Logic and Games
SS 2009

Prof. Dr. Erich Grädel
Łukasz Kaiser, Tobias Ganzow

Mathematische Grundlagen der Informatik
RWTH Aachen

Contents

4 Basic Concepts of Mathematical Game Theory

Up to now we considered finite or infinite games

- with two players,
- played on finite or infinite graphs,
- with perfect information (the players know the whole game, the history of the play and the actual position),
- with qualitative (win or loss) winning conditions (zero-sum games),
- with ω-regular winning conditions (or Borel winning conditions) specified in a suitable logic or by automata, and
- with asynchronous interaction (turn-based games).

Those games are used for verification or to evaluate logic formulae.

In this section we move to concurrent multi-player games in which players get real-valued *payoffs*. The games will still have perfect information and additionally throughout this chapter we assume that the set of possible plays is *finite*, so there exist only finitely many strategies for each of the players.

4.1 Games in Strategic Form

Definition 4.1. A *game in strategic form* is described by a tuple $\Gamma = (N, (S_i)_{i\in N}, (p_i)_{i\in N})$ where

- $N = \{1, \dots, n\}$ is a finite set of players
- S_i is a set of *strategies* for Player i
- $p_i : S \to \mathbb{R}$ is a *payoff function* for Player i

and $S := S_1 \times \cdots \times S_n$ is the set of *strategy profiles*. Γ is called a *zero-sum game* if $\sum_{i\in N} p_i(s) = 0$ for all $s \in S$.

The number $p_i(s_1, \ldots, s_n)$ is called the *value* or *utility* of the strategy profile $(s_1, \ldots, s_n)$ for Player i. The intuition for zero-sum games is that the game is a closed system.

Many important notions can best be explained by two-player games, but are defined for arbitrary multi-player games.

In the sequel, we will use the following notation: Let Γ be a game. Then $S_{-i} := S_1 \times \cdots \times S_{i-1} \times S_{i+1} \times \cdots \times S_n$ is the set of all strategy profiles for the players except i. For $s \in S_i$ and $s_{-i} \in S_{-i}$, (s, s_{-i}) is the strategy profile where Player i chooses the strategy s and the other players choose s_{-i}.

Definition 4.2. Let $s, s' \in S_i$. Then s *dominates* s' if

- for all $s_{-i} \in S_{-i}$ we have $p_i(s, s_{-i}) \geq p_i(s', s_{-i})$, and
- there exists $s_{-i} \in S_{-i}$ such that $p_i(s, s_{-i}) > p_i(s', s_{-i})$.

A strategy s is *dominant* if it dominates some other strategy of the player.

Definition 4.3. An *equilibrium in dominant strategies* is a strategy profile $(s_1, \ldots, s_n) \in S$ such that all s_i are dominant strategies.

Definition 4.4. A strategy $s \in S_i$ is a *best response* to $s_{-i} \in S_{-i}$ if $p_i(s, s_{-i}) \geq p_i(s', s_{-i})$ for all $s' \in S_i$.

Remark 4.5. A dominant strategy is a best response for all strategy profiles of the other players.

Example 4.6.

- Prisoner's Dilemma:

  Two suspects are arrested, but there is insufficient evidence for a conviction. Both prisoners are questioned separately, and are offered the same deal: if one testifies for the prosecution against the other and the other remains silent, the betrayer goes free and the silent accomplice receives the full 10-year sentence. If both stay silent, both prisoners are sentenced to only one year in jail for a minor charge. If both betray each other, each receives a five-year sentence. So this dilemma poses the question: How should the prisoners act?

	stay silent	betray
stay silent	$(-1,-1)$	$(-10,0)$
betray	$(0,-10)$	$(-5,-5)$

An entry (a,b) at position i,j of the matrix means that if profile (i,j) is chosen, Player 1 (who chooses the rows) receives payoff a and Player 2 (who chooses the columns) receives payoff b.

Betraying is a dominant strategy for every player, call this strategy b. Therefore, (b,b) is an equilibrium in dominant strategies. Problem: The payoff $(-5,-5)$ of the dominant equilibrium is not optimal.

The Prisoner's Dilemma is an important metaphor for many decision situations, and there exists extensive literature concerned with the problem. Especially interesting is the situation, where the Prisoner's Dilemma is played infinitely often.

- Battle of the sexes:

	meat	fish
red wine	$(2,1)$	$(0,0)$
white wine	$(0,0)$	$(1,2)$

There are no dominant strategies, and thus there is no dominant equilibrium. The pairs (red wine, meat) and (white wine, fish) are distinguished since every player plays with a best response against the strategy of the other player: No player would change his strategy unilaterally.

Definition 4.7. A strategy profile $s = (s_1, \ldots, s_n) \in S$ is a *Nash equilibrium* in Γ if

$$p_i(\underbrace{s_i, s_{-i}}_{s}) \geq p_i(s', s_{-i})$$

holds for all $i \in N$ and all strategies $s_i' \in S_i$, i.e., for every Player i, s_i is a best response for s_{-i}.

Is there a Nash equilibrium in every game? Yes, but not necessarily in pure strategies!

Example 4.8. Rock, paper, scissors:

	rock	scissors	paper
rock	$(0,0)$	$(1,-1)$	$(-1,1)$
scissors	$(-1,1)$	$(0,0)$	$(1,-1)$
paper	$(1,-1)$	$(-1,1)$	$(0,0)$

There are no dominant strategies and no Nash equilibria: For every pair (f,g) of strategies one of the players can change to a better strategy. Note that this game is a zero-sum game. But there is a reasonable strategy to win this game: Randomly pick one of the three actions with equal probability.

This observation leads us to the notion of mixed strategies, where the players are allowed to randomise over strategies.

Definition 4.9. A *mixed strategy* of Player i in Γ is a probability distribution $\mu_i : S_i \to [0,1]$ on S_i where $\sum_{s\in S_i} \mu(s) = 1$.
$\Delta(S_i)$ denotes the set of probability distributions on S_i. $\Delta(S) := \Delta(S_1) \times \cdots \times \Delta(S_n)$ is the set of all strategy profiles in mixed strategies. The expected payoff is $\widehat{p}_i : \Delta(S) \to \mathbb{R}$,

$$\widehat{p}_i(\mu_1, \ldots, \mu_n) = \sum_{(s_1,\ldots,s_n)\in S} \left(\prod_{j\in N} \mu_j(s_j) \right) \cdot p_i(s_1, \ldots, s_n)$$

For every game $\Gamma = (N, (S_i)_{i\in N}, (p_i)_{i\in N})$ we define the *mixed expansion* $\widehat{\Gamma} = (N, (\Delta(S_i))_{i\in N}, (\widehat{p}_i)_{i\in N})$.

Definition 4.10. A *Nash equilibrium of* Γ *in mixed strategies* is a Nash equilibrium in $\widehat{\Gamma}$, i.e. a Nash equilibrium in Γ in mixed strategies is a mixed strategy profile $\overline{\mu} = (\mu_1, \ldots, \mu_n) \in \Delta(S)$ such that, for every player i and every $\mu_i' \in \Delta(S)$, $\widehat{p}_i(\mu_i, \mu_{-i}) \geq \widehat{p}_i(\mu_i', \mu_{-i})$.

Theorem 4.11 (Nash)**.** Every finite game Γ in strategic form has at least one Nash equilibrium in mixed strategies.

To prove this theorem, we will use a well-known fixed-point theorem.

Theorem 4.12 (Brouwer's fixed-point theorem)**.** Let $X \subseteq \mathbb{R}^n$ be compact (i.e., closed and bounded) and convex. Then every continuous function $f : X \to X$ has a fixed point.

Proof (of Theorem 4.11). Let $\Gamma = (N, (S_i)_{i\in N}, (p_i)_{i\in N})$. Every mixed strategy of Player i is a tuple $\mu_i = (\mu_{i,s})_{s\in S_i} \in [0,1]^{|S_i|}$ such that $\sum_{s\in S_i} \mu_{i,s} = 1$. Thus, $\Delta(S_i) \subseteq [0,1]^{|S_i|}$ is a compact and convex set, and the same applies to $\Delta(S) = \Delta(S_1) \times \cdots \times \Delta(S_n)$ for $N = \{1,\ldots,n\}$. For every $i \in N$, every pure strategy $s \in S_i$ and every mixed strategy profile $\overline{\mu} \in \Delta(S)$ let

$$g_{i,s}(\overline{\mu}) := \max\left(\widehat{p}_i(s, \overline{\mu}_{-i}) - \widehat{p}_i(\overline{\mu}), 0\right)$$

be the gain of Player i if he unilaterally changes from the mixed profile $\overline{\mu}$ to the pure strategy s (only if this is reasonable).

Note that if $g_{i,s}(\overline{\mu}) = 0$ for all i and all $s \in S_i$, then $\overline{\mu}$ is a Nash equilibrium. We define the function

$$\begin{aligned} f : \Delta(S) &\to \Delta(S) \\ \overline{\mu} &\mapsto f(\overline{\mu}) = (\nu_1, \ldots, \nu_n) \end{aligned}$$

where $\nu_i : S_i \to [0,1]$ is a mixed strategy defined by

$$\nu_{i,s} = \frac{\mu_{i,s} + g_{i,s}(\overline{\mu})}{1 + \sum_{s\in S_i} g_{i,s}(\overline{\mu})}.$$

For every Player i and all $s \in S_i$, $\overline{\mu} \mapsto \nu_{i,s}$ is continuous since $\widehat{p}_i$ is continuous and thus $g_{i,s}$, too. $f(\overline{\mu}) = (\nu_1, \ldots, \nu_n)$ is in $\Delta(S)$: Every $\nu_i = (\nu_{i,s})_{s\in S_i}$ is in $\Delta(S_i)$ since

$$\sum_{s\in S_i} \nu_{i,s} = \frac{\sum_{s\in S_i} \mu_{i,s} + \sum_{s\in S_i} g_{i,s}(\overline{\mu})}{1 + \sum_{s\in S_i} g_{i,s}(\overline{\mu})} = \frac{1 + \sum_{s\in S_i} g_{i,s}(\overline{\mu})}{1 + \sum_{s\in S_i} g_{i,s}(\overline{\mu})} = 1.$$

By the Brouwer fixed point theorem f has a fixed point. Thus, there is a $\overline{\mu} \in \Delta(S)$ such that

$$\mu_{i,s} = \frac{\mu_{i,s} + g_{i,s}(\overline{\mu})}{1 + \sum_{s \in S_i} g_{i,s}(\overline{\mu})}$$

for all i and all s.

Case 1: There is a Player i such that $\sum_{s \in S_i} g_{i,s}(\overline{\mu}) > 0$.
Multiplying both sides of the fraction above by the denominator, we get $\mu_{i,s} \cdot \sum_{s \in S_i} g_{i,s}(\mu) = g_{i,s}(\overline{\mu})$. This implies $\mu_{i,s} = 0 \;\Leftrightarrow\; g_{i,s}(\overline{\mu}) = 0$, and thus $g_{i,s}(\overline{\mu}) > 0$ for all $s \in S_i$ where $\mu_{i,s} > 0$.

But this leads to a contradiction: $g_{i,s}(\overline{\mu}) > 0$ means that it is profitable for Player i to switch from (μ_i, μ_{-i}) to (s, μ_{-i}). This cannot be true for all s where $\mu_{i,s} > 0$ since the payoff for (μ_i, μ_{-i}) is the mean of the payoffs (s, μ_{-i}) with arbitrary $\mu_{i,s}$. However, the mean cannot be smaller than all components:

$$\begin{aligned}
\widehat{p}_i(\mu_i, \mu_{-i}) &= \sum_{s \in S_i} \mu_{i,s} \cdot \widehat{p}_i(s, \mu_{-i}) \\
&= \sum_{\substack{s \in S_i \\ \mu_{i,s} > 0}} \mu_{i,s} \cdot \widehat{p}_i(s, \mu_{-i}) \\
&> \sum_{\substack{s \in S_i \\ \mu_{i,s} > 0}} \mu_{i,s} \cdot \widehat{p}_i(\mu_i, \mu_{-i}) \\
&= \widehat{p}_i(\mu_i, \mu_{-i})
\end{aligned}$$

which is a contradiction.

Case 2*:* $g_{i,s}(\overline{\mu}) = 0$ for all i and all $s \in S_i$, but this already means that $\overline{\mu}$ is a Nash equilibrium as stated before. Q.E.D.

The *support* of a mixed strategy $\mu_i \in \Delta(S_i)$ is $\mathrm{supp}(\mu_i) = \{s \in S_i : \mu_i(s) > 0\}$.

Theorem 4.13. Let $\mu^* = (\mu_1, \dots, \mu_n)$ be a Nash equilibrium in mixed strategies of a game Γ. Then for every Player i and every pure strategy $s, s' \in \mathrm{supp}(\mu_i)$

$$\widehat{p}_i(s, \mu_{-i}) = \widehat{p}_i(s', \mu_{-i}).$$

Proof. Assume $\widehat{p}_i(s, \mu_{-i}) > \widehat{p}_i(s', \mu_{-i})$. Then Player i could achieve a higher payoff against μ_{-i} if she played s instead of s': Define $\tilde{\mu}_i \in \Delta(S_i)$ as follows:

- $\tilde{\mu}_i(s) = \mu_i(s) + \mu_i(s')$,
- $\tilde{\mu}_i(s') = 0$,
- $\tilde{\mu}_i(t) = \mu_i(t)$ for all $t \in S_i - \{s, s'\}$.

Then

$$\begin{aligned}\widehat{p}_i(\tilde{\mu}_i, \mu_{-i}) &= \widehat{p}_i(\mu_i, \mu_{-i}) + \underbrace{\mu_i(s')}_{>0} \cdot \underbrace{\left(\widehat{p}_i(s, \mu_{-i}) - \widehat{p}_i(s', \mu_{-i})\right)}_{>0} \\ &> \widehat{p}_i(\mu_i, \mu_{-i})\end{aligned}$$

which contradicts the fact that μ is a Nash equilibrium. Q.E.D.

We want to apply Nash's theorem to two-person games. First, we note that in every game $\Gamma = (\{0,1\}, (S_0, S_1), (p_0, p_1))$

$$\max_{f \in \Delta(S_0)} \min_{g \in \Delta(S_1)} p_0(f, g) \leq \min_{g \in \Delta(S_1)} \max_{f \in \Delta(S_0)} p_0(f, g).$$

The maximal payoff which one player can enforce cannot exceed the minimal payoff the other player has to cede. This is a special case of the general observation that for every function $f : X \times Y \to \mathbb{R}$

$$\sup_x \inf_y h(x, y) \leq \inf_y \sup_x h(x, y).$$

(For all x', y: $h(x', y) \leq \sup_x h(x, y)$. Thus $\inf_y h(x', y) \leq \inf_y \sup_x h(x, y)$ and $\sup_x \inf_y h(x, y) \leq \inf_y \sup_x h(x, y)$.)

Remark 4.14. Another well-known special case is

$$\exists x \, \forall y \, Rxy \models \forall y \, \exists x \, Rxy.$$

Example 4.15. Consider the following two-player "traveller" game $\Gamma = (\{1,2\}, (S_1, S_2), (p_1, p_2))$ with $S_1 = S_2 = \{2, \ldots, 100\}$ and

$$p_1(x,y) = \begin{cases} x+2 & \text{if } x < y, \\ y-2 & \text{if } y < x, \\ x & \text{if } x = y, \end{cases}$$

$$p_2(x,y) = \begin{cases} x-2 & \text{if } x < y, \\ y+2 & \text{if } y < x, \\ y & \text{if } x = y. \end{cases}$$

Let's play this game! These are the results from the lecture in 2009:

$$2, 49, 49, 50, 51, 92, 97, 98, 99, 99, 100.$$

But what are the Nash equilibria? Observe that the only pure-strategy Nash equilibrium is $(2,2)$ since for each (i,j) with $i \neq j$ the player that has chosen the greater number, say i, can do better by switching to $j-1$, and also, for every (i,i) with $i > 2$ each player can do better by playing $i-1$ (and getting the payoff $i+1$ then). But would you really expect such a good payoff playing 2? Look at how others played: 97 seems to be much better against what people do in most cases!

Theorem 4.16 (v. Neumann, Morgenstern)**.**
Let $\Gamma = (\{0,1\}, (S_0, S_1), (p, -p))$ be a two-person zero-sum game. For every Nash equilibrium (f^*, g^*) in mixed strategies

$$\max_{f \in \Delta(S_0)} \min_{g \in \Delta(S_1)} p(f,g) = p(f^*, g^*) = \min_{g \in \Delta(S_1)} \max_{f \in \Delta(S_0)} p(f,g).$$

In particular, all Nash equilibria have the same payoff which is called the *value* of the game. Furthermore, both players have optimal strategies to realise this value.

Proof. Since (f^*, g^*) is a Nash equilibrium, for all $f \in \Delta(S_0)$, $g \in \Delta(S_1)$

$$p(f^*, g) \geq p(f^*, g^*) \geq p(f, g^*).$$

Thus

$$\min_{g \in \Delta(S_1)} p(f^*, g) = p(f^*, g^*) = \max_{f \in \Delta(S_1)} p(f, g^*).$$

So

$$\max_{f \in \Delta(S_0)} \min_{g \in \Delta(S_1)} p(f,g) \geq p(f^*, g^*) \geq \min_{g \in \Delta(S_1)} \max_{f \in \Delta(S_0)} p(f,g)$$

and

$$\max_{f \in \Delta(S_0)} \min_{g \in \Delta(S_1)} p(f,g) \leq \min_{g \in \Delta(S_1)} \max_{f \in \Delta(S_0)} p(f,g)$$

imply the claim. Q.E.D.

4.2 Iterated Elimination of Dominated Strategies

Besides Nash equilibria, the iterated elimination of dominated strategies is a promising solution concept for strategic games which is inspired by the following ideas. Assuming that each player behaves rational in the sense that he will not play a strategy that is dominated by another one, dominated strategies may be eliminated. Assuming further that it is common knowledge among the players that each player behaves rational, and thus discards some of her strategies, such elimination steps may be iterated as it is possible that some other strategies become dominated due to the elimination of previously dominated strategies. Iterating these elimination steps eventually yields a fixed point where no strategies are dominated.

Example 4.17.

$$\begin{array}{c|c}
\begin{array}{ccc} & L & R \\ T & (1,0,1) & \boxed{(1,1,0)} \\ B & (1,1,1) & (0,0,1) \end{array} & \begin{array}{cc} L & R \\ (1,0,1) & (0,1,0) \\ (1,1,1) & (1,0,0) \end{array} \\
X & Y
\end{array}$$

Player 1 picks rows, Player 2 picks columns, and Player 3 picks matrices.

- No row dominates the other (for Player 1);
- no column dominates the other (for Player 2);
- matrix X dominates matrix Y (for Player 3).

Thus, matrix Y is eliminated.

- In the remaining game, the upper row dominates the lower one (for Player 1).

Thus, the lower row is eliminated.

- Of the remaining two possibilities, Player 2 picks the better one.

The only remaining profile is (T, R, X).

There are different variants of strategy elimination that have to be considered:

- dominance by *pure* or *mixed* strategies;
- (*weak*) dominance or *strict* dominance;
- dominance by strategies in the *local* subgame or by strategies in the *global* game.

The possible combinations of these parameters give rise to eight different operators for strategy elimination that will be defined more formally in the following.

Let $\Gamma = (N, (S_i)_{i\in N}, (p_i)_{i\in N})$ such that S_i is finite for every Player i. A subgame is defined by $T = (T_1, \ldots, T_n)$ with $T_i \subseteq S_i$ for all i. Let $\mu_i \in \Delta(S_i)$, and $s_i \in S_i$. We define two notions of dominance:

(1) Dominance with respect to T:
$\mu_i >_T s_i$ if and only if
  - $p_i(\mu_i, t_{-i}) \geq p_i(s_i, t_{-i})$ for all $t_{-i} \in T_{-i}$
  - $p_i(\mu_i, t_{-i}) > p_i(s_i, t_{-i})$ for some $t_{-i} \in T_{-i}$.

(2) Strict dominance with respect to T:
$\mu_i \gg_T s_i$ if and only if $p_i(\mu_i, t_{-i}) > p_i(s_i, t_{-i})$ for all $t_{-i} \in T_{-i}$.

We obtain the following operators on $T = (T_1, \ldots, T_n)$, $T_i \subseteq S_i$, that are defined component-wise:

$$
\begin{aligned}
\mathrm{ML}(T)_i &:= \{t_i \in T_i : \neg\exists \mu_i \in \Delta(T_i)\ \mu_i >_T t_i\},\\
\mathrm{MG}(T)_i &:= \{t_i \in T_i : \neg\exists \mu_i \in \Delta(S_i)\ \mu_i >_T t_i\},\\
\mathrm{PL}(T)_i &:= \{t_i \in T_i : \neg\exists t_i' \in T_i\ t_i' >_T t_i\}, \text{and}\\
\mathrm{PG}(T)_i &:= \{t_i \in T_i : \neg\exists s_i \in S_i\ s_i >_T t_i\}.
\end{aligned}
$$

MLS, MGS, PLS, PGS are defined analogously with $\gg_T$ instead of $>_T$. For all T we have the following obvious inclusions:

- Every M-operator eliminates more strategies than the corresponding P-operator.
- Every operator considering (weak) dominance eliminates more strategies than the corresponding operator considering strict dominance.
- With dominance in global games more strategies are eliminated than with dominance in local games.

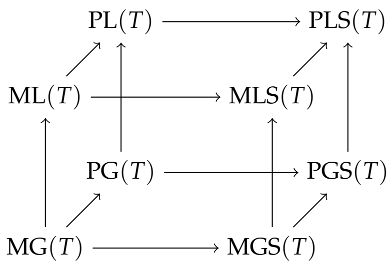

Figure 4.1. Inclusions between the eight strategy elimination operators

Each of these operators is deflationary, i.e. $F(T) \subseteq T$ for every T and every operator F. We iterate an operator beginning with $T = S$, i.e. $F^0 := S$ and $F^{\alpha+1} := F(F^\alpha)$. Obviously, $F^0 \supseteq F^1 \supseteq \cdots \supseteq F^\alpha \supseteq F^{\alpha+1}$. Since S is finite, we will reach a fixed point F^α such that $F^\alpha = F^{\alpha+1} =: F^\infty$. We expect that for the eight fixed points MG^∞, ML^∞, etc. the same inclusions hold as for the operators $\mathrm{MG}(T)$, $\mathrm{ML}(T)$, etc. But this is not the case: For the following game $\Gamma = (\{0,1\}, (S_0, S_1), (p_0, p_1))$ we have $\mathrm{ML}^\infty \not\subseteq \mathrm{PL}^\infty$.

	X	Y	Z
A	$(2,1)$	$(0,1)$	$(1,0)$
B	$(0,1)$	$(2,1)$	$(1,0)$
C	$(1,1)$	$(1,0)$	$(0,0)$
D	$(1,0)$	$(0,1)$	$(0,0)$

We have:

- Z is dominated by X and Y.
- D is dominated by A.
- C is dominated by $\frac{1}{2}A + \frac{1}{2}B$.

Thus:

$$\mathrm{ML}(S) = \mathrm{ML}^1 = (\{A, B\}, \{X, Y\}) \subset \mathrm{PL}(S) = \mathrm{PL}^1 = (\{A, B, C\}, \{X, Y\}).$$

$\mathrm{ML}(\mathrm{ML}^1) = \mathrm{ML}^1$ since in the following game there are no dominated strategies:

	X	Y
A	$(2, 1)$	$(0, 1)$
B	$(0, 1)$	$(2, 1)$

$\mathrm{PL}(\mathrm{PL}^1) = (\{A, B, C\}, \{X\}) = \mathrm{PL}^2 \subsetneq \mathrm{PL}^1$ since Y is dominated by X (here we need the presence of C). Since B and C are now dominated by A, we have $\mathrm{PL}^3 = (\{A\}, \{X\}) = \mathrm{PL}^\infty$. Thus, $\mathrm{PL}^\infty \subsetneq \mathrm{ML}^\infty$ although ML is the stronger operator.

We are interested in the inclusions of the fixed points of the different operators. But we only know the inclusions for the operators. So the question arises under which assumptions can we prove, for two deflationary operators F and G on S, the following claim:

If $F(T) \subseteq G(T)$ for all T, then $F^\infty \subseteq G^\infty$?

The obvious proof strategy is induction over α: We have $F^0 = G^0 = S$, and if $F^\alpha \subseteq G^\alpha$, then

$$F^{\alpha+1} = F(F^\alpha) \subseteq G(F^\alpha)$$
$$F(G^\alpha) \subseteq G(G^\alpha) = G^{\alpha+1}$$

If we can show one of the inclusions $F(F^\alpha) \subseteq F(G^\alpha)$ or $G(F^\alpha) \subseteq G(G^\alpha)$, then we have proven the claim. These inclusions hold if the operators are monotone: $H : S \to S$ is monotone if $T \subseteq T'$ implies $H(T) \subseteq H(T')$. Thus, we have shown:

Lemma 4.18. Let $F, G : \mathcal{P}(S) \to \mathcal{P}(S)$ be two deflationary operators such that $F(T) \subseteq G(T)$ for all $T \subseteq S$. If either F or G is monotone, then $F^\infty \subseteq G^\infty$.

Corollary 4.19. PL and ML are not monotone on every game.

Which operators are monotone? Obviously, MGS and PGS are monotone: If $\mu_i \gg_T s_i$ and $T' \subseteq T$, then also $\mu_i \gg_{T'} s_i$. Let $T' \subseteq T$ and $s_i \in \text{PGS}(T')_i$. Thus, there is no $\mu_i \in S_i$ such that $\mu_i \gg_{T'} s_i$, and there is also no $\mu_i \in S_i$ such that $\mu_i \gg_T s_i$ and we have $s_i \in \text{PGS}(T)_i$. The reasoning for MGS is analogous if we replace S_i by $\Delta(S_i)$.

MLS and PLS are not monotone. Consider the following simple game:

	X
A	$(1,0)$
B	$(0,0)$

$$\text{MLS}(\{A,B\},\{X\}) = \text{PLS}(\{A,B\},\{X\}) = (\{A\},\{X\}) \text{ and}$$
$$\text{MLS}(\{B\},\{X\}) = \text{PLS}(\{B\},\{X\}) = (\{B\},\{X\}),$$

but $(\{B\},\{X\}) \not\subseteq (\{A\},\{X\})$.

Thus, none of the local operators (those which only consider dominant strategies in the current subgame) is monotone. We will see that also MG and PG are not monotone in general. The monotonicity of the global operators MGS and PGS will allow us to prove the expected inclusions $\text{ML}^\infty \subseteq \text{MLS}^\infty \subseteq \text{PLS}^\infty$ and $\text{PL}^\infty \subseteq \text{PLS}^\infty$ between the local operators. To this end, we will show that the fixed points of the local and corresponding global operators coincide (although the operators are different).

Lemma 4.20. $\text{MGS}^\infty = \text{MLS}^\infty$ and $\text{PGS}^\infty = \text{PLS}^\infty$.

Proof. We will only prove $\text{PGS}^\infty = \text{PLS}^\infty$. Since $\text{PGS}(T) \subseteq \text{PLS}(T)$ for all T and PGS is monotone, we have $\text{PGS}^\infty \subseteq \text{PLS}^\infty$. Now we will prove by induction that $\mathrm{PLS}^\alpha \subseteq \mathrm{PGS}^\alpha$ for all α. Only the induction step $\alpha \mapsto \alpha + 1$ has to be considered: Let $s_i \in \mathrm{PLS}_i^{\alpha+1}$. Therefore, $s_i \in \mathrm{PLS}_i^\alpha$ and there is no $s_i' \in \mathrm{PLS}_i^\alpha$ such that $s_i' \gg_{\mathrm{PLS}^\alpha} s_i$. Assume $s_i \notin \mathrm{PGS}_i^{\alpha+1}$, i.e.

$$A = \{s_i' \in S_i : s_i' \gg_{\mathrm{PGS}^\alpha} s_i\} \neq \emptyset$$

(note: By induction hypothesis $\mathrm{PGS}^\alpha = \mathrm{PLS}^\alpha$). Pick an $s_i^* \in A$ which is maximal with respect to $\gg_{\mathrm{PLS}^\alpha}$. Claim: $s_i^* \in \mathrm{PLS}^\alpha$. Otherwise, there exists a $\beta \leq \alpha$ and an $s_{i'} \in S_i$ with $s_i' \gg_{\mathrm{PLS}^\beta} s_{i^*}$. Since $\mathrm{PLS}^\beta \supseteq \mathrm{PLS}^\alpha$, it follows that $s_i' \gg_{\mathrm{PLS}^\alpha} s_i^* \gg_{\mathrm{PLS}^\alpha} s_i$. Therefore, $s_i' \in A$ and s_i^* is not maximal with respect to $\gg_{\mathrm{PLS}^\alpha}$ in A. Contradiction.

But if $s_i^* \in \mathrm{PLS}^\alpha$ and $s_i^* \gg_{\mathrm{PLS}^\alpha} s_i$, then $s_i \notin \mathrm{PLS}^{\alpha+1}$ which again constitutes a contradiction.

The reasoning for MGS^∞ and MLS^∞ is analogous. Q.E.D.

Corollary 4.21. $\mathrm{MLS}^\infty \subseteq \mathrm{PLS}^\infty$.

Lemma 4.22. $\mathrm{MG}^\infty = \mathrm{ML}^\infty$ and $\mathrm{PG}^\infty = \mathrm{PL}^\infty$.

Proof. We will only prove $\mathrm{PG}^\infty = \mathrm{PL}^\infty$ by proving $\mathrm{PG}^\alpha = \mathrm{PL}^\alpha$ for all α by induction. Let $\mathrm{PG}^\alpha = \mathrm{PL}^\alpha$ and $s_i \in \mathrm{PG}_i^{\alpha+1}$. Then $s_i \in \mathrm{PG}_i^\alpha = \mathrm{PL}_i^\alpha$ and hence there is no $s_i' \in S_i$ such that $s_i' >_{\mathrm{PG}^\alpha} s_i$. Thus, there is no $s_i' \in \mathrm{PL}_i^\alpha$ such that $s_i' >_{\mathrm{PL}^\alpha} s_i$ and $s_i \in \mathrm{PL}^{\alpha+1}$. So, $\mathrm{PG}^{\alpha+1} \subseteq \mathrm{PL}^{\alpha+1}$.

Now, let $s_i \in \mathrm{PL}_i^{\alpha+1}$. Again we have $s_i \in \mathrm{PL}_i^\alpha = \mathrm{PG}_i^\alpha$. Assume $s_i \notin \mathrm{PG}_i^{\alpha+1}$. Then

$$A = \{s_i' \in S_i : s_i' >_{\mathrm{PL}^\alpha} s_i\} \neq \emptyset.$$

For every $\beta \leq \alpha$ let $A^\beta = A \cap \mathrm{PL}_i^\beta$. Pick the maximal β such that $A^\beta \neq \emptyset$ and a $s_i^* \in A^\beta$ which is maximal with respect to $>_{\mathrm{PL}^\beta}$.

Claim: $\beta = \alpha$. Otherwise, $s_i \notin \mathrm{PL}_i^{\beta+1}$. Then there exists an $s_i' \in \mathrm{PL}_i^\beta$ with $s_i' >_{\mathrm{PL}^\beta} s_i^*$. Since $\mathrm{PL}^\beta \supseteq \mathrm{PL}^\alpha$ and $s_i^* >_{\mathrm{PL}^\alpha} s_i$, we have $s_i' >_{\mathrm{PL}^\alpha} s_i$, i.e. $s_i' \in A^\beta$ which contradicts the choice of s_i^*. Therefore, $s_i^* \in \mathrm{PL}_i^\alpha$. Since $s_i^* >_{\mathrm{PL}^\alpha} s_i$, we have $s_i \notin \mathrm{PL}_i^{\alpha+1}$. Contradiction, hence the assumption is wrong, and we have $s_i \in \mathrm{PG}^{\alpha+1}$. Altogether $\mathrm{PG}^\alpha = \mathrm{PL}^\alpha$. Again, the reasoning for $\mathrm{MG}^\infty = \mathrm{ML}^\infty$ is analogous. Q.E.D.

Corollary 4.23. $\mathrm{PL}^\infty \subseteq \mathrm{PLS}^\infty$ and $\mathrm{ML}^\infty \subseteq \mathrm{MLS}^\infty$.

Proof. We have $\mathrm{PL}^\infty = \mathrm{PG}^\infty \subseteq \mathrm{PGS}^\infty = \mathrm{PLS}^\infty$ where the inclusion $\mathrm{PG}^\infty \subseteq \mathrm{PGS}^\infty$ holds because $\mathrm{PG}(T) \subseteq \mathrm{PGS}(T)$ for any T and PGS is monotone. Analogously, we have $\mathrm{ML}^\infty = \mathrm{MG}^\infty \subseteq \mathrm{MGS}^\infty = \mathrm{MLS}^\infty$. Q.E.D.

This implies that MG and PG cannot be monotone. Otherwise, we would have $\mathrm{ML}^\infty = \mathrm{PL}^\infty$. But we know that this is wrong.

4.3 Beliefs and Rationalisability

Let $\Gamma = (N, (S_i)_{i\in N}, (p_i)_{i\in N})$ be a game. A *belief* of Player i is a probability distribution over S_{-i}.

Remark 4.24. A belief is not necessarily a product of independent probability distributions over the individual S_j ($j \neq i$). A player may believe that the other players play correlated.

A strategy $s_i \in S_i$ is called a *best response to a belief* $\gamma \in \Delta(S_{-i})$ if $\widehat{p}_i(s_i, \gamma) \geq \widehat{p}_i(s_i', \gamma)$ for all $s_i' \in S_i$. Conversely, $s_i \in S_i$ is *never a best response* if s_i is not a best response for any $\gamma \in \Delta(S_{-i})$.

Lemma 4.25. For every game $\Gamma = (N, (S_i)_{i\in N}, (p_i)_{i\in N})$ and every $s_i \in S_i$, s_i is never a best response if and only if there exists a mixed strategy $\mu_i \in \Delta(S_i)$ such that $\mu_i \gg_S s_i$.

Proof. If $\mu_i \gg_S s_i$, then $\widehat{p}_i(\mu_i, s_{-i}) > \widehat{p}_i(s_i, s_{-i})$ for all $s_{-i} \in S_{-i}$. Thus, $\widehat{p}_i(\mu_i, \gamma) > \widehat{p}_i(s_i, \gamma)$ for all $\gamma \in \Delta(S_{-i})$. Then, for every belief $\gamma \in \Delta(S_{-i})$, there exists an $s_i' \in \mathrm{supp}(\mu_i)$ such that $\widehat{p}_i(s_i', \gamma) > \widehat{p}_i(s_i, \gamma)$. Therefore, s_i is never a best response.

Conversely, let $s_i^* \in S_i$ be never a best response in Γ. We define a two-person zero-sum game $\Gamma' = (\{0,1\}, (T_0, T_1), (p, -p))$ where $T_0 = S_i - \{s_i^*\}$, $T_1 = S_{-i}$ and $p(s_i, s_{-i}) = p_i(s_i, s_{-i}) - p_i(s_i^*, s_{-i})$. Since s_i^* is never a best response, for every mixed strategy $\mu_1 \in \Delta(T_1) = \Delta(S_{-i})$ there is a strategy $s_0 \in T_0 = S_i - \{s_i^*\}$ such that $\widehat{p}_i(s_0, \mu_1) > \widehat{p}_i(s_i^*, \mu_1)$ (in Γ), i.e. $p(s_0, \mu_1) > 0$ (in Γ'). So, in Γ'

$$\min_{\mu_1 \in \Delta(T_1)} \max_{s_0 \in T_0} p(s_0, \mu_1) > 0,$$

and therefore

$$\min_{\mu_1 \in \Delta(T_1)} \max_{\mu_0 \in \Delta(T_0)} p(\mu_0, \mu_1) > 0.$$

By Nash's Theorem, there is a Nash equilibrium (μ_0^*, μ_1^*) in Γ'. By von Neumann and Morgenstern we have

$$\min_{\mu_1 \in \Delta(T_1)} \max_{s_0 \in \Delta(T_0)} p(\mu_0, \mu_1) = p(\mu_0^*, \mu_1^*)$$
$$= \max_{s_0 \in \Delta(T_0)} \min_{\mu_1 \in \Delta(T_1)} p(\mu_0, \mu_1) > 0.$$

Thus, $0 < p(\mu_0^*, \mu_1^*) \leq p(\mu_0^*, \mu_1)$ for all $\mu_1 \in \Delta(T_1) = \Delta(S_{-i})$. So, we have in Γ $\widehat{p}_i(\mu_0^*, s_{-i}) > p_i(s_i^*, s_{-i})$ for all $s_{-i} \in S_{-i}$ which means $\mu_0^* \gg_S s_i^*$. Q.E.D.

Definition 4.26. Let $\Gamma = (N, (S_i)_{i \in N}, (p_i)_{i \in N})$ be a game. A strategy $s_i \in S_i$ is *rationalisable* in Γ if for any Player j there exists a set $T_j \subseteq S_j$ such that

- $s_i \in T_i$, and
- every $s_j \in T_j$ (for all j) is a best response to a belief $\gamma_j \in \Delta(S_{-j})$ where $\text{supp}(\gamma_j) \subseteq T_{-j}$.

Theorem 4.27. For every finite game Γ we have: s_i is rationalisable if and only if $s_i \in \text{MLS}_i^\infty$. This means, the rationalisable strategies are exactly those surviving iterated elimination of strategies that are strictly dominated by mixed strategies.

Proof. Let $s_i \in S_i$ be rationalisable by $T = (T_1, \ldots, T_n)$. We show $T \subseteq \text{MLS}^\infty$. We will use the monotonicity of MGS and the fact that $\text{MLS}^\infty = \text{MGS}^\infty$. This implies $\text{MGS}^\infty = \text{gfp}(\text{MGS})$ and hence, MGS^∞ contains all other fixed points. It remains to show that $\text{MGS}(T) = T$. Every $s_j \in T_j$ is a best response (among the strategies in S_j) to a belief γ with $\text{supp}(\gamma) \subseteq T_{-j}$. This means that there exists no mixed strategy $\mu_j \in \Delta(S_j)$ such that $\mu_j \gg_T s_j$. Therefore, s_j is not eliminated by MGS: $\text{MGS}(T) = T$.

Conversely, we have to show that every strategy $s_i \in \mathrm{MLS}_i^\infty$ is rationalisable by MLS^∞. Since $\mathrm{MLS}^\infty = \mathrm{MGS}^\infty$, we have $\mathrm{MGS}(\mathrm{MLS}^\infty) = \mathrm{MLS}^\infty$. Thus, for every $s_i \in \mathrm{MLS}_i^\infty$ there is no mixed strategy $\mu_i \in \Delta(S_i)$ such that $\mu_i \gg_{\mathrm{MLS}^\infty} s_i$. So, s_i is a best response to a belief in MLS_i^∞. Q.E.D.

Intuitively, the concept of rationalisability is based on the idea that every player keeps those strategies that are a best response to a possible combined rational action of his opponents. As the following example shows, it is essential to also consider correlated actions of the players.

Example 4.28. Consider the following cooperative game in which every player receives the same payoff:

	L	R	L	R	L	R	L	R
T	8	0	4	0	0	0	3	3
B	0	0	0	4	0	8	3	3
	1		2		3		4	

Matrix 2 is not strictly dominated. Otherwise there were $p, q \in [0,1]$ with $p + q \leq 1$ and

$$8 \cdot p + 3 \cdot (1 - p - q) > 4 \text{ and}$$
$$8 \cdot q + 3 \cdot (1 - p - q) > 4.$$

This implies $2 \cdot (p+q) + 6 > 8$, i.e. $2 \cdot (p+q) > 2$, which is impossible.

So, matrix 2 must be a best response to a belief $\gamma \in \Delta(\{T,B\} \times \{L,R\})$. Indeed, the best responses to $\gamma = \frac{1}{2} \cdot ((T,L) + (B,R))$ are matrices 1, 2 or 3.

On the other hand, matrix 2 is not a best response to a belief of independent actions $\gamma \in \Delta(\{T,B\}) \times \Delta(\{L,R\})$. Otherwise, if matrix 2 was be a best response to $\gamma = (p \cdot T + (1-p) \cdot B, q \cdot L + (1-q) \cdot R)$, we would have that

$$4pq + 4 \cdot (1-p) \cdot (1-q) \geq \max\{8pq, 8 \cdot (1-p) \cdot (1-q), 3\}.$$

We can simplify the left side: $4pq + 4 \cdot (1-p) \cdot (1-q) = 8pq - 4p - 4q + 4$. Obviously, this term has to be greater than each of the terms from which we chose the maximum:

$$8pq - 4p - 4q + 4 \geq 8pq \Rightarrow p + q \geq 1$$

and

$$8pq - 4p - 4q + 4 \geq 8 \cdot (1-p) \cdot (1-q) \Rightarrow p + q \leq 1.$$

So we have $p + q = 1$, or $q = 1 - p$. But this allows us to substitute q by $p - 1$, and we get

$$8pq - 4p - 4q + 4 = 8p \cdot (1-p).$$

However, this term must still be greater or equal than 3, so we get

$$\begin{aligned} 8p \cdot (1-p) &\geq 3 \\ \Leftrightarrow p \cdot (1-p) &\geq \frac{3}{8}, \end{aligned}$$

which is impossible since $\max(p \cdot (1-p)) = \frac{1}{4}$ (see Figure 4.2).

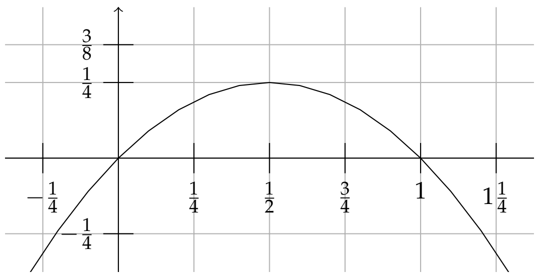

Figure 4.2. Graph of the function $p \mapsto p \cdot (1-p)$

4.4 Games in Extensive Form

A *game in extensive form* (with perfect information) is described by a game tree. For two-person games this is a special case of the games on graphs which we considered in the earlier chapters. The generalisation to n-person games is obvious: $\mathcal{G} = (V, V_1, \ldots, V_n, E, p_1, \ldots, p_n)$ where

(V, E) is a directed tree (with root node w), $V = V_1 \uplus \cdots \uplus V_n$, and the payoff function $p_i : \text{Plays}(\mathcal{G}) \to \mathbb{R}$ for Player i, where Plays($\mathcal{G}$) is the set of paths through (V, E) beginning in the root node, which are either infinite or end in a terminal node.

A strategy for Player i in $\mathcal{G}$ is a function $f : \{v \in V_i : vE \neq \emptyset\} \to V$ such that $f(v) \in vE$. S_i is the set of all strategies for Player i. If all players $1, \ldots, n$ each fix a strategy $f_i \in S_i$, then this defines a unique play $f_1\hat{}\cdots\hat{}f_n \in \text{Plays}(\mathcal{G})$.

We say that $\mathcal{G}$ has *finite horizon* if the depth of the game tree (the length of the plays) is finite.

For every game $\mathcal{G}$ in extensive form, we can construct a game $S(\mathcal{G}) = (N, (S_i)_{i\in N}, (p_i)_{i\in N})$ with $N = \{1, \ldots, n\}$ and $p_i(f_1, \ldots, f_n) = p_i(f_1\hat{}\cdots\hat{}f_n)$. Hence, we can apply all solution concepts for strategic games (Nash equilibria, iterated elimination of dominated strategies, etc.) to games in extensive form. First, we will discuss Nash equilibria in extensive games.

Example 4.29. Consider the game $\mathcal{G}$ (of finite horizon) depicted in Figure 4.3 presented as (a) an extensive-form game and as (b) a strategic-form game. The game has two Nash equilibria:

- The natural solution (b, d) where both players win.
- The second solution (a, c) which seems to be irrational since both players pick an action with which they lose.

What seems irrational about the second solution is the following observation. If Player 0 picks a, it does not matter which strategy her opponent chooses since the position v is never reached. Certainly, if Player 0 switches from a to b, and Player 1 still responds with c, the payoff of Player 0 does not increase. But this threat is not credible since if v is reached after action a, then action d is better for Player 1 than c. Hence, Player 0 has an incentive to switch from a to b.

This example shows that the solution concept of Nash equilibria is not sufficient for games in extensive form since they do not take the sequential structure into account. Before we introduce a stronger notion of equilibrium, we will need some more notation: Let $\mathcal{G}$ be a game in extensive form and v a position of $\mathcal{G}$. $\mathcal{G} \restriction_v$ denotes the *subgame* of $\mathcal{G}$

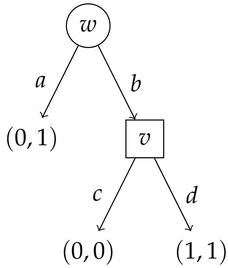

(a) extensive form

	c	d
a	$(0,1)$	$(0,1)$
b	$(0,0)$	$(1,1)$

(b) strategic form

Figure 4.3. A game of finite horizon

beginning in v (defined by the subtree of $\mathcal{G}$ rooted at v). Payoffs: Let h_v be the unique path from w to v in $\mathcal{G}$. Then $p_i^{\mathcal{G}\restriction_v}(\pi) = p_i^{\mathcal{G}}(h_v \cdot \pi)$. For every strategy f of Player i in $\mathcal{G}$ let $f\restriction_v$ be the restriction of f to $\mathcal{G}\restriction_v$.

Definition 4.30. A *subgame perfect equilibrium* of $\mathcal{G}$ is a strategy profile $(f_1, \dots, f_n)$ such that, for every position v, $(f_1\restriction_v, \dots, f_n\restriction_v)$ is a Nash equilibrium of $\mathcal{G}\restriction_v$. In particular, $(f_1, \dots, f_n)$ itself is a Nash equilibrium.

In the example above, only the natural solution (b, d) is a subgame perfect equilibrium. The second Nash equilibrium (a, c) is not a subgame perfect equilibrium since $(a\restriction_v, c\restriction_v)$ is not a Nash equilibrium in $\mathcal{G}\restriction_v$.

Let $\mathcal{G}$ be a game in extensive form, $f = (f_1, \dots, f_n)$ be a strategy profile, and v a position in $\mathcal{G}$. We denote by $\widetilde{f}(v)$ the play in $\mathcal{G}\restriction_v$ that is uniquely determined by $f_1 \dots, f_n$.

Lemma 4.31. Let $\mathcal{G}$ be a game in extensive form with finite horizon. A strategy profile $f = (f_1, \dots, f_n)$ is a subgame perfect equilibrium of $\mathcal{G}$ if and only if for every Player i, every $v \in V_i$, and every $w \in vE$: $p_i(\widetilde{f}(v)) \geq p_i(\widetilde{f}(w))$.

Proof. Let f be a subgame perfect equilibrium. If $p_i(\widetilde{f}(w)) > p_i(\widetilde{f}(v))$ for some $v \in V_i$, $w \in vE$, then it would be better for Player i in $\mathcal{G}\restriction_v$ to change her strategy in v from f_i to f_i' with

$$f_i'(u) = \begin{cases} f_i(u) & \text{if } u \neq v \\ w & \text{if } u = w\,. \end{cases}$$

This is a contradiction.

Conversely, if f is not a subgame perfect equilibrium, then there is a Player i, a position $v_0 \in V_i$ and a strategy $f_i' \neq f_i$ such that it is better for Player i in $\mathcal{G}\restriction_{v_0}$ to switch from f_i to f_i' against f_{-i}. Let $g := (f_i', f_{-i})$. We have $q := p_i(\widetilde{g}(v_0)) > p_i(\widetilde{f}(v_0))$. We consider the path $\widetilde{g}(v_0) = v_0 \dots v_t$ and pick a maximal $m < t$ with $p_i(\widetilde{g}(v_0)) > p_i(\widetilde{f}(v_m))$. Choose $v = v_m$ and $w = v_{m+1} \in vE$. Claim: $p_i(\widetilde{f}(v)) < p_i(\widetilde{f}(w))$ (see Figure 4.4):

$$p_i(\widetilde{f}(v)) = p_i(\widetilde{f}(v_m)) < p_i(\widetilde{g}(v_m)) = q$$

$$p_i(\widetilde{f}(w)) = p_i(\widetilde{f}(v_{m+1})) \geq p_i(\widetilde{g}(v_{m+1})) = q$$

Q.E.D.

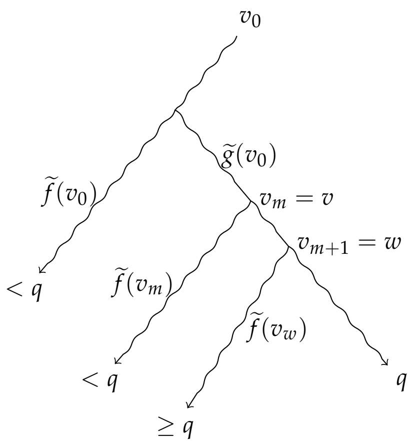

Figure 4.4. $p_i(\widetilde{f}(v)) < p_i(\widetilde{f}(w))$

If f is not a subgame perfect equilibrium, then we find a subgame $\mathcal{G}\restriction_v$ such that there is a profitable deviation from f_i in $\mathcal{G}\restriction_v$, which only differs from f_i in the first move.

In extensive games with finite horizon we can directly define the payoff at the terminal nodes (the leaves of the game tree). We obtain a payoff function $p_i : T \to \mathbb{R}$ for $i = 1, \ldots, n$ where $T = \{v \in V : vE = \emptyset\}$.

Backwards induction: For finite games in extensive form we define a strategy profile $f = (f_1, \ldots, f_n)$ and values $u_i(v)$ for all positions v and every Player i by backwards induction:

- For terminal nodes $t \in T$ we do not need to define f, and $u_i(t) := p_i(t)$.
- Let $v \in V \setminus T$ such that all $u_i(w)$ for all i and all $w \in vE$ are already defined. For i with $v \in V_i$ define $f_i(v) = w$ for some w with $u_i(w) = \max\{u_i(w') : w' \in vE\}$ and $u_j(v) := u_j(f_i(v))$ for all j.

We have $p_i(\widetilde{f}(v)) = u_i(v)$ for every i and every v.

Theorem 4.32. The strategy profile defined by backwards induction is a subgame perfect equilibrium.

Proof. Let $f'_i \neq f_i$. Then there is a node $v_0 \in V_i$ with minimal height in the game tree such that $f'_i(v) \neq f_i(v)$. Especially, for every $w \in vE$, $\widetilde{(f'_i, f_{-i})}(w) = \widetilde{f}(w)$. For $w = f'_i(v)$ we have

$$\begin{aligned} p_i(\widetilde{(f'_i, f_{-i})}(v)) &= p_i(\widetilde{(f'_i, f_{-i})}(w)) \\ &= p_i(\widetilde{f}(w)) \\ &= u_i(w) \leq \max_{w' \in vE}\{u_i(w')\} \\ &= u_i(v) \\ &= p_i(\widetilde{f}(v)). \end{aligned}$$

Therefore, $f\restriction_v$ is a Nash equilibrium in $\mathcal{G}\restriction_v$. Q.E.D.

Corollary 4.33. Every finite game in extensive form has a subgame perfect equilibrium (and thus a Nash equilibrium) in pure strategies.