

Approximability and Parameterized Complexity of Minmax Values^{*}

Kristoffer Arnsfelt Hansen^{**}, Thomas Dueholm Hansen, Peter Bro Miltersen,
and Troels Bjerre Sørensen

Department of Computer Science, University of Aarhus, Denmark
{arnsfelt,tdh,bromille,trold}@cs.au.dk

Abstract. We consider approximating the minmax value of a multi-player game in strategic form. Tightening recent bounds by Borgs *et al.*, we observe that approximating the value with a precision of $\epsilon \log n$ digits (for any constant $\epsilon > 0$) is **NP**-hard, where n is the size of the game. On the other hand, approximating the value with a precision of $c \log \log n$ digits (for any constant $c \geq 1$) can be done in quasi-polynomial time. We consider the parameterized complexity of the problem, with the parameter being the number of pure strategies k of the player for which the minmax value is computed. We show that if there are three players, $k = 2$ and there are only two possible rational payoffs, the minmax value is a rational number and can be computed *exactly* in linear time. In the general case, we show that the value can be approximated with any polynomial number of digits of accuracy in time $n^{O(k)}$. On the other hand, we show that minmax value approximation is **W[1]**-hard and hence not likely to be fixed parameter tractable. Concretely, we show that if **k-CLIQUE** requires time $n^{\Omega(k)}$ then so does minmax value computation.

1 Introduction

A game G in *strategic form* between l players is given by a set of players $\{1, \dots, l\}$ and for each player j a finite *strategy space* S_j and a *utility function* $u_j : S_1 \times S_2 \times \dots \times S_l \rightarrow \mathbf{R}$. In this paper, only the utility function for Player 1 is relevant. When the size of S_j is n_j , we shall refer to the game as an $n_1 \times n_2 \times \dots \times n_l$ game. The *minmax* (or *threat*) *value* of G for Player 1 is given by $\min_{\sigma_{-1} \in \Delta^{(l-1)}} \max_{a \in S_1} E[u_1(a, \sigma_{-1})]$ where $\Delta^{(l-1)}$ is the set of mixed, but uncorrelated, strategy profiles for players $2, \dots, l$. A profile σ_{-1} achieving the minimum in the expression is called an *optimal minmax profile* or an *optimal threat*. The *maxmin* (or *security*) *value* of G for Player 1 is given by $\max_{\sigma_1 \in \Delta} \min_{a_2, \dots, a_l} E[u_1(\sigma_1, a_2, \dots, a_l)]$ where Δ is the set of mixed strategies for Player 1.

^{*} Work supported by *Center for Algorithmic Game Theory*, funded by the Carlsberg Foundation.

^{**} Supported by a postdoc fellowship from the Carlsberg Foundation.

The minmax value of a finite two-player game is a fundamental notion of game theory. Its mathematical and computational properties are extremely well-studied and well-understood, being intimately tied to the theory of linear programming. In particular, the duality theorem of linear programs implies that the minmax value equals the maxmin value. Also, the computation of the minmax value of a two-player game in strategic form is essentially equivalent to solving linear programs and can therefore be done in polynomial time (although a strongly polynomial time algorithm remains an open problem).

Minmax values of *multi*-player games are much less well-studied, although these values are arguably also of fundamental interest to game theory. Most importantly, the minmax value plays a pivotal role in the statement and proof of the so-called *folk theorems* that characterize the Nash equilibria of infinitely repeated games. Additionally, the minmax value is the equilibrium payoff of the so-called *team-maxmin equilibria* studied by von Stengel and Koller [15]. For a multi-player game, the maxmin value may be strictly smaller than the minmax value. Computation of the maxmin value easily reduces to the two-player case and can therefore be done efficiently using linear programming. Rather surprisingly, computation of the minmax value of a multi-player game in strategic form was not studied until very recently, where Borgs *et al.* [1] (motivated by computational aspects of the folk theorem) showed that approximating the minmax value of a three-player game within a certain inverse polynomial additive error is **NP**-hard. Our starting point is this important paper.

Given the fundamental nature of the notion of the minmax value, it is important to understand when the **NP**-hardness result can be circumvented by considering special cases or asking for weaker approximations. The purpose of this paper is to provide a number of results along these lines. First, we observe that the inapproximability result of Borgs *et al.* can be tightened and matched with a positive result, using standard techniques.

Theorem 1. *For any constant $\epsilon > 0$, approximating the minmax value of an $n \times n \times n$ game with 0-1 payoffs within additive error $1/n^\epsilon$ is **NP**-hard. On the other hand there is an algorithm that, given a parameter $\epsilon > 0$ and a game in strategic form with l players each having n strategies and all payoffs being between 0 and 1, approximates the minmax value for Player 1 from above with additive error at most ϵ in time $n^{O(l(\log n)/\epsilon^2)}$.*

This suggests the following important problem: *Can the minmax value of a three-player game with payoffs normalized to $[0, 1]$ be approximated within a non-trivial additive constant (say 0.001 or even 0.499) in polynomial (rather than quasi-polynomial) time?* We leave this problem open.

It is of interest to know when the minmax value can be computed exactly. A prerequisite for this is that it is rational. For three-player games, we characterize when the minmax value for Player 1 can be an irrational number, in terms of the number of strategies of Player 1 and the number of distinct (rational) payoffs.

For the special case where the value is guaranteed to be rational we present an optimal linear time algorithm for exactly computing the minmax value¹.

Theorem 2. *Consider $k \times n \times n$ three-player games with only l distinct rational payoffs. When either $k \geq 2$ and $l \geq 3$ or $k \geq 3$ and $l \geq 2$ there exists a game such that the minmax value for Player 1 is irrational. Otherwise, when $k = 2$ and $l = 2$ the minmax value for Player 1 is a rational number and we can compute it exactly in time $O(n^2)$ (on a unit cost random access machine).*

Thus having observed that the case of few strategies of Player 1 may be easier than the general case, we apply the approach of *parameterized complexity* [8], considering the number of strategies k of Player 1 as the parameter. Combining a classical result of Shapley and Snow [14] with Renegar’s decision procedure for the first order theory of the reals [11,12,13] gives rise to a support enumeration algorithm for finding the minmax value and we show the following.

Theorem 3. *Given a $k \times n \times \dots \times n$ l -player game G with rational payoffs and a rational number α so that (G, α) has combined bit complexity L , we can decide in time $L^{O(1)} k^{O(kl)} n^{kl}$ (on a Turing machine) if the minmax value of G for Player 1 is at most α . Using the terminology of fixed parameter complexity theory, considering k the parameter, this problem is in $\mathbf{W[P]}$, and for the case of 0-1 payoffs in $\mathbf{W[1]}$.*

In particular, if l and k are constants, the complexity is polynomial, and we can approximate the minmax value with any polynomial number of bits of accuracy in polynomial time by using the decision procedure in a binary search. As the exponent in the above complexity bound depends linearly on k with impractical bounds for large k as consequence, we next ask if the problem of approximating the minmax value for Player 1 in a three-player game is *fixed parameter tractable*, i.e., if an algorithm solving the problem in time $f(k)n^c$ exists, where f is any function and c is a constant not depending on k . We provide a reduction from k -CLIQUE that gives negative evidence.

Theorem 4. *Deciding k -CLIQUE in a graph with n vertices reduces in polynomial time to approximating the minmax value for Player 1 within $1/(4k^2)$ in a three-player $2k \times kn \times kn$ game with payoffs 0 and 1.*

Downey and Fellows [7] proved that the k -CLIQUE problem is complete for the class $\mathbf{W[1]}$, and hence it immediately follows that the problem of approximating the minmax value within $1/k^2$ for Player 1 in a $k \times n \times n$ game with k being the parameter is hard for $\mathbf{W[1]}$, even when all payoffs are 0 or 1. Combining this with Theorem 3, we in fact have that the 0-1 case is $\mathbf{W[1]}$ -complete. Readers

¹ As the algorithms of Theorem 1 and Theorem 2 are very simple, we express their complexity in the unit cost random access machine model. E.g., by “linear time” we mean a linear number of atomic operations in the number of real payoffs of the input. On the other hand, the algorithm of Theorem 3 uses sophisticated algorithms from the literature as subroutines and its complexity is better expressed in the Turing machine model, and in terms of bit complexity.

not well-versed in the theory of parameterized complexity may find the following consequence of the reduction more appealing: The minmax value of a $k \times n \times n$ three-player game with 0-1 payoffs cannot be approximated in time $n^{o(k)}$, unless k -CLIQUE can be solved in time $n^{o(k)}$. If k -CLIQUE could be solved in time $n^{o(k)}$ then as proved by Chen *et al.* [3] it would follow that all problems in the class **SNP** (e.g. 3-SAT) could be solved in time $2^{o(n)}$. Thus, under the assumption that all of **SNP** cannot be solved in time $2^{o(n)}$, the algorithm of Theorem 3 is essentially optimal for the case of 0-1 payoffs, in the sense that its complexity is $n^{O(k)}$ and $n^{\Omega(k)}$ is a lower bound.

2 Proofs

2.1 Proof of Theorem 1

We first prove the hardness claim. Borgs *et al.* [1, Theorem 1] showed hardness of approximation with additive error $3/n^2$. Now consider, for a positive integer $c \geq 2$ the following “padding” construction: Given an $n \times n \times n$ game G with strategy space S_i of Player i and utility function u_1 for Player 1. Let $n' = n^c$ and define the $n' \times n' \times n'$ game G' with strategy space $S'_i = S_i \times \{1, \dots, n^{c-1}\}$ and utility function for Player 1 being $u'_1((x, a_1), (y, a_2), (z, a_3)) = u_1(x, y, z)$. In words, G' is simply G with each strategy copied n^{c-1} times. Now, G' and G clearly have the same minmax value. Also, for a given $\epsilon > 0$, by picking c to be a large enough constant, we can ensure that $1/(n')^\epsilon < 3/n^2$, so approximating the minmax value of G within $3/n^2$ reduces to approximating the minmax value of G' within $1/(n')^\epsilon$, which concludes the proof of hardness (we remark that this simple padding argument also yields a somewhat simpler proof of Lemma 7.1 of Chen, Teng and Valiant [4]).

We now proceed with the positive approximation result. We only show the result for the case of three players; the general case being very similar. For the proof, we will use the following theorem by Lipton and Young [10, Theorem 2]:

Theorem 5. *For a two-player zero-sum $n \times n$ game with payoffs in $[0, 1]$, there is a simple strategy for each player that guarantees a payoff within ϵ of the value of the game. Here, a simple strategy is one that mixes uniformly on a multiset of $\lceil \ln n / (2\epsilon^2) \rceil$ pure strategies.*

Now consider a given 3-player game G and consider the optimal threat strategy profile (σ_2, σ_3) of Players 2 and 3 against Player 1. Consider σ_3 as fixed and look at the resulting two-player game G' between Player 1 (maximizer) and Player 2 (minimizer). Clearly, this game has value equal to the minmax value for Player 1 in G . Applying Theorem 5, there is a simple strategy σ'_2 for Player 2 that guarantees this value within ϵ . Fix σ'_2 to this strategy and look at the resulting two-player game G'' between Player 1 and Player 3. By construction of σ'_2 , this game has value at most ϵ larger than the value of G' . Applying Theorem 5 again, there is a simple strategy σ'_3 for Player 3 that guarantees this value within ϵ . Thus, if Player 2 and Player 3 play the profile (σ'_2, σ'_3) in the original game, they are guaranteed the minmax value of G plus at most 2ϵ .

Now, given some ϵ' , we let $\epsilon = \epsilon'/2$ and approximate the threat value of Player 1 within ϵ' by exhaustively searching through all pairs of simple strategies for Player 2 and Player 3, finding the best response of Player 1 to each of them and the associated payoff, and returning the lowest such payoff. This completes the proof of the theorem.

It is natural to ask if one can get any non-trivial approximation by considering strategies that mix uniformly over only a constant size multiset, as this would lead to a polynomial time approximation algorithm rather than a quasi-polynomial one. Unfortunately, the answer is negative: For given n and s , let m be maximal such that $\binom{m}{s}^2 \leq n$. Then

$$n < \binom{m+1}{s}^2 \leq \binom{2m}{s}^2 \leq \left(\frac{2me}{s}\right)^{2s}.$$

Consider the $\binom{m}{s}^2 \times m \times m$ game G defined as follows. For every two subsets T_2, T_3 of pure strategies of size s for Player 2 and Player 3 there is a pure strategy, a_{T_2, T_3} for Player 1 so that $u_1(a_{T_2, T_3}, a_2, a_3) = 1$ for $a_2 \in T_2, a_3 \in T_3$ and $u_1(a_{T_2, T_3}, a_2, a_3) = 0$ for $a_2 \notin T_2$ or $a_3 \notin T_3$. If Player 2 and Player 3 both play a uniform mix on their entire strategy spaces, Player 1 can ensure payoff at most $(\frac{s}{m})^2$. On the other hand, if Player 2 and Player 3 play mixed strategies of support size at most s then Player 1 has a reply ensuring payoff 1. We can now, in a similar way as in the construction in the beginning of this section, construct a padded version of the game, obtaining an $n \times n \times n$ game G' such that the minmax value for Player 1 is at most $(\frac{s}{m})^2 < \frac{(2e)^2}{n^{1/s}}$, but for every strategy profile for Player 2 and Player 3 of support size at most s , Player 1 can ensure payoff 1. Thus to approximate the minmax value within some constant $c < \frac{1}{2} - \frac{2e^2}{n^{1/s}}$, we must have $s > \frac{\ln n}{\ln(\frac{4e^2}{1-2c})}$.

2.2 Proof of Theorem 2

First, we give the claimed examples of games for which the minmax value for Player 1 is irrational. We describe each game by a matrix for each action of Player 1, where rows and columns correspond to the actions of Player 2 and Player 3, respectively. That is, we let $u_1(i, j, k) = A_i(j, k)$.

The first game is a $2 \times 2 \times 2$ game where there are 3 distinct payoffs, given by the following matrices.

$$A_1 = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \quad A_2 = \begin{bmatrix} 0 & 0 \\ 0 & 2 \end{bmatrix}$$

It is easy to see that the minmax strategy profile for Player 2 and Player 3 is the profile where both row 1 and column 1 are played with probability $2 - \sqrt{2}$. This results in the minmax value $6 - 4\sqrt{2}$ for Player 1.

The second game is a $3 \times 2 \times 2$ game where there are 2 distinct payoffs, given by the following matrices.

$$A_1 = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \quad A_2 = \begin{bmatrix} 0 & 0 \\ 1 & 1 \end{bmatrix} \quad A_3 = \begin{bmatrix} 0 & 1 \\ 0 & 1 \end{bmatrix}$$

It is easy to see that the minmax strategy profile for Player 2 and Player 3 is the profile where both row 1 and column 1 are played with probability $\frac{\sqrt{5}-1}{2}$. This results in the minmax value $\frac{3-\sqrt{5}}{2}$ for Player 1.

We now examine the special case where there are only two distinct payoffs, and Player 1 only has two possible actions. For this case, we show that the minmax value is a rational number, and that the optimal threat can be computed in linear time. We assume without loss of generality that the two possible payoffs are 0 and 1. The proof is a case analysis, where each case can be identified and solved in linear time, assuming that none of the previous cases apply. Case 1 and 2 are the trivial cases where either side has a pure optimal strategy.

Case 1: $\exists i \forall j, k : u_1(i, j, k) = 1$. Player 1 has a “safe” action, i , such that no matter what Players 2 and 3 do, Player 1 achieves value 1. Any strategy profile for Players 2 and 3 is an optimal threat, with minmax value 1.

Case 2: $\exists j, k \forall i : u_1(i, j, k) = 0$. The strategy profile (j, k) is an optimal threat, with minmax value 0.

Case 1 and case 2 can easily be identified and solved in linear time. Notice that when we are not in case 2, we have that $\forall j, k \exists i : u_1(i, j, k) = 1$, and therefore $u_1(i, j, k) = 0 \Rightarrow u_1(i', j, k) = 1$ where $i' \neq i$. This means that Player 1 has a maxmin (security) value of at least $\frac{1}{2}$, which can be achieved by a uniform mix of the two strategies. As the minmax value is at least the maxmin value, any threat with minmax value $\frac{1}{2}$ will be optimal. This is exactly what can be achieved in the next two cases:

Case 3: $\exists j \forall i \exists k : u_1(i, j, k) = 0$. Player 2 has a pure strategy, such that Player 3 can play matching pennies against Player 1. Let k and k' be the strategies of Player 3 achieving payoff 0 against i and i' respectively. Then $(j, (\frac{1}{2}k, \frac{1}{2}k'))$ is an optimal threat, with minmax value of $\frac{1}{2}$.

Case 4: $\exists k \forall i \exists j : u_1(i, j, k) = 0$. Player 3 has a pure strategy, such that Player 2 can play matching pennies against Player 1. Let j and j' be the strategies of Player 2 achieving payoff 0 against i and i' respectively. Then $((\frac{1}{2}j, \frac{1}{2}j'), k)$ is an optimal threat, with minmax value of $\frac{1}{2}$.

Case 3 and case 4 can again easily be identified and solved in linear time.

Case 5: None of the above. The negation of case 1 implies $\forall i \exists j, k : u_1(i, j, k) = 0$. The negation of cases 2, 3 and 4 implies $u_1(i, j, k) = 0 \Rightarrow \forall j', k' : u_1(i', j', k) = u_1(i', j, k') = 1$, where $i' \neq i$. That is, any strategy of Player 2 or 3 can achieve payoff 0 against at most one of Player 1's strategies, and Players 2 and 3 must agree on which strategy of Player 1 to try to get payoff 0 against. If they disagree, the payoff is 1 no matter what Player 1 does. The best they can hope for is therefore $\min_{p, q \in [0, 1]} \max\{1 - pq, 1 - (1 - p)(1 - q)\}$, which gives a lower bound on the minmax value of $\frac{3}{4}$. This value can be achieved in this case in the following way: let $u_1(i, j, k) = u_1(i', j', k') = 0$. Then $((\frac{1}{2}j, \frac{1}{2}j'), (\frac{1}{2}k, \frac{1}{2}k'))$ is an optimal threat, with minmax value of $\frac{3}{4}$.

2.3 Proof of Theorem 3

We prove the theorem for three players, the general case is similar. Shapley and Snow [14] showed that every $k \times n$ zero-sum game has a minmax mixed strategy for Player 2 of support at most k , i.e., using at most k pure strategies. We claim that from this it follows that in every $k \times n \times n$ game there are mixed strategies for Player 2 and Player 3 of support at most k so that the resulting strategy profile σ_{-1} yields the minmax value for Player 1 when Player 1 chooses a best response. Indeed, consider the actual minmax strategy profile $\sigma_{-1}^* = (\sigma_2^*, \sigma_3^*)$. If we consider σ_3^* fixed and consider the resulting two-player game between Player 1 and Player 2, it is clear that σ_2^* is a minmax strategy of this game and that Player 2 will still guarantee the minmax payoff by playing the minmax strategy σ_2^* of support k which is guaranteed to exist by Shapley and Snow's result. Similarly, we may replace σ_3^* with a strategy of support k without changing the payoff resulting when Player 1 plays a best response.

Our algorithm is a support enumeration algorithm which exhaustively examines each possible support of size k for Player 2 and Player 3. From the above observation it follows that the minmax value of the game is the minimum of the minmax value of each of the resulting $k \times k \times k$ subgames. Therefore, we only have to explain how to compare the minmax value of such a subgame to a given α , and we will be done. For this, we appeal to classical results on the first order theory of the reals.

The decision procedure for the first order theory of the reals due to Renegar [11,12,13] can decide a sentence with $\omega - 1$ quantifier alternations, the k th block of variables being of size n_k , containing m atomic predicates and involving only polynomials of degree at most d with integer coefficients of maximum bit length L using $L(\log L)2^{O(\log^* L)}(md)^{2^{O(\omega)}} \prod_k n_k$ bit operations² and $(md)^{O(\sum_k n_k)}$ evaluations of the Boolean formula of the atomic predicates. We claim that from this it follows that given a $k \times k \times k$ game G with rational payoffs and a rational number α so that (G, α) has combined bit complexity L , we can decide in time $L(\log L)2^{O(\log^* L)}k^{O(k)}$ (on a Turing machine) if the minmax value of G for Player 1 is at most α . We can assume that the payoffs and α are integers at the expense of increasing the bitlength of every number to at most the combined bitlength of the original problem. Define the following polynomials in $2k$ variables.

$$p_l(x_1, \dots, x_{2k}) = \sum_{i=1}^k \sum_{j=k+1}^{2k} u_1(l, i, j - k)x_i x_j \quad , \quad r_i(x_1, \dots, x_{2k}) = x_i$$

$$q_1(x_1, \dots, x_{2k}) = \sum_{i=1}^k x_i \quad , \quad q_2(x_1, \dots, x_{2k}) = \sum_{i=k+1}^{2k} x_i$$

The sentence we must decide is then

$$(\exists x \in \mathbf{R}^{2k}) [p_1(x) \leq \alpha \wedge \dots \wedge p_k(x) \leq \alpha \wedge q_1(x) = 1 \wedge q_2(x) = 1 \\ \wedge r_1(x) \geq 0 \wedge \dots \wedge r_{2k}(x) \geq 0] .$$

² The bound stated here is the improvement of the bound stated by Renegar due to the recent breakthrough in integer multiplication due to Fürer [9].

For this sentence we have $\omega = 1$, $m = 3k+2$, $d = 2$ and $n_1 = 2k$, and the sentence can thus be decided in the claimed running time using Renegar's procedure. For the support enumeration algorithm this decision procedure must be invoked for $\binom{n}{k}^2$ different $k \times k \times k$ subgames, and the claimed time bound of the statement of the theorem follows.

Next, we show how to use this algorithm to show $\mathbf{W}[\mathbf{P}]$ and $\mathbf{W}[\mathbf{1}]$ membership of the two versions of the problem. We use the framework of afpt-programs of Chen, Flum and Grohe [5] and Buss and Islam [2] to do this. To transform the algorithm into an afpt-program showing that the decision problem is in the class $\mathbf{W}[\mathbf{P}]$, we simply replace the enumeration by existential steps guessing the sets of indices of size k giving the support of the strategies of Player 2 and Player 3. In the remainder of this section we will show that for the special case of 0-1 payoffs the decision problem is in the class $\mathbf{W}[\mathbf{1}]$. The idea is to precompute, for every possible $k \times k \times k$ game with 0-1 payoffs, whether the minmax value for Player 1 is at most α . As in the $\mathbf{W}[\mathbf{P}]$ case, indices of the support of the strategies are guessed, but now the $k \times k \times k$ subgame is used as an index in the precomputed table. To see that this can be turned into an appropriate afpt-program, we will formally define the relations used.

Assume that the payoffs of Player 1 are given as a k -tuple of $n \times n$ 0-1 matrices (U^1, \dots, U^k) . Define a unary relation A over k -tuples of $k \times k$ 0-1 matrices as follows: $(M^1, \dots, M^k) \in A$ if and only if the minmax value for Player 1 in the $k \times k \times k$ subgame given by (M^1, \dots, M^k) is at most α .

Define a 6-ary relation B having as first argument a k -tuple of $k \times k$ 0-1 matrices and with the next 3 arguments being indices from 1 to k and the two last arguments being indices from 1 to n as follows.

$$((M^1, \dots, M^k), l, a, b, i, j) \in B \quad \text{if and only if} \quad M_{ab}^l = U_{ij}^l .$$

The algorithm first computes the relations A and B . In the guessing steps the algorithm guesses a k -tuple of matrices (M^1, \dots, M^k) and indices i_1, \dots, i_k and j_1, \dots, j_k . The final checks the algorithm must perform are $(M^1, \dots, M^k) \in A$ and $((M^1, \dots, M^k), l, a, b, i_a, j_b) \in B$ for all $l, a, b \in \{1, \dots, k\}$. The number of steps used for guessing the indices and the final checks is a function depending only on the parameter k , as required.

As discussed by Buss and Islam [2] we can in a generic way transform an algorithm utilizing a constant number of relations into one that only utilizes a single binary relation, thereby obtaining an afpt-algorithm showing that the decision problem is in $\mathbf{W}[\mathbf{1}]$.

2.4 Proof of Theorem 4

Before starting the proof, we remark that the reduction is based on similar ideas as the reduction proving \mathbf{NP} -hardness by Borgs *et al.* However, they reduce from 3-COLORING rather than k -CLIQUE, and in their coloring based games, we don't see how to restrict the strategy space of Player 1 to a small number of strategies, so as to obtain fixed-parameter intractability.

We now describe the reduction. We assume throughout the proof that $k \geq 5$. Given an undirected graph $G = (V, E)$, with $|V| = n$. We construct a $2k \times kn \times kn$ game from G in the following way. Let $S_1 = \{1, \dots, k\} \times \{2, 3\}$ be the strategy space of Player 1 and $S_2 = S_3 = \{1, \dots, k\} \times V$ be the strategy spaces of Player 2 and Player 3. We define the payoff of Player 1 as:

$$u_1((x_1, i), (x_2, v_2), (x_3, v_3)) = \begin{cases} 1 & \text{if } x_1 = x_i \\ 1 & \text{if } x_2 = x_3 \text{ and } v_2 \neq v_3 \\ 1 & \text{if } x_2 \neq x_3 \text{ and } v_2 = v_3 \\ 1 & \text{if } v_2 \neq v_3 \text{ and } (v_2, v_3) \notin E \\ 0 & \text{otherwise} \end{cases}$$

As Player 2 and Player 3 try to minimize the payoff of Player 1 we shall refer to them as *bullies*. One can think of the game as the bullies each choosing a label and a vertex of G . Player 1 then chooses one of the bullies and tries to guess his label. If he guesses correctly he will get a payoff of 1. If not, he will get a payoff of 0, unless the bullies do one of the following:

- (i) Choose the same label, but different vertices.
- (ii) Choose different labels, but the same vertex.
- (iii) Choose a pair of distinct vertices that does not correspond to an edge.

In these cases he will get a payoff of 1. The intuition behind the proof is that the bullies will be able to avoid these three cases if the graph contains a k -clique, thereby better punishing Player 1. Formally, we shall show that the minmax value for Player 1 is exactly $\frac{1}{k}$ if G contains a k -clique and at least $\frac{1}{k} + \frac{1}{4k^2}$ if G does not contain a k -clique.

First, we notice that if G contains a k -clique, the bullies can bring down the payoff of Player 1 to $1/k$ by choosing a vertex from the k -clique uniformly at random and agreeing on a labeling of the vertices. For given strategies of the bullies, let p_{\max} be the highest probability (over j) of any of the bullies choosing a label $j \in \{1, \dots, k\}$. Player 1 will always be able to get a payoff of p_{\max} by choosing j and the corresponding player. It follows that the minmax value of the game is $1/k$ when G contains a k -clique, as desired.

Next, we consider the case when G contains no k -clique. Assume to the contrary that the bullies can force Player 1 to get a payoff less than $1/k + 1/(4k^2)$. For the rest of the proof, consider a fixed strategy profile of the bullies with this property. We have already seen that in this case, $p_{\max} < 1/k + 1/(4k^2)$.

Consider the case where Player 1 always chooses Player 2 and guesses a label uniformly at random. In this case Player 1 will always guess the correct label with probability $1/k$ independently of the actions of the bullies. Let p be the probability of either (i), (ii) or (iii) happening. We have:

$$\frac{1}{k} + \left(1 - \frac{1}{k}\right) p < \frac{1}{k} + \frac{1}{4k^2} \Rightarrow p < \frac{1}{4k(k-1)} . \quad (1)$$

In particular, the probability of (i) happening is less than $1/(4k(k-1))$. Let p_{\min} be the minimum probability assigned to any label by either of the bullies.

We have $p_{\min} \geq 1 - (k-1)p_{\max} > 3/(4k) + 1/(4k^2)$. Let $(x, v)_i$ denote the event that Player i chooses the label x and the vertex v . We will use \cdot as a wildcard, such that $(x, \cdot)_i$ denotes the event that Player i chooses the label x and $(\cdot, v)_i$ denotes the event that Player i chooses the vertex v . For $v, w \in V$ we see that:

$$\begin{aligned}
\frac{1}{4k(k-1)} &> \Pr[x_2 = x_3 \text{ and } v_2 \neq v_3] \\
&= \sum_{j=1}^k \sum_{v \neq w} \Pr[(\cdot, v)_2 \mid (j, \cdot)_2] \Pr[(j, \cdot)_2] \Pr[(\cdot, w)_3 \mid (j, \cdot)_3] \Pr[(j, \cdot)_3] \\
&\geq p_{\min}^2 \sum_{j=1}^k \sum_{v \neq w} \Pr[(\cdot, v)_2 \mid (j, \cdot)_2] \Pr[(\cdot, w)_3 \mid (j, \cdot)_3] \\
&= p_{\min}^2 \sum_{j=1}^k \left(1 - \sum_{v \in V} \Pr[(\cdot, v)_2 \mid (j, \cdot)_2] \Pr[(\cdot, v)_3 \mid (j, \cdot)_3] \right) \\
&= kp_{\min}^2 - p_{\min}^2 \sum_{j=1}^k \sum_{v \in V} \Pr[(\cdot, v)_2 \mid (j, \cdot)_2] \Pr[(\cdot, v)_3 \mid (j, \cdot)_3] .
\end{aligned}$$

Here, the first inequality follows from (1). The second inequality is by definition of p_{\min} .

We now estimate the probability that Player 2 and Player 3 choose the same vertex, conditioned on the fact that they choose the same label l . We have that for all l in $\{1, \dots, k\}$:

$$\begin{aligned}
&\sum_{v \in V} \Pr[(\cdot, v)_2 \mid (l, \cdot)_2] \Pr[(\cdot, v)_3 \mid (l, \cdot)_3] \\
&> k - \frac{1}{4k(k-1)p_{\min}^2} - \sum_{j \neq l} \sum_{v \in V} \Pr[(\cdot, v)_2 \mid (j, \cdot)_2] \Pr[(\cdot, v)_3 \mid (j, \cdot)_3] \\
&\geq 1 - \frac{1}{4k(k-1)p_{\min}^2} > 1 - \frac{1}{4k(k-1) \left(\frac{3}{4k} + \frac{1}{4k^2} \right)^2} \\
&= 1 - \frac{4k^3}{(k-1)(3k+1)^2} > \frac{1}{2} .
\end{aligned}$$

The second inequality follows from observing that all terms in the outer sum are bounded from above by 1. The last inequality follows from $k \geq 5$.

For fixed l , let v_i^l be the vertex chosen with highest probability by player i given that he chooses label l and let $q_i^l = \Pr[(\cdot, v_i^l)_i \mid (l, \cdot)_i]$ be that probability. Since $\sum_{v \in V} \Pr[(\cdot, v)_2 \mid (l, \cdot)_2] \Pr[(\cdot, v)_3 \mid (l, \cdot)_3] > \frac{1}{2}$ and $(\Pr[(\cdot, v)_2 \mid (l, \cdot)_2])_v$ and $(\Pr[(\cdot, v)_3 \mid (l, \cdot)_3])_v$ are probability distributions, we have

$$\forall i \in \{2, 3\}, \forall l \in \{1, \dots, k\} : q_i^l > \frac{1}{2} .$$

Since q_2^l and q_3^l are both strictly bigger than $\frac{1}{2}$, and the probability that Player 2 and Player 3 choose the same vertex is strictly bigger than $\frac{1}{2}$, it is easy

to see that we must have $v_2^l = v_3^l$ and we will therefore simply refer to this vertex as v^l . That is, for every label j the bullies agree on some vertex v^j that they both choose with high probability when choosing j . For $j, l \in \{1, \dots, k\}$, it will either be the case that there exists some $v^j = v^l$, with $j \neq l$, or that all the v^j 's are distinct. In the first case Player 1 will, with high probability, get a payoff of 1 when one of the bullies chooses label j and the other chooses label l (case (ii)). In the second case there will exist a pair of distinct labels j and l , such that there is no edge between v^j and v^l , since the graph doesn't contain a k -clique. Hence, this will cause Player 1 to get a payoff of 1, with high probability, when one of the bullies chooses label j and the other chooses label l (case (iii)). In both cases we get that the probability that (ii) or (iii) holds is at least

$$\begin{aligned} & \sum_{i=0}^1 \Pr[(\cdot, v^j)_{2+i} \mid (j, \cdot)_{2+i}] \Pr[(j, \cdot)_{2+i}] \Pr[(\cdot, v^l)_{3-i} \mid (l, \cdot)_{3-i}] \Pr[(l, \cdot)_{3-i}] \\ & > 2p_{\min}^2 \left(\frac{1}{2}\right)^2 > \frac{1}{2} \left(\frac{3}{4k} + \frac{1}{4k^2}\right)^2 = \frac{(3k+1)^2}{32k^4} > \frac{1}{4k(k-1)}. \end{aligned}$$

The last inequality follows from $k \geq 5$. This contradicts (1) which states that (i), (ii) or (iii) happens with probability less than $1/(4k(k-1))$. Thus, we have completed our proof by contradiction.

3 Conclusions and open problems

As mentioned above, an important open problem is achieving a non-trivial approximation of the minmax value of an $n \times n \times n$ game in polynomial time, rather than quasi-polynomial time. Another interesting question comes from the following notions: The *threat point* of a game is defined to be the vector of minmax values for each of its players. We may consider approximating the threat point of a three player game where one of the players has few strategies. For this, we have to consider the problem of approximating the threat value for Player 1 in a three-player $n \times k \times n$ game. That is, it is now one of the two ‘‘bullies’’ rather than the threatened player that has few strategies. We observe that for constant $\epsilon > 0$ such an approximation can be done efficiently by simply discretizing the mixed strategy space of the player with few strategies using a lattice with all simplex points having distance at most ϵ to some lattice point, and then for each lattice point solving the game for the remaining two players using linear programming. Combining this with Theorem 3 gives us the following corollary:

Corollary 1. *There is an algorithm that, given a $k \times n \times n$ game with 0-1 payoffs and an $\epsilon > 0$, computes the threat point within additive error ϵ in time $(n/\epsilon)^{O(k)}$.*

The discretization technique gives algorithms with poor dependence on the desired additive approximation ϵ . We leave as an open problem if the minmax value of an $n \times k \times n$ 0-1 game can be approximated within ϵ in time $(n \log(1/\epsilon))^{O(k)}$.

Acknowledgements

We would like to thank Bernhard von Stengel and Daniel Andersson for valuable help.

References

1. C. Borgs, J. Chayes, N. Immerlica, A. T. Kalai, V. Mirrokni, and C. Papadimitriou. The myth of the folk theorem. In *Proceedings of the 40th annual ACM Symposium on Theory of Computing (STOC'08)*, pages 365–372, 2008.
2. J. F. Buss and T. Islam. Simplifying the weft hierarchy. *Theoretical Computer Science*, 351(3):303–313, 2006.
3. J. Chen, B. Chor, M. Fellows, X. Huang, D. W. Juedes, I. A. Kanj, and G. Xia. Tight lower bounds for certain parameterized NP-hard problems. *Information and Computation*, 201(2):216–231, 2005.
4. X. Chen, S.-H. Teng, and P. Valiant. The approximation complexity of win-lose games. In *Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'07)*, pages 159–168, 2007.
5. Y. Chen, J. Flum, and M. Grohe. Machine-based methods in parameterized complexity theory. *Theoretical Computer Science*, 339(2–3):167–199, 2005.
6. R. G. Downey and M. R. Fellows. Fixed-parameter tractability and completeness I: Basic results. *SIAM Journal on Computing*, 24(4):873–921, 1995.
7. R. G. Downey and M. R. Fellows. Fixed-parameter tractability and completeness II: On completeness for $W[1]$. *Theoretical Computer Science*, 141(1–2):109–131, 1995.
8. R. G. Downey and M. R. Fellows. *Parameterized complexity*. Monographs in computer science. Springer, New York, 1999.
9. M. Fürer. Faster integer multiplication. In *Proceedings of the 39th Annual ACM Symposium on Theory of Computing (STOC'07)*, pages 57–66, 2007.
10. R. J. Lipton and N. E. Young. Simple strategies for large zero-sum games with applications to complexity theory. In *Proceedings of the 26th annual ACM Symposium on Theory of Computing (STOC'94)*, pages 734–740, 1994.
11. J. Renegar. On the computational complexity and geometry of the first-order theory of the reals, part I: Introduction. Preliminaries. The geometry of semi-algebraic sets. The decision problem for the existential theory of the reals. *Journal of Symbolic Computation*, 13(3):255–299, 1992.
12. J. Renegar. On the computational complexity and geometry of the first-order theory of the reals, part II: The general decision problem. Preliminaries for quantifier elimination. *Journal of Symbolic Computation*, 13(3):301–327, 1992.
13. J. Renegar. On the computational complexity and geometry of the first-order theory of the reals. part III: Quantifier elimination. *Journal of Symbolic Computation*, 13(3):329–352, 1992.
14. L. S. Shapley and R. N. Snow. Basic solutions of discrete games. In *Contributions to the Theory of Games*, number 24 in Annals of Mathematics Studies, pages 27–35. Princeton University Press, 1950.
15. B. von Stengel and D. Koller. Team-maxmin equilibria. *Games and Economic Behavior*, 21:309–321, 1997.