On Strictly Competitive Multi-Player Games

Felix Brandt Felix Fischer
Computer Science Department
University of Munich
80538 Munich, Germany
{brandtf,fischerf}@tcs.ifi.lmu.de

Yoav Shoham
Computer Science Department
Stanford University
Stanford CA 94305, USA
shoham@cs.stanford.edu

Abstract

We embark on an initial study of a new class of strategic (normal-form) games, so-called ranking games, in which the payoff to each agent solely depends on his position in a ranking of the agents induced by their actions. This definition is motivated by the observation that in many strategic situations such as parlour games, competitive economic scenarios, and some social choice settings, players are merely interested in performing optimal *relative* to their opponents rather than in absolute measures. A simple but important subclass of ranking games are *single-winner* games where in any outcome one agent wins and all other players lose. We investigate the computational complexity of a variety of common game-theoretic solution concepts in ranking games and deliver hardness results for iterated weak dominance and mixed Nash equilibria when there are more than two players and pure Nash equilibria when the number of players is unbounded. This dashes hope that multi-player ranking games can be solved efficiently, despite the structural restrictions of these games.

1 Introduction

A well-studied subclass of games in game theory consists of strictly competitive games for two players, *i.e.*, games where the interests of both players are diametrically opposed (such as in Chess). These games admit a unique rational solution (the minimax solution) that can be efficiently computed (von Neumann, 1928).[1] Unfortunately, things get much more complicated if there are more than two players. To begin with, the notion of strict competitiveness in multi-player games is not unequivocal. The extension of the common definition for two-player games, which says that the sum of payoffs in all outcomes has to be constant, is meaningless in multi-player games because *any* game can be transformed into a constant-sum game by adding an extra player (with only one action at his disposal) who absorbs the payoffs of the other players.

In this paper, we put forward a new class of multi-player games, called *ranking games*, in which the payoff to each agent depends solely on his position in a ranking of the agents induced by their actions. The formal definition allows each agent to specify his individual preferences over ranks so that

- higher ranks are weakly preferred,
- being first is strictly preferred over being last, and
- agents are indifferent over other players' ranks.

This definition is motivated by the observation that in many games of strategy or competitive economic scenarios, players are merely interested in performing optimal *relative* to their competitors. Besides, one can

[1] However, in the case of chess, the enormous size of the game in normal-form prohibits the efficient computation of an exact solution.

also think of social choice settings where agents strive to determine a complete hierarchy among themselves based on individual preferences that satisfy the conditions listed above.

When moving away from two-player constant-sum games, there are numerous applicable solution concepts. From a computational perspective, an important property of any solution concept is the computational effort required to determine the solution, simply because the intractability of a solution concept renders it useless for large problem instances that do not exhibit additional structure. We study the computational complexity of a variety of common game-theoretic solution concepts in ranking games and deliver hardness results for iterated weak dominance and mixed Nash equilibria when there are more than two players and pure Nash equilibria in games with many players. This dashes hope that multi-player ranking games can be solved efficiently, despite the structural restrictions of these games.

Remarkably, all hardness results hold for *arbitrary* preferences over ranks as long as they meet the requirements listed above. In particular, even simple subclasses like single-winner games (where players only care about winning) or single-loser games (where players only want to avoid losing) are hard to solve.

2 Related Work

Most of the research on game playing in Artificial Intelligence (AI) has focused on two-player games (see, *e.g.*, Marsland & Schaeffer, 1990). As a matter of fact, "in AI, 'games' are usually of a rather specialized kind—what game theorists call deterministic, turn-taking, two-player, zero-sum games of perfect information" (Russell & Norvig, 2003, p. 161). A notable exception are complete information *extensive-form* games, a class of multi-player games for which efficient Nash equilibrium search algorithms have been investigated by the AI community (*e.g.*, Luckhardt & Irani, 1986; Sturtevant, 2004). In extensive-form games, players move consecutively and a *pure* Nash equilibrium is guaranteed to exist (see, *e.g.*, Myerson, 1997). Therefore, the computational complexity of finding equilibria strongly depends on the actual representation of the game (see Section 4.3). Normal-form games are more general than extensive-form games because every extensive-form game can be mapped to a corresponding normal-form game (with potentially exponential blowup), while the opposite is not the case.

In game theory, several classes of "strictly competitive" games have been proposed that maintain some of the nice properties of two-player constant-sum games. For example, Aumann (1961) defines *almost strictly competitive* games as games where a unique value can be obtained by playing strategies from a certain set. Moulin & Vial (1978) introduce a class of games that are strategically equivalent to constant-sum games. The notion of strict competitiveness we consider is remotely related to the notion of *spite* defined by Brandt, Sandholm, & Shoham (2005), where agents aim at maximizing their payoff relative to the payoff of all other agents.

3 Definitions

3.1 Game-Theoretic Foundations

An accepted way to model situations of conflict and social interaction is by means of a *normal-form game* (see, *e.g.*, Myerson, 1997).

Definition 1 (Normal-form game) *A* game in normal-form *is a tuple* $\Gamma = (N, (A_i)_{i \in N}, (p_i)_{i \in N})$ *where* N *is a set of* players *and for each player* $i \in N$, A_i *is a nonempty set of* actions *available to player* i, *and* $p_i : (\times_{i \in N} A_i) \to \mathbb{R}$ *is a function mapping each action profile of the game (*i.e.*, combination of actions) to a real-valued* payoff *for player* i.

1	3		1	2
2	1		3	3

Table 1: Three-player single-winner game. Player 1 chooses rows, player 2 chooses columns, and player 3 chooses matrices. The four dashed boxes denote Nash equilibria.

A combination of actions $s \in \bigtimes_{i \in N} A_i$ is also called a profile of *pure strategies*. This concept can be generalized to *mixed strategy profiles* $s \in S = \bigtimes_{i \in N} S_i$, by letting players randomize over their actions. We have S_i denote the set of probability distributions over player i's actions, or *mixed strategies* available to player i. In the following, we further write $n = |N|$ for the number of players in a game, s_i for the ith strategy in profile s, s_{-i} for the vector of all strategies in s but s_i, and s_i^k for the probability of player i's kth action in strategy profile s. Two-player games are also called *bimatrix games*, and games with rational payoffs are called *rational games*.

3.2 Ranking Games

The situations of social interaction this paper is concerned with are such that outcomes are related to a ranking of the players, *i.e.*, an ordering of the players according to how well they have done in the game relative to one another. We assume that players generally prefer higher ranks over lower ones and that they are indifferent to the ranks of other players. Moreover, we hypothesize that the players entertain qualitative preferences over *lotteries* or probability distributions over ranks (*cf.* von Neumann & Morgenstern, 1947). For example, one player may prefer to be ranked second to having a fifty-fifty chance of being ranked first and being ranked third, whereas another player may judge quite differently. Thus, we arrive at the following definition of the *rank payoff* to a player.

Definition 2 (Rank payoff) *The* rank payoff *of a player i is defined as vector $r_i = (r_i^1, r_i^2, \ldots, r_i^n) \in \mathbb{R}^n$ so that*

$$r_i^k \geq r_i^{k+1} \quad \textit{for all } k \in \{1,2,\ldots,n-1\}, \qquad \textit{and } r_i^1 > r_i^n$$

*(*i.e.*, higher ranks are weakly preferred, and for at least one rank the preference is strict). For convenience, we assume rank payoffs to be normalized so that $r_i^1 = 1$ and $r_i^n = 0$.*

Intuitively, r_i^k represents player i's payoff for being ranked in kth. Building on Definition 2, defining ranking games is straightforward.

Definition 3 (Ranking game) *A* ranking game *is a game where for any strategy profile $s \in S$ there is a permutation $(\pi_1, \pi_2, \ldots, \pi_n)$ of the players so that $p_i(s) = r_i^{\pi_i}$ for all $i \in N$.*

A *binary ranking game* is one where each rank payoff vector only consists of zeros and ones, *i.e.*, each player is equally satisfied up to a certain rank. An important subclass of binary ranking games are games where winning is the only goal of all players.

Definition 4 (single-winner game) *A* Single-winner game *is a ranking game where $r_i = (1, 0, \ldots, 0)$ for all $i \in N$.*

In other words, the outcome space in single-winner games is partitioned into n blocks. When considering mixed strategies, the expected payoff in a single-winner ranking game equals the probability of winning the game. Similar to single-winner games, we can define *single-loser games* (like "musical chairs") as games where all $r_i = (1, \ldots, 1, 0)$.

An example single-winner game with three players is given in Table 1. A convenient way to represent these games is to just denote the index of the winning player for each outcome. Nash equilibria are marked by dashed boxes where a box that spans two outcomes denotes an equilibrium where one player mixes uniformly between his actions.[2] Curiously, there is a fifth equilibrium in this game where all players randomize their actions according to the golden ratio $\phi = (1+\sqrt{5})/2$.

4 Solving Ranking Games

Over the years, game theory has produced a number of solution concepts that identify reasonable or desirable strategy profiles in a given game (see, *e.g.*, Myerson, 1997). The key question of this paper is whether the rather restricted structure of ranking games allows us to compute instances of common solution concepts more efficiently than in general games. For this reason, we focus on solution concepts that are known to be intractable for general games, namely (mixed) *Nash equilibria* (Chen & Deng, 2005; Daskalakis, Goldberg, & Papadimitriou, 2006), *iterated weak dominance* (Conitzer & Sandholm, 2005), and *pure Nash equilibria* in graphical normal form (Gottlob, Greco, & Scarcello, 2005) or circuit form games (Schoenebeck & Vadhan, 2006). We do not cover solution concepts for which efficient algorithms are known to exist such as iterated strong dominance (Conitzer & Sandholm, 2005) or correlated equilibria (Papadimitriou, 2005).

Given the current state of complexity theory (see, *e.g.*, Papadimitriou, 1994), we cannot prove the *actual* hardness of most algorithmic problems, but merely give *evidence* for their hardness. Showing the NP-completeness (or PPAD-completeness) of a problem is commonly regarded as a very strong argument for hardness because it relates the problem to a large class of problems for which no efficient algorithm is known (despite enormous efforts to find such algorithms). When in the following we refer to the hardness of a game we mean the computational hardness of solving the game using a particular solution concept.

4.1 Mixed Nash Equilibria

One of the best-known solution concepts is Nash equilibrium (Nash, 1951). In a Nash equilibrium, no player is able to increase his payoff by *unilaterally* changing his strategy.

Definition 5 (Nash equilibrium) *A strategy profile $s \in S$ is called a* Nash equilibrium *if for each player $i \in N$ and each strategy $s'_i \in S_i$,*

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

A Nash equilibrium is called pure if it is a pure strategy profile.

Let us first consider ranking games with only two players. According to Definition 3, two-player ranking games are games with outcomes $(1,0)$ and $(0,1)$ and thus represent a special subclass of constant-sum games. Nash equilibria of constant-sum games can be found by Linear Programming (Vajda, 1956), for which there is a polynomial time algorithm (Khachiyan, 1979).

For more than two players, we argue by showing that three-player ranking games are at least as hard to solve as general rational bimatrix games. This is sufficient for proving hardness, because n-player ranking games are at least as hard as $(n-1)$-player ranking games (by adding an extra player who only has a single

[2] It seems as if every single-winner game has a non-pure equilibrium, *i.e.*, an equilibrium in which at least one player randomizes. However, this claim has so far tenaciously resisted proof.

action and is ranked last in all outcomes). A key concept in our proof is that of a Nash homomorphism, a notion introduced by Abbott, Kane, & Valiant (2005). We generalize their definition to more than two players.

Definition 6 (Nash homomorphism) *A* Nash homomorphism *is a mapping h from a set of games into a set of games, such that there exists a polynomial-time computable function f that, when given an equilibrium of $h(\Gamma)$, returns an equilibrium of Γ.*

A very simple Nash homomorphism, henceforth called *scale homomorphism*, is one where the payoff of each player is scaled using a positive linear transformation. It is well-known that Nash equilibria are invariant under this kind of operation. A slightly more sophisticated mapping, where outcomes of a bimatrix game are mapped to corresponding three-player subgames, so-called simple cubes, is defined next.

Definition 7 (Simple cube substitution (SCS)) *Let h be a mapping from a set of two-player games to a set of three-player games that replaces every outcome $o = (p_1, p_2, \dots, p_n)$ of the original game Γ with a corresponding three-player subgame $\Gamma'(o)$ of the form*

$o_1(o)$	$o_2(o)$
$o_2(o)$	$o_1(o)$

$o_2(o)$	$o_1(o)$
$o_1(o)$	$o_2(o)$

.

h is called a simple cube substitution (SCS) *if for every o*

- *$\Gamma'(o)$ is a constant-sum game (*i.e., $\sum_i p_i(o_1(o)) = \sum_i p_i(o_2(o))$*),*
- *the average of player i's payoff in $o_1(o)$ and $o_2(o)$ equals p_i, and*
- *there is at least one player that prefers $o_1(o)$ over $o_2(o)$ and one that prefers $o_2(o)$ over $o_1(o)$.*

Lemma 1 *SCS is a Nash homomorphism.*

Proof: First of all, observe that

- there exists a strategy profile s that is a Nash equilibrium of $\Gamma'(o)$ for all o (namely the one where all players randomize uniformly over all actions), and
- s_i a maximin strategy for player i, guaranteeing him at least the equilibrium payoff.

Let h be a SCS, and let a^1, a^2 denote the pair of actions of $h(\Gamma)$ corresponding to an action a of Γ. We claim that for an arbitrary game Γ, an equilibrium s of $h(\Gamma)$ can be mapped to an equilibrium $f(s)$ of Γ by adding the probabilities of the actions of a particular subgame, *i.e.*, $f(s)_i^k = s_i^{2k-1} + s_i^{2k}$. For this, we will argue that a strategy profile s of $h(\Gamma)$ in which the sum of payoffs obtained from a^1 and a^2 is different from the payoff obtained from a in Γ under strategy profile $f(s)$ cannot be a Nash equilibrium (of course, either a^1 or a^2, and hence also a, has to be played with positive probability for this to be possible).

To see that this suffices to show that h is a Nash homomorphism, assume that $f(s)$ is not a Nash equilibrium of Γ, *i.e.*, $p_i(f(s)_{-i}, a_i) > p_i(f(s))$ for some action a_i of some player i, and that the payoff from each action b_i of player i in Γ equals the payoff from b_i^1 and b_i^2 in $h(\Gamma)$. Then, we also have $p_i(s_{-i}, s_i') > p_i(s)$ where s_i' is the strategy that uniformly distributes all weight on actions a^1 and a^2. This means, however, that s is no Nash equilibrium of $h(\Gamma)$.

First, consider a strategy profile s such that for some player i the payoff from actions a_i^1 and a_i^2 in $h(\Gamma)$ is less than the payoff from a_i in Γ. As we have noted above, the latter equals the equilibrium payoff and the

Outcome		Scaled outcome		Ranking subgame			
$(0,0)$	$\longmapsto$	$\left(\frac{1}{2}, \frac{1}{2}\right)$	$\longmapsto$	$[1,3,2]$	$[2,3,1]$	$[2,3,1]$	$[1,3,2]$
				$[2,3,1]$	$[1,3,2]$	$[1,3,2]$	$[2,3,1]$
$(1,0)$	$\longmapsto$	$\left(1, \frac{1}{2}\right)$	$\longmapsto$	$[1,3,2]$	$[1,2,3]$	$[1,2,3]$	$[1,3,2]$
				$[1,2,3]$	$[1,3,2]$	$[1,3,2]$	$[1,2,3]$
$(0,1)$	$\longmapsto$	$\left(\frac{1}{2}, 1\right)$	$\longmapsto$	$[2,3,1]$	$[1,2,3]$	$[1,2,3]$	$[2,3,1]$
				$[1,2,3]$	$[2,3,1]$	$[2,3,1]$	$[1,2,3]$

Table 2: Simple cube substitution mapping from binary bimatrix games to three-player single-loser games

maximin payoff of all subgames at a_i^1 and a_i^2. Player i can thus get a higher payoff by equally distributing the weight on a_i^1 and a_i^2, so that s is not a Nash equilibrium.

In turn, assume that the payoff player i gets from actions a_i^1 and a_i^2 under strategy profile s is higher than the payoff from a_i under $f(s)$. Furthermore, by the previous observation, the payoff player i gets from another pair of actions b_i^1, b_i^2 cannot be smaller than the payoff from b_i. Hence, the overall payoff of player i under s in $h(\Gamma)$ is strictly greater than that under $f(s)$ in Γ. Since every single subgame of $h(\Gamma)$ is constant-sum, there has to be some other player $j \neq i$ who receives strictly less payoff in $h(\Gamma)$ than in Γ, and at least one pair of actions a_j^1, a_j^2 for which this is the case. This means, however, that player j can play the (relative) maximin strategy for a_j^1, a_j^2 to increase his payoff. Again, s cannot be a Nash equilibrium. □

Based on the scale homomorphism and SCS, we now show that there exist Nash homomorphisms mapping rational bimatrix games to three-player ranking games.

Lemma 2 *For any given rank payoff profile, there exists a Nash homomorphism from the set of rational bimatrix games to the set of three-player ranking games.*

Proof: It has been shown by Abbott, Kane, & Valiant (2005) that there is a Nash homomorphism from rational bimatrix games to bimatrix games with payoffs 0 and 1 (called *binary* games in the following). Since a composition of Nash homomorphisms is again a Nash homomorphism, we only need to provide a homomorphism from binary bimatrix games to three-player ranking games. Furthermore, there is no need to map instances of binary games containing outcome $(1,1)$, which is Pareto-dominant and constitutes a pure Nash equilibrium wherever it occurs in a binary game (no player can benefit from deviating). Consequently, such instances are easy to solve and need not be considered in our mapping.

Let $(1, r_i^2, 0)$ be the rank payoff of player i, and let $[i, j, k]$ denote the outcome where player i is ranked first, j is ranked second, and k is ranked last. First of all, consider ranking games where $r_i^2 < 1$ for some player $i \in N$ (this is the set of all ranking games *except* single-loser games). Without loss of generality, let $i = 1$. Then, a Nash homomorphism from binary bimatrix games to the aforementioned class of games can be obtained by first scaling the payoffs according to $(p_1, p_2) \longmapsto \big((1 - r_1^2)p_1 + r_1^2, p_2\big)$, and then adding a third player who only has a single action and whose payoff depends on p_1 and p_2 (but is otherwise irrelevant). Obviously, the latter is also a Nash homomorphism. Outcomes $(0,0)$, $(1,0)$, and $(0,1)$ are hence mapped according to

$$\begin{array}{ccccc} (0,0) & \longmapsto & (r_1^2, 0) & \longmapsto & [3,1,2] \\ (1,0) & \longmapsto & (1,0) & \longmapsto & [1,3,2] \\ (0,1) & \longmapsto & (r_1^2, 1) & \longmapsto & [2,1,3]. \end{array}$$

Interestingly, three-player *single-loser* games with only a single action for some player $i \in N$ are easy to solve because

- there either is a Pareto-dominant outcome (*i.e.*, one where i is ranked last, such that the other players both receive payoff 1), or
- the game is a constant-sum game (*i.e.*, i is *not* ranked last in any outcome, such that the payoffs of the other players always sum up to 1).

Nevertheless, binary games can be mapped to single-loser games if the additional player is able to choose between *two* different actions. We claim that the mapping given in Table 2 represents a SCS from the set of binary bimatrix games to three-player single-loser games. First of all, each payoff p_i of player i in the original binary bimatrix game is transformed according to the scale homomorphism $(p_1, p_2) \longmapsto \big((1+p_1)/2, (1+p_2)/2\big)$. Next, we replace outcomes of the resulting game by three-player single-loser subgames according to the mapping shown in Table 2. It can easily be verified that this mapping satisfies the conditions of Definition 7 and thus resembles a Nash homomorphism. $\square$

We are now ready to present the main result of this section concerning the hardness of computing Nash equilibria of ranking games. Since every normal-form game is guaranteed to possess a Nash equilibrium in mixed strategies (Nash, 1951), the decision problem is trivial. However, the associated *search problem* turned out to be not at all trivial. In fact, it has recently been shown to be PPAD-complete (Chen & Deng, 2005; Daskalakis, Goldberg, & Papadimitriou, 2006). TFNP (total functions in NP) is the class of search problems guaranteed to have a solution. As Daskalakis, Goldberg, & Papadimitriou (2006) put it, "this is precisely NP with an added emphasis on finding a witness." PPAD is a certain subclass of TFNP that is believed not to be contained in P. For this reason, the PPAD-completeness of a particular problem can be seen as strong evidence that there is no efficient algorithm for solving it (*cf.* Daskalakis, Goldberg, & Papadimitriou, 2006).

Theorem 1 *Computing a* Nash equilibrium *of a ranking game with more than two players is PPAD-complete. If there are only two players, equilibria can be found in polynomial time.*

Proof: According to Lemma 2, ranking games are at least as hard to solve as general two player games. Since we already know that solving general two-player games is PPAD-complete (Chen & Deng, 2005), and ranking games cannot be harder than general games, this completes the proof. $\square$

4.2 Iterated Weak Dominance

We will now move to another solution concept, namely the elimination of weakly dominated actions.

Definition 8 (Weak Dominance) *An action $d_i \in A_i$ is said to be* weakly dominated *by strategy $s_i \in S_i$ if*

$$p_i(b_{-i}, d_i) \leq \sum_{a_i \in A_i} s_i(a_i) p_i(b_{-i}, a_i), \qquad \text{for all } b_{-i} \in A_{-i}, \text{ and}$$
$$p_i(\hat{b}_{-i}, d_i) < \sum_{a_i \in A_i} s_i(a_i) p_i(\hat{b}_{-i}, a_i), \qquad \text{for at least one } \hat{b}_{-i} \in A_{-i}.$$

After one or more dominated actions have been removed from the game, other actions may become dominated that were not dominated in the original game, and may themselves be removed. In general, the result of such an iterative process depends on the order in which actions are eliminated, since the elimination of an action may render an action of another player undominated. If only one action remains for each player, we say that the game can be solved by means of iterated weak dominance.

	a_2			
	2	1		
	$\vdots$		$\ddots$	
	2			1
	1	2		
			$\ddots$	
	$\vdots$			
				2
a_1	1		$\cdots$	1

Table 3: Iterated weak dominance solvability in two-player ranking games

Definition 9 *We say that a game is* solvable *by iterated weak dominance if there is some path of eliminations that leaves exactly one action per player. The corresponding computational problem of deciding whether a given game is solvable will be called IWD-SOLVABLE.*

If there are only two players, we can decide IWD-SOLVABLE in polynomial time, which is seen as follows. First of all, we observe that in binary games dominance by a mixed strategy always implies dominance by a pure strategy, so we only have to consider dominance by pure strategies.

Consider a path of iterated weak dominance that ends in a single action profile (a_1,a_2), and without loss of generality assume that player 1 (*i.e.*, the row player) wins in this profile. This implies that player 1 must win in any action profile (a_1,a_2') for $a_2' \in A_2$. For a contradiction, consider the particular action a_2^1 such that player 2 wins in (a_1,a_2^1) and a_2^1 is eliminated last on the path that solves the game. Clearly, a_2^1 could not be eliminated in this case. An elimination by player 1 would also eliminate a_1, while an elimination by player 2 would require another action a_2^2 such that player 2 also wins in (a_1,a_2^2), which contradicts the assumption that a_2^1 is eliminated last.

We thus claim that IWD-SOLVABLE for ranking games with two players can be decided by finding a unique action a_1 of player 1 by which he always wins, and a unique action a_2 of player 2 by which he wins for a maximum number of actions of player 1. This situation is shown Table 3. If such actions do not exist or are not unique, the game cannot be solved by means of iterated weak dominance. If they do exist, we can use a_1 to eliminate all actions a_1' such that player 2 does not win in (a_1',a_2), whereafter a_2 can eliminate all other actions of player 2, until finally a_1 eliminates player 1's remaining strategies and solves the game. Obviously, this can be done in polynomial time.[3]

In order to tackle IWD-SOLVABLE for more than two players, we introduce two additional computational problems related to iterated weak dominance.

Definition 10 *Given an action e, IWD-ELIMINABLE asks whether there is some path of iterated weak dominance elimination that eliminates e. Given a pair of actions e^1 and e^2, IWD-PAIR-ELIMINABLE asks whether there is some path of iterated weak dominance that eliminates both e^1 and e^2.*

We proceed to show hardness of IWD-SOLVABLE for ranking games with more than two players by first showing hardness of IWD-PAIR-ELIMINABLE, and then reducing it to IWD-SOLVABLE.

Lemma 3 *IWD-PAIR-ELIMINABLE is NP-complete for any ranking game with at least three players, even if one player only has a single action, and the two actions to be eliminated belong to the same player.*

[3]Since two-player ranking games are a subclass of constant-sum games, weak dominance and *nice weak dominance* (Marx & Swinkels, 1997) coincide, making iterated weak dominance order independent *up to payoff-equivalent action profiles*. This fact is mirrored by Table 3, since there cannot be a row of 1s and a column of 2s in the same matrix.

$(0,0) \mapsto$	$[3,2,1]$	$[3,1,2]$		$(1,0) \mapsto$	$[1,2,3]$	$[3,1,2]$
	$[3,1,2]$	$[3,2,1]$			$[3,1,2]$	$[1,2,3]$
$(0,1) \mapsto$	$[3,2,1]$	$[2,1,3]$		$(1,1) \mapsto$	$[1,2,3]$	$[2,1,3]$
	$[2,1,3]$	$[3,2,1]$			$[2,1,3]$	$[1,2,3]$

Table 4: Dominance-preserving mapping from binary bimatrix games to three-player ranking games

Proof: Membership in NP is immediate. We can simply guess a sequence of eliminations and then verify in polynomial time that this sequence is valid and eliminates e^1 and e^2.

To show *hardness*, we reduce IWD-ELIMINABLE for games with two players and payoffs 0 and 1, which has recently been shown to be NP-hard (Conitzer & Sandholm, 2005), to IWD-PAIR-ELIMINABLE for ranking games. A game Γ of the former class is mapped to a ranking game Γ' as follows:

- Γ' features the two players of Γ, denoted 1 and 2, and an additional player 3.
- Each action a_i^j of player $i \in \{1,2\}$ in Γ is mapped to *two* actions $a_i^{j,1}$ and $a_i^{j,2}$ in Γ'. Player 3 only has a single action.
- Payoffs of Γ are mapped to rankings of Γ' according to the mapping in Table 4. Again, $[i,j,k]$ denotes the outcome where player i is ranked first, j is ranked second, and k is ranked last.

We claim that for any class of ranking game, *i.e.*, irrespective of the rank payoffs $r_i = (1, r_i^2, 0)$, a particular action a^j in Γ can be eliminated by means of iterated weak dominance if and only if it is possible to eliminate both $a^{j,1}$ and $a^{j,2}$ in Γ' *on a single path*. Without loss of generality, we assume that e belongs to player 1. In the following, we exploit two properties of the outcome mapping in Table 4:

1. If an action $a^{j,1}$ can be eliminated by some other action $a^{k,1}$, then $a^{j,2}$ could at the same time be eliminated by $a^{k,2}$, if $a^{k,2}$ has not been eliminated before. This particularly means that under *non-iterated* weak dominance, $a^{j,1}$ can be eliminated if and only if $a^{j,2}$ can be eliminated.
2. Every pair of a non-eliminable action a^j and another action a^k satisfies one of two conditions. Either, a^j is as least as good as a^k at any index (*i.e.*, action of the other player). Or, a^j is strictly worse than a^k at some index, and strictly better than a^k at another index.

Assume there exists a sequence of eliminations that finally eliminates e in Γ. Then, by Property 1, an arbitrary element of the sequence where a^j eliminates a^k, can be mapped to a pair of successive eliminations in Γ' where $b^{j,1}$ eliminates $a^{k,1}$ and $a^{j,2}$ eliminates $a^{k,2}$. This results in a sequence of eliminations in Γ' ending in the elimination of both e^1 and e^2, e's corresponding actions in Γ'.

Conversely, assume that e cannot be eliminated in Γ, and consider a sequence of eliminations in Γ' leading to the elimination of e^1 and then e^2. We will argue that such a sequence cannot exist. Since e is non-eliminable, Property 2 holds for any action a_1^j restricted to the actions of player 2 that have not been eliminated before e^2. If e is as least as good as a_1^j at any remaining action a_2^k, then, by construction of the payoff mapping and restricted to the remaining actions of Γ', one of e^1 or e^2 is at least as good as any other action of player 1. That means they cannot both be eliminated, a contradiction.

Hence, there must be a pair of actions a_2^k and a_2^ℓ such that e is strictly better than a_1^j at a_2^k and strictly worse than a_1^j at a_2^ℓ. Without loss of generality, we assume that a_2^ℓ is the only index where e is strictly worse. Then, for both e^1 and e^2 to be eliminable, one of $a_2^{\ell,1}$ and $a_2^{\ell,2}$ must have been eliminated before. (Observe

a_3^1

		a_2^1	a_2^2	a_2^3
e^1	Γ	$[3,1,2]$	$[2,1,3]$	$[2,1,3]$
e^2		$[3,1,2]$	$[2,1,3]$	$[2,1,3]$
		$\vdots$ $[2,1,3]$ $\vdots$	$\vdots$ $[2,1,3]$ $\vdots$	$\vdots$ $[2,1,3]$ $\vdots$
a_1^1	$\cdots$ $[1,3,2]$ $\cdots$	$[2,1,3]$	$[2,3,1]$	$[3,1,2]$
a_1^2	$\cdots$ $[1,3,2]$ $\cdots$	$[2,1,3]$	$[3,1,2]$	$[2,1,3]$
a_1^3	$\cdots$ $[2,3,1]$ $\cdots$	$[3,1,2]$	$[1,3,2]$	$[3,1,2]$

a_3^2

		a_2^1	a_2^2	a_2^3
	$\vdots$ $[1,2,3]$ $\vdots$	$\vdots$ $[3,2,1]$ $\vdots$	$\vdots$ $[1,2,3]$ $\vdots$	$\vdots$ $[1,2,3]$ $\vdots$
a_1^1	$\cdots$ $[1,2,3]$ $\cdots$	$[1,2,3]$	$[3,2,1]$	$[1,2,3]$
a_1^2	$\cdots$ $[3,2,1]$ $\cdots$	$[1,2,3]$	$[1,2,3]$	$[3,2,1]$
a_1^3	$\cdots$ $[1,2,3]$ $\cdots$	$[1,2,3]$	$[3,2,1]$	$[3,2,1]$

Table 5: Three-player ranking game Γ' used in the proof of Theorem 2

that this elimination further requires $r_1^1 = r_1^2$.) On the other hand, it must not be possible to eliminate both $a_2^{\ell,1}$ and $a_2^{\ell,2}$, since otherwise by Property 1, a_2^ℓ could be eliminated in Γ, whereafter a_1^j could eliminate e. We thus get dominance according to Property 2, similar to the one for e described above. Hence, there again has to be an action $a_1^m \neq e$ such that exactly one of $a_1^{m,1}$ and $a_1^{m,2}$ has been eliminated (and the other one could not have been eliminated). This condition can be traced backwards through the sequence of eliminations that lead to the elimination of e^2. The first elimination in this sequence, however, is in terms of non-iterated dominance, and by Property 1 there can be no pair of actions such that exactly one of them can be eliminated. This is a contradiction. $\square$

We are now ready to state the main result of this section.

Theorem 2 *Deciding whether a ranking game with more than two players is solvable by* iterated weak dominance *is NP-complete. When there are only two players, this can be decided in polynomial time.*

Proof: Membership in NP is immediate. We can simply guess a sequence of eliminations and then verify that this sequence is valid and leaves only one action per player.

For *hardness*, we reduce IWD-PAIR-ELIMINABLE for ranking games with three players, where one of the players has only one action, to IWD-SOLVABLE for ranking games with three players. Therefore, an instance Γ of the former class is mapped to an instance Γ' of the latter as follows:

- All players' actions from Γ are also part of the new instance Γ', including the two actions e^1 and e^2 to be eliminated. The payoffs (*i.e.*, rankings) for the corresponding outcomes remain the same.
- We further add two additional actions a_1^1 and a_1^2 for player 1, two actions a_2^1 and a_2^2 for player 2 and one action for player 3, who now has actions a_3^1 and a_3^2. The payoffs for the outcomes induced by these new actions are given in Table 5.

We claim that, for arbitrary values of r_i^2, e^1 and e^2 can be eliminated in Γ if and only if Γ' can be solved by means of iterated weak dominance.

Assume IWD-PAIR-ELIMINABLE for Γ has a solution. Then, the same sequence of eliminations that eliminates both e^1 and e^2 can also be executed in Γ', because player 1 is ranked equally in all rows of Γ at a_2^i, and player 2 is ranked equally in all columns of Γ at a_1^i for $i = 1,2,3$. As soon as e^1 and e^2 have been eliminated, let a_2^2 be eliminated by a_2^1, which is strictly preferred at (a_1^2, a_3^1) and ranks player 2 equally at any other position. Next, use a_1^1 to eliminate all other rows, which are strictly worse at either (a_2^1, a_3^2) or (a_2^3, a_3^2)

and strictly better at no position. Finally, let a_3^2 be eliminated by a_3^1, being strictly better at a_2^3, and solve the game by eliminating the remaining actions of player 2 by a_2^1.

Conversely, assume that there exists no path of iterated weak dominance elimination that eliminates both e^1 and e^2. We will argue that, as long as either e^1 or e^2 is still there, (i) the newly added actions cannot eliminate any of the original actions and (ii) cannot themselves be eliminated (either by original or new actions). As we have seen above, this also means that the newly added actions have no influence on eliminations between actions of Γ. As for player 1, the newly added actions are strictly worse than any of the original ones at (a_2^1, a_3^2), and strictly better at either (a_2^2, a_3^2) or (a_2^3, a_3^2). a_1^1 is strictly better than a_1^2 and a_1^3 at (a_2^3, a_3^2), and strictly worse at either (a_2^2, a_3^2) or (a_2^2, a_3^1). a_1^2 is strictly better or worse than a_1^3 at the original actions of player 2 and at a_3^1 and a_3^2, respectively. Analogously, for player 2, the newly added actions are strictly worse than any of the original ones at (a_1^3, a_3^1), and strictly better at either (a_1^1, a_3^1) or (a_1^2, a_3^1). a_2^1 is strictly better than a_2^2 and a_2^3 at either (a_1^1, a_3^1) or (a_1^2, a_3^1), and strictly worse at both (e^1, a_3^1) and (e^2, a_3^1). a_2^2 is strictly better or worse than a_2^3 at (a_1^1, a_3^1) and (a_1^2, a_3^1), respectively. Finally, a_3^1 is strictly better than a_3^2 at (a_1^1, a_2^3), and strictly worse at (a_1^2, a_2^3).

This completes the reduction, showing NP-hardness of IWD-SOLVABLE for ranking games with at least three players. □

4.3 Pure Nash Equilibria in Games with Many Players

A very important subset of Nash equilibria are those where players do not have to randomize, *i.e.*, every player deterministically chooses one particular action. These so-called *pure* Nash equilibria (*cf.* Definition 5) can be found efficiently by simply checking every action profile. However, as the number of players increases, the number of profiles to check (as well as the normal-form representation of the game) grows exponentially. An interesting question is whether pure equilibria can be computed efficiently given a *concise* representation of a game (using space polynomial in n). For some concise representations like graphical games with bounded neighborhood, where the payoff of any player only depends on a constant number of neighbors (Gottlob, Greco, & Scarcello, 2005), or circuit games, where the outcome function is computed by a Boolean circuit of polynomial size (Schoenebeck & Vadhan, 2006), deciding the existence of a pure equilibrium has been shown to be NP-complete.

It turns out that graphical games are of very limited use for representing ranking games. If two players are not connected by the neighborhood relation, either directly or via a common player in their neighborhood, then their payoffs are completely independent from each other. For a single-winner game with the reasonable restriction that every player wins in at least one outcome, this implies that there must be one designated player who decides which player wins the game. Similar properties hold for arbitrary ranking games.

We proceed by showing NP-completeness of deciding whether there is a pure Nash equilibrium in ranking games with *efficiently computable outcome functions* which is one of the most general representations of multi-player games one might think of. Please note that in contrast to Theorems 1 and 2, we now fix the number of actions and let the number of players increase.

Theorem 3 *Deciding the existence of a pure Nash equilibrium in a ranking game with many players and a polynomial-time computable outcome function is NP-complete, even if the players only have two actions at their disposal.*

Proof: Since we can check in polynomial time whether a particular player strictly prefers one rank over another, *membership* in NP is immediate. We can guess an action profile s and verify in polynomial time whether s is a Nash equilibrium. For the latter, we check for each player $i \in N$ and for each action $a \in A_i$ whether $p_i(s_{-i}, a) \leq p_i(s)$.

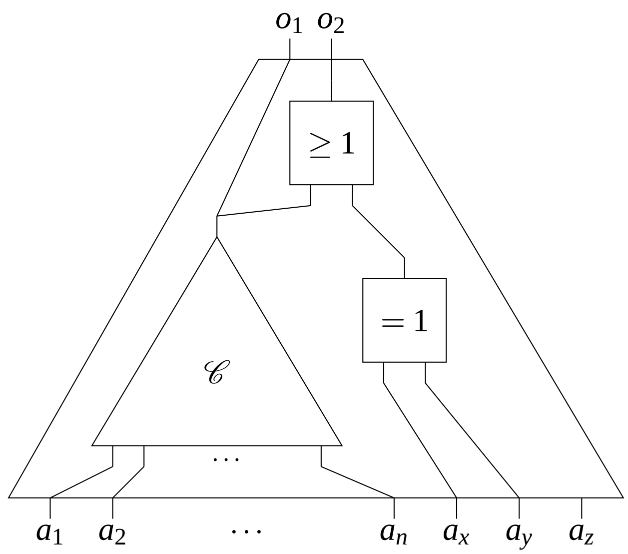

Figure 1: Structure of the Boolean circuit used in the proof of Theorem 3

For *hardness*, recall that circuit satisfiability (CSAT), *i.e.*, deciding whether for a given Boolean circuit $\mathscr{C}$ with n inputs and 1 output there exists an input such that the output is *true*, is NP-complete (see, *e.g.*, Papadimitriou, 1994). We define a game Γ in circuit form for a Boolean circuit $\mathscr{C}$, providing a polynomial-time reduction of satisfiability of $\mathscr{C}$ to the problem of finding a pure Nash equilibrium in Γ.

Let n be the number of inputs of $\mathscr{C}$. We define game Γ with $n+3$ players as follows:

- Let $N = \{1,\ldots,n\} \cup \{x,y,z\}$, and $A_i = \{0,1\}$ for all $i \in N$.
- The outcome function of Γ is computed by a Boolean circuit that takes $n+3$ bits of input $i = (a_1,\ldots,a_n,a_x,a_y,a_z)$, corresponding to the actions of the players in N, and computes 2 bits of output $o = (o_1,o_2)$, given by $o_1 = \mathscr{C}(a_1,\ldots,a_n)$ and $o_2 = (o_1\ OR\ (a_x\ XOR\ a_y))$. Player z is incapable of influencing the outcome of the game.

  Figure 1 illustrates the structure of the circuit.
- The possible outputs of the circuit are identified with permutations (*i.e.*, rankings) of the players in N such that
  - the permutation π_{00} corresponding to $o = (0,0)$ ranks x first, z second, and y last,
  - the permutation π_{01} corresponding to $o = (0,1)$ ranks y first, z second, and x last, and
  - the permutation π_{11} corresponding to $o = (1,1)$ ranks z first, x second, and y last.

  All other players are ranked equally in all three permutations. It should be noted that no matter how permutations are actually encoded as strings of binary values, the encoding of the above permutations can always be computed using a polynomial number of gates.

We claim that, for arbitrary rank payoffs r, Γ has a pure Nash equilibrium if and only if $\mathscr{C}$ is satisfiable. This is seen as follows:

- By construction, the outcome of the game is π_{11} if and only if players $1,\ldots,n$ play a satisfying assignment of $\mathscr{C}$.

- Given an action profile resulting in outcome π_{11}, only a player in $\{1,\ldots,n\}$ could possibly change the outcome of the game by changing his action. However, these players are ranked equally in all the possible outcomes, so none of them can get a higher payoff by doing so. Thus, π_{11} is a Nash equilibrium.
- By construction, both x and y can change between outcomes π_{00} and π_{01} by changing their individual action.
- Since every player strictly prefers being ranked first over being ranked last, x strictly prefers outcome π_{00} over π_{01}, while y strictly prefers π_{01} over π_{00}. Thus, neither π_{00} nor π_{01} are Nash equilibria, since one of x and y could change his action to get a higher payoff.

This completes the reduction. □

5 Conclusion

We proposed a new class of games, so-called ranking games, that model settings in which players are merely interested in performing at least as good as their opponents. Despite the structural simplicity of these games, various solution concepts turned out to be hard to compute, namely mixed equilibria and iterated weak dominance in games with more than two players and pure equilibria in games with an unbounded number of players. As a consequence, the mentioned solution concepts appear to be of limited use in large instances of ranking games that do not possess additional structure. This underlines the importance of alternative, efficiently computable, solution concepts for ranking games.

Acknowledgements

The authors thank Paul Harrenstein, Markus Holzer, Samuel Ieong, Eugene Nudelman, and Rob Powers for valuable comments.

This material is based upon work supported by the Deutsche Forschungsgemeinschaft under grants BR 2312/1-1 and BR 2312/3-1, and by the National Science Foundation under ITR grant IIS-0205633.

References

Abbott, T.; Kane, D.; and Valiant, P. 2005. On the complexity of two-player win-lose games. In *Proceedings of the 46th Symposium on Foundations of Computer Science (FOCS)*, 113–122. IEEE Computer Society Press.

Aumann, R. 1961. Almost strictly competitive games. *Journal of the Society of Industrial and Applied Mathematics* 9(4):544–550.

Brandt, F.; Sandholm, T.; and Shoham, Y. 2005. Spiteful bidding in sealed-bid auctions. In Gmytrasiewicz, P., and Parsons, S., eds., *Proceedings of the 7th IJCAI Workshop on Game Theoretic and Decision Theoretic Agents (GTDT)*.

Chen, X., and Deng, X. 2005. Settling the complexity of 2-player Nash-equilibrium. Technical Report TR05-140, Electronic Colloquium on Computational Complexity (ECCC).

Conitzer, V., and Sandholm, T. 2005. Complexiy of (iterated) dominance. In *Proceedings of the 6th ACM Conference on Electronic Commerce (ACM-EC)*, 88–97. ACM Press.

Daskalakis, C.; Goldberg, P.; and Papadimitriou, C. 2006. The complexity of computing a Nash equilibrium. In *Proceedings of the 38th Annual ACM Symposium on the Theory of Computing (STOC)*. ACM Press.

Gottlob, G.; Greco, G.; and Scarcello, F. 2005. Pure Nash equilibria: Hard and easy games. *Journal of Artificial Intelligence Research* 24:195–220.

Khachiyan, L. 1979. A polynomial algorithm in linear programming. *Soviet Mathemathics Doklady* 20:191–194.

Luckhardt, C., and Irani, K. 1986. An algorithmic solution of *n*-person games. In *Proceedings of the 5th National Conference on Artificial Intelligence (AAAI)*, 158–162. AAAI Press.

Marsland, A. T., and Schaeffer, J., eds. 1990. *Computers, Chess, and Cognition*. Springer-Verlag.

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

Moulin, H., and Vial, J.-P. 1978. Strategically zero-sum games: The class of games whose completely mixed equilibria cannot be improved upon. *International Journal of Game Theory* 7(3–4):201–221.

Myerson, R. B. 1997. *Game Theory: Analysis of Conflict*. Harvard University Press.

Nash, J. F. 1951. Non-cooperative games. *Annals of Mathematics* 54(2):286–295.

von Neumann, J., and Morgenstern, O. 1947. *The Theory of Games and Economic Behavior*. Princeton University Press, 2nd edition.

von Neumann, J. 1928. Zur Theorie der Gesellschaftspiele. *Mathematische Annalen* 100:295–320.

Papadimitriou, C. H. 1994. *Computational Complexity*. Addison-Wesley.

Papadimitriou, C. H. 2005. Computing correlated equilibria in multi-player games. In *Proceedings of the 37th Annual ACM Symposium on the Theory of Computing (STOC)*, 49–56. ACM Press.

Russell, S. J., and Norvig, P. 2003. *Artificial Intelligence. A Modern Approach*. Prentice Hall, 2nd edition.

Schoenebeck, G., and Vadhan, S. 2006. The computational complexity of Nash equilibria in concisely represented games. In *Proceedings of the 7th ACM Conference on Electronic Commerce (ACM-EC)*. ACM Press.

Sturtevant, N. 2004. Current challenges in multi-player game search. In *Proceedings of the 4th International Conference on Computers and Games (CG)*, volume 3846 of *Lecture Notes in Computer Science (LNCS)*. Springer-Verlag.

Vajda, S. 1956. *Theory of Games and Linear Programming*. Wiley.