Ranking Games that have Competitiveness-based Strategies*

Leslie Ann Goldberg
Dept. of Computer Science
University of Liverpool, U.K.
L.A.Goldberg@liverpool.ac.uk

Paul W. Goldberg
Dept. of Computer Science
University of Liverpool, U.K.
P.W.Goldberg@liverpool.ac.uk

Piotr Krysta
Dept. of Computer Science
University of Liverpool, U.K.
P.Krysta@liverpool.ac.uk

Carmine Ventre
Dept. of Computer Science
University of Liverpool, U.K.
Carmine.Ventre@liverpool.ac.uk

ABSTRACT

This paper studies —from the perspective of efficient computation— a type of competition that is widespread throughout the plant and animal kingdoms, higher education, politics and artificial contests. In this setting, an agent gains utility from his relative performance (on some measurable criterion) against other agents, as opposed to his absolute performance. We model this situation using ranking games in which each strategy corresponds to a level of competitiveness, and incurs an upfront cost that is higher for more competitive strategies. We study the Nash equilibria of these games, and polynomial-time algorithms for computing them. For games in which there is no tie between agents' levels of competitiveness we give a polynomial-time algorithm for computing an exact equilibrium in the 2-player case, and a characterization of Nash equilibria that shows an interesting parallel between these games and unrestricted 2-player games in normal form. When ties are allowed, via a reduction from these games to a subclass of anonymous games, we give polynomial-time approximation schemes for two special cases: constant-sized set of strategies, and constant number of players. The latter result is improved to a fully polynomial-time approximation scheme when the constant number of players only compete to win the game, i.e. to be ranked first.

*This work was supported by EPSRC grant EP/G069239/1 "Efficient Decentralised Approaches in Algorithmic Game Theory."

EC'10, June 7–11, 2010, Cambridge, Massachusetts, USA.

Categories and Subject Descriptors

F.2 [**Analysis of Algorithms and Problem Complexity**]; J.4 [**Social and Behavioral Sciences**]: Economics

General Terms

Algorithms, Economics, Theory

Keywords

Anonymous games, Computation of exact and approximate Nash equilibria

1. INTRODUCTION

We consider a class of games in which each action available to a player has two associated quantities — the effort (or cost) of making that action, and a "return" that the player gains — return is a monotonically increasing function of effort. Players derive utility, not so much from maximising the return, as from getting a higher return than other players — given a pure profile, players are ranked according to their returns, and a player's payoff is a function of his rank, minus his effort. By way of illustration, consider a set of athletes who are training for a race. Each athlete spends time and effort in training, and has an increasing function that maps this upfront cost to performance (speed in the race). His total utility is the value of the prize that he wins, minus the cost of making this initial effort. Note that the prize is awarded based on his speed relative to other competitors, with no consideration given to his speed taken in isolation.

A *ranking game* [2] is a game in which the outcome of any pure-strategy profile is a rank-ordering of the players, and a player's utility is a decreasing function of his position in the rank-ordering. In [2], the value to player i of the k-th position in a ranking depends on both i and k — in this paper we assume it just depends on k. But the more significant restriction that we introduce, is the notion of actions that correspond to competitiveness — a more competitive action has a higher upfront cost but may result in a higher position in the ranking.

In our model, the socially optimal behavior for the players is to all play at minimal level of competitiveness (this is due

	# players	# prizes	# actions	Result
Return-symmetric games	O(1)	O(1)	any	PTAS (Thm. 6)
	any	any	O(1)	PTAS (Thm. 7)
	O(1)	1	any	FPTAS (Thm. 8)
	any	any	2	Exact Pure (Thm. 5)
Games without ties	2	2	any	Exact (Thm. 3)
Linear-prize games	any	# players	any	Exact (Thm. 4)

Table 1: Our algorithmic contributions in a nutshell.

to the assumption that any behavior will lead to a ranking of the players, and the award of associated prizes for each rank position; and we noted above that values of positions in the ranking does not depend on which players obtained them). However, in a Nash equilibrium, players will typically be more competitive (as in for example [14], in the context of a continuous game). Competitiveness amongst the players is supposed to result in a positive externality — in the context of spectator sports, the spectators prefer to see a well-run race, or in the context of the DARPA Grand Challenge (or other similar contests), the competitiveness is supposed to generate research progress. Ultimately our hope is that this work will be applicable to the problem of allocating a prize fund amongst a set of separate awards, so as to maximise the competitiveness. A related line of research [13, 4] has addressed from a game-theoretic perspective, the impact of the point scoring system on offensive versus defensive play in the UK Premier League; our concern here is slightly different, being focused on highly competitive versus weakly competitive play.

1.1 The computational perspective

The PPAD-completeness results for unrestricted normal-form games [8, 5] are understood to indicate that it is hard to find a Nash equilibrium, and they motivate the study of more specific types of game that both model a form of competition in the real world, and are computationally easier to solve. The existence of a fast, robust algorithm makes a solution concept more plausible, and holds out the hope that a solution may indeed be obtained by a natural, decentralised process.

One motivation for studying this subclass of ranking games is that while unrestricted ranking games are hard [2] (even for just 3 players), here we obtain some encouraging positive results. Furthermore, for a large number of players, note that (in contrast with unrestricted ranking games) our restriction to competitiveness-based strategies admits a concise representation of games of many players. Indeed, this conciseness applies to the more general setting where there are many players with player-dependent values for the positions in the ranking.

Under some circumstances, Nash equilibria of these games can be studied without loss of generality by considering a further restriction in which any player's strategies have the same set of returns (but not necessarily the same costs) as any other player. These games, that we call "return-symmetric" are a class of *anonymous games*. Anonymous games are games where a player's payoff depends on his own action and the distribution of actions taken by the other players, but not the identities of players who chose each action. Anonymous games admit PTAS's [7, 9] but may be PPAD-complete to solve exactly. The reason why we do not universally apply this reduction to return-symmetric games is that it leads to a blowup in the number of strategies per player, and obscures the interesting special case in which players may not tie for a (shared) position in the ranking.

1.2 Our Contribution

In addition to the return-symmetric games noted above, we consider an alternative special case that is in a sense the opposite – the case of games without ties, which may arise when a tie-breaking rule has been specified in advance, or equivalently when all return values are distinct. Table 1 gives our algorithmic results for both classes of games, including a special class of games without ties (see Section 3.3). We show in Section 3.4 that any competitiveness-based ranking game can be reduced to a return-symmetric game, thus implying that Nash equilibria of these games can be considered without loss of generality, so long as we are not restricting to a fixed number of strategies per player. We note that in the literature, algorithmic results for return-symmetric games are known [9, 7] although with the focus on the more general class of anonymous games. In particular, [9] provides a PTAS in the case of constant number of strategies. However, since our PTAS is tailored to a more specific class of games we are able to give a conceptually simpler and more efficient algorithm with the same approximation guarantee. [7] improves the efficiency of the PTAS in [9] for anonymous games in which players have only 2 strategies. Under this hypothesis, return-symmetric games admit pure equilibria that can be computed efficiently as from Theorem 5.

In addition to the algorithmic results we also give a characterization of Nash equilibria for games without ties and a single prize (see Theorem 1). Such a characterization allows to show an interesting parallel with general 2-player normal form games: a Nash equilibrium for games without ties can be computed in polynomial-time, given knowledge of its support (see Theorem 2).

2. MODEL, NOTATION AND SOME ILLUSTRATIVE EXAMPLES

We work in a classical game-theoretic setting of a finite number of players, each with a finite number of actions[1], and

[1]It may also be interesting to model this kind of competition

we consider the problem of computing Nash equilibria, and approximate Nash equilibria, for these games. Section 2.3 gives the background definitions of Nash and approximate Nash equilibrium. First, we mention a couple of items of terminology in Section 2.1, and the in Section 2.2 we specify in detail the class of games that we study, and introduce some notation. Section 2.4 shows some examples to illustrate various technical issues.

2.1 Terminology

By a *prize* we refer to the reward that a player gains from obtaining a specified position in the ranking, and this relates directly to the standard usage of "first prize", "second prize" etc in competitions.

We say that an action is "stronger" or "more competitive" than another one, if its return (called also strength) is higher. Any pair of actions are comparable in this sense, whether or not they belong to the same player.

2.2 Definition and notation

In a ranking game with competitiveness-based strategies we have d players. Player i has n actions; they will be denoted a^i_1,...,a^i_n in increasing order of competitiveness. Each action has two numeric associated quantities: a cost and a return (strength). The cost of a^i_j will be denoted c^i_j and the return (strength) r^i_j; for $j < n$ we have $c^i_j < c^i_{j+1}$ and $r^i_j < r^i_{j+1}$, by the assumption that the actions are indexed in increasing order of competitiveness. Note that *strict* monotonicity of costs and returns is without loss of generality: if two different actions (i) have the same cost and different returns then the most rewarding dominates the other and (ii) have the same return and different costs then the cheaper will dominate the more expensive. We let u_k be the value of the k-th prize. Prizes are monotonically decreasing: $u_k \geq u_{k+1}$.

The outcome of the game consists in a ranking of the players according to the returns. A player whose position in the ranking is k gets awarded the prize u_k. However, in the event of a tie (where two or more players obtain the same return and are ranked equal) the prizes that would result from tie-breaking are shared. For example when two players choose actions with the same highest return in the game so as to be both ranked first, then each will be awarded half of the first prize plus half of the second prize.

The payoff to a player will be the value of the prize he is awarded, minus the cost of the action selected by that player.

2.3 Exact and approximate Nash equilibria

Here we give the definitions of Nash equilibrium and approximate Nash equilibrium, also some of the notation we use throughout. Let S_i be the set of player i's pure strategies; $S_i = (a^i_j)_j$. Let $S = S_1 \times \ldots \times S_d$ be the set of pure-strategy profiles, where we recall d denotes the number of players. It is convenient to define $S_{-i} = S_1 \times \ldots \times S_{i-1} \times S_{i+1} \times \ldots \times S_d$ as the set of pure-strategy profiles of all players but i.

A *mixed strategy* for player i is a distribution on S_i, that is, real numbers $x^i_j \geq 0$ for each strategy $a_j \in S_i$ such that $\sum_{j\in S_i} x^i_j = 1$. If d is the number of players in a game, a set of d mixed strategies (one for each player) is a *mixed strategy profile*. u^i_s denotes the utility to player i in strategy profile s. A mixed strategy profile $\{x^i_j\}_{j\in S_i}, i = 1,\ldots,d$, is called a *(mixed) Nash equilibrium* if, for each i, $\sum_{s\in S} u^i_s x_s$ is maximised over all mixed strategies of i —where for a strategy profile $s = (s_1,\ldots,s_d) \in S$, we denote by x_s the product $x^1_{s_1} \cdot x^2_{s_2} \cdots x^d_{s_d}$. (The notation x_s naturally extends to strategy profiles $s \in S_{-i}$.) That is, a Nash equilibrium is a set of mixed strategies from which no player has a unilateral incentive to deviate. It is well-known (see, e.g., [17]) that the following is an equivalent condition for a set of mixed strategies to be a Nash equilibrium:

$$\sum_{s\in S_{-i}} u^i_{js} x_s > \sum_{s\in S_{-i}} u^i_{j's} x_s \Longrightarrow x^i_{j'} = 0. \qquad (1)$$

The summation $\sum_{s\in S_{-i}} u^i_{js} x_s$ in the above equation is the expected utility of player i if i plays pure strategy j and the other players use the mixed strategies $\{x^{i'}_j\}_{j\in S_{i'}}, i' \neq i$. Nash's theorem [16] asserts that *every game has a Nash equilibrium*.

We say that a set of mixed strategies x is an *ϵ-approximately well supported Nash equilibrium*, or *ϵ-Nash equilibrium* for short, if, for each i, the following holds:

$$\sum_{s\in S_{-i}} u^i_{js} x_s > \sum_{s\in S_{-i}} u^i_{j's} x_s + \epsilon \Longrightarrow x^i_{j'} = 0. \qquad (2)$$

Condition (2) relaxes that in (1) by allowing a strategy to have positive probability in the presence of another strategy whose expected payoff is better by at most ϵ.

2.4 Some examples

We consider some examples that should be helpful in understanding the model and issues arising. Example 1 shows that the games we consider do not always have *pure* Nash equilibria.

EXAMPLE 1. *Consider two players, for $i = 1, 2$ player i has two actions a^i_1 and a^i_2. Suppose the row player (player 1) is stronger than the column player in the sense that the column player only wins by playing a^2_2 while the row player plays a^1_1 — this can be achieved by setting $r^2_1 = 2$, $r^1_1 = 3$, $r^2_2 = 4$, $r^1_2 = 5$. Assume the costs are $c^i_1 = 0$, $c^i_2 = \frac{1}{2}$ for both players $i = 1, 2$. We have a single prize worth 1, i.e., $u_1 = 1$ and $u_2 = 0$. We have payoff matrix:*

	a^2_1	a^2_2
a^1_1	$(1, 0)$	$(0, \frac{1}{2})$
a^1_2	$(\frac{1}{2}, 0)$	$(\frac{1}{2}, -\frac{1}{2})$

It is easily checked that this game has no pure Nash equilibrium and that the unique equilibrium is the one in which both players mix uniformly.

Example 2 is a kind of anonymous game with binary actions [1], although [1] studies a continuum of players. The example shows that in this kind of game, there may be multiple equilibria, and the number of equilibria may be exponential in the number of players.

EXAMPLE 2. *Consider a symmetric game with an even number $d \geq 4$ of players; a single prize worth 1 unit; each player i has two actions a_1 and a_2 with costs $c_1 = 0$ and $c_2 = c$ (we do not have a superscript to identify a player,*

with *continuous games*, but here we discretise the set of actions available to the players.

since the games are symmetric). The prize will be shared between players who use a_2, or all players if they all use a_1.

Notice first that for $c \in \left(0, \frac{1}{d}\right]$, there is a pure equilibrium in which all players play a_2. For $c \in \left(\frac{1}{d}, 1\right)$ there is, by symmetry, a fully-mixed Nash equilibrium where all players play a_2 with the same probability. This can be seen by the following argument. Suppose each player, but the first, plays a_2 with probability p. The first player has an incentive to play a_2 which is decreasing in p. In particular, for $p = 1$ player 1 has no incentive to play a_2, while for $p = 0$ player 1 has an incentive to play a_2. Then, by continuity, there exists a value of p, say p^*, for which player 1 is indifferent between a_1 and a_2. In a profile in which all players play a_2 with probability p^* all players are indifferent by symmetry.

Now put $c = \frac{2}{d} - \epsilon$, where $\epsilon < \frac{2}{d^2+2d}$. We claim that there are also pure *Nash equilibria where any subset of size $\frac{d}{2}$ play pure a_2 and the others play pure a_1. A player playing a_2 obtains utility $-c + \frac{2}{d} > 0$; no incentive to switch to a_1. A player playing a_1 obtains utility 0, and by switching to a_2 would obtain utility $-c + 1/(\frac{d}{2}+1) = \epsilon - \frac{2}{d} + \frac{2}{d+2} < 0$. Indeed there are also many mixed equilibria where a subset of the players play pure a_1 and the other players all use the same probabilities.*

Observe that there are no Nash equilibria where players may mix with different probabilities — two such players would both be indifferent between a_1 and a_2, but their expected payoffs from playing a_2 would have to differ.

The following example shows that in a game of this type, there may be a single Nash equilibrium where some but not all players have pure strategies.

Example 3. *Consider Example 1 with an additional 3rd player (player 3) with actions a_1^3 and a_2^3 and costs $c_1^3 = 0$, $c_2^3 = 0.9$. Player 3's returns are chosen such that he can only win if he plays a_2^3 while the others play a_1^1 and a_1^2 (for example, choose $r_1^3 = 0$, $r_2^3 = 3.5$). Then player 3 plays pure a_1^3 while players 1 and 2 behave as in Example 1. It is not hard to convince oneself that players 1 and 2 assign enough probability to their "competitive" actions a_2^1 and a_2^2 that player 3 is too unlikely to win, to have an incentive to compete by playing a_2^3.*

The following example shows that there is no bound on the price of anarchy and on the price of stability in these games.

Example 4. *Consider a symmetric game with 2 players; a single prize worth 1 unit; each player i has two actions a_1 and a_2 with costs $c_1 = 0$ and $c_2 = 1/2 - \epsilon$, where $\epsilon > 0$ is a small given number. The payoff for playing (a_2, a_2) is ϵ for both players, and it is higher than the payoffs obtained by deviating from the strategy a_2: player 1 has a payoff of 0 for strategy profile (a_1, a_2) and so does player 2 for strategy profile (a_2, a_1). Thus, strategy profile (a_2, a_2) is a pure Nash equilibrium and its social welfare is 2ϵ. Now, notice that a_2 is actually a strictly dominant strategy for both players thus implying that no other action profile is a Nash equilibrium. The action profile that maximises the social welfare is when both players play their less competitive actions, (a_1, a_1), and its social welfare is 1. Thus, the price of anarchy in this game is $1/(2\epsilon)$. By uniqueness of Nash equilibrium it is also the price of stability of this game. Because ϵ can be chosen arbitrarily small, both price of anarchy and price of stability are unbounded. Note that this game is essentially the Prisoner's Dilemma in which a_1 is the collaborating strategy and a_2 is the defecting one.*

Finally, we note that certain zero-sum games do not belong to the class of games we study here. Consider Matching Pennies. In order for a ranking-games with competitiveness-based strategies to be zero-sum, the costs of any player's strategies must be all the same. In that case, the stronger strategy will dominate the weaker. But, Matching Pennies has no dominating strategies. That observation could apply to other games like rock-paper-scissors that are zero-sum and have no dominating strategies.

3. ALGORITHMS AND PROOFS

We start by noting some preprocessing steps that establish some useful assumptions that we can make without loss of generality. We continue in Section 3.2 by considering separately the special case where players cannot tie for a position in the ranking; this case would arise in competitions that have a tie-breaking rule, and note that a tie-breaking rule could be expressed by adjusting the return values r_j^i so that they are all distinct and select the appropriate winner. The reason for a focus on the tie-free case is that the analysis is simpler and the Nash equilibria turn out to have a special structure. Section 3.3 applies a result of [10] for polymatrix games, to the special case where prize values decrease linearly as a function of rank position.

In Section 3.4 we study the more general case where players may tie for a position in the ranking. We show that we can focus without loss of generality on an anonymous subclass of these games. Pure Nash equilibria of these games are studied in Section 3.4.2. In Sections 3.4.3 and 3.4.4 we give polynomial-time approximation schemes for fixed number of players and strategies, respectively. Finally, in Section 3.4.5 we give a fully polynomial-time approximation scheme for the case of constantly-many players and a single prize.

3.1 Preprocessing

By linear rescaling we may assume that the top prize has a value of 1 and the bottom prize has a value of 0, i.e. $u_1 = 1$ and $u_d = 0$. We may furthermore additively rescale individual players' payoffs such that the weakest pure strategy a_1^i of every player i has a cost of zero, i.e. $c_1^i = 0$ for all i. (To do this, add c_1^i to *all* payoffs of i.) We also assume that no player has an action with a cost greater than 1, since such an action would be dominated by a_1^i.

Note also that the numerical values of the returns r_j^i may be modified without affecting the payoffs and Nash equilibria of the game, provided only that the modification does not affect which are greater than which (in which case the ranking of the players is preserved). However it is usually convenient to specify numerical r_j^i values when describing a game.

Finally, we state a useful fact that will be used to obtain polynomial-time algorithms that return approximate Nash equilibria.

Observation 1. *Let k be an integer and let $\epsilon = 1/k$. Given a probability vector $\mathbf{x}$, it is possible to define a probability vector $\tilde{\mathbf{x}}$, called an ϵ-rounding of $\mathbf{x}$, in which each entry is equal to the corresponding entry in $\mathbf{x}$ rounded either up or down to the nearest non-negative integer multiple of ϵ.*

PROOF. Arrange the quantities in $\mathbf{x}$ to be rounded in some arbitrary order. Round the first one up to the nearest integer multiple of ϵ. We follow the fixed order, and when rounding each subsequent quantity we aim to maintain the property that the total rounding error from rounding up, minus the total rounding error from rounding down, should be strictly less than ϵ in absolute value. It is not hard to see that one can always round each value either up or down so as to maintain this invariant. Provided that $1/\epsilon$ is an integer, we must end with fractions that sum to 1. $\square$

3.2 Games without ties and single prize

We begin by observing that we may restrict our attention to actions' costs that are *strictly* less than 1.

OBSERVATION 2. *Assume player i has an action a^i_j such that $c^i_j = 1$. This action is* weakly *dominated by c^i_1. Therefore, we can eliminate a^i_j from the game at the price of eliminating some potential Nash equilibria.*

Assuming that costs are strictly less than 1, we show that Nash equilibria of games without ties have a nice structure when there is a single prize.

THEOREM 1. *Suppose there is a single prize and assume that actions' costs are less than 1. If no two actions have the same strength (so that ties are impossible) then in any Nash equilibrium*

1. *There is just one player with positive expected payoff; all others have payoff zero.*
2. *The player with positive payoff is the one with the strongest action with a cost of less than 1.*

PROOF. As from the preprocessing steps noted above assume without loss of generality that the single prize has value 1, the costs of all actions lie in the range $[0, 1]$, and each player has an action with cost 0. Let $\mathcal{N}$ be a Nash equilibrium. For each player i let w_i be the weakest action of i that lies in the support of $\mathcal{N}$; thus, i has positive probability of using w_i, and all other actions that i uses with positive probability are stronger than w_i.

For any action a, let $r(a)$ denote the return of that action (the "strength" that is used to rank players). Let p be the player whose weakest action is stronger than all other players' weakest actions, thus $r(w_p) > r(w_{p'})$ for all $p' \neq p$.

Note that for any player $p' \neq p$, the expected payoff to p' from using action $w_{p'}$ is non-positive, $w_{p'}$ cannot win since p guarantees to play a stronger action. But, p' gives positive probability to $w_{p'}$, so no other action available to p' can have higher expected payoff. p' has non-positive expected payoff, and under the assumption (that we may adopt from preprocessing) that players all have a 0-cost action, p''s payoff must in fact be zero.

The second part of the theorem characterizes p in terms of the game (without reference to any Nash equilibria) and can be seen for the following reason. Let p'' be the player with the strongest action having a cost of less than 1. Then p'' can guarantee a positive payoff by using that action. Hence, p'' cannot be one of the players who have expected payoff 0 in a Nash equilibrium, so $p = p''$. $\square$

The following theorem also shows how a Nash equilibrium may be efficiently computed for tie-free games, provided that we know the support of the Nash equilibrium. This shows an interesting parallel between these games, and general 2-player normal form games, especially in conjunction with the subsequent observation that the solution is in rational numbers.

THEOREM 2. *For games with any number of players, actions' costs less than 1 and a single prize where ties are impossible, a Nash equilibrium can be computed in polynomial-time if we are given the support of a solution.*

PROOF. Given a game $\mathcal{G}$, suppose that we remove the pure strategies that are not in the support of some (unknown) Nash equilibrium. The resulting game $\mathcal{G}'$ has a fully-mixed equilibrium $\mathcal{N}$, thus any two strategies that belong to a player have the same expected payoff in $\mathcal{N}$. Our general approach is to compute the probabilities x^i_j in descending order of strength of the associated actions a^i_j.

Let a^i_j be the strongest action in $\mathcal{G}'$ (i.e. having the highest return). Player i's expected payoff is $1-c^i_j$ and by Theorem 1 all other players have expected payoff 0.

Let $a^{i'}_{j'}$ be the second-strongest action in $\mathcal{G}'$; we may assume $i' \neq i$ since if $i' = i$ then a^i_j would be strictly dominated by $a^{i'}_{j'}$. Its expected payoff to i' is $-c^{i'}_{j'}+(1-x^i_j)$, which by Theorem 1 is 0, so we have an expression for x^i_j. Consider the third-strongest action $a^{i''}_{j''}$, whose payoff is given by $-c^{i''}_{j''} + (1 - x^i_j)(1 - x^{i'}_{j'})$ (assuming $i'' \neq i$) which gives us an expression for $x^{i'}_{j'}$.

Generally, the r-th strongest action a^α_β has expected payoff $-c^\alpha_\beta + (1 - \sigma_{i_1}) \cdots (1 - \sigma_{i_r})$ where σ_{i_j} (for player i_j) is the sum of probabilities $x^k_{i_j}$ for actions stronger than the r-th strongest action.

The probabilities for each player's weakest actions will be obtained from the equations that ensure that for every player i, the values x^i_j sum to 1 (are a probability distribution). $\square$

OBSERVATION 3. *For games where ties are impossible, if all action costs are rational numbers then the solution is also in rational numbers.*

This is immediate from the expressions in the above proof that give the values x^i_j.

3.2.1 *Solving 2-player games exactly*

Ranking games (as in [2]) with actions that do not have the upfront costs c^i_j we consider here, are zero-sum, so they can be solved efficiently in the 2-player case. Our games are not zero-sum, but we do have an alternative polynomial-time algorithm to solve them in the 2-player case.

THEOREM 3. *2-player ranking games that have competitiveness-based strategies and are without ties can be solved exactly in polynomial time.*

PROOF. As before, assume a single prize of 1 unit and action costs in $[0, 1]$, which may be assumed by the preprocessing noted earlier.

We can (in polynomial time) compute exact solutions of 2-player games of this type as follows. We start by eliminating certain dominated strategies. Specifically, suppose that for strategies a^i_j and a^i_{j+1}, the set of opponent's strategies that they win against, is the same. Then a^i_{j+1} can be eliminated. Rename the strategies of this game $(a_1, \ldots a_n)$

for the row player and $(a'_1, \ldots a'_{n'})$ for the column player. Assume without loss of generality that it is the row player who has the weakest strategy, thus a'_1 wins against a_1. In general when $n = n'$, the strategies, arranged in ascending order of strength are $a_1, a'_1, a_2, a'_2, \ldots a_n, a'_n$.

Suppose that in some Nash equilibrium $\mathcal{N}$ the row player does not use strategy a_j for some $j > 1$ (that is, the player plays a_j with probability 0). Then the column player does not use strategy a'_j (which is the cheapest one that wins against a_j) since a'_j would now be dominated by a'_{j-1}. For a similar reason, the row player will not use a_{j+1}, the cheapest strategy that wins against a'_j, so the column player will not use a'_{j+1}, and so on. This shows that (in a Nash equilibrium) the strategies in either player's support must be either a prefix of the sequence of his strategies or a prefix of all his strategies but the weakest, with strategies arranged in ascending order of strength.

We can now try to solve for all such supports, since there are polynomially-many of them. Recall that a 2-player game can be solved efficiently in polynomial time if we are told the support of a solution, since it reduces to a linear program. □

The main property used to show Theorem 3 above (i.e., if in any solution a player does not play a certain strategy s then the other one does not play the strategy that "just beats" s) breaks down when ties are allowed. Simply consider two consecutive strategies of player 1, a^1_i and a^1_{i+1}, and of player 2, a^2_j and a^2_{j+1} such that returns of a^1_i and a^2_j (resp. a^1_{i+1} and a^2_{j+1}) are the same. In this case it is false that in any solution if player 1 does not play a^1_i then player 2 does not play a^2_{j+1} as such a strategy is used not just to beat a^1_i but also to share with a^1_{i+1}. This shows that for general games the supports of a Nash equilibrium can be anything and we have to use different approaches to obtain polynomial-time algorithms for them.

3.3 Exact algorithm for linear-prize ranking games

Consider a d-player n-strategies-per-player ranking game $\mathcal{G}$ with competitiveness-based strategies without ties in which the prize for ranking j-th is a linear function $a - jb$, for some value of a and b. We call $\mathcal{G}$ a *linear-prize ranking game.* We claim that we can represent $\mathcal{G}$ as a polymatrix game [10]. A polymatrix game can be represented as a graph: players are the vertices and a player's payoff depends on the actions of his neighbors. The edges are 2-player zero-sum games and once all players have chosen a strategy, the payoff of each player is the sum of the payoffs from each game played with each neighbor.

We can express $\mathcal{G}$ as a polymatrix game as follows. We define a complete $(d+1)$-vertex graph where the additional vertex encodes an external player N, the nature. Edge (i, j), for i and j players in $\mathcal{G}$, is an n by n 2-player zero-sum game (shifted by a constant) punishing who is ranking worst between i and j. So, the entry (k, l) will have payoff 0 for player i and $-b$ for player j if and only if $r^i_k > r^j_l$. Edge (i, N) is an n by 1 2-player zero-sum game in which the nature "gains" what player i is paying in effort minus $a - b$. So the entry k will have payoff $a - b - c^i_k$ for player i and payoff $c^i_k - a + b$ for the nature. Note that once all players of $\mathcal{G}$ choose a strategy, the sum of all the payoffs to a player of these games is the prize he is awarded (from the games with other players of $\mathcal{G}$ we subtract b for any player who beats him and add $a - b$ from the game with the nature) minus the cost for the strategy he chooses (from the game with the nature). The following thus immediately follows from [10].

THEOREM 4. *There is a polynomial-time algorithm that computes a Nash equilibrium for linear-prize ranking games.*

3.4 Games where ties are possible

In this section we consider general competitiveness-based ranking games. We begin by showing that we can study without loss of generality Nash equilibria of competitiveness-based ranking games in which returns are symmetric. That is, all players have n actions a_1,...,a_n. As above player i has costs $c^i_1 < \ldots < c^i_n$ for those actions but, differently from above, returns are player-independent and denoted as $r_1 < \ldots < r_n$.

3.4.1 Reduction to return-symmetric games

The reduction preserves Nash equilibria of the original game (and thus is a Nash homomorphism) and is presented for the case of 2-player games. The argument generalises to more players, although the number of strategies per player would increase by a factor of d, the number of players.

Suppose that player 1 has action a^1_j and player 2 has no action with return equal to r^1_j. We give player 2 a dominated strategy with return r^1_j — if a^2_k is the weakest strategy of player 2 that has higher return than a^1_j, give player 2 an additional strategy with cost c^2_k and return r^1_j. If player 2 does not have a stronger strategy than a^1_j, give player 2 an additional strategy with cost 1 and return r^1_j.

Hence we can assume that each player has n strategies $a_1, \ldots, a_n$ where for each player i, a_j has cost c^i_j that depends on i (and non-decreasing as a function of j), while the return of a_j may be set to r_j, thus a_j is stronger than a_k for $j > k$. Suppose we solve this game, and now we have to recover a solution to the original game before the dominated strategies were added. To do this, each player just has to replace their usage of any dominated strategy by the corresponding dominating one — this raises the question of whether the other player may be given an incentive to deviate as a consequence, however the reason why this does not happen, is that if it did, it would mean that the dominated strategy being used by the first player, was in fact strictly dominated — the opponents' behaviour gives him a positive incentive to switch.

3.4.2 Return-symmetric games and pure equilibria

Unlike games without ties, for which 2-player 2-strategy games might not possess pure equilibria (see Example 1), return-symmetric games in which players have only 2 strategies do have pure Nash equilibria (for any number of players and any number of prizes).

THEOREM 5. *2-strategy competitiveness-based return-symmetric ranking games do have pure Nash equilibria (any number of players; any action costs for individual players). Furthermore, a pure Nash equilibrium can be found in polynomial time.*

PROOF. Assume for contradiction that a 2-strategy return-symmetric competitiveness-based ranking game does not have a pure Nash Equilibrium. This implies that in the Nash dynamics graph any node has an outgoing edge and then the

Nash dynamics graph has a cycle. Let $s_1 \to \ldots \to s_k = s_1$ be such a cycle. Each edge (s_j, s_{j+1}) encodes the fact that some player, denoted by i_j, is better off in s_{j+1} than in s_j, i.e., $u^{i_j}(s_j) - c^{i_j}(s_j) < u^{i_j}(s_{j+1}) - c^{i_j}(s_{j+1})$ where $u^i(s)$ (resp. $c^i(s)$) is the prize given to (resp. the cost experienced by) player i for outcome s.

Let L and H be the low-competitive and high-competitive strategies, respectively. Consider the point in the cycle where the number of players using H reaches its highest value. The player who moves to attain that point in the cycle, will never prefer to return to strategy L, since at other points in the cycle, the payoff for H can only increase (the prize is shared between fewer players using H).

The above indicates how to quickly find a pure Nash equilibrium — the players using H should have costs for playing H that are lower than those of the players using L. (Recall that the costs of playing L can be assumed to be 0 for all players.) So there are just $n-1$ pure-strategy profiles that actually need to be checked. □

A related result is known for symmetric games in which players have only 2 strategies: These games always have a pure Nash equilibrium [6]. Above theorem concerns games that are anonymous but not symmetric (as costs are player-specific). However, let us notice that the arguments used in [3] to prove the result about symmetric games appear to be similar to ours. It is easy to see that these arguments fail when two players have three strategies available, as shown by the next example.

Example 5. *We have 2 players and 3 actions, namely a_1, a_2 and a_3 ordered increasingly by return, i.e., $r_1 < r_2 < r_3$. The prizes are $u_1 = 1$ and $u_2 = 0$, $u_3 = 0$. Costs are $c_1^i = 0$ for $i = 1, 2$, $c_2^1 = \frac{2}{3}, c_3^1 = \frac{4}{5}$ and $c_2^2 = \frac{1}{3}, c_3^2 = \frac{2}{3}$. Thus, we have payoff matrix*

	a_1	a_2	a_3
a_1	$(\frac{1}{2}, \frac{1}{2})$	$(0, \frac{2}{3})$	$(0, \frac{1}{3})$
a_2	$(\frac{1}{3}, 0)$	$(-\frac{1}{6}, \frac{1}{6})$	$(-\frac{2}{3}, \frac{1}{3})$
a_3	$(\frac{1}{5}, 0)$	$(\frac{1}{5}, -\frac{1}{3})$	$(-\frac{3}{10}, -\frac{1}{6})$

It is easily checked that this game has no pure Nash equilibrium. The unique Nash equilibrium of the game is $(\frac{2}{3}, 0, \frac{1}{3})$ for player 1 and $(\frac{2}{5}, \frac{3}{5}, 0)$ for player 2.

3.4.3 PTAS for constant number of players; variable size n of pure-strategy sets

Let d be the number of players and prizes; by the preprocessing step of Section 3.1 we may assume that the first prize is worth 1 and the bottom prize is worth 0; also that all costs are in the range $[0, 1]$.

Theorem 6. *For any $\epsilon > 0$, if the number of players d is a constant then Algorithm 1 returns a 3ϵ-Nash equilibrium in time polynomial in the input size.*

Proof. Consider the case in which $\epsilon = 1/k$ for some integer k. Observe that after Step 2 each player has some number of strategies with the same cost but different returns. Thus, the cost rounding introduces a number of dominated strategies and after Step 3 each player has only a constant number $k+1$ of actions available. This and the fact that costs lie in the interval $[0, 1]$ imply that the number of strategies considered in Step 5 is $(k+1)^{\frac{1}{\delta}}$ per player. This is a constant and since d is constant we can perform the brute

Algorithm 1: PTAS for constant number of players

1 Let k be an integer constant, let $\epsilon = \frac{1}{k}$, and let $\delta = \epsilon/(k+1)$.
2 For each player, round each cost c_j^i down to the nearest non-negative integer multiple of $1/k$.
3 Eliminate dominated strategies.
4 **for** *all players i* **do**
5   Consider every strategy $\{x_1^i, \ldots, x_{k+1}^i\}$ in which each x_j^i is a non-negative integer multiple of δ.
6 Do a brute-force search over all the possible choice of strategies as above by all the d players.
7 Return an ϵ-Nash equilibrium of the reduced game.

force search in Step 6 efficiently in time $(k+1)^{\frac{d}{\delta}}$. Notice furthermore that in Step 7 an ϵ-Nash equilibrium is always found. To see this, consider a δ-rounding (as from Observation 1) of each probability vector of a Nash equilibrium of the game and notice that such a probability vector is checked by the algorithm. How badly a player can do by restricting probabilities to be non-negative integer multiples of δ? For each of his $k+1$ actions, the probability with which he chooses the action may be "wrong" by an additive δ. If he is wrong, he loses a payoff which is upper bounded by 1 and so his overall regret is at most $(k+1)\delta = \epsilon$. We now show that the ϵ-Nash equilibrium $\mathcal{N}$ of the reduced game is an $\epsilon + 2/k$-Nash equilibrium of the original game. To see this consider a player, say i, that in his best response assigns positive probability to an action x which is not in the reduced game. We know that the algorithm considers the smallest non-negative integer multiple of ϵ larger than x. Call it b and let π_z denote the expected utility of player i when playing action $z \in S_i$. By construction we know that $\pi_x \leq \pi_b + 1/k$. (Indeed, by playing x player i can get at most the same expected prize that he gets in b but has to pay $1/k$ less.) Then if b is in the support of $\mathcal{N}$ then the maximum regret that i has by not playing x is $1/k$. On the other hand, consider the case in which b is not in the support of $\mathcal{N}$. Since $\mathcal{N}$ is an ϵ-Nash equilibrium then by (2) we know that $\pi_b < \pi_c + \epsilon$ for an action c in the support of $\mathcal{N}$, thus implying that the maximum regret of i for not playing x is $2/k = 2\epsilon$.

When ϵ is not the inverse of an integer, simply run the algorithm with an $\epsilon' < \epsilon$ which is the inverse of an integer. Since an ϵ'-Nash equilibrium is an ϵ-Nash equilibrium the theorem follows. □

3.4.4 PTAS for many players who share a fixed set of strategies

Assume we have a return-symmetric d-player game, each of whom have strategies $a_1, \ldots, a_n$, where n is a constant. This is an anonymous game, and we can directly apply a result of Daskalakis and Papadimitriou [9] that it has a PTAS.

For the more specific class of games that we are concerned with here, we can describe a conceptually simpler PTAS as follows. In a return-symmetric game actions have a player-independent return and thus players differentiate by their cost vectors. We call the *type* of a player his cost vector. Thus in a return-symmetric game, if two players have the same type they have the same cost vector and same action returns.

Algorithm 2: PTAS for return-symmetric games with constant number of actions

1 Let $\epsilon = 1/k$ for some integer constant k, and let $\delta = \epsilon/n$ and $l = 1/\delta$.
2 For each player, round each cost c^i_j down to the nearest non-negative integer multiple of ϵ.
3 **for** *all players* i **do**
4 Consider every strategy $\{x^i_1, \ldots, x^i_{l+1}\}$ in which each x^i_j is a non-negative integer multiple of δ.
5 Do a brute-force search over all the possible choice of strategies as above by all the n^{l+1} player types.
6 Return an ϵ-Nash equilibrium of the reduced game.

THEOREM 7. *For any $\epsilon > 0$, if the number of actions n is a constant then Algorithm 2 returns a 2ϵ-Nash equilibrium in time polynomial in the input size.*

PROOF. Consider first the case in which ϵ is the inverse of an integer. Since costs lie in the range $[0, 1]$ and as n is constant, after Step 2 is executed there is a constant number n^{l+1} of distinct types. This allows the brute force search in Step 5 to be done efficiently in time $(l + 1)^{n^{2(l+1)}}$. Notice that the algorithm always finds an ϵ-Nash equilibrium in the last step. Consider a δ-rounding (as from Observation 1) of each probability vector of a Nash equilibrium of the game and notice that such a probability vector is checked by the algorithm. We next show that such a probability vector is an ϵ-Nash equilibrium. A player can be playing a strategy with a probability that is δ away from the probability he should have used for his best response. Thus for each of his strategies he can lose at most δ times the payoff he gets for that strategy. Since payoffs are upper bounded by 1, the maximum regret is $n\delta = \epsilon$. The proof concludes by noting that the cost rounding of Step 2 implies an extra regret of at most ϵ.

Whenever ϵ is not the inverse of an integer, then we feed the algorithm with an $\epsilon' < \epsilon$ which is the inverse of an integer. The theorem thus follows. □

Note that this algorithm is oblivious in the sense of Daskalakis and Papadimitriou [11].

3.4.5 A fully polynomial-time approximation scheme for constant number of players and single prize

We will present this for the 2-player case for simplicity. The arguments we use easily extend to the case of constantly-many players.

Let $\mathcal{G}$ be a return-symmetric game with 2 players. In this case, the expected payoff of player i from playing strategy j is given by the expected prize he gets minus the cost c^i_j to play action a_j. As for the expected prize, player 1 wins (and gets the prize of 1) as long as player 2 plays a strategy weaker than a_j and shares (and thus gets a prize of $\frac{1}{2}$) when player 2 plays action a_j. (Similar formula applies to player 2's payoffs.) More formally, let x^i_j be the probability that player i plays strategy j and let i' be the player different from i in the game. Then the expected payoff of player i for playing j, denoted as π^i_j, is given by

$$\pi^i_j = \sum_{k=1}^{j-1} x^{i'}_k + \frac{1}{2} x^{i'}_j - c^i_j. \tag{3}$$

To compute a Nash equilibrium, we need to find real values $x^1_1, \ldots, x^1_n, x^2_1, \ldots, x^2_n$ that satisfy

$$x^i_j \geq 0 \ \forall i, j \qquad \sum_j x^i_j = 1 \ i = 1, 2 \tag{4}$$

saying that for $i = 1, 2$, the values $\{x^i_j\}_j$ are a probability distribution; for $i = 1, 2$ and $j > 1$ the following are also true

$$\begin{array}{lcl} \pi^i_j > \max_{k=1,\ldots,j-1}\{\pi^i_k\} & \Longrightarrow & x^i_1 = \ldots = x^i_{j-1} = 0, \\ \pi^i_j < \max_{k=1,\ldots,j-1}\{\pi^i_k\} & \Longrightarrow & x^i_j = 0. \end{array} \tag{5}$$

LEMMA 1. *The values x^i_j satisfy (4,5) if and only if they are a Nash equilibrium.*

PROOF. The sets $\{x^1_j\}_j$ and $\{x^2_j\}_j$ are constrained by (4) to be probability distributions.

The expressions (5) are equivalent to the definition of Nash equilibrium constraints (1). Indeed, if action a_j dominates all previous weaker actions then any of these actions must not be in the support. Similarly, when a_j is dominated by a weaker action, a_j will not be in the support. Note that if a_j dominates the weaker actions but is dominated by a stronger action a_q then the probability of playing a_j will be set to 0 when π^i_q is compared to $\max_{k=1,\ldots,j,\ldots,q-1}\{\pi^i_k\}$. □

Consequently we have reduced the problem to satisfying the constraints (4,5). We now define variables in addition to x's and π's with the aim to make the constraints (4,5) depend on a constant number of "local" variables. This is needed to define our FPTAS. Let σ^i_j be the partial sum $\sum_{\ell=1}^{j} x^i_\ell$. We can now express Equation (4) as follows:

$$\begin{array}{ccc} \sigma^i_1 = x^i_1 & 0 \leq \sigma^i_j \leq 1 & \sigma^i_{j-1} + x^i_j = \sigma^i_j \\ & \sigma^i_n = 1 \qquad 0 \leq x^i_j \leq 1, & \end{array} \tag{6}$$

and rewrite (3) as

$$\pi^i_1 = \frac{1}{2} x^{i'}_1 - c^i_1 \qquad \pi^i_j = \sigma^{i'}_{j-1} + \frac{1}{2} x^{i'}_j - c^i_j. \tag{7}$$

Additionally, let α^i_j be the maximum expected payoff player i can get by playing one of the first j actions, i.e., $\alpha^i_j = \max_{k=1,\ldots,j}\{\pi^i_k\}$. We can now define

$$\alpha^i_1 = \pi^i_1 \qquad \alpha^i_j = \max\{\alpha^i_{j-1}, \pi^i_j\} \tag{8}$$

and express Equation (5) as follows:

$$\begin{array}{lcl} \pi^i_j > \alpha^i_{j-1} & \Longrightarrow & \sigma^i_{j-1} = 0, \\ \pi^i_j < \alpha^i_{j-1} & \Longrightarrow & x^i_j = 0. \end{array} \tag{9}$$

OBSERVATION 4. *The values x^i_j, σ^i_j, α^i_j and π^i_j satisfy (6, 7,8,9) if and only if x^i_j are a Nash equilibrium.*

Now consider the sequence

$$\mathcal{S} = (\pi^1_j, \pi^2_j, x^1_j, x^2_j, \alpha^1_j, \alpha^2_j, \sigma^1_j, \sigma^2_j)_{j=1,\ldots,n}.$$

Constraints in (6) involve 3 variables that are at distance at most 9 in $\mathcal{S}$ (namely, for $j > 1$, σ^i_{j-1} is followed by 8 elements of $\mathcal{S}$ –including x^i_j– and then by σ^i_j). It is easy to check that the same happens also for the other constraints and conclude then that the following holds.

OBSERVATION 5. *For any $j = 1, \ldots, n$ all constraints in (6,7,8,9) apply to at most 9 consecutive elements of $\mathcal{S}$.*

The algorithm. For $\epsilon > 0$ according to (2) we relax the constraints of (9) as follows:

$$\begin{array}{lcl} \pi_j^i > \alpha_{j-1}^i + \epsilon & \Longrightarrow & \sigma_{j-1}^i = 0, \\ \pi_j^i < \alpha_{j-1}^i - \epsilon & \Longrightarrow & x_j^i = 0. \end{array} \qquad (10)$$

Let $\mathcal{S}_i$ be the sequence of 9 consecutive elements of $\mathcal{S}$ that begins at the i-th element of $\mathcal{S}$. Let $\mathcal{E}_i$ be the set of expressions in (6,7,8,10) that relate elements of $\mathcal{S}_i$ with each other; by Observation 5 the union of the sets $\mathcal{E}_i$ is all constraints (6,7,8,10).

The algorithm (see Algorithm 3) works its way through the sequence $\mathcal{S}$ left-to-right, for each $\mathcal{S}_i$ identifies a subset of $([0,1])^9$ representing possible values of those quantities that form part of an approximate Nash equilibrium. Then it sweeps through the sequence right-to-left identifying allowable values for previous elements. The parameter ϵ controls quality of approximation; the algorithm computes values for terms in $\mathcal{S}$ that are integer multiples of $\epsilon/2$.

Algorithm 3: FPTAS for return-symmetric games with 2 players

---

1 Let $\epsilon = 1/k$ for some integer constant k.

2 For each player, round each cost c_j^i down to the nearest non-negative integer multiple of ϵ.

3 For $1 \leq i \leq 8n - 8$, let D_i be the set of all 9-dimensional vectors of non-negative integer multiples of $\epsilon/2$ for x's and σ's values and of integer multiples of $\epsilon/2$ for π's and α's values that satisfy $\mathcal{E}_i$.

4 For $i > 1$ (in ascending order) discard from D_i any vector whose first 8 entries are different from the last 8 entries of all vectors in D_{i-1}.

5 Let $\mathbf{s}_{8n-8}$ be a point in D_{8n-8}. For $1 \leq i < 8n - 8$ (in descending order) let $\mathbf{s}_i$ be a point in D_i chosen so that its last 8 coordinates are the first 8 coordinates of $\mathbf{s}_{i+1}$.

6 Let $\mathbf{s}$ be the vector of length $8n$ such that $\mathbf{s}_i$ is the i-th sequence of 9 consecutive coordinates of $\mathbf{s}$. Set x_j^i to the entry of $\mathbf{s}$ that corresponds to the position of x_j^i in $\mathcal{S}$.

Theorem 8. *There is a FPTAS for competitiveness-based ranking games with a constant number of players and a single prize.*

Next we show above theorem in relation to Algorithm 3 and the case of 2 players by providing approximation guarantee, correctness and runtime analysis of the algorithm in case ϵ is the inverse of an integer. (Similarly to above, if this is not the case we simply run the algorithm with an $\epsilon' < \epsilon$ which is inverse of an integer.)

Proposition 1 (Approximation guarantee). *Algorithm 3 computes values x_j^i that correspond to an $\frac{n+4}{2}\epsilon$-Nash equilibrium.*

Proof. The x_j^i satisfy (6,7,8,10), where (10) simply rewrites the definition of ϵ-Nash equilibrium (2) thus implying that we are losing an additive ϵ. Another additive loss of ϵ is due to the cost rounding. Furthermore, we are restricting to probability distributions whose values are non-negative integer multiples of $\epsilon/2$. Thus, a player may be forced to play a strategy with a probability that is far at most $\epsilon/2$ from the probability of his best response. This will impose an additional loss of $n\epsilon/2$ in the worst case (this is because we have n actions and on each of them the best response is at most $\epsilon/2$ away while the payoffs are upper bounded by 1). □

Proposition 2 (Correctness). *The algorithm always finds such an* $\mathbf{s}$.

Proof. Consider a Nash equilibrium $\mathcal{N}$ and associated vector $\mathbf{s}$. Take an $\epsilon/2$-rounding of the probabilities vectors x's in $\mathbf{s}$ and round the remaining elements of $\mathbf{s}$ to the nearest non-negative integer multiple of $\epsilon/2$. Call this new sequence $\tilde{\mathbf{s}}$ and observe that $\tilde{\mathbf{s}}$ is considered by the algorithm. It is thus enough to show that such a sequence satisfies all the constraints that the algorithm imposes on the output. Towards this end, we let $\pi_j^i, x_j^i, \alpha_j^i$ and σ_j^i the values of $\mathbf{s}$ and let $\tilde{\pi}_j^i, \tilde{x}_j^i, \tilde{\alpha}_j^i$ and $\tilde{\sigma}_j^i$ the corresponding rounded values of $\tilde{\mathbf{s}}$.

By Observation 1 we have that the values $\tilde{x}$'s and $\tilde{\sigma}$'s of $\tilde{\mathbf{s}}$ satisfy (6). Moreover, since $\tilde{\sigma}$'s and $\tilde{x}$'s are integer multiples of $\epsilon/2$ and as costs are rounded to integer multiples of ϵ we have that that $\tilde{\pi}$'s will be integer multiples of $\epsilon/2$. This immediately implies that $\tilde{\mathbf{s}}$ satisfies constraints (7) and (8) as well. As for constraint (10) we show that $\tilde{\pi}_j^i > \tilde{\alpha}_{j-1}^i + \epsilon \Rightarrow \tilde{\sigma}_{j-1}^i = 0$. (Very same arguments can be used to show the other condition of (10).) Note that because of the rounding we have $|y_j^i - \tilde{y}_j^i| < \epsilon/2$ for $y \in \{\pi, \alpha\}$. Therefore $\tilde{\pi}_j^i > \tilde{\alpha}_j^i + \epsilon$ implies that $\pi_j^i > \alpha_j^i$ and as $\mathbf{s}$ is a Nash equilibrium, by Observation 4 and Equation (9), we have $\sigma_{j-1}^i = 0$. But then by the way we define the rounding of σ's we have that $\tilde{\sigma}_{j-1}^i = 0$. □

Runtime: The sets D_i are of size $O((4/\epsilon)^9)$, so the runtime of the algorithm is indeed polynomial in n and $1/\epsilon$, as required for a FPTAS.

4. CONCLUSIONS AND FURTHER WORK

The FPTAS we provide above, is very analogous to the algorithm of [15] for solving tree-structured graphical games. They give a similar forward-and-backward dynamic programming approach to solving these games; their algorithm takes exponential time for exact equilibria [12] but a similar quantization of real-valued quantities leads to a FPTAS.

Our FPTAS can be used to compute *exact* equilibria in certain cases. When a game with constantly-many players and a single prize has payoffs that are multiple of some $\epsilon > 0$ then we can compute exact Nash equilibria in time polynomial in the size of the input and $1/\epsilon$ by simply using the FPTAS. This observation raises the open problem of determining if there is a polynomial-time algorithm to solve 2-player (competitiveness-based ranking) games in general when ties are possible and the prize is shared in the event of a tie.

A number of other concrete open problems have been raised by the current results, as for example, a complete understanding of the complexity of computing Nash equilibria for competitiveness-based ranking games. More generally, in situations where multiple equilibria may exist, we would like to know whether a specific equilibrium is selected by some natural decentralised dynamic process.

5. REFERENCES

[1] M. Blonski. "Anonymous Games with Binary Actions", *Games and Economic Behavior* **28**(2), pp. 171-80, 1999.

[2] F. Brandt, F. Fischer, P. Harrenstein and Y. Shoham. "Ranking Games", *Artificial Intelligence* **173**(3), pp. 221-239, 2009.

[3] F. Brandt, F. Fischer and M. Holzer. "Symmetries and the complexity of pure Nash equilibrium", *Journal of Computer and System Sciences* **75**(3), pp. 163-177, 2009.

[4] I. Brochas and J.D. Carrillo. Do the "three-point victory" and "golden goal" rules make soccer more exciting? *Journal of Sports Economics*, **5**, pp. 169-185, 2004.

[5] X. Chen, X. Deng and S-H Teng, "Settling the Complexity of Computing Two-Player Nash Equilibria", *Journal of the ACM* **56**(3), 2009.

[6] S.-F. Cheng, D.M. Reeves, Y. Vorobeychik and M.P. Wellman. "Notes on equilibria in symmetric games", *Proc. 6th International Workshop on Game Theoretic and Decision Theoretic Agents (GTDT)*, pp. 23-29, 2004.

[7] C. Daskalakis. "An Efficient PTAS for Two-Strategy Anonymous Games", *Proc. 4th Workshop on Internet & Network Economics (WINE)*, pp. 186-197, 2008.

[8] C. Daskalakis, P.W. Goldberg and C.H. Papadimitriou. "The Complexity of Computing a Nash Equilibrium." *SIAM Journal on Computing* **39**(1), pp. 195-259, 2009.

[9] C. Daskalakis and C.H. Papadimitriou. "Discretized Multinomial Distributions and Nash Equilibria in Anonymous Games", *Proc. of 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pp. 25-34, 2008.

[10] C. Daskalakis and C.H. Papadimitriou. "On a Network Generalization of the Minmax Theorem", *Proc. of 36th International Colloquium on Automata, Languages and Programming (ICALP)*, pp. 423-434, 2009.

[11] C. Daskalakis and C.H. Papadimitriou. "On Oblivious PTAS's for Nash Equilibrium", *Proc. of 41st ACM Symposium On Theory of Computing (STOC)*, pp. 75-84, 2009.

[12] E. Elkind, L.A. Goldberg and P.W. Goldberg. "Nash Equilibria in Graphical Games on Trees Revisited", *Proc. 7th ACM Conference on Electronic Commerce (ACM EC)*, pp. 100-109, 2006.

[13] K.K. Haugen. "Point Score Systems and Competitive Imbalance in Professional Soccer", *Journal of Sports Economics* **9**(2), pp. 191-210, 2008.

[14] E. Hopkins and T. Kornienko. "Running to Keep in the Same Place: Consumer Choices as a Game of Status" *American Economic Review* **94**(4), pp. 1085-1107, 2004.

[15] M. Kearns, M. Littman and S. Singh. "Graphical Models for Game Theory" *Proc. of the 17th Conference in Uncertainty in Artificial Intelligence (UAI)*, pp. 253-260, 2001.

[16] J. Nash. "Noncooperative Games," *Annals of Mathematics* **54**, pp. 289-295, 1951.

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