

A Game-Theoretic Analysis of Strictly Competitive Multiagent Scenarios*

Felix Brandt Felix Fischer Paul Harrenstein

Computer Science Department
University of Munich
80538 Munich, Germany

{brandtf, fischerf, harrenst}@tcs.ifi.lmu.de

Yoav Shoham

Computer Science Department
Stanford University
Stanford CA 94305, USA

shoham@cs.stanford.edu

Abstract

This paper is a comparative study of game-theoretic solution concepts in strictly competitive multiagent scenarios, as commonly encountered in the context of parlor games, competitive economic situations, and some social choice settings. We model these scenarios as *ranking games* in which every outcome is a ranking of the players, with higher ranks being preferred over lower ones. Rather than confining our attention to one particular solution concept, we give matching upper and lower bounds for various comparative ratios of solution concepts within ranking games. The solution concepts we consider in this context are security level strategies (*maximin*), *Nash equilibrium*, and *correlated equilibrium*. Additionally, we also examine *quasi-strict equilibrium*, an equilibrium refinement proposed by Harsanyi, which remedies some apparent shortcomings of Nash equilibrium when applied to ranking games. In particular, we compute the *price of cautiousness*, *i.e.*, the worst-possible loss an agent may incur by playing maximin instead of the worst (quasi-strict) Nash equilibrium, the *mediation value*, *i.e.*, the ratio between the social welfare obtained in the best correlated equilibrium and 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.

1 Introduction

Consider the following three-player game. Alice, Bob, and Charlie independently and simultaneously are to decide whether to raise their hand or not. Alice wins if the number of players raising their hand is odd, whereas Bob wins if it is even and positive. Should nobody raise his hand, Charlie wins. What would you recommend Alice to do?

Clearly, this question lies at the heart of game theory, and game-theoretic solution concepts should be called upon when

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

trying to give a sound answer (see Section 3 for formal definitions of the concepts used in the following paragraphs). In the game described above there can be just one winner; all the other players lose. As such it is an instance of a subclass of *ranking games*, which were recently introduced as models of strictly competitive multi-player scenarios [Brandt *et al.*, 2006]. Outcomes of a ranking game are related to *rankings* of the players, *i.e.*, orderings of the players according to how well they have done in the game relative to one another. Players are assumed to generally prefer higher ranks over lower ones and to be indifferent to the ranks of other players. Formally, ranking games are defined as normal-form games in which the payoff functions represent the preferences of the players regarding lotteries over rankings. In this paper, we conduct a comparative study of game-theoretic solution concepts in ranking games.

It is well-known that *two-player* strictly competitive games admit a unique rational solution (the maximin solution), *i.e.*, a set of (possibly randomized) strategies for each player so that each player is best off playing one of the recommended strategies. Unfortunately, solution concepts for ranking games with more than two players are less appealing due to a lack of normative power. Nash equilibria, for example, which are defined as profiles of strategies that are mutual best responses to each other, may not be unique. Indeed, the game described above possesses numerous Nash equilibria: Raising her hand, not raising her hand, and mixing uniformly between both actions are all optimal strategies for Alice in some equilibrium. The only pure, *i.e.*, non-randomized, equilibrium of the game tells Alice not to raise her hand based on the belief that Bob will raise his hand and Charlie will not (see Figure 1 for an illustration). This assumption, however, is unreasonably strong. 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 weakness is due to the indifference of losers, which is inherent to ranking games. In fact, we argue that *pure* Nash equilibria are particularly weak solutions of such games and 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.

Returning to the example given at the beginning of this section, it is still unclear which strategy Alice should adopt

in order to maximize her chances of winning. We consider three solution concepts in addition to Nash equilibria: maximin strategies, quasi-strict equilibria, and correlated equilibria. By playing her maximin strategy, Alice can *guarantee* a certain chance 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). We will see that this equivalence is no mere coincidence, since in any single-winner game where a player has just two actions, the payoff in his worst quasi-strict equilibrium equals his (positive) security level. However, none of the aforementioned solution concepts offers a solution for multi-player ranking games that is as obviously right as maximin is for strictly competitive two-player games. We nevertheless facilitate the analysis of ranking games by evaluating the following comparative ratios:

- the *price of cautiousness*, *i.e.*, the worst-possible loss an agent may face when playing maximin instead of the worst Nash equilibrium,
- the *price of cautiousness for quasi-strict equilibria*, *i.e.*, the worst-possible loss an agent may face when playing maximin instead of the worst quasi-strict equilibrium,
- the *mediation value*, *i.e.*, the ratio between the social welfare obtainable in the best correlated equilibrium and 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. The interesting question is how these values unfold for games with more than two players.

The remainder of this paper is organized as follows. After reviewing related work in Section 2, we formally introduce ranking games and game-theoretic solution concepts in Section 3. Section 4 discusses a weakness of the Nash equilibrium concept that is characteristic for ranking games. Sections 5 and 6 introduce and evaluate the price of cautiousness and the value of correlation, respectively. The paper concludes with Section 7.

2 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 complete information *extensive-form* games, a class of multi-player games for which efficient Nash equilibrium search algorithms have been investigated by the AI community [*e.g.*, Luckhardt 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]. 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, while the opposite is not the case.

Ranking games were introduced by Brandt *et al.* [2006], who also showed that finding Nash equilibria of ranking games with more than two players is just as hard as for general games and thus unlikely to be feasible in polynomial time. This further underlines the importance of alternative solution concepts such as maximin strategies and correlated equilibria which can both be computed efficiently via linear programming.

Most work on comparative ratios in computational game theory has been inspired by the literature on the *price of anarchy* [Koutsoupias and Papadimitriou, 1999], *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* and the *enforcement value*) were introduced by Ashlagi *et al.* [2005]. To our knowledge, Tennenholtz [2002] was the first to conduct a numerical comparison of Nash equilibrium payoff and the security level. This work is inspired by an intriguing example game due to Aumann [1985] where the only Nash equilibrium yields each player no more than his security level, but the equilibrium strategies are actually different from the maximin strategies. In other words, the equilibrium merely yields security level payoffs but fails to guarantee them.

3 Preliminaries

3.1 Ranking Games

An accepted way to model situations of conflict and social interaction is by means of a *normal-form game* [see, *e.g.*, Myerson, 1991].

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

		c_1		c_2	
		b_1	b_2	b_1	b_2
a_1	3	1	1	2	
a_2	1	2	2	1	

Figure 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.

Unless stated otherwise, we will henceforth assume that every player has at least two different actions. A combination of actions $s \in A = \times_{i \in N} A_i$ is also called a profile of *pure strategies*. This concept can be generalized to *mixed strategy profiles* $s \in S = \times_{i \in N} S_i$, by letting players randomize over their actions. Here, $S_i = \Delta(A_i)$ denotes the set of probability distributions over player i 's actions, or *mixed strategies* available to player i . Payoff functions naturally extend to mixed strategy profiles, and we will frequently write $p_i(s)$ for the *expected* payoff of player i , and $p(s) = \sum_{i \in N} p_i(s)$ for the *social welfare*, under profile s . In the following, we further write $n = |N|$ for the number of players in a game, A_{-i} and S_{-i} for the set of action or strategy profiles for all players but i , s_i for the i th strategy in profile s , s_{-i} for the vector of all strategies in s but s_i , and $s_i(a)$ for the probability assigned to action a by player i in strategy profile s .

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

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

$$r_i^k \geq r_i^{k+1} \quad \text{for all } k \in \{1, 2, \dots, n-1\}, \quad \text{and } r_i^1 > r_i^n.$$

For convenience, we assume rank payoffs to be normalized so that $r_i^1 = 1$ and $r_i^n = 0$.

In other words, higher ranks are weakly preferred, and for at least one rank the preference is strict. Intuitively, r_i^k represents player i 's payoff for being ranked in k th. Building on Definition 2, defining ranking games is straightforward.

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

A *binary ranking game* is one where each rank payoff vector only consists of zeros and ones. An important subclass of binary ranking games are *single-winner games*, *i.e.*, games where $r_i = (1, 0, \dots, 0)$ for all $i \in N$. When considering mixed strategies, the expected payoff in a single-winner game equals the probability of winning. An example single-winner game with three players—the game introduced at the beginning of this paper—is given in Figure 1. A convenient way of representing these games is to just denote the index of the winning player for each outcome. For general ranking games, we will sometimes write $[i_1, i_2, \dots, i_n]$ to denote the outcome where player i_1 is ranked first, i_2 is ranked second, and so forth.

3.2 Solution Concepts

Over the years, game theory has produced a number of solution concepts that identify reasonable or desirable strategy profiles in a given game. Perhaps the most cautious way for a player to play a game is to try to maximize his own payoff *regardless* of which strategies the other player choose, *i.e.*, even when the other players (collaboratively) try to minimize his payoff. Such a strategy is called a *maximin strategy*, and the corresponding (guaranteed minimum) payoff is called the *maximin payoff* or *security level* of that player.

Definition 4 (Maximin strategy) A strategy $s_i \in S_i$ is called a maximin strategy for player $i \in N$ if

$$s_i \in \arg \max_{r_i \in S_i} \min_{t_{-i} \in S} p_i(r_i, t_{-i}).$$

$v_i = \max_{r_i \in S_i} \min_{t_{-i} \in S} p_i(r_i, t_{-i})$ is called the security level for 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 players 2 and 3 is zero.

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

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

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

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

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.

A weakness of Nash equilibrium as a normative solution concept (besides the multiplicity of equilibria) is that players may be indifferent between actions they play with non-zero probability and actions they do not play at all. For example, in the pure Nash equilibrium of the game in Figure 1, players 2 and 3 might as well deviate without decreasing their chances of winning the game. *Quasi-strict equilibrium* as introduced by Harsanyi [1973]¹ tries to alleviate this phenomenon by demanding that *every* best response be played with positive probability. (It follows from the definition of Nash equilibrium that every action played with positive probability yields the same expected payoff.)

Definition 6 (Quasi-strict Nash equilibrium) A Nash equilibrium $s \in S$ is called quasi-strict if for all $i \in N$ and all $a, b \in A_i$ with $s_i(a) > 0$ and $s_i(b) = 0$, $p_i(s_{-i}, a) > p_i(s_{-i}, b)$.

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

Figure 1 shows a quasi-strict equilibrium of the game between Alice, Bob, and Charlie.² While quasi-strict equilibria have recently been shown to always exist in two-player games [Norde, 1999], this is not the case for games with more than two players (see Footnote 3).

Nash equilibrium assumes that players randomize between 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. The corresponding equilibrium concept is defined as follows.

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

$$\sum_{s_{-i} \in A_{-i}} \mu(s) (p_i(s) - p(s_{-i}, a_i)) \geq 0.$$

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 s is chosen according to this distribution, and every player $i \in N$ is only informed about his own action $s_i \in A_i$, it is optimal for i to play s_i , given that the other players play s_{-i} . Correlated equilibrium is based upon the assumption that there exists a trustworthy party who can recommend behavior but cannot enforce it.

It can easily be seen from the definition that the Nash equilibria of any game form a subset of the correlated equilibria, with the additional property of being a product of strategies for the individual players. The existence result by Nash [1951] thus carries over to correlated equilibria. Again consider the game of Figure 1. It is easily verified that 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 it is the case for the game in Figure 1.

4 Equilibrium Points in Ranking Games

As we have already seen in Section 1, the stability of some Nash equilibria in ranking games is questionable because losing players are assumed to play certain strategies even though they could as well play *any* other strategy without decreasing their payoff. By definition, there is at least one player—the

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

	c_1		c_2	
	b_1	b_2	b_1	b_2
a_1	2	1	3	1
a_2	1	2	1	1

Figure 2: Three-player single-winner game. Dashed boxes denote all Nash equilibria (one player may mix arbitrarily in boxes that span two outcomes).

one ranked lowest—in any outcome, who receives his minimum payoff of zero and therefore has no incentive to actually play that particular action. As a consequence, all *pure* equilibria are weak in this sense, especially in single-winner games where $n - 1$ players are indifferent over which action to play. Quasi-strict equilibrium mitigates this phenomenon by additionally requiring that actions played with positive probability yield strictly more payoff than non-equilibrium actions. Thus, quasi-strict equilibrium can be used to formally illustrate the weakness of pure Nash equilibrium.

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.*

There is at least one quasi-strict equilibrium in every *two-player* game (and thus also in every two-player ranking game) [Norde, 1999]. In games with more than two players, there may be no quasi-strict equilibrium. Figure 2 shows that this even holds for single-winner games.³

It appears as if most ranking games possess non-pure equilibria, 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 e in which player 1 wins. Since player 2 must be incapable of increasing his payoff by deviating from e , player 1 has to win no matter which action the second player chooses.

³There are 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 2 might be of independent interest.

As a consequence, the strategies in e 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 $e = (a_1, b_1, c_1)$ 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 e 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 profiles e_1 where player 2 mixes uniformly between e and (a_1, b_2, c_1) and e_2 where player 3 mixes uniformly between e and (a_1, b_1, c_2) . Since player 1 does not win in (a_2, b_1, c_1) , he will not deviate from either e_1 or e_2 even when 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 e_1 *not* to be an equilibrium. Analogously, for e_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. Conversely assume that e is *not* semi-strict. When any of the action profiles in $E = \{(a_2, b_1, c_1), (a_1, b_2, c_1), (a_1, b_1, c_2)\}$ 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 E not to contain any pure equilibria, there must be (exactly) one player for every profile in E who deviates to a profile in $D = \{(a_2, b_2, c_1), (a_2, b_1, c_2), (a_1, b_2, c_2)\}$ because the game is a single-winner game and because e is not semi-strict. This implies two facts: First, action profile $e' = (a_2, b_2, c_2)$ is a pure equilibrium because no player will deviate from e' to any profile in D . Second, the player who wins in e' strictly prefers the equilibrium outcome over the corresponding action profile in D , implying that e' 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 zero payoff, 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 were unable to find a single-winner game that only contains pure equilibria, even when employing a computer program that checked tens of thousands of games. However, a general existence result has so far tenaciously resisted proof.

5 The Price of Cautiousness

Despite its conceptual elegance and simplicity, Nash equilibrium has been criticized on various grounds. In the common

case of multiple equilibria, it is unclear which one to play; coalitions might benefit from jointly deviating; and recent complexity-theoretic results indicate that there might exist no polynomial-time algorithm for finding Nash equilibria [Chen and Deng, 2006]. Adding the indifference of players, which is particularly problematic in ranking games, a compelling question is how much worse a player can be off when reverting to the most defensive choice—his maximin strategy—instead of hoping for an equilibrium outcome. We refer to this value by the *price of cautiousness*. In the following, let \mathcal{G} denote the set of all normal-form games and for $\Gamma \in \mathcal{G}$ let $N(\Gamma)$ denote the set of Nash equilibria of Γ .

Definition 8 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, thus capturing the worst-case loss the player may experience by playing his maximin strategy instead of a Nash equilibrium. For a player whose security level is 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 already mentioned in Section 1, the price of cautiousness in two-player ranking games is 1 due to the Minimax Theorem [von Neumann and Morgenstern, 1947]. In general ranking games, the price of cautiousness is unbounded. The proof of the following theorem is omitted for reasons of limited space.

Theorem 2 Let \mathcal{R} be the class of ranking games with more than two players that involve player i . Then, $PC_i(\mathcal{R}) = \infty$, even if \mathcal{R} only contains games without weakly dominated actions. \square

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

Theorem 3 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. By the requirement that the security level must be strictly positive, we have that for every opponent action profile s_{-i} there must be some action 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, resulting in a price of cautiousness of at most k .

	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)

Figure 3: Three-player ranking game Γ_1 used in the proof of Theorem 3

For a matching lower bound, again consider the single winner game of Figure 2. 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) and yield payoff 1 for player 1, twice as much as his security level of 0.5. For this, we look at the possible strategies for player 3. If player 3 plays c_1 , the game reduces to the well-known matching pennies game 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 must 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_1 of Figure 3, which is a ranking game for rank payoff vectors $r_1 = r_2 = (1, 0, 0)$ and $r_3 = (1, 1, 0)$, 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_1 is weakly dominated and that $v_1(\Gamma_1) = 0.5$. On the other hand, we will argue that all Nash equilibria of Γ_1 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. For 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 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, if player 3 randomizes between actions c_1 and c_2 , he must again get the same payoff from both of these actions. This can only be the case if either player 1 plays a_1 and player 2 plays b_2 , or if player 1 randomizes between a_1 and a_2 and player 2 plays b_1 . In the former case, player 2 will deviate to b_1 . In the latter case, player 1 will deviate to a_2 .

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 can obtain in Nash equilibrium can be at most k times his security level. The proof relies on equilibria in which the payoff of at least one player is 1. As we have already pointed out in Section 4, such equilibria (like pure equilibria) are particularly weak. We therefore also study the price of cautiousness with respect to *quasi-strict equilibria*.

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 with respect to quasi-strict equilibria for player i in Γ is defined as

$$PC_i^{QS}(\Gamma) = \frac{\min \{ p_i(s) \mid s \in N_{QS}(\Gamma) \}}{v_i(\Gamma)},$$

where $N_{QS}(\Gamma)$ denotes the set of quasi-strict equilibria in Γ .

As before, $PC_i^{QS}(C) = \sup_{\Gamma \in C} PC_i^{QS}(\Gamma)$.

Returning to the binary ranking game of Figure 3 and its generalizations, it turns out that player 2 can do nothing about the fact that he always loses in every Nash equilibrium. As a consequence, all Nash equilibria where every action profile with payoff $(1, 0, 1)$ is played with positive probability are quasi-strict, and the price of cautiousness in binary ranking games remains k when restricting attention to quasi-strict equilibria. In single-winner games, on the other hand, a slight decrease in the price of cautiousness can be witnessed. This is due to the fact that there can be no quasi-strict equilibrium in which only one player wins (see also Fact 1).

Theorem 4 Let \mathcal{R}_b be the class of single-winner games with more than two players involving a player i with exactly k actions. Then, $PC_i^{QS}(\mathcal{R}_b) = k - 1$.

Proof: Like in the proof of Theorem 3, an upper bound for the price of cautiousness can be found by letting the numerator and denominator take their maximum and minimum, respectively. As before, the lowest positive security value is $1/k$ for a player with k actions. The argument for a useful upper bound on the payoff in a quasi-strict equilibrium is slightly more delicate. We start by observing that the existence of a quasi-strict equilibrium in which a player (say, player 1) obtains payoff 1 implies that this player has a *winning action*, i.e., an action which always yields payoff 1 regardless

of the other players' actions. This is seen as follows. In a single-winner game, a payoff of 1 for player 1 means that all other players get payoff zero. In a quasi-strict equilibrium, all players have to receive strictly more payoff for equilibrium actions than for actions that are not contained in the equilibrium's support. For this reason, all losing players have to randomize over *all* their actions in a quasi-strict equilibrium in which player 1 wins. This implies that player 1 must have an action that guarantees him a win, and thus his security level is 1.

Since a maximum security level is useless for finding a reasonable upper bound, we restrict our attention to games where no player has a security level of 1. According to our previous argument, there can be no quasi-strict equilibrium in such games where only one player wins. We claim that the highest payoff less than 1 that player 1 may obtain in the worst equilibrium is $(k-1)/k$. First, we observe that we can restrict our attention to equilibria e that do not contain action profile $b \in \times_{i \neq 1} A_i$ by all players except player 1, so that player 1 wins no matter which action he chooses. Whenever b is part of an equilibrium, there must be another equilibrium where b is not played, but that is otherwise identical. Obviously, player 1 cannot get more payoff in this new equilibrium than in the original one. Now assume for contradiction that the payoff to player 1 in e is greater than $(k-1)/k$. For any action a_i that player 1 plays in equilibrium, the sum of probabilities that the other players put on all actions profile $b \in \times_{i \neq 1} A_i$ such that player 1 wins in action profile (a_i, b) must be greater than $(k-1)/k$. Let $Z_j \subseteq \times_{i \neq 1} A_i$ denote the set of all remaining action profiles, *i.e.*, those combinations of actions by other players where player 1 loses. Clearly, the sum of probabilities for all action profiles in Z_j must be strictly less than $1/k$. On the other hand, since player 1 loses at least once for every action profile of the other players, the union of all sets Z_j equals the set of all action profiles played in e , and the probabilities of these actions must sum up to 1, yielding a contradiction.

As for a matching lower bound, consider the single-winner game involving Alice, Bob, and Charlie that is shown in Figure 1. Alice's payoff in the quasi-strict equilibrium marked by the dotted rectangles is 0.5, while her security level of 0.5 implies that there cannot be an equilibrium with lower payoff. For $k > 2$, we instead use 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

$$p((a_i, b_j, c_\ell)) = \begin{cases} (0, 1, 0) & \text{if } \ell = 1 \text{ and } i \neq j \\ (0, 0, 1) & \text{if } \ell = 2 \text{ and } i = j \\ (1, 0, 0) & \text{otherwise.} \end{cases}$$

It is easily verified that the strategy profile where player 3 plays c_2 and players 1 and 2 randomize uniformly between all of their actions is a quasi-strict Nash equilibrium and in fact the only Nash equilibrium of this game. The payoff of player 1 in this equilibrium is $(k-1)/k$, $(k-1)$ times his security level of $1/k$. \square

Applying this theorem to a single-winner game which contains a quasi-strict equilibrium, a player with only two actions at his disposal will not obtain more payoff than his (positive) security level in some quasi-strict equilibrium.

6 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 $v_X(\Gamma) = \max\{p(s) \mid s \in X\}$.

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 $C \subseteq \mathcal{G}$ of games, we further write $MV(C) = \sup_{\Gamma \in C} MV(\Gamma)$ and $EV(C) = \sup_{\Gamma \in C} EV(\Gamma)$.

Ashlagi *et al.* [2005] have shown that both the mediation value and the enforcement value cannot be bounded for any class of 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, but 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 $r_i^k = r_j^k$ for all $i, j \in N$ and $1 \leq k \leq n$, are constant-sum games. Hence, the social welfare is the same in every outcome so that both the mediation value and the enforcement value are 1. This particularly 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 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. For the former, we show that for any $n \geq 3$ there is a ranking game where all Nash equilibria have social welfare 1 while there is a correlated equilibrium with social welfare $n-1$. In particular, we

exploit the fact that a Nash equilibrium has to be a product of strategies for the individual players and design a game where one of the players strictly prefers a designated action given that the other players play a strategy profile involving an outcome with high social welfare, while the same is not the case for a certain correlated strategy. The proof itself is rather involved, and is omitted for reasons of limited space.

Theorem 5 *Let \mathcal{R}' be the class of ranking games with more than two players such that at least one player has more than two actions when there are only three players. Then, $MV(\mathcal{R}') = n - 1$.* \square

In order to match the upper bound of the enforcement value, we design a ranking game that has social welfare $n - 1$ for a single action profile and social welfare $1 + \epsilon$ for all others. To show that there exists no correlated equilibrium with social welfare larger than $1 + \epsilon$, the problem of finding a social welfare maximizing correlated equilibrium is written as a linear program and then transformed into its dual. Since the dual constitutes a minimization problem, it suffices to find a feasible solution with objective value $1 + \epsilon$. We again omit the details of the proof.

Theorem 6 *Let \mathcal{R} be the class of ranking games with more than two players. Then, $EV(\mathcal{R}) = n - 1$, even if \mathcal{R} only contains games without weakly dominated actions.* \square

7 Conclusion

We have quantified and bounded comparative ratios between various solution concepts in ranking games. It turned out that playing one's maximin strategy in binary ranking games with only few actions available 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. Moreover, maximin strategies can be computed in polynomial time while all known algorithms for computing Nash equilibria have exponential worst-case complexity.

In the second part of the paper, we have 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 have the advantage of being polynomial-time computable and do not suffer from the equilibrium selection problem since the equilibrium to be played is selected by a mediator.

References

- I. Ashlagi, D. Monderer, and M. Tennenholtz. On the value of correlation. In *Proc. of 21st UAI Conference*, pages 34–41. AUAI Press, 2005.
- R. Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1:67–96, 1974.
- R. Aumann. On the non-transferable utility value: A comment on the Roth-Shafer examples. *Econometrica*, 53(3):667–678, 1985.
- F. Brandt, F. Fischer, and Y. Shoham. On strictly competitive multi-player games. In Y. Gil and R. Mooney, editors, *Proc. of 21st AAAI Conference*, pages 605–612. AAAI Press, 2006.
- X. Chen and X. Deng. Settling the complexity of 2-player Nash-equilibrium. In *Proc. of 47th FOCS Symposium*. IEEE Press, 2006. To Appear.
- R. Cubitt and R. Sugden. Rationally justifiable play and the theory of non-cooperative games. *Economic Journal*, 104(425):798–803, 1994.
- J. C. Harsanyi. Oddness of the number of equilibrium points: A new proof. *International Journal of Game Theory*, 2:235–250, 1973.
- E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *Proc. of 16th STACS*, volume 1563 of *LNCS*, pages 404–413. Springer, 1999.
- C. Luckhardt and K. Irani. An algorithmic solution of n -person games. In *Proc. of 5th AAAI Conference*, pages 158–162. AAAI Press, 1986.
- A. T. Marsland and J. Schaeffer, editors. *Computers, Chess, and Cognition*. Springer, 1990.
- R. B. Myerson. *Game Theory: Analysis of Conflict*. Harvard University Press, 1991.
- J. F. Nash. Non-cooperative games. *Annals of Mathematics*, 54(2):286–295, 1951.
- H. Norde. Bimatrix games have quasi-strict equilibria. *Mathematical Programming*, 85:35–49, 1999.
- S. J. Russell and P. Norvig. *Artificial Intelligence. A Modern Approach*. Prentice Hall, 2nd edition, 2003.
- T. Sandholm, K. Larson, M. Andersson, O. Shehory, and F. Tohmé. Coalition structure generation with worst case guarantees. *Artificial Intelligence*, 111(1–2):209–238, 1999.
- N. Sturtevant. Current challenges in multi-player game search. In *Proc. of 4th International Conference on Computers and Games (CG)*, volume 3846 of *LNCS*. Springer, 2004.
- M. Tennenholtz. Competitive safety analysis: Robust decision-making in multi-agent systems. *Journal of Artificial Intelligence Research*, 17:363–378, 2002.
- E. van Damme. *Refinements of the Nash Equilibrium Concept*. Springer, 1983.
- J. von Neumann and O. Morgenstern. *The Theory of Games and Economic Behavior*. Princeton University Press, 2nd edition, 1947.