Intro

- Mark Voorneveld, mark.voorneveld@hhs.se

- Main topic: extensive form games (= extensive games = games in extensive form).

- Parts from chapters 6, 11, and 12 in Osborne and Rubinstein (OsRu).

- I refer explicitly to the definitions/results in OsRu in these slides.

- Before going into more advanced stuff, we'll be formalizing some of the things you have seen in Jörgen's part of the course.

- Disadvantage: elaborate notation.

Def. 89.1 An *extensive form game with perfect information*

$$\Gamma = \langle N, H, P, (\succsim_i)_{i \in N} \rangle$$

consists of

☒ nonempty, finite set N of *players*,

☒ set H of *histories*, summarizing the sequence of actions taken so far:

- initial history: $\emptyset \in H$,
- H contains the initial segments of all histories: if $(a^k)_{k=1,\ldots,K} \in H$ (possibly $K = \infty$) and $L < K$ (fewer actions), then $(a^k)_{k=1,\ldots,L} \in H$,
- if an infinite seq $(a^k)_{k \in \mathbb{N}}$ satisfies that all its initial sequences $(a^k)_{k=1,\ldots,K}$ with $K \in \mathbb{N}$ are in H, then so is the infinite sequence itself.

A history is *terminal* if it is infinite or there is no “longer” history:

$(a^k)_{k=1,\dots,K}$ is terminal if there is no a^{K+1} such that $(a^k)_{k=1,\dots,K+1} \in H$.

The set of terminal histories is denoted by $Z \subseteq H$.

☒ *player function* $P : H \setminus Z \to N$ assigning to each nonterminal history a player whose turn it is to take an action,

☒ for each player $i \in N$ a *preference relation* $\succsim_i$ over terminal histories. Often assumed to satisfy the vNM axioms, in which case we represent them by a Bernouilli function u_i on Z.

Example 1.

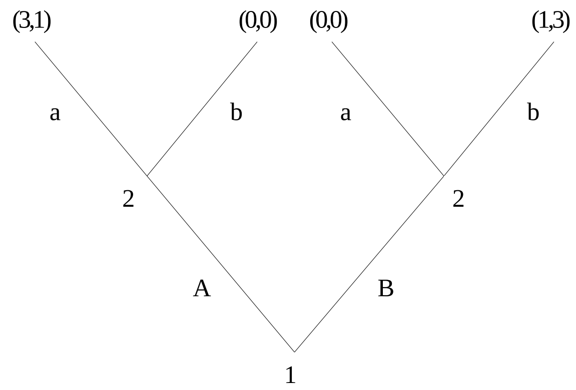

$$N = \{1, 2\}$$

$$H = \{\underbrace{\emptyset}_{\text{length 0}}, \underbrace{A, B}_{\text{length 1}}, \underbrace{(A,a), (A,b), (B,a), (B,b)}_{\text{length 2}}\}$$

$$Z = \{(A,a), (A,b), (B,a), (B,b)\}$$

$$P(\emptyset) = 1, P(A) = P(B) = 2$$

In terminal history (A, a), player 1 gets "payoff" 3 and player 2 gets payoff 1, etc.

Other terminology/notational conventions:

☒ If the set of histories H is a finite set, the game is called *finite*.

☒ If every history $h \in H$ is a finite sequence, the game has a *finite horizon*.

☒ The game's *horizon* is the length of its longest history.The example above has horizon/length 2.

☒ So: finite game $\Rightarrow$ finite horizon (what about the converse?!)

☒ Let $h = (a^1, \ldots, a^K) \in H$ and let a^{K+1} be such that $(a^1, \ldots, a^K, a^{K+1}) \in H$. The latter history is sometimes abbreviated as (h, a^{K+1}).

☒ For nonterminal history $h \in H \setminus Z$, we define

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

as the set of actions that player $P(h)$ can choose in history h.

☒ **Def. 92.1:** a *pure strategy* s_i of player $i \in N$ assigns to each history $\{h \in H \setminus Z : P(h) = i\}$ where it is i's turn to choose, a feasible action $s_i(h) \in A(h)$. The set of pure strategies is denoted by S_i.

In Example 1, pl 1 chooses after $\emptyset$ between A or B: two pure strategies, for convenience denoted as A and B.

Pl 2 chooses after histories A and B, namely between a and b: four pure strategies, denoted aa, ab, ba, bb, where, for instance ab is shorthand for

$$s_2(A) = a, s_2(B) = b.$$

☒ Each pure strategy profile $s = (s_i)_{i \in N} \in S = \times_{i \in N} S_i$ determines a unique *outcome*/terminal history $O(s)$.

☒ **Def. 93.1:** a *(pure) Nash equilibrium* of an extensive form game with perfect information is a pure-strategy profile s^* such that no player can profitably deviate to another strategy:

$$\text{for all } i \in N \text{ and all pure strategies } s_i \text{ of pl } i: \qquad O(s^*) \succsim_i O(s^*_{-i}, s_i).$$

☒ This is just a pure Nash equilibrium of the associated strategic form game (**def. 94.1**) where the pure strategies are those of the extensive form game and the preferences/Bernouilli functions are defined over the outcomes induced by the strategies.

The pure Nash equilibria of the game in Example 1 are just the pure Nash equilibria of the strategic game

	aa	ab	ba	bb
A	$3,1$	$3,1$	$0,0$	$0,0$
B	$0,0$	$1,3$	$0,0$	$1,3$

Reduced games and equivalent strategies:

Example 2.

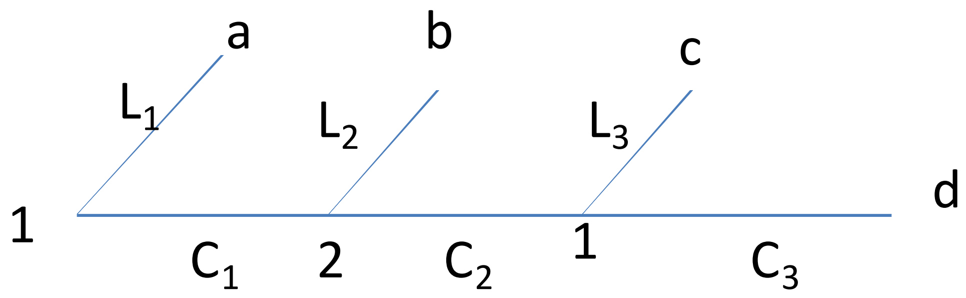

	L_2	C_2
L_1L_3	a	a
L_1C_3	a	a
C_1L_3	b	c
C_1C_3	b	d

⊠ Strategies L_1L_3 and L_1C_3 are outcome-equivalent: the choice between L_3 or C_3 in history (C_1, C_2) is irrelevant, since this history is not reached anyway.

⊠ Strategies s_i and t_i of pl. $i \in N$ are *outcome-equivalent* if — given the strategies of the remaining players — they give rise to the same outcome:

$$\text{for all } s_{-i}: \qquad O(s_i, s_{-i}) = O(t_i, s_{-i}).$$

⊠ So C_1L_3 and C_1C_3 are not outcome-equivalent: the outcomes are different if 2 chooses C_2. But if outcomes c and d were to give the same payoff, the difference in the outcomes is irrelevant. We call the strategies payoff-equivalent.

⊠ Strategies s_i and t_i of pl. $i \in N$ are *payoff-equivalent* if — given the strategies of the remaining players — each player is indifferent between the outcomes generated by s_i and t_i:

$$\begin{aligned} \text{for all } s_{-i}, \text{ for all } j &\in N: \qquad O(s_i, s_{-i}) \sim_j O(t_i, s_{-i}) \\ (\text{or } u_j(O(s_i, s_{-i})) &= u_j(O(t_i, s_{-i}))). \end{aligned}$$

☒ **Def. 95.1:** the *reduced strategic form* of Γ is obtained from its strategic form by replacing each set of equivalent strategies by a single "representative" one.

If c and d give different payoffs, the reduced strategic form is

$$\begin{array}{ccc} & L_2 & C_2 \\ L_1 & a & a \\ C_1 L_3 & b & c \\ C_1 C_3 & b & d \end{array}$$

If c and d give identical payoffs, the reduced strategic form is

$$\begin{array}{ccc} & L_2 & C_2 \\ L_1 & a & a \\ C_1 & b & c \end{array}$$

(Notice: "the" strategic form is somewhat ambiguous: e.g. the names of the strategies are arbitrary...)

Subgame perfection

Example 3.

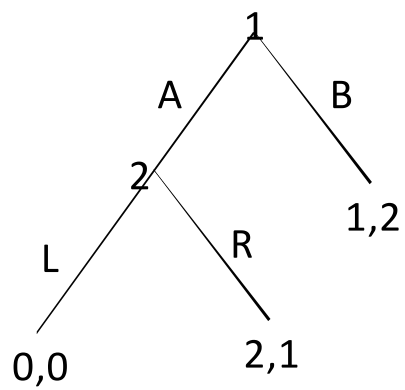

Strategic form:

	L	R
A	$0,0$	$2,1$
B	$1,2$	$1,2$

Nash equilibria: (B,L) and (A,R). But if pl 2 is called upon to play, would 2 choose R?

Subgame perfection requires that players play equilibria in each subgame. In games with perfect information, a subgame is simply a game that starts from an arbitrary history:

Def. 97.1: A *subgame* of $\Gamma = \langle N, H, P, (\succsim_i)_{i \in N} \rangle$ after history $h \in H$ is the game $\Gamma(h) = \langle N, H_{|h}, P_{|h}, (\succsim_{i|h})_{i \in N} \rangle$ with

☒ $H_{|h} = \{h' \mid (h, h') \in H\}$ (histories following h)

☒ for all $h' \in H_{|h}$: $\quad P_{|h}(h') = P(h, h')$ (player assignments the same)

☒ for all terminal histories $h', h'' \in H_{|h}$: $\quad h' \succsim_{i|h} h'' \Leftrightarrow (h, h') \succsim_i (h, h'')$ (preferences the same)

Def. 97.2: A strategy combination $s^* = (s_i^*)_{i \in N}$ is a *subgame perfect equilibrium* of $\Gamma = \langle N, H, P, (\succsim_i)_{i \in N} \rangle$ if it induces a Nash equilibrium in each subgame: for each $h \in H \setminus Z$, $s^*_{|h}$ (the strategy profile s^* restricted to the subgame $\Gamma(h)$) is a Nash equilibrium of $\Gamma(h)$.

Lemma 98.2 (One-deviation property) Let $\Gamma = \langle N, H, P, (\succsim_i)_{i \in N} \rangle$ have finite horizon. The following two claims are equivalent:

(1) s^* is a subgame perfect equilibrium.

(2) in each subgame, the first player to move cannot profitably deviate by changing only the initial action.

Proof sketch: Clearly (1)$\Rightarrow$(2). Conversely, assume (2) holds. If s^* is not subgame perfect, choose a subgame and a profitable deviation by some pl. i with as few deviations as possible. Consider the longest history h after which i deviates. Then $\Gamma(h)$ is a subgame with a profitable deviation in the initial node.

Prop. 99.2 (backwards induction) Every finite extensive form game with perfect information has a subgame perfect equilibrium (in pure strategies).

Proof sketch: Induction on length of subgames. In each subgame, given equilibrium behavior after the initial move, the one-deviation property assures that we only need to check that the first player chooses optimally.

Q: why finite games? Isn't a finite horizon enough?

That's a pretty nice result: equilibrium existence in pure strategies! No need for different types of randomization (mixed/behavioral).

Example 91.1 Dividing 2 indivisible objects. Pl. 1 proposes, pl. 2 accepts or rejects.

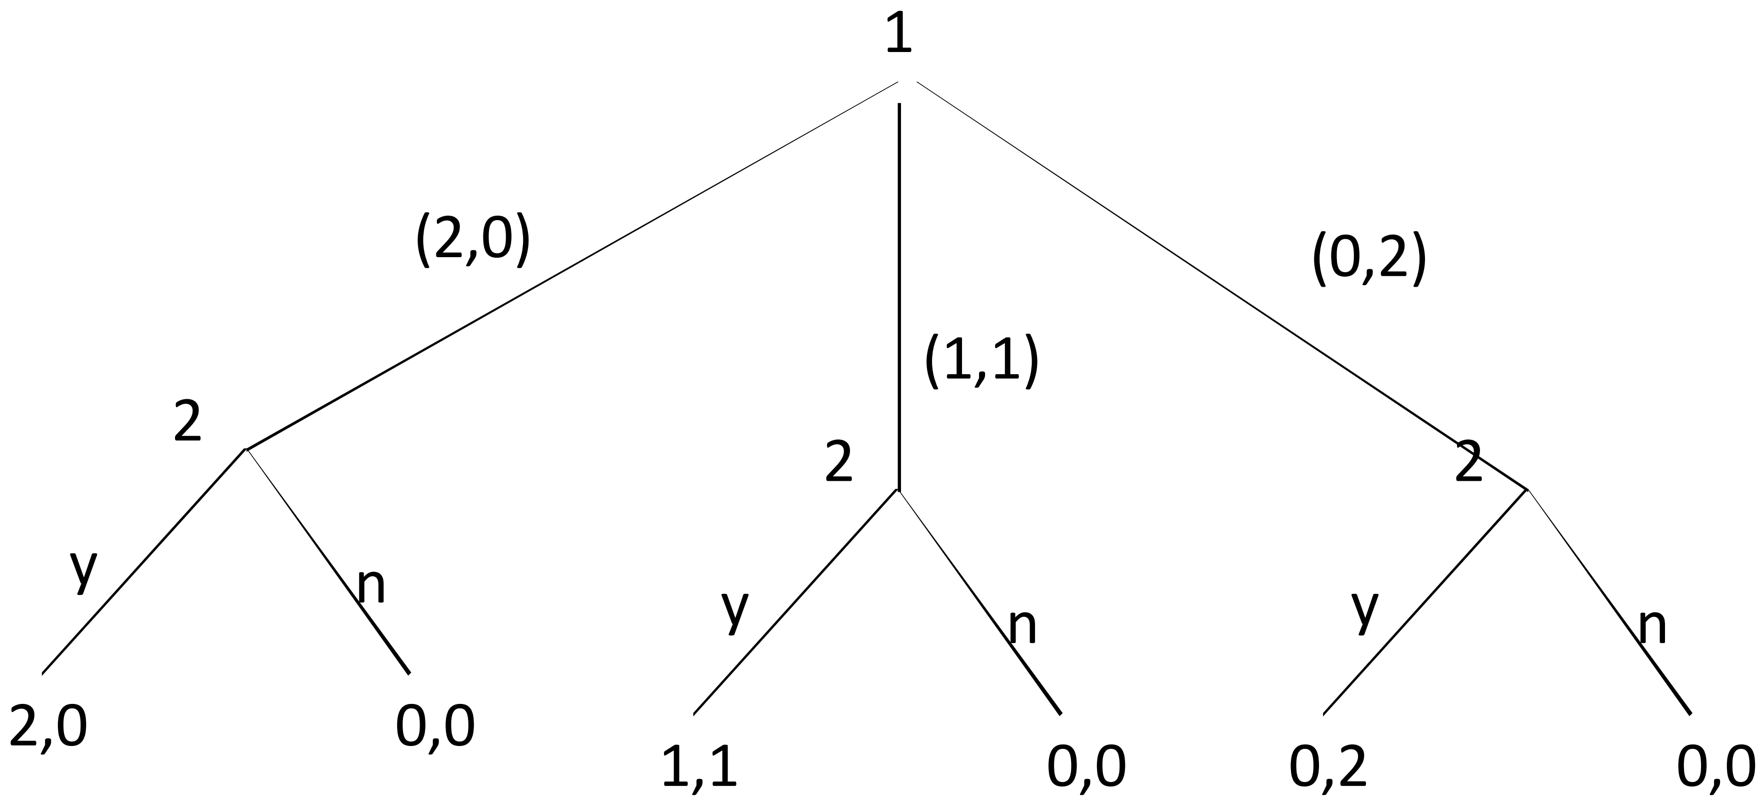

Subgame perfect equilibria?

Section 6.3: two extensions

Extension 1: observable chance moves

Example 4: toss a coin to decide which of two extensive form games with perfect information will be played.

One-deviation property and backward-induction result still hold.

Extension 2: simultaneous moves

Example 5: Player 1 chooses to stay home and end the game with payoffs $(2,2)$ or to play a battle of the sexes game:

	$Bach$	$Stravinsky$
$Bach$	$3,1$	$0,0$
$Stravinsky$	$0,0$	$1,3$

One-deviation property still holds; backward-induction result doesn't (pure equilibria need not exist — Matching Pennies)

Backwards induction and iterated elimination of weakly dominated strategies (IEWDS)

1. If $\Gamma = \langle N, H, P, (\succsim_i)_{i \in N} \rangle$ is such that no player is indifferent between any of the terminal nodes (e.g. all payoffs distinct), then there is a unique subgame perfect equilibrium. There is a process of iterated elimination of weakly dominated strategies (IEWDS) under which this subgame perfect equilibrium survives. (Idea: follow backwards induction).

2. Other orders of elimination may eliminate all subgame perfect equilibria.

3. If there are ties in the payoffs, IEWDS may eliminate all subgame perfect equilibria or result in outcomes that aren't even equilibria (p. 109-110).

Illustration of IEWDS to keep subgame perfect equilibrium

Example 6.

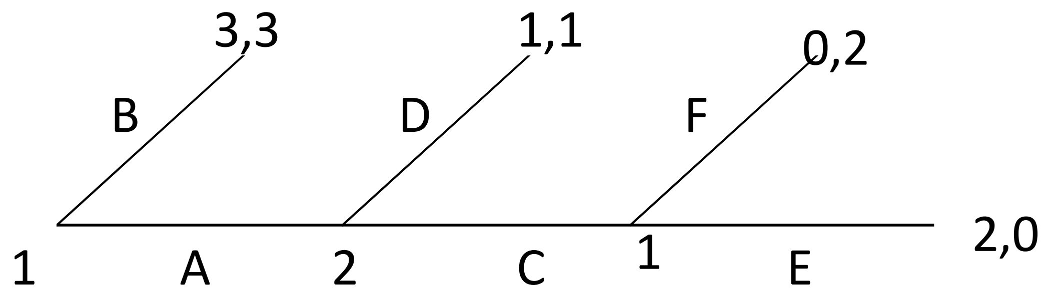

Strategic form:

	C	D
AE	2, 0	1, 1
AF	0, 2	1, 1
BE	3, 3	3, 3
BF	3, 3	3, 3

Backwards induction: AF weakly dominated by AE, then C weakly dominated by D, then AE weakly dominated by BE.

Other order eliminates the subgame perfect equilibrium: AE weakly dominated by BE, then D weakly dominated by C, then AF weakly dominated by BE.

Forward induction

Recall example 5: Player 1 chooses to stay home and end the game with payoffs $(2, 2)$ or to play a battle of the sexes game:

	$Bach$	$Stravinsky$
$Bach$	$3, 1$	$0, 0$
$Stravinsky$	$0, 0$	$1, 3$

Reduced strategic form:

	$Bach$	$Stravinsky$
$Home$	$2, 2$	$2, 2$
$Bach$	$3, 1$	$0, 0$
$Stravinsky$	$0, 0$	$1, 3$

Subgame perfect equilibria?

IEWDS?

Forward induction bases reasoning on what happened *earlier* in the game:

- If pl 2 needs to make a choice, that pl 1 gave up payoff 2.

- That makes sense only if he chooses $Bach$ in the battle-of-the-sexes game.

- Pl 2's best response is also to choose $Bach$.

This type of reasoning applies to a limited class of games, though.