THE LOGIC IN COMPUTER SCIENCE COLUMN

BY

YURI GUREVICH

Computer Science and Engineering
University of Michigan, Ann Arbor, MI 48109, USA
gurevich@umich.edu

A TUTORIAL FOR COMPUTER SCIENTISTS ON FINITE EXTENSIVE GAMES WITH PERFECT INFORMATION

Krzysztof R. Apt
CWI, Amsterdam, The Netherlands
MIMUW, University of Warsaw, Poland
apt@cwi.nl

Sunil Simon
Department of CSE, IIT Kanpur, India
simon@cse.iitk.ac.in

Abstract

We provide a self-contained introduction to finite extensive games with perfect information. In these games players proceed in turns having, at each stage, finitely many moves to their disposal, each play always ends, and in each play the players have complete knowledge of the previously made moves. Almost all discussed results are well-known, but often they are not presented in an optimal form. Also, they usually appear in the literature aimed at economists or mathematicians, so the algorithmic or logical aspects are underrepresented.

Contents

1 Introduction

2 Preliminaries on strategic games

3 Preliminaries on strictly competitive games

4 Extensive games

5 Subgame perfect equilibria

5.1 Definition and examples . . . .
5.2 Backward induction . . . .
5.3 Special classes of extensive games . . . .

6 Backward induction and common knowledge of rationality

7 Weak dominance and backward induction

8 Weak dominance and strictly competitive games

9 Weak acyclicity

10 Win or lose and chess-like games

11 Conclusions

1 Introduction

In computer science for a long time the most commonly studied games have been infinite two-player games (see, e.g., [1] for an account of some of the most popular classes). With the advent of algorithmic game theory various classes of games studied by economists became subject of interest of computer scientists, as well. These games usually involve an arbitrary finite number of players. Among them one the most common ones are strategic games, in which the players select their strategies simultaneously. They have been covered in several books and surveys.

However, in our view a systematic account of another most popular class of games, extensive games with perfect information, is missing. It is true that they are extensively discussed in several books, mostly written for theoretical economists. However, in the introductory texts technical results are usually omitted and illustrated by examples (e.g., [5]). In turn, in the advanced textbooks the presentation is often difficult to follow since these games are introduced as a special case of the extensive games with imperfect information, which leads to involved notation (e.g., [9]). An exception is [16] which devotes a separate chapter to extensive games with perfect information.

From the point of view of computer science the main results are usually not presented in an optimal way. For example, the backward induction is often introduced in a verbose way, or formulated in a way that hides its algorithmic aspect. This way the optimal result that relates it to the set of *all* subgame perfect equilibria (Theorem 7) is often missed. We explain here that it is actually a nondeterministic algorithm that can be even presented as a parallel algorithm. We also discuss an important article [3] that formalizes common knowledge of rationality in finite extensive games with perfect information and relates it to backward induction.

Further, Theorem 17 that clarifies the relation between backward induction and the iterated elimination of weakly dominated strategies, is either only illustrated by an example (e.g., [5]) or only a sketch of a proof is provided (e.g., [16]).

Also, some more recent results, Theorem 19, due to [7], and Theorem 20, due to [12], merit in our opinion some attention among computer scientists. Finally, the so-called Zermelo theorem about chess-like games is in our view often proved in a too elaborate way, starting with the exposition in the classic [22].

These considerations motivated us to write a tutorial presentation of finite extensive games with perfect information aimed at computer scientists. Often these games are discussed by introducing strategic games first. We follow this approach, as well, as it allows us to view extensive games as a subclass of strategic games for which some additional notions can be defined and for which additional results hold.

In our presentation we shall often refer to the account given in [16] that comes closest to our ideal. We shall strengthen some of their results, provide more detailed proofs of some of them, and add some new results. Also, we shall recall their natural notion of a *reduced strategy* that in our view merits more attention.

2 Preliminaries on strategic games

To discuss extensive games it is convenient to introduce first strategic games. A ***strategic game*** for $n \geq 1$ players consists of:

- a set of players $\{1, \ldots, n\}$,

and for each player i

- a non-empty (possibly infinite) set S_i of ***strategies***,
- a ***payoff function*** $p_i : S_1 \times \cdots \times S_n \rightarrow \mathbb{R}$.

We denote it by $(S_1, \ldots, S_n, p_1, \ldots, p_n)$. So the set of players is implicit in this notation.

Strategic games are studied under the following basic assumptions:

- players choose their strategies *simultaneously*; subsequently each player receives a payoff from the resulting joint strategy,
- each player is ***rational***, which means that his objective is to maximize his payoff,
- players have ***common knowledge*** of the game and of each others' rationality.[1]

[1]Intuitively, common knowledge of some fact means that everybody knows it, everybody knows that everybody knows it, etc. It is discussed in the context of extensive games in Section 6.

Finite two-player games are usually represented in the form called a ***bimatrix***, where one assumes that the players choose independently a row or a column. Each entry represents the resulting payoffs to the row and column players. The following examples of two-player games will be relevant in the subsequent discussion.

Example 1. The following game is called ***Prisoner's Dilemma***. The strategies C and D stand for 'cooperate' and 'defect':

	C	D
C	$2, 2$	$0, 3$
D	$3, 0$	$1, 1$

The following game is called ***Matching Pennies***. The strategies H and T stand for 'head' and 'tail':

	H	T
H	$1, -1$	$-1, 1$
T	$-1, 1$	$1, -1$

□

Fix a strategic game $H := (S_1, \ldots, S_n, p_1, \ldots, p_n)$. Let $S = S_1 \times \cdots \times S_n$. We call each element $s \in S$ a ***joint strategy (of players** $1, \ldots, n$**)***, denote the ith element of s by s_i, and abbreviate the sequence $(s_j)_{j \neq i}$ to s_{-i}. We write (s'_i, s_{-i}) to denote the joint strategy in which player's i strategy is s'_i and each other player's j strategy is s_j. Occasionally we write (s_i, s_{-i}) instead of s. Finally, we abbreviate the Cartesian product $\times_{j \neq i} S_j$ to S_{-i}.

Given a joint strategy s, we denote the sequence $(p_1(s), \ldots, p_n(s))$ by $p(s)$ and call it an ***outcome*** of the game. We say that H ***has*** k ***outcomes*** if $|\{p(s) \mid s \in S\}| = k$ and call a game ***trivial*** if it has one outcome. We say that two joint strategies s and t are ***payoff equivalent*** if $p(s) = p(t)$.

We call a strategy s_i of player i a ***best response*** to a joint strategy s_{-i} of the other players if

$$\forall s'_i \in S_i : p_i(s_i, s_{-i}) \geq p_i(s'_i, s_{-i}).$$

Next, we call a joint strategy s a ***(pure) Nash equilibrium*** if for each player i, s_i is a best response to s_{-i}, that is, if

$$\forall i \in \{1, \ldots, n\} : \forall s'_i \in S_i \; p_i(s_i, s_{-i}) \geq p_i(s'_i, s_{-i}).$$

So a joint strategy is a Nash equilibrium if no player can achieve a higher payoff by *unilaterally* switching to another strategy. Intuitively, a Nash equilibrium is a situation in which each player is a posteriori satisfied with his choice. It is often used to predict the outcomes of the strategic games.

It is easy to check that there is a unique Nash equilibrium in Prisoner's Dilemma which is (D, D), while the Matching Pennies game has no Nash equilibria.

A relevant question is whether we can identify natural subclasses of strategic games where a Nash equilibrium is guaranteed to exist. Below, we describe two such classes which are well studied in game theory. Fix till the end of this section a strategic game $H = (S_1, \ldots, S_n, p_1, \ldots, p_n)$.

We say that a pair of joint strategies (s, s') forms a ***profitable deviation*** if there exists a player i such that $s'_{-i} = s_{-i}$ and $p_i(s') > p_i(s)$. If such a pair (s, s') exists, then we say that player i can ***profitably deviate*** from s to s' and denote this by $s \rightarrow s'$. An ***improvement path*** is a maximal sequence (i.e., a sequence that cannot be extended to the right) of joint strategies in which each consecutive pair is a profitable deviation. By an ***improvement sequence*** we mean a prefix of an improvement path.

We say that H has the ***finite improvement property*** (***FIP*** in short), if every improvement path is finite. Clearly, if an improvement path is finite, then its last element is a Nash equilibrium. So if H has the FIP, then it is guaranteed to have Nash equilibrium, which explains the interest in this notion. A trivial example of a game that has the FIP is the Prisoner's Dilemma game.[2]

However, the FIP is a very strong property and several natural games with a Nash equilibrium fail to satisfy it. Young in [23] and independently Milchtaich in [14] proposed a weakening of this condition and introduced the following natural class of games. We say that H is ***weakly acyclic*** if for any joint strategy s, there exists a finite improvement path that starts at s. Consequently, every weakly acyclic game has a Nash equilibrium.

We call the function $P : S \rightarrow \mathbb{R}$ a ***weak potential*** for H if

$\forall s$: if s is not a Nash equilibrium, then for some
profitable deviation $s \rightarrow s'$, $P(s) < P(s')$.

The following natural characterization of finite weakly acyclic games was established in [15].

Theorem 1 (Weakly acyclic)**.** *A finite game is weakly acyclic iff it has a weak potential.*

Sometimes it is convenient to assume that a weak potential is a function to a strict linear ordering (that can be subsequently mapped to $\mathbb{R}$).

One way to find a Nash equilibrium in a strategic game is by using a concept of dominance. In the context of extensive games the most relevant is the notion of weak dominance. By a ***subgame*** of a strategic game H we mean a game obtained from H by removing some strategies.

[2] A more interesting class of games that have the FIP are the congestion games, see [18].

Consider two strategies s_i and s'_i of player i. We say that s_i ***weakly dominates*** s'_i (or equivalently, that s'_i is ***weakly dominated by*** s_i) in H if

$$\forall s_{-i} \in S_{-i} : p_i(s_i, s_{-i}) \geq p_i(s'_i, s_{-i}) \text{ and } \exists s_{-i} \in S_{-i} : p_i(s_i, s_{-i}) > p_i(s'_i, s_{-i}).$$

Denote by H^1 a subgame of H obtained by the elimination of some (not necessarily all) strategies that are weakly dominated in H, and put $H^0 := H$ and $H^{k+1} := (H^k)^1$, where $k \geq 1$. Note that H^k is not uniquely defined, since we do not stipulate which strategies are removed at each stage.

Abbreviate the phrase 'iterated elimination of weakly dominated strategies' to IEWDS. We say that each H^k is obtained from H ***by an IEWDS***. If for some k, some subgame H^k is a trivial game we say that H ***can be solved by an IEWDS***. The relevant result (that we shall not prove) is that if a finite strategic game H can be solved by an IEWDS then every remaining joint strategy is a Nash equilibrium of H. We shall illustrate it in Example 7 in Section 4.

The following lemma will be needed in Section 7.

Lemma 2. *Let $H := (S_1, \ldots, S_n, p_1, \ldots, p_n)$ be a finite strategic game and let $H^k := (S_1^k, \ldots, S_n^k, p_1, \ldots, p_n)$, where $k \geq 1$. Then*

$$\forall i \in \{1, \ldots, n\}\ \forall s_i \in S_i\ \exists t_i \in S_i^k\ \forall s_{-i} \in S_{-i}^k : p_i(t_i, s_{-i}) \geq p_i(s_i, s_{-i}).$$

Proof. We proceed by induction. Take some $i \in \{1, \ldots, n\}$ and $s_i \in S_i$. Suppose $k = 1$. If $s_i \in S_i^1$, then we are done, so assume that $s_i \notin S_i^1$. H is finite and the relation 'weakly dominates' is transitive, so some strategy t_i from H weakly dominates s_i in H and is not weakly dominated in H, and thus is in H^1.

Suppose the claim holds for some $k > 1$. By the induction hypothesis for some $u_i \in S_i^k$ we have $p_i(u_i, s_{-i}) \geq p_i(s_i, s_{-i})$ for all $s_{-i} \in S_{-i}^k$. If $u_i \notin S_i^{k+1}$, then for the same reasons as above some strategy t_i from H^{k+1} weakly dominates u_i in H^k and consequently $p_i(t_i, s_{-i}) \geq p_i(s_i, s_{-i})$ for all $s_{-i} \in S_{-i}^k$. □

Finally, we introduce the following condition defined in [13]. We say that a strategic game $(S_1, \ldots, S_n, p_1, \ldots, p_n)$ satisfies the ***transference of decisionmaker indifference (TDI)*** condition if:

$$\begin{aligned} &\forall i \in \{1, \ldots, n\}\ \forall r_i, t_i \in S_i\ \forall s_{-i} \in S_{-i} : \\ &p_i(r_i, s_{-i}) = p_i(t_i, s_{-i}) \rightarrow p(r_i, s_{-i}) = p(t_i, s_{-i}). \end{aligned}$$

Informally, this condition states that whenever for some player i two of his strategies r_i and t_i are indifferent w.r.t. some joint strategy s_{-i} of the other players then this indifference extends to all players.

In the next section we shall introduce a natural class of strategic games that satisfy the TDI condition.

3 Preliminaries on strictly competitive games

Sections 8 and 10 concern specific extensive games that involve two players. It is convenient to introduce them first as strategic games. A strategic two-player game is called ***strictly competitive*** if

$$\forall i \in \{1, 2\}\ \forall s, t \in S : p_i(s) \geq p_i(t) \text{ iff } p_{-i}(s) \leq p_{-i}(t).$$

(As there are here just two players, $-i$ denotes the opponent of player i, so $p_{-i}(s)$ and $p_{-i}(t)$ are here numbers.) Note that every strictly competitive game satisfies the TDI condition as the definition implies that

$$\forall i \in \{1, 2\}\ \forall s, t \in S : p_i(s) = p_i(t) \text{ iff } p_{-i}(s) = p_{-i}(t). \tag{1}$$

A two-player game is called ***zero-sum*** if

$$\forall s \in S : p_1(s) + p_2(s) = 0.$$

An example is the Matching Pennies game. Clearly every zero-sum game is strictly competitive.

In Section 10 we shall discuss two classes of zero-sum games. A zero-sum game is called a ***win or lose game*** if the only possible outcomes are $(1, -1)$ and $(-1, 1)$, with 1 interpreted as a *win* and -1 as *losing*. Finally, a zero-sum game is called a ***chess-like game*** if the only possible outcomes are $(1, -1)$, $(0, 0)$, and $(-1, 1)$, with 0 interpreted as a *draw*.

The following results about strictly competitive games will be needed in Section 8.

Lemma 3. *Consider a strictly competitive strategic game H with a Nash equilibrium s. Suppose that for some $i \in \{1, 2\}$, t_i weakly dominates s_i. Then (t_i, s_{-i}) is also a Nash equilibrium.*

Proof. Let $H := (S_1, S_2, p_1, p_2)$. Take a strategy s'_i of player i. By the assumptions about s and t_i

$$p_i(t_i, s_{-i}) = p_i(s_i, s_{-i}) \geq p_i(s'_i, s_{-i}).$$

Next, take a strategy s'_{-i} of player $-i$. By (1) and the fact that s is a Nash equilibrium

$$p_{-i}(t_i, s_{-i}) = p_{-i}(s_i, s_{-i}) \geq p_{-i}(s_i, s'_{-i}).$$

This establishes the claim. □

Given a finite strategic game $H := (S_1, \ldots, S_n, p_1, \ldots, p_n)$ we define for each player i

$$maxmin_i(H) := \max_{s_i \in S_i} \min_{s_{-i} \in S_{-i}} p_i(s_i, s_{-i}).$$

We call any strategy s_i^* such that $\min_{s_{-i} \in S_{-i}} p_i(s_i^*, s_{-i}) = maxmin_i(H)$ a ***security strategy*** for player i in H.

The following result goes back to [22], where it was established for zero-sum games. The formulation below is from [16, pages 22-23]).

Theorem 4 (Minimax)**.** *Suppose that s is a Nash equilibrium of a strictly competitive strategic game H. Then each s_i is a security strategy for player i and $p(s) = (maxmin_1(H), maxmin_2(H))$.*

Corollary 5. *Consider a finite strictly competitive strategic game H that has a Nash equilibrium. Then H^1 has a Nash equilibrium, as well, and for all $i \in \{1, 2\}$, $maxmin_i(H) = maxmin_i(H^1)$.*

(The notation H^1 was introduced in Section 2.)

Proof. We first prove that some Nash equilibrium of H is also a joint strategy of H^1. Let s be a Nash equilibrium of H. Suppose first that only one strategy from s, say s_i, is not a strategy in H^1. The game H is finite and the relation 'weakly dominates' is transitive so some strategy t_i weakly dominates s_i and is not weakly dominated. Thus (t_i, s_{-i}) is a joint strategy in H^1, which by Lemma 3 is a Nash equilibrium in H.

Suppose now that none of the strategies from s are strategies in H^1. By the argument just made we conclude that for some joint strategy t in H^1 first (t_i, s_{-i}) is a Nash equilibrium in H and then that t is a Nash equilibrium in H.

We conclude that a joint strategy is both a Nash equilibrium in H and in H^1. The other claim then follows by the Minimax Theorem 4. □

Lemma 6. *Consider a strictly competitive strategic game H that has a Nash equilibrium and has two outcomes. Let H^1 be the result of removing from H all weakly dominated strategies. Then H^1 is a trivial game.*

Proof. Let s^* be a Nash equilibrium of $H = (S_1, S_2, p_1, p_2)$ and s' a joint strategy such that $p(s^*)$ and $p(s')$ are the two outcomes in H. By condition (1) from Section 2 $p_1(s^*) \neq p_1(s')$ and $p_2(s^*) \neq p_2(s')$. H is strictly competitive, so for some i both $p_i(s^*) > p_i(s')$ and $p_{-i}(s') > p_{-i}(s^*)$.

First we show that $p_i(s_i^*, s_{-i}) = p_i(s^*)$ for all $s_{-i} \in S_{-i}$. Suppose otherwise. Take s_{-i} such that $p_i(s_i^*, s_{-i}) \neq p_i(s^*)$. Then $p_i(s_i^*, s_{-i}) = p_i(s')$, so by (1) $p_{-i}(s_i^*, s_{-i}) = p_{-i}(s') > p_{-i}(s^*)$, which contradicts the fact that s^* is a Nash equilibrium.

Hence by the choice of i for all $s_{-i} \in S_{-i}$

$$p_i(s_i^*, s_{-i}) = p_i(s^*) \geq p_i(s_i', s_{-i})$$

and

$$p_i(s_i^*, s_{-i}') = p_i(s^*) > p_i(s').$$

So s_i^* weakly dominates s_i'. This implies that H^1 is a trivial game. □

4 Extensive games

After these preliminaries we now focus on the subject of this tutorial.

A ***rooted tree*** (from now on, just a ***tree***) is a connected directed graph (i.e., such that the undirected version is connected) with a unique node with the in-degree 0, called the ***root***, and in which every other node has the in-degree 1. A ***leaf*** is a node with the out-degree 0. We denote a tree by (V, E, v_0), where V is a non-empty set of nodes, E is a possibly empty set of directed edges, and v_0 is the root. In drawings the edges will be directed downwards.

An ***extensive game with perfect information*** (in short, just an ***extensive game***) for $n \geq 1$ players consists of:

- a set of players $\{1, \ldots, n\}$,
- a ***game tree*** $T := (V, E, v_0)$; we denote its set of leaves by Z,
- a ***turn function*** $turn : V \setminus Z \to \{1, \ldots, n\}$,
- the ***outcome functions*** $o_i : Z \to \mathbb{R}$, for each player i.

We denote it by $(T, turn, o_1, \ldots, o_n)$.

As in the case of strategic games we assume that each player is rational (which now means that his objective is to maximize his outcome in the game) and that the players have common knowledge of the game and of each others' rationality.

A node w is called a ***child*** of v in T if $(v, w) \in E$. A node in T is called a ***preleaf*** if all its children are leaves. We say that an extensive game is ***finite*** if its game tree is finite. In what follows we limit our attention to finite extensive games.

The function *turn* determines at each non-leaf node which player should move. The edges of T represent possible ***moves*** in the considered game, while for a node $v \in V \setminus Z$ the set of its children $C(v) := \{w \mid (v, w) \in E\}$ represents possible ***actions*** of player $turn(v)$ at v.

In the figures below we identify the actions with the labels we put on the edges and thus identify each action with the corresponding move. For convenience we do not assume the labels to be unique, but it will not lead to confusion. Further, we annotate the non-leaf nodes with the identity of the player whose turn it is to move and the name of the node. Finally, we annotate each leaf node with the corresponding sequence of the values of the o_i functions.

Example 2. Consider the Prisoner's Dilemma and Matching Pennies games from Example 1. Suppose the players move sequentially with the row player moving first. The game trees of the resulting extensive games are depicted in Figures 1 and 2 below. The thick lines in the second drawing will be explained later. □

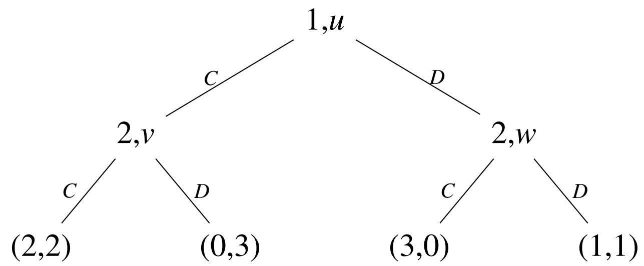

Figure 1: Extensive form of the Prisoner's Dilemma game

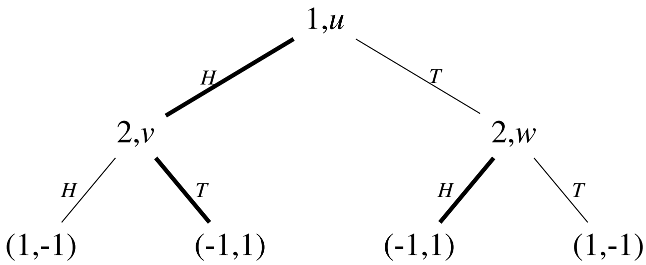

Figure 2: Extensive form of the Matching Pennies game

Example 3. The following two-player game is called the ***Ultimatum game***. Player 1 moves first and selects a number $x \in \{0, 1, \ldots, 100\}$ intepreted as a percentage of some good to be shared, leaving the fraction of $(100 - x)\%$ for the other player. Player 2 either *accepts* this decision, the outcome is then $(x, 100 - x)$, or *rejects* it, the outcome is then $(0, 0)$. The game tree is depicted in Figure 3, where the action of player 1 is a number from the set $\{0, 1, \ldots, 100\}$, and the actions of player 2 are A and R.

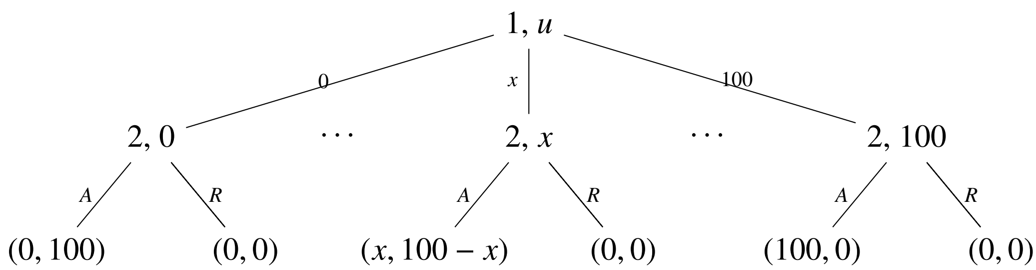

Figure 3: The Ultimatum game

□

Next we introduce strategies in extensive games. Consider a finite extensive game $G := (T, turn, o_1, \ldots, o_n)$. Let $V_i := \{v \in V \setminus Z \mid turn(v) = i\}$. So V_i is the set of nodes at which player i moves. Its elements are called the ***decision nodes*** of player i. A ***strategy*** for player i is a function $s_i : V_i \rightarrow V$, such that $(v, s_i(v)) \in E$ for all $v \in V_i$. ***Joint strategies*** are defined as in strategic games. When the game tree consists of just one node, each strategy is the empty function, denoted by $\emptyset$, and there is only one joint strategy, namely the n-tuple of these functions.

Each joint strategy different from $(\emptyset, \ldots, \emptyset)$ assigns a unique child to every node in $V \setminus Z$. In fact, we can identify joint strategies with such assignments. Each joint strategy $s = (s_1, \ldots, s_n)$ determines a rooted path $play(s) := (v_1, \ldots, v_m)$ in T defined inductively as follows:

- v_1 is the root of T,
- if $v_k \notin Z$, then $v_{k+1} := s_i(v_k)$, where $turn(v_k) = i$.

Informally, given a joint strategy s, we can view $play(s)$ as the resulting *play* of the game.

G is finite, so for each joint strategy s the rooted path $play(s)$ is finite. Denote by $leaf(s)$ the last element of $play(s)$. We call then $(o_1(leaf(s)), \ldots, o_n(leaf(s)))$ the ***outcome*** of the game G when each player i pursues his strategy s_i and abbreviate it to $o(leaf(s))$.

Example 4. Let us return to the extensive form of the Matching Pennies game from Example 2. The strategies for player 1 are: H and T, while the strategies for player 2 are: HH, HT, TH, and TT, where for instance TH stands for a strategy that selects T at the node v and H at the node w. In Figure 2 thick lines correspond with the joint strategy (H, TH). Here $play(H, TH) = (u, v, (-1, 1))$, where we identify each leaf with the corresponding outcome, and $o(leaf(H, TH)) = (-1, 1)$. □

With each finite extensive game $G := (T, turn, o_1, \ldots, o_n)$ we associate a strategic game $\Gamma(G) := (S_1, \ldots, S_n, p_1, \ldots, p_n)$ defined as follows:

- S_i is the set of strategies of player i in G,
- $p_i(s) := o_i(leaf(s))$.

We call $\Gamma(G)$ the ***strategic form*** of G.

All notions introduced in the context of strategic games can now be reused in the context of finite extensive games simply by referring to the corresponding strategic form. This way we obtain the notions of a best response, Nash equilibrium, extensive games that have the FIP, are weakly acyclic, etc.

Example 5. The strategic form of the extensive form of the Matching Pennies game from Example 2 differs from the initial Matching Pennies game from Example 1 and looks as follows:

	HH	HT	TH	TT
H	$1, -1$	$1, -1$	$-1, \ 1$	$-1, \ 1$
T	$-1, \ 1$	$1, -1$	$-1, \ 1$	$1, -1$

Note that this game has two Nash equilibria: (H, TH) and (T, TH). The first one is depicted in Figure 2 by thick lines. By definition, these are the Nash equilibria of the extensive form of the Matching Pennies game.

The intuitive reason that there are two Nash equilibria is that no matter which out of his two strategies the first player selects, the second player can always secure the payoff 1 for himself. Of course, the decision which player moves first affects both the sets of strategies and the sets of Nash equilibria.

One can also easily check that the extensive form of the Prisoner's Dilemma game given in Figure 1 has one Nash equilibrium, (D, DD). □

Example 6. Consider now the Ultimatum game from Example 3. Each strategy for player 1 is a number from $\{0, 1, \ldots, 100\}$, while each strategy for player 2 is a function from $\{0, 1, \ldots, 100\}$ to $\{A, R\}$.

It is easy to check that $(100, s_2)$, where for $y \in \{0, 1, \ldots, 100\}$, $s_2(y) = R$ is a Nash equilibrium with the outcome $(0, 0)$ and that all other Nash equilibria are of the form (x, s_2), where $s_2(x) = A$ and $s_2(y) = R$ for $y > x$, and $x, y \in \{0, 1, \ldots, 100\}$, with the corresponding outcome $(x, 100 - x)$. □

Concepts such as Nash equilibrium can be defined directly, without a detour through the strategic games. However, introducing strategic games first allows us to view finite extensive games as a special class of strategic games and allows us to conclude that some results established for the strategic games, for instance the Weakly Acyclic Theorem 1, also hold for all finite extensive games.

As we shall see, for extensive games, due to their structure, additional results hold. Further, their structure suggests a new equilibrium notion that is meaningful only for these games. Finally, when discussing iterated elimination of weakly dominated strategies in an extensive game, one needs to reason about its strategic form.

All examples in this section are instances of the so-called ***Stackelberg competition***. In such games a *leader* moves first and a *follower*, having observed the resulting action, moves second. Of course, there are other natural extensive games, in particular multi-player games and games with infinite game trees.

For extensive games that are not Stackelberg competition games it is legitimate to question the notion of a strategy. Namely, one would expect that when a player following a strategy makes a move to a node u, then all his subsequent moves should take place in the subtree rooted at u. However, the definition of a strategy does not stipulate it as it is 'overdefined'. A natural revision was introduced in [16, page 94].

Given a node w in the game tree consider the path $v = v_1, \ldots, v_k = w$ to it from the root v. Let

$$[w]_i := \{(v_j, v_{j+1}) \mid j \in \{1, \ldots, k-1\} \text{ and } turn(v_j) = i\}.$$

Informally, $[w]_i$ is the set of the moves of player i that lie on the path from the root to w.

We call rs_i a ***reduced strategy*** of player i if it is a maximal subset of a strategy s_i (recall that each strategy is a function, so a set of pairs of nodes) that satisfies the following property:

$$\text{if } (u, w) \in rs_i, \text{ then } [w]_i \subseteq rs_i.$$

Intuitively, rs_i is a reduced strategy of player i if according to it each of his moves follows his earlier moves in the game.

Just like the joint strategies, each joint reduced strategy determines a play of the game. This allows us to define the outcome of the game when each player pursues his reduced strategy. Associate now with each joint reduced strategy r the following set of original joint strategies:

$$Str(r) := \{s \mid \forall i \in \{1, \ldots n\} : r_i \subseteq s_i\}.$$

One can show that the sets $Str(r)$, where r is a Nash equilibrium in the reduced strategies, form a partition of the set of original Nash equilibria. So each Nash equilibrium in the reduced strategies is a convenient representation of a set of Nash equilibria in the original strategies.

Example 7. Consider the classic centipede game due to [17]. In Figure 4 we present a version of the game from [16, page 107]. C and S represent the actions 'continue' and 'stop'.

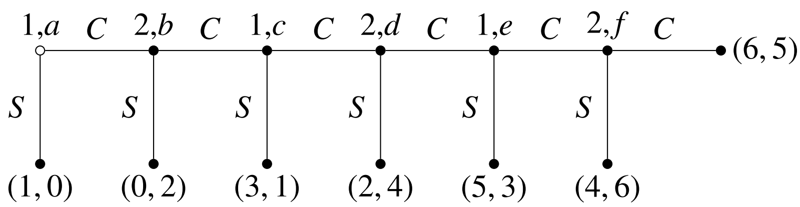

Figure 4: A six-period version of the centipede game

Each player has three decision nodes and two actions at each of them. So each player has eight strategies. In contrast, both players have four reduced strategies. These are for player 1: aS, $aCcS$, $aCcCeS$, $aCcCeC$, and

for player 2: bS, $bCdS$, $bCdCfS$, $bCdCfC$,

where $aCcCeS$, is a shorthand for $\{(a, C), (c, C), (e, S)\}$, etc. The strategic form corresponding to the joint reduced strategies is given in Figure 5.

	bS	$bCdS$	$bCdCfS$	$bCdCfC$
aS	$1, 0$	$1, 0$	$1, 0$	$1, 0$
$aCcS$	$0, 2$	$3, 1$	$3, 1$	$3, 1$
$aCcCeS$	$0, 2$	$2, 4$	$5, 3$	$5, 3$
$aCcCeC$	$0, 2$	$2, 4$	$4, 6$	$6, 5$

Figure 5: The strategic form of the centipede game that uses reduced strategies

This game has a unique Nash equilibrium, namely (aS, bS), with the outcome $(1, 0)$. It can be obtained by solving the game by an IEWDS in six steps, by repeatedly eliminating the rightmost column and the lowest row. In contrast, the strategic form corresponding to the original extensive game has several Nash equilibria. According to the notation introduced above, these Nash equilibria form the set $Str((aS, bS))$. □

5 Subgame perfect equilibria

5.1 Definition and examples

Example 6 suggests that the concept of a Nash equilibrium is not informative enough to predict outcomes of extensive games: it results in too many scenarios, some of which are obviously inferior to all players. Another issue is the problem of so-called 'not credible threat' as illustrated in the following example.

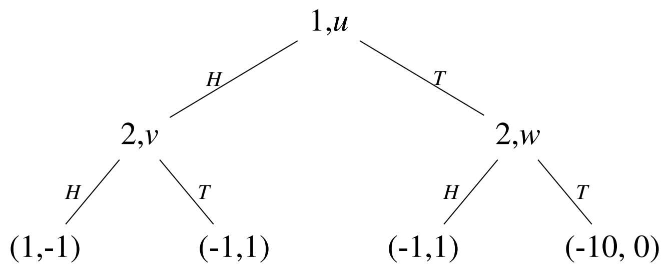

Figure 6: A modification of the Matching Pennies game

Example 8. Consider the extensive game given in Figure 6. This game has three Nash equilibria: (H, TH), (T, TH), and (H, TT). However, the last equilibrium is not plausible: if player 1 chooses T, then player 2 should select H and not T: the 'threat' of player 2 to choose T at the node 2 is not credible. □

Motivated by the issue of non-credible threats Selten introduced in [20] a stronger equilibrium concept. To define it we need to use strategies instead of the restricted strategies.

Consider an extensive game $G := (T, turn, o_1, \ldots, o_n)$ and a non-leaf node w of T. Denote by T^w the subtree of T rooted at w. The ***subgame of*** G ***rooted at the node*** w, denoted by G^w, is defined as follows:

- its set of players is $\{1, \ldots, n\}$,
- its tree is T^w,
- its turn and payoff functions are the restrictions of the corresponding functions of G to the nodes of T^w.

So the notion of a subgame has a different meaning for the strategic and for the extensive games. Note that some players may 'drop out' in G^w, in the sense that at no node of T^w it is their turn to move. Still, to keep the notation simple, it is convenient to admit in G^w all original players in G. Each strategy s_i of player i in G uniquely determines his strategy s_i^w in G^w. Given a joint strategy $s = (s_1, \ldots, s_n)$ of G we denote by s^w the joint strategy $(s_1^w, \ldots, s_n^w)$ in G^w.

A joint strategy s of G is called a ***subgame perfect equilibrium*** in G if for every node w of T, the joint strategy s^w of G^w is a Nash equilibrium in G^w. Informally s is subgame perfect equilibrium in G if it induces a Nash equilibrium in every subgame of G.

It is straightforward to check that all Nash equilibria in the extensive forms of the Prisoner's Dilemma and Matching Pennies games are also subgame perfect equilibria. The modified Matching Pennies game given in Example 8 has two subgame perfect equilibria: (H, TH) and (T, TH). Note that the Nash equilibrium (H, TT) that involves a non-credible threat is not a subgame perfect equilibrium.

Example 9. Return now to the Ultimatum game from Example 3. In Example 6 we noticed that this game has several Nash equilibria. However, only two of them are subgame perfect equilibria. They are depicted in Figures 7 and 8. In the first equilibrium player 1 selects 100 and player 2 accepts all offers, while in the second equilibrium player 1 selects 99 and player 2 accepts all offers except 100. These equilibria are more intuitive than the remaining Nash equilibria and provide natural insights into this game. □

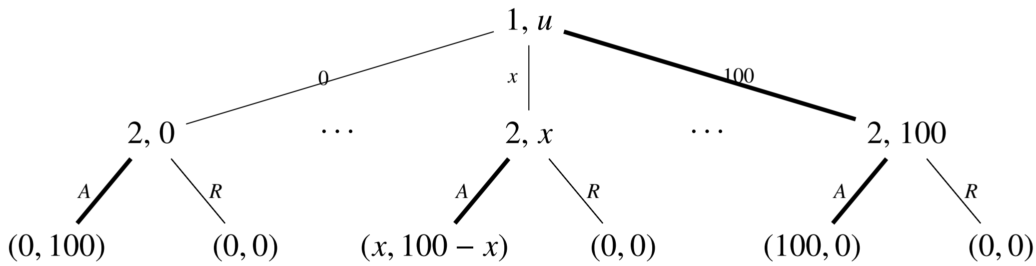

Figure 7: A subgame perfect equilibrium in the Ultimatum game

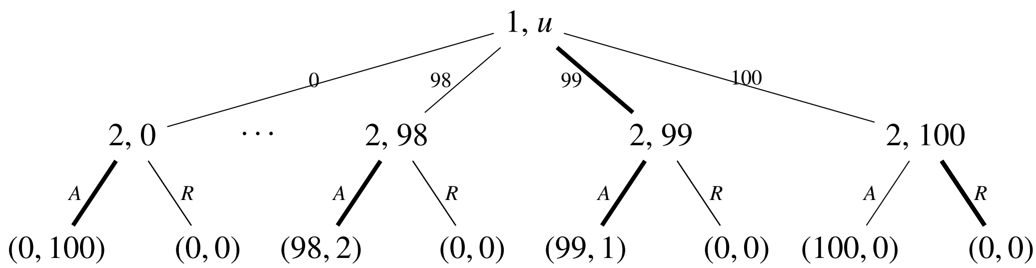

Figure 8: Another subgame perfect equilibrium in the Ultimatum game

It should be noted that the Ultimatum game has been extensively analysed in experimental economics. It has been observed that in practice, people do not often play a Nash equilibrium or a subgame perfect equilibrium.

5.2 Backward induction

We now show that for a finite extensive game G over T, one can construct a subgame perfect equilibrium using the iterative procedure described in Algorithm 1 called the ***backward induction algorithm***. Since each loop iteration of Algorithm 1 modifies the underlying game tree, we use the notation $C(v, T)$ to denote the set of children of node v in the current version of the tree T.

Note that Algorithm 1 always terminates but in general need not have a unique outcome due to the presence of the **choose** statements. Each execution (i.e., each selection of values in the **choose** statements) constructs a unique joint strategy s and an extension of the functions $o_1, \ldots, o_n$ to all nodes of the game tree.

Example 10. Consider the Ultimatum game from Example 3. Algorithm 1 generates two possible outputs that correspond to Figures 7 and 8. For the second outcome the corresponding extensions of the outcome functions to all nodes are given in Figure 9. □

Algorithm 1:

Input: A finite extensive game $G := (T, turn, o_1, \ldots, o_n)$ with $T = (V, E, v_0)$

Output: A subgame perfect equilibrium s in G and extensions of the functions $o_1, \ldots, o_n$ to all nodes of T such that $o(v_0) = o(leaf(s))$

```
while |V| > 1 do
    choose v ∈ V that is a preleaf of T;
    i := turn(v);
    choose w ∈ C(v, T) such that o_i(w) is maximal;
    s_i(v) := w;
    o(v) := o(w);
    V := V \ C(v, T);
    E := E ∩ (V × V);
    T := (V, E, v_0)
```

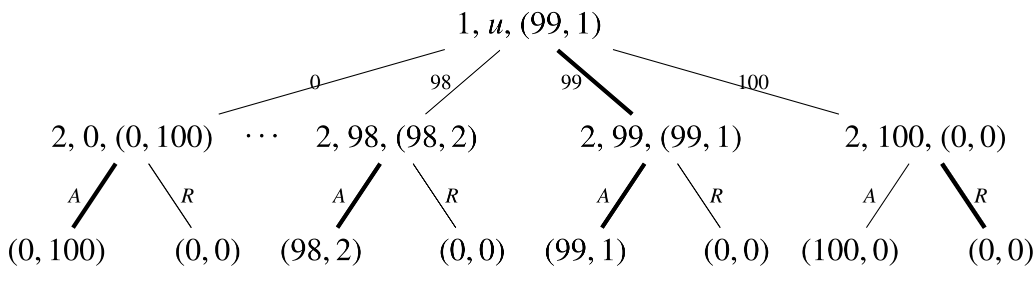

Figure 9: The backward induction algorithm and the Ultimatum game

The following characterisation result makes use of the nondeterminism present in the algorithm.

Theorem 7. *For every finite extensive game all possible executions of the backward induction algorithm generate precisely all subgame perfect equilibria.*

To establish this result we shall need a preparatory lemma, called the 'one deviation property' (see, e.g., [16, page 98]). Recall that for a function $f : X \rightarrow Y$ (with $X \neq \emptyset$), $\text{argmax}_{x \in X} f(x) := \{y \in X \mid f(y) = \max_{x \in X} f(x)\}$.

Lemma 8. *Let G be a finite extensive game over the game tree T. A joint strategy s is a subgame perfect equilibrium in G iff for all non-leaf nodes u in T*

$$s_i(u) \in \text{argmax}_{x \in C(u)} o_i(leaf(s^x)), \textit{ where } i = turn(u).$$

In words, this condition states that for all non-leaf nodes u in T and $i = turn(u)$, $s_i(u)$ selects a child x of u for which $o_i(leaf(s^x))$ is maximal.

For a proof see, e.g, [16, pages 98-99] or a more detailed presentation in the appendix of [2].

Corollary 9 ([2], Corollary 7)**.** *Let G be a finite extensive game over the game tree T with the root v. A joint strategy s is a subgame perfect equilibrium in G iff for all $u \in C(v)$*

- $s_i(v) \in \text{argmax}_{x \in C(v)} o_i(leaf(s^x))$, *where* $i = turn(v)$,
- s^u *is a subgame perfect equilibrium in the subgame* G^u.

Intuitively, the first condition states that among the subgames rooted at the children of the root v, the one determined by the first move in the game G yields the best outcome for the player who moved.

Proof of Theorem 7. The proof proceeds by induction on $height(T)$, defined as the number of edges in the longest path in the tree T. The base case when $height(T) = 1$ is straightforward. Suppose $height(T) > 1$. Let $C(v) = \{w_1, \dots w_k\}$.
($\Rightarrow$) Consider a joint strategy s in G together with some extensions of the functions $o_1, \dots, o_n$ to the nodes of the game tree of G that are generated by an execution of Algorithm 1.

Fix an arbitrary $l \in \{1, \dots, k\}$. Delete in this execution the **while** loop iterations that do not involve the nodes of the game tree T^{w_l} of the subgame G^{w_l} and use at the beginning the game tree T^{w_l} instead of T. This way we obtain an execution of Algorithm 1 applied to the game G^{w_l} that generates a joint strategy s^{w_l} in G^{w_l} together with some extensions of the functions $o_1, \dots, o_n$ to the nodes of the game tree T^{w_l}. By the induction hypothesis s^{w_l} is a subgame perfect equilibrium in G^{w_l} and $o(w_l) = o(leaf(s^{w_l}))$.

Consider now the last iteration of the **while** loop in the original execution of Algorithm 1. At this stage $V = \{v_0, w_1, \ldots, w_k\}$, so before line **2** we have $C(v_0, T) = \{w_1, \ldots, w_k\}$. After line **3** we have $i = turn(v_0)$ and after line **4** w is such that $w \in \{w_1, \ldots, w_k\}$ and $o_i(w) \geq o_i(w_l)$ for all $l \in \{1, \ldots, k\}$.

By the previous conclusion for all $l \in \{1, \ldots, k\}$ we have $o_i(w_l) = o_i(leaf(s^{w_l}))$, so for all $l \in \{1, \ldots, k\}$ we have $o_i(leaf(s^w)) \geq o_i(leaf(s^{w_l}))$. After line **5** we have $s_i(v_0) = w$, so by Lemma 8 s is a subgame perfect equilibrium. Finally, after line **6** we have $o(v_0) = o(leaf(s^w)) = o(leaf(s))$.

($\Leftarrow$) Suppose that s is a subgame perfect equilibrium in G. We show that it can be generated by Algorithm 1 together with the extensions of the functions $o_1, \ldots, o_n$ to all nodes of T such that $o(v_0) = o(leaf(s))$.

Fix an arbitrary $l \in \{1, \ldots, k\}$. By the assumption on s, the joint strategy s^{w_l} is a subgame perfect equilibrium in the subgame G^{w_l}, so by the induction hypothesis some execution of Algorithm 1 applied to the subgame G^{w_l} generates s^{w_l} together with the extensions of the functions $o_1, \ldots, o_n$ to the nodes of the game tree of G^{w_l} such that $o(w_l) = o(leaf(s^{w_l}))$.

Using these k executions of Algorithm 1 applied to the subgames $G^{w_1}, \ldots, G^{w_k}$ we construct the desired execution of Algorithm 1 applied to the game G as follows. First we 'glue' these k executions into one but using at the beginning of the execution the game tree T instead of the game tree of G^{w_1} and using at the beginning of each subsequent execution the current tree T instead of the game tree of the considered subgame.

After these k executions glued together $V = \{v_0, w_1, \ldots, w_k\}$, so before line **2** we have $C(v, T) = \{w_1, \ldots, w_k\}$, in line **2**, v_0 is selected and after line **3** we have $i = turn(v_0)$.

By the induction hypothesis for all $l \in \{1, \ldots, k\}$ we have $o(w_l) = o(leaf(s^{w_l}))$, so by Lemma 8 $w = s_i(v_0)$ is a node from $\{w_1, \ldots, w_k\}$ such that $o_i(w)$ is maximal. So in line **4** we can select this node w, which ensures that the assignment in line **5** is consistent with the original joint strategy s. Further, the assignment in line **6** ensures that $o(v_0) = o(w) = o(leaf(s^w)) = o(leaf(s))$. □

Corollary 10 ([11]). *Every finite extensive game (with perfect information) has a subgame perfect equilibrium (and hence a Nash equilibrium).*

We presented backward induction as a nondeterministic algorithm, but one can go even further. Exploiting the fact that the children of each node can be dealt with independently, we can present it as an algorithm that uses nested parallelism. In such an algorithm there is no need to modify the game tree. Given a non-leaf node v we define a nondeterministic program $Seq(v)$ by

$i := turn(v)$;
choose $w \in C(v)$ such that $o_i(w)$ is maximal;
$s_i(v) := w$;
$o(v) := o(w)$

Then for a preleaf node v we define $Comp(v)$ as $Seq(v)$ and for each node v that is neither a preleaf or a leaf we define $Comp(v)$ by

$$[||_{w \in C(v)} Comp(w)]; Seq(w)$$

where $[||_{w \in C(v)} Comp(w)]$ stands for a parallel composition of the programs $Comp(w)$ for $w \in C(v)$. So each such node v is processed with only after its children have been processed and these children are processed in an arbitrary order.

Then $Comp(v_0)$ is the desired parallel version of the backward induction algorithm.

5.3 Special classes of extensive games

It is natural to study conditions under which an extensive game has a unique subgame perfect equilibrium. The following property was introduced in [4]. We say that an extensive game is ***without relevant ties*** if for all non-leaf nodes u in T the function o_i, where $turn(u) = i$, is injective on the leaves of T^u. This is more general than saying that a game is ***generic***, which means that each o_i is an injective function.

Corollary 11. *Every finite extensive game without relevant ties has a unique subgame perfect equilibrium.*

In particular, every finite generic extensive game has a unique subgame perfect equilibrium.

Proof. If a game is without relevant ties, then so is every subgame of it. This allows us to proceed by induction on the *height* of the game tree. Let G be a finite extensive game without relevant ties over a game tree T. If $height(T) = 0$ the claim clearly holds. Suppose that $height(T) > 0$. Let v be the root of T and $i = turn(v)$.

By the induction hypothesis for each $w \in C(v)$ there is exactly one subgame perfect equilibrium t^w in G^w. Let $t = \times_{w \in C(v)} t^w$. Then for different $w, w' \in C(v)$, $leaf(t^w)$ and $leaf(t^{w'})$ are different leaves of the game tree of G. Since G is without relevant ties, $o_i(leaf(t^w)) \neq o_i(leaf(t^{w'}))$.

This means that the function $g : C(v) \to \mathbb{R}$ defined by $g(w) := o_i(leaf(t^w))$ is injective. Consequently the set $\text{argmax}_{w \in C(v)} o_i(leaf(t^w))$ has a unique element and hence by Corollary 9, G has exactly one subgame perfect equilibrium. □

Note that the centipede game from Example 7 is generic, so by Corollary 11 it has exactly one subgame perfect equilibrium. To determine it we can use the observation there established, namely that in every Nash equilibrium both players select S at the nodes a and b, respectively. Indeed, by the structure of the game this observation also holds for every subgame. It follows that in the unique subgame perfect equilibrium both players select S at all non-leaf nodes. This counterintuitive form of the subgame perfect equilibrium in this game is sometimes used to question the adequacy of this solution concept.

It is also natural to study conditions under which the subgame perfect equilibria are payoff equivalent. The following theorem is implicit in [13]. The TDI condition was introduced in Section 2 when discussing strategic games.

Theorem 12. *In every finite extensive game that satisfies the TDI condition all subgame perfect equilibria are payoff equivalent.*

Proof. Consider a finite extensive game $G := (T, turn, o_1, \ldots, o_n)$ that satisfies the TDI condition. We proceed by induction on the number of nodes in the game tree. The claim holds when the game tree has just one node, since there is then only one subgame perfect equilibrium. Suppose the game tree has more than one node.

Consider two subgame perfect equilibria s and t in G. Take a preleaf v in T. Suppose $s_i(v) = w_1$ and $t_i(v) = w_2$, where $i = turn(v)$. By Corollary 9 $w_1, w_2 \in \text{argmax}_{x \in C(v)} o_i(x)$.

Case 1. $leaf(s) = w_1$ and $leaf(t) = w_2$.

Take a strategy s'_i that differs from s_i only for the node v to which it assigns w_2. Then $leaf(s'_i, s_{-i}) = w_2$, so $o_i(leaf(s)) = o_i(w_1) = o_i(w_2) = o_i(leaf(s'_i, s_{-i}))$. Hence by the TDI property $o(leaf(s)) = o(leaf(s'_i, s_{-i}))$, so $o(leaf(s)) = o(leaf(t))$.

Case 2. $leaf(s) \neq w_1$ or $leaf(t) \neq w_2$.

Without loss of generality suppose that $leaf(t) \neq w_2$. Consider the game $G' := (T', turn, o_1, \ldots, o_n)$ obtained from G by setting $o(v)$ to $o(w_1)$ and by removing all the children of v. So in G the node v is a preleaf, while in G' it is a leaf with the outcome $o(w_1)$. G' also satisfies the TDI condition since all its outcomes are also outcomes of G.

Let s' and t' be joint strategies in G' obtained from s and t by dropping v from the domains of s_i and t_i. Then both s' and t' are subgame perfect equilibria in G'. (We leave the proof of this fact to the reader.)

We have $o(leaf(s)) = o(leaf(s'))$ and by assumption the node v does not lie on the path $play(t)$, so $leaf(t) = leaf(t')$. Hence $o(leaf(s)) = o(leaf(t))$ by the induction hypothesis. □

6 Backward induction and common knowledge of rationality

Recall that player's rationality in an extensive game means that his objective is to maximize his outcome in the game. Backward induction is a natural procedure and it is natural to inquire whether it can be justified by appealing to players' rationality.

In this section we discuss Aumann's result [3] that for a natural class of extensive games common knowledge of players' rationality implies that the game reaches the backward induction outcome.

To formulate this result we introduce first Aumann's approach to modeling knowledge in the context of extensive games. Fix a finite extensive game with no relevant ties $G := (T, turn, o_1, \ldots, o_n)$ with $T = (V, E, v_0)$. Let $S_1, \ldots, S_n$ be the respective sets of strategies of players $1, \ldots, n$.

A ***knowledge system*** for G consists of

- a non-empty set Ω of ***states***,
- a function $\mathbf{s} : \Omega \to S_1 \times \cdots \times S_n$,
- for each player i a partition P_i of Ω.

One possible interpretation of a state is that it represents a 'situation', in which complete information about the players' strategies is available. This information is provided by means of the function $\mathbf{s}$.

Given a knowledge system, player i does not know the actual state ω but he knows the element of the partition P_i to which ω belongs. This interpretation suggests the following assumption.

Define for player i the function $\mathbf{s}_i : \Omega \to S_i$ by

$$\mathbf{s}_i(\omega) := s_i, \text{ where } s_i \text{ is the } i\text{th component of } \mathbf{s}(\omega).$$

Then we assume that for each player i the function $\mathbf{s}_i$ is constant on each element of the partition P_i. Intuitively, it means that in the assumed knowledge system each player knows his strategy.

We first introduce concepts that do not rely on the function $\mathbf{s}$. By an ***event*** we mean a subset of Ω. For an event E and player i we define the event K_iE by

$$K_iE := \bigcup \{L \in P_i \mid L \subseteq E\}.$$

Intuitively, K_iE is the event that player i *knows* E.

Next, we define

$$KE := K^1 E := \bigcap_{i=1}^{n} K_i E,$$

inductively for $k \geq 1$

$$K^{k+1} E := K^k E,$$

and finally

$$CKE := \bigcap_{k=1}^{\infty} K^k E.$$

Intuitively, KE is the event that all players *know* the event E and CKE is the event that there is *common knowledge* of the event E among all players.

Using the function $\mathbf{s}$ we now formalize the notion that player i is rational. To start with, given a node v at which player i moves, his strategy t_i, and the function $\mathbf{s}_{-i} : \Omega \to S_{-i}$ defined in the expected way, we denote by

$$[o_i(leaf((\mathbf{s}_{-i}, t_i)^v)) > o_i(leaf(\mathbf{s}^v))]$$

the event

$$\{\omega \in \Omega \mid o_i(leaf((\mathbf{s}_{-i}(\omega), t_i)^v)) > o_i(leaf(\mathbf{s}(\omega)^v)\}.$$

It states that in the subgame G^v the outcome for player i is higher if he selects strategy t_i instead of the strategy he chooses according to $\mathbf{s}$.

Similarly, for a joint strategy t we denote by

$$[\mathbf{s} = t]$$

the event

$$\{\omega \in \Omega \mid \mathbf{s}(\omega) = t\}.$$

Recall now that for a player i we denoted by V_i the set of nodes at which he moves. We define

$$R_i := \bigcap_{v \in V_i} \bigcap_{t_i \in S_i} \neg K_i [o_i(leaf((\mathbf{s}_{-i}, t_i)^v)) > o_i(leaf(\mathbf{s}^v))].$$

where $\neg$ denotes complementation w.r.t. Ω. Intuitively, this event states that for all nodes v at which player i moves and all his strategies t_i, player i does not know whether t_i would yield a higher outcome than the strategy he chooses according to $\mathbf{s}$. So R_i is the event formalizing that player i is *rational*.

Finally, we define

$$R := \bigcap_{i=1}^{n} R_i.$$

Intuitively, R is the event that each player rational.

We still need to formalize the event that the outcome of the game is prescribed by the backward induction. To this end Aumann assumes that the game is generic, so that the game has a unique subgame perfect equilibrium, but thanks to Corollary 11 it suffices to assume that the game is without relevant ties. Then the backward induction has a unique outcome which is the subgame perfect equilibrium of the game. Denote the latter by s^*. The intended event I is then defined by

$$I := [\mathbf{s} = s^*].$$

We can now state the main result of [3].

Theorem 13. *Consider a finite extensive game G without relevant ties. Then*

$$CKR \subseteq I.$$

This inclusion formalizes the announced statement that common knowledge of players' rationality implies that the backward induction yields the outcome of the game.

The proof of the theorem relies on a number of simple properties of the operators K_i and CK that we list without proof in the following lemma.

Lemma 14.

(i) $CKE = K_i CKE$.

(ii) *If* $E \subseteq F$, *then* $K_i E \subseteq K_i F$.

(iii) $K_i E \cap K_i F = K_i(E \cap F)$.

(iv) $CKE \subseteq E$.

(v) $K_i \neg K_i E = \neg K_i E$.

(vi) $K_i E \subseteq E$.

Given a joint strategy s and a node v that is not a leaf, we define

$$s(v) := s_i(v),$$

where $i = turn(v)$. So if s is the used joint strategy, then $s(v)$ is the move resulting from it at node v. For each such node v we define the function $\mathbf{s}(v) : \Omega \rightarrow V$ in the expected way.

We shall also need the following two observations concerning players' knowledge the proofs of which we omit.

Lemma 15.

(i) For all $t_i \in S_i$, $[\mathbf{s}_i = t_i] \subseteq K_i[\mathbf{s}_i = t_i]$.

(ii) For all $v \in V_i$, $I^v \subseteq K_i I^v$, *where* $I^v := [\mathbf{s}(v) = s^*(v)]$.

Intuitively, (i) states that if player i chooses the strategy t_i he knows that he chooses it and (ii) states that if player i chooses the move $s^*(v)$ at the node v, then he knows this. Note that by Lemma 14(vi) we can replace in (i) and (ii) $\subseteq$ by $=$.

Proof of Theorem 13.

We have $I = \bigcap_{v \in V \setminus Z} I^v$, so it suffices to prove that for all $v \in V \setminus Z$, $CKR \subseteq I^v$. Given two nodes v and w we write $w < v$ if w is a (possibly indirect) descendant of v.

We proceed by induction w.r.t. $<$. Take a node v and suppose that $CKR \subseteq I^w$ for all $w < v$. Let $i = turn(v)$. For a joint strategy s denote by $s^{<v}$ the joint strategy s^v with the pair $(v, s_i(v))$ removed from s_i, and define the function $\mathbf{s}^{<v} : \Omega \to S_1 \times \cdots \times S_n$ by

$$\mathbf{s}^{<v}(\omega) := \mathbf{s}(\omega)^{<v}.$$

By Lemma 14(i) and (ii) and the induction hypothesis $CKR = K_i CKR \subseteq K_i I^w$ for all $w < v$, so by Lemma 14(iii)

$$CKR \subseteq \bigcap_{w<v} K_i I^w = K_i \bigcap_{w<v} I^w = K_i[\mathbf{s}^{<v} = (s^*)^{<v}]. \tag{2}$$

Also, by Lemma 14(iv) and the definition of R_i with t_i set to s_i^*

$$CKR \subseteq R \subseteq R_i \subseteq \neg K_i[o_i(leaf((\mathbf{s}_{-i}, s_i^*)^v)) > o_i(leaf(\mathbf{s}^v))]. \tag{3}$$

Further, by Lemma 14(iii) and the fact that since $i = turn(v)$,

$$\begin{aligned}
& K_i[\mathbf{s}^{<v} = (s^*)^{<v}] \cap K_i[o_i(leaf((\mathbf{s}_{-i}, s_i^*)^v)) > o_i(leaf(\mathbf{s}^v))] \\
= \ & K_i[\mathbf{s}^{<v} = (s^*)^{<v} \wedge o_i(leaf((\mathbf{s}_{-i}, s_i^*)^v)) > o_i(leaf(\mathbf{s}^v))] \\
= \ & K_i[\mathbf{s}^{<v} = (s^*)^{<v} \wedge o_i(leaf((s^*)^v)) > o_i(leaf((s_{-i}^*, \mathbf{s}_i)^v))] \\
= \ & K_i[\mathbf{s}^{<v} = (s^*)^{<v}] \cap K_i[o_i(leaf((s^*)^v)) > o_i(leaf((s_{-i}^*, \mathbf{s}_i)^v))],
\end{aligned}$$

so by taking complement w.r.t. $K_i[\mathbf{s}^{<v} = (s^*)^{<v}]$

$$\begin{aligned}
& K_i[\mathbf{s}^{<v} = (s^*)^{<v}] \cap \neg K_i[o_i(leaf((\mathbf{s}_{-i}, s_i^*)^v)) > o_i(leaf(\mathbf{s}^v))] \\
= \ & K_i[\mathbf{s}^{<v} = (s^*)^{<v}] \cap \neg K_i[o_i(leaf((s^*)^v)) > o_i(leaf((s_{-i}^*, \mathbf{s}_i)^v))] \\
\subseteq \ & \neg K_i[o_i(leaf((s^*)^v)) > o_i(leaf((s_{-i}^*, \mathbf{s}_i)^v))].
\end{aligned} \tag{4}$$

Finally, by (2)–(4), the fact that for each node v, s^v is a unique subgame perfect equilibrium in G^v, Lemma 15, and Lemma 14(v) and (vi)

$$\begin{aligned}
& CKR \subseteq K_i[\mathbf{s}^{<v} = (s^*)^{<v}] \cap \neg K_i[o_i(leaf((\mathbf{s}_{-i}, s_i^*)^v)) > o_i(leaf(\mathbf{s}^v))] \\
\subseteq \ & \neg K_i[o_i(leaf((s^*)^v)) > o_i(leaf((s_{-i}^*, \mathbf{s}_i)^v))] = \neg K_i[\mathbf{s}(v) \neq s^*(v)] \\
= \ & \neg K_i \neg I^v = \neg K_i \neg K_i I^v = \neg\neg K_i I^v = K_i I^v \subseteq I^v,
\end{aligned}$$

as desired. ◻

We conclude by the following observation of [3] showing that non-trivial knowledge systems can be easily constructed.

Note 16. *For every finite extensive game without relevant ties there exists a knowledge system such that $CKR \neq \emptyset$.*

Proof. It suffices to choose Ω to be a singleton set $\{\omega\}$ and set $\mathbf{s}(\omega) := s^*$, where s^* is the unique subgame perfect equilibrium of the considered game. Then $CKR = \Omega$. ◻

Aumann's paper dealt with concepts also studied by philosophers and psychologists. As a result it became highly influential and attracted wide attention. In particular Stalnaker pointed out in [21] that Aumann's notion of rationality involves reasoning about situations (nodes) that the agent knows will never be reached and constructed a model in which common knowledge of players' rationality does not imply that the game reaches the backward induction outcome.

The apparent contradiction between Aumann's and Stalnaker's results was clarified by Halpern in [8]. The difference can be explained by adding to Aumann's knowledge system for an extensive game one more parameter, a function

$$f : \Omega \times V \setminus Z \to \Omega,$$

that for a given state ω and a non-leaf node v yields a state ω' that is 'nearest' (in a well-defined sense) to v and is such that v is *reached* in ω', i.e., is such that v lies on $play(\mathbf{s}(\omega'))$.

Then according to Stalnaker, a player i is *substantively rational* in a state ω, if for each node $v \in V_i$ he is rational in the state $\omega' = f(\omega, v)$, where the latter means that

$$\omega' \in \bigcap_{t_i \in S_i} \neg K_i[o_i(leaf((\mathbf{s}_{-i}, t_i)^v)) > o_i(leaf(\mathbf{s}^v))].$$

Stalnaker's model refers to substantive rationality and not rationality.

7 Weak dominance and backward induction

Iterated elimination of weakly dominated strategies is defined for strategic games, so it can be also applied to the strategic forms of extensive games. For the class of finite extensive games discussed in the previous section this procedure is closely related to backward induction. The aim of this section is to clarify this relation.

The following notion will be needed. Consider a node w in the game tree of an extensive game G such that $turn(w) = i$. We say that a strategy s_i of player i ***can reach*** w if for some s_{-i} the node w lies on the path $play(s_i, s_{-i})$.

We begin by introducing Algorithm 2 that is a modification of the backward induction algorithm 1 from Section 5 in which the input and output are modified and line 7 is added.

Algorithm 2:

Input: A finite extensive game with no relevant ties $G := (T, turn, o_1, \ldots, o_n)$ with $T = (V, E, v_0)$ and $\Gamma(G) = (S_1, \ldots, S_n, p_1, \ldots, p_n)$.

Output: The subgame perfect equilibrium s in G, extensions of the functions $o_1, \ldots, o_n$ to all nodes of T such that $o(v_0) = o(leaf(s))$, and a trivial strategic game $(S_1, \ldots, S_n, p_1, \ldots, p_n)$ that includes s.

1 **while** $|V| > 1$ **do**
2   **choose** $v \in V$ that is a preleaf of T;
3     $i := turn(v)$;
4   **choose** $w \in C(v, T)$ such that $o_i(w)$ is maximal;
5     $s_i(v) := w$;
6     $o(v) := o(w)$;
7     $S_i := S_i \setminus \{s'_i \in S_i \mid s'_i$ can reach v and $s'_i(v) \neq w\}$;
8     $V := V \setminus C(v, T)$;
9     $E := E \cap (V \times V)$;
10     $T := (V, E, v_0)$

The following theorem makes precise the mentioned relation between two concepts.

Theorem 17. *Consider a finite extensive game G without relevant ties and Algorithm 2 applied to it.*

(i) Each strategy removed in line 7 is weakly dominated in the current version of the strategic game.

(ii) The strategic game $(S_1, \ldots, S_n, p_1, \ldots, p_n)$ which is generated upon termination of the algorithm is trivial and includes the subgame perfect equilibrium of G.

Proof. Below s is the subgame perfect equilibrium of the game G. Consider the assertion $I(u)$ defined by

$$I(u) \equiv \forall t \in S : [\text{if } u \text{ appears in } play(t) \text{ then } o(leaf(s^u)) = o(leaf(t^u))],$$

where $S = S_1 \times \cdots \times S_n$.

First notice that if during an execution of the algorithm for some node u the assertion $I(u)$ becomes true, then it remains true. The reason is that the considered set S of joint strategies never increases.

We now show that after each loop iteration $I(v)$ holds, where v is the node being dealt with in current loop iteration and S refers to the current set of joint strategies. Fix an execution of the algorithm. We proceed by induction on the order in which the nodes of T are selected in line **2**. Consider first a preleaf v in the original game tree T. Take $t \in S$ such that v appears in $play(t)$ and let $i = turn(v)$. Then

$$\begin{array}{rl} & o(leaf(t^v)) \\ = & \quad \{ \text{ by line } \mathbf{7}\ t_i(v) = w \} \\ & o(w) \\ = & \quad \{ \text{ by line } \mathbf{5}\ s_i(v) = w \} \\ & o(leaf(s^v)). \end{array}$$

Next, consider a node v in the original game tree T selected in line **2** that is neither a preleaf nor a leaf and consider the program state after the current loop iteration. Then both $i = turn(v)$ and $s_i(v) = w$.

By the order in which the nodes are selected in line **2**, all nodes $u \in C(v)$ have been dealt with in earlier loop iterations. So by the induction hypothesis $I(u)$ holds for $u \in C(v)$. In particular $I(w)$ holds. Take $t \in S$ such that v appears in $play(t)$. Then

$$\begin{array}{rl} & o(leaf(t^v)) \\ = & \quad \{ \text{ by line } \mathbf{7}\ t_i(v) = w \} \\ & o(leaf(t^w)) \\ = & \quad \{ \ I(w)) \ \} \\ & o(leaf(s^w)) \\ = & \quad \{ \text{ by line } \mathbf{5}\ s_i(v) = w \} \\ & o(leaf(s^v)). \end{array}$$

(i) Fix the program state after an arbitrary loop iteration of the algorithm with the current values of $v, w, i, S_1, \ldots, S_n$. In particular $turn(v) = i$.

We first prove that for $u \in C(v)$, $u \neq w$

$$o_i(leaf(s^w)) > o_i(leaf(s^u)). \tag{5}$$

Let t_i be the strategy of player i that differs from s_i only for the node v to which it assigns u. We have $s_i(v) = w$ and by definition s^v is a Nash equilibrium in the subgame G^v, so

$$o_i(leaf(s^w)) = o_i(leaf(s^v)) \geq o_i(leaf(t_i^v, s_{-i}^v)) = o_i(leaf(t_i^u, s_{-i}^u)) = o_i(leaf(s^u)).$$

But $turn(v) = i$, $u, w \in C(v)$, $u \neq w$, and G^v is without relevant ties, so (5) follows.

Take now a strategy s'_i removed in line **7** and suppose $s'_i(v) = u$. Consider the strategy t_i that differs from s'_i only for the node v to which it assigns w (i.e., $s_i(v)$). We claim that after line **7** t_i weakly dominates s'_i in the current version of $(S_1, \ldots, S_n, p_1, \ldots, p_n)$.

Take some $t_{-i} \in S_{-i}$. If $v \notin play(t_i, t_{-i})$, then $v \notin play(s'_i, t_{-i})$, so $play(t_i, t_{-i}) = play(s'_i, t_{-i})$ and hence

$$p_i(t_i, t_{-i}) = o_i(leaf(t_i, t_{-i}) = o_i(leaf(s'_i, t_{-i}) = p_i(s'_i, t_{-i}).$$

If $v \in play(t_i, t_{-i})$, then also $w \in play(t_i, t_{-i})$, $v \in play(s'_i, t_{-i})$ and $u \in play(s'_i, t_{-i})$, where, recall, $s'_i(v) = u$. Since $u, w \in C(v)$, these two nodes have been dealt with in earlier loop iterations. So both $I(u)$ and $I(w)$ hold.

Hence

$$p_i(t_i, t_{-i}) = o_i(leaf(t_i, t_{-i})) = o_i(leaf((t_i, t_{-i})^w)) = o_i(leaf(s^w))$$

and

$$p_i(s'_i, t_{-i}) = o_i(leaf(s'_i, t_{-i})) = o_i(leaf((s'_i, t_{-i})^u)) = o_i(leaf(s^u)).$$

So $p_i(t_i, t_{-i}) > p_i(s'_i, t_{-i})$ by (5).

Now, t_i does not need to be a strategy from S_i but thanks to Lemma 2 we can conclude that a strategy t'_i in S_i exists that weakly dominates s'_i in the current version of $(S_1, \ldots, S_n, p_1, \ldots, p_n)$.

(*ii*) Upon termination of the algorithm $I(v_0)$, i.e., $\forall t \in S : o(leaf(s)) = o(leaf(t))$ holds. This means that upon termination the final game $(S_1, \ldots, S_n, p_1, \ldots, p_n)$ is trivial. Further, for each player each strategy removed in line **7** differs from his strategy in the subgame equilibrium, which means that the final game includes this equilibrium. □

This theorem shows that every finite extensive game G without relevant ties can be solved by an IEWDS. Recall from Corollary 11 that such a game has a unique subgame perfect equilibrium. However, not all instances of the IEWDS behave the desired way. The following example, taken from [16, page 109], shows that some generic extensive games can be solved by an IEWDS that removes the unique subgame perfect equilibrium. This explains why in line **7** in Algorithm 2 only specific weakly dominated strategies are removed.

Example 11. Consider the two-player generic extensive game and its associated strategic game given in Figure 10. In the figure, the nodes are labelled with the player whose turn it is to move.

Consider now an IEWDS that consists of the following sequence of elimination of weakly dominated strategies: AE, D, AF. The resulting trivial subgame has two joint strategies (BE, C), (BF, C). So this instance of the IEWDS eliminates (BE, D), which is the unique subgame perfect equilibrium. □

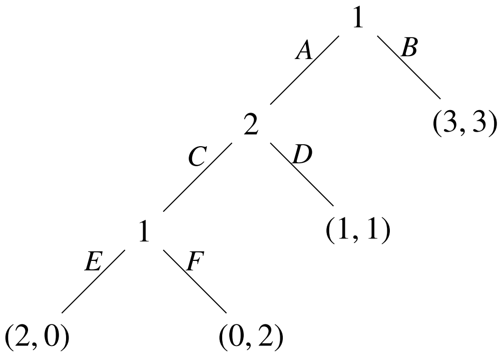

	C	D
AE	2, 0	1, 1
AF	0, 2	1, 1
BE	3, 3	3, 3
BF	3, 3	3, 3

Figure 10: An extensive game (left) and its associated strategic game (right).

8 Weak dominance and strictly competitive games

We now continue an account of iterated elimination of weakly dominated strategies. In Theorem 17 we showed that each finite extensive game without relevant ties can be solved by an IEWDS. A natural question is whether we can extend this result to arbitrary finite extensive games. The following example taken from [16, pages 109-110] shows that this fails to be the case already for two-player games.

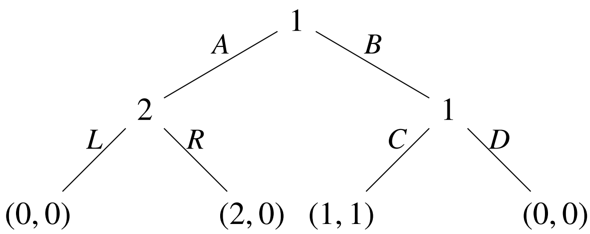

	L	R
AC	0, 0	2, 0
AD	0, 0	2, 0
BC	1, 1	1, 1
BD	0, 0	0, 0

Figure 11: An extensive game (left) and its associated strategic game (right).

Example 12. Consider the two-player extensive game and its associated strategic game given in Figure 11. In the game tree, the nodes are labelled with the player whose turn it is to move. For this game there is just one instance of IEWDS which consists of eliminating the strategy *BD*. The resulting subgame is not trivial, so no instance of IEWDS can solve this extensive game. □

On the other hand, as shown in [7], finite extensive zero-sum games can be solved by an IEWDS in which at each step all weakly dominated strategies are removed. The aim of this section is to present this result for the larger class of finite extensive strictly competitive games for which the same proof remains valid.

From now on, given a strategic game H we denote by H^1 a subgame of H obtained by the elimination of *all* strategies that are weakly dominated in H, and put $H^0 := H$ and $H^{k+1} := (H^k)^1$, where $k \geq 1$. So, in contrast to Sections 2 and 3 each H^k is now uniquely defined.

Below for a strategic game H we denote by H_i the set of strategies of player i. Also, for a finite extensive game G we write $\Gamma^k(G)$ instead of $(\Gamma(G))^k$, $\Gamma_i(G)$ instead of $(\Gamma(G))_i$, and $\Gamma_i^k(G)$ instead of $(\Gamma^k(G))_i$. In particular $\Gamma^0(G) = \Gamma(G)$.

Further, for a finite strictly competitive strategic game $H = (S_1, S_2, p_1, p_2)$ we define for each player i

$$
\begin{aligned}
&p_i^{max}(H) := \max_{s \in S} p_i(s),\\
&win_i(H) := \{s_i \in S_i \mid \forall s_{-i} \in S_{-i}\ p_i(s_i, s_{-i}) = p_i^{max}(H)\},\\
&lose_{-i}(H) = \{s_{-i} \in S_{-i} \mid \exists s_i \in S_i\ p_i(s_i, s_{-i}) = p_i^{max}(H)\}.
\end{aligned}
$$

So $p_i^{max}(H)$ is the maximal payoff player i can receive in the game H, $win_i(H)$ is the set of strategies of player i for which he always gets $p_i^{max}(H)$, while $lose_{-i}(H)$ is the set of strategies of player $-i$ for which his opponent i can get his maximally possible payoff $p_i^{max}(H)$.

The following lemma, with a rather involved proof, is crucial.

Lemma 18. *Let G be a finite strictly competitive extensive game. For all $i \in \{1, 2\}$ and for all $k \geq 0$, if $win_i(\Gamma^k(G)) = \emptyset$ then $lose_{-i}(\Gamma^k(G)) \cap \Gamma_{-i}^{k+2}(G) = \emptyset$.*

This lemma implies that if for all $i \in \{1, 2\}$, $win_i(\Gamma^k(G)) = \emptyset$ then two further rounds of eliminations of weakly dominated strategies remove from $\Gamma^k(G)$ at least two outcomes.

Proof. Fix i and k and suppose $win_i(\Gamma^k(G)) = \emptyset$. So for all $s_i \in \Gamma_i^k(G)$ we have $\min_{s_{-i} \in \Gamma_{-i}^k(G)} p_i(s_i, s_{-i}) < p_i^{max}(H)$, and hence $maxmin_i(\Gamma^k(G)) < p_i^{max}(\Gamma^k(G))$.

By Corollary 10 the strategic game $\Gamma(G)$ has a Nash equilibrium. By the repeated application of Corollary 5 we have $maxmin_i(\Gamma(G)) = maxmin_i(\Gamma^k(G))$. Therefore

$$maxmin_i(\Gamma(G)) < p_i^{max}(\Gamma^k(G)). \tag{6}$$

Take now $s_{-i} \in lose_{-i}(\Gamma^k(G))$. We need to prove that $s_{-i} \notin \Gamma_{-i}^{k+2}(G)$.

For some $s_i \in \Gamma_i^k(G)$ we have $p_i(s_i, s_{-i}) = p_i^{max}(\Gamma^k(G))$. By Lemma 2 we can assume that $s_i \in \Gamma_i^{k+1}(G)$. Consider now the path $play(s_i, s_{-i})$. Then

- by (6) for the first node u lying on $play(s_i, s_{-i})$ (so the root),

  $maxmin_i(\Gamma(G^u)) < p_i^{max}(\Gamma^k(G))$,

- for the last node u lying on $play(s_i, s_{-i})$ (so the leaf),

  $p_i^{max}(\Gamma^k(G)) = p_i(s_i, s_{-i}) = o_i(u) = maxmin_i(\Gamma(G^u))$.

So for some adjacent nodes u, w lying on the path $play(s_i, s_{-i})$

$$maxmin_i(\Gamma(G^u)) < p_i^{max}(\Gamma^k(G)) \le maxmin_i(\Gamma(G^w)). \tag{7}$$

Further, if for some adjacent nodes u', w' lying on the path $play(s_i, s_{-i})$ we have $turn(u') = i$ and $p_i^{max}(\Gamma^k(G)) \le maxmin_i(\Gamma(G^{w'}))$, then $maxmin_i(\Gamma(G^{u'})) = maxmin_i(\Gamma(G^{w'}))$. So $turn(u) = -i$ and $s_{-i}(u) = w$.

If $s_{-i} \notin \Gamma_{-i}^{k+1}(G)$ then $s_{-i} \notin \Gamma_{-i}^{k+2}(G)$. So suppose $s_{-i} \in \Gamma_{-i}^{k+1}(G)$. We prove that then s_{-i} is weakly dominated in $\Gamma^{k+1}(G)$. The dominating strategy is obtained in two steps.

By Corollary 10 the game $\Gamma(G^u)$ has a Nash equilibrium s^*. First, we introduce the strategy $t_{-i} \in \Gamma_{-i}(G)$ defined as follows:

$$t_{-i}(x) := \begin{cases} s_{-i}(x) & \text{if } x \text{ not in } T^u, \\ s_{-i}^*(x) & \text{if } x \text{ in } T^u, \end{cases}$$

where $turn(x) = -i$ and T is the game tree of G.

We now establish two claims relating t_{-i} to s_{-i}.

Claim 1. $\forall s_i' \in \Gamma_i^{k+1}(G) : p_{-i}(s_i', t_{-i}) \ge p_{-i}(s_i', s_{-i})$.

Proof. Suppose by contradiction that there exists $s_i' \in \Gamma_i^{k+1}(G)$ such that $p_{-i}(s_i', t_{-i}) < p_{-i}(s_i', s_{-i})$. The strategy t_{-i} differs from s_{-i} only on the nodes in T^u, so the difference in the payoffs implies that u appears both in $play(s_i', t_{-i})$ and $play(s_i', s_{-i})$. This implies

$$maxmin_{-i}(\Gamma(G^u)) \le p_{-i}((s_i')^u, s_{-i}^*) = p_{-i}(s_i', t_{-i}) < p_{-i}(s_i', s_{-i}). \tag{8}$$

By Theorem 4 s_{-i}^* is a security strategy of player $-i$ in the game $\Gamma(G^u)$. Further, the node u appears in $play(s_i', s_{-i})$, so $s'' := (s_i', s_{-i})^u$ is a joint strategy in G^u. This and (8) imply

$$p_{-i}(s^*) = maxmin_{-i}(\Gamma(G^u)) < p_{-i}(s_i', s_{-i}) = p_{-i}(s''),$$

so by (1), the fact that G^u is strictly competitive, and Theorem 4

$$p_i(s_i', s_{-i}) = p_i(s'') < p_i(s^*) = maxmin_i(\Gamma(G^u)). \tag{9}$$

Next we introduce the strategy $t_i \in \Gamma_i(G)$ defined as follows (recall that $w = s_{-i}(u)$):

$$t_i(x) := \begin{cases} s_i'(x) & \text{if } x \text{ not in } T^w, \\ t_i^*(x) & \text{if } x \text{ in } T^w, \end{cases}$$

where $turn(x) = i$ and t_i^* is a security strategy of player i in the game $\Gamma(G^w)$.

We now establish two claims relating t_i to s'_i:

$$\forall s'_{-i} \in \Gamma^k_{-i}(G) : p_i(t_i, s'_{-i}) \geq p_i(s'_i, s'_{-i}), \tag{10}$$

$$p_i(t_i, s_{-i}) > p_i(s'_i, s_{-i}). \tag{11}$$

To establish (10) consider any strategy $s'_{-i} \in \Gamma^k_{-i}(G)$. By the definition of t_i, if w does not appear in $play(t_i, s'_{-i})$ then $p_i(t_i, s'_{-i}) = p_i(s'_i, s'_{-i})$. So suppose w appears in $play(t_i, s'_{-i})$. By the definition of t_i, (7), and the fact that both s'_i and s'_{-i} are strategies in $\Gamma^k(G)$

$$p_i(t_i, s'_{-i}) = p_i(t^*_i, (s'_{-i})^w) \geq maxmin_i(\Gamma(G^w)) \geq p_i^{max}(\Gamma^k(G)) \geq p_i(s'_i, s'_{-i}).$$

To establish (11) recall that we noted already that u appears in $play(s'_i, s_{-i})$. Since $turn(u) = -i$ and $s_{-i}(u) = w$, also w appears in $play(s'_i, s_{-i})$. The strategy t_i differs from s'_i only on the nodes in T^w, so w apears in $play(t_i, s_{-i})$, as well. Therefore by the definition of t_i, (7) and (9)

$$p_i(t_i, s_{-i}) = p_i(t^*_i, (s_{-i})^w) \geq maxmin_i(\Gamma(G^w)) > maxmin_i(\Gamma(G^u)) > p_i(s'_i, s_{-i}).$$

By Lemma 2 there exists $t'_i \in \Gamma^k_i(G)$ such that $p_i(t'_i, s'_{-i}) \geq p_i(t_i, s'_{-i})$ for all $s'_{-i} \in \Gamma^k_{-i}(G)$. This, together with (10) and (11) implies that s'_i is weakly dominated by t'_i in $\Gamma^k(G)$. Hence $s'_i \notin \Gamma^{k+1}_i(G)$, which yields a contradiction. □

Claim 2. $p_{-i}(s_i, t_{-i}) > p_{-i}(s_i, s_{-i})$.

Proof. The node u appears in $play(s_i, s_{-i})$, so by Theorem 4 and (7)

$$p_i(s^*) = maxmin_i(\Gamma(G^u)) < p_i^{max}(\Gamma^k(G)) = p_i(s_i, s_{-i}) = p_i(s^u_i, s^u_{-i}),$$

and consequently by Theorem 4(i) and the fact that G^u is strictly competitive

$$maxmin_{-i}(\Gamma(G^u)) = p_{-i}(s^*) > p_{-i}(s^u_i, s^u_{-i}) = p_{-i}(s_i, s_{-i}). \tag{12}$$

Further, the strategy t_{-i} differs from s_{-i} only on the nodes in T^u, so u appears in $play(s_i, t_{-i})$, as well. Hence

$$p_{-i}(s_i, t_{-i}) = p_{-i}(s^u_i, s^*_{-i}) \geq maxmin_{-i}(\Gamma(G^u)). \tag{13}$$

Combining (12) and (13) we get the claim. □

By Lemma 2 there exists $t'_{-i} \in \Gamma^{k+1}_{-i}(G)$ such that $p_{-i}(s'_i, t'_{-i}) \geq p_{-i}(s'_i, t_{-i})$ for all $s'_i \in \Gamma^{k+1}_i(G)$. We conclude by Claims 1 and 2 that s_{-i} is weakly dominated by t'_{-i} in $\Gamma^{k+1}(G)$. Therefore $s_{-i} \notin \Gamma^{k+2}_{-i}(G)$, as desired. □

We can now establish the announced result.

Theorem 19. *Let G be a finite strictly competitive extensive game with at most m outcomes. Then $\Gamma^{m-1}(G)$ is a trivial game.*

Proof. We prove a stronger claim, namely that for all $m \geq 1$ and $k \geq 0$ if $\Gamma^k(G)$ has at most m outcomes, then $\Gamma^{k+m-1}(G)$ is a trivial game.

We proceed by induction on m. We can assume that $m > 1$. For $m = 2$ the claim follows by Lemma 6. Take $m > 2$.

Case 1. For some $i \in \{1, 2\}$, $win_i(\Gamma^k(G)) \neq \emptyset$.

For player i every strategy $s_i \in win_i(\Gamma^k(G))$ weakly dominates all strategies $s'_i \notin win_i(\Gamma^k(G))$. So in $\Gamma^{k+1}(G)$ the set of strategies of player i equals $win_i(\Gamma^k(G))$ and hence $p_i^{max}(\Gamma^k(G))$ is his unique payoff in this game. By (1) $\Gamma^{k+1}(G)$, and hence also $\Gamma^{k+m-1}(G)$, is a trivial game.

Case 2. For all $i \in \{1, 2\}$, $win_i(\Gamma^k(G)) = \emptyset$.

Take joint strategies s and t such that $p_1(s) = p_1^{\max}(\Gamma^k(G))$ and $p_2(t) = p_2^{\max}(\Gamma^k(G))$. Since $m > 1$ (1) implies that the outcomes $(p_1(s), p_2(s))$ and $(p_1(t), p_2(t))$ are different.

We have $s_2 \in lose_2(\Gamma^k(G))$ and $t_1 \in lose_1(\Gamma^k(G))$. Hence by Lemma 18 for no joint strategy s' in $\Gamma^{k+2}(G)$ we have $p_1(s') = p_1^{\max}(\Gamma^k(G))$ or $p_2(s') = p_2^{\max}(\Gamma^k(G))$.

So $\Gamma^{k+2}(G)$ has at most $m-2$ outcomes. By the induction hypothesis $\Gamma^{k+m-1}(G)$ is a trivial game. □

9 Weak acyclicity

By Theorem 10 every finite extensive game has a Nash equilibrium. A natural question is whether we can strengthen this result to show that finite extensive games have the FIP. The following example adapted from [15] shows that even for simplest extensive games the answer is negative.

Example 13. Consider the extensive form game given in Figure 1. Following the convention introduced in Example 4, the strategies for player 1 are C and D, while the strategies of player 2 are CC, CD, DC and DD.

Then the following improvement sequence generates an infinite improvement path in this game:

$$\begin{array}{ccccccccc} (D, \underline{DC}) & \rightarrow & (\underline{D}, CD) & \rightarrow & (C, \underline{CD}) & \rightarrow & (\underline{C}, DC) & \rightarrow & (D, \underline{DC}) \\ (3,0) & & (1,1) & & (2,2) & & (0,3) & & (3,0) \end{array}$$

For convenience of the reader in each joint strategy we underlined the strategy which is not a best response and listed the corresponding outcomes. □

However, a weaker result, due to [12], does hold. It implies that every finite extensive game has a Nash equilibrium, a result established earlier, in Corollary 10.

Theorem 20. *Every finite extensive game is weakly acyclic.*

Proof. We prove the claim by defining a weak potential. Take a finite extensive game $G := (T, turn, o_1, \ldots, o_n)$, with $T := (V, E, v_0)$ and let S be the set of joint strategies.

Consider first a function $R : S \times V \to \{0, 1\}$ defined as follows:

$$R(s, v) := \begin{cases} 1 \text{ if } s_i^v \text{ is a best response to } s_{-i}^v \text{ in the subgame } G^v \\ 0 \text{ otherwise,} \end{cases}$$

where $i = turn(v)$.

Let now $L := (v_1, \ldots, v_k)$ be a list of the nodes from V such that each node appears after all of its children in T. For example,

$$((2, 2), (0, 3), (3, 0), b, (1, 1), c, a)$$

is such a list of the nodes of the tree from Figure 1, where we identify each leaf with the corresponding outcome.

Finally define the function $P : S \to \{0, 1\}^k$ by putting

$$P(s) := (R(s, v_1), \ldots, R(s, v_k)),$$

where, recall, $L = (v_1, \ldots, v_k)$, and consider the strict lexicographic ordering $<_{lex}$ over $\{0, 1\}^k$. We show that P is a weak potential w.r.t. this ordering and appeal to the Weakly Acyclic Theorem 1.

So consider a joint strategy s in G that is not a Nash equilibrium. Take a player i such that s_i is not a best response to s_{-i}. Let t_i be a best response of player i to s_{-i} and let $t = (t_i, s_{-i})$. Define s'_i as the modification of t_i such that its values on the nodes not lying on $play(t)$ are the values provided by s_i. More formally, for all nodes v such that $turn(v) = i$ we put

$$s'_i(v) := \begin{cases} t_i(v) & \text{if } v \text{ lies on } play(t) \\ s_i(v) & \text{otherwise} \end{cases}$$

Let $s' = (s'_i, s_{-i})$. Then $play(s') = play(t)$, so $o(leaf(s')) = o(leaf(t))$, and hence s'_i is also a best response of player i to s_{-i}. Since s_i is not a best response to s_{-i}, for some v from $play(s')$ such that $turn(v) = i$ player's i strategy s_i^v is not a best response to s_{-i}^v in the subgame G^v. Take the last such node v from $play(s')$.

So $R(v, s) = 0$ while $R(v, s') = 1$, since $(s'_i)^v$ is a best response to s_{-i}^v in the subgame G^v. Further, by the choice of v and the definition of s'_i we have that $R(w, s') = R(w, s)$ for all nodes w that precede v on the list L. We conclude that $P(s) <_{lex} P(s')$. $\square$

10 Win or lose and chess-like games

In this section we discuss two classes of zero-sum extensive games introduced in Section 2. We begin with the win or lose games. Given such a game G we say that a strategy s_i of player i is ***winning*** if

$$\forall s_{-i} \in S_{-i} : o_i(leaf(s_i, s_{-i})) = 1,$$

and denote the (possibly empty) set of such strategies by $win_i(G)$.

The Matching Pennies game shows that in strategic win or lose games winning strategies may not exist. For finite win or lose extensive games the situation changes.

Theorem 21. *Let G be a finite win or lose extensive game. For all players i we have $win_i(G) \neq \emptyset$ iff $win_{-i}(G) = \emptyset$.*

Proof. Call the players white and black and call a finite win or lose extensive game *white* if the white player has a winning strategy in it and analogously for *black*. We prove that every such game is white or black. Clearly, these alternatives are mutually exclusive.

We proceed by induction on the number of nodes in the game tree. The claim clearly holds when the game tree has just one node. Consider a game G with the game tree T with more than one node. By the induction hypothesis for every child u of the root of T the subgame G^u is white or black.

Without loss of generality assume that in G the white player moves first. We claim that the game G is black if for every child u of the root of T the subgame G^u is black and otherwise that it is white. Indeed, in the first case no matter what is the first move of the white player he loses the game if the black player pursues his winning strategy in the resulting subgame, and otherwise the white player wins the game if he starts by selecting the move that leads to a white subgame and subsequently pursues in this subgame his winning strategy.

Note that we did not assume that the players alternate their moves. □

Next we consider chess-like games. We say that a strategy s_i of player i in such a game G ***guarantees him at least a draw*** if

$$\forall s_{-i} \in S_{-i} : o_i(leaf(s_i, s_{-i})) \geq 0,$$

and denote the (possibly empty) set of such strategies by $draw_i(G)$. The set $win_i(G)$ is defined as above.

Theorem 22. *Let G be a finite chess-like extensive game. We have*

$$win_1(G) \neq \emptyset \text{ or } win_2(G) \neq \emptyset \text{ or } (draw_1(G) \neq \emptyset \text{ and } draw_2(G) \neq \emptyset).$$

We reproduce a proof given in [2].

Proof. We introduce the following abbreviations:

- W_1 for $win_1(G) \neq \emptyset$,
- D_2 for $draw_2(G) \neq \emptyset$,
- W_2 for $win_2(G) \neq \emptyset$,
- D_1 for $draw_1(G) \neq \emptyset$.

Let G_1 and G_2 be the modifications of G in which each outcome $(0, 0)$ is replaced for G_1 by $(-1, 1)$ and for G_2 by $(1, -1)$. Then $win_1(G_1) = win_1(G)$, $win_2(G_1) = draw_2(G)$, $win_1(G_2) = draw_1(G)$, and $win_2(G_2) = win_2(G)$.

Hence by Theorem 21 applied to the games G_1 and G_2 we have $W_1 \vee D_2$ and $W_2 \vee D_1$, so $(W_1 \wedge W_2) \vee (W_1 \wedge D_1) \vee (D_2 \wedge W_2) \vee (D_2 \wedge D_1)$, which implies $W_1 \vee W_2 \vee (D_2 \wedge D_1)$, since $\neg(W_1 \wedge W_2)$, $(W_1 \wedge D_1) \equiv W_1$, and $(D_2 \wedge W_2) \equiv W_2$. □

The above result is attributed to [24]. However, in [19] it was pointed out that the paper contains only the idea and the corresponding result is not formally stated, and that the first rigorous statement of the result and its proof seems to have been provided in [10]. This result is stated in [22, page 125] and proved using backward induction (apparently the first use of it in the literature on game theory). In [6] a proof is provided that does not rely on backward induction and argument also covers chess-like games in which infinite plays, interpreted as draw, are allowed. As noticed in [2] Theorems 21 and 22 also hold for infinite extensive games in which every play is finite.

11 Conclusions

The aim of this tutorial was to provide a self-contained introduction to finite extensive games with perfect information aimed at computer scientists. Our objective was to provide a systematic presentation of the most important results concerning this class of games that in our view could be of interest to computer scientists.

In [2] we argued that the next most natural class of extensive games is the one in which every play is finite (in the set theory terminology the game trees are then *well-founded*). In such a class of games one can in particular consider *behavioural strategies* in the extensive games considered here, according to which a move consists of a probability distribution over the finite set of children of a given node.

Many textbooks on game theory rather choose as the next class extensive games with imperfect information. In these games players do not need to know the previous moves made by the other players. An example is the Battleship game in which the first move for each player consists of a secret placing of the fleet on the grid. An interested reader is referred to Part III of [16].

Acknowledgements

We would like to thank Ruben Brokkelkamp, Marcin Dziubiński, R. Ramanujam, and an anonymous referee for useful comments on the first version of this paper. The second author was partially supported by the grant MTR/2018/001244.

References

[1] K. R. Apt and E. Grädel, editors. *Lectures in Game Theory for Computer Scientists*. Cambridge University Press, 2011.

[2] K.R. Apt and S. Simon. Well-founded extensive games with perfect information. In *Proceedings 18th Conference on Theoretical Aspects of Rationality and Knowledge, TARK 2021*, volume 335, pages 7–21. Electronic Proceedings in Theoretical Computer Science (EPTCS), 2021.

[3] R. Aumann. Backward induction and common knowledge of rationality. *Games and Economic Behavior*, 8:6–19, 1991.

[4] P. Battigalli. On rationalizability in extensive games. *Journal of Economic Theory*, 74:40–61, 1997.

[5] P.K. Dutta. *Strategies and Games*. MIT Press, Cambridge, MA, 2001.

[6] C. Ewerhart. Backward induction and the game-theoretic analysis of chess. *Games and Economic Behaviour*, 39:206–214, 2002.

[7] C. Ewerhart. Iterated weak dominance in strictly competitive games of perfect information. *Journal of Economic Theory*, 107(2):474–482, 2002.

[8] J. Halpern. Substantive rationality and backward induction. *Games and Economic Behavior*, 37(2):425–435, 2001.

[9] G.A. Jehle and P.J. Reny. *Advanced Microeconomic Theory*. Addison Wesley, New York, NY, second edition, 2001.

[10] L. Kalmár. Zur theorie der abstrakten spiele. *Acta Universitatis Szegediensis/Sectio Scientiarum Mathematicarum*, 4:65–85, 1928/29.

[11] H. W. Kuhn. Extensive games. *Proc. of the National Academy of Sciences*, 36:570–576, 1950.

[12] N.S. Kukushkin. Perfect information and potential games. *Games and Economic Behavior*, 38(2):306–317, 2002.

[13] L.M. Marx and J.M. Swinkels. Order independence for iterated weak dominance. *Games and Economic Behaviour*, 18:219–245, 1997.

[14] I. Milchtaich. Congestion games with player-specific payoff functions. *Games and Economic Behaviour*, 13:111–124, 1996.

[15] I. Milchtaich. Schedulers, potentials and weak potentials in weakly acyclic games. Working paper 2013-03, Bar-Ilan University, Department of Economics, 2013.

[16] M.J. Osborne and A. Rubinstein. *A Course in Game Theory*. The MIT Press, 1994.

[17] R. Rosenthal. Games of perfect information, predatory pricing and the chain-store paradox. *Journal of Economic Theory*, 25(1):92–100, 1981.

[18] R. W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, (2):65–67, 1973.

[19] U. Schwalbe and P. Walker. Zermelo and the early history of game theory. *Games and Economic Behavior*, 34(1):123–137, 2001.

[20] R. Selten. Spieltheoretische behandlung eines oligopolmodells mit nachfrageträgheit. *Zeitschrift für die gesamte Staatswisenschaft*, 121:301–324 and 667–689, 1965.

[21] R. Stalnaker. Knowledge, belief and counterfactual reasoning in games. *Economics and Philosophy*, 12(2):133–163, 1996.

[22] J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior (60th Anniversary Commemorative Edition)*. Princeton Classic Editions. Princeton University Press, 2004.

[23] H. Peyton Young. The evolution of conventions. *Econometrica*, 61(1):57–84, 1993.

[24] E. Zermelo. Über eine anwendung der mengenlehre auf die theorie des schachspiels. In *Proc. of The Fifth International Congress of Mathematicians*, pages 501–504. Cambridge University Press, 1913.