Logic and Games
WS 2015/2016

Prof. Dr. Erich Grädel
Notes and Revisions by Matthias Voit

Mathematische Grundlagen der Informatik
RWTH Aachen

This work is licensed under:
http://creativecommons.org/licenses/by-nc-nd/3.0/de/
Dieses Werk ist lizenziert unter:
http://creativecommons.org/licenses/by-nc-nd/3.0/de/

© 2016 Mathematische Grundlagen der Informatik, RWTH Aachen.
http://www.logic.rwth-aachen.de

Contents

1 Reachability Games and First-Order Logic 1
1.1 Model Checking 1
1.2 Model Checking Games for Modal Logic 2
1.3 Reachability and Safety Games 5
1.4 Games as an Algorithmic Construct: Alternating Algorithms . 10
1.5 Model Checking Games for First-Order Logic 20

2 Parity Games and Fixed-Point Logics 25
2.1 Parity Games 25
2.2 Algorithms for parity games 30
2.3 Fixed-Point Logics 35
2.4 Model Checking Games for Fixed-Point Logics 37
2.5 Defining Winning Regions in Parity Games 42

3 Infinite Games 45
3.1 Determinacy 45
3.2 Gale-Stewart Games 47
3.3 Topology 53
3.4 Determined Games 59
3.5 Muller Games and Game Reductions 61
3.6 Complexity 74

4 Basic Concepts of Mathematical Game Theory 79
4.1 Games in Strategic Form 79
4.2 Nash equilibria 81
4.3 Two-person zero-sum games 85
4.4 Regret minimization 86
4.5 Iterated Elimination of Dominated Strategies 89
4.6 Beliefs and Rationalisability 95

4.7 Games in Extensive Form 98
4.8 Subgame-perfect equilibria in infinite games 102

Appendix A 111
4.9 Cardinal Numbers 119

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

- $N = \{1, \ldots, n\}$ 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 every 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 $s' \in S_i$.

Obviously, a dominant strategy is a best response to all strategy profiles of the other players.

Example 4.5. **The 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 repeatedly, possibly infinitely often.

Example 4.6. **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 or her strategy unilaterally.

4.2 Nash equilibria

Definition 4.7. A strategy profile $s = (s_1, \ldots, s_n) \in S$ is a *Nash equilibrium* in Γ if for all $i \in N$ and all strategies $s'_i \in S_i$

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

Thus, in a Nash equilibrium, every player plays with a best response to the profile of his opponents, and thus has no incentive to deviate unilaterally to a different strategy. Is there a Nash equilibrium in every game? The following example shows that this is not always the case, at least not 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.

Although there are no Nash equilibria in pure strategies in rock, paper, scissors, there is of course an obvious good method to play 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 (so that $\sum_{s\in S_i} \mu(s) = 1$).
$\Delta(S_i)$ 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 $\mu_i' \in \Delta(S)$, $\widehat{p}_i(\mu_i,\mu_{-i}) \geq \widehat{p}_i(\mu_i',\mu_{-i})$.

Nash equilibria (in mixed strategies) provide the arguably most important solution concept in classical game theory (although, as we shall point out later, this concept is not without problems). An important

reason for the success of Nash equilibrium as a solution concept is the fact that every finite game has one. To prove this, we shall use a well-known classical fixed-point theorem.

Theorem 4.11 (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.

We do not prove this here but remark that, interestingly, the Brouwer Fixed-Point Theorem can itself be proved via a game-theoretic result, namely the determinacy of HEX.

Theorem 4.12 (Nash)**.** Every finite game Γ in strategic form has at least one Nash equilibrium in mixed strategies.

Proof. 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 $\operatorname{supp}(\mu_i) = \{s \in S_i : \mu_i(s) > 0\}$.

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

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

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.

4.3 Two-person zero-sum games

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 from mathematical logic is that $\exists x\, \forall y\, Rxy \models \forall y\, \exists x\, Rxy$.

Theorem 4.15 (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.

4.4 Regret minimization

To motivate the concept of regret minimization we consider

Example 4.16. **Traveller's Dilemma.** This is a symmetric two-player game $\Gamma = (\{1,2\}, (S_1, S_2), (p_1, p_2))$ with $S_1 = S_2 = \{2, \ldots, 100\}$ and

$$p_1(x,y) = \begin{cases} x+2 & \text{if } x < y, \\ y-2 & \text{if } y < x, \\ x & \text{if } x = y, \end{cases} \qquad p_2(x,y) = p_1(y,x)$$

The only Nash equilibrium in pure strategies 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). Also most other solution concepts from game theory (such as the iterated elimination of dominated strategies discussed in the next section) suggest that the players should choose 2.

However, experiments show that people (even game theorists!) tend to select large numbers, in the range between 90 and 100; moreover they seem right to do so, since they perform much better in these experiments than those who follow what game theory proposes and select strategy 2.

The question arises whether there are alternative solution concepts that justify the choice of large strategies in the Traveller's Dilemma, and if yes, which one. A relatively recent proposal that seems to achieve this is *regret minimization*. When a player uses this concept, he wants to minimize the lost payoff (which he would "regret") due to not playing with the best response to the strategies of the other players.

This idea was formulated in the context of decision theory, concerned with the choices of individual agents rather than the interaction of different agents as in game theory. Accordingly, the payoff is determined by a binary function $p : S \times Z \to \mathbb{R}$, where S is the set of strategies of the player we are considering, and Z is an abstract set of possible *states*.

Before we can introduce regret minimization, we need several definitions. In state $z \in Z$, the maximal payoff for our player is

$$p^*(z) := \max_{s \in S} p(s,z),$$

and if the player chooses the strategy $s \in S$, he will miss the following payoff:

$$\text{regret}_p(s,z) := p^*(z) - p(s,z).$$

The overall maximal regret for the strategy s is

$$\text{maxreg}_p(s) := \max_{z \in Z} \text{regret}_p(s,z).$$

Now, the decision with respect to regret minimization would be: *Choose $s \in S$ such that* $\text{maxreg}_p(s)$ *is minimal.*

Let us reconsider Example 4.16. Since it belongs to game theory, Z is the set of strategy profiles of the other players. We claim that exactly the strategies $s \in \{96, \dots, 100\}$ minimize the maximal regret. To see this, note that for those s, we have that $\text{maxreg}_p(s) = 3$, since

- if $t \leq s$, then $p(s,t) \geq t-2$ and $p^*(t) \leq t+1$, thus $\text{regret}_p(s,t) = p^*(t) - p(s,t) \leq t+1-(t-2) = 3$,
- if $t > s$, then $p(s,t) = s+2$ and $p^*(t) \leq 101$, thus $\text{regret}_p(s,t) \leq 101-(s+2) = 99-s \leq 3$,

and on the other hand,

- $\text{regret}_p(96,100) = 101-98 = 3$,
- for $s \in \{97, \dots, 100\}$, $\text{regret}_p(s,96) = 97-94 = 3$.

Also, for $s \leq 95$, we have that $\text{maxreg}_p(s) \geq 4$, as $\text{maxreg}_p(s) \geq \text{regret}_p(s,100) = 101-(s+2) = 99-s \geq 4$.

Consequently, regret minimization suggests a strategy s with $96 \leq s \leq 100$. We will now iterate this idea. If both players eliminate strategies which do not minimize the regret, we obtain a subgame with strategies $\{96, \dots, 100\}$. In this game, we have that

- $\text{maxreg}_p(97) = 2$, since
  - $\text{regret}_p(97,100) = 101-99 = 2$,
  - $\text{regret}_p(97,99) = 100-99 = 1$,
  - $\text{regret}_p(97,98) = 99-99 = 0$,
  - $\text{regret}_p(97,97) = 98-97 = 1$,
  - $\text{regret}_p(97,96) = 96-95 = 2$.
- $\text{maxreg}_p(100) \geq \text{regret}_p(100,99) = 100-97 = 3$.
- $\text{maxreg}_p(99) \geq \text{regret}_p(99,98) = 99-96 = 3$.
- $\text{maxreg}_p(98) \geq \text{regret}_p(98,97) = 98-95 = 3$.

- $\text{maxreg}_p(96) \geq \text{regret}_p(96,100) = 101-98 = 3$.

Hence, 97 is the unique strategy which minimizes the regret in this subgame and thus is the choice of a player who assumes that his opponent wants to minimize his regret as well.

4.5 Iterated Elimination of Dominated Strategies

Besides Nash equilibria and (iterated) regret minimization, 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.

$$\begin{array}{c|cc|cc}
 & L & R & L & R \\
T & (1,0,1) & \boxed{(1,1,0)} & (1,0,1) & (0,1,0) \\
B & (1,1,1) & (0,0,1) & (1,1,1) & (1,0,0) \\
 & \multicolumn{2}{c|}{\underbrace{\qquad\qquad}_{X}} & \multicolumn{2}{c}{\underbrace{\qquad\qquad}_{Y}}
\end{array}$$

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.

$$\begin{tikzcd}
 & \mathrm{PL}(T) \arrow[rr] & & \mathrm{PLS}(T) \\
\mathrm{ML}(T) \arrow[ur] \arrow[rr] & & \mathrm{MLS}(T) \arrow[ur] & \\
 & \mathrm{PG}(T) \arrow[uu] \arrow[rr] & & \mathrm{PGS}(T) \arrow[uu] \\
\mathrm{MG}(T) \arrow[uu] \arrow[ur] \arrow[rr] & & \mathrm{MGS}(T) \arrow[uu] \arrow[ur] &
\end{tikzcd}$$

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 $\mathrm{ML}^\infty \nsubseteq \mathrm{PL}^\infty$.

	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	$(0,1)$	$(2,1)$

$\mathrm{PL}(\mathrm{PL}^1) = (\{A,B,C\},\{X\}) = \mathrm{PL}^2 \subsetneqq \mathrm{PL}^1$ 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, $\mathrm{PL}^\infty \subsetneqq \mathrm{ML}^\infty$ 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 $s_i \in \mathrm{PGS}(T')_i$. 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 $s_i \in \mathrm{PGS}(T)_i$. 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	$(0,0)$

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

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 $\mathrm{ML}^\infty \subseteq \mathrm{MLS}^\infty \subseteq \mathrm{PLS}^\infty$ and $\mathrm{PL}^\infty \subseteq \mathrm{PLS}^\infty$ 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. $\mathrm{MGS}^\infty = \mathrm{MLS}^\infty$ and $\mathrm{PGS}^\infty = \mathrm{PLS}^\infty$.

Proof. We will only prove $\mathrm{PGS}^\infty = \mathrm{PLS}^\infty$. Since $\mathrm{PGS}(T) \subseteq \mathrm{PLS}(T)$ for all T and PGS is monotone, we have $\mathrm{PGS}^\infty \subseteq \mathrm{PLS}^\infty$. 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.

4.6 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 $s_i \in \text{MLS}_i^\infty$ is rationalisable by MLS^∞. Since $\text{MLS}^\infty = \text{MGS}^\infty$, we have $\text{MGS}(\text{MLS}^\infty) = \text{MLS}^\infty$. Thus, for every $s_i \in \text{MLS}_i^\infty$ there is no mixed strategy $\mu_i \in \Delta(S_i)$ such that $\mu_i \gg_{\text{MLS}^\infty} s_i$. 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:

$$\begin{array}{c|cc|cc|cc|cc}
 & L & R & L & R & L & R & L & R \\
T & 8 & 0 & 4 & 0 & 0 & 0 & 3 & 3 \\
B & 0 & 0 & 0 & 4 & 0 & 8 & 3 & 3 \\
 & \multicolumn{2}{c}{\underbrace{}_{1}} & \multicolumn{2}{c}{\underbrace{}_{2}} & \multicolumn{2}{c}{\underbrace{}_{3}} & \multicolumn{2}{c}{\underbrace{}_{4}}
\end{array}$$

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 were 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 $1 - p$, 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).

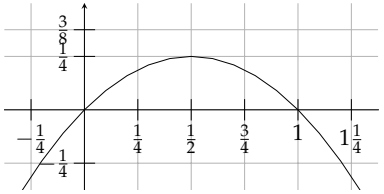

Figure 4.2. Graph of the function $p \mapsto p \cdot (1-p)$

4.7 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 $\text{Plays}(\mathcal{G})$ 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.

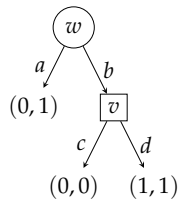

(a) extensive form

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

(b) strategic form

Figure 4.3. A game of finite horizon

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 $\mathcal{G}$ 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):

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

$$p_i(\widetilde{f}(w)) = p_i(\widetilde{f}(v_{m+1})) \geq p_i(\widetilde{g}(v_{m+1})) = q \qquad \text{Q.E.D.}$$

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.

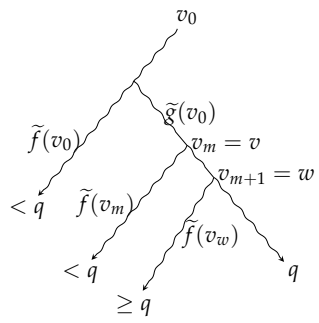

Figure 4.4. $p_i(\widetilde{f}(v)) < p_i(\widetilde{f}(w))$

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 $i = 1, \dots, n$ where $T = \{v \in V : vE = \emptyset\}$.

Backwards induction: For finite games in extensive form we define a strategy profile $f = (f_1, \dots, f_n)$ 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

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

$$= u_i(v)$$
$$= p_i(\widetilde{f}(v)).$$

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.

4.8 Subgame-perfect equilibria in infinite games

We now consider cases of infinite games in extensive form, for which we can establish the existence of subgame-perfect equilibria. Generalizing the model of infinite two-person zero-sum games on graphs, we consider multi-player, turn-based games on graphs with arbitrary (not necesssarily antagonistic) qualitative objectives.

Definition 4.34. An *infinite (turn-based, qualitative) multiplayer game* is a tuple $\mathcal{G} = (N, V, (V_i)_{i\in N}, E, \Omega, (\mathrm{Win}_i)_{i\in N})$ where N is a finite set of *players*, (V, E) is a (finite or infinite) directed graph $(V_i)_{i\in N}$ is a partition of V into the position sets for each player, $\Omega : V \to C$ is a colouring of the positions by some finite set C of colours, and $\mathrm{Win}_i \subseteq C^\omega$ is the winning condition for Player i.

For the sake of simplicity, we assume that $uE := \{v \in V : (u, v) \in E\} \neq \emptyset$ for all $u \in V$, i.e. each vertex of G has at least one outgoing edge. We call $\mathcal{G}$ a *zero-sum game* if the sets Win_i define a partition of C^ω.

A *play* of $\mathcal{G}$ is an infinite path through the graph (V, E), and a *history* is a finite initial segment of a play. We say that a play π is *won* by Player i if $\Omega(\pi) \in \mathrm{Win}_i$. A *(pure) strategy of Player i in* $\mathcal{G}$ is a function $f : V^*V_i \to V$ assigning to each sequence xv ending in a position v of Player i a next position $f(xv) \in vE$. We say that a play $\pi = \pi(0)\pi(1)\ldots$ of $\mathcal{G}$ is *consistent* with a strategy f of Player i if $\pi(k+1) = f(\pi(0)\ldots\pi(k))$ for all $k < \omega$ with $\pi(k) \in V_i$. A *strategy profile* of $\mathcal{G}$ is a tuple $(f_i)_{i\in N}$ where f_i is a strategy of Player i.

It is sometimes convenient to designate an initial vertex $v_0 \in V$ of the game. We call the tuple $(\mathcal{G}, v_0)$ an *initialized infinite multiplayer game*. A *play (history) of* $(\mathcal{G}, v_0)$ is a play (history) of $\mathcal{G}$ starting with v_0.

A strategy (strategy profile) of $(\mathcal{G}, v_0)$ is just a strategy (strategy profile) of $\mathcal{G}$. A strategy f of some player i in $(\mathcal{G}, v_0)$ is *winning* if every play of $(\mathcal{G}, v_0)$ consistent with σ is won by player i. A strategy profile $(f_i)_{i\in N}$ of $(\mathcal{G}, v_0)$ determines a unique play of $(\mathcal{G}, v_0)$ consistent with each f_i, called the *outcome of* $(f_i)_{i\in N}$ and denoted by $\langle (f_i)_{i\in N}\rangle$ or, in the case that the initial vertex is not understood from the context, $\langle (f_i)_{i\in N}\rangle_{v_0}$. In the following we will often use the term *game* to denote an *(initialized) infinite multiplayer game* according to Definition 4.34.

For turn-based (non-stochastic) games with qualitative winning conditions, mixed strategies play no relevant role. Nash equilibria in pure strategies take the following form:

A strategy profile $(f_i)_{i\in N}$ of a game $(\mathcal{G}, v_0)$ is a *Nash equilibrium* if for every player i and all her possible strategies f'_i in $(\mathcal{G}, v_0)$ the play $\langle f'_i, (f_j)_{j\in N\setminus\{i\}}\rangle$ is won by player i only if the play $\langle (f_j)_{j\in N}\rangle$ is also won by her.

Despite the importance and popularity of Nash equilibria, there are several problems with this solution concept, in particular for games that extend over time. This is due to the fact that Nash equilibria do not take into account the sequential nature of games and all the consequences of this. After any initial segment of a play, the players face a new situation and may change their strategies. Choices made because of a threat by the other players may no longer be rational, because the opponents have lost their power of retaliation in the remaining play.

Example 4.35. Consider a two-player Büchi game with its arena depicted in Figure 4.5; round vertices are controlled by player 1; boxed vertices are controlled by player 2; both players win if and only if vertex 3 is visited (infinitely often); the initial vertex is 1. Intuitively, the only rational outcome of this game should be the play 123^ω. However, the game has two Nash equilibria:

(1) Player 1 moves from vertex 1 to vertex 2, and player 2 moves from vertex 2 to vertex 3. Hence, both players win.
(2) Player 1 moves from vertex 1 to vertex 4, and player 2 moves from vertex 2 to vertex 5. Both players lose.

The second equilibrium certainly does not describe a rational behaviour. Indeed both players move according to a strategy that is always

losing (whatever the other player does), and once player 1 has moved from vertex 1 to vertex 2, then the rational behaviour of player 2 would be to change her strategy and move to vertex 3 instead of vertex 5 as this is then the only way for her to win.

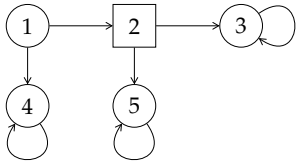

Figure 4.5. A two-player Büchi game.

This example can be modified in many ways. Indeed we can construct games with Nash equilibria in which every players moves infinitely often according to a losing strategy, and only has a chance to win if she deviates from the equilibrium strategy. The following is an instructive example with quantitative objectives.

Example 4.36. Let $\mathcal{G}_n$ be an n-player game with positions $0, \dots, n-1$. Position n is the initial position, and position 0 is the terminal position. Player i moves at position i and has two options. Either she loops at position i (and stays in control) or moves to position $i-1$ (handing control to the next player). For each player, the value of a play π is $n/|\pi|$. Hence, for all players, the shortest possible play has value 1, and all infinite plays have value 0. Obviously, the rational behaviour for each player i is to move from i to $i-1$. This strategy profile, which is of course a Nash equilibrium, gives value 1 to all players. However, the 'most stupid' strategy profile, where each player loops forever at his position, i.e. moves forever according to a losing strategy, is also a Nash equilibrium.

For a game $\mathcal{G} = (N, V, (V_i)_{i\in N}, E, \Omega, (\mathrm{Win}_i)_{i\in\Pi})$ and a history h of $\mathcal{G}$, let the game $\mathcal{G}|_h = (N, V, (V_i)_{i\in N}, E, \Omega, (\mathrm{Win}_i|_h)_{i\in N})$ be defined by $\mathrm{Win}_i|_h = \{\alpha \in C^\omega : \Omega(h)\cdot\alpha \in \mathrm{Win}_i\}$. For an initialized game $(\mathcal{G}, v_0)$ and a history hv of $(\mathcal{G}, v_0)$, we call the initialized game $(\mathcal{G}|_h, v)$ the *subgame of* $(\mathcal{G}, v_0)$ *with history* hv. For a strategy f of Player i in $\mathcal{G}$, let

$f|_h : V^*V_i \to V$ be defined by $f|_h(xv) := f(hxv)$. Obviously, $f|_h$ is a strategy of Player i in $\mathcal{G}|_h$.

Recall that a strategy profile $(f_i)_{i\in N}$ is a *subgame perfect equilibrium (SPE)* if $(f_i|_h)_{i\in N}$ is a Nash equilibrium of $(\mathcal{G}|_h, v)$ for every history hv of $(\mathcal{G}, v_0)$.

Example 4.37. Consider again the game described in Example 4.35. The Nash equilibrium where Player 1 moves from vertex 1 to vertex 4 and Player 2 moves from vertex 2 to vertex 5 is not a subgame perfect equilibrium since moving from vertex 2 to vertex 5 is not optimal for Player 2 after the play has reached vertex 2. On the other hand, the Nash equilibrium where Player 1 moves from vertex 1 to vertex 2 and Player 2 moves from vertex 2 to vertex 3 is also a subgame perfect equilibrium.

The first step in the analysis of subgame perfect equilibria for *infinite* duration games is the notion of subgame-perfect determinacy. While the notion of subgame perfect equilibrium makes sense for more general classes of extensive games, the notion of subgame-perfect determinacy applies only to games with qualitative winning conditions.

Definition 4.38. A game $(\mathcal{G}, v_0)$ is *subgame-perfect determined* if there exists a strategy profile $(f_i)_{i\in N}$ such that for each history hv of the game one of the strategies $f_i|_h$ is a winning strategy in $(\mathcal{G}|_h, v)$.

Proposition 4.39. Let $(\mathcal{G}, v_0)$ be a qualitative zero-sum game such that every subgame is determined. Then $(\mathcal{G}, v_0)$ is subgame-perfect determined.

Proof. Let $(\mathcal{G}, v_0)$ be a multiplayer game such that, for every history hv, there exists a strategy f_i^h for some player i, which is winning in $(\mathcal{G}|_h, v)$. We have to combine these strategies in an appropriate way to strategies f_i. (Let us point out that the trivial combination, namely $f_i(hv) := f_i^h(v)$ does not work in general.) We say that a decomposition $h = h_1 \cdot h_2$ is *good* for player i w.r.t. vertex v if $f_i^{h_1}|_{h_2}$ is winning in $(\mathcal{G}|_h, v)$. If the strategy f_i^h is winning in $(\mathcal{G}|_h, v)$, then the decomposition $h = h \cdot \varepsilon$ is good w.r.t. v, so a good decomposition exists.

For each history hv, if f_i^h is winning in $(\mathcal{G}|_h, v)$, we choose the good

(w.r.t. vertex v) decomposition $h = h_1 h_2$ with minimal h_1, and put

$$f_i(hv) := f_i^{h_1}(h_2 v)\,.$$

Otherwise, we set $f_i(hv) := f_i^h(v)$.

It remains to show that for each history hv of $(\mathcal{G}, v_0)$ the strategy $f_i|_h$ is winning in $(\mathcal{G}|_h, v)$ whenever the strategy f_i^h is. Hence, assume that f_i^h is winning in $(\mathcal{G}|_h, v)$, and let $\pi = \pi(0)\pi(1)\ldots$ be a play starting in $\pi(0) = v$ and consistent with $f_i|_h$. We need to show that π is won by player i in $(\mathcal{G}|_h, v)$.

First, we claim that for each $k < \omega$ there exists a decomposition of the form $h\pi(0)\ldots\pi(k-1) = h_1 \cdot (h_2\pi(0)\ldots\pi(k-1))$ that is good for player i w.r.t. $\pi(k)$. This is obviously true for $k = 0$. Now, for $k > 0$, assume that there exists a decomposition $h\pi(0)\ldots\pi(k-2) = h_1 \cdot (h_2\pi(0)\ldots\pi(k-2))$ that is good for player i w.r.t. $\pi(k-1)$ and with h_1 being minimal. Then $\pi(k) = f_i(h\pi(0)\ldots\pi(k-1)) = f^{h_1}(h_2\pi(0)\ldots\pi(k-1)$, and $h\pi(0)\ldots\pi(k-1) = h_1(h_2\pi(0)\ldots\pi(k-1))$ is a decomposition that is good w.r.t. $\pi(k)$.

Now consider the sequence $h_1^0, h_1^1, \ldots$ of prefixes of the good decompositions $h\pi(0)\ldots\pi(k-1) = h_1^k h_2^k \pi(0)\ldots\pi(k-1)$ (w.r.t. $\pi(k)$) with each h_1^k being minimal. Then we have $h_1^0 \succeq h_1^1 \succeq \ldots$, since for each $k > 0$ the decomposition $h\pi(0)\ldots\pi(k-1) = h_1^{k-1} h_2^{k-1}\pi(0)\ldots\pi(k-1)$ is also good for player i w.r.t. $\pi(k)$. As $\prec$ is well-founded, there must exist $k < \omega$ such that $h_1 := h_1^k = h_1^l$ for each $k \leq l < \omega$. Hence, we have that the play $\pi(k)\pi(k+1)\ldots$ is consistent with $f_i^{h_1}|_{h_2\pi(0)\ldots\pi(k-1)}$, which is a winning strategy in $(\mathcal{G}|_{h\pi(0)\ldots\pi(k-1)}, \pi(k))$. So the play $h\pi$ is won by player i in $(\mathcal{G}, v_0)$, which implies that the play π is won by player i in $(\mathcal{G}|_h, v)$. Q.E.D.

We say that a class of winning conditions is closed under taking subgames, if for every condition $X \subseteq C^\omega$ in the class, and every $h \in C^*$, also $X|_h := \{x \in C^\omega : hx \in X\}$ belongs to the class. Since Borel winning conditions are closed under taking subgames, it follows that any two-player zero-sum game with Borel winning condition is subgame-perfect determined.

Corollary 4.40. Let $(\mathcal{G}, v_0)$ be a two-player zero-sum Borel game. Then $(\mathcal{G}, v_0)$ is subgame-perfect determined.

Multiplayer games are usually not zero-sum games. Indeed when we have many players the assumption that the winning conditions of the players form a partition of the set of plays is very restrictive and unnatural. We now drop this assumption and establish general conditions under which a multiplayer game admits a subgame perfect equilibrium. In fact we will relate the existence of subgame perfect equilibria with the determinacy of associated two-player games. In particular, it will follow that every multiplayer game with Borel winning conditions has a subgame perfect equilibrium.

In the rest of this subsection, we are only concerned with the *existence* of equilibria, not with their complexity. Thus, without loss of generality, we tacitly assume that the arena of the game under consideration is a tree or a forest with the initial vertex as one of its root. The justification for this assumption is that we can always replace the arena of an arbitrary game by its unravelling from the initial vertex, ending up in an equivalent game.

Definition 4.41. Let $\mathcal{G} = (N, V, (V_i)_{i\in N}, E, \Omega, (\mathrm{Win}_i)_{i\in N})$ be a multiplayer game (played on a forest), with winning conditions $\mathrm{Win}_i \subseteq C^\omega$. The associated class $\mathrm{Two}(\mathcal{G})$ of two-player zero-sum games is obtained as follows:

(1) For each player i, $\mathrm{Two}(\mathcal{G})$ contains the game $\mathcal{G}_i$ where player i plays $\mathcal{G}$, with his winning condition Win_i, against the coalition of all other players, with winning condition $C^\omega \setminus \mathrm{Win}_i$.
(2) Close the class under taking subgames (i.e. consider plays after initial histories).
(3) Close the class under taking subgraphs (i.e. admit deletion of positions and moves).

Note that the order in which the operations (1), (2), and (3) are applied has no effect on the class $\mathrm{Two}(\mathcal{G})$.

Theorem 4.42. Let $(\mathcal{G}, v_0)$ be a multiplayer game such that every game in $\mathrm{Two}(\mathcal{G})$ is determined. Then $(\mathcal{G}, v_0)$ has a subgame perfect equilibrium.

Proof. Let $\mathcal{G} = (N, V, (V_i)_{i\in N}, E, \Omega, (\mathrm{Win}_i)_{i\in N})$ be a multiplayer game such that every game in $\mathrm{Two}(\mathcal{G})$ is determined. For each ordinal α we define a set $E^\alpha \subseteq E$ beginning with $E^0 = E$ and

$$E^\lambda = \bigcap_{\alpha<\lambda} E^\alpha$$

for limit ordinals λ. To define $E^{\alpha+1}$ from E^α, we consider for each player $i \in N$ the two-player zero-sum game $\mathcal{G}_i^\alpha = (V, V_i, E^\alpha, \Omega, \mathrm{Win}_i)$ where player i plays, with his winning condition Win_i against the coalition of all other players (with winning condition $C^\omega \setminus \mathrm{Win}_i$). Every subgame of $\mathcal{G}_i^\alpha$ belongs to $\mathrm{Two}(\mathcal{G})$ and is therefore determined. Hence we can use Proposition 4.39 to fix a subgame perfect equilibrium $(f_i^\alpha, f_{-i}^\alpha)$ of $(\mathcal{G}, v_0)$ where f_i^α is a strategy of player i and f_{-i}^α is a strategy of the coalition. Moreover, as the arena of $\mathcal{G}^\alpha$ is a forest, these strategies can be assumed to be positional. Let X_i^α be the set of all $v \in V$ such that f_i^α is winning in $(\mathcal{G}_i^\alpha|_h, v)$ for the unique maximal history h of $\mathcal{G}$ leading to v. For vertices $v \in V_i \cap X_i^\alpha$ we delete all outgoing edges except the one taken by the strategy f_i^α, i.e. we define

$$E^{\alpha+1} = E^\alpha \setminus \bigcup_{i\in N} \{(u,v) \in E : u \in V_i \cap X_i^\alpha \text{ and } v \neq f_i^\alpha(u)\}.$$

Obviously, the sequence $(E^\alpha)_{\alpha\in\mathrm{On}}$ is non-increasing. Thus we can fix the least ordinal γ with $E^\gamma = E^{\gamma+1}$ and define $f_i = f_i^\gamma$ and $f_{-i} = f_{-i}^\gamma$. Moreover, for each player $j \neq i$ let $f_{j,i}$ be the positional strategy of player j in $\mathcal{G}$ that is induced by f_{-i}.

Intuitively, Player i's equilibrium strategy g_i is as follows: Player i plays f_i as long as no other player deviates. Whenever some player $j \neq i$ deviates from her equilibrium strategy f_j, player i switches to $f_{i,j}$. Formally, define for each vertex $v \in V$ the player $p(v)$ who has to be "punished" at vertex v where $p(v) = \perp$ if nobody has to be punished. If the game has just started, no player should be punished. Thus we let

$$p(v) = \perp \text{ if } v \text{ is a root.}$$

At vertex v with predecessor u, the same player has to be punished as

at vertex u as long as the player whose turn it was at vertex u did not deviate from her prescribed strategy. Thus for $u \in V_i$ and $v \in uE$ we let

$$p(v) = \begin{cases} \perp & \text{if } p(u) = \perp \text{ and } v = f_i(u), \\ p(u) & \text{if } p(u) \neq i,\ p(u) \neq \perp \text{ and } v = f_{i,p(u)}(u), \\ i & \text{otherwise.} \end{cases}$$

Now, for each player $i \in N$ we can define the equilibrium strategy g_i by setting

$$g_i(v) = \begin{cases} f_i(v) & \text{if } p(v) = \perp \text{ or } p(v) = i, \\ f_{i,p(v)}(v) & \text{otherwise} \end{cases}$$

for each $v \in V$.

It remains to show that $(g_i)_{i\in N}$ is a subgame perfect equilibrium of $(\mathcal{G}, v_0)$. First note that f_i is winning in $(\mathcal{G}_i^\gamma|_h, v)$ if f_i^α is winning in $(\mathcal{G}_i^\alpha|_h, v)$ for some ordinal α because if f_i^α is winning in $(\mathcal{G}_i^\alpha|_h, v)$ every play of $(\mathcal{G}_i^{\alpha+1}|_h, v)$ is consistent with f_i^α and therefore won by player i. As $E^\gamma \subseteq E^{\alpha+1}$, this also holds for every play of $(\mathcal{G}_i^\gamma|_h, v)$. Now let v be any vertex of $\mathcal{G}$ with h the unique maximal history of $\mathcal{G}$ leading to v. We claim that $(g_j)_{j\in N}$ is a Nash equilibrium of $(\mathcal{G}|_h, v)$. Towards this, let g' be any strategy of any player $i \in N$ in $\mathcal{G}$; let $\pi = \langle (g_j)_{j\in N}\rangle_v$, and let $\pi' = \langle g', (g_j)_{j\in N\setminus\{i\}}\rangle_v$. We show that $h\pi$ is won by player i or that $h\pi'$ is not won by player i. The claim is trivial if $\pi = \pi'$. Thus assume that $\pi \neq \pi'$ and fix the least $k < \omega$ such that $\pi(k+1) \neq \pi'(k+1)$. Clearly, $\pi(k) \in V_i$ and $g'(\pi(k)) \neq g_i(\pi(k))$. Without loss of generality, let $k = 0$. We distinguish the following two cases:

- f_i is winning in $(\mathcal{G}_i^\gamma|_h, v)$. By the definition of each g_j, π is a play of $(\mathcal{G}_i^\gamma|_h, v)$. We claim that π is consistent with f_i, which implies that $h\pi$ is won by player i. Otherwise fix the least $l < \omega$ such that $\pi(l) \in V_i$ and $f_i(\pi(l)) \neq \pi(l+1)$. As f_i is winning in $(\mathcal{G}_i^\gamma|_h, v)$, f_i is also winning in $(\mathcal{G}_i^\gamma|_{h\pi(0)\dots\pi(l-1)}, \pi(l))$. But then $(\pi(l), \pi(l+1)) \in E^\gamma \setminus E^{\gamma+1}$, a contradiction to $E^\gamma = E^{\gamma+1}$.
- f_i is not winning in $(\mathcal{G}_i^\gamma|_h, v)$. Hence f_{-i} is winning in $(\mathcal{G}_i^\gamma|_h, v)$. As $g'(v) \neq g_i(v)$, player i has deviated, and it is the case that

$\pi' = \langle g', (f_{j,i})_{j \in N \setminus \{i\}} \rangle_v$. We claim that π' is a play of $(\mathcal{G}_i^\gamma|_h, v)$. As f_{-i} is winning in $(\mathcal{G}_i^\gamma|_h, v)$, this implies that $h\pi'$ is not won by player i. Otherwise fix the least $l < \omega$ such that $(\pi'(l), \pi'(l+1)) \notin E^\gamma$ together with the ordinal α such that $(\pi'(l), \pi'(l+1)) \in E^\alpha \setminus E^{\alpha+1}$. Clearly, $\pi'(l) \in V_i$. Thus f_i^α is winning in $(\mathcal{G}_i^\alpha|_{h\pi'(0)\ldots\pi'(l-1)}, \pi'(l))$, which implies that f_i is winning in $(\mathcal{G}_i^\gamma|_{h\pi'(0)\ldots\pi'(l-1)}, \pi'(l))$. As π' is consistent with f_{-i}, this means that f_{-i} is not winning in $(\mathcal{G}_i^\gamma|_h, v)$, a contradiction.

It follows that $(g_j)_{j \in N} = (g_j|_h)_{j \in N}$ is a Nash equilibrium of $(\mathcal{G}|_h, v)$ for every history hv of $(\mathcal{G}, v_0)$, hence $(g_j)_{j \in N}$ is a subgame perfect equilibrium of $(\mathcal{G}, v_0)$. Q.E.D.

Corollary 4.43. Every multiplayer game with Borel winning conditions has a subgame perfect equilibrium.

O course this also implies that every multiplayer game with Borel winning conditions has a Nash equilibrium. Indeed, for the existence of Nash equilibria, a slightly weaker condition suffices. Let $\mathrm{Two}(\mathcal{G})_{\mathrm{Nash}}$ be defined in the same way as $\mathrm{Two}(\mathcal{G})$ but without closure under subgraphs.

Corollary 4.44. If every game in $\mathrm{Two}(\mathcal{G})_{\mathrm{Nash}}$ is determined, then $\mathcal{G}$ has a Nash equilibrium.