

Ranking Games

Felix Brandt^{a,*} Felix Fischer^a Paul Harrenstein^a
Yoav Shoham^b

^a*Institut für Informatik, Universität München
Oettingenstr. 67, 80538 München, Germany*

^b*Computer Science Department, Stanford University
353 Serra Mall, Stanford CA 94305, USA*

Abstract

The outcomes of many strategic situations such as parlor games or competitive economic scenarios are rankings of the participants, with higher ranks generally at least as desirable as lower ranks. Here we define *ranking games* as a class of n -player normal-form games with a payoff structure reflecting the players' von Neumann-Morgenstern preferences over their individual ranks. 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 equilibrium when there are more than two players, and for pure Nash equilibrium when the number of players is unbounded but the game is described succinctly. This dashes hope that multi-player ranking games can be solved efficiently, despite their profound structural restrictions. Based on these findings, we provide matching upper and lower bounds for three comparative ratios, each of which relates two different solution concepts: *the price of cautiousness*, *the mediation value*, and *the enforcement value*.

Keywords: Multi-Agent Systems, Game Theory, Strict Competitiveness, n -Player Games, Solution Concepts, Computational Complexity

1 Introduction

The situations studied by the theory of games may involve different levels of antagonism. On the one end of the spectrum are games of pure coordination, on the other

* Corresponding author; Telephone: +49 89 2180 9406; Fax: +49 89 2180 9338

Email addresses: brandtf@tcs.ifi.lmu.de (Felix Brandt),
fischerf@tcs.ifi.lmu.de (Felix Fischer), harrenst@tcs.ifi.lmu.de (Paul
Harrenstein), shoham@cs.stanford.edu (Yoav Shoham).

those in which the players' interests are diametrically opposed. In this paper, we put forward a new class of competitive multi-player games whose outcomes are *rankings* of the players, *i.e.*, orderings of the players representing how well they have done in the game relative to one another. We assume players to weakly prefer a higher rank over a lower one and to be indifferent as to the other players' ranks. This type of situation is commonly encountered in parlor games, competitions, patent races, competitive resource allocation domains, social choice settings, or any other strategic situation where players are merely interested in performing optimal *relative* to their opponents rather than in absolute measures. Formally, ranking games are defined as normal-form games in which the payoff function represents the players' von Neumann-Morgenstern preferences over lotteries over rankings. A noteworthy special case of particular relevance to game playing in AI are *single-winner games* where in any outcome one player wins and all others lose.

While two-player ranking games form a subclass of zero-sum games, no such relationship holds for ranking games with more than two players. Moreover, whereas the notion of a ranking is most natural in multi-player settings, this seems to be less so for the requirement that the sum of payoffs in all outcomes be constant, as *any* game can be transformed into a constant-sum game by merely introducing an additional player (with only one action at his disposal) who absorbs the payoffs of the other players (von Neumann and Morgenstern, 1947).

As with games in which both contrary and common interests prevail, it turns out that solving ranking games tends to become considerably more complicated as soon as more than two players are involved. The maximin solution does not unequivocally extend to general n -player games and numerous alternative solution concepts have been proposed to cope with this type of situation. None of them, however, seems to be as compelling as maximin is for two-player zero-sum games. In this paper we study and compare the properties of a variety of solution concepts in ranking games. The results of this paper fall into two different categories. First, we investigate the complexity of a number of computational problems related to common solution concepts in ranking games, particularly *Nash equilibrium* and *iterated weak dominance*. Second, we study a number of *comparative ratios* in ranking games, each of which relates two different solution concepts: *the price of cautiousness*, *the mediation value*, and *the enforcement value*.

The computational effort required to determine a solution is obviously a very important property of any solution concept. If computing a solution is intractable, the solution concept is rendered virtually useless for large problem instances that do not exhibit additional structure. The importance of this aspect has by no means escaped the attention of game theorists. In an interview with Eric van Damme (1998), Robert Aumann claimed: "My own viewpoint is that, *inter alia*, a solution concept must be calculable, otherwise you are not going to use it." It has subsequently been argued that this still holds if one subscribes to a purely descriptive view of solution concepts: "I believe that the complexity of equilibria is of fundamental importance

in game theory, and not just a computer scientist’s afterthought. Intractability of an equilibrium concept would make it implausible as a model of behavior” (Papadimitriou, 2005). In computational complexity theory, the distinction between tractable and intractable problems is typically one between membership in the class P of problems that can be solved in time polynomial in the size of the problem instance versus hardness for the class NP of problems a solution of which can be verified efficiently. A third class that will play an important role in the context of this paper is PPAD. Problems in PPAD are guaranteed to possess a solution, and emphasis is put on actually finding it. Given the current state of complexity theory, we cannot prove the *actual* intractability of most algorithmic problems, but merely give *evidence* for their intractability. NP-hardness of a problem is commonly regarded as very strong evidence against computational tractability 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. To some extent, the same reasoning can also be applied to PPAD-hardness.

We study the computational complexity of common game-theoretic solution concepts in ranking games and deliver NP-hardness and PPAD-hardness results, respectively, for iterated weak dominance and (mixed) Nash equilibria when there are more than two players, and an NP-hardness result for pure Nash equilibria in games with an unbounded number of players. This dashes hope that multi-player ranking games can be solved efficiently, despite their profound structural restrictions. Remarkably, all hardness results hold for arbitrary preferences over ranks, provided they meet the requirements listed above. Accordingly, even very restricted subclasses of ranking games such as *single-winner games*—in which players only care about winning—or *single-loser games*—in which players merely wish not to be ranked last—are computationally hard to solve.

By contrast, *maximin strategies* (von Neumann, 1928) as well as *correlated equilibria* (Aumann, 1974) are known to be computationally easy via linear programming for any class of games. Against the potency of these concepts, however, other objections can be brought in. Playing a maximin strategy is extremely defensive and a player may have to forfeit a considerable amount of payoff in order to guarantee his security level. Correlation, on the other hand, may not be feasible in all practical applications, and may fail to provide an improvement of social welfare in restricted classes of games (Moulin and Vial, 1978). Thus, we come to consider the following comparative ratios in an effort to facilitate the quantitative analysis of solution concepts in ranking games:

- the *price of cautiousness*, *i.e.*, the ratio between an agent’s minimum payoff in a Nash equilibrium and his security level
- the *mediation value*, *i.e.*, the ratio between the social welfare obtainable in the best correlated equilibrium vs. the best Nash equilibrium, and
- the *enforcement value*, *i.e.*, the ratio between the highest obtainable social welfare and that of the best correlated equilibrium.

Each of these values obviously equals 1 in the case of *two-player* ranking games, as these form a subclass of constant-sum games. Accordingly, the interesting question to ask concerns the bounds of these values for ranking games with more than two players.

2 Introductory Example

To illustrate the issues addressed in this paper, consider a situation in which Alice, Bob, and Charlie are to choose a winner from among themselves by means of the following protocol. Each of them is either to raise or not to raise their hand; they are to do so simultaneously and independently of one another. Alice wins if the number of hands raised, including her own, is odd, whereas Bob is victorious if this number equals two. Should nobody raise their hand, Charlie wins. The normal-form of this game is shown in Figure 1. What course of action would you recommend to Alice? There is a Nash equilibrium in which Alice raises her hand, another one in which she does not raise her hand, and still another one in which she randomizes uniformly between these two options. In the only pure, *i.e.*, non-randomized, equilibrium of the game, Alice does not raise her hand. If the latter were to occur, we must assume that Alice believes that Bob will raise his hand and Charlie will not. This assumption, however, is unreasonably strong as no such beliefs can be derived from the mere description of the game. Moreover, both Bob and Charlie may deviate from their respective strategies to *any* other strategy without decreasing their chances of winning. After all, they cannot do any worse than losing.

This points at a weakness of pure Nash equilibrium as a solution concept inherent in ranking games, since in any outcome some player has to be ranked last. On the other hand, it is very well possible that *all* the actions in the support of a *mixed* equilibrium yield each player a strictly higher expected payoff than any action not in the support, mitigating the phenomenon mentioned above. In other words, such equilibria can be *quasi-strict*, a property no pure equilibrium in a ranking game has. While quasi-strict equilibria may fail to exist in ranking games with more than two players (see Figure 4), we conjecture, and prove for certain sub-cases, that every single-winner game possesses at least one *non-pure* equilibrium, *i.e.*, an equilibrium where at least one player randomizes. We note without proof that this property fails to hold for general ranking games.

Returning to our example, it is unclear which strategy would maximize Alice's chances of winning. By playing her maximin strategy, Alice can *guarantee* a particular probability of winning, her so-called *security level*, no matter which actions her opponents choose. Alice's security level in this particular game is 0.5 and can be obtained by randomizing uniformly between both actions. The same expected payoff is achieved in the worst quasi-strict equilibrium of the game where Alice and Charlie randomize uniformly and Bob invariably raises his hand (see Figure 1).

		c^1			c^2	
		b^1	b^2		b^1	b^2
a^1	3	1		1	2	
a^2	1	2		2	1	

Fig. 1. Three-player single-winner game. Alice (1) chooses row a^1 or a^2 , Bob (2) chooses column b^1 or b^2 , and Charlie (3) chooses matrix c^1 or c^2 . Outcomes are denoted by the winner’s index. The dashed square marks the only pure Nash equilibrium. Dotted rectangles mark a quasi-strict equilibrium in which Alice and Charlie randomize uniformly over their respective actions.

3 Related Work

Game playing research in AI has largely focused on two-player games (see, *e.g.*, Marsland and 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 and Norvig, 2003, p. 161). Notable exceptions include *cooperative games* in the context of coalition formation (see, *e.g.*, Sandholm et al., 1999) and perfect 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 and Irani, 1986; Sturtevant, 2004). In extensive-form games, players move consecutively and a *pure* (so-called subgame perfect) Nash equilibrium is guaranteed to exist (see, *e.g.*, Myerson, 1991). Therefore, the computational complexity of finding equilibria strongly depends on the actual representation of the game (see also Section 6.3). Normal-form games are more general than (perfect-information) 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 proposals for broader classes of games that maintain some of the quintessential properties of two-player constant-sum games have been made. Aumann (1961) defines *almost strictly competitive* games as the class of two-player games in which a pair of strategies is an equilibrium point, *i.e.*, no player can increase his payoff by unilaterally changing his strategy, if and only if it is a so-called twisted equilibrium point, *i.e.*, no player can decrease the payoff of his opponent. These games permit a set of optimal strategies for each player and a unique value that is obtained whenever a pair of such strategies is played. Moulin and Vial (1978) call a game *strategically zero-sum* if it is best-response equivalent to a zero-sum game. In the case of two players, and only in this case, one obtains exactly the class of games for which no completely mixed equilibrium can be improved upon by a correlated equilibrium. A game is *unilaterally competitive*, as defined by Kats and Thisse (1992), if any deviation by a player that (weakly) increases his own payoff must (weakly) decrease the payoffs of all other players. Unilaterally competitive

games retain several interesting properties of two-player constant-sum games in the n -player case: all equilibria yield the same payoffs, equilibrium strategies are interchangeable, and, the set of equilibria is convex provided that some mild conditions hold. It was later shown by Wolf (1999) that *pure* Nash equilibria of n -player unilaterally competitive games are always profiles of maximin strategies. When there are just two players, all of the above classes contain constant-sum games and thus two-player ranking games. Neither is contained in the other in the n -player case. The notion of competitiveness as embodied in ranking games is remotely related to *spitefulness* (Morgan et al., 2003; Brandt et al., 2007), where agents aim at maximizing their payoff relative to the payoff of all other agents.

Most work on comparative ratios in game theory has been inspired by the literature on the *price of anarchy* (Koutsoupias and Papadimitriou, 1999; Roughgarden, 2005), *i.e.*, the ratio between the highest obtainable social welfare and that of the best Nash equilibrium. Similar ratios for correlated equilibria, the *value of mediation*, *i.e.*, the ratio between the social welfare obtainable in the best correlated equilibrium vs. the best Nash equilibrium and the *enforcement value*, *i.e.*, the ratio between the highest obtainable social welfare and that of the best correlated equilibrium, were introduced by Ashlagi et al. (2005). It is known that the mediation value of strategically zero-sum games is 1 and that of almost strictly competitive games is greater than 1, showing that correlation can be beneficial even in games of strict antagonism (Raghavan, 2002). To the best of our knowledge, Tennenholtz (2002) was the first to conduct a quantitative comparison of Nash equilibrium payoffs and security levels. This work is inspired by an intriguing example game due to Aumann (1985), in which the only Nash equilibrium yields each player no more than his security level although the equilibrium strategies are different from the maximin strategies. In other words, the equilibrium strategies yield security level payoffs without guaranteeing them.

4 The Model

A *game form* is a quadruple $(N, (A_i)_{i \in N}, \Omega, g)$, where N is a finite non-empty set of players, A_i a finite and non-empty set of *actions* available to player i , Ω a set of outcomes, and $g: \times_{i \in N} A_i \rightarrow \Omega$ an *outcome function* mapping each action profile to an outcome in Ω . The set $\times_{i \in N} A_i$ of action profiles is denoted by A . We assume that each player entertain preferences over lotteries over Ω that comply with the von Neumann-Morgenstern axioms (von Neumann and Morgenstern, 1947). Thus, the preferences of each player i can be represented by a real valued payoff function p_i on Ω . We arrive at the following definition of a *normal-form game*.

Definition 1 (Game in normal form) A game in normal form Γ is given by a quintuple $(N, (A_i)_{i \in N}, \Omega, g, (p_i)_{i \in N})$ where $(N, (A_i)_{i \in N}, \Omega, g)$ is a game form and each $p_i: \Omega \rightarrow \mathbb{R}$ is a real valued payoff function.

We generally assume the payoff functions p_i to be extended so as to apply directly to action profiles $a \in A$ by setting $p_i(a) = p_i(g(a))$.

We say a game is *rational* if for all $i \in N$ and all $a \in A$, $p_i(a) \in \mathbb{Q}$. A game is *binary* if for all $i \in N$ and all $a \in A$, $p_i(a) \in \{0, 1\}$. A game with two players will also be referred to as a *bimatrix game*. Unless stated otherwise, we will henceforth assume that every player has at least two actions. Subscripts will be used to identify the player to which an action belongs, superscripts to index the actions of a particular player. For example, we write a_i for a typical action of player i and a_i^j for his j th action. For better readability, we also use lower case roman letters from the beginning of the alphabet to denote the players' actions in such a way that $a^j = a_1^j$, $b^j = a_2^j$, $c^j = a_3^j$, and so forth.

The concept of an action profile can be generalized to that of a *mixed strategy profile* by letting players randomize over their actions. We have S_i denote the set $\Delta(A_i)$ of probability distributions over player i 's actions, the *mixed strategies* available to player i , and S the set $\times_{i \in N} S_i$ of mixed strategy profiles with s as typical element. Payoff functions naturally extend to mixed strategy profiles, and we will frequently write $p_i(s)$ for the *expected* payoff of player i , and $p(s)$ for the *social welfare* $\sum_{i \in N} p_i(s)$ under the strategy profile s . We have n stand for the number $|N|$ of players. In the following, A_{-i} and S_{-i} denote the set of action profiles for all players but i and the set of strategy profiles for all players but i , respectively. We use s_i for the i th strategy in profile s and s_{-i} for the vector of all strategies in s but s_i . Furthermore, $s(a_i)$ and $s_i(a_i)$ stand for the probability player i assigns to action a_i in strategy profile s or strategy s_i , respectively. The *pure* strategy s_i such that $s_i(a_i) = 1$ we also denote by a_i whenever this causes no confusion. Moreover, we use (s_{-i}, t_i) to refer to the strategy profile obtained from s by replacing s_i by t_i . For better readability we usually avoid double parentheses and write, e.g., $p(s_{-i}, t_i)$ instead of $p((s_{-i}, t_i))$.

4.1 Rankings and Ranking Games

A *ranking game* is a normal-form game whose outcomes are *rankings* of its players. A ranking indicates how well each player has done relative to the other players in the game. Formally, a ranking $r = [r_1, \dots, r_n]$ is an ordering of the players in N in which player r_1 is ranked first, player r_2 ranked second, and so forth, with player r_n ranked last. Obviously, this limits the number of possible outcomes to $n!$ irrespective of the the number of actions the players have at their disposal. The set of rankings over a set N of players we denote by R_N . A game form $(N, (A_i)_{i \in N}, \Omega, g)$ is a *ranking game form* if the set of outcomes is given by the set of rankings of the players, i.e., if $\Omega = R_N$.

We assume that all players weakly prefer higher ranks over lower ranks, and strictly

prefer being ranked first to being ranked last. Furthermore, each player is assumed to be indifferent as to the ranks of the other players. Even so, a player may prefer to be ranked second for certain to having a fifty-fifty chance of being ranked first or being ranked third, whereas other players may judge quite differently. Accordingly, we have a *rank payoff function* $p_i: R_N \rightarrow \mathbb{R}$ represent player i 's von Neumann-Morgenstern preferences over lotteries over R_N . For technical convenience, we normalize the payoffs to the unit interval $[0, 1]$. Formally, a rank payoff function p_i over R_N satisfies the following three conditions for all rankings $r, r' \in R_N$:

- (i) $p_i(r) \geq p_i(r')$, if $r_k = r'_m = i$ and $k \leq m$,
- (ii) $p_i(r) = 1$, if $i = r_1$, and
- (iii) $p_i(r) = 0$, if $i = r_n$.

We are now in a position to formally define the concept of a *ranking game*.

Definition 2 (Ranking game) *A normal form game $\Gamma = (N, (A_i)_{i \in N}, \Omega, g, (p_i)_{i \in N})$ is a ranking game if Ω is the set R_N of rankings over N and each $p_i: R_N \rightarrow \mathbb{R}$ is a rank payoff function over R_N .*

Condition (i) above implies that a player's payoff for a ranking r only depends on the rank assigned to him in r . Accordingly, for $1 \leq k \leq n$, we have p_i^k denote the *unique* payoff player i obtains in any ranking r in which i is ranked k th. The rank payoff function of player i can then conveniently and compactly be represented by his *rank payoff vector* $\vec{p}_i = (p_i^1, \dots, p_i^n)$.

In a *binary ranking game*, a player is completely satisfied up to a certain rank, and not satisfied at all for any lower rank. The expected payoff of a player given a strategy profile can then be taken as his chances of being satisfied. Thus, the use of expected utility, and thus randomized strategies, is justified without relying on the von Neumann-Morgenstern axioms (see also Aumann, 1987). An interesting subclass of binary ranking games are so-called *single-winner games*, in which all players are only interested in being ranked first. Formally, a single-winner game is a ranking game in which $\vec{p}_i = (1, 0, \dots, 0)$ for all $i \in N$. When considering mixed strategies, the expected payoff in a single-winner ranking game equals the probability of winning. Analogous to single-winner games, we can define *single-loser games* as ranking games in which the players' only concern is not to be ranked last, as for instance in a round of musical chairs. Formally, single-loser games are ranking games where $\vec{p}_i = (1, \dots, 1, 0)$ for each player i . For an example illustrating the definitions of a ranking game form and a ranking game the reader is referred to Figures 2 and 3, respectively.

At this point, a remark as to the relationship between ranking games and n -player constant-sum games is in order. By virtue of conditions (ii) and (iii), two-player ranking games constitute a subclass of zero-sum games. If more than two players are involved, however, any such relation with n -person constant-sum games no

		c^1		c^2	
		b^1	b^2	b^1	b^2
a^1	[1, 3, 2]	[2, 1, 3]	[3, 2, 1]	[3, 1, 2]	
a^2	[2, 3, 1]	[3, 2, 1]	[2, 1, 3]	[1, 3, 2]	

Fig. 2. A $2 \times 2 \times 2$ ranking game form. One player chooses rows, another columns, and a third matrices. Each combination of actions results in a ranking. For example, action profile (a^2, b^2, c^2) leads to the row player 1 being ranked first, the matrix player 3 second and the column player 2 third.

		c^1		c^2	
		b^1	b^2	b^1	b^2
a^1	(1, 0, 1)	($\frac{1}{2}$, 1, 0)	(0, 0, 1)	($\frac{1}{2}$, 0, 1)	
a^2	(0, 1, 1)	(0, 0, 1)	($\frac{1}{2}$, 1, 0)	(1, 0, 1)	

Fig. 3. A ranking game associated with the ranking game form depicted in Figure 2. The rank payoff for the three players are given by $\vec{p}_1 = (1, \frac{1}{2}, 0)$, $\vec{p}_2 = (1, 0, 0)$ and $\vec{p}_3 = (1, 1, 0)$.

longer holds. A strategic game can be converted to a zero-sum game *via* positive affine transformations only if all outcomes of the game lie on an $(n-1)$ -dimensional hyperplane in the n -dimensional outcome space. Clearly, there are ranking games (with non-identical rank payoff vectors and more than two players) for which this is not the case. For example, consider a three-player ranking game with rank payoff vectors $\vec{p}_1 = \vec{p}_2 = (1, 0, 0)$ and $\vec{p}_3 = (1, 1, 0)$ that has among its outcomes the rankings [1, 2, 3], [2, 1, 3], [3, 1, 2], and [1, 3, 2]. As a consequence, ranking games are no subclass of (the games that can be transformed into) constant-sum games. It is readily appreciated that the opposite inclusion does not hold either.

5 Solution Concepts

In this section we review a number of well-known solution concepts and prove some properties specific to ranking games.

On a normative interpretation, the solution concepts game theory has produced identify reasonable, desirable, or otherwise significant strategy profiles in games. Perhaps the most cautious way for a player to proceed is to ensure his *security level* by playing his *maximin strategy*, the strategy that maximizes his payoff in case the other players were to conspire against him and try to minimize his payoff.

Definition 3 (Maximin strategy and security level) *A strategy $s_i^* \in S_i$ is called a*

maximin strategy for player $i \in N$ if

$$s_i^* \in \operatorname{argmax}_{s_i \in S_i} \min_{s_{-i} \in S_{-i}} p_i(s_i, s_{-i}).$$

The value $v_i = \max_{s_i \in S_i} \min_{s_{-i} \in S_{-i}} p_i(s_i, s_{-i})$ is called the security level of player i .

Given a particular game Γ , we will write $v_i(\Gamma)$ for the security level of player i in Γ . In the game of Figure 1, Alice can achieve her security level of 0.5 by uniform randomization over her actions, *i.e.*, by raising her hand with probability 0.5. The security level for both Bob and Charlie is zero.

We will now move on to the next solution concept, namely that of the iterated elimination of weakly dominated actions.

Definition 4 (Weak Dominance) An action $d_i \in A_i$ is said to be weakly dominated by strategy $s_i \in S_i$ if

$$p_i(a_{-i}, d_i) \leq p_i(a_{-i}, s_i) \quad \text{for all } a \in A,$$

and

$$p_i(a_{-i}, d_i) < p_i(a_{-i}, s_i) \quad \text{for some } a \in A.$$

After one or more dominated actions have been removed, other actions may become dominated that were not dominated previously, and may subsequently be removed. In general, the result of such an iterative elimination process depends on the order in which actions are eliminated, since the elimination of an action of some player may render an action of another player undominated. We say that a game is *solvable* by iterated weak dominance if there is some sequence of eliminations that leaves exactly one action per player.

Perhaps the best-known solution concept is Nash equilibrium (Nash, 1951), which identifies strategy profiles in which no player could increase his payoff by *unilaterally* deviating and playing another strategy. A Nash equilibrium is therefore often called a strategy profile of *mutual best responses*.

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.

Nash (1951) has shown that every normal-form game possesses at least one equilibrium. There are infinitely many Nash equilibria in the single-winner game of Figure 1, the only pure equilibrium is denoted by a dashed square.

		c^1			c^2	
		b^1	b^2		b^1	b^2
a^1	2	1		3	1	
a^2	1	2		1	1	

Fig. 4. Three-player single-winner game without quasi-strict equilibria. Dashed boxes mark all Nash equilibria (one player may mix arbitrarily in boxes that span two outcomes).

A weakness of Nash equilibrium as a normative solution concept is that, given particular strategies of the other players, a player may be indifferent between an action he plays with non-zero probability and an action he does not play at all. For example, in the pure Nash equilibrium of the game in Figure 1, players 2 and 3 might just as well play *any other* strategy without decreasing their chances of winning. To alleviate the effects of this phenomenon, Harsanyi (1973) proposed to impose the additional requirement that *every best response* be played with positive probability. Any Nash equilibrium that also satisfies this latter restriction is called a *quasi-strict equilibrium*.¹

Definition 6 (Quasi-strict equilibrium) A Nash equilibrium $s^* \in S$ is called quasi-strict equilibrium if for all $i \in N$ and all $a_i, a'_i \in A_i$ with $s^*(a_i) > 0$ and $s^*(a'_i) = 0$,

$$p_i(s_{-i}^*, a_i) > p_i(s_{-i}^*, a'_i).$$

Figure 1 shows a quasi-strict equilibrium of the game among Alice, Bob and Charlie.² While quasi-strict equilibria have been shown to always exist in two-player games (Norde, 1999), this is not generally the case for games with more than two players. Figure 4 shows that quasi-strict equilibria need not even exist in single-winner games.³

In ranking games, the stability of some Nash equilibria is especially questionable because they prescribe losing players to play certain strategies even though they could just as well play *any* other strategy without a chance of decreasing their payoff. In each outcome of a ranking game, there is at least one player that is ranked lowest and accordingly receives the minimum payoff of zero. Consequently, any such player has no incentive to actually play the action prescribed by the Nash

¹ Harsanyi originally referred to quasi-strict equilibrium as “quasi-strong”. However, this term has been dropped to distinguish the concept from Aumann’s strong equilibrium (Aumann, 1959).

² Observe that Charlie plays a weakly dominated action with positive probability in this equilibrium.

³ There are only few examples in the literature for games without quasi-strict equilibria (essentially there is one example by van Damme (1983) and another one by Cubitt and Sugden (1994)). For this reason, the game depicted in Figure 4 might be of independent interest.

equilibrium. It follows that all *pure* equilibria are weak in this sense. This problem is especially urgent in single-winner games, where all players but the winner are indifferent over which action to play. Quasi-strict equilibrium can be used to formally illustrate this weakness.

Fact 1 *Quasi-strict equilibria in ranking games are never pure, i.e., in any quasi-strict equilibrium there is at least one player who randomizes over some of his actions.*

Although ranking games may have pure Nash equilibria, it seems as if most of them possess non-pure equilibria as well, *i.e.*, mixed strategy equilibria where at least one player randomizes. We prove this claim for three subclasses of ranking games.

Theorem 1 *The following classes of ranking games always possess at least one non-pure equilibrium:*

- (i) *two-player ranking games,*
- (ii) *three-player single-winner games where each player has two actions, and*
- (iii) *n-player single-winner games where the security level of at least two players is positive.*

Proof: Statement (i) follows from Fact 1 and the existence result by Norde (1999). For reasons of completeness, we give a simple alternative proof. Assume for contradiction that there is a two-player ranking game that only possesses pure equilibria and consider, without loss of generality, a pure equilibrium s^* in which player 1 wins. Since player 2 must be incapable of increasing his payoff by deviating from s^* , player 1 has to win no matter which action the second player chooses. As a consequence, the strategies in s^* remain in equilibrium even if player 2's strategy is replaced with an arbitrary randomization among his actions.

As for (ii), consider a three-player single winner game with actions $A_1 = \{a^1, a^2\}$, $A_2 = \{b^1, b^2\}$, and $A_3 = \{c^1, c^2\}$. Assume for contradiction that there are only pure equilibria in the game and consider, without loss of generality, a pure equilibrium $s^* = (a^1, b^1, c^1)$ in which player 1 wins. In the following, we say that a pure equilibrium is *semi-strict* if at least one player strictly prefers his equilibrium action over all his other actions given that the other players play their equilibrium actions. In single-winner games, this player has to be the winner in the pure equilibrium. We first show that if s^* is semi-strict, *i.e.*, player 1 does not win in action profile (a^2, b^1, c^1) , then there must exist a non-pure equilibrium. For this, consider the strategy profile $s^1 = (a^1, s_2^1, c^1)$, where s_2^1 is the uniform mixture of player 2's actions b^1 and b^2 , along with the strategy profile $s^2 = (a^1, b^1, s_3^2)$, where s_3^2 is the uniform mixture of the actions c^1 and c^2 of player 3. Since player 1 does not win in (a^2, b^1, c^1) , he has no incentive to deviate from either s^1 or s^2 even if he wins in (a^2, b^2, c^1) and (a^2, b^1, c^2) . Consequently, player 3 must win in (a^1, b^2, c^2) in order for s^1 *not* to be an equilibrium. Analogously, for s^2 *not* to be an equilibrium,

player 2 has to win in the same action profile (a^1, b^2, c^2) , contradicting the assumption that the game is a single-winner game. Thus, the existence of a semi-strict pure equilibrium implies that of a non-pure equilibrium. Now assume that s^* is *not* semi-strict. When any of the action profiles in $B = \{(a^2, b^1, c^1), (a^1, b^2, c^1), (a^1, b^1, c^2)\}$ is a pure equilibrium, this also yields a non-pure equilibrium because two pure equilibria that only differ by the action of a single player can be combined into infinitely many mixed equilibria. For B not to contain any pure equilibria, there must be (exactly) one player for every profile in B who deviates to a profile in $C = \{(a^2, b^2, c^1), (a^2, b^1, c^2), (a^1, b^2, c^2)\}$ because the game is a single-winner game and because s^* is not semi-strict. Moreover, either player 1 or player 2 wins in (a^2, b^2, c^1) , player 2 or player 3 in (a^1, b^2, c^2) , and player 1 or player 3 in (a^2, b^1, c^2) . This implies two facts. First, the action profile $s^3 = (a^2, b^2, c^2)$ is a pure equilibrium because no player will deviate from s^3 to any profile in C . Second, the player who wins in s^3 strictly prefers the equilibrium outcome over the corresponding action profile in C , implying that s^3 is semi-strict. The above observation that every semi-strict equilibrium also yields a non-pure equilibrium completes the proof.

As for (iii), recall that the payoff a player obtains in equilibrium must be at least his security level. Thus, a positive security level for player i rules out all equilibria in which player i receives payoff zero, in particular all pure equilibria in which he does not win. If there are two players with positive security levels, both of them have to win with positive probability in any equilibrium of the game. In single-winner games, this can only be the case in a non-pure equilibrium. \square

We conjecture that this existence result in fact applies to the class of all single-winner games. It does *not* extend, however, to general ranking games. Starting from the three-player game of van Damme (1983) that possesses no quasi-strict equilibrium and adding actions that are strongly dominated, it is possible to construct a ranking game with five players that only has pure equilibria.

In Nash equilibrium the players randomize among their actions *independently* from each other. Aumann (1974) introduced the notion of a *correlated strategy*, where players are allowed to coordinate their actions by means of a device or agent that randomly selects one of several action profiles and recommends the actions of this profile to the respective players. Formally, the set of correlated strategies is defined as $\Delta(A_1 \times \cdots \times A_n)$. The corresponding equilibrium concept is then defined as follows.

Definition 7 (Correlated equilibrium) *A correlated strategy $\mu \in \Delta(A)$ is called a correlated equilibrium if for all $i \in N$ and all $a_i^*, a_i \in A_i$,*

$$\sum_{a_{-i} \in A_{-i}} \mu(a_{-i}, a_i^*) (p_i(a_{-i}, a_i^*) - p_i(a_{-i}, a_i)) \geq 0.$$

In other words, a correlated equilibrium of a game is a probability distribution μ over the set of action profiles, such that, if a particular action profile $a^* \in A$ is chosen according to this distribution, and every player $i \in N$ is only informed about his own action a_i^* , it is optimal in expectation for i to play a_i^* , given that he only knows the conditional distribution over values of a_{-i}^* . Correlated equilibrium assumes the existence of a trusted third party who can recommend behavior but cannot enforce it.

It can easily be seen that every Nash equilibrium naturally corresponds to a correlated equilibrium. Nash's existence result thus carries over to correlated equilibria. Again consider the game of Figure 1. The correlated strategy that assigns probability 0.25 each to action profiles (a^1, b^1, c^1) , (a^1, b^2, c^1) , (a^2, b^1, c^1) , and (a^2, b^1, c^2) is a correlated equilibrium in which the expected payoff is 0.5 for player 1 and 0.25 for players 2 and 3. In this particular case, the correlated equilibrium is a convex combination of Nash equilibria, and correlation can be achieved by means of a publicly observable random variable. Perhaps surprisingly, Aumann (1974) has shown that in general the (expected) social welfare of a correlated equilibrium may exceed that of every Nash equilibrium, and that correlated equilibrium payoffs may in fact be outside the convex hull of the Nash equilibrium payoffs. This is of course not possible if social welfare is identical in all outcomes, as is the case in our example.

6 Solving Ranking Games

The question we will try to answer in this section is whether the rather specific payoff structure of ranking games makes it possible 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 and Deng, 2006; Daskalakis et al., 2006), *iterated weak dominance* (Conitzer and Sandholm, 2005), and *pure Nash equilibria* in circuit form games (Schoenebeck and Vadhan, 2006). Graphical games, in which pure Nash equilibria are also known to be intractable (Gottlob et al., 2005), 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 alone decides which player wins the game. Similar properties hold for arbitrary ranking games. For iterated strong dominance (Conitzer and Sandholm, 2005) or correlated equilibria (Papadimitriou, 2005) efficient algorithms exist for general games, and *a fortiori* also for ranking games. Thus there is no further need to consider these solution concepts here. When in the following we refer to the hardness of a game we mean NP-hardness or PPAD-hardness of solving the game using a particular solution concept.

6.1 Mixed Nash Equilibria

Let us first consider Nash equilibria of games with a bounded number of players. Two-player ranking games only allow outcomes $(1, 0)$ and $(0, 1)$ and thus constitute a subclass of constant-sum games. Nash equilibria of constant-sum games can be found by linear programming (see, *e.g.*, Vajda, 1956), for which there is a polynomial time algorithm (Khachiyan, 1979).

To prove hardness for the case with more than two players, it suffices to show that three-player ranking games are at least as hard to solve as general rational bimatrix games. To appreciate this, observe that any n -player ranking game can be turned into an $(n + 1)$ -player ranking game by adding a player who has only one action at his disposal and who is invariably ranked last, keeping relative rankings of the other players intact. Nash equilibria of the $(n + 1)$ -player game then naturally correspond to Nash equilibria of the n -player game. A key concept in our proof is that of a Nash homomorphism, a notion introduced by Abbott et al. (2005). We generalize their definition to games with more than two players.

Definition 8 (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 a game Γ and an equilibrium s^* of $h(\Gamma)$, returns an equilibrium $f(s^*)$ of Γ .*

Obviously, the composition of two Nash homomorphisms is again a Nash homomorphism. Furthermore, any sequence of polynomially many Nash homomorphisms that maps some class of games to another class of games provides us with a polynomial-time reduction from the problem of finding Nash equilibria in the former class to finding Nash equilibria in the latter. Any efficient, *i.e.*, polynomial-time, algorithm for the latter directly leads to an efficient algorithm for the former. On the other hand, hardness of the latter implies hardness of the former.

A very simple example of a Nash homomorphism is the one that scales the payoff of each player by means of a positive affine transformation. It is well-known that Nash equilibria are invariant under this kind of mapping, and f can be taken to be the identity. We will now combine this Nash homomorphism with a more sophisticated function, which maps payoff profiles of a two-player binary game to corresponding three-player subgames with two actions for each player, and obtain Nash homomorphisms from rational bimatrix games to three-player ranking games with different rank payoff profiles.

Lemma 1 *For every rank payoff profile, there exists a Nash homomorphism from the set of rational bimatrix games to the set of three-player ranking games.*

Proof: Abbott et al. (2005) have shown 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, outcome $(1, 1)$ is Pareto-dominant and therefore constitutes a pure Nash equilibrium in any binary game (no player can benefit from deviating). Instances containing such an outcome are easy to solve and need not be considered in our mapping.

In the following, we denote by $(1, p_i^2, 0)$ the rank payoff vector of player i , and by $[i, j, k]$ the outcome where player i is ranked first, j is ranked second, and k is ranked last. First of all, consider ranking games where $p_i^2 < 1$ for some player $i \in N$, *i.e.*, the class 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 transforming the payoffs according to

$$(x_1, x_2) \mapsto ((1 - p_1^2)x_1 + p_1^2, x_2)$$

and then adding a third player who only has a single action and whose payoff is chosen such that the resulting game is a ranking game (but is otherwise irrelevant). We obtain the following mapping, which is obviously a Nash homomorphism:

$$\begin{aligned} (0, 0) &\mapsto (p_1^2, 0) \mapsto [3, 1, 2] \\ (1, 0) &\mapsto (1, 0) \mapsto [1, 3, 2] \\ (0, 1) &\mapsto (p_1^2, 1) \mapsto [2, 1, 3]. \end{aligned}$$

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

- there is an outcome in which i is ranked last and the other two players both receive their maximum payoff of 1 (*i.e.*, a Pareto-dominant outcome), or
- i is *not* ranked last in any outcome, such that the payoffs of the other two players always sum up to 1 and the game is equivalent to a two-player constant-sum game.

If the third player is able to choose between *two* different actions, however, binary games can be mapped to single-loser games. For this, consider the mapping from binary bimatrix games to three-player single-loser games shown in Figure 5. As a first step, binary bimatrix games are mapped to three-player constant-sum games according to

$$(x_1, x_2) \mapsto \left(\frac{1}{2}(x_1 + 1), \frac{1}{2}(x_2 + 1), 1 - \frac{1}{2}(x_1 + x_2)\right).$$

The first two players and their respective sets of actions are the same as in the original game, the third player only has one action c . It is again obvious that this

Outcome	Constant-sum outcome	Ranking subgame
(0, 0)	$\mapsto (\frac{1}{2}, \frac{1}{2}, 1)$	$\mapsto \begin{array}{ c } \hline (1, 0, 1) (0, 1, 1) \\ \hline (0, 1, 1) (1, 0, 1) \\ \hline \end{array} \quad \begin{array}{ c } \hline (0, 1, 1) (1, 0, 1) \\ \hline (1, 0, 1) (0, 1, 1) \\ \hline \end{array}$
(1, 0)	$\mapsto (1, \frac{1}{2}, \frac{1}{2})$	$\mapsto \begin{array}{ c } \hline (1, 0, 1) (1, 1, 0) \\ \hline (1, 1, 0) (1, 0, 1) \\ \hline \end{array} \quad \begin{array}{ c } \hline (1, 1, 0) (1, 0, 1) \\ \hline (1, 0, 1) (1, 1, 0) \\ \hline \end{array}$
(0, 1)	$\mapsto (\frac{1}{2}, 1, \frac{1}{2})$	$\mapsto \begin{array}{ c } \hline (0, 1, 1) (1, 1, 0) \\ \hline (1, 1, 0) (0, 1, 1) \\ \hline \end{array} \quad \begin{array}{ c } \hline (1, 1, 0) (0, 1, 1) \\ \hline (0, 1, 1) (1, 1, 0) \\ \hline \end{array}$

Fig. 5. Mapping from binary bimatrix games to three-player single-loser games

constitutes a Nash homomorphism. Next, outcomes of the three-player constant-sum game are replaced by three-player single-loser subgames. Let Γ be a binary game, and denote by Γ' and Γ'' the three-player constant-sum game and the three-player single-loser game, respectively, obtained by applying the two steps of the mapping in Figure 5 to Γ . We further write p'_i , and p''_i for the payoff function of player i in Γ' and Γ'' , respectively, and a_i^1 and a_i^2 for the two actions of player i in Γ'' corresponding to an action a_i in Γ' .

The second part of the mapping in Figure 5 is chosen such that for all strategy profiles s , all players i and all actions $a_i \in A_i$ in Γ' we have

$$\frac{1}{2}p''_i(a_i^1, s_{-i}) + \frac{1}{2}p''_i(a_i^2, s_{-i}) = p'_i(a_i, f(s)_{-i}), \quad (1)$$

where for each strategy profile s of Γ'' , $f(s)$ is the strategy profile in Γ' such that for each player $i \in \{1, 2, 3\}$ and each action $a_i \in A_i$

$$f(s)(a_i) = s_i(a_i^1) + s_i(a_i^2).$$

An important consequence of this fact is that each player can guarantee his payoff in Γ'' , for any strategy profile of the other players, to be at least as high as his payoff under the corresponding strategy profile in Γ' , by distributing the weight on a_i uniformly on a_i^1 and a_i^2 .

Let s^* be a Nash equilibrium in Γ'' . We first prove that for every player $i \in \{1, 2, 3\}$ and each action a_i of player i in Γ' ,

$$s^*(a_i^1)p''_i(a_i^1, s^*_{-i}) + s^*(a_i^2)p''_i(a_i^2, s^*_{-i}) = (f(s^*)(a_i))p'_i(a_i, f(s^*)_{-i}). \quad (2)$$

Recall that we write $s(a_i)$ for the probability of action a_i in strategy profile s , so $f(s^*)(a_i)$ is the probability with which a_i is played in strategy profile $f(s^*)$ of Γ' .

The above equation thus states that the expected joint payoff from a_i^1 and a_i^2 in equilibrium s^* equals that from a_i under the corresponding strategy profile $f(s^*)$ of Γ' . To see this, first assume for contradiction that for some player i and some action $a_i \in A_i$,

$$s^*(a_i^1)p_i''(a_i^1, s_{-i}^*) + s^*(a_i^2)p_i''(a_i^2, s_{-i}^*) < (f(s^*)(a_i))p_i'(a_i, f(s^*)_{-i}),$$

i.e., that the expected joint payoff from a_i^1 and a_i^2 in Γ'' is strictly smaller than the expected payoff from a_i in Γ' . Define s_i to be the strategy of player i in Γ'' such that $s_i(a_i^1) = s_i(a_i^2) = \frac{1}{2}(s^*(a_i^1) + s^*(a_i^2))$ and $s_i(a_i) = s^*(a_i)$ for all actions $a_i' \in A_i$ distinct from a_i^1 and a_i^2 . It then holds that

$$\begin{aligned} & s^*(a_i^1)p_i''(a_i^1, s_{-i}^*) + s^*(a_i^2)p_i''(a_i^2, s_{-i}^*) \\ & < (f(s^*)(a_i))p_i'(a_i, f(s^*)_{-i}) \\ & = (s^*(a_i^1) + s^*(a_i^2))p_i'(a_i, f(s^*)_{-i}) \\ & = (s^*(a_i^1) + s^*(a_i^2))(\frac{1}{2}p_i''(a_i^1, s_{-i}^*) + \frac{1}{2}p_i''(a_i^2, s_{-i}^*)) \\ & = \frac{1}{2}(s^*(a_i^1) + s^*(a_i^2))p_i''(a_i^1, s_{-i}^*) + \frac{1}{2}(s^*(a_i^1) + s^*(a_i^2))p_i''(a_i^2, s_{-i}^*) \\ & = s_i(a_i^1)p_i''(a_i^1, s_{-i}^*) + s_i(a_i^2)p_i''(a_i^2, s_{-i}^*). \end{aligned}$$

The second and last step follow from the definition of f and s_i , respectively. The third step follows from (1). We conclude that player i obtains a higher payoff by playing s_i instead of s_i^* , contradicting the assumption that s^* is a Nash equilibrium. In particular we have shown that for all $i \in N$ and every $a_i \in A_i$,

$$s^*(a_i^1)p_i''(a_i^1, s_{-i}^*) + s^*(a_i^2)p_i''(a_i^2, s_{-i}^*) \geq (f(s^*)(a_i))p_i'(a_i, f(s^*)_{-i}). \quad (3)$$

Now assume, again for contradiction, that for some player i and some action $a_i \in A_i$,

$$s^*(a_i^1)p_i''(a_i^1, s_{-i}^*) + s^*(a_i^2)p_i''(a_i^2, s_{-i}^*) > (f(s^*)(a_i))p_i'(a_i, f(s^*)_{-i}),$$

i.e., that the expected joint payoff to i from a_i^1 and a_i^2 in Γ'' is strictly greater under s^* than the expected payoff from a_i in Γ' . It follows from (3) that the expected payoff player i receives from *any* action under $f(s^*)$ cannot be greater than the expected joint payoff from the corresponding pair of actions under s^* , and thus $p_i''(s^*) > p_i'(f(s^*))$. Since Γ' and Γ'' are both constant-sum games, there exists some player $j \neq i$ who receives strictly less payoff under s^* in Γ'' than under $f(s^*)$ in Γ' . In particular, there has to be an action $a_j \in A_j$ such that

$$s^*(a_j^1)p_j''(a_j^1, s_{-j}^*) + s^*(a_j^2)p_j''(a_j^2, s_{-j}^*) < (f(s^*)(a_j))p_j'(a_j, f(s^*)_{-j}),$$

contradicting (3).

We are now ready to prove that the mapping in Figure 5 is indeed a Nash homomorphism. To this end, let s^* be a Nash equilibrium of Γ'' , and assume for a contradiction that $f(s^*)$ is not a Nash equilibrium of Γ' . Then there has to be a player i and

some action $a_i \in A_i$ such that $p'_i(a_i, f(s^*)_{-i}) > p'_i(f(s^*))$. Define s_i to be the strategy of i in Γ'' such that $s_i(a_i^1) = s_i(a_i^2) = \frac{1}{2}$. Then, by (1), $p''_i(s_i, s_{-i}) = p'_i(a_i, f(s^*)_{-i})$. It further follows from (2) that for all players i , $p''_i(s^*) = p'_i(f(s^*))$. Thus,

$$p''_i(s^*) = p'_i(f(s^*)) < p'_i(a_i, f(s^*)_{-i}) = p''_i(s'_i, s^*_{-i}),$$

contradicting the assumption that s^* is a Nash equilibrium in Γ'' . \square

The ground has now been cleared 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 as to the existence of Nash equilibria is trivial. However, the associated *search problem* turns out to be not at all trivial. In fact, it has recently been shown to be PPAD-complete for general bimatrix games (Chen and Deng, 2006; Daskalakis et al., 2006). TFNP (for “total functions in NP”) is the class of search problems guaranteed to have a solution. As Daskalakis et al. (2006) put it, “this is precisely NP with an added emphasis on finding a witness.” TFNP is further divided into subclasses based on the mathematical argument used to establish the existence of a solution. PPAD (for “polynomial parity argument, directed version”) is one such subclass that is believed not to be contained in P. For this reason, the PPAD-hardness of a particular problem can be seen as “a rather compelling argument for intractability” (Papadimitriou, 2007, p. 39).

Theorem 2 *Computing a Nash equilibrium of a ranking game with more than two players is PPAD-hard for any rank payoff profile. If there are only two players, equilibria can be found in polynomial time.*

Proof: According to Lemma 1, ranking games with more than two players are at least as hard to solve as general two-player games. We already know that solving general games is PPAD-hard in the two-player case (Chen and Deng, 2006).

Two-player ranking games, on the other hand, form a subclass of two-player zero-sum games, in which Nash equilibria can be found efficiently via linear programming. \square

6.2 Iterated Weak Dominance

We now turn to iterated weak dominance. If there are only two players, the problem of deciding whether a ranking game can be solved via iterated weak dominance is tractable.

Theorem 3 *For two-player ranking games, iterated weak dominance solvability can be decided in polynomial time.*

	a_2^*		
	2	1	
	\vdots		\ddots
	2		1
	1	2	
	\vdots		\ddots
			2
a_1^*	1	\dots	1

Fig. 6. Iterated weak dominance solvability in two-player ranking games

Proof: First we recall that if an action in a binary game is weakly dominated by a mixed strategy, it is also dominated by a pure strategy (Conitzer and Sandholm, 2005). Accordingly, we only have to consider dominance by pure strategies. Now consider a path of iterated weak dominance that ends in a single action profile (a_1^*, a_2^*) . Without loss of generality we may assume that player 1 (*i.e.*, the row player) is the winner in this profile. This implies that player 1 wins in (a_1^*, a_2) for any $a_2 \in A_2$, *i.e.*, in the entire row. For a contradiction, assume the opposite and 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. It is easy to see that a_2^1 could not have been eliminated in this case. An elimination by player 1 would also eliminate a_1^* , while an elimination by player 2 could only take place via another action a_2^2 such that player 2 also wins in (a_1^*, a_2^2) , contradicting the assumption that a_2^1 is eliminated last. We now claim that a ranking game with two players is solvable by iterated weak dominance if and only if there exists a unique action a_1^* of player 1 by which he always wins, and an action a_2^* of player 2 by which he wins for a strictly maximal set of actions of player 1. More precisely, the latter property means that there exists a set of actions of player 1 against which player 2 always wins when playing a_2^* and loses in at least one case for every other action he might play. This is illustrated in Figure 6, and can be verified efficiently by ordering the aforementioned sets of actions of player 1 according to strict inclusion. If the ordering does not have a maximal element, the game cannot be solved by means of iterated weak dominance. If it does, we can use a_1^* to eliminate all actions $a_1 \in A_1$ such that player 2 does not win in (a_1, a_2^*) , whereupon a_2^* can eliminate all other actions of player 2, until finally a_1^* eliminates player 1's remaining actions and solves the game.⁴ \square

⁴ Since two-player ranking games are a subclass of constant-sum games, weak dominance and *nice weak dominance* (Marx and Swinkels, 1997) coincide, making iterated weak dominance order independent *up to payoff-equivalent action profiles*. This fact is mirrored by Figure 6, since there cannot be a row of 1s and a column of 2s in the same matrix.

Theorem 4 For ranking games with more than two players, and for any rank payoff profile, deciding iterated weak dominance solvability is NP-complete.

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 solves the game.

For *hardness*, we first reduce *eliminability* in binary bimatrix games, which asks whether there exists a sequence of eliminations that contains a given action and has recently been shown to be NP-hard (Conitzer and Sandholm, 2005), to the same problem in 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 by 1 and 2, and an additional player 3.
- Players 1 and 2 have the same actions as in Γ , player 3 has two actions c^1 and c^2 .
- Payoffs of Γ are mapped to rankings of Γ' according to

$$\begin{array}{llll} (0, 0) \mapsto & \boxed{[3, 2, 1]} & \boxed{[3, 1, 2]} & (1, 0) \mapsto \boxed{[1, 2, 3]} & \boxed{[3, 1, 2]} \\ (0, 1) \mapsto & \boxed{[3, 2, 1]} & \boxed{[2, 1, 3]} & (1, 1) \mapsto \boxed{[1, 2, 3]} & \boxed{[2, 1, 3]}. \end{array}$$

In the following, we write p and p' for the payoff functions of Γ and Γ' , respectively.

First observe that we can restrict our attention to dominance by pure strategies. This property holds for binary games by Lemma 1 of Conitzer and Sandholm (2005), and thus also for actions of player 3, who receives a payoff of either 0 or 1 in any outcome. For players 1 and 2 we can essentially apply the same argument, because each of them can obtain only two different payoffs for any fixed action profile of the remaining two players.

We now claim that irrespective of the rank payoffs $p_i = (1, p_i^2, 0)$, and for any subsets of the actions of players 1 and 2, a particular action of these players is dominated in the restriction of Γ' to these subsets if and only if the corresponding action is dominated in the restriction of Γ to the same subsets. To see this, observe that if player 3 plays c^1 , then for any action profile $(a_1, a_2) \in A_1 \times A_2$, player 1 receives the same payoff he would receive for the corresponding action profile in Γ , *i.e.*, $p'_1(a_1, a_2, c^1) = p_1(a_1, a_2)$, whereas player 2 receives a payoff of p_2^2 . If on the other hand player 3 plays c^2 , then player 1 obtains a payoff of p_1^2 , and the payoff of player 2 for any action profile $(a_1, a_2) \in A_1 \times A_2$ is the same as that for the corresponding profile in Γ , *i.e.*, $p'_2(a_1, a_2, c^2) = p_2(a_1, a_2)$. Moreover, the implication from left to right still holds if one of the actions of player 3 is removed, because this leaves one of players 1 and 2 indifferent between all of his remaining actions but does not have any effect on dominance between actions of the other player. We have thus established a direct correspondence between sequences of eliminations in Γ and Γ' , which in turn implies NP-hardness of deciding whether a particular

	c^1			c^2		
	b^1	\dots	b^k	b^1	\dots	b^k
a^1	$[\cdot, 2, \cdot]$	\dots	$[\cdot, 2, \cdot]$	$[\cdot, 1, \cdot]$	\dots	$[\cdot, 1, \cdot]$
\vdots	\vdots	\ddots	\vdots	\vdots	\ddots	\vdots
a^m	$[\cdot, 2, \cdot]$	\dots	$[\cdot, 2, \cdot]$	$[\cdot, 1, \cdot]$	\dots	$[\cdot, 1, \cdot]$
a^{m+1}	$[3, 2, 1]$	\dots	$[3, 2, 1]$	$[2, 1, 3]$	\dots	$[2, 1, 3]$

Fig. 7. Three-player ranking game Γ' used in the proof of Theorem 4

action of a ranking game with at least three players can be eliminated.

It also follows from the above that Γ can be solved by iterated weak dominance if Γ' can. The implication in the other direction does not hold, however, because it may not always be possible to eliminate an action of player 3. To this end, assume without loss of generality that some player of Γ' has at least two actions, and that this player is player 1. Otherwise both Γ and Γ' are trivially solvable. We augment Γ' by introducing to player 1's action set $A_1 = \{a^1, \dots, a^m\}$ an additional action a^{m+1} of player 1 such that for every action b^j of player 2, $g(a^{m+1}, b^j, c^1) = [3, 2, 1]$ and $g(a^{m+1}, b^j, c^2) = [2, 1, 3]$. The structure of the resulting game is shown in Figure 7.

It is easily verified that the above arguments about Γ' still apply, because player 1 never receives a higher payoff from a^{m+1} than from any other action, and player 2 is indifferent between all of his actions when player 1 plays a^{m+1} . Now assume that Γ can be solved. Without loss of generality we may assume that (a^1, b^1) is the remaining action profile. Clearly, for Γ to be solvable, player 1 must be ranked first in some outcome of Γ' , and it must hold that $p_1(a^1, b^1) = 1$ or $p_2(a^1, b^1) = 1$. We distinguish two cases. If $p_1(a^1, b^1) = p_2(a^1, b^1) = 1$, then Γ' can be solved by performing the eliminations that lead to the solution of Γ , followed by the elimination of c^2 and a^{m+1} . Otherwise we can start by eliminating a^{m+1} , which is dominated by the action for which player 1 is sometimes ranked first, and proceed with the eliminations that solve Γ . In the two action profiles that then remain Player 3 is ranked first and last, respectively, and he can eliminate one of his actions to solve Γ' . \square

6.3 Pure Nash Equilibria in Games with Many Players

We now consider the situation in which players do not randomize but choose their actions deterministically. Nash equilibria in pure strategies can be found efficiently by simply checking every action profile. As the number of players increases, however, 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 *succinct* representation of a game that only uses space polynomial in n .

We proceed to show that this is most likely not the case. More precisely, we show 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 2 and 4, we keep the number of actions fixed and let the number of players grow.

Theorem 5 *For ranking games with an unbounded number of players and a polynomial-time computable outcome function, and for any rank payoff profile, deciding the existence of a pure Nash equilibrium is NP-complete, even if the players have only 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 $a_i \in A_i$ whether $p_i(s_{-i}, a_i) \leq p_i(s)$.

For *hardness*, recall that circuit satisfiability (CSAT), *i.e.*, deciding whether for a given Boolean circuit φ 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 φ , providing a polynomial-time reduction of satisfiability of φ to the problem of finding a pure Nash equilibrium in Γ .

Let m be the number of inputs of φ . We define game Γ with $m+2$ players as follows:

- Let $N = \{1, \dots, m\} \cup \{x, y\}$, and $A_i = \{0, 1\}$ for all $i \in N$.
- The outcome function of Γ is computed by a Boolean circuit that takes $m+2$ bits of input $i = (a_1, \dots, a_m, a_x, a_y)$, corresponding to the actions of the players in N , and computes two bits of output $o = (o_1, o_2)$, given by $o_1 = \varphi(a_1, \dots, a_m)$ and $o_2 = (o_1 \text{ OR } (a_x \text{ XOR } a_y))$.
- 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 $o = (0, 0)$ and the permutation π_{11} corresponding to $o = (1, 1)$ rank x first and y last,
 - the permutation π_{01} corresponding to $o = (0, 1)$ ranks y first, and x last, and
 - all other players are ranked in the same order 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, Γ has a pure Nash equilibrium if and only if φ is satisfiable. This can be seen as follows:

- If (a_1, \dots, a_m) is a satisfying assignment of φ , only a player in $\{1, \dots, m\}$ could possibly change the outcome of the game by changing his action. However,

these players are ranked in the same order in all the possible outcomes, so none of them can get a higher payoff by doing so. Thus, every action profile $a = (a_1, \dots, a_m, a_x, a_y)$ where (a_1, \dots, a_m) satisfies φ is a Nash equilibrium.

- If in turn (a_1, \dots, a_m) is not a satisfying assignment of φ , both x and y are able to switch 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, $a = (a_1, \dots, a_m, a_x, a_y)$ cannot be a Nash equilibrium in this case, since either x or y could play a different action to get a higher payoff. \square

7 Comparative Ratios

Despite its conceptual elegance and simplicity, Nash equilibrium has been criticized on various grounds (see, *e.g.*, Luce and Raiffa, 1957, for a discussion). In the common case of multiple equilibria, it is unclear which one should be selected. Also, coalitions might benefit from jointly deviating, and there might exist no polynomial-time, *i.e.*, efficient, algorithms for finding Nash equilibria, a problem we discussed in the previous section. Moreover, players may be utterly indifferent among equilibrium and non-equilibrium strategies, which we saw is pervasive in ranking games.

7.1 The Price of Cautiousness

A compelling question is how much worse off a player can be when if he were to revert to his most defensive course of action—his maximin strategy—instead of hoping for an equilibrium outcome. This difference in payoff can be represented by a numerical value which we refer to as the *price of cautiousness*. In what follows, let \mathcal{G} denote the class of all normal-form games, and for $\Gamma \in \mathcal{G}$, let $N(\Gamma)$ be the set of Nash equilibria of Γ . Recall that $v_i(\Gamma)$ denotes player i 's security level in game Γ .

Definition 9 Let Γ be a normal-form game with non-negative payoffs, $i \in N$ a player such that $v_i(\Gamma) > 0$. The price of cautiousness for player i in Γ is defined as

$$PC_i(\Gamma) = \frac{\min \{ p_i(s) \mid s \in N(\Gamma) \}}{v_i(\Gamma)}.$$

For any class $C \subseteq \mathcal{G}$ of games involving player i , we further write $PC_i(C) = \sup_{\Gamma \in C} PC_i(\Gamma)$. In other words, the price of cautiousness of a player is the ratio between his minimum payoff in a Nash equilibrium and his security level. It thus captures the worst-case loss the player may incur by playing his maximin strategy

	c^1		c^2	
	b^1	b^2	b^1	b^2
a^1	$(0, 1, 1)$	$(1, 0, 0)$	$(\epsilon, 1, 0)$	$(\epsilon, 0, 1)$
a^2	$(1, 0, 1)$	$(0, 1, 1)$	$(\epsilon, 1, 0)$	$(\epsilon, 0, 1)$

Fig. 8. Three-player ranking game Γ_1 used in the proof of Theorem 6

instead of a Nash equilibrium.⁵ For a player whose security level equals his minimum payoff of zero, *every* strategy is a maximin strategy. Since we are mainly interested in a comparison of normative solution concepts, we will only consider games where the security level of at least one player is positive.

As we have already mentioned in Section 1, the price of cautiousness in two-player ranking games equals 1 in virtue of the Minimax Theorem of von Neumann (1928). In general ranking games, however, the price of cautiousness is unbounded.

Theorem 6 *Let \mathcal{R} be the class of ranking games with more than two players that involve player i . Then, the price of cautiousness is unbounded, i.e., $PC_i(\mathcal{R}) = \infty$, even if \mathcal{R} only contains games without weakly dominated actions.*

Proof: Consider the game Γ_1 of Figure 8, which is a ranking game for rank payoff vectors $\vec{p}_1 = (1, \epsilon, 0)$, $\vec{p}_2 = (1, 0, 0)$, and $\vec{p}_3 = (1, 1, 0)$, and rankings $[2, 3, 1]$, $[1, 3, 2]$, $[1, 2, 3]$, $[2, 1, 3]$, and $[3, 1, 2]$. It is easily verified that none of the actions of Γ_1 is weakly dominated and that $v_1(\Gamma_1) = \epsilon$. Let further $s = (s_1, s_2, c^1)$ be the strategy profile where s_1 and s_2 are uniform mixtures of a^1 and a^2 , and of b^1 and b^2 , respectively. We will argue that s is the only Nash equilibrium of Γ_1 . For this, consider the possible strategies of player 3. If player 3 plays c^1 , the game reduces to the well-known matching pennies game for players 1 and 2, the only Nash equilibrium being the one described above. If on the other hand player 3 plays c^2 , action b^1 strongly dominates b^2 . If b^1 is played, however, player 3 will deviate to c^1 to get a higher payoff. Finally, if player 3 randomizes between actions c^1 and c^2 , the payoff obtained from both of these actions must be the same. This can only be the case if either player 1 plays a^1 and player 2 randomizes between b^1 and b^2 , or if player 1 plays a^2 and player 2 plays b^2 . In the former case, player 2 will deviate to b^1 . In the latter case, player 1 will deviate to a^1 . Since the payoff of player 1 in the above equilibrium is 0.5, we have $PC(\Gamma_1) = 0.5/\epsilon \rightarrow \infty$ for $\epsilon \rightarrow 0$. \square

We proceed to show that, due to their structural limitations, the price of cautiousness in *binary* ranking games is bounded from above by the number of actions of the respective player. We also derive a matching lower bound.

⁵ In our context, the choice of whether to use the worst or the best equilibrium when defining the price of cautiousness is merely a matter of taste. All results in this section still hold when the best equilibrium is used instead of the worst one.

Theorem 7 *Let \mathcal{R}_b be the class of binary ranking games with more than two players involving a player i with exactly k actions. Then, $PC_i(\mathcal{R}_b) = k$, even if \mathcal{R}_b only contains single-winner games or games without weakly dominated actions.*

Proof: By definition, the price of cautiousness takes its maximum for maximum payoff in a Nash equilibrium, which is bounded by 1 in a ranking game, and minimum security level. It being required that the security level is strictly positive, for every opponent action profile s_{-i} there is some action $a_i \in A_i$ such that $p_i(a_i, s_{-i}) > 0$, i.e., $p_i(a_i, s_{-i}) = 1$. It is then easily verified that player i can ensure a security level of $1/k$ by uniform randomization over his k actions. This results in a price of cautiousness of at most k .

For a matching lower bound, again consider the single winner game depicted in Figure 4. We will argue that all Nash equilibria of this game are mixtures of the action profiles (a^2, b^1, c^2) , (a^2, b^2, c^2) and (a^1, b^2, c^2) . Each of these equilibria yields payoff 1 for player 1, twice as much as his security level of 0.5. To appreciate this, consider the strategies that are possible for player 3. If player 3 plays c^1 , the game reduces to the well-known game of matching pennies for players 1 and 2, in which they will randomize uniformly over both of their actions. In this case, player 3 will deviate to c^2 . If player 3 plays c^2 , we immediately obtain the equilibria described above. Finally, if player 3 randomizes between actions c^1 and c^2 , the payoff obtained from both of these actions should be the same. This can only be the case if either player 1 plays a^2 and player 2 randomizes between b^1 and b^2 , or if player 1 randomizes between a^1 and a^2 and player 2 plays b^2 . In the former case, player 2 will play b^2 , causing player 1 to deviate to a^1 . In the latter case, player 1 will play a^1 , causing player 2 to deviate to b^1 .

The above construction can be generalized to $k > 2$ by virtue of a single-winner game with actions $A_1 = \{a^1, \dots, a^k\}$, $A_2 = \{b^1, \dots, b^k\}$, and $A_3 = \{c^1, c^2\}$, and payoffs

$$p(a^i, b^j, c^\ell) = \begin{cases} (0, 1, 0) & \text{if } \ell = 1 \text{ and } i \neq k - j + 1 \\ (0, 0, 1) & \text{if } \ell = 2 \text{ and } i = j = 1 \\ (1, 0, 0) & \text{otherwise.} \end{cases}$$

It is easily verified that the security level of player 1 in this game is $1/k$ while, by the same arguments as above, his payoff in every Nash equilibrium equals 1. This shows tightness of the upper bound of k on the price of cautiousness for single-winner games.

Now consider the game Γ_2 of Figure 9, which is a ranking game for rank payoff vectors $\vec{p}_1 = \vec{p}_2 = (1, 0, 0)$ and $\vec{p}_3 = (1, 1, 0)$, and rankings $[2, 3, 1]$, $[1, 2, 3]$, $[2, 1, 3]$, and $[1, 3, 2]$. It is easily verified that none of the actions of Γ_2 is weakly dominated and that $v_1(\Gamma_2) = 0.5$. On the other hand, we will argue that all Nash equilibria of Γ_2 are mixtures of action profiles (a^2, b^1, c^2) and (a^2, b^2, c^2) , corresponding to a payoff of 1 for player 1. To see this, we again look at the possible strategies for player 3. If player 3 plays c^1 , players 1 and 2 will again randomize uniformly over

	c^1		c^2	
	b^1	b^2	b^1	b^2
a^1	(0, 1, 1)	(1, 0, 0)	(0, 1, 0)	(1, 0, 0)
a^2	(1, 0, 0)	(0, 1, 0)	(1, 0, 1)	(1, 0, 1)

Fig. 9. Three-player ranking game Γ_2 used in the proof of Theorem 7

both of their actions, causing player 3 to deviate to c^2 . If player 3 plays c^2 , we immediately obtain the equilibria described above. Finally, assume that player 3 randomizes between actions c^1 and c^2 , and let α denote the probability with which player 1 plays a^1 . Again, player 3 must be indifferent between c^1 and c^2 , which can only hold for $0.5 \leq \alpha \leq 1$. In this case, however, player 2 will deviate to b^1 .

This construction can be generalized to $k > 2$ by virtue of a game with actions $A_1 = \{a^1, \dots, a^k\}$, $A_2 = \{b^1, \dots, b^k\}$, and $A_3 = \{c^1, c^2\}$, and payoffs

$$p(a^i, b^j, c^\ell) = \begin{cases} (0, 1, 1) & \text{if } i = j = \ell = 1 \\ (1, 0, 0) & \text{if } \ell = 1 \text{ and } i = k - j + 1 \\ & \text{or } \ell = 2, i = 1 \text{ and } j > 1 \\ (1, 0, 1) & \text{if } \ell = 2 \text{ and } j > 2 \\ (0, 1, 0) & \text{otherwise.} \end{cases}$$

Again, it is easily verified that the security level of player 1 in this game is $1/k$ while, by the same arguments as above, his payoff is 1 in every Nash equilibrium. Thus, the upper bound of k for the price of cautiousness is tight as well for binary ranking games without weakly dominated actions. \square

Informally, the previous theorem states that the payoff a player with k actions can obtain in Nash equilibrium can be at most k times his security level.

7.2 The Value of Correlation

We will now turn to the question whether, and by which amount, *social welfare* can be improved by allowing players in a ranking game to *correlate* their actions. Just as the payoff of a player in any Nash equilibrium is at least his security level, social welfare in the best correlated equilibrium is at least as high as social welfare in the best Nash equilibrium. In order to quantify the value of correlation in strategic games with non-negative payoffs, Ashlagi et al. (2005) recently introduced the *mediation value* of a game as the ratio between the maximum social welfare in a correlated versus that in a Nash equilibrium, and the *enforcement value* as the ratio between the maximum social welfare in any outcome versus that in a correlated equilibrium. Whenever social welfare, *i.e.*, the sum of all players' payoffs, is used

as a measure of global satisfaction, one implicitly assumes the inter-agent comparability of payoffs. While this assumption is controversial, social welfare is nevertheless commonly used in the definitions of comparative ratios such as the price of anarchy (Koutsoupias and Papadimitriou, 1999). For $\Gamma \in \mathcal{G}$ and $X \subseteq \Delta(S)$, let $C(\Gamma)$ denote the set of correlated equilibria of Γ and let $v_X(\Gamma) = \max\{p(s) \mid s \in X\}$. Recall that $N(\Gamma)$ denotes the set of Nash equilibria in Γ .

Definition 10 *Let Γ be a normal-form game with non-negative payoffs. The mediation value $MV(\Gamma)$ and the enforcement value $EV(\Gamma)$ of Γ are defined as*

$$MV(\Gamma) = \frac{v_{C(\Gamma)}(\Gamma)}{v_{N(\Gamma)}(\Gamma)} \quad \text{and} \quad EV(\Gamma) = \frac{v_S(\Gamma)}{v_{C(\Gamma)}(\Gamma)}.$$

If both numerator and denominator are 0 for one of the values, the respective value is defined to be 1. If only the denominator is 0, the value is defined to be ∞ . For any class $\mathcal{C} \subseteq \mathcal{G}$ of games, we further write $MV(\mathcal{C}) = \sup_{\Gamma \in \mathcal{C}} MV(\Gamma)$ and $EV(\mathcal{C}) = \sup_{\Gamma \in \mathcal{C}} EV(\Gamma)$.

Ashlagi et al. (2005) have shown that both the mediation value and the enforcement value cannot be bounded for games with an arbitrary payoff structure, as soon as there are more than two players or some player has more than two actions. This holds even if payoffs are normalized to the interval $[0, 1]$. Ranking games also satisfy this normalization criterion, and here social welfare is also strictly positive for every outcome of the game. Ranking games with identical rank payoff vectors for all players, *i.e.*, ones where $p_i^k = p_j^k$ for all $i, j \in N$ and $1 \leq k \leq n$, are constant-sum games. Hence, social welfare is the same in every outcome so that both the mediation value and the enforcement value are 1. This in particular concerns all ranking games with two players. In general, social welfare in an arbitrary outcome of a ranking game is bounded by $n - 1$ from above and by 1 from below. Since the Nash and correlated equilibrium payoffs must lie in the convex hull of the feasible payoffs of the game, we obtain trivial lower and upper bounds of 1 and $n - 1$, respectively, on both the mediation and the enforcement value. It turns out that the upper bound of $n - 1$ is tight for both the mediation value and the enforcement value.

Theorem 8 *Let \mathcal{R}' be the class of ranking games with $n > 2$ players, such that in games with only three players at least one player has more than two actions. Then, $MV(\mathcal{R}') = n - 1$.*

Proof: It suffices to show that for any of the above cases there is a ranking game with mediation value $n - 1$. For $n = 3$, consider the game Γ_3 of Figure 10, which is a ranking game for rank payoff vectors $\vec{p}_1 = \vec{p}_3 = (1, 0, 0)$ and $\vec{p}_2 = (1, 1, 0)$. First of all, we will show that every Nash equilibrium of this game has social welfare 1, by showing that there are no Nash equilibria where c^1 or c^2 are played with positive probability. Assume for contradiction that s^* is such an equilibrium. The strategy

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

Fig. 10. Three-player ranking game Γ_3 used in the proof of Theorem 8

played by player 3 in s^* must either be (i) c^1 or c^2 as a pure strategy, (ii) a mixture of c^1 and c^3 or between c^2 and c^3 , or (iii) a mixture where both c^1 and c^2 are played with positive probability. If player 3 plays a pure strategy, the game reduces to a two-player game for players 1 and 2. In the case of c^1 , this game has the unique equilibrium (a^1, b^1) , which in turn causes player 3 to deviate to c^2 . In the case of c^2 , the unique equilibrium is (a^2, b^2) , causing player 3 to deviate to c^1 . Now assume that player 3 mixes between c^1 and c^3 , and let α and β denote the probabilities with which players 1 and 2 play a^1 and b^1 , respectively. Since player 3's payoff from c^1 and c^3 must be the same in such an equilibrium, we must either have $\alpha = \beta = 1$, in which case player 3 will deviate to c^2 , or $0 \leq \alpha \leq 0.5$ and $0 \leq \beta \leq 0.5$, causing player 2 to deviate to b^1 . Analogously, if player 3 mixes between c^2 and c^3 , we must either have $\alpha = \beta = 0$, in which case player 3 will deviate to c^1 , or $0.5 \leq \alpha \leq 1$ and $0.5 \leq \beta \leq 1$, causing player 2 to deviate to b^2 . Finally, if both c^1 and c^2 are played with positive probability, we must have $\alpha + \beta = 1$ for player 3 to get an identical payoff of $\alpha\beta \leq 1/4$ from both c^1 and c^2 . In this case, however, player 3 can deviate to c^3 for a strictly greater payoff of $1 - 2\alpha\beta$. Thus, a strategy profile s^* as described above cannot exist.

Now let t^* be the correlated strategy where action profiles (a^1, b^1, c^1) , (a^2, b^2, c^1) , (a^1, b^1, c^2) , and (a^2, b^2, c^2) are played with probability 0.25 each. This correlation can for example be achieved by tossing two coins independently. Players 1 and 2 observe the first coin toss and play a^1 and b^1 , respectively, if the coin falls on heads, and a^2 and b^2 otherwise. Player 3 observes the second coin toss and plays c^1 if the coin falls on heads and c^2 otherwise. The expected payoff for player 2 under t^* is 1, so he cannot gain by changing his action. If player 1 observes heads, he knows that player 2 will play b^1 , and that player 3 will play c^1 and c^2 with probability 0.5 each. He is thus indifferent between a^1 and a^2 . Player 3 knows that players 1 and 2 will play (a^1, b^1) and (a^2, b^2) with probability 0.5 each, so he is indifferent between c^1 and c^2 and strictly prefers both of them to c^3 . Hence, none of the players has an incentive to deviate, t^* is a correlated equilibrium. Moreover, the social welfare under t^* is 2, and thus $MV(\Gamma_3) = 2$.

Now consider the four-player game Γ_4 of Figure 11, which is a ranking game for rank payoffs $\vec{p}_1 = \vec{p}_3 = (1, 0, 0, 0)$, $\vec{p}_2 = (1, 1, 0, 0)$, and $\vec{p}_4 = (1, 1, 1, 0)$, and rankings $[1, 2, 4, 3]$, $[1, 3, 2, 4]$, $[3, 2, 4, 1]$, $[2, 3, 1, 4]$, and $[4, 1, 2, 3]$. It is easily verified that none of the action profiles with social welfare 2 is a Nash equilibrium. Furthermore, player 4 strictly prefers action d^2 over d^1 as soon as one of the remaining

		c^1			c^2	
		b^1	b^2		b^1	b^2
a^1	(1, 1, 0, 1)	(1, 0, 0, 0)		(0, 1, 1, 1)	(0, 1, 0, 0)	d^1
a^2	(0, 1, 0, 0)	(0, 1, 1, 1)		(1, 0, 0, 0)	(1, 1, 0, 1)	
a^1	(0, 0, 0, 1)	(0, 0, 0, 1)		(0, 0, 0, 1)	(0, 0, 0, 1)	d^2
a^2	(0, 0, 0, 1)	(0, 0, 0, 1)		(0, 0, 0, 1)	(0, 0, 0, 1)	

Fig. 11. Four-player ranking game Γ_4 used in the proof of Theorem 8

action profiles for players 1 to 3 (*i.e.*, those in the upper half of the game where the social welfare is 1) is played with positive probability. Hence, d^1 is not played with positive probability in any Nash equilibrium of Γ_4 , and every Nash equilibrium of Γ_4 has social welfare 1. In turn, consider the correlated strategy μ^* where actions profiles (a^1, b^1, c^1, d^1) , (a^2, b^2, c^1, d^1) , (a^1, b^1, c^2, d^1) , and (a^2, b^2, c^2, d^1) are played with probability 0.25 each. It is easily verified that none of the players can increase his payoff by unilaterally deviating from μ^* . Hence, μ^* is a correlated equilibrium with social welfare 3, and $MV(\Gamma_4) = 3$.

For $n > 4$, we can restrict our attention to games where the additional players only have a single action. We return to the game Γ_4 of Figure 11 and transform it into a game Γ_4^n with $n > 4$ players by assigning to players $5, \dots, n$ a payoff of 1 in the four action profiles (a^1, b^1, c^1, d^1) , (a^2, b^2, c^1, d^1) , (a^1, b^1, c^2, d^1) , and (a^2, b^2, c^2, d^1) that constitute the correlated equilibrium with maximum social welfare, and a payoff of zero in all other action profiles. Since the additional players cannot influence the outcome of the game, this construction does not affect the equilibria of the game. To see that the resulting game is a ranking game, consider the rank payoff vectors $\vec{p}_1 = \vec{p}_3 = (1, 0, 0, \dots, 0)$, $\vec{p}_2 = (1, 1, \dots, 0)$, $r_m^k = 1$ if $k \leq m - 1$ and 0 otherwise, for $m \geq 4$. It is easily verified that we can retain the original payoffs of players 1 to 4 and at the same time assign a payoff of 0 or 1, respectively, to players 5 to n by ranking the latter according to their index and placing either no other players or exactly one other player behind them in the overall ranking. More precisely, Γ_4^n is a ranking game by virtue of the above rank payoffs and rankings $[1, 2, 4, 5, \dots, n, 3]$, $[1, 3, 2, 4, 5, \dots, n]$, $[3, 2, 4, 5, \dots, n, 1]$, $[2, 3, 1, 4, 5, \dots, n]$, and $[4, 1, 2, 3, 5, \dots, n]$. Furthermore, $MV(\Gamma_4^n) = n - 1$. \square

Theorem 9 *Let \mathcal{R} be the class of ranking games with $n > 2$ players. Then, $EV(\mathcal{R}) = n - 1$, even if \mathcal{R} only contains games without weakly dominated actions.*

Proof: It suffices to show that for any $n \geq 3$ there is a ranking game with enforcement value $n - 1$ in which no action is weakly dominated. Consider the ranking game Γ_5 of Figure 12, which is a ranking game by virtue of rank payoff vectors $\vec{p}_1 = (1, 1, 0)$, $\vec{p}_2 = (1, 0, 0)$, and $\vec{p}_3 = (1, \epsilon, 0)$ and rankings $[1, 2, 3]$, $[2, 3, 1]$,

	c^1		c^2	
	b^1	b^2	b^1	b^2
a^1	$(1, 0, 0)$	$(0, 1, \epsilon)$	$(0, 0, 1)$	$(1, 0, \epsilon)$
a^2	$(0, 1, \epsilon)$	$(1, 0, 1)$	$(1, 0, \epsilon)$	$(1, 0, \epsilon)$

Fig. 12. Three-player ranking game Γ_5 used in the proof of Theorem 9

$[3, 1, 2]$, $[3, 2, 1]$, and $[1, 3, 2]$. Obviously, all of the actions of Γ_5 are undominated and $v_S(\Gamma_5) = 2$. It remains to be shown that the social welfare in any correlated equilibrium of Γ_5 is at most $(1 + \epsilon)$, such that $v_{C(\Gamma_5)}(\Gamma_5) \rightarrow 1$ and $EV(\Gamma_5) \rightarrow 2$ for $\epsilon \rightarrow 0$.

Finding a correlated equilibrium that maximizes social welfare constitutes a *linear programming* problem constrained by the inequalities of Definition 7 and the probability constraints $\sum_{s \in A} \mu(a) = 1$ and $\mu(a) \geq 0$ for all $a \in A$. Feasibility of this problem is a direct consequence of Nash's existence theorem. Boundedness follows from boundedness of the quantity being maximized. To derive an upper bound for social welfare in a correlated equilibrium of Γ_5 , we will transform the above linear program into its dual. Since the primal is feasible and bounded, the primal and the dual will have the same optimal value, in our case the maximum social welfare in a correlated equilibrium. The latter constitutes a minimization problem and finding a feasible solution with objective value v shows that the optimal value cannot be greater than v . Since there are three players with two actions each, the primal has six constraints of the form $\sum_{a_{-i} \in A_{-i}} \mu(a_{-i}, a_i^*) (p_i(a_{-i}, a_i^*) - p_i(a_{-i}, a_i)) \geq 0$. For $j \in \{1, 2\}$, let x_j, y_j, z_j , denote the variable of the dual associated with the constraint for the j th action of player 1, 2, and 3, respectively. Furthermore, let v denote the variable of the dual associated with constraint $\sum_{a \in A} \mu(a) = 1$ of the primal. Then the dual reads

$$\begin{aligned}
& \text{minimize } v \\
& \text{subject to } -x_1 + y_1 + z_1 + v \geq 1, \\
& \quad x_2 - y_1 + v \geq 1 + \epsilon, \\
& \quad x_1 - y_2 + v \geq 1 + \epsilon, \\
& \quad -x_2 + y_2 + (\epsilon - 1)z_1 + v \geq 2, \\
& \quad x_1 - z_2 + v \geq 1, \\
& \quad -x_2 + v \geq 1 + \epsilon, \\
& \quad v \geq 1 + \epsilon, \\
& \quad (1 - \epsilon)z_2 + v \geq 1 + \epsilon, \\
& \quad x_1 \geq 0, x_2 \geq 0, y_1 \geq 0, y_2 \geq 0, z_1 \geq 0, \text{ and } z_2 \geq 0.
\end{aligned}$$

Now let $x_2 = y_1 = z_2 = 0$, $x_1 = y_2 = (\epsilon - 1)^2/\epsilon$, $z_1 = (1 - 2\epsilon)/\epsilon$, and $v = 1 + \epsilon$, and observe that for every $\epsilon > 0$, this is a feasible solution with objective value $1 + \epsilon$. However, the objective value of any feasible solution to the dual is an upper bound

for that of the optimal solution, which in turn equals $v_{C(\Gamma_3)}(\Gamma_3)$.

The above construction can easily be generalized to games Γ_5^n with $n > 4$ by adding additional players that receive payoff 1 in action profile a if $a_1 = a^2$, $a_2 = b^2$, and $a_3 = c^1$, and payoff 0 otherwise. This can for example be achieved by means of rank payoff vectors $\vec{p}_1 = (1, 0, \dots, 0)$, $\vec{p}_2 = (1, 1, 0, \dots, 0)$, $\vec{p}_3 = (1, \epsilon, 0, \dots, 0)$, and $\vec{p}_m^k = 1$ if $k \leq m - 1$ and 0 otherwise for $m \geq 4$. By the same arguments as in the proof of Theorem 8, this does not affect the maximum social welfare achievable in a correlated equilibrium. It is thus easily verified that $EV(\Gamma_5^{k_1 \times \dots \times k_4}) \rightarrow n - 1$ for $\epsilon \rightarrow 0$. \square

8 Conclusion

We proposed a new class of strategic games, so-called ranking games, which model settings in which players are merely interested in outperforming their opponents. Despite the structural simplicity of these games, various solution concepts turned out to be just as hard to compute as in general normal-form games. In particular we obtained hardness results for mixed Nash equilibria and iterated weak dominance in games with more than two players and pure Nash 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 such as maximin strategies or correlated equilibrium.

Based on these findings, we have quantified and bounded comparative ratios of various solution concepts in ranking games. It turned out that playing one's maximin strategy in binary ranking games with only few actions might be a prudent choice, not only because this strategy guarantees a certain payoff even when playing against irrational opponents, but also because of the limited price of cautiousness and the inherent weakness of Nash equilibria in ranking games.

We also investigated the relationship between correlated and Nash equilibria. While correlation can never *decrease* social welfare, it is an important question which (especially competitive) scenarios permit an *increase*. In scenarios with many players and asymmetric preferences over ranks (*i.e.*, non-identical rank payoff vectors) overall satisfaction can be improved substantially by allowing players to correlate their actions. Furthermore, correlated equilibria do not suffer from the equilibrium selection problem since the equilibrium to be played is selected by a mediator.

Acknowledgements

The authors thank Vincent Conitzer, Markus Holzer, Samuel Jeong, Eugene Nudelman, and Rob Powers, and the anonymous referees for valuable comments. This article 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. Research on this topic was initiated during a post-doctoral stay of the first author at Stanford University.

Preliminary versions of parts of this work have appeared in the proceedings of the 21st National Conference on Artificial Intelligence (AAAI) and the 20th International Joint Conference on Artificial Intelligence (IJCAI), and have been presented at the 17th International Conference on Game Theory.

References

- Abbott, T., Kane, D., Valiant, P., 2005. On the complexity of two-player win-lose games. In: Proceedings of the 46th Symposium on Foundations of Computer Science (FOCS). IEEE Computer Society Press, pp. 113–122.
- Ashlagi, I., Monderer, D., Tennenholtz, M., 2005. On the value of correlation. In: Proceedings of the 21st Annual Conference on Uncertainty in Artificial Intelligence (UAI). AUAI Press, pp. 34–41.
- Aumann, R. J., 1959. Acceptable points in general n -person games. In: Tucker, A. W., Luce, R. D. (Eds.), Contributions to the Theory of Games IV. Vol. 40 of Annals of Mathematics Studies. Princeton University Press, pp. 287–324.
- Aumann, R. J., 1961. Almost strictly competitive games. *Journal of the Society of Industrial and Applied Mathematics* 9 (4), 544–550.
- Aumann, R. J., 1974. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics* 1, 67–96.
- Aumann, R. J., 1985. On the non-transferable utility value: A comment on the Roth-Shafer examples. *Econometrica* 53 (3), 667–678.
- Aumann, R. J., 1987. Game theory. In: Eatwell, J., Milgate, M., Newman, P. (Eds.), The New Palgrave: A Dictionary of Economics. Vol. 2. MacMillan, pp. 460–482.
- Brandt, F., Sandholm, T., Shoham, Y., 2007. Spiteful bidding in sealed-bid auctions. In: Veloso, M. (Ed.), Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI). pp. 1207–1214.
- Chen, X., Deng, X., 2006. Settling the complexity of 2-player Nash-equilibrium. In: Proceedings of the 47th Symposium on Foundations of Computer Science (FOCS). IEEE Press, pp. 261–272.
- Conitzer, V., Sandholm, T., 2005. Complexity of (iterated) dominance. In: Proceedings of the 6th ACM Conference on Electronic Commerce (ACM-EC). ACM Press, pp. 88–97.

- Cubitt, R., Sugden, R., 1994. Rationally justifiable play and the theory of non-cooperative games. *Economic Journal* 104 (425), 798–803.
- Daskalakis, C., Goldberg, P., 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, pp. 71–78.
- Gottlob, G., Greco, G., Scarcello, F., 2005. Pure Nash equilibria: Hard and easy games. *Journal of Artificial Intelligence Research* 24, 195–220.
- Harsanyi, J. C., 1973. Oddness of the number of equilibrium points: A new proof. *International Journal of Game Theory* 2, 235–250.
- Kats, A., Thisse, J.-F., 1992. Unilaterally competitive games. *International Journal of Game Theory* 21, 291–299.
- Khachiyan, L., 1979. A polynomial algorithm in linear programming. *Soviet Mathematics Doklady* 20, 191–194.
- Koutsoupias, E., Papadimitriou, C., 1999. Worst-case equilibria. In: *Proceedings of the 16th International Symposium on Theoretical Aspects of Computer Science (STACS)*. Vol. 1563 of *Lecture Notes in Computer Science (LNCS)*. Springer-Verlag, pp. 404–413.
- Luce, R. D., Raiffa, H., 1957. *Games and Decisions: Introduction and Critical Survey*. John Wiley & Sons Inc.
- Luckhardt, C., Irani, K., 1986. An algorithmic solution of n -person games. In: *Proceedings of the 5th National Conference on Artificial Intelligence (AAAI)*. AAAI Press, pp. 158–162.
- Marsland, A. T., Schaeffer, J. (Eds.), 1990. *Computers, Chess, and Cognition*. Springer-Verlag.
- Marx, L. M., Swinkels, J. M., 1997. Order independence for iterated weak dominance. *Games and Economic Behavior* 18, 219–245.
- Morgan, J., Steiglitz, K., Reis, G., 2003. The spite motive and equilibrium behavior in auctions. *Contributions to Economic Analysis & Policy* 2 (1), 1102–1127.
- Moulin, H., 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., 1991. *Game Theory: Analysis of Conflict*. Harvard University Press.
- Nash, J. F., 1951. Non-cooperative games. *Annals of Mathematics* 54 (2), 286–295.
- Norde, H., 1999. Bimatrix games have quasi-strict equilibria. *Mathematical Programming* 85, 35–49.
- 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)*. ACM Press, pp. 49–56.
- Papadimitriou, C. H., 2007. The complexity of finding Nash equilibria. In: Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V. (Eds.), *Algorithmic Game Theory*. Cambridge University Press, Ch. 2, pp. 29–78.
- Raghavan, T. E. S., 2002. Non-zero-sum two-person games. In: Aumann, R. J., Hart, S. (Eds.), *Handbook of Game Theory with Economic Applications*. Vol.

- III. North-Holland, Ch. 17, pp. 1687–1721.
- Roughgarden, T., 2005. *Selfish Routing and the Price of Anarchy*. MIT Press.
- Russell, S. J., Norvig, P., 2003. *Artificial Intelligence. A Modern Approach*, 2nd Edition. Prentice Hall.
- Sandholm, T., Larson, K., Andersson, M., Shehory, O., Tohmé, F., 1999. Coalition structure generation with worst case guarantees. *Artificial Intelligence* 111 (1–2), 209–238.
- Schoenebeck, G., 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, pp. 270–279.
- Sturtevant, N., 2004. Current challenges in multi-player game search. In: *Proceedings of the 4th International Conference on Computers and Games (CG)*. Vol. 3846 of *Lecture Notes in Computer Science (LNCS)*. Springer-Verlag, pp. 285–300.
- Tennenholtz, M., 2002. Competitive safety analysis: Robust decision-making in multi-agent systems. *Journal of Artificial Intelligence Research* 17, 363–378.
- Vajda, S., 1956. *Theory of Games and Linear Programming*. John Wiley & Sons Inc.
- van Damme, E., 1983. *Refinements of the Nash Equilibrium Concept*. Springer-Verlag.
- van Damme, E., 1998. On the state of the art in game theory: An interview with Robert Aumann. *Games and Economic Behavior* 24, 181–210.
- von Neumann, J., 1928. Zur Theorie der Gesellschaftspiele. *Mathematische Annalen* 100, 295–320.
- von Neumann, J., Morgenstern, O., 1947. *The Theory of Games and Economic Behavior*, 2nd Edition. Princeton University Press.
- Wolf, O. D., 1999. Optimal strategies in n-person unilaterally competitive games. Discussion Paper 9949, Center for Operations Research and Econometrics, Université catholique de Louvain.