

Well-Supported vs. Approximate Nash Equilibria: Query Complexity of Large Games*

Xi Chen^{†1}, Yu Cheng^{‡2}, and Bo Tang^{§3}

- 1 Columbia University, New York, USA
xichen@cs.columbia.edu
- 2 University of Southern California, Los Angeles, USA
yu.cheng.1@usc.edu
- 3 Oxford University, Oxford, United Kingdom
tangbonk1@gmail.com

Abstract

In this paper we present a generic reduction from the problem of finding an ϵ -well-supported Nash equilibrium (WSNE) to that of finding an $\Theta(\epsilon)$ -approximate Nash equilibrium (ANE), in large games with n players and a bounded number of strategies for each player. Our reduction complements the existing literature on relations between WSNE and ANE, and can be applied to extend hardness results on WSNE to similar results on ANE. This allows one to focus on WSNE first, which is in general easier to analyze and control in hardness constructions.

As an application we prove a $2^{\Omega(n/\log n)}$ lower bound on the randomized query complexity of finding an ϵ -ANE in binary-action n -player games, for some constant $\epsilon > 0$. This answers an open problem posed by Hart and Nisan [23] and Babichenko [2], and is very close to the trivial upper bound of 2^n . Previously for WSNE, Babichenko [2] showed a $2^{\Omega(n)}$ lower bound on the randomized query complexity of finding an ϵ -WSNE for some constant $\epsilon > 0$. Our result follows directly from combining [2] and our new reduction from WSNE to ANE.

1998 ACM Subject Classification F.2 Analysis of Algorithms and Problem Complexity

Keywords and phrases Equilibrium Computation, Query Complexity, Large Games

Digital Object Identifier 10.4230/LIPIcs.ITCS.2017.57

1 Introduction

The celebrated theorem of Nash [29] states that every finite game has an equilibrium point. The solution concept of Nash equilibrium (NE) has been tremendously influential in economics and social sciences ever since (e.g. see [24]). The complexity and efficient approximation of NE have been studied intensively during the past decade, and much progress has been made (e.g., see [27, 1, 6, 26, 34, 13, 9, 16, 28, 14, 4, 10, 32, 7, 12, 15, 33, 11, 5]).

In this paper, we study the randomized query complexity of finding an ϵ -approximate Nash equilibrium (ANE) in large games, for some constant $\epsilon > 0$. Given a game \mathcal{G} with n players and α actions for each player, we index the players by the set $[n] = \{1, \dots, n\}$ and index the actions by the set $[\alpha] = \{1, \dots, \alpha\}$. Recall that an ϵ -ANE of \mathcal{G} is a mixed strategy profile $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$, where each $\mathbf{x}_i \in [0, 1]^\alpha$ sums to 1 and is an ϵ -best response

* This work was done in part while the authors were visiting the Simons Institute.

[†] Xi Chen is supported in part by NSF grants CCF-1149257 and CCF-1423100.

[‡] Yu Cheng is supported in part by Shang-Hua Teng's Simons Investigator Award.

[§] Bo Tang is supported in part by ERC grant 321171.



of player i to other players' strategies \mathbf{x}_{-i} ¹ (see Section 2 for the formal definitions). Since the notion of ANE is additive, we always assume that payoff functions of games considered in this paper take values between 0 and 1.

We consider the payoff query model, where an oracle algorithm with unlimited computational power is given an approximation parameter ϵ , the number of players n and the number of actions α in an unknown game \mathcal{G} , and needs to find an ϵ -ANE of \mathcal{G} . The algorithm has oracle access to the payoff functions of players in \mathcal{G} : For each round, the algorithm can adaptively query a pure strategy profile $\mathbf{a} \in [\alpha]^n$, and receives the payoff of each player with respect to \mathbf{a} . We are interested in the number of queries needed by any randomized oracle algorithm for this task. Note that a trivial upper bound is α^n by simply querying all the pure strategy profiles.

1.1 Prior Results and Related Work

The query complexity of (approximate) Nash equilibria and related solution concepts has received considerable attention recently, e.g., see [17, 23, 18, 19, 2, 3, 20, 33]. Below we review results that are most relevant to our work.

The query complexity of (approximate) correlated equilibria (CE)² is well understood. For the payoff query model considered here, randomized algorithms exist (e.g., regret-minimizing algorithms [22, 21, 8]) for finding an ϵ -CE using $\text{poly}(1/\epsilon, \alpha, n)$ many queries. It turns out that both randomization and approximation are necessary. [3] showed that every deterministic algorithm that finds an exact CE requires exponentially many queries in n . [23] then showed that the same exponential lower bound holds for any deterministic algorithm for (1/2)-CE and any randomized algorithm for exact CE. For the stronger (expected payoff) query model, where the oracle returns the expected payoffs of any mixed strategy profile³, [30] and [25] obtained a deterministic algorithm that computes an exact CE in polynomial time using polynomially many queries (both in α and n).

Now turning to the harder, but perhaps more interesting, problem of approximating Nash equilibria under the payoff query model, the deterministic lower bound of [23] for (1/2)-CE directly implies the same bound for (1/2)-ANE, because any ϵ -ANE by definition is an ϵ -CE as well. For the randomized query complexity, Babichenko [2] showed that any randomized algorithm requires $2^{\Omega(n)}$ queries to find an ϵ -well-supported Nash equilibrium (WSNE), in a binary-action, n -player game. Recall that an ϵ -WSNE of a game is a mixed strategy profile \mathbf{x} in which the probability of player i playing action j is positive only when action j is an ϵ -best response with respect to \mathbf{x}_{-i} . By definition, an ϵ -WSNE is also an ϵ -ANE but the inverse is not true. Following a well-known connection between WSNE and ANE [13] (and using random samples to approximate expected payoffs), [2] showed that the same $2^{\Omega(n)}$ bound holds for the randomized query complexity of ϵ -ANE, but only when $\epsilon = O(1/n)$. Before our work, the randomized query complexity of ϵ -ANE in large games remains an open problem when $\epsilon > 0$ is a constant.

In this paper we prove a $2^{\Omega(n/\log n)}$ lower bound on the randomized query complexity of finding an ϵ -ANE, for some constant $\epsilon > 0$. Subsequently, Rubinstein [33] showed a tight $2^{\Omega(n)}$ lower bound on the randomized query complexity of ϵ -ANE, using more sophisticated machinery in coding theory to remove the $\log n$ factor in the exponent. However, our

¹ We use $\mathbf{x}_{-i} := (\mathbf{x}_1, \dots, \mathbf{x}_{i-1}, \mathbf{x}_{i+1}, \dots, \mathbf{x}_n)$ to denote the strategies of players other than i in \mathbf{x} .

² An ϵ -correlated equilibrium is a probability distribution over pure strategy profiles, such that any player unilaterally deviating from strategies drawn from it can increase her expected payoff by no more than ϵ .

³ Such an oracle can be implemented in polynomial time for many classes of succinct games; see [30].

reduction from ANE to WSNE is much simpler and could be used to obtain hardness results on ANE in other applications (e.g. [31]). It remains an interesting open question if there is a more efficient reduction from WSNE to ANE, without paying the extra factor of $\log n$.

1.2 Our Results

For binary-action, n -player games, we show that $2^{\Omega(n/\log n)}$ queries are required for any randomized algorithm to find an ϵ -ANE, for some constant $\epsilon > 0$. To state the result we use $\text{QC}_p(\text{ANE}(n, \epsilon))$, for some $p > 0$, to denote the smallest T such that there exists a randomized oracle algorithm that uses no more than T queries and outputs an ϵ -ANE with probability at least p , given any unknown binary-action, n -player game. Our main result is the following lower bound on $\text{QC}_p(\text{ANE}(n, \epsilon))$:

► **Theorem 1 (Main).** *There exist two constants $\epsilon > 0$ and $c > 0$ such that*

$$\text{QC}_p(\text{ANE}(n, \epsilon)) = 2^{\Omega(n/\log n)}, \quad \text{where } p = 2^{-cn/\log n}.$$

Our lower bound answers an open problem posed in [23] and in [2]. Our result shows that, in terms of their query complexities, finding an ϵ -ANE is almost as hard as finding an ϵ -WSNE in a large game, even for constant $\epsilon > 0$. It also directly implies the following corollary regarding the rate of convergence of k -queries dynamics (see [2] for the definition).

► **Corollary 2.** *There exist constants $\epsilon, c > 0$ such that no k -queries dynamic can converge to an ϵ -ANE in $2^{\Omega(n/\log n)}/k$ steps with probability at least $2^{-cn/\log n}$ for every binary-action and n -player game.*

Our proof of Theorem 1 relies on a polynomial-time reduction⁴ from the problem of finding an ϵ -WSNE to that of finding an $(\epsilon' = \Omega(\epsilon))$ -ANE in a *succinct game* with a fixed number of actions. As defined in [30], an α -action succinct game is a pair (n, U) , where n is the number of players and U is a (multi-output) Boolean circuit that, given a pure strategy profile $\mathbf{a} \in [\alpha]^n$ (encoded in binary), outputs the payoffs of all n players with respect to \mathbf{a} .

► **Theorem 3.** *Let $\epsilon \geq 0$ and $\alpha \in \mathbb{N}$ be two constants. The problem of finding an ϵ -WSNE is polynomial-time reducible to that of finding an $\epsilon/(4\alpha)$ -ANE, both in α -action succinct games.*

1.3 Approximate vs. Well-Supported Nash Equilibria

Let $\text{QC}_p(\text{WSNE}(n, \epsilon))$ denote the smallest T such that there exists a randomized oracle algorithm that uses no more than T queries and outputs an ϵ -WSNE with probability at least p , given any unknown binary-action, n -player game. Babichenko [2] showed that

► **Theorem 4 ([2]).** *There exist constants $\epsilon, c > 0$ such that*

$$\text{QC}_p(\text{WSNE}(n, \epsilon)) = 2^{\Omega(n)}, \quad \text{where } p = 2^{-cn}.$$

Given Theorem 4, the same exponential lower bound follows directly for the randomized query complexity of ϵ -ANE, for certain small enough constant $\epsilon > 0$, if

⁴ Recall that a polynomial-time reduction from a total search problem A to a total search problem B is a pair (f, g) of polynomial-time computable functions such that: 1) for every input instance x of A , $f(x)$ is an input instance of B ; and 2) for every solution y to $f(x)$ in B , $g(y)$ is a solution to x in A .

Given oracle access to \mathcal{G} and any ϵ' -ANE of \mathcal{G} , where $\epsilon' = c(\alpha) \cdot \epsilon$ for some constant $c > 0$ that only depends on α , there is a query-efficient procedure that outputs an ϵ -WSNE of \mathcal{G} .

However, the best such procedure known is the following result from [13]. The parameters are subsequently improved in [2], where the number of queries needed is also analyzed:

Given oracle access to \mathcal{G} and any $\epsilon^2/(16n)$ -ANE of \mathcal{G} , there is a procedure that outputs an ϵ -WSNE of \mathcal{G} using $\text{poly}(\alpha, n, 1/\epsilon)$ payoff queries, where n denotes the number of players.

The procedure is very natural: For each player, reallocate probabilities on actions with a relatively low expected payoff to a best-response action. Using Theorem 4, such a procedure implies the same exponential lower bound for ϵ -ANE [2] but only when ϵ is $O(1/n)$.

Before our work, no better procedure is known. By definition, an ANE poses a slightly weaker condition on each player compared to that of a WSNE. More specifically, given the mixed strategies of other players \mathbf{x}_{-i} , for an ϵ -WSNE, \mathbf{x}_i must be supported on actions that are ϵ -best responses to \mathbf{x}_{-i} , while in an ϵ -ANE, \mathbf{x}_i can be any mixed strategy that yields an overall ϵ -best response to \mathbf{x}_{-i} . For example, \mathbf{x}_i may allocate $1 - \epsilon$ probability on best-response actions while putting ϵ probability on any other actions. This makes WSNE much easier to analyze and control in hardness reductions, which is why it played a critical role in characterizing the complexity of Nash equilibria, starting with the work of [13], later in [9] and subsequent works. The reason that Babichenko's lower bound (Theorem 4) does not hold for ϵ -ANE is that, if every player places a tiny probability on a suboptimal action, in aggregate there are always some players who play suboptimally, which makes the outcome quite unpredictable.

1.4 Our Approach

We prove Theorem 1 via a *query-efficient* reduction from the problem of finding an ϵ -WSNE to that of finding an $\Theta(\epsilon)$ -ANE:

Given any α -action, n -player game \mathcal{G} and any parameter $\epsilon > 0$, one can define a new α -action game \mathcal{G}' with a slightly larger set of $O(\alpha^2 \log(n/\epsilon) \cdot n)$ players such that

1. *To answer each payoff query on \mathcal{G}' , it suffices to make αn payoff queries on \mathcal{G} ;*
2. *There is a procedure that, given any ϵ -ANE \mathbf{x} of \mathcal{G}' , outputs a $(4\alpha\epsilon)$ -WSNE \mathbf{y} of \mathcal{G} , with no payoff oracle access to \mathcal{G} or \mathcal{G}' .*

Our reduction is presented in Section 3. Theorem 1 then follows immediately from the lower bound of [2] on the randomized query complexity of WSNE (in Theorem 4). Theorem 3 follows from the fact that: 1) the payoff entries of \mathcal{G}' are easy to compute; and 2) the procedure to obtain \mathbf{y} from \mathbf{x} runs in time polynomial in the length of the binary representation of \mathbf{x} , when the number of actions α is bounded. We first give the intuition behind our reduction.

Recall that in the procedure of [13] and [2], an ϵ -WSNE is obtained from an ϵ' -ANE with $\epsilon' = \epsilon^2/(16n)$ by reallocating probabilities on actions with relatively low expected payoff (formally, actions with payoff $\Omega(\epsilon)$ lower than the best response) to best-response actions. From the definition of ANE, no player can have probability more than $O(\epsilon'/\epsilon) = O(\epsilon/n)$ on actions with low payoff in any ϵ' -ANE. Thus, the procedure changes the expected payoff of each player on each action by at most $n \cdot O(\epsilon/n) = O(\epsilon)$ since it changes the mixed strategy of each player by $O(\epsilon/n)$. It follows that the new mixed strategy profile is an ϵ -WSNE. The blow up of a factor of n from ϵ' to ϵ is precisely due to the cumulative impact on a player's expected payoff imposed by small changes to all other players' mixed strategies.

Our reduction from WSNE to ANE overcomes this obstacle by constructing from \mathcal{G} a new and slightly larger game \mathcal{G}' with $O(n \log n)$ players, where each player i in the original n -player game \mathcal{G} is simulated by a group of $O(\log n)$ players indexed by (i, j) in the new game \mathcal{G}' , and we use the *majority* strategy in group i to decide the strategy of player i in the original game. The payoff function of the player (i, j) in \mathcal{G}' is exactly the same as that of player i in \mathcal{G} , but is now defined with respect to the *aggregate* action (by plurality voting) of each group of players in \mathcal{G}' .

We show that an ϵ -WSNE of \mathcal{G} can be recovered from any ϵ' -ANE of \mathcal{G}' , where $\epsilon' = \Omega(\epsilon)$, by (1) computing the distribution of the majority action of each group and (2) truncating the small entries in each distribution. Intuitively, by focusing on the aggregate behavior of each group of $O(\log n)$ *independent* players in \mathcal{G}' , we make sure that the mixed strategies obtained from Step (1) are highly concentrated on actions with close-to-best expected payoffs, and actions with low payoffs can only appear as the majority action of a group with probability $O(\epsilon/n)$. Therefore, in Step (2) we only need to truncate entries with probability $O(\epsilon/n)$, and the remaining positive entries would correspond to close-to-best actions. We can also control the effect of this truncation at the same time, because when the number of actions is bounded, the aggregate behavior of each group changes by at most $O(\epsilon/n)$, which allows us to show that the result is an ϵ -WSNE of the original game \mathcal{G} .

1.5 Organization

The rest of the paper is organized as follows. We first give formal definitions of ANE and WSNE in Section 2. In Section 3 we present the reduction from WSNE to ANE for large games, and then use it to prove Theorem 1 and Theorem 3 in Section 4.

2 Preliminaries

A game \mathcal{G} is a triple (n, α, \mathbf{u}) , where n is the number of players, α is the number of actions for each player, and $\mathbf{u} = (u_1, \dots, u_n)$ are the payoff functions, one for each player. We always use $[n] = \{1, \dots, n\}$ to denote the set of players and $[\alpha] = \{1, \dots, \alpha\}$ to denote the set of actions for each player. Since we are interested in additive approximations, each u_i maps $[\alpha]^n$ to $[0, 1]$.

Let Δ_α denote the set of probability distributions over $[\alpha]$. A mixed strategy profile of \mathcal{G} is then a tuple $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$ of mixed strategies, where $\mathbf{x}_i \in \Delta_\alpha$ denotes the mixed strategy of player i . Given \mathbf{x} , we use \mathbf{x}_{-i} to denote the tuple of mixed strategies of all players other than i . As a shorthand, we write $u_i(\mathbf{x})$ to denote the expected payoff of player i with respect to \mathbf{x} , and write $u_i(a, \mathbf{x}_{-i})$ to denote the expected payoff of player i playing action $a \in [\alpha]$ with respect to \mathbf{x}_{-i} :

$$u_i(\mathbf{x}) = \mathbb{E}_{\mathbf{a} \sim \mathbf{x}} [u_i(\mathbf{a})] \quad \text{and} \quad u_i(a, \mathbf{x}_{-i}) = \mathbb{E}_{\mathbf{b} \sim \mathbf{x}_{-i}} [u_i(a, \mathbf{b})].$$

Next we define approximate and well-supported Nash equilibria.

► **Definition 5.** Given $\epsilon > 0$, an ϵ -approximate Nash equilibrium of an α -action and n -player game $\mathcal{G}(n, \alpha, \mathbf{u})$ is a mixed strategy profile $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$ such that for every player $i \in [n]$:

$$u_i(\mathbf{x}) \geq u_i(a', \mathbf{x}_{-i}) - \epsilon, \quad \text{for all } a' \in [\alpha].$$

► **Definition 6.** Given $\epsilon > 0$, an ϵ -well-supported Nash equilibrium of $\mathcal{G}(n, \alpha, \mathbf{u})$ is a mixed strategy profile $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$ such that for all $i \in [n]$ and action a in the support of \mathbf{x}_i :

$$u_i(a, \mathbf{x}_{-i}) \geq u_i(a', \mathbf{x}_{-i}) - \epsilon, \quad \text{for all } a' \in [\alpha].$$

Finally, we give a formal definition of succinct games [30].

► **Definition 7.** An α -action succinct game is a pair (n, U) , where n is the number of players and U is a (multi-output) Boolean circuit that, given any pure strategy profile $\mathbf{a} \in [\alpha]^n$ (encoded in binary), outputs the payoffs of all n players with respect to \mathbf{a} in the game. Note that the input size of (n, U) is the size of the circuit U .

3 A Reduction from WSNE to ANE

Given an α -action, n -player game $\mathcal{G}(n, \alpha, \mathbf{u})$ and a parameter $\epsilon \in (0, 1)$, we now define a new game $\mathcal{G}'(sn, \alpha, \mathbf{u}')$ with sn players, where $s = 2\alpha^2 \cdot \lceil \ln(n/\epsilon) \rceil$. We prove that given an ϵ -ANE \mathbf{x} of the new game \mathcal{G}' , one can compute a $(4\alpha\epsilon)$ -WSNE \mathbf{y} of \mathcal{G} without making any payoff queries to \mathcal{G} or \mathcal{G}' .

For each player $i \in [n]$ in \mathcal{G} , we introduce a group of s players in \mathcal{G}' , indexed by (i, j) with $j \in [s]$, and use $u'_{i,j}$ to denote the payoff function of player (i, j) . Given any pure strategy profile $\mathbf{a} = (a_{i,j} : i \in [n], j \in [s])$, we define the payoff $u'_{i,j}(\mathbf{a})$ of player (i, j) as follows. First, for each $i \in [n]$, let $\bar{a}_i \in [\alpha]$ denote the *majority* action played by the i -th group (players (i, j) , $j \in [s]$) in the pure strategy profile \mathbf{a} (break ties by choosing the action with the smallest index). Write $\bar{\mathbf{a}} = (\bar{a}_1, \dots, \bar{a}_n)$. Next, the payoff of player (i, j) under \mathbf{a} is defined as

$$u'_{i,j}(\mathbf{a}) = u_i(a_{i,j}, \bar{\mathbf{a}}_{-i}). \quad (1)$$

This completes the definition of \mathcal{G}' . The next lemma follows from the definition.

► **Lemma 8.** *To answer a payoff query on \mathcal{G}' , it suffices to make αn queries on \mathcal{G} .*

Proof. By the definition of \mathcal{G}' , $u'_{i,j}(\mathbf{a})$'s for all (i, j) , are determined by

$$(u_i(a', \bar{\mathbf{a}}_{-i}) : i \in [n], a' \in [\alpha]),$$

for which αn payoff queries on \mathcal{G} suffice. ◀

We conclude our reduction by proving the following lemma:

► **Lemma 9.** *Given any ϵ -ANE \mathbf{x} of \mathcal{G}' , one can compute a $(4\alpha\epsilon)$ -WSNE \mathbf{y} of \mathcal{G} without making any payoff queries on \mathcal{G} or \mathcal{G}' . Moreover, when α is a constant, the computation of \mathbf{y} from \mathbf{x} can be done in time polynomial in the number of bits needed in the binary representation of \mathbf{x} and $1/\epsilon$.*

Proof. Let $\mathbf{x} = (x_{i,j})$ be an ϵ -ANE of \mathcal{G}' . For each group i and action $k \in [\alpha]$, let

$$\bar{x}_{i,k} = \Pr_{\mathbf{a} \sim \mathbf{x}} [\bar{a}_i = k]. \quad (2)$$

Recall that \bar{a}_i is the majority action played by players (i, j) , $j \in [s]$, in the pure strategy profile \mathbf{a} . By definition, each $\bar{\mathbf{x}}_i = (\bar{x}_{i,1}, \dots, \bar{x}_{i,\alpha})$ is a probability distribution over $[\alpha]$.

Next, we define a mixed strategy $\mathbf{y} = (\mathbf{y}_1, \dots, \mathbf{y}_n)$ of \mathcal{G} , and show that \mathbf{y} is a $(4\alpha\epsilon)$ -WSNE. We zero out entries smaller than ϵ/n in $\bar{\mathbf{x}}_i$ and rescale it, formally

$$c_{i,k} = \begin{cases} \bar{x}_{i,k} & \text{if } \bar{x}_{i,k} \leq \epsilon/n \\ 0 & \text{otherwise} \end{cases} \quad \text{and} \quad y_{i,k} = \frac{\bar{x}_{i,k} - c_{i,k}}{1 - \sum_{j \in [\alpha]} c_{i,j}}. \quad (3)$$

It is easy to verify that $\mathbf{y}_i = (y_{i,1}, \dots, y_{i,\alpha})$ is indeed a probability distribution over $[\alpha]$.

Now assume for contradiction that \mathbf{y} is not a $(4\alpha\epsilon)$ -WSNE, i.e. for some player $i \in [n]$ there exists an action $\ell \in [\alpha]$ such that $y_{i,\ell} > 0$ but

$$\max_{k \in [\alpha]} u_i(k, \mathbf{y}_{-i}) > u_i(\ell, \mathbf{y}_{-i}) + 4\alpha\epsilon. \quad (4)$$

But note that, the total variation distance between $\bar{\mathbf{x}}_j$ and \mathbf{y}_j for each $j \in [n]$ is at most $\alpha\epsilon/n$. So by coupling and applying union bound, we have that

$$|u_i(k, \bar{\mathbf{x}}_{-i}) - u_i(k, \mathbf{y}_{-i})| \leq (n-1) \cdot (\alpha\epsilon/n) < \alpha\epsilon, \quad \text{for all } k \in [\alpha]. \quad (5)$$

It then follows from (4) and (5) that

$$\max_{k \in [\alpha]} u_i(k, \bar{\mathbf{x}}_{-i}) > u_i(\ell, \bar{\mathbf{x}}_{-i}) + 2\alpha\epsilon. \quad (6)$$

By the definition (1) of the payoff function $u'_{i,j}$, we have

$$u'_{i,j}(k, \mathbf{x}_{-i}) = u_i(k, \bar{\mathbf{x}}_{-i}), \quad \text{for all } j \in [s] \text{ and } k \in [\alpha]. \quad (7)$$

Combining (6) and (7), we have that for every player (i, j) , $j \in [s]$:

$$\max_{k \in [\alpha]} u'_{i,j}(k, \mathbf{x}_{-i}) - u'_{i,j}(\ell, \mathbf{x}_{-i}) \geq 2\alpha\epsilon.$$

Since \mathbf{x} is an ϵ -ANE of \mathcal{G}' , $x_{i,j,\ell} \leq 1/(2\alpha)$. By Hoeffding bound and our choice of s ,

$$\begin{aligned} \bar{x}_{i,\ell} &= \Pr[\ell \text{ is the majority action among players } (i, j), j \in [s]] \\ &\leq \Pr[\text{the number of players } (i, j) \text{ playing } \ell \text{ is at least } s/\alpha] \leq e^{-s/(2\alpha^2)} \leq \epsilon/n. \end{aligned}$$

By (3), this implies that $y_{i,\ell} = 0$, which contradicts our assumption and proves that \mathbf{y} is indeed a $(4\alpha\epsilon)$ -WSNE of \mathcal{G} .

It is clear that from the definition of \mathbf{y} , the computation of \mathbf{y} from \mathbf{x} does not require any payoff queries. For the running time, when α is a constant, to compute $\bar{x}_{i,k}$ in (2) one needs to go through

$$\alpha^s = \alpha^{2\alpha^2 \cdot \lceil \ln(n/\epsilon) \rceil} = (n/\epsilon)^{O(1)}$$

many pure strategy profiles of players (i, j) , $j \in [s]$. Therefore \mathbf{y} can be computed in time polynomial in the number of bits needed in the binary representation of \mathbf{x} and $1/\epsilon$. ◀

4 Proofs of Theorems 1 and 3

We use our query-efficient reduction to prove Theorem 1 and Theorem 3.

Proof of Theorem 1. By Theorem 4 there exist constants $\epsilon', c' > 0$ such that

$$\text{QC}_{p'}(\mathbf{WSNE}(n', \epsilon)) = 2^{\Omega(n')}, \quad \text{where } p' = 2^{-c'n'}.$$

Let $n = 8n' \cdot \lceil \ln(n'/\epsilon') \rceil$ and $\epsilon = 8\epsilon'$. It follows from Lemma 8 and Lemma 9 that

$$\text{QC}_{p'}(\mathbf{ANE}(n, \epsilon)) \geq \text{QC}_{p'}(\mathbf{WSNE}(n', \epsilon)) = 2^{\Omega(n')}.$$

The theorem then follows from $n' = \Omega(n/\log n)$. ◀

Proof of Theorem 3. Using Lemma 9, it suffices to prove that, given any α -action succinct game $\mathcal{G} = (n, U)$, we can construct in polynomial time a Boolean circuit U' that implements the payoff functions of players in \mathcal{G}' . This can be done by following the definition of \mathcal{G}' in the previous section since the payoffs of a pure strategy profile \mathbf{a} in \mathcal{G}' only depends (in a straight-forward fashion) on the payoffs of $\alpha n = O(n)$ easy-to-compute profiles of \mathcal{G} . ◀

References

- 1 T. Abbott, D. Kane, and P. Valiant. On the complexity of two-player win-lose games. In *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*, pages 113–122, 2005.
- 2 Y. Babichenko. Query complexity of approximate Nash equilibria. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, pages 535–544, 2014.
- 3 Y. Babichenko and S. Barman. Query complexity of correlated equilibrium. *ACM Transactions on Economics and Computation*, 3(4):22:1–22:9, 2015.
- 4 Y. Babichenko, C.H. Papadimitriou, and A. Rubinstein. Can almost everybody be almost happy? PCP for PPAD and the inapproximability of Nash. In *Proceedings of the 7th Annual Innovations in Theoretical Computer Science*, 2015.
- 5 Y. Babichenko and A. Rubinstein. Communication complexity of approximate Nash equilibria. *ArXiv e-prints*, 2016.
- 6 I. Bárány, S. Vempala, and A. Vetta. Nash equilibria in random games. In *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*, pages 123–131, 2005.
- 7 S. Barman. Approximating Nash equilibria and dense bipartite subgraphs via an approximate version of Caratheodory’s theorem. In *Proceedings of the 47th Annual ACM Symposium on Theory of Computing*, pages 361–369, 2015.
- 8 A. Blum and Y. Monsoor. *From External to Internal Regret*. University of Chicago Press, 2007.
- 9 X. Chen, X. Deng, and S.-H. Teng. Settling the complexity of computing two-player Nash equilibria. *Journal of the ACM*, 56(3):1–57, 2009.
- 10 X. Chen, D. Durfee, and A. Orfanou. On the complexity of Nash equilibria in anonymous games. In *Proceedings of the 46th ACM Symposium on Theory of Computing*, pages 381–390, 2015.
- 11 Y. Cheng, I. Diakonikolas, and A. Stewart. Playing anonymous games using simple strategies. In *Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms*, 2017. To appear.
- 12 C. Daskalakis, A. De, G. Kamath, and C. Tzamos. A size-free CLT for Poisson multinomials and its applications. In *Proceedings of the 48th Annual ACM Symposium on Theory of Computing*, pages 1074–1086, 2016.
- 13 C. Daskalakis, P.W. Goldberg, and C.H. Papadimitriou. The complexity of computing a Nash equilibrium. *SIAM Journal on Computing*, 39(1), 2009.
- 14 C. Daskalakis and C.H. Papadimitriou. Approximate Nash equilibria in anonymous games. *Journal of Economic Theory*, 156:207–245, 2015.
- 15 I. Diakonikolas, D.M. Kane, and A. Stewart. The Fourier transform of Poisson multinomial distributions and its algorithmic applications. In *Proceedings of the 48th Annual ACM Symposium on Theory of Computing*, pages 1060–1073, 2016.
- 16 K. Etessami and M. Yannakakis. On the complexity of Nash equilibria and other fixed points. *SIAM Journal on Computing*, 39(6):2531–2597, 2010.
- 17 J. Fearnley, M. Gairing, P. Goldberg, and R. Savani. Learning equilibria of games via payoff queries. In *Proceedings of the 14th ACM Conference on Electronic Commerce*, pages 397–414, 2013.
- 18 J. Fearnley and R. Savani. Finding approximate Nash equilibria of bimatrix games via payoff queries. In *Proceedings of the 15th ACM Conference on Economics and Computation*, pages 657–674, 2014.
- 19 P.W. Goldberg and A. Roth. Bounds for the query complexity of approximate equilibria. In *Proceedings of the 15th ACM Conference on Economics and Computation*, pages 639–656, 2014.

- 20 P.W. Goldberg and S. Turchetta. Query complexity of approximate equilibria in anonymous games. In *Proceedings of the 11th Conference on Web and Internet Economics*, 2015.
- 21 S. Hart. Adaptive heuristics. *Econometrica*, 73(5):1401–1430, 2005.
- 22 S. Hart and A. Mas-Colell. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68(5):1127–1150, 2000.
- 23 S. Hart and N. Nisan. The query complexity of correlated equilibria. In *Proceedings of the 6th International Symposium on Algorithmic Game Theory*, 2013.
- 24 C.A. Holt and A.E. Roth. The Nash equilibrium: A perspective. *Proceedings of the National Academy of Sciences*, 101(12):3999–4002, 2004.
- 25 A.X. Jiang and K. Leyton-Brown. Polynomial-time computation of exact correlated equilibrium in compact games. In *Proceedings of the 12th ACM Conference on Electronic Commerce*, 2011.
- 26 R. Kannan and T. Theobald. Games of fixed rank: A hierarchy of bimatrix games. In *Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1124–1132, 2007.
- 27 R.J. Lipton, E. Markakis, and A. Mehta. Playing large games using simple strategies. In *Proceedings of the 4th ACM Conference on Electronic Commerce*, pages 36–41, 2003.
- 28 R. Mehta. Constant rank bimatrix games are PPAD-hard. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, pages 545–554, 2014.
- 29 J.F. Nash. Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences*, 36(1):48–49, 1950.
- 30 C.H. Papadimitriou and T. Roughgarden. Computing correlated equilibria in multi-player games. *Journal of the ACM*, 55(3), 2008.
- 31 T. Roughgarden and O. Weinstein. On the communication complexity of approximate fixed points. In *Proceedings of the 57th Annual IEEE Symposium on Foundations of Computer Science*, pages 229–238, 2016.
- 32 A. Rubinfeld. Inapproximability of Nash equilibrium. In *Proceedings of the 46th ACM Symposium on Theory of Computing*, pages 409–418, 2015.
- 33 A. Rubinfeld. Settling the complexity of computing approximate two-player Nash equilibria. In *Proceedings of the 57th Annual IEEE Symposium on Foundations of Computer Science*, pages 258–265, 2016.
- 34 H. Tsaknakis and P.G. Spirakis. An optimization approach for approximate Nash equilibria. In *Proceedings of the 3rd International Workshop on Internet and Network Economics*, pages 42–56, 2007.