not seen

- Combinatorial game theory
  A mathematical theory that studies two-player games which have a position in which the players take turns changing in defined ways or moves to achieve a defined winning condition.
- Game trees with imperfect information
  A player does not know exactly what actions other players took up to that point. Technically, there exists at least one information set with more than one node.

What we have not seen

- Combinatorial game theory
  A mathematical theory that studies two-player games which have a position in which the players take turns changing in defined ways or moves to achieve a defined winning condition.
- Game trees with imperfect information
  A player does not know exactly what actions other players took up to that point. Technically, there exists at least one information set with more than one node.
- Games and logic
  Important examples are semantic games used to define truth, back-and-forth games used to compare structures, and dialogue games to express (and perhaps explain) formal proofs.

What we have not seen

- Combinatorial game theory
  A mathematical theory that studies two-player games which have a position in which the players take turns changing in defined ways or moves to achieve a defined winning condition.
- Game trees with imperfect information
  A player does not know exactly what actions other players took up to that point. Technically, there exists at least one information set with more than one node.
- Games and logic
  Important examples are semantic games used to define truth, back-and-forth games used to compare structures, and dialogue games to express (and perhaps explain) formal proofs.
- others ...

M.J. Osborne and A. Rubinstein.

A Course in Game Theory.

MIT Press, 1994.

Ken Binmore.

Fun and Games - A Text on Game Theory.

D. C. Heath & Co., 1991.

Roger B. Myerson.

Game Theory: Analysis of Conflict.

Harvard University Press, September 1997.

R.D. Luce and H. Raiffa.

Games and Decisions.

J. Wiley, New York, 1957.

J.F. Nash.

Equilibrium points in n-person games.

Proceedings of the National Academy of Sciences of the United States of America, 36:48–49, 1950.

Andrés Perea.

Rationality in extensive form games.

Boston : Kluwer Academic Publishers, 2001.

E. Grädel, W. Thomas, and T. Wilke, editors.

Automata, Logics, and Infinite Games: A Guide to Current Research [outcome of a Dagstuhl seminar, February 2001], volume 2500 of *Lecture Notes in Computer Science*.

Springer, 2002.