NORMAL FORM GAMES: invariance and refinements
DYNAMIC GAMES: extensive form

(slides from Nax-Pradelski)

Heinrich H Nax & Heiko Rauhut

hnax@ethz.ch

March 19, 2019: Lecture 5

ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Plan

- **Normal form games**
  - Equilibrium invariance
  - Equilibrium refinements
- **Dynamic games**
  - Extensive form games
  - Incomplete information
  - Sub-game perfection

Nash's equilibrium existence theorem

Theorem (Nash 1951)

Every finite game has at least one [Nash] equilibrium in mixed strategies.

Cook book: How to find mixed Nash equilibria

- Find all pure strategy NE.

Check whether there is an equilibrium in which row mixes between several of her strategies:

- Identify candidates:
  - If there is such an equilibrium then each of these strategies must yield the same expected payoff given column's equilibrium strategy.
  - Write down these payoffs and solve for column's equilibrium mix.
  - Reverse: Look at the strategies that column is mixing on and solve for row's equilibrium mix.
- Check candidates:
  - The equilibrium mix we found must indeed involve the strategies for row we started with.
  - All probabilities we found must indeed be probabilities (between 0 and 1).
  - Neither player has a positive deviation.

Battle of the Sexes revisited

PLAYERS The players are the two students $N = \{row, column\}$.

STRATEGIES Row chooses from $S_{row} = \{Cafe, Pub\}$
Column chooses from $S_{column} = \{Cafe, Pub\}$.

PAYOFFS For example, $u_{row}(Cafe, Cafe) = 4$. The following matrix summarises:

	$Cafe(q)$	$Pub(1-q)$	Expected
$Cafe(p)$	$\underline{4}, \underline{3}$	$1, 1$	$4q+(1-q)$
$Pub(1-p)$	$0, 0$	$\underline{3}, \underline{4}$	$3(1-q)$
Expected	$3p$	$p+4(1-p)$	

Column chooses $q = 1$ whenever $3p > p + 4(1-p) \Leftrightarrow 6p > 4 \Leftrightarrow p > \frac{2}{3}$.

Row chooses $p = 1$ whenever $4q + (1-q) > 3(1-q) \Leftrightarrow 6q > 2 \Leftrightarrow q > \frac{1}{3}$.

Battle of the Sexes: Best-reply graph

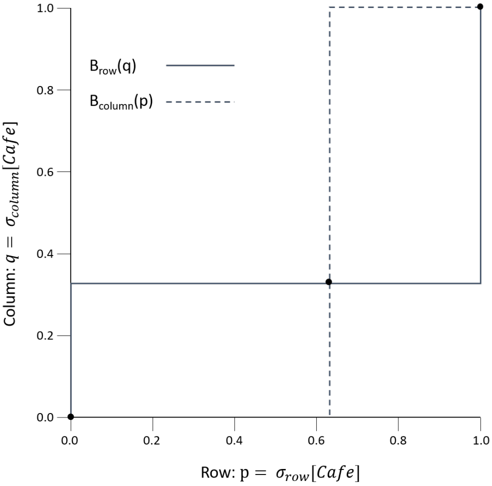

There is a mixed Nash equilibrium with $p = \frac{2}{3}$ and $q = \frac{1}{3}$.

Battle of the Sexes: Expected payoff

	$Cafe(1/3)$	$Pub(2/3)$	Expected
$Cafe(2/3)$	$4, 3$	$1, 1$	$4{\cdot}1/3{+}2/3$
$Pub(1/3)$	$0, 0$	$3, 4$	$3{\cdot}2/3$
Expected	$3 \cdot 2/3$	$2/3 + 4 \cdot 1/3$	

Frequency of play:

	$Cafe(1/3)$	$Pub(2/3)$
$Cafe(2/3)$	$2/9$	$4/9$
$Pub(1/3)$	$1/9$	$2/9$

Expected utility to row player: 2

Expected utility to column player: 2

Example

	L	R
T	$0, 0$	$\underline{3}, \underline{5}$
B	$\underline{2}, \underline{2}$	$\underline{3}, 0$

There are two pure-strategy Nash equilibria, at (B, L) and (T, R).

If row player places probability p on T and probability $1 - p$ on B.

$\Rightarrow$ Column player's best reply is to play L if $2(1 - p) \geq 5p$, i.e., $p \leq \frac{2}{7}$.

If column player places probability q on L and $(1 - q)$ on R.

$\Rightarrow$ B is a best reply. T is only a best reply to $q = 0$.

The best-reply graph

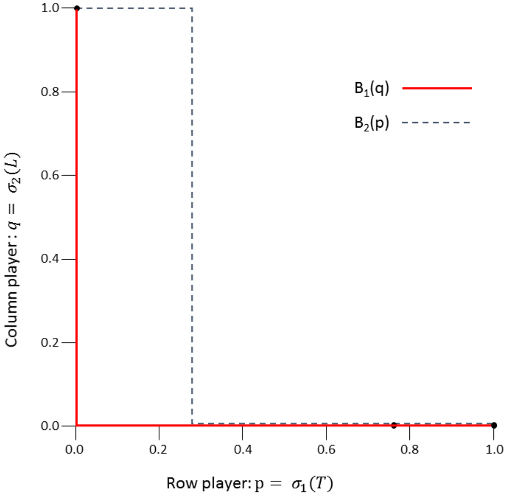

There is a *continuum* of mixed equilibria at $\frac{2}{7} \leq p \leq 1$, all with $q = 0$.

Example: Expected payoffs of mixed NEs

	L	R
T	$0, 0$	$\underline{3}, \underline{5}$
B	$\underline{2}, \underline{2}$	$\underline{3}, 0$

Frequency of play:

	$Cafe(0)$	$Pub(1)$
$Cafe(p > 2/7)$	0	p
$Pub(1 - p)$	0	$1 - p$

Expected utility to row player: 3

Expected utility to column player: $5 \cdot p \in (10/7 \approx 1.4, 5]$

Weakly and strictly dominated strategies

	L	R
T	$0, 0$	$\underline{3}, \underline{5}$
B	$\underline{2}, \underline{2}$	$\underline{3}, 0$

Note that T is *weakly dominated* by B.

- A weakly dominated pure strategy may play a part in a mixed (or pure) Nash equilibrium.
- A strictly dominated pure strategy cannot play a part in a Nash equilibrium!
  - Any mixed strategy which places positive weight on a strictly dominated pure strategy is itself strictly dominated. This can be seen by moving weight away from the dominated strategy.

Odd number of Nash equilibria

Theorem (Wilson, 1970)

Generically, any finite normal form game has an odd number of Nash equilibria.

“Generically” = if you slightly change payoffs the set of Nash equilibria does not change.

Returning to our example

	L	R
T	$0, 0$	$\underline{3}, \underline{5}$
B	$\underline{2}, \underline{2}$	$\underline{3}, 0$

There are two pure-strategy Nash equilibria, at (B, L) and (T, R).
There is a *continuum* of mixed equilibria at $\frac{2}{7} \leq p \leq 1$, all with $q = 0$.

The best-reply graph

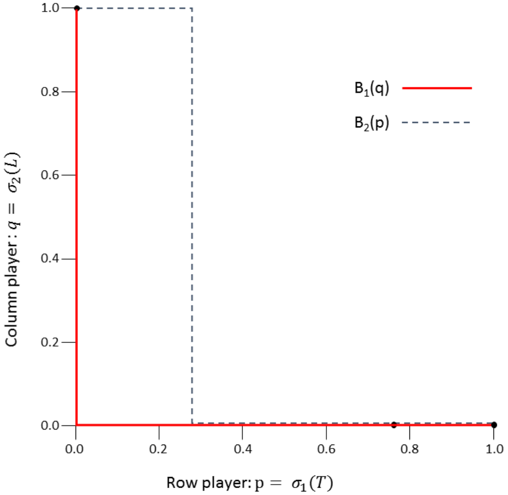

There is a *continuum* of mixed equilibria at $\frac{2}{7} \leq p \leq 1$, all with $q = 0$.

Example: Expected utility of mixed NEs

	L	R
T	$0, 0$	$\underline{3\ \ .1}, \underline{5}$
B	$\underline{2}, \underline{2}$	$\underline{3}, 0$

There are two pure-strategy Nash equilibria, at (B, L) and (T, R).

If row player places probability p on T and probability $1 - p$ on B.

$\Rightarrow$ Column player's best reply is to play L if $2(1 - p) \geq 5p$, i.e., $p \leq \frac{2}{7}$.

If column player places probability q on L and $(1 - q)$ on R.

$\Rightarrow$ Row player's best reply is to play T if $3.1(1 - q) \geq 2q + 3(1 - q)$, i.e., $q \leq 0.1/2.1$.

The unique mixed strategy equilibrium is where $p = 2/7$ and $q = 0.1/2.1$.

The best-reply graph

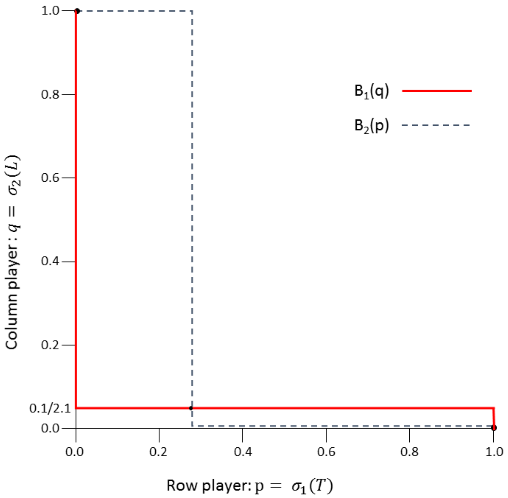

There is a an odd number of equilibria.

Coordination game

	Email	*Fax*
Email	$\underline{5}, \underline{5}$	$1, 1$
Fax	$0, 0$	$\underline{3}, \underline{4}$

The two pure Nash equilibria are $\{Email, Email\}$ and $\{Fax, Fax\}$.

The unique mixed equilibrium is given by row player playing $\sigma_1 = (1/2, 1/2)$ and column player playing $\sigma_2 = (2/7, 5/7)$

Invariance of Nash equilibria

Proposition

Any two games G, G' which differ only by a positive affine transformation of each player's payoff function have the same set of Nash equilibria.

Adding a constant c to all payoffs of some player i which are associated with any fixed pure combination s_i for the other players sustains the set of Nash equilibria.

Coordination game

Now apply the transformation $u' = 2 + 3 \cdot u$ to the row player's payoffs:

	Email	*Fax*
Email	$\underline{5}, \underline{5}$	$1, 1$
Fax	$0, 0$	$\underline{3}, \underline{4}$

	Email	*Fax*
Email	$\underline{17}, \underline{5}$	$5, 1$
Fax	$2, 0$	$\underline{11}, \underline{4}$

The two pure Nash equilibria remain $\{Email, Email\}$ and $\{Fax, Fax\}$.

The unique mixed equilibrium is again given by row player playing $\sigma_1 = (1/2, 1/2)$ and column player playing $\sigma_2 = (2/7, 5/7)$

Some remarks on Nash equilibrium

Nash equilibrium is a very powerful concept since it exists (in finite settings)!

But there are often a multitude of equilibria. Therefore game theorists ask which equilibria are more or less likely to be observed.

We will focus next on a static refinements, strict and perfect equilibrium.

Later we will talk about dynamic refinements.

Strict Nash equilibria

Definition: Strict Nash Equilibrium

A *strict Nash equilibrium* is a profile σ^* such that,

$$U_i(\sigma_i^*, \sigma_{-i}^*) > U_i(\sigma_i, \sigma_{-i}^*) \text{ for all } \sigma_i \text{ and } i.$$

Perfect equilibrium or "trembling hand" perfection

Selten: 'Select these equilibria which are robust to small "trembles" in the player's strategy choices'

Definition: ε-perfection

Given any $\varepsilon \in (0,1)$, a strategy profile σ is ε-perfect if it is interior ($x_{ih} > 0$ for all $i \in N$ and $h \in S_i$) and such that:

$$h \notin \beta_i(x) \Rightarrow x_{ih} \leq \varepsilon$$

Definition: Perfect equilibrium

A strategy profile σ is perfect if it is the limit of some sequence of ε_t-perfect strategy profiles x^t with $\varepsilon_t \to 0$.

Perfect equilibrium or “trembling hand” perfection

Example:

	L	R
T	$1,1$	$1,0$
B	$1,0$	$0,0$

There are two pure Nash equilibibra B,L and T,L. The mixed equilibrium is such that column player plays L and row player plays any interior mix.

Only T,L is perfect.

Note that T,L is not strict.

Perfect equilibrium or “trembling hand” perfection

Proposition (Selten 1975)

For every finite game there exists at least one perfect equilibrium. The set of perfect equilibria is a subset of the set of Nash equilibria.

Proposition

Every strict equilibrium is perfect.

Dynamic games

Many situations (games) are characterized by sequential decisions and information about prior moves

- Market entrant vs. incumbent (think BlackBerry vs. Apple iPhone)
- Chess
- ...

When such a game is written in strategic form, important information about timing and information is lost.

Solution:

- Extensive form games (via game trees)
- Discussion of timing and information
- New equilibrium concepts

Example: perfect information

Battle of the sexes:

	a	b
A	3, 1	0, 0
B	0, 0	1, 3

What if row player (player 1) can decide first?

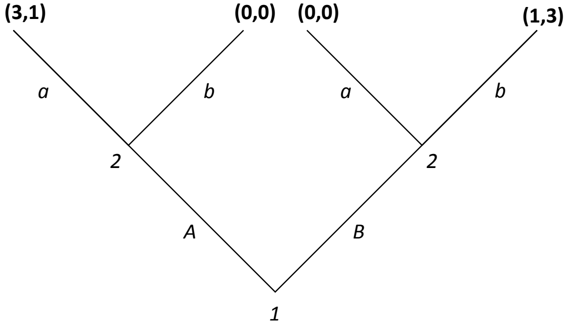

Example: perfect information

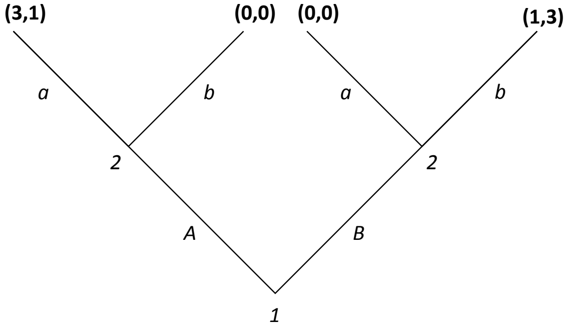

What would you do as player 1, A or B?
What would you do as player 2 if player 1 played A, a or b?
What would you do as player 2 if player 1 played B, a or b?

Example: perfect information

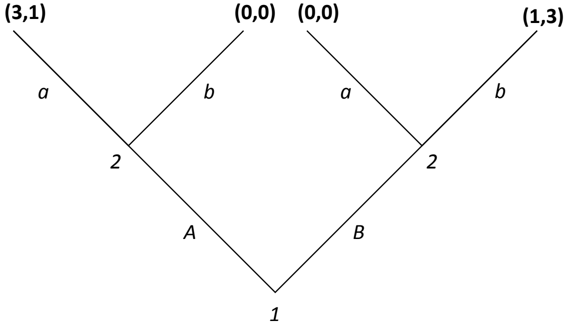

Player 2 would like to commit that if player 1 plays A he will play b (in order to make player 1 play B).
But fighting is not time consistent. Once player 1 played A it is not rational for player 2 to play b.
The expected outcome is A followed by a for payoffs $(3, 1)$.
This is called **backward induction**. It results in a **subgame perfect equilibrium**. More later!

Example: imperfect information

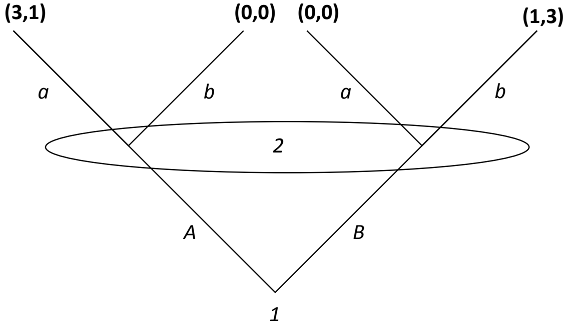

What would you do as player 1, A or B?
What would you do as player 2, a or b?

Timing and information matters!

Extensive form game: Definition

An **extensive-form game** is defined by:

- **Players**, $N = \{1, \dots, n\}$, with typical player $i \in N$. Note: *Nature* can be one of the players.
- Basic structure is a tree, the **game tree** with nodes $a \in A$. Let a_0 be the root of the tree.
- Nodes are game states which are either
  - **Decision nodes** where some player chooses an action
  - **Chance nodes** where nature plays according to some probability distribution

Representation

Extensive form

- Directed graph with single initial node; edges represent moves
- Probabilities on edges represent Nature moves
- Nodes that the player in question cannot distinguish (information sets) are circled together (or connected by dashed line)

Extensive form → normal form

- A strategy is a player's complete plan of action, listing move at every information set of the player
- Different extensive form games may have same normal form (loss of information on timing and information)

Question: What is the number of a player's strategies?
Product of the number of actions available at each of his information sets.

Subgames (Selten 1965, 1975)

Given a node *a* in the game tree consider the subtree rooted at *a*. *a* is the root of a subgame if

- *a* is the only node in its information set
- if a node is contained in the subgame then all its successors are contained in the subgame
- every information set in the game either consists entirely of successor nodes to *a* or contains no successor node to *a*.

If a node *a* is a subroot, then each player, when making a choice at any information set in the game, knows whether *a* has been reached or not. Hence if *a* has been reached it is as if a “new” game has started.

Subgame examples

How many subgames does the game have?

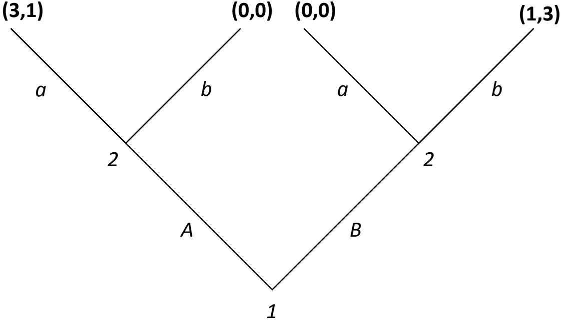

Which strategies does each player have?

Strategies player 1: $\{A, B\}$

Strategies player 2: $\{(a,a),(a,b),(b,a),(b,b)\}$

Subgame examples

How many subgames does the game have?

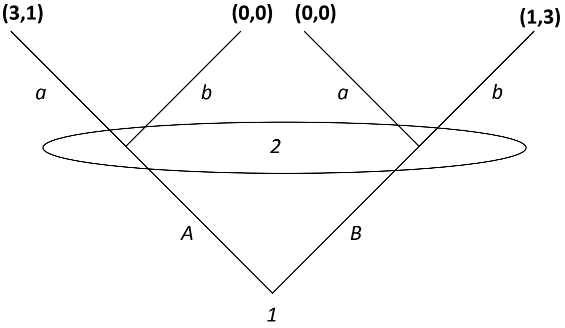

Which strategies does each player have?

Strategies player 1: $\{A, B\}$

Strategies player 2: $\{a, b\}$

Subgame examples: Equivalence to normal form

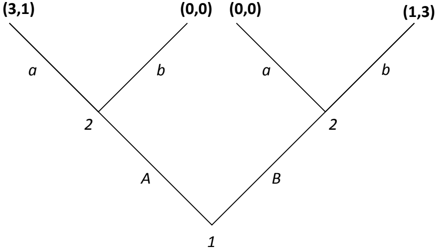

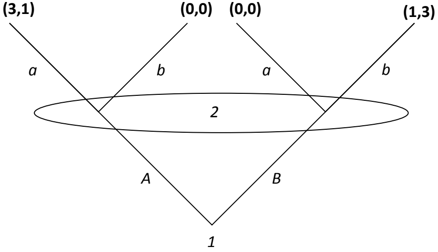

	a,a	a,b	b,a	b,b
A	$3,1$	$3,1$	$0,0$	$0,0$
B	$0,0$	$1,3$	$0,0$	$1,3$

	a	b
A	$3,1$	$0,0$
B	$0,0$	$1,3$

where columns strategies are of the form *strategy against A, strategy against B*

Strategies in extensive games

PURE STRATEGY s_i One move for each information set of the player.

MIXED STRATEGY σ_i Any probability distribution x_i over the set of pure strategies S_i.

BEHAVIOR STRATEGY y_i Select randomly at each information set the move to be made (can delay coin-toss until getting there).

Behavior strategies are special case of a mixed strategy: moves are made with independent probabilities at information sets.

Pure strategies are special case of a behavior strategy.

Example (imperfect recall)

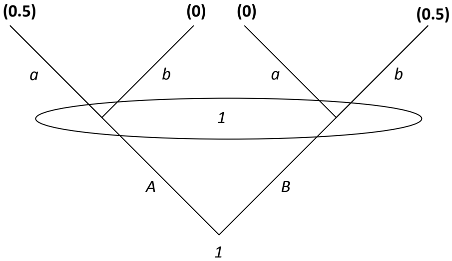

There is one player who has "forgotten" his first move when his second move comes up. (For example: did he lock the door before leaving or not?)

The indicated outcome, with probabilities in brackets, results from the mixed strategy, $\frac{1}{2}Aa + \frac{1}{2}Bb$.

$\Rightarrow$ There exists no behavior strategy that induces this outcome.

The player exhibits "poor memory" / "imperfect recall".

Perfect recall

Perfect recall (Kuhn 1950)

Player i in an extensive form game has *perfect recall* if for every information set h of player i, all nodes in h are preceded by the same sequence of moves of player i.

Kuhn's theorem

Definition: Realization equivalent

A mixed strategy σ_i is *realization equivalent* with a behavior strategy y_i if the realization probabilities under the profile σ_i, σ_{-i} are the same as those under y_i, σ_{-i} for all profiles σ.

Kuhn's theorem

Consider a player i in an extensive form with perfect recall. For every mixed strategy σ_i there exists a realization-equivalent behavior strategy y_i.

Kuhn's Theorem - proof (not part of exam)

Given: mixed strategy σ

Wanted: realization equivalent behavior strategy y

Idea: $y =$ observed behavior under σ

$y(c) =$ observed probability $\sigma(c)$ of making move c.

What is $\sigma(c)$?

Look at sequence ending in c, here Ibc.

$\sigma[lbc]$ = probability of lbc under $\sigma = \sigma(l, b, c)$.

Sequence lb leading to info set h

$$\mu[lb] = \sigma(l, b, c) + \sigma(l, b, d)$$

$$\Rightarrow \sigma[lb] = \sigma[lbc] + \sigma[lbd]$$

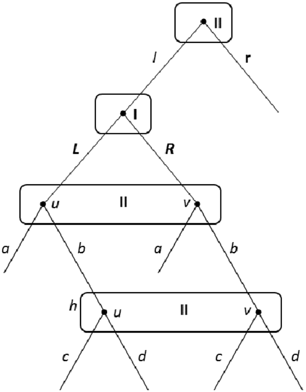

Kuhn's Theorem - proof (not part of exam)

$$\Rightarrow \sigma(c) = \frac{\sigma[lbc]}{\sigma[lb]} =: y(c)$$

$$\Rightarrow \sigma(b) = \frac{\sigma[l]}{\sigma[l]} =: y(b)$$

first info set: $\sigma[\emptyset] = 1 = \sigma[l] + \sigma[r]$

$$\sigma(l) = \frac{\sigma[l]}{\sigma[\emptyset]} =: y(l)$$

$$\Rightarrow y(l)y(b)y(c) = \frac{\sigma[l]}{\sigma[\emptyset]} \cdot \frac{\sigma[l]}{\sigma[l]} \cdot \frac{\sigma[lbc]}{\sigma[lb]} = \sigma[lbc]$$

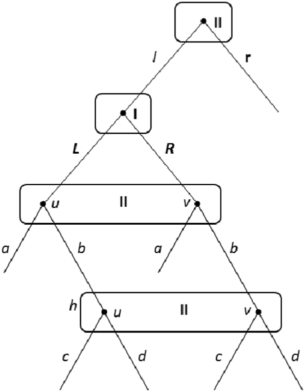

$\Rightarrow$ $\boxed{y \text{ equivalent to } \sigma}$

Subgame perfect equilibrium

Definition: subgame perfect equilibrium (Selten 1965)

A behavior strategy profile in an extensive form game is a *subgame perfect equilibrium* if for each subgame the restricted strategy is a Nash equilibrium of the subgame.

Theorem

Every finite game with prefect recall has at least one subgame perfect equilibrium. Generic such games have a unique subgame perfect equilibrium.

Generic = with probability 1 when payoffs are drawn from continuous independent distributions.

Example: An Outside-option game

Reconsider the battle-of-sexes game (BS game), but player 1 can decide if she joins the game before.

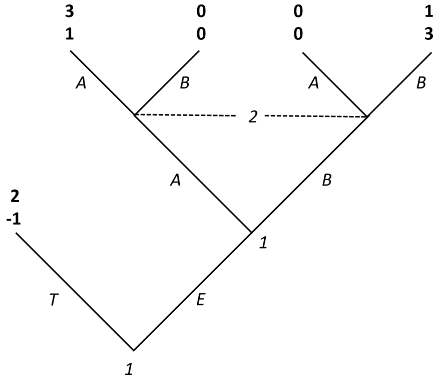

- What are the subgames?
- What are the subgame perfect equilibria?

Example: An Outside-option game

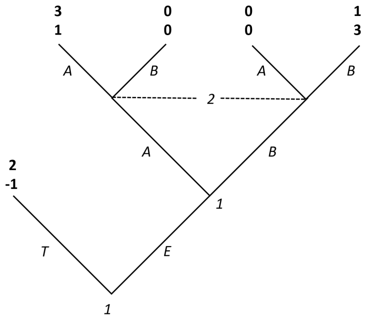

If player 1 decides to enter the BS subgame, player 2 will know that player 1 joint, but will not know her next move.

There exist three subgame perfect equilibria, one for each Nash equilibria of the BS game:

- $S = \{EA, A\}$ Player 1 earns 3, Player 2 earns 1.
- $S = \{TB, B\}$ Player 1 earns 2, Player 2 earns -1.
- $S = \{T(3/4 \cdot A + 1/4 \cdot B), (1/4 \cdot A + 3/4 \cdot B)\}$ Player 1 earns 2, Player 2 earns -1.

Cook-book: Backward induction

"Reasoning backwards in time":

- First consider the last time a decision might be made and choose what to do (that is, find Nash equilibria) at that time
- Using the former information, consider what to do at the second-to-last time a decision might be made
- ...
- This process terminates at the beginning of the game, the found behavior strategies are subgame prefect equilibria

Example: The Centiped game (Rosenthal)

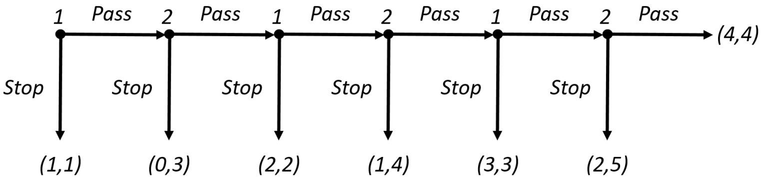

What is the unique subgame perfect equilibirum?

Stop at all nodes.

But in experiments most subjects *Pass* initially: a "trust bubble" forms.

Palacios-Huerta & Volij:

- Chess masters stop right away; students do not…
- …unless they are told they are playing chess masters.

THANKS EVERYBODY
See you next week!
And keep checking the website for new materials as we progress:
http://www.gametheory.online