Query Complexity of Approximate Equilibria in Anonymous Games

Paul W. Goldberg and Stefano Turchetta(✉)

Department of Computer Science, University of Oxford, Oxford, UK
{paul.goldberg,stefano.turchetta}@cs.ox.ac.uk

Abstract. We study the computation of equilibria of two-strategy anonymous games, via algorithms that may proceed via a sequence of adaptive queries to the game's payoff function, assumed to be unknown initially. The general topic we consider is *query complexity*, that is, how many queries are necessary or sufficient to compute an exact or approximate Nash equilibrium.

We show that exact equilibria cannot be found via query-efficient algorithms. We also give an example of a 2-strategy, 3-player anonymous game that does not have any exact Nash equilibrium in rational numbers. Our main result is a new randomized query-efficient algorithm that finds a $O(n^{-1/4})$-approximate Nash equilibrium querying $\tilde{O}(n^{3/2})$ payoffs and runs in time $\tilde{O}(n^{3/2})$. This improves on the running time of pre-existing algorithms for approximate equilibria of anonymous games, and is the first one to obtain an inverse polynomial approximation in poly-time. We also show how this can be used to get an efficient PTAS. Furthermore, we prove that $\Omega(n \log n)$ payoffs must be queried in order to find any ϵ-well-supported Nash equilibrium, even by randomized algorithms.

1 Preliminaries

This paper studies two-strategy *anonymous* games, in which a large number of players n share two pure strategies, and the payoff to a player depends on the number of players who use each strategy, but not their identities. Due to the this property, these games have a polynomial-size representation. Daskalakis and Papadimitriou [13] consider anonymous games and graphical games to be the two most important classes of concisely-represented multi-player games. Anonymous games appear frequently in practice, for example in voting systems, traffic routing, or auction settings. Although they have polynomial-sized representations, the representation may still be inconveniently large, making it desirable to work with algorithms that do not require all the data on a particular game of interest.

Query complexity is motivated in part by the observation that a game's entire payoff function may be syntactically cumbersome. It also leads to new results

P.W.Goldberg – Supported by EPSRC project EP/K01000X/1.

S. Turchetta – Work based on MSc. thesis at the Technische Universität München.

E. Markakis and G. Schäfer (Eds.): WINE 2015, LNCS 9470, pp. 357–369, 2015.
DOI: 10.1007/978-3-662-48995-6_26

that distinguish the difficulty of alternative solution concepts. We assume that an algorithm has black-box access to the payoff function, via queries that specify an anonymized profile and return one or more of the players' payoffs.

1.1 Definitions and Notation

Anonymous Games. A k-strategy anonymous game is a tuple $(n, k, \{u^i_j\}_{i\in[n],j\in[k]})$ that consists of n players, k pure strategies per player, and a utility function $u^i_j : \{0, \dots, n-1\} \longrightarrow [0, 1]$ for each player $i \in [n]$ (where we use $[n]$ to denote the set $\{1, \dots, n\}$) and every strategy $j \in [k]$, whose input is the number of other players who play strategy one if $k = 2$. The number of payoffs stored by a 2-strategy game is $2n^2$ (generally, $O(n^k)$). As indicated by u^i_j's codomain, we make a standard assumption that all payoffs are normalized into the interval $[0, 1]$.

For all $i \in [n]$, let X_i be a random indicator variable being equal to one if and only if player i plays strategy one. For 2-strategy games, a mixed strategy for i is represented by the probability $p_i := \mathbb{E}[X_i]$ that player i plays strategy one. Let $X_{-i} := \sum_{\ell\in[n]\setminus\{i\}} X_\ell$ be the sum of all the random variables but X_i. The expected utility obtained by player $i \in [n]$ for playing pure strategy $j \in \{1, 2\}$ against X_{-i} is

$$\mathbb{E}[u^i_j(X_{-i})] := \sum_{x=0}^{n-1} u^i_j(x) \cdot \Pr[X_{-i} = x].$$

If i is playing a mixed strategy (i.e., $p_i \in (0, 1)$) her expected payoff simply consists of a weighted average, i.e., $\mathbb{E}[u^i(X)] := p_i \cdot \mathbb{E}[u^i_1(X_{-i})] + (1 - p_i) \cdot \mathbb{E}[u^i_2(X_{-i})]$, where $X := (X_i, X_{-i})$. It is known that $\mathbb{E}[u^i_j(X_{-i})]$, which involves computing the p.m.f. of X_{-i} – a Poisson Binomial Distribution – can be computed in polynomial time (see e.g., [13]).

Exact and Approximate Nash Equilibria. With the above notation, we say that X_i is a best-response if and only if $\mathbb{E}[u^i(X)] \geq \mathbb{E}[u^i_j(X_{-i})]$ for all $j \in \{1, 2\}$. A *Nash equilibrium* (NE) requires the players to be best-responding to each other; therefore, the above best-response condition must hold for every $i \in [n]$. This can be also viewed as no player having an incentive to deviate from her strategy. We consider a relaxation of NE, the notion of an *ϵ-approximate Nash equilibrium* (ϵ-NE), where every player's incentive to deviate is at most $\epsilon > 0$. We say that $(X_i)_{i\in[n]}$, which represents a mixed-strategy profile, constitutes an ϵ-NE if for all $i \in [n]$ and all $j \in \{1, 2\}$,

$$\mathbb{E}[u^i(X)] + \epsilon \geq \mathbb{E}[u^i_j(X_{-i})].$$

This definition, however, does not prohibit allocating a small amount of probability to arbitrarily bad strategies. An *ϵ-approximate well-supported Nash equilibrium* (ϵ-WSNE) addresses this issue by forcing every player to place a positive amount of probability solely on ϵ-approximate best-responses, i.e., $(X_i)_{i\in[n]}$ constitutes an ϵ-WSNE if for all $i \in [n]$,

$$\mathbb{E}[u_1^i(X_{-i})] + \epsilon < \mathbb{E}[u_2^i(X_{-i})] \Longrightarrow p_i = 0, \text{ and}$$
$$\mathbb{E}[u_2^i(X_{-i})] + \epsilon < \mathbb{E}[u_1^i(X_{-i})] \Longrightarrow p_i = 1.$$

Although an ϵ-WSNE is also an ϵ-NE, the converse need not be true.

Query-Efficiency and Payoff Query Models. Our general interest is in polynomial-time algorithms that find solutions of anonymous games, while checking just a small fraction of the $2n^2$ payoffs of an n-player, 2-strategy game. The basic kind of query is a *single-payoff query* which receives as input a player $i \in [n]$, a strategy $j \in \{1, 2\}$, and the number $x \in \{0, \ldots, n-1\}$ of players playing strategy one, and it returns the corresponding payoff $u_j^i(x)$. The *query complexity* of an algorithm is the expected number of single-payoff queries that it needs in the worst case. Hence, an algorithm is query-efficient if its query complexity is $o(n^2)$.

A *profile query* (used in [15]) consists of an action profile $(a_1, \ldots, a_n) \in \{1, 2\}^n$ as input and outputs the payoffs that *every* player i obtains according to that profile. Clearly, a profile query can be simulated using n single-payoff queries. Finally, an *all-players query* consists of a pair (x, j) for $x \in \{0, \ldots, n-1\}$, $j \in \{1, 2\}$, and the response to (x, j) consists of the values $u_j^i(x)$ for all $i \in [n]$. We will consider the cost of a query to be equal to the number of payoffs it returns; hence, a profile or an all-players query costs n single-payoff queries. We find that an algorithm being constrained to utilize profile queries may incur a linear loss in query-efficiency[1]. Therefore, we focus on single-payoff and all-players queries, which better exploit the symmetries of anonymous games.

1.2 Related Work

In the last decade, there has been interest in the complexity of computing approximate Nash equilibria. A main reason is the **PPAD**-completeness results for computing an exact NE, for normal-form games [5,8] (the latter paper extends the hardness also to an FPTAS), and recently also for anonymous games with 7 strategies [6]. The **FIXP**-completeness results of [14] for multiplayer games show an algebraic obstacle to the task of writing down a useful description of an exact equilibrium. On the other hand, there exists a subexponential-time algorithm to find an ϵ-NE in normal-form games [20], and one important open question regards the existence of a PTAS for bimatrix games.

Daskalakis and Papadimitriou proved that anonymous games admit a PTAS and provided several improvements of its running time over the past few years. Their first algorithm [9] concerns two-strategy games and is based upon the quantization of the strategy space into nearby multiples of ϵ. This result was also

[1] Due to space constraints, we defer this discussion to the full version of the paper (http://arxiv.org/abs/1412.6455).

extended to the multi-strategy case [10]. Daskalakis [7] subsequently gave an efficient PTAS whose running time is $\text{poly}(n) \cdot (1/\epsilon)^{O(1/\epsilon^2)}$, which relies on a better understanding of the structure of ϵ-equilibria in two-strategy anonymous games: There exists an ϵ-WSNE where either a small number of the players – at most $O(1/\epsilon^3)$ – randomize and the others play pure strategies, or whoever randomizes plays the same mixed strategy. Furthermore, Daskalakis and Papadimitriou [11] proved a lower bound on the running time needed by any *oblivious* algorithm, which lets the latter algorithm be essentially optimal. In the same article, they show that the lower bound can be broken by utilizing a non-oblivious algorithm, which has the currently best-known running time for finding an ϵ-equilibrium in two-strategy anonymous games of $O(\text{poly}(n) \cdot (1/\epsilon)^{O(\log^2(1/\epsilon))})$. A complete proof is in [12].

In Sect. 3 we present a bound for λ-Lipschitz games, in which λ is a parameter limiting the rate at which $u^i_j(x)$ changes as x changes. Any λ-Lipschitz k-strategy anonymous game is guaranteed to have an ϵ-approximate *pure* Nash equilibrium, with $\epsilon = O(\lambda k)$ [1,13]. The convergence rate to a Nash equilibrium of best-reply dynamics in the context of two-strategy Lipschitz anonymous games is studied by [2,19]. Moreover, Brandt et al. [4] showed that finding a pure equilibrium in anonymous games is easy if the number of strategies is constant w.r.t. the number of players n, and hard as soon as there is a linear dependence.

In the last two years, several researchers obtained bounds for the query complexity for approximate equilibria in different game settings, which we briefly survey. Fearnley et al. [15] presented the first series of results: they studied bimatrix games, graphical games, and congestion games on graphs. Similar to our negative result for exact equilibria of anonymous games, it was shown that a Nash equilibrium in a bimatrix game with k strategies per player requires k^2 queries, even in zero-sum games. However, more positive results arise if we move to ϵ-approximate Nash equilibria. Approximate equilibria of bimatrix games were studied in more detail in [16].

The query complexity of equilibria of n-player games – a setting where payoff functions are exponentially-large – was analyzed in [3,17,18]. Hart and Nisan [18] showed that exponentially many deterministic queries are required to find a $\frac{1}{2}$-approximate correlated equilibrium (CE) and that any randomized algorithm that finds an exact CE needs $2^{\Omega(n)}$ expected cost. Notice that lower bounds on correlated equilibria automatically apply to Nash equilibria. Goldberg and Roth [17] investigated in more detail the randomized query complexity of ϵ-CE and of the more demanding ϵ-well-supported CE. Babichenko [3] proved an exponential-in-n randomized lower bound for finding an ϵ-WSNE in n-player, k-strategy games, for constant $k = 10^4$ and $\epsilon = 10^{-8}$. These exponential lower bounds do not hold in anonymous games, which can be fully revealed with a polynomial number of queries.

1.3 Our Results and Their Significance

Query-efficiency seems to serve as a criterion for distinguishing exact from approximate equilibrium computation. It applies to games having exponentially-large representations [18], also for games having poly-sized representations (e.g. bimatrix games [15]). Here we extend this finding to the important class of anonymous games. We prove that even in two-strategy anonymous games, an exact Nash equilibrium demands querying the payoff function exhaustively, even with the most powerful query model (Theorem 1). Alongside this, we provide an example of a three-player, two-strategy anonymous game whose unique Nash equilibrium needs all players to randomize with an irrational amount of probability (Theorem 2), answering a question posed in [13]. These results motivate our subsequent focus on approximate equilibria.

We exhibit a simple query-efficient algorithm that finds an approximate pure Nash equilibrium in Lipschitz games (Algorithm 1; Theorem 3), which will be used by our main algorithm for anonymous games.

Our main result (Theorem 4) is a new randomized approximation scheme[2] for anonymous games that differs conceptually from previous ones and offers new performance guarantees. It is query-efficient (using $o(n^2)$ queries) and has improved computational efficiency. It is the first PTAS for anonymous games that is polynomial in a setting where n and $1/\epsilon$ are polynomially related. In particular, its runtime is polynomial in n in a setting where $1/\epsilon$ may grow in proportion to $n^{1/4}$ and also has an improved polynomial dependence on n for all $\epsilon \geq n^{-1/4}$. In more detail, for any $\epsilon \geq n^{-1/4}$, the algorithm adaptively finds a $O(\epsilon)$-NE with $\tilde{O}(\sqrt{n})$ (where we use $\tilde{O}(\cdot)$ to hide polylogarithmic factors of the argument) all-players queries (i.e., $\tilde{O}\left(n^{3/2}\right)$ single payoffs) and runs in time $\tilde{O}\left(n^{3/2}\right)$. The best-known algorithm of [13] runs in time $O(\text{poly}(n) \cdot (1/\epsilon)^{O(\log^2(1/\epsilon))})$, where $\text{poly}(n) \geq O(n^7)$.

In addition to this, we derive a randomized logarithmic lower bound on the number of all-players queries needed to find any non-trivial ϵ-WSNE in two-strategy anonymous games (Theorem 5).

2 Exact Nash Equilibria

We lower-bound the number of single-payoff queries (the least constrained query model) needed to find an exact NE in an anonymous game. We exhibit games in which any algorithm must query most of the payoffs in order to determine what strategies form a NE. Difficult games are ones that only possess NE where $\Omega(n)$ players must randomize.

Example 1. Let G be the following two-strategy, n-player anonymous game. Let n be even, and let $\delta = 1/n^2$. Half of the players have a utility function as shown by the top side (a) of Fig. 1, and the remaining half as at (b).

Theorem 1. *A deterministic single-payoff query-algorithm may need to query $\Omega(n^2)$ payoffs in order to find an exact Nash equilibrium of an n-player, two-strategy anonymous game.*

[2] To make Theorem 4 easier to read, we state it only for the best attainable approximation (i.e., $n^{-1/4}$); however, it is possible to set parameters to get any approximation $\epsilon \geq n^{-1/4}$. For details, see the proof of Theorem 4.

x	0	1	2	...	$n-2$	$n-1$
$u_1^i(x)$	$\frac{1}{2} - \left(\frac{n}{2} - \frac{1}{2}\right)\delta$	$u_1^i(0) - \left(\frac{n}{2} - \frac{3}{2}\right)\delta$	$u_1^i(1) - \left(\frac{n}{2} - \frac{5}{2}\right)\delta$	...	$u_1^i(n-3) + \left(\frac{n}{2} - \frac{3}{2}\right)\delta$	$u_1^i(n-2) + \left(\frac{n}{2} - \frac{1}{2}\right)\delta$
$u_2^i(x)$	$\frac{1}{2}$	$u_1^i(0)$	$u_1^i(1)$	...	$u_1^i(n-3)$	$u_1^i(n-2)$

(a) Payoff table for "majority-seeking" player i

x	0	1	2	...	$n-2$	$n-1$
$u_1^i(x)$	$\frac{1}{2} + \left(\frac{n}{2} - \frac{1}{2}\right)\delta$	$u_1^i(0) + \left(\frac{n}{2} - \frac{3}{2}\right)\delta$	$u_1^i(1) + \left(\frac{n}{2} - \frac{5}{2}\right)\delta$	...	$u_1^i(n-3) - \left(\frac{n}{2} - \frac{3}{2}\right)\delta$	$u_1^i(n-2) - \left(\frac{n}{2} - \frac{1}{2}\right)\delta$
$u_2^i(x)$	$\frac{1}{2}$	$u_1^i(0)$	$u_1^i(1)$	...	$u_1^i(n-3)$	$u_1^i(n-2)$

(b) Payoff table for "minority-seeking" player i

Fig. 1. Majority-minority game G's payoffs. There are $\frac{n}{2}$ majority-seeking players and $\frac{n}{2}$ minority-seeking players. x denotes the number of players other than i who play 1.

The proof of Theorem 1 (in the full version of the paper) shows that in *any* NE of G, at least $n/2$ players must use mixed strategies. Consequently the distribution of the number of players using either strategy has support $\geq n/2$, so for a typical player it is necessary to check $n/2$ of his payoffs.

2.1 A Game Whose Solution Must Have Irrational Numbers

Daskalakis and Papadimitriou [13] note as an open problem, the question of whether there is a 2-strategy anonymous game whose Nash equilibria require players to mix with irrational probabilities. The following example shows that such a game does indeed exist, even with just 3 players. In the context of this paper, it is a further motivation for our focus on approximate rather than exact Nash equilibria.

Example 2. Consider the following anonymous game represented in normal-form in Fig. 2. It can be checked that the game satisfies the anonymity condition. In the unique equilibrium, the row, the column, and the matrix players must randomize respectively with probabilities

$$p_r = \frac{1}{12}(\sqrt{241} - 7),\ p_c = \frac{1}{16}(\sqrt{241} - 7),\ p_m = \frac{1}{36}(23 - \sqrt{241}).$$

Matrix 1:

	1	2
1	$(1, 0, 1)$	$(1, \frac{1}{2}, 0)$
2	$(0, 0, 0)$	$(\frac{1}{2}, \frac{1}{4}, 0)$

Matrix 2:

	1	2
1	$(1, 0, 0)$	$(0, \frac{1}{4}, \frac{1}{2})$
2	$(\frac{1}{2}, 1, \frac{1}{2})$	$(1, 0, 1)$

Fig. 2. The three-player two-strategy anonymous game in normal form. A payoff tuple (a, b, c) represents the row, the column, and the matrix players' payoff, respectively.

Theorem 2. *There exists a three-player, two-strategy anonymous game that has a unique Nash equilibrium where all the players must randomize with irrational probabilities.*

We show in the full version of the paper that Example 2 is a game that does indeed satisfy the conditions of Theorem 2.

3 Lipschitz Games

Lipschitz games are anonymous games where every player's utility function is Lipschitz-continuous, in the sense that for all $i \in [n]$, all $j \in \{1, 2\}$, and all $x, y \in \{0, \ldots, n-1\}$, $\left|u_j^i(x) - u_j^i(y)\right| \leq \lambda \left|x - y\right|$, where $\lambda \geq 0$ is the Lipschitz constant. For games satisfying a Lipschitz condition with a small value of λ, we obtain a positive result (that we apply in the next section) for approximation and query complexity.

Definition 1. Let $(x \in \{0, \ldots, n-1\}, j \in \{1, 2\})$ be the input for an all-players query. For $\delta \geq 0$, a δ-*accurate* all-players query returns a tuple of values $(f_j^1(x), \ldots, f_j^n(x))$ such that for all $i \in [n]$, $|u_j^i(x) - f_j^i(x)| \leq \delta$, i.e., they are within an additive δ of the correct payoffs $(u_j^1(x), \ldots, u_j^n(x))$.

Theorem 3. *Let G be an n-player, two-strategy λ-Lipschitz anonymous game. Algorithm 1 finds a pure-strategy $3(\lambda + \delta)$-WSNE with $4 \log n$ δ-accurate all-players payoff queries.*

The proof (in the full version of the paper) shows how a solution can be found via a binary search on $\{0, \ldots, n-1\}$. Existence of pure approximate equilibria is known already by [13] in the context of k-strategy games. Their proof reduces the problem to finding a Brouwer fixed point. Theorem 3 is used in the next section as part of an algorithm for general anonymous games.

4 General Two-Strategy Anonymous Games

First, we present our main result (Theorem 4). Next, we prove a lower bound on the number of queries that any randomized algorithm needs to make to find any ϵ-WSNE.

4.1 Upper Bound

Before going into technical lemmas, we provide an informal overview of the algorithmic approach. Suppose we are to solve an n-player game G. The first idea is to *smooth* every player's utility function, so that it becomes λ-Lipschitz continuous for some λ. We smooth a utility function by requiring every player to use some amount of randomness. Specifically, for some small ζ we make every player place probability either ζ or $1 - \zeta$ onto strategy one. Consequently, the expected payoff for player i is obtained by averaging her payoff values w.r.t. a sum of two binomial distributions, consisting of a discrete bell-shaped distribution whose standard deviation is at least $\zeta\sqrt{n}$.

We construct the smooth game $\bar{G}$ in the following manner. The payoff received in $\bar{G}$ by player i when x other players are playing strategy one is given by the expected payoff received in G by player i when x other players play one with probability $1-\zeta$ and $n-1-x$ other players play one with probability ζ. This creates a λ-Lipschitz game $\bar{G}$ with $\lambda = O\left(1/\zeta\sqrt{n}\right)$.

Due to dealing with a two-strategy Lipschitz game, we can use the bisection method of Algorithm 1. If we were allowed to query $\bar{G}$ directly, a logarithmic number of all-players queries would suffice. Unfortunately, this is not the case; thus, we need to simulate a query to $\bar{G}$ with a small number of queries to the original game G. Those queries are randomly sampled from the mixed anonymous profile above, and we take enough samples to ensure we get good estimates of the payoffs in $\bar{G}$ with sufficiently high probability.

Thus, we are able to find an approximate pure Nash equilibrium of $\bar{G}$ with $\tilde{O}(\sqrt{n})$ all-players queries. This equilibrium is mapped back to G by letting the players who play strategy one in $\bar{G}$, play it with probability $1-\zeta$ in G, and the ones who play strategy two in $\bar{G}$ place probability ζ on strategy one in G. The quality of the approximation is proportional to $\left(\zeta + (\zeta\sqrt{n})^{-1}\right)$.

Before presenting our main algorithm (Algorithm 2) and proving its efficiency, we state the following lemmas (proven in the full version and used in the proof of Theorem 4).

Lemma 1 [13]**.** Let X, Y be two random variables over $\{0,\ldots,n\}$ such that $\|X-Y\|_{\mathrm{TV}} \leq \delta$ (where $\|X-Y\|_{\mathrm{TV}}$ denotes the total variation distance between X and Y, i.e., $1/2 \cdot \sum_{x=0}^{n} |\Pr[X=x] - \Pr[Y=x]|$). Let $f : \{0,\ldots,n\} \longrightarrow [0,1]$. Then,

$$\sum_{x=0}^{n} f(x) \cdot (\Pr[X=x] - \Pr[Y=x]) \leq 2\delta.$$

Lemma 2 (Simulation of a query to $\bar{G}$ (Algorithm 2)). Let $\delta, \tau > 0$. Let X be the sum of $n-1$ Bernoulli random variables representing a mixed anonymous profile of an n-player game G. Suppose we want to estimate, with additive error δ, the expected payoffs $\mathbb{E}[u_j^i(X)]$ for all $i \in [n], j \in \{1,2\}$. This can be done with probability $\geq 1-\tau$ using $(1/2\delta^2) \cdot \log(4n/\tau)$ all-players queries.

Lemma 3. *Let $X^{(j,n)} := \sum_{i\in[n]} X_i$ denote the sum of n independent* 0-1 *random variables such that $\mathbb{E}[X_i] = 1-\zeta$ for all $i \in [j]$, and $\mathbb{E}[X_i] = \zeta$ for all $i \in [n]\setminus[j]$. Then, for all $j \in [n]$, we have that*

$$\left\|X^{(j-1,n)} - X^{(j,n)}\right\|_{\mathrm{TV}} \leq O\left(\frac{1}{\zeta\sqrt{n}}\right).$$

Definition 2. Let $G = (n, 2, \{u_j^i\}_{i\in[n], j\in\{1,2\}})$ be an anonymous game. For $\zeta > 0$, the ζ-smoothed version of G is a game $\bar{G} = (n, 2, \{\bar{u}_j^i\}_{i\in[n], j\in\{1,2\}})$ defined as follows. Let $X_{-i}^{(x)} := \sum_{j\neq i} X_i$ denote the sum of $n-1$ Bernoulli random variables where x of them have expectation equal to $1-\zeta$, and the remaining ones have expectation equal to ζ. The payoff $\bar{u}^i_j(x)$ obtained by every player $i \in [n]$ for playing strategy $j \in \{1, 2\}$ against $x \in \{0, \ldots, n-1\}$ is

$$\bar{u}^i_j(x) := \sum_{y=0}^{n-1} u^i_j(y) \cdot \Pr\left[X^{(x)}_{-i} = y\right] = \mathbb{E}\left[u^i_j\left(X^{(x)}_{-i}\right)\right].$$

Theorem 4. *Let* $G = (n, 2, \{u^i_j\}_{i\in[n],j\in\{1,2\}})$ *be an anonymous game. For* ϵ *satisfying* $1/\epsilon = O(n^{1/4})$*, Algorithm 2 can be used to find (with probability* $\geq \frac{3}{4}$*) an* ϵ*-NE of* G*, using* $O(\sqrt{n} \cdot \log^2 n)$ *all-players queries (hence,* $O(n^{3/2} \cdot \log^2 n)$ *single-payoff queries) in time* $O(n^{3/2} \cdot \log^2 n)$*.*

Algorithm 1. Approximate NE Lipschitz

Data: δ-accurate query access to utility function $\bar{u}$ of n-player λ-Lipschitz game $\bar{G}$.
Result: pure-strategy $3(\delta + \lambda)$-NE of $\bar{G}$.
begin
  Let $BR_1(i)$ be the number of players whose best response (as derived from the δ-accurate queries) is 1 when i of the other players play 1 and $n - 1 - i$ of the other players play 2.
  Define $\phi(i) = BR_1(i) - i$. `// by construction,` $\phi(0) \geq 0$ `// and` $\phi(n-1) \leq 0$
  If $BR_1(0) = 0$, **return** all-1's profile.
  If $BR_1(n-1) = n$, **return** all-2's profile.
  Otherwise, `// In this case,` $\phi(0) > 0$ `and` $\phi(n-1) \leq 0$
  Find, via binary search, x such that $\phi(x) > 0$ and $\phi(x+1) \leq 0$.
  Construct pure profile $\bar{p}$ as follows:
  For each player i, if $\bar{u}^i_1(x) - \bar{u}^i_2(x) > 2\delta$, let i play 1, and if $\bar{u}^i_2(x) - \bar{u}^i_1(x) > 2\delta$, let i play 2. (The $\bar{u}^i_j$'s are δ-accurate.) Remaining players are allocated either 1 or 2, subject to the constraint that x or $x + 1$ players in total play 1.
  return $\bar{p}$.
end

Proof. Set ζ equal to ϵ and let $\bar{G}$ be the ζ-smoothed version of G. We claim that $\bar{G}$ is a λ-Lipschitz game for $\lambda = O\left((\zeta\sqrt{n})^{-1}\right)$. Let $X^{(x)}_{-i}$ be as in Definition 2. By Lemma 3, $\left\|X^{(x-1)}_{-i} - X^{(x)}_{-i}\right\|_{\text{TV}} \leq O\left(\frac{1}{\zeta\sqrt{n}}\right)$ for all $x \in [n-1]$. Then by Lemma 1, we have

$$\left|\bar{u}^i_j(x-1) - \bar{u}^i_j(x)\right| \leq O\left(\frac{1}{\zeta\sqrt{n}}\right).$$

Theorem 3 shows that Algorithm 1 finds a pure-strategy $3(\lambda + \delta)$-WSNE of $\bar{G}$, using $O(\log n)$ δ-accurate all-players queries. Thus, Algorithm 1 finds a $O(\frac{1}{\zeta\sqrt{n}} + \delta)$-WSNE of $\bar{G}$, where δ is the additive accuracy of queries.

Algorithm 2. Approximate NE general

Data: ϵ; query access to utility function u of n-player anonymous game G; parameters τ (failure probability), δ (accuracy of queries).
Result: $O(\epsilon)$-NE of G.
begin
 Set $\zeta = \epsilon$. Let $\bar{G}$ be the ζ-smoothed version of G, as in Definition 2.
 `// By Lemma 1 and Lemma 3 it follows that`
 `//` $\bar{G}$ `is` λ`-Lipschitz for` $\lambda = O(1/\zeta\sqrt{n})$.
 Apply Algorithm 1 to $\bar{G}$, simulating each all-players δ-accurate query to $\bar{G}$ using multiple queries according to Lemma 2.
 Let $\bar{p}$ be the obtained pure profile solution to $\bar{G}$.
 Construct p by replacing probabilities of 0 in $\bar{p}$ with ζ and probabilities of 1 with $1 - \zeta$.
 return p.
end

Despite not being allowed to query $\bar{G}$ directly, we can simulate any δ-accurate query to $\bar{G}$ with a set of randomized all-players queries to G. This is done in the body of Algorithm 2. By Lemma 2, for $\tau > 0$, $(1/2\delta^2)\log(4n/\tau)$ randomized queries to G correctly simulate a δ-accurate query to $\bar{G}$ with probability $\geq 1-\tau$.

In total, the algorithm makes $O\left(\log n \cdot (1/\delta^2) \cdot \log(n/\tau)\right)$ all-players payoff queries to G. With a union bound over the $4\log n$ simulated queries to $\bar{G}$, this works with probability $1 - 4\tau \log n$.

Once we find this pure-strategy $O\left(\frac{1}{\zeta\sqrt{n}} + \delta\right)$-WSNE of $\bar{G}$, the last part of Algorithm 2 maps the pure output profile to a mixed one where whoever plays 1 in $\bar{G}$ places probability $(1-\zeta)$ on 1, and whoever plays 2 in $\bar{G}$ places probability ζ on 1. It is easy to verify that the regret experienced by player i (that is, the difference in payoff between i's payoff and i's best-response payoff) in G is at most ζ more than the one she experiences in $\bar{G}$.

The extra additive ζ to the regret of players means that we have an ϵ-NE of G with $\epsilon = O(\zeta+\delta+\frac{1}{\zeta\sqrt{n}})$. The query complexity thus is $O(\log n \cdot (1/\delta^2) \cdot \log(n/\tau))$.

Setting $\delta = 1/\sqrt[4]{n}$, $\zeta = 1/\sqrt[4]{n}$, $\tau = 1/16\log n$, we find an $O(1/\sqrt[4]{n})$-Nash equilibrium using $O(\sqrt{n} \cdot \log^2 n)$ all-players queries with probability at least 3/4. We remark that the above parameters can be chosen to satisfy any given approximation guarantee $\epsilon \geq n^{-1/4}$, i.e., simply find solutions to the equation $\epsilon = \zeta + \delta + (\zeta\sqrt{n})^{-1}$. This allows for a family of algorithms parameterized by ϵ, for $\epsilon \in [n^{-1/4}, 1)$, thus an approximation scheme.

The runtime is equal to the number of single-payoff queries and can be calculated as follows. Calculating the value of $\phi(i)$ in Algorithm 1 takes $O(n\sqrt{n}\log n)$. We make $O(\sqrt{n}\log n)$ queries to G to simulate one in $\bar{G}$, and once we gather all the information, we need an additional linear time to count the number of players whose best response is 1. The fact that the above part is performed at every step of the binary search implies a total running time of $O(n^{3/2} \cdot \log^2 n)$

for Algorithm 1. Algorithm 2 simply invokes Algorithm 1 and only needs linear time to construct the profile p; thus, it runs in the same time. □

4.2 Lower Bound

We use the minimax principle and thus define a distribution over instances that will lead to the lower bound on query complexity, for any deterministic algorithm. We specify a distribution over certain games that possess a unique pure Nash equilibrium. The n players that participate in any of these games are partitioned into $\log n$ groups, which are numbered from 1 to $\log n$. Group i's equilibrium strategy depends on what all the previous groups $\{1, \ldots, i-1\}$ play at equilibrium. Hence, finding out what the last group should play leads to a lower bound of $\Omega(\log n)$ all-players queries.

Lemma 4. *Let $\mathcal{G}_n$ be the class of n-player two-strategy anonymous games such that $u_1^i(x) = 1 - u_2^i(x)$ and $u_1^i(x) \in \{0, 1\}$, for all $i \in [n], x \in \{0, \ldots, n-1\}$. Then, there exists a distribution $\mathcal{D}_n$ over $\mathcal{G}_n$ such that every G drawn from $\mathcal{D}_n$ has a unique (pure-strategy) ϵ-WSNE.*

Theorem 5. *Let $\mathcal{G}_n$ be defined as in Lemma 4. Then, for any $\epsilon \in [0, 1)$, any randomized all-players query algorithm must make $\Omega(\log n)$ queries to find an ϵ-WSNE of $\mathcal{G}_n$ in the worst case.*

5 Conclusions and Further Work

Our interest in the query complexity of anonymous games has resulted in an algorithm that has an improved runtime-efficiency guarantee, although limited to when the number of strategies k is equal to 2. Algorithm 2 (Theorem 4) finds an ϵ-NE faster than the PTAS of [13], for any $\epsilon \geq 1/\sqrt[4]{n}$. In particular, for $\epsilon = 1/\sqrt[4]{n}$, their algorithm runs in subexponential time, while ours is just $\tilde{O}(n^{3/2})$; however, our ϵ-NE is not well-supported.

An immediate question is whether we can obtain sharper bounds on the query complexity of two-strategy games. There are ways to potentially strengthen the results. First, our lower bound holds for well-supported equilibria; it would be interesting to know whether a logarithmic number of queries is also needed to find an ϵ-NE for $\epsilon < \frac{1}{2}$. We believe this is the case at least for small values of ϵ. Second, the ϵ-NE found by our algorithm are not well-supported since all players are forced to randomize. Is there a query-efficient algorithm that finds an ϵ-WSNE? Third, we may think of generalizing the algorithm to the (constant) k-strategy case by letting every player be obliged to place probability either $\frac{\zeta}{k}$ or $1 - \frac{k-1}{k}\zeta$ and obtain a similar smooth utility function. However, in this case we cannot use a bisection algorithm to find a fixed point of the smooth game. As a consequence, the query complexity might be strictly larger.

References

1. Azrieli, Y., Shmaya, E.: Lipschitz games. Math. Oper. Res. **38**(2), 350–357 (2013)
2. Babichenko, Y.: Best-reply dynamics in large binary-choice anonymous games. Games Econ. Behav. **81**(1), 130–144 (2013)
3. Babichenko, Y.: Query complexity of approximate Nash equilibria. In: Proceedings of the 46th Annual ACM Symposium on Theory of Computing. pp. 535–544. STOC 2014. ACM, USA (2014)
4. Brandt, F., Fischer, F., Holzer, M.: Symmetries and the complexity of pure Nash equilibrium. J. Comput. Syst. Sci. **75**, 163–177 (2009)
5. Chen, X., Deng, X., Teng, S.: Settling the complexity of computing two-player Nash equilibria. J. ACM **56**(3), 1–57 (2009)
6. Chen, X., Durfee, D., Orfanou, A.: On the complexity of Nash equilibria in anonymous games. In: Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015, ACM, pp. 381–390 (2015)
7. Daskalakis, C.: An efficient PTAS for two-strategy anonymous games. In: Papadimitriou, C., Zhang, S. (eds.) WINE 2008. LNCS, vol. 5385, pp. 186–197. Springer, Heidelberg (2008)
8. Daskalakis, C., Goldberg, P.W., Papadimitriou, C.H.: The complexity of computing a Nash equilibrium. SIAM J. Comput. **39**(1), 195–259 (2009)
9. Daskalakis, C., Papadimitriou, C.H.: Computing equilibria in anonymous games. In: Proceedings of the 48th Symposium on Foundations of Computer Science (FOCS), pp. 83–93 (2007)
10. Daskalakis, C., Papadimitriou, C.H.: Discretized multinomial distributions and Nash equilibria in anonymous games. In: Proceedings of the 49th Symposium on Foundations of Computer Science (FOCS), pp. 25–34 (2008)
11. Daskalakis, C., Papadimitriou, C.H.: On oblivious PTAS's for Nash equilibrium. In: Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, pp. 75–84. ACM, USA (2009)
12. Daskalakis, C., Papadimitriou, C.H.: Sparse covers for sums of indicators (2013). CoRR abs/1306.1265
13. Daskalakis, C., Papadimitriou, C.H.: Approximate Nash equilibria in anonymous games. J. Econ. Theory **156**, 207–245 (2015)
14. Etessami, K., Yannakakis, M.: On the complexity of Nash equilibria and other fixed points. SIAM J. Comput. **39**(6), 2531–2597 (2010)
15. Fearnley, J., Gairing, M., Goldberg, P.W., Savani, R.: Learning equilibria of games via payoff queries. In: Proceedings of the 14th ACM Conference on Electronic Commerce, EC 2013, pp. 397–414. ACM, USA (2013)
16. Fearnley, J., Savani, R.: Finding approximate Nash equilibria of bimatrix games via payoff queries. In: Proceedings of the Fifteenth ACM Conference on Economics and Computation, EC 2014, pp. 657–674. ACM, USA (2014)
17. Goldberg, P.W., Roth, A.: Bounds for the query complexity of approximate equilibria. In: Proceedings of the Fifteenth ACM Conference on Economics and Computation, EC 2014, pp. 639–656. ACM, USA (2014)
18. Hart, S., Nisan, N.: The query complexity of correlated equilibria. In: Vöckling, B. (ed.) SAGT 2013. LNCS, vol. 8146, p. 268. Springer, Heidelberg (2013)

19. Kash, I.A., Friedman, E.J., Halpern, J.Y.: Multiagent learning in large anonymous games. In: Eighth International Conference on Autonomous Agents and Multiagent Systems (AAMAS), pp. 765–772 (2009)
20. Lipton, R.J., Markakis, E., Mehta, A.: Playing large games using simple strategies. In: Proceedings of the 4th ACM Conference on Electronic Commerce, EC 2003, pp. 36–41. ACM, USA (2003)